17 96

LOMBARD ODIER
INVESTMENT MANAGERS

# MARS: the benefits of using range-based measures to forecast volatility

– For professional investor use only –

05/23

May 2023

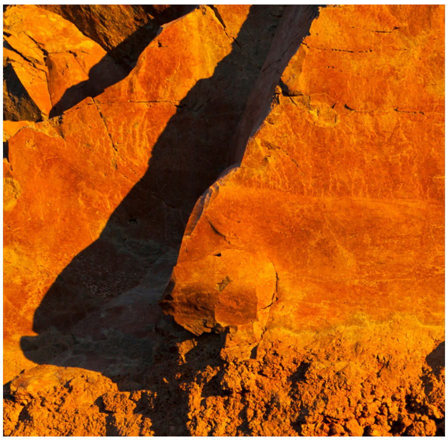

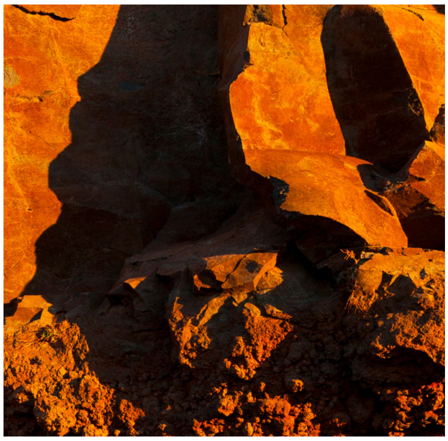

Range-based measures can improve the accuracy of volatility forecasts.

p.08

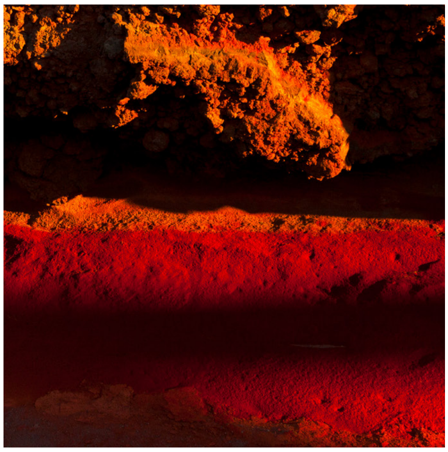



**For professional investor use only. Please read important information at the end of this document.**


This document has been prepared by:

**François Chareyron,** ARPM
Portfolio Manager

**Julien Royer,** PhD
Quantitative Analyst

**Gabriele Grignani**
Quantitative Analyst Intern


For further information please visit
[Professional | Lombard Odier Investment Managers]

**MARS**, our **M**ulti-**A**sset **R**esearch **S**eries that brings quantitative research back to Earth.

# Introduction

Range-based volatility measures present numerous advantages when estimating the volatility process of financial time series. As discussed in the previous paper of our Multi-Asset Research Series,[1] these estimators capture intraday patterns without the burden of handling the full intraday path. A natural question, however, arises as a consequence: are range-based measures useful to generate more accurate volatility forecasts?

Risk-based portfolio construction is directly linked to the ability to build forecasts of volatility and translate them into investment rules. In standard settings, forecasts' quality is evaluated by comparing forecasts with the ex-post realisation of the studied variable through a loss function: for example the mean square error. Studying the loss functions for different competing models is also the basis for model selection. Evaluating volatility forecasts is, however, more complex. As the true volatility process is a latent, unobservable process, one cannot directly compute the loss inherent in the forecasting exercise.

In this white paper, we present how to evaluate volatility forecasts by using a (noisy) proxy and discuss the advantages in using range-based volatility to forecast volatility in lieu of the classical close-to-close estimator.

The first section introduces how to incorporate the range-based measures into standard volatility models to generate forecasts. The second section discusses the choice of the scoring function to compute the loss series of the forecasts when the true process is unobservable. This notably enables the comparison of different volatility models, highlighting the additional information carried by intraday data when forecasting daily volatility. The third section presents how range-based exogenous variables can help build more accurate forecasts for longer horizons. The final section concludes and surveys potential future research perspectives.

