Realized Factor Models for Vast Dimensional Covariance Estimation

Karim Bannouh* Martin Martens† Roel Oomen‡ Dick van Dijk§

PRELIMINARY, INCOMPLETE, COMMENTS SOLICITED

June 9, 2009

Abstract

We introduce a novel approach for estimating vast dimensional covariance matrices of asset returns by combining a linear factor model structure with the use of high- *and* low-frequency data. Specifically, we propose the use of "liquid" factors – i.e. factors that can be observed free of noise at high frequency – to estimate the factor covariance matrix and idiosyncratic risk with high precision from intra-day data whereas the individual assets' factor exposures are estimated from low frequency data to counter the impact of non-synchronicity between illiquid stocks and highly liquid factors. Our theoretical and simulation results illustrate that the performance of this "mixed-frequency" factor model is excellent: it compares favorably to the Hayashi and Yoshida (2005) covariance estimator (in a bi-variate setting) and the realized covariance estimator in the presence of market microstructure noise and non-synchronous trading. In empirical applications for the S&P500, S&P400 and S&P600 stock universes and using highly liquid ETFs as proxies for the Fama and French (1992) style and industry factors, we find that the mixed-frequency factor model delivers better tracking errors and Value-at-Risk forecasts compared to the realized covariance. In contrast to the realized covariance the performance of the "mixed-frequency factor model" is robust across sampling frequencies and forecast weighting schemes.

*Econometric Institute, Erasmus University Rotterdam (corresponding author, e-mail: `bannouh@ese.eur.nl`)

†Department of Finance, Erasmus University Rotterdam

‡Deutsche Bank, London and Department of Quantitative Economics, University of Amsterdam

§Econometric Institute, Erasmus University Rotterdam

1 Introduction

Accurate measures and forecasts of asset return covariances are important for risk management and portfolio management. Recent academic research in these areas has focused on two different issues. First, intra-day data has been shown to render more precise measures and forecasts of daily asset return volatilities and covariances. Second, for the practically relevant case of portfolios consisting of a large number of assets, several studies have considered factor structures to tackle the "curse of dimensionality". In this paper we put forward a novel approach for accurate measurement and forecasting of the covariance matrix of vast dimensional portfolios by combining the use of high- and low-frequency data with a linear factor structure. Specifically, we develop a "mixed-frequency" factor model (MFFM), where high-frequency data on liquid factors is used for precise estimation of the factor covariance matrix whereas the factor loadings are estimated from low-frequency data. The first aspect exploits the benefits of intra-day data, while the second aspect is a conservative choice to avoid the possibility that factor loading estimates are biased towards zero due to non-synchronous trading patterns between relatively frequently traded factors and stocks that are less liquid in general.

In recent years, a substantial body of literature has emerged on the use of financial high-frequency data for obtaining more accurate measures and forecasts of financial risk. In a univariate setting, high-frequency data is generally very useful for the purpose of volatility estimation, see e.g. Andersen et al. (2006) for a recent review. Yet, in a multivariate setting, especially when the number of assets is large, the benefits of high-frequency data are less clear-cut. The obvious estimator of the covariance between two assets, the so-called realized covariance (RC), is computed by summing the cross-products of their intra-day returns, see e.g. Barndorff-Nielsen and Shephard (2004). This has several drawbacks, however. First, the sensitivity of RC to non-synchronous trading and market microstructure noise reduces its efficiency. Second, it is susceptible to spurious intra-day dependence. Third, it can produce unstable covariance matrices, particularly when the dimension is relatively high compared to the number of observations and this may lead to error maximization in portfolio construction.

Several recent studies have revisited the use of factor models for covariance estimation in case of a large number of assets, in order to reduce the dimensionality of the problem, see e.g. Chan et al. (1999) and Fan et al. (2008). The factor model approach may substantially improve over the sample (realized) covariance matrix in particular when the portfolio optimization problem requires the inverse of the covariance matrix, as shown by Fan et al. (2008). Obviously this is due to the fact that in the factor model approach only the factor covariance matrix needs to be inverted, which typically is of much lower dimension. In addition, using the covariance matrix based on a factor structure reduces the problem of error maximization for portfolio construction applications, see Jagannathan and Ma (2003).

In applications of factor models, typically all ingredients (including the factor covariance matrix, idiosyncratic risk and factor loadings) are estimated using returns at a daily or even lower sampling frequency. In this paper we propose a "mixed-frequency" factor model for estimating the daily covariance matrix for a vast number of assets, which aims to exploit the benefits of high-frequency data and a factor structure. In this MFFM, the factor loadings are obtained in the conventional way by linear regression using a history of daily stock- and factor-returns. However, the factor covariance matrix and residual variances are calculated with high precision from intra-day data. The motivation for this particular mix of frequencies is as follows. First, nowadays highly liquid financial contracts are available as proxies for the most commonly used factors. Natural candidates for such "factor proxies" include index futures contracts and exchange-traded-funds (ETF) covering a range of asset classes, industries, styles and segments of the market and for which high-quality intra-day data is plentiful. Given that these contracts are highly liquid, their covariances can be estimated with high precision from intra-day data. Second, although intra-day data may also be available for individual stocks, these are generally less liquid than index futures and ETFs. The realized covariance between a stock and a factor proxy thus could be heavily affected by the non-synchronicity in their trading patterns. We therefore do not use intra-day returns to estimate the factor loadings, we propose to use the more conservative daily sampling frequency. Third, the relative illiquidity of individual stocks is much less of a problem for estimating their individual variances. Hence, we do use intra-day data to estimate the idiosyncratic variance in the factor model.

The "mixed-frequency" factor model methodology has several advantages over the realized covariance matrix. First, the advantages of dimension reduction in the context of the factor model based purely on daily data continue to hold in the MFFM. Second, the MFFM makes efficient use of high-frequency factor data while bypassing potentially severe biases induced by microstructure noise for the individual assets. This applies in particular to the issue of non-synchronous trading which induces a bias towards zero in the realized covariance. By using liquid ETFs to proxy factors, we exploit the fact that these assets trade much more frequently than individual stocks, reducing the non-synchronicity problem. For the same reason we estimate the factor loadings using daily returns data. Third, we can easily expand the number of assets in the MFFM approach while this is more difficult with the RC matrix for which the inverse does not exist when the number of assets exceeds the number of return observations per asset.

In terms of theory, we show that under the assumption that we use the correct factor specification and constant factor loadings, the covariance estimates from our MFFM are substantially more efficient than those computed using the Hayashi and Yoshida (2005) estimator on tick-data, which in turn is more efficient than the RC when considering pair-wise covariances. In extensive Monte Carlo simulations, we further explore the performance of the MFFM compared to the RC in less ideal circumstances. The data generating process (DGP) is a factor structure, in which non-synchronous trading is implemented by assuming that trades occur according to a Poisson arrival process. Realistic trading intensities are obtained by estimating the arrival probabilities using the average number of trades per day of individual S&P500 constituents and ETFs that cover the Fama and French (1992) style and industry factors. We analyze the impact of market microstructure noise by calibrating the "noise ratio" of Oomen (2006) which relates the level of noise to the number of intra-day observations of an asset. In addition, we allow for several magnitudes of estimation error in the factor loadings. In all cases MFFM produces more accurate estimates of the covariance matrix than the RC. Random (cross-sectionally independent) noise of several magnitudes on the factor loadings has only a minor impact on the performance of the MFFM.

To evaluate the empirical performance of MFFM we consider the S&P500 (large cap), S&P400 (mid cap) and S&P600 (small cap) stock universes. It is quite unique in the literature to consider such a large universe, see also Engle et al. (2008). For empirical data we obviously do not have the true covariance matrix to judge the quality of the MFFM estimator and the realized covariance. Hence, we need applications for which the covariance matrix is needed. We consider two different applications, one in risk management for which the covariance matrix is needed and one in portfolio optimization where the inverse of the covariance matrix is the key input. First, for the risk management application we analyze the empirical performance of RC and MFFM by forecasting the Value-at-Risk (VaR) of equal-weighted portfolios based on the covariance matrix. Second, we use the inverse of the covariance matrix to calculate the optimal portfolio weights in a minimum tracking error application where the objective is to construct portfolios that stay as close as possible to a benchmark by minimizing the standard deviation of daily return differences between the portfolio and the benchmark, known as tracking error. Chan et al. (1999) for monthly data compare the sample covariance matrix with several factor models based on constructing minimum variance and minimum tracking error portfolios. The problem with minimum variance portfolios is that they primarily focus on low beta stocks to minimize systematic risk by eliminating as much as possible the market factor from the portfolio. Hence, there is little difference between the various approaches due to the dominant market factor. In contrast, for minimizing the tracking error the market factor plays a limited role and it becomes more important to consider other factors, such as style and industry factors. Cavaglia et al. (2000) illustrate the increasing importance of industry factors and Chan et al. (1999) conclude that especially size and industries are important factors to consider. Together with the fact that in recent years several ETFs have become very liquid we decide to use ETF proxies for the Fama and French (1992) factors and industries in the empirical applications.

For the S&P500, S&P400 and S&P600 stock universes MFFM manages a tracking error of 4.0%, 5.2% and 6.1% per annum, respectively, when using all stocks. An additional benchmark that we include is the performance of naïvely diversified equal-weighted 1/N portfolios which manage a tracking error of 5.1%, 6.0% and 6.5% for the large, mid and small caps. In each S&P universe the performance of MFFM is better than that of the equal-weighted portfolio and RC. The best results for the realized covariance matrix for the S&P500, S&P400 and S&P600 are 4.2%, 5.5% and 9.5%, with, in contrast to the MFFM, results depending heavily on the forecast weighting scheme and sampling frequency used. Hence, we find that the differences between RC and MFFM increase with the level of non-synchronous trading in the individual stocks, which is relatively small for S&P500 large caps (8,272 trades per day on average) but substantial for the S&P600 small caps (1,411 trades per day).

An additional draw back of using RC is that the inverse does not exist at sampling frequencies lower than 30 minutes whereas MFFM produces robust and better results at all considered sampling frequencies.

For the S&P500 constituents we find that MFFM produces better out-of-sample VaR forecasts compared to RC. The null hypothesis of accurate (un)conditional coverage and time-series independence is rejected much more frequently based on likelihood ratio tests for RC than for MFFM forecasts. For example, in the case of portfolios consisting of 25 stocks and 5% VaR levels we find that when using MFFM the null of accurate conditional coverage (which is equivalent to testing for unconditional coverage and time-series independence simultaneously) is rejected for less than 6% of the portfolios at the 5% significance level when using a sampling frequency between 5-130 minutes. In contrast, when using the same portfolios but now RC for the VaR forecast, the null is rejected for at least 17% of the portfolios. Similar to the results in the minimum tracking error application we find that the results for RC depend severely on the sampling frequency and forecast scheme applied, whereas the MFFM performance is substantially more robust regarding the choice of sampling frequency and forecast weights.

The remainder of this paper is structured as follows. In Section 2 we provide theoretical examples which illustrate the merits of the "mixed-frequency factor model" methodology. Section 3 contains a Monte Carlo study. In Section 4 we apply MFFM to empirical data by analyzing VaR forecasts and forming minimum tracking error portfolios. We conclude in Section 5.

2 The Mixed-Frequency Factor Model

Consider a linear factor structure on asset$-i$ returns:

$$r_i = \beta_i' f + \varepsilon_i \tag{1}$$

where r_i and ε_i are scalars, and β_i and f are $K \times 1$ vectors. The covariance between asset i and asset j can be expressed as:

$$\gamma_{ij} \equiv cov(r_i, r_j) = \beta_i' \Lambda \beta_j + \sigma_{ij} \tag{2}$$

where $\Lambda = E(ff')$ and $\sigma_{ij} = E(\varepsilon_i \varepsilon_j)$. Throughout, we consider a "strict" factor structure in the spirit of Ross (1976), i.e. we assume that the factor structure exhausts the dependence among the assets such that $\sigma_{ij} = 0$.[1] With estimated quantities of β and Λ, the covariance estimator is then equal to:

$$\widehat{\gamma}_{ij} = \widehat{\beta}_i' \widehat{\Lambda} \widehat{\beta}_j \qquad \text{for } i \neq j \tag{3}$$

The properties of this covariance estimator are characterized in the theorem below, where we use the notation $\widehat{X} = X + X^{\varepsilon}$.

[1] Approximate factor models where σ_{ij} can be non-zero but small are considered in Chamberlain and Rothschild (1983), Ingersoll (1984) and Connor and Korajczyk (1994).

