

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

A Score-Driven Conditional Correlation Model for Noisy and Asynchronous Data: an Application to High-Frequency Covariance Dynamics

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Buccheri, G., Bormetti, G., Corsi, F., Lillo, F. (2021). A Score-Driven Conditional Correlation Model for Noisy and Asynchronous Data: an Application to High-Frequency Covariance Dynamics. JOURNAL OF BUSINESS & ECONOMIC STATISTICS, 39(4), 920-936 [10.1080/07350015.2020.1739530].

Availability:

This version is available at: <https://hdl.handle.net/11585/746692> since: 2023-02-14

Published:

DOI: <http://doi.org/10.1080/07350015.2020.1739530>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Buccheri, G., Bormetti, G., Corsi, F., & Lillo, F. (2021). A score-driven conditional correlation model for noisy and asynchronous data: An application to high-frequency covariance dynamics. Journal of Business and Economic Statistics, 39(4), 920-936

The final published version is available online at
<https://dx.doi.org/10.1080/07350015.2020.1739530>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

A Score-Driven Conditional Correlation Model for Noisy and Asynchronous Data: an Application to High-Frequency Covariance Dynamics*

Giuseppe Buccheri¹, Giacomo Bormetti², Fulvio Corsi^{3,4}, and Fabrizio Lillo^{2,5}

¹University of Rome Tor Vergata, Italy

²University of Bologna, Italy

³University of Pisa, Italy

⁴City University of London, UK

⁵CADS, Human Technopole, Milan, Italy

March 2, 2020

Abstract

The analysis of the intraday dynamics of covariances among high-frequency returns is challenging due to asynchronous trading and market microstructure noise. Both effects lead to significant data reduction and may severely affect the estimation of the covariances if traditional methods for low-frequency data are employed. We propose to model intraday log-prices through a multivariate local-level model with score-driven covariance matrices and to treat asynchronicity as a missing value problem. The main advantages of this approach are: (i) all available data are used when filtering the covariances, (ii) market microstructure noise is taken into account, (iii) estimation is performed by standard maximum likelihood. Our empirical analysis, performed on 1-second NYSE data, shows that opening hours are dominated by idiosyncratic risk and that a market factor progressively emerges in the second part of the day. The method can be used as a nowcasting tool for high-frequency data, allowing to study the real-time response of covariances to macro-news announcements and to build intraday portfolios with very short optimization horizons.

Keywords: Intraday Correlations; Dynamic Dependencies; Asynchronicity; Microstructure Noise

JEL codes: C58; D53; D81

*Corresponding author: fabrizio.lillo@unibo.it. We wish to thank the Editor, the Associate Editor and the Reviewers for extensive comments that greatly strengthened the article. We are particularly grateful for suggestions we have received from Maria Elvira Mancino, Davide Delle Monache, Ivan Petrella, Fabrizio Venditti, Giampiero Gallo, Davide Pirino and participants to the IAAE 2017 conference in Sapporo, the 10th SoFiE conference in New York, the VIECO 2017 conference in Wien and the workshop on “Score-Driven Time-Series models” in Cambridge in March 2019.

1 Introduction

A large class of conditional covariance models have been proposed in the econometric literature and their use is widespread in risk and portfolio management at daily or lower frequencies. Popular multivariate dynamic time-series models include the class of multivariate extensions of the univariate GARCH model of Engle (1982) and Bollerslev (1986) and the Dynamic Conditional Correlation (DCC) model of Engle (2002). A drawback of these models is that they are misspecified if data are recorded with observational noise and require synchronization in case data are irregularly spaced. As a consequence, they cannot be straightforwardly applied to intraday data, since high-frequency prices are contaminated by microstructure noise and assets are traded asynchronously. Both effects may lead to ignore a large portion of the data and to significantly underestimate the correlations. The problem of estimating and forecasting intraday volatilities and correlations is, however, of crucial importance in high-frequency finance. For instance, an high-frequency trader is interested in rebalancing the portfolio on an intraday basis, and thus she needs accurate short-term covariance forecasts. Similarly, the study of the intraday dependencies of financial assets is useful to examine the reaction of the market to external information and has a theoretical relevance in market microstructure research.

We contribute to the literature on intraday covariance estimation by proposing a modeling strategy that can handle both asynchronous trading and microstructure effects. High-frequency log-prices are modeled through a conditionally normal local-level model where the efficient log-prices are affected by measurement errors and the covariances of both the efficient returns and the noise are time-varying. In this state-space representation, asynchronous trading can be treated as a standard missing value problem. Essentially, this means that a multivariate time-series with asynchronous observations is regarded as a synchronized time-series with missing values. The dynamics of the time-varying parameters are driven by the score of the conditional density (Creal et al. 2013, Harvey 2013). The latter can be computed in closed form using the Kalman filter (Creal et al. 2008, Delle Monache et al. 2016, Delle Monache et al. 2019), and thus estimation proceeds via standard maximum likelihood.

The main advantage of the proposed approach is that it allows to model the covariances of the latent efficient returns using all the available *observed prices*. In standard conditional covariance models, the covariances of the *observed returns* are instead modeled. While in low frequency data observed price models are essentially equivalent to observed return models, in high-frequency data the latter are subject to several distortions and are affected by data reduction. For instance, in order to deal with asynchronous trading, it is a standard practice to impute the missing trades by previous-tick interpolation. The latter produces a large number of zero returns, which are known to jeopardize the inference of the covariances. This effect is analogous to the downward bias of high-frequency sample covariances, the well-known “Epps effect” (Epps 1979), which arises when using previous-tick interpolation (see Hayashi and Yoshida 2005 and references therein). Our approach, treating asynchronicity as a missing value problem, avoids the introduction of artificial zero returns and thus it is not subject to bias. As it will be shown in our simulation study, alternative methods to deal with asynchronous trading in observed return models, e.g. the missing value approach of Lucas et al. (2016), lead to significant data reduction.

Microstructure effects, which are generally modeled as measurement errors, constitute an additional source of distortion in the inference of standard conditional covariance models. As in error-in-variables models, they lead to a bias in the estimation of the parameters. By modeling the covariances of the latent efficient returns, we are guaranteed that the recovered covariance estimates are not affected by microstructure noise. Compared to standard conditional covariance models, the proposed approach is thus specifically designed to deal with high-frequency data and can easily be employed to construct intraday portfolios. To this end, we discuss two alternative parameterizations of the covariance matrix leading to positive-definite estimates.

Since Andersen and Bollerslev (1997) and Tsay (2005), it is known that intraday volatilities have the typical U-shape, being larger at the opening and closing hours of the trading day. In contrast, due to the aforementioned difficulties, the intraday behavior of correlations has received less attention. Notably exceptions are given by the work of Bibinger et al. (2019), who proposed a nonparametric spot covariance estimator for intraday data, and that of Koopman et al. (2018), which is based on dynamic copula models. The main difference between our approach and the dynamic copula of Koopman et al. (2018) is that we model the dependencies of the latent efficient returns instead of the observed returns. As a consequence, our approach does not suffer from data reduction when applied to high-frequency data and it is robust to microstructure noise. Multivariate GARCH generalizations have been proposed, among others, by Engle and Kroner (1995), Tse and Tsui (2002), van der Weide (2002), Alexander (2002), Engle (2002), Creal et al. (2011). As underlined above, asynchronicity and market microstructure effects can lead to several unwanted features in the inference of these models. In our simulation and empirical study, we examine in detail the impact of these two effects on the DCC model of Engle (2002) and the t -GAS model of Creal et al. (2011).

Score-driven models are a general class of observation-driven models (Cox 1981) where the dynamics of the time-varying parameters are driven by the score of the conditional likelihood. They have been successfully applied in the recent econometric literature (see e.g. Creal et al. 2011, Creal et al. 2014 and Oh and Patton 2018). One of the main advantages of this class of models is that the conditional log-likelihood can be written in closed form, and thus estimation is performed by standard maximum likelihood. In a linear-Gaussian state-space representation, the score of the system matrices can be computed through the Kalman filter, as described by Creal et al. (2008), Delle Monache et al. (2016) and Delle Monache et al. (2019). We build upon these results to estimate and forecast the intraday covariances in our state-space representation.

The Monte-Carlo analysis presented in this paper has three main goals. First, we investigate the finite sample properties of the maximum likelihood estimator. We find that it remains unbiased even in case many observations are missing. Then, we use the model as a filter for a misspecified DGP and compare it to standard dynamic models employed in a low-frequency setting. We find that, in the presence of measurement errors and asynchronous observations, standard methods are subject to a downward bias. The local-level model performs significantly better, as it exploits all the available data and provides robustness to measurement errors. Finally, we investigate the performance of the model in the presence of fat-tails and asynchronicity and/or measurement errors. To this end, we use the t -GAS model of Creal et al. (2011) to simulate the intraday log-prices and the covariances. After randomly censoring the log-prices, we estimate

both the local-level model and the t -GAS. We find that, as the number of missing observations increases, the local-level model provides lower in-sample and out-of-sample average losses. The t -GAS, being a model for the observed returns, is indeed subject to large data reduction in the presence of asynchronous observations. A similar result is obtained when we add a measurement error to the simulated log-prices. We conclude that, in a high-frequency setting, the use of the local-level model is preferable to that of correctly specified observed return models, even under extreme fat-tails.

We apply the model to transaction data of 10 NYSE stocks. The in-sample analysis based on the AIC reveals that the local-level model fits the data significantly better than standard conditional covariance models for low-frequency data. We find that correlations have an increasing pattern. The rate of increase is larger during the first two hours, then correlations increase at a slower rate and tend to decrease during the last 15 minutes of the trading day. By analyzing the dynamics of the largest eigenvalues of the correlation matrix, we interpret this phenomenon as the emergence of a market factor that progressively explains a larger fraction of the total variance of the market.

The local-level model, being robust to asynchronous trading, can be estimated at ultra high frequencies (1-second in our application) and thus provides a description of the dynamics of the covariances at very small time scales. This allows us to study the real-time response of correlations to macro-news announcements. Once macro-news arrive on the market, they are instantaneously captured by the model, even if very few assets are traded at that moment. The method can thus be employed as a nowcasting tool for high-frequency data. This interesting feature shares some similarities with the macroeconomic literature on nowcasting, where dynamic factor models are used to update the forecasts of macroeconomic variables based on mixed-frequency observations (see e.g. Giannone et al. 2008, Creal et al. 2014 and Delle Monache et al. 2016). In the second part of the empirical analysis, we assess the performance of the model as a nowcasting tool for high-frequency data. We construct intraday out-of-sample portfolios with short investment horizons and find that the portfolios built through the local-level model feature significantly lower ex-post risk. Furthermore, the local-level model provides reliable intraday Value-at-Risk assessment, whereas the competing approaches largely fail in assessing the risk of high-frequency returns.

We collect in an online appendix several extensions of the proposed methodology. First, in order to avoid the detrimental effect of jumps, which may occur at high frequencies, we propose a modification to the standard Gaussian algorithm that provides robust estimates of both the latent efficient log-prices and the covariances. Inspired by Harvey and Luati (2014) and Buccheri et al. (2019a), we replace the Gaussian scores driving the dynamics of the time-varying parameters with the scores of a Student- t density. The implementation of the alternative update scheme serves also as a robustness check for our Monte-Carlo and empirical results. Second, we discuss alternative parameterizations of the covariances which reduce the number of parameters when the cross-section dimension increases. Finally, we show how intraday volatility patterns can be included in the model.

The remaining part of the paper is organized as follows: in Section (2) we describe the methodology in its full generality, including the parameterization of the correlations and the estimation method; Section (3) discusses the results of Monte-Carlo experiments; in Section (4) we provide empirical evidence on the

advantages of the model over standard techniques and study the intraday dynamic behavior of covariances; Section (5) concludes. Model extensions are presented in the online supplemental material.

