



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*

September 21, 2005

Abstract

This paper investigates the merits of high-frequency intraday data when forming minimum variance portfolios and minimum tracking error 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, which strikes a balance between variance and bias in covariance matrix estimates due to market microstructure effects such as non-synchronous trading and bid-ask bounce. The optimal sampling frequency typically ranges between 30- and 65-minutes, considerably lower than the popular five-minute frequency. We also examine how bias-correction procedures, based on the addition of leads and lags and on scaling, and a variance-reduction technique, based on sub-sampling, affect the performance.

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

JEL Classification Code: G11.

*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, such that covariances or correlations between returns on different assets are 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 consider the statistical accuracy of (co-)variance forecasts only. 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 data instead of daily data. Similarly, Liu (2004) constructs 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 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 address the optimal sampling frequency issue for constructing mean-variance efficient portfolios from the individual constituents of the S&P 100 index. In particular, we consider 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. In addition, we examine how different bias- and variance-reduction techniques affect the choice of sampling frequency. First, we explore the added value of subsampling as proposed by Aït-Sahalia *et al.* (2005). Subsampling 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. 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. Third, and finally, we consider the suggestion of Fleming *et al.* (2003) for correcting the bias in realized (co-)variances by means of scaling.

Our main findings are as follows. For both minimum variance and minimum tracking error portfolios, we find that 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 typically ranges between 30 and 65

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.

minutes, considerably lower than the popular five-minute frequency. The same result occurs for minimum variance portfolios subject to a target return. In contrast, for the minimum tracking error portfolio subject to a target return the optimal sampling frequency appears to be much higher at 1- to 2-minutes. These findings are robust to the use of subsampling and lead-lag bias correction. Both of these techniques marginally improve the out-of-sample 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. Finally, the bias correction procedure of Fleming *et al.* (2003) runs into problems as the resulting bias-adjusted conditional covariance matrices are often not positive definite.

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. In Section 4 the results are discussed. Finally, Section 5 concludes.

2 Data

The data set was obtained from Price-Data.com² 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 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

²<http://www.price-data.com/>

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

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

where $r_{t,c-o}$ is the $(N \times N)$ vector of close-to-open (overnight) returns from day $t-1$ (close) to day t (open),³ 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, which are denoted 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 all sampling frequencies such that $390/I$ is integer, that is 1,2,3,5,10,15,30,65 and 130 minutes. 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 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 exactly at the end of each (intraday) interval. The average covariance using one-minute returns is 0.826, whereas

³For obvious reasons the overnight return from 10 to 17 September, 2001 (the first trading day after 9/11) has been dropped.

for daily data it is almost double at 1.568. Third, the variance of the realized (co-)variance 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.⁴

- 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 Aït-Sahalia *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 subsample is 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. Summary statistics for the realized (co-)variances that are obtained with this subsampling procedure are presented in Panel B of Table 1. In general the effects 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. Note that the average realized (co-)variances are not affected, that is subsampling does not affect the bias.⁵

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

⁴An exception is the variance at the one- and two-minute frequencies, where also the variance increases due to the increased importance of bid-ask bounce.

⁵Zhang *et al.* (2005) suggest a bias-correction procedure based on combining the realized variance obtained with subsampling with the realized variance obtained with the highest available sampling frequency.

(2005, 2006). In particular, let $\Gamma_{t,\Delta,l}$ denote the l -th order autocovariance matrix of intraday Δ -period returns, that is

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

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

$$V_{t,\Delta,b} = V_{t,\Delta} + \sum_{l=1}^q d_l (\Gamma_{t,\Delta,l} + \Gamma'_{t,\Delta,l}), \quad (2)$$

where $V_{t,\Delta}$ is given by (1) and the weights d_l for the leads and lags are taken to be either $d_l = 1$ for all $l = 1, \dots, q$ or $d_l = 1 - l/(q+1)$. The equal-weighting scheme is commonly used for estimating market betas of illiquid stocks and was suggested by Zhou (1996) in the context of realized variance. The use of the Bartlett-weights $d_l = 1 - l/(q+1)$, on the other hand, guarantees that the realized covariance matrix $V_{t,\Delta,b}$ is positive definite, see Newey and West (1987).