Theorem 2.1 *Assuming (i) $E(\sigma_{ij})=0$, (ii) $E(\beta^\varepsilon)=0$, (iii) $E(\Lambda^\varepsilon)=0$, and (iv) $\beta^\varepsilon \perp \Lambda^\varepsilon$ element-by-element, then we have the following result for $i \neq j$:*

$$E(\widehat{\gamma}_{ij}) = \gamma_{ij} \tag{4}$$

and

$$\begin{aligned} V(\widehat{\gamma}_{ij}) &= \beta_i'\Lambda\Sigma_{\beta,j}\Lambda'\beta_i + \beta_j'\Lambda'\Sigma_{\beta,i}\Lambda\beta_j + tr(\Sigma_{\beta,i}\Lambda\Sigma_{\beta,j}\Lambda') + g(\beta_i\beta_i', \beta_j\beta_j', \Phi) \\ &\quad + g(\beta_i\beta_i', \Sigma_{\beta,j}, \Phi) + g(\beta_j\beta_j', \Sigma_{\beta,i}, \Phi) + g(\Sigma_{\beta,i}, \Sigma_{\beta,j}, \Phi) \end{aligned} \tag{5}$$

where $\Sigma_{\beta,i} = V(\widehat{\beta}_i)$ and $\Phi = E(vech(\Lambda^\varepsilon)vech(\Lambda^\varepsilon)')$ and

$$g(A,B,\Phi) = \sum_{m,n,p,q}^{N} A_{mp}B_{nq}\Phi_{f(p,n),f(q,m)}$$

and $f(p,q) = N(\min\{p,q\}-1) + \frac{1}{2}(\min\{p,q\} - \min\{p,q\}^2) + \max\{p,q\}$.

Proof See Appendix A. ■

We now specialize this general setting to one where the factor factors are observed at "high" frequency and the individual assets' returns at "low" frequency, i.e. define :

F : low-frequency factor return observations ($T \times K$),
$\mathcal{F}$: high-frequency factor return observations ($M \times K$),
R_i : low-frequency asset$-i$ return observations ($T \times 1$),
$\mathcal{R}_i$: high-frequency asset$-i$ return observations ($N_i \times 1$),
τ_i : time-stamps associated with $\mathcal{R}_i$ ($N_i \times 1$).

Assumption N the factor returns $\mathcal{F}$ are jointly Normal with zero mean, serially uncorrelated and observed without friction. The factor covariance matrix is estimated as $\widehat{\Lambda} = \mathcal{F}'\mathcal{F}$.

Assumption O the asset return dynamics at low frequency are governed by a linear factor model as in Eq. (1) with i.i.d. Normal residuals. The factor loadings are estimated using linear regression $\widehat{\beta}_i = (F'F)^{-1}F'R_i$.

Corollary 2.2 *Let assumption **N**, **O**, and those in Theorem 2.1 hold. Then:*

$$V(\widehat{\gamma}_{ij}) = \frac{A}{T} + \frac{B+C}{M}, \tag{6}$$

where

$$\begin{aligned} A &= \sigma_j^2\beta_i'\Lambda\beta_i + \sigma_i^2\beta_j'\Lambda\beta_j + \sigma_i^2\sigma_j^2\frac{K}{T}, \\ B &= \sum_{m,n,p,q}^{K} \beta_i(m)\beta_i(p)\beta_j(n)\beta_j(q)(\Lambda_{pq}\Lambda_{nm} + \Lambda_{pm}\Lambda_{nq}), \\ C &= \sum_{m,n,p,q}^{K} (\beta_i(m)\beta_i(p)\Sigma_{\beta,j}(n,q) + \beta_j(m)\beta_j(p)\Sigma_{\beta,i}(n,q) + \Sigma_{\beta,i}(m,p)\Sigma_{\beta,j}(n,q))(\Lambda_{pq}\Lambda_{nm} + \Lambda_{pm}\Lambda_{nq}). \end{aligned}$$

Figure 1: Comparison of MFFM to Hayashi-Yoshida in term of ln MSE

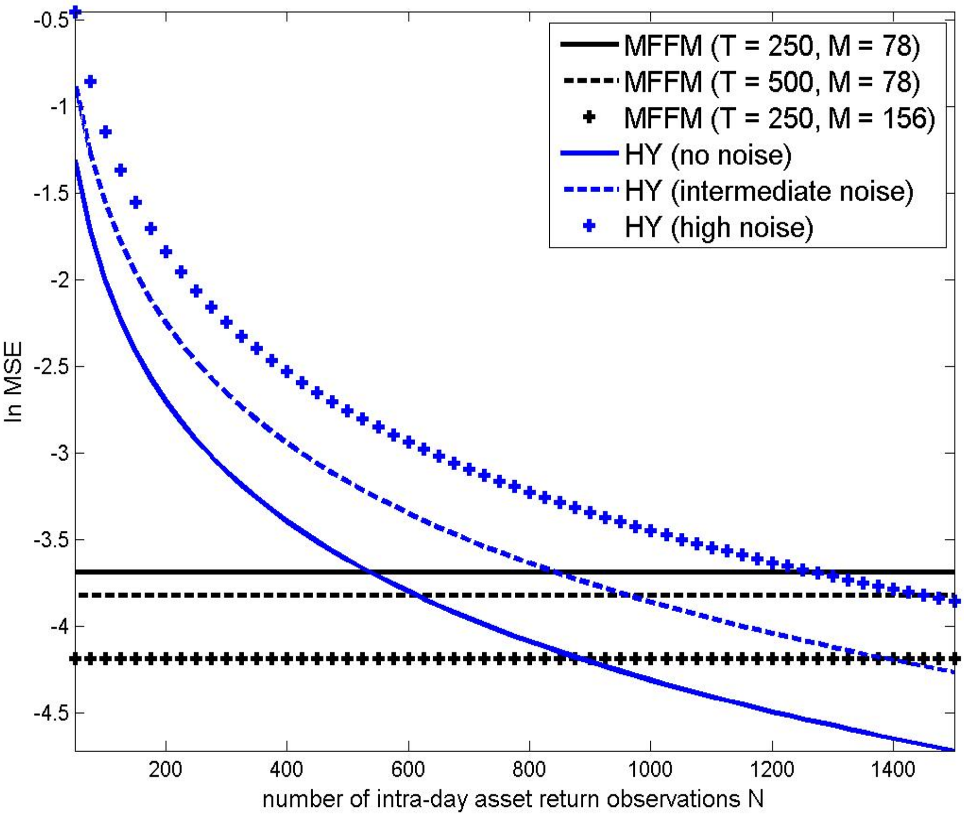

Proof See Appendix A. ■

The above setting describes the mixed-frequency factor model, namely factors are observed at high frequency free of micro-structure noise and the factor covariance matrix can be estimated with maximal precision using the conventional realized covariance. Individual asset returns, on the other hand, are only well described by a linear factor model with i.i.d. innovations when sampled at low frequency and this limits the precision at which we can measure the factor loadings. The above corollary highlights this mechanism: the variance of the mixed-frequency estimator can be attributed to one component relating to the measurement error in factor loadings (i.e. A/T) and another component quantifying the measurement error in the factor covariance matrix (i.e. $(B+C)/M$).[2]

To illustrate the efficiency gains that can be attained with this mixed-frequency factor model, assume that asset$-i$ intra-day price observations ($\mathcal{R}_i$) arrive according to a Poisson process with intensity $\lambda_i = E(N_i)$. Further, assume that prices are contaminated with i.i.d. microstructure noise

[2]Note that in some circumstances β is (assumed to be) known so that $V(\widehat{\gamma}_{ij}) = B/M$, see e.g. Grinold and Kahn (2000, Ch. 3).

with variance $\xi_i^2 = \pi_i \gamma_i^2/\lambda_i$. Then the covariance between asset i and j can be estimated using the Hayashi and Yoshida (2005) estimator:

$$HY_{ij} = \sum_{p=1}^{N_i} \sum_{q \in S_p} \mathcal{R}_i(p)\mathcal{R}_j(q), \tag{7}$$

where $S_p = \{q|(\tau_i(p-1), \tau_i(p)) \cap (\tau_j(q-1), \tau_j(q)) \neq \emptyset\}$. In this setting, the estimator is unbiased, i.e. $E(HY_{ij}) = \gamma_{ij}$, and has a variance (see Griffin and Oomen 2006) equal to:

$$V(HY_{ij}) = 2\gamma_i^2\gamma_j^2 \frac{\lambda_i + \lambda_j}{\lambda_i\lambda_j} + 2\frac{\gamma_{ij}^2}{\lambda_i + \lambda_j}\left(\frac{\lambda_j}{\lambda_i} + \frac{\lambda_i}{\lambda_j}\right) + 2\gamma_i^2\xi_j^2 + 2\gamma_j^2\xi_i^2 + 4\xi_i^2\xi_j^2\frac{\lambda_i\lambda_j}{\lambda_i + \lambda_j}. \tag{8}$$

Regarding the specification of the factor structure, we consider $K = 5$ factors, with factor loadings $\beta_i = (0.5, -0.1, 0, 0.2, 0.6)'$, $\beta_j = (0.7, -0.2, -0.3, 0.4, 0.2)'$, and factor covariance matrix $\Lambda = I_K + \frac{1}{2}(1 - I_K)$. The specific or idiosyncratic risk component is $\sigma_h^2 = \beta_h' \Lambda \beta_h$ for[3] $h \in \{i, j\}$ so that $\rho_{ij} \approx 40\%$ with:

$$V(r) = (\beta_i, \beta_j)'\Lambda(\beta_i, \beta_j) + \Sigma = \begin{pmatrix} 2.075 & 0.765 \\ 0.765 & 1.584 \end{pmatrix}$$

In Figure 1 we compare the variance of the two competing estimators, namely (i) the MFFM estimator based on T "daily" asset and factor returns to get the loadings and M "clean intra-day" factor returns to get the factor covariance and (ii) the HY estimator based on N asynchronous and noisy intra-day returns for assets i and j. From the results, we see that for reasonable scenarios, the MFFM comfortably outperforms the HY estimator unless a large number of intra-day return observations on the individual assets is available. For instance, using 5-minute ($M = 78$) factor returns to estimate the 5×5 factor covariance matrix and 1 year ($T = 250$) of daily asset returns to estimate the 5×1 factor loading vector β, the MFFM delivers better estimates unless the HY estimator has access to more than 500 clean or 1250 noisy intra-day (asynchronous) observations. Griffin and Oomen (2006) show that the efficiency of the realized covariance estimator – in this bi-variate setting, using synchronized data – is typically inferior to the HY estimator. From this it follows that the MFFM also compares favorable to RC. Finally, as already pointed out above, an additional advantage of the MFFM is that it delivers stable and positive definite covariance matrices (unlike HY and in some instances RC). The next section further explores the properties of the MFFM and RC in an extensive simulation study.

3 Monte Carlo Simulation

In this section we analyze the properties of the "mixed-frequency factor model" and the realized covariance estimator by means of simulation. The level of non-synchronous trading, market microstructure noise and estimation error in the MFFM factor loadings are based on estimates from empirical data.

[3]This value of the idiosyncratic or asset-specific variance implies a factor-regression R^2 of 50% which is empirically reasonable.

3.1 Covariance models

3.1.1 Realized covariance

The realized covariance is an efficient estimator of the latent integrated covariance. RC converges in probability to the integrated covariance in the absence of noise, see Barndorff-Nielsen and Shephard (2004). In the bi-variate case RC can be computed using cross-products of intra-day returns and in general it is defined as the outer-product of intra-day returns:

$$RC = \mathcal{R}'\mathcal{R}$$

where $\mathcal{R}$ denotes the matrix of $T \times N$ intra-day returns for the N index constituents. For increasing sampling frequencies empirical asset price data tends to be more contaminated with market microstructure noise and the impact of assets trading non-synchronously becomes more severe which causes covariance estimates to be biased towards zero.

When applying RC to a small number of assets the number of return observations is sufficiently large compared to the number of assets when sampling at moderate intra-day frequencies such as 30 minutes or hourly. When the number of assets N increases to the order of hundreds, however, we have to resort to higher frequencies to obtain more intra-day observations than the number of assets. Since the impacts of market microstructure noise and non-synchronous trading are severe at these ultra high frequencies, using the realized covariance becomes problematic and no longer provides a consistent estimate of the integrated covariance, see e.g. Bandi and Russell (2005).

3.1.2 Mixed-frequency factor model

To circumvent the problems regarding market microstructure noise and to be able to handle a vast amount of assets we propose a hybrid approach which combines the merits of factor models and high-frequency data. Besides being able to handle a vast amount of assets, the ETFs that we propose as factors for the "mixed-frequency factor model" are substantially more liquid and less noisy than individual stocks. To avoid beta estimates being possibly biased towards zero due to non-synchronous stock and factor prices we estimate these by linear regression using daily returns. The "mixed-frequency factor model" estimator uses low-frequency β estimates combined with high-frequency estimates of the factor covariance matrix and residual variances:

$$MFFM = \beta'\Lambda\beta + \text{diag}(\varepsilon'\varepsilon) \tag{9}$$

where $\Lambda = \mathcal{F}'\mathcal{F}$ is the $K \times K$ realized covariance matrix obtained using intra-day returns on the Fama and French (1992) factors, β denotes the $K \times N$ matrix of factor loadings, and the $T \times N$ idiosyncratic residuals are obtained using $\varepsilon = \mathcal{R} - \mathcal{F}\beta$.

