



HAL
open science

Do We Need High Frequency Data to Forecast Variances?

Denisa Banulescu-Radu, Christophe Hurlin, Bertrand Candelon, Sébastien Laurent

► **To cite this version:**

Denisa Banulescu-Radu, Christophe Hurlin, Bertrand Candelon, Sébastien Laurent. Do We Need High Frequency Data to Forecast Variances?. *Annals of Economics and Statistics*, 2016, 123/124, pp.135–174. 10.15609/annaeconstat2009.123-124.0135 . hal-01448237

HAL Id: hal-01448237

<https://amu.hal.science/hal-01448237>

Submitted on 9 Feb 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Do We Need High Frequency Data to Forecast Variances?

Denisa BANULESCU-RADU

University of Orléans (LEO, UMRS CNRS 7332)

Bertrand CANDELON

Maastricht University

Christophe HURLIN

University of Orléans (LEO, UMRS CNRS 7332)

Sébastien LAURENT

*Aix-Marseille School of Economics, CNRS and EHESS,
Aix-Marseille Graduate School of Management, France*

In this paper we study various MIDAS models for which the future daily variance is directly related to past observations of intraday predictors. Our goal is to determine if there exists an optimal sampling frequency in terms of variance prediction. Via Monte Carlo simulations we show that in a world without microstructure noise, the best model is the one using the highest available frequency for the predictors. However, in the presence of microstructure noise, the use of very high-frequency predictors may be problematic, leading to poor variance forecasts. The empirical application focuses on two highly liquid assets (i.e., Microsoft and S&P 500). We show that, when using raw intraday squared log-returns for the explanatory variable, there is a “high-frequency wall” – or frequency limit – above which MIDAS-RV forecasts deteriorate or stop improving. An improvement can be obtained when using intraday squared log-returns sampled at a higher frequency, provided they are pre-filtered to account for the presence of jumps, intraday diurnal pattern and/or microstructure noise. Finally, we compare the MIDAS model to other competing variance models including GARCH, GAS, HAR-RV and HAR-RV-J models. We find that the MIDAS model – when it is applied on filtered data – provides equivalent or even better variance forecasts than these models.*

I. Introduction

Volatility has been, and will remain, one of the most active and successful research topics in financial econometrics and economic forecasting. Nowadays, due to the widespread progress of financial activities, different financial agents have developed relatively different expectations and views on forecasting volatility. High-frequency traders are certainly interested in computing ultra-high-frequency volatility forecasts (e.g., second-by-second volatility forecasts), while financial regulators, for instance, still require daily up to bi-weekly volatility forecasts. In addition, due to the widespread availability of high-quality financial data sampled at very high-frequency, and the rich information it is supposed to contain, traders might be tempted to use it as predictors of volatility, regardless of the volatility forecasts horizon. However, high-frequency data is costly, processing it requires supplementary human capital, and nothing guarantees that very high-frequency data improves the quality of daily/weekly/bi-weekly volatility forecasts.

*JEL: C22, C53, G12 / KEY WORDS: Variance Forecasting, MIDAS, High-Frequency Data.

This paper focuses on longer horizon variance forecasts (i.e., daily/weekly/bi-weekly variance forecasts needed, inter alia, in the context of portfolio management, option pricing and financial market regulation), and, more precisely, on the type and frequency of predictors used to obtain these forecasts. To this end, we use the mixed data sampling (henceforth MIDAS) regression model, introduced in GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2004], as it makes it possible for instance, to forecast a measure of the *daily* variance (e.g., realized variance) by considering past *intraday* log-returns. In their seminal paper, GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006] consider various MIDAS regressions with different daily (squared returns, absolute returns, realized variance, realized power and return range) and intradaily regressors (squared returns, absolute returns), to examine whether one specification dominates the others. The goal of our study is different in that it consists in determining for a *given* intradaily predictor whether a *sampling frequency* (or range of frequencies) dominates the others.¹ The objective is then to identify the best sampling frequency, using out-of-sample forecast evaluation criteria.

This issue is not straightforward. On the one hand, not including the readily available high-frequency observations to perform variance forecasts implies a loss of information due to temporal aggregation. On the other hand, if the sampling frequency of the predictors is increased too much, market microstructure noise (bid-ask bounce, screen fighting, jumps, and irregular or missing data) may lead to less accurate variance forecasts.

This question is different in nature from the well-documented discussion about the optimal sampling frequency of the returns used to compute realized estimators of daily variance (see HANSEN, P. R., and A. LUNDE [2004]; AIT-SAHALIA, Y., and L. MANCINI [2008]; GARCIA, R., and N. MEDDAHI [2006]; GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006], among others). Our goal consists in focusing on the optimal sampling frequency for the purpose of variance *prediction*, and not for variance *measurement*. Consider a MIDAS variance model whose aim is to predict a measure of variance over some future horizon. This variance measure is typically a realized measure (realized variance, realized kernel etc.), based on intradaily returns sampled at a certain frequency (one minute, five minutes, etc.). In order to forecast variance, we consider intradaily predictors (squared returns, intradaily bipower variation, etc.) sampled at a frequency that may be different from that used to compute the realized measure of variance. As indicated above, the upcoming discussion concerns only the sampling frequency of the predictors.² With this in mind, we consider MIDAS models in which we directly project

1. In this study, we limit our analysis to the MIDAS specifications in which the future variance is directly related to past observations of intraday predictors, as in GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006]. An alternative consists in using high-frequency data to compute daily realized measures (realized variance, two-scale estimator, realized kernel, etc.) which are, in a second step, introduced as predictors into a MIDAS regression model, as in GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006] and GHYSELS, E., and A. SINKO [2011].

2. In a related paper, GHYSELS, E., and A. SINKO [2011] also examine the optimal sampling for the purpose of variance prediction, but in a slightly different context. The authors observe that, generally, discussions about the impact of microstructure have mostly focused on measurement. GHYSELS, E., and A. SINKO [2011] focus instead on the prediction of variance, by using as regressors variance measures that are contaminated by market microstructure noise. Their dependent variable is defined as the two scales estimator of the weekly variance AIT-SAHALIA, Y., P. A. MYKLAND, and L. ZHANG [2005], and computed from the 5-minute, 1-minute or 2-second returns. One of the main differences with our study is that the authors consider various MIDAS specifications for which the predictors also correspond to realized estimators (plain vanilla, two scales estimator; ZHOU, B., [1996], etc.), constructed using different sampling frequencies (from two seconds to ten minutes). Thus, high-frequency data is aggregated into daily realized measures, which are then used as predictors of future variance. In contrast, our goal is to analyze the direct impact of the intra-

future realized variance onto high-frequency regressors, as in GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006].³

To address these issues, we investigate the sensitivity of MIDAS variance models to the choice of the sampling frequency of predictors on real data, i.e., log-returns of the S&P 500 index and Microsoft over the period from October 29, 2004 to December 31, 2008. Thus, we apply simple MIDAS specifications in which daily realized variance is predicted by past intradaily squared log-returns sampled at a frequency ranging from one minute to one day. The variance forecasts are compared based on the robust loss function proposed by PATTON, A. J. [2011] and the model confidence set (MCS) test introduced by HANSEN, P. R., A. LUNDE, and J. M. NASON [2011]. This test identifies, among the set of competing models (i.e., sampling frequencies), the subset of models that are equivalent in terms of forecasting ability and which outperform all the other models for a given confidence level. The empirical results obtained for these two assets allow us to draw some interesting conclusions. First, when using raw intraday returns, variance forecasts are not statistically different for sampling frequencies of the predictors ranging from five minutes to one hour. Moreover, it turns out that very high-frequency regressors (i.e., higher than five minutes) do not provide useful information to improve the variance forecasts because the loss function increases. The shape of the loss function indicates the presence of a “high-frequency wall”, i.e., a limit frequency beyond which the quality of the forecasts deteriorates or stops improving. This result is due to the presence of microstructure noise, jumps and intraday diurnal patterns in the regressors. When the MIDAS regression model is applied to filtered data LEE, S. S., and P. A. MYKLAND [2008], BOUDT, K., C. CROUX, and S. LAURENT [2011], LAHAYE, J., S. LAURENT, and C. J. NEELY [2011], the conclusion in favor of the use of the highest available frequency remains valid. This point is crucial and indicates that the mixing frequency may require the use of filtered series. Indeed, the weighting scheme in MIDAS models does not allow us, by itself, to underweight the observations affected by jumps or other market microstructure noise. These results are robust to the choice of the variance measure (realized variance, realized kernel), the forecasting horizon and the sample period (calm/crisis).

Finally, we compare the performance of the MIDAS model (considering filtered or unfiltered regressors) to other competing variance models, namely the GARCH(1,1), the Student Generalized Autoregressive Score (GAS) model CREAL, D. D., S. J. KOOPMAN, and A. LUCAS [2013], the Heterogeneous Autoregressive Realized Variance-based (HAR-RV) model CORSI, F. [2009] and the HAR-RV adjusted for jumps ANDERSEN, T. G., T. BOLLERSLEV, and F. X. DIEBOLD [2007]. We show that MIDAS models provide comparable or even better variance forecasts when filtered high-frequency data is used.

The paper is structured as follows. SECTION II introduces the notations, the MIDAS model and the sampling frequency puzzle. SECTION III presents the main results of a Monte Carlo simulation study. In SECTION IV, we perform an empirical analysis and study the influence of the jumps and the intraday diurnal pattern on the MIDAS performances. We also compare MIDAS to other competing variance models. SECTION V concludes.

daily predictors on the variance forecasts, and ultimately to evaluate the usefulness of the mixing of frequencies in this context.

3. Surprisingly, GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006] find that the forecasts directly using high-frequency data do not outperform those based on daily regressors (although the daily regressors are themselves obtained through the aggregation of high-frequency data). One related question is to understand whether this result depends on the sampling frequency of the high-frequency data.

II. Modeling Strategies

II.1. Notation

To set the notation, let p_t denote the price for a financial asset sampled at daily frequency, and the corresponding daily return be defined by $r_{t,t-1} \equiv \log(p_t) - \log(p_{t-1})$. The equally spaced series of continuously compounded returns is assumed to be observed m times per day (or to have a horizon of $1/m$), and be computed as $r_{t,t-1/m}^{(m)} \equiv \log(p_t) - \log(p_{t-1/m})$, where $t = 1/m, 2/m, \dots$. Throughout the analysis, we consider that the trading day spans the time period from 9:30 am to 16:00 pm, covering for instance $m = 390$ 1-minute equally spaced intervals and $m = 78$ 5-minute equally spaced intervals. $r_{t,t-1/78}$ corresponds to the last 5-minute return of the day $t - 1$, $r_{t-1/78,t-2/78}$ corresponds to the return of the penultimate 5-minute period of day $t - 1$, and so on.

II.2. MIDAS variance Models and Sampling Frequency

MIDAS models for variance predictions have been introduced in a number of recent studies, including GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2005], GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006], GHYSELS, E., and A. SINKO [2011], GHYSELS, E., and R. VALKANOV [2012], CHEN, X., and E. GHYSELS [2011], among others.

The general specification of the MIDAS variance model is given by:

$$\sigma_{t+H,t}^2 = \mu_{H,m_1} + \phi_{H,m_1} \Omega_{H,m_1}(L^{1/m_1}) X_{t,t-1/m_1}^{(m_1)} + \varepsilon_t, \quad (1)$$

where $\sigma_{t+H,t}^2$ is a measure of variance evaluated over some future horizon H , and $X_{t,t-1/m_1}^{(m_1)}$ denotes an intradaily regressor sampled at frequency m_1 . The distributed lag polynomial is defined as:

$$\Omega_{H,m_1}(L^{1/m_1}) = \sum_{k=0}^{k_{max}} L^{k/m_1} \omega_{H,m_1}(k, \theta_{H,m_1}), \quad (2)$$

where $\omega_{H,m_1}(k, \theta_{H,m_1})$ corresponds to the lag coefficient associated with $X_{t,t-1/m_1}^{(m_1)}$, θ_{H,m_1} is a finite set of parameters, L is the lag operator such that $L^{1/m_1} X_{t,t-1/m_1}^{(m_1)} = X_{t-1/m_1,t-2/m_1}^{(m_1)}$, and k_{max} denotes the maximum number of lagged coefficients. In this specification, the low-frequency variance (for instance, daily variance if $H = 1$, weekly variance if $H = 5$, etc.) is predicted by the right-side intradaily forecasting factors, which are sampled at a high-frequency m_1 (for instance, five minutes if $m_1 = 78$). Several intradaily regressors can be considered with this aim in mind (e.g., intradaily squared returns, intradaily bipower variation). Following GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006] we consider the intradaily squared returns $r_{t,t-1/m_1}^{(m_1)^2}$, while other alternatives will be used to appraise the robustness of our results.⁴

4. Another alternative would be to accommodate large volumes of data in a parsimonious way by using intraday aggregate measures of variance (e.g., 5-minute realized variance, etc.) as regressors. However, the aim of this paper

Since $\sigma_{t+H,t}^2$ is unobservable, we rely on a proxy. For simplicity, we adopt the realized variance ANDERSEN, T. G., and T. BOLLERSLEV [1998a], defined for the period t to $t + H$ as follows:⁵

$$RV_{t+H,t}^{(m_2)} = I_{H,m_2}(L^{1/m_2})r_{t+H,t+H-1/m_2}^{(m_2)^2} \quad (3)$$

where the distributed lag polynomial in L^{1/m_2} is defined such that $I_{H,m_2}(L^{1/m_2}) = \sum_{j=0}^{Hm_2} L^{j/m_2}$, and m_2 accounts for the sampling frequency of the squared returns used to compute the realized variance. Notice that the frequencies m_1 and m_2 may be different. In fact, the choice of m_2 is related to the variance *measurement* issue (see HANSEN, P. R., and A. LUNDE [2004]; AIT-SAHALIA, Y., and L. MANCINI [2008]; GARCIA, R., and N. MEDDAHI [2006]; GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006], among others), i.e., the consistency of the estimator defined by the realized measure.⁶ On the other hand, the choice of the sampling frequency of the predictors, m_1 is related to the variance *prediction* issue. Indeed, this choice determines the regressors in the MIDAS variance model, and, as a consequence, its forecasting abilities.

Under these assumptions, the MIDAS-RV regression becomes:

$$RV_{t+H,t}^{(m_2)} = \mu_{H,m_1} + \phi_{H,m_1} \Omega_{H,m_1}(L^{1/m_1})r_{t,t-1/m_1}^{(m_1)^2} + \varepsilon_t \quad (4)$$

One advantage of this specification is that it preserves the information contained in high-frequency data GHYSELS, E., and R. VALKANOV [2012] without computing daily aggregates such as realized variance for the regressors. In this context, we aim at determining the influence of the sampling frequency m_1 on the forecasting performances of the MIDAS model. In a related study, GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006] compare several MIDAS specifications based on *different* intradaily or daily variance regressors (e.g., squared returns, absolute returns, realized variance, realized power, and range). The rationale is similar here, except that we consider the *same intradaily regressor*, i.e., $X_{t,t-1/m_1}^{(m_1)}$, for *various sampling frequencies*. For instance, we compare various MIDAS models where the same predictor is sampled at one minute ($m_1 = 390$), two minutes ($m_1 = 195$), five minutes ($m_1 = 78$), and so on. The question is whether increasing the sampling frequency m_1 systematically improves the quality of the variance forecasts, and ultimately whether we need high-frequency data in order to forecast daily variances.

is rather to observe the behavior of high-frequency data when estimating/forecasting variance and not to propose the best variance model and/or predictor of variance.

5. A large number of alternative estimators (e.g., realized bipower variation, realized kernel, etc.) that deal with issues such as jumps and other market microstructure noise, have been proposed, especially by BARNDORFF-NIELSEN, O. E., and N. SHEPHARD [2004a], BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, and N. SHEPHARD [2008], ZHANG, L. [2006], HANSEN, P. R., and G. HOREL [2009], *inter alia*. Some of these will be considered in the section devoted to the robustness analysis of our findings.

6. Since this study is not meant to determine the optimal sampling frequency m_2 , in the rest of the paper, the daily RV will always be computed by summing up *5-minute* squared returns (i.e., $m_2 = 78$), as recommended by ANDERSEN, T. G., and T. BOLLERSLEV [1997a]. Since the RV is derived from observed trades and/or quotes that are contaminated with market microstructure noise (which can lead to an autocorrelation of the intraday returns), the RV might become biased and inconsistent for the integrated variance. Nevertheless, ANDERSEN, T. G., and T. BOLLERSLEV [1997a] argue that it is common practice to use moderate-frequency intraday returns, such as 5-minute returns, as this partially offsets the bias.