Panels C and D of Table 1 present characteristics of $V_{t,\Delta,b}$ with $q = 1$ and $d_l = 1$ and $d_l = 1 - q/(l+1)$, respectively. As expected, the bias in both realized variances and realized covariances is reduced for all frequencies, in particular when the equal-weighting scheme is applied. For example, the average realized variance based on 1-minute returns is reduced to 7.556, only slightly higher than the average daily squared return of 7.386. Similarly, the average realized covariance at the 1-minute frequency is increased to 1.223, which comes much closer to the average cross-product of daily returns (1.586) than the standard case. Note, however, that the reduction in bias generally comes at the cost of increased variance, which also is more pronounced when $d_l = 1$. An exception is the 1-minute frequency where not only the average variance is reduced and much closer to the average daily squared return, but at the same time the variance is reduced from 597 to 501. In general, whether the net effect on the MSE is positive or negative will depend on the data as well as the empirical application.

3 Methodology

The benefits of high-frequency intraday data and the optimal way to employ these will be gauged by their application 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 optimization problems for each day t :

$$\begin{aligned} \min_{w_t} w_t' \Sigma_t w_t \\ \text{s.t. } w_t' \iota = 1 \end{aligned} \quad (3)$$

and

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

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 is the $(N \times 1)$ vector with conditional expected returns for the individual stocks, that is $\mu_t \equiv \mathbb{E}[r_t | \mathcal{I}_{t-1}]$, 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 $\Sigma_t \equiv \mathbb{E}[(r_t - \mu_t)(r_t - \mu_t)' | \mathcal{I}_{t-1}]$. We return to these below. The solution to the problem in (3), the weights for the fully invested minimum variance portfolio, is given by

$$w_{t,\text{MVP}} = \frac{\Sigma_t^{-1} \iota}{\iota' \Sigma_t^{-1} \iota}. \quad (5)$$

For the solution of the problem in (4) first weights for the maximum Sharpe ratio portfolio are computed as

$$w_{t,\text{MSR}} = \frac{\Sigma_t^{-1} \mu_t}{\mu_t' \Sigma_t^{-1} \mu_t} \quad (6)$$

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

$$w_{t,P} = \frac{\mu_{t,\text{MSR}} - \mu_P}{\mu_{t,\text{MSR}} - \mu_{t,\text{MVP}}} w_{t,\text{MVP}} + \frac{\mu_P - \mu_{t,\text{MVP}}}{\mu_{t,\text{MSR}} - \mu_{t,\text{MVP}}} w_{t,\text{MSR}} \quad (7)$$

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. Subsequently the performance of the portfolios is evaluated using the actual stock returns. For the minimum variance portfolio, we consider the actual standard deviation, and for the target return portfolios we monitor the actual return, standard deviation, and Sharpe ratio.

In addition the above analysis is repeated for stock returns in excess of the S&P 500 futures returns. The solution to the problem in (3) 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 (4) 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, solves this problem.

Implementation of the portfolio construction methods discussed above requires estimates or forecasts of the vector of conditional mean returns μ_t and the conditional covariance matrix Σ_t . 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 return in the complete out-of-sample period.⁶ 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 closely follow Fleming *et al.* (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

$$\hat{\Sigma}_t = \sum_{k=1}^{\infty} \Omega_{t-k} \otimes r_{t-k} r'_{t-k} \quad (8)$$

where Ω_{t-k} is a symmetric $(N \times N)$ matrix of weights, and \otimes denotes element-by-element multiplication. Fleming *et al.* (2001, 2003) choose the weighting scheme $\Omega_{t-k} = \alpha \exp(-\alpha k) \iota \iota'$, such that (8) can be rewritten as

$$\hat{\Sigma}_t = \exp(-\alpha) \hat{\Sigma}_{t-1} + \alpha \exp(-\alpha) r_{t-1} r'_{t-1}. \quad (9)$$

⁶As 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).