3.2 Data generating process

As inputs for the simulation we estimate the exposures to the Fama and French factors[4] for the S&P500 constituents[5] using daily return data over the period 1/1/1998 to 31/12/2007 (2514 observations). To obtain estimates of the factor loadings we regress the excess returns of the stocks on the daily returns of the Fama and French factors,

$$R_{i,t} - R_{f,t} = \beta_{i,M}(R_{M,t} - R_{f,t}) + \beta_{i,SMB}R_{SMB,t} + \beta_{i,HML}R_{HML,t} + \varepsilon_{i,t}$$

where $R_{i,t} - R_{f,t}$ is the excess return over the risk-free rate on stock $i = 1, \ldots, 500$ and $\beta_i = [\beta_{i,M} \quad \beta_{i,SMB} \quad \beta_{i,HML}]'$ are the estimated exposures to the market, size (Small Minus Big) and value (High Minus Low) factor, respectively. Let $F = [R_M - R_f \;\; R_{SMB} \;\; R_{HML}]$ denote the matrix of daily factor returns. Estimated regression errors are obtained using $e_i = R_i - F\beta_i$ so that an unbiased estimator of the variance error of the regression is $\sigma^2_{e_i} = \frac{e_i' e_i}{T-K}$. To account for estimation errors in the factor loadings we calculate the standard errors of the factor loadings $\sigma_{b_i} = \sigma_{e_i}\sqrt{f_{jj}}$ where f_{jj} is the j-th diagonal element of $(F'F)^{-1}$. Figure 2 illustrates the regression R^2's (a) and estimated Fama and French (1992) factor loadings (b-d). The standard errors of the factor loadings are displayed in Figure 3.

Using the estimated daily factor covariance matrix $\Lambda = E(F'F)$, factor loadings β and residual variances σ^2_e we construct the covariance matrix for the DGP: $\Sigma = \beta'\Lambda\beta + \text{diag}(\sigma^2_e)$. We simulate second-by-second (23,400 seconds in a NYSE trading day) factor data $\mathcal{F}$ from a multivariate Gaussian distribution, $\mathcal{F} \sim \Phi(0, \Lambda\Delta)$ and high-frequency idiosyncratic residuals $\varepsilon \sim N(0, \text{diag}(\sigma^2_e\Delta))$ with time-step $\Delta = 1/23,400$. A simulated sample of high-frequency stock returns is obtained using $\mathcal{R} = \mathcal{F}\beta + \varepsilon$.

3.3 Non-synchronicity, microstructure noise and estimation error in factor loadings

We implement non-synchronous trading by assuming trades arrive following a Poisson process. The Poisson trading intensities are estimated by means of the the average number of trades per day of stocks and factors from empirical data, see Table 1.

In practice high-frequency financial asset prices are contaminated by market microstructure noise. We set the level of noise for individual assets by calibrating the "noise" ratio of Oomen (2006). Let Y denote the uncontaminated log-price series of an individual asset with integrated variance IV (the variance in the DGP) and M intra-day observations. We assume we observe $Z = Y + \eta$ where $\eta \sim N(0, \frac{IV}{4M})$. Hence the level of noise in this setup is related to the trading intensity and liquidity, it is common to assume that assets which trade frequently are more liquid and thus less contaminated

[4] $http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html$
[5] $http://wrds.wharton.upenn.edu/$

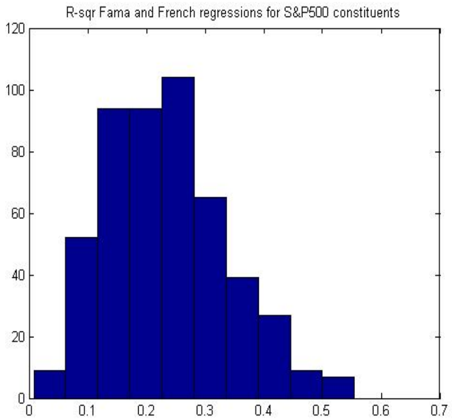

(a) R^2

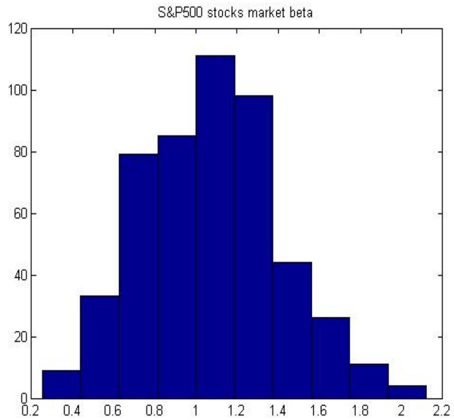

(b) Market beta's

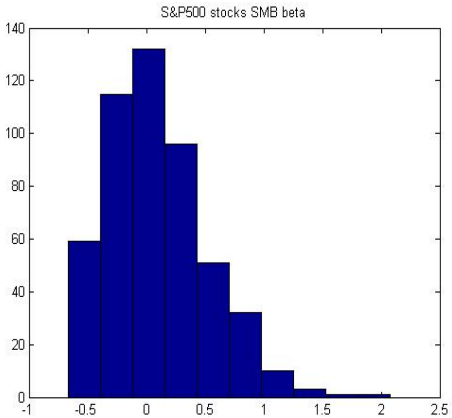

(c) SMB beta

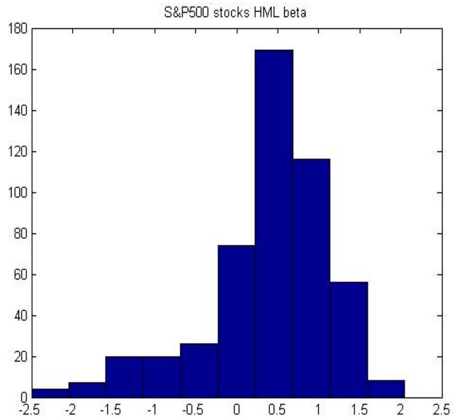

(d) HML beta

Figure 2: This Figure displays histograms of R^2's and factor loadings based on Fama and French regressions for S&P500 stocks using daily data over the period January 1998 to December 2007.

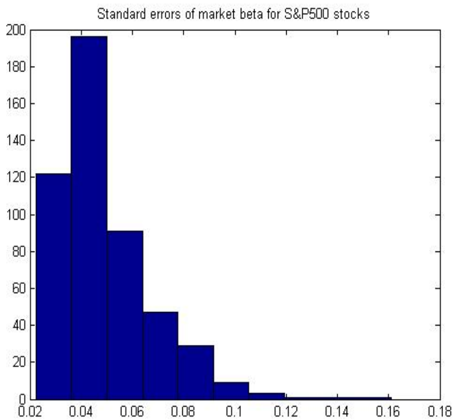

(a) Std. error market beta

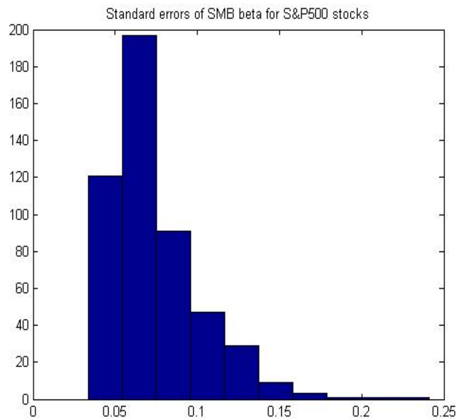

(b) Std. error SMB beta

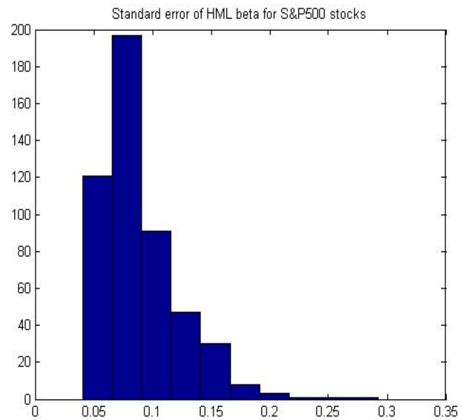

(c) Std. error HML beta

Figure 3: Standard errors of the Fama and French factor loadings based on daily data over the period January 1998 to December 2007.

Table 1: Description of ETF contracts

ticker	description	sector / style classification	# trades per day
XLE.A	Energy Sector SPDR Fund	Energy	64,110
XLB.A	Materials Sector SPDR Fund	Materials	22,423
XLI.A	Industrial Sector SPDR Fund	Industrials	12,235
XLY.A	Consumer Discretionary Sector SPDR Fund	Consumer Discretionary	11,198
XLP.A	Consumer Staples Sector SPDR Fund	Consumer Staples	5,550
XLV.A	Health Care Sector SPDR Fund	Health Care	6,353
XLF.A	Financial Sector SPDR Fund	Financials	146,853
XLK.A	Technology Sector SPDR Fund	Information Technology	9,245
IYZ.N	iShares Telecommunications Sector Fund	Telecommunications	930
XLU.A	Utilities Sector SPDR Fund	Utilities	11,544
SPY.A	SPDR Trust Series 1	Large Cap	300,104
IWM.A	iShares Russell 2000 Index Fund	Small Cap	163,148
IVE.N	S&P 500 Value Index Fund	Value	3,201
IVW.N	S&P 500 Growth Index Fund	Growth	4,526
	Average across ETFs		54,387
	Average across S&P400 constituents		2,912
	Average across S&P500 constituents		8,272
	Average across S&P600 constituents		1,411

Note: This table lists the ETF contracts used in the empirical analysis, together with the average number of trades per day over the period November 2006 through May 2008. The "SMB" ("HML") factor is specified as IWM.A - SPY.A (IVE.N - IVW.N).

by microstructure noise than assets which trade infrequently, see also Aït-Sahalia and Yu (2009) for related discussions.

We include measurement errors in the factor loadings by contaminating the DGP factor loadings with two different magnitudes of additive noise. The magnitude of the estimation errors is based on the sample size used for estimating the factor loading, i.e. we fix the standard errors of the factor loadings estimated using a sample of 10 years, see Figure 3, and then also scale these standard errors to match with a smaller sample size of 1 year to allow for higher levels of estimation error.[6]

[6] Detailed results for an infinite sample size (no estimation error) and for a sample size of 4 months (high estimation error) are available on request.

- **Case I**: full sample size (10 years, low noise), $\omega_i \sim N(0, \sigma^2_{b_i})$
- **Case II**: sample size of 1 year (intermediate noise), $\omega_i \sim N(0, 10\sigma^2_{b_i})$

The factor loadings contaminated with different levels of estimation error are then computed using $\tilde{\beta}_i = \beta_i + \omega_i$.

We use the microstructure noise contaminated non-synchronous intra-day stock returns $\mathcal{R}$ and factor returns $\mathcal{F}$ to obtain the realized factor covariance matrix $\Lambda = \mathcal{F}'\mathcal{F}$ and realized covariance $RC = \mathcal{R}'\mathcal{R}$. The residual intra-day returns ε are computed using $\varepsilon = \mathcal{R} - \mathcal{F}\tilde{\beta}$ which we use to obtain the residual variances diag($\varepsilon'\varepsilon$) and then compute the MFFM estimate using Equation (9).

3.4 Error statistics and results

We compare the realized covariance and the MFFM covariance matrix estimators based on their distance to the DGP covariance matrix Σ. Define an error matrix $X = \Sigma - \widehat{\Sigma}$, where $\widehat{\Sigma}$ is the covariance matrix estimated using RC or the MFFM. The traditional Frobenius norm is a measure of the quality of the covariance matrix based on the Euclidian distance of the covariance matrix estimates relative to the covariance matrix in the DGP. It is defined as the square-root of the sum over the squared elements in the error matrix X. This is equivalent to the square of the trace of the product of X and its conjugate transpose,

$$\|X\| = \sum_{i=1}^{M} \sum_{j=1}^{N} |x_{ij}|^2 = \sqrt{tr(XX')} \tag{10}$$

where x_{ij} denotes the matrix element of X on row i and column j, with $i = 1, \ldots, M$ and $j = 1, \ldots, N$. For the analysis we evaluate the (off-diagonal) covariance and (diagonal) variance elements separately.

