Cardiff Economics Working Papers



Working Paper No. E2026/12

Learn the measure, estimate the moment: machine-learned drivers in dynamic conditional correlation models

Yongdeng Xu

September 2026

ISSN 1749-6010

Cardiff Business School
Cardiff University
Colum Drive
Cardiff CF10 3EU
United Kingdom
t: +44 (0)29 2087 4000
f: +44 (0)29 2087 4419
business.cardiff.ac.uk

This working paper is produced for discussion purpose. These working papers are expected to be published in due course, in revised form.
Cardiff Economics Working Papers are available online from:
http://econpapers.repec.org/paper/cdfwpaper/ and
https://carbsecon.com/wp/index_abstracts.html
Enquiries: EconWP@cardiff.ac.uk

Learn the measure, estimate the moment: machine-learned drivers in dynamic conditional correlation models

Yongdeng Xu*

Cardiff University

September 16, 2026

Abstract

Realized measures are biased for the moments of returns. We propose a two-step approach to forecasting large correlation matrices: machine-learned forecasts of realized measures enter a dynamic conditional correlation model as drivers, and their weights are estimated by quasi-maximum likelihood on returns. The correlation recursion is run in a matrix-logarithm parametrization, so every forecast is a valid correlation matrix. For Dow Jones stocks, a learned volatility driver and this recursion improve on a dynamic conditional correlation model with realized drivers at horizons of one, five and 22 days. For S&P 500 stocks, the model fits monthly returns better than machine-learned projections of realized correlations, has no invalid forecast in any month, and trades less. The criterion that selects the weight on a realized driver matters as much as the model, and practitioners can obtain valid forecasts calibrated to returns with standard software.

Keywords: DCC-GARCH-X, Realized measures, Covariance forecasting, Forecast evaluation, Positive definiteness, Log-correlation coordinates, Nonlinear shrinkage.

JEL Classification: C32, C53, C58, G11, G17

*Cardiff Business School, Cardiff University, CF10 3EU, UK. Email: xuy16@cardiff.ac.uk.

1 Introduction

The conditional covariance matrix of asset returns is the input of portfolio selection and risk management, and at the dimension of a broad equity index its correlation block carries most of the parameters: 500 stocks have 124,750 distinct correlations against 500 variances. Forecasting that block well, and keeping the forecast a valid correlation matrix, is therefore the harder half of the problem.

The literature has approached the correlation block from two directions. The first models it as a recursion. Engle (2002) introduced the dynamic conditional correlation (DCC) model, driven by the outer product of standardized returns. Shephard and Sheppard (2010) and Noureldin et al. (2012) developed HEAVY models, which let realized measures computed from intraday returns drive the conditional moments. DCC and the DCC-HEAVY model of Bauwens and Xu (2023), in which the lagged realized correlation drives the correlation recursion, are special cases of what we call the DCC-GARCH-X model: variables known at the previous date drive the recursions of the variances, as in the GARCH-X model (see, e.g., Han and Kristensen, 2014), and of the correlations, and the two models differ in these drivers. Engle et al. (2019) extended DCC to several hundred assets by estimating it with composite likelihood (Pakel et al., 2021) around a nonlinearly shrunk target (Ledoit and Wolf, 2020), and De Nard et al. (2022) added intraday information to that model. Archakov and Hansen (2021) provided a parametrization of correlation matrices by the off-diagonal elements of the matrix logarithm, and Archakov et al. (2025) built a multivariate realized GARCH model on it. The second direction projects the realized measure directly. Bollerslev et al. (2026) fit machine-learned projections of monthly realized correlations on 25 realized features, pooled over more than 80,000 pairs of S&P 500 stocks, and assemble the fitted pairs into a forecast matrix.

However, both directions take the realized measure as the moment they forecast, and this assumption is too restrictive in three respects. Firstly, a matrix assembled from pairwise projections need not be positive definite. Bollerslev et al. (2026) report an indefinite forecast in 6.7% of their evaluation months; on our S&P 500 panel, whose sample and measure differ from theirs, their SHAR-F-Exp projection is indefinite in 83.6% of the 183 evaluation months and their LASSO in 27.3%. Secondly, the realized correlation is biased for the correlation of returns. Over 144 months, the mean five-minute realized correlation of S&P 500 stocks is 0.301 against

0.352 for daily returns and 0.339 for monthly returns, and the realized variance is 1.38 times the variance of monthly returns. A forecast fitted to the measure, and evaluated against it, inherits that bias. Thirdly, a projection has no state: the measurement error in each month's realized correlation enters its forecast with weight one.

In this paper, we propose a two-step approach that learns the measure and estimates the moment. In each block of a DCC-GARCH-X model, a learner is fitted outside the likelihood on the realized measure, and a short quasi-maximum likelihood (QML) step, re-estimated at every forecast origin on the returns themselves, maps the fitted driver into the conditional moment of returns. In correlation space the DCC-GARCH-X recursion contains DCC, DCC-HEAVY and the projection of Bollerslev et al. (2026) as special cases; we run it in the log-correlation coordinates of Archakov and Hansen (2021), where the forecast is positive definite at any dimension, for any driver and any parameter value. The level of the correlation block comes from a target estimated on the returns, and the realized measure enters only through its deviations from its own mean, on that panel only through the structure of those deviations.

A simulation in which the true correlation is known shows why the criterion that selects the weight on the driver matters. The returns likelihood selects a weight within 0.04 of the weight selected by an oracle that knows the true correlation, in all 24 combinations of dimension, grid, noise and driver. With microstructure noise and stale prices, that weight falls from 0.63 at 30 stocks to 0.20 at 300 for a projection driver on the fifteen-minute grid, while a criterion computed against the realized correlation selects between 0.77 and 0.96 at every dimension, and at 300 stocks the links it selects are indefinite in every evaluation month.

We apply the approach to Dow Jones stocks at the daily horizon and to S&P 500 stocks at the monthly horizon, where a daily realized correlation matrix is singular by construction. On Dow Jones stocks at horizons of 1, 5 and 22 days, both changes that the approach introduces improve the forecast. Linking the volatility block to a learned driver lowers QLIKE against the returns by 1.16 at one day, 1.34 at five days and 3.29 at 22 days, and running the correlation recursion in coordinates lowers it by 0.29 to 0.30 at one day, with every p value below 0.003. The model that combines the two attains a QLIKE of 52.01 at one day and 55.46 at 22 days, against 53.47 and 58.81 for DCC-GARCH-X with the lagged realized variance and correlation as drivers, with p values below 0.001. At 22 days its minimum variance portfolios also have

a lower realized return standard deviation, 10.67% against 12.38% a year ($t = -2.25$ for the squared returns). On S&P 500 stocks at the monthly horizon, both coordinate links attain a lower returns likelihood than the corrected LASSO of Bollerslev et al. (2026), by 11.36 and 11.82 per day ($t = -3.09$ and -3.27), and are better in 69% and 70% of the 183 evaluation months, and QLIKE against the realized correlation ranks the forecasts in the same order. No link is indefinite in any month, while the failures of their projections grow with the dimension of the assembled matrix. The links also trade less, with a turnover of 1.70 and 1.75 against 2.93. Most of the gain over the projections comes from the level of the target, which alone is better than the corrected LASSO by 10.51 per day, and the rest from the structure that the drivers supply. When the drivers move only the structure of the correlations, the learned coordinates improve on the target by 1.30 per day ($t = -5.55$) and on the lagged realized correlation by 0.46 ($t = -2.95$), and by 0.66 on the fifteen-minute grid ($t = -4.76$). These gains hold with 36-month windows and with the lagged realized variance as the volatility driver.

It is worth noting that our approach is closely related to three models, but there are some differences. Firstly, relative to Bollerslev et al. (2026), the realized measure, or a projection of it, enters as the driver of a recursion whose weight is estimated on returns, so the forecast carries a state and discounts measurement error, and the coordinate forecast is positive definite without their convexity correction; with the LASSO as our learner, the learned coordinates remain better than their corrected LASSO, so the comparison does not rest on boosting. Secondly, relative to DCC-GARCH-X models with realized drivers (Bauwens and Xu, 2023), the drivers of both blocks are learned from many candidates rather than fixed at the lagged realized variance and correlation. Thirdly, relative to Engle et al. (2019) and Archakov et al. (2025), the drivers are fitted outside the likelihood, so the dimension of the candidate set does not enter the quasi-likelihood, and the coordinate recursion is estimated at several hundred stocks.

Overall, a realized measure is better used as a driver than as a target. A forecast fitted to the measure inherits its level and its measurement error, while a weight selected on the returns likelihood discounts both, and at the dimension of the S&P 500 a learned driver adds to the forecast through the structure of the correlations, while their level is taken from the returns. The criterion that selects the weight on a realized driver matters as much as the model.

The remainder of the paper is organized as follows. Section 2 introduces the model and

its estimation, and shows what measurement error implies for the weight on a realized driver. Section 3 describes the data, documents the level of realized measures against the moments of returns, and gives the drivers of each panel and the forecast design. Section 4 presents the results for the Dow Jones panel and the S&P 500 panel. Section 5 concludes. Section S1 of the supplementary appendix (SA) states the conditions under which the two-step estimator is consistent, and the SA contains the proofs, the simulation and additional results.

