Macro-Driven VaR Forecasts:
From very High to Very Low-Frequency Data

Yves Dominicy
SBS-EM, ECARES, Université libre de Bruxelles

Harry Vander Elst
SBS-EM, ECARES, Université libre de Bruxelles

November 2015

ECARES working paper 2015-41

ECARES
ULB - CP 114/04
50, F.D. Roosevelt Ave., B-1050 Brussels BELGIUM
www.ecares.org

Macro-driven VaR forecasts: From very high to very low-frequency data*

Yves Dominicy† **Harry Vander Elst**‡

Université libre de Bruxelles,
Solvay Brussels School of Economics and Management, ECARES

Abstract

This paper studies in some details the joint-use of high-frequency data and economic variables to model financial returns and volatility. We extend the Realized LGARCH model by allowing for a time-varying intercept, which responds to changes in macroeconomic variables in a MIDAS framework and allows macroeconomic information to be included directly into the estimation and forecast procedure. Using more than 10 years of high-frequency transactions for 55 U.S. stocks, we argue that the combination of low-frequency exogenous economic indicators with high-frequency financial data improves our ability to forecast the volatility of returns, their full multi-step ahead conditional distribution and the multi-period Value-at-Risk. We document that nominal corporate profits and term spreads generate accurate risk measures forecasts at horizons beyond two business weeks.

Keywords: Realized LGARCH, Value-at-Risk, density forecasts, realized measures of volatility.

JEL Classification: C22, C53, C58, G17

*Acknowledgments. Yves Dominicy gratefully acknowledges the Fonds National de la Recherche Scientifique (FNRS), Communauté française de Belgique, for support via a Mandat de Chargé de Recherches FNRS. Harry Vander Elst gratefully acknowledges financial support from a FRESH grant from FNRS. We would like to thank Torben Andersen, Griselda Deelstra, Siegfried Hörmann, Asger Lunde, Marco Geraci, Davy Paindaveine, and David Veredas for insightful comments. We are grateful to Christian Conrad, Karin Loch, Farrukh Javed and Hossein Asgharian for sharing codes for the estimation of the GARCH-MIDAS model and to Asger Lunde for providing cleaned high-frequency data. All remaining errors and inaccuracies are ours.

†**Corresponding author**: Y. Dominicy at Université libre de Bruxelles, Solvay Brussels School of Economics and Management, ECARES, Brussels, Belgium, email: yves.dominicy@ulb.ac.be.

‡H. Vander Elst at Université libre de Bruxelles, Solvay Brussels School of Economics and Management, ECARES, Brussels, Belgium, email: harryvanderelst@gmail.com.

1 Introduction

This paper empirically validates the joint-use of high-frequency financial transactions and economic data to forecast the conditional distribution of returns and some specific interesting quantities such as the conditional volatility and the multi-period Value-at-Risk (VaR).[1] Our framework relies on the use of conditionally heteroskedastic models, introduced in the seminal contributions of Engle (1982) and Bollerslev (1986). In particular, we combine insights from Hansen et al. (2012) and Engle et al. (2013) and propose a new Realized LGARCH model with a time-varying intercept. Doing so, we tackle the following research questions: Can we immediately account for the time-varying nature of some parameters in the model? Can we link their changes to the global economic environment? How sensitive is the approach to the choice of economic indicators? What are the gains in terms of volatility, density, and precision of risk measure forecasts?

In order to answer these questions we extend the classical Realized LGARCH of Hansen et al. (2012) allowing the intercept to change through data-driven adjustments. We show that our model can be represented as a two component model akin to the additive model of Engle and Lee (1999) and argue that filtering the low-frequency economic indicators using MIDAS techniques à la Ghysels et al. (2005) provides a convenient framework to model the slow-moving component of volatility. We find that using economic variables captures changes in the intercept of the model and allows to provide more dynamic forecasts. Our new model accommodates modifications in the economic environment and allows to construct more precise forecasts over horizons beyond two business weeks. We find out that nominal corporate profits and term spreads provide the best results in terms of volatility, density, and VaR forecasting.

The relation between stock market volatility and exogenous economic variables has been extensively studied over the past three decades. The two seminal papers are the contributions of Officer (1973) and Schwert (1989) who first argued that stock market volatility changes because of broader changes in the underlying economic environment. A subject of particular interest concerns the joint-behavior of volatility and the economy with a special emphasis on their lead-lag pattern and on the cyclical behavior of stock market volatility. Many contributions, including Hamilton and Lin (1996), Perez-Quiros and Timmermann (2000), Christiansen et al. (2012), Corradi et al. (2013), and Paye (2012), reported evidence of the counter-cyclical behavior of the long-run market volatility. Hence, if the economy shrinks, stock market volatility is expected to increase in response to concerns about future market conditions. Recently, Engle et al. (2013), Asgharian et al. (2013), and Conrad and Loch (2014) confirmed this conjecture using GARCH-MIDAS models and provided more details on the forecasting implications. In particular, Conrad and Loch (2014) showed that the GARCH-MIDAS provides accurate long-run volatility forecasts. Using our new model, we confirm the empirical findings of the literature and argue that the combination of high-frequency data and economic indicators provide accurate forecasts at long horizons not only for the

[1] By multi-period VaR we refer to the VaR computed over a period of several days.

conditional volatility, but also for the whole conditional density of returns and the multi-period VaR.

The usefulness of high-frequency data to model and forecast volatility is well-known both from the statistical and applied point of view (see for instance Andersen et al. (2003), Fleming et al. (2003), and Christoffersen et al. (2014)). The informational content of high-frequency data is exploited through the lens of realized measures of volatility. These are recent nonparametric volatility measures based on high-frequency transactions that have earned great success over the last decade. The most simple realized measure of volatility is known as the Realized Variance (RV) and is simply obtained by summing up the intraday squared returns (see Andersen, Bollerslev, Diebold and Labys (2001)). Despite its simplicity, RV are sensitive to micro-structure noise and jumps and this has paved the way for many refined versions of RV, for further details on it we refer the reader to McAleer and Medeiros (2008). The main insight is that realized measures provide an accurate signal of the true latent volatility process of returns and can be used as an input in many models. Engle (2002) was the first to include realized measures in the volatility equation of GARCH models introducing the GARCH-X models. Later, Hansen et al. (2012) proposed the class of Realized GARCH models that generalizes the GARCH-X by including a measurement equation for the realized measure of volatility. Hansen and Huang (2015), Hansen et al. (2014), and Vander Elst (2015) completed the class of Realized GARCH models with the Realized EGARCH, the Realized Beta GARCH, and the FloGARCH, respectively. Shephard and Sheppard (2010) proposed the HEAVY model that also focuses on modeling the conditional volatility of returns. On slightly different grounds, Corsi (2009) proposed the HAR-RV model that provides a convenient framework to directly forecast the realized measure of volatility. Our model lies in the family of Realized GARCH models and focuses on the conditional volatility of returns.

Besides precise volatility forecasts, risk management also requires knowledge of the full conditional distribution of returns at different horizons to compute risk measures and to price financial derivatives. The multi-period ahead distribution is barely ever available in closed form and numerical techniques have to be used. Being complete, our model allows to simulate a large amount of paths for the future daily returns from which estimates can be obtained. We provide the details of a Monte Carlo simulation that allows to construct multi-step ahead conditional density estimates and multi-period VaR. Similar works to ours include Giot and Laurent (2004), Clements et al. (2008), and Brownlees and Gallo (2009) who provided a framework based on high-frequency data to compute the VaR. We take a step further by computing the VaR and the conditional density over horizons ranging from 1 day to 8 business weeks. Additionally, we rely on Amisano and Giacomini (2007) and Maheu and McCurdy (2011) to assess the ability of our models to forecast the conditional density of log-returns. In particular, Maheu and McCurdy (2011) provide a framework to compare the ability of several models to provide multi-step ahead conditional density forecasts. We use their approach in combination with the Model Confidence Set (MCS) of Hansen et al. (2011).

Along this paper, we will divide the time horizon into fixed periods (e.g. a week, a month, a quarter, etc.) and denote by N_t the number of trading days in the period t. For instance, $N_t = 5$ for a week,

$N_t = 22$ for a month, $N_t = 65$ for a quarter, and so on. Unless explicitly stated otherwise, a variable recorded on day i in period t will be written as $y_{i,t}$, for $i = 1, ..., N_t$ and $t = 1, ..., T$.

The rest of the paper is structured as follows. Section 2 describes the dataset used in this paper. Section 3 presents the new model used in our empirical investigation, namely the Realized LGARCH-MIDAS. Additionally, we provide estimation results for our dataset of stocks and analyse the smooth component of volatility. We give more details about the forecasting methodology in Section 4. In particular, we present the framework to construct volatility and density forecasts based on Monte Carlo simulations. Sections 5 provides the results of our investigation and gives more details on the tools used to evaluate the forecasts. Concluding remarks are given in Section 6.

2 High-frequency data and economic variables

We consider two types of data for our empirical investigation: high-frequency returns and economic variables. We use high-frequency data from the Trade and Quote (TAQ) database for 55 very liquid stocks covering a period from January 4, 2000 to March 30, 2012. Daily transaction prices are cleaned following the procedure of Barndorff-Nielsen et al. (2009). We construct daily realized kernels and sub-sampled RV, which are respectively used as input for the model and as proxy to evaluate forecasts. We need the baseline RV to compute both the realized kernels and the sub-sampled RV

$$RV_{i,t} = \sum_{j=1}^{T_{i,t}} r_{j,i,t}^2,$$

where $r_{j,i,t}$ denotes the high-frequency log-return computed over interval j, of day i, during the period t, and $T_{i,t}$ denotes the amount of high-frequency log-returns on day i during the period t.

The RV computed at the highest frequency are known to be sensitive to market micro-structure noise and to produce upward biases. Zhang et al. (2005) argued that using sub-sampled RV dampens the micro-structure biases. Following Zhang et al. (2005), the sub-sampled RV is formally defined as the sample mean of 5 minute RV computed over non-overlapping sparse grids of returns as

$$ssRV_{i,t}^{5m} = \frac{1}{S} \sum_{s=1}^{S} RV_{s,i,t}^{5m},$$

where $RV_{s,i,t}^{5m}$ denotes the 5-minute RV computed from S non-overlapping samples on day i, in period t. The sub-sampled RV will be used as proxy in order to evaluate the quality of the forecasts produced by our model.

Additionally, we use the realized kernels of of Barndorff-Nielsen et al. (2008) as input for our models. Realized kernels are as well insensitive to micro-structure noise and are computed on day i, during period t, as

$$RK_{i,t} = \sum_{h=-H}^{H} k\left(\frac{h}{H+1}\right)\gamma_h, \quad \text{where} \quad \gamma_h = \sum_{j=|h|+1}^{n} r_{j,i,t} r_{j-|h|,i,t},$$

is called the h-th realized autocovariance. The kernel weight function is denoted here by $k(\frac{h}{H+1})$ with $\frac{h}{H+1} \in [0,1]$ for $h = -H, \ldots, -1, 0, 1, \ldots, H$, and $H \in \mathbb{N}$.

Table 1 provides summary statistics about the stocks used in this paper. The first two columns report daily averages about the amount of prices and non-duplicate prices available in our dataset. From the third column it can be seen that most of the selected stocks are actively traded with an average median duration of about 6 seconds. It appears that the amount of non-duplicate prices can go up to half of the amount of available sample. However, the effects are limited by the short intervals during which prices repeat. Indeed in column four of the table, the amount of zero returns lasting above 10 seconds remains small compared to the total amount of transaction prices.

Additionally, we provide some statistics for daily squared returns (r^2), 5-minute sub-sampled realized variances ($ssRV^{5min}$), and realized kernels (RK), and this for each asset in the columns five to eight of the Table 1, respectively. We compute the annualized version of these volatility measures ($Avol$) by taking the square root of the mean of 250 times the considered measure. The Avol numbers are in general a bit higher for the squared close-to-close returns than for the realized measures because close-to-close returns contain overnight information, which is not included in the realized measures we consider.

