

de Pooter, Michiel; Martens, Martin; van Dijk, Dick

Working Paper

Predicting the Daily Covariance Matrix for S&P 100 Stocks Using Intraday Data - But Which Frequency to Use?

Tinbergen Institute Discussion Paper, No. 05-089/4

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: de Pooter, Michiel; Martens, Martin; van Dijk, Dick (2005) : Predicting the Daily Covariance Matrix for S&P 100 Stocks Using Intraday Data - But Which Frequency to Use?, Tinbergen Institute Discussion Paper, No. 05-089/4, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/86328>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



TI 2005-089/4

Tinbergen Institute Discussion Paper

Predicting the Daily Covariance Matrix for S&P 100 Stocks Using Intraday Data – But Which Frequency to Use?

Michiel de Pooter

Martin Martens

Dick van Dijk

¹ *Econometric Institute, Faculty of Economics, Erasmus University Rotterdam, and Tinbergen Institute.*

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam

Roetersstraat 31

1018 WB Amsterdam

The Netherlands

Tel.: +31(0)20 551 3500

Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam

Burg. Oudlaan 50

3062 PA Rotterdam

The Netherlands

Tel.: +31(0)10 408 8900

Fax: +31(0)10 408 9031

Please send questions and/or remarks of non-scientific nature to driessen@tinbergen.nl.

Most TI discussion papers can be downloaded at <http://www.tinbergen.nl>.

Predicting the Daily Covariance Matrix for S&P 100 Stocks Using Intraday Data - But Which Frequency to Use?*

Michiel de Pooter[†]

*Tinbergen Institute
Erasmus University Rotterdam*

Martin Martens[‡]

*Econometric Institute
Erasmus University Rotterdam*

Dick van Dijk[§]

*Econometric Institute
Erasmus University Rotterdam*

January 3, 2006

Abstract

This paper investigates the merits of high-frequency intraday data when forming mean-variance efficient stock portfolios with daily rebalancing from the individual constituents of the S&P 100 index. We focus on the issue of determining the optimal sampling frequency as judged by the performance of these portfolios. The optimal sampling frequency ranges between 30 and 65 minutes, considerably lower than the popular five-minute frequency, which typically is motivated by the aim of striking a balance between the variance and bias in covariance matrix estimates due to market microstructure effects such as non-synchronous trading and bid-ask bounce. Bias-correction procedures, based on combining low-frequency and high-frequency covariance matrix estimates and on the addition of leads and lags do not substantially affect the optimal sampling frequency or the portfolio performance. Our findings are also robust to the presence of transaction costs and to the portfolio rebalancing frequency.

Key words: realized volatility, high-frequency data, bias-correction, volatility timing, mean-variance analysis, tracking error.

JEL Classification Code: G11.

*We thank Federico Bandi, the editors Michael McAleer and Essie Maasoumi, and two anonymous referees for helpful comments and suggestions. Any remaining errors are ours alone.

[†]Tinbergen Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands, e-mail: depooter@few.eur.nl

[‡]Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands, e-mail: mmartens@few.eur.nl

[§]Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands, e-mail: djvandijk@few.eur.nl (corresponding author)

1 Introduction

The work of Andersen and Bollerslev (1998) has triggered a vast amount of research on the use of high-frequency data to measure, model and forecast volatility of financial asset returns. Most empirical studies on this topic of ‘realized volatility’ focus exclusively on the variance of individual asset returns, see Andersen, Bollerslev, Diebold and Ebens (2001), Andersen, Bollerslev, Diebold and Labys (2001), Areal and Taylor (2002), Thomakos and Wong (2003), Martens *et al.* (2004), Pong *et al.* (2004), and Koopman *et al.* (2005), among others. Many financial applications such as risk management and portfolio construction, however, require estimates or forecasts of the entire covariance matrix, making covariances or correlations between returns on different assets at least as important. Yet only limited (empirical) research has addressed the merits of high-frequency data for potential economic or forecasting gains in a multivariate context. Andersen *et al.* (2003) use a vector autoregressive (VAR) framework for the daily realized variances and covariance of two exchange rates (DEM/USD and YEN/USD) based on 30-minute returns, but they only consider the statistical accuracy of (co-)variance forecasts. Fleming *et al.* (2003) use five-minute returns on three actively traded futures contracts (S&P 500 index, Treasury bonds, and gold) to show that a mean-variance efficient investor would be willing to pay 50 to 200 basis points per annum for being able to use daily covariance matrix forecasts based on high-frequency intraday returns instead of daily returns. Similarly, Liu (2004) constructs and assesses the performance of the minimum variance portfolio and the minimum tracking error portfolio (tracking the S&P 500 index) using five-minute returns for the 30 Dow Jones index constituents.

These three studies have in common that they motivate the selected intraday sampling frequency as a trade-off between accuracy and potential biases due to market microstructure effects. The sensitivity of the results to the choice of sampling frequency used in constructing realized covariances is not investigated though. Martens (2004) demonstrates that non-trading, non-synchronous trading, and bid-ask bounce are indeed crucial determinants of the optimal sampling frequency that minimizes the Mean Squared Error (MSE) for measuring, and hence forecasting, the covariance matrix. The MSE is the sum of the squared bias and the variance of the realized (co-)variance. High sampling frequencies lead to a potentially large upward

bias in realized variances due to bid-ask bounce and to a substantial downward bias in realized covariances due to non-synchronous trading. On the other hand, the variance of both realized variances and realized covariances usually decreases with higher sampling frequencies. As the degree of non-trading, non-synchronous trading, and bid-ask bounce varies widely across assets, the appropriate sampling frequency in a particular application needs to be investigated carefully.¹

In this study we examine the economic significance of determining the optimal sampling frequency, in the context of constructing mean-variance efficient portfolios from the individual constituents of the S&P 100 index. Our analysis builds on the framework developed in Fleming *et al.* (2001, 2003). In particular, we consider a risk-averse investor who constructs minimum variance portfolios and minimum tracking error portfolios with daily rebalancing, where portfolio risk is minimized either globally or subject to a fixed target return. We focus on pure volatility-timing strategies, in the sense that the portfolio weights are determined exclusively by forecasts of the daily conditional covariance matrix, which in turn is constructed using the realized covariance matrix with the sampling frequency of intraday returns ranging from one minute to 130 minutes. The economic value of using the optimal sampling frequency is assessed by comparing portfolio performance across this range of sampling frequencies. In particular, we consider the fee that the investor would be willing to pay to switch from one frequency to another.

We also examine how different bias- and variance-reduction techniques affect the choice of sampling frequency. First, we explore the usefulness of the two time-scales estimator proposed by Zhang *et al.* (2005),² which combines the realized covariance matrix constructed using subsampling at a certain frequency and the realized covariance matrix constructed using the highest possible sampling frequency. Sub-sampling makes use of the fact that, for example, five-minute returns for a trading session starting at 9:30 could not only be measured using the intervals 9:30-9:35, 9:35-9:40, ..., but also using 9:31-9:36, 9:36-9:41, ..., etc. As explained in more

¹The issue of sampling frequency in the presence of market microstructure noise has also been investigated in the context of univariate realized volatility measurement, see Aït-Sahalia *et al.* (2005), Bandi and Russell (2005a,b), Zhang *et al.* (2005), and Hansen and Lunde (2006), among others.

²Zhang *et al.* (2005) focus solely on estimating the variance but here we apply their approach to covariances as well, as suggested by Zhang (2005).

detail below, subsampling can be used to reduce the variance of the realized covariance estimator. However, subsampling still renders a biased estimate of the true integrated covariance matrix with the bias being a function of the covariance matrix of the microstructure noise component in the intraday returns. The two time-scales estimator attempts to correct for this bias. Second, following the idea of Scholes and Williams (1977) for estimating (illiquid) stock betas, we investigate the merits of using leads and lags in measuring the realized covariances. For all these methods we also consider the effects of transaction costs and the holding period or portfolio rebalancing frequency.

Our main findings are as follows. For both minimum variance and minimum tracking error portfolios, using daily conditional covariance matrix forecasts based on high-frequency intraday returns instead of daily returns considerably improves portfolio performance. For the global minimum risk portfolios, the optimal sampling frequency for the S&P 100 constituents ranges between 30 and 65 minutes, considerably lower than the popular five-minute frequency. The same result occurs for minimum variance portfolios subject to a target return. Here, the Sharpe ratio increases from 0.6 to 0.8 going from daily to intraday returns, and a risk-averse investor would be willing to pay between 150 and 400 basis points per year to capture this gain in portfolio performance. In contrast, for the minimum tracking error portfolio subject to a target return the optimal sampling frequency appears to be much higher at one- to two-minutes. The performance gains are substantial, with the information ratio increasing from 0.1 to 0.4. The fee a risk-averse investor might pay for this enhanced performance ranges between 100 and 180 basis points per year.

The above findings are robust to the use of the two time-scales estimator and the lead-lag bias correction procedure. Both of these techniques marginally improve the performance for the minimum variance portfolios and the minimum tracking error portfolios. However, selecting the appropriate sampling frequency appears to be much more important than choosing between different bias- and variance-reduction techniques for the realized covariance matrices.

For the target return portfolios we find that turnover is lower when using intraday data, hence in the presence of transaction costs an investor is willing to pay even more for covariance forecasts based on high-frequency data. The opposite is true for the target excess return portfolios. Lowering the rebalancing frequency from daily

to weekly or monthly obviously reduces transaction costs, while at the same time having a similar or even better performance. Reducing the rebalancing frequency in the presence of transaction costs is especially beneficial for the minimum tracking error portfolios based on high-frequency data.

As noted above, the issue of the optimal sampling frequency in the presence of market microstructure noise has been investigated quite heavily, but in the context of univariate realized volatility measurement, see the references in footnote 1. In concurrent and independent work Bandi *et al.* (2005) derive a theoretical expression for the optimal sampling frequency to compute realized covariances. Applying this expression to our data we find an average³ optimal sampling frequency of 1.4 minutes, nowhere near the optimal 65 minutes. Apparently the assumptions under which Bandi *et al.* (2005) derive their optimal sampling frequency estimator are not appropriate for the S&P100 constituents.

The remainder of this paper is organized as follows. Section 2 describes the data and the construction of the realized covariances. The mean-variance methodology is presented in Section 3. Results are discussed in Section 4. Section 5 concludes.

2 Data

The data set was obtained from Price-Data.com⁴ and consists of open, high, low, and close transaction prices at the one-minute sampling frequency for the June 2004 S&P 100 index constituents, covering the period from April 16, 1997 until June 18, 2004 (1804 trading days). We disregard stocks for which the price series start at a later date, leaving 78 stocks for the analysis. The appendix provides a list of ticker symbols and company names. The data also comprise all (tick-by-tick) transaction prices of the S&P 500 index futures from April 16, 1997, through May 27, 2004. We follow the conventional practice of using the futures contract with the largest trading volume. This typically is the contract nearest to maturity, until a week before maturity when the next nearest contract takes over. Since the stock files miss

³The approach of Bandi *et al.* (2005) considers each pair of stocks separately such that the optimal sampling frequency may differ across (co)variances. The resulting realized covariance matrix then is not guaranteed to be positive definite. Hence it is logical to employ an average optimal frequency to be used for all stock pairs.

