UNIVERSITY OF CAMBRIDGE
Faculty of Economics

Cambridge Working Paper Economics

Cambridge Working Paper Economics: 1315

MODELING DYNAMIC DIURNAL PATTERNS IN HIGH-FREQUENCY FINANCIAL DATA

Ryoko Ito

Updated 24 May 2016

We introduce the spline-DCS model with a dynamic cubic spline as a way of capturing periodic behavior in financial data that evolves over time. Our empirical application provides evidence for changing diurnal patterns in the high-frequency financial data we study. We illustrate that this generalization can lead to an improvement in the quality of the fit of the model to the empirical distribution of data, especially in the tail region, for an extended out-of-sample period. Moreover, it can lead to a substantial improvement in predicting intra-day volume proportions, which is useful for Volume-Weighted Average Price strategies. Our novel approach gives new insights into regular trading behavior and how it responds to changing market conditions.

Modeling Dynamic Diurnal Patterns in High-Frequency Financial Data

Ryoko Ito[1]

Cambridge Working Papers in Economics CWPE1315
This version: May 24, 2016

Abstract

We introduce the spline-DCS model with a dynamic cubic spline as a way of capturing periodic behavior in financial data that evolves over time. Our empirical application provides evidence for changing diurnal patterns in the high-frequency financial data we study. We illustrate that this generalization can lead to an improvement in the quality of the fit of the model to the empirical distribution of data, especially in the tail region, for an extended out-of-sample period. Moreover, it can lead to a substantial improvement in predicting intra-day volume proportions, which is useful for Volume-Weighted Average Price strategies. Our novel approach gives new insights into regular trading behavior and how it responds to changing market conditions.

Keywords: order slicing; price impact; spline; volume prediction; score; seasonality

JEL classification: C22, C51, C53, C58, G01, G12

1. Introduction

Intra-day periodicity caused by periodic trading patterns is a stylized feature of high-frequency financial data. Popular methods for capturing this effect include the use of the Fourier series or a deterministic spline and computing a sample moment for each intra-day bin.[2] It is commonly assumed in the literature that the shape of periodicity is the same every day. It is also common to estimate the intra-day periodic component first and *diurnally adjust* data before estimating other non-periodic components. But such a two-step procedure can render the asymptotic properties of statistical tests invalid.

[1]Nuffield College and the Department of Economics, Oxford University. Email: ryoko.ito@economics.ox.ac.uk. The author thanks Andrew Harvey for his helpful comments on this research and guiding her academic progress. The author also thanks Philipp Andres, Michele Caivano, Adam Clements, Oliver Linton, Donald Robertson, Stephen Thiele, and the participants of the Score Workshop in 2013 at Tinbergen Institute, especially Siem Jan Koopman and Andre Lucas, for providing thoughtful comments on this paper. Finally, the author thanks the Keynes Fund and the Stevenson Fund (of the Faculty of Economics, Cambridge University) for funding travels to present this paper, and the International Monetary Fund, the Cambridge Trust, and the Royal Economic Society for funding her PhD.

[2]See, for instance, Andersen and Bollerslev (1998), Engle and Russell (1998), Zhang et al. (2001), Campbell and Diebold (2005), Engle and Rangel (2008), Brownlees et al. (2011), and Engle and Sokalska (2012).

In this paper, we extend the spline-DCS model of Ito (2016) to capture periodic behavior that evolves over time, and introduce it as a new tool for studying and understanding intra-day seasonality in finance. The periodic component is estimated easily and simultaneously with all of the other components of the model by the method of maximum likelihood (ML). We apply the model to the trade volume of the IBM stock traded on the New York Stock Exchange (NYSE). We find statistically significant evidence for changing diurnal patterns in the data we study, and illustrate the empirical merit of this generalization. We find that allowing the pattern of periodicity to change can improve the quality of the fit of the model to the empirical distribution of the data, especially in the tail region, for an extended out-of-sample period (up to 100 days). Moreover, using the slicing loss function of Brownlees et al. (2011), we show that our model can significantly improve the prediction of intra-day volume proportions, which is a key ingredient in high-frequency trading algorithms that aim to minimize the price impact[3] of a given transaction and to achieve the Volume-Weighted Average Price (VWAP) benchmark.[4]

Our model successfully captures diurnal U-shaped patterns in our data. But how does this pattern change over time? How does it change with macroeconomic conditions? We find that the amplitude (or the height) of the U-shape increases in the latter half of 2007, which roughly coincides with the onset of the 2007-2008 financial crisis. This may be a reflection of an increase in the amount of important news transmitted over night, considering that many important policy and financial decisions were made at night during the crisis. The profitability of the firm was also at risk and its revenue declined during the U.S. economic downturn in 2008 and 2009. We find that the amplitude of the U-shape reverts to a level comparable to the pre-crisis period by the end of 2013.

Does the U-shape become flatter over time? This is a reasonable guess since more trades in recent years are executed automatically and continuously throughout the day in high frequency by computers with algorithmic trading schemes. This means that traders don't "take lunch" like they used to. However, contrary to this plausible hypothesis, we find that the U-shape become more exacerbated over time. We find that transactions at the beginning and the end of the day have grown over the past decade to account for a larger share of total daily volume. Moreover, we find that the shape of periodicity does not change much around midday but evolves the most in the morning and the afternoon trading hours. This phenomenon can be explained by the fact that high-frequency traders typically try to minimize price impact by slicing a given order into small transaction sizes and executing them at different times of the day. An effective way of placing transactions to achieve this objective is to mimic the pattern of market activity and disguise themselves in the market, since a given transaction is less likely to move market price when the market

[3]The impact of the size of a given transaction on the execution price. It is also called the market impact.

[4]It is widely used by traders as a guarantee to clients that their orders will be executed at that target.

is attracting a high volume of transactions. The rise of algorithmic traders that exploit intra-day patterns of the market in this way can exacerbate the periodic behavior.

The structure of this paper is as follows. Section 2 defines the dynamic cubic spline in the spline-DCS framework. Section 3 describes our trade volume data. The in-sample and out-of-sample estimation results are reported in Sections 4 and 5. Section 6 concludes by laying out possible extensions for further research. In particular, we note that our model can be generalized further to introduce the version with a weekly spline as a new way of capturing the day-of-the-week effect in financial data.

2. Spline-DCS with dynamic spline

2.1. Notations

We adhere to the notations of Ito (2016) throughout this paper. We use the time subscript $\cdot_{t,\tau}$ to denote the τ-th intra-day bin on the t-th trading day for $\tau = 1, \ldots, I$ and $t = 1, \ldots, T$, where $I, T \in \mathbb{N}_{>0}$. We set $\cdot_{t+1,0} = \cdot_{t,I}$ for all t. We use the set notation, $\Psi_{T,I} = \{(t,\tau) \in \{1,2,\ldots,T\} \times \{1,2,\ldots,I\}\}$. $\mathcal{F}_{t,\tau}$ denotes the set of information available at time $(t,\tau) \in \Psi_{T,I}$. Any $\mathcal{F}_{1,1}$-measurable random variables are almost surely constant. We simply write $\cdot_{t,\tau}$ instead of $\cdot_{t,\tau|t,\tau-1}$ even when the variable is conditional on $\mathcal{F}_{t,\tau-1}$. We denote the estimated quantities by $\widehat{\cdot}$ and forecast quantities by $\widetilde{\cdot}$.

2.2. Spline-DCS

The version of the spline-DCS we study in this paper is

$$y_{t,\tau} = \varepsilon_{t,\tau} \exp(\lambda_{t,\tau}), \quad \lambda_{t,\tau} = \omega + \mu_{t,\tau} + \eta_{t,\tau} + s_{t,\tau}, \quad \varepsilon_{t,\tau} \sim \text{i.i.d. } F^*(\cdot;\theta^*)$$

for $(t,\tau) \in \Psi_{T,I}$. $F^*(\cdot;\theta^*)$ denotes the cumulative distribution function (c.d.f.) of $\varepsilon_{t,\tau}$ with the vector of distribution parameters denoted by θ^*.[5] $u_{t,\tau}$ is the score of the distribution, F^*. The non-periodic non-stationary component is $\mu_{t,\tau} = \beta_{t,\tau-1} + \mu_{t,\tau-1} + \kappa_\mu u_{t,\tau-1}$. This has the dynamics of an integrated random walk with time-varying drift, where the drift is given by $\beta_{t,\tau} = \beta_{t,\tau-1} + \kappa_\beta u_{t,\tau-1}$. The initial value, $\beta_{1,1}$, of $\beta_{t,\tau}$ is treated as unknown and estimated with other parameters of the model. We find that $\beta_{1,1}$ is insignificant in the estimation section (Section 4). The scale of trade volume should increase in response to positive news, giving $\kappa_\beta, \kappa_\mu > 0$. The drift term disappears if $\kappa_\beta = \beta_{1,1} = 0$. We set $\mu_{1,1} = 0$ so that $\omega \in \mathbb{R}$ is identified. The non-periodic stationary component is $\eta_{t,\tau} = \eta^{(1)}_{t,\tau} + \eta^{(2)}_{t,\tau}$ with

$$\eta^{(1)}_{t,\tau} = \phi^{(1)}_1 \eta^{(1)}_{t,\tau-1} + \phi^{(1)}_2 \eta^{(1)}_{t,\tau-2} + \phi^{(1)}_3 \eta^{(1)}_{t,\tau-3} + \kappa^{(1)}_\eta u_{t,\tau-1} + \kappa^{(1)}_{\eta,\mathrm{a}} \mathrm{sign}(-r_{t,\tau-1})(u_{t,\tau-1} + \nu\xi),$$
$$\eta^{(2)}_{t,\tau} = \phi^{(2)}_1 \eta^{(2)}_{t,\tau-1} + \kappa^{(2)}_\eta u_{t,\tau-1} + \kappa^{(2)}_{\eta,\mathrm{a}} \mathrm{sign}(-r_{t,\tau-1})(u_{t,\tau-1} + \nu\xi).$$

[5] We do not include the probability mass, $p = \mathbb{P}(y_{t,\tau} = 0)$, of zero-valued observations in the model since the number of zero-valued observations in our data is negligible. (See Section 3.)

