Journal of Empirical Finance 16 (2009) 686–700

Contents lists available at ScienceDirect

Journal of Empirical Finance

journal homepage: www.elsevier.com/locate/jempfin

Which power variation predicts volatility well?☆

Eric Ghysels [a,b,*], Bumjean Sohn [b]

[a] Department of Economics, University of North Carolina at Chapel Hill, United States
[b] Department of Finance, Kenan-Flagler Business School, University of North Carolina at Chapel Hill, United States

ARTICLE INFO

Article history:
Received 11 September 2007
Received in revised form 29 December 2008
Accepted 5 March 2009
Available online 2 April 2009

JEL classification:
C22
C53
G17

Keywords:
Stock Market Volatility
Forecasting
Power variation
MIDAS regressions

ABSTRACT

We estimate MIDAS regressions with various (bi)power variations to predict future volatility – measured via increments in quadratic variation. Instead of pre-determining the (bi)power variation we parameterize it and estimate the intra-daily return power transformation that optimally predicts future increments in quadratic variation. We find that the longer the prediction horizon, the smaller the optimal power transformation.

1. Introduction

The prediction of stock market volatility is a topic that demands attention from both academics and practitioners. Asset allocation, risk management, pricing of assets, are only a few examples where volatility modeling plays a key role. Not surprisingly, a lot has been written about the topic. The seminal work by Engle on ARCH-type models has generated a fertile research area over the last two decades, and many financial institutions trade billions of dollars on a daily basis – where trading is based one way or another on volatility predictions. This applies both at the microstructure level – where limit order book dynamics are volatility driven – as well as at the macro level, where aggregate economic variables affect the long run predictions of market fluctuations (see e.g. Engle et al., 2006). The current paper brings together insights from various existing results: (1) the original findings regarding persistence in power-transformed returns noted by Ding et al. (1993), (2) the extensive literature on using high frequency financial data to predict volatility using data-driven measures, and (3) the Mixed Data Sampling (henceforth MIDAS) regression models used by Ghysels et al. (2006) and Forsberg and Ghysels (2007) to formulate regression-based prediction models.

Ding et al. (1993) examine autocorrelations of power transformations of daily returns (i.e. $|r_t|^\delta$) and find that they are the strongest when δ is around 1. More precisely, Ding et al. (1993) use S&P500 stock market daily closing price index for a period from Jan 3, 1928 to Aug 30, 1991. On the other hand, Ding and Granger (1996) use daily foreign exchange rate returns for the DM/$ from Jan 1971 to March 1992 and they find that the power transformation of the absolute return series with $\delta=\frac{1}{4}$ appear to have strong and persistent autocorrelation. These findings prompted them, and others, to consider ARCH-type models based on absolute returns, instead of squared returns.

The drawback of ARCH-type models based on power transformations is that they predict future $|r_{t+i}|^\delta$ for some $i>0$ and $\delta<2$, based on past $|r_{t-j}|^\delta$ with $j>0$ and the *same* δ. Often we are not interested in $|r_{t+i}|^\delta$ per se, but rather future $|r_{t+i}|^2$, that is volatility measured

☆ We are grateful to two referees and the editor, Franz Palm for their invaluable comments and help.
* Corresponding author. Department of Economics, University of North Carolina at Chapel Hill, Gardner Hall CB 3305, Chapel Hill NC 27599, United States.
E-mail addresses: eghysels@email.unc.edu (E. Ghysels), sohnb@kenan-flagler.unc.edu (B. Sohn).

0927-5398/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jempfin.2009.03.002

with squared returns. The question is whether we can use past $|r_{t-j}|^{\delta}$ with $\delta \neq 2$ to predict future volatility. In the current data-rich environment of high frequency financial series the question can be rephrased as follows: can we predict future increments in *quadratic* variation – measured as the sum of intra-daily squared returns, *i.e.* realized variance (henceforth RV) – with past realized power variation – where the latter are say daily sums of *δ-powered absolute intra-daily returns* (see Barndorff-Nielsen and Shephard (2003)) for discussion of power variation measures). The purpose of this paper is to examine whether there are indeed gains – in terms of volatility forecasting – to be made by considering past power variations where one *estimates* the parameter δ as part of the prediction problem.

The topic of the paper is most directly related to the work of Forsberg and Ghysels (2007) who study the prediction problem of future realized volatility using past absolute power variation, i.e. $\delta = 1$. Hence, Forsberg and Ghysels (2007) do not estimate which power variation to use. In a sense, what is being investigated in the paper are the gains from letting the data decide what power variation is optimal in the context of a MIDAS regression prediction of future increments in quadratic variation. There are costs and benefits to the analysis pursued here. For example, if the optimal value is close to one, then the approach in Forsberg and Ghysels (2007) may still dominate as estimation uncertainty may undo all gains from considering the explicit estimation of δ. On the other hand, if δ is indeed far away from one, then despite sampling error, we may be better off to estimate the optimal power – where optimal means yielding the best predictability in a least squares sense. It is the purpose of the paper to examine these issues.

The paper is organized as follows. Section 2 introduces various volatility measures of our interest and their asymptotic properties. Section 3 investigates various correlation structure between *RV* and the measures introduced in Section 3 to motivate the empirical model building in Section 4. Section 4 introduces MIDAS regression framework, presents estimation results, and comparison of forecasting performance among the models. Section 5 concludes the paper.

2. Volatility measures

We start by specifying a returns process and assume that the true underlying log-price process p_t follows a Brownian semimartingale:

$$p_t = p_0 + \int_0^t a_u \mathrm{d}u + \int_0^t \sigma_{u-} \mathrm{d}W_u \tag{2.1}$$

where W is a standard Brownian motion, a is a predictable and has locally bounded sample paths, and the spot volatility has cádlág sample paths. Both a and σ can have jumps, intraday seasonality and long memory. We define the (discrete) daily log return as

$$r_{t,t-1} = lnP_t - lnP_{t-1} = p_t - p_{t-1} \tag{2.2}$$

where the t refers to daily sampling (henceforth we will refer to the time index t as daily sampling). The intra-daily return is then denoted

$$r_{t,j}^M = p_{t-j/M} - p_{t-(j-1)/M} \tag{2.3}$$

where $1/M$ is the (intra-daily) sampling frequency. For example, when dealing with typical stock market data we will use $M = 78$ corresponding to a five-minute sampling frequency. It is possible to consistently estimate quadratic variation over some period $[t,t+H]$ by summing squared intra-daily returns, yielding the so called *realized variance*, namely:

$$RV_{t,t+H}^M = \sum_{j=1}^{MH} \left(r_{t,j}^M\right)^2 \tag{2.4}$$

When the sampling frequency increases, *i.e.* $M \to \infty$, then the realized variance converges to the quadratic variation. More specifically,

$$RV_{t,t+H}^M \to \int_t^{t+H} \sigma_u^2 \mathrm{d}u \tag{2.5}$$

where the convergence is in probability, locally uniform in time.

The second and third measures of volatility are based on sums of powers and product of powers of absolute returns. The former is called *realized* δ-th *order power variation* and is defined as

$$RPV_{t,t+H}^{[\delta]} = M^{-1+\delta/2} \sum_{j=1}^{MH} |r_{t,j}^M|^{\delta} \tag{2.6}$$

and the latter is called *realized* δ_1,δ_2-th *order bipower variation* and is defined as

$$BPV_{t,t+H}^{[\delta_1,\delta_2]} = M^{-1+(\delta_1+\delta_2)/2} \sum_{j=1}^{MH} |r_{t,j}^M|^{\delta_1} |r_{t,j+1}^M|^{\delta_2} \tag{2.7}$$

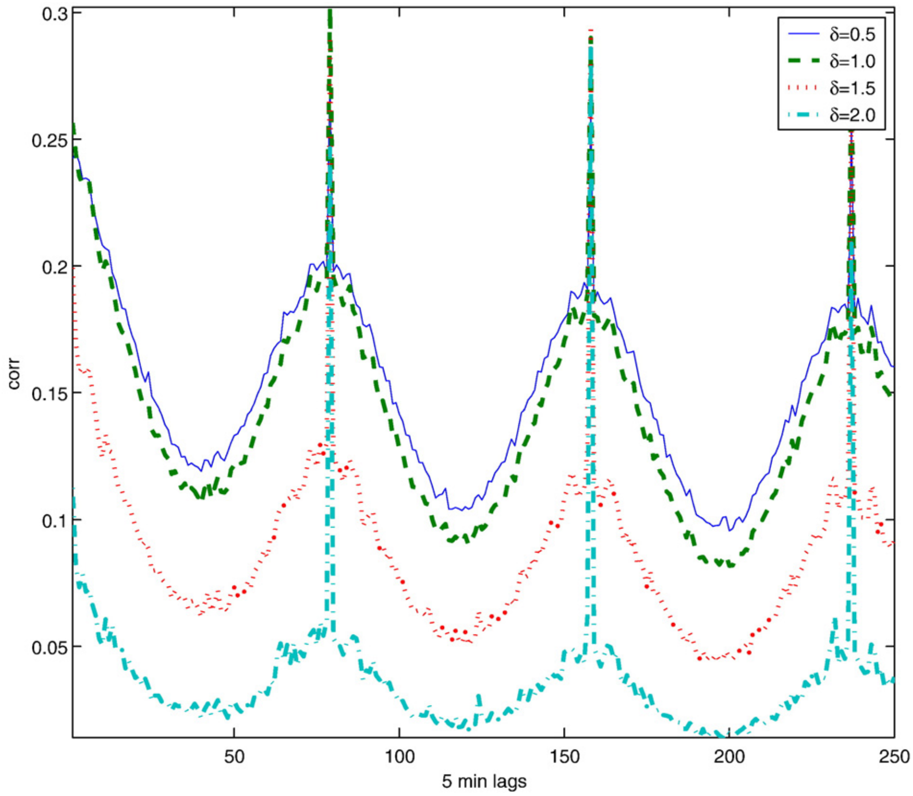

Fig. 1. Autocorrelation of $|r_t^{(5\ min)}|^{\delta}$. Our dataset consists of five-minute intra-daily returns of the Dow Jones Composite (DJ) over ten year period from April 1, 1993 to October 31, 2003. From this dataset, we compute autocorrelation function of power-transformed absolute (five-minute) return series. Note that there are 79 five-minute return data in a day.