⁴<http://www.price-data.com/>

April 9, 2003, and the futures files miss March 30, 2003 and May 3, 2004, this leaves 1788 common trading days from April 16, 1997, through May 27, 2004.

For each day t , we divide the trading session on the NYSE, which runs from 9:30 EST until 16:00 EST (390 minutes), into I intervals of equal length $h \equiv 1/I$, normalizing the daily interval to unity for ease of notation. For example, $I = 78$ for the five-minute sampling frequency. Let p_{t-1+ih} denote the $(N \times 1)$ vector of log close transaction prices and let $r_{t-1+ih,h} \equiv p_{t-1+ih} - p_{t-1+(i-1)h}$ denote the $(N \times 1)$ vector of returns for the i^{th} intraday period on day t , for $i = 2, \dots, I$, where $N = 78$ is the number of stocks. The return for the first intraday period, $r_{t-1+h,h}$, is defined as the difference between the log close and open transaction prices during that interval. The realized covariance matrix $V_{t,h}$ is defined as

$$V_{t,h} = r_{t,c-o}r'_{t,c-o} + \sum_{i=1}^I r_{t-1+ih,h}r'_{t-1+ih,h} \quad (1)$$

where $r_{t,c-o}$ is the $(N \times 1)$ vector of close-to-open (overnight) returns from day $t-1$ (close) to day t (open).⁵ Martens (2002) documents that the overnight volatility represents an important fraction of total daily volatility, hence incorporating the cross-product of overnight returns as in (1) is important for accurately measuring (co-)variances, see also Fleming *et al.* (2003) and Hansen and Lunde (2005) for discussion. For the daily frequency the realized (co-)variance matrix V_t is defined as the outer product of the daily (close-to-close) returns, denoted by r_t , that is $V_t = r_t r'_t$.

Table 1, Panel A, illustrates some characteristic features of the daily realized variances and covariances by showing the mean (across stocks and across trading days) and variance for sampling frequencies of $390h = 1, 2, 3, 5, 10, 15, 30, 65$ and 130 minutes, such that in all cases the corresponding I intra-day intervals completely cover the 390-minute trading day. Several familiar patterns arise. First, the average realized variance increases with the sampling frequency (except for frequencies below 30 minutes). Bid-ask bounce induces negative autocorrelations in returns when prices are sampled more frequently leading to an upward bias in the realized variance. For example, the average variance using daily returns is 7.386 (corresponding to an

⁵For obvious reasons the overnight return from 10 to 17 September, 2001 (the first trading day after 9/11) has been dropped.

annualized standard deviation of about 43%), whereas it is 9.494 for one-minute returns. Second, the average realized covariance decreases monotonically with the sampling frequency, where this downward bias can be attributed to non-synchronous trading, i.e. not every stock trades in each (intraday) interval or exactly at the end of each interval. The average covariance using one-minute returns is 0.826, whereas for daily data it is almost double at 1.568. Third, the variance of the realized (co-)variances becomes smaller for higher frequencies, simply because more data points are used. Hence in general for realized (co-)variances the bias increases and the variance decreases for higher sampling frequencies.⁶

- insert Table 1 about here -

One way to reduce the variance of realized covariances, given a particular sampling frequency, is to employ subsampling as first suggested in Zhang *et al.* (2005) in this context. In particular, the grid of x -minute intervals can be laid over the trading day in x different ways. For example, for the three-minute frequency rather than starting with the interval 9:30-9:33 one could also start with 9:31-9:34 or 9:32-9:35. In this way three ‘subsamples’ are created and each of these can be used to compute the realized covariance matrix. The final realized covariance matrix is then taken to be the average across subsamples. A practical problem with this procedure is how to treat the loose ends at the start and the end of the trading session. Here the start of the day is added to the overnight return, while the end of the day is omitted. The covariances measured during the trading session are proportionally inflated for the missing part of the trading session. Unreported summary statistics for the realized (co-)variances that are obtained with this procedure show that, in general, the effects of subsampling are ambiguous. There is a minor reduction in the variance of the realized covariances for the two- to 30-minute frequencies, but an increase in the variance of the realized variances, which becomes quite substantial for the lower sampling frequencies.

Zhang *et al.* (2005) suggest a bias-correction procedure for the subsampling estimator as described above using the realized covariance matrix obtained with the highest available sampling frequency. The essential argument is that subsampling

⁶An exception is the realized variance at the one- and two-minute frequencies, where also the variance increases due to the increased importance of bid-ask bounce.

still renders a biased estimate of the true integrated volatility with the bias being a function of the (co-)variance of the noise in the return processes. In fact, the realized (co-)variance estimator using the highest possible frequency consistently estimates this noise (co-)variance and can therefore be used to reduce and potentially even eliminate the bias of the subsampling estimator. Based on this idea, the two time-scales estimator $V_{t,h}^{TTS}$ is obtained as

$$V_{t,h}^{TTS} = \frac{I_{Max}}{I_{Max} - 1} \left(V_{t,h}^{SubS} - \frac{I}{I_{Max}} V_{t,h}^{Max} \right), \quad (2)$$

where $V_{t,h}^{SubS}$ is the subsampling estimator using I returns over intervals of $390h$ minutes as described above, and $V_{t,h}^{Max}$ is the realized covariance matrix based on the highest possible sampling frequency with I_{Max} intraday return observations. In our case this is the one-minute frequency such that $I_{Max} = 390$.

Summary statistics for the realized (co-)variances that are obtained with the two time-scales estimator are presented in Panel B of Table 1. The bias-correction procedure appears to work quite well, especially for the higher sampling frequencies, in the sense that the mean realized (co-)variances get much closer to the mean values obtained with daily return observations. Note that this comes at the cost of increased variance of the realized (co-)variances though, except for the realized variances at the one- and two-minute frequencies.

Finally, we examine whether the downward bias in the realized covariances can be reduced by adding lead and lagged covariances to the contemporaneous cross-product of returns in the spirit of Scholes and Williams (1977) and Cohen *et al.* (1983). Similarly, this might reduce the upward bias in the realized variance due to the negative autocorrelations in high-frequency returns, see Hansen and Lunde (2005, 2006). In particular, let $\Gamma_{t,h,l}$ denote the l -th order cross-covariance matrix of intraday h -period returns, that is

$$\Gamma_{t,h,l} = \sum_{i=1}^{I-l} r_{t-1+ih,h} r'_{t-1+(i-l)h,h}.$$

The realized covariance matrix with lead and lags is then obtained as

$$V_{t,h}^{LL} = V_{t,h} + \sum_{l=1}^q d_l (\Gamma_{t,h,l} + \Gamma'_{t,h,l}), \quad (3)$$

where $V_{t,h}$ is given by (1) and the weights d_l for the leads and lags are taken to be $d_l = 1 - l/(q + 1)$. The use of these Bartlett-kernel weights guarantees that the

realized covariance matrix $V_{t,h}^{LL}$ is positive definite, see Newey and West (1987) and Andrews (1991) for discussion of alternative weighting schemes that also achieve this objective. Precisely for this reason we do not consider the equal-weighting scheme ($d_l = 1$ for all $l = 1, \dots, q$), as commonly used for estimating market betas of illiquid stocks and suggested by Zhou (1996) and Hansen *et al.* (2005) in the context of realized variance and covariances, respectively. Barndorff-Nielsen *et al.* (2004) demonstrate that the Bartlett-kernel estimator (3) and the subsampling estimator of Zhang *et al.* (2005) are almost identical.

Panel C of Table 1 present characteristics of $V_{t,h}^{LL}$ with $q = 1$.⁷ As expected, the bias in both realized variances and realized covariances is reduced for all frequencies, although to a lesser extent compared to the two time-scales estimator. For example, the average realized variance based on one-minute returns is reduced to 8.525, still considerably higher than the average daily squared return of 7.386. Similarly, the average realized covariance at the one-minute frequency is increased to 1.025, which comes closer to the average cross-product of daily returns (1.568) than the standard case. Note, however, that again the reduction in bias generally comes at the cost of increased variance. An exception is the one-minute frequency where not only the average variance is reduced and closer to the average daily squared return, but at the same time the variance is reduced from 597 to 533.

3 Methodology

3.1 Volatility-timing strategies

The benefits of high-frequency intraday data and the optimal way to employ these will be gauged by their economic value in the context of portfolio construction. In particular, we consider volatility timing strategies within the framework of conditional mean-variance analysis. We construct the minimum variance portfolio as well as the portfolio that minimizes variance given a set target return, which is denoted μ_P , allowing for daily rebalancing. To be precise, we solve the following two opti-

⁷We experimented with alternative values for q , which led to qualitatively similar findings. Detailed results are available upon request. The issue of determining the optimal value of q is beyond the scope of this paper and is left for future research.

mization problems for each day t :

$$\begin{aligned} \min_{w_t} w_t' \Sigma_t w_t \\ \text{s.t. } w_t' \iota = 1 \end{aligned} \quad (4)$$

and

$$\begin{aligned} \min_{w_t} w_t' \Sigma_t w_t \\ \text{s.t. } w_t' \mu_t = \mu_P \quad \text{and} \quad w_t' \iota = 1 \end{aligned} \quad (5)$$

where w_t is the $(N \times 1)$ vector of portfolio weights, and ι denotes an $(N \times 1)$ vector of ones. In addition, μ_t is the $(N \times 1)$ vector with conditional expected returns for the individual stocks, that is $\mu_t \equiv \mathbf{E}[r_t | \mathcal{I}_{t-1}]$, where \mathcal{I}_{t-1} denotes the information set available at the end of day $t - 1$. Similarly, Σ_t is the $(N \times N)$ conditional covariance matrix, that is $\Sigma_t \equiv \mathbf{E}[(r_t - \mu_t)(r_t - \mu_t)' | \mathcal{I}_{t-1}]$. In order to concentrate on the use of high-frequency data for estimating and forecasting (co-)variances, we assume that μ_t is constant and, moreover, set it equal to the average returns in the complete out-of-sample period.⁸ Hence, we consider pure volatility-timing strategies, in the sense that the portfolio weights are determined exclusively by forecasts of the daily conditional covariance matrix Σ_t . We return to these in Section 3.4 below.