1 Chareyron, F., and Royer, J. (2023). A primer on range-based volatility estimators. *Lombard Odier Investment Managers - Multi Asset Research Series.*

# From a measure to a forecast

Similarly to the realised volatility measure (see for example Andersen et al. (2021)), range-based measures are model-free estimators that consistently estimate the path of the latent volatility process without the requirement to specify its dynamic. While this approach has the advantage of curbing misspecification risk, it limits the ability to generate forecasts as there is no explicit link between the estimator and its previous values.

To remediate this issue, different approaches have been proposed: from directly leveraging the persistence of the proposed noisy estimator (as in Corsi (2009)), to combining the model-free estimator with prominent conditional volatility models. In this white paper, we will focus on the latter, as such models remain the workhorse for financial econometrics applications. Nevertheless, the forecast evaluation procedure detailed in the following sections remains valid independently of the model considered.

Conditional volatility models directly describe the dynamic of the volatility process, and it is known that they are difficult to outperform in a volatility forecasting exercise. Amongst the myriad of different specifications, two have particularly caught the attention of practitioners as they allow capturing stylised facts of financial returns while remaining very parsimonious: the GARCH equation introduced by Bollerslev (1986) and the Exponential GARCH (EGARCH) of Nelson (1991). In the former, the variance dynamic is given by:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{1}$$

where $\varepsilon_t = \sigma_t\eta_t$ denotes the daily returns and $\eta_t$ is a white noise. In this setting, a 1-day ahead forecast for the volatility can be readily made once the parameters of the model have been estimated:

$$\sigma_{t+1|\mathrm{t}}^2 = \mathrm{E}[\sigma_{t+1}^2|\mathcal{F}_t] = \omega + (\alpha+\beta)\sigma_t^2.$$

However, this model ignores potential additional information that is not captured by close-to-close returns. To circumvent this issue, practitioners augmented the conditional variance equation (1) by including exogenous variables, yielding the GARCH-X equation (see e.g. Han and Kristensen (2014) and Francq and Thieu (2019))

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 + \gamma X_{t-1}$$

where $X$ denotes a positive exogenous variable. Examples of financial metrics used as an exogenous variable include bid-ask spreads (Bollerslev and Melvin, 1994), futures open interest (Girma and Mougoué, 2002) or trading volumes (Lamoureux and Lastrapes, 1990). Setting $X$ as the range-based measure of volatility thus easily augments the volatility forecast with intraday patterns captured by the range-based measure.

Similarly, the EGARCH model of Nelson, in which the volatility dynamic is given by

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 + \gamma X_{t-1}$$

where the use of the logarithm alleviates the positivity constraint on the coefficients and can be augmented by the range-based measure as proposed by Brandt and Jones (2006). When setting $X$ as the range-based measure computing using the approach of Garman and Klass (1980), the EGARCH-X equation thus becomes

$$\ln\sigma_t^2 = \omega + \alpha(|\eta_{t-1}| - \mathrm{E}[|\eta_{t-1}|]) + \beta\ln\sigma_{t-1}^2 + \gamma(\ln X_{t-1} - \mathrm{E}[\ln X_{t-1}]).$$

Assuming that the innovations $\eta_t$ are a sequence of independent and identically distributed centered Gaussian variables with unit variance, the forecast for the variance is given by

$$\sigma_{t+1|\mathrm{t}}^2 = \mathrm{E}[\sigma_{t+1}^2|\mathcal{F}_t] = \Lambda\sigma_t^{2\beta}$$

where, denoting $\Phi$ the cumulative density function of a normally distributed variable,

$$\Lambda = \exp(\omega)\exp\left(-\alpha\sqrt{\frac{2}{\pi}}\right)\exp\left(\frac{\alpha^2}{2}\right)[2\Phi(\alpha)].$$

In the remainder of this note, we will consider eight competing models, four for the GARCH and their counterparts in the EGARCH framework. First, the standard implementation relying solely on daily squared returns as variance estimate. Second, *Model*-RB denotes the conditional volatility model augmented by the Garman-Klass range-based measure. Third, *Model*-RV denotes the conditional volatility model augmented by the realised volatility measure computed using 5-minute returns, and finally, *Model*-X denotes the conditional volatility model augmented by both the range-based and realised volatility measures. The different models are fitted on daily S&P 500 returns ranging from 2000 to 2019.