Another key feature of MIDAS models is that they provide a parsimonious specification. This property is particularly important in our context, as the inclusion of high-frequency data might imply a significant increase in the number of lagged forecasting variables, and hence in the number of unrestricted parameters to be estimated GHYSELS, E., and R. VALKANOV [2012]. For instance, running unrestricted regressions based on the intraday information over the last 30 days implies estimating 30×390 parameters for a 1-minute regressor, 30×78 parameters for a 5-minute regressor, and so on. Nevertheless, the MIDAS model directly projects future variance onto a large number of high-frequency lagged regressors while considering a small number of parameters. The trick consists in using a suitable parametrization for the weights $\omega_{H,m_1}(k, \theta_{H,m_1})$ to circumvent the problem of parameter proliferation. Therefore, as noted by GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006], the parametrization $\omega_{H,m_1}(k, \theta_{H,m_1})$ becomes a key ingredient in a MIDAS regression.

Two specifications of the weight function are generally considered, namely the exponential Almon lag and the Beta lag GHYSELS, E., A. SINKO, and R. VALKANOV [2007]. These specifications have several interesting features: *i*) the distributed lag polynomial is tightly parameterized and prevents the proliferation of parameters as well as additional pre-testing or lag-selection procedures;⁷ *ii*) the coefficients are positive, which guarantees non-negative weights and consequently non-negative variance forecasts; *iii*) the sum of the data-driven weights is normalized to one in order to identify the scale parameter ϕ_{H,m_1} . There is no clear theoretical a priori for assuming that one specification is better than the other. However, CHEN, Y.-C., and W.-J. TSAY. [2011] and FRALE, C., and L. MONTEFORTE. [2011] find that the Beta function is more suitable for an important number of time lags, as Almon can be very computationally demanding in such a context. For this reason, we adopt the Beta lag polynomial

$$\omega_{H,m_1}(k, \theta_{H,m_1}) = \frac{f(k / k^{\max}, \theta_1; \theta_2)}{\sum_{j=0}^{k^{\max}} f(j / k^{\max}, \theta_1; \theta_2)}, \quad (5)$$

where $\theta_{H,m_1} = (\theta_1, \theta_2)'$ is a vector of positive parameters, $f(z, a, b) = z^{a-1}(1-z)^{b-1} / B(a, b)$, with $B(\cdot)$ the Beta function defined as $B(a, b) = \Gamma(a)\Gamma(b) / \Gamma(a+b)$, and $\Gamma(\cdot)$ representing the Gamma function. Depending on the value of the parameter θ_1 , this weight function can take many shapes, including flat weights, gradually declining weights and hump-shaped patterns. The second parameter, θ_2 , determines the decreasing speed of the weighting shape. The smaller the parameter θ_2 , the smoother the weighting scheme. In other words, θ_2 determines the proportion of the total weight associated with the more recent past observations.

III. Monte Carlo Simulation Study

Before starting the empirical application, we proceed by a Monte Carlo simulation study to analyze the influence of the sampling frequency of the regressors on the predictive abilities of

7. The selection of k^{\max} can be done by considering a large value and letting the weights vanish.

MIDAS-RV models.⁸ Considering a noise-free diffusion process, we generate returns series at different sampling frequencies and daily realized variance measures, using the same set of continuous-time structural parameters. We apply simple MIDAS specifications in which daily realized variance is predicted by past intradaily squared log-returns sampled at a frequency ranging from one minute to one day. Sequences of $T = 500$ daily realized variance forecasts, $\{\widehat{RV}_{t+1,t}^{(m_2)}\}_{t=1}^T$, are then computed for each MIDAS-RV specification.⁹

In order to compare the variance forecasts, we use a loss function defined as a general function of the variance forecasts and the true variance. In the simulation framework, the variance can be measured by the daily integrated variance, $IV_{t,t-1} = \int_{t-1}^t \sigma^2(\tau) d\tau$. However, in practice, the integrated variance is not observable and we have to use a proxy. To reproduce the real conditions of the MIDAS-RV models, we also consider a variance proxy in the simulation study, i.e., the realized variance $RV_{t+H,t}^{(m_2)}$. However, it is well known that the use of a proxy may distort the ranking of models based on loss functions. ANDERSEN, T. G., and T. BOLLERSLEV [1998a] and ANDERSEN T. G., T. BOLLERSLEV, and N. MEDDAHI [2005] show that the comparison of losses – even based on a conditionally unbiased proxy – may lead to a different outcome than the one obtained if the true latent variable had been used. More recently, HANSEN, P. R., and A. LUNDE [2006a], PATTON, A. J., and K. SHEPPARD [2009], PATTON, A. J. [2011], LAURENT, S., J. V. ROMBOUTS, and F. VIOLANTE [2013] have also insisted on the possible distortions observed in the ranking of variance/volatility forecasts induced by the use of a noisy proxy.¹⁰ For these reasons, we adopt the family of robust and homogeneous loss functions proposed by PATTON, A. J. [2011], i.e.,

$$L(\widehat{\sigma}^2, \sigma^2; b) = \begin{cases} \frac{1}{(b+1)(b+2)} (\widehat{\sigma}^{2(b+2)} - \sigma^{2(b+2)} - \frac{1}{b+1} \widehat{\sigma}^{2(b+1)} (\widehat{\sigma}^2 - \sigma^2)), & \text{for } b \notin \{-1, -2\} \\ \sigma^2 - \widehat{\sigma}^2 + \widehat{\sigma}^2 \log \frac{\widehat{\sigma}^2}{\sigma^2}, & \text{for } b = -1 \\ \frac{\widehat{\sigma}^2}{\sigma^2} - \log \frac{\widehat{\sigma}^2}{\sigma^2} - 1, & \text{for } b = -2 \end{cases} \quad (6)$$

with b a scalar parameter, σ^2 a measure of the true variance (i.e., the realized variance in our case) and $\widehat{\sigma}^2$ the predicted variance measure. This loss function encompasses in particular the MSE and the QLIKE loss functions when $b = 0$ and $b = -2$, respectively. In the empirical application we only consider the results of the QLIKE function. Nevertheless, in order to assess the robustness of our results, the three values for the parameter b are considered in the simulation analysis.

8. For ease of exposition, all the results related to this study are reported in a Web Appendix.

9. The Monte Carlo simulation exercise is based on 10,000 replications.

10. The robustness of the ranking forecasts also has an impact on the statistical inference used to assess the predictive accuracy. If the loss function ensures consistency of the ranking, the variability of the variance proxy is only likely to reduce the power of the test, but not its asymptotic size, which means that for a robust loss function it is always possible to recover the true ranking asymptotically. For more details, see LAURENT, S., J. V. ROMBOUTS, and F. VIOLANTE [2013].

Evaluating the influence of the sampling frequency of the predictors on the predictive abilities of the MIDAS-RV model reduces to determining the sign of the derivative of the average loss function given by:

$$L_{m_1} = T^{-1} \frac{\partial \sum_{t=1}^T L(\widehat{RV}_{t+1,t}^{(m_2)}, RV_{t+1,t}^{(m_2)}; b)}{\partial m_1} \quad (7)$$

Since the sign of this derivative cannot be determined analytically, we proceed by numerical analysis.

Therefore, the variance forecasts are compared based on the robust loss function proposed by PATTON, A. J. [2011] and the model confidence set (MCS) test introduced by HANSEN, P. R., A. LUNDE, and J. M. NASON [2011].¹¹ This test aims at identifying, among the set of competing models (i.e., sampling frequencies), the subset of models that are equivalent in terms of forecasting ability and which outperform all the other models for a given confidence level. Several results stand out. First, we observe that a higher sampling frequency for the regressors leads to attributing higher weights to the most recent observations of the regressors. Second, we show that increasing the frequency of the regressors always improves the forecasting abilities of the MIDAS model. The average loss increases when the regressors are sampled less frequently, regardless of the choice of the loss function. Nevertheless, opting for very high frequency regressors is not optimal in the presence of microstructure noise. Therefore, the choice of the optimal sampling frequency m_1 for the predictors in MIDAS-RV models is not obvious.

In the Monte Carlo simulation study we considered two cases where the use of the highest available frequency does not necessarily improve the quality of the fit or the predictions, and discovered that a “high-frequency wall” might exist (i.e., a frequency limit above which MIDAS-RV forecasts deteriorate or do not improve). The two following features of the DGP were analyzed in order to justify the presence of this “high-frequency wall”:

- the process is not a pure continuous-time model, but rather a model where the conditional variance is constant by chunks of, for instance, one or five minutes;
- and/or the presence of microstructure noise.

For instance, we know that very a high-frequency price process is likely to be contaminated by microstructure effects arising from market frictions.¹² Therefore, an important part of our simulation study consists in using the pre-averaging method PODOLSKIJ, M., M. VETTER, ET AL. [2009], JACOD, J., Y. LI, P. A. MYKLAND, and M. PODOLSKIJ, and M. VETTER [2009] as a powerful technique to robustify variance estimators to the presence of microstructure noise.¹³ FIGURE I displays the average QLIKE for 500 out-of-sample one-step-ahead forecasts (over 10,000 replications) as a function of the sampling frequency m_1 . Three models are considered. The (blue) solid line corresponds to the case where the MIDAS-RV is estimated on non-contaminated log-returns. The (red) dashed line corresponds to the case where the MIDAS-RV is

11. For more details, see Appendix A.

12. To contaminate the log-returns by noise, a normal random variable with mean 0 and variance $10^{-3} \times IQ$ is added to every 1-second log-return, knowing that the daily integrated quarticity is defined as follows: $IQ_{t,t-1} = \int_{t-1}^t \sigma^4(\tau) d\tau$. Including different noise structures (i.e., i.i.d noise vs. autocorrelated noise) remains an interesting direction for future research.

13. More technical details on this method are provided in the next section.

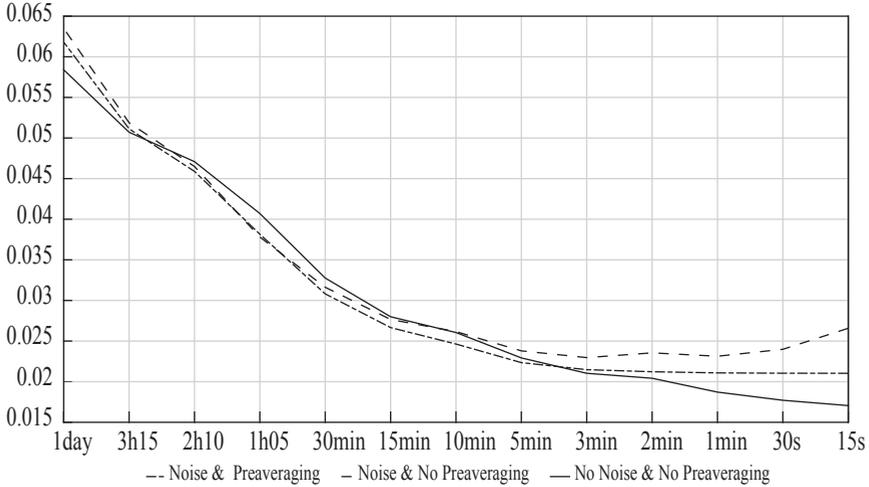


FIGURE 1. – MIDAS-RV average QLIKE

Note: This figure displays the average QLIKE (y-axis) associated with the MIDAS-RV forecasts based on various sampling frequencies (m_1) of the predictors (x-axis). Three different MIDAS specifications are considered: (i) the regressors are not contaminated by noise, (ii) the regressors are contaminated by noise, (iii) the regressors are contaminated by noise but based on pre-averaged returns. In this experiment the variance is assumed to be constant within each 1-minute interval.

estimated on contaminated log-returns. The results clearly suggest that microstructure noise deteriorates the fit of MIDAS-RV models when using very high-frequency returns. The best variance forecasts are obtained for a frequency of the regressors between 15 seconds and five minutes, but the average QLIKE is about 40% greater than in the case without noise.

The (black) dashed line corresponds to the case where the MIDAS-RV is estimated on contaminated but pre-averaged log-returns. Pre-averaging proves to be useful in the context of MIDAS models, especially when relying on data sampled at frequencies higher than 15 minutes. Interestingly, for frequencies between 15 seconds and five minutes, the QLIKE of this model is stable and does not blow up, as in the case of the MIDAS-RV model estimated on contaminated log-returns.

IV. Empirical Application

In addition to microstructure noise, it has also been largely documented that high-frequency log-returns are characterized by the presence of strong intraday diurnal pattern in variance and jumps. Intraday diurnal pattern or periodicity, as it is called by several authors WOOD, R. A., T. H. McINISH, and J. K. ORD [1985], HARRIS, L. [1986], ANDERSEN, T. G., and T. BOLLERSLEV [1997b], ANDERSEN, T. G., and T. BOLLERSLEV [1998b], HECQ, A., S. LAURENT, and F. C. PALM [2012], can be defined as the cyclical pattern of variance within the trading day, i.e., the fact that variance is typically more important at the opening and closing of the trading day and lower in the middle of the day, while jumps correspond to large discontinuities in prices. Unlike intraday pattern, jumps are not regular (most of the time they appear as the result of unexpected news) and are known to greatly affect variance estimates and forecasts. For more details about the properties and the detection of jumps, see BATES, D. S. [1996], BARNDORFF-NIELSEN, O. E.,

and N. SHEPHARD [2004b], BARNDORFF-NIELSEN, O. E., and N. SHEPHARD [2006], LEE, S. S., and P. A. MYKLAND [2008], BOUDT, K., C. CROUX, and S. LAURENT [2011], LAHAYE, J., S. LAURENT, and C. J. NEELY [2011], among many others. In the empirical application, we propose to investigate the impact of these two additional features of the data on MIDAS-RV forecasting models in an application on two highly liquid assets, one exchange-traded fund (ETF) and one quoted share. The use of an ETF is justified by the increasing importance of these assets in the fund management industry.¹⁴

IV.1. Data

The dataset consists of tick-by-tick prices and quotations from the NYSE Trade and Quote (TAQ) database for Microsoft (MSFT) and one ETF (provided by SPDR ETFs) that tracks the S&P 500 index, spanning the period from September 2, 2004 to December 31, 2008. The price and quote series are reported every trading day from 9:30 am to 4:00 pm and rigorously cleaned using a set of baseline rules proposed by BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, and N. SHEPHARD [2009]. In order to avoid the effect of variance that comes from the overnight or holiday closures, all the variables are computed by using open-to-close data and focusing hence only on the effective trading day variance. The equally spaced intraday returns are subsequently derived from the high-frequency price series. The dataset thus contains 1,101 trading days with 390/78/39/26/13/6/2/1 observations per day of respectively 1-minute/5-minute/10-minute/15-minute/30-minute/1h05/3h15/1-day log-returns.

To compute the variance forecasts, we consider a rolling sample estimation scheme. The parameter estimates are updated every 50 days. For a fair comparison of the MIDAS models, the lag order k^{max} is fixed such that the information used to estimate the parameters covers a period of 70 days, regardless of the sampling frequency of the regressors. For instance, for a 5-minute regressor we use a k^{max} equal to 78×70 lags, where 78 represents the number of 5-minute intervals within a trading day.

Finally, the out-of-sample sample covers two years, i.e., 2007 and 2008. To test the robustness of the results on the state of financial markets, the sample is split into two periods. The first one corresponds to the relatively calm variance period of 2007, and the second to the financial crisis of 2008 (the end of this period corresponding to the peak of the subprime crisis).

IV.2. Optimal Sampling Frequency for MIDAS-RV on Raw Data

We first consider one-step-ahead forecasts of MIDAS-RV models estimated on raw data, sampled at different frequencies m_1 ranging between one minute to one day. Three forecasting horizons are considered for the endogenous variable $RV_{t+H,t}^{(m_2)}$, i.e., one day ($H = 1$), one week ($H = 5$) and two weeks ($H = 10$).

TABLE I reports the MCS test for both the calm and crisis periods. For each horizon H , the average QLIKE is reported along with the p -value of the MCS test. The entries in bold correspond to the best models selected by the MCS procedure. The striking result is that the loss function does not smoothly decrease with the sampling frequency and seems to indicate the

14. At the end of August 2015, 4 264 ETFs worldwide were managing USD 2 727 bn. According to the Greenwich Associates U.S. Exchange-Traded Funds Study, in 2014, 14% of the assets in the fund management industry were managed by ETFs.