The solution to the problem in (4), the weights for the fully invested minimum variance portfolio, is given by

$$w_{t,\text{MVP}} = \frac{\Sigma_t^{-1} \iota}{\iota' \Sigma_t^{-1} \iota}. \quad (6)$$

For the solution of the problem in (5) first weights for the maximum Sharpe ratio portfolio are computed as

$$w_{t,\text{MSR}} = \frac{\Sigma_t^{-1} \mu_t}{\mu_t' \Sigma_t^{-1} \mu_t} \quad (7)$$

and the weights for the target return portfolio are then provided by

$$w_{t,P} = \frac{\mu_{t,\text{MSR}} - \mu_P}{\mu_{t,\text{MSR}} - \mu_{t,\text{MVP}}} w_{t,\text{MVP}} + \frac{\mu_P - \mu_{t,\text{MVP}}}{\mu_{t,\text{MSR}} - \mu_{t,\text{MVP}}} w_{t,\text{MSR}} \quad (8)$$

⁸As explained below, we require part of the sample period to initialize the conditional covariance matrix estimates, which in our case equals 122 trading days. This implies that the effective sample period available for portfolio construction and evaluation runs from October 8, 1997 until May 27, 2004 (1666 trading days).

where $\mu_{t,\text{MVP}} = w'_{t,\text{MVP}}\mu_t$ and $\mu_{t,\text{MSR}} = w'_{t,\text{MSR}}\mu_t$ are the expected returns on the minimum variance portfolio and the maximum Sharpe ratio portfolio, respectively.

In addition the above analysis is repeated using the conditional mean and covariance matrix for stock returns in excess of the S&P 500 futures returns. The solution to the problem in (4) then determines the minimum tracking error portfolio, i.e. the portfolio of the 78 S&P 100 stocks that tracks the S&P 500 index most closely. Similarly the solution to the problem in equation (5) then minimizes the tracking error given a certain target level of active return (i.e. portfolio return in excess of the S&P 500 return). The use of minimum tracking error portfolios is motivated by the analysis in Chan *et al.* (1999) who demonstrate that based on minimum variance portfolios it is difficult to distinguish between different covariance matrix estimates in the presence of a dominant (market) factor. Eliminating the dominant factor, in this case by switching to tracking error portfolios, largely solves this problem.

3.2 The economic value of volatility timing

The performance of the portfolios in the different volatility timing strategies is evaluated using the ex-post daily stock returns r_t . For the minimum variance portfolio we consider the standard deviation, and for the target return portfolios we monitor the mean return, standard deviation, and Sharpe ratio, all based on ex-post returns. Similarly, for the minimum tracking error portfolio we consider the tracking error, and for the target active return portfolios we monitor the mean excess return, tracking error, and information ratio (excess return divided by tracking error), based on the ex-post daily excess returns.

Following Fleming *et al.* (2001, 2003), for the target return portfolios we assess the economic value of the different covariance matrix estimators in volatility timing strategies by determining the maximum performance fee a risk-averse investor would be willing to pay to switch from using one covariance matrix estimator to another. In particular, we assume the investor has a quadratic utility function given by

$$U(r_{t,P}) = W_0 \left(1 + r_{t,P} - \frac{\gamma}{2(1 + \gamma)} (1 + r_{t,P})^2 \right), \quad (9)$$

where $r_{t,P} = w'_{t,P}r_t$ is the ex-post portfolio return, γ is the investor's relative risk aversion and W_0 is initial wealth. In order to compare two volatility timing strategies based on different covariance matrix estimators with portfolio returns denoted as

r_{t,P_1} and r_{t,P_2} , we determine the maximum amount the investor is willing to pay to switch from the first strategy to the second. That is, we determine the value of Δ such that

$$\sum_{t=1}^T U(r_{t,P_1}) = \sum_{t=1}^T U(r_{t,P_2} - \Delta). \quad (10)$$

We interpret Δ as a performance fee and report estimates in terms of basis points on an annual basis for $\gamma = 1$ and 10.

3.3 Transaction costs and rebalancing frequency

With daily rebalancing, the turnover of the volatility timing strategies is considerable, as shown in detail below. Hence, transaction costs play a non-trivial role and should be considered in evaluating the (relative) performance of different strategies. We handle this issue as follows.

After rebalancing on day $t - 1$, the i -th stock has been given a weight $w_{i,t-1}$ in the portfolio, $i = 1, \dots, N$. The return on the i -th stock on day t is denoted as $r_{i,t}$ such that the portfolio return is $r_{t,P} = \sum_{i=1}^N w_{i,t-1} r_{i,t}$. At the moment just before rebalancing, denoted as t^- , the actual weight of the i -th stock in the portfolio therefore has changed to $w_{i,t^-} = w_{i,t-1} \frac{1+r_{i,t}}{1+r_{t,P}}$. The new weight $w_{i,t}$ for stock i follows from solving the investor's optimization problem, using time t information. The change in weight, or the required rebalancing, at time t thus is equal to $w_{i,t} - w_{i,t^-}$. We assume that transaction costs amount to a fixed percentage c on each traded dollar for any stock. Setting the initial wealth W_0 equal to 1 for simplicity, total transaction costs at time t are equal to

$$c_t = c \sum_{i=1}^N |w_{i,t} - w_{i,t^-}|,$$

such that the net portfolio return is given by $r_{t,P} - c_t$.

We report results for transaction cost levels between 2% and 20%, expressed in annualized percentage points. Note that this would be the reduction in the annualized portfolio return if the entire portfolio would have to be traded every day during a whole year, that is $\sum_{i=1}^N |w_{i,t} - w_{i,t^-}| = 1$ for $t = 1, \dots, T$.⁹

A closely related issue is that of the portfolio rebalancing frequency. If daily turnover is substantial, transaction costs may eat away a considerable part of the

⁹Fleming *et al.* (2003) use a similar approach when assessing the effect of transaction costs.

portfolio performance and it may be better to rebalance the portfolio less frequently. In fact, given a certain level of transaction costs c one may attempt to determine the optimal rebalancing frequency, where a trade-off has to be made between updating the portfolio weights using the most recent covariance matrix information and incurring higher transaction costs. We consider this challenging problem to be beyond the scope of this paper, though. We do provide some insight into the effect of the rebalancing frequency, by considering the portfolio performance if the holding period is set equal to a week or a month (or five and 21 trading days, respectively), as follows. We construct a new portfolio every day, but this is held on to for the next five (21) days. Hence, at any point in time the strategies effectively hold five (21) minimum variance portfolios, for example, each formed one day apart. To handle the problems concerned with overlapping returns, we calculate the overall return on day t as the average of all the portfolios that are held at that time.¹⁰

3.4 Conditional covariance matrix estimators

Implementation of the portfolio construction methods discussed above requires estimates or forecasts of the conditional covariance matrix Σ_t . We closely follow Fleming *et al.* (2001, 2003) by using rolling volatility estimators for Σ_t , building on the work by Foster and Nelson (1996) and Andreou and Ghysels (2002). The general rolling conditional covariance matrix estimator based on daily data is of the form

$$\widehat{\Sigma}_t = \sum_{k=1}^{\infty} \Omega_{t-k} \odot r_{t-k} r'_{t-k} \quad (11)$$

where Ω_{t-k} is a symmetric ($N \times N$) matrix of weights, and \odot denotes element-by-element multiplication. The weighting scheme is taken to be $\Omega_{t-k} = \alpha \exp(-\alpha k) \iota \iota'$, such that (11) can be rewritten as

$$\widehat{\Sigma}_t = \exp(-\alpha) \widehat{\Sigma}_{t-1} + \alpha \exp(-\alpha) r_{t-1} r'_{t-1}. \quad (12)$$

This choice is consistent with Foster and Nelson (1996) in that exponentially weighted estimators generally produce the smallest asymptotic MSE. In addition using a single

¹⁰Note that our approach here differs from Fleming *et al.* (2003). Our method of holding multiple portfolios simultaneously is commonly applied in the literature on stock selection, see Jegadeesh and Titman (1993) and Rouwenhorst (1998), among many others.

parameter (α) to control the rate at which the weights decay with lag length guarantees that $\widehat{\Sigma}_t$ is positive definite. One way of interpreting this weighting scheme is as a restricted multivariate GARCH model.¹¹ The optimal in-sample decay rate can therefore be estimated using (quasi) maximum likelihood for the model

$$r_t = \widehat{\Sigma}_t^{1/2} z_t \quad (13)$$

where $z_t \sim NID(0, I)$ and $\widehat{\Sigma}_t$ is given by (12). We estimate α using observations for the sample period October 8, 1997 until May 27, 2004 (1666 trading days). The reason for not using the sample from the first available day, April 16, 1997, onwards is that the covariance matrix estimate $\widehat{\Sigma}_t$ needs to be initialized. We use the first 122 observations as ‘burn-in’ period.

Given that the portfolios that subsequently are constructed using the weights $w_{MVP,t}$ from (6) and $w_{P,t}$ from (8) are evaluated over the same period that is used for estimating α , this raises the issue of data snooping. However, as noted by Fleming *et al.* (2001), the statistical loss function used to estimate the decay parameter is rather different from the methods used to evaluate the performance of the various portfolios. Hence, look-ahead bias probably is not too big a problem. We return to this issue in Section 4.4.

Andersen *et al.* (2003) and Barndorff-Nielsen and Shephard (2004) show that intraday returns can be used to construct (co-)variance estimates that are more efficient than those based on daily returns. Sticking to the concept of rolling estimators and facilitating a direct comparison between daily and intraday data, it is most natural to replace the daily update $r_{t-1}r'_{t-1}$ in (12) by the realized covariance matrix $V_{t-1,h}$, that is, the conditional covariance matrix is estimated using high-frequency data as

$$\widehat{\Sigma}_{t,h} = \exp(-\alpha_h)\widehat{\Sigma}_{t,h} + \alpha_h \exp(-\alpha_h)V_{t-1,h} \quad (14)$$

where α_h can again be estimated by means of maximum likelihood for the model (13), but now using $\widehat{\Sigma}_{t,h}$ instead of $\widehat{\Sigma}_t$. In addition to the realized covariance matrix $V_{t-1,h}$ obtained from the ‘basic’ form given in (1), we implement (14) using the

¹¹Fleming *et al.* (2003) show that actually using the (unrestricted) multivariate GARCH model leads to a better fit of the data as expected, but the covariance matrix forecasts result in worse portfolios than those obtained from the rolling covariance estimator. They cite the smoothness of the rolling estimator as the main reason for this.

two time-scales estimator $V_{t-1,h}^{TTS}$ given in (2) and the lead-lag corrected estimator $V_{t-1,h}^{LL}$ given in (3). As mentioned before, we examine different sampling frequencies to construct the realized covariance matrix $V_{t-1,h}$, dividing the 390-minute NYSE trading session in (nonoverlapping) intervals of 1, 2, 3, 5, 10, 15, 30, 65 or 130 minutes.