Barndorff-Nielsen et al. (2005) derive the probability limit of these measures when the underlying return process follows Eq. (2.1):

$$RPV_{t,t+H}^{[\delta]} \rightarrow \mu_{\delta} \int_t^{t+H} \sigma_u^{\delta} \mathrm{d}u \tag{2.8}$$

$$BPV_{t,t+H}^{[\delta_1,\delta_2]} \rightarrow \mu_{\delta_1}\mu_{\delta_2} \int_t^{t+H} \sigma_u^{\delta_1+\delta_2} \mathrm{d}u \tag{2.9}$$

where $\mu_{\delta} = E(|u|^{\delta})$ and $u \sim N(0,1)$. Barndorff-Nielsen and Shephard (2004) introduce an interesting feature of these volatility measures: They are robust to rare jumps. Using this property, they propose a method to separate *QV* into its continuous component and jump component which is estimated by the difference between *RV* and *BPV*. The theoretical advantages where used in various empirical settings. For example, Andersen et al. (2006) find that almost all of the predictability in daily, weekly, and monthly return volatilities comes from the continuous component. Ghysels et al. (2006) and Forsberg and Ghysels (2007) show that regressors involving *RPV* with δ fixed at 1 are better at predicting future volatility and this for a variety of reasons: (1) desirable population predictability features, (2) better sampling error behavior and (3) immunity to jumps. The major difference between this paper and the previous work is that we take an empirical approach to the determination of the power variation.[1]

3. Empirical correlation analysis

The dataset we use consists of five-minute intraday returns of the Dow Jones Composite (DJ) over ten year period, from April 1, 1993 to October 31, 2003. We follow a data filter procedure advocated by Andersen et al. (2001) to clean the original raw data. The final dataset contains 2669 trading days with 79 observations per day for a total of 210,851 five-minute returns. From this dataset, we compute daily returns, daily realized variance, daily power variations, and so on.

[1] In order to derive central limit theorems for these volatility measures, we need to impose some stronger assumptions both on the return process and the functional form of return transformation. As was mentioned before, Barndorff-Nielsen et al. (2005) outline the necessary assumptions and proofs for deriving the asymptotic distributions of these measures. We assume in our empirical work that these conditions apply.

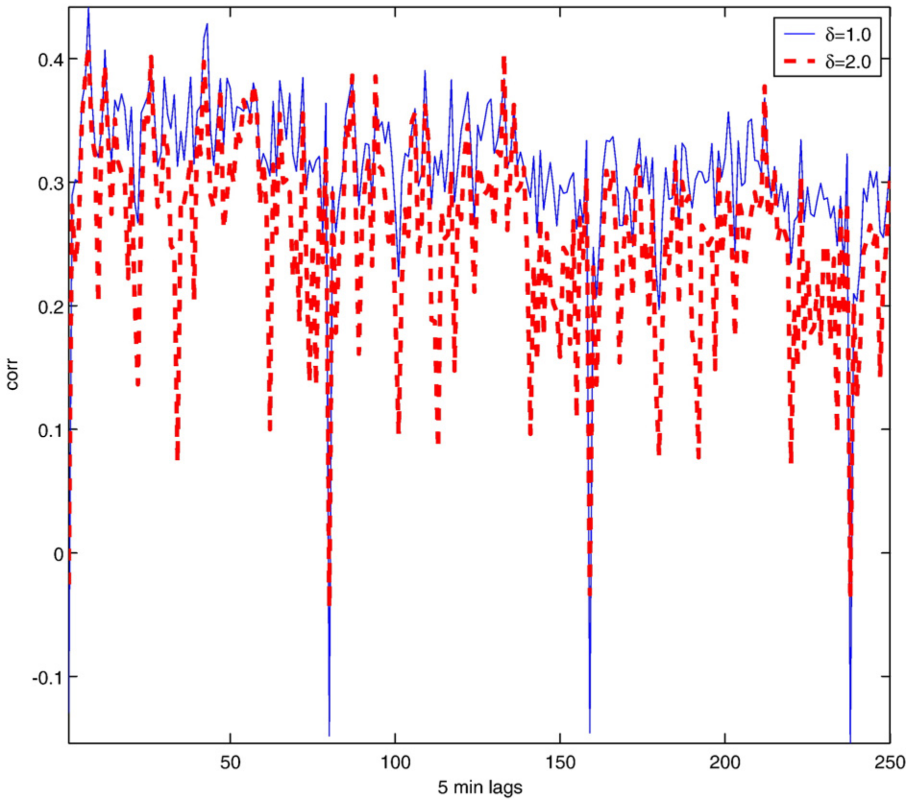

Fig. 2. Cross-correlation between daily $RV_{t,t+1}$ and lags of $|r_t^{(5\ min)}|^{\delta}$. Our dataset consists of five-minute intra-daily returns of the Dow Jones Composite (DJ) over ten year period from April 1, 1993 to October 31, 2003. From this dataset, daily $RV_{t,t+1}$ series are computed as defined in Eq. (2.4) and the cross-correlation between daily $RV_{t,t+1}$ and power-transformed absolute five-minute lagged returns are calculated as follows;

$$Corr(RV^M_{t,t+1}, |r^M_{t,i}|^{\delta}) \quad (5.2)$$

where $r^M_{t,i}$ is defined in Eq. (2.3) and i denotes the lag number. Note that there are 79 five-minute returns in a day.

What matters most for our analysis is the cross-correlation between RV and lagged RPV's or RV and lagged BPV's. Since the autoregressive models have been a tradition in volatility modeling, autocorrelation structures of returns and return volatility measures have been intensively studied both from a theoretical and empirical perspective. However, far less attention has been paid to the cross-correlation structure between target RV and other volatility measures such as RPV and BPV. Forsberg and Ghysels (2007) provide a theoretical univariate regression analysis to show that RPV with δ fixed 1 is a better predictor than RV in population. That is, $\int_{t-1}^{t}\sigma_s ds$ does a better job at predicting $\int_t^{t+1}\sigma_s^2 ds$ than a lag of QV itself, *i.e.* $\int_{t-1}^{t}\sigma_s^2 ds$. They show this by deriving theoretically the cross-correlation function under the assumption that the σ_t process follows a non-Gaussian Ornstein–Uhlenbeck process.

Since our dataset consists of 5 min returns, we investigate cross-correlation between daily RV and volatility measures of various frequency. That is, we will examine the cross-correlation between daily RV and lags of (1) power transformation of absolute 5 min returns, *i.e.* $|r^M_{t,j}|^{\delta}$ (2) power transformation of absolute daily returns, *i.e.* $|r_{t,t-1}|^{\delta}$ and (3) daily realized δ-th order power variations, *i.e.* $RPV^{[\delta]}_{t,t+1}$. The forecasting horizon is one to four weeks, which means that the regressand will be $RV_{t,t+H}$ with $H \in \{5,10,15,20\}$.[2] However, since the multi-horizon RV's are the summation of short-horizon RV's, it is enough to investigate the correlation of daily RV with each of the volatility measures.

Figs. 1 and 2 show the autocorrelation and cross-correlation structure of power transformations of 5 min absolute returns, respectively. Fig. 1 clearly shows repetitive U-shaped daily pattern, which is also often reported in high frequency data in FX market. Although it appears much weaker, Fig. 2 also shows the repetitive daily pattern. This is not surprising. In fact, intraday seasonal patterns in financial markets have been widely documented in the literature since the work of Wood et al. (1985). Various observations are worth noting. First, one should note that overall level of cross-correlation is quite low and the structure is not smooth. Second, the series with $\delta=1$ shows better behavior when compared with the one with $\delta=2$, by which we mean that cross-correlation of $|r^M_{t,j}|$ with RV is less bumpy and overall higher than that of $|r^M_{t,j}|^2$. Actually, these two were used in Ghysels et al. (2006) as regressors

[2] Although we present weekly RV as $RPV_{t,t+5}$, it is not always a sum of five daily RV's. To eliminate any seasonality in day of a week, we aggregate RV's of a calendar week to construct a weekly RV. Other RV's over different horizon are defined similarly.

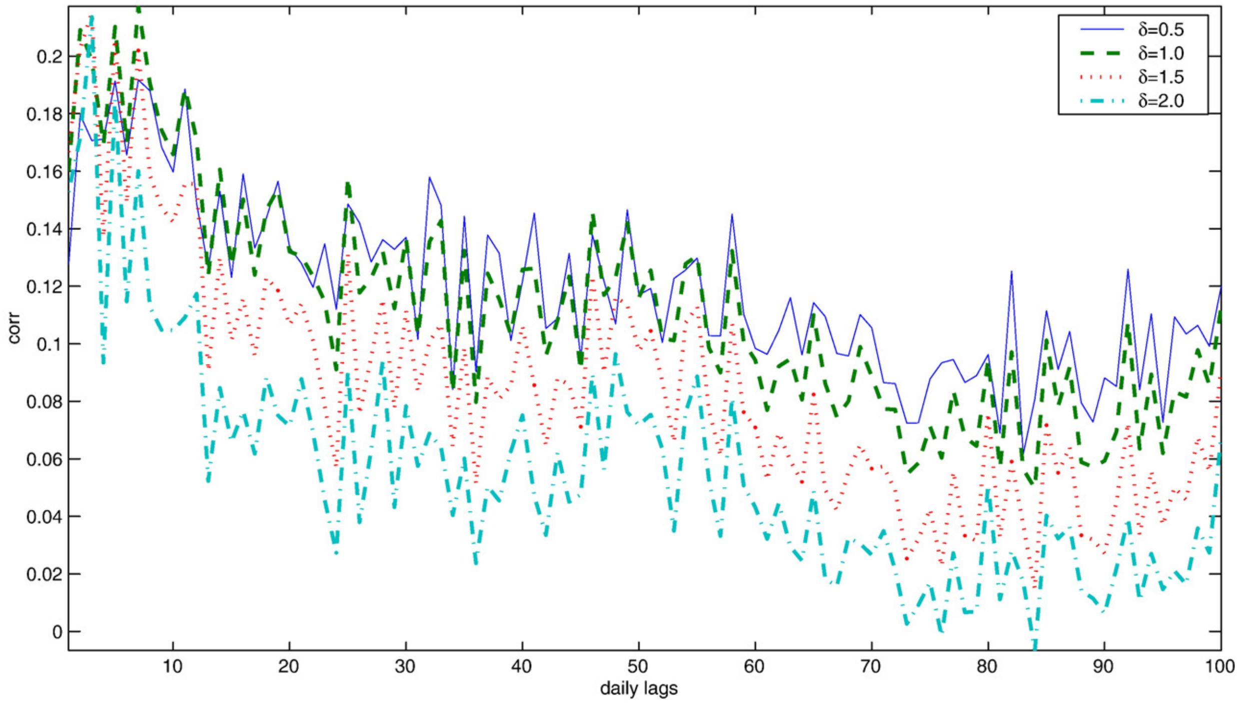

Fig. 3. Autocorrelation of $|r_t^{(day)}|^{\delta}$. Our dataset consists of five-minute intra-daily returns of the Dow Jones Composite (DJ) over ten year period from April 1, 1993 to October 31, 2003. From this dataset, we construct daily return series and compute autocorrelation function of power-transformed absolute (daily) return series.