TABLE I. – MIDAS sampling frequency puzzle

Panel A: Calm period (2007)													
	H=1				H=5				H=10				
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT		
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	
1min	0.2423	0.0945	0.2161	0.0157	0.2203	0.0663	0.0971	0.0932	0.2357	0.2327	0.1083	0.0095	
5min	0.2152	1.0000	0.1369	1.0000	0.1904	1.0000	0.0823	0.9635	0.2061	0.7692	0.0861	0.1837	
10min	0.2203	0.3562	0.1412	0.6192	0.1968	0.3828	0.0878	0.4825	0.2013	1.0000	0.0950	0.1770	
15min	0.2265	0.0945	0.1447	0.5429	0.2060	0.0663	0.0881	0.4825	0.2245	0.2327	0.0948	0.0985	
30min	0.2152	0.9989	0.1407	0.6192	0.1920	0.8240	0.0833	0.9635	0.2140	0.6019	0.0874	0.1837	
1h05	0.2254	0.3562	0.1447	0.6192	0.2017	0.3828	0.0933	0.4825	0.2187	0.6019	0.0930	0.1837	
3h15	0.2713	0.0232	0.1449	0.6192	0.2471	0.0663	0.0815	1.0000	0.2518	0.2327	0.0732	1.0000	
1day	0.2672	0.0945	0.1627	0.5429	0.2359	0.0663	0.0971	0.4825	0.2529	0.2327	0.0943	0.1837	

Panel B: Crisis period (2008)													
	H=1				H=5				H=10				
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT		
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	
1min	0.2185	0.1704	0.8640	0.1235	0.2337	0.1648	0.7399	0.1036	0.5308	0.0651	0.2231	0.1622	
5min	0.2033	1.0000	0.2891	0.1998	0.2182	1.0000	0.7152	0.1079	0.4144	0.1915	0.1914	1.0000	
10min	0.2166	0.1704	0.1848	1.0000	0.2340	0.1648	0.6791	0.1079	0.3245	0.4715	0.2839	0.1160	
15min	0.2172	0.1704	0.1960	0.5976	0.2306	0.1648	0.1820	1.0000	0.3146	1.0000	0.9670	0.0218	
30min	0.2317	0.0138	0.7694	0.1998	0.2662	0.0406	0.2397	0.1746	0.3408	0.1915	0.2670	0.1160	
1h05	0.2428	0.0138	0.2566	0.1998	0.3081	0.0406	0.2958	0.1079	0.3850	0.1915	0.5070	0.1160	
3h15	0.2468	0.1704	0.2612	0.1998	0.3437	0.0075	0.2485	0.1746	0.4352	0.0651	0.2213	0.3685	
1day	0.3305	0.0138	0.2539	0.1998	0.3190	0.0283	0.2719	0.1746	0.3911	0.1915	0.2419	0.1622	

Note: This table presents the MCS test results for the S&P 500 and Microsoft. The results are reported for three forecasting horizons, i.e., one day ($H = 1$), one week ($H = 5$) and two weeks ($H = 10$). The QLIKE is reported along with the p -value of the MCS test. The confidence level for the MCS test is set to $\alpha = 25\%$ and 10,000 bootstrap resamples are used, with block length of five observations, to obtain the distribution under the null of equal predictive accuracy. The set of the competing models includes eight MIDAS specifications with regressors sampled at a frequency ranging from one minute to one day.

presence of a “high-frequency wall”. In particular, the use of very high-frequency regressors leads to a deterioration in the quality of variance forecasts.

Consider the S&P 500 during the calm period (Panel A). The loss function has a convex shape and its minimum is reached for a predictor sampled at five minutes, whatever the forecasting horizon considered. Using 1-minute log-returns leads to a deterioration in the quality of the variance forecasts. This deterioration is statistically significant because MIDAS-RV estimated on 1-minute log-returns does not belong to the MCS set of optimal models. For the crisis period, the MCS test selects the 5-minute frequency as optimal for $H = 1$ and $H = 5$, and 10- and 15-minute frequencies for the two-week horizon. For Microsoft all the models but that estimated on 1-minute log-returns are found to be statistically equivalent and superior during the calm period (panel A) for $H = 1$. All in all, the previous results question the usefulness of very high-frequency data in the context of daily/weekly/bi-weekly MIDAS-RV models.

IV.3. Breaking the Wall

We can suggest several explanations for the existence of this “high-frequency wall”: i.e., an underlying DGP whose conditional variance is constant by pieces of e.g., one or five minutes; or the presence of intraday diurnal pattern, jumps and microstructure noise. While no solution for breaking this wall might exist for the first one, filtering the raw lag-returns might help to improve the performance of MIDAS-RV in the presence of intraday pattern, jumps and microstructure noise. This is precisely the purpose of this section.

IV.3.1. Intraday Diurnal Pattern

FIGURE 2 illustrates the intraday diurnal pattern in the variance for the S&P 500 and Microsoft series by plotting the average squared log-returns for each 1-minute, 5-minute, 30-minute and 1h05 interval, respectively. A clear U-shaped pattern is identifiable, as first noted by WOOD, R. A., T. H. MCINISH, and J. K. ORD [1985], suggesting that the variance is systematically high at the opening, declines to a low point at midday and then increases at the end of the trading day.

To estimate the intraday pattern in variance we rely on the non-parametric weighted standard deviation (WSD) of BOUDT, K., C. CROUX, and S. LAURENT [2011], a non-parametric estimator that is robust to additive jumps. If r_i denotes a raw return (sampled at a certain frequency), the corresponding intraday pattern adjusted return is obtained by dividing r_i by \hat{f}_i^{WSD} , i.e., r_i / \hat{f}_i^{WSD} , where \hat{f}_i^{WSD} is the estimated WSD of BOUDT, K., C. CROUX, and S. LAURENT [2011] for the i^{th} return.

IV.3.2. Jumps

To filter out the jumps in the regressors of the MIDAS-RV model, we first apply a modified version of the jump test of LEE, S. S., and P. A. MYKLAND [2008] proposed by BOUDT, K., C. CROUX, and S. LAURENT [2011]. More specifically, we assume that the log-price process $\log p(s)$ follows a *Brownian SemiMartingale with Finite Activity Jumps* (BSMFJA) diffusion $dp(s) = \mu(s)ds + \sigma(s)dw(s) + \kappa(s)dq(s)$, where $\mu(s)$ is the drift, $\sigma(s)$ is the spot volatility, $w(s)$ is a standard Brownian motion, the occurrence of jumps is governed by a finite activity counting process $q(s)$ and the size of the jumps is given by $\kappa(s)$.

The idea behind the jump test of LEE, S. S., and P. A. MYKLAND [2008] is that in the absence of jumps, instantaneous returns are increments of Brownian motion and, therefore, standardized returns that are too large to plausibly come from a standard Brownian motion must reflect jumps. In their original paper, LEE, S. S., and P. A. MYKLAND [2008] standardize every intraday return r_i by a robust estimate of the spot volatility, denoted \hat{s}_i , that assumes that the volatility is constant on a local window spanning between several hours to one or two days before or around the tested return. Their original statistic for jumps is $J_i = \frac{|r_i|}{\hat{s}_i}$, where \hat{s}_i is the averaged bi-power variation belonging to the local window. To control for the size of the multiple jump tests LEE, S. S., and P. A. MYKLAND [2008] use the extreme value theory result that the maximum of n i.i.d. realizations of the absolute value of a standard normal random variable is asymptotically (for $n \rightarrow \infty$) Gumbel distributed. More specifically, in the absence of jumps the probability that the maximum of any set of n J-statistics exceeds $g_{n,\alpha} = -\log(-\log(1 - \alpha))b_n + c_n$, with $b_n = 1 / \sqrt{2 \log n}$ and $c_n = (2 \log n)^{1/2} - [\log \pi + \log(\log n)] / [2(2 \log n)^{1/2}]$ is about α .

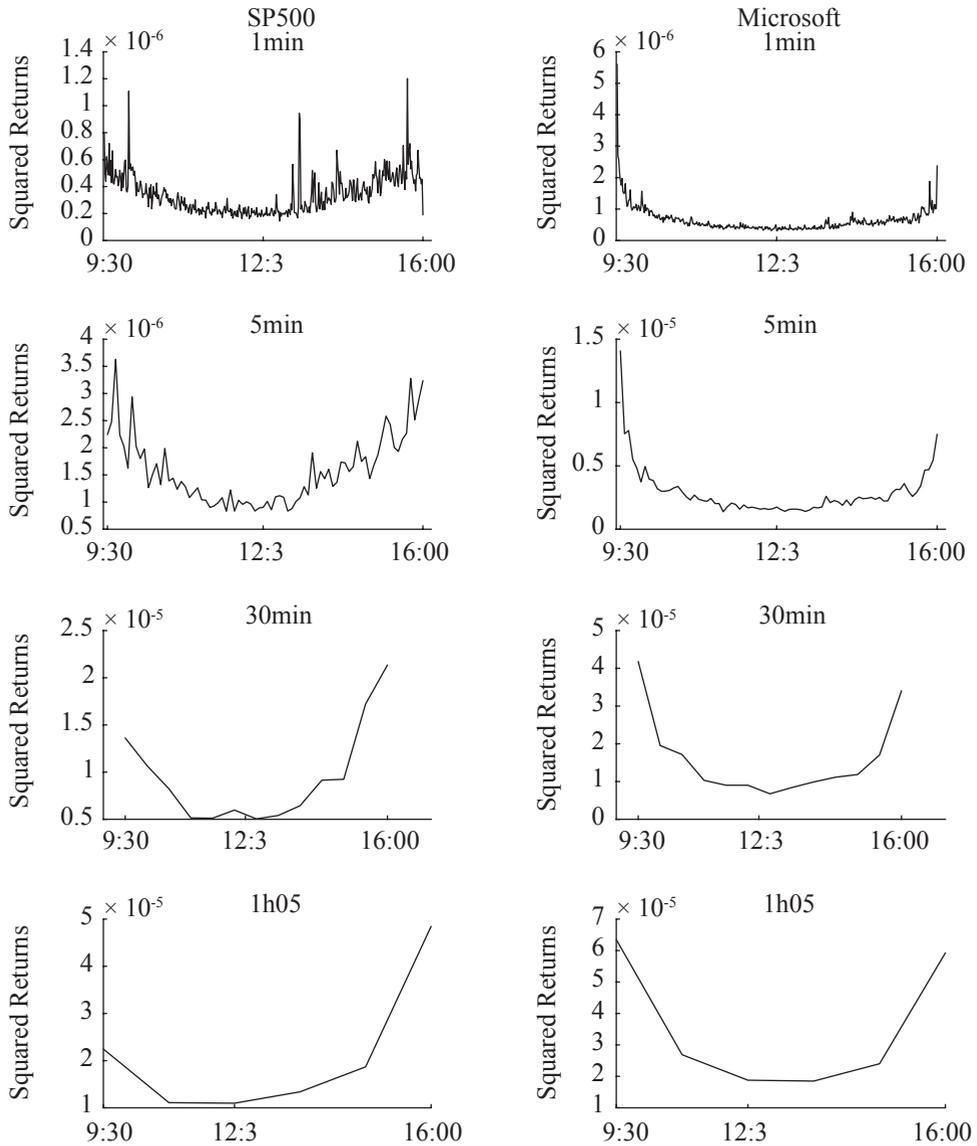


FIGURE 2. – Intraday diurnal pattern

Note: This figure displays the average squared log-returns for each 1-minute, 5-minute, 30-minute and 1h05 interval, for the S&P 500 and Microsoft.

LEE, S. S., and P. A. MYKLAND [2008]’s proposal is that all returns for which the J test statistic exceeds this threshold $g_{n,\alpha}$ should be declared to be affected by jumps. In the application, we set $\alpha = 1\%$ and n to the total number of observations in the sample.

However, for such long windows, the assumption of constant volatility is at odds with the overwhelming empirical evidence that the intraday variation in market activity causes intraday volatility to be strongly time-varying and even displays discontinuities (see FIGURE 2). For this reason, we implemented the modified version proposed by BOUDT, K., C. CROUX, and S. LAURENT

[2011] that accounts for the presence of intraday diurnal pattern, i.e., $FJ_i^{WSD} = \frac{|r_i|}{\hat{f}_i^{WSD} \hat{s}_i^{WSD}}$,

where \hat{f}_i^{WSD} is the estimated WSD of BOUDT, K., C. CROUX, and S. LAURENT [2011] for the i^{th} return (which is standardized such that its square has mean one in the local window).

Periodicity and jumps adjusted returns are computed as $(r_i / \hat{f}_i^{WSD}) \times I(FJ_i^{WSD} < g_{n,1\%}) + I(FJ_i^{WSD} > g_{n,1\%})$, where $I(\cdot)$ is an indicator function.

IV.3.3. Microstructure noise

Another disadvantage of using high-frequency data is that, at these frequencies, the true price process is likely to be contaminated by microstructure effects arising from market frictions, such as the bid-ask bounce or the discreteness of prices. This phenomenon produces spurious variations in asset prices and induces autocorrelation in high-frequency log-returns (see HANSEN, P. R., and A. LUNDE [2006b]; ZHOU, B. [1996]; AÏT-SAHALIA, Y., P. A. MYKLAND, and L. ZHANG [2005]; AÏT-SAHALIA, Y., and J. YU [2009]). The consequence of this noise on the realized variance is known (i.e., it is upward biased) but its impact on MIDAS-RV has not been investigated so far. We argue that this noise has an impact on the optimal frequency of the variance predictors when relying on raw data, and therefore leads to a loss of information.

To account for the presence of microstructure noise in the context of non-parametric variance estimators, it is standard practice to pre-filter high-frequency log-returns using the pre-averaging technique introduced by PODOLSKIJ, M., M. VETTER, ET AL. [2009] and JACOD, J., Y. LI, P. A. MYKLAND, and M. PODOLSKIJ, and M. VETTER [2009].¹⁵ To the best of our knowledge, pre-averaging has never been used in the context of MIDAS models.

Instead of noisy intraday returns (r_t), the authors suggest using pre-averaged returns (\tilde{r}_t) which, by the law of large numbers, asymptotically lose the noise component. More precisely, \tilde{r}_t is approximated by an average of staggered returns r_t in a neighborhood of t , the noise being thus averaged away. The pre-averaging approach depends on a bandwidth parameter, or window length, that increases with the sample and indicates the weighting scheme to be put into effect. The order of the window size is chosen to lead to optimal convergence rates ($n^{-1/4}$).

The balanced pre-averaging has been applied on 1- and 5-minute returns previously filtered for intraday diurnal pattern and jumps, since it delivers the best rate of convergence, according to CHRISTENSEN, K., S. KINNEBROCK, and M. PODOLSKIJ [2010].

IV.3.4. Results

FIGURE 3 displays the filtered 1-minute return series for the S&P 500 and Microsoft, adjusted for the intraday diurnal pattern and jumps. The correction procedure purges the intraday pattern, identifies and smoothes the jumps, but preserves the variance dynamics. The procedure is also applied to the 5-minute return series.

15. For more details on this technique, see Appendix B.

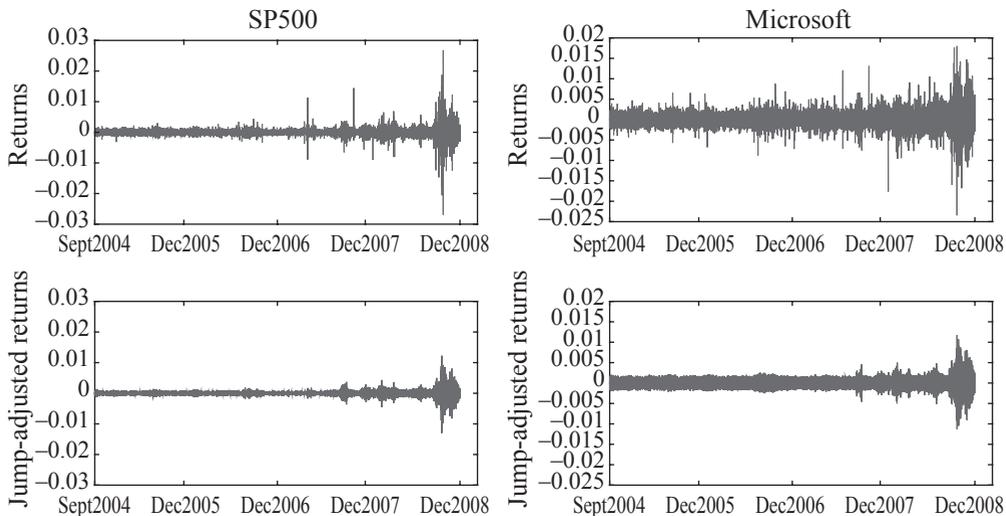


FIGURE 3. – Intraday returns and intraday jump-adjusted returns

Note: This figure displays the 1-minute intraday return series (in red) and the 1-minute return series filtered for intraday pattern and jumps (in blue) LEE, S. S., and P. A. MYKLAND [2008], BOUDT, K., C. CROUX, and S. LAURENT [2011].

The MCS test is subsequently applied on the MIDAS out-of-sample obtained with both filtered and unfiltered data. The results are reported in TABLE II. We notice a significant improvement in the MIDAS variance forecasts when using high-frequency predictors filtered for intraday diurnal pattern and jumps.

For the calm period (panel A), the S&P 500 forecasts obtained with filtered (both for jumps and diurnal pattern) 1-minute predictors always belong to the set of superior forecasting models as identified by MCS. During the crisis period, similar results are obtained with filtered 5-minute regressors for short horizons ($H = 1$ or $H = 5$). These results prove the importance of using the filtered data, especially for short forecasting horizons.