We close this section by noting that the conditional covariance matrix estimate $\widehat{\Sigma}_{t,h}$ obtained from (14) may suffer from the biases in the realized covariance matrix $V_{t-1,h}$ due to market microstructure effects. For that reason, Fleming *et al.* (2003) propose a bias-correction method based on scaling the elements of $\widehat{\Sigma}_{t,h}$ with factors determined from the contemporaneous estimates of the daily-returns-based rolling estimator $\widehat{\Sigma}_t$ obtained from (12), see also Hansen and Lunde (2005, 2006). Although we considered this approach in our analysis, we were unable to obtain satisfactory results. The key problem with this bias adjustment procedure is that it adjusts each individual element in the covariance matrix separately, with a possibly different correction factor. Hence, whereas the unadjusted covariance matrix $\widehat{\Sigma}_{t,h}$ obtained from (14) is guaranteed to be positive definite, this does not hold for the bias-adjusted matrix. In the empirical application considered in Fleming *et al.* (2003) concerning three highly-liquid future contracts, this issue turns out not to be relevant, but for our application to the 78 S&P 100 stocks we ran into this problem quite frequently due to the large number of stocks. We tried to address this in several different ways but to no avail.

4 Results

4.1 Optimal decay rates

Table 2 shows the optimal decay rates α and α_h that maximize the likelihood of the model in equations (13) with (12) for daily returns and with (14) for intraday returns at the different sampling frequencies considered. Starting with total returns (as opposed to returns in excess of the S&P 500 return) and the standard case (no two time-scales, no lead-lag correction), the optimal decay parameter increases monotonically from 0.0070 for daily data to 0.2106 for the one-minute frequency. This pattern implies that the update $V_{t-1,h}$ in (14) is given more weight when it is measured, presumably more accurately, at higher sampling frequencies. Fleming

et al. (2003) report decay parameters of 0.031 and 0.064 for daily returns and five-minute returns, respectively, for the three liquid futures contracts they consider. The lower decay parameters at these frequencies obtained here for the 78 S&P 100 stocks are likely to be caused by having relatively more noise in the intra-day returns data and a well-known phenomenon in multivariate GARCH models (for daily returns) that the larger the number of assets, the lower the decay parameter, see Engle and Sheppard (2001) and Hafner and Franses (2003) for discussion.

- insert Table 2 about here -

The two time-scales estimator and the lead-lag correction reduce the bias but increase the variance of the realized covariance matrix for a particular sampling frequency. It appears that for both methods the latter is more important here, given that the decay parameters are lower for the corrected covariance matrices compared to the standard case. Note that the log-likelihood is improved, however, except when using the lead-lag correction for the lowest sampling frequencies. The decay parameters in Panel B, considering excess returns, are in general slightly higher in all instances, but otherwise the findings correspond to those for the total returns.

4.2 Portfolio performance

Table 3 shows the performance of the overall minimum variance portfolio, with weights defined in (6), and the minimum variance portfolio given an annualized target return of 10%,¹² with weights given by (8). For the overall minimum variance portfolio the optimal sampling frequency turns out to be 65 minutes in the standard case. The annualized standard deviation of 12.16% compares favorably to the 14.00% for daily data. For the popular five-minute frequency the standard deviation is 12.68%, clearly above the minimum. Also for the target return portfolios the 65-minute frequency is optimal, resulting in a Sharpe ratio of 0.786 compared to 0.596 for daily returns and 0.626 for five-minute returns. In terms of the performance fees (Δ_γ), an investor with low relative risk aversion ($\gamma = 1$) would be willing to pay 155 basis points per year to switch from the covariance matrix estimate based on

¹²We examined the sensitivity of our results to the target return level by varying μ_P between 2% and 18%. These alternative target return levels led to qualitatively similar conclusions as those reported below. Detailed results are therefore not shown here, but are available on request.

daily returns to the realized covariance matrix obtained with 65-minute returns. An investor with high relative risk aversion ($\gamma = 10$) would be willing to pay even 399 basis points.

- insert Table 3 about here -

The results in Panel B of Table 3, using the two time-scales estimator, show only a marginal improvement for the overall minimum variance portfolio with a standard deviation of 12.10% compared to 12.16% before, both at the 65-minute sampling frequency. The same conclusion holds for all other frequencies except 15 minutes. For the target return portfolios, however, the results are ambiguous, in the sense that for sampling frequencies of 10 minutes and higher the two time-scales estimator leads to a higher Sharpe ratio but for lower sampling frequencies portfolio performance worsens. At the optimal frequency of 130 minutes the Sharpe ratio is lower at 0.711, compared to 0.786 for the 65-minute frequency in the standard case. The performance fees Δ_γ show the same pattern.

The lead-lag correction in (3) leads to a higher optimal sampling frequency of 30 minutes for the minimum variance portfolio. The 12.01% annualized standard deviation is slightly better than the 12.16% and 12.10% at the optimal 65-minute frequency in the standard and two time-scales cases, respectively. In fact, the lead-lag bias-correction leads to a reduction in volatility of the minimum variance portfolio at all frequencies, such that the 10-minute sampling frequency now leads to approximately the same level of volatility as the optimal 65-minute frequency in the standard case. Hence, using the lead-lag correction allows for a substantially higher sampling frequency before the increased noise level due to the use of leads and lags offsets this advantage. For the target return portfolios, the optimal sampling frequency remains at 65 minutes as in the standard case, although the corresponding Sharpe ratio is somewhat higher (0.797 compared to 0.786). The same applies to the performance fees Δ_γ , which increase to 168 and 412 basis points per year for low and high relative risk aversion, respectively (compared to 155 and 399).

The performance of the minimum tracking error portfolios is shown in Table 4. Using the standard realized covariance matrix, the tracking error is minimized at 4.43% using the 30-minute frequency compared to 4.75% for daily data and 4.92% at the popular five-minute frequency. The two time-scales estimator provides a further

improvement with the minimum tracking error equal to 4.18% at the 65-minute frequency. Finally, using one lead and lag results in a higher optimal sampling frequency of 15 minutes as for the minimum variance portfolio, with a marginally lower tracking error at 4.35%. Hence here we do observe that bias-correction further improves the performance.

- insert Table 4 about here -

Table 4 also demonstrates that for the active portfolio manager with an annualized target excess return of five percent the optimal sampling frequency is much higher than for total returns. The ex-post information ratio (excess return divided by tracking error) is optimal for the two-minute frequency in the standard case at 0.436 compared to an information ratio of 0.110 at the daily frequency. Risk-averse investors would be willing to pay between 151 and 181 basis points per year to make use of the two-minute frequency realized covariance matrix. The optimal frequency using one lead and one lag is even the one-minute frequency, but it results in a slightly lower information ratio of 0.406 and slightly lower performance fees of 136 and 161 basis points. The two time-scales estimator also results in an optimal frequency of two minutes but with an information ratio of 0.414 and performance fees of 138 and 172 basis points, below the optimum in the standard case. Comparing the information ratios at other frequencies with the corresponding results in the standard case, again we find ambiguous results. The information ratio declines for sampling frequencies of 15 minutes and higher, but it increases for lower sampling frequencies, while the same is observed for the performance fees Δ_γ . Note that this pattern is the complete opposite of that found for the target return portfolios in Panel B of Table 3.

In sum, the general conclusion from Tables 3 and 4 when computing the minimum variance portfolio or minimum tracking error portfolio is that the two time-scales estimator and the lead-lag bias correction marginally improve the out-of-sample performance. We emphasize, however, that selecting the appropriate sampling frequency appears to be much more important than choosing between different bias- and variance-reduction techniques for the realized covariance matrices. For example, the reduction in volatility of the minimum variance portfolio when going from the popular five-minute frequency to the optimal 65-minute frequency in the standard

case (from 12.68% to 12.16%) is more than three times as large as the additional reduction achieved by applying the lead-lag bias correction at the 30-minute frequency (which further reduces volatility to 12.01%).

In general we would like to express a warning note on the target return results in Tables 3 and 4. The actual return pattern at the various frequencies is anything but smooth and hence subject to a certain degree of ‘luck’. Obviously these results depend both on the quality of the expected (excess) returns and the covariance matrix forecasts, making a direct comparison of the quality of the covariance forecasts more difficult than is the case for the minimum variance and minimum tracking error portfolios.

4.3 Transaction costs and rebalancing frequency

Table 3 shows that daily turnover in the volatility timing strategies is considerable, ranging between 12 and 17 percent for most sampling frequencies. In case the ‘standard’ realized covariance matrix or the lead-lag correction is used at the one-minute sampling frequency, turnover increases even to around 25%. For the tracking error portfolios in Table 4, turnover is below 10% for sampling frequencies below five minutes, but rapidly increases when returns are sampled more frequently. Given that the optimal sampling frequency for the target excess return portfolios was found to be two minutes, transaction costs may be substantial and should be taken into account when assessing the portfolio performance.

Panel A of Table 5 shows the performance of the target return portfolios with daily rebalancing for transaction cost levels (c in the first column) between 2% and 20%, in annualized percentage points as explained before. Results are shown for portfolios based on covariance matrix estimates obtained with daily returns and with intraday returns at the optimal frequency, which turns out to be 65 minutes irrespective of the transaction costs level. No bias-corrections are applied to the realized covariance matrix in this case. As expected, transaction costs reduce the portfolio return while the portfolio variance is largely unaffected, leading to a monotonic decline of the Sharpe ratio as the level of transaction costs increases. Note that the reduction in returns and Sharpe ratio is larger for the portfolios based on covariance matrix estimates obtained with daily returns. Therefore the difference in Sharpe ratios with the portfolios based on high-frequency intraday returns actually

becomes larger, such that the performance fees increase to 245 and 487 basis points for $\gamma = 1$ and 10 in case transaction costs amount to 20%. This is not surprising of course, given that daily turnover for these portfolios equals 16.8 and 13.2%, respectively (see Table 3).

- insert Table 5 about here -

Panel A of Table 6 reveals that transaction costs have more dramatic effects for the target excess return portfolios. Daily turnover for the portfolio based on covariance matrix estimates obtained with the optimal two-minute returns is more than four times as high as for the portfolio based on daily returns, at 22.4% compared to 5.2%. The reduction in the mean active return and the information ratio therefore is much more pronounced for the intraday returns based strategy, such that the performance fee actually becomes negative for transaction costs in excess of 10%. Hence, if transaction cost levels are considerable, it does not pay off to use high-frequency intraday returns to estimate the covariance matrix. Also note that the optimal sampling frequency becomes lower at three minutes for high levels of transaction costs.

