USING HIGH-FREQUENCY DATA IN DYNAMIC PORTFOLIO CHOICE*

Federico M. Bandi,† Jeffrey R. Russell,‡ and Yinghua Zhu§

First draft: May 20, 2005
This draft: March 15, 2006

Abstract

This paper evaluates the economic benefit of methods that have been suggested to optimally sample (in an MSE sense) high-frequency return data for the purpose of realized variance/covariance estimation in the presence of market microstructure noise (Bandi and Russell, 2004, 2005a). We compare certainty equivalents derived from volatility-timing trading strategies relying on optimally-sampled realized variances and covariances, on realized variances and covariances obtained by sampling every 5 minutes, and on realized variances and covariances obtained by sampling every 15 minutes. In our sample, we show that a risk-averse investor who is given the option of choosing variance/covariance forecasts derived from MSE-based optimal sampling methods versus forecasts obtained from 5- and 15-minute intervals (as generally proposed in the literature) would be willing to pay up to about 80 basis points per year to achieve the level of utility that is guaranteed by optimal sampling. We find that the gains yielded by optimal sampling are economically large, statistically significant, and robust to realistic transaction costs.

Keywords: realized covariance, realized variance, market microstructure noise, dynamic portfolio choice.

JEL Classification: G12, C14, C22

---

*We are grateful to Essie Maasoumi and Mike McAleer (the editors), two anonymous referees, Torben Andersen, Michiel de Pooter, Rob Engle, Ken West, and the participants at the CIRANO-CIREQ conference "Financial Econometrics," Montreal, Canada, May 20-21, 2005, and the "Financial Engineering and Risk Management Workshop," Shanghai, China, July 2-4, 2005, for helpful comments and discussions.

†Address: Graduate School of Business, The University of Chicago, office 408, 5807 South Woodlawn Avenue, Chicago, IL 60637. E-mail: federico.bandi@gsb.uchicago.edu.

‡Address: Graduate School of Business, The University of Chicago, office 452, 5807 South Woodlawn Avenue, Chicago, IL 60637. E-mail: jeffrey.russell@gsb.uchicago.edu.

§Address: Department of Statistics, University of Chicago, Eckhart 124, 5734 South University Avenue, Chicago, IL 60637. E-mail: julia.yinghua.zhu@gmail.com.

1 Introduction

Large intra-day financial data sets potentially provide a rich source of information about asset price dynamics. Specifically, nonparametric variance/covariance measures constructed by summing intra-daily return data (i.e., realized variances and covariances) have the potential to provide very accurate estimates of the underlying quadratic variation and covariation (see, e.g., Andersen et al. (2004) and Barndorff-Nielsen and Shephard (2004)). These measures, however, have been shown to be sensitive to market microstructure noise inherent in the observed asset prices.[1]

This paper considers the importance of these market microstructure influences from the perspective of an asset-allocation decision problem. The paper shows that "careful" extraction of the information contained in high-frequency asset prices results in realized variance/covariance measures that yield utility benefits.

We "carefully" extract information from noisy high-frequency asset price data by optimally sampling continuously-compounded returns for the purpose of realized variance/covariance estimation. While very frequent sampling results in variance/covariance estimates that become swamped by market microstructure noise effects, infrequent sampling leads to imprecise variance/covariance estimates. Bandi and Russell (2004, 2005a) quantify these effects and provide an optimal sampling theory (in an MSE sense) for realized variance/covariance estimates.

We construct optimally-sampled daily variance/covariance estimates as in Bandi and Russell (2004, 2005a) as well as estimates obtained by using conventional (in the existing literature) 5- and 15-minute intervals. From each of these series, we derive one-day-ahead forecasts of the variance/covariance matrix. A conditional mean-variance investor can use these forecasts to optimally rebalance his/her portfolio each period. As in West et al. (1993) and Fleming et al. (2001, 2003), we compare the investor's long-run utility for optimal portfolio weights constructed from each forecast.

In our sample, we show that there are non-negligible utility gains to the investor associated with optimally-sampled realized variance/covariance matrices relative to the ad-hoc estimates. Specifically, given different values of risk-aversion, the representative investor would be willing to pay between about 3 basis points and about 62 basis points per year to use optimally-sampled realized variances/covariances versus realized variances/covariances relying on 5-minute intervals. The same investor

[1]Papers that study the impact of market microstructure noise effects on realized variance estimates are Bandi and Russell (2004, 2006), Bollen and Inder (2002), Corsi et al. (2001), Curci and Corsi (2004), Hansen and Lunde (2006), Oomen (2005, 2006) and Zhang et al. (2004), *inter alia.* Bandi and Russell (2005b) and Barndorff-Nielsen and Shephard (2005) survey this literature. Recent papers that study the impact of market microstructure noise effects on realized covariance estimates are Bandi and Russell (2005a) and Martens (2004). Martens (2004) provides a discussion of the early literature.

would be willing to pay between about 4 basis points and about 80 basis points per year to use optimally-sampled realized variances/covariances rather than realized variances/covariances based on 15-minute intervals. These gains can be economically large and are statistically significant. Furthermore, we show that the gains are robust to realistic transactions costs.

In concurrent and independent work, De Pooter et al. (2006) also study optimal sampling for the purpose of variance/covariance estimation. They select a fixed optimal frequency intended to maximize the utility that a risk-averse investor receives from alternative volatility-timing trading strategies. They find optimal sampling frequencies ranging between 30 and 65 minutes. A direct comparison between our results and the interesting results in De Pooter et al. (2006) is not obvious. Our approaches are different. Our method is designed to select a *time-varying* optimal sampling frequency for *each* entry in the variance/covariance matrix based on MSE criteria. Subsequently, the economic gains yielded by MSE-based optimal sampling are evaluated by comparing the utility gains provided by optimally-sampled realized variances/covariances and realized variances/covariances based on fixed intervals. The method in De Pooter et al. (2006) is designed to select a *constant* optimal frequency for the purpose of directly optimizing an economic criterion, rather than the MSE of the variance/covariance estimator. Despite these methodological differences, both papers emphasize the economic importance of explicitly accounting for market microstructure noise in constructing variance/covariance forecasts using high-frequency asset price data.

This work is organized as follows. Section 2 discusses our assumed price formation mechanism allowing for market microstructure contaminations. In section 3 we illustrate our identification procedure to separate the moments of the *unobserved* market microstructure noise components from the moments of the *unobserved* efficient price processes. Section 4 presents the economic metric that we use to evaluate alternative high-frequency variance/covariance forecasts. Section 5 discusses the data. Section 6 presents estimates of the noise moments of the assets in our sample. Section 7 discusses estimates of the optimal sampling frequencies for the purpose of realized variance/covariance estimation. Section 8 reports the realized variance/covariance estimates. Section 9 discusses the economic benefit of optimal sampling. The feasibility of the reported economic gains in the presence of realistic transaction costs is evaluated in Section 10. Section 11 concludes.

2 Market microstructure noise and asset prices

We use the notation $\widetilde{p}_{ih}$ to represent a k-vector of logarithmic price processes sampled at fixed time intervals h (a trading day, for instance). In the presence of market microstructure noise, the *observed* vector $\widetilde{p}_{ih}$ can be expressed as $\widetilde{p}_{ih} = p_{ih} + \eta_{ih}$, $i = 1, 2, 3, .., n$, where p_{ih} represents an *unobserved* k-vector of efficient prices, η_{ih} denotes an *unobserved* k-vector of noise contaminations, and n is the number of time intervals in the sample. We use the notations $\widetilde{p}_{(u)ih}$, $p_{(u)ih,}$ and $\eta_{(u)ih}$ with $u = 1, ..., k$ to denote individual components of the previously-defined k-vectors. In other words, $\widetilde{p}_{(u)ih}$ represents the observed price of the u-th asset of interest.

Each period can be further divided into M sub-periods of size $\delta = h/M$. The observed high-frequency continuously-compounded return vector over sub-periods can then be defined as $\widetilde{r}_{j,i} = \widetilde{p}_{(i-1)h+j\delta} - \widetilde{p}_{(i-1)h+(j-1)\delta}$ with $j = 1, 2, 3, .., M$, where $\delta = h/M$ is the time distance over which continuously-compounded returns are computed. Hence, $\widetilde{r}_{j,i}$ is the j-th intra-period return vector in period i. This return vector can be expressed as $\widetilde{r}_{j,i} = r_{j,i} + \varepsilon_{j,i}$, where $r_{j,i}$ and $\varepsilon_{j,i} = \eta_{(i-1)h+j\delta} - \eta_{(i-1)h+(j-1)\delta}$ correspond to the unobserved efficient return vector and the unobserved noise return vector over the j-th sub-interval, respectively. For simplicity, in what follows we state results for the period $i = 1$ without loss of generality.

We impose the following assumptions on the price processes and market microstructure noise contaminations. For a discussion of these assumptions, we refer the reader to Bandi and Russell (2004, 2005a).

Assumption 1. (The efficient logarithmic price vector process.)

(1) $p_t = m_t$, *where* m_t *is a multivariate stochastic volatility continuous local martingale defined as* $\int_0^t \Theta(s)dW_s$ *with* $\{W_t : t \geq 0\}$ *representing a standard vector Brownian motion with dimension* q.

(2) *The instantaneous volatility process* $\Theta(s)$ *has elements that are all cádlág.*

(3) *The instantaneous covariance* $\sum(s) = \Theta(s)\Theta'(s)$ *with generic element* $\sum_{(u)(l)}(s)$ *(where* $u \leq k, l \leq k$*) is such that* $\int_0^t \sum_{(u)(u)}(s)ds < \infty$ *for all* $t < \infty$.

Denoting the u-th row of the matrix $\Theta(s)$ as $(\sigma_{(u)1}(s), \sigma_{(u)2}(s), ..., \sigma_{(u)q}(s))$, we define the quadratic covariation process of $p_{(u)}$ and $p_{(l)}$ over h as

$$C_{(u)(l)} = \int_0^h \sum_{(u)(l)} (s)ds = \int_0^h \sum_{b=1}^q \sigma_{(u)b}(s)\sigma_{(l)b}(s)ds.$$

If $u = l$, then $C_{(u)(u)} = \int_0^h \sum_{b=1}^q \sigma^2_{(u)b}(s)ds = V_{(u)}$ which is the quadratic variation of p_u. In the literature, $C = \int_0^h \sum(s)ds$ is typically estimated by calculating $\widehat{C} = \sum_{j=1}^{M} \widetilde{r}_{j,i}\widetilde{r}^{'}_{j,i}$ (see, for example, Andersen et al. (2004) and Barndorff-Nielsen and Shephard (2004)).

Assumption 2. (The microstructure noise vector.)

(1) *The microstructure noise vectors in the price process $\eta^{'}_j s$ have mean zero and are strictly stationary with joint density $f_M(.)$.*

(2) *The variance of $\varepsilon_j = \eta_j - \eta_{j-1}$ is $O(1)$ for all M.*

(3) *The frictions $\eta^{'}_j s$ are independent of the price process for all M.*

In the presence of noise conforming with Assumption 2, Bandi and Russell (2005a) obtain $\sum_{j=1}^{M} \widetilde{r}_{j,i}\widetilde{r}^{'}_{j,i} \underset{M\to\infty}{\overset{p}{\nrightarrow}} \int_0^h \sum(s)ds$. The summing of an increasing number of cross-products of observed returns induces the summing of cross-products of noise terms. This sum diverges in probability, thereby rendering $\widehat{C}$ inconsistent for C.

The presence of noise determines a bias-variance trade-off. Realized covariance estimates calculated on the basis of high sampling frequencies (i.e., small δ values) are bound to contain a non-negligible noise component. Realized covariance estimates calculated on the basis of low sampling frequencies (i.e., large δ values) are expected to be less biased but noisier. Under Assumption 1 and 2, Bandi and Russell (2004, 2005a) characterize the bias-variance trade-off in terms of the conditional (on $\sum$) MSE of $\widehat{C}$. They also discuss methods to empirically evaluate the optimal sampling frequency δ. In what follows, we construct $\widehat{C}$ by optimally-sampling return data for the purpose of minimizing the MSE of each of the matrix's entry. Subsequently, we consider simple transformations of $\widehat{C}$ which are more robust to non-syncronous trading.

3 Separating the efficient price and noise covariances

We show that a straightforward approach to optimal sampling can entail considerable economic gains. To this end, we make a further assumption.

Assumption 3. ($MA(1)$ assumption.)

$$\eta_{(u)} \perp\!\!\!\perp \eta_{(u)-j}, \eta_{(u)} \perp\!\!\!\perp \eta_{(l)-j} \qquad \forall u, l, j \neq 0. \tag{1}$$

The $MA(1)$ assumption simplifies the MSE expressions in Bandi and Russell (2004, 2005a) and renders the MSEs easier to evaluate. The empirical relevance of the $MA(1)$ model is discussed in Bandi and Russell (2004, 2005b). Hansen and Lunde (2006) provide empirical support for the validity of this assumption at the frequencies at which we will sample continuously-compounded returns for the purpose of realized variance and covariance estimation.