Finally, we consider the ratio between the standard deviations of squared returns and sub-sampled realized variances computed over the full sample. This ratio provides a signal to noise type of information and is always higher than one, suggesting that realized volatility are less volatile than squared returns. This well known fact confirms that realized measures are less noisy than daily squared returns.

We also rely on a set of U.S. macroeconomic indicators from 1997(Q1) to 2011(Q4): the real GDP growth (GDP), the industrial production growth (IP), the nominal corporate profits growth (CP), the national activity index of the Chicago FED (NAI), the new orders index of the Institute for Supply Management (NO), and the 10-years yields vs 3-months yields term spreads (TS).[2] Table 2 provides summary statistics for the considered series. In Table 2, the first four columns report the estimated mean, the estimated standard deviation, the skewness and kurtosis coefficients. The fifth column displays the first lag serial correlation. The last two columns present the sample cross-correlations of the indicators with the Real GDP and NAI, which provide insight about the correlation of the indicators in our sample with the real business cycle. A striking observation is that most of the variables are highly correlated with the NAI suggesting that this real business indicator is akin to a common factor to these variables except for corporate profits and term spreads. As we will see in Section 5, these two provide the best forecasts. Since we use these variables to capture changes in the long-run volatility, they are sampled on quarterly basis. The data were obtained from the Federal Reserve Bank of St. Louis. The variables were selected on the basis of empirical results of Conrad and Loch (2014) who validated their good properties for forecasting volatility.

[2]The difference between the sample of economic variables and the financial data is due to the 3 years of initial filtering for the long-run volatility prior to the first observed return. More details are provided in Section 3.

Table 1: Summary statistics of the high-frequency transactions.

	# Trans.	*# non-dupl. Trans.*	*dura.*	*# ret* $= 0$ ($> 10sec$)	*Avol-*r^2	*Avol-*$ssRV^{5m}$	*Avol-RK*	*Avol-*RV^{all}	*Ratio*
AA	4099	2472	4.873	243	45.62	37.76	38.02	51.99	2.43
ABT	3146	1811	5.343	260	25.67	24.28	24.49	29.39	2.79
AES	1845	872	8.236	192	64.53	54.31	52.39	68.81	3.21
AIG	4249	2544	4.371	252	73.15	58.52	55.47	132.81	1.43
AKS	1889	1145	22.75	122	67.61	61.53	59.32	72.95	2.09
AMD	3372	1840	4.806	232	66.14	53.73	52.85	89.00	4.04
AAPL	6489	4538	2.266	155	49.36	40.52	39.56	72.08	11.22
AXP	3579	1983	4.439	275	41.31	35.01	35.88	37.63	2.18
BA	3306	2015	4.993	227	33.14	28.62	28.83	33.15	2.77
BAC	7197	4707	3.524	261	53.92	40.26	41.01	67.03	2.84
BMY	3223	1768	4.817	267	30.23	27.95	28.13	35.23	1.46
BSX	2552	1290	7.646	219	41.62	37.09	36.09	55.76	3.67
C	7245	4167	2.811	314	57.72	45.79	45.32	131.36	2.53
CAG	1749	859	9.191	197	23.35	23.85	22.34	33.77	3.43
CAT	3505	2300	5.440	198	35.39	30.61	30.66	33.17	2.24
CHK	3174	1956	19.50	159	51.80	50.57	43.89	58.53	1.91
CSX	2638	1594	9.361	189	36.59	32.68	32.02	35.05	2.71
D	1926	1074	8.490	201	23.31	23.15	23.16	25.26	2.05
DD	3131	1760	4.953	260	30.92	28.26	28.38	32.76	1.79
DIS	3381	1762	4.727	282	34.46	29.80	29.91	43.01	2.17
DOW	3097	1734	5.629	249	39.26	33.69	33.55	39.01	2.61
EMC	3581	1871	4.195	280	50.95	42.47	42.27	65.31	2.31
FCX	4219	3025	11.33	130	52.44	43.99	42.62	57.34	2.18
GE	6906	4457	2.642	294	34.18	30.45	30.69	51.40	3.45
GIS	2050	1159	8.730	200	19.35	19.45	19.11	22.38	2.37
GLW	3455	1886	4.651	262	60.21	49.73	48.63	73.74	1.78
HAL	4109	2275	4.292	263	49.34	42.35	42.48	48.21	2.78
HD	3978	2153	3.873	292	35.95	30.64	30.82	38.42	3.20
IBM	4357	2852	3.682	237	28.56	24.55	25.18	28.69	5.13
INTC	9535	5460	1.787	219	42.69	35.13	34.01	70.37	4.18
JCP	2490	1404	8.405	209	45.51	40.87	39.76	45.32	2.48
JNJ	3898	2269	4.291	275	20.96	19.60	19.81	23.69	3.51
JPM	5240	2986	3.628	284	45.58	37.40	37.98	41.28	3.16
KEY	2514	1328	8.063	218	52.69	46.52	45.67	62.77	2.35
KO	3241	1686	4.731	275	23.22	21.04	21.41	25.50	2.69
MCD	3239	1790	5.112	262	26.14	25.15	24.92	33.98	2.45
MDT	2965	1578	5.299	273	28.07	25.29	25.16	28.20	2.09
MMM	3037	1889	5.457	208	25.57	23.54	23.91	23.99	2.44
MO	3956	2407	4.374	260	27.19	24.39	23.70	38.82	1.84
MSFT	9359	5447	1.829	236	33.77	27.80	27.15	52.43	2.01
NBR	2471	1338	7.825	224	48.52	40.76	41.10	42.26	1.37
NEM	3521	2059	6.858	228	41.80	35.70	35.06	38.63	2.49
ORCL	6802	3613	2.639	250	46.09	39.42	37.56	89.36	4.06
PFE	5614	3352	3.070	288	28.12	25.56	25.72	46.33	3.56
SLB	4632	3040	3.772	209	40.13	35.45	36.05	36.53	2.63
SPY	7671	4372	3.675	197	21.82	17.94	17.53	20.44	2.74
TJX	2135	1064	8.819	223	32.87	30.95	30.28	33.72	2.61
USB	3136	1661	6.675	257	39.91	36.17	36.20	41.57	2.35
UTX	2961	1783	5.778	216	30.26	26.05	26.41	26.86	3.81
VLO	3484	2218	11.26	171	44.11	36.35	36.67	38.52	2.41
WFC	5020	2888	4.099	286	44.52	36.16	36.95	40.74	2.38
WMT	4241	2403	3.754	277	26.36	24.51	24.70	29.39	2.17
WY	2176	1329	6.800	169	36.43	32.33	32.12	33.51	1.80
XOM	5565	3292	3.364	275	26.89	24.03	24.68	27.27	4.17
XRX	2155	1080	7.897	205	51.47	43.07	39.45	83.52	2.40

Summary statistics based on the cleaned dataset of high-frequency transactions for the panel of stocks considered. The sample covers the period from January 4, 2000 to March 30, 2012. The first two columns report daily averages about the amount of prices and non-duplicate prices available in our dataset. We also provide the median duration and the average amount of zero returns lasting above 10 seconds. Avol provides a proxy of the annual volatility and is computed by taking the square root of 252 times the mean of either squared returns or the realized measure. r^2 represents the squared returns, $ssRV^{5m}$ denotes the 5-min realized variances, RK the realized kernels, and RV^{all} the realized variances computed using all transaction prices. The last column shows the ratio between the standard deviations of r^2 and sub-sampled realized variances computed over the full sample.

Table 2: Summary statistics of the economic indicators.

	$\hat{\mu}$	$\hat{\sigma}$	*Skewness*	*Kurtosis*	$\rho(1)$	*GDP*	*NAI*
GDP	2.622	2.258	−1.141	6.048	0.501	1.000	0.774
IP	1.669	5.300	−1.773	6.793	0.700	0.777	0.874
CP	14.186	36.585	1.872	9.969	−0.025	−0.091	−0.104
NAI	−0.242	0.841	−2.013	7.353	0.847	0.774	1.000
NO	54.732	7.223	−0.948	4.861	0.713	0.651	0.781
TS	1.749	1.236	−0.087	1.632	0.928	−0.141	−0.232

Summary statistics for the panel of economic indicators considered. The sample covers the period from from 1997(Q1) to 2011(Q4) and observations are recorded every quarter leading to a sample of 60 observations. The first four columns report the estimated mean, standard deviation, skewness and kurtosis. The fifth column report the first lag serial correlation and the last two columns report sample cross-correlations of each indicator with Real GDP and NAI, providing insight on the correlation of the indicators with the real business cycle.

3 A joint model for returns, volatilities, and economic variables

The log-returns $r_{i,t}$ are computed as $r_{i,t} = \ln(\frac{P_{i,t}}{P_{i-1,t}})$, where $P_{i,t}$ denotes the price at time i in period t. We consider daily close-to-close log-return and assume a conditionally heteroskedastic process

$$r_{i,t} = \mu + \sqrt{h_{i,t}}\, z_{i,t}, \quad \forall i = 1, \ldots, N_t, \quad \text{and} \quad \forall t = 1, \ldots, T, \tag{1}$$

where $z_{i,t} \sim i.i.d.(0,1)$. Moreover, we denote by $V[r_{i,t}|\mathcal{F}_{i-1,t}] = h_{i,t} > 0$ the conditional variance, and by $E[r_{i,t}|\mathcal{F}_{i-1,t}] = \mu \in \mathbb{R}$ the conditional expectation, where $\mathcal{F}_{i-1,t}$ is a σ-field representing the available information up to day $i-1$ during period t. Equation (1) is referred to as the return equation and will be equivalent for all the considered models.

The conditional variance $h_{i,t}$ captures the total volatility at time i and is modeled using a Realized LGARCH model with a time-varying intercept

$$\breve{h}_{i,t} = (1 - \alpha - \beta - \gamma\varphi)\breve{\tau}_t + \alpha\breve{r}^2_{i-1,t} + \beta\breve{h}_{i-1,t} + \gamma(\breve{x}_{i-1,t} - \xi), \tag{2}$$

$$\breve{x}_{i,t} = \xi + \varphi\breve{h}_{i,t} + \delta(z_{i,t}) + u_{i,t}, \tag{3}$$

where we use the notation $\breve{y}_t = \log y_t$. Equations (2) and (3) are respectively called the GARCH and the measurement equation. We set $\breve{r}^2_{i,t} = \log(r_{i,t} - \mu)^2 - \bar{z}$ and $\bar{z} = E[\log z^2_{i,t}]$, which is a bias correction.[3] Furthermore, $u_{i,t} \sim i.i.d.\ (0, \sigma^2_u)$, with $\sigma^2_u > 0$, is a sequence of random variables independent from the residuals $z_{i,t}$ of equation (1). The variable $x_{i,t}$ denotes a realized measure of volatility and if $\breve{x}_{i,t}$ is a conditionally unbiased estimator of $\breve{h}_{i,t}$, then we have $\xi = 0$ and $\varphi = 1$. Here, $\delta(z_{i,t})$ is called the leverage function, which can generate an asymmetric response in volatility to return shocks and, without loss of generality, we have $E[\delta(z_{i,t})|\mathcal{F}_{i-1,t}] = 0$. A convenient choice for the leverage function is to use Hermite polynomials.[4] Given the simple structure, the model is easy to estimate and interpret. Moreover, the

[3] From Equation (1): $\log(r_{i,t} - \mu)^2 = \breve{h}_{i,t} + \log(z_{i,t})^2$. Assuming a standard normal distribution leads to $\bar{z} \approx -1.2704$.

[4] The quadratic specification $\delta(z_t) = \delta_1 z_t + \delta_2(z^2_t - 1)$ is usually used, with δ_1 and $\delta_2 \in \mathbb{R}$.

fact that $(1-\alpha-\beta-\gamma\varphi)\ \breve{\tau}_t$ is time-varying allows to capture changes in the intercept of the Realized LGARCH model and closely relates to the local variance targeting technique for GARCH models, see Francq et al. (2011). Models with a time-varying intercept and parameters were also analysed by Amado and Teräsvirta (2013), and Amado and Teräsvirta (2014).