2 The model

We first state the DCC-GARCH-X model, since the benchmarks and the other forecasts in correlation space are special cases of it. The volatility block and the correlation block are then defined in Sections 2.2 and 2.3, the target of the correlation block in Section 2.4, the estimation in Section 2.5, and the weight on the driver in Section 2.6.

2.1 The benchmark: DCC-GARCH-X

Let $r_t = (r_{1t}, \ldots, r_{kt})'$ denote the $k \times 1$ vector of returns on day t, and let $r_{(j)t}$ denote its j-th intraday return, $j = 1, \ldots, n$. With six and a half hours of trading and five-minute returns, $n = 78$, and adding the overnight return gives 79 returns a day. The conditional covariance matrix is written as

$$H_t = D_t R_t D_t, \qquad D_t = \operatorname{diag}\left(h_{1t}^{1/2}, \ldots, h_{kt}^{1/2}\right), \tag{1}$$

where h_{it} is the conditional variance of asset i and R_t the conditional correlation matrix. The DCC-GARCH-X model specifies

$$h_{it} = \omega_i + \beta_i h_{i,t-1} + \delta_i' x_{i,t-1}, \tag{2}$$

$$Q_t = (1 - \beta_R)\bar{R} + \alpha_R(N_{t-1} - \bar{N}) + \beta_R Q_{t-1}, \qquad R_t = Q_t^{*-1} Q_t \, Q_t^{*-1}, \tag{3}$$

where $x_{i,t-1}$ is a vector of drivers of the variance of asset i known at the end of day $t-1$, N_{t-1} a $k \times k$ matrix driver of the correlation with mean $\bar{N}$, $\bar{R}$ the correlation matrix of the standardized returns $\varepsilon_t = D_t^{-1} r_t$, and $Q_t^* = \operatorname{diag}\left(q_{11,t}^{1/2}, \ldots, q_{kk,t}^{1/2}\right)$ holds the square roots of the diagonal elements of Q_t. In (2), β_i summarizes the persistence of the conditional variance and δ_i represents how informative the drivers are about it. In (3), α_R measures the weight of the

driver, β_R the persistence of the correlation, and the driver enters in deviation from its mean, so that the level of R_t is set by the returns through $\bar{R}$ rather than by the driver. With the realized variance as the driver, the variance equation is stationary if $\beta_i + \delta_i \kappa_i < 1$, with κ_i the ratio of the mean realized variance to the mean conditional variance, and with the squared return, for which $\kappa_i = 1$, this is the GARCH condition; the correlation recursion is stationary if $\alpha_R + \beta_R < 1$.

The model contains the leading specifications as special cases, obtained by the choice of the drivers and restrictions on the parameters. With $x_{i,t-1} = r^2_{i,t-1}$ and $N_{t-1} = \varepsilon_{t-1}\varepsilon'_{t-1}$, whose mean $\bar{N}$ is set equal to the target $\bar{R}$, it is the DCC-GARCH model of Engle (2002), and with $\bar{R}$ shrunk nonlinearly it is the model of Engle et al. (2019). With $x_{i,t-1} = \mathrm{RV}_{i,t-1}$, the realized variance of asset i, and $N_{t-1} = \mathrm{RC}_{t-1}$, the $k \times k$ realized correlation matrix, it is the DCC-HEAVY model of Bauwens and Xu (2023) for the conditional moments of returns, whose variance equation is the HEAVY-r equation of Shephard and Sheppard (2010). With $N_{t-1} = P_t$, the projection of Bollerslev et al. (2026) formed at the end of period $t-1$, $\alpha_R = 1$, $\beta_R = 0$ and $\bar{R} = \bar{N}$, the forecast is the projection itself, and with $\alpha_R = \beta_R = 0$ it is the target $\bar{R}$. When N_{t-1} is a correlation matrix, as the realized correlation and the projection are, $N_{t-1} - \bar{N}$ has a zero diagonal, so Q_t, started at $\bar{R}$, keeps a unit diagonal and $R_t = Q_t$. Our benchmark, which we call DCC-GARCH-X in what follows, takes the lagged realized variance and correlation as its drivers, and Section 3.4 adds, on the daily panel, an unrestricted version that uses every candidate of Section 3.3 linearly.

2.2 The volatility block

We replace the drivers $x_{i,t-1}$ of (2) with a single learned one, in two steps. The first is a boosted fit of the log realized variance on a vector $z_{i,t-1}$ of candidate features dated $t-1$ or earlier,

$$\hat{f} = \underset{f \in \mathcal{H}}{\operatorname{argmin}} \sum_t \sum_i \big(\log \mathrm{RV}_{it} - f(z_{i,t-1})\big)^2 + \Omega(f), \qquad \hat{d}_{it} = \exp\{\hat{f}(z_{i,t-1})\}, \tag{4}$$

where $\mathcal{H}$ is a class of tree ensembles and Ω its regularization. The second enters the fitted driver into the variance recursion,

$$h_{it} = \omega_i + \beta_i h_{i,t-1} + \delta_i \hat{d}_{it}, \tag{5}$$

estimated by QML on the returns, with $\omega_i > 0$ and $\beta_i, \delta_i \geq 0$. The exponential in (4) makes $\hat{d}_{it}$ positive, so (5) delivers $h_{it} > 0$ along every path without signing any coefficient of the learner. The benchmark is the point $\hat{d}_{it} = \mathrm{RV}_{i,t-1}$ of the same recursion. The second step is needed whether or not one wants a recursion: exponentiating a squared-error fit on the logarithm returns a conditional geometric mean of the realized variance (Gneiting, 2011), and the realized variance is not the variance of returns (Section 3.2), so δ_i maps one into the other and is estimated rather than assumed.

The learner can be fitted for each stock and horizon or pooled across stocks, and the recursion (5) can be estimated stock by stock or pooled across stocks with common (β, δ) and a targeted intercept,

$$\omega_i = (1 - \beta)\, \mu_i - \delta\, \bar{d}_i, \tag{6}$$

where μ_i is the mean squared daily return and $\bar{d}_i$ the mean driver over the estimation window. The unconditional level of h_{it} is then the level of returns, and $\omega_i > 0$ for every stock requires $\beta + \delta \max_i(\bar{d}_i/\mu_i) < 1$, which we impose when the intercept is targeted. Section 3.3 gives the features and the learners of each panel.

2.3 The correlation block

The correlation block has a recursion, a driver and a space. Let the driver N_{t-1} of (3) be one of three $k \times k$ correlation matrices known at the end of period $t-1$: the lagged realized correlation RC_{t-1}, as in the benchmark, the projection P_t of Bollerslev et al. (2026), assembled from pairwise forecasts and passed through their convexity correction, or the inverse of learned log-correlation coordinates. Since $R_t = Q_t$ for such a driver, the recursion in correlation space is

$$R_t = (1 - \beta)\bar{R} + \alpha\, \mathcal{S}(N_{t-1} - \bar{N}) + \beta R_{t-1}, \tag{7}$$

where $\bar{N}$ is the mean of the driver over the target window and $\mathcal{S}$ is a linear map on the deviations, defined below, which keeps their zero diagonal and is the identity in (3). At $\alpha = 1$, $\beta = 0$ and $\mathcal{S}$ the identity the forecast is the driver's deviations added to the target, and at $\alpha = \beta = 0$ it is the target alone. Its positive definiteness, however, is a condition on (α, β) and on the driver rather than a guarantee.

We therefore run the recursion in the log-correlation coordinates of Archakov and Hansen

(2021). Let $\gamma = \text{vecl}(\log R)$ denote the $k(k-1)/2$ coordinates, where vecl stacks the strict lower triangle of a symmetric matrix into a vector. The map from correlation matrices to $\mathbb{R}^{k(k-1)/2}$ is a bijection, and its inverse $\mathcal{A}^{-1}$ returns a correlation matrix for any real argument, so the recursion

$$\hat{\gamma}_t = (1-\beta)\bar{\gamma} + \alpha\,\mathcal{S}(m_{t-1} - \bar{m}) + \beta\hat{\gamma}_{t-1}, \qquad R_t = \mathcal{A}^{-1}(\hat{\gamma}_t), \tag{8}$$

gives a forecast that is positive definite at every k, in every period and for every (α, β), where m_{t-1} is the driver in coordinates, $\bar{m}$ its mean over the target window and $\bar{\gamma} = \text{vecl}(\log \bar{R})$ the coordinates of the target. With the lagged realized coordinates as the driver, (8) is the benchmark recursion run in coordinates. The projection of Bollerslev et al. (2026) has no coordinates when it is indefinite, so it enters in correlation space only, and the results of the recursion in correlation space are reported in the SA.

The learned driver. The learned coordinates $\hat{m}_t$ are boosted fits of the realized coordinates on features dated $t-1$. We use gradient-boosted trees rather than a linear learner such as LASSO, so that the fit can capture nonlinear relations and interactions among the candidate features; robustness checks show that the S&P 500 results are close with LASSO as the learner. Since $\gamma_{ij} = \sum_\ell \log(\lambda_\ell)\, u_{\ell i} u_{\ell j}$ over the eigenvalues λ_ℓ and eigenvectors u_ℓ of the realized correlation, the coordinate of one pair depends on all k eigenvalues and eigenvectors, so the features of a pair can include, besides its own history, the rank-r reconstructions $\sum_{\ell \le r} \log(\lambda_\ell)\, u_{\ell i} u_{\ell j}$. Section 3.3 gives the features and the learners of each panel.