Under Assumptions 1, 2, and 3, Bandi and Russell (2005a) show that the approximate optimal frequency at which to sample continuously-compounded returns to compute $\widehat{C}_{(u)(l)}$ is

$$M_{(u)(l)}^{(opt)} \approx \left(\frac{Q_{(u)(l)}}{2E^2(\varepsilon_{(u)}\varepsilon_{(l)})} \right)^{1/3}, \tag{2}$$

where $Q_{(u)(l)} = \int_0^h (\sum_{(u)(u)}(s) \sum_{(l)(l)}(s) + \sum_{(u)(l)}^2(s))ds$. When $u = l$ (in the realized variance case), the expression in Eq. (2) specializes to

$$M_{(u)(u)}^{(opt)} \approx \left(\frac{\int_0^h \sum_{(u)(u)}^2(s)ds}{E^2(\varepsilon_{(u)}^2)} \right)^{1/3}, \tag{3}$$

see Bandi and Russell (2004).

The ratios in Eq. (2) and in Eq. (3) are, effectively, signal-to-noise ratios relating moments of the *unobserved* efficient prices to moments of the *unobserved* noises in returns. The larger the latter relative to the former, the smaller M should be to avoid substantial contaminations.

The numerator and the denominator of these ratios can be evaluated. Specifically, the return noise moments $E(\varepsilon_{(u)}^2)$ and $E(\varepsilon_{(u)}\varepsilon_{(l)})$ can be consistently estimated (over any period of interest) by computing sample moments of the observed continuously-compounded return data sampled at the highest frequencies. Under simple boundedness conditions that the microstructure noise moments are expected to satisfy, Bandi and Russell (2004, 2005a) show that

$$\frac{\sum_{j=1}^{M} \widetilde{r}_{(u)j}^2}{M} \underset{M\to\infty}{\overset{p}{\to}} E(\varepsilon_{(u)}^2) \tag{4}$$

and

$$\frac{\sum_{j=1}^{M} \widetilde{r}_{(u)j}\widetilde{r}_{(l)j}}{M} \underset{M\to\infty}{\overset{p}{\to}} E(\varepsilon_{(u)}\varepsilon_{(l)}). \tag{5}$$

If the price contaminations are iid across periods, the estimators in Eq. (4) and Eq. (5) can be further averaged across periods. "Fairly" unbiased estimates of the efficient price moment $Q_{(u)(l)}$ can be obtained by computing

$$\widehat{Q}_{(u)(l)} = \overline{M} \sum_{j=1}^{\overline{M}} \widetilde{r}^2_{(u)j} \widetilde{r}^2_{(l)j} - \overline{M} \sum_{j=1}^{\overline{M}-1} \widetilde{r}_{(u)j} \widetilde{r}_{(l)j} \widetilde{r}_{(u)j+1} \widetilde{r}_{(l)j+1} \tag{6}$$

with $\frac{h}{\overline{M}} = 15$ minutes.[2] Barndorff-Nielsen and Shephard (2004) provide a discussion of the estimator $\widehat{Q}_{(u)(l)}$.

4 Realized covariation and asset allocation

We use the methodology suggested by West et al. (1993) and Fleming et al. (2001, 2003) to evaluate the economic benefit of optimal sampling in the context of an asset allocation strategy relying on volatility timing. Specifically, we compare the utility obtained by virtue of covariance forecasts based on optimal sampling to the utility obtained through covariance forecasts constructed using the ubiquitous (in the extant literature) 5- and 15-minute intervals.

Let R^f and R_{t+1} be the risk-free return and the return vector on k risky assets over a day $t, t+1$, respectively. Define $\boldsymbol{\mu}_t = E_t(R_{t+1})$ and $\Phi_t = E_t\left[(R_{t+1} - \boldsymbol{\mu}_t)(R_{t+1} - \boldsymbol{\mu}_t)'\right]$, the conditional expected value and the conditional covariance matrix of R_{t+1}. We consider a mean-variance investor who solves the program

$$\min_{w_t} w_t' \Phi_t w_t, \tag{7}$$

subject to

$$w_t' \boldsymbol{\mu}_t + (1 - w_t' \mathbf{1}_k) R^f = \mu_p, \tag{8}$$

where w_t is a k-vector of portfolio weights, μ_p is a target expected return on the portfolio, and $\mathbf{1}_k$ is a $k \times 1$ unit vector. The solution to this program is

$$w_t = \frac{\left(\mu_p - R^f\right) \Phi_t^{-1} \left(\boldsymbol{\mu}_t - R^f \mathbf{1}_k\right)}{\left(\boldsymbol{\mu}_t - R^f \mathbf{1}_k\right)' \Phi_t^{-1} \left(\boldsymbol{\mu}_t - R^f \mathbf{1}_k\right)}. \tag{9}$$

In what follows, we estimate Φ_t using one-day-ahead forecasts $\left(\widetilde{C}_t\right)$ given a time series of daily C estimates. To evaluate the value of volatility timing using alternative covariance forecasts, we compute three estimates of C for every day t using (i)

[2] The simulations in Bandi and Russell (2006) show that the use of credible sampling frequencies for the integrated quarticity $\widehat{Q}_{(u)(u)}$ do not have any meaningful impact on the MSE (and optimal sampling frequency) of the realized variance estimator. In particular, the 15 minute frequency works well in practise.

optimally-sampled realized covariances, (*ii*) realized covariances constructed using 5-minute returns, and (*iii*) realized covariances using 15-minute returns. By virtue of the criterion in Eq. (7) and Eq. (8), and given sensible choices of R^f, μ_p and $\boldsymbol{\mu}_t$, each one-day-ahead forecast leads to the determination of a daily portfolio weight w_t. The time series of daily portfolio weights then leads to daily portfolio returns. Finally, it is reasonable to employ the investor's long-run mean-variance utility as a metric to evaluate the economic benefit of alternative covariance forecasts $\widetilde{C}_t$, i.e.,

$$AU^* = \overline{R}^p - \frac{\lambda}{2}\frac{1}{m}\sum_{t=1}^{m}(R^p_{t+1} - \overline{R}^p)^2, \tag{10}$$

where

$$R^p_{t+1} = R^f + w'_t(R_{t+1} - R^f\mathbf{1}_k) \tag{11}$$

is the return on the portfolio with estimated weights w_t, $\overline{R}^p = \frac{1}{m}\sum_{t=1}^{m} R^p_{t+1}$ is the sample mean of the portfolio returns across $m \leq n$ days, and λ is a coefficient of risk-aversion. In order to avoid contaminations induced by noisy first moment estimation, and in light of our criterion which sets a constant target return for each period, we conduct the comparison by simply looking at the variance component of AU^*, namely

$$AU = \frac{\lambda}{2}\frac{1}{m}\sum_{t=1}^{m}(R^p_{t+1} - \overline{R}^p)^2 \tag{12}$$

(see Engle and Colacito, 2005, for further justifications of this approach). Below, we use three conventional values of λ, namely 2, 7, and 10.

We interpret the difference between $AU^{5\,\min}$ and AU^{opt}, for example, as the fee that the investor would be willing to pay to switch from covariance forecasts based on 5-minute sampling $\left(\widetilde{C}_t^{5\,\min}\right)$ to covariance forecasts based on optimally-sampled realized covariances $\left(\widetilde{C}_t^{opt}\right)$.

5 The data

We study the stocks Merrill Lynch, MER, SBC Communications, SBC, and EXXON Mobile Corporation, XOM (XON, prior to November 1999). We choose these stocks because they represent median and extreme features of the S&P 100 stocks, as summarized by the ratio between noise variance and variance of the efficient price (i.e., $\frac{E(\varepsilon^2_{(.)})}{V_{(.)}}$), for the data used in Bandi and Russell (2006).

We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. The data are provided by the Trade and Quote (TAQ) database. Our final data set consists of bid-ask quote updates.[3] We remove quote updates and bid-ask spreads larger than 10%. We also remove trading days without quote updates after 3:15 pm (possibly due to holidays) and trading days that have fewer than 50 quote updates after applying the previously-described filter.

Our portfolio choice problem involves MER, SBC, XOM, and a risk-less asset. However, in what follows, we report descriptive results only for SBC and XOM. The corresponding results pertaining to MER are qualitatively similar and are omitted for conciseness.

Fig. 1 (panel a) presents the histograms of the durations (in seconds) between quote updates for SBC and XOM. Fig. 1 (panel b) presents the average duration times (in seconds) for both stocks and every year in our sample. The decreasing trend in the average durations is apparent.

We generate equally-spaced continuously-compounded returns using midquotes and the previous tick method. The procedure is standard. First, we define a grid of intervals (in seconds) that spans the trading day. At each point on the grid, we define the price as the prevailing logarithmic price. We then take first differences of the logarithmic prices to construct continuously-compounded returns. The daily realized variance/covariance estimates and the estimator in Eq. (6) are constructed using fixed calendar time intervals. The daily noise second moments are estimated using quote-to-quote continuously-compounded returns. The daily noise covariances are estimated using equispaced returns over intervals equal to two times the largest (among the relevant two stocks) daily average duration between quote updates.

6 Estimating the noise components of the observed return series

For each trading day in our sample, we estimate the standard deviations of the stocks' noise components by taking the square root of the estimator in Eq. (4). Fig. 2 reports the time series of the estimated noise standard deviations of SBC and XOM. For both stocks, the magnitude of the standard deviations has been largely decreasing since 1993. This decrease suggests that higher frequencies may be needed for the purpose of realized variance estimation in the last part of the sample (see Eq. (3) above). The downward trend of the estimated noise standard deviations can be partly due to residual efficient price variance components in the estimated noise variance at the beginning of the sample. In effect, in light of the asymptotic justification for the consistency result underlying Eq. (4), the estimates of the noise second moments are

[3] The quote data also contains observations indicating revisions in the depth. These observations carry no additional price information and are therefore excluded.

expected to be less biased in the presence of a large number of intradaily returns. During the first four years in our sample, the average duration is relatively high, thereby resulting in a smaller number of available quotes for each day and, possibly, larger biases (see Fig. 2). Fig. 3 reports the time series of the estimated noise correlations obtained by using the estimators in Eq. (4) and Eq. (5). The estimates are largely positive. Their level is fairly high around 1997. This is likely to be a true effect in the data. Possible biases in our noise moment estimates would likely not explain the higher correlation levels around 1997. There are two potential sources of bias in the noise correlation estimates. First, the noise covariances and the noise variances might contain residual components due to the efficient price dynamics. Second, the noise covariance estimates might be biased due to non-synchronous trading. Non-synchronous trading and the potential (positive) biases in the variance estimates would, if anything, drive the noise correlations towards zero. However, the noise covariances might be biased upward by virtue of contaminations induced by largely positive efficient price covariances as documented in Section 8. This effect is unlikely to explain the higher noise correlation levels in 1997 since the correlation between noise covariances and efficient price covariances (as measured by using low frequency estimates) is fairly low. As in the case of the estimated noise second moments, the noise covariance and correlation estimates have the potential to be more biased at the beginning of our sample due to the corresponding long durations. Thus, we exclude the first four years of data from the portfolio allocation experiment in Section 9.

7 Estimating the optimal sampling frequencies

Figs. 4 through 6 contain the time-series of the daily optimal sampling frequencies (in minutes) for the realized variances of SBC and XOM and their realized covariance. The three series of optimal frequencies are highly positively correlated with the corresponding estimated daily moments of the noise components, namely $E\left(\varepsilon^2_{(SBC)}\right)$, $E\left(\varepsilon^2_{(XOM)}\right)$, and $E\left(\varepsilon_{(SBC)}\varepsilon_{(XOM)}\right)$. The corresponding correlations are 0.769, 0.707, and 0.687. Both the variability and the magnitude of the optimal frequencies decrease over time as implied by noise moments whose variability and magnitude decrease over time.