To illustrate the differences of behaviour of those models, we present in figure 1 the one-day ahead volatility forecasts obtained by fitting those models on S&P 500 returns.

**FIG. 1 ONE-DAY AHEAD S&P 500 VOLATILITY FORECASTS FOR COMPETING (E)GARCH(-X) MODELS.**

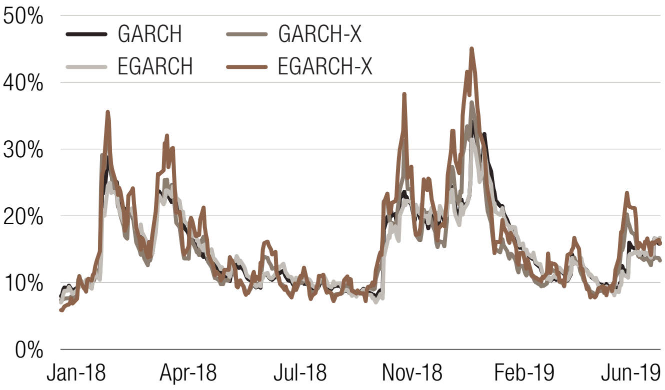


Source: LOIM. Data as of June 2019.

Table 1 presents the value of the parameters in the eight competing models. Interestingly, the coefficient of the exogenous variable is always different from zero, highlighting the relevance of including either the range-based measure or the realised volatility as exogenous variables. The last line of the table presents the result of the test for the hypothesis:

- $H_0: \gamma_{RB} = 0$ for *Model-RB*;
- $H_0: \gamma_{RV} = 0$ for *Model-RV*;
- $H_0: \gamma_{RB} = 0$ for *Model-X*.

Note that in the GARCH case, the positivity of the parameters implies a boundary restriction that renders the asymptotic distribution of the test statistic non-standard and we follow Francq and Zakoïan (2009) to obtain the correct threshold. Notably, when including both range-based and realised volatility measures, the range-based covariate becomes insignificant as it is redundant with the other intraday based measure. Nevertheless, as noted in the previous edition of the Multi-Asset Research Series, the difficulty of obtaining trustworthy realised volatility measures provides an argument in favour of the range-based version.

**TAB. 1 PARAMETERS FOR THE DIFFERENT MODELS ESTIMATED BY QUASI-MAXIMUM LIKELIHOOD AND LIKELIHOOD-RATIO TEST FOR THE SIGNIFICANCE OF THE EXOGENOUS VARIABLE**

| | | GARCH | GARCH-RB | GARCH-RV | GARCH-X | EGARCH | EGARCH-RB | EGARCH-RV | EGARCH-X |
|---|---|---|---|---|---|---|---|---|---|
| Parameters | $\omega$ | 0.022 | 0.038 | 0.061 | 0.061 | 0.009 | 0.203 | 0.309 | 0.320 |
| | $\alpha$ | 0.114 | 0.022 | 0.000 | 0.000 | 0.225 | -0.006 | -0.131 | -0.139 |
| | $\beta$ | 0.869 | 0.707 | 0.433 | 0.433 | 0.975 | 0.690 | 0.383 | 0.370 |
| | $\gamma_{RB}$ | – | 0.271 | – | 0.000 | – | 0.297 | – | 0.036 |
| | $\gamma_{RV}$ | – | – | 0.567 | 0.567 | – | – | 0.588 | 0.567 |
| Test | | | ✓ | ✓ | ✗ | | ✓ | ✓ | ✗ |

Source: LOIM. A check mark indicates that the coefficient in the test hypothesis is significantly different from 0 at the 95%-confidence level according to the classical likelihood ratio statistics.