Besides, this model can be represented as a two-component model. We define $\breve{g}_{i,t}$ as the component capturing the short-term dynamics of volatility and $\breve{\tau}_t$ as the component representing the slow component or the long-term variance. Writing $\breve{g}_{i,t} = \breve{h}_{i,t} - \breve{\tau}_t$, the dynamics of $\breve{g}_{i,t}$ are given by

$$\breve{g}_{i,t} = \alpha(\breve{r}^2_{i-1,t} - \breve{\tau}_t) + \beta\breve{g}_{i-1,t} + \gamma(\breve{x}_{i-1,t} - \xi - \varphi\breve{\tau}_t). \tag{4}$$

Hence, our specification appears to be similar in spirit to the model of Engle and Lee (1999) who decompose the conditional variance as a sum $h_t = \tau_t + g_t$. However, we use logarithms in such a way to avoid negative values for h_t, which will be convenient for Monte Carlo simulations.

Different approaches can be analysed to model $\breve{\tau}_t$ including Splines such as in Engle and Rangel (2008). Our approach consists of filtering exogenous low-frequency economic variables with the recent MIDAS tools presented in Ghysels et al. (2005)

$$\breve{\tau}_t = m + \theta \sum_{k=1}^{K} \phi_k(\omega_1, \omega_2) X_{t-k}, \tag{5}$$

where X_t represents economic variables. The joint model is naturally named the Realized LGARCH-MIDAS model. The component $\breve{\tau}_t$ remains constant between data releases and is updated every time a new value is published. The parameter θ is useful to document counter-cyclical patterns in the secular volatility, see Engle et al. (2013), Asgharian et al. (2013), and Conrad and Loch (2014). Notice that if θ is equal to 0, we return to a classical Realized LGARCH model.

The number K represents the periods over which we smooth the economic variables. We choose the same specification for all models. To the best of our knowledge, there is no efficient procedure on model selection allowing to use the usual information criterion with our new model. It could be argued that since inference relies on standard QMLE, classical model selection tools could be used. Nonetheless, it would imply testing a very large amount of different models, searching for the optimal amount of lags in the filter at the appropriate sampling frequency for the exogenous data. Instead, we pick a specification, i.e. K=12 and quarterly data, that has proved to be successful in the literature for similar problems, see for instance Conrad and Loch (2014).

The filter ϕ_k is parametrized with a Beta lag polynomial structure[5]

$$\phi_k(\omega_1, \omega_2) = \frac{(k/(K+1))^{\omega_1 - 1}(1 - k/(K+1))^{\omega_2 - 1}}{\sum_{j=1}^{K}(j/(K+1))^{\omega_1 - 1}(1 - j/(K+1))^{\omega_2 - 1}}, \quad \text{with} \quad \sum_{k=1}^{K} \phi_k(\omega_1, \omega_2) = 1,$$

where ω_1, $\omega_2 > 0$, $0 < \frac{k}{K+1} < 1$, and $\phi_k \geq 0$ for $k = 1, \ldots, K$. Its shape is controlled by the coefficients ω_1 and ω_2. The Beta lag polynomial structure is very flexible for modeling different weighting schemes,

[5] Ghysels et al. (2005) suggest different filters for the implementation of MIDAS techniques.

although limited to uni-modal shapes.[6] The restriction $\omega_1 = 1$ and $\omega_2 > 1$ ensures a decaying pattern in the filter with maximum weight at the observation of the last period.

The parameter space is given by $\Theta = (\mu, \xi, \alpha, \beta, \gamma, \varphi, \delta_1, \delta_2, m, \theta, \omega_2)$ with the following conditions on the parameters: μ and $\xi \in \mathbb{R}$, α, β, γ and $\varphi > 0$ with $\alpha + \beta + \gamma\varphi < 1$, δ_1, δ_2 and $m \in \mathbb{R}$, and θ and $\omega_2 > 0$. The parameters of the Realized LGARCH-MIDAS are estimated using Quasi-Maximum Likelihood (henceforth QMLE – see Bollerslev and Wooldridge (1992)). Consistently with the literature (see Engle (2002), Hansen et al. (2012), Hansen and Huang (2015) and Vander Elst (2015)), we use a Gaussian distribution for our QMLE and the log-likelihood function to be maximized is given by [7]

$$l(r, x; \mathbf{\Theta}) = -\frac{1}{2} \sum_{t=1}^{T} \sum_{i=1}^{N_t} [2\log(2\pi) + \log(h_{i,t}) + (r_{i,t} - \mu)^2 / h_{i,t} + \log(\sigma_u^2) + u_{i,t}^2 / \sigma_u^2].$$

Before analysing some estimation results, we make the following remark. The fact that realized measures of volatility are included in the GARCH equation is known to improve greatly the performances of models in terms of forecasting. Nonetheless, the advantages are only striking for short-term horizon forecasts calling for additional refinements of the current models. Using a time varying intercept in the model is expected to improve the longer horizons forecasts as the intercept will adapt automatically to the current market conditions. Nonetheless, it could be argued that using a short rolling-window with the Realized LGARCH model could glean more reactive forecasts and lead to equivalent results, the cost being a smaller sample to estimate the model. The goodness of the Realized LGARCH-MIDAS lies in the fact that all the observations are used while an automated procedure allows for data-driven changes in the long-term volatility.

Table 3 contains detailed estimation results for the Realized LGARCH-MIDAS implemented with nominal corporate profits. Additionally, we provide in Figure 1 the volatility for the S&P500 tracker (SPY) estimated from the Realized LGARCH-MIDAS and implemented with the 6 different economic indicators.

The most striking observation from Table 3 concerns the parameter α, which is systematically found to be close to zero. This feature has been documented by Hansen et al. (2012) who pointed out that squared returns are not useful when realized measures of volatility are used in the model.

Second, the intercept parameter ξ of the measurement equation is negative for all the stocks suggesting that the $\breve{x}_{i,t}$ is not a conditionally unbiased estimator of $\breve{h}_{i,t}$. One of the reason relates to the use of logarithmic volatilities and was documented in Hansen and Huang (2015) and Vander Elst (2015).

It is noteworthy to mention that the parameter φ is on average close to 1 suggesting that the differences between the conditional and the realized volatility are actually captured by the intercept. The coefficients δ_1 and δ_2 of the leverage function have the expected sign and the only stocks to produce counter

[6] We refer the reader to Ghysels et al. (2006) for further details regarding the patterns one can obtain with the Beta lag structure.

[7] One could criticize our choice of the Gaussian distribution, but we can justified it by two arguments. On the one hand, it is well known that a conditional Gaussian model can display unconditional heavy tails because of changes in the volatility. On the other hand, Andersen, Bollerslev, Diebold and Ebens (2001) pointed out that log-returns standardized by realized volatility measures are close to Gaussian.

intuitive results are BSX and GIS. Additionally, the parameter θ is negative for all stocks (except for VLO) which confirms the counter-cyclical impact of nominal corporate profits, i.e. if profits increase, one expects positive returns and a decline in the long-run volatility. Finally, the constraint on the parameters $\pi = \alpha + \beta + \gamma\varphi < 1$ turns out to be respected for all stocks. Nevertheless, the element π that represents the persistence of the short-term dynamics $\breve{g}_{i,t}$ is contained in the interval $[0.95; 1]$ for most of the stocks pointing to potentially long-memory features.

Two observations from Figure 1 close this section. First, a careful visual inspection of the plots enables us to see that the choice of the economic indicator has almost no impact on the in-sample volatility paths. The differences will become clearer when forecasting volatility or conditional VaR. Second, it appears that the long-run volatility component implied real GDP growth, industrial production growth, real business cycles (NAI) and new orders is slightly lagging the total volatility process. On the contrary, nominal corporate profits seem to provide a contemporaneous component while term spreads clearly generate a leading and smoother long-run component.

4 Forecasting methodology

This section provides details on the methodology used to produce the different forecasts. We compare 8 different models, the Realized LGARCH-MIDAS model with the 6 different economic indicators separately, the Realized LGARCH, and the GARCH(1,1). The Realized LGARCH-MIDAS is specified using $K = 12$ lags of quarterly data, i.e. 3 years of variables to filter the long-term component of the model. The relevance of our approach is confirmed by the results in Section 5. For this section, we will use the notation $h_{i+k,t|i} = E\left[h_{i+k,t}|\mathcal{F}_{i,t}\right]$.[8]

4.1 Conditional volatility forecasting

We conduct a pseudo out-of-sample forecasting exercise and compare the precision of the 8 competing models for forecasts ranging from 1 day to 8 weeks (40 days) ahead. The Realized LGARCH-MIDAS models are implemented on the basis of an expanding window and the initial sample covers the period from January 4, 2000 to December 26, 2003 (i.e. 1000 observations). On the contrary, the Realized LGARCH and the GARCH(1,1) are implemented using a rolling-window containing 4 years of data, i.e. 1000 observations. The two benchmark models could also be implemented using an expanding window but it would probably lead to less accurate forecasts due to possible changes in the parameters.

Since the Realized LGARCH-MIDAS is a complete model, we can construct multi-step ahead forecasts for the latent volatility process. Iterative forecasts are constructed from the model for the k-step ahead observation $\breve{h}_{i+k,t} = \breve{g}_{i+k,t} + \breve{\tau}_t$. The total log-volatility forecast is composed of two elements that will be treated separately. On the one hand, $\breve{\tau}_t$ is considered as a random walk and will be regarded as constant

[8]For the sake of simplicity, we omit the fact that if $i + k > N_t$, the volatility falls in the next period $t + 1$ and we will still write forecasts as $h_{i+k,t|i}$.

Table 3: Estimation results for the Realized LGARCH-MIDAS with nominal corporate profits.