- insert Table 6 about here -

Transaction costs may be reduced by rebalancing the portfolio less frequently. The effects on portfolio performance are shown in Panels B and C of Table 5 for the target return portfolios. First note that, as expected, the volatility of the portfolio increases when the portfolio holding period increases, but only slightly. Somewhat surprisingly, the portfolio return increases considerably and comes much closer to the target return of 10% when the rebalancing frequency decreases. This corresponds with the findings of Fleming *et al.* (2003). Turning to the effects of transaction costs, we find that the reduction in returns and Sharpe ratio is indeed much less pronounced when rebalancing the portfolio weekly or monthly rather than daily. Again, turnover is higher for the portfolios based on covariance matrix estimates obtained with daily returns, such that the maximum fee investors are willing to pay to switch to covariance matrix estimates obtained with intraday returns increases with the level of transaction costs. Also note that the magnitude of the performance fee Δ_γ declines when the rebalancing frequency becomes lower. This is due to the fact

that the improvement in performance when going from daily to weekly or monthly rebalancing is relatively larger for the portfolio based on daily returns.

In order to assess the economic value of rebalancing less frequently more directly, we compute the performance fee Δ_γ that an investor is willing to pay to switch from daily rebalancing to weekly (or monthly) rebalancing for a given level of transaction costs. For the daily and weekly rebalanced portfolios based on intraday returns at the optimal sampling frequencies and annualized transaction costs equal to 10%, we find that Δ_γ is equal to 94 and 47 basis points for $\gamma = 1$ and 10, respectively. These performance fees even increase to 128 and 64 points when comparing the daily and monthly rebalancing frequencies.

Finally, the benefits of rebalancing less frequently become very clear from Panels B and C of Table 6 for the target excess return portfolio. Although the active return is still reduced due to transaction costs in case of weekly or monthly rebalancing, it remains higher for the portfolio based on intraday returns than for the daily returns portfolio even in case of transaction costs up to 20%. Given that the levels of ex-post tracking error do not differ very much, the IR remains higher as well, and investors are willing to pay considerable fees to make use of the high-frequency returns portfolio.

Again we compute the performance fee Δ_γ using portfolios with daily and weekly (or monthly) holding periods for a given level of transaction costs to evaluate the economic gains from rebalancing less frequently directly. With annualized transaction costs equal to 10%, we find that an investor is willing to pay 200 basis points to switch from daily to weekly rebalancing for both low and high relative risk aversion. Comparing the daily and monthly rebalancing frequencies, Δ_γ is equal to 280 and 276 basis points for $\gamma = 1$ and 10, respectively. It would be interesting to include the rebalancing frequency in the portfolio optimization problem. Obviously this is difficult to achieve and beyond the scope of this paper.

4.4 Genuine out-of-sample forecasting

Fleming *et al.* (2001, 2003) suggest that determining the decay parameters α and α_h in (12) and (14), respectively, using maximum likelihood on the full sample does not lead to serious data snooping problems because the final evaluation criterion (maximizing return or minimizing risk) differs from the likelihood objective function. To

test the validity of this argument, and to test a true out-of-sample strategy, we proceed as follows. First we find the decay parameters that maximize the performance of the various portfolios over the first 250 days following the initial burn-in period, i.e. the values of α and α_h that minimize the (relative) variance or maximizes the Sharpe (or information) ratio. These decay parameters are then used to estimate the conditional covariance matrices $\widehat{\Sigma}_t$ and $\widehat{\Sigma}_{t,h}$ for the first day following the in-sample period, for which optimal portfolio weights are then constructed using (6) and (8). This procedure is repeated using an expanding in-sample estimation window where each time a new observation is added. This not only implies that the decay parameter varies over time, but also that the portfolio performance thus obtained is truly out-of-sample. Since we lose an additional 250 days at the start of the sample, for comparison we re-estimated the decay parameter using maximum likelihood for the shorter sample of 1416 trading days and constructed the corresponding portfolio weights and performance.

- insert Table 7 about here -

The results are presented in Table 7. For both the minimum variance and minimum tracking error portfolios the results are re-assuring, in the sense that the optimal sampling frequency is still 65 and 30 minutes, respectively. Also the performance itself is similar to that of the standard case. By contrast, for the target return portfolios the results do change considerably. In the total return case the optimal sampling frequency is now 10 minutes instead of 65, and the Sharpe ratio has deteriorated from 0.640 to 0.554. In the excess return case the optimal sampling frequency is now one minute instead of two, but with a better information ratio at 0.457 versus 0.373. Perhaps most revealing, the optimal decay parameters are much lower when determined using in-sample portfolio performance than when estimated with maximum likelihood (except for the target return portfolios, when performance is measured by the Sharpe ratio). This holds especially for the higher sampling frequencies. To verify that this is not an artefact of using different decay rates over time, we also did a datasnooping exercise with a constant decay parameter equal to the value that maximizes performance (rather than the log-likelihood) over the entire out-of-sample period. These results (not reported here) confirm that performance-based decay rates are much lower than the ones based on the log-

likelihood. In addition, this enhances the performance at those frequencies. Hence the log-likelihood procedure tends to give too much weight to the updates. A logical explanation for this is that the noise pattern of the updates suits the log-likelihood when standardizing equally noise daily returns, but more smoothing is needed (lower decay parameters) for forecasting the covariance matrices.

5 Conclusion

Existing studies that use high-frequency intra-day data to measure and forecast the daily covariance matrix make ad-hoc choices with regard to the sampling frequency. The presence of bid-ask bounce and non-synchronous trading creates a trade-off between higher sampling frequencies leading to lower variances of the (co-)variance measures due to having more data, and lower sampling frequencies reducing the impact of these market microstructure effects. Popular ad-hoc choices to strike a balance between the resulting bias and variance of the realized covariance estimates are the five- and 30-minute sampling frequencies.

In this study we show that choosing the optimal sampling frequency is crucial for the out-of-sample performance of portfolios constructed using realized covariances. Even for the relatively liquid stocks that comprise the S&P 100 index the optimum is more likely to be in the neighbourhood of an hour rather than five or 30 minutes.

We also investigated the use of bias- and variance-reduction methods for computing the realized covariances. Both the two time-scales estimator and the lead-lag bias-correction procedure result in a marginal improvement over the standard realized covariance matrix estimator at the same frequency. Transaction costs were shown to affect portfolio performance considerably, and in particular they imply that rebalancing the portfolio less frequently may be beneficial.

Several interesting topics for further research come to mind. First, it would be interesting to explore other ways to correct for biases in realized covariances due to non-synchronous trading. Second, it may be worthwhile to allow the sampling frequency to vary over time, to take into account changes in trading intensity. The S&P stocks considered here, for example, were traded much more frequently at the end of the sample period than in the beginning, see also Bandi *et al.* (2005). Third, Andersen *et al.* (2003) suggest that with more and more assets eventually

a factor model will be needed, see Andersen, Bollerslev, Diebold and Ebens (2001) and Hafner *et al.* (2005) for additional motivation and discussion and Bollerslev and Zhang (2003) for an application using the Fama-French three-factor model. Fourth, it would be interesting to examine the effects of restrictions on the portfolio weights, which we did not consider here. As shown by Jagannathan and Ma (2003), imposing short-selling constraints and a maximum weight constraint, for example, may enhance portfolio performance, even if the restrictions are wrong. Finally, in this study we considered the popular approach that makes use of artificially constructed equidistant prices in calendar time, in part because the empirical data set was constrained to the close of each minute rather than all transaction prices. It would be interesting to see empirical work on the scale of this paper with many stocks that considers transaction time sampling rather than calendar time sampling, for example using the covariance estimators of Harris *et al.* (1995), De Jong and Nijman (1997) and Hayashi and Yoshida (2005). Martens (2004) provides an overview and comparison in a simulation setting, while Hansen *et al.* (2005) discuss theoretical issues related to such estimators including bias correction procedures.

References

- Aït-Sahalia, Y., P. Mykland and L. Zhang (2005), How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise, *Review of Financial Studies* **18**, 351–416.
- Andersen, T.G. and T. Bollerslev (1998), Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts, *International Economic Review* **39**, 885–905.
- Andersen, T.G., T. Bollerslev, F.X. Diebold, and H. Ebens (2001), The Distribution of Realized Stock Return Volatility, *Journal of Financial Economics* **61**, 43–76.
- Andersen T.G., T. Bollerslev, F.X. Diebold and P. Labys (2001), The Distribution of Realized Exchange Rate Volatility, *Journal of the American Statistical Association* **96**, 42–55.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2003), Modeling and Forecasting Realized Volatility, *Econometrica* **71**, 579–625.
- Andreou E. and E. Ghysels (2002), Rolling Sample Volatility Estimators: Some New Theoretical, Simulation and Empirical Results, *Journal of Business and Economic Statistics* **20**, 363–376.
- Andrews, D.W.K. (1991), Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation, *Econometrica* **59**, 817–858.
- Areal, N.M.P.C. and S.J. Taylor (2002), The Realized Volatility of FTSE-100 Futures Prices, *Journal of Futures Markets* **22**, 627–48.
- Bandi, F.M. and J.R. Russell (2005a), Microstructure Noise, Realized Variance, and Optimal Sampling, Graduate School of Business, University of Chicago, working paper.
- Bandi, F.M. and J.R. Russell (2005b), Separating Microstructure Noise from Volatility, *Journal of Financial Economics*, to appear.
- Bandi, F.M., J.R. Russell and Y. Zhu (2005), Using High-Frequency Data in Dynamic Portfolio Choice, *Econometric Reviews*, this issue.
- Barndorff-Nielsen, O.E. and N. Shephard (2004), Econometric Analysis of Realised Covariation: High Frequency Covariance, Regression and Correlation in Financial Economics, *Econometrica* **72**, 885–925.
- Barndorff-Nielsen, O.E., P.R. Hansen, A. Lunde and N. Shephard (2004), Regular and Modified Kernel-Based Estimators of Integrated Variance: The Case of Independent Noise, Oxford Financial Research Centre Economics Series 2004-FE-20.
- Bollerslev, T. and B.Y.B. Zhang (2003), Measuring and Modeling Systematic Risk in Factor Pricing Models Using High-Frequency Data, *Journal of Empirical Finance* **10**, 533–558.
- Chan, L.K.C., J. Karceski, and J. Lakonishok (1999), On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model, *Review of Financial Studies* **12**, 937–674.
- Cohen, K.J., G.A. Hawawini, S.F. Maier, R.A. Schwartz, and D.K. Whitcomb (1983), Friction in the Trading Process and the Estimation of Systematic Risk, *Journal of Financial Economics* **12**, 263–278.

