# Realized Multivariate GARCH with Factors

Murilo Getlinger Coelho

2022

**Abstract**

Forecasting second moments of asset returns is essential in portfolio selection. In a multivariate setting, the dimensionality of the problem and the precision of predictions are the main concerns. We propose a new methodology for forecasting covariance matrices joining two extant approaches in the literature: intraday data to enhance predictive ability and factors to reduce the dimensionality. We assume a multivariate realized GARCH model for the factors and a set of multivariate realized GARCH between each stock and the factors. We compare our methodology empirically with the standard literature by optimizing a portfolio on the S&P500 stocks universe.
**Keywords:** Financial Volatility, Realized GARCH, High Frequency Data, Multivariate Modeling, Correlation Matrix, Factors
**JEL Classification:** G11, G17, C32, C53, C58

## 1 Introduction

Predicting asset returns volatility is crucial in asset allocation. For many years researchers have been dedicating themselves to create models that enable better predictive ability. In a multivariate setting, the dimensionality of the problem and the precision of forecasts are the main concerns since the univariate GARCH structure cannot be generalized in a straightforward manner.

In this paper we propose a new methodology for forecasting covariance matrices joining two extant approaches in the literature: intraday data as timely information about the underlying volatility to enhance predictive ability and factors to reduce the dimensionality and make estimation feasible. Our main contribution is to construct a model using factors models with realized multivariate GARCH, and to compare different estimation approaches of the covariance matrix of asset returns to find the best performance in terms of portfolio allocation.

Usually, the literature solves these problems assuming partially constant conditional covariance matrices, minimizing the number of parameters and guaranteeing semi-positivity. The first multivariate GARCH implementations had important simplifications. Bollerslev (1990)'s CCC-GARCH assumes that the conditional correlation matrix is time-invariant. This approach reduces the number of parameters significantly, only one matrix has to be estimated and guarantees semi-posivitity by construction. The ECCC-GARCH proposed by Jeantheau (1998) expands the previous model by allowing the autocorrelation structure to change over time. The number of parameters is a bit larger since there is a need to estimate the autocorrelation parameters for each asset. Positivity, again, comes from the construction of the conditional correlation matrix.

Engle (2002) introduced the DCC-GARCH using as assumption that the conditional correlation matrix has a dynamic process over time. The parameters are for the initial level of the covariance matrix and for the dynamic relation of each entry. The semi-positivity is guaranteed because to estimate the conditional correlation matrix the inputs are the previous matrix and the standard squared residual estimator. Some authors began to use shrinkage methods to the sample covariance matrix in order to correct in-sample bias, resulting, most importantly, on the DCC-NL model in Engle et al. (2019). The DCC-NL is the state-of-the-art on multivariate GARCH modeling without the use of realized measures.

A complementary branch of the literature moved towards enhancing the predictive ability of the models. The most used advancement on this front is based on leveraging high frequency data; the literature started to rely on realized measures, using the observed volatility of returns as a timely information about the underlying volatility. The Realized GARCH combines realized measures with GARCH models to predict conditional volatilities and covariances in a simple and efficient manner. See Hansen et al. (2012) for univariate modeling and Archakov et al. (2020) for a multivariate version. The Multivariate Realized GARCH model works well with a small number of assets. When trying to estimate it with high dimension data, the curse of dimensionality comes in full force, the number of parameters grows restrictively fast.

In this paper we will solve the dimensionality problem with factors as in Engle et al. (1990). The idea is to project the volatility of a multitude of assets in a few dimensions, making the estimation feasible even on high dimensional multivariate GARCH models. To enhance the predictive ability we will also use realized measures as in Archakov et al. (2020). In particular, since we need high-frequency data to estimate the realized measures of the factors, we will use the five Fama-French factors plus momentum that Aït-Sahalia et al. (2020) compute.[1] Factor structures are a natural solution in the finance literature. This approach has a long pedigree (Barigozzi and Hallin (2016), Barigozzi and Hallin (2017) and Herskovic et al. (2016)).

To empirically assess our approach, we compute the Global Minimum Variance (GMV) portfolio using US stocks from 2002 to 2017. We focus on the GMV portfolio because it does not require the estimation of the expected stock returns. The latter are very hard to estimate and hence could undermine our empirical evaluation (Black (1993) and Fabozzi et al. (2002)). We use five-minute returns on S&P500 stock constituents from the PiTrading database, apart from the high-frequency factors of Aït-Sahalia et al. (2020). The five-minute sampling interval provides a sufficient number of intraday returns to ensure precise weekly realized measures, though low enough to alleviate market microstructure noise. Since we will not rebalance the portfolio every day, we will implement our empirical application using weekly returns (instead of daily returns) and forecast the weekly conditional correlation matrix.

The results are mixed. Our proposed approaches performed well compared to the naive $1/N$ allocation, but failed to produce portfolios with smaller variance than the extant models in the literature. The Model Confidence Set procedure (Hansen et al. (2011)) indicates that the DCC-NL outperforms the others when looking at the realized volatility of the GMV portfolios. Computational challenges on the estimation of the more complex models might partially explain the difference. The estimated turnover of the portfolios suggest, as one would expect, that there exist a sweet spot for the trade off between reducing the volatility and incurring in transaction costs. The approach with the lowest mean turnover ratio is the Realized Multivariate GARCH with Factors restricting the idiosyncratic matrix. The one with the highest Sharpe ratio also is the DCC-NL, but Ledoit and Wolf's (2008) test does not reject the null of equal ratios.

The next section provides the theoretical framework of the model and the estimation procedure, section 3 describes the data used in the paper. Section 4 details the empirical application, section 5 has the empirical results and finally in section 6 we have the conclusion.

# 2 Model

## 2.1 Notation

The subscript $i$ indexes assets that vary from 1 to $N$, where $N$ is the dimension of investable assets. The subscript $j$ indexes factors, varying from 1 to $J$, where $J$ is the number of utilized factors. The subscript $t$ indexes weeks from 1 to $T$.

We are also using the following notation: $x_{i,t}$ denotes the observed asset returns for asset $i$ at date $t$; $f_{j,t}$ the observed factor returns for factor $j$ at date $t$ and can be stacked into $F_t \equiv (f_{1,t}, \dots, f_{J,t})'$; $\beta_{t+1|t}^{(ij)}$ the exposure of asset $i$ to factor $j$ at date $t+1$ conditional on date $t$. $Z_t^{FF} \equiv Cov_t(F_t)$ is the realized covariance matrix of the factors at date $t$ and $H_{t+1|t}^{FF} \equiv Cov_t(F_{t+1})$ is the covariance matrix of the factors at date $t+1$ conditional on date $t$. Similarly, $Z_t^{iF} \equiv Cov_t(x_{i,t}, F_t)$ is the realized covariance matrix of factors and asset $i$ at date $t$ and $H_{t+1|t}^{iF} \equiv Cov_t(x_{i,t+1}, F_{t+1})$ is the covariance matrix of factors and asset $i$ at date $t+1$ conditional on date $t$.

[1]Available to download on https://dachxiu.chicagobooth.edu/download/Factors_96_17_V2.0.zip

## 2.2 General Model

The objective is to estimate the conditional covariance matrix of asset returns. Our model is based on Archakov et al. (2020) with some important modifications. They use a block relation to estimate the covariance matrix of returns, and their approach works well for a small number of assets. Our model utilizes factors to reduce dimensionality and estimate in a feasible manner the covariance of asset returns.