The level of the driver. A deviation of the driver from its mean shifts every correlation together through its mean over the pairs, and moves the pairs relative to one another through the remainder, its structure. The map $\mathcal{S}$ decides which of the two enters. It is either the identity or the map that removes the mean over the $k(k-1)/2$ pairs, $\mathcal{S}(A) = A - \bar{a}\,(\iota\iota' - I_k)$ in correlation space, where $\bar{a}$ is the mean off-diagonal element of A and ι is a vector of ones, and $\mathcal{S}(v) = v - \bar{v}\,\iota$ in coordinates, where $\bar{v}$ is the mean element of v, so that the driver moves only the structure and the level of the forecast comes from the target. The second form suits a measure whose bias changes over time, since a deviation from a trailing mean removes a constant bias but not a changing one, and a large cross-section, where a common shift moves every correlation at once and the inverse of the coordinates amplifies it. Section 3.3 states the form used on each panel.

In both spaces, α is the weight the likelihood places on the information in the driver, β the

persistence of the state and $1-\beta$ the weight on the target. Two details of (8) are not cosmetic. Firstly, the target is the coordinate of the target correlation, not the average of coordinates: averaging inside the logarithm and inverting is biased upward, and the inverse amplifies the common direction by a factor of order k. Secondly, the driver enters in deviation from its own mean rather than from the state's, since the intercept of a targeted recursion is $(1-\beta)\bar{\gamma} - \alpha\bar{m}$.

2.4 The target

The target $\bar{R}$ is estimated from returns, as the correlation of the standardized returns over a trailing window. When the number of stocks is not small relative to the number of return vectors in the window, a sample correlation matrix has a dispersed spectrum, so we shrink the covariance of the standardized returns nonlinearly with the analytical formula of Ledoit and Wolf (2020) before normalizing it. This is the target of DCC-NL (Engle et al., 2019). We do not use the trailing mean realized correlation as the target, for the reason documented in Section 3.2, and Section 3.3 gives the window of each panel.

2.5 Estimation

The drivers are fitted outside the likelihood, so what remains to be estimated in each block is a map from a fitted driver to the conditional moment of returns. We estimate it by QML on the returns of the trailing window, and re-estimate it at every forecast origin. On the daily panel the correlation parameters solve

$$(\hat{\alpha}, \hat{\beta}) = \underset{\alpha,\beta}{\operatorname{argmin}} \sum_{t} \Big[\log|R_t| + \varepsilon_t' R_t^{-1} \varepsilon_t \Big], \tag{9}$$

where ε_t are the returns standardized by the variances of (2) with the lagged realized variance as the driver, as in the benchmark. On the monthly panel the correlation is held fixed within a month, so the criterion becomes

$$(\hat{\alpha}, \hat{\beta}) = \underset{\alpha,\beta}{\operatorname{argmin}} \sum_{m} \Big[n_m \log|R_m| + \operatorname{tr}\big(R_m^{-1} S_m\big) \Big], \qquad S_m = \sum_{d \in m} \varepsilon_d \varepsilon_d', \tag{10}$$

where n_m is the number of trading days in month m. Both criteria are minimized subject to $\alpha, \beta \geq 0$ and $\alpha + \beta \leq 0.999$. The volatility parameters are estimated in the same way, by the

Gaussian quasi-likelihood of the returns standardized by h_{it}.

Three properties of this estimation step matter for what follows. Firstly, the recursion nests its driver: at $\alpha = 1$, $\beta = 0$, $\mathcal{S}$ the identity and a target equal to the driver's mean, (7) returns the driver, so the linked model contains the projection it is built on. Secondly, QML calibrates the level: with a free intercept the first-order condition sets the mean of r_{it}^2/h_{it} close to one over the window, and with the targeted intercept of (6) the unconditional level of h_{it} is that of the squared returns by construction, so the model maps the measure to the moment and corrects its level. This is a statement about the unconditional level and not a claim of conditional unbiasedness, which would require a correctly specified model and does not pass through the square roots and ratios that form correlations. Thirdly, the driver is generated rather than observed, the generated-regressor problem of Pagan (1984) in a filtering context: a first-stage error at any date reaches every later state through the recursion. Section S1 of the SA shows that, under regularity conditions and a first stage that converges in mean square, this effect does not accumulate: the QML estimator and the filtered state are consistent, in correlation space and in coordinates; if, in addition, the learned driver converges at the last dates of the window, the one-step forecast converges in probability to the forecast of the link at ϕ_0. The result is a consistency statement only, and the first stage enters the distribution of $\hat{\alpha}$ at first order, so we report $\hat{\alpha}$ throughout as a diagnostic of how much the returns value the driver, without standard errors. The forecast comparisons of Section 4 do not rest on these results: they compare methods out of sample in rolling windows, which evaluates each method inclusive of its estimation step (Giacomini and White, 2006).

2.6 The weight on the driver

What weight should the returns give the driver? Write the deviation of a realized correlation from its mean as $x_t - \bar{x} = (\rho_t^0 - \bar{\rho}^0) + u_t$, with ρ_t^0 the true correlation of a pair, $\mathrm{Var}(\rho_t^0) = \sigma_0^2$ and u_t a measurement error with mean zero, independent of ρ_t^0, with variance σ_u^2 of order $(Dn)^{-1}$ for D days of n intraday returns.

Proposition 1 (Signal-to-noise)**.** *The mean squared error of $\hat{\rho}_t = \bar{\rho}^0 + \alpha(x_t - \bar{x})$ against ρ_t^0 is $(1-\alpha)^2\sigma_0^2 + \alpha^2\sigma_u^2$, which is minimized at $\alpha^* = \sigma_0^2/(\sigma_0^2 + \sigma_u^2) < 1$ whenever $\sigma_u^2 > 0$; a constant bias in the level of x_t does not change α^*.*

For a matrix the cost of the error is larger. Unstructured errors across the $k(k-1)/2$ pairs have a spectral norm of order $2\sigma_u\sqrt{k}$, so the driver can move the smallest eigenvalue of the forecast by $2\alpha\sigma_u\sqrt{k}$, and the right weight falls with k at a given pairwise signal-to-noise ratio, which a criterion that averages squared errors over pairs does not see. A simulation in which the true correlation is known confirms this; Table 1 reports the weight each rule selects. With k between 30 and 300 and realized measures with and without microstructure noise, the returns likelihood (10) selects a weight within 0.04 of that of the QLIKE oracle, which knows the true correlation, in all 24 cells, and for a projection driver on the fifteen-minute grid with noise that weight falls from 0.63 at $k = 30$ to 0.20 at $k = 300$. The squared error against the realized correlation of Bollerslev et al. (2026) selects between 0.77 and 0.96 for that driver at every k, and at $k = 300$ the links it selects are indefinite in every evaluation month. Fitting the measure pays in two cases: without noise at $k \leq 100$ the projection attains the lowest true QLIKE (0.55 against 0.87 at $k = 30$ and $n = 27$), and at $k = 300$ the lowest true risk of the minimum variance portfolio (13.69 against 14.26 with noise and $n = 78$) (SA).

Forecasts at horizons longer than one period average the correlation matrices of the path, and in coordinates we invert the expected coordinate path. The volatility driver exponentiates a fit on logarithms, which returns a conditional geometric mean; it can be rescaled by the smearing factor of Duan (1983), and the remaining difference in level is absorbed by the link's parameters. The learned coordinates enter through their deviations from their own mean with a weight estimated on returns, so a constant difference in level does not reach the forecast.

3 Data and forecast design

3.1 Data

We use five-minute bars from Kibot for every stock that has been an S&P 500 constituent, 862 stocks, over the period between 02/01/1998 and 24/08/2026. Prices are split- and dividend-adjusted, and the 16:00 bar is the closing auction, which we use as the closing price. For a grid of M minutes the intraday prices are the 09:30 open, the closes of the M-minute intervals and the 16:00 close, and the overnight return from the previous close is added, so a day has 79 returns on the five-minute grid and 27 on the fifteen-minute grid of Bollerslev et al. (2026). The daily realized covariance is the sum of the outer products of these returns.