In the estimation section, we find that the inclusion of the drift term in $\mu_{t,\tau}$ and the third lag in $\eta^{(1)}_{t,\tau}$ can capture the dynamics of the data well.

The role of each component is such that $\mu_{t,\tau}$ should be less sensitive to changes in $u_{t,\tau-1}$ than $\eta^{(1)}_{t,\tau}$, which should be, in turn, less sensitive than $\eta^{(2)}_{t,\tau}$. This gives $|\kappa_\mu| < |\kappa^{(1)}_\eta| < |\kappa^{(2)}_\eta|$. $r_{t,\tau}$ denotes the equity returns of IBM, which is computed as 100 times the first difference in the logarithms of prices. The asymmetry coefficients, $\kappa^{(1)}_{\eta,a}$ and $\kappa^{(2)}_{\eta,a}$, can be either positive or negative. We set $\eta^{(j)}_{1,1} = \mathbb{E}[\eta^{(j)}_{t,\tau}] = 0$ for $j = 1, 2$.[6]

2.2.1. Dynamic spline

The periodic component, $s_{t,\tau}$, captures the pattern of intra-day periodicity. Recall from the exposition by Ito (2016) that the static daily spline ($s_{t,\tau} = s_\tau$) is

$$s_\tau = \sum_{j=1}^{k} \mathbb{1}_{\{\tau\in[\tau_{j-1},\tau_j]\}}\, \boldsymbol{z}_j(\tau)\cdot\boldsymbol{\gamma}, \qquad \tau = 1,\dots,I, \tag{1}$$

where $\boldsymbol{z}_j : [\tau_{j-1},\tau_j]^{k+1} \to \mathbb{R}^{k+1}$ for $j = 1,\dots k$ is a $(k+1)$-dimensional vector of deterministic functions that conveys all information about the *polynomial order*, *continuity*, and *zero-sum conditions* of the spline. $\tau_0 < \cdots < \tau_k$ are the coordinates of the knots along the time axis with $\tau_0 = 1$, $\tau_k = I$, and $\tau_j \in \{2,\dots,I-1\}$ for $j = 1,\dots,k-1$. The y-coordinates (height) of the knots are denoted by $\boldsymbol{\gamma} = (\gamma_0,\dots,\gamma_k)^\top$. Then we have $k+1$ knots for some $k \in \mathbb{N}_{>0}$ such that $k < I$. The derivation of $\boldsymbol{z}_j(\tau)$ is the same as Ito (2016). We capture the overnight effect by allowing for $\gamma_0 \neq \gamma_k$.

The static spline in (1) becomes dynamic by letting the height of the knots, $\boldsymbol{\gamma}$, to change over time. We do this by re-defining (1) as

$$s_{t,\tau} = \sum_{j=1}^{k} \mathbb{1}_{\{\tau\in[\tau_{j-1},\tau_j]\}}\, \boldsymbol{z}_j(\tau)\cdot\boldsymbol{\gamma}_{t,\tau}, \qquad \boldsymbol{\gamma}_{t,\tau} = \boldsymbol{\gamma}_{t,\tau-1} + \boldsymbol{\kappa}^*\cdot u_{t,\tau-1}, \tag{2}$$

where $\boldsymbol{\kappa}^* = (\kappa^*_0,\dots,\kappa^*_k)^\top$ is a vector of parameters. $s_{t,\tau}$ needs to sum to zero over one complete period for the parameters to be identified. That is, it must satisfy $\sum_{\tau=1}^{I} s_{t,\tau} = \boldsymbol{w}_*\cdot\boldsymbol{\gamma}_{t,\tau} = 0$ for $t = 1,\dots,T$, where $\boldsymbol{w}_*$ is defined by Ito (2016). The construction of $\boldsymbol{z}_j(\tau)$ ensures that this constraint holds, but we also need to set $\boldsymbol{w}_*\cdot\boldsymbol{\gamma}_{1,1} = 0$ and $\boldsymbol{w}_*\cdot\boldsymbol{\kappa}^* = 0$. We can impose these conditions on $\boldsymbol{\gamma}_{1,1}$ and $\boldsymbol{\kappa}^*$ by setting $\gamma_{k;1,1} = -\sum_{i=0}^{k-1} w_{*i}\gamma_{i;1,1}/w_{*k}$ and $\kappa^*_k = -\sum_{i=0}^{k-1} w_{*i}\kappa^*_i/w_{*k}$, where $\gamma_{i;1,1}$ and w_{*i} denote the ith element of $\boldsymbol{\gamma}_{1,1}$ and $\boldsymbol{w}_*$, respectively.[7]

The location of knots, $\tau_0,\dots,\tau_k$ and the size of k depend on the empirical shape of diurnal patterns and the number of intra-day observations. The first knot and the last knot are placed at 9.30am and 4pm. Following the rule of thumb discussed by Ito (2016), we found placing $\gamma_1\dots,\gamma_{k-1}$ at 10am, 11am, 12.30pm, 2pm, and 3.30pm works well.

[6] This assumes that $y_{t,\tau}$ and $r_{t,\tau}$ are independent, and that $r_{t,\tau}$ is symmetric around zero.

[7] Harvey and Koopman (1993) use a set of contemporaneous Gaussian disturbances to drive the dynamics of the spline instead of the lagged score. This means that these identification restrictions on $\boldsymbol{\gamma}_{1,1}$ and $\boldsymbol{\kappa}^*$ are different from the formulation given by the authors.

Window	In-sample	Out-of-sample
1	Mon 7 Jan 2002 - Fri 1 Apr 2005	Mon 4 Apr 2005 - Fri 19 Aug 2005
2	Mon 8 Jan 2007 - Fri 31 Dec 2010	Mon 3 Jan 2011 - Fri 20 May 2011
3	Mon 9 Jan 2012 - Tue 31 Dec 2013	Thu 2 Jan 2014 - Fri 21 Mar 2014

Table 1: The list of sampling windows. The sample includes weekdays only. Each day includes the hours between 9.30am and 4pm in the NY local time.

	In-sample		Out-of-sample	
Window	Work days	Holidays	Work days	Holidays
1	813	32	98	2
2	1001	38	98	2
3	498	18	55	2

Table 2: The number of days included in each sampling window.

Window	Mean	S.D.	Skew	% zeros
1	168,026	152,512	5.4	0.0%
2	179,544	170,656	6.2	0.0%
3	88,908	101,612	10.1	0.0%

Table 3: Sample statistics for each in-sample window. The rightmost column is the percentage of observations that are zero-valued.

Thus, we have $k+1=7$ knots in total. The shape of the spline up to 12.30pm captures busy trading hours in the morning, between 12.30pm and 2.30pm captures quiet lunch hours, and after 2.30pm captures an acceleration in trading activities before the market closes. There is little to no improvement in the quality of the fit of the model to the data when the number of knots per day increases from this specification.

2.3. Maximum likelihood estimation

All of the parameters of the model are estimated by ML. Ito (2016) outlines the estimation procedure. We set $\beta_{t,\tau}=\beta_{t,\tau-1}$, $\mu_{t,\tau}=\mu_{t,\tau-1}$, $\eta^{(1)}_{t,\tau}=\eta^{(1)}_{t,\tau-1}$, $\eta^{(2)}_{t,\tau}=\eta^{(2)}_{t,\tau-1}$, and $\boldsymbol{\gamma}_{t,\tau}=\boldsymbol{\gamma}_{t,\tau-1}$ for all $\tau=1,\ldots,I$ if the t-th sampling day is a public holiday. The joint log-likelihood function is computed only for the days the market was open. The computing time taken for the maximum likelihood estimation procedure to converge was about 5 to 10 minutes.

As regard computing the analytical standard errors of the ML estimator, in addition to the recursive equations derived by Ito (2016) for computing the Fisher information matrix, the following recursions arising from the dynamic spline are needed:

$$\frac{\partial s_{t,\tau}}{\partial \boldsymbol{\vartheta}_i}=\sum_{j=1}^{k}\mathbb{1}_{\{\tau\in[\tau_{j-1},\tau_j]\}}\boldsymbol{z}_{j,l}(\tau)\frac{\partial \boldsymbol{\gamma}_{t,\tau}}{\partial \boldsymbol{\vartheta}_i},$$

$$\frac{\partial \gamma_{l;t,\tau}}{\partial \boldsymbol{\vartheta}_i}=\frac{\partial \gamma_{l;t,\tau-1}}{\partial \boldsymbol{\vartheta}_i}+u_{t,\tau-1}\mathbb{1}_{\{\boldsymbol{\vartheta}_i=\kappa^*_l\}}+\boldsymbol{\kappa}^*\frac{\partial u_{t,\tau-1}}{\partial \boldsymbol{\vartheta}_i}+\mathbb{1}_{\{\boldsymbol{\vartheta}_i=\gamma_{l;1,1},\,(t,\tau)=(1,1)\}},\quad l=1,\ldots,k.$$

Here, $\boldsymbol{\vartheta}$ is the vector of the constant parameters of the model and $\boldsymbol{\vartheta}_i$ denotes the i-th element of $\boldsymbol{\vartheta}$.

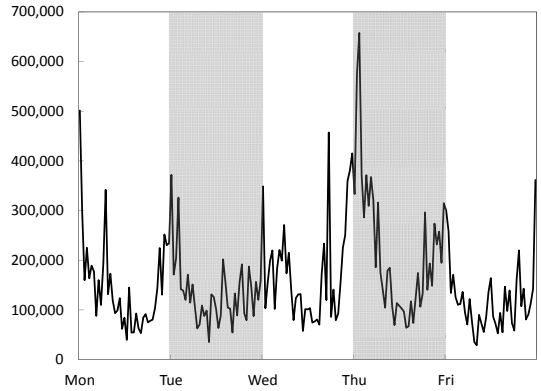

Figure 1: IBM10m between Monday 7 January 2002 and Friday 11 January 2002. Each day covers between 9.30am and 4pm in the NY local time.