Let the return for each asset be:

$$x_{it+1} = \mu_i + \sum_{j=1}^{q} \beta_{t+1|t}^{(ij)} f_{jt+1} + \varepsilon_{it+1} \tag{1}$$

From the factor structure of $x_{it}$:

$$Var_t(x_{it+1}) = B'_{t+1|t} H_{t+1|t}^{FF} B_{t+1|t} + Var_t(\varepsilon_{it+1}) \tag{2}$$

where $B_{t+1|t}$ is the matrix stacking all factor betas:

$$B_{t+1|t} = \begin{bmatrix} \beta_{t+1|t}^{(i1)} & \beta_{t+1|t}^{(i2)} & \dots & \beta_{t+1|t}^{(iJ)} \end{bmatrix}'$$

and

$$\beta_{t+1|t}^{(ij)} = \frac{Cov_t(x_{it+1}, f_{jt+1})}{Var_t(f_{jt+1})} \tag{3}$$

Let's start with the common factors $f_{jt}$ that we will model as a realized GARCH:

$$\begin{aligned} H_{t+1|t}^{FF} &= \omega + \delta H_{t|t-1}^{FF} + \gamma Z_t^{FF} \\ Z_t^{FF} &= \zeta + \Phi H_{t|t-1}^{FF} + v_t^{FF} \end{aligned}$$

where $Z_t^{FF}$ is the realized covariance matrix of the factors at date $t$.

To close the model we need a structure to the covariance matrix between $x_{it}$ and $F_t$. This will also be done through a Realized GARCH:

$$\begin{aligned} H_{t+1|t}^{iF} &= \omega_i + \delta_i H_{t|t-1}^{iF} + \gamma_i Z_t^{iF} \\ Z_t^{iF} &= \zeta_i + \Phi_i H_{t|t-1}^{iF} + v_t^{iF} \end{aligned}$$

Next, we will present implementations to estimate the conditional covariance matrix imposing different restrictions. In the first approach we estimate a Multivariate Realized GARCH for the covariances between asset returns and factor returns to find the beta implied in the GARCH structure. Here the betas are derived from the conditional variances and covariances, the realized measures are used to forecast the conditional covariance matrix through the Realized Multivariate GARCH. The second approach imposes more structure by estimating a diagonal covariance matrix for the idiosyncratic innovations. This makes the GARCH of residuals easier to estimate but can imply a loss in predictive ability.

As for the third approach, we now have time-varying betas (through realized measures) and estimate the residuals using (1). Here we do not impose any restrictions on the covariance matrix of the residuals. In the final approach we use static betas and an exact factor structure, losing precision if the populational parameters vary much over time. We estimate a univariate GARCH model for each residual in the static-beta regressions. If the populational betas have enough variance over time, this approach should perform worse than the previous, but if they do not vary much over time the performance can be better as we will add noise estimating with intraday data.

## 2.3 GARCH-implied beta with unconstrained residual estimation

We will estimate each beta using the forecast of two different Realized Multivariate GARCH models. The first one is for the factors from where we will estimate $Var_t(f_{jt+1})$ and the second one is for the relation between asset $i$ and the factors and we will use to estimate $Cov_t(x_{it+1}, f_{jt+1})$. From these two models we use equation (3) to get one set of betas for each time period as implied in (1) and (3).

To finish the covariance matrix forecast of the idiosyncratic part we get the residuals from these estimated betas and run a multivariate GARCH model for asset returns using non linear and linear shrinkage methods. Then we use (2) to get the final covariance matrix forecast.

## 2.4 GARCH-implied beta with constrained residual estimation

Here we impose more structure on the estimation. To estimate the betas we use the same method as the previous approach, the difference being on the estimation of the covariance matrix of the idiosyncratic part. Now we impose a diagonal covariance matrix, i.e., the correlation of asset returns must come only through the factors, the idiosyncratic part is assumed to be truly idiosyncratic.

The rest of the estimation is exactly the same: we use the betas, the forecast of the covariance matrix of factor returns and the covariance matrix (diagonal now) of the idiosyncratic part to estimate the covariance matrix of asset returns through equation (2).

Since we are imposing more structure we might lose accuracy on the estimation, but if the factors can fully explain the covariance, estimating a full matrix for the idiosyncratic part (as we do in the first approach, even with shrinkage) might imply in a larger variance of the estimator.

## 2.5 Realized Betas

In this approach we estimate weekly betas using intraday data for each asset and each factor:

$$x_{it+1} = \mu_i + \sum_{j=1}^{q} \beta_{t+1|t}^{(ij)} f_{jt+1} + \varepsilon_{it+1}$$

When we estimate the realized beta we are only seeing one realization of the conditional beta, and not the conditional beta itself. Assuming that the conditional beta has a mean reversion to the unconditional beta over time, a not too strong assumption, a better estimator is a convex combination between the realized beta and the OLS beta. So, the final beta will be 2/3 static beta (estimated without look-ahead bias) and 1/3 realized beta. With this we are trying to use recent information to improve the estimation of each beta.

Then we extract the residuals generated from this approach and estimate a univariate GARCH to forecast the idiosyncratic volatility for each asset. Then again we use a Realized Multivariate GARCH to estimate the covariance matrix of the factors and use (2) with the combined betas to find the covariance matrix of the assets. Compared to the previous approaches, we expect a worse performance since we are not using the Realized Multivariate GARCH for the assets and thus leaving behind potentially relevant information.

## 2.6 Static Betas

In this final approach we will estimate constant betas for each stock:

$$x_{it+1} = \mu_i + \sum_{j=1}^{q} \beta^{(ij)} f_{jt+1} + \varepsilon_{it+1}$$

The estimation is simple: we run a regression for each asset on the factors and extract the residuals. From these residuals we estimate a univariate GARCH and forecast the idiosyncratic volatility. Then, to forecast the volatility from the factors we use a Realized Multivariate GARCH only for the factors and the final covariance matrix is estimated as proposed in equation (2). Compared to the last approach we expect a worse forecasting performance based on the less accurate weekly beta estimation, assuming the betas vary sufficiently over time.

## 2.7 Global Minimum Variance Portfolio

The Markovitz optimal solution to the asset allocation problem requires modeling the expected return of assets. Modeling the expected returns is considerably hard (Black (1993) and Fabozzi et al. (2002)) and it is not useful in our setting since there is no advantage if we are looking at the covariance matrix forecasting. As such, to assess the predictive ability of our proposed approaches, we use the forecasts to find the Global Minimum Variance portfolio (GMV).

Let $\Sigma_t^m$ be the conditional covariance matrix of asset returns through approach $m$ at time $t$. To find the minimum variance portfolio we just need to find weights $w_i^m$ such that:

$$W_t^m = \text{argmin}_W \quad W'\Sigma_t^m W$$

where

$$W = [w_1 \, w_2 \, \ldots \, w_N]'$$

The feasible solution is:

$$\widehat{W}_t^m = \frac{\left(\hat{\Sigma}_t^m\right)^{-1} \mathbf{1}}{\mathbf{1}' \left(\hat{\Sigma}_t^m\right)^{-1} \mathbf{1}}$$

where $\hat{\Sigma}_t^m$ is the estimate of the conditional covariance matrix of asset returns.