- De Jong, F. and T. Nijman (1997), High Frequency Analysis of Lead-Lag Relationships between Financial Markets, *Journal of Empirical Finance* **4**, 259–277.
- Engle, R.F. and K. Sheppard (2001), Theoretical and Empirical Properties of Dynamic Conditional Correlation MVGARCH, UCSD Working Paper No. 2001-15.
- Fleming, J., C. Kirby and B. Ostdiek (2001), The Economic Value of Volatility Timing, *Journal of Finance* **56**, 329–352.
- Fleming, J., C. Kirby and B. Ostdiek (2003), The Economic Value of Volatility Timing Using “Realized” Volatility, *Journal of Financial Economics* **67**, 473–509.
- Foster, D.P. and D.B. Nelson (1996), Continuous Record Asymptotics for Rolling Sample Variance Estimators, *Econometrica* **64**, 139–174.
- Hafner, C.M. and P.H. Franses (2003), A Generalized Dynamic Conditional Correlation Model for Many Assets, Econometric Institute Report No. 2003-18.
- Hafner, C.M., D. van Dijk and P.H. Franses (2005), Semiparametric Modelling of Correlation Dynamics, in: T. Fomby, C. Hill and D. Terrell (eds.), *Advances in Econometrics Vol. 20/A - Econometric Analysis of Financial and Economic Time Series*, Amsterdam: Elsevier, to appear.
- Hansen, P.R. and A. Lunde (2005), A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data, *Journal of Financial Econometrics* **3**, 525–554.
- Hansen, P.R. and A. Lunde (2006), Realized Variance and Market Microstructure Noise, *Journal of Business and Economic Statistics*, to appear.
- Hansen, P.R., A. Lunde and V. Voev (2005), Realized Covariance Estimation Using High-Frequency Data in the Presence of Noise, Stanford University, working paper.
- Harris, F.H. deB., T.H. McInish, G.L. Shoemith and R.A. Wood (1995), Cointegration, Error Correction, and Price Discovery on Informationally Linked Security Markets, *Journal of Financial and Quantitative Analysis* **30**, 563–579.
- Hayashi, T. and N. Yoshida (2005), On Covariance Estimation of Non-Synchronously Observed Diffusion Processes, *Bernoulli* **11**, 359–379.
- Jagannathan, R. and T. Ma (2003), Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps, *Journal of Finance* **58**, 1651–1683.
- Jegadeesh, N., and S. Titman, 1993, Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency, *Journal of Finance* **48**, 65–91.
- Koopman, S.J., B. Jungbacker and E. Hol (2005), Forecasting Daily Variability of the S&P 100 Stock Index Using Historical, Realised and Implied Volatility Measurements, *Journal of Empirical Finance* **12**, 445–475.
- Liu, Q. (2004), On portfolio optimization: How do we Benefit from High-Frequency Data, working paper, University of Hawaii.
- Martens, M. (2002), Measuring and Forecasting S&P 500 Index-Futures Volatility Using High-Frequency Data, *Journal of Futures Markets* **22**, 497–518.
- Martens, M. (2004), Estimating Unbiased and Precise Realized Covariances, working paper, Erasmus University Rotterdam.

- Martens, M., D. van Dijk and M. de Pooter (2004), Modeling and Forecasting S&P 500 Volatility: Long Memory, Structural Breaks and Nonlinearity, Tinbergen Institute Discussion Paper No. 067/4.
- Newey, W.K. and K.D. West (1987), A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* **55**, 703–708.
- Pong, S., M.B. Shackleton, S.J. Taylor, and X. Xu (2004), Forecasting Currency Volatility: A Comparison of Implied Volatilities and AR(FI)MA Models, *Journal of Banking and Finance* **28**, 2541–2563.
- Rouwenhorst, K.G., 1998, International Momentum Strategies, *Journal of Finance* **53**, 267–284.
- Scholes, M., and J. Williams (1977), Estimating Betas from Nonsynchronous Data, *Journal of Financial Economics* **5**, 309–327.
- Thomakos, D.D. and T. Wang (2003), Realized Volatility in the Futures Markets, *Journal of Empirical Finance* **10**, 321–353.
- Zhang, L. (2005), Estimating Covariation: Epps Effect and Microstructure Noise, University of Illinois at Chicago, working paper.
- Zhang, L., P. Mykland and Y. Aït-Sahalia (2005), A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data, *Journal of the American Statistical Association*, to appear.
- Zhou, B. (1996), High-frequency data and volatility in foreign-exchange rates, *Journal of Business and Economic Statistics* **14**, 45–52.

Appendix: S&P 100 constituents on June 18, 2004

The 100 constituents of the S&P 100 index on June 18, 2004. The 78 stocks marked with a * are included in the analysis. For these stocks there is a complete set of one-minute open-high-low-close prices from April 16, 1997, through May 27, 2004 (1788 trading days).

Symbol	Issue name	Symbol	Issue name
AA*	ALCOA INC	IBM*	INTL BUS MACHINE
AEP*	AMER ELEC PWR	INTC*	INTEL CORP
AES*	THE AES CORP	IP	INTL PAPER CO
AIG*	AMER INTL GROUP	JNJ*	JOHNSON&JOHNSON
ALL*	ALLSTATE CP	JPM*	JP MORGAN CHASE
AMGN*	AMGEN	KO*	COCA COLA CO
AOL	AOL TIME WARNER	LEH*	LEHMAN BROS
ATI	ALLEGHENY TECH	LTD*	LIMITED BRANDS
AVP	AVON PRODS INC	LU*	LUCENT TECH
AXP*	AMER EXPRESS CO	MAY*	MAY DEPT STORES
BA*	BOEING CO	MCD*	MCDONALDS CORP
BAC*	BANK OF AMERICA	MDT*	MEDTRONIC INC
BAX*	BAXTER INTL INC	MEDI	MEDIMMUNE INC
BCC*	BOISE CASCADE	MER*	MERRILL LYNCH
BDK*	BLACK & DECKER	MMM*	3M COMPANY
BHI*	BAKER HUGHES INC	MO*	ALTRIA GROUP
BMY*	BRISTOL MYERS SQ	MRK*	MERCK & CO
BNI*	BURL NTHN SANTA	MSFT*	MICROSOFT CP
BUD*	ANHEUSER BUSCH	MWD	MORGAN STANLEY
C*	CITIGROUP	NSC*	NORFOLK SOUTHERN
CCU*	CLEAR CHANNEL	NSM*	NATL SEMICONDUCT
CI*	CIGNA CORP	NXTL*	NEXTEL COMMS
CL*	COLGATE PALMOLIV	ONE*	BANK ONE CORP
CPB*	CAMPBELL SOUP CO	ORCL*	ORACLE CORP
CSC	COMPUTER SCIENCE	PEP*	PEPSICO INC
CSCO*	CISCO SYSTEMS	PFE*	PFIZER INC
DAL*	DELTA AIR LINES	PG	PROCTER & GAMBLE
DD*	DU PONT CO	ROK*	ROCKWELL AUTOMAT
DIS*	WALT DISNEY CO	RSH	RADIOSHACK
DOW	DOW CHEMICAL CO	RTN	RAYTHEON CO
EK*	EASTMAN KODAK	S*	SEARS ROEBUCK
EMC*	EMC CORP	SBC*	SBC COMMS
EP	EL PASO CORP	SLB*	SCHLUMBERGER LTD
ETR*	ENTERGY CP	SLE*	SARA LEE CORP
EXC	EXELON CORP	SO*	SOUTHERN CO
F	FORD MOTOR CO	T*	AT&T CORP
FDX	FEDEX CORP	TOY*	TOYS R US CORP
G*	GILLETTE CO	TXN*	TEXAS INSTRUMENT
GD*	GENERAL DYNAMICS	TYC*	TYCO INTL
GE*	GENERAL ELEC CO	UIS*	UNISYS CORP
GM*	GENERAL MOTORS	USB	US BANCORP
GS	GOLDM SACHS GRP	UTX*	UNITED TECH CP
HAL*	HALLIBURTON CO	VIAb	VIACOM CL B
HCA	HCA INC	VZ	VERIZON COMMS
HD*	HOME DEPOT INC	WFC*	WELLS FARGO & CO
HET*	HARRAHS ENTER	WMB*	WILLIAMS COS INC
HIG*	HARTFORD FINL	WMT*	WAL-MART STORES
HNZ*	H J HEINZ CO	WY	WEYERHAEUSER CO
HON*	HONEYWELL INTL	XOM	EXXON MOBIL
HPQ*	HEWLETT-PACKARD	XRX*	XEROX CORP

Table 1: Mean and variance of the realized (co-)variance

Frequency	Realized Variance		Realized Covariance	
	Mean	Variance	Mean	Variance
Daily	7.386	1763	1.568	93.58
<u>Panel A: Standard</u>				
130 minutes	7.369	689.7	1.394	31.94
65 minutes	7.324	624.3	1.357	22.46
30 minutes	7.311	563.5	1.316	16.48
15 minutes	7.463	545.8	1.307	14.46
10 minutes	7.614	547.2	1.305	13.25
5 minutes	7.912	531.7	1.239	11.49
3 minutes	8.193	527.6	1.136	10.34
2 minutes	8.525	537.1	1.025	9.60
1 minute	9.494	597.0	0.826	8.73
<u>Panel B: Two time-scales</u>				
130 minutes	7.836	1090.6	1.462	37.05
65 minutes	7.295	760.9	1.418	22.94
30 minutes	7.150	606.4	1.414	16.82
15 minutes	7.233	570.1	1.420	14.97
10 minutes	7.315	562.8	1.407	13.90
5 minutes	7.446	554.9	1.361	12.24
3 minutes	7.500	534.5	1.297	11.32
2 minutes	7.524	521.4	1.231	10.80
<u>Panel C: 1 lead and 1 lag</u>				
130 minutes	7.422	758.6	1.418	36.73
65 minutes	7.368	667.0	1.389	26.27
30 minutes	7.329	595.6	1.351	18.97
15 minutes	7.342	552.7	1.332	15.51
10 minutes	7.420	545.6	1.334	14.23
5 minutes	7.611	532.3	1.300	12.67
3 minutes	7.812	538.3	1.257	11.56
2 minutes	8.020	536.3	1.193	10.70
1 minute	8.525	533.4	1.025	9.47

Notes: The table shows mean and variance of the realized (co-)variances at various sampling frequencies for 78 constituents of the S&P100 index from April 16, 1997, through May 27, 2004 (1788 trading days). For the realized variance the mean reflects the average taken over all 78 stocks and over all 1788 trading days. The variance is the average taken over the 78 sample variances of the realized variances. For the realized covariance the mean reflects the average taken over all 3003 pairs of stocks and over all 1788 trading days. The variance is the average taken over the 3003 sample variances of the realized covariances. In Panel A the “standard” realized covariance matrix $V_{t-1,h}$ given in (1) is used. Panel B is based on the two time-scales estimator $V_{t-1,h}^{TTS}$ given in (2), while Panel C shows results for the lead-lag corrected estimator $V_{t-1,h}^{LL}$ given in (3), with Bartlett-kernel weights $d_l = 1 - l/(q + 1)$ and $q = 1$.