8 Estimating the realized variance/covariance process

Define $\widehat{C}^{(\cdot)}$ as

$$\left|\begin{array}{c|c} \displaystyle\sum_{j=1}^{M^{(\cdot)}_{(SBC)}} \widetilde{r}^2_{(SBC)j} & \blacklozenge \\ \hline \displaystyle\sum_{j=1}^{M^{(\cdot)}_{(SBC)(XOM)}} \widetilde{r}_{(SBC)j}\widetilde{r}_{(XOM)j} & \displaystyle\sum_{j=1}^{M^{(\cdot)}_{(XOM)}} \widetilde{r}^2_{(XOM)j} \end{array}\right| .$$

The three-asset extension including MER, which will be considered in Section 9, will be obvious in light of our discussion in this section. The optimally-sampled $\widehat{C}^{(\cdot)}$, $\widehat{C}^{(opt)}$, is obtained using $M^{(opt)}_{(SBC)(XOM)}$, $M^{(opt)}_{(SBC)}$, and $M^{(opt)}_{(XOM)}$ where $M^{(opt)}_{(SBC)(XOM)}$ is defined in Eq. (2) while $M^{(opt)}_{(SBC)}$ and $M^{(opt)}_{(XOM)}$ are defined in Eq. (3). The 5-minute $\widehat{C}^{(\cdot)}$, $\widehat{C}^{(5\,\text{min})}$, is obtained using $M^{(5\text{ min})}_{(SBC)(XOM)} = M^{(5\text{ min})}_{(SBC)} = M^{(5\text{ min})}_{(XOM)} = 72(= 6 \times 60/5)$. The 15-minute $\widehat{C}^{(\cdot)}$, $\widehat{C}^{(15\,\text{min})}$, is obtained using $M^{(15\text{ min})}_{(SBC)(XOM)} = M^{(15\text{ min})}_{(SBC)} = M^{(15\text{ min})}_{(XOM)} = 24(= 6 \times 60/15)$. Now define $\widehat{C}^{(\cdot)}_v$ as

$$\left|\begin{array}{c|c} \displaystyle\sum_{j=1}^{M^{(\cdot)}_{(SBC)}} \widetilde{r}^2_{(SBC)j} & \blacklozenge \\ \hline \displaystyle\sum_{j=1}^{M^{(\cdot)}_{(SBC)(XOM)}} \widetilde{r}_{(SBC)j}\sum_{s=-v}^{v}\widetilde{r}_{(XOM)j-s} & \displaystyle\sum_{j=1}^{M^{(\cdot)}_{(XOM)}} \widetilde{r}^2_{(XOM)j} \end{array}\right|$$

with $v = 0, 1$, and 2. When $v = 0$, $\widehat{C}^{(\cdot)}_{v=0} = \widehat{C}^{(\cdot)}$ as represented above. When $v \neq 0$, the realized covariance estimator is adjusted using leads and lags to correct for the possibility of non-syncronous trading. The justification behind this adjustment is well-known (see, for example, Cohen et al., 1983). When the efficient return process is a local martingale, the off-diagonal elements of $\widehat{C}^{(\cdot)}_v$ are "virtually" unbiased for the true covariation over the period provided v is large enough. When v is small, the potential lack of price updates is bound to induce biases.

For each trading day we compute $\widehat{C}^{(\cdot)}_v$ with (.) = (*opt*), (5 min), (15 min), and $v = 0, 1, 2$. We constrain the absolute value of the realized correlation to be at most 1 by imposing the condition

$$\overline{C}_{v(SBC)(XOM)} = sign(\widehat{C}_{v(SBC)(XOM)}) \times \min\left(|\widehat{C}_{v(SBC)(XOM)}|, \sqrt{\widehat{V}_{(SBC)}\widehat{V}_{(XOM)}}\right). \tag{13}$$

Fig. 7 shows the three time series of daily realized correlations based on optimally-sampled realized variances and covariances and variances/covariances constructed using 5- and 15-minute intervals, namely $\frac{\overline{C}^{(opt)}_{v=0(SBC)(XOM)}}{\sqrt{\widehat{V}^{(opt)}_{(SBC)}\widehat{V}^{(opt)}_{(XOM)}}}$, $\frac{\overline{C}^{(5\,\text{min})}_{v=0(SBC)(XOM)}}{\sqrt{\widehat{V}^{(5\,\text{min})}_{(SBC)}\widehat{V}^{(5\,\text{min})}_{(XOM)}}}$, and $\frac{\overline{C}^{(15\,\text{min})}_{v=0(SBC)(XOM)}}{\sqrt{\widehat{V}^{(15\,\text{min})}_{(SBC)}\widehat{V}^{(15\,\text{min})}_{(XOM)}}}$, respectively. Figure 8 reports the three time series of daily optimally-sampled realized correlations constructed using no leads and lags, one lead and one lag, and two leads and two lags, namely $\frac{\overline{C}^{(opt)}_{v=0(SBC)(XOM)}}{\sqrt{\widehat{V}^{(opt)}_{(SBC)}\widehat{V}^{(opt)}_{(XOM)}}}$, $\frac{\overline{C}^{(opt)}_{v=1(SBC)(XOM)}}{\sqrt{\widehat{V}^{(opt)}_{(SBC)}\widehat{V}^{(opt)}_{(XOM)}}}$, and $\frac{\overline{C}^{(opt)}_{v=2(SBC)(XOM)}}{\sqrt{\widehat{V}^{(opt)}_{(SBC)}\widehat{V}^{(opt)}_{(XOM)}}}$, respectively. Graphically, the introduction of one lead and one lag ($v = 1$) appears to provide a correction for the downward bias induced by non-syncronous trading when $v = 0$. The introduction of additional leads and lags (i.e., the case $v = 2$) does not appear to have an impact on the level of the time-varying correlations. Since the additional leads and lags only inflate the correlations' variance we prefer not to over-correct and set $v = 1$ in what follows.[4]

9 The economic benefit of optimal-sampling

9.1 Obtaining the realized variance/covariance forecasts

Implementation of the criterion in Eq. (12) requires the computation of one-day-ahead variance/covariance forecasts. For the time series of estimated variances and covariances obtained by optimal sampling, 5-minute sampling, and 15-minute sampling, we construct out-of-sample forecasts on the basis of univariate ARFIMA models. As in Bandi and Russell (2006), we set the orders of the autoregressive and moving average representations equal to 2. The ARMA parameters are re-estimated in real time. Residual diagnostics demonstrate that the (2,2) model is sufficient for all subsamples of data. The fractional d parameter is fixed at the estimated value obtained using the full sample. We employ the traditional GPH estimator to estimate the d parameter (Geweke and Porter-Hudak, 1983). Table 1 contains the estimated d values for the 18 series contained in $\overline{C}^{(.)}_{v=1}$ with $(.) = (opt)$, $(5\,\text{min})$, $(15\,\text{min})$. Since the fractional d parameter is based on the full sample, our empirical exercise is not genuinely out-of-sample. While it is well-known that this parameter can hardly be estimated efficiently, thereby justifying our approach, we find that alternative choices of the d parameter do not affect our results. For instance, we find that setting $d = 0.45$ for all variance/covariance measures does not change our utility gains (reported below).

As shown above, since the average durations are high prior to 1997, the estimated moments of the noise and the optimal sampling frequencies might be biased in the

[4]It is of interest to derive the optimal sampling interval of the realized covariance estimator as a function of the number of leads and lags used to correct for non-syncronous trading. This study is beyond the scope of the present work and is left for future research. Importantly, the use of different sampling frequencies for realized variances and covariances does not guarantee positive-definiteness of the resulting variance/covariance matrix. While, empirically, this is never a problem in our case, the study of methods that guarantee positive-definiteness for large, optimally-sampled realized variance/covariance matrices is also an important topic for future research.

early part of our sample. To circumvent this issue, we forecast using only post-1997 data. We utilize almost a year worth of data, i.e., 200 daily observations, to construct the first forecast. The total number of out-of-sample forecasts m is equal to 1,537.

Before computing the daily variance/covariance forecasts, we correct the realized variance/covariance estimates to obtain daily measures, rather than measures based on a 6-hour period. For each trading day in our sample, the correction factor

$$\zeta_{(u)} = \frac{1/n \sum_{i=1}^{n} \widetilde{r}_{(u)i}^{2}}{1/n \sum_{i=1}^{n} \widehat{V}_{(u)i}}, \tag{14}$$

where $\widetilde{r}_{(u)i}$ denotes the return on stock u for day i, is applied to the u-th diagonal element of $\overline{C}_{i,v=1}$ giving $\widehat{V}_{(u)i}^{*} = \zeta_{(u)} \widehat{V}_{(u)i}$. Equivalently, the correction factor

$$\zeta_{(u,l)} = \frac{1/n \sum_{i=1}^{n} \widetilde{r}_{(u)i} \widetilde{r}_{(l)i}}{1/n \sum_{i=1}^{n} \overline{C}_{(u,l)i}} \tag{15}$$

is applied to the (u, l) off-diagonal element of $\overline{C}_{i,v=1}$ giving $\overline{C}_{(u,l)i,v=1}^{*} = \zeta_{(u,l)} \overline{C}_{(u,l)i,v=1}$. These corrections guarantee that the average realized variances and covariances equal their counterparts estimated from daily returns over the sample.

Fig. 9 shows the time series plots of the "corrected" variance forecasts for SBC (panel a) and XOM (panel b) using optimal sampling, 5-minute intervals, and 15-minute intervals. Fig. 10 contains the time series plots of the corrected SBC/XOM covariance forecasts based on the three sampling intervals.

9.2 Forecasting regressions

Before turning to asset-allocation, we implement a forecasting exercise in the spirit of Andersen et al. (2003). We evaluate the forecasting power of optimally-sampled realized covariances, of covariances constructed using 5-minute intervals, and of covariances constructed using 15-minute intervals. For each pair of stocks, we regress each covariance measure on one-step-ahead forecasts obtained by virtue of all three measures. The results are reported in Table 2 through 10. We find that (*1*) optimally-sampled realized covariances help forecast both realized covariances sampled every 5 minutes and realized covariances sampled every 15 minutes, (*2*) the forecasting power of realized covariances sampled optimally is superior to the forecasting power of realized covariances sampled every 15 minutes *even* when forecasting realized covariances sampled every 15 minutes, and (*3*) the forecasting power of optimally-sampled realized covariances is generally superior and, in a few instances, very similar to the forecasting power of realized covariances sampled every 5 minutes. Furthermore, comparing the R^2 of each regression suggests that realized covariances sampled optimally are less noisy, and hence, more predictable, than realized covariances sampled every 5 minutes or every 15 minutes.

9.3 Volatility timing and dynamic portfolio choice

We implement the criterion in Eq. (12) by setting R^f equal to 0.03 (converted to daily values by dividing by 250). We consider three targets μ_p, namely 0.09, 0.12, and 0.15. For all times t, we set the components of the 3×1 vector $\boldsymbol{\mu}_t = E_t(R_{t+1})$ equal to the sample means of the returns on the three risky assets over the forecasting horizon (October 1997 - December 2003). In other words, we concentrate on volatility timing and abstract from the issues that would be posed by expected stock return predictability. The sample means of the daily returns are equal to 0.000248 (SBC), 0.000374 (XOM), and 0.000746 (MER). For all times t, the conditional covariance matrix is computed as an out-of-sample forecast from an ARFIMA(2,2) model as described above. For illustration, in Fig. 11 we plot the daily portfolio weights constructed using optimal intervals and 5-minute intervals for a target portfolio with $\mu_p = 0.12$.

As in West et al. (1993) and Fleming et al. (2001, 2003), we interpret the difference between the average utility computed on the basis of optimally-sampled realized covariances and that computed on the basis of alternative realized covariances as the return that the investor would sacrifice to switch to optimally-sampled covariance estimates. Table 11 contains the results for three levels of risk-aversion and three target expected returns. The difference between average utilities is shown as an annualized fee. When the target is 0.09, the investor would pay between 15.6 basis points (when $\lambda = 10$) and 3.1 basis points (when $\lambda = 2$) per year to use optimally-sampled realized covariances versus realized covariances relying on 5-minute intervals. The same investor would pay marginally more (19.7 basis points when $\lambda = 10$, 3.94 basis points when $\lambda = 2$) to use optimally-sampled realized covariances rather than realized covariances based on 15-minute intervals. When the target is 0.12, the investor would pay between 35.1 basis points (when $\lambda = 10$) and 7 basis points (when $\lambda = 2$) per year to use optimally-sampled realized covariances versus realized covariances relying on 5-minute intervals. Again, the investor would give up more (44.3 basis points when $\lambda = 10$, 8.86 basis points when $\lambda = 2$) to use optimally-sampled realized covariances rather than realized covariances based on 15-minute intervals. Finally, when the target is 0.15, the investor would pay between 62.38 basis points (when $\lambda = 10$) and 12.47 basis points (when $\lambda = 2$) per year to use optimally-sampled realized covariances versus realized covariances relying on 5-minute intervals. As earlier, the investor would pay more (78.82 basis points when $\lambda = 10$, 15.76 basis points when $\lambda = 2$) to use optimally-sampled realized covariances rather than realized covariances based on 15-minute intervals.[5]

In order to evaluate the incremental benefit of sampling only the realized covariances optimally, we compare the average utility computed on the basis of optimally-sampled realized covariances and 5-minute (15 minute) realized variances to the average utility computed on the basis of 5 minute (15 minute) realized variances and covariances. The results are in Table 12. The gains are smaller in this case and, as we will show in the next subsection, statistically insignificant.