## 2.8 Portfolio Turnover

In a portfolio selection setting, one of the main issues is the transaction costs surrounding the rebalancing of the assets. The usual measure is the portfolio turnover, defined as:

$$\text{Turnover Ratio}_\tau = \sum_{i=1}^{N} \left| w_{i,\tau}^{\text{BOP}} - w_{i,\tau-1}^{\text{EOP}} \right|$$

where $\tau$ is any rebalancing date and $w_{i,\tau}^{\text{BOP}}$ is the weight of asset $i$ in the portfolio in the beginning of date $\tau$ and $w_{i,\tau}^{\text{BOP}}$ is the weight in the end of period $\tau$. This measures what percentage of the portfolio had to be reallocated in the rebalancing window.

It is important to look at the portfolio turnover to see if the gains in forecasting the covariance matrix come with a large turnover cost. One relevant question is how to measure the tradeoff between turnover cost and reducing the volatility of the final portfolio.

## 2.9 Model Confidence Set

To compare the predictive ability of the different approaches through the GMV we shall use the Model Confidence Set (MCS), proposed by Hansen et al. (2011). We will use as loss function the realized variance of the proposed portfolio of each strategy. The portfolios considered will be rebalanced every five trading days.

The MCS works as follows. Let $l_{m,t}$ be the realized variance of the minimum variance portfolio of approach $m = 1 \ldots M$ at time $t$. Define:

$$d_{mk,t} = l_{m,t} - l_{k,t}, \quad 1 \leq m \neq k \leq M, \quad t = 1, \ldots, T \tag{4}$$

as the difference between the loss functions of approaches $m$ and $k$.

These are the hypotheses of interest:

$$
\begin{aligned}
H_0 \quad &: \quad E[\bar{d}_{mk}] = 0 \quad \text{for all } 1 \leq m \neq k \leq M \\
H_1 \quad &: \quad E[\bar{d}_{mk}] \neq 0 \quad \text{for some } 1 \leq m \neq k \leq M
\end{aligned}
$$

where $\bar{d}_{mk} = n^{-1} \sum_{t=1}^{T} d_{mk,t}$ and $n$ is the sample size.

The test statistic is built using bootstraping methods to generate artificial samples and estimate the standard deviation, and thus:

$$
t_{mk} = \frac{\bar{d}_{mk}}{\sqrt{\widehat{Var}(\bar{d}_{mk})}} \quad \text{for } 1 \leq m \neq k \leq M
$$

If, by any chance, any of the approaches has a realized portfolio variance statistically significantly larger than the others the test will recognize that and the re-run the test without this approach. This is done until there is no statistically significant difference between the realized portfolio variance of the remaining approaches.

## 2.10 Ledoit and Wolf test

Another interesting question is whether the portfolios constructed have significant Sharpe ratio differences when trying to minimize the variance. If two portfolios have the same variance but one has a statistically significant larger Shrape ratio, then that method should be selected.

To test this we will consider the Ledoit and Wolf test proposed in Ledoit and Wolf (2008) to compare Sharpe ratios. Their test is built on the difference between Sharpe ratios of two portfolios. Let $\Delta_{m,k} \equiv \text{Sh}_m - \text{Sh}_k$. For the null hypothesis $H_0 \,:\, \Delta_{m,k} = 0$ the test statistic is:

$$
t^S_{m,k} = -\frac{|\hat{\Delta}|}{s\left(\hat{\Delta}\right)}
$$

where $s\left(\hat{\Delta}\right)$ is estimated using $\sqrt{T}\left(\hat{\Delta} - \Delta\right) \xrightarrow{d} N\left(0; \nabla' f(v) \Phi \nabla f(\hat{v})\right)$, i.e., $s\left(\hat{\Delta}\right) = \sqrt{\frac{\nabla' f(v) \hat{\Phi} \nabla f(\hat{v})}{T}}$, with

$$
\nabla' f(a, b, c, d) = \left( \frac{c}{(c - a^2)^{1.5}}, -\frac{d}{(d - b^2)^{1.5}}, -\frac{1}{2} \frac{a}{(c - a^2)^{1.5}}, \frac{1}{2} \frac{b}{(d - b^2)^{1.5}} \right)
$$

and

$$
\Phi = \lim_{T \to \infty} \frac{1}{T} \sum_{s=1}^{T} \sum_{t=1}^{T} E[y_s y'_t], \text{ where } y'_t = \left(r_{tm} - E[r_{tm}], r_{tk} - E[r_{tk}], r^2_{tm} - E[r^2_{tm}], r^2_{tk} - E[r^2_{tk}]\right)
$$

$r_{tm}$ and $r_{tk}$ are the returns of portfolios $m$ and $k$ respectively. The estimator of $\Phi$ is obtained using HAC inference and block bootstrap.

# 3 Data

The data come from two sources. The first is the high-frequency factors constructed by Aït-Sahalia et al. (2020) available at Dacheng Xiu's website. They use data from NYSE, AMEX, and NASDAQ stock markets for 1996–2017 to build six high-frequency factors. In particular, they contemplate a momentum factor (MOM) plus the five Fama-French factors based on size (SMB), book-to-market ratio (HML), profitability (RMW), investment strategy (CMA), and market return (MKT).

The remaining data come directly from high-frequency stocks data from PiTrading, we will use five minute data of the stocks in the S&P500 index. The stocks change over time as some firms might disappear or lose importance in the market. Both samples range from the December thirtieth 2002 to February seventeenth 2017 amounting to 711 weeks.

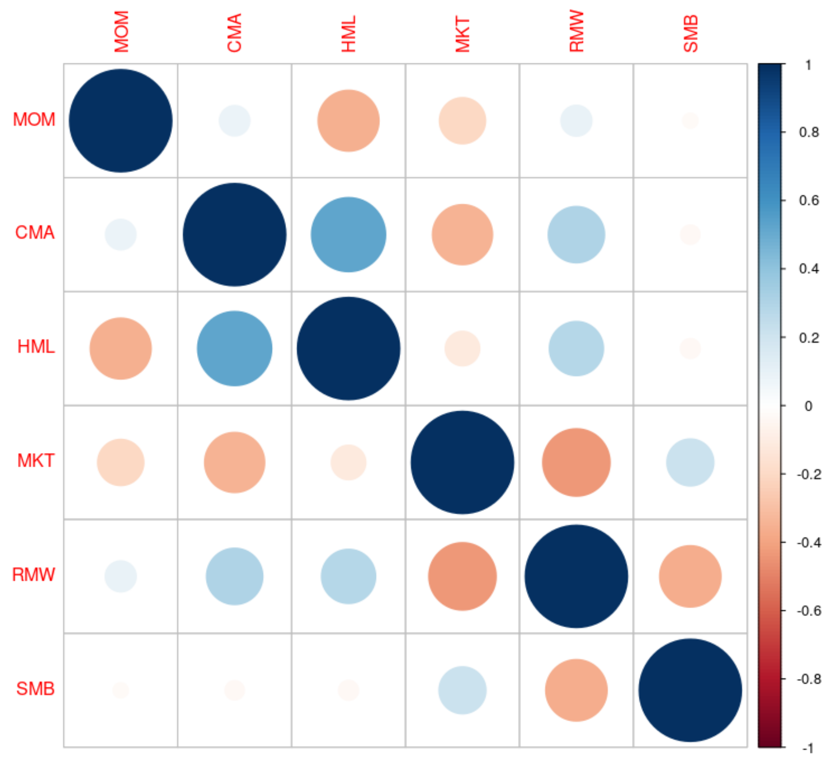


Figure 1: The correlation matrx of Aït-Sahalia et. al factors

The use of empirically chosen high-frequency factors could pose some problem as they are not necessarily orthogonal. As we can see in Figure 1 Aït-Sahalia et al. factors don't have such a high correlation. Another option would be to use sectorial and economic ETFs as factors, but their correlation is significantly higher than the Aït-Sahalia et al. factors.[2]