3. Data characteristics

We apply the model to trade volume of the IBM stock on the New York Stock Exchange (NYSE). Trade volume is as defined by the number of traded shares. This data was obtained from www.kibot.com. Each sampling day covers the hours the market is open, which is between 9.30am and 4pm in the New York (NY) local time. The sampling frequency is 10 minutes. There are 39 observations per trading day. For convenience, we call our time series IBM10m. We consider three discrete sampling windows given in Table 1. The first sampling window is before the 2007-2008 financial crisis. The second window includes the onset of the crisis, and the third window is roughly in the post-crisis period and when the U.S. economy showed signs of recovery. Table 2 lists for each window the number of weekdays when NYSE was open, as well as the number of public holidays when the market was closed. The number of zero-valued observations are negligible in IBM10m. (See Table 3.)

Figure 1 gives a snapshot of IBM10m between Monday 7 January 2002 and Friday 11 January 2002. The volume fluctuates and clearly exhibits diurnal U-shaped patterns. The left column of Figure 2 shows that intra-day periodic patterns are also reflected in the sample autocorrelation of volume, which peaks at the multiples of the 39th lag. Finally, the skewness statistics in Table 3 and the middle and the right columns of Figure 2 show that our series are right-skewed and heavy-tailed.

4. Estimation results

Table 4 shows the estimated coefficients when $s_{t,\tau}$ is the dynamic spline. It tabulates the results for the full model defined in Section 2, as well as the final model that imposes coefficient restrictions. In this section, first, we discuss the significance of the dynamic parameters, $\boldsymbol{\kappa}^*$, of the spline. We then discuss the model selection procedure, the significance of other parameters, and the quality of the fit of the model to the data. The results reported here are not sensitive to the choice of initial parameter values that initializes

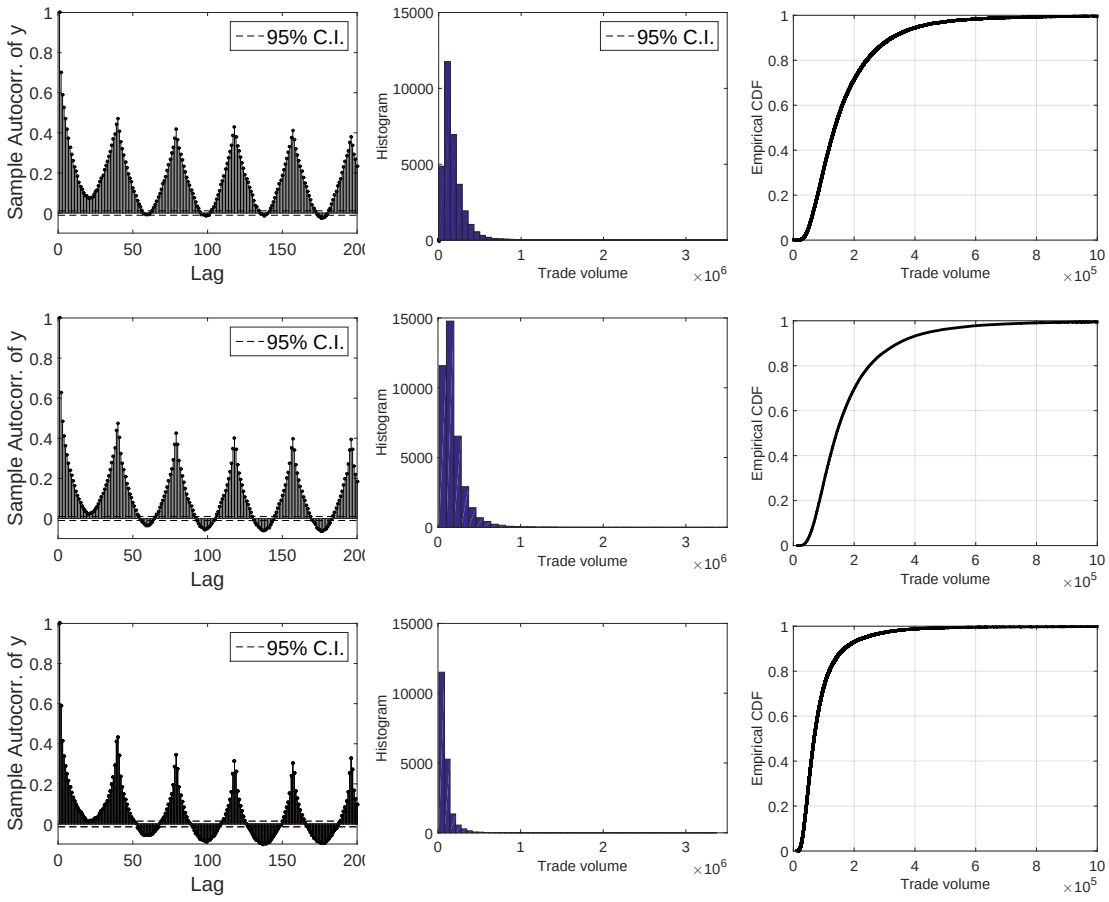

Figure 2: The autocorrelation function (left column), the histogram (middle column), and the empirical c.d.f. (right column) of IBM10m in the in-sample window 1 (top row), 2 (middle row), and 3 (bottom row).

the ML optimization.

Window	1 (2002-2005)				2 (2007-2010)				3 (2012-2013)			
Model	Full		Final		Full		Final		Full		Final	
Coef	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.
$\beta_{1,1}$	8.6E-05	2.7E-05	—	—	-3.1E-05	1.2E-05	—	—	-1.1E-05	2.8E-05	—	—
κ_β	2.4E-07	6.0E-08	2.6E-08	7.8E-09	2.2E-07	3.1E-08	1.9E-07	3.0E-08	6.1E-08	9.7E-08	1.6E-07	8.0E-08
κ_μ	1.1E-05	1.8E-04	3.5E-04	1.1E-04	1.2E-03	1.5E-04	1.3E-03	1.7E-04	-5.9E-04	1.3E-04	1.5E-03	3.7E-04
$\phi_1^{(1)}$	0.527	0.042	0.603	0.040	0.577	0.036	0.579	0.037	0.649	0.046	0.553	0.040
$\phi_2^{(1)}$	0.332	0.043	0.381	0.040	0.288	0.036	0.287	0.036	0.314	0.046	0.241	0.040
$\phi_3^{(1)}$	0.124	0.045	—	—	0.118	0.038	0.116	0.038	0.017	0.054	0.182	0.040
$\kappa_\eta^{(1)}$	0.020	0.001	0.018	0.001	0.019	0.001	0.019	0.001	0.024	0.002	0.024	0.001
$\kappa_{a,\eta}^{(1)}$	1.0E-03	2.7E-04	9.6E-04	1.7E-04	-1.0E-04	1.9E-04	—	—	2.7E-04	2.3E-04	—	—
$\phi_1^{(2)}$	0.752	0.015	0.716	0.015	0.701	0.019	0.698	0.019	0.680	0.040	0.739	0.032
$\kappa_\eta^{(2)}$	0.032	0.001	0.032	0.001	0.024	0.001	0.024	0.001	0.031	0.002	0.032	0.001
$\kappa_{a,\eta}^{(2)}$	-1.1E-04	3.2E-04	—	—	4.2E-04	2.2E-04	3.4E-04	1.4E-04	2.0E-04	2.7E-04	4.5E-04	1.8E-04
$\gamma_{0;1,1}$	0.975	0.030	0.972	0.022	0.985	0.045	0.878	0.011	1.017	0.032	0.993	0.038
$\gamma_{1;1,1}$	0.554	0.019	0.661	0.023	0.525	0.029	0.457	0.007	0.535	0.020	0.498	0.021
$\gamma_{2;1,1}$	0.133	0.018	0.163	0.016	0.088	0.014	0.060	0.007	0.162	0.022	0.120	0.013
$\gamma_{3;1,1}$	-0.399	0.016	-0.395	0.012	-0.407	0.008	-0.404	0.006	-0.355	0.022	-0.362	0.021
$\gamma_{4;1,1}$	-0.301	0.016	-0.330	0.013	-0.284	0.014	-0.254	0.006	-0.315	0.021	-0.286	0.015
$\gamma_{5;1,1}$	0.285	0.017	0.230	0.017	0.322	0.017	0.359	0.007	0.209	0.026	0.247	0.017
κ_0^*	1.2E-04	6.7E-05	9.0E-05	3.3E-05	-8.2E-04	1.1E-04	-8.0E-04	1.0E-04	1.3E-04	1.2E-04	-5.6E-04	1.6E-04
κ_1^*	-5.1E-06	4.2E-05	2.0E-04	3.9E-05	-5.2E-04	7.0E-05	-5.1E-04	6.9E-05	-5.0E-05	7.5E-05	-2.8E-04	1.0E-04
κ_2^*	4.6E-05	3.8E-05	8.7E-05	2.4E-05	-2.2E-04	5.1E-05	-2.1E-04	4.9E-05	-1.4E-04	7.9E-05	-6.8E-05	8.4E-05
κ_3^*	-1.7E-05	3.4E-05	-5.3E-06	1.7E-05	2.7E-05	4.4E-05	4.5E-05	4.2E-05	-1.6E-04	8.1E-05	2.9E-04	9.7E-05
κ_4^*	-1.1E-05	3.5E-05	-6.3E-05	2.1E-05	2.3E-04	4.9E-05	2.4E-04	4.8E-05	7.2E-05	7.7E-05	1.3E-04	8.9E-05
κ_5^*	-2.2E-05	3.9E-05	-1.2E-04	2.8E-05	2.9E-04	5.5E-05	2.4E-04	5.2E-05	2.4E-04	9.4E-05	-2.0E-04	9.5E-05
ω	11.743	0.070	12.025	0.054	11.519	0.071	11.586	0.045	10.946	0.003	11.233	0.128
ν	2.604	0.062	2.674	0.064	3.071	0.095	3.076	0.095	2.631	0.111	2.653	0.122
ξ	3.161	0.132	3.038	0.127	3.404	0.233	3.394	0.232	4.543	0.439	4.489	0.482
ζ	2.185	0.085	2.097	0.081	1.783	0.077	1.779	0.077	2.217	0.138	2.183	0.144
AIC	23.562		23.557		23.5471		23.5470		22.095		22.094	
SIC	23.568		23.563		23.553		23.552		22.106		22.104	
Loglike	-388,209		-388,136		-477,508		-477,509		-222,725		-222,715	

Table 4: The estimated coefficients

Window	AIC		SIC	
	Static S	Dyn. S	Static S	Dyn. S
1 (2002-2005)	23.561	23.557	23.566	23.563
2 (2007-2010)	23.552	23.547	23.556	23.552
3 (2012-2013)	22.099	22.094	22.106	22.104

Window	Static S (Restricted)	Dyn. S (Unrestricted)	Likelihood ratio stat	χ^2_6 p-value
1 (2002-2005)	-388,207	-388,136	141	0.000
2 (2007-2010)	-477,622	-477,509	226	0.000
3 (2012-2013)	-222,770	-222,715	111	0.000

Table 5: The static spline versus the dynamic spline. The second table shows the likelihood ratio statistics to test the null of static daily spline ($H_0 : \boldsymbol{\kappa}^* = 0$) against the alternative of dynamic daily spine ($H_1 : \boldsymbol{\kappa}^* \neq \mathbf{0}$). This assumes that the statistic asymptotically follows the chi-squared distribution with six degrees of freedom. The final model specification given in Table 4 is used.