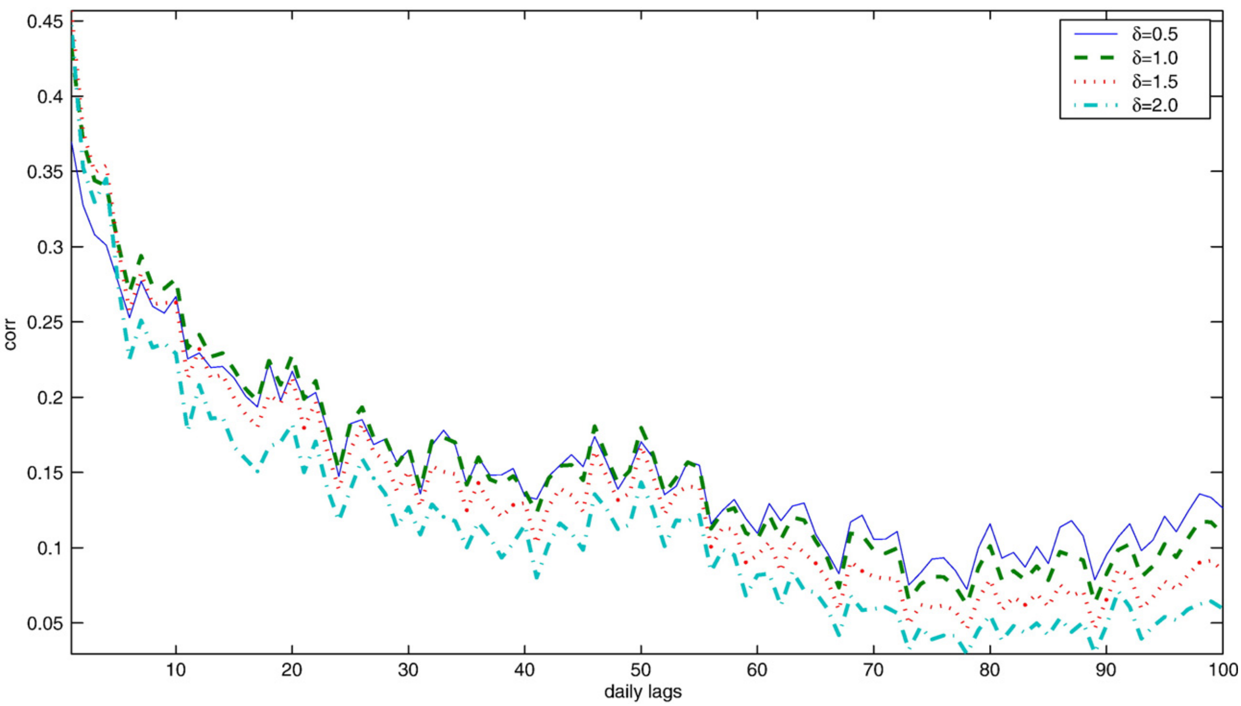

Fig. 4. Cross-correlation between daily $RV_{t,t+1}$ and lags of $|r_t^{(day)}|^{\delta}$. Our dataset consists of five-minute intra-daily returns of the Dow Jones Composite (DJ) over ten year period from April 1, 1993 to October 31, 2003. From this dataset, daily $RV_{t,t+1}$ series are computed as defined in Eq. (2.4) and the cross-correlation between daily $RV_{t,t+1}$ and power-transformed absolute daily lagged returns are calculated as follows;

$$Corr(RV^M_{t,t+1}, |r_{t-i,t-i-1}|^{\delta}) \tag{5.3}$$

where $r_{t,t-1}$ is defined in Eq. (2.2) and i denotes the lag number.

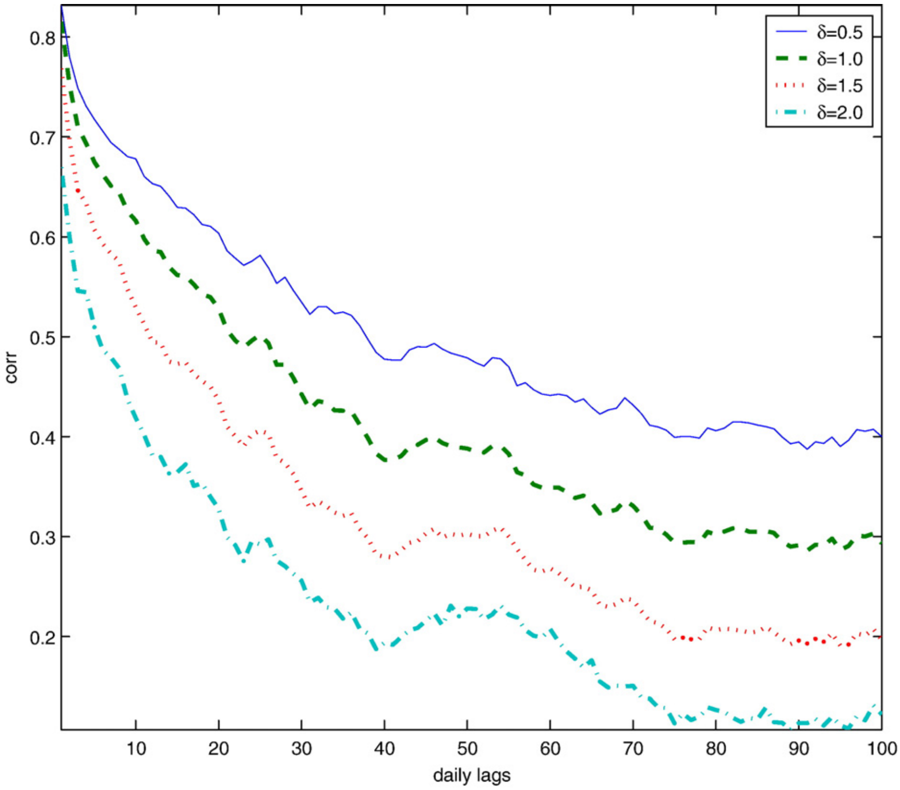

Fig. 5. autocorrelation of daily $RPV^{[\delta]}_{t,t+1}$. Our dataset consists of five-minute intra-daily returns of the Dow Jones Composite (DJ) over ten year period from April 1, 1993 to October 31, 2003. From this dataset, we compute (daily) $RPV^{[\delta]}_{t,t+1}$ and autocorrelation function of this volatility measure.

and $|r^M_{t,j}|$ performed better, indeed, as a predictor. However, as it is already pointed out, overall cross-correlation level is too low and that resulted in no gain in direct modeling of high frequency data.

Figs. 3 and 4 present the autocorrelation and cross-correlation structure of power transformation of daily absolute returns. Autocorrelation of (daily) $|r_{t,t-1}|^\delta$ was extensively studied in Ding et al. (1993). They look into daily return series of S&P500 for a long period from 1928 to 1991 and found that $|r_{t,t-1}|^\delta$ series with $\delta \approx 1$ features the strongest autocorrelation. Overall, the autocorrelation structure shown in Fig. 3 is much weaker than that of Ding et al. (1993). Also, except for the first 10 lags, it's hard to tell whether the series with $\delta = 1$ or $\delta = 0.5$ has the strongest autocorrelation structure. Surprisingly, the cross-correlation structure in Fig. 4 features nice properties; the first 10 lags of power-transformed (daily) absolute returns are strongly correlated with RV, and the cross-correlation decays quite smoothly over the remaining lags. An interesting observation is that $|r_{t,t-1}|^2$ appears to have the weakest cross-correlation with RV throughout. This is also true with 5 min return series as one can verify from Fig. 2.

Lastly, Figs. 5 and 6 plot autocorrelation and cross-correlation of $RPV^{[\delta]}_{t,t+1}$. From Fig. 5 one can clearly see that $RPV^{[0.5]}_{t,t+1}$ shows the strongest autocorrelation from lag 1 to lag 100 whereas $RPV^{[2]}_{t,t+1}$ the weakest. Fig. 6 presents the similar pattern except for the first 10 lags. The cross-correlation for the first 10 lags are also presented in Table 1. Although the rank for the strength of cross-correlation changes considerably over the first 10 lags, $RPV^{[2]}_{t,t+1}$ has the weakest cross-correlation throughout without any exception and $RPV^{[1]}_{t,t+1}$ shows the strongest with a couple of exceptions. However, beyond the 10th lag, $RPV^{[0.5]}_{t,t+1}$ features the strongest cross-correlation with $RPV_{t,t+1}$. This result is rather surprising since $RPV^{[2]}_{t,t+1}$ is, in fact, $RV_{t,t+1}$. This is consistent with the empirical results of Ghysels et al. (2006) where lagged $RPV^{[1]}_{t,t+1}$'s are shown to do a better job at predicting RV's than lagged $RV_{t,t+1}$'s.

The findings that the cross-correlations are almost constant across lags and that the autocorrelations very slowly die out raise the question whether this finding is a true feature of the high-frequency data or whether it results from structural changes over the sample period considered. Indeed, several authors have tested for breaks in volatility (see e.g. Andreou and Ghysels (2002)) and the presence of breaks has been suggested as resulting in long memory (see e.g. Diebold and Inoue (2001), Gourieroux and Jasiak (2001) and Granger and Hyung (2004)).[3] We will not elaborate on this as it would take us far beyond the topic of the present paper.