Figure 2 reveals that market and momentum factors have the widest distribution with the largest range of returns. Our five factors have average returns close to zero, an important feature required for the Realized GARCH model: we can justify the use of squared daily returns without worrying about the drift. The only factor that has a negative average weekly return is momentum (-0.02%). The others have slightly positive returns, the larger being the market factor (0.15%). Looking at the median returns, momentum, market and SMB portfolios have positive median returns, whereas the other three factors have negative.

In Figure 3 we see that the histogram of returns of the stocks is considerably more disperse, because of idiosyncratic shocks. Table 1 shows the annualized volatility for the stocks and for thew factors in the whole sample. While the volatility of the factors range from 5% to 16% per year, the average of the volatilities of the stocks is 31%, with one stock having 71% over the whole period.

To construct our realized measures we use five-minute asset and factor returns. We split the data into sets of five consecutive nonoverlapping trading days and find the covariance matrices of each time frame. With this we have 711 time frames, and for each stock in every window we have one realized covariance matrix between the stock and the factors.

For an asset to be eligible it must have less than 5% of the estimation period with missing data. As we can see from Figure 4 the sample starts with less than 430 eligible stocks and comes to almost 490 in the in the beginning of 2017.

[2]This comparison was done looking at the financial, energy, utilities, industrial, health care and technology sector ETFs and growth, value, small cap and momentum ETFs for the S&P500.