Table 2: Optimal decay parameters

Frequency	Standard		Two time-scales		1 lead, 1 lag	
	α	LogL	α	LogL	α	LogL
<u>Panel A: Total Returns</u>						
Daily	0.0070	-300,492	0.0070	-300,492	0.0070	-300,492
130 minutes	0.0119	-276,376	0.0129	-275,476	0.0111	-276,939
65 minutes	0.0149	-274,580	0.0156	-273,578	0.0137	-274,782
30 minutes	0.0204	-273,747	0.0200	-272,644	0.0179	-273,519
15 minutes	0.0273	-273,802	0.0256	-272,707	0.0231	-273,186
10 minutes	0.0329	-274,121	0.0293	-273,071	0.0273	-273,261
5 minutes	0.0481	-275,004	0.0356	-274,024	0.0375	-273,774
3 minutes	0.0678	-275,975	0.0386	-274,939	0.0493	-274,407
2 minutes	0.1025	-276,985	0.0385	-275,729	0.0643	-275,164
1 minute	0.2106	-278,971			0.1255	-276,723
<u>Panel B: Excess Returns</u>						
Daily	0.0070	-298,480	0.0070	-298,480	0.0070	-298,480
130 minutes	0.0119	-274,455	0.0130	-273,449	0.0112	-275,033
65 minutes	0.0151	-272,712	0.0158	-271,507	0.0138	-272,907
30 minutes	0.0208	-271,926	0.0204	-270,572	0.0183	-271,669
15 minutes	0.0282	-272,100	0.0263	-270,702	0.0238	-271,381
10 minutes	0.0342	-272,516	0.0304	-271,133	0.0282	-271,522
5 minutes	0.0514	-273,562	0.0373	-272,238	0.0393	-272,157
3 minutes	0.0757	-274,663	0.0406	-273,293	0.0529	-272,925
2 minutes	0.1178	-275,675	0.0402	-274,192	0.0713	-273,751
1 minute	0.2468	-277,537			0.1440	-275,393

Notes: The table shows the decay rates (α) that maximize the likelihood of the model in (13) and (12) for daily data and (13) and (14) for intraday data. In Panel A the model is estimated for total returns, whereas in Panel B the model is estimated for excess returns (stock returns minus S&P500 returns). The second and third column show the optimal decay rates and accompanying log-likelihood values when the covariance updates are based on the standard realized (co-)variances, the fourth and fifth column when the updates are based on the two time-scales estimator, and the final two columns when 1 lead and 1 lag of the (co-)variances are added to the contemporaneous (realized) covariances.

Table 3: Out-of-sample performance - total returns

Frequency	Target return portfolio						Min. variance portfolio	
	μ_P	σ_P	SR	Δ_1	Δ_{10}	TO	σ_{MVP}	TO
Daily	8.84	14.83	0.596			16.8	14.00	16.4
<u>Panel A: Standard</u>								
130 minutes	10.24	13.13	0.780	163.3	376.9	13.5	12.46	12.9
65 minutes	10.12	12.87	0.786	154.5	399.1	13.2	12.16	12.5
30 minutes	8.16	12.99	0.628	-42.8	187.4	13.4	12.17	12.7
15 minutes	8.82	13.12	0.673	22.0	237.3	13.8	12.21	13.0
10 minutes	8.29	13.30	0.623	-33.9	160.1	14.2	12.38	13.4
5 minutes	8.56	13.69	0.626	-11.5	135.4	15.6	12.68	14.7
3 minutes	8.48	13.92	0.610	-22.7	95.6	17.4	12.84	16.5
2 minutes	7.67	14.18	0.541	-107.6	-22.3	21.1	13.06	20.0
1 minute	7.89	14.44	0.546	-89.6	-38.3	28.6	13.33	26.9
<u>Panel B: Two time-scales</u>								
130 minutes	9.27	13.04	0.711	67.8	292.6	13.2	12.26	12.7
65 minutes	8.82	12.88	0.685	25.0	268.3	12.4	12.10	11.9
30 minutes	8.16	12.88	0.633	-41.4	201.8	12.1	12.13	11.6
15 minutes	8.42	13.04	0.646	-17.3	208.0	12.4	12.24	11.8
10 minutes	8.78	13.17	0.667	17.3	226.2	12.6	12.33	12.0
5 minutes	9.18	13.49	0.680	52.1	222.5	12.9	12.57	12.2
3 minutes	9.14	13.73	0.665	44.9	186.2	12.8	12.73	12.0
2 minutes	9.21	13.91	0.662	49.8	168.3	12.4	12.88	11.6
1 minute								
<u>Panel C: 1 lead and 1 lag</u>								
130 minutes	10.23	13.11	0.780	162.8	378.8	13.6	12.44	13.0
65 minutes	10.25	12.87	0.797	168.0	412.9	13.3	12.16	12.6
30 minutes	9.07	12.81	0.708	50.8	302.7	13.2	12.01	12.6
15 minutes	8.41	12.89	0.653	-16.3	226.3	13.4	12.05	12.7
10 minutes	7.81	13.01	0.600	-78.2	149.6	13.7	12.15	12.9
5 minutes	8.54	13.28	0.643	-8.9	187.0	14.4	12.35	13.6
3 minutes	8.50	13.55	0.627	-15.9	147.6	15.6	12.57	14.8
2 minutes	8.05	13.80	0.583	-64.6	68.6	17.1	12.75	16.2
1 minute	8.06	14.10	0.571	-68.2	27.0	23.2	12.99	22.0

Notes: The table shows the out-of-sample performance of the overall minimum variance portfolio, with weights given in (6), and the minimum variance portfolio given a target level of return of 10%, with weights given in (8), constructed using rolling covariance matrix forecasts based on various sampling frequencies and based on different ways of measuring the realized covariance matrix (standard, two time-scales, and 1 lead and 1 lag). For the target return portfolios, we report the mean return (μ_P) and standard deviation (σ_P) in annualized percentage points, the Sharpe ratio (SR), the annualized basis points fee (Δ_γ) an investor with quadratic utility and constant relative risk aversion of γ would pay to switch from the daily returns covariance matrix estimate to the intraday returns of the optimal portfolios, and average daily turnover (TO) in percentage points. For the minimum variance portfolios, we report the standard deviation (σ_{MVP}) in annualized percentage points average daily turnover (TO) in percentage points.

Table 4: Out-of-sample performance - excess returns

Frequency	Target active return portfolio						Min. TE portfolio	
	μ_P	TE_P	IR	Δ_1	Δ_{10}	TO	TE_{MTE}	TO
Daily	0.52	4.77	0.110			5.2	4.75	5.2
<u>Panel A: Standard</u>								
130 minutes	0.45	4.55	0.099	-6.4	3.0	6.2	4.53	6.2
65 minutes	0.28	4.51	0.063	-22.7	-11.9	6.7	4.48	6.7
30 minutes	0.09	4.48	0.019	-42.4	-30.1	7.6	4.43	7.5
15 minutes	0.57	4.52	0.127	6.2	16.9	8.8	4.46	8.8
10 minutes	0.75	4.67	0.161	23.2	27.8	9.8	4.59	9.8
5 minutes	0.89	4.99	0.179	35.7	26.3	12.6	4.92	12.6
3 minutes	1.85	5.23	0.353	129.9	109.3	16.6	5.18	16.6
2 minutes	2.36	5.42	0.436	180.5	150.8	22.4	5.40	22.5
1 minute	1.39	5.79	0.241	81.7	33.3	36.2	5.78	36.3
<u>Panel B: Two time-scales</u>								
130 minutes	0.54	4.26	0.126	3.7	24.4	4.4	4.24	4.4
65 minutes	0.44	4.21	0.105	-5.7	17.1	4.4	4.18	4.3
30 minutes	0.49	4.22	0.116	-0.9	21.5	4.7	4.19	4.6
15 minutes	0.20	4.32	0.046	-30.4	-11.8	5.5	4.28	5.4
10 minutes	0.56	4.43	0.126	5.0	19.1	6.2	4.38	6.1
5 minutes	1.13	4.78	0.236	60.5	60.1	7.9	4.72	7.9
3 minutes	1.81	5.16	0.351	126.9	109.6	9.5	5.11	9.5
2 minutes	2.28	5.51	0.414	172.1	138.0	11.1	5.45	11.1
1 minute								
<u>Panel C: 1 lead and 1 lag</u>								
130 minutes	0.18	4.58	0.038	-34.0	-25.9	6.2	4.56	6.1
65 minutes	0.32	4.50	0.071	-19.0	-7.8	6.3	4.48	6.3
30 minutes	0.41	4.44	0.092	-9.9	3.9	6.8	4.40	6.8
15 minutes	0.62	4.40	0.141	11.4	26.8	7.5	4.35	7.5
10 minutes	0.63	4.48	0.140	11.7	24.1	8.1	4.42	8.0
5 minutes	0.72	4.58	0.156	20.1	28.1	9.4	4.51	9.4
3 minutes	1.57	4.77	0.329	104.7	104.7	11.4	4.71	11.3
2 minutes	1.95	4.97	0.394	142.2	133.7	13.7	4.91	13.7
1 minute	2.16	5.33	0.406	161.3	135.9	22.2	5.31	22.3

Notes: The table shows the out-of-sample performance of the overall minimum tracking error portfolio, with weights given in (6), and the minimum tracking error portfolio given a target level of return of 5%, with weights given in (8), constructed using rolling covariance matrix forecasts based on various sampling frequencies and based on different ways of measuring the realized covariance matrix (standard, two time-scales, and 1 lead and 1 lag). For the target active return portfolios, we report the mean active return (μ_P) and tracking error (TE_P) in annualized percentage points, the information ratio (IR), the annualized basis points fee (Δ_γ) an investor with quadratic utility and constant relative risk aversion of γ would pay to switch from the daily returns covariance matrix estimate to the intraday returns of the optimal portfolios, and average daily turnover (TO) in percentage points. For the minimum tracking error portfolios, we report the tracking error (TE_{MTE}) in annualized percentage points average daily turnover (TO) in percentage points.