Figure 4 illustrates performance of the RC and the MFFM when prices are non-synchronous but not contaminated by market microstructure noise. The covariance results illustrate that the MFFM has an excellent performance and is very robust across sampling frequencies and, in contrast to RC, its performance is not affected by non-synchronicity. Non-synchronicity, however, does affect the MFFM variance estimates which may seem counter intuitive at first thought as non-synchronicity usually affects the covariances and not so much the variances. The reason for the upward bias in the MFFM variances is caused by a mismatch between the very frequently observed factor returns and less frequently observed stock returns, which results in an additional quadratic bias term in the MFFM diagonal.[7] The mismatch between liquid factors and less liquid stocks disappears when sampling at the 5min sampling frequency and lower frequencies because the mismatch becomes smaller when sampling less frequently. Also note that when increasing the number of assets in a portfolio, the variance elements play a more limited role whereas the covariances become more dominant. Consider the simulation with 500 stocks, in this case we only have 500 variances in contrast to 249,500 covariance entries.

[7] An informal derivation of this quadratic bias term is available upon request.

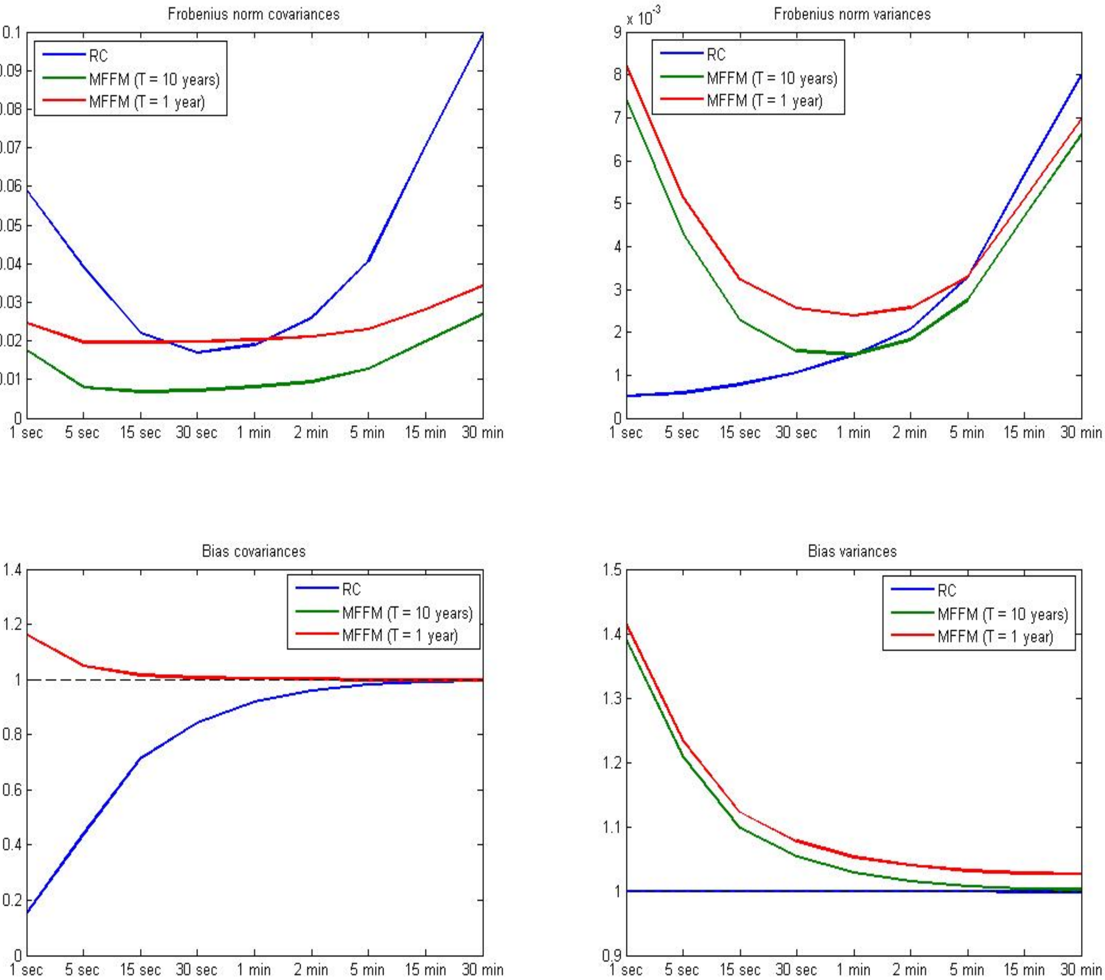

Figure 4: Error metrics for variances and covariances, high liquidity, non-synchronicity, no-noise. 10,000 repetitions.

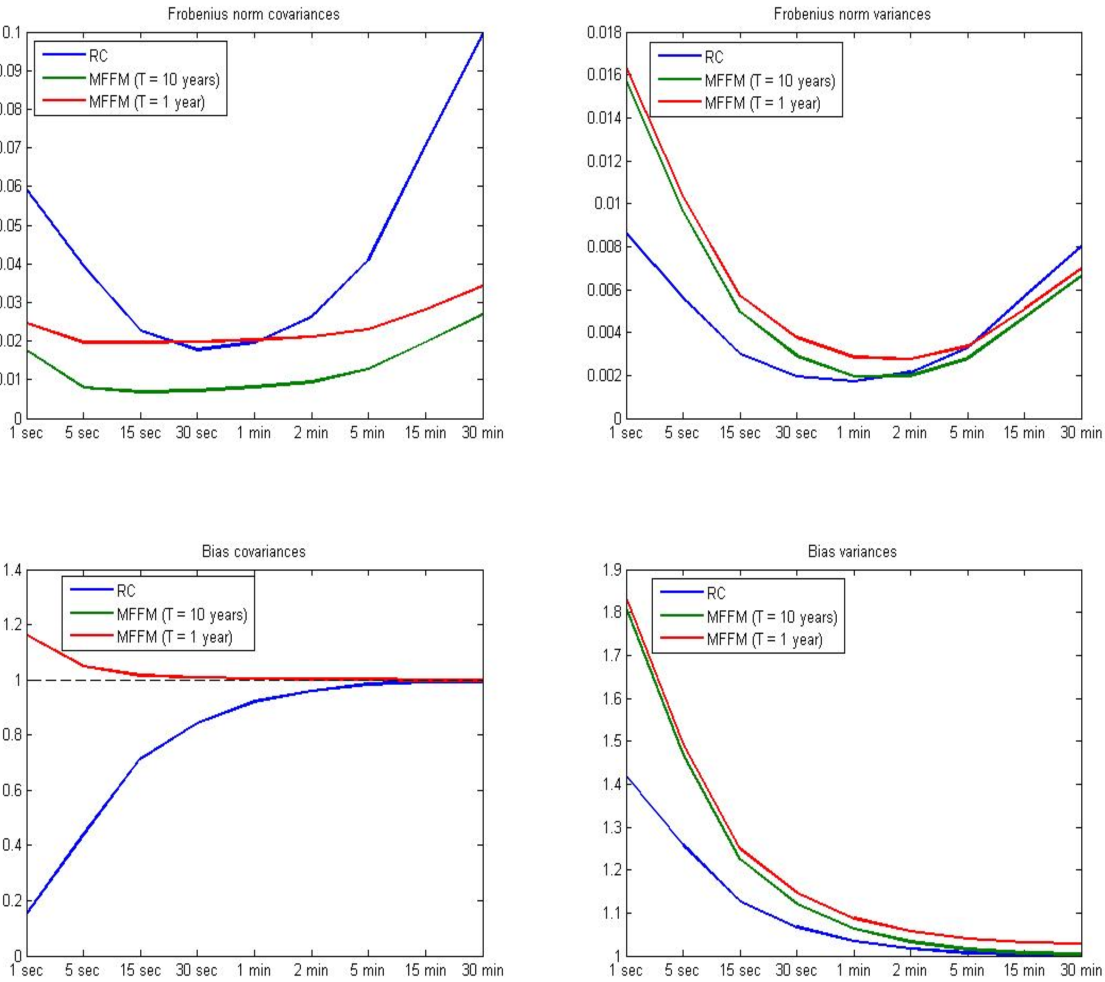

Figure 5: Error metrics for variances and covariances, high liquidity, non-synchronicity and noise. 10,000 repetitions

Figure 5 illustrates the more practically relevant case where prices are non-synchronous but also contaminated by additive market microstructure noise. Market microstructure noise does not deteriorate the performance of both covariance estimators as the noise is (assumed to be) cross-sectionally independent.[8] However, the noise does affect the variances computed with the MFFM (through the residual variances) and RC. For both estimators the diagonal elements perform fairly similar at the 5min and lower sampling frequencies while the MFFM covariances are substantially more efficient.

4 Empirical applications

In this section we evaluate the empirical out-of-sample performance of the "mixed-frequency factor-model" and the realized covariance estimator by constructing minimum tracking error portfolios and Value-at-Risk forecasts. For the minimum tracking error application we evaluate the performance of vast portfolios consisting of all stocks in each of the S&P stock universes. For VaR forecasts we consider the performance across several portfolio dimensions, specifically, we consider 100 random subsets of the stock universes that contain $S = \{25, 50, 100, 250\}$ stocks each. The selected subsets remain fixed over the sample period, across sampling frequencies and forecast weighting schemes and are equivalent for the competing estimators. We first discuss the data, the realized covariance and realized factor models, the forecasting methodology and the results.

4.1 Data

We apply the "mixed-frequency factor-model" empirically to the three S&P stock universes covering different market capitalization stocks. Our data comprise the S&P500 large caps, S&P400 mid caps and S&P600 small cap constituents.[9] We sample National Best Bid Best Offer (NBBO) mid-points, originating from NYSE and NASDAQ only, at the 15-seconds sampling frequency.

Because we use NBBO mid-points the impact of market microstructure noise is somewhat limited for the most liquid S&P500 constituents. We compute the mean autocorrelation function (ACF) of 15sec intra-day returns across all stocks in each of the three universes and across the 19 ETFs and find that on average there is some small negative first-order autocorrelation for S&P500 and S&P600

[8] In related (future) work we analyze a broader and more complex set of market microstructure frictions such as serial dependence, cross-sectional dependence, noise correlated with the efficient price, etc. We also propose methods for eliminating the variance bias in the MFFM arising from the mismatch between liquid factors and illiquid stocks as discussed above. Further we will compare the bias-adjusted MFFM versions to bias-adjusted RC estimators, such as the kernel lead-lag estimators analyzed in Barndorff-Nielsen et al. (2008), which are expected to handle non-synchronicity and noise better than RC. These market microstructure related issues are beyond the scope of the current paper.

[9] We only use those stocks that were constituents in a single index during the sample 1/5/2004 - 25/5/2008. Although we use high-frequency data from 1/11/2006 onward, we also need an additional 2.5 years of daily constituent data for estimating the factor loadings. This leaves 455 large-caps, 353 mid-caps and 515 small-caps.

stocks and the ACF for S&P400 and ETFs is positive at lag one. The mean ACFs for the three stock universes and ETFs are plotted in Figure 6 with 5%/95% bounds for stocks and 15%/85% bounds for ETFs.

The out-of-sample period we use, Jan 2007 – May 2008, is a period with very high financial market volatility compared to the preceding years and this affects the results as we will discuss in more detail in the results sections.

4.2 Realized Covariance

The RC and MFFM estimator in the simulation study are only suitable in the case of a constant covariance matrix in the DGP. For empirical data, however, it is well known that volatility and covariances are time-varying and we incorporate these features in the models.

In the portfolio Value-at-Risk forecasting exercise with N stocks we use the traditional RC estimator,

$$RC_t = \mathcal{R}_t' \mathcal{R}_t \tag{11}$$

where R_t is the $T \times N$ matrix of (intra-day) stock returns on day t.

In the minimum tracking error application we employ intra-day excess stock returns net of the benchmark, the *active* realized covariance estimator is calculated using

$$RC_t = (\mathcal{R}_t - \mathcal{R}_{Mt} \mathbf{1}_N)'(\mathcal{R}_t - \mathcal{R}_{Mt} \mathbf{1}_N) \tag{12}$$

where $\mathcal{R}_{Mt}$ is the $T \times 1$ vector of intra-day returns on the corresponding market factor, and $\mathbf{1}_N$ is a N-dimensional row vector of ones.

4.3 Mixed-frequency factor models

In the portfolio Value-at-Risk forecasting application for the MFFM we employ the Fama and French (1992) 3-factor model and a 12-factor model where the market factor is replaced by 10 industry factors. In the minimum tracking error application we only evaluate the performance of the 12 factor model because the stock returns considered there are excess returns over the market factor.

For the MFFM approach we need estimates of the factor loadings. We obtain estimates of conditional time-varying factor loadings using a simple moving window history of 2.5 years (632 days) of daily *close-to-close* return data[10] we re-estimate the betas daily using OLS. For the Fama and French 3-factor model the regressions we run are

$$R_{it} = F_t \beta_t + \varepsilon_{it} \tag{13}$$

[10]In earlier studies on factor models the number of observations used for estimating betas is usually 3 or 5 years, here we use 2.5 years as using a longer history would limit the number of constituents that survived our sample period and therefore reduces the dimension of the covariance matrix.