2 Framework

2.1 Model

Let $t \in [0, S]$ and denote by $X_t = (X_t^{(1)}, \dots, X_t^{(n)})'$ an $n \times 1$ vector of intraday efficient log-prices. We consider T equally-spaced observation times $0 \leq t_0 < t_1 < \dots < t_{T-1} \leq S$ and propose to model X_{t_i} , $i = 0, \dots, T-1$, as a random walk with heteroskedastic innovations:

$$X_{t_i} = X_{t_{i-1}} + \eta_{t_i}, \quad \eta_{t_i} \sim (0, Q_{t_i}) \quad (1)$$

The time-varying covariance matrix Q_{t_i} describes the dynamics of volatilities and correlations of the efficient returns and is the main object of interest of this paper. The efficient log-price X_{t_i} is unobservable because of market microstructure effects (e.g. bid-ask bounce). Let Y_{t_i} be the $n \times 1$ vector of observed log-prices. We write:

$$Y_{t_i} = X_{t_i} + \epsilon_{t_i}, \quad \epsilon_{t_i} \sim (0, H_{t_i}) \quad (2)$$

where ϵ_{t_i} is an uncorrelated measurement error term representing market microstructure effects. The two innovations η_{t_i} and ϵ_{t_i} are assumed to be independent from each others. The covariances in H_{t_i} are allowed to vary over time to capture potential dynamic effects in microstructure noise. For instance, the bid-ask spread has a well-known intraday pattern, being larger at the opening hours and then declining throughout the day (McInish and Wood 1992). The model in Eq. (1), (2) is at the basis of traditional market microstructure analysis of trading frictions, asymmetric information and inventory control (Roll 1984, Hasbrouck 1993, Madhavan 2000).

2.2 Time-varying covariances

Eq. (1) describes the martingale dynamics of the efficient log-price process, whereas Eq. (2) is the associated measurement equation. We can re-write both equations as:

$$Y_t = X_t + \epsilon_t, \quad \epsilon_t \sim (0, H_t) \quad (3)$$

$$X_t = X_{t-1} + \eta_t, \quad \eta_t \sim (0, Q_t) \quad (4)$$

where, without loss of generality, we have set $t_i - t_{i-1} = 1$. The model described by Eq. (3), (4) is known as *local-level* model (Harvey 1991, Durbin and Koopman 2012). If the two covariance matrices H_t and Q_t are constant, the local-level model can be estimated by quasi-maximum likelihood through the Kalman filter. This is the route of Corsi et al. (2015) and Shephard and Xiu (2017), who propose a quasi-maximum likelihood estimator of the integrated covariance of high-frequency log-prices. We are interested here in a different problem, namely the dynamic modeling of the covariances of both the noise and the efficient returns.

A convenient method to handle the model described by Eq. (3), (4) is to assume that the disturbance terms are conditionally normal (Harvey 1991). Let \mathcal{F}_{t-1} be the σ -field generated by the observations of the log-price process up to time $t - 1$. The conditionally normal local-level model reads:

$$Y_t = X_t + \epsilon_t, \quad \epsilon_t | \mathcal{F}_{t-1} \sim \text{NID}(0, H_t) \quad (5)$$

$$X_t = X_{t-1} + \eta_t, \quad \eta_t | \mathcal{F}_{t-1} \sim \text{NID}(0, Q_t) \quad (6)$$

Conditionally on the information available at time $t - 1$, the distributions of ϵ_t and η_t are thus normal, with known covariance matrices H_t and Q_t . The Kalman filter can be applied to compute the likelihood in the usual prediction error form. As discussed by Harvey (1991), conditionally normal state-space models allow to “inject” nonlinear dynamics in the model while still preserving the possibility of applying the standard Kalman filter. In our empirical framework, this is a substantial advantage, given that the Kalman filter easily handles missing values and allows to reconstruct the efficient price based on all the available observations.

We now need to specify the law of motion of H_t and Q_t . Score-driven models (Creal et al. 2013, Harvey 2013) are a general class of observation-driven models (Cox 1981). In score-driven models, parameters are updated based on the score of the conditional density. The GARCH model of Bollerslev (1986), the EGARCH model of Nelson (1991) and the ACD model of Engle and Russell (1998) are examples of models that can be recovered in this general framework. In addition, score-driven models satisfy an information theoretic criterion, as shown by Blasques et al. (2015). Motivated by the flexibility of score-driven models, we choose to model the dynamics of H_t and Q_t based on the score of the conditional density.

Similarly to the DCC model of Engle (2002), we decompose the covariance matrix Q_t as:

$$Q_t = D_t R_t D_t \quad (7)$$

where $D_t = \text{diag}(Q_t)^{1/2}$ is a diagonal matrix of standard deviations and R_t is a correlation matrix. For parsimony, we assume that the covariance matrix H_t of the noise is diagonal. Such assumption can be relaxed at the expense of increasing considerably the number of time-varying parameters. However, as pointed out by Corsi et al. (2015), in the static case the off-diagonal elements of H_t are found to be close to zero, and thus they assume a diagonal matrix. A similar assumption is made by Shephard and Xiu (2017).

Let f_t denote the vector of time-varying parameters. We write:

$$f_t = \begin{pmatrix} \log(\text{diag}(H_t)) \\ \log(\text{diag}(D_t^2)) \\ \phi_t \end{pmatrix} \quad (8)$$

where ϕ_t is a $q \times 1$ vector depending on the parameterization of the correlation matrix R_t . The latter will be discussed in Section (2.3). The number of components of f_t is thus $k = 2n + q$.

The update rule in score-driven models is given by:

$$f_{t+1} = \omega + A s_t + B f_t \quad (9)$$

where s_t is the scaled score vector:

$$s_t = (\mathcal{I}_{t|t-1})^{-1} \nabla_t, \quad \nabla_t = \left[\frac{\partial \log p(Y_t | f_t, \mathcal{F}_{t-1}, \Theta)}{\partial f_t'} \right]', \quad \mathcal{I}_{t|t-1} = E[\nabla_t \nabla_t'] \quad (10)$$

The $k \times 1$ vector ω and the $k \times k$ matrices A, B are included in the vector Θ of static parameters to be estimated. Let $n_t \leq n$ denote the number of assets traded at time t . The conditional log-likelihood is given by:

$$\log p(Y_t | f_t, \mathcal{F}_{t-1}, \Theta) = \frac{n_t}{2} \log(2\pi) - \frac{1}{2} (\log |F_t| + v_t' F_t^{-1} v_t) \quad (11)$$

where v_t and F_t are defined in Section (2) of the online appendix.

As shown by Delle Monache et al. (2019), the score $\nabla_t = \partial \log p(Y_t | f_t, \mathcal{F}_{t-1}, \Theta) / \partial f_t$ and the Fisher information matrix $\mathcal{I}_{t|t-1} = E_{t-1}[\nabla_t \nabla_t']$ can be computed as:

$$\nabla_t = -\frac{1}{2} \left[\dot{F}_t' (\mathbb{I}_{n_t} \otimes F_t^{-1}) \text{vec}(\mathbb{I}_{n_t} - v_t v_t' F_t^{-1}) + 2 \dot{v}_t' F_t^{-1} v_t \right] \quad (12)$$

$$\mathcal{I}_{t|t-1} = \frac{1}{2} \left[\dot{F}_t' (F_t^{-1} \otimes F_t^{-1}) \dot{F}_t + 2 \dot{v}_t' F_t^{-1} \dot{v}_t \right] \quad (13)$$

Together with v_t and F_t , the computation of ∇_t and $\mathcal{I}_{t|t-1}$ requires \dot{v}_t and \dot{F}_t , which denote derivatives of v_t and F_t with respect to the time-varying parameter vector f_t . They can be computed through the expressions reported in Section (2) of the online appendix.

By running in parallel the Kalman filter and the filter in Eq. (9), we update the time-varying parameters and compute the conditional log-likelihood in Eq. (11). The static parameters Θ are estimated by optimizing the log-likelihood function through standard quasi-Newton algorithms:

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmax}} \sum_{t=1}^T \log p(Y_t | f_t, \mathcal{F}_{t-1}, \Theta) \quad (14)$$

Restrictions on the structure of A, B are discussed in the empirical application in Section (4). In Section (1.1) of the online appendix, we discuss an alternative update scheme which improves over the above method in the presence of jumps. The robust filter is simply obtained by replacing the Gaussian scores appearing in the update of the efficient price and of the covariances with the scores of a Student- t density.

2.3 Parameterization of R_t

To have a full model specification, we need to choose a parameterization for the correlation matrix R_t . We restrict our attention to parameterizations that guarantee a positive-definite correlation matrix. One possibility is to use hyperspherical coordinates, as in Creal et al. (2011). Another possibility is given by the equicorrelation parameterization of Engle and Kelly (2012).

In the first case, we write the correlation matrix as $R_t = Z_t' Z_t$. The matrix Z_t has the form:

$$Z_t = \begin{pmatrix} 1 & c_{12} & c_{13} & \dots & c_{1n} \\ 0 & s_{12} & c_{23}s_{13} & \dots & c_{2n}s_{1n} \\ 0 & 0 & s_{23}s_{13} & \dots & c_{3n}s_{2n}s_{1n} \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & \prod_{k=1}^{n-1} s_{kn} \end{pmatrix} \quad (15)$$

where $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$. Note that the time index has been omitted for ease of notation. The i -th column of Z_t contains the hyperspherical coordinates of a vector of unit norm in an i -th dimensional subspace of \mathbb{R}^n , which is parameterized by $i - 1$ angles. We have therefore $n(n - 1)/2$ angles, equal to the number of correlations in R_t . We set $\phi_t = [\theta_{12,t}, \theta_{13,t}, \dots, \theta_{n-1,n,t}]'$ in Eq. (8), so that $q = n(n - 1)/2$. The number of time-varying parameters is thus $k = 2n + n(n - 1)/2$.

In the equicorrelation parameterization, the correlation matrix R_t is written as:

$$R_t = (1 - \rho_t)\mathbb{I}_n + \rho_t\mathcal{J}_n \quad (16)$$

where \mathcal{J}_n denotes an $n \times n$ matrix of ones. The correlation matrix R_t is positive-definite if and only if the parameter ρ_t satisfies the constraint $-1/(n - 1) \leq \rho_t \leq 1$. One possibility to guarantee this constraint is to write:

$$\rho_t = \frac{1}{2} \left[\left(1 - \frac{1}{n - 1}\right) + \left(1 + \frac{1}{n - 1}\right) \tanh(\theta_t) \right] \quad (17)$$

as in Koopman et al. (2018). We set $\phi_t = \theta_t$, so that $q = 1$. The number of time-varying parameters is thus $k = 2n + 1$. The equicorrelation parameterization is very parsimonious, as it assumes the same correlation for all couples of assets. Notwithstanding, it has been proven to be effective in several empirical problems (see discussions in Engle and Kelly 2012).

3 Monte-Carlo analysis

3.1 Finite-sample properties

We first study the finite sample properties of the maximum likelihood estimator through Monte-Carlo simulations. We set $n = 10$, the same number of assets in our empirical application in Section (4). The number of time-varying parameters is thus $k = 2n + n(n - 1)/2 = 65$ in the parameterization with hyperspherical coordinates and $k = 2n + 1 = 21$ in the equicorrelation parameterization. In the first experiment, we assume that the static parameters ω , A , B have the following structure:

$$\omega = \begin{pmatrix} \omega^h \\ \omega^d \\ \omega^r \end{pmatrix}, \quad A = \text{diag} \begin{pmatrix} A^h \\ A^d \\ A^r \end{pmatrix}, \quad B = \text{diag} \begin{pmatrix} B^h \\ B^d \\ B^r \end{pmatrix} \quad (18)$$

where ω^h , A^h , B^h are $n \times 1$ vectors driving the variances of the noise, ω^d , A^d , B^d are $n \times 1$ vectors driving the variances of the efficient returns and ω^r , A^r , B^r are $q \times 1$ vectors driving the correlations. We further constraint the elements in each vector to be equal among each other. The number of static parameters is therefore equal to nine. In particular, we set: $\omega^h = -0.0461\iota_n$, $\omega^d = -0.0461\iota_n$, $\omega^r = 0.0185\iota_q$, $A^h = A^d = 0.04\iota_n$, $A^r = 0.04\iota_q$, $B^h = B^d = 0.98\iota_n$, $B^r = 0.98\iota_q$. The values of ω^h , ω^d are set in such a way that the signal-to-noise ratio (defined as the ratio between the unconditional standard deviation of the efficient returns and that of the measurement error) is similar to that found on empirical data, which is on average equal to one (see Table 5).