4.1. Significance of dynamic periodicity

The estimates of $\boldsymbol{\kappa}^*$ are close to zero in Table 4. Using the Akaike information criterion (AIC), the Schwarz information criterion (SIC), and the likelihood ratio test, we can assess the significance of $\boldsymbol{\kappa}^*$. If they are significant, the data provides evidence for dynamic periodic patterns. The null is the static spline ($H_0 : \boldsymbol{\kappa}^* = \mathbf{0}$) and the alternative is the dynamic spline ($H_1 : \boldsymbol{\kappa}^* \neq \mathbf{0}$). The statistics are tabulated in Table 5. We find statistical evidence for dynamic periodicity in our data. Both SIC and AIC are in favor of the dynamic spline over the static one. The likelihood ratio test also rejects the null at the 1% significance level.[8]

Figure 3 shows a three dimensional visualization of the estimated dynamic spline. We find that $\widehat{s}_{t,\tau}$ successfully captures diurnal U-shaped patterns. The pattern of periodicity varies over time.

The level of the spline at the beginning of the day is different from the level at the end of the day on any two consecutive trading days, reflecting the overnight effect. For the first in-sample window, the level of the spline at the start of the day is higher than the level at the end of the day, whereas the opposite is true for the second and the third windows.

It is interesting to observe in the in-sample window 2 that the amplitude of the U-shape increases in the latter half of this period. In particular, the height of the spline at

[8]The likelihood ratio test should be interpreted with care as it is assumed that we do not have a boundary value problem of the type studied in Chernoff (1954) under the null. We assume this given that the elements of $\boldsymbol{\kappa}^*$ can be in any quadrant. If this assumption is violated, standard theory based on the statistic being asymptotically chi-squared would lead to the region of acceptance being too large (see Harvey (1991, p. 236)). Since H_0 is rejected in our case, this should imply that H_0 will be rejected at the same significance level even if we have the boundary value problem.

(a) The in-sample window 1 (Jan. 2002 - Apr. 2005)

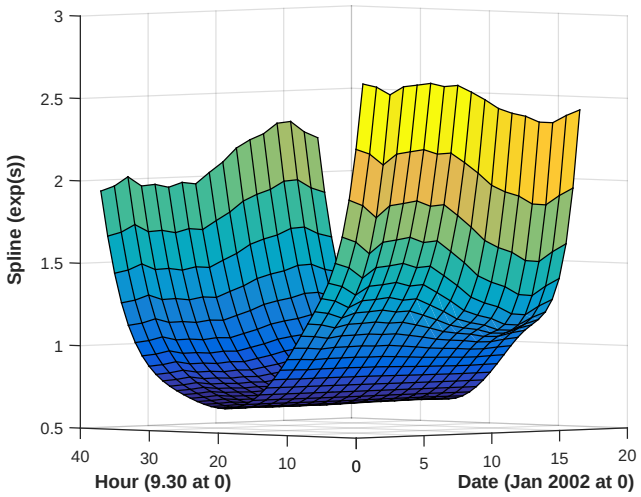

(b) The in-sample window 2 (Jan. 2007 - Dec. 2010)

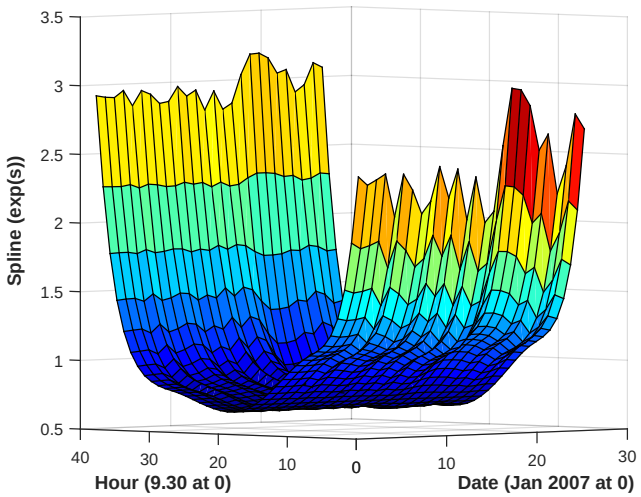

(c) The in-sample window 3 (Jan. 2012 - Dec. 2013)

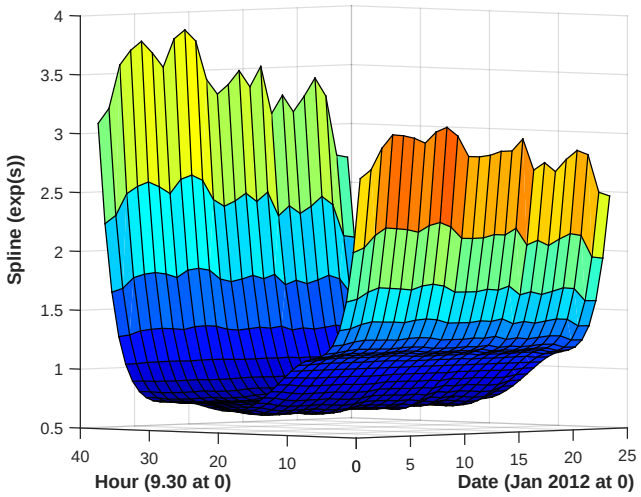

Figure 3: $\exp(\widehat{s}_{t,\tau})$ for IBM10m when $s_{t,\tau}$ is the dynamic spline. The x-axis and the y-axis are trading time between 9.30am-4pm (9.30am at the origin) and the sampling days (the first day of the sampling period at the origin).

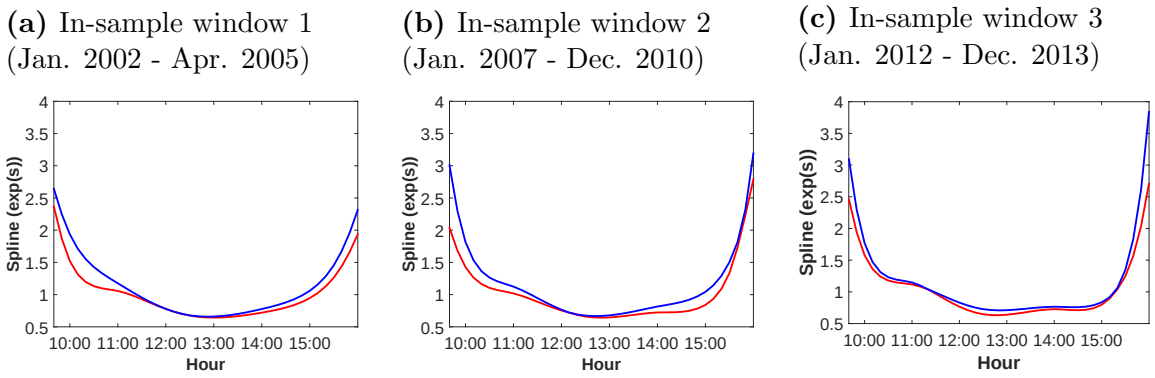

Figure 4: The min-max range of $\exp(\widehat{s}_{t,\tau})$ at each intra-day bin for the specified in-sample windows. The spline-DCS with the dynamic spline. Intra-day hours between 9.30am and 4pm in the NY local time along the x-axis.

the beginning of the day increases by a lot in the middle of the in-sample window 2. This roughly coincides with the beginning of the 2007-2008 financial crisis in the U.S. It could be a reflection of changes in the amount of news transmitted at night, especially since many policy decisions were made at night during the crisis. The profitability of the firm was also at risk and its revenue declined during the U.S. economic downturn in 2008 and 2009. (See Figure 11 in Appendix A.) The relatively large flow of overnight news and the anticipation of it may have intensified trading activities at the beginning and the end of each day. The amplitude of the U-shape declines to a level comparable to the pre-crisis period by the end of 2013.[9]

Figure 4 shows that the shape of the spline evolves the most in the morning and the afternoon, but does not change much around midday. Moreover, transactions at the beginning and the end of the day have grown over the past decade to account for a larger share of total daily volume by the end of 2013. This could be due to high-frequency algorithmic trading, which gained popularity in the past decade. High-frequency traders typically try to minimize the price impact of orders by slicing a given order into small transaction sizes and spreading them throughout the day. An effective way of placing transactions to achieve this objective is to mimic the pattern of market activity and disguise themselves in the market. The rise of algorithmic traders that exploit the pattern of periodicity can exacerbate periodic behavior.