9.4 The statistical significance of the economic gains

For any target μ_p and any sampling method, one can define R^p_{t+1} as in Eq. (11). Denote by

$$R^{p(s,d)}_{t+1} = R^f + w^{(s,d)\prime}_t (R_{t+1} - R^f \mathbf{1}_3) \qquad s = opt, 5\,\text{min}, 15\,\text{min}\,,\; d = 0.09, 0.12, 0.15,$$

the return on the portfolio obtained by using a specific sampling scheme (optimal, 5 minutes, or 15 minutes) for the realized variances/covariances and imposing a certain target return (0.09, 0.12, or 0.15). Now define, for example,

$$a^{5,opt,d}_{t+1} = \left(R^{p(opt,d)}_{t+1} - \overline{R}^{p(opt,d)}\right)^2 - \left(R^{p(5,d)}_{t+1} - \overline{R}^{p(5,d)}\right)^2.$$

When comparing a portfolio obtained by sampling optimally to a portfolio obtained by sampling every 5 minutes for a specific target expected return d, assessing the statistical significance of the economic gains can be conducted by testing whether the mean of $a^{5,opt,d}_{t+1}$ is larger than (or equal to) zero against the alternative that the mean is smaller than zero. For each target return d we define the vector

$$A^d_{t+1} = \left(a^{5,opt,d}_{t+1}, a^{15,opt,d}_{t+1}\right)'$$

and write the regression model

$$A^d_{t+1} = \delta^d \mathbf{1}_2 + \epsilon_{t+1},$$

where δ^d is a scalar parameter. We then perform the one-sided test $H_0 : \delta^d \geq 0$ against $H_A : \delta^d < 0$ for each d. The parameter δ^d is estimated by GMM using a Bartlett HAC covariance matrix. A similar approach is adopted in Engle and Colacito (2005). The t-statistics of the three tests are all equal to about -1.73 yielding a one-sided

[5] The risk-aversion coefficient simply re-scales the portfolio variances but does not affect the portfolio holdings directly. Hence, the gains/losses are monotonic in the value of this coefficient.

p-value of 0.041 implying statistical significance of the economic gains at conventional levels.

It should be noted that optimizing the covariances while sampling the realized variances using 5- or 15-minute intervals is not sufficient, in our sample, to obtain statistically significant economic gains. The test statistics are equal to about -0.45 for all three targets, giving a p-value of about 0.65. In sum, it is the combination of optimally-sampled variances and covariances which provides statistically significant economic gains over the 5- and 15-minute counterparts for the data used in this study.

10 The feasibility of the economic gains

The assumed daily rebalancing raises the issue of whether realistic transaction costs have the potential to affect the practical feasibility of our findings. To this extent, we analyze the impact of alternative transaction costs on the reported utility gains. This exercise is conducted by assuming a fixed, daily, rebalancing frequency. Choice of the optimal rebalancing time given reasonable transaction costs is a topic of interest but beyond the scopes of the current paper. Let ρ denote the transaction cost per traded dollar. For simplicity, assume ρ is fixed across assets. The portfolio return net of transaction costs can then be expressed as:

$$netR_{t+1}^{p(s,d)} = R_{t+1}^{p(s,d)} - \rho \sum_{j=1}^{4} (1+R_{t+1}^{j}) |\Delta\omega_{t+1}^{j(s,d)}| \qquad s = opt, 5\,\text{min}, 15\,\text{min}\,,\ d = 0.09, 0.12, 0.15,$$

where $R_{t+1}^{p(s,d)}$ is the unadjusted porfolio return obtained by sampling with the method s for a target return d, R_{t+1}^{j} is the unadjusted return on asset j, and $\Delta\omega_{t+1}^{j(s,d)}$ is equal to $\omega_{t+1}^{j(s,d)} - \omega_{t}^{j(s,d)}$ with $\omega_{t}^{j(s,d)}$ defined as the portfolio weight on asset j for a specific target return d and a certain sampling method s. The cost ρ is set equal to 0.00005, 0.0001, and 0.0025. These correspond to a 0.5 cent half spread on a 100 dollar stock, a 1 cent half spread on a 100 dollar stock, and a 2.5 cent half spread on a 10 dollar stock.

In Table 13 we report the results for a target return equal to 0.12. We find that the presence of transaction costs does not affect the reported gains in any meaningful way. This result is due to the fact that transaction costs affect the variance of the portfolio returns only slightly (i.e., the variance of the unadjusted portfolio returns dominates the variance induced by re-balancing) and our economic criterion effectively compares portfolio variances obtained by virtue of alternative variance/covariance estimates. Interestingly, using our data, the fee that a risk-averse investor would pay to switch from realized variances/covariances based on fixed frequencies to optimally-sampled realized variances/covariances increases slightly when allowing for transaction costs. The corresponding results for a target return equal to 0.09 and 0.15 are virtually identical, i.e., in all cases higher transaction costs yield slightly higher utility gains.

11 Conclusions

In the absence of market microstructure noise, very precise realized variance and covariance estimates can be obtained. In practice, one must confront the fact that observed financial prices contain non-negligible market microstructure effects that lead to contaminations of the estimates. While the statistical importance of these contaminations can be reduced by optimally-sampling asset prices, this paper establishes the economic importance of optimal sampling procedure. We study the forecasting power of alternative realized variance measures in the context of an important economic metric, namely the long-run utility of a conditional mean-variance investor rebalancing his/her portfolio each period. We show that constructing variance/covariance forecasts using optimally-sampled realized variances and covariances is a superior strategy to constructing variance/covariance forecasts based on realized variance and covariance estimates obtained from arbitrarily fixed intervals, as typically suggested in practise. The economic gains are non-negligible in our sample. Specifically, we find that a conditional mean-variance investor would be willing to pay up to about 80 basis points per year to switch to variance forecasts constructed on the basis of optimally-sampled realized variance/covariance estimates. We find that these economic gains are statistically significant and robust to realistic transaction costs.

As Andersen et al. (2005) point out in their recent review paper on volatility forecasting "the use of realized volatility measures for forecasting is still in its infancy and many issues must be explored in future work." We believe that the ultimate criterion to evaluate alternative forecasts should be based on a relevant economic metric. This paper focused on the classical realized variance/covariance estimator evaluated in the framework of a traditional portfolio choice problem. We showed that the optimization of the finite sample properties of the realized variance/covariance estimator results in important economic gains. Similarly, we expect the optimization of the finite sample properties of the recently-proposed integrated variance estimators, as advocated by Bandi and Russell (2005c) in the context of the kernel-based estimates of Barndorff-Nielsen et al. (2005) and Zhang et al. (2004) for instance, to be beneficial from an economic standpoint. In addition, new economic metrics should be proposed. Much work remains to be done along these lines.

References

[1] Andersen, T.G., T. Bollerslev, P.F. Christoffersen, and F.X. Diebold (2005). Volatility and correlation forecasting. In G. Elliott, C.W.J. Granger, and Allan Timmermann (eds.), *Handbook of Economic Forecasting*. Amsterdam: North-Holland, forthcoming.

[2] Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (2003). Modeling and forecasting realized volatility. Econometrica, 71, 579-625.

[3] Andersen, T.G., T. Bollerslev, F.X. Diebold, and J. Wu (2004). Realized beta: persistence and predictability. *Working paper.*

[4] Bandi, F. M., and J.R. Russell (2004). Microstructure noise, realized variance, and optimal sampling. *Working paper.*

[5] Bandi, F. M., and J.R. Russell (2005a). Realized covariation, realized beta, and microstructure noise. *Working paper.*

[6] Bandi, F. M., and J.R. Russell (2005b). Volatility. In J. R. Birge and Vadim Linetsky (eds.), *Handbook of Financial Engineering*. Amsterdam: North-Holland, forthcoming.

[7] Bandi, F. M., and J.R. Russell (2005c). Market microstructure noise, realized variance, and the limitations of asymptotic approximations: a solution. *Working paper.*

[8] Bandi, F. M., and J.R. Russell (2006). Separating microstructure noise from volatility. Journal of Financial Economics, 79, 655-692.

[9] Barndorff-Nielsen, O. E. and N. Shephard (2004). Econometric analysis of realized covariation: high-frequency based covariance, regression and correlation in financial economics. Econometrica, 72, 885-925.

[10] Barndorff-Nielsen, O. E. and N. Shephard (2005). Variation, jumps, market frictions and high-frequency data in financial econometrics. *Invited lecture, World Congress of the Econometric Society, 2005.*

[11] Barndorff-Nielsen, O. E., P. Hansen, A. Lunde, and N. Shephard (2005). Regular and modified kernel-based estimators of integrated variance: the case with independent noise. *Working paper.*

[12] Bollen, B. and B. Inder (2002). Estimating daily volatility in financial markets utilizing intraday data. Journal of Empirical Finance, 9, 551-562.

[13] Cohen, K.J., G.A. Hawanini, S.F. Maier, R.A. Schwartz, and D.K. Whitcomb (1983). Friction in the trading process and estimation of systematic risk. Journal of Financial Economics, 12, 263-278.

[14] Corsi, F., G. Zumbach, U.A. Müller, and M. Dacorogna (2001). Consistent high-precision volatility from high-frequency data. Economic Notes, 30, 183-204.

[15] Curci, G. and F. Corsi (2004). A discrete sine transform approach for realized volatility measurement. *Working paper.*

[16] De Pooter, M., M. Martens, and Dick van Dijk (2006). Predicting the daily covariance matrix for S&P100 stocks using intraday data: but which frequency to use? *Econometric Reviews,* this issue.

[17] Engle, R. and R. Colacito (2005). Testing and valuing dynamic correlations for asset allocation. *Working paper.*

[18] Fleming, J., C. Kirby, and B. Ostdiek (2001). The economic value of volatility timing. Journal of Finance, 56, 329-352.

[19] Fleming, J., C. Kirby, and B. Ostdiek (2003). The economic value of volatility timing using realized volatility. Journal of Financial Economics, 67, 473-509.

[20] Geweke, J. and S. Porter-Hudak (1983). The estimation and application of long memory time series models. Journal of Time Series Analysis, 4, 221-238.

[21] Hansen, P.R., and A. Lunde (2006). Realized variance and market microstructure noise (with discussions). Journal of Business and Economic Statistics, forthcoming.

[22] Martens, M. (2004). Estimating unbiased and precise realized covariances. *Working paper.*

[23] Oomen, R.C.A. (2005). Properties of bias-corrected realized variance under alternative sampling schemes. Journal of Financial Econometrics, 4, 555-577.

[24] Oomen, R.C.A. (2006). Properties of realized variance under alternative sampling schemes. Journal of Business and Economic Statistics, forthcoming.

[25] West, K.D., H.J. Edison, and D. Cho (1993). A utility-based comparison of some models of exchange rate volatility. Journal of International Economics, 35, 23-45.

[26] Zhang, L., P. Mykland and Y. Ait-Sahalia (2004). A tale of two time scales: determining integrated volatility with noisy high-frequency data. Journal of the American Statistical Association, forthcoming.

(a)

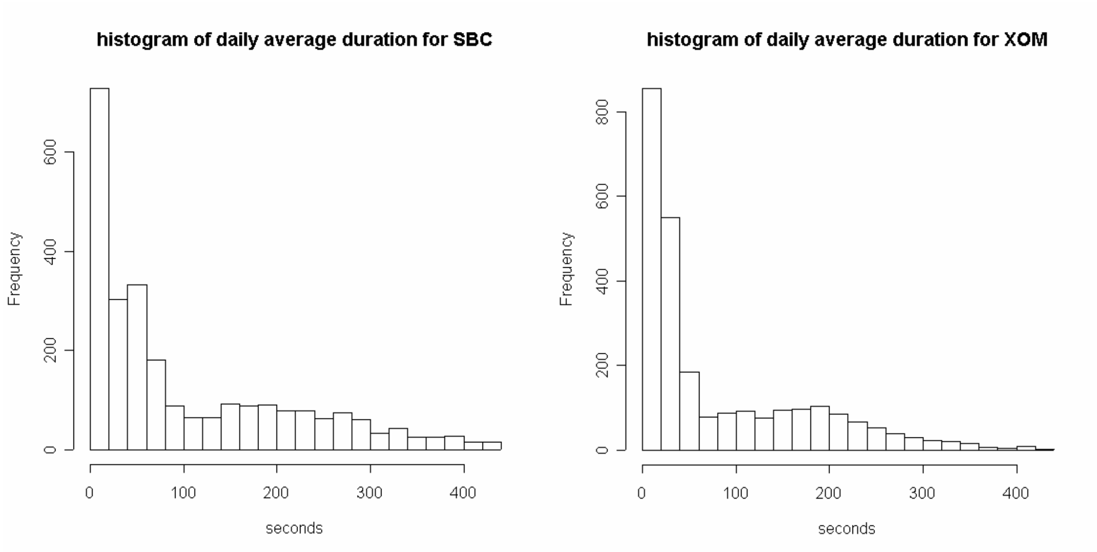

(b)