	α	β	γ	ξ	φ	δ_1	δ_2	θ	σ_u^2	π	$Av.\breve{g}$
AA	0.02	0.57	0.36	−0.50	1.03	−0.06	0.07	−0.03	0.15	0.95	−0.07
ABT	0.01	0.61	0.36	−0.18	0.95	−0.04	0.05	−0.04	0.18	0.96	−0.04
AES	0.01	0.56	0.48	−0.13	0.82	−0.06	0.05	−0.05	0.26	0.96	−0.07
AIG	0.01	0.44	0.60	−0.32	0.86	−0.04	0.05	−0.02	0.22	0.97	−0.18
AKS	0.02	0.65	0.22	−1.16	1.23	−0.05	0.09	−0.01	0.32	0.94	−0.08
AMD	0.02	0.59	0.32	−0.74	1.05	−0.02	0.03	−0.03	0.19	0.95	−0.10
AAPL	0.02	0.50	0.43	−0.57	1.00	−0.06	0.01	−0.04	0.21	0.95	−0.16
AXP	0.01	0.51	0.44	−0.38	0.99	−0.05	0.07	−0.05	0.17	0.96	−0.05
BA	0.01	0.55	0.40	−0.38	0.97	−0.04	0.07	−0.03	0.16	0.95	−0.02
BAC	0.02	0.41	0.58	−0.35	0.92	−0.06	0.05	−0.05	0.18	0.97	−0.08
BMY	0.01	0.59	0.34	−0.32	1.00	−0.05	0.04	−0.04	0.21	0.94	−0.07
BSX	0.02	0.72	0.21	−0.56	1.04	0.01	0.08	−0.03	0.29	0.96	−0.14
C	0.03	0.44	0.54	−0.25	0.91	−0.07	0.07	−0.05	0.17	0.96	−0.10
CAG	0.00	0.54	0.43	−0.17	0.87	−0.04	0.06	−0.04	0.24	0.92	−0.01
CAT	0.01	0.52	0.37	−0.64	1.13	−0.05	0.05	−0.02	0.16	0.95	−0.04
CHK	0.02	0.63	0.31	−0.43	0.99	−0.06	0.09	−0.03	0.26	0.96	−0.09
CSX	0.01	0.63	0.31	−0.50	1.03	−0.03	0.08	−0.03	0.19	0.96	−0.03
D	0.02	0.56	0.34	−0.19	1.09	−0.05	0.06	−0.03	0.19	0.95	−0.06
DD	0.01	0.51	0.47	−0.16	0.93	−0.06	0.05	−0.04	0.15	0.95	−0.03
DIS	0.01	0.55	0.42	−0.29	0.93	−0.06	0.05	−0.04	0.17	0.95	−0.05
DOW	0.01	0.62	0.37	−0.28	0.91	−0.04	0.06	−0.03	0.18	0.97	−0.06
EMC	0.02	0.54	0.47	−0.09	0.84	−0.04	0.04	−0.06	0.17	0.95	−0.08
FCX	0.02	0.58	0.33	−0.73	1.07	−0.07	0.08	−0.02	0.20	0.95	−0.04
GE	0.02	0.50	0.46	−0.29	0.96	−0.05	0.05	−0.05	0.17	0.96	−0.05
GIS	0.01	0.54	0.34	−0.20	1.12	0.00	0.05	−0.03	0.20	0.93	−0.03
GLW	0.02	0.51	0.44	−0.45	0.95	−0.06	0.08	−0.04	0.19	0.95	−0.08
HAL	0.02	0.60	0.39	−0.11	0.87	−0.07	0.02	−0.04	0.16	0.96	−0.09
HD	0.02	0.49	0.53	−0.14	0.83	−0.03	0.03	−0.05	0.16	0.94	−0.04
IBM	0.01	0.48	0.49	−0.29	0.93	−0.06	0.05	−0.04	0.15	0.95	−0.04
INTC	0.02	0.44	0.57	−0.23	0.86	−0.04	0.02	−0.05	0.13	0.94	−0.02
JCP	0.01	0.60	0.32	−0.67	1.10	−0.04	0.09	−0.03	0.21	0.95	−0.02
JNJ	0.02	0.58	0.39	−0.15	0.94	−0.03	0.06	−0.04	0.17	0.97	−0.12
JPM	0.02	0.46	0.50	−0.29	0.94	−0.06	0.07	−0.06	0.16	0.95	−0.08
KEY	0.03	0.55	0.44	−0.21	0.92	−0.03	0.07	−0.04	0.22	0.98	−0.18
KO	0.01	0.54	0.45	−0.15	0.90	−0.04	0.06	−0.04	0.16	0.96	−0.05
MCD	0.01	0.65	0.30	−0.28	1.03	−0.05	0.08	−0.03	0.18	0.97	−0.06
MDT	0.00	0.52	0.37	−0.48	1.11	−0.04	0.04	−0.03	0.20	0.93	−0.02
MMM	0.01	0.53	0.37	−0.37	1.09	−0.04	0.05	−0.03	0.18	0.94	−0.03
MO	0.01	0.59	0.43	−0.24	0.81	−0.05	0.07	−0.05	0.23	0.94	−0.02
MSFT	0.01	0.46	0.54	−0.26	0.87	−0.03	0.03	−0.05	0.14	0.93	−0.03
NBR	0.02	0.62	0.32	−0.50	1.03	−0.06	0.06	−0.02	0.16	0.97	−0.04
NEM	0.01	0.59	0.31	−0.73	1.15	−0.04	0.08	−0.02	0.14	0.95	−0.03
ORCL	0.02	0.46	0.54	−0.17	0.87	−0.05	0.06	−0.04	0.16	0.94	−0.05
PFE	0.02	0.54	0.37	−0.32	1.03	−0.05	0.05	−0.03	0.16	0.94	−0.06
SLB	0.02	0.57	0.36	−0.41	1.06	−0.06	0.05	−0.02	0.13	0.96	−0.04
SPY	0.01	0.42	0.51	−0.50	1.00	−0.13	0.03	−0.04	0.15	0.95	−0.04
TJX	0.01	0.55	0.37	−0.28	0.99	−0.01	0.08	−0.04	0.21	0.92	−0.01
USB	0.01	0.50	0.51	−0.11	0.89	−0.04	0.08	−0.05	0.21	0.96	−0.06
UTX	0.01	0.53	0.48	−0.24	0.86	−0.05	0.02	−0.04	0.18	0.95	−0.01
VLO	0.02	0.63	0.30	−0.75	1.10	−0.05	0.09	0.00	0.20	0.97	−0.06
WFC	0.02	0.53	0.48	−0.16	0.90	−0.05	0.06	−0.05	0.17	0.98	−0.11
WMT	0.02	0.60	0.34	−0.21	1.01	−0.03	0.06	−0.04	0.15	0.96	−0.05
WY	0.01	0.59	0.36	−0.39	1.01	−0.04	0.07	−0.03	0.16	0.96	−0.05
XOM	0.01	0.51	0.41	−0.28	1.05	−0.10	0.05	−0.03	0.14	0.95	−0.02
XRX	0.03	0.62	0.36	−0.22	0.81	−0.02	0.05	−0.06	0.25	0.95	−0.17

Estimation results of the parameters $\alpha, \beta, \gamma, \xi, \varphi, \delta_1, \delta_2, \theta$ for the 55 stocks of our dataset over the full sample covering the period from January 4, 2000 to March 30, 2012. The elements $\pi = \alpha + \beta + \gamma\varphi$ and $Av.\breve{g}$ represent the persistence and the sample mean of the dynamic component, respectively.

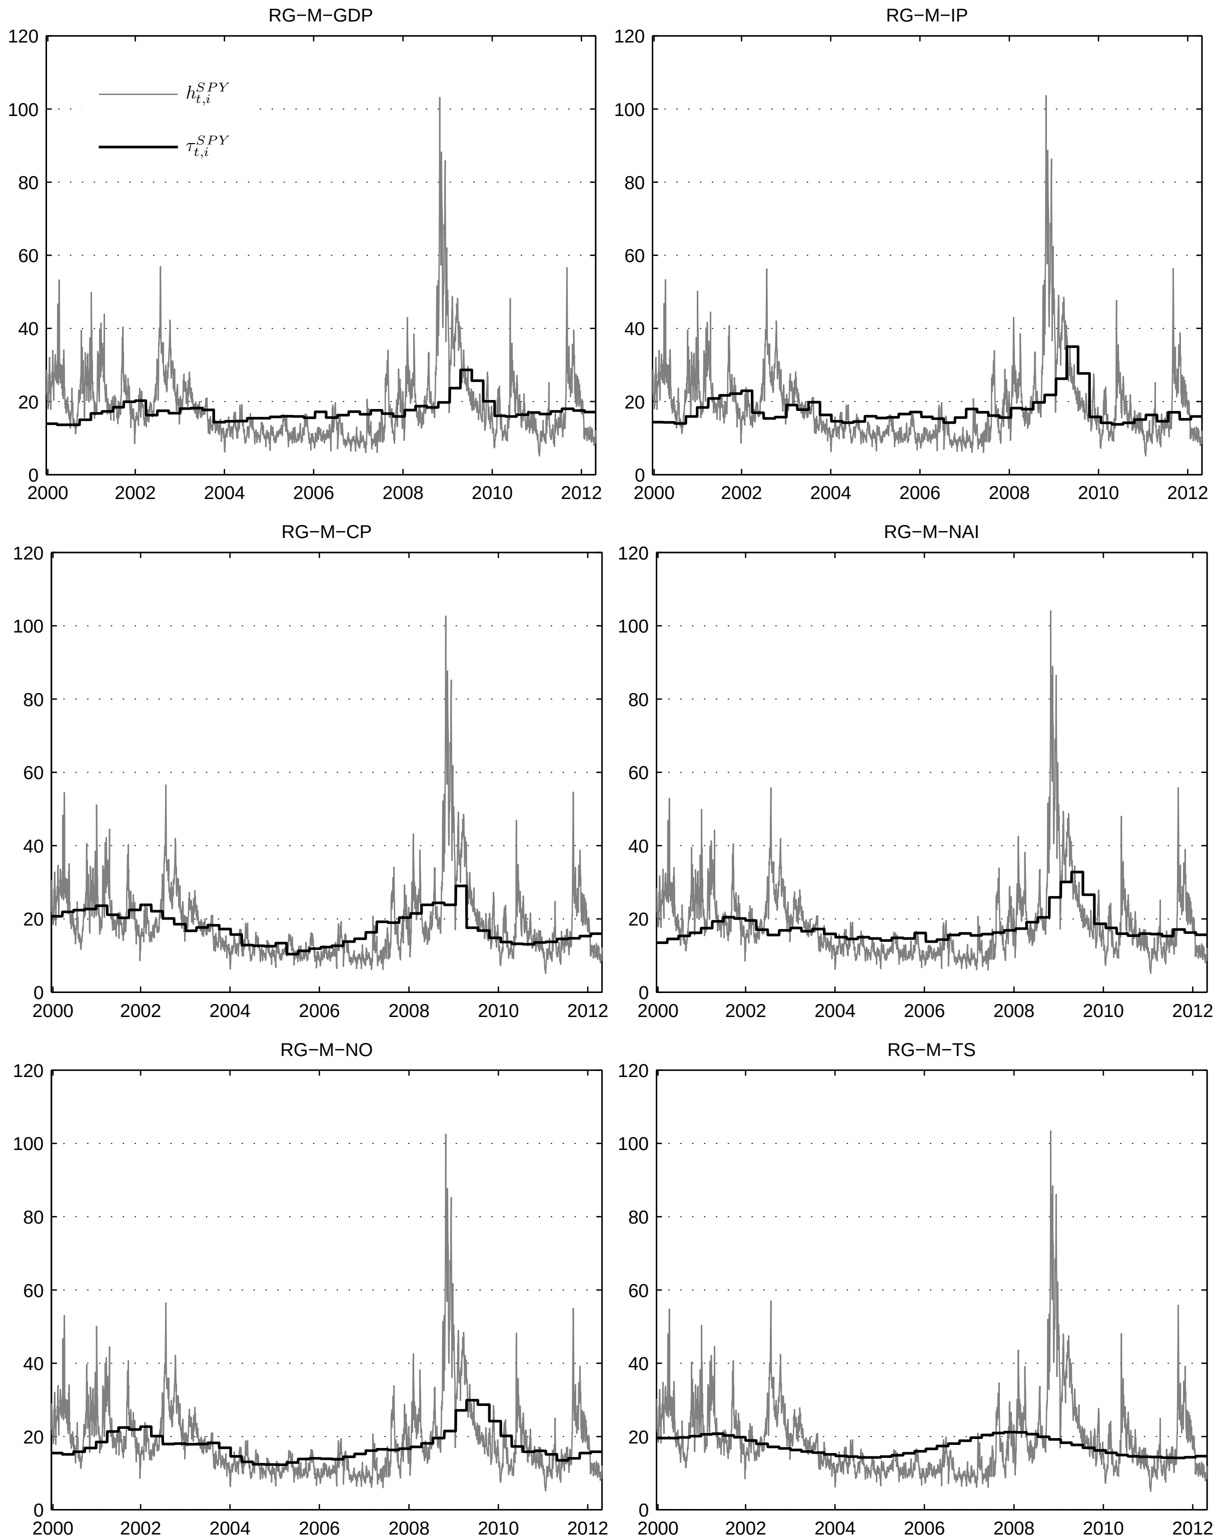

Figure 1: The 6 figures provide the volatility path estimated from the Realized LGARCH-MIDAS model for SPY and implemented with the 6 different considered indicators (sample for returns: January 4, 2000 to March 30, 2012 – sample for economic variables: 1997(Q1) to 2011(Q4)). The grey line represents the annualized conditional volatility $h_{i,t}$ and the black thick line, the annualized component $\tau_{i,t}$.

for all horizons even though some dates may fall in the next period. Different approaches could be used to produce more dynamic forecasts but this is out of the scope of this work. The random walk approach is expected to generate more persistent forecasts. In contrasts, a classical time series model, like an AR(1) process, is expected to imply quicker mean-reversion in the forecasts than the random walk. Our results in Section 5 support this naïve procedure.