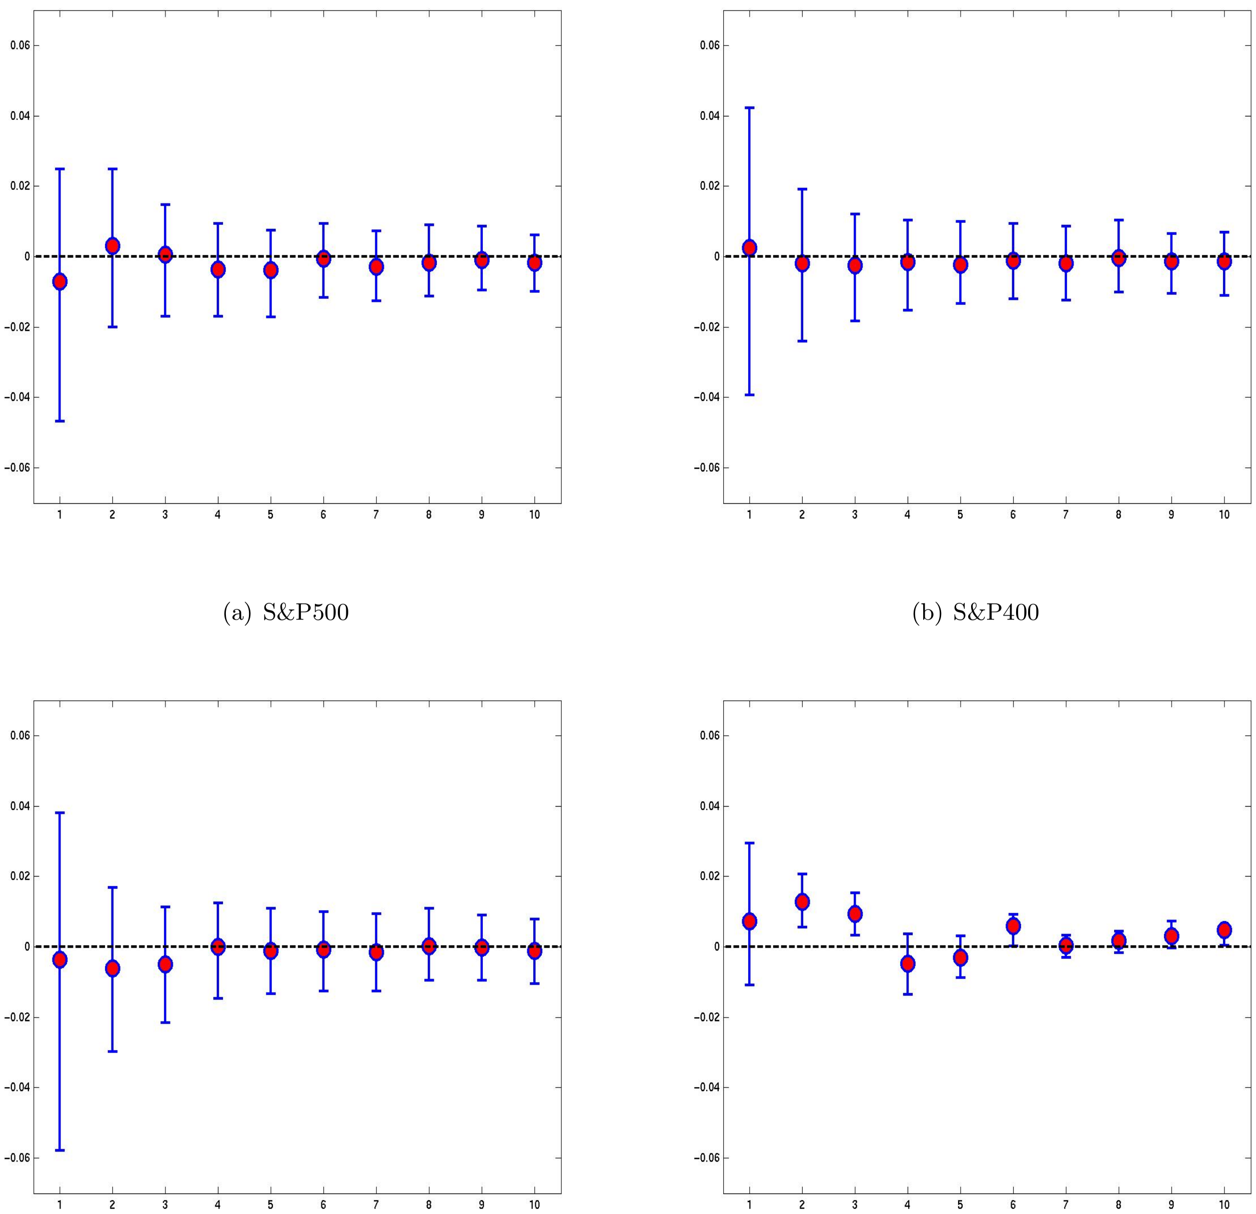

(a) S&P500

(b) S&P400

(c) S&P600

(d) ETFs

Figure 6: This Figure displays the mean autocorrelation function of 15 sec. returns for S&P500 (a), S&P400 (b), S&P600 (c) stocks with 5%-95% bounds. The autocorrelation function across 19 ETFs (d) is plotted with 15%-85% bounds.

where R_{it} is a 632×1 vector of daily returns on stock i and $F_t = [R_{Mt} \quad SMB_t \quad HML_t]$ is a 632×3 matrix of factor returns on the market, size (Small Minus Big) and value (High Minus Low) factors. The residuals needed to compute residual variances are computed using

$$\varepsilon_t = \mathcal{R}_t - \mathcal{F}_t \hat{\beta}_{t|t-1} \tag{14}$$

and the MFFM estimator covariance matrix estimate for day t is:

$$MFFM_t = \hat{\beta}'_{t|t-1} \Lambda_t \hat{\beta}_{t|t-1} + diag(\varepsilon'_t \varepsilon_t). \tag{15}$$

where $\Lambda_t = \mathcal{F}'_t \mathcal{F}_t$. For the minimum tracking error application we run similar regressions but the dependent variable is now a vector of stock returns in excess of the returns on the market factor,

$$R_{it} - R_{Mt} = F_t \beta_t + \varepsilon_{it} \tag{16}$$

where $F_t = [SMB_t \quad HML_t \quad I_1 \quad \ldots \quad I_{10}]$ is the 632×12 matrix containing returns on the size (SMB), value (HML) and 10 industry factors. The motivation to use 10 industry sectors as factors is that many stocks have exposures to several industries. Apple for example is active in many industries and therefore it is reasonable to estimate the exposures to all industry factors, see Grinold and Kahn (2000), page 60. The factor loading estimates are filtered one day ahead by using only data available at day $t-1$ for the betas on day t.

Given the filtered betas we calculate the *active* intra-day residuals that are needed for estimating the residual variances,

$$\varepsilon_t = \mathcal{R}_t - \mathcal{R}_{M,t} - \mathcal{F}_t \hat{\beta}_{t|t-1} \tag{17}$$

and compute the MFFM estimator for day t using Equation (15).

4.4 Covariance matrix forecasts

Given a sequence of unconditional covariance matrix estimates $\{\widehat{\Sigma}_t\}_{t=1}^{T}$, we produce forecasts based on an exponentially weighted moving average (EWMA) scheme motivated by the work of Foster and Nelson (1996) and Andreou and Ghysels (2002),

$$\Sigma_{t|t-1} = \alpha \Sigma_{t-1|t-2} + (1-\alpha)\widehat{\Sigma}_{t-1} \tag{18}$$

where α is a fixed scalar decay parameter, $\Sigma_{t|t-1}$ is the conditional rolling sample covariance matrix forecast for day t and $\widehat{\Sigma}_{t-1}$ is the unconditional covariance matrix estimate on day $t-1$, this innovation is either the RC or the MFFM on day $t-1$. For the forecasts we consider several weighting schemes with $\alpha \in \{0.94, 0.75, 0.50, 0.25\}$. The $\alpha = 0.94$ decay parameter is the optimal decay parameter for daily data documented by RiskMetrics. We use smaller decay parameters as a sensitivity analysis which allows us to interpret the performance when using a higher weight on more recent data. This provides more insight in the performance of the covariance estimator itself rather than the 'smoothed' forecast. Less 'smoothing' can be important also from an economic point of view, as a large level of "smoothing" will imply that the forecast adjusts less rapidly to important changes in (co)variance dynamics, which for example occur at turning points between periods of high and low volatility.

We forecast the covariance matrices one day ahead and for this EWMA scheme we use nov/dec 2006 as burn-in period for the covariance dynamics and exclude these months in the performance evaluations below. The out-of-sample period is 1/1/2007 - 25/5/2008 (351 days).

4.5 Minimum tracking error portfolios

Given the one day ahead EWMA forecasts of the covariance matrices we construct minimum TE portfolios by calculating the standard fully-invested minimum variance portfolios (when using the *active* covariance matrix as we do here, then the minimum TE portfolio is the minimum variance portfolio):

$$w_t = \frac{\Sigma_{t|t-1}^{-1}\iota}{\iota'\Sigma_{t|t-1}^{-1}\iota} \tag{19}$$

where ι is a $N \times 1$ vector of ones and Σ_t is the EWMA conditional covariance matrix forecast of RC or MFFM. The daily minimum TE portfolio returns are obtained by computing $R_{Pt} = w_t' r_t$ where r_t is the vector of daily stock returns. We calculate the *ex-post* tracking error using daily returns $\text{TE} = \text{Std}(R_P - R_M)$ and compare the results for the RC and the MFFM.

4.6 Minimum tracking error results

Table 2 illustrates the performance in terms of annualized minimum tracking errors for the S&P500 large caps. Consistent with the simulation results we find that the MFFM covariance matrix estimator is robust across sampling frequencies indicating that, in contrast to RC, the factor covariance matrix can be estimated at very high frequencies as the level of market microstructure noise and non-synchronicity in the factors is relatively small compared to that in individual stocks. At each of the considered sampling frequencies and forecast weights, the MFFM produces better results than RC. The difference between the MFFM and RC are small when we "snoop" on the sampling frequency and forecasts weights for RC but the differences are substantial on average across these settings. The covariance matrices considered here have a dimension of 455×455 and we find that at sampling frequencies of 30min and lower the RC is not well-conditioned and therefore not invertible, we indicate this with "NA".

Reconciling the levels of tracking errors with those in earlier studies:
Comparing our tracking error results with earlier studies that report tracking errors might suggest that our tracking errors are too large. We ran the following checks to find out what causes this seemingly high levels of tracking errors. The answer is two-fold. First, we are looking at a sample period with very high financial market volatility compared to earlier samples. Second, most of the earlier studies calculate the tracking errors based on monthly data rather than daily data. To

Table 2: Annualized tracking errors S&P500 (large cap) stocks

RC								
α	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.043	0.042	0.044	0.045	0.050	NA	NA	NA
0.75	0.043	0.045	0.049	0.062	0.081	NA	NA	NA
0.50	0.044	0.049	0.063	0.091	0.123	NA	NA	NA
0.25	0.049	0.056	0.087	0.129	NA	NA	NA	NA
MFFM								
α	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.040	0.040	0.040	0.040	0.040	0.040	0.040	0.040
0.75	0.040	0.040	0.040	0.040	0.040	0.040	0.040	0.041
0.50	0.040	0.040	0.040	0.040	0.040	0.040	0.040	0.043
0.25	0.040	0.040	0.040	0.040	0.040	0.041	0.041	0.046

Note: This table reports the *ex-post* annualized minimum tracking errors in percentages using 455 of the S&P500 constituents. The tracking errors are based on RC and MFFM to forecast the active covariance matrix one day ahead using EWMA forecasts over the sample period 1/1/2007 - 25/5/2008 with decay parameter α. For the MFFM we use a 12-factor model specification (size, value, and 10 industry factors). The "NA" entries indicate that the RC is not-invertible at certain combinations of sampling frequencies and weighting schemes.

illustrate this we first calculate the daily tracking errors obtained with a naïve 1/N portfolio over the sample 1/1/2004 – 25/5/2008. Most studies report, see e.g. Chan et al. (1999), tracking errors of around 3% p.a. for 1/N portfolios based on monthly returns for computing tracking errors. The tracking error we calculate for the 1/N with daily returns over 1/1/2004 – 25/5/2008 is 4.1% p.a. Now aggregating the daily returns to monthly returns and then computing the tracking error results in a tracking error of 3.3% p.a. which is close to the tracking errors of 1/N portfolios in earlier studies given also that our sample period can be classified as a high volatility sample. To illustrate the difference between the pre-2007 and post-2007 sample we calculate the tracking error over the sample 1/1/2004 – 31/12/2006 and find 3.7% with daily data and 2.9% with monthly data. Next we compute the tracking error of the 1/N portfolio using all 500 stocks over the sample 1/1/2007 – 25/5/2008 and find a tracking error of 4.8% with daily data and 3.5% with monthly data. Hence, the fact that we find a a tracking error of 5.1% based on daily data for the 1/N portfolio which uses 455 stocks is perfectly in line with the 4.8% when using all 500 stocks. Hence, the tracking errors seem high due to the use of daily rather than monthly data for computing tracking errors and the fact that financial market volatility was relatively high during the out-of-sample period, see also Figure 7 in the Appendix which displays daily tracking errors of the 1/N portfolio over the sample 1/1/2004 –

Table 3: Annualized tracking errors in percentages S&P400 (mid cap) stocks

RC								
a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.055	0.055	0.057	0.058	0.071	0.117	NA	NA
0.75	0.059	0.063	0.067	0.076	0.116	NA	NA	NA
0.50	0.065	0.074	0.085	0.111	NA	NA	NA	NA
0.25	0.072	0.090	0.111	0.168	NA	NA	NA	NA
MFFM								
a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.053	0.053	0.053	0.053	0.053	0.053	0.053	0.053
0.75	0.053	0.053	0.052	0.052	0.053	0.052	0.052	0.053
0.50	0.053	0.053	0.052	0.053	0.053	0.053	0.053	0.054
0.25	0.053	0.053	0.053	0.053	0.053	0.054	0.054	0.058

Note: This table reports the *ex-post* annualized minimum tracking errors in percentages. The tracking errors are based on RC and MFFM to forecast the active covariance matrix one day ahead using EWMA forecasts over the sample period 1/1/2007 - 25/5/2008 with decay parameter α. For the MFFM we use a 12-factor model specification (size, value, and 10 industry factors). The "NA" entries indicate that the RC is not-invertible at certain combinations of sampling frequencies and weighting schemes.