FIGURES 4 and 5 display the average QLIKE for the calm and crisis periods, and the three forecasting horizons, for both the S&P 500 and Microsoft. First, we remark that the gains related to the use of filtered data for intraday diurnal pattern are generally lower than the gains related to data filtered both for intraday pattern and jumps. Second, considering filtered data during a relatively calm period, we obtain a loss function which smoothly decreases with the sampling frequency m_1 as in the Monte Carlo experiment. Note that when the loss function pattern becomes flat, or decreases slowly (as result of the filtering techniques), the variance forecasts have statistically the same quality. This means that very high-frequency data predictors (e.g., 1-min data) are equally preferred to lower frequency predictors (e.g., 5-min to 1h05 data). However, the higher costs of very high-frequency data might discourage traders from using it, since variance forecasts are similar.

To complete our analysis, we also perform MIDAS variance forecasts based on regressors filtered for intraday diurnal pattern, jumps and microstructure noise (through the pre-averaging technique). The results are available in Appendix C and Appendix D. For instance, we observe that 1-minute pre-averaged regressors improve Microsoft variance forecasts during both the calm and crisis periods. These results apply also for the S&P 500 variance forecasts

TABLE II. – MIDAS sampling frequency puzzle: intraday pattern and jump adjustments

Panel A: Calm period (2007)													
	H=1				H=5				H=10				
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT		
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	
1min	0.2423	0.1403	0.2161	0.0084	0.2203	0.1032	0.0971	0.1282	0.2357	0.1433	0.1083	0.0119	
1min	0.2320	0.1403	0.1521	0.0584	0.2078	0.5087	0.0993	0.0944	0.2228	0.6283	0.1104	0.0105	
Per.Adj													
1min	0.2141	0.9962	0.1631	0.0084	0.1823	1.0000	0.1105	0.0083	0.2053	0.9284	0.1196	0.0018	
Jumps.													
Adj													
5min	0.2152	0.9482	0.1369	1.0000	0.1904	0.7536	0.0823	0.9896	0.2061	0.9284	0.0861	0.2639	
RAW													
5min	0.2138	1.0000	0.1399	0.7404	0.1884	0.8198	0.0831	0.9896	0.2063	0.9284	0.0873	0.2639	
Per.Adj													
5min	0.2152	0.9962	0.1380	0.7762	0.1850	0.8198	0.0830	0.9896	0.2029	0.9324	0.0888	0.2639	
Jumps.													
Adj													
10min	0.2203	0.3497	0.1412	0.7404	0.1968	0.5087	0.0878	0.5991	0.2013	1.0000	0.0950	0.2542	
15min	0.2265	0.1403	0.1447	0.6187	0.2060	0.1032	0.0881	0.5603	0.2245	0.2596	0.0948	0.1450	
30min	0.2152	0.9962	0.1407	0.7413	0.1920	0.8198	0.0833	0.9896	0.2140	0.8609	0.0874	0.2639	
1h05	0.2254	0.3497	0.1447	0.7404	0.2017	0.5589	0.0933	0.5603	0.2187	0.6557	0.0930	0.2639	
3h15	0.2713	0.0414	0.1449	0.7404	0.2471	0.1032	0.0815	1.0000	0.2518	0.1433	0.0732	1.0000	
1day	0.2672	0.1403	0.1627	0.6187	0.2359	0.1032	0.0971	0.5991	0.2529	0.2596	0.0943	0.2639	
Panel B: Crisis period (2008)													
	H=1				H=5				H=10				
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT		
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	
1min	0.2185	0.0205	0.8640	0.0189	0.2337	0.0354	0.7399	0.0113	0.5308	0.0944	0.2231	0.0616	
1min	0.2112	0.0390	0.1635	0.1167	0.2227	0.3369	0.1465	0.1439	0.2975	0.0999	0.1679	1.0000	
Per.Adj													
1min	0.2042	0.0390	0.1587	1.0000	0.2026	0.4552	0.1362	1.0000	0.3172	0.0999	0.1758	0.6321	
Jumps.													
Adj													
5min	0.2033	0.0390	0.2891	0.0189	0.2182	0.3369	0.7152	0.0113	0.4144	0.0999	0.1914	0.2172	
RAW													
5min	0.2008	0.0390	0.1673	0.1041	0.2142	0.4552	0.1534	0.1439	0.2737	1.0000	0.4426	0.0145	
Per.Adj													
5min	0.1842	1.0000	0.1598	0.8721	0.1956	1.0000	0.1528	0.1439	0.3073	0.0999	0.1735	0.6321	
Jumps.													
Adj													
10min	0.2166	0.0205	0.1848	0.0189	0.2340	0.0354	0.6791	0.0113	0.3245	0.0999	0.2839	0.0240	
15min	0.2172	0.0390	0.1960	0.1041	0.2306	0.0354	0.1820	0.0309	0.3146	0.0999	0.9670	0.0145	
30min	0.2317	0.0204	0.7694	0.0189	0.2662	0.0212	0.2397	0.0113	0.3408	0.0999	0.2670	0.0145	
1h05	0.2428	0.0204	0.2566	0.0189	0.3081	0.0204	0.2958	0.0113	0.3850	0.0999	0.5070	0.0145	
3h15	0.2468	0.0205	0.2612	0.0189	0.3437	0.0037	0.2485	0.0177	0.4352	0.0403	0.2213	0.2172	
1day	0.3305	0.0204	0.2539	0.0189	0.3190	0.0120	0.2719	0.0113	0.3911	0.0999	0.2419	0.0616	

Note: This table presents the MCS test results obtained for two assets (the S&P 500 and Microsoft) during both calm and crisis periods. The results are reported for three forecasting horizons, namely one day ($H = 1$), one week ($H = 5$) and two weeks ($H = 10$). For each we present the average value of the QLIKE loss function along with the corresponding p -value resulting from the MCS test. The confidence level for the MCS test is set to $\alpha = 25\%$ and 10,000 bootstrap resamples are used, with block length of five observations, to obtain the distribution under the null of equal predictive accuracy. The set of the competing variance models includes eight MIDAS specifications with regressors sampled at a frequency ranging from one minute to one day, as well as four MIDAS models with 1- and 5-minute regressors adjusted for intraday diurnal pattern and both intraday pattern and jumps.

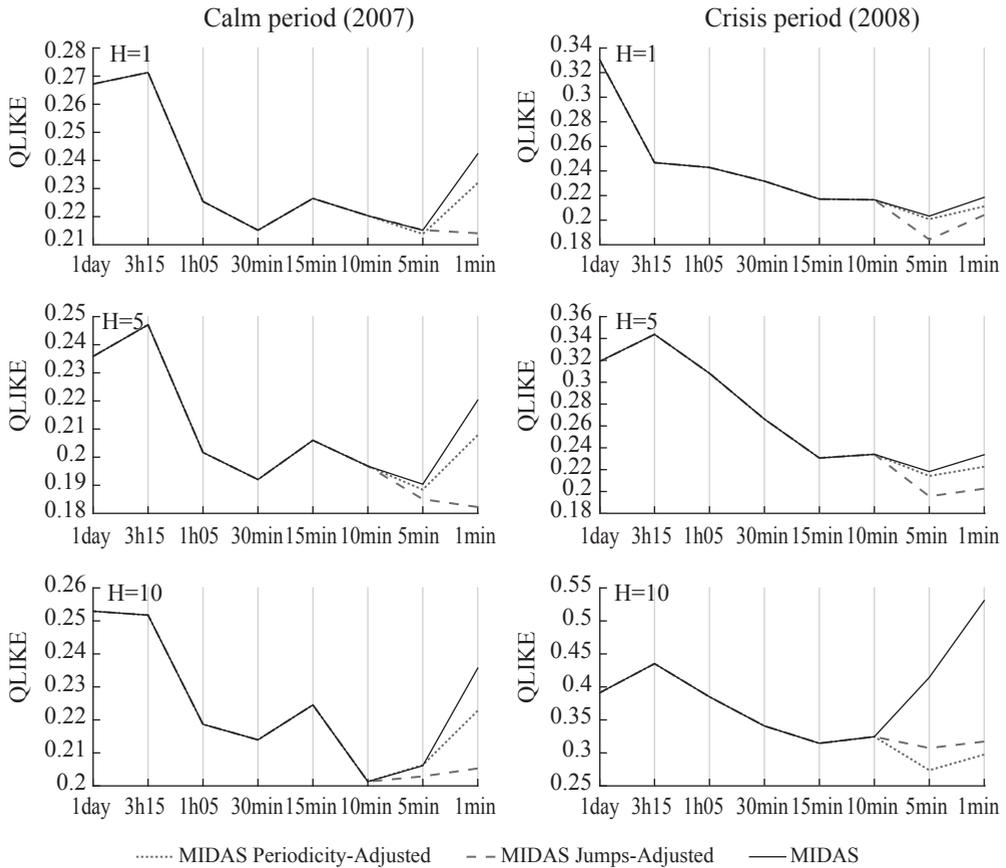


FIGURE 4. – S&P 500 average QLIKE

Note: The S&P 500 - This figure displays the average QLIKE for the MIDAS-RV forecasts for various sampling frequencies (m_t) of the predictors, and for the three forecasting horizons. The left panel corresponds to the calm period (2007) and the right panel to the crisis period (2008). The solid blue line corresponds to the MIDAS-RV model with raw data sampled at frequencies 1-min to 1-day. The green and red dotted lines correspond respectively to the MIDAS-RV models on intraday diurnal pattern, and both jumps and intraday pattern filtered log-returns sampled at frequencies between 1-min and 5-min.

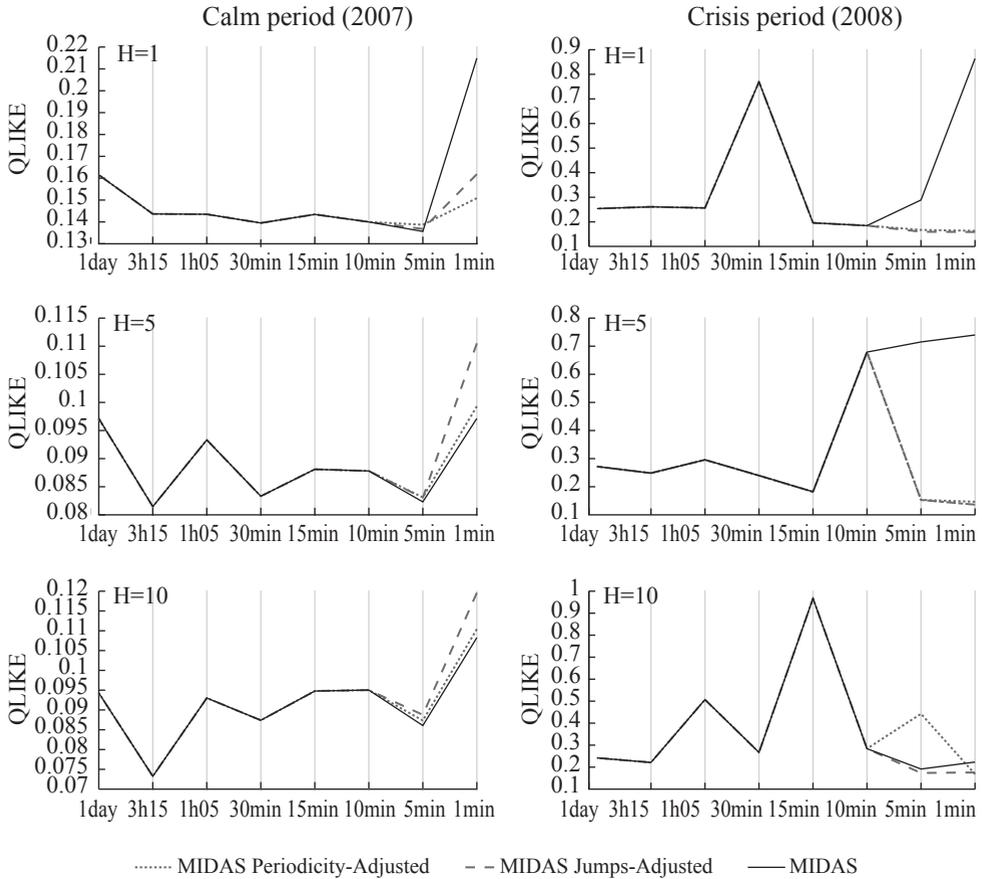


FIGURE 5. – Microsoft average QLIKE

Note: See FIGURE 4.

at short forecasting horizons. They become more puzzling for long forecasting horizons (e.g., two weeks), as well as for the crisis period. For instance, the QLIKE functions for weekly and bi-weekly Microsoft variance forecasts (left-middle and left-bottom of FIGURE 5) show that the high-frequency wall persists even after filtering.¹⁶ However, we observe that the 1-min raw MIDAS model is never included in the subset of best models, according to the MCS test, regardless of the period under analysis.

IV.4. Robustness Check

In this section we examine the robustness of our results. Three main exercises are considered, namely: (i) the identification of the pure effect of sampling frequency; (ii) the comparison of

16. As noted by the anonymous referee, this is not a problem per se, since our objective is to identify the problem, even if standard filtering techniques are not able to solve it.

the performances of MIDAS-RV models with existing variance forecasting models; (iii) the logarithmic MIDAS-RV setup.

IV.4.1. *The Pure Effect of Sampling Frequency*

The first robustness check exercise aims at highlighting the pure effect of sampling frequency. To achieve this goal, we first estimated a MIDAS-RV model with regressors sampled at one minute and then imposed the weights obtained previously to lower frequency regressors of variance. These weights would not be optimal, but it is a good way to isolate the differences in predictability attributed to the difference in sampling frequency versus the change in optimal weights, with respect to different data samples.¹⁷

TABLE III. – Same weights

	ϕ_{H,m_t}		QLIKE	
	S&P 500	MSFT	S&P 500	MSFT
1min	334.4542	332.3547	0.1547	0.1064
5min	73.1739	76.4844	0.1513	0.1124
10min	36.6549	40.2871	0.1600	0.1230
15min	25.5069	27.8444	0.1596	0.1285
30min	12.9734	14.0814	0.1581	0.1457
1h05	5.7102	6.3409	0.1891	0.1908
3h15	1.6805	1.9722	0.2866	0.3689
1day	0.8570	0.8085	0.3401	1.1116

Note: This table presents the estimates for the scale parameter, ϕ_{H,m_t} , and the QLIKE loss function for the daily MIDAS-RV model with regressors sampled at frequencies ranging from one minute to one day. The weights from the model using 1-minute returns are imposed on lower frequency returns.

TABLE III presents the estimates for the scale parameter, ϕ_{H,m_t} , and the QLIKE loss function for the daily MIDAS-RV model with regressors sampled at frequency ranging from one minute to one day. As previously indicated, the weights from the model using 1-minute regressors are imposed on lower frequency regressors. We observe the decreasing pattern of the parameter estimates for the two assets. The same declining pattern is also noticed for the QLIKE loss function. The more frequent information we use, the more we gain in precision.

IV.4.2. *MIDAS and Other Competing Variance Models*

In this second robustness check exercise, we compare the predictive accuracy of the MIDAS-RV forecasts with those obtained for four widely used variance models based on daily and/or intradaily data, i.e., the GARCH model, the Generalized Autoregressive Score (GAS), the Heterogeneous Autoregressive Realized Variance-based model (HAR-RV) and the HAR-RV adjusted for jumps (HAR-RV-J).

