# A Stochastic Price Duration Model for Estimating High-Frequency Volatility

Denis Pelletier*

Wei Wei†‡

February 1, 2018

**Abstract**

We propose a class of stochastic price duration models to estimate high-frequency volatility. Price duration is directly linked to volatility from the passage time theory for Brownian motions and it possesses several advantages over returns for estimating volatility. We employ price durations in a parametric model where stochastic volatility dynamics is directly specified. Our parametric approach allows us to estimate intraday spot volatility and incorporate additional information such as trade durations. In the simulation study we find our proposed models compare favorably to other duration-based estimators and return-based estimators.

*JEL Codes*: C41, C51, C58, G1

*Keywords*: Price durations, Stochastic volatility, High-frequency data

*Department of Economics, North Carolina State University, Raleigh, NC 27695, USA, email: denis_pelletier@ncsu.edu

†CREATES, Department of Economics and Business Economics, Aarhus University

‡Department of Econometrics and Business Statistics, Monash University, email: wei.wei2@monash.edu

# 1 Introduction

Volatility plays a pivotal role in modern day financial economics. Since asset prices are generally considered to be driven by Brownian motions, the natural way to estimate volatility is to look at how much price changes in a given time interval. The passage time theory for Brownian Motion offers an alternative approach: one can look at how long it takes for the price to change by a given amount. Let price duration refer to the time it takes for the logarithmic price to travel the distance $\delta$. If volatility is constant, the expectation of the price duration is inversely proportional to the variance. Intuitively, if volatility is high, price will be changing quickly and the expected price duration will be relatively short.

In financial markets, volatility is not constant. The time-varying nature of volatility has spurred a large literature on its estimation using returns. There are roughly three classes of return-based volatility estimation, namely GARCH, stochastic volatility (SV) and realized volatility (RV). GARCH-type models[1] assume that volatility is some deterministic function of past returns. In the SV-type models, volatility is assumed to follow a stochastic process. The availability of high-frequency financial data has popularized the RV estimator, which uses returns sampled at shorter horizons (such as 5 minutes) to measure integrated variance over a longer horizon, typically a day. The RV approach allows volatility to be stochastic without specifying any parametric form.

The literature on duration-based volatility estimation is considerably smaller. Cho and Frees (1988) are the first to use passage times of Brownian Motion to estimate volatility, assuming it is constant. Engle and Russell (1998) propose the autoregressive conditional duration (ACD) model for durations between trades. They apply the model to price durations by treating the price arrival times as a point process, and linking the

[1]See Hansen and Lunde (2005) for a list of 330 specifications in the GARCH universe and their evaluation.

price arrival intensity to volatility. ACD is similar to GARCH—the volatility traced out from price intensity only depends on past price durations. Tse and Yang (2012) adopt an augmented ACD specification to model price durations and estimate high-frequency volatility in what they call the ACD-ICV method. They find that ACD-ICV outperforms many version of RV methods in Monte Carlo exercises. Andersen, Dobrev, and Schaumburg (2009) introduce a family of nonparametric volatility estimation using different types of passage times. Their duration-based volatility (DV) estimators is a natural dual approach to realized volatility: both are nonparametric, assume volatility is stochastic, and focus on estimating integrated variance. Their duration-based estimator compares favorably to many robust RV type estimators.

In this paper, we proposes a class of stochastic price duration (SPD) models to estimate high-frequency volatility parametrically. The SPD models specify a stochastic process for volatility directly as the traditional SV models, but use price durations instead of returns to infer the latent volatility. Our models focus on spot volatility, which is important for intraday dynamics (see e.g. Zu and Boswijk, 2014), although integrated variance is easily obtained from the estimates.

We consider two SPD models: the baseline model *SPD0*, and the extended model *SPD1*. In the baseline model, logarithmic variance follows an Ornstein-Uhlenbeck (OU) diffusion, which leads to an AR(1) process when discretized. The baseline model can be augmented by incorporating additional information. In particular, we consider modeling trade durations together with price durations. The asymmetric information models by Easley and O'Hara (1987) suggest that trades durations have an interdependent relationship with volatility. Pelletier and Zheng (2012) and Wei and Pelletier (2013) propose a bivariate OU process for the logarithmic volatility and conditional trade duration. We adopt a simplified version of their bivariate OU process for the *SPD1* model.

Our propose SPD models have several advantages over competitors. First, compared to return-based approaches, duration-based approaches, including the ACD-ICV method and the DV estimator, possess some natural robustness to jumps and price discreteness. A jump might shorten the price duration, but the amount by which the price change exceeds the threshold does not matter. Price discreteness is an important source of market microstructure noise and it results in zero returns that complicates volatility estimation. Duration-based methods measure time conditional on a pre-specified price change hence they are less affected. Second, compared to non-parametric approaches, the SPD models utilize the persistence of volatility and estimate intraday volatility. Also, we can extend the model to incorporate additional information, such as trade durations. Last, compared to the ACD-ICV method, the SPD models specify volatility dynamics directly rather than through the condition duration. This allows simple interpretation of the volatility dynamics and simple computation of integrated variance or other power-variations.

The rest of this paper is organized as follows: Section 2 describes the model specification. Section 3 discusses the estimation procedure. Section 4 conducts simulation studies and discusses the selection of price thresholds. Section 5 presents empirical results using IBM high-frequency trade data in the year 2011. Section 6 concludes.

# 2 Model Specifications

In this section, we first build our baseline model *SPD0* and the extended model *SPD1*. Next, we discuss how prominent features in high frequency data—in particular diurnal patterns, finite activity jumps, and market microstructure noise—affect the price duration models. Last, we compare our models to the price duration model based on ACD.

## 2.1 The *SPD0* Model

We start by assuming that the logarithmic asset price $y_t$ follows

$$dy_t = \sqrt{V_t} dW_t^y, \tag{1}$$

where $\sqrt{V_t}$ is the latent instantaneous volatility and

$$d\log V_t = -\kappa^v(\log V_t - \mu^v)dt + \sigma_v dW_t^v. \tag{2}$$

The standard Brownian motions $W_t^y$ and $W_t^v$ are assumed to be independent, i.e., there is no leverage effect. The logarithmic variance, $\log V_t$, follows an OU diffusion; this specification is a continuous-time version of the popular lognormal autoregressive volatility model, see e.g. Andersen, Bollerslev, and Diebold (2010). The OU diffusion has three parameters: $\kappa^v$ measures the persistence, $\sigma_v$ the variability, and $\mu^v$ the mean of log-variance. If $\kappa^v > 0$, log-variance is stationary with marginal distribution $\log V_t \sim N(\mu^v, \sigma_v^2/2\kappa^v)$, and the long-run mean of $V_t$ equals $\exp(\mu^v + \sigma_v^2/4\kappa^v)$.

Given the continuous-time model (1), return-based estimators fix a time interval $\Delta t$ and measure the change in $y$ during this time interval. Let $r_{t+\Delta t} = y_{t+\Delta t} - y_t$, the continuous compounded return follows a mixed normal distribution, $r_{t+\Delta t} \sim N(0, \int_t^{t+\Delta t} V_s ds)$. If $\Delta t$ is small, $\int_t^{t+\Delta t} V_s ds$ can be approximated by $V_t\Delta t$. Hence, we can use the equation $r_{t+\Delta t} \overset{d}{=} \sqrt{V_t \Delta t}\epsilon_{t+\Delta t}$, where $\epsilon_{t+\Delta t} \sim N(0,1)$, to draw inference on volatility from returns.

In contrast, duration-based estimators fix a price threshold $\delta$ and measure the change in $t$ when $y$ increases or decreases by the amount $\delta$. The change in $t$ is the price duration. Specifically, the $i$+1th price duration, $\tau_{i+1}$, is defined by $\tau_{i+1} = \inf\{s > 0 | \, |y_{t_i+s} - y_{t_i}| \geq \delta\}$, where $t_i$ denotes the time when the change in $y_t$ crosses the threshold for the $i$th time.

The sequence $\{t_i\}_{i=0}^N$ partitions the time line $[t_0, t_N]$ into $N$ intervals, while each interval corresponds to a price duration, i.e., $\tau_{i+1} = t_{i+1} - t_i$.