25/5/2008 after applying the RiskMetrics "smoother". We do not calculate the tracking errors with monthly data as we only have a sample of 17 months.

For the S&P400, see Table 3, we decrease the naïve 1/N tracking error of 6.0% to 5.5% with RC and this result depends heavily on the forecast weighting scheme and sampling frequency. Using the MFFM further decreases the tracking error to 5.2% with results being robust. As expected, the tracking errors have increased for the S&P400 mid caps compared to the S&P500 large caps, see also Table 1 for the average number of trades per day in each S&P universe. The fact that we have less available stocks to track the benchmark and the higher levels of non-synchronicity and microstructure noise in individual stocks explain this result.

For the S&P600 small caps, where the level of non-synchronicity plays a more important role than for mid- and large-caps, we find larger tracking errors, as expected. The MFFM manages a tracking error of 6.1% and comfortably outperforms the 1/N portfolio which achieves a tracking error 6.5%. Similar to the results for the mid- and large-caps the "best" RC result, being a tracking error of 9.5%, is not robust across sampling frequencies and forecast weighting schemes, in fact it is substantially worse than the naïve 1/N portfolio. Also the optimal sampling frequency of 15sec points towards this sampling frequency being optimal because the impact of non-synchronicity is largest at the 15sec frequency. Hence, the bias towards zero in covariances is especially large at the 15sec frequency

Table 4: Annualized tracking errors in percentages S&P600 (small-cap) stocks

RC

α	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.095	0.097	0.097	0.105	0.112	NA	NA	NA
0.750	0.103	0.107	0.118	0.138	0.182	NA	NA	NA
0.500	0.111	0.122	0.161	0.201	NA	NA	NA	NA
0.250	0.121	0.146	0.220	NA	NA	NA	NA	NA
MFFM								
α	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.062	0.062	0.061	0.061	0.061	0.061	0.061	0.062
0.75	0.062	0.062	0.061	0.061	0.062	0.062	0.062	0.063
0.5	0.062	0.062	0.061	0.061	0.062	0.062	0.063	0.064
0.25	0.062	0.062	0.061	0.062	0.062	0.063	0.064	0.068

Note: This table reports the *ex-post* annualized minimum tracking errors in percentages. The tracking errors are based on RC and MFFM to forecast the active covariance matrix one day ahead using EWMA forecasts over the sample period 1/1/2007 - 25/5/2008 with decay parameter α. For the MFFM we use a 12-factor model specification (size, value, and 10 industry factors). The "NA" entries indicate that the RC is not-invertible at certain combinations of sampling frequencies and weighting schemes.

causing it to be optimal due to (limited) shrinkage towards the equal weighted portfolio which can be obtained by setting all covariances equal to zero.[11]

Remark that outperforming the equally-weighted portfolio is not necessarily an easy task. DeMiguel et al. (2009) analyze various advanced methods consisting of Bayesian estimation, shrinkage, robust allocation etc. and find that none of the 14 models they implement can consistently outperform the 1/N portfolio. Hence, the fact that the MFFM consistently outperforms the 1/N and RC portfolios is encouraging support. Further, the results in Madhavan and Yang (2003) illustrate that using the sample (realized) covariance matrix for unrestricted optimization, results in a performance that is worse than the equal-weighted portfolio. Imposing a no short-sales restriction, however, halves the tracking error relative to not using such constraint in their study. This is consistent with the results in Chan et al. (1999) who impose the no short-sales constraint and an upper-bound of 2% on the weights to obtain better results with the sample (realized) covariance matrix.

[11]Setting all covariances equal to zero makes the portfolio equal-weighted but not necessarily fully-invested. Using identical variances and setting covariances to zero makes the portfolio equal-weighted and fully-invested.

4.7 Value-at-Risk Forecasts

We evaluate the empirical out-of-sample performance of Value-at-Risk (VaR) forecasts for portfolios consisting of subsets of the S&P500 universe. The VaR of a portfolio can be interpreted as a loss in terms of a portfolio return for which we are $100(1-\alpha)\%$ certain that it will not be exceeded. This measure is appealing for practitioners who wish to communicate the down-side risk of a portfolio in terms of a single number.

In the VaR application we assume the portfolio manager holds an equal-weighted portfolio of $S \in \{25, 50, 100, 250\}$ stocks and computes one-day ahead EWMA covariance matrix forecasts using Equation (18) to forecast her portfolio volatility,

$$\widehat{\sigma}_{P,t} = \sqrt{w' \Sigma_{t|t-1} w} \tag{20}$$

which we use to calculate the one day ahead VaR prediction. We assume that the distribution of portfolio returns is approximately Gaussian and heteroscedastic, i.e. $R_{P,t} \sim N(0, \sigma^2_{P,t})$ such that the VaR forecast is $\widehat{VaR}_{t|t-1} = Z_\alpha \widehat{\sigma}_{P,t}$, where Z_α is obtained from the left tail of the normal distribution, i.e. $Z_{0.01} = -2.33$ and $Z_{0.05} = -1.645$.

We test for correct (un)conditional coverage and independence using the approach of Christoffersen (1998). Given the one-day ahead $100(1-\alpha)\%$ portfolio VaR forecasts, we keep track of the ex-post number of exceedances using an indicator variable,

$$X_t = \begin{cases} 1 & \text{if } R_{P,t} \leq VaR_{t|t-1} \\ 0 & \text{elsewhere} \end{cases} \tag{21}$$

Let $n = \sum_{t=1}^{T} X_t$ denote the number of exceedances in a sample of size T, we can then test the null of correct unconditional coverage using the likelihood under the null, $L_0 = \alpha^n (1-\alpha)^{T-n}$ against the alternative $L_1 = (\frac{n}{T})^n (1 - \frac{n}{T})^{T-n}$, by using a likelihood ratio test statistic $LR_{uc} = 2[\log(L_1) - \log(L_0)] \sim \chi^2_1$.

To test for (time-series) dependence in exceedances we calculate the likelihood under the alternative of first-order dependence, $L_2 = (1-\pi_{01})^{N_{00}} \pi_{01}^{N_{01}} (1-\pi_{11})^{N_{10}} \pi_{11}^{N_{11}}$, where π_{ij} is the ex-post probability of observing state j after state i. The number of states i followed by state j is denoted $N_{ij} = \pi_{ij} T$. Given L_2 we can now test the null of time-series independence in exceedances $LR_{ind} = 2[\log(L_2) - \log(L_1)] \sim \chi^2_1$. Aggregating the likelihood ratio test statistics for unconditional coverage and independence is equivalent to testing for conditional coverage, $LR_{cc} = LR_{uc} + LR_{ind} \sim \chi^2_2$.

4.8 Portfolio Value-at-Risk Forecasting Results

Table 5 and 6 summarize the percentage of 5% VaR and 1% VaR exceedances obtained by forecasting the portfolio VaR using the MFFM 12-factor model and RC, respectively. The results for the Fama and French (1992) 3-factor model are collected in the Appendix, the 12-factor model produces less exceedances, but the 3-factor model also easily outperforms RC. For each subset size of the universe we consider {25,50,100,250} we use 100 randomly selected portfolios that remain fixed over the sample period, across forecasting schemes and are of course equivalent for the MFFM and RC.

We find that the MFFM 3- and 12-factor models produce better VaR forecasts compared to RC across all intra-day sampling frequencies, forecasting schemes and different VaR levels. A somewhat counter intuitive result is that in general the number of exceedances increases with the number of stocks in the portfolio. This is explained by the fact that in our sample period market volatility is (almost monotonically) increasing. When using more stocks the portfolio volatility (ex-ante and ex-post) decreases, ceteris paribus. However, the ex-ante volatility decreases more rapidly than the ex-post volatility which for the market as a whole and for random subsets is only increasing during the sample, hence the difference between predicted and ex-post volatility increases with the number of stocks in a portfolio and thus the number of exceedances increases as well causing more rejections of accurate (un)conditional coverage and time-series independence. For this reason we only report results of the statistical likelihood ratio tests on (un)conditional coverage and independence for the 100 portfolios of 25 stocks and only report tests based on a 5% significance level.[12]

[12]Results for portfolios of 50, 100 and 250 stocks, and tests based on the 1% significance level are qualitatively similar in the sense that MFFM produces better results than RC but the tests for accurate (un)conditional coverage and time-series independence are of course rejected more frequently as follows directly from Table 5 and 6. These results are available upon request.

Table 5: Percentage of Portfolio VaR exceedances over the sample period 1/1/2007 – 25/5/2008

MFFM 12-F Model

MFFM (5% VaR exceedances, 25 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.069	0.064	0.060	0.060	0.062	0.061	0.063	0.081
0.75	0.073	0.068	0.065	0.066	0.067	0.067	0.070	0.094
0.5	0.080	0.076	0.072	0.074	0.074	0.077	0.082	0.117
0.25	0.084	0.079	0.077	0.079	0.078	0.085	0.089	0.142

MFFM (1% VaR exceedances, 25 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.027	0.024	0.022	0.023	0.023	0.024	0.023	0.036
0.750	0.025	0.022	0.021	0.022	0.022	0.024	0.023	0.048
0.500	0.029	0.026	0.026	0.027	0.027	0.031	0.034	0.063
0.250	0.033	0.030	0.030	0.032	0.032	0.039	0.044	0.084

MFFM (5% VaR exceedances, 50 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.081	0.073	0.067	0.067	0.070	0.069	0.071	0.084
0.75	0.087	0.081	0.075	0.076	0.078	0.078	0.081	0.101
0.5	0.093	0.087	0.082	0.084	0.085	0.087	0.092	0.125
0.25	0.095	0.089	0.086	0.088	0.087	0.094	0.098	0.150

MFFM (1% VaR exceedances, 50 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.035	0.031	0.028	0.028	0.028	0.029	0.028	0.039
0.750	0.032	0.028	0.026	0.027	0.027	0.030	0.032	0.051
0.500	0.038	0.033	0.032	0.034	0.033	0.038	0.043	0.070
0.250	0.042	0.037	0.036	0.039	0.039	0.047	0.055	0.095

MFFM (5% VaR exceedances, 100 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.090	0.081	0.073	0.073	0.076	0.075	0.079	0.085
0.75	0.096	0.089	0.082	0.084	0.085	0.084	0.087	0.103
0.5	0.101	0.096	0.091	0.092	0.091	0.094	0.097	0.131
0.25	0.101	0.095	0.092	0.094	0.091	0.099	0.103	0.154

MFFM (1% VaR exceedances, 100 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.041	0.037	0.033	0.033	0.033	0.033	0.033	0.041
0.750	0.038	0.033	0.030	0.032	0.030	0.033	0.038	0.055
0.500	0.043	0.037	0.035	0.037	0.036	0.042	0.050	0.075
0.250	0.048	0.042	0.042	0.045	0.044	0.054	0.063	0.103

MFFM (5% VaR exceedances, 250 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.096	0.086	0.077	0.077	0.080	0.080	0.083	0.086
0.75	0.102	0.095	0.090	0.092	0.092	0.089	0.093	0.105
0.5	0.106	0.100	0.097	0.097	0.096	0.097	0.100	0.136
0.25	0.104	0.100	0.096	0.096	0.092	0.100	0.105	0.160