The Dow Jones panel consists of 36 stocks that have been constituents of the index, with 5,352 trading days between 03/01/2005 and 21/08/2026. Its daily realized covariance covers the trading period only, and we rescale it to the level of close-to-close returns with the transformation of Sheppard and Xu (2019), as in Bauwens and Xu (2025),

$$\mathrm{RC}_t^* = \Lambda\,\mathrm{RC}_t\,\Lambda', \qquad \Lambda = \bar{\Sigma}^{1/2}\,\bar{M}^{-1/2}, \tag{11}$$

where $\bar{\Sigma}$ is the mean outer product of the daily returns and $\bar{M}$ the mean realized covariance over an estimation window, and the square roots are symmetric. The mean of the rescaled realized covariance then equals the mean outer product of the returns, correlations included, and every rescaled matrix is positive definite. Bauwens and Xu (2025) find that rescaling the trading-period realized covariance in this way gives more economic value to daily covariance forecasts than adding the outer product of the overnight return, whose estimate of the overnight covariance is noisy. We estimate Λ on the 3,000 days of each estimation window, and for the realized proxy of the evaluation on the 3,000 days before each day, so neither uses later data. Over the 2,352 evaluation days the mean correlation of the rescaled realized covariance is 0.356, against 0.347 for daily returns and 0.325 when only the variances are rescaled (Hansen and Lunde, 2005). Every model of the Dow Jones panel, and every realized criterion, uses the rescaled measure. The S&P 500 panel carries forward the cross-section of our replication of Bollerslev et al. (2026) (SA): of its 412 stocks, 357 have coverage through August 2026 after two stocks whose price series splice across a corporate action are removed. The realized covariance of a month is the sum of the daily realized covariances over the days on which every stock of a balanced sub-panel of 315 stocks is valid, so it is a Gram matrix, and its correlation and the coordinates of that correlation describe the same object; the eigenvalues of the correlation are floored at 0.02 before the logarithm is taken. Since the daily realized covariance already includes the overnight return, we apply (11) to the monthly matrix, which is the form Bauwens and Xu (2025) use for a measure augmented by the overnight return, although at the daily horizon they find it less valuable than rescaling a trading-period measure. $\bar{M}$ is the mean monthly realized covariance of the 60 months before the month and $\bar{\Sigma}$ the mean over the same months of the sum of the daily return outer products; at 315 stocks the daily second moment is shrunk nonlinearly (Ledoit and Wolf, 2020), as the target of Section 2.4 is. Bauwens and Xu (2025) establish the

benefit of the rescaling for daily forecasts, and we assume that it carries over to the monthly horizon. Over the 183 evaluation months the rescaling raises the mean realized correlation from 0.254 to 0.283 on the five-minute grid and from 0.264 to 0.282 on the fifteen-minute grid. The projections of Bollerslev et al. (2026) are fitted to the rescaled measure, and every realized criterion uses it. The drivers of the links are built from the measure before rescaling: they enter only through the structure of their deviations from their mean, with the level taken from the target (Section 3.3), and with the rescaled measure as the driver the likelihood puts a mean weight of at most 0.001 on the lagged realized correlation and on the learned coordinates (SA).

On the five-minute grid a stock-day is invalid when any of its returns is missing, and the missing returns are concentrated in the first three after the open. We impute up to eight missing returns a day by previous-tick sampling: a missing grid price takes the last observed one, so the missing returns are zero except the one that closes the gap, which is the day's close-to-close return minus the returns observed. With this rule 319 of the 357 stocks have valid returns on at least 98% of days, against 139 without it.

3.2 The level of realized measures

Table 2 compares, on the S&P 500 panel, four second-moment matrices of the same stocks and months: the realized covariance from five-minute and from fifteen-minute returns, and the outer products of daily and of monthly returns. A stock enters a month only when its return and both realized variances are valid on every trading day of the month, and every correlation is computed over the months in which both stocks enter.

Insert Table 2 here

Table 2 shows that the realized correlation is below the correlation of returns on every grid and in every period. Over 144 months from 2009 to 2020 and 354 stocks, the mean five-minute realized correlation is 0.301 and the fifteen-minute one 0.308, against 0.352 for daily returns and 0.339 for monthly returns. The gap to monthly returns is -0.048 in 2009–2014 and -0.028 in 2015–2020, and on the extended five-minute panel of 315 stocks it is -0.038 over the 212 months from 2009 to 2026 and -0.055 over the 68 months since January 2021. The realized variance is above the variance of returns by a similar margin: 1.38 times the variance of monthly returns and 1.10 times that of daily returns at five minutes, and 1.33 and 1.07 times at fifteen minutes.

Both directions are what the microstructure literature predicts. Microstructure noise inflates the realized variance (Hansen and Lunde, 2006), which shrinks the realized correlation, and non-synchronous trading attenuates realized covariances (Epps, 1979), more so on the finer grid, which is why the five-minute realized correlation is below the fifteen-minute one. The daily Dow Jones panel of Section 3.1 shows the same direction: over its 2,352 evaluation days the correlation of the pooled realized covariance is 0.325, against 0.347 for daily returns. The implication for forecasting is direct. A forecast whose level is read off the realized measure forecasts the measure: a trailing mean realized correlation used as the target, or a projection of the realized correlation used as the forecast, carries the downward bias of 0.03 to 0.06 in the mean correlation that Table 2 measures. This is why the target of Section 2.4 is estimated from returns and the realized measure enters only through its deviations, and on the monthly panel only through their structure (Section 3.3).

3.3 Drivers and targets

Dow Jones, daily. The volatility learner (4) has 43 candidates: the stock's realized variance and its weekly and monthly means, its squared return, its realized variance on days of negative return, its squared overnight return, the market realized variance, a negative-market indicator, and the realized variances of the other 35 stocks. There is one learner for each stock and horizon, fitted directly on the logarithm of the h-day average of the realized variance, and at $h > 1$ the recursion (5) is iterated with the one-day forecast as the driver of the first day and the h-day forecast as the driver of the remaining days. The learned variance is rescaled by a smearing factor (Duan, 1983) and a level calibration. Each of the 630 coordinates has its own learner, fitted directly at each horizon on ten candidates: the pair's own coordinate and its weekly and monthly means, the mean coordinates of the pairs that share each of the two stocks, the cross-sectional mean coordinate, the two realized variances, the market realized variance and a negative-market indicator. The map $\mathcal{S}$ is the identity, and the target is the sample correlation of the standardized returns of the estimation window, 3,000 days against 36 stocks.

S&P 500, monthly. The two steps of Section 2.2 are taken at the monthly frequency. The target of (4) is the month's realized variance per day, $\mathrm{RV}_{i,m}$, and $h_{i,m}$ is the conditional variance of a daily return in month m. The thirteen features are the logarithms of the stock's realized

variance over the last day, the last week, the month, the quarter and the year, the logarithm of the downside share of its realized variance, its monthly return and absolute return, the logarithm of its mean squared daily return, the cross-sectional means of the log realized variance over the month and the quarter and of the monthly return, and the stock's log realized variance relative to the cross-sectional mean; the learner is pooled over stocks and refitted every January on the trailing 60 months, and the recursion is pooled across stocks with common (β, δ) and the targeted intercept (6). One learner of the coordinates is pooled over the 49,455 pairs and refitted annually with the rolling scheme of Bollerslev et al. (2026), four training years, one validation year and one testing year. Its features are the pair's own coordinate and its means over the last three and twelve months, the rank-r reconstructions for $r = 1$, 3 and 10, the diagonal elements of $\log \mathrm{RC}_{t-1}$ at i and j, the mean coordinates of the pairs containing i, of those containing j and of all pairs, the mean and the smallest log eigenvalue, and the log realized variances of the two stocks and their cross-sectional mean. The map $\mathcal{S}$ removes the mean over the pairs, for two reasons. Firstly, the realized correlation is biased for the correlation of returns by an amount that changes over time, between -0.028 and -0.055 in the periods of Section 3.2. Secondly, a common shift moves all 49,455 correlations at once. We fixed the rule for choosing between the two forms before computing the results of the structure-only form (SA). The target is the correlation of the standardized daily returns $\varepsilon_d = D_m^{-1} r_d$ of the trailing $L = 48$ months, about 1,000 vectors against 315 stocks, and at that concentration ratio, 315/1,000, it is shrunk nonlinearly (Section 2.4).

3.4 Forecast design

Dow Jones, daily. Parameters are estimated on a rolling window of 3,000 days, with the QML parameters and the learners re-estimated every five days and the learners' hyperparameters re-tuned every twenty, and forecasts are made at horizons of 1, 5 and 22 days for 2,352 origins, with evaluation days from 10/04/2017 to 21/08/2026. The realized proxy at $h > 1$ is the average realized covariance over the h days. In every model with a learned driver the link parameters are estimated on the learner's predictions cross-fitted within the window, from five contiguous blocks with 22 days dropped on either side of the block being predicted, so that the driver the likelihood weighs has the quality of an out-of-sample prediction. The benchmarks are DCC-

GARCH-X, with the lagged realized variance and correlation as the drivers, and the unrestricted DCC-GARCH-X. At one day the unrestricted version enters the 43 candidates of Section 3.3 in (2) with non-negative loadings estimated by QML and no restriction on their sum; at $h > 1$ it forecasts the variance by the exponential of a least-squares projection of the h-day average log realized variance on the candidates, rescaled as in Section 3.3; and at every horizon its correlation driver is the inverse of least-squares projections of the realized coordinates on their ten candidates (Section 3.3), which enters (7). The recursions on the lagged realized variance and correlation, in DCC-GARCH-X and in our models, are iterated at $h > 1$ with the expected paths of these drivers from a multiplicative error equation for the realized variance (Shephard and Sheppard, 2010) and a recursion for the realized correlation (Bauwens and Xu, 2023). In the horizon-matched version of each benchmark, least-squares forecasts at each horizon of the log realized variance and of the realized coordinates, on their lagged values for DCC-GARCH-X and on the candidates for the unrestricted version, enter the links (5) and (7) in place of the learned drivers, exponentiated or inverted as those are. The least-squares drivers of the benchmarks enter their links as in-sample fitted values.