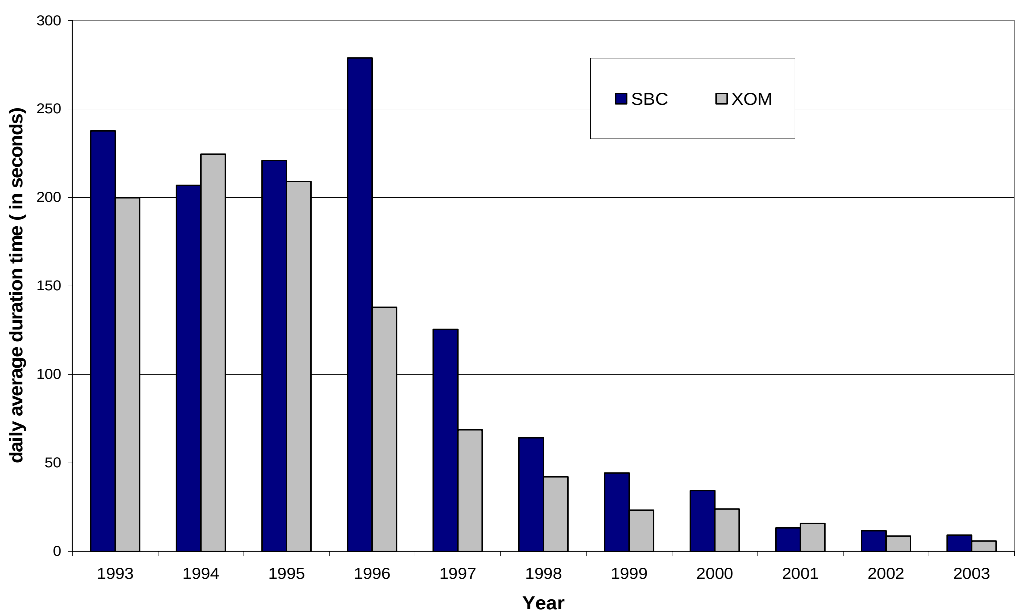

Fig 1. The durations of SBC and XOM in seconds. Panel (a) shows the histogram of the durations (in seconds) between quote updates for SBC Communications (SBC) and EXXON Mobile Corporation (XOM). Panel (b) shows the average durations (in seconds) for both stocks and every year in our sample. The data are provided by the TAQ database. We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST.

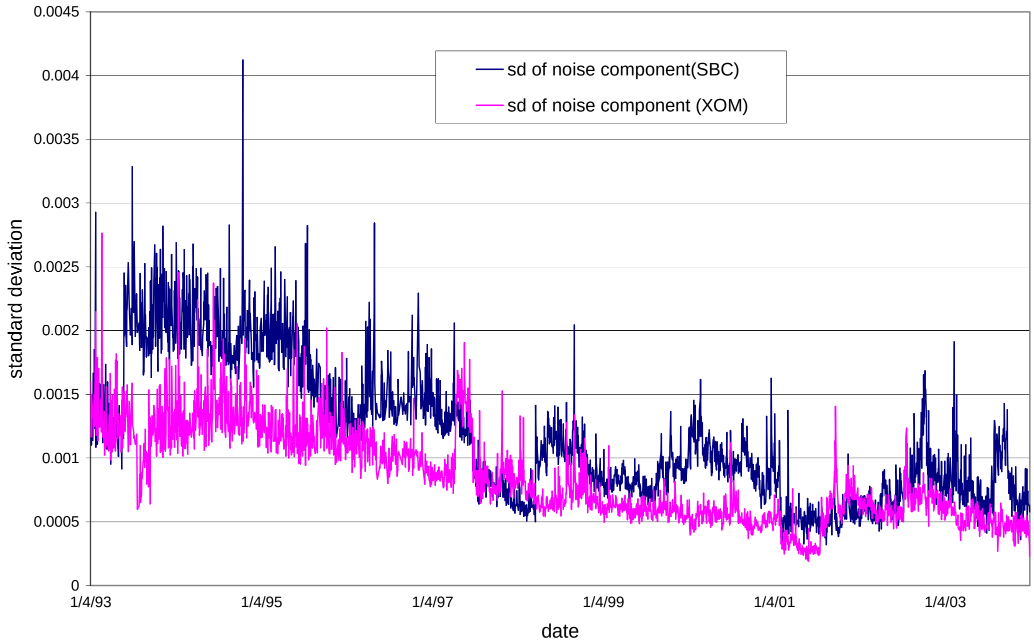

Fig 2. Daily time series of the standard deviations of the noise components of the logarithmic mid-quotes of SBC Communications (SBC) and EXXON Mobile Corporation (XOM). The standard deviations of the noise components are obtained as the square root of the sample second moment of the quote-to-quote continuously-compounded returns. The data are provided by the TAQ database. We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST.

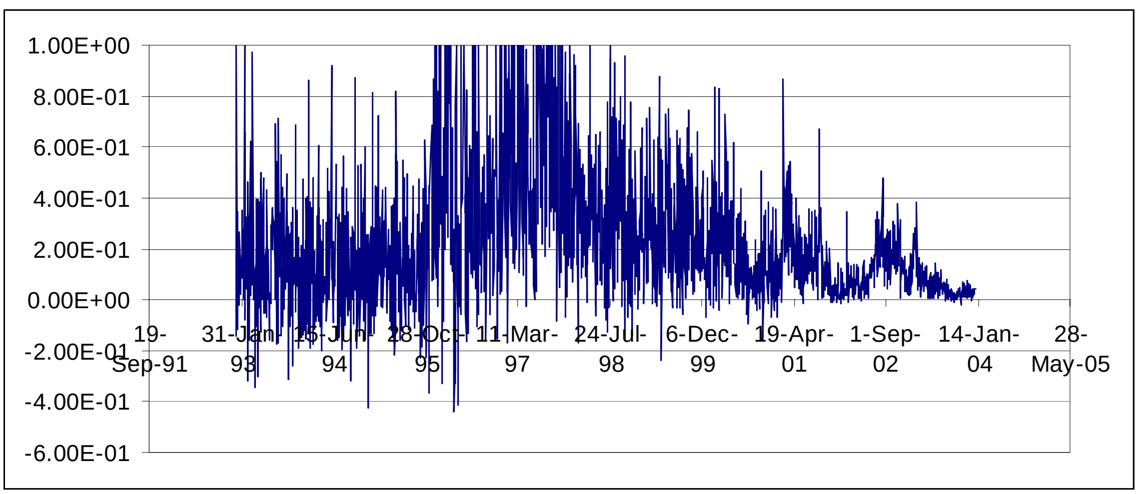

Fig 3. Daily time series of the correlation of the noise components of the logarithmic mid-quotes of SBC Communications (SBC) and EXXON Mobile Corporation (XOM). The covariances of the noise components are obtained as the sample first cross-moments of the continuously-compounded returns sampled at twice the largest (among the two stocks) average daily duration. The standard deviations of the noise components are obtained as the square root of the sample second moments of the quote-to-quote continuously-compounded returns. The data are provided by the TAQ database. We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST.

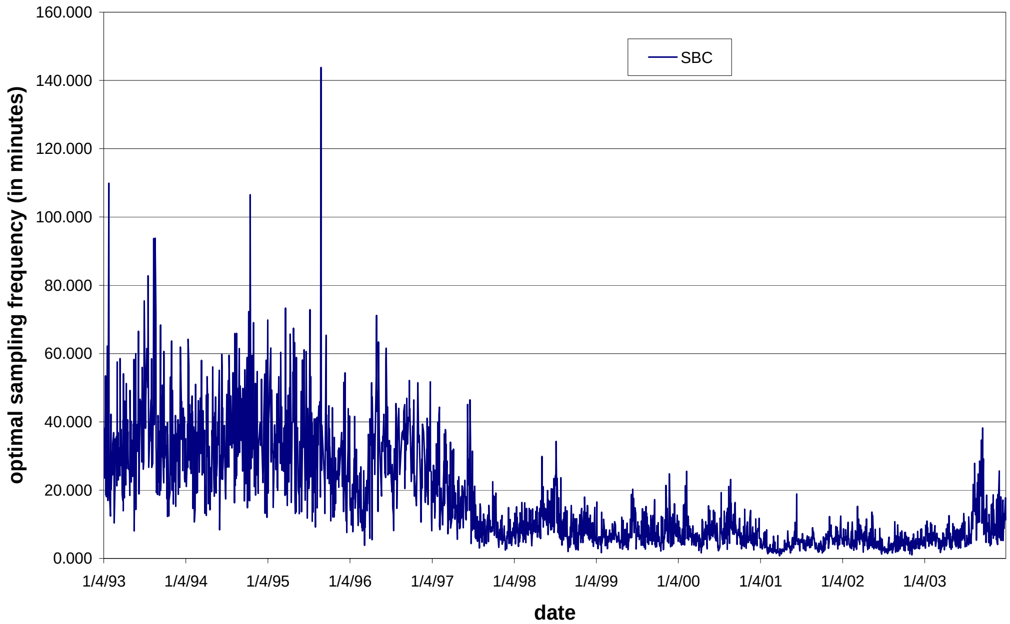

Fig 4. Time series of the daily optimal (in an MSE sense) sampling intervals for the realized variance of SBC Communications (SBC). The plot shows the optimal intervals (in minutes) at which to sample intra-daily continuously-compounded returns for the purpose of realized variance estimation. The data are provided by the TAQ database. We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST.

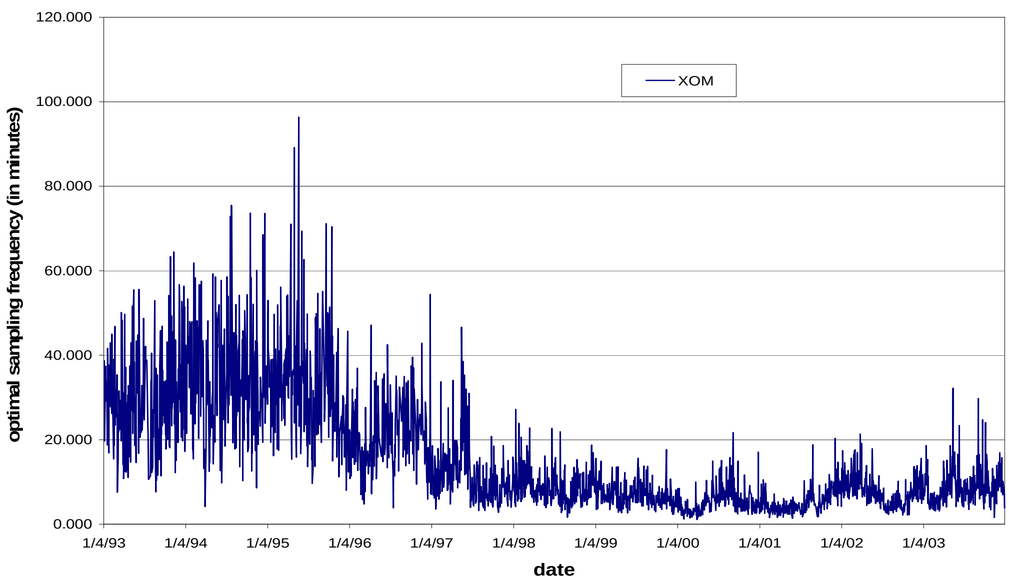

Fig 5. Time series of the daily optimal (in an MSE sense) sampling intervals for the realized variance of EXXON Mobile Corporation (XOM). The plot shows the optimal intervals (in minutes) at which to sample intra-daily continuously-compounded returns for the purpose of realized variance estimation. The data are provided by the TAQ database. We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST.

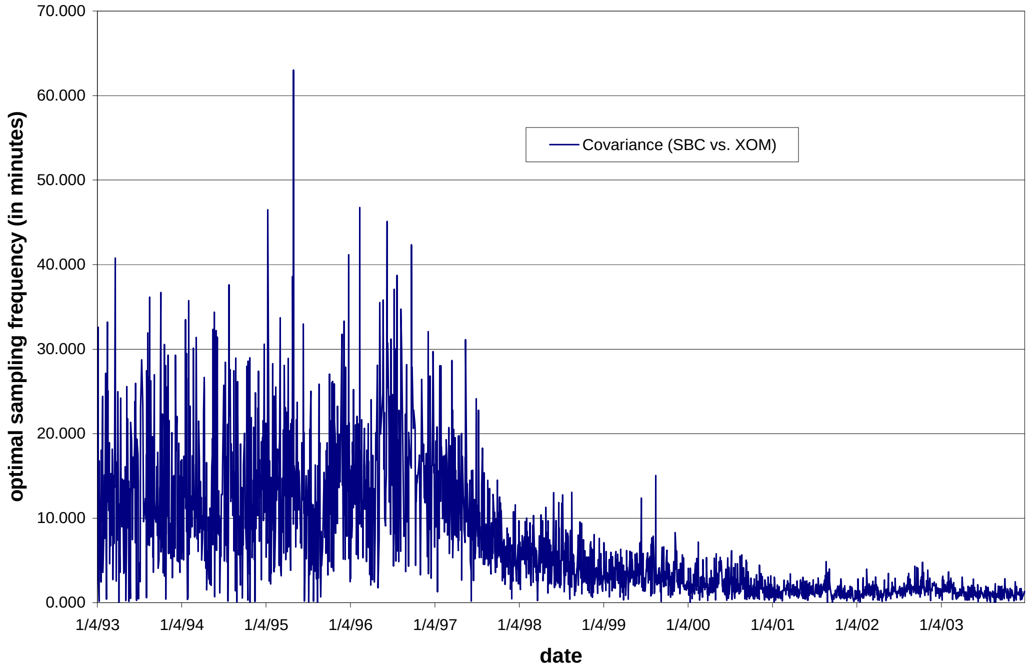

Fig 6. Time series of the daily optimal (in an MSE sense) sampling intervals for the realized covariance of SBC Communications (SBC) and EXXON Mobile Corporation (XOM). The plot shows the optimal intervals (in minutes) at which to sample intra-daily continuously-compounded returns for the purpose of realized covariance estimation. The data are provided by the TAQ database. We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST.