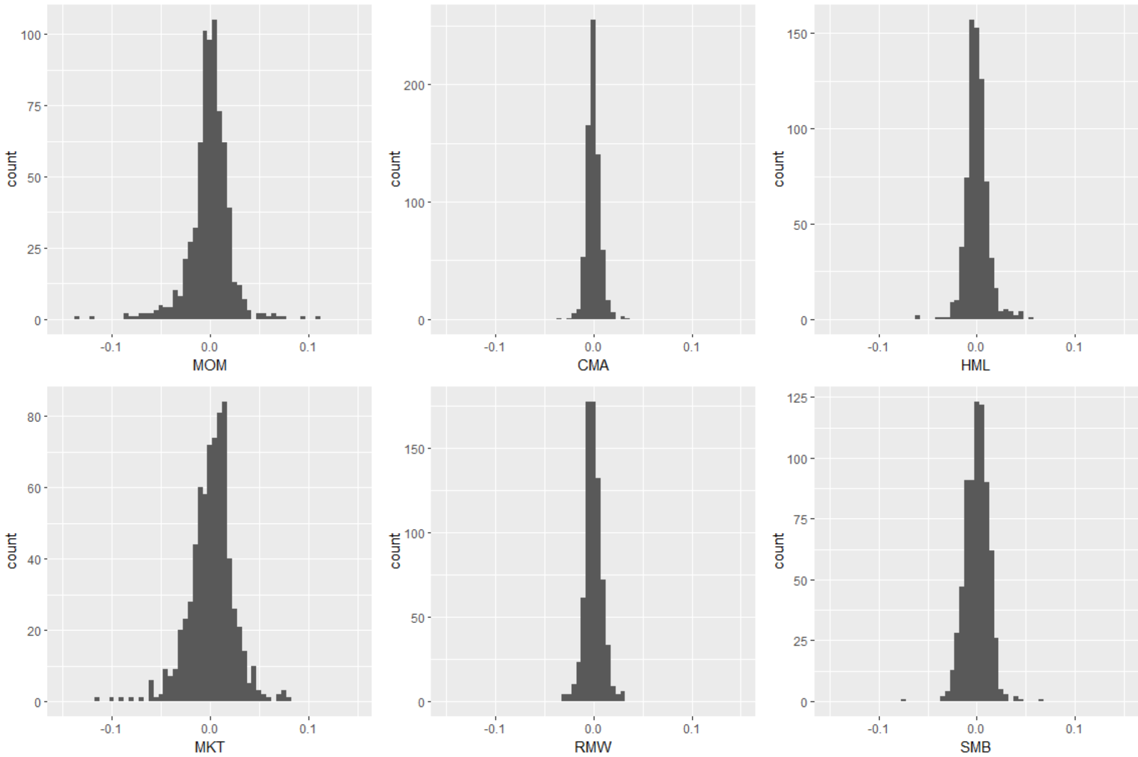


Figure 2: Aït-Sahalia et. al high frequency factors return distribution

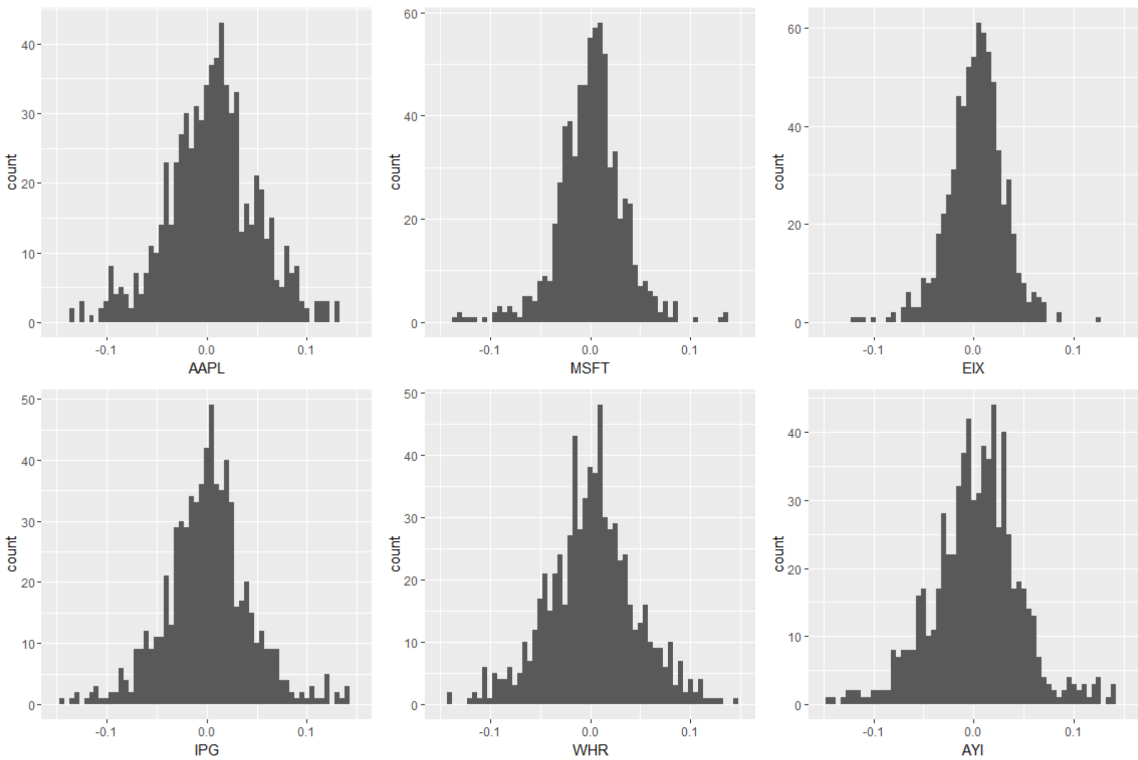


Figure 3: Distribution of returns of six S&P500 stocks

| | Annualized Volatility |
|---|---|
| Momentum | 14.74% |
| Conservative minus aggressive | 4.73% |
| High minus low | 7.99% |
| Market | 16.27% |
| Robust minus weak | 6.26% |
| Small minus big | 8.59% |
| Stocks' min | 14.15% |
| Stocks' median | 30.01% |
| Stocks' mean | 31.45% |
| Stocks' max | 71.13% |

Table 1: Annualized Volatilities of Factors and Stocks

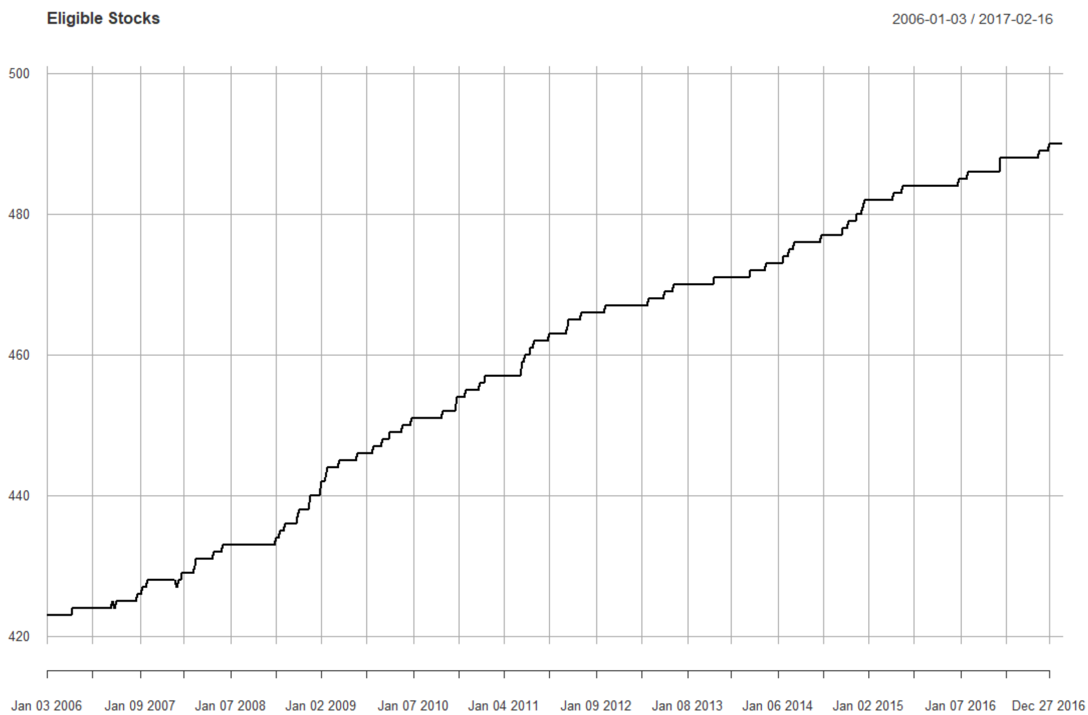


Figure 4: Number of eligible stocks over time

# 4 Empirical Application

In our empirical application we will compute the Global Minimum Variance portfolio for each conditional covariance forecast. To construct a portfolio we will use the forecast of each approach to calculate the optimal weights and keep the weights fixed for five trading days. After the fifth day we will use the newly estimated conditional covariance matrix to compute new weights and use them for the next five days. This will be done throughout the whole sample.

For the estimation we will use a rolling window of 150 weeks. Using all 711 weeks in our sample we remain with 561 weeks to forecast the conditional covariance matrix and build GMV portfolios to compare the predictive ability of the approaches.

To compare our method to standard methods in the literature we will use the $1/N$ naive portfolio and the GMV built with the DCC-NL forecast of the conditional covariance matrix. We expect every model to have a better forecasting performance than the $1/N$ diversification, as it assumes covariance zero between returns.

As for the other benchmark, the DCC-NL was proposed by Engle et al. (2019) and the idea is to use nonlinear shrinkage on the sample covariance matrix to correct the in-sample bias. The use of DCC-NL has been widespread in recent literature as a powerful forecasting tool without the use of intraday data (see De Nard et al. (2019) and Ding et al. (2021)). We expect some benefit on predictive ability when using more information to forecast the conditional covariance matrix as we do in the last two approaches.

## 4.1 Computational Challenges

When implementing this model we were faced with a multitude of computational issues. The treatment of the database and the Static and Realized models were straight forward, but the Realized Multivariate GARCH posed some issues: the model requires minimization of a function with 126 parameters for the factors and 168 parameters for the relation between each stock and the factors.

The factor structure made the estimation feasible, but the number of parameters is still very large. For $n$ assets the number of required parameters is $3n(n+1)$. Computationally this optimization is extremely costly as there is no closed form for the hessian matrix. We utilized an adaptive differential evolution with radius limited sampling to arrive to estimate the parameters.[3]

The estimation of many of the models, most likely, did not reach a global optimum. In many cases the optimizer ran for a couple of times and did not arrive in a feasible estimation for the parameters, the resulting betas were orders of magnitude above what any would expect for a factor loading. Facing this problem, we restricted the optimizer to look for solutions with betas with a magnitude bellow five.

# 5 Does any model beat DCC-NL?

While the proposed estimation approaches managed to forecast the conditional covariance matrices better than the $1/N$ strategy, the DCC-NL proved to be a better alternative when looking at the portfolio returns' standard deviations. As we can see by Table 2, the GARCH-implied beta with unconstrained residual estimation and non linear shrinkage for the idiosyncratic conditional covariance matrix has the best performance among the proposed approaches. This probably happens due to imprecise estimation of the Realized Multivariate GARCH parameters for the relation between each stock and factors (as discussed in the computational challenges section). The nonlinear shrinkage can correct part of this estimation error and provide a better result.

The result of the Realized model implies that the betas do not vary enough over time to justify the use of realized measures. This approach has a slightly higher standard deviation compared to the Static one.

[3]In the end, the estimation took many days to run and was based on a combination of R and Julia programming languages. The optimizer was implemented in Julia and the method utilized was an adaptive differential evolution with radius limited sampling. We tried a multitude of optimizing techniques and this one proved to be the most reliable.