S&P 500, monthly. Forecasts are made for one month ahead. The projections of Bollerslev et al. (2026) are refitted on the rescaled measure of the balanced sub-panel with their scheme of four training years, one validation year and one testing year, from the testing year 2005; the learned coordinates start in February 2005 on the five-minute grid and in February 2007 on the fifteen-minute grid. The link parameters are re-estimated at every origin on the trailing 60 months, and on the fifteen-minute grid the first 12 origins have between 48 and 59 months, since its learned coordinates start in February 2007. The evaluation sample runs from February 2011 to April 2026, 183 months, split at January 2021.

3.5 Evaluation criteria

We use statistical criteria computed against returns and against the realized measure, a validity criterion and an economic one, and each is stated with the direction in which a superior model moves it. The returns likelihood is the Gaussian quasi-likelihood of the evaluation period's standardized returns under the forecast correlation; a superior model attains a lower value. QLIKE of the forecast covariance against the realized covariance is robust to noise in the

proxy (Laurent et al., 2013); a superior model attains a lower value. Because the rescaling matches the level of a realized proxy to that of returns only on average over its estimation window, on the daily panel QLIKE is also computed against the mean outer product of the daily returns over the h days, which is conditionally unbiased for the covariance of returns and so ranks forecasts consistently (Patton, 2011), with less power; at $h = 1$ it is the Gaussian quasi-likelihood of the day's return, and a superior model attains a lower value. QLIKE of the forecast correlation against the realized correlation and the pooled R^2 of Bollerslev et al. (2026) against their HAR projection are criteria computed against the measure, and we report them because the comparison with the likelihood is part of the argument of Section 2.6; a superior model attains a lower QLIKE and a higher R^2. Validity is measured by the share of periods in which the forecast is not positive definite and by its smallest eigenvalue; a superior model has fewer indefinite periods. Differences in losses are tested with the test of Diebold and Mariano (1995) using Newey–West standard errors. The economic criterion is the minimum variance portfolio formed from the forecast and held for one holding period, a month on the S&P 500 panel and h days, non-overlapping, on the Dow Jones panel; its realized risk is computed from the rescaled realized covariance of the holding period and annualized, and a superior model produces an optimal portfolio with lower risk. Since the rescaled covariance takes its average from the second moment of returns, and the target of the links is estimated on returns, we also report the standard deviation of the realized return of each portfolio over the holding period, which uses no realized measure, and test the difference in the squared portfolio returns; a superior model attains a lower standard deviation. The turnover of the portfolio is the mean sum of the absolute changes in its weights from one holding period to the next; a superior model trades less.

4 Empirical results

Both panels use the model of Section 2 with the drivers and the design of Section 3. On the Dow Jones panel we test the learned volatility driver and the coordinate recursion against DCC-GARCH-X, and on the S&P 500 panel the correlation block against the projections of Bollerslev et al. (2026).

4.1 Dow Jones, daily

Table 3 reports QLIKE against the returns and against the realized covariance at horizons of 1, 5 and 22 days for the coordinate recursions and four benchmarks, and Panel B of that table tests one change at a time against the returns. Each model is labelled by its two drivers: RV or ML-RV in the volatility block, the lagged realized variance or the learned variance, and RC or ML-RC in the correlation block, the lagged realized correlation or the learned coordinates. Every driver enters through a link estimated by QML, a learned driver on the learner's cross-fitted predictions (Section 3.4). The correlation-space versions of the recursions and the models with a learned correlation driver are reported in the SA; on this panel the learned correlation driver is worse than the lagged realized correlation (SA).

Insert Table 3 here

4.1.1 Does the learned volatility driver help?

It does, at every horizon. Replacing the lagged realized variance in the variance equation by the learned driver lowers QLIKE against the returns by 1.16 at one day, 1.34 at five days and 3.29 at 22 days in the coordinate recursion on the lagged realized correlation ($p < 0.001$), and by similar margins in the three other specifications (SA). The gain grows with the horizon because the iterated forecast on the lagged realized variance over-predicts the variance of returns more as the horizon lengthens: the mean squared return over the h days is 0.95, 0.87 and 0.68 of its forecast variance at 1, 5 and 22 days, against 1.06, 0.98 and 0.92 for the linked forecast, whose recursion is driven by a direct forecast of the h-day average realized variance (SA).

4.1.2 Does the coordinate recursion help?

It does. With the lagged realized correlation as the driver, running the recursion in coordinates rather than in correlation space lowers QLIKE against the returns by 0.30 at one day, 0.19 at five days and 0.07 at 22 days with RV volatility, and by 0.29, 0.21 and 0.11 with linked volatility, with every p value below 0.003. The gain is not a matter of validity, since at $k = 36$ the evaluation's eigenvalue floor binds for no forecast in either space, and 83% to 94% of it survives a shift of both spaces to a common level (SA).

4.1.3 How does the model compare with the benchmarks?

The model with linked volatility and a coordinate recursion on the lagged realized correlation attains the lowest QLIKE against the returns in Table 3 at every horizon, 52.01, 53.41 and 55.46 at 1, 5 and 22 days, against 53.47, 54.95 and 58.81 for DCC-GARCH-X and 53.64, 55.55 and 56.92 for its horizon-matched version.[1] The differences to both are significant at every horizon, with every p value below 0.001 (Table 3, Panel B). At one day the model is also better than the unrestricted DCC-GARCH-X, by 0.94 ($p < 0.001$). The forecasts of the unrestricted DCC-GARCH-X at longer horizons and of its horizon-matched version are near-singular at some origins (Table 3), so their comparisons with the model are not informative. Against the realized covariance the model also attains the lowest QLIKE at every horizon, 47.11, 48.76 and 51.53, against 48.13, 50.08 and 55.43 for DCC-GARCH-X. Overall, the two changes that the approach introduces on the daily panel, a learned volatility driver and a coordinate recursion, both improve the forecast, and together they improve on DCC-GARCH-X by 1.46 at one day, 1.54 at five days and 3.36 at 22 days.

4.1.4 What is the economic value?

Table 4 reports the minimum variance portfolios of DCC-GARCH-X, its horizon-matched version and the coordinate recursions at horizons of 1, 5 and 22 days. Against the rescaled realized covariance, the minimum variance portfolios of the coordinate ML-RV / RC model are less risky than those of DCC-GARCH-X: 10.42%, 10.81% and 11.33% a year at 1, 5 and 22 days, against 10.58%, 11.06% and 11.97%, significant at five and 22 days ($t = -2.54$ and -7.60) but not at one day ($t = -1.84$). In the returns of the portfolios, which use no realized measure, the standard deviation is 13.58%, 11.81% and 10.67% against 13.74%, 12.17% and 12.38%, and the difference in the squared returns is significant only at 22 days ($t = -2.25$). At one and five days the model's portfolios also trade less, with a turnover of 0.86 and 0.95 against 1.11 and 1.36. At 22 days an investor with a relative risk aversion of five would pay 3.87% a year to switch from DCC-GARCH-X to the model (Fleming et al., 2003), but the paired t statistic of the return differences is 1.55. At the same horizon the horizon-matched DCC-GARCH-X is

[1]The inverse $\mathcal{A}^{-1}$ is nonlinear, so by Jensen's inequality the plug-in forecast $\mathcal{A}^{-1}(\mathrm{E}_t\hat{\gamma}_{t+s})$ differs from $\mathrm{E}_t\mathcal{A}^{-1}(\hat{\gamma}_{t+s})$. For the models in Table 3, averaging the inverse over simulated draws of the driver's forecast path lowers QLIKE at 22 days by at most 0.09 under either criterion, and their ranking is unchanged (SA).

less risky than DCC-GARCH-X by 0.21 points ($t = -1.70$), with a return standard deviation of 11.26% against 10.67% for the model, and the fee for switching to it is higher, 4.21% a year (paired $t = 2.07$), a difference that rests on one holding period (SA).

Insert Table 4 here

4.2 S&P 500, monthly

Table 5 reports the forecasts for the 315 stocks on the five-minute grid over the 183 months from February 2011 to April 2026. Besides the two coordinate links, with the lagged realized correlation (RC) and with the learned coordinates (ML-RC) as the driver, it includes the shrunk target alone, the lagged realized correlation, its link in correlation space, and the projections of Bollerslev et al. (2026), fitted to the realized covariance rescaled by (11), with and without their convexity correction, all of which are special cases of the recursion in correlation space, (3) or (7). Every criterion computed against a realized measure uses the rescaled measure. Every link uses the linked volatility block, whose forecasts are closest to the level of returns (1.007, against 1.045 for RV volatility and 1.183 for the learner alone; SA), and its driver enters through its structure (Section 3.3). The correlation-space links and further variants are reported in the SA.