We started the investigation of cross-correlation between daily RV and volatility measures of various frequency in search of optimal candidates for regressors in MIDAS regression models for RV. Comparing the three correlograms of Figs. 2, 4 and 6, it is quite obvious that the volatility measure that aggregates high frequency information as sum of power transformation of absolute returns will perform well in predicting future RV. We did not look into the case with BPV, but as the asymptotic results in Eqs. (2.8)

[3] For recent surveys, see also Andreou and Ghysels (2009) and Banerjee and Urga (2005).

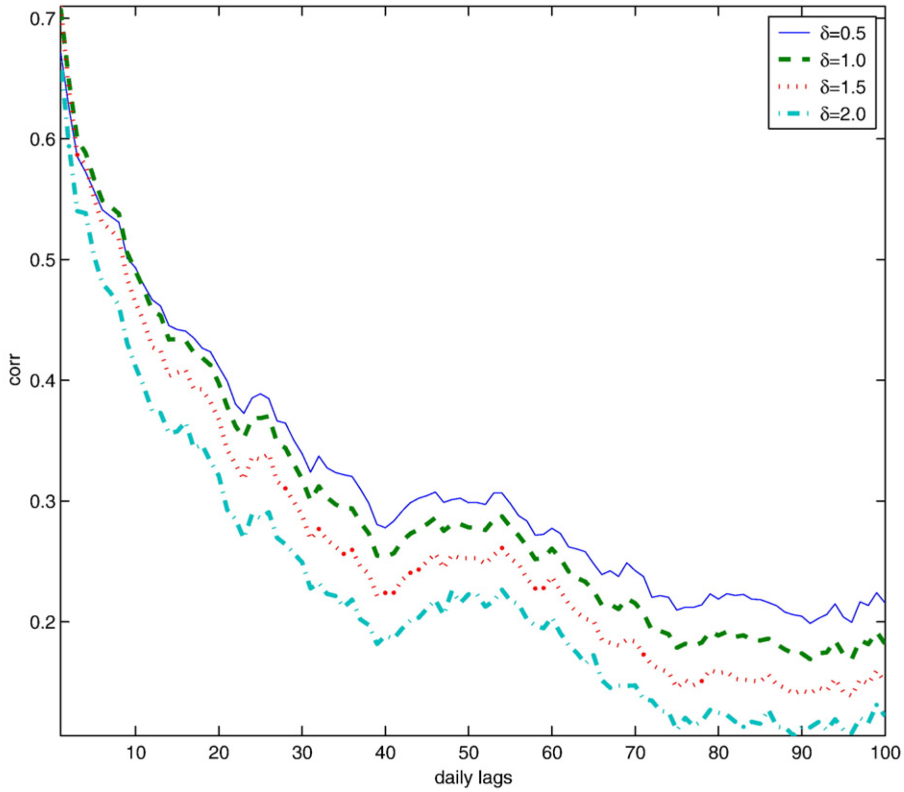

Fig. 6. Cross-correlation between daily $RPV_{t,t+1}$ and lags of daily $RPV^{[\delta]}_{t,t+1}$. Our dataset consists of five-minute intra-daily returns of the Dow Jones Composite (DJ) over ten year period from April 1, 1993 to October 31, 2003. From this dataset, (daily) $RV_{t,t+1}$ series and (daily) $RV^{[\delta]}_{t,t+1}$ series are computed as defined in Eqs. (2.4) and (2.6). Then, the cross-correlation between these two volatility measures are calculated as follows;

$$Corr(RV_{t,t+1}, RPV^{[\delta]}_{t-i,t-i+1}) \tag{5.4}$$

where i denotes the lag number.

and (2.9) suggest, BPV encompasses what RPV is measuring in the limit although the finite sample behavior might be quite different. Consequently, we will use RPV and BPV as regressors for MIDAS regression model for target RV in the next section.

4. MIDAS models of conditional variance

In this section, we will briefly introduce the MIDAS regression models of conditional variance using the regressors examined in the previous section. Comparing both in-sample and out-of-sample forecasting performance, we will examine whether there is

Table 1
Cross-correlation between daily $RV_{t,t+1}$ and lags of daily $RPV^{[\delta]}_{t,t+1}$.

δ	Lag 1		Lag 2		Lag 3		Lag 4		Lag 5		Lag 6		Lag 7		Lag 8		Lag 9		Lag 10	
	XCorr	Rank	XCorr	Rank	XCorr	Rank	XCorr	Rank	XCorr	Rank	XCorr	Rank	XCorr	Rank	XCorr	Rank	XCorr	Rank	XCorr	Rank
0.5	0.6719	3	0.6253	3	0.5853	3	0.5725	3	0.5572	2	0.5413	2	0.5359	2	0.5309	2	0.5011	2	0.4933	1
1.0	0.7080	2	0.6475	1	0.5995	1	0.5884	1	0.5671	1	0.5490	1	0.5440	1	0.5381	1	0.5058	1	0.4906	2
1.5	0.7101	1	0.6398	2	0.5868	2	0.5786	2	0.5510	3	0.5303	3	0.5249	3	0.5171	3	0.4842	3	0.4643	3
2.0	0.6644	4	0.5939	4	0.5402	4	0.5385	4	0.5033	4	0.4805	4	0.4726	4	0.4618	4	0.4307	4	0.4115	4

Our dataset consists of five-minute intraday returns of the Dow Jones Composite (DJ) over a ten year period from April 1, 1993 to October 31, 2003. We compute daily $RV_{t,t+1}$ and $RV^{[\delta]}_{t,t+1}$ as defined in Eqs. (2.4) and (2.6) for various values of δ. XCorr refers to the cross-correlation between daily $RV_{t,t+1}$ and lagged daily $RV^{[\delta]}_{t,t+1}$;

$$XCorr = Corr\left(RV_{t,t+1}, RPV^{[\delta]}_{t-k,t+1-k}\right)$$

where k is the lag number. The strength of cross-correlation for a given lag number is ranked across the choice of δ.

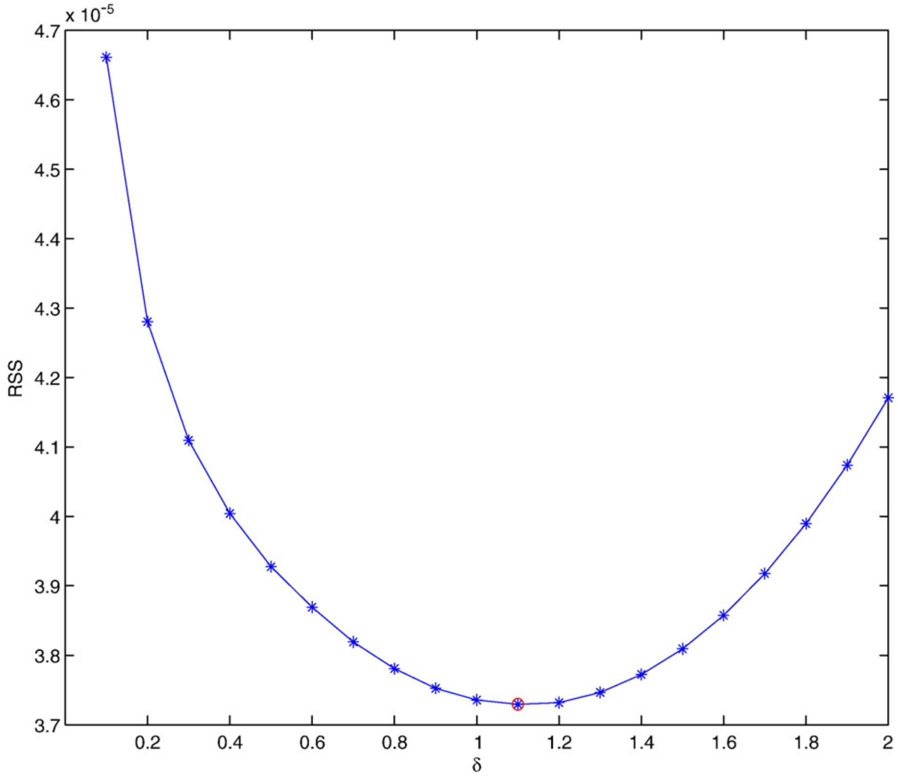

Fig. 7. Least RSS over choice of δ when $H=10$. In search of appropriate starting value for δ in estimation of MIDAS-RV model with $RPV^{[\delta]}_{t,t+1}$ regressors, we profile *RSS* from the estimation with δ fixed at a set of specific values. The figure shows the results when the forecast horizon is two weeks; that is, the target is two-week *RV*.

any gain in the prediction of *RV* with *RPV* and *BPV* when we set power transformation parameters free; we let the data choose the optimal value of δ in $RPV^{[\delta]}$, and δ_1 and δ_2 in $BPV^{[\delta_1,\delta_2]}$.

4.1. MIDAS regression models

Mixed Data Sampling (MIDAS) regression models have recently been introduced by Ghysels et al. (2002). With MIDAS regressions it is possible to use data sampled at different frequencies for the dependent and independent variables. This particular property of MIDAS regression is especially relevant since we want to model the realized variance over a range of horizons and we will use variables at the daily frequency as regressors.

The MIDAS-RV model is the standard MIDAS model with RV as dependent variable and where the $RV_{t,t+H}$ refers to the fact that we model the increment of RV from t to $t+H$ where H is the prediction horizon in days. Throughout the empirical analysis, we refer to the multi-period realized variances for $H=5$,10,15 and 20 as weekly, bi-weekly, tri-weekly and monthly realized variances.[4] The model can be written as:

$$RV_{t,t+H}=\mu+\phi\sum_{k=1}^{k_{max}}w_kX_{t-k,t+1-k}+\varepsilon_t \tag{4.1}$$

where the weights w_k are chosen by some function and sum up to one. $X_{t-k,t+1-k}$ denotes the regressor from $(t-k)$ to $(t+1-k)$. Specifically, as was discussed in section 3, we will adopt $RPV^{[\delta]}_{t,t+1}$ and $BPV^{[\delta_1,\delta_2]}_{t,t+1}$ as regressors. Throughout this paper we use $k_{\max}=50$.[5] For the weights w_k in Eq. (4.1) there are several functions available in the literature. Here, we use the lag polynomial structure introduced in Ghysels et al. (2002) and use a specification based on the Beta function, that is

$$w(k;\theta)=\frac{(k/k_{max})^{\theta_1-1}(1-k/k_{max})^{\theta_2-1}}{\sum_{j=1}^{k_{max}}(j/k_{max})^{\theta_1-1}(1-j/k_{max})^{\theta_2-1}} \tag{4.2}$$

[4] Usually in the literature a trading month is defined to be 22 days, in which case we would use $H=22$ for the month. However, we use $H=20$ to define four weeks of prediction horizon, and we will refer to this horizon as the four weeks horizon and monthly horizon interchangeably. Also, note that these *RV*'s are non-overlapping.