We approximate $V_t$ by a piecewise constant process: the instantaneous variance in the interval $[t_i, t_{i+1}]$ is approximated by $V_i$, the variance at the left end point of the interval. The distribution of $\tau_{i+1}$ is then obtained from the passage time theory for Brownian motion (see Andersen, Dobrev, and Schaumburg, 2009):

$$\tau_{i+1} \stackrel{d}{=} \frac{\delta^2}{V_i}\eta_{i+1}, \tag{3}$$

with the pdf of $\eta_{i+1}$ given by

$$p(\eta_{i+1}) = \sum_{k=-\infty}^{\infty} \frac{2(1+4k)}{\sqrt{2\pi}\eta_{i+1}^{3/2}} e^{-\frac{(1+4k)^2}{2\eta_{i+1}}}. \tag{4}$$

The random variable $\eta_{i+1}$ is the standardized price duration with volatility and price threshold both equal to 1. In passage time theory, $\eta$ is referred to as the first exit time; it records the time when a standard Brownian motion first exits the band $[-1, 1]$.

The inverse relationship between price duration and volatility, as evident in equation (3), yields the basis for estimating volatility from price durations. Parametric inferences requires another element—the transitional dynamics of volatility. Discretizing the OU diffusion in equation (2) using price durations, we obtain

$$\log V_{i+1} = \left(1 - e^{-\kappa^v \tau_{i+1}}\right) \mu^v + e^{(-\kappa^v \tau_{i+1})} \log V_i + u_{i+1}^v, \tag{5}$$

where

$$u_{i+1}^v \sim N\left(0, \frac{\sigma_v^2}{2\kappa_v}\left(1 - e^{(-2\kappa^v \tau_{i+1})}\right)\right).$$

Equation (3) and (5) form the baseline model *SPD0*. It is a non-linear and non-Gaussian state space model where (3) is the observation equation and (5) is the transition equation.

## 2.2 The *SPD1* Model

Asset prices are not recorded continuously in time. Instead, trades or quote revisions occur in random time intervals. Market microstructure theory suggest thats the time intervals between trades, defined as trade durations, is interdependent with volatility. When the volatility of an asset is high, trades tend to happen more frequently, and vice versa. To exploit this relationship, Pelletier and Zheng (2012) and Wei and Pelletier (2013) propose a joint model for returns and trade durations. We adopt their framework and adapt it to model price durations and trade durations jointly in the *SPD1* model.

### 2.2.1 Stochastic Trade Duration

The trade duration $D_{j+1}$ is defined as the time interval between a trade that occurred at $t_j$ and the next trade at $t_{j+1}$. Without loss of generality, we assume that

$$D_{j+1} = \lambda_{t_j} \varepsilon_{j+1}, \tag{6}$$

where $\lambda_{t_j}$ denote the conditional expectation of $D_{j+1}$ given the information set available at $t_j$, i.e., $E(D_{j+1}|I_{t_j}) = \lambda_{t_j}$, and $\varepsilon_{j+1}$ is a random variable with mean equal to 1. If $\lambda_{t_j}$ is constant and $\varepsilon_{j+1}$ is i.i.d. exponential, the trade arrivals are characterized by a Poisson process.

In financial markets, trade arrivals are often clustered, and trade durations autocorrelated. To capture the dynamics in trade durations, Engle and Russell (1998) adopt an autoregressive structure for the conditional duration, $\lambda_{t_j}$, and generalize the distribution of $\varepsilon_{j+1}$ to Weibull in their ACD model. Bauwens and Veredas (2004) consider a stochastic autoregressive model for $\log \lambda_{t_j}$, and a Gamma or Weibull distribution for $\varepsilon_{j+1}$ in their stochastic conditional duration (SCD) model[2]. We specify a OU diffusion for the logarithmic conditional duration:

$$d \log \lambda_t = -\kappa^\lambda (\log \lambda_t - \mu^\lambda) dt + \sigma_\lambda dW_t^\lambda,$$

where $\kappa^\lambda$, $\mu^\lambda$ and $\sigma_\lambda$ denote the mean-reversion speed, the mean, and the volatility of the OU diffusion, respectively. This continuous-time specification allows easy discretization at any sampling intervals, in contrast to the ACD or SCD model. For $\varepsilon_{j+1}$, we adopt the Gamma distribution, which has an attractive feature that the sum of i.i.d. Gamma random variables remains Gamma-distributed. Specifically, $\varepsilon_{j+1} \sim \Gamma(d_s, 1/d_s)$, where $d_s$ is the shape parameter. The scale parameter is fixed at $1/d_s$ to ensure that the mean of $\varepsilon_{j+1}$ is equal to 1, and the conditional mean of $D_{j+1}$ is separated from the mean of $\varepsilon_{j+1}$.

Suppose that $M$ trades happened in a time interval with length $\tau$, we are interested in the distribution of $\tau$ given $M$ and $\lambda_t$. If the conditional duration $\lambda_t$ is constant within the time interval, $\tau$ is the sum of $M$ i.i.d. random variables each distributed as $\Gamma(d_s, \lambda_t/d_s)$, hence $\tau \sim \Gamma(d_s M, \lambda_t/d_s)$. Moreover, the average trade duration, defined by $d^a = \tau/M$, is also Gamma-distributed, $d^a \sim \Gamma(d_s M, \lambda_t/d_s M)$. Using the scaling property of Gamma distribution, we can write $d^a$ as $\lambda_t$ multiplied by a random variable with a $\Gamma(d_s M, 1/d_s M)$ distribution.

In general, if we observe $M_{i+1}$ trades in the time interval $[t_i, t_{i+1}]$ with $\tau_{i+1} = t_{i+1} - t_i$, the average trade duration $d_{i+1}^a = \tau_{i+1}/M_{i+1}$ can be written as

$$d_{i+1}^a = \lambda_i e_{i+1}, \tag{7}$$

[2] Bauwens and Veredas (2004) also apply the model to price durations and volume durations although they do not directly specify or measure volatility.

where $e_{i+1} \sim \Gamma(d_s M_{i+1}, 1/d_s M_{i+1})$, and $\lambda_i$ is the conditional duration at time $t_i$. It is easy to see that $E(d^a_{i+1}|I_{t_i}) = \lambda_i$.

### 2.2.2 Modeling Price Durations and Trade Durations Jointly

To create persistence and interdependence between volatility and trade duration, we model the logarithm of $\lambda_t$ and $V_t$ using a bivariate OU diffusion. Let $x_t = (\log(V_t), \log(\lambda_t))'$, $x_t$ solves:

$$dx_t = -\Psi(x_t - \mu^x)dt + S_x dW^x_t, \tag{8}$$

where $\Psi = \text{diag}(\kappa^v, \kappa^\lambda)$, and $\mu^x = (\mu^v, \mu^\lambda)$, and $S_x = \text{diag}(\sigma_v, \sigma_\lambda)$. $W^x_t$ is a Brownian motion in $\mathcal{R}^2$ with $dW^v_t dW^\lambda_t = \rho dt$, where $\rho$ is the instantaneous correlation. The instantaneous covariance matrix is given by

$$\Sigma_x = S_x \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} S_x = \begin{pmatrix} \sigma_v^2 & \rho\sigma_v\sigma_\lambda \\ \rho\sigma_v\sigma_\lambda & \sigma_\lambda^2 \end{pmatrix}.$$

The observables for the *SPD1* model are price durations, $\{\tau_i\}_{i=1}^N$, and average trade duration, $\{d^a_i\}_{i=1}^N$. We discretized the bivariate OU diffusion using price durations and assume that volatility and conditional durations are constant within each price duration. This construction leads to a bivariate non-linear and non-Gaussian state space model. The observation equations are

$$\begin{aligned} \tau_{i+1} &= \frac{\delta^2}{V_i}\eta_{i+1} \\ d^a_{i+1} &= \lambda_i e_{i+1}, \end{aligned} \tag{9}$$

where the distribution of $\eta_{i+1}$ is given in (4) and $e_{i+1} \sim \Gamma(d_s M_{i+1}, 1/d_s M_{i+1})$. The transition equation is

$$x_{i+1} = (I_2 - e^{-\Psi\tau_{i+1}})\mu^x + e^{-\Psi\tau_{i+1}}x_i + u_{i+1}, \tag{10}$$

where

$$\begin{aligned} &u_{i+1} \sim N(0, \Sigma_{i+1}) \\ &vec(\Sigma_{i+1}) = (\Psi \oplus \Psi)^{-1}(I_2 - e^{-(\Psi\oplus\Psi)\tau_{i+1}})vec(\Sigma_x), \end{aligned}$$

and $e^{-\Psi\tau_{i+1}}$ is a matrix exponential. Provided $\Psi = \text{diag}(\kappa^v, \kappa^\lambda)$, $e^{-\Psi\tau_{i+1}}$ simply equals to $\text{diag}(e^{-\kappa^v\tau_{i+1}}, e^{-\kappa^\lambda\tau_{i+1}})$. The transition equation comes from the strong solution to equation (8). We use this exact discretization scheme rather than an Euler Scheme because the discretization error in the Euler scheme depends on $\tau_{i+1}$.