On the other hand, forecasts for $\breve{g}_{i+k,t}$ can be produced easily by using iteratively the model for the k-step ahead observation

$$\breve{g}_{i+k,t} = \alpha(\breve{r}^2_{i+k-1,t} - \breve{\tau}_t) + \beta\breve{g}_{i+k-1,t} + \gamma(\breve{x}_{i+k-1,t} - \xi - \varphi\breve{\tau}_t), \tag{6}$$

$$\breve{x}_{i+k,t} = \xi + \varphi\breve{h}_{i+k,t} + \delta(z_{i+k,t}) + u_{i+k,t}, \tag{7}$$

where $\delta(z_{i+k,t}) + u_{i+k,t}$ is a martingale difference sequence. By re-arranging the different variables, we obtain a zero-mean AR(1) representation for the dynamic component

$$\breve{g}_{i+k,t} = \mathcal{A}\,\breve{g}_{i+k-1,t} \;+\; \mathcal{B}\,\boldsymbol{\epsilon}_{i+k-1,t}, \tag{8}$$

where $\mathcal{A} = \alpha + \beta + \gamma\varphi$, $\mathcal{B} = \begin{bmatrix}\gamma & \gamma & \alpha\end{bmatrix}$, and $\boldsymbol{\epsilon}_{i+k-1,t} = \begin{bmatrix}\delta(z_{i+k-1,t}) & u_{i+k-1,t} & \log(z_{i+k-1,t})^2 - \bar{z}\end{bmatrix}'$. Moreover, $\mathcal{B}\,\boldsymbol{\epsilon}_{i+k-1,t}$ is also a martingale difference sequence. Hence, the k-step ahead forecast can easily be computed as $\breve{h}_{i+k,t|i} = \breve{g}_{i+k,t|i} + \breve{\tau}_t$. However, by Jensen's inequality, the use of logarithms implies that $E[\log h_{i+k,t}|\mathcal{F}_{i-1,t}] \leq \log E[h_{i+k,t}|\mathcal{F}_{i-1,t}]$. We cope with this issue by using Monte Carlo simulations to compute direct forecasts for $h_{i+k,t}$. The residuals of the model can be generated from

$$\zeta^{(n)}_{i+k,t} := \begin{pmatrix} z^{(n)}_{i+k,t} \\ u^{(n)}_{i+k,t} \end{pmatrix} \sim \mathcal{N}_2\left(0, \begin{bmatrix} 1 & 0 \\ 0 & \sigma^2_u \end{bmatrix}\right),$$

where $k = 1, ..., H$, $n = 1, \ldots, N$ and $\sigma^2_u > 0$. Using the sequence of simulated residuals, a path for the volatility for the next H steps can be constructed by using equation (8) and $\breve{h}^{(n)}_{i+k,t} = \breve{g}^{(n)}_{i+k,t} + \breve{\tau}_t$. Consistent estimates of $h_{i+k,t}$ can be obtained at each horizon from $\frac{1}{N}\sum_{n=1}^{N}\exp(\breve{h}^{(n)}_{i+k,t})$. An equivalent procedure is applied to the Realized LGARCH following insights from Hansen et al. (2014) and we use the usual recursions to produce forecasts for the GARCH(1,1). In our procedure, we will use $N = 1000$ simulations for each horizon.

4.2 Conditional density and VaR forecasting

The Realized LGARCH-MIDAS has the flexibility to provide forecasts for the full conditional density of returns from which the conditional VaR can be extracted. Because returns are not independent and volatility is time-varying, there is generally no analytical expression for the multi-step ahead conditional density and numerical techniques have to be used. In the context of multi-period VaR forecasting, Christoffersen (2003) suggested to use either standard Monte Carlo simulations or filtered historical simulations. Our approach relies on Monte Carlo simulations. The conditional distribution of multi-step ahead returns comes as a by-product of the procedure described in Subsection 4.1. Simulations for the volatility can be obtained using equation (8) and $\breve{h}_{i+k,t} = \breve{g}_{i+k,t} + \breve{\tau}_t$. Denoting the k-step ahead n-th simulated volatility by $h^{(n)}_{i+k,t} = \exp(\breve{h}^{(n)}_{i+k,t})$, we obtain

$$r^{(n)}_{i+k,t} = \mu + \sqrt{h^{(n)}_{i+k,t}}\, z^{(n)}_{i+k,t}, \qquad k = 1, \dots, H, \quad \text{and} \quad n = 1, \dots, N. \tag{9}$$

From the series of simulated daily log-returns, we can estimate the conditional density and extract conditional moments and quantiles of interest such as the VaR. Moreover, the simulated series allow to conveniently compute returns over a longer period by aggregating the returns at different horizons. We denote the cumulative returns as $r^{(n),hor}_{i,t} = \sum_{d=1}^{hor} r^{(n)}_{i+d,t}$ and obtain N vectors of returns with horizons ranging from 1 to 40 days. For instance, for $hor = 2$, we have that $r^{(n),2}_{i,t} = r^{(n)}_{i+1,t} + r^{(n)}_{i+2,t}$. Those series constitute the basis to describe the conditional density of returns on longer horizons and will be used to compute the VaR over several periods. An equivalent procedure is applied to the Realized LGARCH and the GARCH(1,1) models which are easy to simulate.

5 Results

We will make repeated use of the MCS of Hansen et al. (2011). The MCS is a convenient tool to statistically compare many competing models. It relies on recursive testing and elimination of poor performing models providing a data-driven optimal reduced set of models that are statistically not distinguishable in terms of forecasting performances. The analysis is performed for different loss functions for all stocks and we report the frequencies at which each model was included in the MCS at 5% level. The MCS procedure is implemented with 10,000 bootstrap replications and using a 5% level of confidence. We evaluate the quality of the forecasts using 5-min sub-sampled realized variances as proxy for the true latent conditional volatility.

5.1 Conditional volatility forecasts evaluation

We evaluate the volatility forecasts using two approaches. First, we provide the R^2 from simple Mincer-Zarnowitz regressions, see Mincer and Zarnowitz (1969), computed from

$$ssRV^{5m}_{h+i,t} = \alpha + \beta\, h^{\mathcal{M}}_{h+i,t|i} + \epsilon_{i,t}, \tag{10}$$

where $h^{\mathcal{M}}_{h+i,t|i}$ denotes the h-step ahead forecasts for model $\mathcal{M}$. Models providing accurate forecasts should be characterized by high R^2 pointing to their ability to explain the variability of the conditional volatility proxy. Additionally, we compare the forecasts at different horizons using the MSE loss function

$$L(ssRV^{5m}_{h+i,t},\, h^{\mathcal{M}}_{h+i,t|i}) = (h^{\mathcal{M}}_{h+i,t|i} - ssRV^{5m}_{h+i,t})^2,$$

and the QLIKE loss function studied by Patton (2011) and which is defined as

$$L(ssRV^{5m}_{h+i,t},\, h^{\mathcal{M}}_{h+i,t|i}) = \log h^{\mathcal{M}}_{h+i,t|i} + \frac{ssRV^{5m}_{h+i,t}}{h^{\mathcal{M}}_{h+i,t|i}}.$$

Both functions are shown by Patton (2011) to be robust and are used as input for the MCS at the different horizons considered. A complete review of useful techniques to compare volatility and correlation forecasts is provided by Patton and Sheppard (2009) and it includes a section covering loss functions in more details.

Figure 2 provides the average R^2 for the different models computed from the 55 stocks. We observe from Figure 2 that the Realized LGARCH provides more accurate forecasts for short horizons. Clearly, the R^2 computed from the Mincer-Zarnowitz regression is higher for horizons ranging from 1 to 7 days. For longer periods, the R^2 of the different models decreases which visually confirms that it is more difficult to forecast at longer horizons. Moreover, the Realized LGARCH model curve seems to decrease at a faster rate and displays strong convexity. In contrast, the Realized LGARCH-MIDAS has slower linear decay and this points to higher performances to forecast volatility at longer horizons. In fact, the Realized LGARCH-MIDAS outperforms the competing models for most of the horizons beyond 5 days. This is especially the case for nominal corporate profits and term spreads that provide uniformly more precise forecasts from horizon 5-7 to 40 steps ahead. The Realized LGARCH-MIDAS implemented with nominal corporate profits provides clearly the most impressive results. Finally, it appears that Realized LGARCH only performs well at short horizons and is outperformed by the GARCH(1,1) for long horizons. In fact, the GARCH(1,1) provides more accurate long-run forecasts than real GDP growth, industrial production growth, and the real business cycle indicator NAI.

Table 4 provides the frequency at which the models are included in the MCS for different loss functions. We display results for a few horizons ranging to 40 steps ahead on the basis of a sample covering the period from January 4, 2000 to March 30, 2012. Panel A and B provide results for the MSE and the QLIKE loss functions, respectively. The results also suggest that the Realized LGARCH provides more precise forecasts at short horizons but is outperformed by the Realized LGARCH-MIDAS at longer horizons, and in particular by the model implemented with nominal corporate profits. For both loss functions, our model provides the best result and is most frequently included in the MCS. Notice that the GARCH(1,1) is also often included in the MCS confirming earlier finding of Hansen and Lunde (2005) on the usefulness of a simple GARCH(1,1) model for volatility forecasting.

5.2 Conditional density forecasts evaluation

We compare the models' ability to forecast the conditional density using the predictive likelihood measure studied in Amisano and Giacomini (2007) and Maheu and McCurdy (2011). We analyse the ability of the different models to produce multi-step ahead density forecasts. Let us modify a bit our notation for the next two subsections in order to make the presentation easier. We denote returns by $r_t,\ t = 1, ..., T$ and the average predictive likelihood over the sample is computed for $k \geq 1$ as

$$D_k = \frac{1}{T - 1000 - 40 + 1} \sum_{t=1000+40-k}^{T-k} \log f_k \left(r_{t+k} | \mathcal{F}_t, \Theta \right),$$