Insert Table 5 here

4.2.1 Which forecast does the returns likelihood prefer?

Against the returns both coordinate links attain a lower likelihood than HAR and the corrected projections of Bollerslev et al. (2026). The difference to their corrected LASSO is -11.36 per day with the lagged realized correlation and -11.82 with the learned coordinates (Newey–West $t = -3.09$ and -3.27), the links are better in 69% and 70% of months, and the differences are -10.45 and -11.09 at the median. Most of this gain comes from the target, which is estimated on returns and alone is better than the corrected LASSO by 10.51 ($t = -2.88$), and the rest from the structure that the drivers supply: the two links improve on the target by 0.84 and 1.30 per day ($t = -6.25$ and -5.55). QLIKE against the rescaled realized correlation ranks the links and the projections as the likelihood does, with 169.43 and 168.64 for the links against 213.68 for the corrected LASSO ($t = -23.97$ and -24.20), while the pooled R^2 of Bollerslev

et al. (2026) favours the projections, 8.8% against −20.0% and −17.1%, as the simulation (SA) leads one to expect for a squared error computed against the measure.

4.2.2 What does the learned driver add?

The learned coordinates improve on the lagged realized correlation as a driver. The returns likelihood is lower by 0.46 per day ($t = -2.95$) on the five-minute grid and by 0.66 ($t = -4.76$) on the fifteen-minute grid, and the link on the learned coordinates is better than the correlation-space link on the lagged realized correlation by 0.78 and 0.81 ($t = -3.44$ and -4.80). The returns also value the learned driver more: the likelihood puts a mean weight of 0.164 on the learned coordinates, with a persistence of 0.14, against 0.025 on the lagged realized coordinates. This result rests on two choices (Sections 3.3 and 3.1). The drivers move only the structure of the correlations, a form fixed by a rule set before its results were computed; with the level of the driver kept, the learned coordinates do not improve significantly on the target (SA). The drivers are also built from the realized measure before rescaling; built from the rescaled measure, their weight is at most 0.001 (SA).

4.2.3 Is the forecast valid?

Both coordinate links are positive definite in every month, as the recursion guarantees, and so is the correlation-space link on the lagged realized correlation; their smallest eigenvalues are 0.071, 0.069 and 0.068. The projections of Bollerslev et al. (2026) are not: SHAR-F-Exp is indefinite in 83.6% of the 183 months and the LASSO in 27.3%, and on the fifteen-minute grid in 63.4% and 20.8%. Table 6 shows that these failures grow with the size of the matrix into which the pairwise forecasts are assembled rather than with $q = k/(Dn)$, the concentration ratio of the realized measure (Marčenko and Pastur, 1967). Up to $k = 100$ SHAR-F-Exp is indefinite in at most 5.7% of months on either grid and the LASSO in at most 0.5%, while at $q = 0.18$ SHAR-F-Exp fails in 2.4% of months on the fifteen-minute grid with $k = 100$, and at $q = 0.19$ in 83.6% on the five-minute grid with $k = 315$. This matches the $2\alpha\sigma_u\sqrt{k}$ term of the matrix argument in Section 2.6: the pairwise error of a projection does not shrink with the intraday grid, so its effect on the smallest eigenvalue grows with k rather than with q.

Insert Table 6 here

The convexity correction of Bollerslev et al. (2026) repairs an invalid forecast, and its trigger of 0.1 is the best of the four triggers tested for the corrected LASSO (119.45 per day, against 122.02 to 127.20); the learned coordinates are better than the corrected LASSO at every trigger (t between -4.23 and -3.27; SA).

4.2.4 What is the economic value?

Table 7 reports the minimum variance portfolios on the five-minute grid. The links trade less: the turnover of their minimum variance portfolios is 1.70 and 1.75, against 2.93 for the corrected LASSO. Against the rescaled realized covariance their portfolios are also less risky, 8.21% and 8.15% a year against 10.15%, a difference of 1.95 and 2.01 percentage points ($t = -11.67$ and -11.98). This covariance takes its average from the shrunk second moment of returns, which the target also estimates, and the standard deviations of the realized portfolio returns do not differ significantly (11.95% against 11.62%, $t = 0.67$ and 0.62 for the squared returns).

Insert Table 7 here

4.2.5 Robustness

Panel B of Table 5 repeats the comparison for the learned coordinates on the fifteen-minute grid of Bollerslev et al. (2026), with 36-month windows for the volatility block, the target and the links, with RV volatilities and with the LASSO as the learner. In every variant the learned coordinates attain a lower returns likelihood than the corrected LASSO, by 8.46 to 13.06 per day (t between -3.27 and -2.00), and than the shrunk target, by 0.14 to 1.35 (t between -6.41 and -2.21), and none is indefinite. Their gain over the lagged realized correlation holds across grids, windows and volatility blocks, by 0.13 to 0.68 per day (t between -5.29 and -2.20), and is not significant with the LASSO learner ($t = -1.02$ and -1.79). At 315 stocks the advantage rests on the nonlinear shrinkage of the target: with the sample correlation as the target, the correlation-space links are worse than the corrected LASSO by 89 to 95 per day (SA).

5 Conclusions

This paper proposes a two-step approach that learns the measure and estimates the moment: a learner is fitted outside the likelihood on the realized measure, and a QML step on the

returns maps the fitted driver into the conditional moment of returns. The model is a DCC-GARCH-X model, which contains DCC, DCC-HEAVY and the projections of Bollerslev et al. (2026) as special cases in correlation space, and we run its correlation recursion in the log-correlation coordinates of Archakov and Hansen (2021), where the forecast is positive definite at any dimension, and its level comes from a target estimated on returns.

The empirical findings are summarized as follows:

1. Realized measures are biased for the moments of returns: over 144 months the five-minute realized correlation of S&P 500 stocks averages 0.301 against 0.339 for monthly returns. In a simulation the returns likelihood selects the weight on a driver that the true correlation selects, while a criterion computed against the realized measure selects between 0.77 and 0.96 at every dimension.

2. On Dow Jones stocks, a learned volatility driver and the coordinate recursion both improve the forecast: linking the volatility block to a learned driver lowers QLIKE against the returns by 1.16 at one day and by 3.29 at 22 days, running the correlation recursion in coordinates lowers it by 0.29 to 0.30 at one day, and together they improve on DCC-GARCH-X with realized drivers by 1.46 at one day and by 3.36 at 22 days; at 22 days the portfolios of the combined model also have a lower realized return standard deviation, 10.67% against 12.38% a year ($t = -2.25$ for the squared returns).

3. On S&P 500 stocks, with the realized covariance rescaled to the level of returns, the coordinate links attain a lower returns likelihood than the corrected LASSO of Bollerslev et al. (2026), by 11.36 and 11.82 per day, and no link forecast is indefinite, against 83.6% of months for their SHAR-F-Exp projection; the links also trade less, with a turnover of 1.70 and 1.75 against 2.93.

4. Most of the gain over the projections comes from the target estimated on returns, and the rest from the structure that the drivers supply: the learned coordinates improve on the target by 1.30 per day and on the lagged realized correlation by 0.46 per day on the five-minute grid and by 0.66 on the fifteen-minute grid, and these gains hold with 36-month windows and with the lagged realized variance as the volatility driver.

On the daily panel the learned correlation driver does not improve on the lagged realized

correlation, and on the monthly panel its gain rests on moving only the structure of the correlations and on the boosted learner (SA). Overall, the criterion that selects the weight on a realized driver matters as much as the model. Thus, market practitioners can readily use the approach, which fits the learners with standard software and estimates a few parameters by QML, to forecast large correlation matrices that are valid by construction in coordinates. For future research, the weight on the driver can be split into separate weights on its level and its structure, or allowed to differ across sectors.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the author used generative AI tools to help program the estimation of the models and to proofread the paper. The author reviewed and checked all code, results and text, and takes full responsibility for the content of the publication.

References

Archakov, I. & Hansen, P.R. (2021). A New Parametrization of Correlation Matrices. *Econometrica*, 89, 1699–1715.

Archakov, I., Hansen, P.R. & Lunde, A. (2025). A Multivariate Realized GARCH Model. *Journal of Econometrics*, doi:10.1016/j.jeconom.2025.106040.

Bauwens, L. & Xu, Y. (2023). DCC- and DECO-HEAVY: Multivariate GARCH Models Based on Realized Variances and Correlations. *International Journal of Forecasting*, 39, 938–955.

Bauwens, L. & Xu, Y. (2025). The Contribution of Realized Variance–Covariance Models to the Economic Value of Volatility Timing. *International Journal of Forecasting*, 41, 1165–1183.

Bollerslev, T., Li, S.Z. & Tang, Y. (2026). Forecasting and Managing Correlation Risks. *Management Science*, Articles in Advance, published online 24 April 2026, doi:10.1287/mnsc.2024.08294.

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, 253–263.

Duan, N. (1983). Smearing Estimate: A Nonparametric Retransformation Method. *Journal of the American Statistical Association*, 78, 605–610.

Engle, R.F. (2002). Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business & Economic Statistics*, 20, 339–350.

Engle, R.F., Ledoit, O. & Wolf, M. (2019). Large Dynamic Covariance Matrices. *Journal of Business & Economic Statistics*, 37, 363–375.

Epps, T.W. (1979). Comovements in Stock Prices in the Very Short Run. *Journal of the American Statistical Association*, 74, 291–298.

Fleming, J., Kirby, C. & Ostdiek, B. (2003). The Economic Value of Volatility Timing Using 'Realized' Volatility. *Journal of Financial Economics*, 67, 473–509.

Giacomini, R. & White, H. (2006). Tests of Conditional Predictive Ability. *Econometrica*, 74, 1545–1578.