## 2.3 Diurnal Patterns

Intraday volatility and trade durations have well known diurnal patterns, see e.g. Andersen and Bollerslev (1997) and Tse and Dong (2014). When the market first opens or is near closing, trades happen more frequently, accompanied by higher volatility. Left undealt with, the deterministic patten can distort our estimates of the parameters and states. We specify two quadratic functions[3] to approximate the diurnal patterns in volatility and condition duration in the *SPD1* model (the procedure for the *SPD0* model naturally follows). Let $V_i^*$ and $\lambda_i^*$ denote the deseasonalized volatility and conditional duration, and $x_i = (\log V_i^*, \log \lambda_i^*)'$ follows the transitional equation (10). The observation

[3]The choice of a quadratic function is a trade-off between better approximation and less parameters to estimate. Higher order approximation may improve the fit, and we leave that to future work.

equation (9) is augmented by

$$
\begin{aligned}
V_i &= V_i^* g_v(t_i) \\
\lambda_i &= \frac{\lambda_i^*}{g_d(t_i)},
\end{aligned} \tag{11}
$$

where

$$
\begin{aligned}
g_v(t) &= a_1(t + a_2)^2 + 1 \\
g_d(t) &= a_3(t + a_4)^2 + 1.
\end{aligned} \tag{12}
$$

This specification produces the U-shaped patten in volatility and the inverse U-shaped patten in the conditional duration. The level of the quadratic function is fixed by setting its minimum to 1, otherwise the mean of the latent process $x_t$ becomes unidentifiable.

### 2.4 Integrated Variance and Finite Activity Jumps

The *SPD0* and *SPD1* models are cast in high frequency and they allow us to understand and estimate the dynamics of intraday spot volatility. Another important volatility measure is the integrate variance over longer time intervals, in particularly a day. Let $IV_k$ denote the integrate variance over day $k$, it is defined by $IV_k = \int_{s \in day(k)} V_s ds$. An estimate of the integrated variance is easily obtained from the spot variance estimates from *SPD0* and *SPD1* model:

$$
\widehat{IV_k} = \sum_{t_i \in day(k)} V_i \tau_{i+1}. \tag{13}
$$

If the logarithmic price process follows the diffusion in equation (1), the integrated variance equals the quadratic variation of $y_t$. However, even with stochastic volatility, diffusions tend to have a hard time explaining the fat tails in financial returns, and adding jumps to the price process is becoming common practice. In an off-the-shelf jump diffusion model, the log-price follows

$$dy_t = \sqrt{V_t} dW_t^y + dJ_t, \tag{14}$$

where $J_t$ is a compound Poisson process with jump intensity $\gamma_J$, and normally distributed jump sizes, $N(\mu_J, \sigma_J^2)$. The quadratic variation in $y_t$ in this model comprises the integrated variance, defined same as before, and jump variation, defined as the sum of squared jumps.

Although the jump term would change its distribution, the price duration possesses a large degree of natural robustness to finite activity jumps, as demonstrated by Andersen, Dobrev, and Schaumburg (2009) and Tse and Yang (2012). When a jump occurs in the price process, it might shorten the price duration, but the amount by which the price change exceeds the threshold does not impact the observables. If jumps are rare in the sense that they have finite activity, only a small proportion of price durations would be affected. In section 3, we show in simulation that price durations are robust to jumps and that the $\widehat{IV}_k$ computed from equation (13) is indeed an estimate of the integrated variance.

## 2.5 Market Microstructure Noise

In the high frequency world, market microstructure—such as price discreteness, bid-ask bounce, and various trading mechanisms—renders the efficient price unobservable. The observed price, either from trades or from quotes, is a noisy version of the efficient price. Let $y_i^o$ denotes the observed log-price at time $t_i$ and $y_i$ the efficient log-price, the microstructure noise, given by $m_i = y_i^o - y_i$, is often assumed to be independent of the price process. As the sampling interval decreases, the magnitude of $m_i$ remained unchanged, whereas the integrated variance decreases, and hence the noise-to-signal ratio increases.

To mitigate the effect of microstructure noise, one can sample sparsely, such that the noise-to-signal ratio becomes negligible. In return-based volatility estimation, common practice is to use 5-min returns. In duration-based estimator, we adjust sampling frequency by adjusting the price threshold $\delta$. Andersen, Dobrev, and Schaumburg (2009) recommend choosing the price threshold to be larger than three times the log-spread. Alternatively, we can set the price threshold so that the average price duration matches the commonly used sampling interval for returns, such as 5 min. In our simulation studies, we also investigate setting $\delta$ according to the level of microstructure noise, which can be estimated using returns samples at the highest frequency, see e.g. Bandi and Russell (2006).

## 2.6 Comparison with the ACD Models

In Engle and Russell (1998) and Tse and Yang (2012), the class of ACD models are used to specify price durations and estimate volatility. Let $\psi_{i+1}$ denote the conditional expectation of the price duration, i.e., $\psi_{i+1} = E(\tau_{i+1}|\mathcal{F}_i)$, an ACD model specifies the distribution for $\tau_{i+1}$ and the dynamics for $\psi_{i+1}$. If $\tau_{i+1}$ is exponentially distributed conditional on $\psi_{i+1}$, the spot variance is given by $V_i = \delta^2/\psi_{i+1}$, and the estimation for integrated variance naturally follows.

The SPD models and the ACD price duration models are different in three aspects. First, while the distribution of $\tau_{i+1}$ is derived from the passage theory of Brownian motion in the SPD models, the ACD models need to specify the distribution ad hoc and if it is not exponential, the estimation of variance becomes more complicated. Second, the conditional duration and hence spot variance is deterministic in the ACD models—it depends on only the past price durations. Allowing variance to be stochastic could generate more realistic dynamics. Third, the SPD models specify the dynamics of volatility directly rather than through conditional durations. This renders simple interpretation of the volatility dynamics and simple computation of the integration variance or other power-variations, which is important for jump identification. The *SPD1* model differs further from the ACD model in that it utilizes both price duration and trade durations.

We specify an ACD model with parametric diurnal pattern to compare the two classes of model directly. Let $V_i = V_i^* g_v(t_i)$, where $g_v(t) = a_1(t + a_2)^2 + 1$, the ACD model is specified as follows,

$$\begin{aligned} \tau_{i+1}^* &= \tau_{i+1} g_v(t_i) = \psi_{i+1}^* \varepsilon_{i+1} \\ \psi_{i+1}^* &= \omega + \alpha \tau_{i+1}^* + \beta \psi_i^*. \end{aligned} \tag{15}$$

Assuming $\varepsilon_{i+1}$ is exponentially distributed with mean equal to 1, the parameters and conditional duration is easily estimated from MLE, and $V_i$ is obtained from $V_i = g_v(t_i)\delta^2/\psi_{i+1}^*$.

# 3 Estimation Procedure

The inference for models with stochastic volatility or stochastic conditional duration is nontrivial since the evaluation of the likelihood involves integrating out the latent variables. We adopt the quasi-maximum likelihood estimation (QMLE) method that is popular in return-based SV models (see e.g. Harvey, Ruiz, and Shephard, 1994 and Ruiz, 1994). The QMLE method approximates the non-linear non-Gaussian state space model by a linear and Gaussian one, then use the Kalman filter to obtain the likelihood.

To apply QMLE to the baseline model *SPD0*, we start by taking the logarithm of equation (3):

$$\log \tau_{i+1} = 2 \log \delta - \log V_i + \log \eta_{i+1}, \tag{16}$$

then approximate $\log \eta_{i+1}$ by a normally distributed variable that has the same mean and variance[4]. In return-based SV models, this step involves taking the logarithm of squared returns, then approximate the logarithm of a chi-squared distribution by a normal distribution. The approximation error affects the efficiency of the QMLE estimator; if the true distribution is far from normal, parameter estimates yielded by QMLE could be highly inefficient. Figure (1) plots the true distributions versus their normal approximations for both logarithmic squared returns and logarithmic price durations. For squared returns, normal distribution is a poor approximation to the logarithm of a chi-squared distribution. This observation has spurred a large literature working on inference methods that evaluate the exact likelihood, including simulated maximum likelihood (Danielsson, 1994), Markov Chain Monte Carlo (Kim, Shephard, and Chib, 1998), and numerically accelerated importance sampling (Koopman, Lucas, and Scharth, 2015). However, these methods are computationally intensive, and it is difficult to estimate data in a long period of time given the sample size of high-frequency data[5]. In contrast, the logarithm of price duration is reasonably approximated by the normal distribution[6]. Hence, for our duration-based models, we gain computational speed from using QMLE without much loss of efficiency.