After simulating the log-prices, we randomly censor the observations to mimic asynchronous trading. In real financial markets, the missing value probability changes throughout the day. It has an inverse U-shape pattern, being lower at the opening and closing hours of the trading day. We use a scheme similar to that of Bucchieri et al. (2019b) to generate missing values. The missing value probability is written as:

$$p_t = \theta_t \xi_t \quad (19)$$

where θ_t is a common seasonal pattern and ξ_t is a stochastic component. The seasonal pattern is given by:

$$\theta_t = \bar{\theta} [b - a(t^2 - t)] \quad (20)$$

where $a = 0.45$, $b = 0.34$ and $\bar{\theta}$ is a normalization ensuring that $1/T \sum_0^T \theta_t = 1$. The values of a and b are obtained by fitting θ_t to the average missing value probability pattern found in the dataset used in our empirical application. Let $F(x)$ denote the cumulative distribution function of a standard normal random variable and let $\{\zeta_t\}$ be an i.i.d. sequence of standard normal random variables. The stochastic component ξ_t is generated as:

$$\xi_t = F(u_t) \quad (21)$$

$$u_t = u_{t-1} + (F^{-1}(\lambda) - u_{t-1})/T + \sigma_u \sqrt{T} \zeta_t \quad (22)$$

where $\sigma_u = 0.25$. The value of λ determines the unconditional mean of the probability of missing values. In our study, we change λ to generate scenarios with different levels of sparsity.

To set the initial values of the time-varying parameters, we estimate a local-level model with constant parameters in the subsample comprising the first 100 observations. This can be done through the EM algorithm, as described by Corsi et al. (2015). We then set f_1 based on Eq. (8), with H_t , D_t , R_t equal to the EM estimates. We simulate $N = 1000$ time-series of $T = 2000$ observations and consider four scenarios characterized by $\lambda = 0, 0.3, 0.5, 0.8$. Table (1) reports the summary statistics of the maximum likelihood estimates in the parameterization with hyperspherical coordinates. Table (2) reports analogous statistics obtained with the equicorrelation parameterization. The results show that, in both parameterizations, the maximum likelihood estimates concentrate around the true parameters. Not surprisingly, missing values lead to an increase of the variance of the estimator. However, even in the highly asynchronous scenario with $\lambda = 0.8$, relative errors remain small for all the parameters in ω , A , B .

In the second experiment, we assume that ω , A , B have the following structure:

$$\omega = \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}, \quad A = \text{diag} \begin{pmatrix} A^h \\ A^d \\ A^r \end{pmatrix}, \quad B = \text{diag} \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \quad (23)$$

Under these random-walk type restrictions, the time-varying parameters do not mean revert to an unconditional level. As shown in Section (4.2), such restriction provides a better description of the dynamics of intraday covariances and it is thus employed to estimate the local-level model on empirical data. We set $A^h = A^d = 0.01\iota_n$, $A^r = 0.01\iota_q$ and report in Table (3) the summary statistics of the maximum likelihood

λ	Mean	Std	Mean	Std	Mean	Std
	$\omega^h = -0.0461\iota_n$		$\omega^d = -0.0461\iota_n$		$\omega^r = 0.0185\iota_q$	
0.0	-0.0461	0.0036	-0.0461	0.0026	0.0185	0.0005
0.3	-0.0484	0.0126	-0.0486	0.0101	0.0203	0.0033
0.5	-0.0492	0.0169	-0.0492	0.0145	0.0215	0.0051
0.8	-0.0477	0.0204	-0.0480	0.0210	0.0219	0.0153
	$A^h = 0.04\iota_n$		$A^d = 0.04\iota_n$		$A^r = 0.04\iota_q$	
0.0	0.0398	0.0017	0.0399	0.0013	0.0400	0.0006
0.3	0.0411	0.0053	0.0388	0.0057	0.0403	0.0032
0.5	0.0414	0.0070	0.0387	0.0075	0.0420	0.0047
0.8	0.0386	0.0118	0.0385	0.0111	0.0426	0.0083
	$B^h = 0.98\iota_n$		$B^d = 0.98\iota_n$		$B^r = 0.98\iota_q$	
0.0	0.9800	0.0016	0.9800	0.0011	0.9800	0.0005
0.3	0.9781	0.0055	0.9792	0.0042	0.9783	0.0032
0.5	0.9779	0.0070	0.9790	0.0072	0.9777	0.0040
0.8	0.9791	0.0107	0.9797	0.0093	0.9750	0.0143

Table 1: Mean and standard deviations of maximum likelihood estimates of the local-level model in the parameterization with hyperspherical coordinates under the restrictions in Eq. (18).

λ	Mean	Std	Mean	Std	Mean	Std
	$\omega^h = -0.0461\iota_n$		$\omega^d = -0.0461\iota_n$		$\omega^r = 0.0185\iota_q$	
0.0	-0.0469	0.0081	-0.0468	0.0073	0.0196	0.0059
0.3	-0.0475	0.0109	-0.0472	0.0102	0.0205	0.0094
0.5	-0.0485	0.0144	-0.0483	0.0137	0.0216	0.0125
0.8	-0.0472	0.0228	-0.0479	0.0215	0.0224	0.0220
	$A^h = 0.04\iota_n$		$A^d = 0.04\iota_n$		$A^r = 0.04\iota_q$	
0.0	0.0400	0.0035	0.0398	0.0035	0.0393	0.0057
0.3	0.0409	0.0048	0.0400	0.0049	0.0375	0.0076
0.5	0.0414	0.0061	0.0398	0.0059	0.0368	0.0094
0.8	0.0386	0.0111	0.0378	0.0100	0.0366	0.0142
	$B^h = 0.98\iota_n$		$B^d = 0.98\iota_n$		$B^r = 0.98\iota_q$	
0.0	0.9796	0.0035	0.9797	0.0031	0.9788	0.0063
0.3	0.9794	0.0047	0.9795	0.0044	0.9775	0.0101
0.5	0.9789	0.0063	0.9790	0.0059	0.9760	0.0138
0.8	0.9793	0.0099	0.9791	0.0093	0.9751	0.0148

Table 2: Mean and standard deviations of maximum likelihood estimates of the local-level model in the parameterization based on the equicorrelation matrix under the restrictions in Eq. (18).

estimates. As in the previous experiment, the parameter estimates concentrate around the true values and the variances increase with the number of missing values. Overall, we conclude that maximum likelihood provides satisfactory parameter estimates in all the considered scenarios.

λ	Mean	Std	Mean	Std	Mean	Std
<u>HYPERSPHERICAL</u>						
	$A^b = 0.01\iota_n$		$A^d = 0.01\iota_n$		$A^r = 0.01\iota_q$	
0.0	0.0100	0.0009	0.0100	0.0005	0.0100	0.0002
0.3	0.0085	0.0040	0.0106	0.0039	0.0099	0.0021
0.5	0.0082	0.0059	0.0120	0.0053	0.0113	0.0027
0.8	0.0080	0.0110	0.0111	0.0102	0.0129	0.0071
<u>EQUICORRELATION</u>						
	$A^b = 0.01\iota_n$		$A^d = 0.01\iota_n$		$A^r = 0.01\iota_q$	
0.0	0.0099	0.0013	0.0099	0.0012	0.0098	0.0025
0.3	0.0117	0.0019	0.0116	0.0018	0.0103	0.0032
0.5	0.0120	0.0028	0.0122	0.0027	0.0109	0.0043
0.8	0.0129	0.0076	0.0122	0.0065	0.0115	0.0090

Table 3: Mean and standard deviations of maximum likelihood estimates of the local-level model under the restrictions in Eq. (23).

3.2 Assessing the effect of noise and asynchronicity

In this section we aim to assess the effect of measurement errors and asynchronous trading on common dynamic correlation models and to show the advantages provided by the proposed modeling strategy. High-frequency log-prices are typically synchronized before being analyzed. The most popular synchronization scheme is previous-tick interpolation, consisting in filling the missing values with the previous available price. Such procedure (and other similar schemes) leads to a downward bias of correlations toward zero. In the literature on realized covariance estimation, the latter is known as the “Epps effect” (Epps 1979). As it is shown here, a similar bias arises when estimating dynamic correlations. The presence of measurement errors is an additional source of bias for the estimated correlations. This is not surprising, as attenuation biases due to measurement errors occur in several econometric and statistics problems, e.g. in *error-in-variables* models.

For simplicity, we consider a bivariate model generated as follows:

$$Y_t = X_t + \epsilon_t, \quad \epsilon_t \sim N(0, H) \quad (24)$$

$$X_t = X_{t-1} + \eta_t, \quad \eta_t \sim N(0, DR_tD) \quad (25)$$

where $H = h\mathbb{I}_2$, $D = d\mathbb{I}_2$ and:

$$R_t = \begin{pmatrix} 1 & \rho_t \\ \rho_t & 1 \end{pmatrix} \quad (26)$$

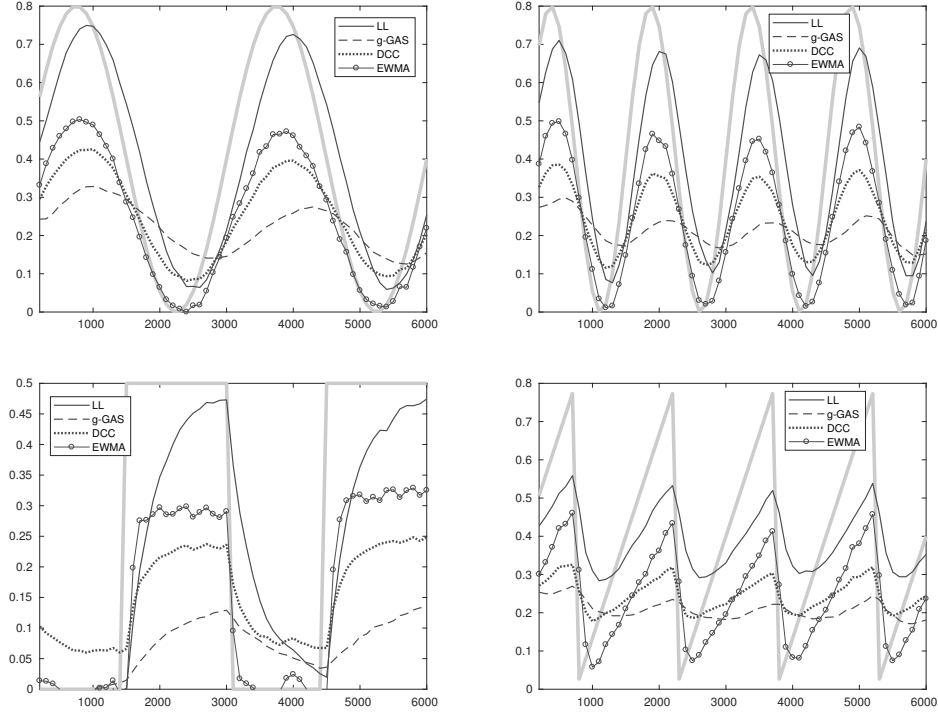


Figure 1: Estimates of LL (solid line), g -GAS (dashed line), DCC (dotted line) and EWMA (solid circle line) averaged over all simulations in the scenario $(\delta = 2, \lambda = 0.5)$. We also show the simulated correlation (grey line).