17. We thank the anonymous referee for this suggestion.

i) The first competing model is the popular GARCH(1,1) model, pioneered by ENGLE, R. F. [1982] and BOLLERSLEV, T. [1986], i.e.:

$$r_{t+1,t} = c + z_{t+1,t} \sqrt{h_{t+1,t}}, \quad (8)$$

$$h_{t+1,t} = \alpha_0 + \alpha_1(r_{t,t-1} - c)^2 + \beta_1 h_{t,t-1}. \quad (9)$$

ii) The second model is the GAS model, recently introduced by HARVEY, A. C. [2013] and CREAL, D. D., S. J. KOOPMAN, and A. LUCAS [2013]. This model is designed to better deal with large outliers. We consider the Student GAS specification where the one step-ahead conditional variance is defined as follows:

$$h_{t+1,t} = w_0 + a_1 u_{t,t-1} h_{t,t-1} + \phi_1 h_{t,t-1}, \quad (10)$$

with $u_{t,t-1} = ((v+1)z_{t,t-1}^2) / (v-2 + z_{t,t-1}^2) - 1$, and $z_t \sim t(0,1,v)$. Notice that for the GARCH and GAS models, the variance forecasts for $H > 1$ are obtained as $\sum_{i=1}^H h_{t+i,t}$ and not directly from $r_{t+H,t}$ as opposed to the MIDAS-RV model.

iii) The third competing model is the HAR-RV model, proposed by CORSI, F. [2009]:

$$RV_{t+1,t} = \alpha_0 + \alpha_1 RV_{t,t-1} + \alpha_2 RV_{t,t-1}^w + \alpha_3 RV_{t,t-1}^m + \varepsilon_{t+1}, \quad (11)$$

where $RV_{t+1,t}$ is the daily realized variance (Eq. 3) and by convention, $RV_{t,t-1}^w = \frac{1}{5} \sum_{i=0}^4 RV_{t-i,t-i-1}$ and $RV_{t,t-1}^m = \frac{1}{22} \sum_{i=0}^{21} RV_{t-i,t-i-1}$. This model is conceived as an additive cascade of different variance components defined over different time horizons of one day, one week (w), and one month (m), respectively. The HAR-RV is therefore a constrained version of the MIDAS-RV model with intradaily squared return regressors and a particular weight structure. Indeed, given the definition of the realized variance, Eq. (11) can be rewritten as a weighted sum of past observations of the intraday squared returns. For more details, see Appendix E.

iv) ANDERSEN, T. G., T. BOLLERSLEV, and F. X. DIEBOLD [2007] extended the classical HAR-RV framework by taking into account the lagged effect of jumps. The HAR-RV-J model (where J stands for jumps) is formally defined as follows:

$$RV_{t+1,t} = \alpha_0 + \alpha_1 RV_{t,t-1} + \alpha_2 RV_{t,t-1}^w + \alpha_3 RV_{t,t-1}^m + \gamma_1 J_{t,t-1} + \gamma_2 J_{t,t-1}^w + \gamma_3 J_{t,t-1}^m + \varepsilon_{t+1} \quad (12)$$

where $J_{t,t-1} = I_t \times (RV_{t,t-1} - BV_{t,t-1})$ is a random variable that is nonzero for the intervals in which jumps occur and zero otherwise, BV is the daily realized bipower variation BARN-DORFF-NIELSEN, O. E., and N. SHEPHARD [2004b] which is defined as:

$$BV_t = \mu_1^{-2} \sum_{l=2}^m |r_{t,l}| |r_{t,l-1}|, \quad (13)$$

with $\mu_1 = \sqrt{(2/\pi)} \approx 0.79788$, and $I_t \equiv I(Z_t > \Phi_{0.999})$, where Z_t is defined as:

$$Z_t = \frac{m^2 (RV_{t,t-1} - BV_{t,t-1}) RV_{t,t-1}^{-1}}{[(\mu_1^{-4} + 2\mu_1^{-2} - 5) \max\{1, TQ_{t,t-1}(m) BV_{t,t-1}^{-2}\}]^{1/2}} \quad (14)$$

with $TQ_{t,t-1}$ the tri-power quarticity, a robust estimator of the integrated variance, and $\Phi_{0.999}$ the 99.9% quantile of the standard normal distribution.¹⁸

The MCS procedure is now applied on 18 models, namely the eight MIDAS models with regressors sampled between 1-min and one day, the six MIDAS specifications with 1- and 5-minute regressors adjusted for intraday diurnal pattern, jumps and/or microstructure noise, and the four competing variance models (HAR-RV, HAR-RV-J, Student GAS, GARCH). The results are summarized in TABLE IV. The broad conclusion is that the MIDAS models provide (at least for these two assets) comparable, or even better, variance forecasts than the other competing models. In terms of the loss function, the models are dominated during the calm period by the MIDAS models, except for the daily Microsoft forecasts. During the calm period, we find that the forecasts provided by different MIDAS specifications are statistically comparable to those from the HAR-RV and HAR-RV-J models. The GAS model provides comparable forecasts only in the case of the S&P 500 for a forecasting horizon of two weeks. During the crisis, the best forecasts are generally provided by the MIDAS models with 1- or 5-minute filtered predictors, and the cluster of superior forecasting models no longer includes the HAR-RV and HAR-RV-J models. For longer horizons, the GAS and the GARCH provide similar results to those obtained with the MIDAS model. These findings confirm the intuition that high-frequency data can be used to successfully forecast variance, provided that these data are filtered for intraday pattern and jumps. For these two assets, MIDAS models outperform in many cases standard variance models such as the GARCH model, or even the HAR-RV, HAR-RV-J or GAS models.

IV.4.3. Logarithmic Framework

The most frequent criticism on both MIDAS and HAR-RV models concerns the absence of some constraints ensuring the positivity of the variance process. A straightforward solution consists in predicting the logarithm of the variance proxy, i.e.,

Log-MIDAS:

$$\log(RV_{t+H,t}^{(m_2)}) = \mu_{H,m_1} + \phi_{H,m_1} \Omega_{H,m_1} (L^{1/m_1}) \log(X_{t,t-1/m_1}^{(m_1)}) + \varepsilon_t. \quad (15)$$

Log-HAR-RV:

$$\log(RV_{t+H,t}) = \alpha_0 + \alpha_1 \log(RV_{t,t-1}) + \alpha_2 \log(RV_{t,t-1}^w) + \alpha_3 \log(RV_{t,t-1}^m) + \varepsilon_{t+1}. \quad (16)$$

Log-HAR-RV-J:

$$\begin{aligned} \log(RV_{t+H,t}) = & \alpha_0 + \alpha_1 \log(RV_{t,t-1}) + \alpha_2 \log(RV_{t,t-1}^w) + \alpha_3 \log(RV_{t,t-1}^m) \\ & + \gamma_1 \log(J_{t,t-1} + 1) + \gamma_2 \log(J_{t,t-1}^w + 1) + \gamma_3 \log(J_{t,t-1}^m + 1) + \varepsilon_{t+1}. \end{aligned} \quad (17)$$

To forecast the logarithm of the realized measure of variance, we follow exactly the same procedure as for the level of variance. Next, in order to compare the log-variance forecasts with the level of variance proxy, the following transformation is required ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, and P. LABYS [2003]:

18. $TQ_t \equiv m \mu_{4/3}^{-3} \sum_{t=3}^m |r_{t,t}|^{4/3} |r_{t,t-1}|^{4/3} |r_{t,t-2}|^{4/3}$, where $\mu_{4/3} \equiv 2^{2/3} \Gamma(7/6) \Gamma(1/2)^{-1}$.

TABLE IV. – Comparing competing variance models

Panel A: Calm period (2007)													
	H=1				H=5				H=10				
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT		
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	
1min	0.2423	0.1617	0.2161	0.0103	0.2203	0.1444	0.0971	0.1726	0.2357	0.1830	0.1083	0.0046	
1min Per.Adj	0.2320	0.1617	0.1521	0.0704	0.2078	0.6375	0.0993	0.1285	0.2228	0.6896	0.1104	0.0039	
1min Jumps. Adj	0.2141	0.9961	0.1631	0.0103	0.1823	1.0000	0.1105	0.0059	0.2053	0.9262	0.1196	0.0031	
1min Jumps. Adj_ Preav	0.2155	0.1617	0.1560	0.0103	0.1900	0.7983	0.1000	0.0059	0.4243	0.1830	0.1054	0.0039	
5min	0.2152	0.9773	0.1369	0.7738	0.1904	0.7983	0.0823	0.9888	0.2061	0.9262	0.0861	0.2702	
5min Per.Adj	0.2138	1.0000	0.1399	0.7738	0.1884	0.8186	0.0831	0.9888	0.2063	0.9262	0.0873	0.2702	
5min Jumps. Adj	0.2152	0.9773	0.1380	0.7738	0.1850	0.8186	0.0830	0.9888	0.2029	0.9320	0.0888	0.2702	
5min Jumps. Adj_ Preav	0.2149	0.9961	0.1560	0.0213	0.1899	0.7983	0.1001	0.0059	0.4286	0.1830	0.1052	0.0039	
10min	0.2203	0.1617	0.1412	0.7738	0.1968	0.6375	0.0878	0.5281	0.2013	1.0000	0.0950	0.2100	
15min	0.2265	0.1617	0.1447	0.3101	0.2060	0.1444	0.0881	0.5148	0.2245	0.3046	0.0948	0.1429	
30min	0.2152	0.9961	0.1407	0.7738	0.1920	0.8186	0.0833	0.9888	0.2140	0.8471	0.0874	0.2702	
1h05	0.2254	0.1617	0.1447	0.6354	0.2017	0.7171	0.0933	0.5148	0.2187	0.6896	0.0930	0.2702	
3h15	0.2713	0.0630	0.1449	0.7738	0.2471	0.1444	0.0815	1.0000	0.2518	0.1830	0.0732	1.0000	
1day	0.2672	0.1617	0.1627	0.3101	0.2359	0.1444	0.0971	0.5148	0.2529	0.1830	0.0943	0.2702	
HAR- RV	0.2176	0.1617	0.1345	1.0000	0.1941	0.7726	0.0868	0.5148	0.2172	0.6896	0.0943	0.2100	
HAR-RV jumps	0.2187	0.1617	0.1359	0.7738	0.1973	0.7171	0.0883	0.5148	0.2226	0.4406	0.0960	0.1429	
GARCH	0.3240	0.0630	0.2208	0.0027	0.2849	0.1444	0.1677	0.0002	0.2935	0.1830	0.1747	>0.0001	
GAS	0.3161	0.0630	0.1884	0.0103	0.2669	0.1444	0.1292	0.0059	0.2688	0.3046	0.1296	0.0046	

Continued on next page

TABLE IV. – Comparing competing variance models (*Continued*)

Panel B: Crisis period (2008)												
	H=1				H=5				H=10			
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT	
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value
1min	0.2185	0.0304	0.8640	0.0265	0.2337	0.0539	0.7399	0.0121	0.5308	0.0315	0.2231	0.0883
1min Per.Adj	0.2112	0.0698	0.1635	0.1823	0.2227	0.4921	0.1465	0.1455	0.2975	0.1162	0.1679	1.0000
1min Jumps. Adj	0.2042	0.0698	0.1587	1.0000	0.2026	0.6140	0.1362	1.0000	0.3172	0.1162	0.1758	0.6477
1min Jumps. Adj_ Preav	0.2430	0.0304	0.1698	0.1823	0.2596	0.0258	0.1725	0.0322	0.3517	0.1162	0.1918	0.4436
5min	0.2033	0.0698	0.2891	0.0265	0.2182	0.4921	0.7152	0.0121	0.4144	0.1162	0.1914	0.3105
5min Per.Adj	0.2008	0.0698	0.1673	0.1339	0.2142	0.6140	0.1534	0.1455	0.2737	0.3766	0.4426	0.0207
5min Jumps. Adj	0.1842	1.0000	0.1598	0.8786	0.1956	1.0000	0.1528	0.1455	0.3073	0.1162	0.1735	0.6477
5min Jumps. Adj_ Preav	0.2418	0.0304	0.1780	0.0271	0.2627	0.0258	0.1757	0.0239	0.3502	0.1162	0.1943	0.3105
10min	0.2166	0.0304	0.1848	0.0265	0.2340	0.0539	0.6791	0.0121	0.3245	0.1162	0.2839	0.0207
15min	0.2172	0.0698	0.1960	0.0271	0.2306	0.0539	0.1820	0.0322	0.3146	0.1162	0.9670	0.0207
30min	0.2317	0.0304	0.7694	0.0265	0.2662	0.0258	0.2397	0.0121	0.3408	0.1162	0.2670	0.0207
1h05	0.2428	0.0304	0.2566	0.0265	0.3081	0.0258	0.2958	0.0121	0.3850	0.0315	0.5070	0.0207
3h15	0.2468	0.0304	0.2612	0.0265	0.3437	0.0058	0.2485	0.0153	0.4352	0.0315	0.2213	0.3105
1day	0.3305	0.0304	0.2539	0.0265	0.3190	0.0162	0.2719	0.0121	0.3911	0.0318	0.2419	0.0359
HAR- RV	0.2102	0.0698	0.1781	0.0271	0.2449	0.0539	0.1754	0.0153	0.3563	0.0318	0.2346	0.0207
HAR-RV jumps	0.2148	0.0698	0.1719	0.1339	0.2472	0.0539	0.1669	0.0717	0.3569	0.0315	0.2205	0.0359
GARCH	0.2291	0.0698	0.2205	0.0265	0.2282	0.4921	0.2041	0.0153	0.2723	0.2778	0.2208	0.3105
GAS	0.2363	0.0304	0.2447	0.0265	0.2173	0.6140	0.2142	0.0153	0.2348	1.0000	0.2226	0.3105

Note: This table presents the MCS test results obtained for two assets (the S&P 500 and Microsoft) during both calm and crisis periods. The results are reported for three forecasting horizons, namely one day ($H = 1$), one week ($H = 5$) and two weeks ($H = 10$). For each we present the average value of the QLIKE loss function along with the corresponding p -value resulting from the MCS test. The confidence level for the MCS test is set to $\alpha = 25\%$ and 10,000 bootstrap resamples are used, with block length of five observations, to obtain the distribution under the null of equal predictive accuracy. The set of the competing variance models includes eight MIDAS specifications with regressors sampled at a frequency ranging from one minute to one day, six MIDAS models with 1- and 5-minute regressors adjusted for intraday diurnal pattern, intraday pattern and jumps, and intraday pattern, jumps and microstructure noise, the HAR-RV, HAR-RV-J, GARCH and GAS models.

$$\widehat{RV}'_{t+H,t} = \exp\left(\log(\widehat{RV}_{t+H,t}) + \frac{1}{2}\text{Var}(e_{t+H,t})\right) \quad (18)$$

where $\log(\widehat{RV}_{t+H,t})$ is the forecast of the log of the realized variance and $\text{Var}(e_{t+H,t})$ is the variance of the forecasting errors.

The results of the MCS-based comparison procedure are reported in TABLE V. Once again, during the calm period, the standard models are generally dominated by the log-MIDAS models. For shorter forecasting horizons (one day and one week), the cluster also includes the log-HAR-RV-J model. Only in the case of the S&P 500, does the GAS model provide statistically comparable forecasts for a horizon of one and two weeks. During the crisis, the daily log-MIDAS-RV with 5-minute regressors pre-filtered for intraday diurnal pattern, jumps and microstructure noise has the smallest QLIKE. For the one-week forecasting horizon the best forecast fit is given by the log-HAR-RV-J model for both Microsoft and the S&P 500. The subset of superior forecasting models (as identified by the MCS) encompasses a smaller number of log-MIDAS specifications than in the calm period, the log-HAR-RV-J model (for one day and one week forecasting horizons) and the Student GAS model (for one and two week-ahead S&P 500 forecasts).

Another robustness check exercise consists in changing the measure of variance to be predicted. It is well-documented that the realized variance estimator may become biased and inconsistent in the presence of market microstructure noise. A large number of alternative proxies of variance (e.g., realized bipower variation, realized kernel, etc.) that deal with issues such as jumps and other market microstructure noise, have consequently been introduced by BARNDORFF-NIELSEN, O. E., and N. SHEPHARD [2004a], ZHANG, L. [2006], BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, and N. SHEPHARD [2008], HANSEN, P. R., and G. HOREL [2009], *inter alia*. To assess the robustness of our results, we also consider the realized kernel BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, and N. SHEPHARD [2008] as the dependent variable in our specifications and obtain similar results (Appendix F). Our results are also robust to the choice of the predictors of variance (e.g., intradaily bipower variation instead of squared intradaily returns). For a synthesis of all these extra findings see Appendix G.

V. Conclusion

This paper analyses the forecasting performance of MIDAS-RV models in which future variances are directly related to past intraday log-returns. These predictors are usually constructed from tick-by-tick data and, consequently, the econometrician needs to choose a sampling frequency. The question we raise is whether high frequency data is needed to forecast longer horizon variances.