We estimate the *SPD1* model using QMLE as well. To linearize the average trade durations, we take the logarithm of equation (7) and approximate $\log e_{i+1}$ by a normal distribution. Since $e_{i+1}$ is distributed as $e_{i+1} \sim \Gamma(d_s M_{i+1}, 1/d_s M_{i+1})$, the mean and

[4]The pdf of $\log \eta_{i+1}$ is not known in closed-form so we compute the mean and variance from simulation.

[5]One exception is Stroud and Johannes (2014), who estimate a high-frequency stochastic volatility model using MCMC with efficient programming in C.

[6]This feature is also shared by range-base estimators, see Alizadeh, Brandt, and Diebold (2002).

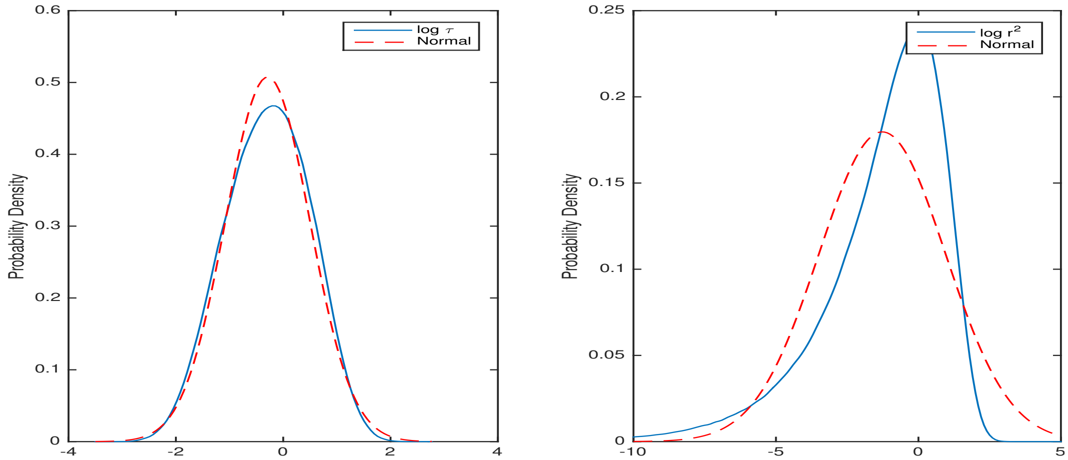


Figure 1: PDF of the true distributions versus their normal approximations. The left panel plots the distribution of logarithmic price durations versus a normal distribution with the same mean and variance. The right panel plots the distribution of squared returns versus its normal approximation.

variance of $\log e_{i+1}$ are given by $\psi(d_s M_{i+1}) - \log(d_s M_{i+1})$ and $\psi_1(d_s M_{i+1})$ respectively, where $\psi(x)$ denotes the digamma function and $\psi_1(x)$ denotes the trigamma function.

After considering the diurnal effect, the observation equation for *SPD1* model becomes

$$\begin{pmatrix} \log \tau_{i+1} \\ \log d^a_{i+1} \end{pmatrix} = \begin{pmatrix} 2\log\delta + E(\log\eta) - \log g_v(t_i) \\ E(\log e_{i+1}) - \log g_d(t_i) \end{pmatrix} + \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} x_i + w_{i+1}, \qquad (17)$$

where

$$w_{i+1} \sim N\left(0, \begin{pmatrix} \mathrm{Var}(\log\eta) & 0 \\ 0 & \psi_1(d_s M_{i+1}) \end{pmatrix}\right),$$

and $x_i = (\log V_i^*, \log\lambda_i^*)'$ and it follows the transition equation (10).

Equation (17) and (10) form the linear state space representation of *SPD1* model, so we can use the Kalman filter to get parameter estimates and smoothed volatility estimates. See de Jong (1989) for the filtering and smoothing procedure with time-varying coefficients.

An important issue in applications is to infer the stochastic volatility. We obtain volatility estimates from the smoothed latent variables. Let $x_{i|N}$ and $P_{i|N}$ denote the projection of $x_i$ on all observations and its mean squared error, i.e., $x_{i|N} = E(x_i|\mathcal{F}_N)$ and $P_{i|N} = \text{MSE}(x_{i|N})$, the smoothed estimate for $V_i$ and $\lambda_i$ are obtained from $\exp(x_{i|N} + P_{i|N}/2)$.

In real data we do not observe the overnight price process. Hence, we assume that each day "starts fresh": the latent OU process starts with its long-run mean and variance each day. This is applied both to the simulation in Section 4 and the empirical study in Section 5.

# 4 Simulation Studies

We use simulations to illustrate the properties of the price duration models, and compare them with popular estimators for integrated variance. We generate latent variance and conditional duration from equation (8), trade durations from equation (6), and the efficient logarithmic price from equation (1), discretized at each trade duration. Diurnal patterns are omitted for simplicity. Parameters are set as follows: $(\kappa^v, \kappa^\lambda) = (5e^{-4}, 6e^{-4})$, $(\mu^v, \mu^\lambda) = (-20, 0.48)$, $(\sigma_v, \sigma_\lambda, \rho) = (0.035, 0.017, -0.8)$, and $d_s = 0.1$. These parameters are taken from our empirical estimates in section 5; the time scale is in seconds. With these parameters, the mean spot variance, $V_0$, equals $3.8e^{-9}$, which corresponds to an annualized volatility of 15%, and trade happens every 1.8 seconds on average. For each simulation, we generate 20 trading days of data—each trading day has 23400 seconds. Given the stochastic nature of trade durations, the total number of trades in each simulation varies.

In section 4.1, we present the parameter and state estimates from the *SPD0* and

*SPD1* models in the absence of microstructure noise or jumps. In section 4.2, we add jumps and microstructure noise to the efficient price, and compare the integrated variance estimated from the price duration models with realized variance and bipower variation.

## 4.1 Simulations without Microstructure Noise or Jumps

The data simulated using trade durations resemble those we obtain from a high-frequency database—we have the trade (or quote) price and its time stamp. Price duration models require setting a price threshold $\delta$ to sample the data. To investigate the effect of different $\delta$, we choose $\delta = (0.0007, 0.001, 0.0014)$ such that the average price duration is calibrated to 3, 5 or 10 minutes. These sampling frequencies are empirically relevant in the high-frequency literature.

Since prices are not observed continuously, the observed price change may "overshoot" and exceed the pre-specified price threshold. Tse and Yang (2012) suggest using the average price changes conditional on the threshold being exceeded to measure volatility instead of the pre-specified threshold. The "actual price range" approach mitigates the overshoot problem, but it is not as robust to noise or jumps as the pre-specified threshold. Hence, we still use the pre-specified threshold in the simulation and empirical studies.

Table 1 reports the estimates and standard deviation of model parameters from 100 simulations. From Table 1, we note that the parameter estimates are fairly robust to the different choice of $\delta$.

Figure 2 and Figure 3 present a sample day of the true spot variance/conditional duration and their estimates. Several comments can be made regarding these figures. First, there are more sampling points when volatility is high, and less points when volatility is low. Hence, the ratio of noise variance over the volatility integrated over the sampling period is kept relatively flat, and this sampling scheme is more efficient than the fixed interval sampling. Second, the volatility estimated from price durations models are able

Table 1: Parameter Estimates

| | True | $\delta = 0.0007$ SPD1 | $\delta = 0.0007$ SPD0 | $\delta = 0.0010$ SPD1 | $\delta = 0.0010$ SPD0 | $\delta = 0.0014$ SPD1 | $\delta = 0.0014$ SPD0 |
|---|---|---|---|---|---|---|---|
| $\Psi_{11}$ | 0.0005 | 0.0006 (0.0001) | 0.0005 (0.0001) | 0.0006 (0.0001) | 0.0006 (0.0001) | 0.0006 (0.0001) | 0.0006 (0.0002) |
| $\Psi_{22}$ | 0.0006 | 0.0007 (0.0001) | | 0.0007 (0.0001) | | 0.0007 (0.0001) | |
| $\mu^v$ | -20.00 | -20.07 (0.09) | -20.07 (0.09) | -19.92 (0.09) | -19.91 (0.09) | -19.75 (0.09) | -19.74 (0.09) |
| $\mu^d$ | 0.4800 | 0.4966 (0.0455) | | 0.4749 (0.0452) | | 0.4413 (0.0442) | |
| $\sigma_v$ | 0.0350 | 0.0310 (0.0015) | 0.0291 (0.0015) | 0.0310 (0.0020) | 0.0299 (0.0022) | 0.0319 (0.0030) | 0.0312 (0.0039) |
| $\sigma_d$ | 0.0170 | 0.0182 (0.0009) | | 0.0170 (0.0009) | | 0.0164 (0.0012) | |
| $\rho$ | -0.8000 | -0.8448 (0.0214) | | -0.8277 (0.0264) | | -0.8168 (0.0339) | |
| $d_s$ | 0.1000 | 0.1533 (0.0058) | | 0.1418 (0.0077) | | 0.1320 (0.0131) | |

Note: this table reports the posterior mean and standard deviation (in parenthesis) of model parameters from 100 simulations.

to capture the main dynamics in spot volatility, and the *SPD1* model tend to trace the true volatility more closely than the *SPD0* or ACD model. Third, due to the large negative correlation between volatility and conditional duration ($\rho = -0.8$), there are more sampling points when the conditional duration is low. The estimates trace changes in the true conditional duration quite well.