The correlation coefficient ρ_t obeys the following dynamic patterns:

1. Sine $\rho_t = b_s \sin(c_s \pi t)$
2. Fast Sine $\rho_t = b_f \sin(c_f \pi t)$
3. Step $\rho_t = \alpha_s - \beta_s(\mathbb{1}_{t < T/4} + \mathbb{1}_{T/2 < t < 3T/4})$
4. Ramp $\rho_t = \frac{1}{b_r} \bmod(t + a_r, c_r)$
5. Model $\rho_t = \exp(h_t) / [1 + \exp(h_t)]$

where h_t is an AR(1) process:

$$h_{t+1} = c_m + b_m h_t + a_m \phi_t, \quad \phi_t \sim N(0, 1)$$

and T is the number of observations. The parameters governing the dynamics of ρ_t are chosen as: $b_s = b_f = 0.4$, $c_s = 4/T$, $c_f = 8/T$, $\alpha_s = 0.25$, $\beta_s = 0.5$, $b_r = 5/16T$, $a_r = T/8$, $c_r = T/4$, $b_m = 0.99$, $c_m = -0.4(1 - b_m)$, $a_m = 0.01$. The variance h of the measurement errors is constant and equal to 0.1 for all the simulated patterns. The variance d^2 of the latent efficient returns is computed based on the signal-to-noise ratio $\delta = d/\sqrt{h}$. We consider three different scenarios with $\delta = 0.5$ (low signal), 1 (moderate signal), 2 (high signal). As it is evident in Section (4), these values are close to those estimated on real data. Once observations have been generated, we apply a censoring scheme similar to that used in the previous

analysis. In particular, we consider three scenarios with $\lambda = 0, 0.3, 0.5$, where λ denotes the unconditional mean of the probability of missing values.

We generate $N = 250$ simulations of $T = 6000$ observations for each dynamic pattern and for each scenario. As benchmarks, we consider:

- The DCC model of Engle (2002)
- An EWMA scheme given by: $\hat{Q}_{t+1} = \gamma \hat{Q}_t + (1 - \gamma)r_t r'_t$, $\gamma = 0.96$
- The g -GAS model of Creal et al. (2011)

The g -GAS is an observed return model for dynamic correlations based on a Gaussian conditional density. The DCC and the EWMA are estimated by synchronizing the data through previous-tick interpolation. Depending on the value of λ , previous-tick leads to a large number of zero returns. In order to attenuate the effect of the noise and of the zero returns, we sample the observations at a lower frequency. This procedure is commonly employed, for instance, when computing the realized covariance. Indeed, at lower frequencies, returns are less affected by microstructure noise and by staleness effects due to the absence of trading. A natural consequence of subsampling is that a significant portion of the data is neglected. However, subsampling greatly improves the correlation estimates of the DCC and the EWMA, which otherwise would be extremely biased toward zeros. In each scenario, the new sampling frequency is set as the one minimizing the MSE in a pre-simulation study with $N = 20$ simulations. The selected values range in the interval between 3 and 6 time units.

The g -GAS model is instead estimated using the missing value approach of Lucas et al. (2016), consisting in computing the score with respect to the marginal density of the available data. As the conditional density is a multivariate Gaussian, the marginals are readily available in closed form.

Both in-sample and out-of-sample losses are computed. To this end, the sample is divided into two subsamples containing the same number of observations. Thus, there are $T_{\text{sub}} = 3000$ observations available for the in-sample estimation of the local-level model and of the g -GAS. In the case of the DCC, the number of available in-sample observations depends on the sampling frequency. The second subsample is used to compute the out-of-sample losses. The parameterization used in the local-level model and in the g -GAS is the one based on hyperspherical coordinates.

Figure (1) shows, in the scenario ($\delta = 2, \lambda = 0.5$), the average of the correlations estimated by the four models. Even by sampling returns at lower frequencies, DCC and EWMA estimates are biased toward zero. This is mainly a consequence of asynchronicity, which is high in this scenario. The g -GAS is not affected by zero-returns because, similarly to the local-level model, asynchronicity is treated as a missing value problem. However, it is not robust to measurement errors and, being a model for the observed returns, it is also subject to data reduction when prices are not synchronized (see Section 3.3 for further details on this effect). As a consequence, even the g -GAS estimates are biased.

Figure (2) shows the average correlations in the scenario ($\delta = 1, \lambda = 0$). The DCC and the EWMA have again a downward bias. In contrast to the previous scenario, the bias is now entirely due to measurement errors because there are no missing values in this scenario. Similarly, the g -GAS significantly underestimates

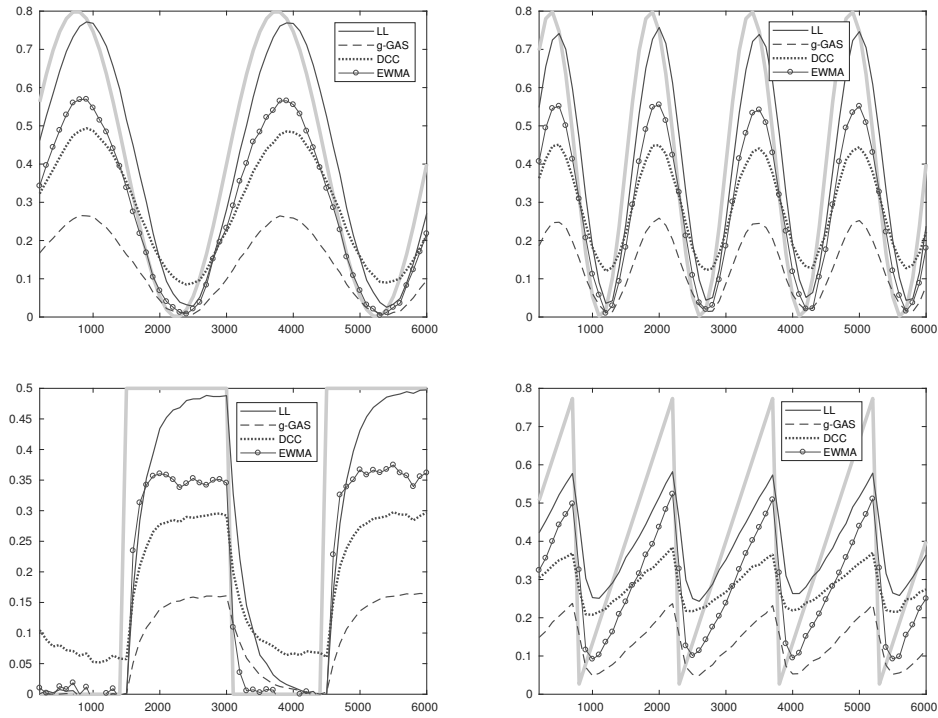


Figure 2: Estimates of LL (solid line), g -GAS (dashed line), DCC (dotted line) and EWMA (solid circle line) averaged over all simulations in the scenario ($\delta = 1$, $\lambda = 0$). We also show the simulated correlation (grey line).

the correlations. The bias is more pronounced in the g -GAS, since the effect of measurement errors in the DCC and EWMA is mitigated by the use of lower aggregation frequencies. Thus, in a high-frequency setting, the use of standard dynamic correlation models leads to data reduction and underestimates (absolute value of) correlations.

Table (4) reports in-sample and out-of-sample MSE of the four models in all the simulated scenarios. The local-level significantly outperforms the benchmark models. In particular, the relative differences increase in patterns “step”, “ramp” and “model”, characterized by frequent changes or stochasticity. Indeed, by exploiting all the available observations, the local-level can track the variation of these patterns with a higher efficiency.

Model	Sine	Fast	Step	Ramp	Model	Sine	Fast	Step	Ramp	Model	Sine	Fast	Step	Ramp	Model
	$\delta = 0.5$					$\delta = 1$					$\delta = 2$				
	$\lambda = 0$														
LL	0.0686	0.0745	0.0330	0.0532	0.0491	0.0183	0.0306	0.0198	0.0446	0.0298	0.0133	0.0213	0.0150	0.0297	0.0218
g -GAS	0.1578	0.1622	0.0954	0.1475	0.1858	0.1105	0.1181	0.0706	0.1110	0.1416	0.0679	0.0741	0.0450	0.0725	0.0903
DCC	0.0811	0.0936	0.0618	0.0828	0.0989	0.0443	0.0564	0.0423	0.0569	0.0636	0.0273	0.0381	0.0308	0.0430	0.0415
EWMA	0.0778	0.0830	0.0619	0.0866	0.0930	0.0461	0.0523	0.0442	0.0597	0.0603	0.0311	0.0379	0.0350	0.0471	0.0421
LL	0.0699	0.0761	0.0442	0.0534	0.0770	0.0210	0.0335	0.0304	0.0449	0.0359	0.0141	0.0222	0.0228	0.0299	0.0232
g -GAS	0.1621	0.1660	0.0967	0.1510	0.1867	0.1144	0.1220	0.0715	0.1126	0.1455	0.0689	0.0741	0.0473	0.0721	0.0904
DCC	0.0852	0.0960	0.0637	0.0834	0.1020	0.0471	0.0588	0.0441	0.0582	0.0699	0.0295	0.0394	0.0340	0.0434	0.0470
EWMA	0.0802	0.0846	0.0616	0.0854	0.0913	0.0478	0.0546	0.0450	0.0608	0.0600	0.0321	0.0396	0.0378	0.0478	0.0403
	$\lambda = 0.3$														
LL	0.0552	0.0728	0.0383	0.0536	0.0496	0.0290	0.0461	0.0276	0.0509	0.0378	0.0210	0.0337	0.0226	0.0418	0.0295
g -GAS	0.1600	0.1693	0.1074	0.1543	0.2021	0.1178	0.1326	0.0882	0.1218	0.1547	0.0784	0.0964	0.0622	0.0877	0.1007
DCC	0.0973	0.1103	0.0697	0.0966	0.1202	0.0625	0.0733	0.0530	0.0700	0.0846	0.0418	0.0529	0.0413	0.0546	0.0593
EWMA	0.0940	0.0986	0.0696	0.1010	0.1139	0.0620	0.0667	0.0539	0.0729	0.0799	0.0439	0.0492	0.0438	0.0573	0.0563
LL	0.0628	0.0765	0.0575	0.0541	0.0729	0.0341	0.0517	0.0429	0.0528	0.0464	0.0237	0.0363	0.0345	0.0433	0.0331
g -GAS	0.1741	0.1841	0.1093	0.1614	0.2076	0.1354	0.1468	0.0884	0.1298	0.1549	0.0915	0.1049	0.0645	0.0938	0.1057
DCC	0.1036	0.1145	0.0711	0.0977	0.1272	0.0691	0.0773	0.0548	0.0727	0.0886	0.0473	0.0570	0.0426	0.0565	0.0707
EWMA	0.0987	0.1051	0.0696	0.1009	0.1170	0.0676	0.0713	0.0543	0.0768	0.0783	0.0484	0.0541	0.0436	0.0602	0.0579
	$\lambda = 0.5$														
LL	0.0620	0.0768	0.0506	0.0539	0.0566	0.0402	0.0637	0.0395	0.0532	0.0467	0.0308	0.0511	0.0324	0.0512	0.0389
g -GAS	0.1696	0.1754	0.1158	0.1567	0.2049	0.1301	0.1441	0.1067	0.1262	0.1597	0.0946	0.1107	0.0933	0.0914	0.1179
DCC	0.1153	0.1255	0.0773	0.1076	0.1335	0.0800	0.0931	0.0629	0.0845	0.1027	0.0573	0.0731	0.0530	0.0684	0.0821
EWMA	0.1100	0.1163	0.0787	0.1136	0.1330	0.0786	0.0846	0.0640	0.0873	0.0987	0.0582	0.0669	0.0543	0.0722	0.0796
LL	0.0735	0.0838	0.0710	0.0565	0.0828	0.0518	0.0746	0.0576	0.0580	0.0597	0.0385	0.0590	0.0483	0.0557	0.0501
g -GAS	0.1893	0.1943	0.1170	0.1686	0.2298	0.1536	0.1581	0.1096	0.1338	0.1746	0.1191	0.1237	0.0945	0.1002	0.1317
DCC	0.1230	0.1294	0.0786	0.1091	0.1536	0.0903	0.0989	0.0631	0.0868	0.1144	0.0684	0.0787	0.0527	0.0710	0.0950
EWMA	0.1178	0.1220	0.0769	0.1177	0.1427	0.0883	0.0922	0.0610	0.0914	0.1016	0.0683	0.0734	0.0514	0.0753	0.0804

Table 4: MSE of LL, g -GAS, DCC and EWMA estimates in the nine scenarios obtained by combining the three missing value scenarios ($\lambda = 0, 0.3, 0.5$) to each of the three signal-to-noise scenarios ($\delta = 0.5, 1, 2$). The first three lines of each scenario show the in-sample results, whereas the last three lines show the out-of-sample results.