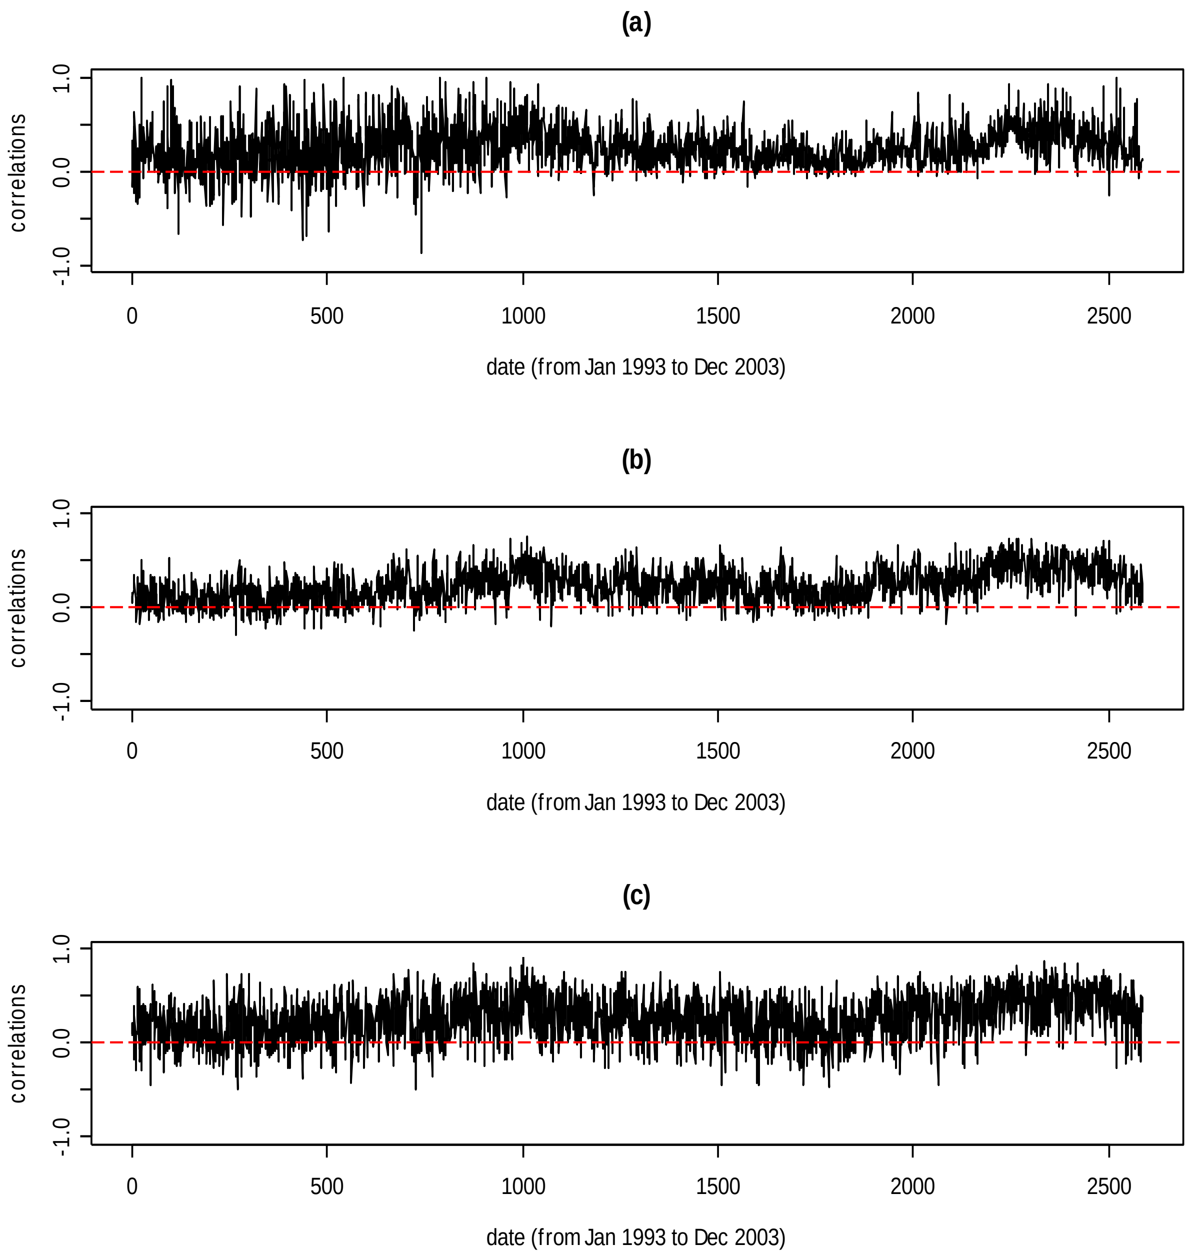

Fig 7. Daily time series of the realized correlations of SBC Communications (SBC) and EXXON Mobile Corporation (XOM). The realized correlations are obtained as the ratio of the realized covariances and the product of the realized standard deviations. Panel (a) contains realized correlations obtained by constructing the realized covariances and the realized standard deviations using an optimal (in an MSE sense) number of intra-daily continuously-compounded returns. Panel (b) contains realized correlations obtained by constructing the realized covariances and the realized standard deviations using 5-minute intervals to sample continuously-compounded returns. Panel (c) contains realized correlations obtained by constructing the realized covariances and the realized standard deviations using 15-minute intervals to sample continuously-compounded returns. The data are provided by the TAQ database. We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST.

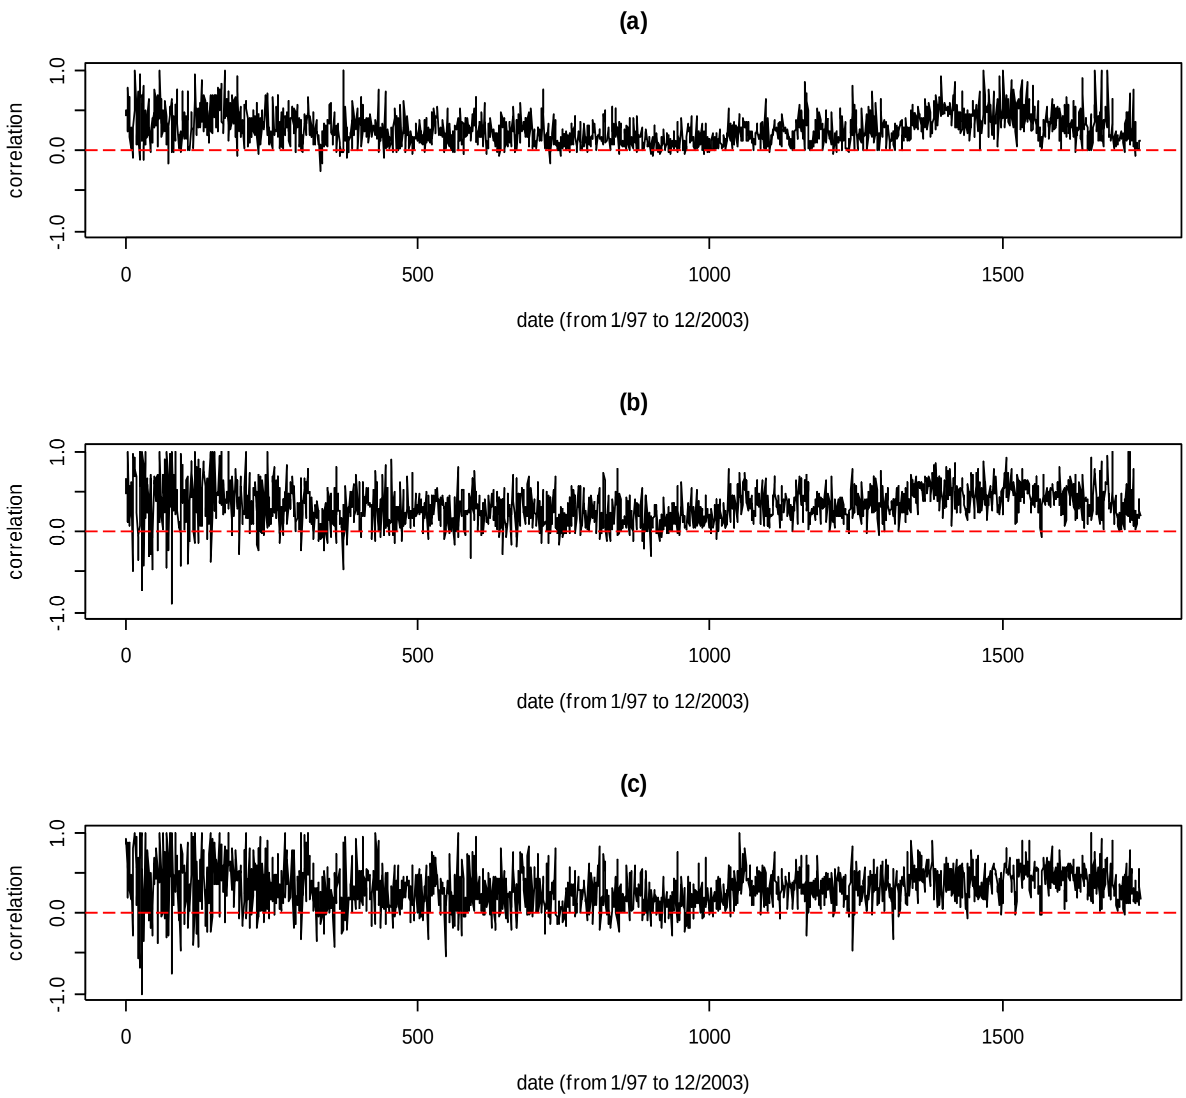

Fig 8. Daily time series of the realized correlations of SBC Communications (SBC) and EXXON Mobile Corporation (XOM). The realized correlations are obtained as the ratio of the realized covariances and the product of the realized standard deviations. We construct the realized covariances and the realized variances by using an optimal (in an MSE sense) number of intra-daily continuously-compounded returns. Panel (a) contains realized correlations obtained by constructing the realized covariances without leads and lags. Panel (b) contains realized correlations obtained by constructing the realized covariances with 1 lead and 1 lag. Panel (c) contains realized correlations obtained by constructing the realized covariances with 2 leads and 2 lags. The data are provided by the TAQ database. We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST.