MFFM (1% VaR exceedances, 250 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.042	0.040	0.035	0.034	0.034	0.034	0.034	0.042
0.750	0.039	0.035	0.033	0.036	0.032	0.034	0.040	0.055
0.500	0.046	0.038	0.037	0.041	0.037	0.045	0.055	0.076
0.250	0.051	0.043	0.044	0.048	0.046	0.057	0.066	0.111

Table 6: Percentage of Portfolio VaR exceedances over the sample period 1/1/2007 – 25/5/2008

Realized Covariance

RC (5% VaR exceedances, 25 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.102	0.086	0.078	0.077	0.079	0.076	0.079	0.071
0.75	0.106	0.090	0.082	0.082	0.083	0.083	0.085	0.085
0.5	0.110	0.094	0.087	0.088	0.088	0.091	0.097	0.108
0.25	0.114	0.098	0.091	0.094	0.092	0.100	0.107	0.136

RC (1% VaR exceedances, 25 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.047	0.037	0.033	0.033	0.033	0.033	0.033	0.027
0.750	0.047	0.035	0.031	0.031	0.031	0.034	0.036	0.038
0.500	0.054	0.042	0.038	0.038	0.038	0.044	0.049	0.058
0.250	0.060	0.047	0.043	0.043	0.043	0.053	0.058	0.085

RC (5% VaR exceedances, 50 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.109	0.093	0.084	0.083	0.085	0.081	0.085	0.073
0.75	0.113	0.098	0.088	0.088	0.089	0.088	0.092	0.088
0.5	0.116	0.100	0.092	0.094	0.094	0.096	0.102	0.112
0.25	0.119	0.103	0.096	0.099	0.098	0.104	0.112	0.140

RC (1% VaR exceedances, 50 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.052	0.041	0.036	0.036	0.035	0.036	0.036	0.028
0.750	0.053	0.038	0.033	0.033	0.032	0.036	0.039	0.040
0.500	0.061	0.046	0.041	0.041	0.040	0.047	0.052	0.060
0.250	0.066	0.052	0.047	0.046	0.046	0.056	0.063	0.089

RC (5% VaR exceedances, 100 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.112	0.098	0.088	0.087	0.089	0.085	0.089	0.075
0.75	0.117	0.101	0.090	0.090	0.092	0.090	0.093	0.089
0.5	0.118	0.104	0.097	0.098	0.097	0.099	0.105	0.114
0.25	0.121	0.104	0.100	0.103	0.101	0.105	0.113	0.143

RC (1% VaR exceedances, 100 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.055	0.043	0.039	0.037	0.038	0.038	0.039	0.029
0.750	0.058	0.041	0.034	0.035	0.034	0.036	0.042	0.043
0.500	0.069	0.049	0.041	0.042	0.041	0.049	0.056	0.063
0.250	0.072	0.057	0.050	0.046	0.047	0.059	0.066	0.094

RC (5% VaR exceedances, 250 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.114	0.101	0.091	0.092	0.091	0.087	0.091	0.074
0.75	0.116	0.106	0.094	0.095	0.097	0.093	0.097	0.090
0.5	0.119	0.106	0.100	0.101	0.100	0.099	0.106	0.116
0.25	0.121	0.105	0.102	0.105	0.104	0.105	0.114	0.145

RC (1% VaR exceedances, 250 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.055	0.044	0.039	0.038	0.038	0.039	0.040	0.027
0.750	0.062	0.042	0.035	0.036	0.033	0.036	0.042	0.042
0.500	0.072	0.053	0.043	0.041	0.040	0.051	0.055	0.064
0.250	0.074	0.058	0.053	0.047	0.047	0.062	0.066	0.099

Table 7: MFFM 12-factor model 5% and 1% VaR test results for accurate (un)conditional coverage and time-series independence, reported is the number of rejections found for the 100 considered portfolios based on a significance level of 5% and portfolios consisting of 25 stocks

Unconditional Coverage

MFFM (5% VaR, number of rejections 25 Stocks)									MFFM (1% VaR, number of rejections 25 Stocks)								
a	15s	1m	5m	15m	30m	65m	130m	C2C	a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	27	5	2	1	3	1	3	68	0.94	80	66	54	59	59	64	60	99
0.75	44	18	4	9	9	8	24	100	0.75	71	47	41	48	56	68	60	100
0.5	70	52	32	44	36	60	85	100	0.5	89	74	73	82	83	95	98	100
0.25	82	71	62	65	66	90	97	100	0.25	97	95	93	95	95	99	100	100

Independence

MFFM (5% VaR, number of rejections 25 Stocks)									MFFM (1% VaR, number of rejections 25 Stocks)								
a	15s	1m	5m	15m	30m	65m	130m	C2C	a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	14	14	2	3	6	1	5	12	0.94	0	0	0	0	0	0	0	0
0.75	37	29	15	14	20	14	20	8	0.75	0	0	0	0	0	0	0	0
0.5	69	61	40	30	40	18	28	1	0.5	0	0	0	0	0	0	0	8
0.25	72	65	52	37	55	18	24	3	0.25	0	0	0	0	0	0	0	53

Conditional coverage

MFFM (5% VaR, number of rejections 25 Stocks)									MFFM (1% VaR, number of rejections 25 Stocks)								
a	15s	1m	5m	15m	30m	65m	130m	C2C	a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	16	14	2	3	6	1	5	25	0.94	68	50	37	34	38	46	38	97
0.75	40	31	15	14	20	14	20	69	0.75	48	33	25	32	28	46	41	100
0.5	80	66	41	33	46	28	53	99	0.5	75	55	58	65	74	91	91	100
0.25	85	76	58	51	60	51	78	100	0.25	95	90	86	92	92	99	100	100

Table 8: RC 5% and 1% VaR test results for accurate (un)conditional coverage and time-series independence, reported is the number of rejections found for the 100 considered portfolios based on a significance level of 5% and portfolios consisting of 25 stocks

Unconditional Coverage

	RC (5% VaR, number of rejections 25 Stocks)									RC (1% VaR, number of rejections 25 Stocks)							
a	15s	1m	5m	15m	30m	65m	130m	C2C	a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	100	95	65	64	75	54	73	24	0.94	100	100	99	99	98	100	100	86
0.75	100	98	79	81	86	82	94	97	0.75	100	99	98	98	99	100	100	100
0.50	100	99	96	97	98	99	100	100	0.50	100	100	100	100	100	100	100	100
0.25	100	100	98	100	99	100	100	100	0.25	100	100	100	100	100	100	100	100

Independence

	RC (5% VaR, number of rejections 25 Stocks)									RC (1% VaR, number of rejections 25 Stocks)							
a	15s	1m	5m	15m	30m	65m	130m	C2C	a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	2	11	19	16	18	12	16	18	0.94	0	0	0	0	0	0	0	0
0.75	10	44	37	26	33	19	26	29	0.75	1	0	0	0	0	0	0	0
0.50	19	42	41	18	27	13	19	12	0.50	3	1	0	0	0	1	0	1
0.25	17	49	40	17	27	9	18	14	0.25	11	1	0	0	0	1	2	70

Conditional coverage

	RC (5% VaR, number of rejections 25 Stocks)									RC (1% VaR, number of rejections 25 Stocks)							
a	15s	1m	5m	15m	30m	65m	130m	C2C	a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	78	45	27	22	27	17	26	18	0.94	100	99	93	93	92	92	94	75
0.75	98	83	59	51	55	56	66	61	0.75	100	95	88	89	92	100	99	100
0.50	100	95	77	74	64	79	95	97	0.50	100	100	99	100	100	100	100	100
0.25	100	97	83	89	82	89	100	100	0.25	100	100	100	100	100	100	100	100

5 Conclusion

Recently there has been great interest in the use of high-frequency data to estimate variances and covariances. On the one hand this results in more accurate covariance estimates, but on the other hand it also brings problems such as non-synchronous trading leading to covariance estimates being biased towards zero. What so far has been lacking is to bring the merits of high-frequency data to factor models. With the introduction of exchange-traded funds important factors are now traded much more actively than individual stocks. For example the S&P500 ETFs (Spiders) have on average traded 36 times more frequently than the average individual stock in the S&P500. In this study we have proposed the Mixed Frequency Factor Model. In particular we use very high-frequency data for ETFs to obtain a very accurate estimate of the factor covariance matrix, and we use daily data to estimate the factor loadings. Thus we avoid non-synchronicity problems inherent in the use of high-frequency data for individual stocks. Furthermore we take advantage of the facts that factor models can easily be applied to vast numbers of assets and that covariance matrices from factor models are less prone to error maximization in portfolio construction problems. In a minimum tracking error application we reduce the tracking errors by using the MFFM rather than RC for computing the covariance matrix. The differences between RC and MFFM increase with the level of non-synchronicity between individual stocks, i.e. we find a larger difference when considering the S&P600 small caps than when we consider the S&P500 large caps. The RC outperforms the naïvely diversified equal-weighted 1/N portfolios when considering large- and mid-caps but fails by a substantial margin when we consider the illiquid S&P600 small caps. The MFFM however comfortably outperforms the 1/N portfolios regardless of the universe considered. For the S&P500 large-caps we also find that MFFM produces better Value-at-Risk forecasts resulting in substantially less exceedances leading to less rejections of accurate (un)conditional coverage and time-series independence. For realized covariance the results in both the empirical applications depend severely on the sampling frequency and the weighting scheme applied to the past daily covariance matrices. In contrast, the performance of the MFFM is robust across sampling frequencies and weighting schemes and consistently outperforms RC and the naïve 1/N portfolios.

References

Aït-Sahalia, Y. and Yu, J.: 2009, High frequency market microstructure noise estimates and liquidity measures, *The Annals of Applied Statistics* **3**(1), 422–457.

Andersen, T. G., Bollerslev, T., Christoffersen, P. and Diebold, F.: 2006, Volatility and correlation forecasting, *in Graham Elliott, Clive W. J. Granger and Allan Timmermann (eds.), Handbook of Economic Forecasting, Elsevier, North-Holland,* p. 777878.

Andreou, E. and Ghysels, E.: 2002, Rolling-sample volatility estimators: Some new theoretical, simulation and empirical results, *Journal of Business & Economic Statistics* **20**(3), 363–376.

Bandi, F. and Russell, J.: 2005, Realized covariation, realized beta, and microstructure noises, *Working paper* .

Barndorff-Nielsen, O., Hansen, P., Lunde, A. and Shephard, N.: 2008, Multivariate realised kernels: consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading.

Barndorff-Nielsen, O. and Shephard, N.: 2004, Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics, *Econometrica* **72**(3), 885–925.

Cavaglia, S., Brightman, C. and Aked, M.: 2000, The increasing importance of industry factors, *Financial Analysts Journal* **56**(4), 41–54.

Chamberlain, G. and Rothschild, M.: 1983, Arbitrage, factor structure, and mean-variance analysis on large asset markets, *Econometrica* **51**(5), 1281–1304.

Chan, L., Karceski, J. and Lakonishok, J.: 1999, On portfolio optimization: Forecasting covariances and choosing the risk model, *Review of Financial Studies* **12**(Winter), 937–974.

Christoffersen, P.: 1998, Evaluating interval forecasts, *International Economic Review* **39**, 841–862.

Connor, G. and Korajczyk, R.: 1994, A test for the number of factors in an approximate factor model, *Journal of Finance* **48**(4), 1263–1291.

DeMiguel, V., Garlappi, L. and Uppal, R.: 2009, Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?, *Review of Financial Studies* **22**, 1915–1953.

Engle, R., Shephard, N. and Sheppard, K.: 2008, Fitting and testing vast dimensional time-varying covariance models, *Working paper* .

Fama, E. and French, K.: 1992, The cross-section of expected stock returns, *Journal of Finance* **47**(2), 427–465.

Fan, J., Fan, Y. and Lv, J.: 2008, High dimensional covariance matrix estimation using a factor model, *Journal of Econometrics* **147**, 186–197.

Foster, D. and Nelson, D.: 1996, Continuous record asymptotics for rolling sample variance estimators, *Econometrica* **64**(3), 139–174.

Griffin, J. and Oomen, R.: 2006, Covariance measurement in the presence of non-synchronous trading and market microstructure noise, *Available at: http://ssrn.com/abstract=912541* .

Grinold, R. and Kahn, R.: 2000, Active portfolio management, a quantitative approach for producing superior returns and controlling risk, *2nd Edition, McGraw-Hill* .

Hayashi, T. and Yoshida, N.: 2005, On covariance estimation of non-synchronously observed diffusion processes, *Bernoulli* **11**, 359—-379.

Ingersoll, J.: 1984, Some results in the theory of arbitrage pricing, *Journal of Finance* **39**(4), 1021–1039.

Jagannathan, R. and Ma, T.: 2003, Risk reduction in large portfolios: Why imposing the wrong constraints helps, *Journal of Finance* **58**(4), 1651–1684.