## 4.2 Monte Carlo Comparisons

We estimate integrated variance using the price duration models and compare them with the realized volatility approach. Latent volatility and conditional duration are simulated as before, but we add jumps and market microstructure noise to simulate real world situations. The efficient logarithmic price follows the jump diffusion model in equation (14) with $(\lambda_J, \mu_J, \sigma_J) = (3.8e^{-5}, 0, 0.005)$. The jump parameters are calibrated such that the jump variation accounts for roughly 20% of the quadratic variation of $y$. The

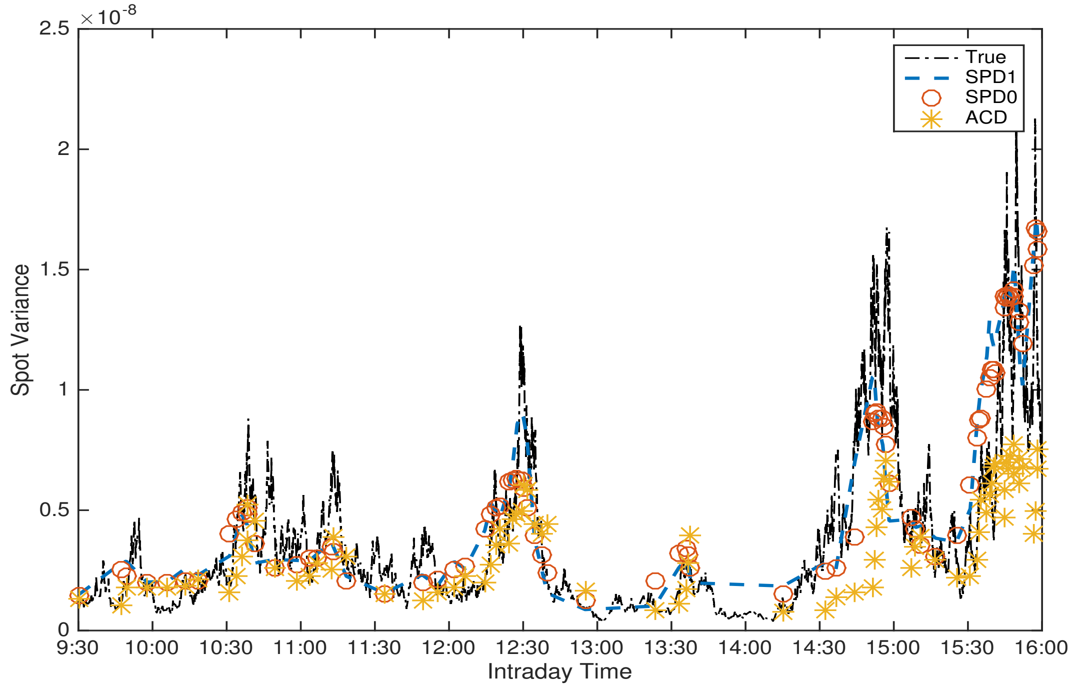


Figure 2: True Spot Variance and its Estimates.

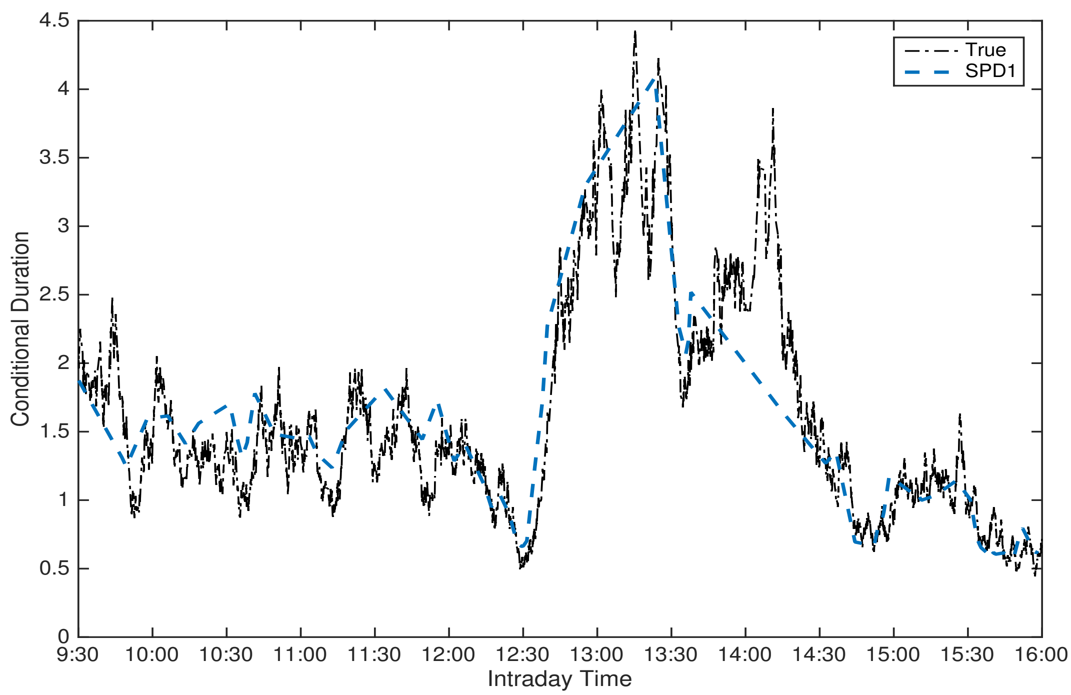


Figure 3: True Conditional Duration and its Estimates from *SPD1*.

observed price is contaminated by i.i.d. microstructure noises,

$$y_i^o = y_i + m_i, \tag{18}$$

where $m_i \sim N(0, \sigma_m^2)$. Three noise levels are considered, $\sigma_m^2 = (5e^{-9},\ 1e^{-8},\ 2e^{-8})$, which corresponds to small, moderate, and large noise situation. Define the Noise-to-Signal Ratio (NSR) by NSR $= \sigma_m/\sqrt{V_0}$, where $V_0$ is the long run mean of spot volatility with $V_0 = 3.9e-9$ as before, our choice of noise level corresponds to NSR $= (1.13, 1.6, 2.26)$.[7]

We conduct $G = 100$ simulations, each with $T = 20$ trading days of observation. For each noise level, we choose three price threshold proportional to the noise volatility: $\delta = (8\sigma_m,\ 10\sigma_m,\ 12\sigma_m)$. The resulting average sampling intervals are list in table. On each day $t$, the smoothed variance estimates from Kalman filter are plugged in equation (13) to obtain integrated variance from model *SPD0* and *SPD1*. Realized variance is computed by the sum of squared returns within the day, $RV_t = \sum_{j=1}^{1/\Delta} r_{t+j\Delta}^2$, and bipower variation $BV_t$ is computed from $BV_t = \frac{\pi}{2}\sum_{j=2}^{1/\Delta} \left|r_{t+(j-1)\Delta}\right| \left|r_{t+j\Delta}\right|$. For each simulation $g$, we set $\Delta$ to the average sampling interval of the duration-based approach. Hence, the duration-based approach has the same number of observations as the return-based approach.

Theoretically, $BV_t$ converges to the integrated variance over day $t$, while $RV_t$ converges to the quadratic variation, which is the sum of integrated variance and jump variation in the jump diffusion model. We use RMSE to compare the performance of the estimated IV, with RMSE $= \sqrt{\sum_{g=1}^{G}\sum_{t=1}^{T}(\widehat{IV}_t^{(g)} - IV_t^{(g)})^2/TG}$. We also use RMSE to compare the spot volatility estimates between the price duration model, computed by RMSE $= \sqrt{\sum_{i=1}^{N^{(g)}}(V_i^{(g)} - \hat{V}_i^{(g)})^2/N^{(g)}}$, where $V_i$ and $\hat{V}_i$ denote the true and estimated spot variance at time $t_i$, respectively.