# Evaluating volatility forecasts

While Table 1 provides arguments in favour of the inclusion of range-based measures as exogenous variables in standard conditional volatility models, the test statistic only gives an in-sample vision of the benefits. To strengthen confidence in the robustness of these results, one must assess the gains in an out-of-sample exercise. Evaluating point forecasts usually means defining a scoring function between the forecast and the ex-post realisation of the variable of interest (see e.g., Gneiting (2012)). This procedure cannot be readily implemented for volatility forecasts as the variable of interest is unobservable.

A common solution is to compute the loss between the forecast and an imperfect proxy of the variable of interest. This, however, complicates the comparison of the losses between two forecasts as they are not measured as a distance to the true latent process. Additionally, volatility comparison can be very sensitive to extreme observations, which renders the use of classical loss functions such as the mean square error (MSE) impractical.

To circumvent these issues, Patton (2011) studied different usual loss functions to determine their robustness. More precisely, a loss function $\mathcal{L}$ is said to be robust if the ranking of any two variance forecasts $h_{1t}$ and $h_{2t}$ is invariant whether the loss is computed with regard to the true conditional variance $h_t$ or some conditionally unbiased variance proxy $\hat{h}_t$:

$$\mathrm{E}[\mathcal{L}(h_t, h_{1,t})] \gtreqless \mathrm{E}[\mathcal{L}(h_t, h_{2,t})] \Leftrightarrow \mathrm{E}[\mathcal{L}(\hat{h}_t, h_{1,t})] \gtreqless \mathrm{E}[\mathcal{L}(\hat{h}_t, h_{2,t})]$$

for any proxy $\hat{h}_t$ such that $\mathrm{E}[\hat{h}_t | \mathcal{F}_{t-1}] = h_t$.

In particular, the author shows the entire subset of robust and homogenous loss functions $\mathcal{L}(\hat{\sigma}^2, h; b)$ is given by the following family:

- $h - \hat{\sigma}^2 + \hat{\sigma}^2 \ln\frac{\hat{\sigma}^2}{h}$ for $b = -1$
- $\frac{\hat{\sigma}^2}{h} - \ln\frac{\hat{\sigma}^2}{h} - 1$ for $b = -2$
- $(b+1)^{-1}(b+2)^{-1}(\hat{\sigma}^{2b+4} - h^{b+2})$ $-(b+1)^{-1}h^{b+1}(\hat{\sigma}^2 - h)$ for $b \notin \{-1, -2\}$.

Interestingly, up to additive and multiplicative constants that do not change the ranking of two losses, the MSE loss function is obtained for $b = 0$ while the QLIKE loss function is obtained for $b = -2$. Such a function is very popular when evaluating volatility forecasts as it is derived from the Gaussian likelihood of the conditional volatility equation:

$$\mathcal{L}(\hat{\sigma}^2, h)_{\mathrm{QLIKE}} = \ln(h) + \frac{\hat{\sigma}^2}{h}.$$

Following Diebold and Mariano (1995) and West (1996), we propose testing for the significance of better forecast ability between our eight selected models. To do so, we fix the training set as data from 2000 to 2017 and our test set as data from 2018 to 2019. We re-estimate our models every 22 points, meaning that 22 consecutive one-day ahead forecasts are produced before re-evaluating the parameters. We select the QLIKE loss function from the set of robust functions derived by Patton (2011) as it is less sensitive to outliers than MSE. However, one should keep in mind that the QLIKE loss function is asymmetrical and as such can tend to favour positively biased forecasts.

Table 2 presents the results of the comparison between the competing models presented in the previous section. A plus sign means that the model in the column is significantly better at forecasting volatility than the model in the row, an equal sign means that one cannot discriminate between the two, and a minus sign means that the model on the row is significantly better than the model in the column.

For example, the EGARCH model augmented with the range-based volatility outperforms both the standard GARCH and EGARCH models, is not significantly different from the GARCH-RB model and is outperformed by the models including both range-based and realised volatility measures as covariates. Overall, it is striking to see how standard models are dominated by models augmented by either the range-based or realised volatility measure. It is, however, difficult to discriminate between the GARCH and the EGARCH frameworks.