Another interesting result is for the GARCH-implied beta with constrained residual estimation. This approach performed worse than the Static betas, the residual correlation seems to be more important than the variation of betas over time. When assuming that the idiosyncratic conditional correlation matrices are diagonal, we are leaving behind important information about the correlation of the shocks, that is, we are not using all the information that we could to enhance portfolio diversification. Clearly, shocks on the volatility can happen simultaneously in different assets through dimensions beyond the ones encapsulated by the factors, specially when we are using weekly data.

| Static | Realized | Implied Linear | Implied Non Linear | Implied Restricted | 1/N | DCC-NL |
|---|---|---|---|---|---|---|
| 15.20% | 15.21% | 13.08% | 12.72% | 16.82% | 21.87% | 11.66% |

Table 2: Portfolios Returns' Standard Deviations

In Figure 5 we can see how the standard deviations of the portfolios changed over time. We can clearly see that the $1/N$ is the worst performing portfolio. Using the MCS test the set of superior models contains only the DCC-NL, i.e., the DCC-NL conditional covariance matrix forecast creates portfolios with variance statistically significantly smaller than any of the other estimators.

The order of elimination of the models was: GARCH-implied beta with constrained residuals, $1/N$, Static Betas, Realized Betas, GARCH-implied beta with unconstrained residual estimation and linear shrinkage and finally GARCH-implied beta with unconstrained residual estimation and non linear shrinkage. We used a 5% confidence level and 5000 bootstraps

Looking at the standard deviations in high volatility regimes and low volatility regimes we arrive in similar conclusions. For high volatility the order of elimination in the MCS algorithm is the same, with DCC-NL the only one in the superior set. For the low volatility regime the order was: $1/N$, GARCH-implied beta with constrained residuals, Static Betas, GARCH-implied beta with unconstrained residual estimation and linear shrinkage, Realized Betas and GARCH-implied beta with unconstrained residual estimation and non linear shrinkage. Again the DCC-NL remains in the superior set. We can see that the performance of the models are similar in both regimes, the best of our approaches being the GARCH-implied beta with unconstrained residual estimation and non linear shrinkage in both.[4]

We can see in Table 3 the results from Ledoit and Wolf's 2008 Sharpe Ratio test. Interestingly, we were not able to reject the null in any of the pairwise tests, i.e., we cannot say that the difference of the Sharpe ratios of the portfolios are statistically different from zero. From this we gather that the return of each portfolio is apparently compensating enough the extra risk, albeit likely because the portfolios were buying market risk.

| | Static | Realized | Implied Linear | Implied Non Linear | Implied Restricted | 1/N | DCC NL |
|---|---|---|---|---|---|---|---|
| Static | | 0.69 | 0.96 | 0.83 | 0.77 | 0.76 | 0.44 |
| Realized | 0.69 | | 0.89 | 0.90 | 0.88 | 0.83 | 0.51 |
| Implied Linear | 0.96 | 0.89 | | 0.40 | 0.81 | 0.78 | 0.39 |
| Implied Non Linear | 0.83 | 0.90 | 0.40 | | 0.98 | 0.95 | 0.50 |
| Implied Restricted | 0.77 | 0.88 | 0.81 | 0.98 | | 0.82 | 0.57 |
| 1/N | 0.76 | 0.83 | 0.78 | 0.95 | 0.82 | | 0.70 |
| DCC-NL | 0.44 | 0.51 | 0.39 | 0.50 | 0.57 | 0.70 | |

Table 3: Ledoit and Wolf's 2008 Sharpe Test p-values

The turnover of the portfolios can be split into two groups. The DCC-NL, GARCH-implied beta with unconstrained estimation with linear and with non linear shrinkage have turnover ratios larger than the other estimators. Those three are also the ones with the lowest Standard Deviations. We can see in Figure 6 that those two groups remain stable over time. In Table 4 we see that on average the estimator that generates a portfolio with the lowest turnover is the GARCH-implied beta with constrained residual estimation, probably because we are ignoring the covariance of the idiosyncratic component. In the other statistics the results remain similar.

[4]We define high volatility regimes those whose standard deviation are above the median for the $1/N$ portfolio, and low volatility for periods when the standard deviation is bellow the median for the $1/N$ portfolio.