The main findings of our study are the following. First, we show in a Monte Carlo simulation study that, to a world without jumps, intraday diurnal pattern in variance and microstructure noise, there is an advantage in using the highest available frequency for the predictors. The information content of very high-frequency data significantly improves the quality of the MIDAS forecasts. Second, when considering two highly liquid assets (namely Microsoft and the S&P 500) contaminated with typical market microstructure noise and intraday diurnal pattern, we find that the use of very high-frequency predictors may become problematic. In particular, we show that there may exist a “high-frequency wall”, i.e., a frequency limit above

TABLE V. – Log version - MCS Test

Panel A: Calm period (2007)												
	H=1				H=5				H=10			
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT	
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value
1min	0.2406	0.3927	0.1553	0.0793	0.2299	0.4631	0.0998	0.0862	0.2454	0.0472	0.1136	0.0186
1min Per.Adj	0.2350	0.8643	0.1546	0.0793	0.2294	0.4631	0.0997	0.3683	0.2057	1.0000	0.1253	0.0069
1min Jumps. Adj	0.2395	0.3927	0.1660	0.0793	0.2118	0.7967	0.1262	0.0034	0.2152	0.9008	0.1263	0.0011
1min Jumps. Adj_ Preav	0.2295	0.6360	0.1656	0.0793	0.2174	0.4805	0.1126	0.0618	1.0253	0.0014	0.1145	0.0186
5min	0.2251	0.8643	0.1403	0.9546	0.1987	0.7967	0.0862	0.8161	0.2149	0.9219	0.0915	0.6301
5min Per.Adj	0.2252	0.8643	0.1419	0.9142	0.2051	0.7967	0.0913	0.6182	0.2178	0.9008	0.0924	0.6301
5min Jumps. Adj	0.2246	0.8643	0.1409	0.9546	0.2097	0.7571	0.0902	0.6182	0.2079	0.9820	0.0948	0.5343
5min Jumps. Adj_ Preav	0.2105	1.0000	0.1517	0.0793	0.1931	1.0000	0.1022	0.0862	0.2305	0.9008	0.1056	0.0246
10min	0.2252	0.8643	0.1432	0.9142	0.2116	0.4805	0.0902	0.6182	0.2294	0.0472	0.0905	0.6301
15min	0.2324	0.4502	0.1451	0.6697	0.2180	0.4631	0.0893	0.5495	0.2289	0.6145	0.0906	0.5343
30min	0.2198	0.8643	0.1369	1.0000	0.2158	0.6218	0.0811	0.8161	0.2118	0.9820	0.0834	0.8619
1h05	0.2348	0.5145	0.1420	0.9546	0.2193	0.4805	0.0837	0.8161	0.2104	0.9820	0.0800	0.8792
3h15	0.2644	0.0592	0.1508	0.6697	0.2498	0.2145	0.0773	1.0000	0.2523	0.0472	0.0785	1.0000
1day	0.2973	0.0592	0.1805	0.0793	0.2385	0.3504	0.1069	0.4387	0.2287	0.9008	0.0880	0.8619
HAR- RV	0.2444	0.1056	0.1528	0.0793	0.2622	0.0124	0.1179	0.0014	0.3308	0.0014	0.1425	0.0007
HAR-RV jumps	0.2205	0.8643	0.1386	0.9546	0.2161	0.4805	0.0922	0.5495	0.2615	0.0375	0.1056	0.0186
GARCH	0.3240	0.0592	0.2208	0.0697	0.2849	0.3504	0.1677	0.0014	0.2935	0.0472	0.1747	0.0011
GAS	0.3161	0.0592	0.1884	0.0793	0.2669	0.3504	0.1292	0.0862	0.2688	0.0472	0.1296	0.0186

Continued on next page

TABLE V. – Log version - MCS Test (Continued)

Panel B: Crisis period (2008)													
	H=1				H=5				H=10				
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT		
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	
1min	0.3628	0.0001	0.1881	0.0953	0.5074	0.0024	0.2244	0.0182	0.6579	0.0011	0.3851	0.0047	
1min Per.Adj	0.3553	0.0001	0.2575	0.0139	0.5223	0.0024	0.2093	0.0182	0.4240	0.1263	0.4045	0.0039	
1min Jumps. Adj	0.3369	0.0001	0.2407	0.0276	0.4799	0.0115	0.3244	0.0182	0.6406	0.0024	0.4123	0.0047	
1min Jumps. Adj_ Preav	0.2926	0.0048	0.2369	0.0139	0.4875	0.0115	0.3434	0.0182	0.6936	0.0024	0.4136	0.0047	
5min	0.2530	0.0048	0.2085	0.0832	0.3928	0.0115	0.2029	0.0182	0.5136	0.0024	0.1682	1.0000	
5min Per.Adj	0.2544	0.0048	0.1712	0.0953	0.4154	0.0115	0.2683	0.0182	0.5072	0.0024	0.1737	0.3670	
5min Jumps. Adj	0.2083	0.0295	0.2011	0.0953	0.3872	0.0115	0.1986	0.0182	0.5550	0.0024	0.3325	0.0047	
5min Jumps. Adj_ Preav	0.1514	1.0000	0.1471	1.0000	0.2404	0.5737	0.2175	0.0182	0.3625	0.1263	0.2851	0.1848	
10min	0.2323	0.0295	0.1699	0.0953	0.2952	0.0115	0.1624	0.0634	0.4903	0.0024	0.3197	0.0047	
15min	0.1930	0.0295	0.1591	0.3878	0.2065	0.9076	0.1946	0.0182	0.4679	0.0580	0.1820	0.2582	
30min	0.1926	0.0295	0.1725	0.0953	0.2992	0.0115	0.2407	0.0182	0.2814	0.3369	0.2797	0.1362	
1h05	0.1857	0.0295	0.2122	0.0276	0.2854	0.1576	0.2594	0.0182	0.3575	0.1263	0.2628	0.1029	
3h15	0.1993	0.0295	0.2173	0.0495	0.2552	0.0115	0.2300	0.0182	0.3809	0.1060	0.2717	0.1810	
1day	0.2765	0.0001	0.2273	0.0139	0.2544	0.0645	0.2042	0.0182	0.3243	0.1263	0.2302	0.1848	
HAR- RV	0.1748	0.1520	0.1662	0.0953	0.2649	0.0115	0.1821	0.0182	0.4382	0.0024	0.2644	0.0047	
HAR-RV jumps	0.1667	0.2427	0.1497	0.8194	0.2039	1.0000	0.1426	1.0000	0.3294	0.1263	0.1965	0.1848	
GARCH	0.2291	0.0048	0.2205	0.0953	0.2282	0.5737	0.2041	0.0182	0.2723	0.2280	0.2208	0.1848	
GAS	0.2363	0.0295	0.2447	0.0276	0.2173	0.9076	0.2142	0.0182	0.2348	1.0000	0.2226	0.1848	

Note: This table presents the MCS test results obtained for the two assets under analysis (the S&P 500 and Microsoft) during both calm and crisis periods. The results are reported for three forecasting horizons, namely one day ($H = 1$), one week ($H = 5$) and two weeks ($H = 10$). For each we present the average value of the QLIKE loss function along with the corresponding p -value resulting from the MCS test. The confidence level for the MCS test is set to $\alpha = 25\%$ and 10,000 bootstrap resamples are used, with block length of five observations, to obtain the distribution under the null of equal predictive accuracy. The set of the competing variance models includes eight log-MIDAS specifications with regressors sampled at a frequency ranging from one minute to one day, six log-MIDAS models with 1- and 5-minute regressors adjusted for intraday diurnal pattern, intraday pattern and jumps, and intraday pattern, jumps and microstructure noise, the log-HAR-RV, log-HAR-RV-J, GARCH and GAS models.

which the MIDAS forecasts may be less accurate, or stop improving. This result clearly illustrates the influence of the jumps and the intraday diurnal pattern on the *prediction* of variance, and not only on its *measurement*. Third, we discuss the potential solutions to combine the gains issued from high-frequency predictors and the negative impact of microstructure noise. A first solution consists in augmenting the MIDAS model by modifying the weighting scheme in order to limit the influence of the contaminated observations. A second solution consists in applying the MIDAS regression model on filtered data. Here we adopt the latter solution and show that estimating MIDAS-RV models on filtered log-returns leads to significantly better out-of-sample forecasts. Finally, we compare the MIDAS model to other competing variance models including the GARCH, GAS, HAR-RV and HAR-RV-J models. The results suggest that, for both assets, MIDAS models yield better forecasts in most cases, and, importantly, never yield inferior forecasts, provided they are applied on filtered data.

A future research direction would be to compare the approach taken in this paper, where realized variance is directly related to past intraday data, as in GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006], with that of GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV [2006] or GHYSELS, E., and A. SINKO [2011], where daily realized measures (that are potentially robust to microstructure noise and jumps) are introduced in a MIDAS-RV model.

ACKNOWLEDGEMENTS

Financial support from the Risk Foundation Chair Dauphine-ENSAE-Groupama “Behavioral and Household Finance, Individual and Collective Risk Attitudes” is gratefully acknowledged. Bertrand Candelon has conducted this project within the research program “Risk Management and Investment Strategies (RMS)” under the aegis of the Europlace Institute of Finance, and Insti7. The authors thank Eric Ghysels and the participants at the Ph.D. course on MIDAS models jointly organized by CESAM, the National Bank of Belgium (NBB) and the Center for Operations Research and Econometrics (CORE) in June 2013. We also thank Gilbert Colletaz and Christophe Boucher for helpful comments on the paper, as well as the participants at the 61st Congress of the French Economic Association in Paris, 2012, at the 6th International Workshop on Methods in International Finance Network in Sydney, 2012, at the 6th International Conference on Computational and Financial Econometrics in Oviedo, 2012, at the 21st Symposium of the Society for Nonlinear Dynamics and Econometrics in Milan, 2013, at the 3rd Spring International Conference of the French Finance Association in Lyon, 2013. The usual disclaimers apply.

Correspondence:

Denisa Banulescu-Radu

University of Orléans (LEO, UMRS CNRS 7332).

E-mail: georgiana.banulescu@univ-orleans.fr

Bertrand Candelon

Maastricht University.

E-mail: candelonb@gmail.com

Christophe Hurlin

University of Orléans (LEO, UMRS CNRS 7332).

E-mail: christophe.hurlin@univ-orleans.fr

Sébastien Laurent

Aix-Marseille University (Aix-Marseille School of Economics), CNRS & EHESS, France

Aix-Marseille Graduate School of Management - IAE, France

E-mail: sebastien.laurent@iae-aix.com

Appendix

A. Model Confidence Set

Let us consider r_t the logarithmic return process where a unit of time corresponds to a trading day, and \mathbb{F}_t the information set at time t . We can then define the conditional mean, $\mu_t = E(r_t | \mathbb{F}_{t-1})$, and the conditional variance, $\sigma_t^2 \equiv \text{Var}(r_t | \mathbb{F}_{t-1})$. We denote by M_0 the set of K models used to forecast variance, and $\sigma_{t,k}^2$, $k = 1, \dots, K$, the j^{th} variance forecast series. The objective of the model confidence set (MCS) of HANSEN, P. R., A. LUNDE, and J. M. NASON [2011] is to identify the cluster of “best” forecasting models, in this case variance models, under a specified loss function, $L(\sigma_t^2, \sigma_{t,k}^2)$.

The relative performance of each pair of forecasts is measured by $d_{t,k,j} = L(\sigma_t^2, \sigma_{t,k}^2) - L(\sigma_t^2, \sigma_{t,j}^2)$, for all $k, j \in M_0$ with $k \neq j$. Under the assumption that $d_{t,k,j}$ is stationary, the null hypothesis of equal predictive ability takes the form

$$H_0 : E[d_{t,k,j}] = 0 \quad \forall k, j \in M_0. \quad (\text{A.1})$$

If the null of equal predictive ability is rejected at a given confidence level α , then the worst performing model is eliminated from M_0 . The iterative testing procedure ends when the first non rejection occurs, or, obviously, if all models but one have been recursively eliminated. The confidence level α is held at each iteration, thus allowing us to construct a $(1 - \alpha)$ -confidence set, $M^* \equiv \{k \in M_0 : E(d_{t,k,j}) \leq 0 \quad \forall j \in M_0\}$, for the best model(s) in M_0 . Finally, the MCS p-value is equal to $p_i = \max(p_{i-1}, p(i))$, $i = 1, \dots, K$, where p_i is the p-value of the test under the null hypothesis $H_{0:M_i}$, i.e., at the i^{th} step of the iteration process. By convention the p-value when there is only one surviving model is $p_K = 1$.

The main advantages of the MCS test are that it does not require that a benchmark be specified and it relies on simple hypotheses making it possible to compute standard asymptotics. However, when K is large the test statistics have a non-standard distribution, the reason for which a bootstrap scheme is employed to obtain the distribution under the null hypothesis.

B. Pre-Averaging Methodology

Let us consider a vector of log-prices p_t and \mathbb{F}_t the information set available at time t . Standard no-arbitrage conditions suggest that log-prices are semi-martingales. However, in practice, market microstructure noise (including bid-ask bounces, gradual response of prices to a block trade, price discreteness, differences in trade sizes or informational content of price changes, strategic inventory control effects, component of the order flow, etc.) create spurious variations in asset prices, thus distorting the efficient price. As a result, we observe at time-points i/n , $i = 1, \dots, n$, a price process \tilde{p} which is distorted by noise, such that:

$$\tilde{p}_t = p_t + \varepsilon_t, \quad (\text{B.1})$$

where ε_t is an *i.i.d* process, stochastically independent from p .

PODOLSKIJ, M., M. VETTER, ET AL. [2009] and CHRISTENSEN, K., S. KINNEBROCK, and M. PODOLSKIJ [2010] and argue that, under zero mean *i.i.d* microstructure noise, the impact of the noise could be reduced by smoothing the observed log-price $X\tilde{p}$. Therefore, they approximate p_t by an average of observations of \tilde{p} in the neighborhood of t , the noise being thus averaged away.

The steps to pursue in order to conduct the pre-averaging are the following:

Step 1. Consider a sequence of integers, k_n , and a number $\theta \in (0, \infty)$, such that:

$$\frac{k_n}{\sqrt{n}} = \theta + o(n^{-1/4}). \quad (\text{B.2})$$

Step 2. Choose a continuous function f on $[0,1]$, piecewise continuously differentiable with a piecewise Lipschitz derivative f' , and which satisfies $f(0) = f(1) = 0$ and $\int_0^1 f^2(s)ds > 0$.¹⁹

Step 3. Compute the noisy high-frequency returns, $\Delta_i^n \tilde{p}$, and the pre-averaged returns, $\bar{\tilde{p}}_i^n$, using f as a weight function for the latter:

$$\Delta_i^n \tilde{p} = \frac{p_i}{n} - \frac{p_{i-1}}{n}, \text{ for } i = 1, 2, \dots, n \quad (\text{B.7})$$

$$\bar{\tilde{p}}_i^n = \sum_{j=1}^{k_n-1} g\left(\frac{j}{k_n}\right) \Delta_{i+j}^n \tilde{p}, \text{ for } i = 0, \dots, n - k_n + 1. \quad (\text{B.8})$$

As shown above, k_n represents the pre-averaging window length and depends on the parameter θ . Considering $k_n = \lceil \theta \sqrt{n} \rceil$, we can show that $\bar{\tilde{p}}_i^n = O_p(n^{-1/4})$, the method thus delivering the best rate of convergence. In our simulation and empirical analysis, we use

¹⁹. The following functions and numbers are associated with \int :

$$\phi_1(s) = \int_s^1 f'(u)f(u-s)du \quad (\text{B.3})$$

$$\phi_2(s) = \int_s^1 f(u)f(u-s)du \quad (\text{B.4})$$

$$\psi_1 = \phi_1(0), \psi_2 = \phi_2(0), \quad (\text{B.5})$$

$$\phi_{11} = \int_0^1 \phi_1^2(s)ds, \phi_{12} = \int_0^1 \phi_1(s)\phi_2(s)ds, \phi_{22} = \int_0^1 \phi_2^2(s)ds \quad (\text{B.6})$$

the same weight function as PODOLSKIJ, M., M. VETTER, ET AL. [2009] and CHRISTENSEN, K., S. KINNEBROCK, and M. PODOLSKIJ [2010], namely:

$$f(x) = \min(x, 1 - x). \quad (\text{B.9})$$

We take:

$$\frac{k_n}{n^{1/2+\delta}} = \theta + o(n^{-1/4 + \delta/2}) \quad (\text{B.10})$$

for some $0 < \delta < 1/2$ and set $\delta = 0.1$ and $\theta = 1$ based on the results of PODOLSKIJ, M., M. VETTER, ET AL. [2009] and CHRISTENSEN, K., S. KINNEBROCK, and M. PODOLSKIJ [2010].