[5] We experimented with other lag lengths, yielding essentially the same results.

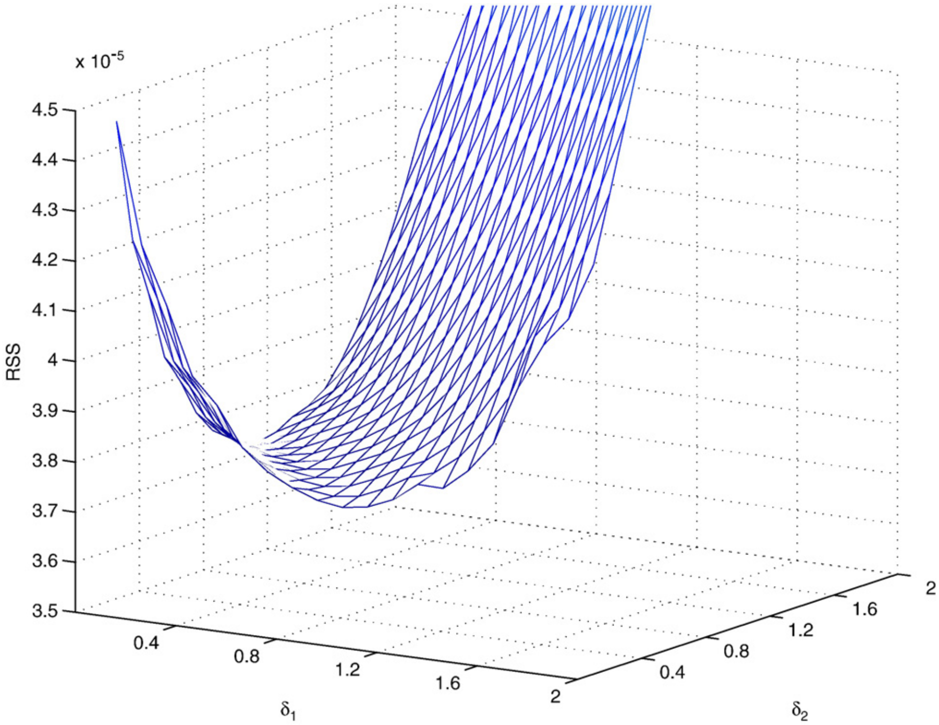

Fig. 8. Least RSS over choice of δ_1 and δ_2 when H = 10. In search of appropriate starting value for δ_1 and δ_2 in estimation of MIDAS-RV model with $BPV_{t,t+1}^{[\delta_1,\delta_2]}$ regressors, we profile *RSS* from the estimation with $\{\delta_1,\delta_2\}$ fixed at a set of specific values. The figure shows the results when the forecast horizon is two weeks; that is, the target is two-week *RV*.

Note that this weighting scheme provides positive weights, which we need for the positivity of the estimated variance. In summary, we consider following specifications:

$$RV_{t,t+H} = \mu + \phi \sum_{k=1}^{50} w_k(\theta_1,\theta_2) RPV_{t-k,t+1-k}^{[\delta]} + \varepsilon_t \tag{4.3}$$

$$RV_{t,t+H} = \mu + \phi \sum_{k=1}^{50} w_k(\theta_1,\theta_2) RPV_{t-k,t+1-k}^{[\delta_1-\delta_2]} + \varepsilon_t \tag{4.4}$$

where $H \in \{5,10,15,20\}$. We fix θ_1 to be 1 in the beta-polynomials.[6] Initially, we estimated the models with both of free parameters, θ_1 and θ_2, but it turns out that $\hat{\theta}_1$ is insignificant and all the optimal weighting schemes are monotonically decaying over the lags. Hence, the parameters to be estimated are μ, ϕ, θ_1, and δ, or δ_1, and δ_2. It should be pointed out that values of power transformation parameters (i.e. δ, δ_1, and δ_2) do vary upon a choice of forecasting horizon H, and it certainly is a important focus of this paper to examine the relationship between the optimal value of δ and forecasting horizon *H*.[7]

4.2. Estimation and empirical results

We split the sample into two parts; in-sample from April 1, 1993 to December 31, 2000 and out-of-sample from January 1, 2001 to October 31, 2003 to check for overfitting of the model. To fit the models specified in Eqs. (4.3) and (4.4) over the in-sample period, we use Non-Linear Least Squares (henceforth NLLS) methodology. For each of forecasting horizons *H*, we choose the

[6] This will give us declining weights in the lag polynomial. For further details about the beta polynomials and their shape and form, see Ghysels et al. (2002).

[7] Another interesting specification motivated by the findings in Section 3 is the following:

$$RV_{t,t+H} = \mu + \phi_1 \sum_{k=1}^{10} w_k^{(1)}(\theta_1,\theta_2) RPV_{t-k,t+1-k}^{[\delta_1]} + \phi_2 \sum_{k=11}^{50} w_k^{(2)}(\theta_3,\theta_4) RPV_{t-k,t+1-k}^{[\delta_2]} + \varepsilon_t$$

Table 1 and Fig. 6 suggest that cross-correlation between the current peiord *RV* and lags of *RPV* $^{[\delta]}$ has a different structure depending on the lag order; for the first 10 lags, the highest correlation obtains for $\delta = 1$ while, for the lags higher than 10, the highest correlation obtains for $\delta = 0.5$. The above specification is designed to exploit this feature concerning the relationship between the regressand and the regressors. It turns out that the above specification suffers overfitting problem. Over the in-sample period, the above MIDAS-RV model with two separately parameterized *RPV*'s outperforms the corresponding simple specification (4.3). However, the reverse relation holds true for out-of-sample period except for case with the shortest forecast horizon (*i.e.* a week) examined where two *RSS*'s come very close.

Table 2
Parameter estimates and relative forecasting performance of MIDAS-RV model with $RPV^{[\delta]}$ regressors.

H	μ	ϕ	θ_2	δ	In-sample				Out-of-sample			
					R^2	RSS (10^{-5})	$\frac{RSS}{RSS_2}$	$\frac{RSS}{RSS_1}$	R^2	RSS (10^{-5})	$\frac{RSS}{RSS_2}$	$\frac{RSS}{RSS_1}$
5	−0.00011	0.07370	29.80089	1.30715	0.60	2.12658	0.928	0.983	0.52	3.79769	0.952	0.949
	(−1.78)	(1.21)	(6.56)	(9.13)								
10	−0.00036	0.04501	18.32377	1.10739	0.59	3.72805	0.894	0.998	0.53	5.99421	0.777	0.984
	(−1.49)	(0.81)	(4.31)	(5.21)								
15	−0.00111	0.01153	8.21744	0.80862	0.54	5.70232	0.952	0.997	0.41	10.97600	0.929	1.002
	(−1.22)	(0.55)	(2.75)	(2.55)								
20	−0.00231	0.00477	6.45418	0.60765	0.57	6.61876	0.852	0.984	0.41	12.92912	0.978	1.038
	(−1.09)	(0.54)	(2.43)	(1.82)								

MIDAS-RV model with $RPV^{[\delta]}$ regressors as defined in Eq. (4.3) is fitted over in-sample dataset (April 1, 1993–December 31, 2000) by NLLS methodology.

$$RV_{t,t+H} = \mu + \phi \sum_{k=1}^{50} w_k(\theta_1, \theta_2)\mathrm{RPV}^{[\delta]}_{t-k,t+1-k} + \varepsilon_t.$$

Then, its out-of-sample forecast performance is evaluated over the period from January 1, 2001 to October 31, 2003. H refers to forecast horizon in days. θ_1 is fixed at 1. Various RSS's are computed for comparison. RSS is the residual sum of squares from fitted MIDAS-RV model with $RPV^{[\delta]}$ regressors. For comparison, RSS_1 and RSS_2 are also computed from the fitted MIDAS-RV model with $RPV^{[1]}$ (*i.e. RAV*) and $RPV^{[2]}$ (*i.e. RV*) as regressors, respectively. R^2 is calculated as $1-RSS/TSS$. Finally, the numbers in the parenthesis are the t-stats computed with HAC standard errors.

regressand $RV_{t,t+H}$ to be non-overlapping, as overlapping $RV_{t,t+H}$'s lead to a persistent regressand and might also lead to spurious results in forecastability of $RV_{t,t+H}$'s. For a given forecast horizon H and a choice of model specification, out-of-sample R^2 and RSS are calculated with the in-sample parameter estimates.

Our models of stock market volatility as specified in Eqs. (4.3)–(4.4) are different from a class of standard ARCH/GARCH models in the sense that they don't have autoregressive structures. To ensure legitimacy of our MIDAS-RV models as a class of forecasting models for stock market volatility, we refer to Ghysels et al. (2006). They estimate similar MIDAS-RV models with $RPV^{[\delta]}$ where δ is fixed at 1 or 2. The same dataset is used to fit the models although the in-sample period differs slightly. The benchmark model taken in Ghysels et al. (2006) is ARFI(5,d) model of Andersen et al. (2003). They show that the MIDAS-RV model perform far better than the benchmark model in forecasting future RV over both in- and out-of-sample, and for all forecasting horizons.