**TAB. 2 FORECASTING ABILITY OF THE COMPETING MODELS BASED ON THE DIEBOLD-MARIANO TEST**

| | GARCH-RB | GARCH-RV | GARCH-X | EGARCH | EGARCH-RB | EGARCH-RV | EGARCH-X |
|---|---|---|---|---|---|---|---|
| GARCH | + | + | + | − | + | + | + |
| GARCH-RB | | + | + | − | = | + | + |
| GARCH-RV | | | + | − | − | = | + |
| GARCH-X | | | | − | − | − | = |
| EGARCH | | | | | + | + | + |
| EGARCH-RB | | | | | | + | + |
| EGARCH-RV | | | | | | | + |

Source: LOIM.

# Forecasting over longer horizons

Although forecasting one-day ahead volatility is essential to some investment strategies, most portfolio constructions require forecasting over a longer horizon, for example at the monthly frequency. Therefore, one must usually generate multiple $k$-day ahead forecasts and aggregate them (see e.g., De Nard et al. (2022)). In that exercise, standard GARCH models suffer from a well-known limitation: the estimated parameters are often found to be such that $\alpha + \beta \cong 1$, the so-called *near-integration* case. As those parameters define the mean-reversion speed of the estimated volatility process, GARCH models exhibit a slow decaying pattern, as noted by Mikosch and Stărică (2004), that may be inadequate to capture the true term structure of volatility.

On that matter, figure 2 shows strong arguments in favour of using intraday-based measures as exogenous variables when forecasting over longer horizons. In the chart, the red line presents the average term structure of the realised volatility from the instantaneous to the 25-day frequency in periods of high (left-hand figure) and low (right-hand figure) volatility.[2] Using the parameters obtained over the whole sample for the different GARCH models presented in Table 1, we compute the forecasted term structure of the competing GARCH models. It can be seen that in both periods, the GARCH (blue line) behaves poorly as the near-integration implies a decay toward the long-term variance that is too slow compared to the empirical decay of the realised variance.

Interestingly, the inclusion of the exogenous variable clearly mitigates this issue as the term structure of the GARCH-RB and GARCH-RV models appear to be more consistent and closer to the empirical behaviour. This improvement could be linked to the ability of the exogenous variable to mitigate the effect of imposing constant parameters in the conditional volatility equation, leading to overly persistent models as suggested by Hillebrand (2005).

**FIG. 2 EMPIRICAL AND FORECASTED TERM STRUCTURE OF ANNUALISED VOLATILITY**

TOP QUINTILE

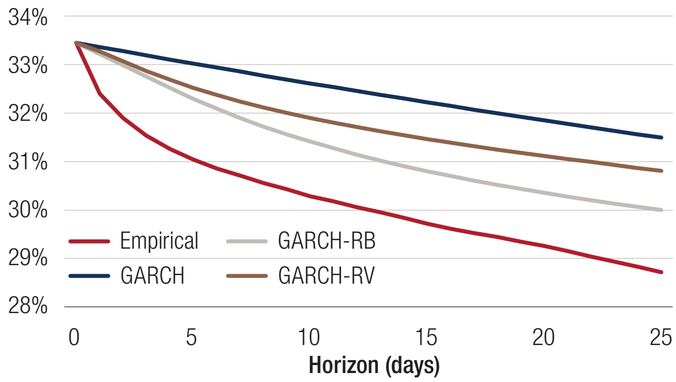


BOTTOM QUINTILE

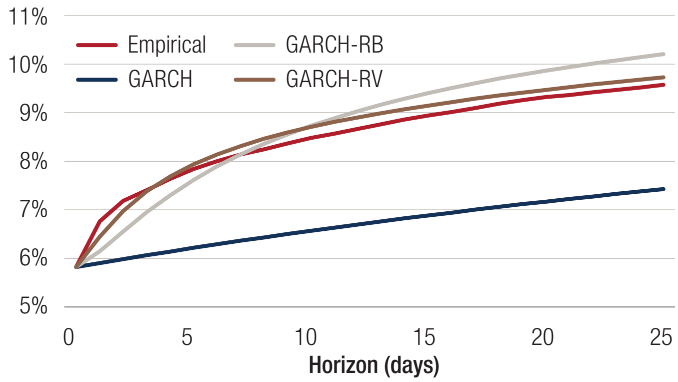


Source: LOIM.