Gneiting, T. (2011). Making and Evaluating Point Forecasts. *Journal of the American Statistical Association*, 106, 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, 416–429.

Hansen, P.R. & Lunde, A. (2005). A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data. *Journal of Financial Econometrics*, 3, 525–554.

Hansen, P.R. & Lunde, A. (2006). Realized Variance and Market Microstructure Noise. *Journal of Business & Economic Statistics*, 24, 127–161.

Laurent, S., Rombouts, J.V.K. & Violante, F. (2013). On Loss Functions and Ranking Forecasting Performances of Multivariate Volatility Models. *Journal of Econometrics*, 173, 1–10.

Ledoit, O. & Wolf, M. (2020). Analytical Nonlinear Shrinkage of Large-Dimensional Covariance Matrices. *Annals of Statistics*, 48, 3043–3065.

Marčenko, V.A. & Pastur, L.A. (1967). Distribution of Eigenvalues for Some Sets of Random Matrices. *Mathematics of the USSR-Sbornik*, 1, 457–483.

Noureldin, D., Shephard, N. & Sheppard, K. (2012). Multivariate High-Frequency-Based Volatility (HEAVY) Models. *Journal of Applied Econometrics*, 27, 907–933.

Pagan, A. (1984). Econometric Issues in the Analysis of Regressions with Generated Regressors. *International Economic Review*, 25, 221–247.

Pakel, C., Shephard, N., Sheppard, K. & Engle, R.F. (2021). Fitting Vast Dimensional Time-Varying Covariance Models. *Journal of Business & Economic Statistics*, 39, 652–668.

Patton, A.J. (2011). Volatility Forecast Comparison Using Imperfect Volatility Proxies. *Journal of Econometrics*, 160, 246–256.

Shephard, N. & Sheppard, K. (2010). Realising the Future: Forecasting with High-Frequency-Based Volatility (HEAVY) Models. *Journal of Applied Econometrics*, 25, 197–231.

Sheppard, K. & Xu, W. (2019). Factor High-Frequency-Based Volatility (HEAVY) Models. *Journal of Financial Econometrics*, 17, 33–65.

Table 1 – Simulation: the weight on the driver selected by each rule

k	n	Cell	Driver	Realized criterion	Returns QML	Oracle, QLIKE	Oracle, Frobenius
30	27	clean	projection	0.91	0.61	0.59	0.91
30	27	clean	lagged realized	0.51	0.33	0.33	0.52
30	27	noisy	projection	0.77	0.63	0.59	1.00
30	27	noisy	lagged realized	0.38	0.35	0.31	0.67
30	78	clean	projection	0.91	0.67	0.66	0.92
30	78	clean	lagged realized	0.73	0.56	0.58	0.72
30	78	noisy	projection	0.87	0.73	0.71	1.00
30	78	noisy	lagged realized	0.57	0.40	0.40	0.80
100	27	clean	projection	0.90	0.33	0.34	0.90
100	27	clean	lagged realized	0.50	0.18	0.18	0.50
100	27	noisy	projection	0.95	0.30	0.30	1.00
100	27	noisy	lagged realized	0.44	0.15	0.15	0.70
100	78	clean	projection	0.93	0.42	0.41	0.93
100	78	clean	lagged realized	0.72	0.34	0.33	0.72
100	78	noisy	projection	0.96	0.40	0.41	1.00
100	78	noisy	lagged realized	0.63	0.29	0.28	0.80
300	27	clean	projection	0.91	0.19	0.20	0.90
300	27	clean	lagged realized	0.50	0.11	0.11	0.50
300	27	noisy	projection	0.96	0.20	0.20	1.00
300	27	noisy	lagged realized	0.41	0.10	0.10	0.70
300	78	clean	projection	0.93	0.28	0.29	0.93
300	78	clean	lagged realized	0.73	0.20	0.20	0.73
300	78	noisy	projection	0.96	0.20	0.20	1.00
300	78	noisy	lagged realized	0.65	0.15	0.15	0.80

Mean α over evaluation months and replications, selected by the pooled squared error against the realized correlation, by the returns likelihood (10), and by two infeasible oracles that minimize QLIKE and the Frobenius distance against the true correlation.

Table 2 – The level of realized measures against the moments of returns

Second moment	Mean correlation	Gap to monthly	Variance / daily	Variance / monthly
354 names, 2009–2020, 144 months				
Five-minute realized	0.301	−0.037	1.10	1.38
Fifteen-minute realized	0.308	−0.031	1.07	1.33
Daily returns	0.352	0.013	1.00	1.20
Monthly returns	0.339	–	0.83	1.00
315 names, 2009–2026, 212 months				
Five-minute realized	0.278	−0.038	1.09	1.31
Fifteen-minute realized	0.287	−0.029	1.06	1.27
Daily returns	0.321	0.006	1.00	1.17
Monthly returns	0.316	–	0.86	1.00
315 names, 2021–2026, 68 months				
Five-minute realized	0.230	−0.055	1.07	1.17
Fifteen-minute realized	0.243	−0.042	1.04	1.13
Daily returns	0.263	−0.022	1.00	1.09
Monthly returns	0.284	–	0.92	1.00

Pooled second moments of the same stocks and months; a stock enters a month only when its return and both realized variances are valid on every trading day. Gap to monthly is the difference from the correlation of monthly returns, and the variance ratios are medians over stocks.

Table 3 – Dow Jones panel: QLIKE and one change at a time

Panel A: QLIKE against returns and against the realized covariance

	Against returns			Against the realized covariance		
Model	$h = 1$	$h = 5$	$h = 22$	$h = 1$	$h = 5$	$h = 22$
Benchmarks						
DCC-GARCH-X	53.47	54.95	58.81	48.13	50.08	55.43
Unrestricted DCC-GARCH-X	52.95	102.35	100.60	47.92	89.71	90.32
DCC-GARCH-X, horizon-matched	53.64	55.55	56.92	48.28	50.27	52.67
Unrestricted, horizon-matched	57.58	57.84	59.04	51.83	52.67	54.80
Coordinates						
RV / RC	53.17	54.76	58.75	47.88	49.90	55.35
ML-RV / RC	52.01	53.41	55.46	47.11	48.76	51.53

Panel B: one change at a time, against returns

	ΔQLIKE against returns (p value)		
Comparison	$h = 1$	$h = 5$	$h = 22$
ML-RV vs RV, volatility driver			
Coordinates, RC	−1.16 (<0.001)	−1.34 (<0.001)	−3.29 (<0.001)
Coordinates vs correlation space			
RV / RC	−0.30 (<0.001)	−0.19 (<0.001)	−0.07 (0.002)
ML-RV / RC	−0.29 (<0.001)	−0.21 (<0.001)	−0.11 (<0.001)
Coordinates, ML-RV / RC, against			
DCC-GARCH-X	−1.46 (<0.001)	−1.54 (<0.001)	−3.36 (<0.001)
Unrestricted DCC-GARCH-X	−0.94 (<0.001)	−48.94 (<0.001)	−45.14 (0.001)
DCC-GARCH-X, horizon-matched	−1.62 (<0.001)	−2.14 (<0.001)	−1.46 (<0.001)
Unrestricted, horizon-matched	−5.56 (<0.001)	−4.43 (<0.001)	−3.58 (0.015)

Panel A: QLIKE against the mean outer product of the daily returns (against returns) and against the mean rescaled realized covariance over the h days. DCC-GARCH-X has the lagged realized variance and correlation as drivers, and the unrestricted models use every candidate linearly (Section 3.4); RV / RC and ML-RV / RC are the coordinate recursion (8) on the lagged realized correlation with RV and with linked volatility. The forecasts of the unrestricted DCC-GARCH-X at $h > 1$ and of its horizon-matched version are near-singular at some origins (SA). Panel B: difference in mean QLIKE against returns, second forecast minus first, with Diebold–Mariano p values in parentheses; a negative value favours the change named first in the first two blocks and the coordinate ML-RV / RC model in the third.