3.3 Comparison with fat-tail return models

High-frequency data may exhibit extreme movements during flash crashes or in correspondence of macro-news announcements. It is therefore important to investigate the quality of the correlation estimates provided by the proposed approach in the presence of outliers. Creal et al. (2011) introduced a score-driven dynamic correlation model, named *t*-GAS, based on a multivariate Student-*t* distribution. The conditional density of the log-returns in the *t*-GAS is:

$$p(r_t|\Sigma_t, \nu) = \frac{\Gamma(\frac{\nu+n}{2})}{\Gamma(\frac{\nu}{2})[(\nu-2)\pi]^{n/2}|\Sigma_t|^{1/2}} \left(1 + \frac{r_t'\Sigma_t r_t}{\nu-2}\right)^{-(\nu+n)/2} \quad (27)$$

The covariances in Σ_t obey the usual score-driven update rule. Different parameterizations for Σ_t are possible, such as the one discussed in Section (2.3) based on hyperspherical coordinates.

In this simulation study, we use Eq. (27) as a data generating process for the log-returns. After computing the log-prices, the censoring scheme described in Section (3.1) is applied to mimic asynchronous trading. We then estimate the correlations using the *t*-GAS and the local-level model. The *t*-GAS is estimated using the missing-value approach of Lucas et al. (2016). As in the case of the *g*-GAS, the latter can be applied straightforwardly, since the marginals of a multivariate Student-*t* density are known in closed form (see e.g. Golam Kibria and Joarder 2017.)

The missing value approach avoids the introduction of artificial zero returns. However, it induces significant data reduction, as only consecutive trades can be used to compute the returns. For instance, if the *i*-th assets is traded at time *t* but not at time *t* − 1, the return $r_{i,t}$ is treated as a missing value and the information related to the price of the *i*-th asset at time *t* is lost. In contrast, at time *t* the local-level model exploits the observation of the price of the *i*-th asset to reconstruct the vector of the *n* efficient prices and to update the covariances.

To illustrate such effect, we simulate the log-prices through a *t*-GAS model with *n* = 10 and, in order to emphasize the effect of fat-tails, we set $\nu = 3$. Ten different scenarios are considered, characterized by an increasing level of sparsity ($\lambda = 0, 0.1, \dots, 0.9$). We perform *N* = 250 simulations of time-series of *T* = 6000 log-prices. As in the previous study, the first subsample of $T_{\text{sub}} = 3000$ log-prices is used to estimate the three models, whereas the second subsample is used to compute the out-of-sample forecasts. As a loss measure, we use the Frobenius distance, defined as:

$$\|\hat{\Sigma}_t - \Sigma_t\|_F = \sqrt{\text{Tr}[(\hat{\Sigma}_t - \Sigma_t)^2]} \quad (28)$$

The parameterization of R_t used in the two models is that based on hyperspherical coordinates, and the static parameters driving the dynamic elements in H_t , D_t and R_t have the restrictions in Eq. (18).

Figure (3) shows on the left the average out-of-sample losses of the two models. There is an evident trade-off between fat-tails and asynchronicity. Not surprisingly, if data are synchronized ($\lambda = 0$), the *t*-GAS performs significantly better than the Gaussian local-level. However, once missing values come into play, the latter tends to improve its relative performance. The average loss of the *t*-GAS quickly increases with λ . In particular, for $\lambda > 0.5$, the Gaussian local-level provides better out-of-sample forecasts. The




Figure 3: Left: average out-of-sample Frobenius losses, as a function of the average probability of missing values λ , of the t -GAS and the local-level model. The DGP is a t -GAS with $\nu = 3$. Right: average out-of-sample Frobenius losses of the same models as a function of the signal-to-noise ratio δ . The DGP is a t -GAS with $\nu = 3$ contaminated by a Student- t distributed measurement error with $\nu_{err} = 3$ and we set $\lambda = 0$.

rapid deterioration of the relative performance of the t -GAS in the presence of missing values is due to the aforementioned effect of data reduction.

We obtain similar results when adding a noise to the simulated log-prices. In a second experiment, we ignore asynchronicity, but contaminate the log-prices with a zero-mean Student- t distributed measurement error with $\nu_{err} = 3$. We consider different levels of the signal-to-noise ratio, from $\delta = 0.1$ to $\delta = 100$. As shown on the right in Figure (3), for $\delta < 50$, the average out-of-sample loss of the Gaussian local-level is significantly lower than that of the t -GAS. This is due to the fact that, in the local-level, we model the covariances of the latent returns instead of the observed returns. Of course, adding noise and asynchronicity together would lead to further increase the relative difference between the local-level model and the t -GAS. We provide additional Monte-Carlo evidences in Section (1.1.1) of the online appendix, where we show the results obtained with the alternative update scheme based on the scores of the Student- t density.

As it is evident from our empirical application, high-frequency data are characterized by high levels of sparsity ($\lambda > 0.8$) and are significantly noisy ($\delta \approx 1$). In light of the results of this section, we conclude that the use of the local-level model is more suited in a high-frequency setting, where these two effects are extremely important.

4 Empirical illustration

4.1 Dataset

Hereafter, to make the exposition clear, we denote by LL_h the local-level model with hyperspherical coordinates and by LL_e the local-level model with the parameterization based on equicorrelations. Our dataset contains 1-second transaction data of 10 most frequently traded NYSE assets in 2014. Data are related to

Symbol	Asset	Avg. duration (sec)	λ_{avg}	Avg. n. of trades	δ_{avg}
XOM	Exxon Mobil Corporation	5.434	0.816	4304	1.178
C	Citigroup Inc	6.135	0.836	3832	1.246
JPM	JPMorgan Chase	6.250	0.840	3743	0.999
HAL	Halliburton Company	6.369	0.843	3690	0.872
CVX	Chevron	6.579	0.848	3553	0.850
DIS	Walt Disney	6.622	0.849	3543	0.846
JNJ	Johnson & Johnson	6.666	0.850	3529	0.809
SLB	Schlumberger Limited	6.802	0.853	3454	0.613
DAL	Delta Air Lines	6.993	0.857	3348	0.766
WMT	Walmart	7.042	0.858	3325	0.698

Table 5: Summary statistics and average signal-to-noise ratio of the high-frequency transaction data employed in the empirical analysis.

transactions from 02-01-2014 to 31-12-2014, including a total of 251 business days. The exchange opens at 9.30 and closes at 16.00 local time, so that the number of seconds per day is $T = 23400$. We perform the standard procedures described by Barndorff-Nielsen et al. (2009) to clean the data. Table (5) shows the summary statistics of the ten assets, including the average duration in seconds between trades, the average probability of missing values λ_{avg} (computed as the average fraction of missing values per day) and the average number of trades per day. We also report the average signal-to-noise ratio δ_{avg} , computed from the matrices D_t , H_t estimated through the LL_h model. Specifically, δ_{avg} is defined as the ratio $D_{t,ii}/\sqrt{H_{t,ii}}$, $i = 1, \dots, 10$, averaged over all the timestamps.

The figures in the table provide a first evidence on the relevance of asynchronous trading and microstructure effects in high-frequency data. The average probability of missing values is greater than 80%, indicating large levels of sparsity even for the liquid assets included in this dataset. The average signal-to-noise is close to one, implying that also microstructure effects play a relevant role. As done in the simulation study in Section (3), we thus aim to show the advantages of the proposed method over standard techniques ignoring these effects.

We point out that the time-resolution of one second is the highest available in our dataset. If millisecond or microsecond data are available, one could assume a higher resolution to exploit all the observed prices. At those resolutions, data are extremely sparse, with a lot of timestamps where the prices of all the assets are missing. If the latter circumstance occurs, the update equations of the Kalman filter simplify considerably, as discussed in Section (2) of the online appendix. Even the computation of the score in Eq. (12) simplifies, since it becomes equal to zero. Therefore, while the number of recursions increase, most of them become extremely simple from a computational point of view.

Restriction type	23 Jan	11 Apr	27 May	6 Aug	4 Sep	16 Oct
Mean-reverting	-6.1237	-7.0196	-6.4519	-4.5173	-4.7355	-10.2073
Random-walk	-6.2291	-7.1361	-6.5344	-4.6845	-4.7890	-10.226

Table 6: AIC ($\times 10^5$) statistics computed on six randomly selected trading days with the mean-reverting restriction based on Eq. (18) and the random-walk restriction based on Eq. (23).

4.2 In-sample analysis

The vector ω and the matrices A , B in the dynamic equation (9) have dimensions $k \times 1$ and $k \times k$, respectively. For $n = 10$, k is equal to 21 in the equicorrelation parameterization, whereas it is equal to 65 in the parameterization based on hyperspherical coordinates. It is therefore necessary to impose some restrictions on the parameter space in order to estimate the model. We perform an AIC test to choose among different restrictions. Specifically, we compare the mean-reverting restriction in Eq. (18) with the random-walk type restriction in Eq. (23). Table (6) reports the AIC statistics of the LL_h model estimated on six randomly selected trading days with the two restrictions. The AIC provided by the random-walk type restriction is always lower than that of the mean-reverting restriction. As it is shown in the next sections, the intuitive reason behind this result is that intraday covariances have non-stationary dynamics which do not mean revert to an unconditional level. A similar result holds for the LL_e model. We therefore implement the random-walk type restriction when applying the local-level model in all the subsequent analyses.

Parameter	t -GAS	LL_h	LL_e
A^h	-	0.0323	0.0179
	-	(0.0137)	(0.0117)
A^d	0.1515	0.0226	0.0210
	(0.0373)	(0.0088)	(0.0093)
A^r	0.0232	0.0048	0.0057
	(0.0109)	(0.0018)	(0.0029)
ν	13.555	-	-
	(4.2204)	-	-

Table 7: Average of maximum likelihood estimates of parameters A^h , A^d , A^r , ν of t -GAS, LL_h and LL_e models. The averages are computed over the whole sample of 251 business days. Standard deviations are indicated in parenthesis.

We estimate the two local-level specifications on each business day of 2014 (the average estimation time on Matlab 2018 is 28.3 min on an i7-2600 CPU at 3.40GHz). In order to initialize time-varying parameters, we proceed as in Section (3.1), i.e. we estimate the static version of the model in a pre-sample including the first 10 minutes of the trading day. The starting values of the time-varying parameters are then set equal to the static estimates. We compare the covariance estimates obtained by the local-level model to the covariances of the t -GAS model. As in Creal et al. (2011), the correlations in the t -GAS are parameterized using the hyperspherical coordinates. In the presence of asynchronous data, the t -GAS can be estimated




Figure 4: Left: daily AIC of LL_h , t -GAS and DCC models. Right: AIC of LL_e normalized by the AIC of the LL_h model.

by previous-tick synchronization or using the missing value approach of Lucas et al. (2016). As discussed in Sections (3.2), (3.3), the first method generates a lot of zero returns, whereas the second implies data reduction, as a consequence of modeling returns rather than prices. We use the missing value approach, as zero returns dramatically jeopardize the estimation of the covariances in the t -GAS. The same random-walk type restrictions on static parameters are implemented for the t -GAS. Table (7) shows the average and the standard deviation of the estimated parameters of the two local-level specifications and of the t -GAS. We note that the degrees of freedom parameter ν of the t -GAS is close to 13, which would suggest that the conditional distribution of high-frequency returns is not fat-tailed. This is most likely due to the misspecification induced by microstructure noise, which affects the inference of the t -GAS parameters.