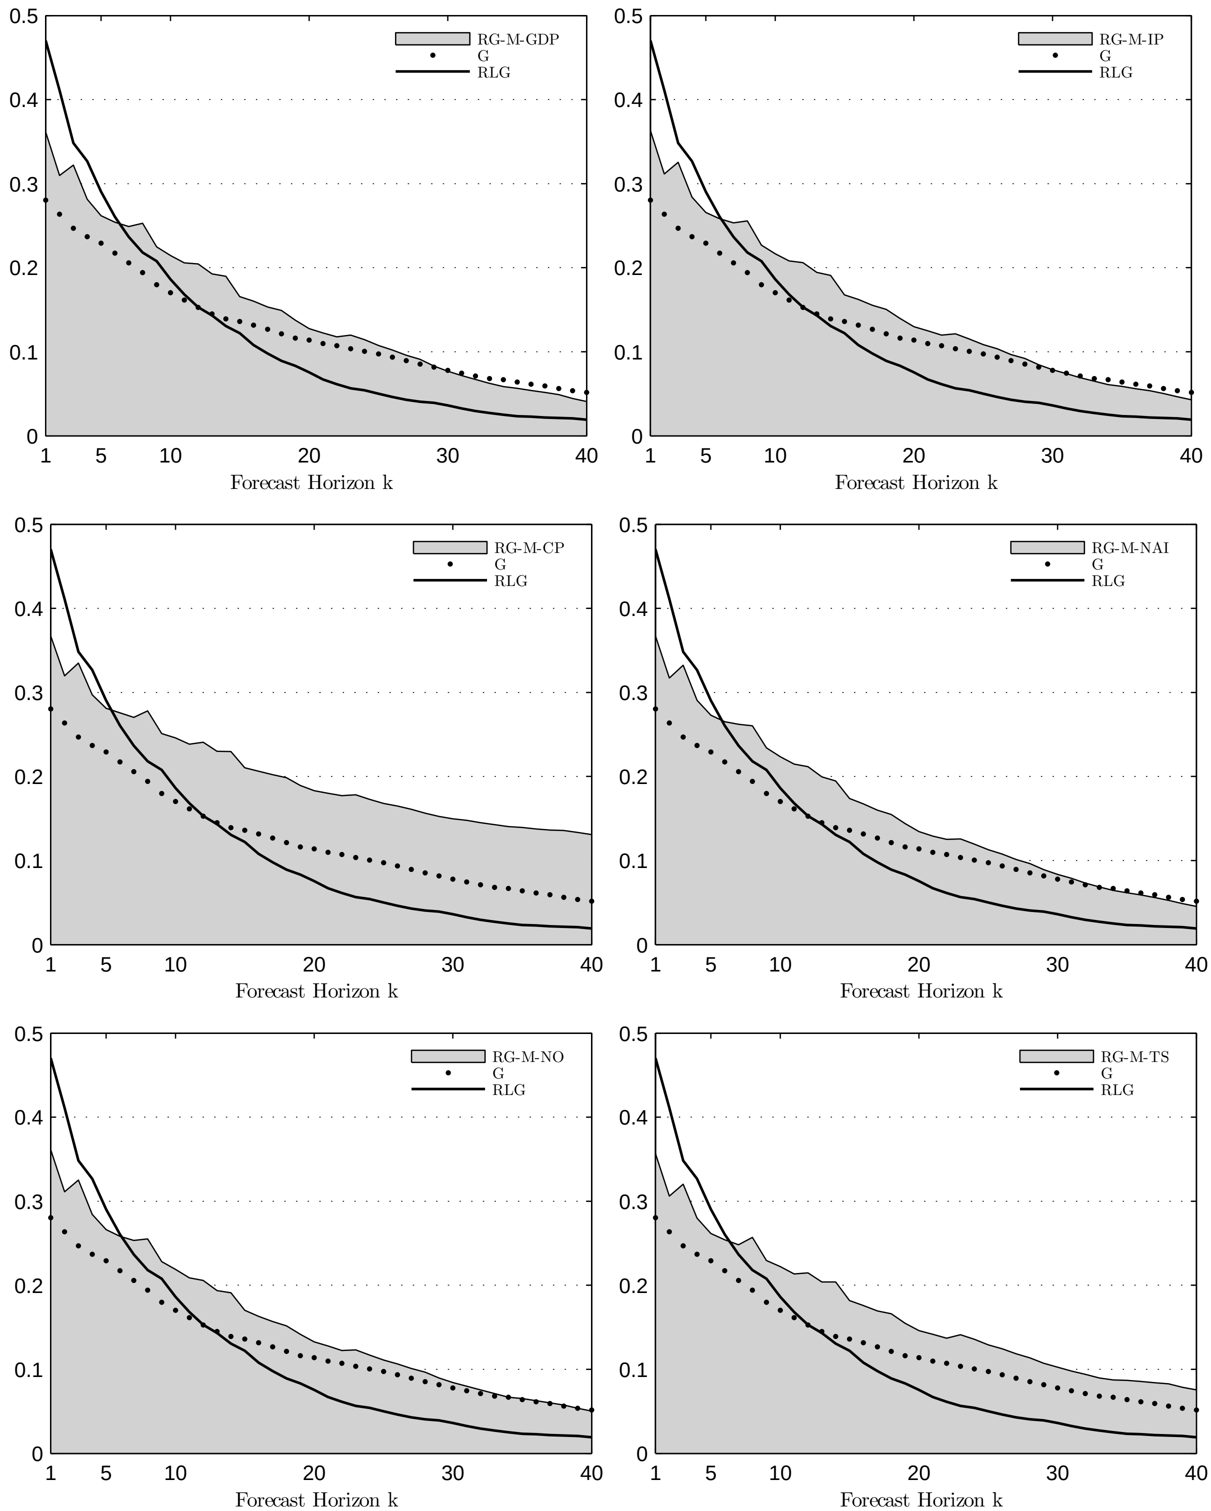

Figure 2: Average Mincer-Zarnowitz R^2 for the GARCH(1,1) (G), Realized LGARCH (RLG) and Realized LGARCH-MIDAS (RG-M-X).

Table 4: Model Confidence Set Results.

Horizon	1	5	10	15	20	25	30	35	40
Panel A: Volatility forecasting – Loss=MSE									
$G(1,1)$	0.22	0.29	0.31	0.36	0.31	0.31	0.31	0.31	0.31
RLG	0.80	0.62	0.25	0.15	0.07	0.05	0.04	0.04	0.02
RLG-M-GDP	0.13	0.07	0.04	0.05	0.07	0.11	0.09	0.07	0.07
RLG-M-IP	0.09	0.05	0.04	0.04	0.04	0.07	0.07	0.05	0.05
RLG-M-CP	0.40	0.58	0.69	0.75	0.69	0.69	0.65	0.64	0.64
RLG-M-NAI	0.07	0.07	0.02	0.04	0.02	0.02	0.02	0.00	0.00
RLG-M-NO	0.05	0.11	0.15	0.13	0.16	0.20	0.18	0.11	0.11
RLG-M-TS	0.20	0.16	0.13	0.20	0.11	0.13	0.05	0.05	0.04
Panel B: Volatility forecasting – Loss=QLIKE									
$G(1,1)$	0.04	0.31	0.42	0.42	0.44	0.45	0.40	0.36	0.38
RLG	0.96	0.42	0.11	0.07	0.02	0.02	0.02	0.02	0.02
RLG-M-GDP	0.04	0.07	0.04	0.02	0.02	0.02	0.04	0.04	0.04
RLG-M-IP	0.00	0.05	0.04	0.04	0.02	0.00	0.02	0.02	0.04
RLG-M-CP	0.02	0.56	0.65	0.69	0.67	0.71	0.73	0.75	0.75
RLG-M-NAI	0.04	0.07	0.04	0.04	0.04	0.04	0.04	0.04	0.04
RLG-M-NO	0.00	0.07	0.04	0.05	0.02	0.02	0.00	0.00	0.00
RLG-M-TS	0.02	0.05	0.05	0.02	0.00	0.00	0.00	0.00	0.00
Panel C: Density forecasting – Loss=-Pred. Density									
$G(1,1)$	0.15	0.29	0.27	0.24	0.20	0.24	0.20	0.16	0.13
RLG	0.96	0.95	0.76	0.27	0.18	0.15	0.09	0.09	0.07
RLG-M-GDP	0.15	0.47	0.44	0.33	0.24	0.20	0.16	0.15	0.09
RLG-M-IP	0.15	0.53	0.51	0.38	0.33	0.29	0.22	0.18	0.20
RLG-M-CP	0.20	0.69	0.78	0.87	0.82	0.84	0.85	0.84	0.85
RLG-M-NAI	0.16	0.60	0.53	0.35	0.22	0.25	0.16	0.15	0.09
RLG-M-NO	0.13	0.51	0.44	0.27	0.20	0.22	0.15	0.13	0.15
RLG-M-TS	0.13	0.38	0.25	0.18	0.13	0.11	0.13	0.11	0.09
Panel D: VaR forecasting – Loss=Tick Loss									
$G(1,1)$	0.85	0.75	0.67	0.62	0.53	0.49	0.47	0.51	0.49
RLG	0.96	0.96	0.67	0.64	0.42	0.29	0.11	0.11	0.05
RLG-M-GDP	0.78	0.78	0.62	0.47	0.49	0.33	0.33	0.27	0.22
RLG-M-IP	0.78	0.73	0.64	0.58	0.47	0.42	0.25	0.24	0.16
RLG-M-CP	0.78	0.75	0.64	0.53	0.35	0.31	0.27	0.18	0.18
RLG-M-NAI	0.58	0.44	0.25	0.24	0.16	0.13	0.13	0.13	0.09
RLG-M-NO	0.56	0.51	0.33	0.22	0.20	0.20	0.15	0.13	0.13
RLG-M-TS	0.78	0.75	0.67	0.65	0.58	0.53	0.55	0.55	0.56

MCS results for different loss functions. The entries of the table correspond to the frequency at which a model is included in the MCS at 5%. We provide results for a few horizons ranging to 40 steps ahead on the basis of a sample covering the period from January 4, 2000 to March 30, 2012. The pseudo out-of-sample exercise relies on a rolling window of size 1000 for the Realized LGARCH (RLG) and the GARCH(1,1) (G(1,1)) and on an expanding window with initial sample of size 1000 for the Realized LGARCH-MIDAS models (RLG-M-X). The abbreviations in the RLG-M-X correspond to the economic variables, see Section 2. Panel A and B provide volatility forecasting results (MSE and QLIKE loss functions). Panel C provides results related to conditional density forecasts while Panel D summarizes results for the VaR forecasts.

where $f_k(\cdot|\mathcal{F}_t, \Theta)$ denotes the k-period ahead predictive density and is evaluated at the realized return r_{t+k}. The numbers 1000 and 40 correspond to the initial sample size and the maximum horizon over which forecasts are produced. It is computed for $k = 1$ using the Gaussian distribution and for $k > 1$ on the basis of Monte Carlo simulation. It is the key input to compute the predictive likelihood and, using N draws, it is evaluated at the observed return r_{t+k}

$$f_k(r_{t+k}|\mathcal{F}_t, \Theta) \approx \frac{1}{N}\sum_{n=1}^{N} f_k\left(r_{t+k}|\ \mu,\ h_{t+k}^{(n)}\right),$$

where $h_{t+k}^{(n)}$ is simulated following the algorithm from Subsection 4.1 and $f_k\left(\cdot|\ \mu,\ h_{t+k}^{(n)}\right)$ is the Gaussian distribution with mean μ and variance $h_{t+k}^{(n)}$. The intuition is that models capturing properly the data features will produce larger values for D_k.

Figure 3 provides the sample mean of the average predictive likelihood for the competing models at different horizons. Results confirm the findings from the previous subsection that the Realized LGARCH is performing well at short horizons but produces poor long-run forecasts. Moreover, it appears clearly that the Realized LGARCH-MIDAS is performing better than the other models at long horizons. In particular, the Realized LGARCH-MIDAS implemented with nominal corporate profits provides the most convincing results displaying a slow level of decay of the average predictive likelihood.

We use the MCS with the time series $-\log f_k(r_{t+k}|\mathcal{F}_t, \Theta)$ as input. Notice that we use a minus in front of the log density in order to translate it to a loss function. Panel C from Table 4 contains the results and confirms the findings from Figure 3. Again, it appears that models including economic variables are performing very well and the Realized LGARCH-MIDAS-CP has clearly more power to forecast at longer horizons.

5.3 Conditional VaR forecasts evaluation

We provide multi-period VaR forecasts for periods ranging from 1 day to 40 days at probability level $p = 1\%$. Results for multi-period VaR forecasts are analysed using the tick loss function usually employed in conditional quantiles regressions, see Komunjer (2005). Keeping our notation from the previous subsection, the tick loss function is defined as

$$TL(p, hor) = \frac{1}{T - hor - 1000 + 1}\sum_{t=1000}^{T-hor}\left(p - \mathbf{1}_{\left\{r_t^{hor} < VaR_t^{hor,p}\right\}}\right)\left(r_t^{hor} - VaR_t^{hor,p}\right), \tag{11}$$

where r_t^{hor} denotes the cumulative returns over $hor = 1, ..., 40$ days, and is computed from daily log-returns $r_t^{hor} = \sum_{d=1}^{hor} r_{t+d} = \ln(P_{t+hor}/P_t)$. Moreover, $VaR_t^{hor,p}$ denotes the VaR forecast at time t for the period of time covering $[t; t + hor]$. Notice that results from this subsection should not be compared with those presented in Subsection 5.2 because we try to forecast the VaR over a period covering several days whereas the average predictive likelihood was informative about density forecasts quality for daily future returns $r_{i+k,t}$. Hence, even though the two problems are connected, differences in the results can appear.