[2] Periods of high and low volatility are obtained by sampling realised volatility paths into 5 groups based on the value of the realised volatility at the first day of the considered path. The top (respectively bottom) quintile thus represents paths with the highest (lowest) volatility.

# Conclusion

In addition to featuring strong statistical properties, as discussed in the previous edition of our Multi-Asset Research Series, range-based volatility measures provide a significant edge when forecasting volatility. Assessing that edge might, however, be more difficult than anticipated as in-sample, standard testing procedures must be modified to tackle statistical issues, while out-of-sample, the computation of forecasts' losses is rendered complex by the latent nature of the volatility process.

In this white paper, we review state of the art techniques to handle such difficulties and illustrate the benefits of volatility measures based on intraday data when forecasting daily or monthly volatility for the S&P 500 over the last 20 years. Validation across instruments and asset classes are excluded for the sake of brevity.

Further research will address the economic value of enhanced volatility forecasts in the context of risk-based portfolios. Indeed, there is no guarantee that a statistically better model necessarily yields economically better outcomes when used for portfolio construction, a topic we will investigate in the next issue of our Multi-Asset Research Series.

**References**

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics, 61*(1), 43-76.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 31*(3), 307-327.

Bollerslev, T., & Melvin, M. (1994). Bid-ask spreads and volatility in the foreign exchange market: An empirical analysis. *Journal of International Economics, 36*(3-4), 355-372.

Brandt, M. W., & Jones, C. S. (2006). Volatility forecasting with range-based EGARCH models. *Journal of Business & Economic Statistics, 24*(4), 470-486.

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

De Nard, G., Engle, R. F., Ledoit, O., & Wolf, M. (2022). Large dynamic covariance matrices: Enhancements based on intraday data. *Journal of Banking & Finance, 138*, 106426.

Diebold, F. X., & Mariano, R. S. (1995). Comparing Predictive Accuracy. *Journal of Business & Economic Statistics, 13*(3), 253-263.

Francq, C., & Thieu, L. (2019). QML inference for volatility models with covariates. *Econometric Theory, 35*(1), 37-72.

Francq, C., & Zakoïan, J.-M. (2009). Testing the nullity of GARCH coefficients: correction of the standard tests and relative efficiency comparisons. *Journal of the American Statistical Association, 104*(485), 131-324.

Garman, M. B., & Klass, M. J. (1980). On the estimation of security price volatilities from historical data. *The Journal of Business, 53*(1), 67-78.

Girma, P., & Mougoué, M. (2002). An empirical examination of the relation between futures spreads volatility, volume, and open interest. *Journal of Futures Markets, 22*(11), 1083-1102.

Gneiting, T. (2012). Making and Evaluating Point Forecasts. *Journal of the American Statistical Association, 106*(494), 746-762.

Han, H., & Kristensen, D. (2014). Asymptotic Theory for the QMLE in GARCH-X Models With Stationary and Nonstationary Covariates. *Journal of Business & Economic Statistics, 32*(3), 416-429.

Hillebrand, E. (2005). Neglecting parameter changes in GARCH models. *Journal of Econometrics, 129*(1), 121-138.

Lamoureux, C. G., & Lastrapes, W. D. (1990). Heteroskedasticity in Stock Return Data: Volume versus GARCH Effects. *The Journal of Finance, 45*(1), 221-229.

Mikosch, T., & Stărică, C. (2004). Nonstationarities in Financial Time Series, the Long-Range Dependence, and the IGARCH Effects. *The Review of Economics and Statistics, 86*(1), 378-390.

Nelson, D. B. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica, 59*(2), 347-370.

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

West, K. D. (1996). Asymptotic Inference about Predictive Ability. *Econometrica, 644*(5), 1067-1084.

## Important information

**For professional investor use only.**