As an additional benchmark, we use the DCC model estimated through the previous-tick interpolation. In order to attenuate microstructure effects, we aggregate the data at the sampling frequency of one minute. At frequencies higher than one minute the estimated correlations are even more biased, while at lower frequencies the inference becomes extremely inefficient due to data reduction.

Figure (4) shows on the left the daily AIC provided by the LL_h , the t -GAS and the DCC. The AIC of the LL_h is significantly lower than that of the other two models. These differences are mainly due to data reduction, and provide a quantitative assessment of the relevance of this effect. On the right, we plot the AIC of the LL_e model normalized by the AIC of the LL_h . We note that the relative difference between LL_h and LL_e is small, meaning that the equicorrelation structure does not lead to an excessive deterioration of the in-sample fit of the model.

Figure (5) shows the estimated intraday volatilities and correlations, averaged over the whole sample of 251 business days. The volatilities are also averaged over the 10 assets, whereas the correlations are averaged over all the couples of assets. The intraday pattern of the volatilities provided by the t -GAS is above that of the two local-level models. The reason is that the estimated volatilities in the t -GAS include the microstructure noise volatility. Instead, the local-level model sets the efficient return volatilities in D_t apart from the noise volatilities in H_t . In the t -GAS, microstructure effects thus lead to an upward bias

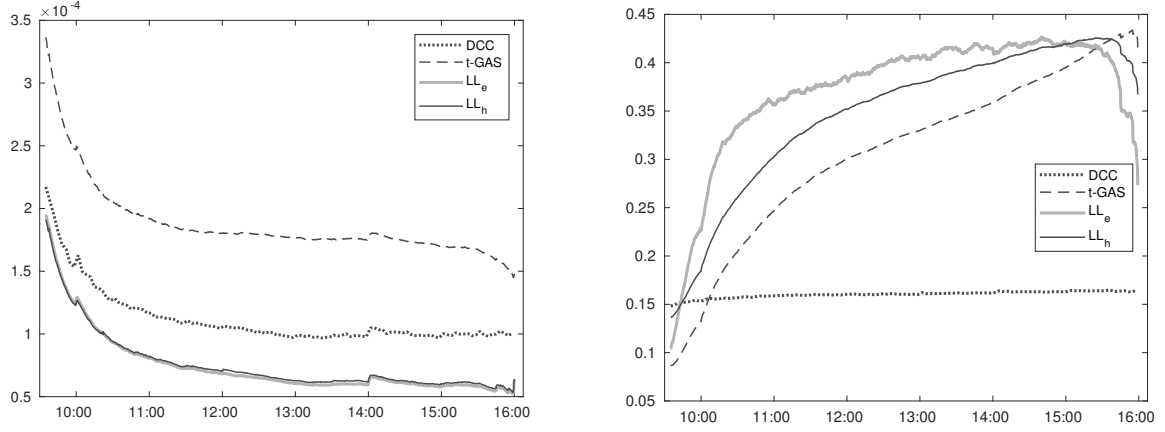


Figure 5: Intraday average volatilities (left) and correlations (right) of LL_h, LL_e, *t*-GAS and DCC models. Volatilities are average over all the 10 assets, whereas correlations are averaged over all couples of assets.

in the estimated variances. The average volatility estimated by the DCC is closer to that of the two local-level models. Microstructure effects in the DCC are indeed attenuated by the lower aggregation frequency. However, previous-tick interpolation leads to correlations which are extremely biased toward zero. The average correlations of the *t*-GAS are significantly less biased compared to DCC correlations, confirming the ability of the missing value approach to avoid the distortions due to the introduction of artificial zero returns. Nevertheless, they suffer from the additional bias due to microstructure effects. Overall, the in-sample analysis provides clear evidence that the proposed approach based on the local-level specification is more suited when modeling intraday covariances.

The intraday patterns of the local-level model with hyperspherical coordinates are close to those provided by the equicorrelation parameterization, in accordance with the result in Figure (4). Larger deviations are observed in the correlations, in particular in the first part of the trading day, where the average equicorrelation increases at a slightly higher rate, and during the last few minutes, where the decrease is more pronounced.

4.3 Intraday dynamics of the covariances

We now study in more detail the estimates provided by the local-level model with the aim to extract meaningful information on the behavior of the intraday covariances. To gain insights on the level of heterogeneity of correlations, we perform the analysis using the parameterization based on hyperspherical coordinates. We first examine the variation of the intraday patterns over time. To this purpose, we take averages of the estimated time-varying parameters over assets or couples of assets. For $t = 1, \dots, 23400$ and $j = 1, \dots, 251$, we compute:

$$\begin{aligned} \tilde{d}_t^j &= \frac{1}{n} \sum_{i=1}^n D_t^j(i, i), & \tilde{h}_t^j &= \frac{1}{n} \sum_{i=1}^n \sqrt{H_t^j(i, i)} \\ \tilde{\delta}_t^j &= \frac{1}{n} \sum_{i=1}^n \frac{D_t^j(i, i)}{\sqrt{H_t^j(i, i)}}, & \tilde{\rho}_t^j &= \frac{2}{n(n-1)} \sum_{p>q} R_t^j(p, q) \end{aligned}$$

where $n = 10$. Similarly, to examine the variation of the intraday patterns among different assets, we take averages of time-varying parameters over time. For $t = 1, \dots, 23400$ and $i, p = 1, \dots, 10, q < p$ we compute:

$$\begin{aligned}\bar{d}_t^i &= \frac{1}{N} \sum_{j=1}^N D_t^j(i, i), & \bar{h}_t^i &= \frac{1}{N} \sum_{j=1}^N \sqrt{H_t^j(i, i)}, \\ \bar{\delta}_t^i &= \frac{1}{N} \sum_{j=1}^N \frac{D_t^j(i, i)}{\sqrt{H_t^j(i, i)}}, & \bar{\rho}_t^{p,q} &= \frac{1}{N} \sum_{j=1}^N R_t^j(p, q)\end{aligned}$$

where $N = 251$. The average intraday pattern can be computed by averaging the variables labeled with a “tilde” over $j = 1, \dots, 251$ or, equivalently, by averaging the variables labeled with a bar over $i, p = 1, \dots, n$ and $q < p$. For each timestamp $t = 1, \dots, 23400$, Figures (6) and (7) show 10%, 90% quantiles of the variables denoted by “tilde” and 10%, 90% quantiles of the variables denoted by bars, respectively. We also report in each figure the average intraday pattern of the time-varying parameters.

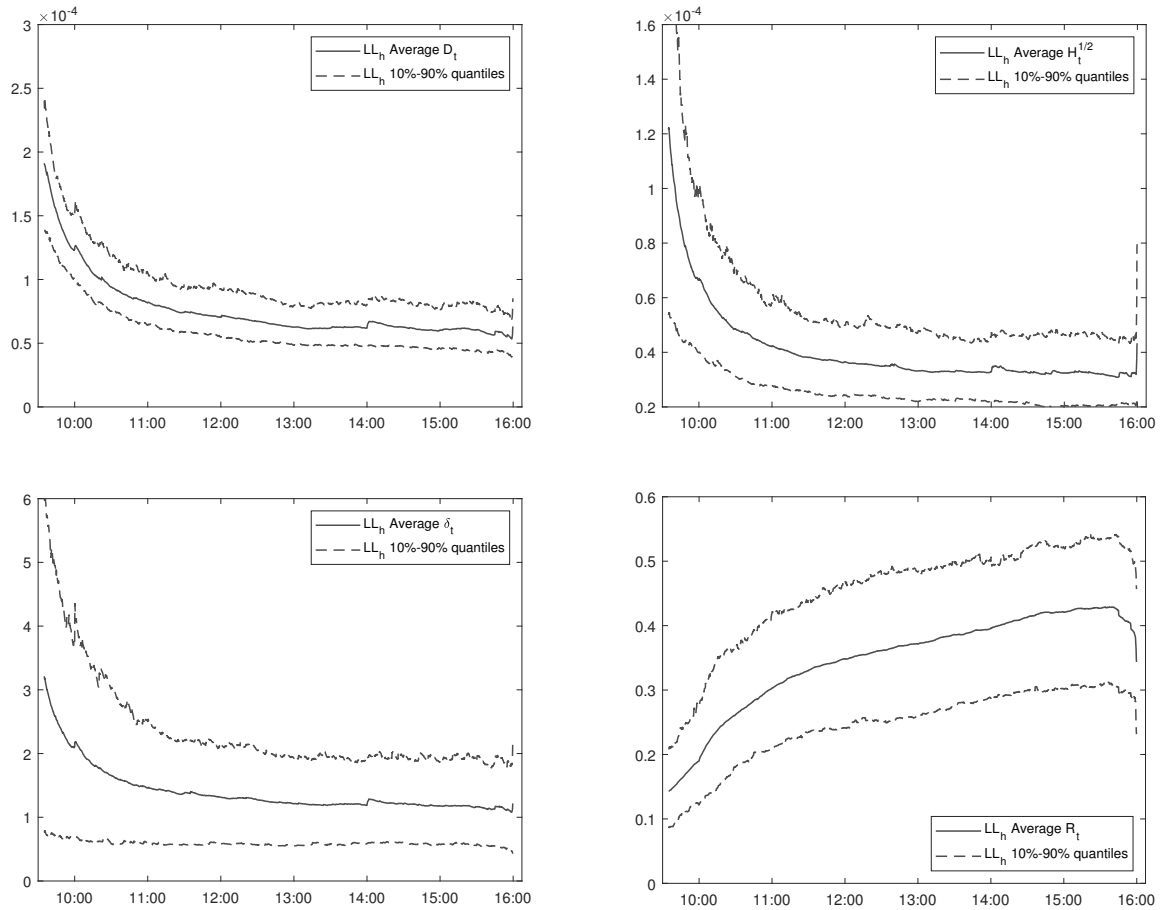


Figure 6: Average intraday patterns of the estimated LL_h time-varying parameters D_t , $H_t^{1/2}$, δ_t and R_t . For each $t = 1, \dots, 23400$, we report 10% and 90% quantiles of the distribution (over days) of averages of D_t , $H_t^{1/2}$, δ_t and R_t (over assets or couples of assets).

The efficient return volatilities in D_t are large at the beginning of the trading day and gradually decline in the remaining part of the day. The average intraday pattern of the microstructure noise volatility has

two regimes: a steep decline from 9:30 to 10:00 and a slow decline from 10:00 until the end of the trading day. This dynamic behavior is close to the typical intraday pattern of bid-ask spreads (see e.g. McInish and Wood 1992), confirming that the local-level model consistently sets apart the efficient log-price process from microstructure effects. The signal-to-noise ratio exhibits a similar pattern: it is larger at the beginning of the day ($\delta \approx 3$), whereas it is lower in the remaining part of the day ($\delta \approx 1$), indicating that prices are significantly affected by microstructure effects during a large portion of the trading day.

Correlations exhibit an interesting increasing pattern. They are low at the beginning of the trading day, implying that the dynamics of log-prices are largely affected by idiosyncratic risk in that part of the day. We then observe a steep increase until 11:00, which is associated to the fast decline of the volatilities. After 11:00, correlations increase, though at a lower rate, until 15:45. Thereafter, we observe a steep decline of all the correlations.

Figure (8) shows the average intraday pattern of the first five eigenvalues of the correlation matrix R_t . We note that the explanatory power of the first eigenvalue, associated to the market factor, progressively increases during the trading day. This implies that, while at the beginning of the day the dynamics are dominated by idiosyncratic risk, a systematic component emerges in the second part of the day. This component is associated to market risk, as all the remaining eigenvalues decrease in time. Such results are in agreement with the empirical findings of Allez and Bouchaud (2011), who employ standard sample correlations to study the intraday evolution of the correlations. An increase of correlations during the trading day is also found by Bibinger et al. (2019) and Koopman et al. (2018).