[7] In realized volatility literature (see e.g. ), NSR is usually defined as $\sigma_m^2/\sqrt{IQ_t}$, where $IQ_t$ stands for integrated quarticity in day $t$. In our simulation $IQ_t$ varies with $t$.

Table 2: Monte Carlo Comparison with Jumps and Small Microstructure Noise

| | Bias | RMSE | Relative RMSE | Spot RMSE |
|---|---|---|---|---|
| | $\sigma_m^2$ = 5e-09, $\delta = 8\ \sigma_m$, $\bar{\tau}$ = 1.7 min | | | |
| SPD1 | $-1.46e-05$ | $2.61e-05$ | 1.251 | $9.10e-09$ |
| SPD0 | $-1.38e-05$ | $2.52e-05$ | 1.206 | $9.20e-09$ |
| ACD | $-2.59e-05$ | $3.78e-05$ | 1.813 | $1.32e-08$ |
| RV | $2.53e-05$ | $5.20e-05$ | 2.491 | |
| BV | $4.77e-06$ | $2.09e-05$ | 1.000 | |
| | $\sigma_m^2$ = 5e-09, $\delta = 10\ \sigma_m$, $\bar{\tau}$ = 2.6 min | | | |
| SPD1 | $-1.20e-05$ | $2.27e-05$ | 0.905 | $8.82e-09$ |
| SPD0 | $-1.12e-05$ | $2.19e-05$ | 0.871 | $8.94e-09$ |
| ACD | $-1.90e-05$ | $3.24e-05$ | 1.290 | $1.35e-08$ |
| RV | $2.57e-05$ | $5.40e-05$ | 2.150 | |
| BV | $5.16e-06$ | $2.51e-05$ | 1.000 | |
| | $\sigma_m^2$ = 5e-09, $\delta = 12\ \sigma_m$, $\bar{\tau}$ = 3.6 min | | | |
| SPD1 | $-9.73e-06$ | $2.03e-05$ | 0.738 | $8.69e-09$ |
| SPD0 | $-8.88e-06$ | $1.94e-05$ | 0.707 | $8.83e-09$ |
| ACD | $-1.34e-05$ | $2.97e-05$ | 1.081 | $1.42e-08$ |
| RV | $2.48e-05$ | $5.44e-05$ | 1.982 | |
| BV | $4.59e-06$ | $2.75e-05$ | 1.000 | |

Note: the first two columns of the this table reports Bias and RMSE of estimated integrated variance. The relative RMSE reports the RMSE of each model compared to Bipower Variation. The Spot RMSE reports the RMSE of spot variance from the three price duration model. These Monte Carlo results are from $G = 100$ simulations, each consists of $T = 20$ days of observations.

Table 3: Monte Carlo Comparison with Jumps and Moderate Microstructure Noise

| | Bias | RMSE | Relative RMSE | Spot RMSE |
|---|---|---|---|---|
| | | $\sigma_m^2$ = 1e-08, $\delta = 8\ \sigma_m$, $\bar{\tau} = 2.9$ min | | |
| SPD1 | $1.43e-06$ | $1.61e-05$ | 0.605 | $7.76e-09$ |
| SPD0 | $2.32e-06$ | $1.56e-05$ | 0.585 | $7.87e-09$ |
| ACD | $-2.59e-05$ | $3.78e-05$ | 1.419 | $1.32e-08$ |
| RV | $2.58e-05$ | $5.37e-05$ | 2.014 | |
| BV | $5.81e-06$ | $2.67e-05$ | 1.000 | |
| | | $\sigma_m^2$ = 1e-08, $\delta = 10\ \sigma_m$, $\bar{\tau} = 4.5$ min | | |
| SPD1 | $9.66e-07$ | $1.49e-05$ | 0.473 | $7.82e-09$ |
| SPD0 | $1.91e-06$ | $1.45e-05$ | 0.461 | $7.97e-09$ |
| ACD | $-1.90e-05$ | $3.24e-05$ | 1.027 | $1.35e-08$ |
| RV | $2.60e-05$ | $5.60e-05$ | 1.774 | |
| BV | $6.05e-06$ | $3.16e-05$ | 1.000 | |
| | | $\sigma_m^2$ = 1e-08, $\delta = 12\ \sigma_m$, $\bar{\tau} = 6.4$ min | | |
| SPD1 | $1.88e-06$ | $1.48e-05$ | 0.438 | $7.70e-09$ |
| SPD0 | $2.92e-06$ | $1.47e-05$ | 0.436 | $7.87e-09$ |
| ACD | $-1.34e-05$ | $2.97e-05$ | 0.880 | $1.42e-08$ |
| RV | $2.55e-05$ | $5.79e-05$ | 1.716 | |
| BV | $4.57e-06$ | $3.37e-05$ | 1.000 | |

Note: the first two columns of the this table reports Bias and RMSE of estimated integrated variance. The relative RMSE reports the RMSE of each model compared to Bipower Variation. The Spot RMSE reports the RMSE of spot variance from the three price duration model. These Monte Carlo results are from $G = 100$ simulations, each consists of $T = 20$ days of observations.

Table 4: Monte Carlo Comparison with Jumps and Large Microstructure Noise

| | Bias | RMSE | Relative RMSE | Spot RMSE |
|---|---|---|---|---|
| | $\sigma_m^2$ = 2e-08, $\delta = 8\ \sigma_m$, $\bar{\tau}$ = 4.9 min | | | |
| SPD1 | $1.82e-05$ | $2.27e-05$ | 0.743 | $7.06e-09$ |
| SPD0 | $1.92e-05$ | $2.32e-05$ | 0.761 | $7.20e-09$ |
| ACD | $-2.59e-05$ | $3.78e-05$ | 1.240 | $1.32e-08$ |
| RV | $2.71e-05$ | $5.72e-05$ | 1.874 | |
| BV | $6.86e-06$ | $3.05e-05$ | 1.000 | |
| | $\sigma_m^2$ = 2e-08, $\delta = 10\ \sigma_m$, $\bar{\tau}$ = 7.9 min | | | |
| SPD1 | $1.43e-05$ | $1.99e-05$ | 0.521 | $7.26e-09$ |
| SPD0 | $1.54e-05$ | $2.05e-05$ | 0.538 | $7.44e-09$ |
| ACD | $-1.90e-05$ | $3.24e-05$ | 0.852 | $1.35e-08$ |
| RV | $2.58e-05$ | $6.09e-05$ | 1.599 | |
| BV | $4.46e-06$ | $3.81e-05$ | 1.000 | |
| | $\sigma_m^2$ = 2e-08, $\delta = 12\ \sigma_m$, $\bar{\tau}$ = 11.6 min | | | |
| SPD1 | $1.22e-05$ | $1.93e-05$ | 0.464 | $7.56e-09$ |
| SPD0 | $1.36e-05$ | $2.02e-05$ | 0.486 | $7.76e-09$ |
| ACD | $-1.34e-05$ | $2.97e-05$ | 0.714 | $1.42e-08$ |
| RV | $2.52e-05$ | $6.58e-05$ | 1.584 | |
| BV | $2.17e-06$ | $4.16e-05$ | 1.000 | |

Note: the first two columns of the this table reports Bias and RMSE of estimated integrated variance. The relative RMSE reports the RMSE of each model compared to Bipower Variation. The Spot RMSE reports the RMSE of spot variance from the three price duration model. These Monte Carlo results are from $G = 100$ simulations, each consists of $T = 20$ days of observations.

Tables 2, 3, and 4 report the Monte Carlo results. For price duration models, the bias has two main components: a positive component introduced by the noise, and a negative component introduced by overshooting (actual price range exceeds the threshold). The negative bias from overshooting dominates in Table 2, where the average sampling intervals corresponding to $\delta = (8\sigma_m, 10\sigma_m, 12\sigma_m)$ are $\bar{\tau} = (1.7, 2.6, 3.6)$ minutes. As the sampling interval increases, the overshooting bias becomes negligible and the SPD models outperform BV. In table 3 and 4, the sampling intervals are larger, and the gains from using SPD models are evident. The ACD model also performs quite well considering it is misspecified. It outperforms BV when the sampling intervals are large.

In terms of the spot volatility estimation, The *SPD1* models has smaller RMSE than the *SPD0* or the ACD model. The gain slightly increases when NSR increases. This is as expected since *SPD0* does not utilize trade durations, and when NSR increases, prices are more contaminated while trade durations are not affected.

# 5 Empirical Results

## 5.1 Data

We apply our model to the milli-second time stamped IBM trade data in the US Equity Data provided by TickData. The sample period is from January 3, 2011 to December 30, 2011 (252 trading days). Following the cleaning procedure proposed by Barndorff-Nielsen, Hansen, Lunde, and Shephard (2009), we delete entries that have a correction indicator other than 0, or have a time stamp outside of the normal opening time, or have abnormal sales conditions. Entries from the first five minutes after opening are deleted to eliminate the price changes due to information accumulated overnight. Lastly, we treat entries within 0.1 second as one observation and use the median price to alleviate possible measurement error in the transaction time.