4.2. Other estimated coefficients and diagnostics

No coefficient restrictions were imposed to obtain the estimated values in Table 4, except that $\kappa_\eta^{(1)} < \kappa_\eta^{(2)}$ was imposed for the in-sample window 3. In the final specification, coefficients are set to zero if they are insignificant at the 5% significance level and AIC and SIC fall by doing so. We set marginally significant coefficients to zero if doing so

[9]It would be interesting to investigate more formally any relationship between indicators of company profitability and the shape of periodicity.

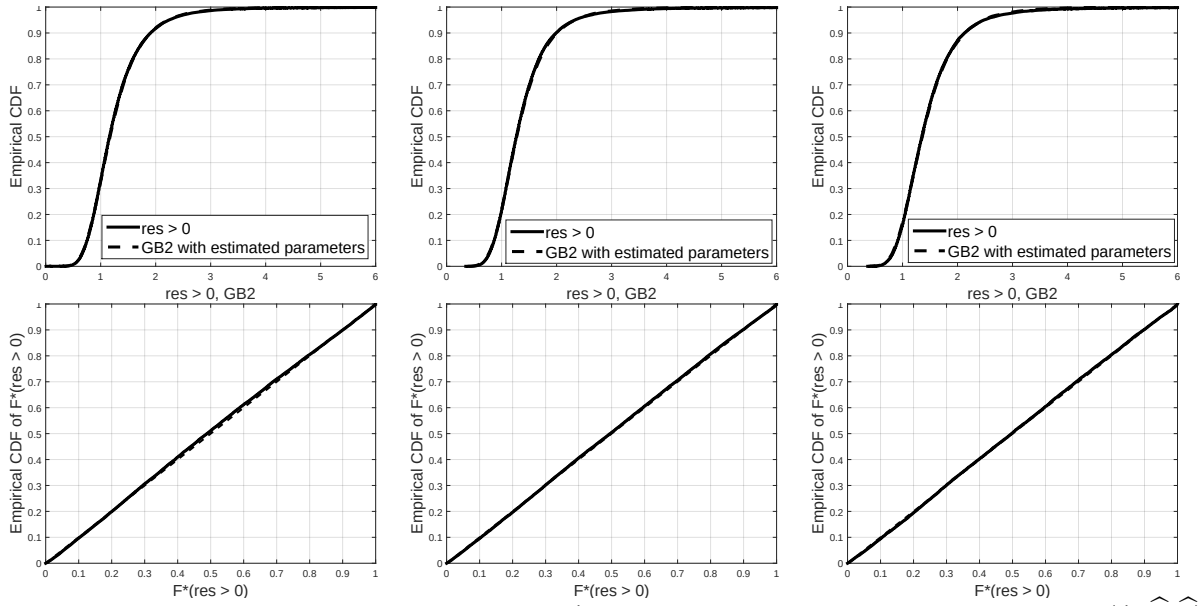

Figure 5: The empirical c.d.f. of non-zero $\widehat{\varepsilon}_{t,\tau}$ against the theoretical c.d.f. of GB2$(\widehat{\nu}, \widehat{\xi}, \widehat{\zeta})$ (left). The empirical c.d.f. of the PIT values of non-zero $\widehat{\varepsilon}_{t,\tau}$ when $F^*(\cdot; \widehat{\boldsymbol{\theta}}^*)$ is GB2$(\widehat{\nu}, \widehat{\xi}, \widehat{\zeta})$ (right). For the in-sample window 1 (left), 2 (center), and 3 (right). The spline-DCS with the dynamic spline.

Window	1	2	3
KS test (p-val.)	0.00	0.02	0.21

Table 6: The Kolmogorov-Smirnov test to assess the null that $\widehat{\varepsilon}_{t,\tau}$ comes from GB2$(\widehat{\nu}, \widehat{\xi}, \widehat{\zeta})$. The table reports the p-values of the statistics.

decreases AIC and SIC.[10] Based on the empirical findings of Section 3, the generalized beta distribution of the second kind (GB2) was chosen and tested as the error distribution, F^*. GB2 has the shape parameters $\theta^* = (\nu, \xi, \zeta)$.

In the final specifications, we have $\widehat{\kappa}_\eta^{(2)} > \widehat{\kappa}_\eta^{(1)} > \widehat{\kappa}_\mu > 0$ so that $\eta_{t,\tau}^{(2)}$ is more sensitive to changes in $u_{t,\tau-1}$ than $\eta_{t,\tau}^{(1)}$, and that $\eta_{t,\tau}^{(1)}$ is more sensitive to $u_{t,\tau-1}$ than $\mu_{t,\tau}$. We also have $0 < \widehat{\phi}_1^{(2)} < 1$ so $\eta_{t,\tau}^{(2)}$ is stationary. For the in-sample window 1, the stationarity of $\eta_{t,\tau}^{(1)}$ in the final model requires $\phi_1^{(1)} + \phi_2^{(1)} < 1$, $-\phi_1^{(1)} + \phi_2^{(1)} < 1$, and $\phi_2^{(1)} > -1$. (See Harvey (1993, p.19).) These conditions are satisfied by $\widehat{\phi}_1^{(1)}$ and $\widehat{\phi}_2^{(1)}$ so that $\eta_{t,\tau}^{(1)}$ is stationary. For the in-sample windows 2 and 3, numerically inspecting the roots of the equation, $1 - \widehat{\phi}_1^{(1)} z - \widehat{\phi}_2^{(1)} z^2 - \widehat{\phi}_3^{(1)} z^3 = 0$, for some real or complex z gave the roots outside the unit disk, suggesting that $\eta_{t,\tau}^{(1)}$ is stationary. The parameter values on the lagged coefficients suggest that $\widehat{\eta}_{t,\tau}^{(1)}$ is more persistent than $\widehat{\eta}_{t,\tau}^{(2)}$ in all of the final specifications. AIC and SIC decreased when we set $\beta_{1,1} = 0$ in all of the cases. For the in-sample windows 2 and 3, $\widehat{\kappa}_{a,\eta}^{(1)}$ was insignificant, suggesting that the asymmetry effect is significant only for the short-run stationary component $\eta_{t,\tau}^{(2)}$. This is consistent with Engle and Lee (1999)

[10] We relied more on AIC and SIC than the t-statistics when selecting the model since we did not formally check the asymptotic properties of the ML estimators in our version of the DCS model.

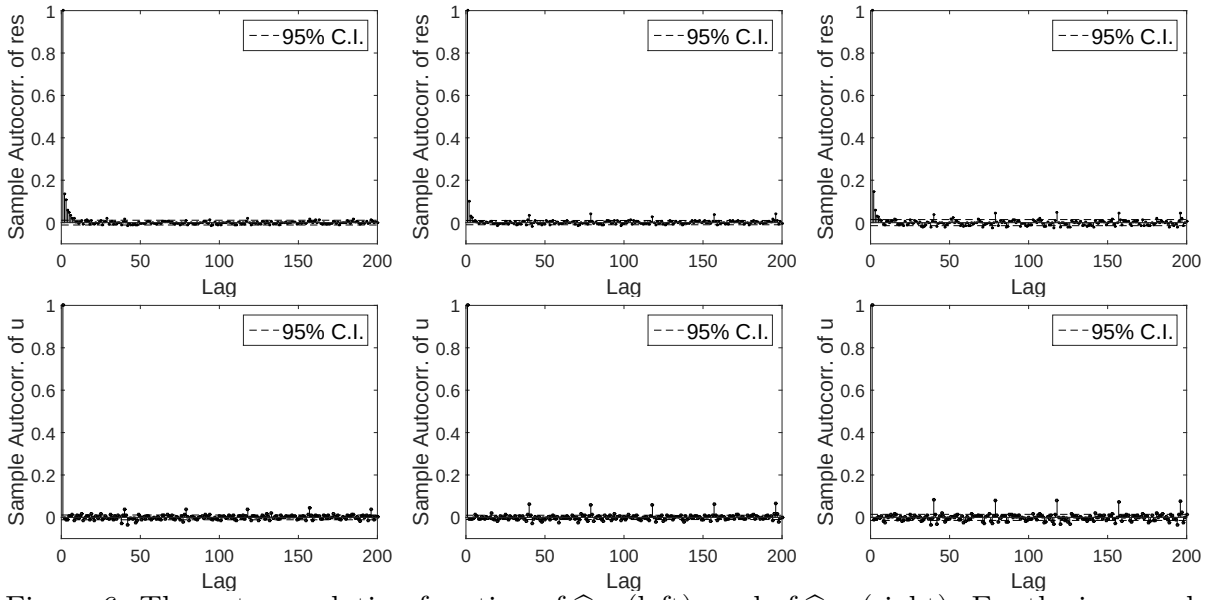

Figure 6: The autocorrelation function of $\widehat{\varepsilon}_{t,\tau}$ (left), and of $\widehat{u}_{t,\tau}$ (right). For the in-sample window 1 (left), 2 (center), and 3 (right). The spline-DCS with the dynamic spline.

and Harvey (2013), who find that the leverage effect in two-component return volatility models is mainly restricted to the short-run volatility component. However, for the in-sample window 1, $\kappa_{a,\eta}^{(1)}$ was significant but not $\kappa_{a,\eta}^{(2)}$, which means that the asymmetry effect is in the low-frequency stationary component, $\eta_{t,\tau}^{(1)}$. In all of the final specifications, the sign of the asymmetry effect is positive, which is consistent with the study of trade volume of selected Exchange Traded Funds by Brownlees et al. (2011).

Figure 5 illustrates the quality of the fit of GB2. The empirical c.d.f. of non-zero $\widehat{\varepsilon}_{t,\tau}$ overlaps the c.d.f. of GB2$(\widehat{\nu},\widehat{\xi},\widehat{\zeta})$ so that these two lines are visually indistinguishable. The closeness of the fit is also seen in the empirical c.d.f. of the probability integral transform (PIT) of non-zero $\widehat{\varepsilon}_{t,\tau}$ when $F^*(\cdot;\widehat{\boldsymbol{\theta}}^*)$ is GB2$(\widehat{\nu},\widehat{\xi},\widehat{\zeta})$. It lies very close to the diagonal, indicating that the PIT values are remarkably close to being standard uniformly distributed (denoted by $U[0,1]$). The Kolmogorov-Smirnov test rejects the null that $\widehat{\varepsilon}_{t,\tau}$ comes from GB$(\widehat{\nu},\widehat{\xi},\widehat{\zeta})$ at the 1% level for the in-sample window 1 and at the 5% level for the in-sample window 2. (See Table 6.) We think that this is partly due to the relative large sample size for these windows. The estimated GB2 parameters give $5 < \widehat{\nu}\widehat{\zeta} < 6$ so that the moments up to the fifth exist. Since $\widehat{\nu}\widehat{\zeta}$ is the upper tail-index of GB2, the estimated GB2 is heavy-tailed.