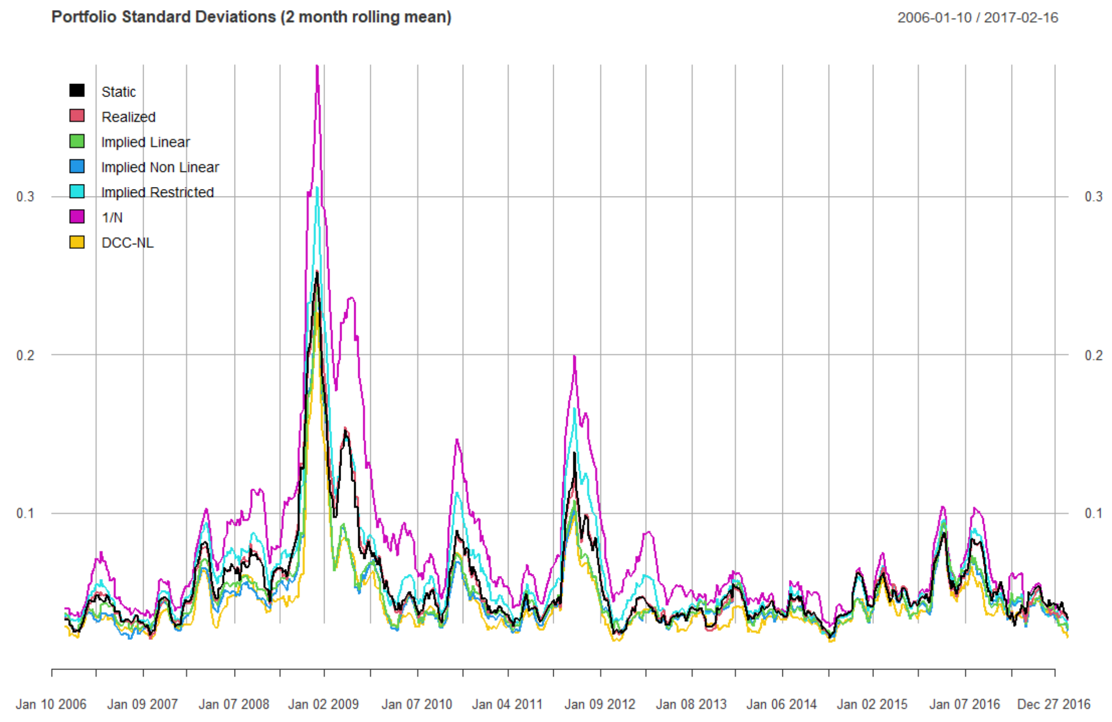


Figure 5: Portfolios' Standard Deviations

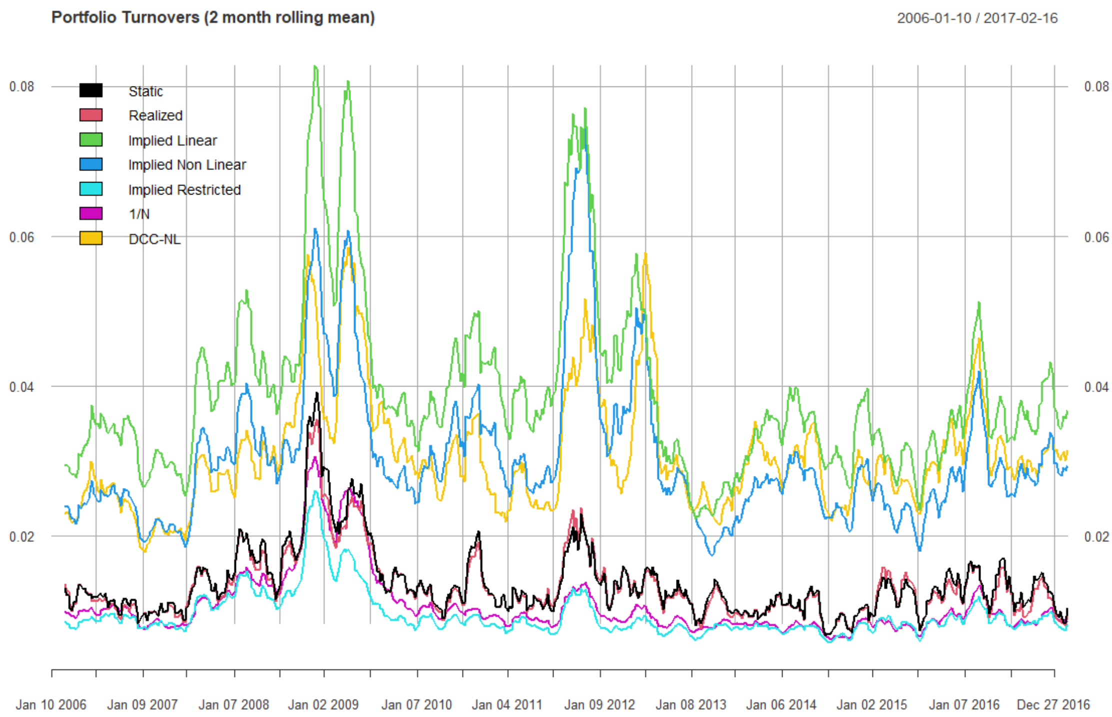


Figure 6: Portfolios' Turnover Ratio

| | Static | Realized | Implied Linear | Implied Non Linear | Implied Restricted | 1/N | DCC NL |
|---|---|---|---|---|---|---|---|
| Min. | 0.35 | 0.35 | 1.65 | 1.39 | 0.36 | 0.42 | 1.30 |
| 1st Q | 0.84 | 0.82 | 2.95 | 2.30 | 0.73 | 0.81 | 2.35 |
| Median | 1.16 | 1.10 | 3.58 | 2.77 | 0.85 | 0.94 | 2.76 |
| Mean | 1.37 | 1.33 | 3.93 | 3.07 | 0.95 | 1.07 | 3.08 |
| 3rd Q | 1.71 | 1.64 | 4.36 | 3.42 | 1.06 | 1.15 | 3.53 |
| Max. | 6.28 | 6.00 | 11.90 | 10.20 | 3.45 | 4.08 | 13.16 |

Table 4: Turnover Ratio

| Elimination Order | Standard Deviation | Standard Deviation (high vol) | Standard Deviation (low vol) | Turnover Ratio |
|---|---|---|---|---|
| 1 | Implied Restricted | Implied Restricted | $1/N$ | Implied Linear |
| 2 | $1/N$ | $1/N$ | Implied Restricted | DCC NL |
| 3 | Static | Static | Static | Implied Non Linear |
| 4 | Realized | Realized | Implied Linear | Static |
| 5 | Implied Linear | Implied Linear | Realized | Realized |
| 6 | Implied Non Linear | Implied Non Linear | Implied Non Linear | $1/N$ |
| Final Set | DCC NL | DCC NL | DCC NL | Implied Restricted |

Table 5: MCS Elimination Order

We can see in Table 5 the elimination order of the models. If we look only at the Standard Deviation, it is clear that the DCC-NL is the best model choice, but one should also consider the turnover of the portfolio as a whole. The conclusion is less obvious since the DCC-NL is the second one to be eliminated and the one that remains in the superior set of models is the GARCH-implied beta with constrained residuals, the one that is eliminated first in the MCS for the Standard Deviation.

Since we have a tradeoff between portfolio turnover and portfolio volatility, it is interesting to look at a measure of how much the turnover raises for a one percent reduction in the volatility comparing to the $1/N$ portfolio. This measure can be seen in Table 6. So, compared to the $1/N$, we see that the Static Beta, Realized Beta and GARCH-implied beta with constrained residual estimation seems to be the best ones, the cost of reducing volatility beyond this point results in possibly large transaction costs. The criterion for the smallest variance is not obvious in this case.

| Static | Realized | Implied Linear | Implied Non Linear | Implied Restricted | DCC NL |
|---|---|---|---|---|---|
| 0.21 | 0.20 | 0.45 | 0.34 | 0.19 | 0.30 |

Table 6: Turnover Gains

Looking at the maximum weight and maximum leverage on a single stock in Tables 7 and 8 we can see that, as expected, the portfolios with highest concentration are the ones with highest turnover ratio. We define maximum weight as the single stock on the portfolio with the largest allocation (long or short) and maximum leverage as the single stock with the largest short allocation. This is relevant as in many scenarios there can be restrictions to leverage or concentration. The relation with the standard deviation is not direct, as we are diversifying considering our forecasts of the conditional covariance matrices. We can see in Figure 7 that the Static Beta, the Realized Beta and the GARCH-implied beta with constrained residual estimation generate portfolios considerably more diversified, with maximum weights under 3% in average, compared to a range between 9% and 10% for the other models.