On some specific days, the time-varying parameters deviate substantially from their average intraday pattern. To this purpose, we plot in Figures (9), (10) the quantities \tilde{d}_t^j , \tilde{h}_t^j , $\tilde{\delta}_t^j$ and $\tilde{\rho}_t^j$ computed in correspondence of two meetings of the Federal Open Market Committee (FOMC), the first on 30-04-2014 and the second on 30-07-2014. Compared to the average intraday patterns, we observe significant deviations in the interval between 14:00 to 15:00, coinciding with the press conference in which economic news are released by the central bank. The local-level model, exploiting all the available 1-second data, instantaneously updates the time-varying parameters, allowing to reconstruct in high resolution the dynamic interaction of the covariances with the external information. FOMC events are characterized by an increase of volatilities at 14:00, followed by a decline in the subsequent minutes. The increase of the microstructure noise volatility at 14:00 is in agreement with the fact that bid-ask spreads are observed to increase during public announcements. After 15:00, volatilities go back to their average pattern. Correlations are also characterized by a significant increase in correspondence of the opening of the FOMC. It is worth pointing out that, due to data reduction, standard return models would not provide such high resolution description of covariance dynamics. The additional empirical results reported in the online appendix in Section (1.1.2) show that, on FOMC days, the covariances behave similarly even under the robust filtering algorithm based on the scores of the Student- t density. Thus, the fast increase in variances and correlations is not due to misspecification, but is an effect of the release of new information.

The fact that the local-level model rapidly responds to new information arriving on the market shares some similarities with the macroeconomic literature on nowcasting, where dynamic factor models estimated on

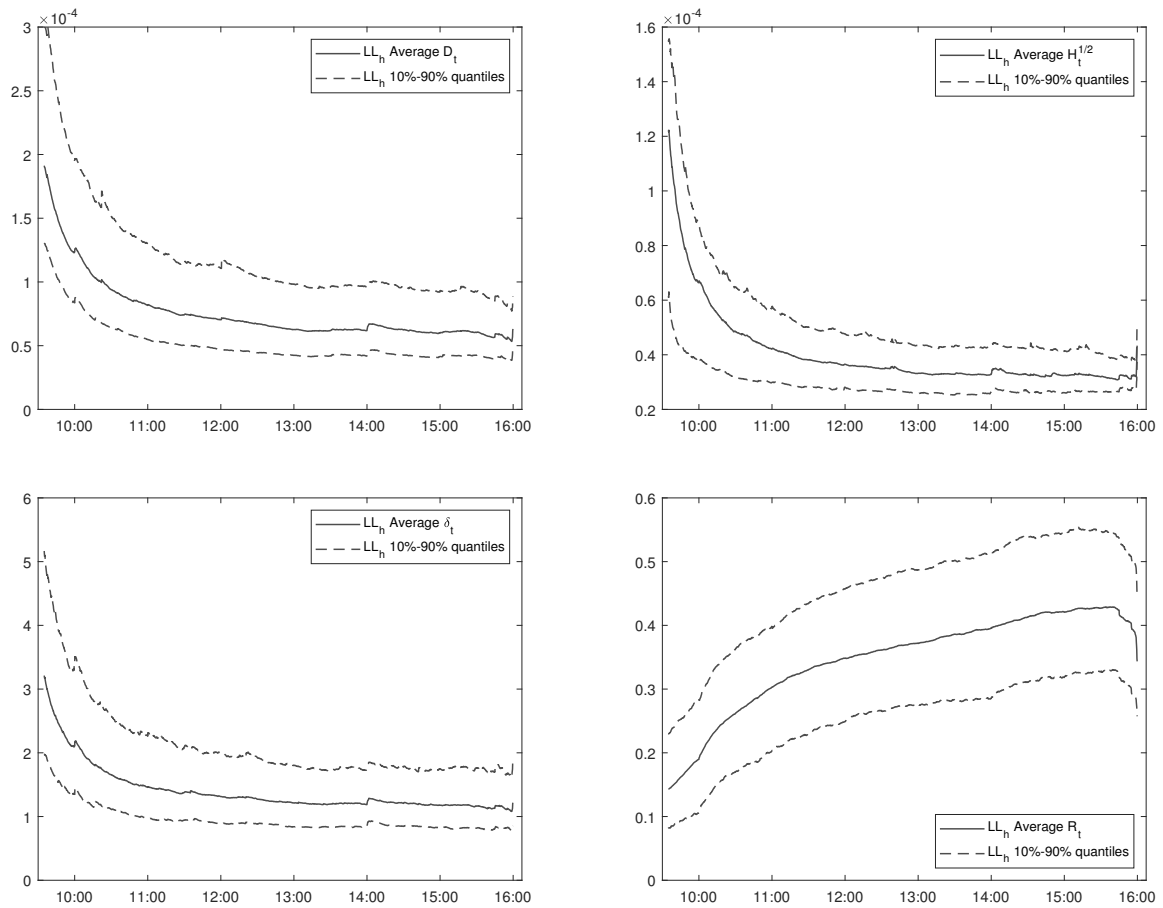


Figure 7: Average intraday patterns of the estimated LL_h time-varying parameters D_t , $H_t^{1/2}$, δ_t and R_t . For each $t = 1, \dots, 23400$, we report 10% and 90% quantiles of the distribution (over assets or couples of assets) of averages of D_t , $H_t^{1/2}$, δ_t and R_t (over days).

mixed-frequency observations are used to update the forecasts of macroeconomic variables (see e.g. Giannone et al. 2008, Creal et al. 2014 and Delle Monache et al. 2016). In a similar fashion, the local-level model can be employed as a nowcasting tool for high-frequency data, as shown in the next section.

4.4 Intraday out-of-sample analysis: portfolio selection and VaR assessment

In this section, we examine the performance of the local-level model as a nowcasting tool for high-frequency data. We perform two empirical tests that are typically employed to assess the quality of covariance forecasts in the financial econometric literature. The main novelty of our empirical setting is that we work on an intraday basis, considering short optimization horizons. Intraday risk and portfolio management have become increasingly popular due to high-frequency automatic trading, which accounts today for about 50% of total trades in the US equity market (see e.g. Karmakar and Paul 2019).

In the first test, we construct intraday Global Minimum Variance (GMV) portfolios and compare the ex-post realized variance of the portfolios constructed through the LL_h and LL_e models with that of the portfolios constructed through alternative methods. We adopt the following strategy: on each day $j = 2, \dots, 251$, we




Figure 8: Average intraday pattern of the first five eigenvalues of the estimated LL_h correlation matrix R_t .

build a sequence of 1-minute portfolios using the parameter estimates recovered on day $j - 1$. We rebalance the portfolios every minute, ending up with 390 portfolios per day. Following Engle and Colacito (2006) and Patton and Sheppard (2009), we choose as “best covariance estimator” the one minimizing the ex-post portfolio variance, computed as:

$$\sigma_j^2 = \sum_{k=1}^{390} (\hat{w}'_{k,j} r_{k,j})^2 \quad (29)$$

where $\hat{w}_{k,j}$ denotes the solution of the GMV problem at minute k of day j and $r_{k,j}$ is the vector of 1-minute returns at minute k of day j . We employ the same dataset of $n = 10$ assets used in the previous analysis. As benchmarks, we consider the t -GAS model, the DCC and the EWMA. The DCC and the EWMA are estimated by synchronizing the prices at the frequency of one minute, whereas the t -GAS is estimated using the missing value approach. The ex-post realized portfolio variances provided by the different methods are compared through the Model Confidence Set (MCS) of Hansen et al. (2011). In particular, we consider a MCS with a confidence level of 90% (denoted by $\mathcal{M}_{90\%}$) and perform a MCS analysis on each day of the dataset, ending up with 250 MCS tests.

Table (8) reports the results of the analysis. The first line shows the average of σ_j^2 , over $j = 2, \dots, 251$. The second and third lines show the number of days in which the corresponding portfolio is included in the MCS and the average p -values. We observe that the two local-level specifications feature significantly lower ex-post portfolio variances. The competing approaches, including the t -GAS, are instead excluded from the MCS in the vast majority of the cases.

As a second test, we perform a VaR assessment in the spirit of Creal et al. (2011). However, we work at the intraday level and compute the 1-minute VaR of the previously constructed GMV portfolios at 99% and 95% confidence levels. To assess the quality of the VaR forecasts, we compute the coverage, defined as the fraction of VaR exceedances within the trading day. In Table (8), we report the average, over the 250 business days, of the coverage provided by each model. We use the Kupiec (1995) test with a significance




Figure 9: We plot \tilde{d}_t^j , \tilde{h}_t^j , $\tilde{\delta}_t^j$, $\tilde{\rho}_t^j$ computed for $j = 82$, corresponding to 30-04-2014

level of 5% to assess the statistical significance of the estimated coverage values. The null hypothesis \mathcal{H}_0^K in the Kupiec test is that the coverage is equal to $1-\text{CL}$, where CL denotes the VaR confidence level. In Table (8) we also report the number of rejections of \mathcal{H}_0^K and the average p -values.

The two local-level specifications provide a good matching to the nominal confidence. The benchmark models clearly fail in capturing the risk of high-frequency asset returns. In particular, we note that the t -GAS provides a very small coverage, meaning that it heavily overestimates the risk. This is not surprising, as we have seen in Figure (5) that the effect of microstructure noise on the t -GAS is to overestimate the variances. Concerning the relative performance between LL_h and LL_e models, in both the portfolio and the VaR tests we find that the LL_h features slightly larger average p -values and that the null assumption in the two tests is rejected a larger number of times for the LL_e model. Overall, our results suggest that, while the in-sample loss due to the equicorrelation assumption is not large, the LL_h model provides significant out-of-sample forecast gains in common intraday financial applications.

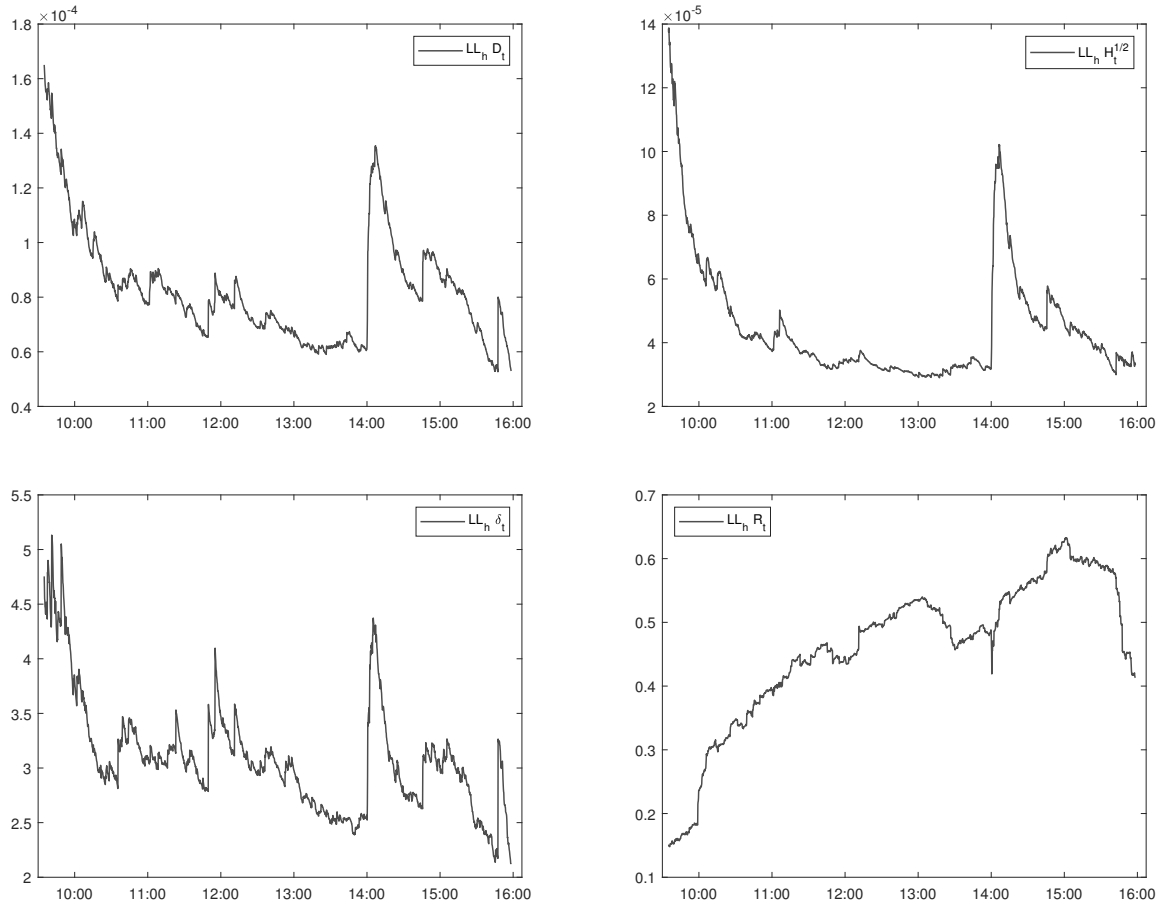


Figure 10: We plot \tilde{d}_t^j , \tilde{h}_t^j , $\tilde{\delta}_t^j$, $\tilde{\rho}_t^j$ computed for $j = 145$, corresponding to 30-07-2014