Figure 4 provides the average losses at different horizons. The first observation is that the loss increases with the length of the period over which the VaR is computed. However, the average exception rate

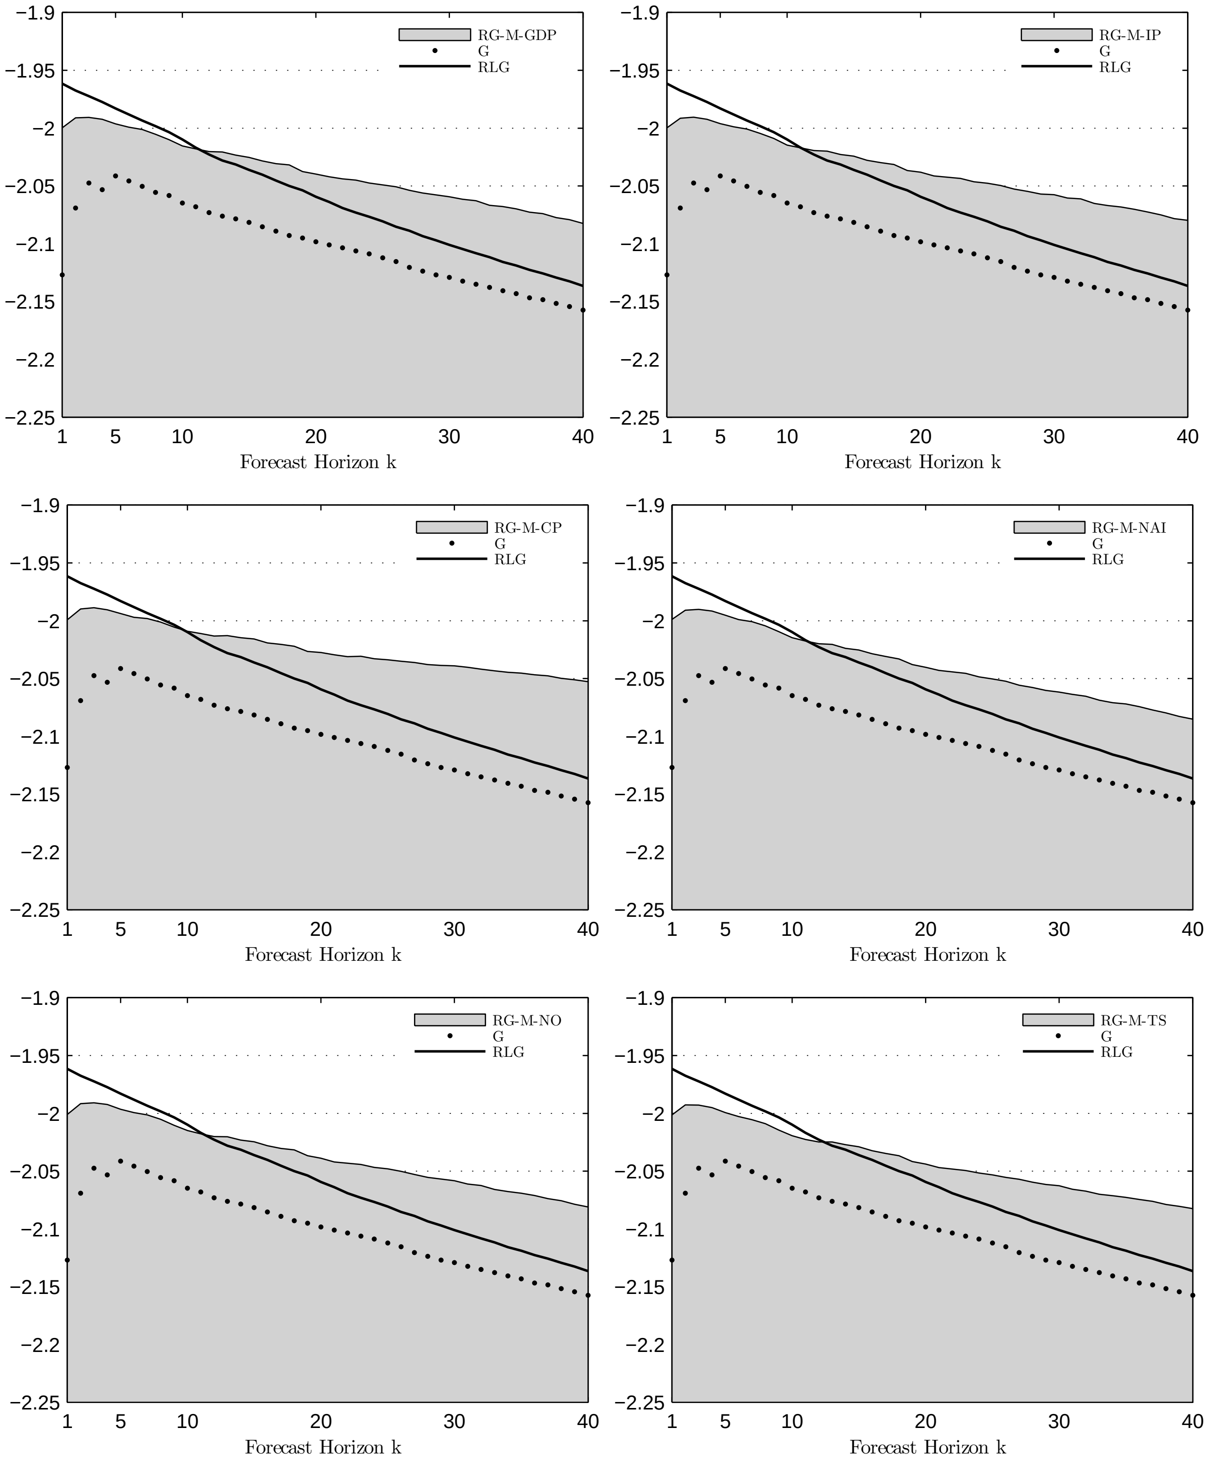

Figure 3: Average predictive likelihood for the GARCH(1,1) (G), Realized LGARCH (RLG) and Realized LGARCH-MIDAS (RG-M-X).

remains constant and the increase in the tick loss only corresponds to larger deviations between returns and the corresponding VaR. Moreover, it clearly appears that real GDP growth, industrial production growth, nominal corporate profits, and term spreads provide the best results among the economic variables. Moreover, they provide more accurate multi-period VaR forecasts than the Realized LGARCH for horizons beyond 10-days. For shorter horizons, it is difficult to draw conclusions from these plots as the three models seem to coincide for most of the periods. Moreover, we find that the Realized LGARCH-MIDAS with term spreads is the most attractive model providing lower losses at long horizons. The other models often provide close results to the GARCH(1,1) model which turns out to be a very attractive model for VaR forecasting. Only the Realized LGARCH-MIDAS-TS systematically outperforms the GARCH(1,1). In contrast, our models based on real business cycle indicators and new orders index provide somewhat disappointing results and seem to be uniformly outperformed both by the Realized LGARCH and the GARCH(1,1) models. In particular, the model based on new orders provide extremely bad results.

We use the Tick losses $\left(p - \mathbf{1}_{\left\{r_t^{hor} < VaR_t^{hor,p}\right\}}\right)\left(r_t^{hor} - VaR_t^{hor,p}\right)$ as input for the MCS to statistically compare the ability of the models to forecast the multi-period VaR. Results are reported in Panel D of Table 4. We find out that the Realized LGARCH model is very difficult to beat at short horizons and is most frequently included in the model confidence set. Nonetheless, as from horizon 10, we find that both the GARCH(1,1) and the the Realized LGARCH-MIDAS implemented with term spread are performing equally well. From horizon 15, the Realized LGARCH-MIDAS-TS is the most attractive model to compute multi-period ahead forecasts. The fact that models with realized measures of volatility perform well at short horizon (e.g. 1 day) is not surprising and confirms earlier findings of Brownlees and Gallo (2009). Nevertheless, it must be acknowledged that the GARCH(1,1) performs impressively well at long horizons suggesting that realized measure of volatility may be only useful at short horizons.