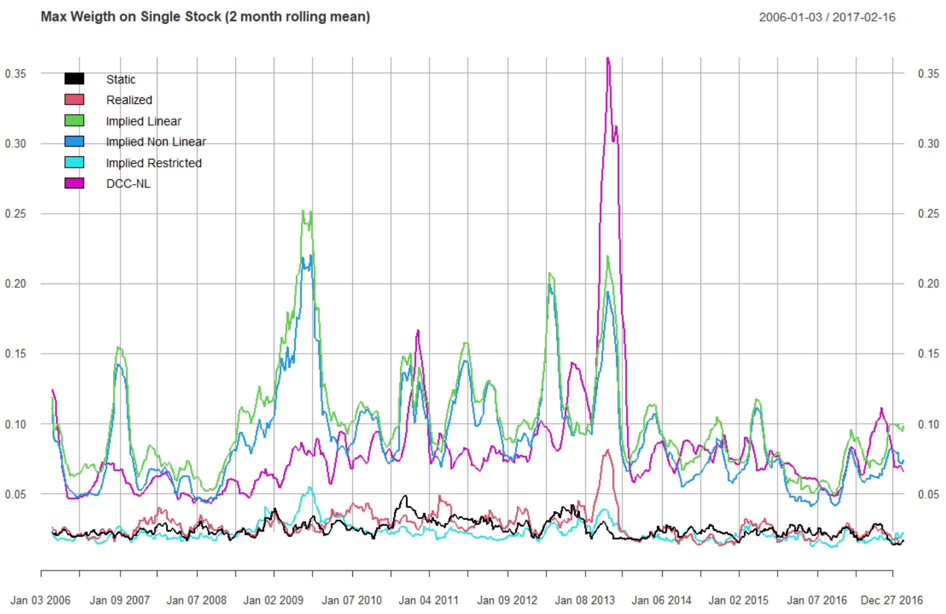


Figure 7: Maximum weight on single stock

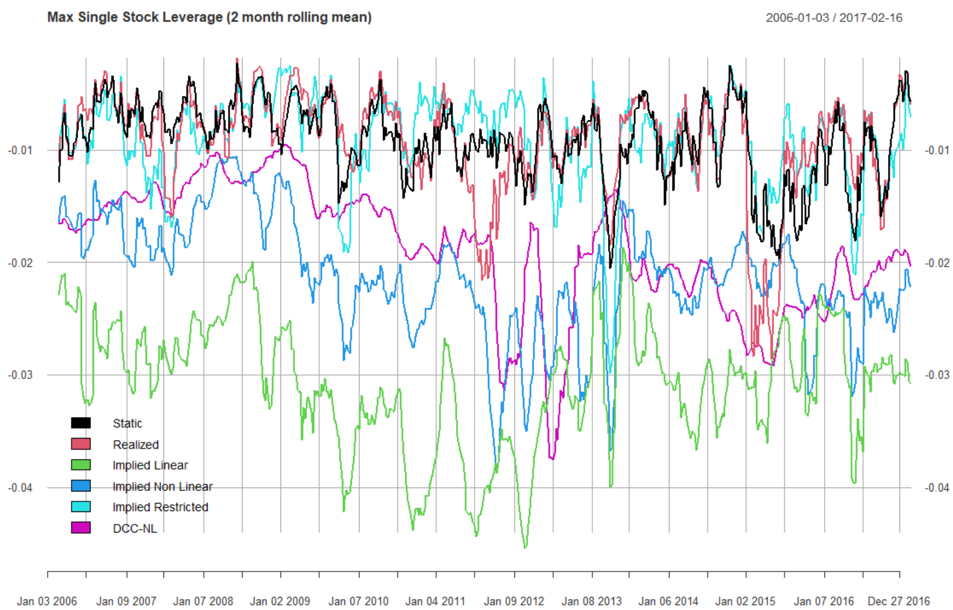


Figure 8: Maximum leverage on single stock

| | Static | Realized | Implied Linear | Implied Non Linear | Implied Restricted | DCC NL |
|---|---|---|---|---|---|---|
| Median | 2.3% | 2.3% | 9.0% | 7.8% | 2.0% | 7.1% |
| Mean | 2.5% | 2.7% | 10.1% | 8.9% | 2.3% | 8.2% |
| Max | 7.9% | 13.4% | 37.1% | 33.2% | 9.3% | 48.0% |

Table 7: Maximum weight on single asset (maximum absolute weight on one stock)

| | Static | Realized | Implied Linear | Implied Non Linear | Implied Restricted | DCC NL |
|---|---|---|---|---|---|---|
| Median | 0.7% | 0.7% | 2.8% | 2.0% | 0.6% | 1.8% |
| Mean | 0.9% | 0.9% | 3.1% | 2.1% | 0.9% | 1.9% |
| Max | 3.8% | 6.9% | 9.0% | 8.3% | 6.5% | 4.6% |

Table 8: Maximum leverage on single asset (maximum short weight on one stock)

# 6 Conclusion

In this paper we have introduced a new way of modeling a Multivariate Realized GARCH fit for a large number of assets. The framework is based on two fronts, the first one using realized measures to enhance the predictive ability of the methods, and the second one to use economic factors to reduce the dimensionality of the estimation and make it feasible.

We base the modeling approach on Archakov et al. (2020) and use their proposal to estimate the relations between factor returns and asset returns. The final conditional covariance matrix comes from the structure imposed by our modified framework.

We applied the Multivariate Realized GARCH with Factors to building Global Minimum Variance portfolios using the S&P500 stocks and compare its results to standard approaches in the literature, the naive $1/N$ portfolio and the DCC-NL. Our proposed methods managed to produce forecasts of the conditional covariance matrices that generated GMVs with lower variance than the $1/N$ portfolio, but the DCC-NL was still the best performing estimator, and the one model selected by MCS. The DCC-NL was also the best performing model when splitting the sample into high a low volatility regimes.

Looking at the turnover of the portfolios, there seems to be a sweet spot. If transactions costs are relevant, one should consider using a different approach as the gains from reducing the standard deviation are not equal among models. Particularly, the approach that yielded the lowest mean turnover ratio if the Realized Multivariate GARCH with Factors restricting the idiosyncratic matrix.

It is also important to note that some models generated portfolios that in practice would be unfeasible, with huge concentration and a large leverage. This happened mostly with the DCC-NL, the GARCH-implied beta with unconstrained residual estimation and linear shrinkage and GARCH-implied beta with unconstrained residual estimation and non linear shrinkage, the ones that provided the lowest Standard Deviation. If the investor has concentration and leverage constraints, as it is common in many places, than the optimal solution might be to use one of the other models.

# References

Archakov, I., Hansen, and A. Lunde (2020). A multivariate realized garch model.

Aït-Sahalia, Y., I. Kalnina, and D. Xiu (2020). High-frequency factor models and regressions. *Journal of Econometrics 216*(1), 86–105.

Barigozzi, M. and M. Hallin (2016). Generalized dynamic factor models and volatilities: recovering the market volatility shocks. *The Econometrics Journal 19*(1), C33–C60.

Barigozzi, M. and M. Hallin (2017). Generalized dynamic factor models and volatilities: estimation and forecasting. *Journal of Econometrics 201*(2), 307–321.

Black, F. (1993). Estimating expected return. *Financial Analysts Journal 49*(5), 36–38.

Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: A multivariate generalized arch model. *The Review of Economics and Statistics 72*(3), 498–505.

De Nard, G., O. Ledoit, and M. Wolf (2019, 01). Factor models for portfolio selection in large dimensions: The good, the better and the ugly. *Journal of Financial Econometrics 19*(2), 236–257.

Ding, Y., Y. Li, and X. Zheng (2021). High dimensional minimum variance portfolio estimation under statistical factor models. *Journal of Econometrics 222*(1, Part B), 502–515. Annals Issue:Financial Econometrics in the Age of the Digital Economy.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business Economic Statistics 20*(3), 339–350.

Engle, R., O. Ledoit, and M. Wolf (2019). Large dynamic covariance matrices. *Journal of Business & Economic Statistics 37*(2), 363–375.

Engle, R., V. K. Ng, and M. Rothschild (1990). Asset pricing with a factor-arch covariance structure: Empirical estimates for treasury bills. *Journal of Econometrics 45*(1), 213–237.

Fabozzi, F. J., F. Gupta, and H. M. Markowitz (2002). The legacy of modern portfolio theory. *The Journal of Investing 11*(3), 7–22.

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

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

Herskovic, B., B. Kelly, H. Lustig, and S. Van Nieuwerburgh (2016). The common factor in idiosyncratic volatility: Quantitative asset pricing implications. *Journal of Financial Economics 119*(2), 249–283.

Jeantheau, T. (1998). Strong consistency of estimators for multivariate arch models. *Econometric Theory 14*(1), 70–86.

Ledoit, O. and M. Wolf (2008). Robust performance hypothesis testing with the sharpe ratio. *Journal of Empirical Finance 15*(5), 850–859.