(a)

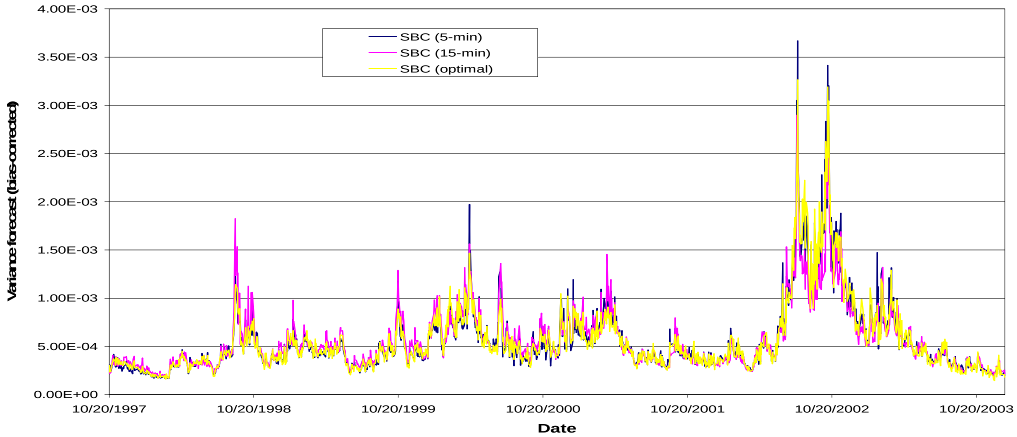

(b)

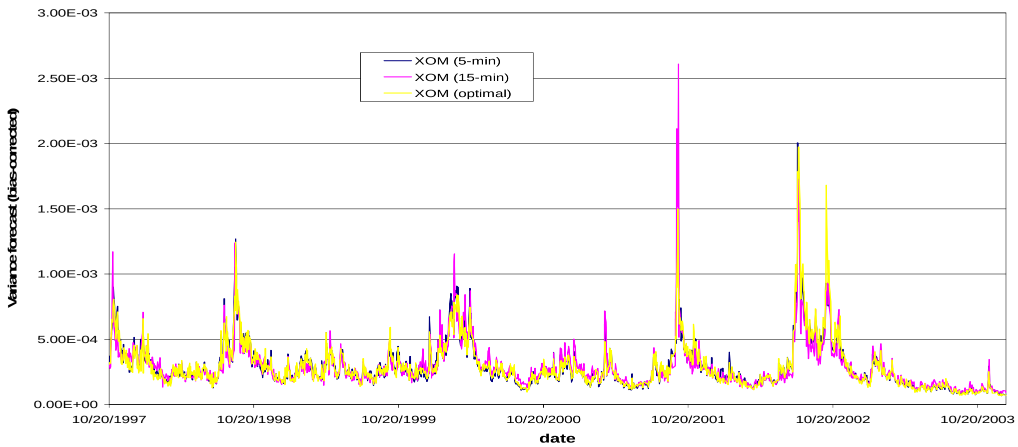

Fig 9. Daily time series of realized variance forecasts for SBC Communications (SBC) and EXXON Mobile Corporations (XOM). Panel (a) contains variance forecasts for SBC based on optimally-sampled (in an MSE sense) realized variances, on realized variances constructed using 5-minute returns, and on realized variances constructed using 15-minute returns. Panel (b) contains variance forecasts for XOM based on optimally-sampled (in an MSE sense) realized variances, on realized variances constructed using 5-minute returns, and on realized variances constructed using 15-minute returns. The optimally-sampled daily realized variances are computed by summing optimally-sampled squared continuously-compounded mid-quote returns. The 5-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. The one-day ahead forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. For SBC the estimated d values are equal to 0.48 in the case of the optimally-sampled realized variances, 0.45 in the case of the 5-minute realized variances, and 0.4 in the case of the 15-minute realized variances. For XOM the estimated d values are equal to 0.48 in the case of the optimally-sampled realized variances, 0.49 in the case of the 5-minute realized variances, and 0.36 in the case of the 15-minute realized variances. The data are provided by the TAQ database. We use SBC and XOM quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

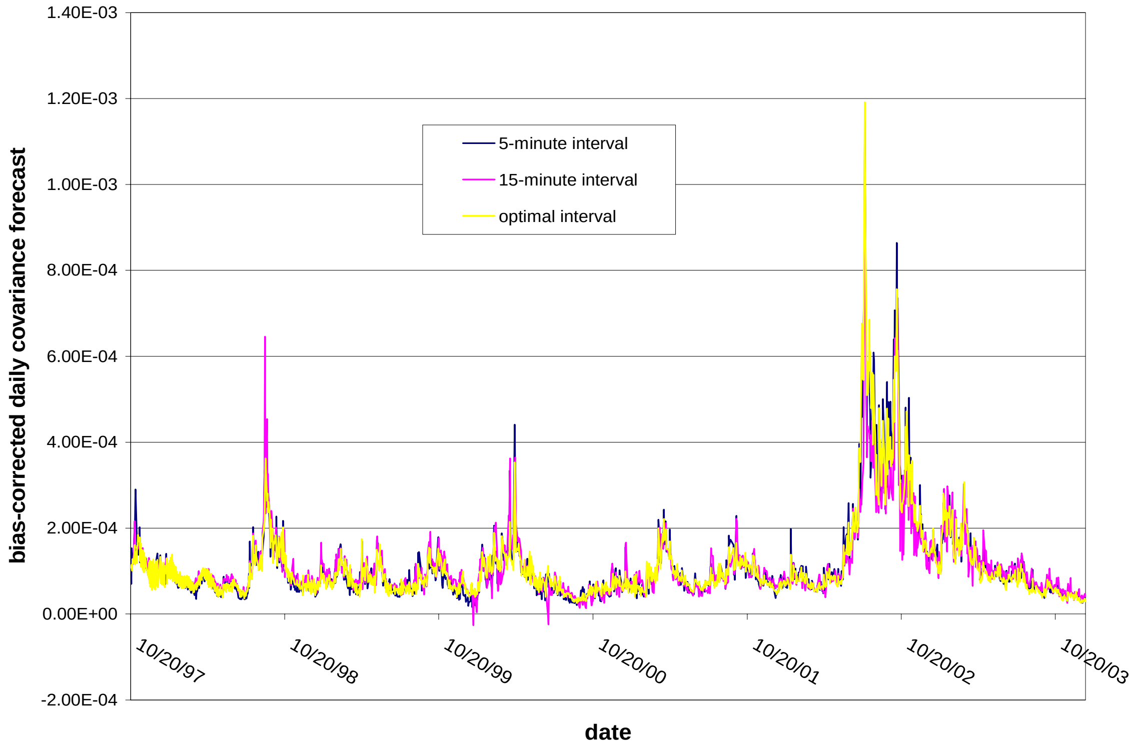

Fig 10. Daily time series of realized covariance forecasts for SBC Communications (SBC) and EXXON Mobile Corporations (XOM). The figure contains covariance forecasts for SBC and XOM based on optimally-sampled (in an MSE sense) realized covariances, on realized covariances constructed using 5-minute returns, and on realized covariances constructed using 15-minute returns. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The optimally-sampled daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. The one-day ahead forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The estimated d values are equal to 0.42 in the case of the optimally-sampled realized covariances, 0.40 in the case of the 5-minute realized covariances, and 0.31 in the case of the 15-minute realized covariances. The data are provided by the TAQ database. We use SBC and XOM quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

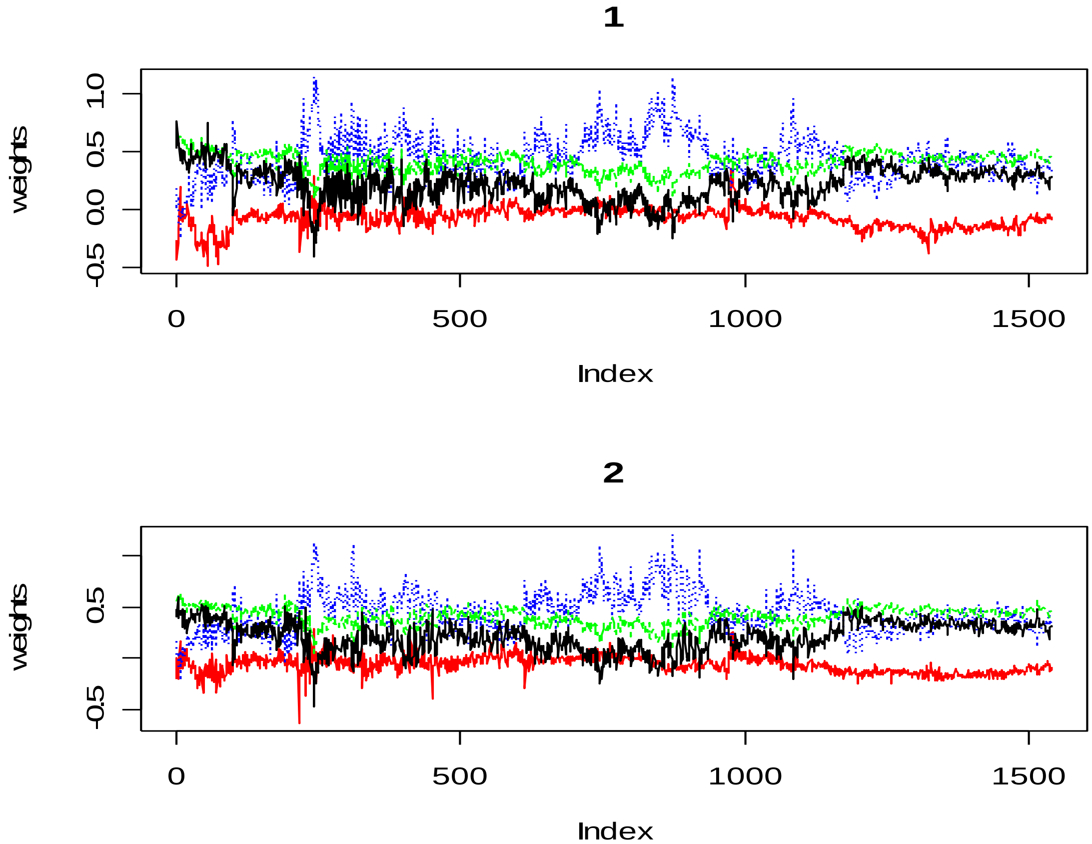

Fig 11. Daily time series of portfolio weights for SBC Communications (SBC), EXXON Mobile Corporation (XOM), Merril Lynch (MER), and the risk-free asset. The portfolio weights are obtained for an investor who minimizes the variance of a portfolio containing SBC, XOM, MER, and the risk-free asset subject to an annual target expected return on the portfolio equal to 12%. The annual risk-free rate is set equal to 3%. We report daily time-series of portfolio weights for SBC (in red), XOM (in blue), MER (in green), and the risk-free asset (in black) based on optimally-sampled realized variances/covariances (Panel 1) and realized variances/covariances constructed using 5-minute returns (Panel 2). In all cases, the realized covariances are constructed using 1 lead and 1 lag. The optimally-sampled daily realized variances/covariances are computed by summing the outer products of the optimally-sampled continuously-compounded vector returns based on mid-quotes. The 5-minute daily realized variances/covariances are computed by summing the outer-products of the continuously-compounded vector returns constructed by sampling mid-quotes every 5 minutes. The one-day ahead variance/covariance forecasts to be used in the portfolio problem are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We employ SBC, XOM, and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We use 200 daily observations to construct the first forecast. The total number of one-day ahead forecasts is equal to 1,537.

Table 1
GPH estimates of the d parameters of the realized variances and covariances of SBC Communications (SBC), EXXON Mobile Corporation (XOM), and Merrill Lynch (MER)[a]

	d estimates
SBC realized variance (5-min)	0.45
SBC realized variance (15-min)	0.40
SBC realized variance (opt)	0.48
XOM realized variance (5-min)	0.49
XOM realized variance (15-min)	0.36
XOM realized variance (opt)	0.48
MER realized variance (5-min)	0.40
MER realized variance (15-min)	0.37
MER realized variance (opt)	0.43
SBC/XOM Realized covariance (5-min)	0.40
SBC/XOM Realized covariance (15-min)	0.31
SBC/XOM Realized covariance (opt)	0.42
SBC/MER Realized covariance (5-min)	0.36
SBC/MER Realized covariance (15-min)	0.32
SBC/MER Realized covariance (opt)	0.39
XOM/MER Realized covariance (5-min)	0.37
XOM/MER Realized covariance (15-min)	0.28
XOM/MER Realized covariance (opt)	0.39

[a]The data are provided by the TAQ database. We use quotes sampled between 10 am and 4 pm from January 1993 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. The optimally-sampled (in an MSE sense) daily realized variances are computed by summing optimally-sampled squared continuously-compounded mid-quote returns. The 5-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. The optimally-sampled (in an MSE sense) daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. We use one lead and one lag in the covariance estimates. The d parameters are estimated using the classical GPH estimator (Geweke and Porter-Hudak, 1983).