Madhavan, A. and Yang, J.: 2003, Practical risk management, *Journal of Portfolio Management* **30**, 73–85.

Oomen, R.: 2006, Properties of realized variance under alternative sampling schemes, *Journal of Business & Economic Statistics* **24**(2), 219–237.

Ross, S.: 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* **13**, 341–360.

A Proofs

Proof of Theorem 2.1 Using the notation $\widehat{X} = X + X^\varepsilon$, we have for $i \neq j$:

$$\widehat{\gamma}_{ij} = \beta_i'\Lambda\beta_j + \beta_i'\Lambda\beta_j^\varepsilon + \beta_i'\Lambda^\varepsilon\beta_j + \beta_i'\Lambda^\varepsilon\beta_j^\varepsilon + \beta_i^{\varepsilon\prime}\Lambda\beta_j + \beta_i^{\varepsilon\prime}\Lambda\beta_j^\varepsilon + \beta_i^{\varepsilon\prime}\Lambda^\varepsilon\beta_j + \beta_i^{\varepsilon\prime}\Lambda^\varepsilon\beta_j^\varepsilon.$$

From assumptions (ii-iv) unbiasedness directly follows, i.e. $E(\widehat{\gamma}_{ij}) = \gamma_{ij}$. Moreover, with assumption (i) we have that $\beta_i^\varepsilon \perp \beta_j^\varepsilon$ so that all terms in the above expression are mutually uncorrelated. Next, using assumptions (ii-iv), we have:

$$\begin{aligned}
V(\beta_i'\Lambda\beta_j^\varepsilon) &= \beta_i'\Lambda\Sigma_{\beta,j}\Lambda'\beta_i \\
V(\beta_i'\Lambda^\varepsilon\beta_j) &= E(\beta_i'\Lambda^\varepsilon\beta_j\beta_j'\Lambda^{\varepsilon\prime}\beta_i) = E(\text{tr}(\beta_i\beta_i'\Lambda^\varepsilon\beta_j\beta_j'\Lambda^{\varepsilon\prime})) = g(\beta_i\beta_i', \beta_j\beta_j', \Phi) \\
V(\beta_i'\Lambda^\varepsilon\beta_j^\varepsilon) &= E(\beta_i'\Lambda^\varepsilon\beta_j^\varepsilon{\beta_j^\varepsilon}'\Lambda^{\varepsilon\prime}\beta_i) = E(\beta_i'\Lambda^\varepsilon\Sigma_{\beta,j}\Lambda^{\varepsilon\prime}\beta_i) = E(\text{tr}(\beta_i\beta_i'\Lambda^\varepsilon\Sigma_{\beta,j}\Lambda^{\varepsilon\prime})) = g(\beta_i\beta_i', \Sigma_{\beta,j}, \Phi) \\
V(\beta_i^{\varepsilon\prime}\Lambda\beta_j) &= \beta_j'\Lambda'\Sigma_{\beta,i}\Lambda\beta_j \\
V(\beta_i^{\varepsilon\prime}\Lambda\beta_j^\varepsilon) &= E(\beta_i^{\varepsilon\prime}\Lambda\beta_j^\varepsilon\beta_j^{\varepsilon\prime}\Lambda'\beta_i^\varepsilon) = E(\beta_i^{\varepsilon\prime}\Lambda\Sigma_{\beta,j}\Lambda'\beta_i^\varepsilon) = \text{tr}(\Sigma_{\beta,i}\Lambda\Sigma_{\beta,j}\Lambda') \\
V(\beta_i^{\varepsilon\prime}\Lambda^\varepsilon\beta_j) &= E(\beta_j'\Lambda^\varepsilon\Sigma_{\beta,i}\Lambda^{\varepsilon\prime}\beta_j) = E(\text{tr}(\beta_j\beta_j'\Lambda^\varepsilon\Sigma_{\beta,i}\Lambda^{\varepsilon\prime})) = g(\beta_j\beta_j', \Sigma_{\beta,i}, \Phi) \\
V(\beta_i^{\varepsilon\prime}\Lambda^\varepsilon\beta_j^\varepsilon) &= E(\beta_i^{\varepsilon\prime}\Lambda^\varepsilon\beta_j^\varepsilon\beta_j^{\varepsilon\prime}\Lambda^{\varepsilon\prime}\beta_i^\varepsilon) = E(\beta_i^{\varepsilon\prime}\Lambda^\varepsilon\Sigma_{\beta,j}\Lambda^{\varepsilon\prime}\beta_i^\varepsilon) = E(\text{tr}(\Sigma_{\beta,i}\Lambda^\varepsilon\Sigma_{\beta,j}\Lambda^{\varepsilon\prime})) = g(\Sigma_{\beta,i}, \Sigma_{\beta,j}, \Phi)
\end{aligned}$$

All terms involving Λ^ε are of the form $E(\text{tr}(AZBZ))$ where A, B, and Z are square symmetric matrices of equal dimension with A and B are fixed and Z random with $E(\text{vech}(Z)\text{vech}(Z)') = \Phi$. Define $\overline{A} = AZ$ and $\overline{B} = BZ$ with

$$\overline{A}_{ij} = \sum_k A_{ik}Z_{kj} \qquad \text{and} \qquad \overline{B}_{ij} = \sum_m B_{im}Z_{mj}.$$

Then

$$E(\text{tr}(AZBZ)) = \text{tr}(E(\overline{A}\,\overline{B})) = \sum_{i,j} E(\overline{A}_{ij}\overline{B}_{ji}) = \sum_{i,j,k,m} A_{ik}B_{jm}E(Z_{kj}Z_{mi}) = \sum_{i,j,k,m} A_{ik}B_{jm}\Phi_{f(k,j),f(m,i)}.$$

■

Proof of Corollary 2.2 Let x be multivariate normal with characteristic function $\ln\phi(\xi) = -\xi'\Sigma\xi/2$. The covariance estimator is unbiased, i.e.

$$\widehat{\sigma}_{mn} \equiv \frac{1}{M}\sum_{i=1}^{M} x_i^{(m)}x_i^{(n)} \qquad \text{where} \qquad E(\widehat{\sigma}_{mn}) = \sigma_{mn},$$

and we have:

$$E(\widehat{\sigma}_{mn}\widehat{\sigma}_{pq}) = \sigma_{mn}\sigma_{pq} + \frac{\sigma_{mp}\sigma_{nq} + \sigma_{mq}\sigma_{np}}{M}.$$

■

B Tables

C Figures

Table 9: Percentage of Portfolio VaR exceedances over the sample period 1/1/2007 – 25/5/2008

MFFM 3-F Model

MFFM (5% VaR exceedances, 25 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.069	0.067	0.065	0.067	0.069	0.067	0.070	0.090
0.75	0.074	0.072	0.070	0.073	0.074	0.075	0.079	0.106
0.5	0.081	0.079	0.078	0.080	0.081	0.085	0.090	0.126
0.25	0.083	0.082	0.081	0.084	0.084	0.090	0.096	0.151

MFFM (1% VaR exceedances, 25 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.029	0.027	0.026	0.027	0.027	0.028	0.027	0.046
0.750	0.026	0.024	0.024	0.025	0.025	0.027	0.029	0.059
0.500	0.029	0.028	0.028	0.030	0.031	0.034	0.040	0.078
0.250	0.033	0.032	0.032	0.035	0.036	0.044	0.051	0.097

MFFM (5% VaR exceedances, 50 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.079	0.075	0.071	0.074	0.077	0.075	0.078	0.094
0.75	0.085	0.083	0.080	0.083	0.083	0.086	0.090	0.112
0.5	0.090	0.088	0.086	0.089	0.089	0.094	0.099	0.133
0.25	0.092	0.090	0.089	0.092	0.091	0.098	0.105	0.158

MFFM (1% VaR exceedances, 50 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.034	0.032	0.030	0.031	0.031	0.032	0.032	0.049
0.750	0.032	0.029	0.028	0.030	0.030	0.033	0.036	0.063
0.500	0.036	0.034	0.034	0.036	0.036	0.041	0.050	0.086
0.250	0.040	0.038	0.039	0.042	0.043	0.054	0.062	0.110

MFFM (5% VaR exceedances, 100 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.086	0.082	0.077	0.080	0.083	0.082	0.085	0.096
0.75	0.092	0.090	0.086	0.089	0.090	0.092	0.095	0.114
0.5	0.096	0.094	0.091	0.095	0.093	0.098	0.103	0.136
0.25	0.097	0.094	0.094	0.097	0.095	0.100	0.110	0.160

MFFM (1% VaR exceedances, 100 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.038	0.036	0.035	0.035	0.035	0.035	0.037	0.052
0.750	0.035	0.033	0.032	0.034	0.033	0.035	0.041	0.067
0.500	0.039	0.037	0.037	0.039	0.040	0.045	0.057	0.093
0.250	0.045	0.043	0.044	0.046	0.050	0.062	0.067	0.119

MFFM (5% VaR exceedances, 250 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.94	0.090	0.087	0.083	0.085	0.087	0.087	0.089	0.097
0.75	0.100	0.096	0.092	0.094	0.095	0.098	0.100	0.113
0.5	0.102	0.098	0.096	0.098	0.097	0.100	0.108	0.138
0.25	0.101	0.097	0.095	0.099	0.098	0.103	0.115	0.165

MFFM (1% VaR exceedances, 250 Stocks)

a	15s	1m	5m	15m	30m	65m	130m	C2C
0.940	0.039	0.038	0.035	0.036	0.036	0.036	0.038	0.055
0.750	0.035	0.034	0.034	0.035	0.033	0.035	0.045	0.070
0.500	0.043	0.039	0.040	0.041	0.043	0.048	0.061	0.098
0.250	0.050	0.046	0.049	0.050	0.054	0.068	0.069	0.124

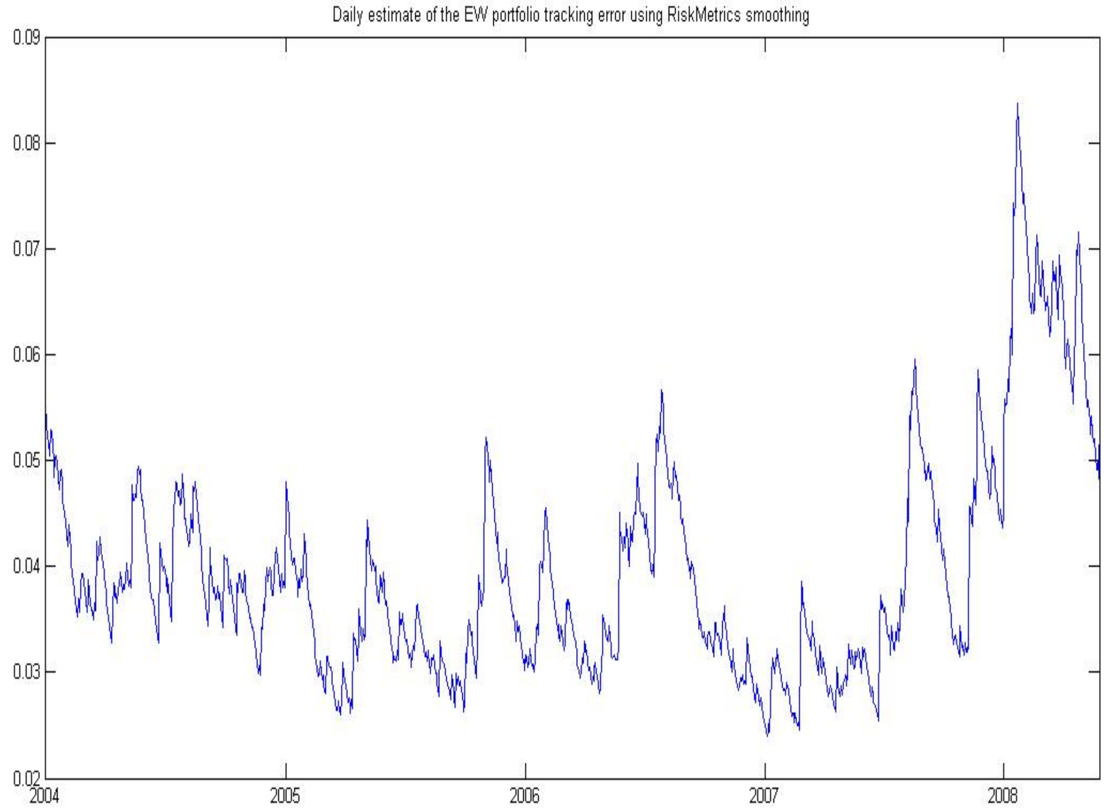

Figure 7: This Figure illustrates the "smoothed" annualized daily tracking error estimates of the equally-weighted S&P500 stocks computed with daily data. For smoothing we use the RiskMetrics persistence of 0.94.