Figure 6 shows that the dynamics of IBM10m are captured fairly well by the model as the estimation residuals $\widehat{\varepsilon}_{t,\tau}$ and the score $\widehat{u}_{t,\tau}$ show little to no marked signs of serial correlation, although there is a small degree of persistence in the first few lags of $\widehat{\varepsilon}_{t,\tau}$. Outliers in $\widehat{\varepsilon}_{t,\tau}$ usually weaken serial correlation, but the small persistence in $\widehat{\varepsilon}_{t,\tau}$ (as opposed to $\widehat{u}_{t,\tau}$) could be due to a couple of successive large observations in $\widehat{\varepsilon}_{t,\tau}$. See Harvey and Lange (2015).[11] The model could not perfectly capture the daily periodic

[11] $\widehat{u}_{t,\tau}$ is more robust to outliers than $\widehat{\varepsilon}_{t,\tau}$ because the score weighs down the effects of large-sized

	$\widehat{\varepsilon}_{t,\tau}$						$\widehat{\varepsilon}^2_{t,\tau}$					
Window	$\widehat{\rho}_1$	$\widehat{\rho}_{day}$	Q_1	Q_{day}	p-val. 1	p-val. day	$\widehat{\rho}_1$	$\widehat{\rho}_{day}$	Q_1	Q_{day}	p-val. 1	p-val. day
1	0.148	0.003	420.807	628.262	0.000	0.000	0.231	-0.003	1033.660	1265.957	0.000	0.000
2	0.135	0.003	574.188	1231.820	0.000	0.000	0.039	0.000	47.542	182.244	0.000	0.000
3	0.100	0.000	390.045	547.160	0.000	0.000	0.041	-0.001	65.933	72.469	0.000	0.000

	$\widehat{u}_{t,\tau}$						$\widehat{u^2}_{t,\tau}$					
Window	$\widehat{\rho}_1$	$\widehat{\rho}_{day}$	Q_1	Q_{day}	p-val. 1	p-val. day	$\widehat{\rho}_1$	$\widehat{\rho}_{day}$	Q_1	Q_{day}	p-val. 1	p-val. day
1	-0.010	0.015	1.873	237.396	0.171	0.000	0.036	0.016	24.535	135.275	0.000	0.000
2	-0.005	0.009	0.807	168.844	0.369	0.000	0.065	0.021	134.920	699.380	0.000	0.000
3	-0.007	0.009	1.799	265.434	0.180	0.000	0.027	0.020	28.141	271.608	0.000	0.000

Table 7: In-sample residual analysis for the spline-DCS with the dynamic cubic spline. Q_l is the Ljung-Box statistic to test the null of no autocorrelation up to the l-th lag.

(a) Out-of-sample window 1. 50 days ahead (top). 100 days ahead (bottom).

(b) Out-of-sample window 2. 50 days ahead (top). 100 days ahead (bottom).

(c) Out-of-sample window 3. 50 days ahead.

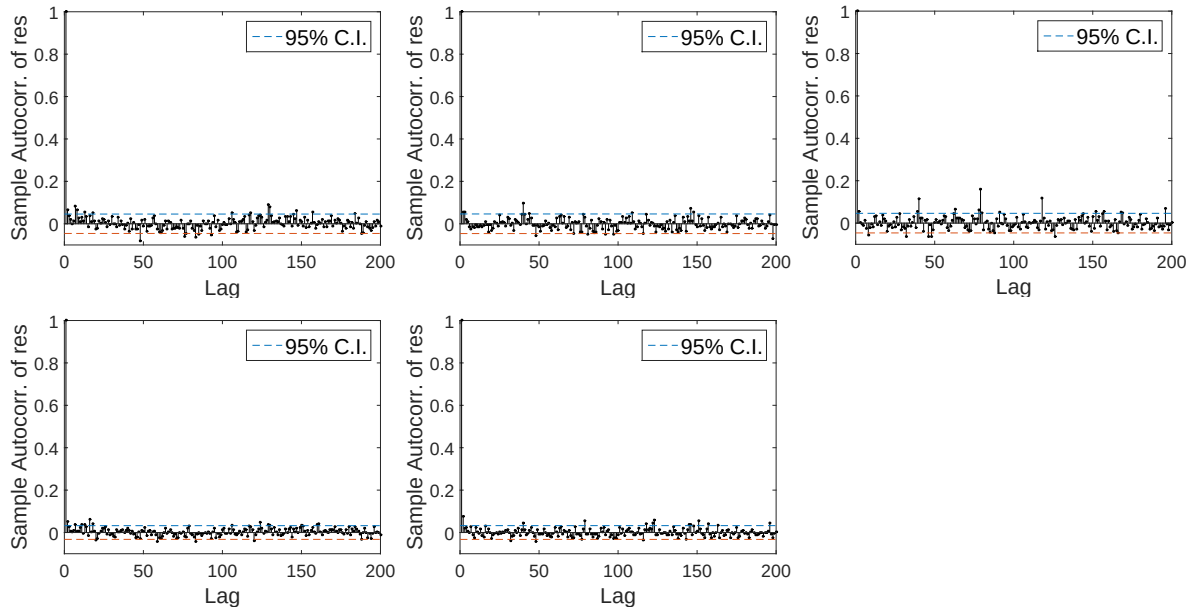

Figure 7: The autocorrelation function of non-zero $\widetilde{\varepsilon}_{t,\tau}$ for the specified out-of-sample window. The spline-DCS with the dynamic spline. The parameter estimates were obtained by fitting the model to the corresponding in-sample window.

patterns, as there are still peaks in the autocorrelation of $\widehat{\varepsilon}_{t,\tau}$ and $\widehat{u}_{t,\tau}$ at the multiples of the 39th lag. But the model does a good job in capturing a substantial part of the periodic patterns. The autocorrelation at these lags could be removed further by including additional information (e.g. overnight price change) to improve how the model captures the overnight effect. The features discussed here are also seen in the Ljung-Box tests reported in Table 7. It is worthwhile noting that the large sample sizes make the Ljung-Box test statistics sensitive to small departures from zero autocorrelation.

observations.

	$\widehat{\varepsilon}_{t,\tau}$						$\widehat{\varepsilon}^2_{t,\tau}$					
Window	$\widehat{\rho}_1$	$\widehat{\rho}_{day}$	Q_1	Q_{day}	p-val. 1	p-val. day	$\widehat{\rho}_1$	$\widehat{\rho}_{day}$	Q_1	Q_{day}	p-val. 1	p-val. day
1	0.050	0.010	9.686	81.429	0.002	0.000	0.043	0.012	7.030	88.390	0.008	0.000
2	0.076	0.014	21.768	66.052	0.000	0.000	0.031	-0.001	3.577	10.001	0.059	0.998
3	0.056	0.026	6.622	79.423	0.010	0.000	0.033	0.011	2.410	36.065	0.121	0.090

Table 8: Out-of-sample residual analysis of one-step ahead forecasts for the spline-DCS with the dynamic cubic spline. Q_l is the Ljung-Box statistic to test the null of no autocorrelation up to the l-th lag. We test $\widetilde{\varepsilon}_{t,\tau}$ up to 100 out-of-sample days ahead

Window	1	2	3
KS test (p-val.)	0.09	0.36	0.00

Table 9: The Kolmogorov-Smirnov test to assess the null that $\widetilde{\varepsilon}_{t,\tau}$ comes from GB2$(\widehat{\nu}, \widehat{\xi}, \widehat{\zeta})$. The table reports the p-values of the statistics. We test $\widetilde{\varepsilon}_{t,\tau}$ up to 100 out-of-sample days ahead.

5. Out-of-sample performance

5.1. One-step ahead forecasts

We use the predictive c.d.f. to assess the stability of the model and the performance of its one-step ahead forecasts over a given out-of-sample period. The notations and the methodology are the same as Ito (2016). We recall some of them here. We use the set notation, $\Psi_h = \{(t,\tau) \in \{T+1,\ldots,T+h\} \times \{1,\ldots,I\}\}$, where $h \in \mathbb{N}_{>0}$ is the number of out-of-sample days.

The predictive c.d.f. is $F^*(\widetilde{\varepsilon}_{t,\tau}; \widehat{\boldsymbol{\theta}}^*)$, where $\widetilde{\varepsilon}_{t,\tau} = y_{t,\tau}/\exp(\widetilde{\lambda}_{t,\tau})$ for all $(t,\tau) \in \Psi_h$. The one-step ahead forecasts, $\widetilde{\lambda}_{t,\tau}$, is obtained by recursively updating $\lambda_{t,\tau}$ at each new out-of-sample observation point $y_{t,\tau}$ for $(t,\tau) \in \Psi_h$ without re-estimating the in-sample parameter values. The parameter values used here are from the estimation results we discussed in Section 4. The predictive c.d.f. gives the PIT values of forecast standardized observations. We use the out-of-sample windows tabulated in Table 1.

Figure 7 shows the autocorrelation function of the forecast standardized observations, $\widetilde{\varepsilon}_{t,\tau}$, when the model is the final specification given in Table 4. They are roughly free of autocorrelation for up to 50 or 100 days ahead. As we have 39 observations per day for IBM10m, $h = 100$ corresponds to 3,900 steps ahead. However, the Ljung-Box statistics detect statistically significant autocorrelation $\widehat{\varepsilon}_{t,\tau}$ as reported in Table 8, when we test $\widetilde{\varepsilon}_{t,\tau}$ up to 100 out-of-sample days ahead. As before, this is partly due to the large sample size.