C. Pre-Averaged MIDAS Regressors – S&P 500

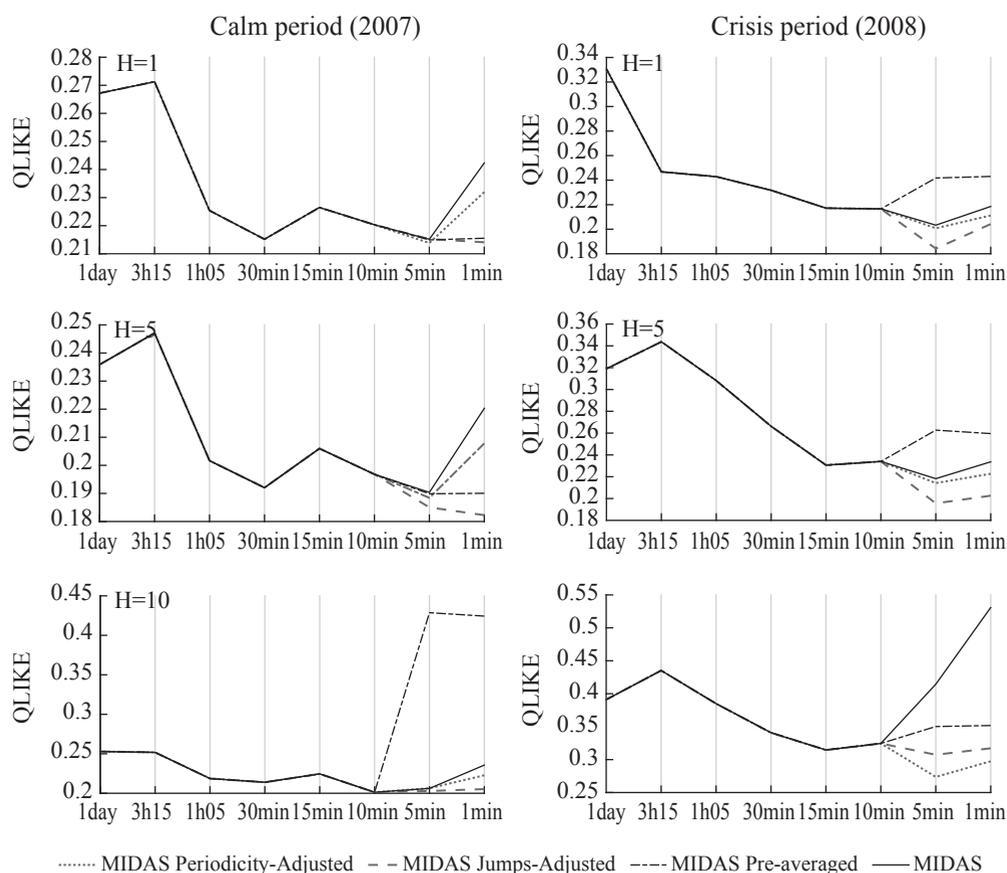


FIGURE C.1. – S&P 500 average QLIKE

Note: The S&P 500 - This figure displays the average QLIKE for the MIDAS-RV forecasts for various sampling frequencies (m_t) of the predictors, and the three forecasting horizons. The left panel corresponds to the calm period (2007) and the right panel to the crisis period (2008). The solid blue line corresponds to the MIDAS-RV model with raw data sampled at frequencies 1-min to 3h15. The green, red and blue dotted lines correspond respectively to the MIDAS-RV models on intraday diurnal pattern, jumps and intraday pattern, and jumps, intraday pattern and noise filtered log-returns, sampled at frequencies between 1-min and 5-min.

D. Pre-Averaged MIDAS Regressors – Microsoft

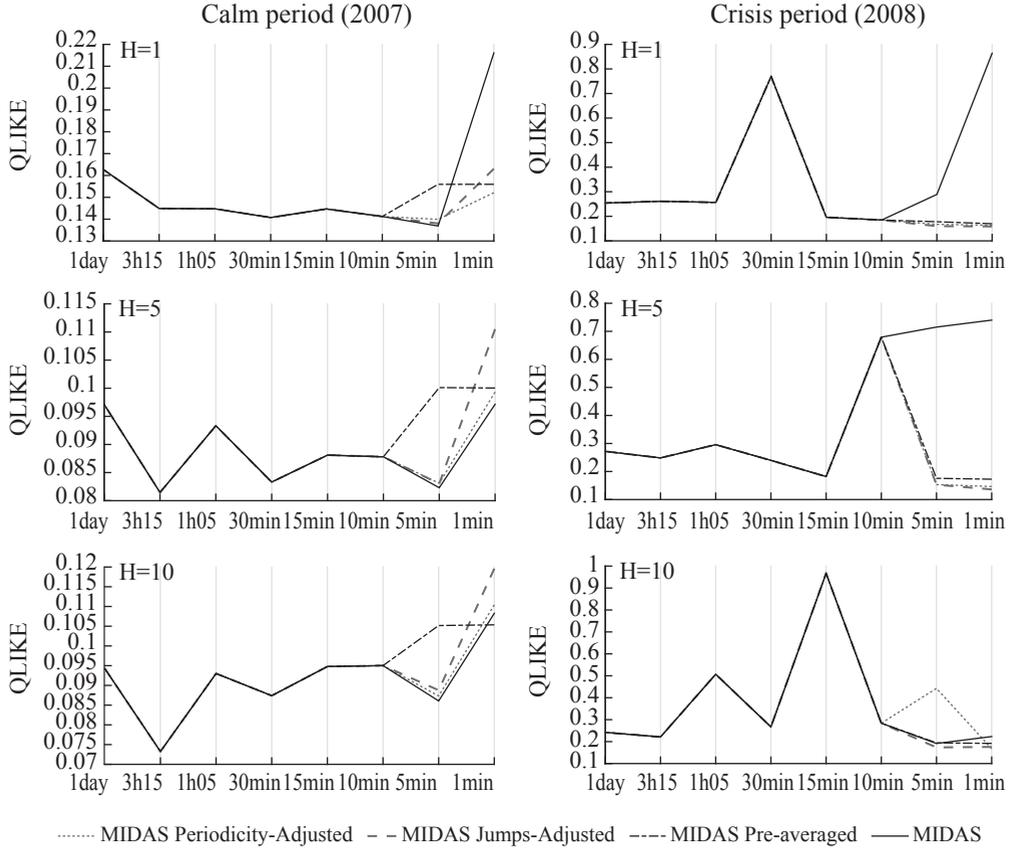


FIGURE D.1. – Microsoft average QLIKE

Note: See FIGURE C.1.

E. HAR-RV versus MIDAS

In this appendix, we show that the HAR-RV model proposed by CORSI, F. [2009] can be written as a weight-constrained form of the MIDAS model with regressors sampled at a frequency m_2 . The HAR-RV model is defined as:

$$RV_{t+1,t}^{(m_2)} = \alpha_0 + \alpha_1 RV_{t,t-1}^{(m_2)} + \alpha_2 RV_{t,t-1}^{(m_2)w} + \alpha_3 RV_{t,t-1}^{(m_2)m} + \varepsilon_{t+1}, \quad (\text{E.1})$$

where $RV_{t+1,t}$ is the daily realized variance given by:

$$RV_{t+1,t}^{(m_2)} = I_{m_2} (L^{1/m_2}) r_{t+1,t+1-1/m_2}^{(m_2)^2}, \quad (\text{E.2})$$

with $I_{m_2}(L^{1/m_2}) = \sum_{j=0}^{m_2-1} L^{j/m_2}$, and m_2 the sampling frequency of the squared returns used to compute the realized variance. By convention

$$RV_{t,t-1}^{(m_2)w} = \frac{1}{5} \sum_{i=0}^4 RV_{t-i,t-i-1}^{(m_2)}, \quad (\text{E.3})$$

and

$$RV_{t,t-1}^{(m_2)m} = \frac{1}{22} \sum_{i=0}^{21} RV_{t-i,t-i-1}^{(m_2)}. \quad (\text{E.4})$$

To complete the explanation, we include Eq. (31), Eq. (32) and Eq. (33) into the definition of the model and obtain:

$$\begin{aligned} RV_{t+1,t}^{(m_2)} &= \alpha_0 + \left(\alpha_1 + \frac{1}{5} \alpha_2 + \frac{1}{22} \alpha_3 \right) RV_{t,t-1}^{(m_2)} + \left(\frac{1}{5} \alpha_2 + \frac{1}{22} \alpha_3 \right) \sum_{i=1}^4 RV_{t-i,t-i-1}^{(m_2)} + \\ &\quad \frac{1}{22} \alpha_3 \sum_{i=5}^{21} RV_{t-i,t-i-1}^{(m_2)} + \varepsilon_{t+1} \\ &= \alpha_0 + \left(\alpha_1 + \frac{1}{5} \alpha_2 + \frac{1}{22} \alpha_3 \right) I_{m_2}(L^{1/m_2}) r_{t,t-1/m_2}^{(m_2)2} \\ &\quad + \left(\frac{1}{5} \alpha_2 + \frac{1}{22} \alpha_3 \right) \sum_{i=1}^4 I_{m_2}(L^{1/m_2}) r_{t-i,t-i-1/m_2}^{(m_2)2} \\ &\quad + \frac{1}{22} \alpha_3 \sum_{i=5}^{21} I_{m_2}(L^{1/m_2}) r_{t-i,t-i-1/m_2}^{(m_2)2} + \varepsilon_{t+1} \square \end{aligned} \quad (\text{E.5})$$

Finally, the HAR-RV model takes the form of a daily MIDAS-RV model with squared return regressors sampled at a frequency m_2 :

$$\begin{aligned} RV_{t+1,t}^{(m_2)} &= \alpha_0 + \left(\alpha_1 + \frac{1}{5} \alpha_2 + \frac{1}{22} \alpha_3 \right) \sum_{j=0}^{m_2-1} L^{j/m_2} r_{t,t-1/m_2}^{(m_2)2} \\ &\quad + \left(\frac{1}{5} \alpha_2 + \frac{1}{22} \alpha_3 \right) \sum_{i=1}^4 \sum_{j=0}^{m_2-1} L^{j/m_2} r_{t-i,t-i-1/m_2}^{(m_2)2} \\ &\quad + \frac{1}{22} \alpha_3 \sum_{i=5}^{21} \sum_{j=0}^{m_2-1} L^{j/m_2} r_{t-i,t-i-1/m_2}^{(m_2)2} + \varepsilon_{t+1}. \end{aligned} \quad (\text{E.6})$$

F. MIDAS-RK Specification

TABLE F.1. – MIDAS-RK specification

Panel A: Calm period (2007)												
	H=1				H=5				H=10			
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT	
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value
1min	0.2488	0.1361	0.2008	0.1436	0.2165	0.2038	0.1096	0.0519	0.2306	0.1975	0.1132	0.1338
1min Per.Adj	0.2403	0.1361	0.2029	0.0985	0.2039	0.5292	0.1115	0.0519	0.2199	0.2987	0.1142	0.1944
1min Jumps. Adj	0.2211	0.9069	0.2164	0.0169	0.1738	1.0000	0.1244	0.0034	0.1960	0.7304	0.1302	0.2673
1min Jumps. Adj_ Preav	0.2209	0.1361	0.1993	0.1436	0.1792	0.6843	0.1087	0.0519	0.4065	0.1975	0.1086	0.2243
5min	0.2204	0.9069	0.1841	0.5802	0.1836	0.6843	0.0914	0.9359	0.2044	0.4371	0.0903	0.3099
5min Per.Adj	0.2192	1.0000	0.1886	0.3606	0.1810	0.6843	0.0929	0.9359	0.1998	0.6074	0.0900	0.3534
5min Jumps. Adj	0.2231	0.1361	0.1870	0.3606	0.1771	0.7382	0.0929	0.9359	0.1944	1.0000	0.0921	0.5883
5min Jumps. Adj_ Preav	0.2200	0.9375	0.1993	0.1436	0.1787	0.7382	0.1090	0.0455	0.4063	0.1975	0.1084	0.1260
10min	0.2254	0.1361	0.1888	0.3606	0.1905	0.5292	0.0993	0.4303	0.2127	0.1975	0.0980	0.2278
15min	0.2313	0.1361	0.1899	0.3606	0.1988	0.2045	0.0967	0.6397	0.2194	0.2987	0.0953	0.2709
30min	0.2228	0.1361	0.1852	0.3606	0.1866	0.6843	0.0928	0.9359	0.2079	0.6074	0.0912	0.4657
1h05	0.2339	0.1361	0.1900	0.3606	0.1946	0.5724	0.1027	0.4303	0.2115	0.4371	0.0963	0.3038
3h15	0.2766	0.0837	0.1857	0.3606	0.2366	0.2038	0.0883	1.0000	0.2417	0.1975	0.0762	0.3093
1day	0.2683	0.1361	0.2042	0.3606	0.2278	0.2045	0.1051	0.4303	0.2439	0.1975	0.0991	0.1999
HAR- RV	0.2242	0.1361	0.1809	1.0000	0.1916	0.5597	0.0979	0.5051	0.2127	0.2987	0.0982	0.2659
HAR-RV jumps	0.2292	0.1361	0.1837	0.3606	0.1953	0.5292	0.1007	0.1194	0.2178	0.1975	0.1030	0.0312
GARCH	0.3195	0.0837	0.2203	0.0985	0.2663	0.2038	0.1265	0.0519	0.2708	0.1975	0.1236	0.0090
GAS	0.3102	0.0837	0.2026	0.3606	0.2466	0.2038	0.1043	0.6397	0.2451	0.1975	0.0954	0.2823

Continued on next page

TABLE F.1. – MIDAS-RK specification (*Continued*)

Panel B: Crisis period (2008)													
	H=1				H=5				H=10				
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT		
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	
1min	0.2342	0.0130	0.3377	0.0497	0.2458	0.0906	1.1943	0.0035	0.3280	0.0787	0.3076	0.0036	
1min Per.Adj	0.2255	0.0269	0.4581	0.0497	0.2353	0.5226	0.5386	0.0227	0.3120	0.0897	0.4533	0.0036	
1min Jumps. Adj	0.2119	0.0269	0.1886	0.8972	0.2155	0.5555	0.1502	1.0000	0.3344	0.0897	0.1797	1.0000	
1min Jumps. Adj_ Preav	0.3830	0.0094	0.1976	0.2843	0.2973	0.0021	0.1858	0.0656	0.4030	0.0787	0.1991	0.4496	
5min	0.2193	0.0269	0.1976	0.0973	0.2301	0.5226	0.1963	0.0656	0.2995	0.0897	0.9667	0.0034	
5min Per.Adj	0.2150	0.0269	0.1932	0.2843	0.2280	0.5555	0.1646	0.1629	0.2873	0.3441	0.4717	0.0036	
5min Jumps. Adj	0.1805	1.0000	0.1878	1.0000	0.2065	1.0000	0.1649	0.1433	0.3243	0.0897	0.1805	0.9300	
5min Jumps. Adj_ Preav	0.2587	0.0094	0.2083	0.0497	0.3011	0.0019	0.1880	0.0518	0.4033	0.0787	0.2012	0.3384	
10min	0.2306	0.0130	0.2083	0.0497	0.2423	0.0906	0.3447	0.0035	0.3428	0.0787	0.3525	0.0034	
15min	0.2319	0.0130	0.3678	0.0497	0.2448	0.0906	0.7383	0.0035	0.3320	0.0787	0.5074	0.0034	
30min	0.2491	0.0094	0.2585	0.0497	0.2887	0.0197	0.2169	0.0227	0.3600	0.0787	0.2417	0.0259	
1h05	0.2606	0.0094	0.2300	0.0497	0.3137	0.0368	0.6801	0.0035	0.4043	0.0787	0.3155	0.0034	
3h15	0.2635	0.0130	0.2847	0.0497	0.3670	0.0019	0.2551	0.0227	0.4520	0.0351	0.2275	0.3384	
1day	0.3943	0.0094	0.2848	0.0497	0.3481	0.0021	0.2806	0.0227	0.4085	0.0787	0.2490	0.1867	
HAR-RV	0.2355	0.0269	0.2022	0.0973	0.2735	0.0906	0.1890	0.0518	0.3950	0.0383	0.2519	0.0036	
HAR-RV jumps	0.2418	0.0130	0.1968	0.2843	0.2742	0.0884	0.1818	0.0656	0.3943	0.0351	0.2397	0.0036	
GARCH	0.2445	0.0269	0.2285	0.0497	0.2385	0.5226	0.1907	0.1433	0.2772	0.3441	0.1976	0.7651	
GAS	0.2533	0.0130	0.2589	0.0497	0.2293	0.5555	0.2112	0.0518	0.2426	1.0000	0.2126	0.3384	

Note: This table presents the MCS test results obtained for the S&P 500 and Microsoft during both calm and crisis periods. The results are reported for three forecasting horizons, namely one day ($H = 1$), one week ($H = 5$) and two weeks ($H = 10$). For each we present the average value of the QLIKE loss function along with the corresponding p -value resulting from the MCS test. The confidence level for the MCS test is set to $\alpha = 25\%$ and 10,000 bootstrap resamples are used, with block length of five observations, to obtain the distribution under the null of equal predictive accuracy. The set of the competing variance models includes eight MIDAS-RK specifications with regressors (squared return) sampled at a frequency ranging from one minute to one day, six MIDAS-RK models with 1- and 5-minute regressors adjusted for intraday diurnal pattern, intraday pattern and jumps, and intraday pattern, jumps and microstructure noise, the HAR-RV, HAR-RV-J, GARCH and GAS models.