Table 5: Summary statistics for IBM in 01/03/2011-12/30/2011

| | Mean | Median | Standard Deviation |
|---|---|---|---|
| Price Duration | 314.9 | 122.8 | 589.4 |
| Number of Trades per price duration | 316.4 | 178.0 | 420.1 |
| Average Trade Duration | 0.849 | 0.721 | 0.575 |

Notes: The durations are reported in seconds.

The price threshold is set to 0.0013 (roughly a 0.13% change in price) so that the average price duration is roughly 5 minutes. This results in a total of 17944 sampling points. We then divide each price duration by the number of trades within that price duration to obtain average trade durations. Summary statistics for the observables is given in Table 5.

## 5.2 Estimation

Table 6 presents the parameter estimates from the SPD models. Each day is assumed to "start fresh": the latent OU process starts with its long-run mean and variance each day. In the *SPD0* and *SPD1* model, the long run mean of spot variance, $V_0$, equal to $3.1e-9$ and $3.5e-9$ respectively, while the half-life of the variance processes are about 29 minutes and 23 minutes. These parameters suggest that there is a fast decaying volatility dynamics at the intraday level. In the *SPD1* model, volatility and conditional durations are highly correlated, with $\rho = -0.8$. The diurnal patterns are pronounced in both volatility and conditional duration. Figure 4 plots the diurnal patterns estimated from the *SPD1* model. Volatility near the beginning of a trading day is almost 5 times as big as its minimum around noon; the conditional trading durations are less than of the conditional trading durations near the middle of the day.

Table 6: Parameter estimates for the *SPD1* and *SPD0* model

| | $\Psi_{11}$ | $\Psi_{22}$ | $\mu^v$ | $\mu^d$ | $\sigma_v$ | $\sigma_d$ | $\rho$ | $d_s$ | $a_1$ | $a_2$ | $a_3$ | $a_4$ |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| SPD1 | 4.94e-04 | 5.56e-04 | -20.07 | 0.486 | 0.035 | 0.017 | -0.807 | 0.065 | 2.28e-08 | -1.24e+04 | 1.42e-08 | -1.10e+04 |
| | (6.20e-05) | (4.47e-05) | (0.06) | (0.021) | (0.002) | (0.001) | (0.011) | (0.002) | (3.02e-09) | (2.87e+02) | (7.72e-10) | (1.26e+02) |
| SPD0 | 3.94e-04 | | -20.13 | | 0.029 | | | | 2.35e-08 | -1.31e+04 | | |
| | (5.44e-05) | | (0.05) | | (0.002) | | | | (2.06e-09) | (2.50e+02) | | |

Notes: the report reports the parameter estimation from *SPD1* and *SPD0* model using milli-second IBM data from 01/03/2011 to 12/30/2011.

Table 7: Parameter estimates for the ACD model

| | $\omega$ | $\alpha$ | $\beta$ | $a_1$ | $a_2$ |
|---|---|---|---|---|---|
| ACD | 1.04 | 0.164 | 0.834 | 2.02e-08 | -1.29e+04 |
| | (0.24) | (0.007) | (0.008) | (1.27e-09) | (1.96e+02) |

Notes: the report reports the parameter estimation from the ACD model using milli-second IBM data from 01/03/2011 to 12/30/2011.

We also estimate the ACD model and the parameters are presented in Table 7. Figure 5 plots the daily integrated variance estimated from *SPD0*, *SPD1,* ACD, and 5-minute BV. The integrated variance in the parametric models is obtained from the smoothed estimates of spot volatility, $\widehat{IV}_t = \sum_{i \in day(t)} V_i \tau_{i+1}$ and the number is reported in terms of annualized volatility. From Figure 5, the volatility estimates trace each other quite closely.

Figure 6 and 7 plots the spot volatility estimates in two sample months. The spot variance estimates from *SPD0* is quite close to *SPD1* and hence omitted. The spot variance from *SPD1* has more large spikes then the ACD model, especially in the high volatility period in August. In the relatively calm month of May, the *SPD1* model also produces larger variations in the volatility estimates. This suggests that stochastic models might be better at capturing the intraday volatility dynamics.

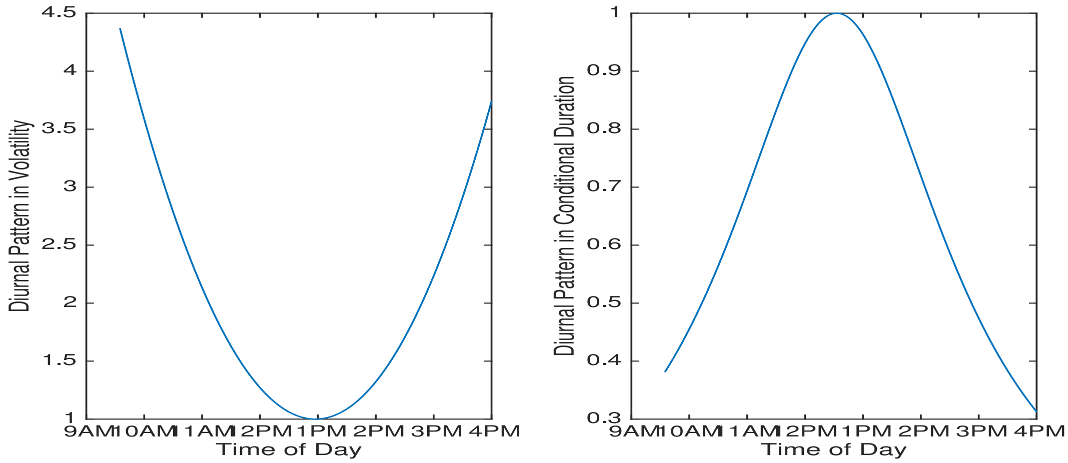


Figure 4: Diurnal patterns in volatility and conditional duration.

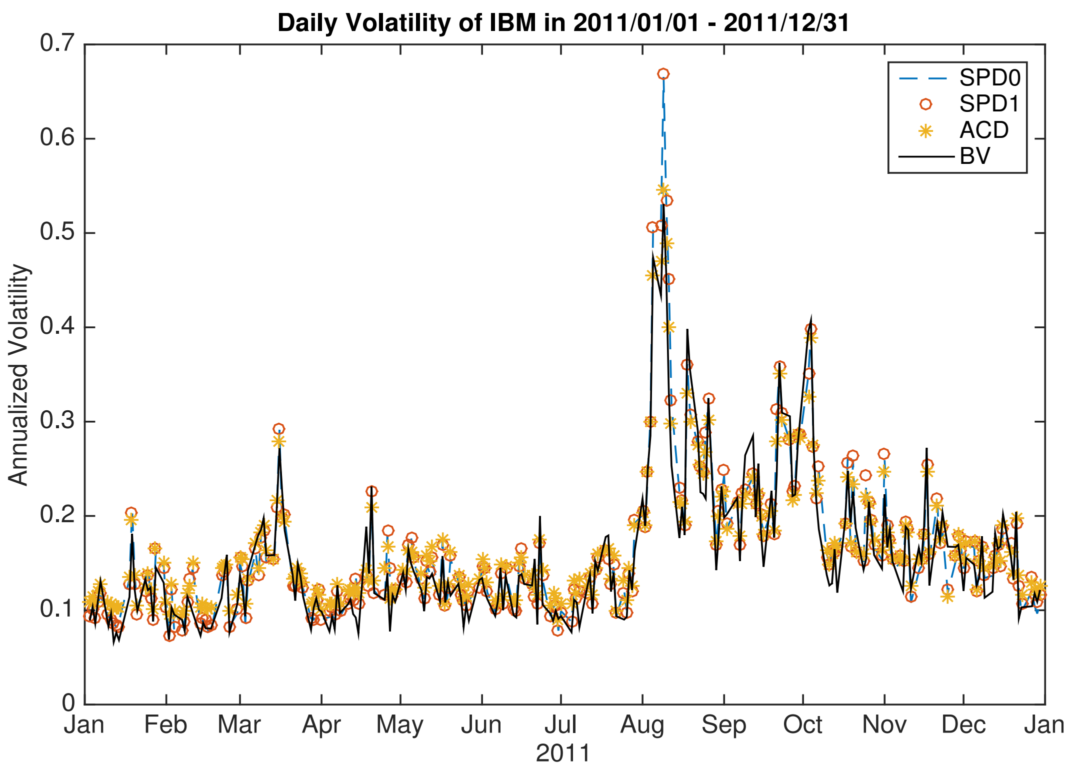


Figure 5: Daily Integrated Variance of IBM in 2011, 252 days.