To ensure the optimality of estimation for the nonlinear parameters, especially δ in $RPV^{[\delta]}$, and δ_1 and δ_2 in $BPV^{[\delta_1,\delta_2]}$, we resorted *first* to profiling before estimation. Namely, we first fix the power transformation parameters in Eqs. (4.3) and (4.4) at specific values and estimate the other parameters. For the model specified in Eqs. (4.3), we start estimation with δ fixed at 0.1, and then increase the value by 0.1 until the δ reaches 2.0. A set of RSS from fitted MIDAS-(2 week)RV is shown in Fig. 7. It turns out that $\delta=1.1$ yields the least RSS as is marked with the circle in the figure. This exercise is repeated over different forecasting horizons, and all of them result in similar U-shaped RSS – δ graphs. However, the δ's that allow the least RSS are different depending on the length of forecasting horizons. The combinations of forecasting horizon, H, and the corresponding optimal $\check{\delta}$ are $H=5$, $\check{\delta}=1.3$ / $H=10$, $\check{\delta}=1.1$ / $H=15$, $\check{\delta}=0.8$ / $H=20$, $\check{\delta}=0.6$. For a given forecasting horizon, we use these values of $\check{\delta}$'s for the starting values of the model estimation. The same exercise is done for the estimation of the MIDAS-RV model specified in Eq. (4.4). The resulting graph of RSS over the grid of fixed δ_1 and δ_2 when $H=10$ is shown in Fig. 8. The combinations of forecasting horizon and the corresponding optimal $\check{\delta}_1$ and $\check{\delta}_2$ are $H=5$, $\check{\delta}_1=0.7$, $\check{\delta}_2=0.7$ / $H=10$, $\check{\delta}_1=1.0$, $\check{\delta}_2=0.1$ / $H=15$, $\check{\delta}_1=0.1$, $\check{\delta}_2=0.8$/$H=20$, $\check{\delta}_1=0.1$, $\check{\delta}_2=0.5$. Again, these values are used for the starting values when we estimate the model in Eq. (4.4). There are two interesting observations in these results. First, in Fig. 8, one can see that the RSS graph looks like an upside down saddle, which means that there are a set of δ_1 and δ_2 values that yield low values of RSS. Interestingly enough, these values are the combinations of δ_1 and δ_2 such that

$$\delta_1 + \delta_2 = \check{\delta}_1 + \check{\delta}_2 \tag{4.5}$$

for given H. The second interesting feature of the results is that, for given H,

$$\check{\delta} \approx \check{\delta}_1 + \check{\delta}_2 \tag{4.6}$$

These interesting features are due to the asymptotic results introduced earlier in Eq. (2.8) and Eq. (2.9).[8] In other words, if the relation stated in Eq. (4.6) holds exactly, then, for an arbitrary time interval $[t,t+1]$, $RPV^{[\check{\delta}]}_{t,t+1}$ and $BPV^{[\check{\delta}_1,\check{\delta}_2]}_{t,t+1}$ converge to the same integrated power transformation of the spot volatility process only differing in a constant scale multiplier, namely,

$$\int_t^{t+1} \sigma_s^{\check{\delta}} \mathrm{d}s = \int_t^{t+1} \sigma_s^{\check{\delta}_1+\check{\delta}_2} \mathrm{d}s \tag{4.7}$$

[8] In this sense, our results implicitly support asymptotic properties of these volatility measures.

Table 3
Parameter estimates and relative forecasting performance of MIDAS-RV model with $BPV^{[\delta_1,\delta_2]}$ regressors.

H	μ	ϕ	θ_2	δ_1	δ_2	In-sample				Out-of-sample			
						R^2	RSS (10^{-5})	$\frac{RSS}{RSS_2}$	$\frac{RSS}{RSS_1}$	R^2	RSS	$\frac{RSS}{RSS_2}$	$\frac{RSS}{RSS_1}$
5	−0.00005	0.17206	30.34535	0.72706	0.70518	0.60	2.10228	0.917	0.972	0.51	3.93361	0.986	0.983
	(−0.98)	(1.22)	(6.98)	(1.03)	(1.00)								
10	−0.00033	0.01206	10.08867	1.00848	0.11194	0.58	3.76496	0.903	1.008	0.50	6.44172	0.835	1.057
	(−1.38)	(0.80)	(3.92)	(1.71)	(0.24)								
15	−0.00087	0.00296	9.75516	0.03678	0.80155	0.54	5.67822	0.948	0.993	0.41	11.04998	0.935	1.009
	(−1.08)	(0.59)	(2.54)	(0.06)	(1.00)								
20	−0.00233	0.00048	7.07805	0.01962	0.52022	0.57	6.55608	0.844	0.975	0.39	13.22403	1.001	1.061
	(−1.02)	(0.68)	(2.23)	(0.03)	(0.55)								

MIDAS-RV model with $BPV^{[\delta_1,\delta_2]}$ regressors as defined in Eq. (4.4) is fitted over in-sample dataset (April 1, 1993 – December 31,2000) by NLLS methodology.

$$RV_{t,t+H} = \mu + \phi \sum_{k=1}^{50} w_k(\theta_1,\theta_2)\mathrm{BPV}^{[\delta_1,\delta_2]}_{t-k,t+1-k} + \varepsilon_t.$$

Then, its out-of-sample forecast performance is evaluated over the period from January 1, 2001 to October 31, 2003. H refers to forecast horizon in days. θ_1 is fixed at 1. Various RSS's are computed for comparison. RSS is the residual sum of squares from fitted MIDAS-RV model with $BPV^{[\delta_1,\delta_2]}$ regressors. For comparison, RSS_1 and RSS_2 are also computed from the fitted MIDAS-RV model with $RPV^{[1]}$ (*i.e. RAV*) and $RPV^{[2]}$ (*i.e. RV*) as regressors, respectively. R^2 is calculated as $1-RSS/TSS$. Finally, the numbers in the parenthesis are the t-stats computed with HAC standard errors.

and, hence as $M\to\infty$,

$$plim\, RPV^{[\check{\delta}]}_{t,t+1} = C\, plim\, BPV^{[\check{\delta}_1,\check{\delta}_2]}_{t,t+1} \quad \forall\, t \tag{4.8}$$

where $C=\mu_{\check{\delta}1}\mu_{\check{\delta}2}/\mu_{\check{\delta}1+\check{\delta}2}$ is a constant. Consequently, the regressors of lagged $RPV^{[\check{\delta}]}$ and $BPV^{[\check{\delta}_1,\check{\delta}_2]}$ span the same space asymptotically when $\check{\delta}=\check{\delta}_1+\check{\delta}_2$.

Using these starting values from the profiling, the estimation results for the MIDAS-RV model in Eq. (4.3) are shown in Table 2 and for Eq. (4.4) in Table 3. In both tables, ϕ's are not statistically significant with moderate t-stats. However, it does not mean that

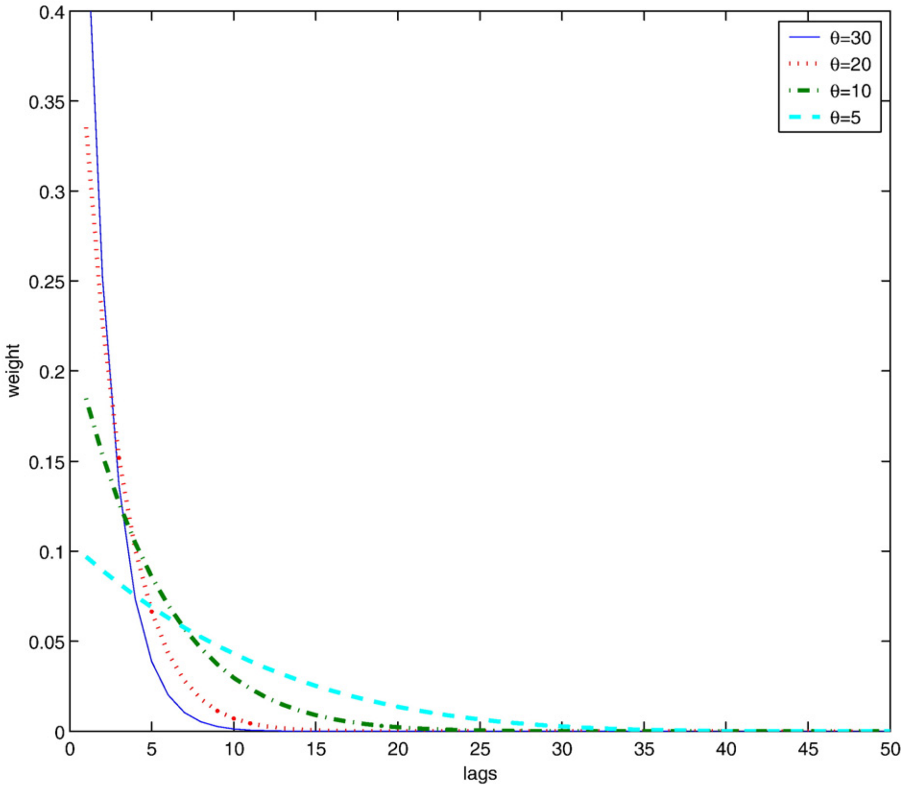

Fig. 9. Beta Weighting Scheme. This figure shows how Beta weighting scheme, as defined in Eq. (4.2), works in practice. Although the original Beta weighting scheme is a function of two θ parameters (*i.e.* θ_1 and θ_2), we fix $\theta_1 = 1$ in all estimations. This results in monotonically decreasing weighting scheme over the lags. Hence, the θ parameter shown in the figure is actually θ_2 in the original specification.