5 Conclusions

In this paper we analyze the problem of intraday covariance modeling with noisy and asynchronous prices and propose a modeling strategy that handles both effects. Specifically, we model intraday data through a conditionally normal local-level model with time-varying covariance matrices. The dynamics of the covariances are driven by the score of the conditional density, allowing to write the log-likelihood in closed form. In this framework, asynchronous trading is treated as a standard missing value problem, which is easily handled by the Kalman filter.

The main advantage of the proposed approach is that we model the covariances of the latent efficient returns instead of the observed returns. In high-frequency data, where observations are asynchronous, models for observed returns are indeed subject to several distortions. On the one hand, if data are synchronized by previous-tick interpolation, a large number of artificial zero returns arises which lead to a downward bias in the estimated correlations. If a missing value approach is adopted, observed return models are instead subject to large data reduction. On the other hand, microstructure noise leads to additional biases in the estimation of the intraday covariances in observed return models.

In our simulation analysis we study the finite sample properties of the maximum likelihood estimator

	LL _h	LL _e	<i>t</i> -GAS	DCC	EWMA
	Ex-post Portfolio Variance				
Avg. variance	0.5428	0.5501	0.5895	0.6987	0.7023
N. of days in $\mathcal{M}_{90\%}$	170	165	98	23	0
Avg. <i>p</i> -value	0.6599	0.6238	0.2621	0.0061	0.0000
	99% VaR				
Avg. coverage	0.0198	0.0202	0.0016	0.4976	0.4812
N. of rejections of \mathcal{H}_0^K	67	82	198	250	250
Avg. <i>p</i> -value	0.2681	0.2151	0.0236	0.000	0.000
	95% VaR				
Avg. coverage	0.0566	0.0605	0.0038	0.4976	0.4923
N. of rejections of \mathcal{H}_0^K	48	52	239	250	250
Avg. <i>p</i> -value	0.2885	0.2885	0.0025	0.000	0.0000

Table 8: Top: Average variance ($\times 10^5$) of 1-minute GMV portfolios constructed through out-of-sample covariance forecasts of LL_h, LL_e, *t*-GAS and DCC models. We report the number of days in which the models are included in $\mathcal{M}_{90\%}$ and the average of the *p*-values of the MCS test. Middle: Average coverage of 1-minute VaR's at 99% CL obtained from the covariance forecasts of GMV portfolios. We report the number of rejections of the Kupiec test with a significance level of 5% and the average *p*-value. Down: As before, but VaR's are computed at 95% CL.

and find that it remains unbiased in scenarios characterized by high levels of sparsity. We then show the advantages of the proposed methodology over standard correlation models for observed returns. In particular, in the presence of missing values, the use of our methodology is preferable even under a misspecified data generating process with a fat-tailed conditional density. The reason lies in the ability of the local-level model to use all the available prices in order to reconstruct the efficient price and the covariance dynamics. In contrast, correctly specified return models are subject to large data reduction, and the quality of their estimates rapidly deteriorates as data become more sparse.

The advantages of the methodology are then assessed on 1-second transaction data of 10 NYSE stocks. The in-sample analysis based on the AIC reveals that benchmark return models are heavily affected by data reduction and that correlations are biased due to market microstructure effects. The analysis of the intraday covariances shows that, while the first trading hours are dominated by idiosyncratic risk, a market risk factor progressively emerges in the second part of the trading day, with the correlations among all couples of assets increasing until the last few minutes of trading. The fact that all the available prices are used to estimate the covariances allows to capture instantaneously fast changes on covariance dynamics due to macro-news announcements, such as those released during FOMC announcements.

Finally, we test the local-level model as a nowcasting tool for high-frequency data. To this purpose, we design an intraday out-of-sample portfolio experiment and compare the ex-post realized variance of portfolios constructed through competing estimators. We found that, in most of the days in our dataset, the variance of portfolios based on local-level forecasts is significantly lower than other portfolio variances. We also find that the local-level model provides reliable intraday Value-at-Risk assessment, whereas standard methods

largely fail in measuring the risk of high-frequency returns.

References

- Alexander, C., 2002. Principal component models for generating large GARCH covariance matrices. *Economic Notes* 31, 337–359.
- Allez, R., Bouchaud, J.P., 2011. Individual and collective stock dynamics: intra-day seasonalities. *New Journal of Physics* 13, 025010.
- Andersen, T., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance* 4, 115–158.
- Barndorff-Nielsen, O.E., Hansen, P.R., Lunde, A., Shephard, N., 2009. Realized kernels in practice: trades and quotes. *The Econometrics Journal* 12, C1–C32.
- Bibinger, M., Hautsch, N., Malec, P., Reiss, M., 2019. Estimating the spot covariation of asset prices—statistical theory and empirical evidence. *Journal of Business & Economic Statistics* 37, 419–435.
- Blasques, F., Koopman, S.J., Lucas, A., 2015. Information-theoretic optimality of observation-driven time series models for continuous responses. *Biometrika* 102, 325.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307–327.
- Buccheri, G., Bormetti, G., Corsi, F., Lillo, F., 2019a. Filtering and smoothing with score-driven models. Working paper .
- Buccheri, G., Livieri, G., Pirino, D., Pollastri, A., 2019b. A closed-formula characterization of the Epps effect. *Quantitative Finance* .
- Corsi, F., Peluso, S., Audrino, F., 2015. Missing in asynchronicity: A Kalman-EM approach for multivariate realized covariance estimation. *Journal of Applied Econometrics* 30, 377–397.
- Cox, D., 1981. Statistical analysis of time series: Some recent developments [with discussion and reply]. *Scandinavian Journal of Statistics* 8, 93–115.
- Creal, D., Koopman, S.J., Lucas, A., 2008. A general framework for observation driven time-varying parameter models. *SSRN Electronic Journal* .
- Creal, D., Koopman, S.J., Lucas, A., 2011. A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. *Journal of Business & Economic Statistics* 29, 552–563.
- Creal, D., Koopman, S.J., Lucas, A., 2013. Generalized autoregressive score models with applications. *Journal of Applied Econometrics* 28, 777–795.

- Creal, D., Schwaab, B., Koopman, S.J., Lucas, A., 2014. Observation-driven mixed-measurement dynamic factor models with an application to credit risk. *The Review of Economics and Statistics* 96, 898–915.
- Delle Monache, D., Petrella, I., Venditti, F., 2016. Adaptive state space models with applications to the business cycle and financial stress. *CEPR Discussion Paper* .
- Delle Monache, D., Petrella, I., Venditti, F., 2019. Price dividend ratio and long-run stock returns: a score driven state space model. *Working paper* .
- Durbin, J., Koopman, S., 2012. *Time Series Analysis by State Space Methods: Second Edition*. Oxford Statistical Science Series, OUP Oxford.
- Engle, R., 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50, 987–1007.
- Engle, R., 2002. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics* 20, 339–50.
- Engle, R., Colacito, R., 2006. Testing and valuing dynamic correlations for asset allocation. *Journal of Business & Economic Statistics* 24, 238–253.
- Engle, R., Kelly, B., 2012. Dynamic equicorrelation. *Journal of Business & Economic Statistics* 30, 212–228.
- Engle, R., Russell, J., 1998. Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica* 66, 1127–1162.
- Engle, R.F., Kroner, K.F., 1995. Multivariate simultaneous generalized ARCH. *Econometric Theory* 11, 122–150.
- Epps, T.W., 1979. Comovements in stock prices in the very short run. *Journal of the American Statistical Association* 74, 291–298.
- Giannone, D., Reichlin, L., Small, D., 2008. Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics* 55, 665 – 676.
- Golam Kibria, B.M., Joarder, A.H., 2017. A short review of multivariate t-distribution. *Journal of Statistical Research* .
- Hansen, P.R., Lunde, A., Nason, J.M., 2011. The model confidence set. *Econometrica* 79, 453–497.
- Harvey, A., 1991. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Harvey, A., Luati, A., 2014. Filtering with heavy tails. *Journal of the American Statistical Association* 109, 1112–1122.
- Harvey, A.C., 2013. *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. Cambridge University Press. *Econometric Society Monographs*.

- Hasbrouck, J., 1993. Assessing the quality of a security market: A new approach to transaction- cost measurement. *The Review of Financial Studies* 6, 191–212.
- Hayashi, T., Yoshida, N., 2005. On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli* 11, 359–379.
- Karmakar, M., Paul, S., 2019. Intraday portfolio risk management using VaR and CVaR: a CGARCH-EVT-copula approach. *International Journal of Forecasting* 35, 699 – 709.
- Koopman, S.J., Lit, R., Lucas, A., Opschoor, A., 2018. Dynamic discrete copula models for high-frequency stock price changes. *Journal of Applied Econometrics* 0. doi:10.1002/jae.2645.
- Kupiec, P.H., 1995. Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives* 3, 73–84.
- Lucas, A., Opschoor, A., Schaumburg, J., 2016. Accounting for missing values in score-driven time-varying parameter models. *Economics Letters* 148, 96 – 98.
- Madhavan, A., 2000. Market microstructure: A survey. *Journal of Financial Markets* 3, 205 – 258.
- McInish, T.H., Wood, R.A., 1992. An analysis of intraday patterns in bid/ask spreads for NYSE stocks. *The Journal of Finance* 47, 753–764.
- Nelson, D.B., 1991. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* 59, 347–370.
- Oh, D.H., Patton, A.J., 2018. Time-varying systemic risk: Evidence from a dynamic copula model of CDS spreads. *Journal of Business & Economic Statistics* 36, 181–195.
- Patton, A.J., Sheppard, K., 2009. *Evaluating Volatility and Correlation Forecasts*. Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 801–838.
- Roll, R., 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance* 39, 1127–1139.
- Shephard, N., Xiu, D., 2017. Econometric analysis of multivariate realised QML: Estimation of the covariation of equity prices under asynchronous trading. *Journal of Econometrics* 201, 19 – 42.
- Tsay, R.S., 2005. *Analysis of financial time series*. Wiley series in probability and statistics, Wiley-Interscience, Hoboken (N.J.).
- Tse, Y.K., Tsui, A.K.C., 2002. A multivariate generalized autoregressive conditional heteroscedasticity model with time-varying correlations. *Journal of Business & Economic Statistics* 20, 351–362.
- van der Weide, R., 2002. GO-GARCH: A multivariate generalized orthogonal GARCH model. *Journal of Applied Econometrics* 17, 549–564.