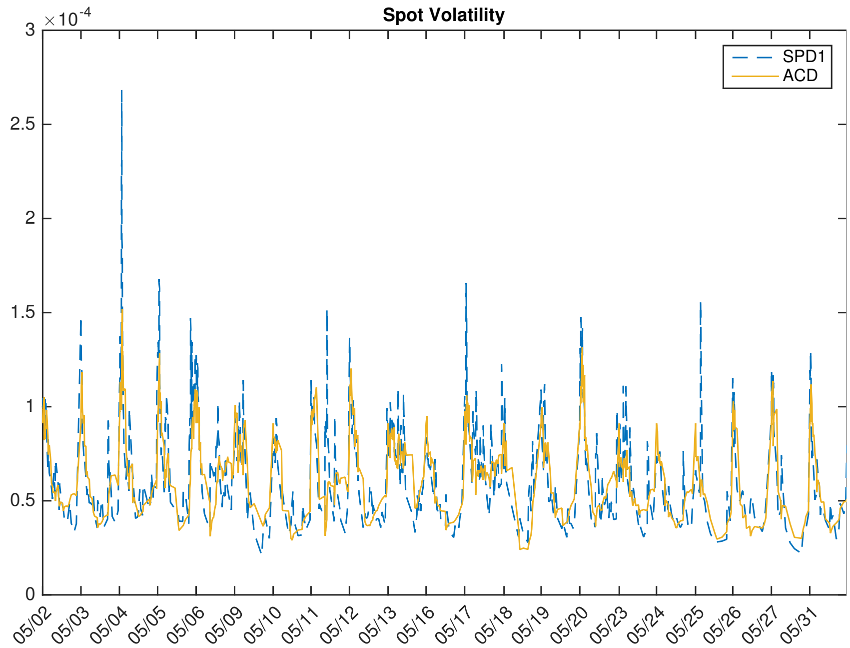


Figure 6: Estimated Spot Variance in May.

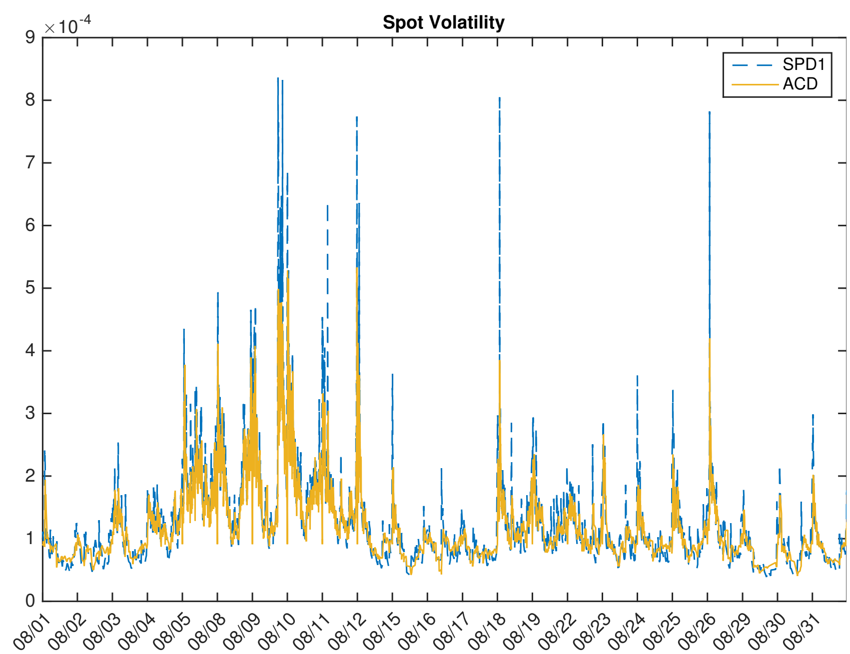


Figure 7: Estimated Spot Variance in August.

# 6 Conclusion

In this paper we present a new parametric model to estimate stochastic volatility based on price durations. This model has several advantages. First, price durations have some robustness to jumps and market microstructure noise, especially the noise from price discreteness. Second, we utilize the persistence of volatility to infer spot volatility and integrated variance over any period of time. Third, we assume that volatility is stochastic, and we obtain the distribution of price durations from the passage theory for Brownian motions. Last, we can conduct inference easily using QMLE without much loss of efficiency since the logarithmic price duration is better approximated by a normal distribution than the logarithmic squared returns.

We also extend the baseline model *SPD0* to the model *SPD1* to incorporate information from trading durations. Monte Carlo studies suggest that *SPD1* outperforms *SPD0* in estimating spot volatility, and both duration-based model performs better than realized volatility in estimating the integrated variance. Our empirical study on IBM suggests the presence of important intraday volatility dynamics.

# References


Alizadeh, S., Brandt, M. W., and Diebold, F. X. (2002). 'Range-Based Estimation of Stochastic Volatility Models', *The Journal of Finance*, 57(3): 1047–1091.

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., Bollerslev, T., and Diebold, F. X. (2010). 'Parametric and Nonparametric Volatility Measurement', in Y. Aït-Sahalia and L. P. Hansen (eds.), *Handbook* *of Financial Econometrics: Tools and Techniques*, vol. 1 of *Handbooks in Finance*, pp. 67 – 137. North-Holland, San Diego.

Andersen, T. G., Dobrev, D., and Schaumburg, E. (2009). 'Duration-Based Volatility Estimation', Global COE Hi-Stat Discussion Paper Series gd08-034, Institute of Economic Research, Hitotsubashi University.

Bandi, F. M., and Russell, J. R. (2006). 'Separating microstructure noise from volatility', *Journal of Financial Economics*, 79(3): 655–692.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2009). 'Realized kernels in practice: trades and quotes', *Econometrics Journal*, 12(3): C1–C32.

Bauwens, L., and Veredas, D. (2004). 'The stochastic conditional duration model: a latent variable model for the analysis of financial durations', *Journal of Econometrics*, 119(2): 381–412.

Cho, D. C., and Frees, E. W. (1988). 'Estimating the Volatility of Discrete Stock Prices', *Journal of Finance*, 43(2): 451–66.

Danielsson, J. (1994). 'Stochastic volatility in asset prices estimation with simulated maximum likelihood', *Journal of Econometrics*, 64(1-2): 375–400.

de Jong, P. (1989). 'Smoothing and Interpolation with the State-Space Model', *Journal of the American Statistical Association*, 84(408): pp. 1085–1088.

Easley, D., and O'Hara, M. (1987). 'Price, trade size, and information in securities markets', *Journal of Financial Economics*, 19(1): 69–90.

Engle, R. F., and Russell, J. R. (1998). 'Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data', *Econometrica*, 66(5): 1127–1162.

Hansen, P. R., and Lunde, A. (2005). 'A forecast comparison of volatility models: does anything beat a GARCH(1,1)?', *Journal of Applied Econometrics*, 20(7): 873–889.

Harvey, A., Ruiz, E., and Shephard, N. (1994). 'Multivariate Stochastic Variance Models', *Review of Economic Studies*, 61(2): 247–64.

Kim, S., Shephard, N., and Chib, S. (1998). 'Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models', *Review of Economic Studies*, 65(3): 361–93.

Koopman, S. J., Lucas, A., and Scharth, M. (2015). 'Numerically Accelerated Importance Sampling for Nonlinear Non-Gaussian State-Space Models', *Journal of Business & Economic Statistics*, 33(1): 114–127.

Pelletier, D., and Zheng, H. (2012). 'Joint Modeling of High-Frequency Price and Duration Data', Discussion paper, North Carolina State University.

Ruiz, E. (1994). 'Quasi-maximum likelihood estimation of stochastic volatility models', *Journal of Econometrics*, 63(1): 289 – 306.

Stroud, J. R., and Johannes, M. S. (2014). 'Bayesian Modeling and Forecasting of 24-Hour High-Frequency Volatility', *Journal of the American Statistical Association*, 109(508): 1368–1384.

Tse, Y.-K., and Dong, Y. (2014). 'Intraday periodicity adjustments of transaction duration and their effects on high-frequency volatility estimation', *Journal of Empirical Finance*, 28: 352 – 361.

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.

Wei, W., and Pelletier, D. (2013). ‘A Jump Diffusion Model for Volatility and Duration’, Discussion paper, North Carolina State University.

Zu, Y., and Boswijk, H. P. (2014). ‘Estimating spot volatility with high-frequency financial data’, *Journal of Econometrics*, 181(2): 117 – 135.