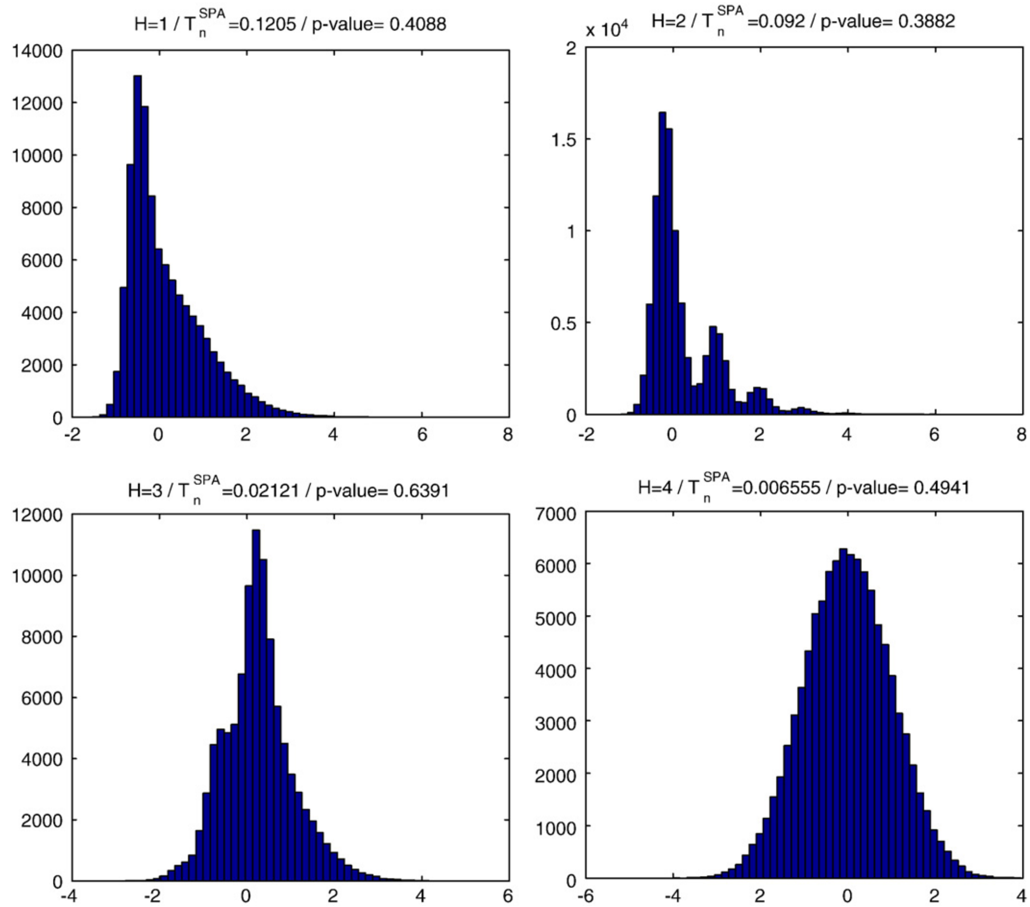

Fig. 10. Tests for Superior Predictive Ability (I). This figure shows the test statistics, their distribution, and the corresponding p-values for the superior predictive ability tests proposed by Hansen (2005). The distributions of test statistics are estimated by the stationary bootstrap of Politis and Romano (1994). In this set of tests over various forecast horizons, the benchmark model is the MIDAS-RV model with $RPV^{[\delta]}$ while the models compared are the MIDAS-RV with $RPV^{[1]}$ and $RPV^{[2]}$ ($=RV$). The test statistic T_n^{SPA} is computed as in Eq. (4.12) and the corresponding p-value is calculated from the empirical distribution.

the regressors of $\mathrm{RPV}^{[\delta]}$ or $BPV^{[\delta_1,\delta_2]}$ are useless in predicting future RV. As was already shown in Ghysels et al. (2006), these MIDAS-RV models consistently outperform ARFI(5,d) model of Andersen et al. (2003) in forecasting future RV's for both in- and out-of-sample and for various horizons. It is weakly identified though in the sense that the standard error is large. This seems to be due to parameterization of weights and power which adds free parameters to the regressors.

To appraise the weighting function in the MIDAS model we refer the reader to Fig. 9. If θ is around 30, the weighting function barely puts any weights beyond the 10th lag. For the θ values of 20, 10 and 5, the maximum lags that the weighting function puts any weights are 15, 20, and 30, respectively. Interestingly enough, the estimated θ_2's (θ_1 is fixed at 1) in both Tables 2 and 3 decrease as the forecast horizon gets longer; also the longer the forecast horizon gets, the larger number of lags get non-zero weights. Obviously, as the forecast horizon gets longer, we are actually modeling lower frequency movements of stock market volatility. In order to model low frequency (say a month or a quarter) variation of market volatility, it is necessary to make use of more lags of high frequency (say a day or a week) regressors since it is hard to capture slowly time-varying low frequency dynamics using regressors spanning a short period. Within this context, we can understand the estimated δ's in Tables 2 and 3; δ in Table 2 and $\delta_1+\delta_2$ in Table 3 become smaller as the forecasting horizon increases. As noted in Table 1 and Fig. 6, $RPV^{[1]}$ and $RPV^{[1.5]}$ showed very strong cross-correlation with RV for the first 10 lags, while cross-correlations with $RPV^{[0.5]}$ dominate after the 10th lag. If one needs to take into account more information from the past to predict the future volatility over longer horizon, one needs to make use of the strong cross-correlation of $RPV^{[0.5]}$ and hence the optimal value of δ will get smaller. However, since the cross-correlation relations examined in Table 1 and Fig. 6 are bivariate relations, we don't expect the estimated δ in Table 2 and $\delta_1+\delta_2$ in Table 3 to exactly match the magnitude of power that yield the highest cross-correlations. With regard to the results for BPV, we can make similar arguments using the relation (4.6).

An interesting comparison regarding forecast performance pertains to the MIDAS-RV model with RPV and BPV as shown in Tables 2 and 3. The additional free parameter does help reduce the forecast errors over the in-sample period, but not the out-of-sample period; the simple MIDAS-RV model with RPV turns out to be more robust. Also, we observe that δ's are strongly identified with small standard errors while δ_1's and δ_2's are weakly identified. This might be due to the asymptotic properties of $RPV^{[\delta]}$ and $BPV^{[\delta_1,\delta_2]}$. As was already discussed, these two measures of volatility essentially converge to the same population quantity when $\delta=\delta_1+\delta_2$. In fact, we can

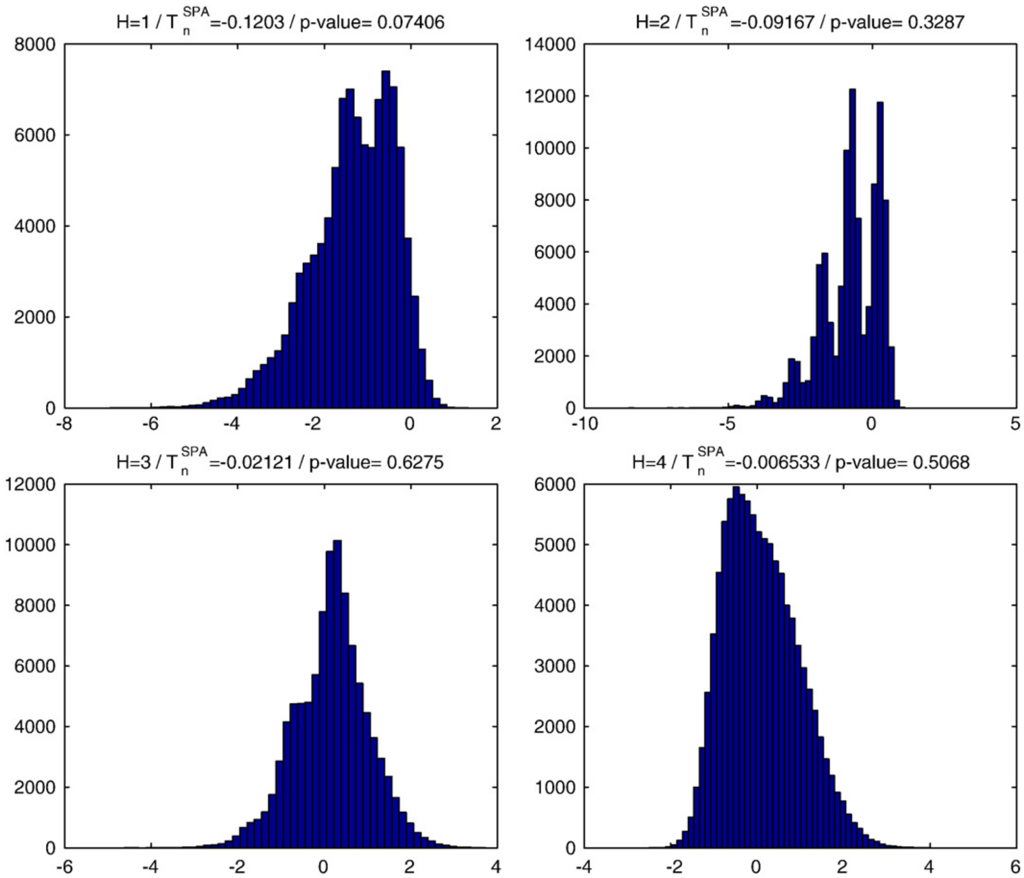

Fig. 11. Tests for Superior Predictive Ability (II). This figure shows the test statistics, their distribution, and the corresponding *p*-values for the superior predictive ability tests proposed by Hansen (2005). The distributions of test statistics are estimated by the stationary bootstrap of Politis and Romano (1994). In this set of tests over various forecast horizons, the benchmark model is the MIDAS-RV model with $RPV^{[1]}$ while the models compared are the MIDAS-RV with $RPV^{[\delta]}$ and $RPV^{[2]}$ ($=RV$). The test statistic T_n^{SPA} is computed as in Eq. (4.12) and the corresponding p-value is calculated from the empirical distribution.

observe that $\hat{\delta} \approx \hat{\delta}_1 + \hat{\delta}_2$ in Tables 2 and 3. Empirically, $RPV^{[\delta_1+\delta_2]}$ and $BPV^{[\delta_1+\delta_2]}$ might differ in many ways. However, asymptotic theory suggests that these two measures of volatility essentially converge to the same thing, which explains weak identification of power transformation parameters of *BPV*.

In order to evaluate the gains of modeling power transformations as free parameters in predicting future volatility measured by *RV*, Tables 2 and 3 provide relative forecasting performance measures; RSS/RRS_2 and RSS/RRS_1. *RSS* is the least residual sum of squares attained with a free parameter of δ in case of Eq. (4.3) and with δ_1 and δ_2 in case of Eq. (4.4). On the other hand, RSS_1 and RSS_2 are the least residual sum of squares attained by the MIDAS-RV models with selected lagged regressors of *realized absolute variation* $RAV = \text{RPV}^{[1]}$ and *realized variance* $RV = \text{RPV}^{[2]}$, respectively. The results from both Tables 2 and 3 suggest that there are some gains in adding free parameters when compared with the MIDAS-RV model with *RV* regressors, but barely any gains when compared with the case with *RAV* regressors. Another notable feature of the results shown in Tables 2 and 3 is that R^2, as a measure of overall fit of the model, falls significantly in the out-of-sample period when $H=15$ or $H=20$.[9] This implies that both MIDAS-RV models show sign of overfitting in long horizon forecast cases.