Figure 8 shows the empirical c.d.f. of the predictive c.d.f. of $\widetilde{\varepsilon}_{t,\tau}$. It illustrates that the PIT values are remarkably close to being $U[0,1]$ for an extended out-of-sample period, especially for the out-of-sample windows 1 and 2. For these windoes, the Kolmogorov-Smirnov tests cannot reject the null that the forecast residuals $\widetilde{\varepsilon}_{t,\tau}$ comes from GB2$(\widehat{\nu}, \widehat{\xi}, \widehat{\zeta})$ at the 5% level, when we test $\widetilde{\varepsilon}_{t,\tau}$ up to 100 out-of-sample days ahead. For the out-of-

(a) Out-of-sample window 1. **(b)** Out-of-sample window 2. **(c)** Out-of-sample window 3.

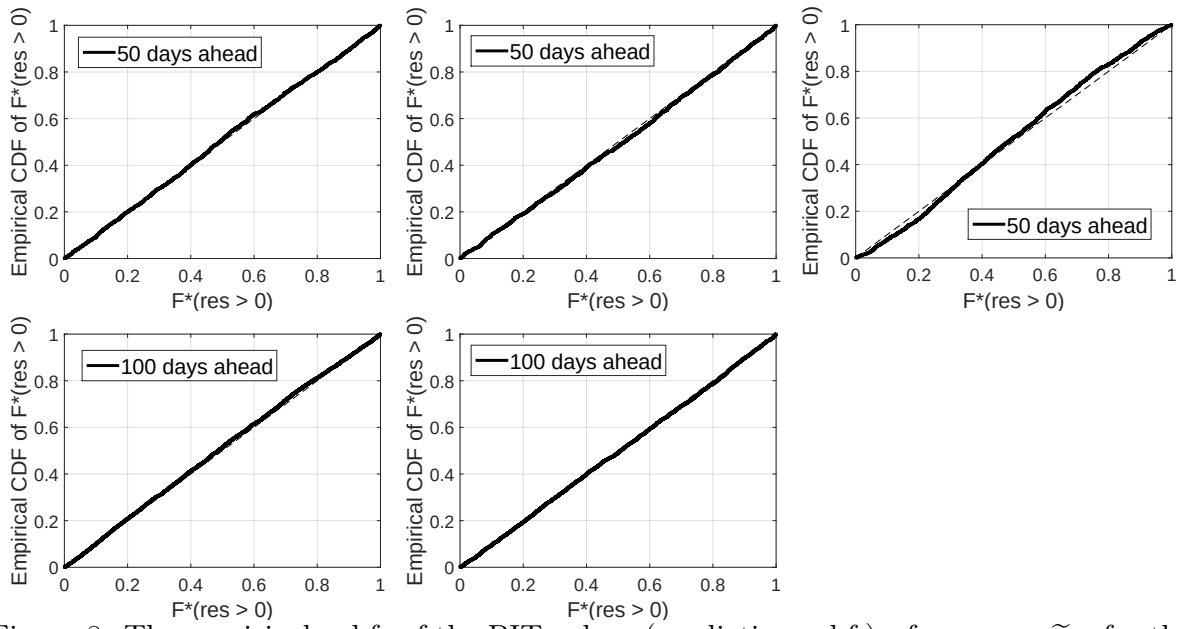

Figure 8: The empirical c.d.f. of the PIT values (predictive c.d.f.) of non-zero $\widetilde{\varepsilon}_{t,\tau}$ for the specified out-of-sample window. The spline-DCS with the dynamic spline.

sample windows 1 and 2, the Kolmogorov-Smirnov tests cannot reject the null that $\widetilde{\varepsilon}_{t,\tau}$ comes from GB2$(\widehat{\nu}, \widehat{\xi}, \widehat{\zeta})$ when we test $\widetilde{\varepsilon}_{t,\tau}$ for up to 100 out-of-sample days of the respective windows. Thus, the estimated GB2 appears to capture the empirical distribution of the out-of-sample data very well.

Figure 9 compares the QQ-plot of $\widetilde{\varepsilon}_{t,\tau}$ for when the spline is either dynamic or static. The model used here is still the final specification given in Table 4, and the static spline is obtained by setting $\boldsymbol{\kappa}^* = 0$. If the tail region of a given QQ-plot lies below the diagonal line, it means that the estimated GB2 (i.e. GB2$(\widehat{\nu}, \widehat{\xi}, \widehat{\zeta})$) understates the size of the tail of the data. For the out-of-sample windows 1 and 3, it is interesting to observe that the QQ-plot for the model with the dynamic spline appears to lie closer to the diagonal line (especially in the tail region) than they do for the model with the static spline. We do not observe any marked difference in the QQ-plot for the out-of-sample window 2. These results suggest that the dynamic spline does either as well as, or better than, the static spline in capturing the empirical distribution of the data.

5.2. Out-of-sample model comparison: static versus dynamic spline

A quantity of great interest in high-frequency volume prediction is the proportion of daily trade volume attributed to each intra-day bin. That is, $w_{t,\tau} = y_{t,\tau} / \sum_{\tau=1}^{I} y_{t,\tau}$ for each $(t, \tau) \in \Psi_h$. Accurate prediction of this quantity helps traders minimize the slippage of each transaction, which is the difference between the expected price and the actual traded price, and to achieve the execution price of transactions for the day to be near the daily

(a) Out-of-sample window 1

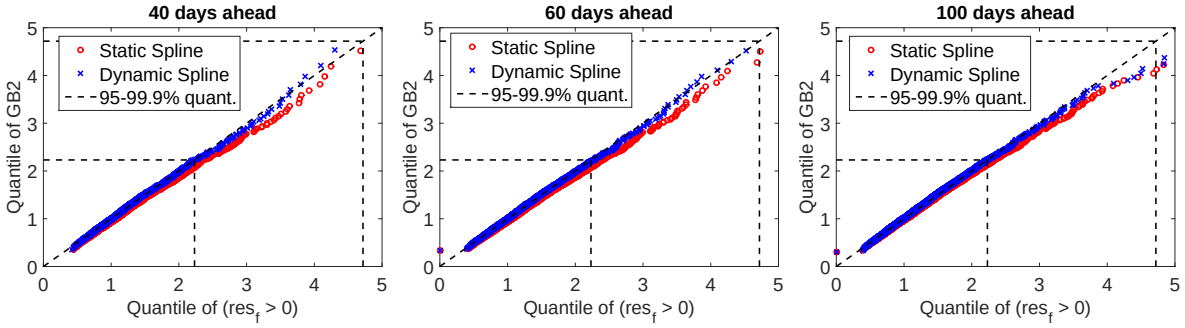

(b) Out-of-sample window 2

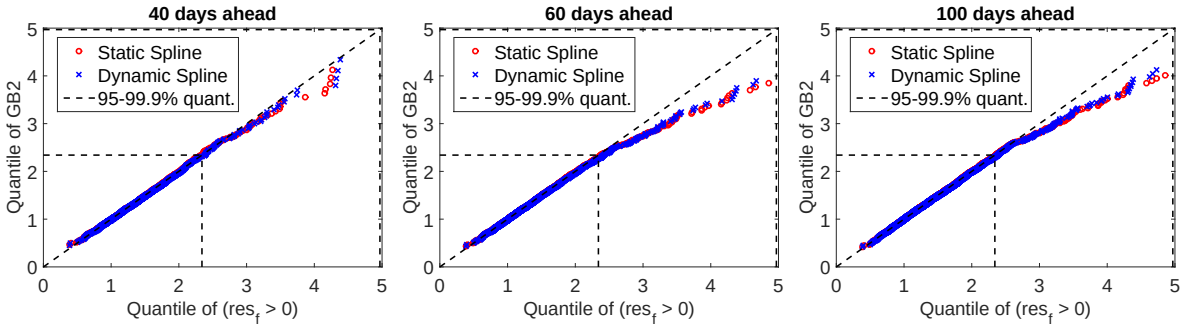

(c) Out-of-sample window 3

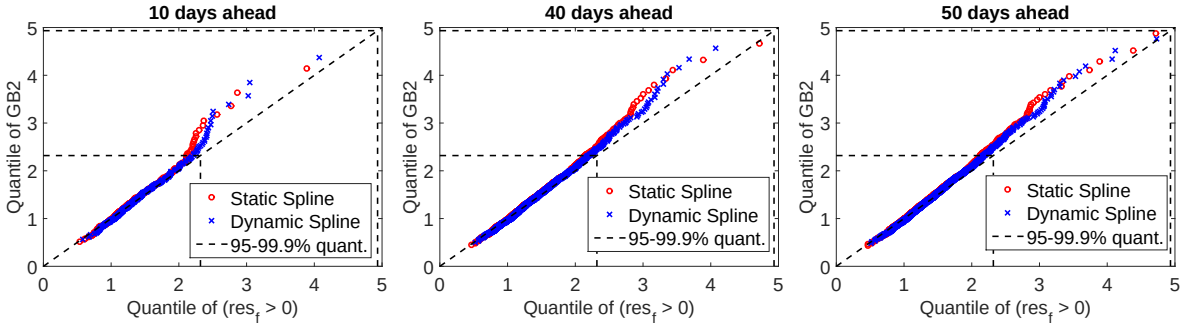

Figure 9: The QQ-plot of non-zero $\widetilde{\varepsilon}_{t,\tau}$ for the specified out-of-sample window. The spline-DCS with the static spline (o, red) and the dynamic spline (x, blue). The off-diagonal dashed lines are the 95% and 99.9% theoretical quantiles of GB2$(\widehat{\nu}, \widehat{\xi}, \widehat{\zeta})$, where the distribution parameters are obtained from estimating the spline-DCS with the static spline.