G. MIDAS with Bipower Variation Return Regressors

TABLE G.1. – MIDAS with bipower variation return regressors

Panel A: Calm period (2007)												
	H=1				H=5				H=10			
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT	
	QLIKE	p-value	QLIKE	p-value	QLIKE	p-value	QLIKE	p-value	QLIKE	p-value	QLIKE	p-value
1min	0.2506	0.1766	0.2468	0.0173	0.2298	0.1358	0.0972	0.0728	0.2417	0.2714	0.1089	0.0073
1min Per.Adj	0.2391	0.2366	0.1534	0.0507	0.2158	0.6589	0.0994	0.0728	0.2307	0.4804	0.1112	0.0073
1min Jumps. Adj	0.2127	0.9454	0.1589	0.0173	0.1816	1.0000	0.1047	0.0180	0.2062	1.0000	0.1149	0.0016
1min Jumps. Adj_ Preav	0.2155	0.2366	0.1560	0.0173	0.1900	0.8142	0.1000	0.0180	0.4248	0.2714	0.1053	0.0073
5min	0.2089	1.0000	0.1363	0.7362	0.1890	0.8142	0.0803	0.7513	0.2095	0.9505	0.0860	0.5805
5min Per.Adj	0.2097	0.9454	0.1379	0.7362	0.1887	0.8142	0.0796	1.0000	0.2086	0.9505	0.0849	0.5805
5min Jumps. Adj	0.2145	0.8360	0.1386	0.7362	0.1885	0.8142	0.0815	0.7513	0.2085	0.9505	0.0881	0.3249
5min Jumps. Adj_ Preav	0.2137	0.9454	0.1566	0.0173	0.1891	0.8142	0.1003	0.0507	0.4424	0.2575	0.1046	0.0191
10min	0.2171	0.2366	0.1397	0.7362	0.1987	0.7734	0.0888	0.1794	0.2161	0.6993	0.0934	0.3249
15min	0.2212	0.2366	0.1507	0.0764	0.2016	0.7734	0.0928	0.0728	0.2230	0.6387	0.0969	0.0191
30min	0.2200	0.2366	0.1499	0.2474	0.1970	0.8142	0.0916	0.0728	0.2302	0.4804	0.0957	0.1700
1h05	0.2625	0.2366	0.1456	0.5887	0.2101	0.6589	0.0964	0.0728	0.2229	0.6993	0.1037	0.0191
3h15	0.3410	0.0160	0.1502	0.3013	0.3065	0.1358	0.0833	0.7513	0.2993	0.2714	0.0766	1.0000
1day	0.3642	0.1766	0.1737	0.0507	0.2769	0.6447	0.1063	0.0728	0.4990	0.0720	0.1032	0.3249
HAR- RV	0.2176	0.2366	0.1345	1.0000	0.1941	0.8142	0.0868	0.0740	0.2172	0.7722	0.0943	0.1349
HAR-RV jumps	0.2187	0.2366	0.1359	0.7362	0.1973	0.7734	0.0883	0.0728	0.2226	0.4804	0.0960	0.0191
GARCH	0.3240	0.1527	0.2208	0.0173	0.2849	0.1358	0.1677	0.0011	0.2935	0.2714	0.1747	0.0002
GAS	0.3161	0.1385	0.1884	0.0173	0.2669	0.1358	0.1292	0.0180	0.2688	0.2714	0.1296	0.0160

Continued on next page

TABLE G.1. – MIDAS with bipower variation return regressors (Continued)

Panel B: Crisis period (2008)													
	H=1				H=5				H=10				
	S&P 500		MSFT		S&P 500		MSFT		S&P 500		MSFT		
	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	QLIKE	<i>p</i> -value	
1min	0.2261	0.0318	0.8523	0.0034	0.2428	0.0132	0.7263	0.0029	0.3122	0.1570	1.2599	0.0066	
1min Per.Adj	0.2200	0.0318	0.1990	0.1032	0.2330	0.0472	0.5763	0.0205	0.3000	0.4979	0.1664	1.0000	
1min Jumps. Adj	0.2095	0.0318	0.1605	1.0000	0.2031	0.5605	0.1343	1.0000	0.3207	0.4979	0.1748	0.5802	
1min Jumps. Adj_ Preav	0.2431	0.0318	0.1717	0.2419	0.2609	0.0132	0.1741	0.0205	0.3524	0.1570	0.1927	0.3513	
5min	0.4821	0.0318	0.1722	0.1032	0.2314	0.4355	0.1586	0.0577	0.8549	0.0389	0.2625	0.0204	
5min Per.Adj	0.1858	0.0318	0.1663	0.2419	0.2025	0.5605	0.1508	0.1618	0.2866	0.4979	0.1702	0.6597	
5min Jumps. Adj	0.1747	1.0000	0.1620	0.8316	0.1937	1.0000	0.1519	0.1170	0.3097	0.4979	0.1721	0.6200	
5min Jumps. Adj_ Preav	0.2465	0.0318	0.1927	0.1032	0.2712	0.0132	0.1848	0.0205	0.3625	0.1570	0.2002	0.1029	
10min	0.2233	0.0318	0.2059	0.0034	0.2293	0.1082	0.3507	0.0029	0.3277	0.1570	0.6099	0.0139	
15min	0.2139	0.0318	0.1743	0.1032	0.2249	0.1082	0.2972	0.0205	0.5412	0.0389	0.3097	0.0139	
30min	0.1937	0.0318	0.2547	0.0034	0.2454	0.0132	0.2603	0.0163	0.8991	0.0389	0.2655	0.0142	
1h05	0.2339	0.0318	0.2638	0.0034	0.2985	0.0132	0.3464	0.0029	0.3637	0.1570	1.0577	0.0126	
3h15	0.2937	0.0318	0.2880	0.0034	0.2843	0.0132	0.2497	0.0088	0.9281	0.0389	0.2327	0.0204	
1day	0.3798	0.0215	0.3876	0.0034	0.4460	0.0132	0.3400	0.0088	0.4760	0.0389	0.3069	0.0142	
HAR- RV	0.2102	0.0318	0.1781	0.1032	0.2449	0.0472	0.1754	0.0205	0.3563	0.0389	0.2346	0.0142	
HAR-RV jumps	0.2148	0.0318	0.1719	0.2419	0.2472	0.0132	0.1669	0.0525	0.3569	0.0389	0.2205	0.0204	
GARCH	0.2291	0.0318	0.2205	0.1032	0.2282	0.4355	0.2041	0.0205	0.2723	0.4979	0.2208	0.1029	
GAS	0.2363	0.0318	0.2447	0.1032	0.2173	0.5605	0.2142	0.0205	0.2348	1.0000	0.2226	0.1029	

Note: This table presents the MCS test results obtained for the S&P 500 and Microsoft during both calm and crisis periods. The results are reported for three forecasting horizons, namely one day ($H = 1$), one week ($H = 5$) and two weeks ($H = 10$). For each we present the average value of the QLIKE loss function along with the corresponding p -value resulting from the MCS test. The confidence level for the MCS test is set to $\alpha = 25\%$ d10,000 bootstrap resamples are used, with block length of five observations, to obtain the distribution under the null of equal predictive accuracy. The set of the competing variance models includes eight MIDAS specifications with regressors (bipower variation) sampled at a frequency ranging from one minute to one day, six MIDAS models with 1- and 5-minute regressors adjusted for intraday diurnal pattern, intraday pattern and jumps, and intraday pattern, jumps and microstructure noise, the HAR-RV, HAR-RV-J, GARCH and GAS models.

References

- AIT-SAHALIA, Y., and L. MANCINI (2008): “Out of Sample Forecasts of Quadratic Variation”, *Journal of Econometrics*, **147**, 17-33. [136, 139]
- AIT-SAHALIA, Y., P. A. MYKLAND, and L. ZHANG (2005): “How often to Sample a Continuous-time Process in the Presence of Market Microstructure Noise”, *Review of Financial Studies*, **18(2)**, 351-416. [136]
- AIT-SAHALIA, Y., and J. YU (2009): “High Frequency Market Microstructure Noise Estimates and Liquidity Measures”, Tech. Rep. 1, *Institute of Mathematical Statistics*. [148]
- ANDERSEN, T. G., and T. BOLLERSLEV (1997a): “Intraday Periodicity and Volatility Persistence in Financial Markets”, *Journal of Empirical Finance*, **4**, 115-158. [139]
- ANDERSEN, T. G., and T. BOLLERSLEV (1997b): “Intraday Periodicity and Volatility Persistence in Financial Markets”, *Journal of Empirical Finance*, **4**, 115-158. [143]
- ANDERSEN, T. G., and T. BOLLERSLEV (1998a): “Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts”, *International Economic Review*, **39(4)**, 885-905. [139, 141]
- ANDERSEN, T. G., and T. BOLLERSLEV (1998b): “Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies”, *Journal of Finance*, **53(1)**, 219-265. [143]
- ANDERSEN, T. G., T. BOLLERSLEV, and F. X. DIEBOLD (2007): “Roughing It Up: Including Jump Components in the Measurement, Modeling, and Forecasting of Return Volatility”, *The Review of Economics and Statistics*, **89(4)**, 701-720. [137, 154]
- ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, and P. LABYS (2003): “Modeling and Forecasting Realized Volatility”, *Econometrica*, **71(2)**, 579-625. [155]
- ANDERSEN T. G., T. BOLLERSLEV, and N. MEDDAHI (2005): “Correcting the Errors: Volatility Forecast Evaluation Using High-Frequency Data and Realized Volatilities”, *Econometrica*, **73**, 279–296. [141]
- BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, and N. SHEPHARD (2008): “Designing Realised Kernels to Measure the ex-Post Variation of Equity Prices in the Presence of Noise”, *Econometrica*, **76**, 1481-536. [139, 158]
- BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, and N. SHEPHARD (2009): “Realized Kernels in Practice: Trades and Quotes”, *The Econometrics Journal*, **12(3)**, C1-C32. [144]
- BARNDORFF-NIELSEN, O. E., and N. SHEPHARD (2004a): “Econometric Analysis of Realised Covariation: High Frequency Based Covariance, Regression and Correlation in Financial Economics”, *Econometrica*, **72**, 885-925. [139, 158]

- BARNDORFF-NIELSEN, O. E., and N. SHEPHARD (2004b): “Power and Bipower Variation with Stochastic Volatility and Jumps (with Discussion)”, *Journal of Financial Econometrics*, **2**, 1-48. [144, 154]
- BARNDORFF-NIELSEN, O. E., and N. SHEPHARD (2006): “Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation”, *Journal of financial Econometrics*, **4(1)**, 1-30. [144]
- BATES, D. S. (1996): “Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options”, *Review of financial studies*, **9(1)**, 69-107. [143]
- BOLLERSLEV, T. (1986): “Generalized Autoregressive Heteroskedasticity”, *Journal of Econometrics*, **31**, 307-327. [154]
- BOUDT, K., C. CROUX, and S. LAURENT (2011): “Robust Estimation of Intraweek Periodicity in Volatility and Jump Detection”, *Journal of Empirical Finance*, **18(2)**, 353-367. [137, 144, 146, 148, 149]
- CHEN, X., and E. GHYSELS (2011): “News-Good or Bad-And Its Impact on Volatility Predictions Over Multiple Horizons”, *Review of Financial Studies*, **24(1)**, 46-81. [138]
- CHEN, Y.-C., and W.-J. TSAY. (2011): “Forecasting Commodity Prices with Mixed-Frequency Data: An OLS-Based Generalized ADL Approach”, *Working paper*. [140]
- CHRISTENSEN, K., S. KINNEBROCK, and M. PODOLSKIJ (2010): “Pre-Averaging Estimators of the ex-Post Covariance Matrix in Noisy Diffusion Models with Non-Synchronous Data”, *Journal of Econometrics*, **159(1)**, 116-133. [148, 163, 164]
- CORSI, F. (2009): “A Simple Approximate Long-Memory Model of Realized Volatility”, *Journal of Financial Econometrics*, **7(2)**, 174-196. [137, 154, 165]
- CREAL, D. D., S. J. KOOPMAN, and A. LUCAS (2013): “Generalized Autoregressive Score Models with Applications”, *Journal of Applied Econometrics*, **28(5)**, 777-795. [137, 154]
- ENGLE, R. F. (1982): “Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation”, *Econometrica*, **45**, 987-1007. [154]
- FRALE, C., and L. MONTEFORTE. (2011): *FaMIDAS: A Mixed Frequency Factor Model with MIDAS Structure*. Banca d'Italia. [140]
- GARCIA, R., and N. MEDDAHI (2006): “Comment on Realized Variance and Market Microstructure Noise”, *Journal of Business & Economic Statistics*, **24**, 184-191. [136, 139]
- GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV (2004): “The MIDAS Touch: Mixed Data Sampling Regression Models”, *Working paper*. [136]
- GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV (2005): “There Is a Risk-Return Trade-Off after All”, *Journal of Financial Economics*, **76(3)**, 509-548. [138]

- GHYSELS, E., P. SANTA-CLARA, and R. VALKANOV (2006): “Predicting Volatility: Getting the Most out of Return Data Sampled at Different Frequencies”, *Journal of Econometrics*, **131(1)**, 59-95. [136, 137, 138, 139, 140, 161]
- GHYSELS, E., and A. SINKO (2011): “Volatility Forecasting and Microstructure Noise”, *Journal of Econometrics*, **160(1)**, 257-271. [136, 138, 161]
- GHYSELS, E., A. SINKO, and R. VALKANOV (2007): “MIDAS Regressions: Further Results and New Directions”, *Econometric Reviews*, **26(1)**, 53-90. [140]
- GHYSELS, E., and R. VALKANOV (2012): “Forecasting Volatility with MIDAS”. In: L. Bauwens, C. Hafner, and S. Laurent (Eds.), *Handbook of Volatility Models and Their Applications*, 383-401. [138, 139, 140]
- HANSEN, P. R., and G. HOREL (2009): “Quadratic Variation by Markov Chains”. *Working Paper*. [139, 158]
- HANSEN, P. R., and A. LUNDE (2004): “Technical Appendix: An Unbiased Measure of Realized Variance”, Brown University, 2004. [136, 139]
- HANSEN, P. R., and A. LUNDE (2006a): “Consistent Ranking of Volatility Models”, *Journal of Econometrics*, **131**, 97-121. [141]
- HANSEN, P. R., and A. LUNDE (2006b): “Realized Variance and Market Microstructure Noise”, *Journal of Business & Economic Statistics*, **24(2)**, 127-161. [148]
- HANSEN, P. R., A. LUNDE, and J. M. NASON (2011): “The Model Confidence Set”, *Econometrica*, **79**, 456-497. [137, 142, 162]
- HARRIS, L. (1986): “A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns”, *Journal of Financial Economics*, **16**, 99-117. [143]
- HARVEY, A. C. (2013): *Dynamic Models for Volatility and Heavy Tails*. Cambridge University Press. [154]
- HECQ, A., S. LAURENT, and F. C. PALM (2012): “Common Intraday Periodicity”, *Journal of Financial Econometrics*, **10(2)**, 325-353. [143]
- JACOD, J., Y. LI, P. A. MYKLAND, and M. PODOLSKIJ, and M. VETTER (2009): “Microstructure Noise in the Continuous Case: The Pre-Averaging Approach”, *Stochastic Processes and Their Applications*, **119(7)**, 2249-2276. [142, 148]
- LAHAYE, J., S. LAURENT, and C. J. NEELY (2011): “Jumps, Cojumps and Macro Announcements”, *Journal of Applied Econometrics*, **26(6)**, 893-921. [137, 144]
- LAURENT, S., J. V. ROMBOUTS, and F. VIOLANTE (2013): “On Loss Functions and Ranking Forecasting Performances of Multivariate Volatility Models”, *Journal of Econometrics*, **173(1)**, 1-10. [141]

- LEE, S. S., and P. A. MYKLAND (2008): “Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics”, *Review of Financial studies*, **21(6)**, 2535-2563. [137, 144, 146, 147, 149]
- PATTON, A. J. (2011): “Volatility Forecast Comparison Using Imperfect Volatility Proxies”, *Journal of Econometrics*, **160(1)**, 246-256. [137, 141, 142]
- PATTON, A. J., and K. SHEPPARD (2009): *Evaluating Volatility and Correlation Forecasts*. In: *Handbook of Financial Time Series*, Springer, 801-838. [141]
- PODOLSKIJ, M., M. VETTER, et al. (2009): “Estimation of Volatility Functionals in the Simultaneous Presence of Microstructure Noise and Jumps”, *Bernoulli*, **15(3)**, 634-658. [142, 148, 163, 164]
- WOOD, R. A., T. H. MCINISH, and J. K. ORD (1985): “An Investigation of Transactions Data for NYSE Stocks”, *The Journal of Finance*, **40(3)**, 723-739. [143, 146]
- ZHANG, L. (2006): “Efficient Estimation of Stochastic Volatility Using Noisy Observations: A Multi-Scale Approach”, *Bernoulli*, **12(6)**, 1019-1043 [139, 158]
- ZHOU, B. (1996): “High-Frequency Data and Volatility in Foreign-Exchange Rates”, *Journal of Business & Economic Statistics*, **14(1)**, 45-52. [136, 148]