Table 4 – Dow Jones panel: minimum variance portfolios

Panel A: portfolios

Forecast	Return (%)	Realized risk (%)	Return s.d. (%)	Sharpe	Turnover
$h = 1$, 2,352 holding periods					
DCC-GARCH-X	8.68	10.58	13.74	0.49	1.11
DCC-GARCH-X, horizon-matched	10.87	11.08	14.06	0.63	1.18
Coordinates, RV / RC	8.11	10.52	13.73	0.44	1.12
Coordinates, ML-RV / RC	8.72	10.42	13.58	0.49	0.86
$h = 5$, 470 holding periods					
DCC-GARCH-X	9.80	11.06	12.17	0.64	1.36
DCC-GARCH-X, horizon-matched	11.08	11.49	12.86	0.71	1.53
Coordinates, RV / RC	9.39	11.02	12.15	0.61	1.36
Coordinates, ML-RV / RC	9.86	10.81	11.81	0.67	0.95
$h = 22$, 106 holding periods					
DCC-GARCH-X	6.39	11.97	12.38	0.35	0.99
DCC-GARCH-X, horizon-matched	9.92	11.77	11.26	0.70	1.36
Coordinates, RV / RC	6.11	11.95	12.42	0.33	0.98
Coordinates, ML-RV / RC	9.26	11.33	10.67	0.68	0.98

Panel B: differences and fees against DCC-GARCH-X

Forecast	Δ realized risk (t)	Δ return s.d. (t)	Fee, $\gamma = 5$ (paired t)
$h = 1$, 2,352 holding periods			
DCC-GARCH-X, horizon-matched	0.50 (8.20)	0.32 (2.23)	1.97 (1.55)
Coordinates, RV / RC	−0.06 (−5.11)	−0.01 (−0.25)	−0.56 (−1.56)
Coordinates, ML-RV / RC	−0.16 (−1.84)	−0.16 (−0.74)	0.15 (0.03)
$h = 5$, 470 holding periods			
DCC-GARCH-X, horizon-matched	0.43 (4.72)	0.69 (2.35)	0.84 (0.97)
Coordinates, RV / RC	−0.04 (−2.91)	−0.03 (−0.71)	−0.40 (−1.39)
Coordinates, ML-RV / RC	−0.25 (−2.54)	−0.36 (−1.53)	0.28 (0.04)
$h = 22$, 106 holding periods			
DCC-GARCH-X, horizon-matched	−0.21 (−1.70)	−1.12 (−1.96)	4.21 (2.07)
Coordinates, RV / RC	−0.02 (−1.81)	0.04 (0.37)	−0.30 (−1.52)
Coordinates, ML-RV / RC	−0.65 (−7.60)	−1.71 (−2.25)	3.87 (1.55)

Panel A: realized risk is the annualized square root of the portfolio variance under the mean rescaled realized covariance of the holding period, and return s.d. the annualized standard deviation of the realized portfolio return, which uses no realized measure; return and Sharpe ratio are annualized, with a risk-free rate of 2%, and turnover is the sum of the absolute changes in the weights between consecutive holding periods. Panel B: differences to DCC-GARCH-X, in percentage points, with the Newey–West t statistics of the differences in realized risk and in squared portfolio returns; the fee is the annual fee, in percent, that an investor with relative risk aversion five would pay to switch from DCC-GARCH-X to the forecast (Fleming et al., 2003), with the paired t statistic of the return differences, and a negative fee means the investor would need to be paid. The unrestricted DCC-GARCH-X and its horizon-matched version, whose forecasts are near-singular at some origins, are reported in SA Table S9.

Table 5 – S&P 500 panel: statistical and validity criteria, and robustness of the learned coordinates

Panel A: statistical and validity criteria, five-minute grid

Forecast	Returns likelihood	QLIKE, covariance	QLIKE, correlation	R^2 vs HAR (%)	Indefinite (%)	Smallest eigenvalue
Shrunk target	108.94	251.01	170.44	−21.9	0.0	0.072
Lagged RC	403.03	485.68	356.80	−3.4	0.0	0.020
HAR	134.77	315.95	228.09	0.0	0.0	0.147
SHAR-F-Exp	–	–	–	5.9	83.6	−1.511
SHAR-F-Exp, corrected	126.89	311.53	224.49	7.4	0.0	0.100
LASSO	–	–	–	7.6	27.3	−0.664
LASSO, corrected	119.45	303.84	213.68	8.8	0.0	0.100
Correlation space, RC	108.42	250.55	170.06	−21.1	0.0	0.068
Coordinates, RC	108.10	250.20	169.43	−20.0	0.0	0.071
Coordinates, ML-RC	107.64	249.75	168.64	−17.1	0.0	0.069

Panel B: robustness of the learned coordinates

Variant	Learned coordinates, returns likelihood difference (t): vs corrected LASSO	vs shrunk target	vs coordinates, RC	Mean α	Indefinite (%)
Five-minute grid (main)	−11.82 (−3.27)	−1.30 (−5.55)	−0.46 (−2.95)	0.164	0.0
Fifteen-minute grid	−13.06 (−2.69)	−1.31 (−6.21)	−0.66 (−4.76)	0.133	0.0
36-month windows	−12.35 (−3.24)	−0.14 (−2.21)	−0.13 (−2.20)	0.016	0.0
RV volatility, five-minute	−8.46 (−2.33)	−1.24 (−5.86)	−0.40 (−2.87)	0.155	0.0
RV volatility, fifteen-minute	−9.90 (−2.00)	−1.35 (−6.41)	−0.68 (−5.29)	0.132	0.0
LASSO learner, five-minute	−11.45 (−3.15)	−0.93 (−6.23)	−0.09 (−1.02)	0.118	0.0
LASSO learner, fifteen-minute	−12.53 (−2.59)	−0.79 (−5.64)	−0.13 (−1.79)	0.081	0.0

Panel A: 183 months, February 2011 to April 2026. Every link uses the ML-RV volatility block; HAR, SHAR-F-Exp and LASSO are the projections of Bollerslev et al. (2026); the uncorrected ones have a likelihood and QLIKE only in the months in which they are positive definite, and are otherwise reported on R^2 and validity (–). The returns likelihood is per day, QLIKE is against the rescaled realized covariance and correlation, and R^2 pools the pairwise squared errors over all 183 months, indefinite ones included. Panel B: difference in the returns likelihood per day of the link on the learned coordinates, with Newey–West t statistics; a negative value favours it. The LASSO rows compare with the boosted run's link on the lagged realized correlation, and with 36-month windows that link nearly coincides with the shrunk target.

Table 6 – S&P 500 panel: validity of the projections of Bollerslev et al. (2026) against the dimension

	Fifteen-minute grid			Five-minute grid		
		Indefinite	$\lambda_{\min} < 0.05$		Indefinite	$\lambda_{\min} < 0.05$
k	q	(%)	(%)	q	(%)	(%)
SHAR-F-Exp						
25	0.04	0.0	0.2	0.02	0.4	0.4
50	0.09	0.1	0.9	0.03	0.7	1.2
100	0.18	2.4	4.0	0.06	5.7	11.0
150	0.27	9.4	17.2	0.09	23.5	34.5
200	0.36	27.3	37.2	0.12	48.9	64.1
250	0.45	44.3	56.2	0.15	70.9	79.9
315	0.56	63.4	73.2	0.19	83.6	90.7
LASSO						
25	0.04	0.0	0.0	0.02	0.0	0.0
50	0.09	0.0	0.0	0.03	0.0	0.0
100	0.18	0.0	0.1	0.06	0.5	1.3
150	0.27	0.8	2.3	0.09	3.2	6.8
200	0.36	4.6	8.9	0.12	10.7	16.2
250	0.45	11.0	17.1	0.15	19.4	24.9
315	0.56	20.8	29.5	0.19	27.3	36.1

Share of months in which the principal submatrix of the stored forecast for a random subset of k stocks is indefinite, or has a smallest eigenvalue below 0.05, averaged over 20 subsets; $q = k/(Dn)$, with D the mean number of days in a month and n the number of intraday returns a day. HAR is positive definite at every k.

Table 7 – S&P 500 panel, five-minute grid: minimum variance portfolios

Panel A: portfolios

Forecast	Return (%)	Realized risk (%)	Return s.d. (%)	Sharpe	Turnover
Shrunk target	12.34	8.29	11.96	0.86	1.67
Lagged RC	13.02	9.10	11.26	0.98	4.15
HAR	13.51	9.92	10.74	1.07	3.07
SHAR-F-Exp, corrected	13.33	10.35	11.51	0.98	3.36
LASSO, corrected	12.65	10.15	11.62	0.92	2.93
Correlation space, RC	12.35	8.23	11.88	0.87	1.71
Coordinates, RC	12.48	8.21	11.95	0.88	1.70
Coordinates, ML-RC	12.59	8.15	11.95	0.89	1.75

Panel B: differences and fees against the corrected LASSO

Forecast	Δ realized risk (t)	Δ return s.d. (t)	Fee, $\gamma = 5$ (paired t)
Shrunk target	−1.87 (−11.29)	0.34 (0.68)	−0.53 (−0.20)
Lagged RC	−1.05 (−10.78)	−0.35 (−0.77)	0.57 (0.33)
HAR	−0.23 (−2.76)	−0.87 (−1.08)	1.37 (0.82)
SHAR-F-Exp, corrected	0.20 (3.76)	−0.10 (−0.19)	0.74 (1.15)
Correlation space, RC	−1.93 (−11.49)	0.26 (0.51)	−0.46 (−0.19)
Coordinates, RC	−1.95 (−11.67)	0.33 (0.67)	−0.38 (−0.11)
Coordinates, ML-RC	−2.01 (−11.98)	0.33 (0.62)	−0.27 (−0.04)

Panel A: realized risk is the annualized square root of the portfolio variance under the month's rescaled realized covariance, which takes its average from the shrunk second moment of returns that the target also estimates, and return s.d. the annualized standard deviation of the realized monthly portfolio return, which uses no realized measure; return and Sharpe ratio are annualized, with a risk-free rate of 2%, and turnover is the sum of the absolute monthly changes in the weights. The uncorrected projections form the same portfolios as the corrected ones. Panel B: differences to the corrected LASSO, in percentage points, with the Newey–West t statistics of the differences in realized risk and in squared monthly returns; the fee is the annual fee, in percent, that an investor with relative risk aversion five would pay to switch from the corrected LASSO to the forecast (Fleming et al., 2003), with the paired t statistic of the monthly return differences, and a negative fee means the investor would need to be paid.