4.3. Superior predictive ability tests

In the previous section, we discussed, in terms of *RSS* ratios, the gains from parameterizing (bi)power variations in predicting future volatility. These ratios provide intuitive perspective to the discussion, but they cannot offer a formal answer to the question whether there are gains in parameterizing (bi)power variation since they are not statistical tests. In the present section, we discuss inference using formal statistical tests proposed by Hansen (2005) called tests for Superior Predictive Ability (henceforth SPA).[10]

[9] It should be pointed out that the deterioration of another measure of fit, the *RSS* ratio, in out-of-sample seems to be independent of the forecast horizon. This is because *RSS* ratio measures the goodness of the conditional forecast relative to that of another conditional forecast whose performance also deeply deteriorates as the forecast horizon gets longer in out-of-sample.

[10] For an empirical implementation, see also Hansen and Lunde (2005).

We would like to compare the predictive abilities of the MIDAS-RV models. We would like to see if the MIDAS-RV model with $RPV^{[\delta]}$ significantly does a better job at predicting future volatility than the MIDAS-RV model with $RPV^{[1]}$ or $RPV^{[2]}$ ($=RV$) where δ's are fixed. There are numerous loss functions and it is not clear which loss function is more appropriate for the evaluation of volatility models, as is discussed by Bollerslev et al. (1994), Diebold and Lopez (1996), and Lopez (2001). Patton (2006) notes that when volatility forecast comparisons are made using imperfect volatility proxies, as we do, the mean squared error has many desirable properties. Hence, in comparing and evaluating volatility models, we adopt the traditional loss function of mean squared error:

$$L^{(m)}_{t,t+H} = \left(\sigma^2_{t,t+H} - h^2_{(m)t,t+H}\right)^2 \tag{4.9}$$

where $\sigma^2_{t,t+H} = \int_t^{t+H} \sigma^2_s ds$ and m refers to the MIDAS-RV model specification that generated volatility forecasts $h^2_{(m)t,t+H}$. In our empirical analysis we substitute the realized variance RV for the latent $\sigma^2_{t,t+H}$.

The parameters of the MIDAS-RV models are estimated using the in-sample return data and these estimates are used to make forecasts over the out-of-sample period. During this period, we calculate the realized variance from intraday returns and obtain $\hat{\sigma}^2_{t,t+H}$ for each forecasting horizon H. Let the model $m=0$ be the benchmark model that is compared to models $m=1,\ldots,K$. For each model and forecast horizon, we define the relative performance variables

$$Y^{(m)}_{t,t+H} = L^{(0)}_{t,t+H} - L^{(m)}_{t,t+H} \tag{4.10}$$

Our null hypothesis is that the benchmark model is as good as any other model in terms of expected loss. This can be formulated as the hypothesis

$$H_0 : E\left[Y^{(m)}_{t,t+H}\right] \leq 0 \quad \textit{for all } m \tag{4.11}$$

The SPA test proposed by Hansen (2005) is based on the test statistic

$$T_n^{\mathrm{SPA}} \equiv \max_{m=1,\ldots,K} \bar{Y}^{(m)} / \hat{\omega}_m \tag{4.12}$$

where n is the number of non-overlapping $Y^{(m)}_{t,t+H}$'s in the out-of-sample period, $\bar{Y}^{(m)} = n^{-1}\sum Y^{(m)}_{t,t+H}$, and $\hat{\omega}^2_m$ is a consistent estimator of $\omega^2_m \equiv lim_{n\to\infty} Var\left(\sqrt{n}\bar{Y}^{(m)}\right)$ for $m=1,\ldots,K$. To obtain estimates of ω^2_m and the distribution of the test statistic T_n^{SPA}, we use the stationary bootstrap of Politis and Romano (1994).[11] We follow Hansen and Lunde (2005) to implement the bootstrap methodology.

First, we would like to check if the MIDAS-RV model with parameterized $RPV^{[\delta]}$ is as good as the MIDAS-RV models with $RPV^{[1]}$ and $RPV^{[2]}$. We take the MIDAS-RV model with parameterized power variation as the benchmark model for the SPA test. Fig. 10 shows the results. For all forecast horizons, the p-values are large which means that we cannot reject the null hypothesis that the benchmark model is as good as the comparison models. Second, we adopt the MIDAS-RV model with $RPV^{[1]}$ as the benchmark model. *RSS* ratios in Table 2 show the possibility that the assumed benchmark model outperforms the MIDAS-RV model with the parameterized power variation in predicting long horizon future volatility. Fig. 11 shows interesting results. For the shortest forecast horizon of a week, the p-value is 7.4%. With fairly generous level of significance (*e.g.* 10%), we can reject the null hypothesis that the MIDAS-RV model with $RPV^{[1]}$ is as good as the MIDAS-RV model with $RPV^{[\delta]}$. However, for all the longer forecast horizon, we cannot reject the null hypothesis. These two sets of SPA tests for various forecast horizons suggest the following; the MIDAS-RV model with parameterized power variation is better than those with fixed δ's when the forecast horizon is short, but, for longer forecast horizons, it is hard to tell which one is significantly better than the other.

5. Conclusions

We investigated the empirical cross-correlation structure between the current RV and lagged RPV's for various power transformations and find that it is quite strong. For the first 10 lags, RPV with $\delta=1$ seems to show the strongest correlation with RV, beyond the 10th lag, it is RPV with $\delta=0.5$. This feature has important implication in predicting volatility over different horizons. As the prediction horizon increases, the optimal value of δ becomes smaller. This suggests that we may want to estimate the power variation that predicts future volatility best.

For both volatility measures RPV and BPV, the gains of letting power transformations be free parameters are quite large when compared with autoregressive models. However, RPV with δ fixed at 1 is good enough that the gains are very marginal for in-sample comparisons and even negative in many of out-of-sample results. These findings reinforce the results of Forsberg and Ghysels (2007).

[11] The procedure involves a dependence parameter q that serves to preserve possible time-dependence in $Y^{(m)}_{t,t+H}$. We used $q=0.5$ and generated $B=100{,}000$ bootstrap re-samples in our empirical analysis.

References

Andersen, T., Bollerslev, T., Diebold, F.X., Labys, P., 2001. The distribution of exchange rate volatility. Journal of American Statistical Association 96, 42–55.
Andersen, T.G., Bollerslev, T., Diebold, F.X., 2006. Roughing it up: including jump components in the measurement, modeling and forecasting of return volatility. Review of Economics and Statistics 89, 701–720.
Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2003. Modeling and forecasting realized volatility. Econometrica 71, 579–625.
Andreou, E., Ghysels, E., 2002. Detecting multiple breaks in financial market volatility dynamics. Journal of Applied Econometrics 17, 579–600.
Andreou, E., Ghysels, E., 2009. Structural breaks in financial time series. In: Andersen, T., Davis, R., Kreiss, J.P., Mikosch, T. (Eds.), Handbook of Financial Time Series. Wiley.
Banerjee, A., Urga, G., 2005. Modelling structural breaks, long memory and stock market volatility: an overview. Journal of Econometrics 129, 1–34.
Barndorff-Nielsen, O.E., Shephard, N., 2003. Realised power variation and stochastic volatility models. Bernoulli 9, 243–265.
Barndorff-Nielsen, O.E., Shephard, N., 2004. Power and bipower variation with stochastic volatility and jumps. Journal of Financial Econometrics 2, 1–48.
Barndorff-Nielsen, O.E., Graversen, S.E., Jacod, J., Shephard, N., 2005. Limit theorems for bipower variation in financial econometrics. Econometric Theory 22, 677–719.
Bollerslev, T., Engle, R.F., Nelson, D.B., 1994. ARCH Models, vol. IV of Handbook of Econometrics. Elsevier Science, Amsterdam.
Diebold, F., Inoue, A., 2001. Long memory and regime switching. Journal of Econometrics 105, 131–159.
Diebold, F.X., Lopez, J.A., 1996. Forecast Evaluation and Combination, vol. 14 of Handbook of Statistics. North-Holland, Amsterdam.
Ding, Z., Granger, C.W.J., 1996. Modeling volatility persistence of speculative returns: a new approach. Journal of Econometrics 73, 185–215.
Ding, Z., Granger, C., Engle, R., 1993. A long memory property of stock market returns and a new model. Journal of Empirical Finance 1, 83–106.
Engle, R., Ghysels, E., and Sohn, B. (2006), "On the Economic Sources of Stock Market Volatility," Unpublished manuscript, NYU and UNC.
Forsberg, L., Ghysels, E., 2007. Why do absolute returns predict volatility so well. Journal of Financial Econometrics 5, 31–67.
Ghysels, E., Santa-Clara, P., Valkanov, R., 2002. The MIDAS Touch: Mixed Data Sampling Regression Models. University of North Carolina at Chapel Hill.
Ghysels, E., Santa-Clara, P., Valkanov, R., 2006. Predicting volatility: getting the most out of return data sampled at different frequencies. Journal of Econometrics 131, 59–95.
Gourieroux, C., Jasiak, J., 2001. Memory and infrequency breaks. Economics Letters 70, 29–41.
Granger, C., Hyung, N., 2004. Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns. Journal of Empirical Finance 11, 399–421.
Hansen, P.R., 2005. A test for superior predictive ability. Journal of Business and Economic Statistics 23, 365–380.
Hansen, P.R., Lunde, A., 2005. A forecast comparison of volatility models: does anything beat a GARCH(1,1). Journal of Applied Econometrics 20, 873–889.
Lopez, J.A., 2001. Evaluation of predictive accuracy of volatility models. Journal of Forecasting 20, 87–109.
Patton, A. (2006), "Volatility Forecast Comparison using Imperfect Volatility Proxies," Unpublished manuscript, Oxford University.
Politis, D.N., Romano, J.P., 1994. The stationary bootstrap. Journal of the American Statistical Association 89, 1303–1313.
Wood, R.A., McInish, T.H., Ord, J.K., 1985. An investigation of transactions data for NYSE stocks. Journal of Finance 40, 723–739.