Table 2
One-day ahead predictive regression of daily optimally-sampled realized covariances using optimally-sampled realized covariances, 5-minute realized covariances, and 15-minute realized covariances. The stocks are SBC and XOM[a]

	Coefficient	T-statistics
Intercept	*-1.68E-05*	*-3.664328*
C^*	*0.753398*	*8.090837*
5mC	*0.362997*	*3.389481*
15mC	*-0.030077*	*-0.375245*
	R^2=50.2%	

[a]The table contains the results of a regression of daily optimally-sampled realized covariances on one-day ahead forecasts constructed using daily optimally-sampled realized covariances (C^*), daily 5-minute realized covariances (5mC), and daily 15-minute realized covariances (15mC). We use quote data for SBC Communications (SBC) and EXXON Mobile Corporation (XOM) over the period between January 1997 and December 2003. The data come from the TAQ database. The optimally-sampled daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized co-variances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use SBC and XOM quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 3
One-day ahead predictive regression of daily 5-minute realized covariances using optimally-sampled realized co-variances, 5-minute realized covariances, and 15-minute realized covariances. The stocks are SBC and XOM[a]

	Coefficient	T-statistics
Intercept	*-1.35E-05*	*-2.647061*
C^*	*0.336372*	*3.259298*
5mC	*0.861853*	*7.261036*
15mC	*-0.140945*	*-1.586617*
	R^2=45.4%	

[a]The table contains the results of a regression of daily 5-minute realized covariances on one-day ahead forecasts constructed using daily optimally-sampled realized covariances (C^*), daily 5-minute realized co-variances (5mC), and daily 15-minute realized covariances (15mC). We use quote data for SBC Communications (SBC) and EXXON Mobile Corporation (XOM) over the period between January 1997 and December 2003. The data come from the TAQ database. The optimally-sampled daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use SBC and XOM quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 4
One-day ahead predictive regression of daily 15-minute realized covariances using optimally-sampled realized covariances, 5-minute realized covariances, and 15-minute realized covariances. The stocks are SBC and XOM[a]

	Coefficient	T-statistics
Intercept	*-9.16E-06*	*-1.421340*
C^*	*0.443386*	*3.396553*
5mC	*0.431432*	*2.873624*
15mC	*0.128799*	*1.146269*
	R^2=29.6%	

[a]The table contains the results of a regression of daily 15-minute realized co-variances on one-day ahead forecasts constructed using daily optimally-sampled realized covariances (C^*), daily 5-minute realized covariances (5mC), and daily 15-minute realized covariances (15mC). We use quote data for SBC Communications (SBC) and EXXON Mobile Corporation (XOM) over the period between January 1997 and December 2003. The data come from the TAQ database. The optimally-sampled daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use SBC and XOM quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 5
One-day ahead predictive regression of daily optimally-sampled realized covariances using optimally-sampled realized covariances, 5-minute realized covariances, and 15-minute realized covariances. The stocks are SBC and MER[a]

	Coefficient	T-statistics
Intercept	*-5.03E-06*	*-0.696253*
C^*	*0.823106*	*11.22222*
5mC	*0.066636*	*0.719954*
15mC	*0.080643*	*1.242095*
	R^2=37.5%	

[a]The table contains the results of a regression of daily optimally-sampled realized covariances on one-day ahead forecasts constructed using daily optimally-sampled realized covariances (C^*), daily 5-minute realized covariances (5mC), and daily 15-minute realized covariances (15mC). We use quote data for SBC Communications (SBC) and Merrill Lynch (MER) over the period between January 1997 and December 2003. The data come from the TAQ database. The optimally-sampled daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use SBC and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 6
One-day ahead predictive regression of daily 5-minute realized covariances using optimally-sampled realized covariances, 5-minute realized covariances, and 15-minute realized covariances. The stocks are SBC and MER[a]

	Coefficient	T-statistics
Intercept	*-8.88E-06*	*-1.017351*
C^*	*0.602449*	*6.796091*
5mC	*0.277230*	*2.478303*
15mC	*0.125552*	*1.600015*
	R^2=30.6%	

[a]The table contains the results of a regression of daily 5-minute realized covariances on one-day ahead forecasts constructed using daily optimally-sampled realized covariances (C^*), daily 5-minute realized covariances (5mC), and daily 15-minute realized covariances (15mC). We use quote data for SBC Communications (SBC) and Merrill Lynch (MER) over the period between January 1997 and December 2003. The data come from the TAQ database. The optimally-sampled daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use SBC and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 7
One-day ahead predictive regression of daily 15-minute realized covariances using optimally-sampled realized covariances, 5-minute realized covariances, and 15-minute realized covariances. The stocks are SBC and MER[a]

	Coefficient	T-statistics
Intercept	*-2.82E-06*	*-0.238178*
C^*	*0.640613*	*5.316054*
5mC	*-0.043071*	*-0.283240*
15mC	*0.367146*	*3.441880*
	R^2=18.4%	

[a]The table contains the results of a regression of daily 15-minute realized covariances on one-day ahead forecasts constructed using daily optimally-sampled realized covariances (C^*), daily 5-minute realized covariances (5mC), and daily 15-minute realized covariances (15mC). We use quote data for SBC Communications (SBC) and Merrill Lynch (MER) over the period between January 1997 and December 2003. The data come from the TAQ database. The optimally-sampled daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use SBC and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 8

One-day ahead predictive regression of daily optimally-sampled realized covariances using optimally-sampled realized covariances, 5-minute realized covariances, and 15-minute realized covariances. The stocks are XOM and MER[a]

	Coefficient	T-statistics
Intercept	-6.50E-06	-1.301899
C^*	0.491266	6.840792
5mC	0.435125	4.715864
15mC	0.057666	1.247487
	R^2=41.3%	

[a]The table contains the results of a regression of daily optimally-sampled realized covariances on one-day ahead forecasts constructed using daily optimally-sampled realized covariances (C^*), daily 5-minute realized covariances (5mC), and daily 15-minute realized covariances (15mC). We use quote data for EXXON Mobile Corporation (XOM) and Merrill Lynch (MER) over the period between January 1997 and December 2003. The data come from the TAQ database. The optimally-sampled daily realized co-variances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use XOM and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 9
One day ahead predictive regression of daily 5-minute realized covariances using optimally-sampled realized covariances, 5-minute realized covariances, and 15-minute realized covariances. The stocks are XOM and MER[a]

	Coefficient	T-statistics
Intercept	-1.02E-05	-1.734732
C^*	0.310644	3.666744
5mC	0.630739	5.794621
15mC	0.076045	1.394490
	R^2=35.7%	

[a]The table contains the results of a regression of daily 5-minute realized covariances on one-day ahead forecasts constructed using daily optimally-sampled realized covariances (C^*), daily 5-minute realized covariances (5mC), and daily 15-minute realized covariances (15mC). We use quote data for EXXON Mobile Corporation (XOM) and Merrill Lynch (MER) over the period between January 1997 and December 2003. The data come from the TAQ database. The optimally-sampled daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use XOM and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 10
One-day ahead predictive regression of daily 15-minute realized covariances using optimally-sampled realized covariances, 5-minute realized covariances, and 15-minute realized covariances. The stocks are XOM and MER[a]

	Coefficient	T-statistics
Intercept	-5.04E-06	-0.777565
C^*	0.402043	4.129779
5mC	0.613572	5.123069
15mC	-0.124649	-1.931235
	R^2=24.3%	

[a]The table contains the results of a regression of daily 15-minute realized covariances on one-day ahead forecasts constructed using daily optimally-sampled realized covariances (C^*), daily 5-minute realized covariances (5mC), and daily 15-minute realized covariances (15mC). We use quote data for EXXON mobile corporation (XOM) and Merrill Lynch (MER) over the period between January 1997 and December 2003. The data come from the TAQ database. The optimally-sampled daily realized co-variances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use XOM and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 11
Annualized fees (in basis points) that a mean-variance investor would be willing to pay to perform volatility-timing using optimally-sampled realized variances and covariances versus 5-minute and 15-minute realized variances and covariances[a]

Target: 9%

Period		λ **=2**	λ **=7**	λ **=10**
10/1997 – 12/2003	*opt/5min*	*3.11*	*10.91*	*15.59*
	opt/15min	*3.94*	*13.79*	*19.70*

Target: 12%

Period		λ **=2**	λ **=7**	λ **=10**
10/1997 – 12/2003	*opt/5min*	*7.01*	*24.56*	*35.09*
	opt/15min	*8.86*	*31.03*	*44.34*

Target: 15%

Period		λ **=2**	λ **=7**	λ **=10**
10/1997 – 12/2003	*opt/5min*	*12.47*	*43.67*	*62.38*
	opt/15min	*15.76*	*55.17*	*78.82*

[a]The table contains the annualized fees (in basis points) that a conditional mean-variance investor with absolute risk-aversion parameter λ =2,7, and 10 would be willing to pay to perform volatility-timing using optimally-sampled (in an MSE sense) realized variances/covariances (opt) versus 5-minute realized variances/covariances (5min) and 15-minute realized variances/covariances (15min). The portfolio weights are obtained by minimizing the variance of a portfolio containing SBC Communications (SBC), EXXON Mobile Corporation (XOM), Merrill Lynch (MER), and the risk-free asset for a given (target) expected return on the portfolio. The annual risk-free rate is set equal to 3%. The utility gains are assessed using the procedure suggested by West et al. (1993) and Fleming et al. (2001, 2003), as described in the main text. The optimally-sampled daily realized variances are computed by summing optimally-sampled squared continuously-compounded mid-quote returns. The 5-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. The optimally-sampled daily realized co-variances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead variance/covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use SBC, XOM, and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 12
Annualized fees (in basis points) that a mean-variance investor would be willing to pay to perform volatility-timing using optimally-sampled realized covariances and 5-minute (15-minute) variances versus 5-minute (15-minute) realized variances and covariances[a]

Target: 9%

Period		λ **=2**	λ **=7**	λ **=10**
10/1997 – 12/2003	*opt/5min*	*2.34*	*8.20*	*11.71*
	opt/15min	*2.15*	*7.53*	*10.75*

Target: 12%

Period		λ **=2**	λ **=7**	λ **=10**
10/1997 – 12/2003	*opt/5min*	*5.27*	*18.46*	*26.37*
	opt/15min	*4.84*	*16.94*	*24.20*

Target: 15%

Period		λ **=2**	λ **=7**	λ **=10**
10/1997 – 12/2003	*opt/5min*	*9.37*	*32.81*	*46.87*
	opt/15min	*8.60*	*30.12*	*43.02*

[a]The table contains the annualized fees (in basis points) that a conditional mean-variance investor with absolute risk-aversion parameter λ =2,7, and 10 would be willing to pay to perform volatility-timing using optimally-sampled (in an MSE sense) realized covariances and 5-minute (15-minute) realized variances (opt) versus 5-minute (15-minute) realized variances/covariances. The portfolio weights are obtained by minimizing the variance of a portfolio containing SBC Communications (SBC), EXXON Mobile Corporation (XOM), Merrill Lynch (MER), and the risk-free asset for a given (target) expected return on the portfolio. The annual risk-free rate is set equal to 3%. The utility gains are assessed using the procedure suggested by West et al. (1993) and Fleming et al. (2001, 2003), as described in the main text. The optimally-sampled daily realized variances are computed by summing optimally sampled squared continuously-compounded mid-quote returns. The 5-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. The optimally-sampled daily realized covariances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead variance/covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use SBC, XOM, and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.

Table 13
Annualized fees (in basis points) that a mean-variance investor would be willing to pay to perform volatility-timing with transaction costs using optimally-sampled realized variances and covariances versus 5-minute and 15-minute realized variances and covariances[a]

Transaction cost (per dollar traded): 0.00005

Period		λ **=2**	λ **=7**	λ **=10**
10/1997 – 12/2003	*opt/5min*	*7.03*	*24.62*	*35.17*
	opt/15min	*8.87*	*31.06*	*44.38*

Transaction cost (per dollar traded): 0.0001

Period		λ **=2**	λ **=7**	λ **=10**
10/1997 – 12/2003	*opt/5min*	*7.05*	*24.67*	*35.25*
	opt/15min	*8.88*	*31.09*	*44.42*

Transaction cost (per dollar traded): 0.0025

Period		λ **=2**	λ **=7**	λ **=10**
10/1997 – 12/2003	*opt/5min*	*7.77*	*27.20*	*38.86*
	opt/15min	*9.51*	*33.30*	*47.57*

[a]The table contains the annualized fees (in basis points) that a conditional mean-variance investor with absolute risk-aversion parameter λ =2,7, and 10 would be willing to pay to perform volatility-timing with transaction costs using optimally-sampled (in an MSE sense) realized variances/covariances (opt) versus 5-minute realized variances/covariances (5min) and 15-minute realized variances/covariances (15min). The portfolio weights are obtained by minimizing the variance of a portfolio containing SBC Communications (SBC), EXXON Mobile Corporation (XOM), Merrill Lynch (MER), and the risk-free asset for a target expected return on the portfolio equal to 12%. The annual risk-free rate is set equal to 3%. The transaction cost (per dollar traded) is fixed across assets and equal to 0.00005, 0.0001, and 0.0025. The utility gains are assessed using the procedure suggested by West et al. (1993) and Fleming et al. (2001, 2003), as described in the main text. The optimally-sampled daily realized variances are computed by summing optimally-sampled squared continuously-compounded mid-quote returns. The 5-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized variances are computed by summing squared continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. The optimally-sampled daily realized co-variances are computed by summing optimally-sampled cross-products of continuously-compounded mid-quote returns. The 5-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 5 minutes. The 15-minute daily realized covariances are computed by summing cross-products of continuously-compounded returns constructed by sampling mid-quotes every 15 minutes. In all cases, the realized covariances are constructed using 1 lead and 1 lag. The one-day ahead variance/covariance forecasts are obtained by virtue of an ARFIMA(2,d,2) model. We use GPH estimates of the d parameter. The data are provided by the TAQ database. We use SBC, XOM, and MER quotes sampled between 10 am and 4 pm from January 1997 to December 2003. The quotes come from two exchanges, the NYSE and the MIDWEST. We employ 200 observations to construct the first forecast. The total number of one-day ahead forecasts is 1,537.