VWAP benchmark. The loss function we consider here is the daily slicing loss function of Brownlees et al. (2011) given by:

$$L^{slicing}\left((y_{T+h,\tau}, \widetilde{y}_{T+h,\tau})_{\tau=1}^{I}\right) = -\sum_{\tau=1}^{I} w_{T+h,\tau} \log \widehat{w}_{T+h,\tau}.$$

This slicing loss is developed by Brownlees et al. (2011) to evaluate the performance of trading strategies that aim to replicate the daily VWAP. The slicing loss is a common term determining the ranking of models by the negative multinomial log-likelihood loss and the Kullback-Leibler loss. We compute the forecast slicing weights (or volume proportions), $\widehat{w}_{T+h,\tau}$, under the dynamic VWAP replication strategy outlined by Brown-

Window	1	2	3
Dynamic VWAP	-21.56***	-17.86***	20.22***

Table 10: Diebold-Mariano statistics to test the null of equal predictive ability against the alternative of different predictive ability, which is a two-sided test with * 10% significance, ** 5% significance, and *** 1% significance. If the alternative is of a one-sided test, a statistically significant negative (positive) value is in favor of the dynamic (static) spline.

lees et al. (2011).[12] The dynamic VWAP updates static-VWAP throughout the day to re-compute proportions forecast for the rest of the day using new intra-day data. The static-VWAP replication strategy is volume proportions forecast for one complete day ahead, but it is not updated during the day. See Brownlees et al. (2011) for the details.

The dynamic VWAP is a function of one-step and multi-step ahead point forecasts of trade volume. With the spline-DCS, the formulae we use for one-step ahead point forecast is:

$$\widetilde{y}_{T+h,\tau+1} \equiv \mathbb{E}[y_{T+h,\tau+1}|\mathcal{F}_{T+h,\tau}] = \exp(\widetilde{\lambda}_{T+h,\tau+1}) \int_0^\infty x f(x;\widehat{\boldsymbol{\theta}})dx. \tag{3}$$

Since the error distribution is fully given by GB2, the integral can be computed analytically. The multi-step predictor of volume for a given day is $\widetilde{y}_{t,\tau+i} = \mathbb{E}\left[\exp(\lambda_{t,\tau+i})|\mathcal{F}_{t,\tau}\right] \int_0^\infty x f(x;\widehat{\boldsymbol{\theta}})dx$ for $i = 1, \ldots, I-\tau$ and $(t,\tau) \in \Psi_H$. The conditional moment is evaluated using the moment generating function (m.g.f.) of the score variable, $u_{t,\tau+i}$, for $i = 1, \ldots, I-\tau$, as well as $\widetilde{\lambda}_{t,\tau}$ computed above because we do not re-estimate constant parameters over a given out-of-sample period. See Brownlees et al. (2011). The m.g.f. of $u_{t,\tau}$ exists and can be evaluated analytically in our case.

The Diebold-Mariano statistics reported in Table 10 are comfortably in favor of the dynamic spline over the static one for the out-of-sample windows 1 and 2. In other words, for these windows, the use of the dynamic spline leads to the dynamic VWAP that substantially improves on the model with the static spline in minimizing the slicing loss and achieving the VWAP target.

5.3. Discussions

In Section 5.1, we observed an improvement in the quality of the fit of the model to the empirical distribution of the data, especially in the tail region. Moreover, in Section 5.2, a model that allows for changing periodic patterns can lead to improved volume proportions forecasts, and outperform the model with the static spline in minimizing the slicing loss function.

These results can be explained by the fact that the dynamic spline can reflect changes in the pattern of morning and afternoon trading activity, and adjust how the model standardizes relatively large-sized observations each day. As Figures 4 and 10 show, the

[12]We do not consider the VWAP-tracking MSE discussed in Bialkowski et al. (2008) as we did not have the price data. Brownlees et al. (2011) prefer the slicing loss, which is less noisy.

(a) Out-of-sample window 1 **(b)** Out-of-sample window 2 **(c)** Out-of-sample window 3

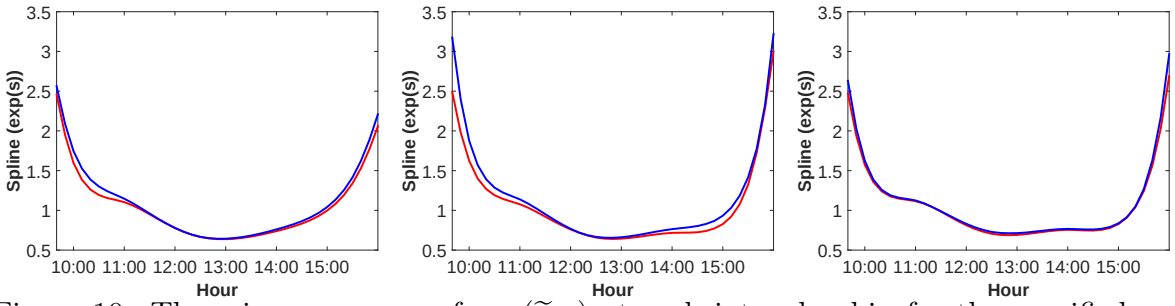

Figure 10: The min-max range of $\exp(\widetilde{s}_{t,\tau})$ at each intra-day bin for the specified out-of-sample windows. The spline-DCS with the dynamic spline. Intra-day hours between 9.30am and 4pm in the NY local time along the x-axis.

pattern of periodicity changes the most in the morning and in the afternoon. Among the out-of-sample windows, the pattern of periodicity changes the most in the out-of-sample windows 1 and 2. They are the windows in which the dynamic spline outperformed the static one in capturing the tail of the data and minimizing the slicing loss.

6. Concluding remarks

This paper generalized the spline-DCS to capture dynamic diurnal patterns in high-frequency financial data. We found statistical evidence for changing diurnal patterns in our data, and that this generalization can improve the quality of the fit of the model to the data, especially in the tail region. Moreover, it can lead to a significantly improved performance in volume proprotions forecasts and minimizing the slicing loss function for assessing the optimality of VWAP replication strategies. The out-of-sample forecast results also show that the in-sample estimation results are stable, and that our model is able to provide good one-step ahead forecasts.

It would be interesting to introduce a weekly spline by letting the periodicity to be complete over one week instead of one day. Then we can allow for the shape of periodicity to depend on the day of the week. This introduces a new way of capturing the day-of-the-week effect in finance. In the literature, it is modelled using dummy variables, which only allows for the level of the data to depend on the day of the week. We can also allow for the weekly spline to evolve over time in the fashion we introduced in this paper.

References

Andersen, T. G. and Bollerslev, T. (1998), "Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Annoucements, and Longer Run Dependencies," *Journal of Finance*, 53, 219–265.

Bialkowski, J., Darolles, S., and Le Fol, G. (2008), "Improving VWAP Strategies: A Dynamic Volume Approach," *Journal of Banking and Finance*, 32, 1709–1722.

Brownlees, C. T., Cipollini, F., and Gallo, G. M. (2011), "Intra-Daily Volume Modelling

and Prediction for Algorithmic Trading," *Journal of Financial Econometrics*, 9, 489–518.

Campbell, S. D. and Diebold, F. X. (2005), "Weather Forecasting for Weather Derivatives," *Journal of American Statistical Association*, 100, 6–16.

Chernoff, H. (1954), "On the Distribution of the Likelihood Ratio," *Annals of Mathematical Statistics*, 25, 573–578.

Engle, R. F. and Lee, G. G. J. (1999), "A Long-Run and Short-Run Component Model of Stock Return Volatility," in *Cointegration, Causality, and Forecasting: A Festschrift in Honour of Clive W.J. Granger.*, eds. Engle, R. F. and White, H., Oxford University Press.

Engle, R. F. and Rangel, J. G. (2008), "The Spline-GARCH Model for Low-Frequency Volatility and Its Global Macroeconomic Causes," *Review of Financial Studies*, 21, 1187–1222.

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

Engle, R. F. and Sokalska, M. E. (2012), "Forecasting Intraday Volatility in the US Equity Market. Multiplicative Component GARCH," *Journal of Financial Econometrics*, 10, 54–83.

Harvey, A. C. (1991), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press.

— (1993), *Time Series Models*, Harvester: Wheatsheaf, 2nd ed.

— (2013), *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*, Econometric Society Monograph, Cambridge University Press.

Harvey, A. C. and Koopman, S. J. (1993), "Forecasting Hourly Electricity Demand Using Time-Varying Splines," *Journal of the American Statistical Association*, 88, 1228–1236.

Harvey, A. C. and Lange, R.-J. (2015), "Modeling the Interactions between Volatility and Returns," Cambridge Working Papers in Economics CWPE1518, University of Cambridge.

Ito, R. (2016), "Spline-DCS for Forecasting Trade Volume in High-Frequency Financial Data," Cambridge Working Papers in Economics CWPE1606, University of Cambridge.

Zhang, M. Y., Russell, J. R., and Tsay, R. S. (2001), "A Nonlinear Autoregressive Conditional Duration Model with Applications to Financial Transaction Data," *Journal of Econometrics*, 104, 179–207.

Appendix A: Selected balance-sheet statistics

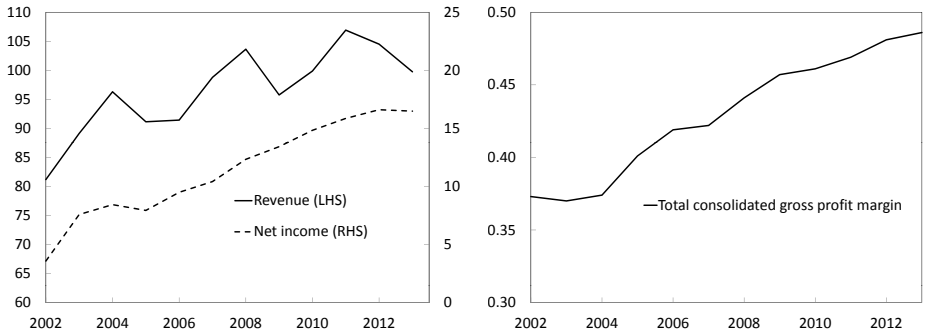

Figure 11: The annual revenue and net income (in $ billions, left), and the consolidated gross profit margin (right) for IBM for the fiscal years ending December. (Note that the left axis of the revenue line starts from $60 billion.) Source: IBM annual report (www.ibm.com).