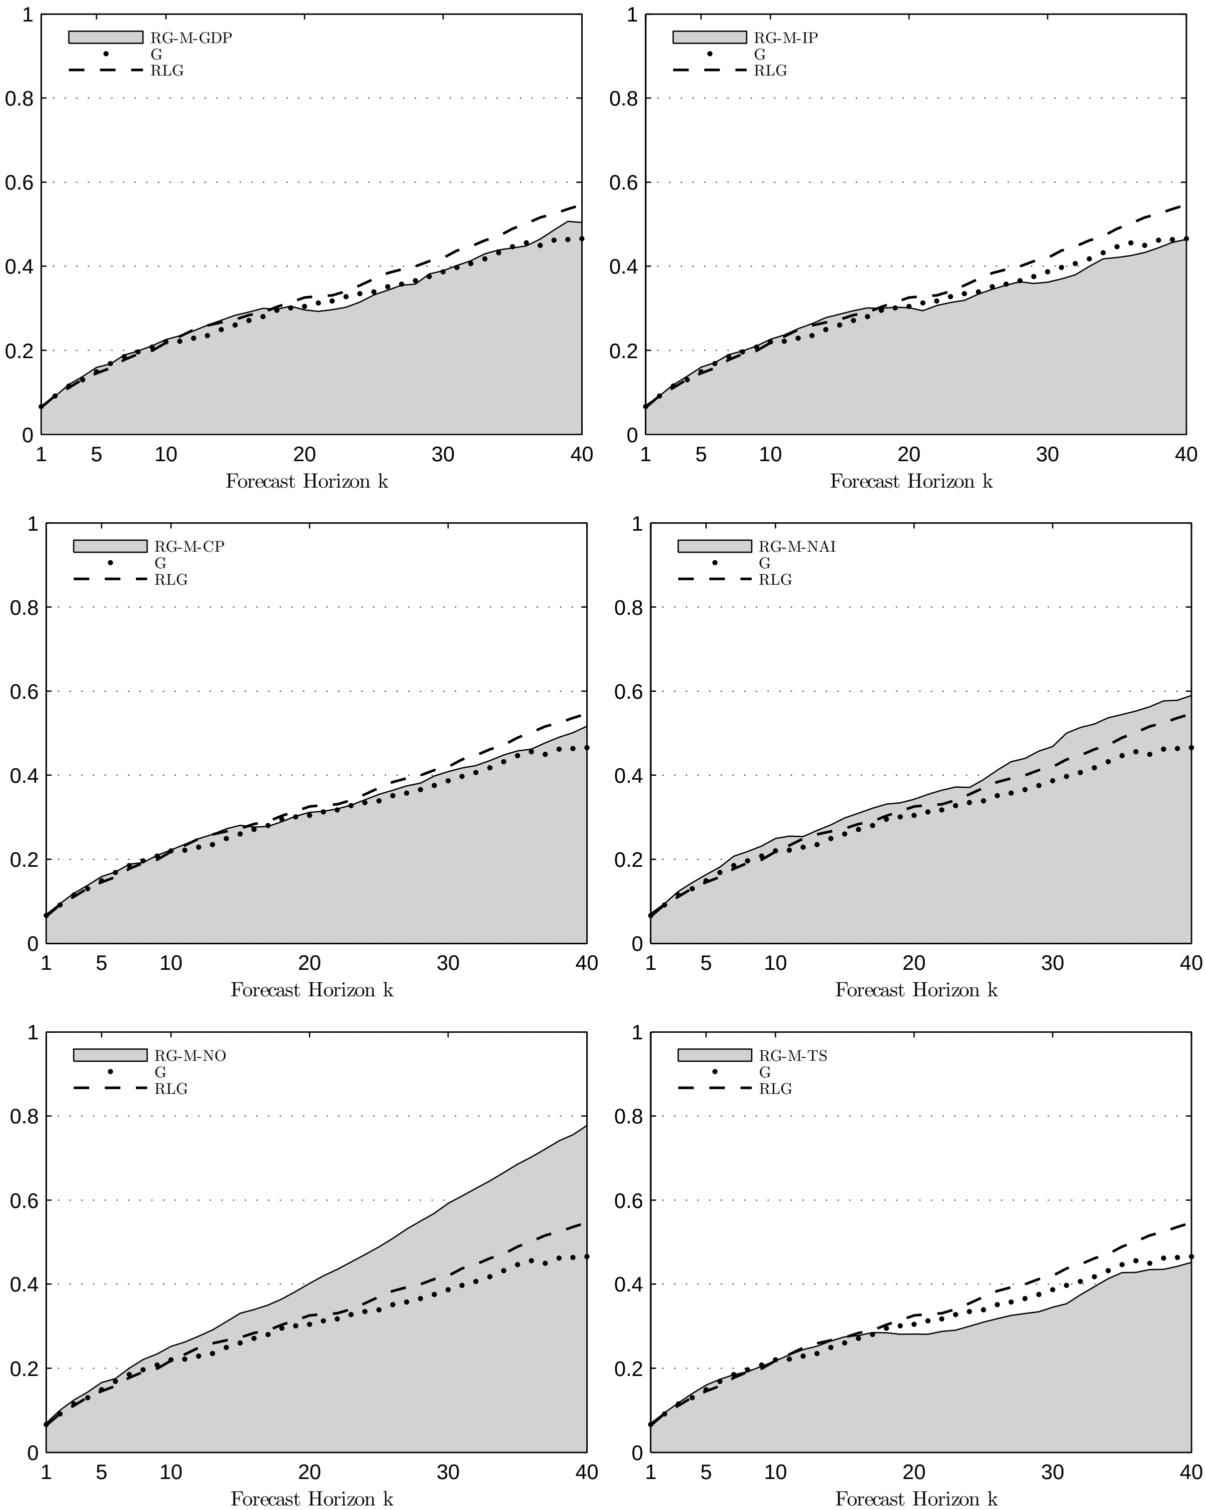

Figure 4: Average tick losses for the GARCH(1,1) (G), Realized LGARCH (RLG) and Realized LGARCH-MIDAS (RG-M-X).

6 Conclusion and outlook on future research

In this paper, we investigate the forecasting power of models relying on both realized measures of volatility and exogenous economic indicators. We analyse their ability to forecast the conditional volatility of returns as well as their full conditional density, and the conditional multi-period VaR. Our framework relies on a refined extension of the Realized LGARCH model and includes a time-varying intercept that accommodates for changes in the long-run volatility of the returns. Hence, it allows macroeconomic information to be included directly into the estimation and forecast procedures. We use MIDAS techniques to filter the low-frequency data and call our model the Realized LGARCH-MIDAS model. We perform an analysis on a panel of 55 U.S. stocks relying on more than 10 years of high-frequency data, and using 6 broad economic indicators of the U.S. economy including real GDP growth, industrial production growth, nominal corporate profits growth, a real business cycle indicators, a new orders index, and term spreads.

We find that the Realized LGARCH-MIDAS model implemented with nominal corporate profits and term spreads provides the best results. It outperforms the Realized LGARCH and the GARCH(1,1) for horizons beyond two business weeks in terms of volatility, conditional density and VaR forecasting. Our model is outperformed by the Realized LGARCH only for short horizons. This finding advocates in favor of our model to compute the regulatory 10 and 15 days-ahead VaR, which are required for most of the financial institutions. We provide more empirical evidence about the usefulness of the models and their good performance by using the MCS to compare a set of loss functions capturing important practical features such as density forecasts precision.

The focus of this paper has been set on comparing many stocks rather than many possible specifications, and it produced already interesting and intriguing results. Nonetheless, there is clear room for further exploration of the setting and many things could be done in terms of variable selection, specification of the model, distributional assumptions and variables combinations. For instance, our working model could be generalized to more flexible distributions and exploring the impact of non-Gaussian innovation distributions constitutes probably a very attractive task for VaR. However if one wishes to perform model selection, try out different distributions or variable combinations, the computational burden could possibly become quickly unmanageable. Moreover, one could as well investigate the impact of using higher-frequencies to update the model more frequently. The recent techniques related to *nowcasting* could be used to provide more timely updates of the infrequently released economic variables and produce more reactive forecasts. Ideally, sub-samples should be used to test for breaks in macro and financial series along with further investigation for spurious correlations. An other interesting development would be to link each asset with its own series of profits. These are just some of the research avenues which are left for future work.

References

Amado, C. and Teräsvirta, T. (2013), 'Modelling volatility by variance decomposition', *Journal of Econometrics* **175**(2), 142–153.

Amado, C. and Teräsvirta, T. (2014), 'Modelling changes in the unconditional variance of long stock return series', *Journal of Empirical Finance* **25**, 15–35.

Amisano, G. and Giacomini, R. (2007), 'Comparing density forecasts via weighted likelihood ratio tests', *Journal of Business & Economic Statistics* **25**(2), 177–190.

Andersen, T. G., Bollerslev, T., Diebold, F. and Ebens, H. (2001), 'The distribution of realized stock return volatility', *Journal of Financial Economics* **61**(1), 43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. and Labys, P. (2001), 'The distribution of realized exchange rate volatility', *Journal of the American Statistical Association* **96**, 42–55.

Andersen, T. G., Bollerslev, T., Diebold, F. and Labys, P. (2003), 'Modeling and forecasting realized volatility', *Econometrica* **71**(2), 579–625.

Asgharian, H., Hou, A. J. and Javed, F. (2013), 'The importance of the macroeconomic variables in forecasting stock return variance: A GARCH-MIDAS approach', *Journal of Forecasting* **32**(7), 600–612.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A. and Shephard, N. (2008), 'Designing realized kernels to measure the ex post variation of equity prices in the presence of noise', *Econometrica* **76**(6), 1481–1536.

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

Bollerslev, T. (1986), 'Generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics* **31(3)**, 307–327.

Bollerslev, T. and Wooldridge, J. M. (1992), 'Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances', *Econometric Reviews* **11**(2), 143–172.

Brownlees, C. T. and Gallo, G. M. (2009), 'Comparison of volatility measures: a risk management perspective', *Journal of Financial Econometrics* **8**(1), 29–56.

Christiansen, C., Schmeling, M. and Schrimpf, A. (2012), 'A comprehensive look at financial volatility prediction by economic variables', *Journal of Applied Econometrics* **27**(6), 956–977.

Christoffersen, P. F. (2003), *Elements of financial risk management*, CA: Elsevier Science.

Christoffersen, P., Feunou, B., Jacobs, K. and Meddahi, N. (2014), 'The economic value of realized volatility: using high-frequency returns for option valuation', *Journal of Financial and Quantitative Analysis* **49**(3), 663–697.

Clements, M. P., Galvão, A. B. and Kim, J. H. (2008), 'Quantile forecasts of daily exchange rate returns from forecasts of realized volatility', *Journal of Empirical Finance* **15**(4), 729–750.

Conrad, C. and Loch, K. (2014), 'Anticipating long-term stock market volatility', *Journal of Applied Econometrics* **Forthcoming**.

Corradi, V., Distaso, W. and Mele, A. (2013), 'Macroeconomic determinants of stock market volatility and volatility premiums', *Journal of Monetary Economics* **60**, 203–220.

Corsi, F. (2009), 'A simple approximate long-memory model of realized volatility', *Journal of Financial Econometrics* **7**(2), 174–196.

Engle, R. (2002), 'New frontiers for arch models', *Journal of Applied Econometrics* **17**(5), 425–446.

Engle, R. F. (1982), 'Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation', *Econometrica* **159**, 987–1007.

Engle, R. F., Ghysels, E. and Sohn, B. (2013), 'Stock market volatility and macroeconomic fundamentals', *Review of Economics and Statistics* **95**(3), 776–797.

Engle, R. F. and Lee, G. G. J. (1999), A permanent and transitory component model of stock return volatility, *in* R. F. Engle and H. White, eds, 'Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W. J. Granger', Oxford University Press, Oxford, pp. 475–497.

Engle, R. F. and Rangel, J. G. (2008), 'The Spline-GARCH model for low-frequency volatility and its global macroeconomic causes', *Review of Financial Studies* **21**(3), 1187–1222.

Fleming, J., Kirby, C. and Ostdiek, B. (2003), 'The economic value of volatility timing using realized volatility', *Journal of Financial Economics* **67**(3), 473–509.

Francq, C., Horvath, L. and Zakoïan, J.-M. (2011), 'Merits and drawbacks of variance targeting in GARCH models', *Journal of Financial Econometrics* **9**(4), 619–656.

Ghysels, E., Santa-Clara, P. and Valkanov, R. (2005), 'There is a risk-return trade-off after all', *Journal of Financial Economics* **76**(3), 509–548.

Ghysels, E., Sinko, A. and Valkanov, R. (2006), 'MIDAS regressions: Further results and new directions', *Econometric Reviews* **26**, 53–90.

Giot, P. and Laurent, S. (2004), 'Modelling daily value-at-risk using realized volatility and ARCH type models', *Journal of Empirical Finance* **11**(3), 379–398.

Hamilton, J. D. and Lin, G. (1996), 'Stock market volatility and the business cycle', *Journal of Applied Econometrics* **11**(5), 573–593.

Hansen, P. R. and Huang, Z. (2015), 'Exponential GARCH modeling with realized measures of volatility', *Journal of Business & Economic Statistics* **Forthcoming**.

Hansen, P. R., Huang, Z. and Shek, H. H. (2012), 'Realized GARCH: a joint model for returns and realized measures of volatility', *Journal of Applied Econometrics* **27**(6), 877–906.

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

Hansen, P. R., Lunde, A. and Nason, J. M. (2011), 'The model confidence set', *Econometrica* **79**(2), 453–497.

Hansen, P. R., Lunde, A. and Voev, V. (2014), 'Realized beta GARCH: a multivariate GARCHmodel with realized measures of volatility', *Journal of Applied Econometrics* **29**(5), 774–799.

Komunjer, I. (2005), 'Quasi-maximum likelihood estimation for conditional quantiles', *Journal of Econometrics* **128**(1), 137–164.

Maheu, J. M. and McCurdy, T. H. (2011), 'Do high-frequency measures of volatility improve forecasts of return distributions?', *Journal of Econometrics* **160**(1), 69–76.

McAleer, M. and Medeiros, M. C. (2008), 'Realized volatility: A review', *Econometric Reviews* **27**(1-3), 10–45.

Mincer, J. A. and Zarnowitz, V. (1969), The evaluation of economic forecasts, *in* 'Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance', NBER, pp. 1–46.

Officer, R. R. (1973), 'The variability of the market factor of the new york stock exchange', *Journal of Business* **46**(3), 434–453.

Patton, A. J. (2011), 'Volatility forecast comparison using imperfect volatility proxies', *Journal of Econometrics* **160**(1), 246–256.

Patton, A. J. and Sheppard, K. (2009), Evaluating volatility and correlation forecasts, *in* 'Handbook of financial time series', Springer, pp. 801–838.

Paye, B. S. (2012), 'Déjà vol: Predictive regressions for aggregate stock market volatility using macroeconomic variables', *Journal of Financial Economics* **106**(3), 527–546.

Perez-Quiros, G. and Timmermann, A. (2000), 'Firm size and cyclical variations in stock returns', *The Journal of Finance* **55**(3), 1229–1262.

Schwert, G. W. (1989), 'Why does stock market volatility change over time?', *The Journal of Finance* **44**(5), 1115–1153.

Shephard, N. and Sheppard, K. (2010), 'Realising the future: forecasting with high-frequency-based volatility (HEAVY) models', *Journal of Applied Econometrics* **25**(2), 197–231.

Vander Elst, H. (2015), 'Flogarch: Realizing long memory and asymmetries in returns volatility', *ECARES WP 2015/12* .

Zhang, L., Mykland, P. and Aït-Sahalia, Y. (2005), 'A tale of two time scales: Determining integrated volatility with noisy high frequency data', *Journal of The American Statistical Association* **100(472)**, 1394–1411.