Table 5: Transaction costs and rebalancing frequency - total returns

c	Daily returns			Intraday returns					
	μ_P	σ_P	SR	μ_P	σ_P	SR	h	Δ_1	Δ_{10}
Panel A: Daily rebalancing									
0	8.84	14.83	0.596	10.12	12.87	0.786	65	154.5	399.1
2	8.49	14.83	0.572	9.84	12.87	0.765	65	162.7	407.0
4	8.12	14.83	0.548	9.56	12.87	0.743	65	171.0	415.1
6	7.75	14.82	0.523	9.27	12.87	0.720	65	179.6	423.5
8	7.37	14.82	0.497	8.98	12.87	0.698	65	188.3	431.9
10	6.98	14.82	0.471	8.68	12.87	0.674	65	197.2	440.6
12	6.58	14.82	0.444	8.37	12.87	0.651	65	206.3	449.5
14	6.18	14.82	0.417	8.06	12.87	0.626	65	215.6	458.6
16	5.76	14.82	0.389	7.74	12.87	0.601	65	225.1	467.9
18	5.33	14.82	0.360	7.41	12.87	0.576	65	234.9	477.4
20	4.90	14.81	0.330	7.08	12.87	0.550	65	244.9	487.2
Panel B: Weekly rebalancing									
0	9.62	14.88	0.646	10.36	13.17	0.787	130	98.4	314.8
2	9.43	14.88	0.634	10.22	13.17	0.776	130	103.1	319.4
4	9.25	14.88	0.621	10.08	13.17	0.766	130	107.9	324.2
6	9.05	14.88	0.608	9.94	13.17	0.755	130	112.8	329.0
8	8.86	14.88	0.595	9.80	13.17	0.744	130	117.8	333.9
10	8.66	14.88	0.582	9.65	13.17	0.732	130	122.9	339.0
12	8.45	14.88	0.568	9.49	13.17	0.721	130	128.1	344.1
14	8.24	14.88	0.554	9.34	13.17	0.709	130	133.5	349.4
16	8.03	14.88	0.540	9.18	13.17	0.697	130	138.9	354.8
18	7.81	14.88	0.525	9.01	13.17	0.684	130	144.5	360.3
20	7.58	14.88	0.510	8.85	13.17	0.672	130	150.3	366.0
Panel C: Monthly rebalancing									
0	10.53	14.95	0.704	10.36	13.31	0.779	130	7.0	217.1
2	10.43	14.95	0.698	10.29	13.31	0.774	130	9.3	219.4
4	10.34	14.95	0.691	10.22	13.31	0.768	130	11.6	221.7
6	10.24	14.95	0.685	10.15	13.31	0.763	130	14.0	224.1
8	10.14	14.95	0.678	10.07	13.31	0.757	130	16.4	226.5
10	10.04	14.95	0.671	9.99	13.31	0.751	130	18.9	229.0
12	9.93	14.95	0.664	9.91	13.31	0.745	130	21.5	231.5
14	9.83	14.95	0.657	9.83	13.30	0.739	130	24.1	234.1
16	9.72	14.95	0.650	9.75	13.30	0.733	130	26.7	236.7
18	9.60	14.95	0.642	9.67	13.30	0.727	130	29.5	239.4
20	9.49	14.95	0.635	9.58	13.30	0.720	130	32.3	242.2

Notes: The table shows the out-of-sample performance of the the minimum variance portfolio given a target level of return of 10%, with weights given in (8), constructed using rolling covariance matrix forecasts based on daily returns and on intraday returns at the sampling frequency that maximized the information ratio, based on the ‘standard’ way of measuring the realized covariance matrix. We report the mean return (μ_P) and standard deviation (σ_P) in annualized percentage points, the Sharpe ratio (SR), and the annualized basis points fee (Δ_γ) an investor with quadratic utility and constant relative risk aversion of γ would pay to switch from the daily returns covariance matrix estimate to the intraday returns of the optimal portfolios. The column headed c indicates the level of transaction costs, expressed in annualized percentage points, which correspond with the reduction in the annualized portfolio return if the entire portfolio would have to be traded every day during the whole year. The column headed h indicates the optimal sampling frequency, expressed as the length of the corresponding return interval in minutes.

Table 6: Transaction costs and rebalancing frequency - excess returns

c	Daily returns			Intraday returns					
	μ_P	TE_P	IR	μ_P	TE_P	IR	h	Δ_1	Δ_{10}
Panel A: Daily rebalancing									
0	0.52	4.77	0.110	2.36	5.42	0.436	2	180.5	150.8
2	0.41	4.77	0.085	1.91	5.42	0.352	2	146.6	116.7
4	0.29	4.77	0.061	1.44	5.42	0.266	2	111.9	81.9
6	0.17	4.77	0.035	0.97	5.42	0.178	2	76.5	46.4
8	0.05	4.77	0.010	0.48	5.43	0.089	2	40.3	10.1
10	-0.08	4.76	-0.017	0.08	5.24	0.015	3	13.5	-7.7
12	-0.21	4.76	-0.044	-0.30	5.24	-0.057	3	-11.3	-32.7
14	-0.34	4.76	-0.072	-0.68	5.24	-0.131	3	-36.7	-58.2
16	-0.48	4.76	-0.100	-1.08	5.24	-0.206	3	-62.7	-84.3
18	-0.61	4.76	-0.129	-1.48	5.24	-0.283	3	-89.3	-111.1
20	-0.76	4.76	-0.159	-1.90	5.24	-0.362	3	-116.6	-138.5
Panel B: Weekly rebalancing									
0	0.23	4.81	0.048	2.36	5.29	0.447	2	211.0	189.3
2	0.17	4.81	0.035	2.16	5.29	0.408	2	196.6	174.9
4	0.11	4.81	0.023	1.96	5.29	0.369	2	182.0	160.2
6	0.05	4.81	0.010	1.74	5.29	0.330	2	167.0	145.2
8	-0.01	4.81	-0.002	1.53	5.29	0.289	2	151.7	129.9
10	-0.08	4.81	-0.016	1.30	5.18	0.251	3	135.8	119.5
12	-0.14	4.81	-0.029	1.13	5.18	0.217	3	124.7	108.3
14	-0.21	4.81	-0.043	0.95	5.18	0.183	3	113.3	96.9
16	-0.27	4.81	-0.057	0.76	5.18	0.147	3	101.7	85.3
18	-0.34	4.81	-0.071	0.57	5.18	0.111	3	89.8	73.3
20	-0.42	4.81	-0.086	0.38	5.18	0.073	3	77.5	61.1
Panel C: Monthly rebalancing									
0	0.22	4.88	0.044	3.45	5.28	0.654	1	321.5	303.0
2	0.19	4.88	0.038	3.36	5.28	0.636	1	315.2	296.7
4	0.16	4.88	0.032	3.27	5.28	0.619	1	308.8	290.3
6	0.13	4.88	0.026	3.17	5.28	0.601	1	302.3	283.8
8	0.10	4.87	0.020	3.07	5.28	0.582	1	295.6	277.1
10	0.07	4.87	0.013	2.97	5.28	0.563	1	288.8	270.3
12	0.03	4.87	0.007	2.87	5.28	0.544	1	281.8	263.3
14	0.00	4.87	0.000	2.77	5.28	0.524	1	274.7	256.2
16	-0.03	4.87	-0.007	2.66	5.28	0.504	1	267.4	248.9
18	-0.07	4.87	-0.014	2.55	5.28	0.483	1	259.9	241.4
20	-0.10	4.87	-0.021	2.44	5.28	0.462	1	252.3	233.7

Notes: The table shows the out-of-sample performance of the the minimum tracking error portfolio given a target level of return of 5%, with weights given in (8), constructed using rolling covariance matrix forecasts based on daily returns and on intraday returns at the sampling frequency that maximized the information ratio, based on the 'standard' way of measuring the realized covariance matrix. We report the mean active return (μ_P) and tracking error (TE_P) in annualized percentage points, the information ratio (IR), the annualized basis points fee (Δ_γ) an investor with quadratic utility and constant relative risk aversion of γ would pay to switch from the daily returns covariance matrix estimate to the intraday returns of the optimal portfolios. The column headed c indicates the level of transaction costs, expressed in annualized percentage points, which correspond with the reduction in the annualized portfolio return if the entire portfolio would have to be traded every day during the whole year. The column headed h indicates the optimal sampling frequency, expressed as the length of the corresponding return interval in minutes.

Table 7: Out-of-sample α 's

Frequency	Target return portfolio				Minimum risk portfolio				
	α_h		SR/IR		α_h		σ_P/TE_P		
	Mean	St.Dev	Perf.	LogL	Mean	St.Dev	Perf.	LogL	
<u>Panel A: Total Returns</u>									
Daily	0.001	0.003	0.507	0.482	0.004	0.002	13.770	13.669	
130 minutes	0.069	0.037	0.534	0.626	0.014	0.002	12.534	12.229	
65 minutes	0.087	0.107	0.262	0.640	0.018	0.003	11.937	11.945	
30 minutes	0.035	0.020	0.436	0.519	0.024	0.003	11.986	11.972	
15 minutes	0.097	0.130	0.261	0.603	0.034	0.005	12.061	12.034	
10 minutes	0.352	0.034	0.554	0.561	0.042	0.006	12.220	12.207	
5 minutes	0.312	0.086	0.434	0.597	0.047	0.005	12.471	12.466	
3 minutes	0.223	0.189	0.471	0.578	0.080	0.030	12.592	12.556	
2 minutes	0.215	0.194	0.452	0.526	0.107	0.076	12.797	12.755	
1 minute	0.361	0.116	0.478	0.516	0.124	0.076	13.002	12.971	
<u>Panel B: Excess returns</u>									
Daily	0.041	0.004	0.357	-0.031	0.006	0.001	4.858	4.871	
130 minutes	0.004	0.006	0.167	0.122	0.008	0.003	4.593	4.581	
65 minutes	0.002	0.002	0.298	0.161	0.008	0.004	4.595	4.554	
30 minutes	0.002	0.009	0.181	0.096	0.011	0.005	4.512	4.477	
15 minutes	0.001	0.001	0.376	0.198	0.012	0.006	4.516	4.483	
10 minutes	0.004	0.012	0.203	0.312	0.012	0.006	4.630	4.629	
5 minutes	0.001	0.002	0.395	0.234	0.010	0.005	4.814	4.949	
3 minutes	0.009	0.010	0.239	0.348	0.010	0.005	4.970	5.220	
2 minutes	0.024	0.041	0.269	0.373	0.011	0.005	5.095	5.514	
1 minute	0.031	0.028	0.453	0.184	0.015	0.006	5.249	5.864	

Notes: The table shows the out-of-sample performance of the overall minimum volatility (tracking error) portfolio, with weights given in (6), and the minimum variance portfolio given an annualized target level of (active) return of 10% (5%), with weights given in (8), constructed using rolling covariance matrix forecasts based on various sampling frequencies and based on the ‘standard’ realized covariance matrix. Panel A shows results for total returns and Panel B for excess returns (stock returns minus S&P 500 returns). The optimal decay parameters are determined by optimizing portfolio performance using an expanding window period (starting with 250 days). Columns 2 and 3, and 6 and 7, report the mean and standard deviation of the resulting estimates of α_h . Columns 4 and 8, headed ‘Perf.’, show the Sharpe ratio and volatility (panel A) or the information ratio and tracking error (panel B) for the resulting portfolios. Columns 5 and 9, headed ‘LogL’, show the SR/IR and σ_P/TE_P for portfolios constructed with decay parameters for the conditional covariance matrix that are estimated by maximizing the log-likelihood over the complete out-of-sample period.