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 parameter (α) to control the rate at which the weights decay with lag length guarantees that $\hat{\Sigma}_t$ is positive definite. One way of interpreting this weighting scheme is as a restricted multivariate GARCH model.⁷ The optimal decay rate can therefore be estimated using (quasi) maximum likelihood for the model

$$r_t = \hat{\Sigma}_t^{1/2} z_t \quad (10)$$

where $z_t \sim NID(0, I)$ and $\hat{\Sigma}_t$ is given by (9). 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 $\hat{\Sigma}_t$ 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 $w_{t,MVP}$ from (5) and $w_{t,P}$ from (7) 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 here 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 the next section.

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 (9) by the realized covariance matrix V_{t-1} , that is, the conditional covariance matrix is estimated using high-frequency data as

$$\hat{\Sigma}_{t,\Delta} = \exp(-\alpha_\Delta)\hat{\Sigma}_{t,\Delta} + \alpha_\Delta \exp(-\alpha_\Delta)V_{t-1,\Delta} \quad (11)$$

where α_Δ can again be estimated by means of maximum likelihood for the model (10), but now using $\hat{\Sigma}_{t,\Delta}$ instead of $\hat{\Sigma}_t$. For $V_{t-1,\Delta}$ in (11), we consider the realized

⁷Fleming *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.

covariance matrix obtained from the ‘basic’ form given in (1), from the subsampling procedure described in Section 2, and from the lead-lag correction given in (2).

Fleming *et al.* (2003) mention three potential biases for the realized covariance matrix $V_{t,\Delta}$ and, hence, for the conditional covariance matrix estimate $\hat{\Sigma}_{t,\Delta}$ obtained from (11). First, the outer product of the vector of overnight returns is an imprecise estimator of the covariance matrix over the non-trading period. Second, during the trading period it uses non-simultaneous price observations across the assets under consideration. Third, the intraday returns exhibit serial correlation induced by price discreteness and bid-ask bounce. They argue that the impact of the last two factors can be limited by choosing the sampling interval appropriately, and propose a bias correction procedure (described below) primarily for the first potential bias. In addition Hansen and Lunde (2005, 2006) demonstrate (for the realized variance) that the same bias correction procedure can in fact be useful for correcting the adverse effects of non-trading and bid-ask bounce as well. The evidence in Table 1 suggests that these problems may be relevant for the S&P 100 stock data at the popular five minute frequency. This provides all the more reason to consider Fleming *et al.*’s bias corrections to the rolling estimator in (11), not just for the noise in the overnight return but also for the downward bias in the estimated covariances.⁸

The bias corrections proposed by Fleming *et al.* (2003) are based on scaling the elements of $\hat{\Sigma}_{t,\Delta}$ with factors determined from the contemporaneous estimates from the daily-returns-based rolling estimator. Specifically, the i th diagonal element of $\hat{\Sigma}_{t,\Delta}$ is replaced by

$$\hat{\sigma}_{it,\Delta}^{*2} = \left(\frac{\sum_{l=1}^q \hat{\sigma}_{i,t-l}^2}{\sum_{l=1}^q \hat{\sigma}_{i,t-l,\Delta}^2} \right) \hat{\sigma}_{it,\Delta}^2 \quad (12)$$

where $\hat{\sigma}_{it}^2$ and $\hat{\sigma}_{it,\Delta}^2$ denote the i th diagonal elements of $\hat{\Sigma}_t$ (based on daily returns) and $\hat{\Sigma}_{t,\Delta}$ (intraday returns), respectively. This way the conditional variance based on intraday returns is adjusted by the average bias relative to the conditional variance based on the daily returns measured over the last q days. Likewise, to correct for

⁸Recall that adding leads and lags to the realized covariance matrix, as in (2), is also aimed at reducing the bias. Hence, the bias correction of Fleming *et al.* (2003) is an alternative to this procedure.

the bias in the realized covariances, the off-diagonal elements of $\hat{\Sigma}_{t,\Delta}$ are replaced by

$$\hat{\sigma}_{ijt,\Delta}^* = \hat{\sigma}_{it,\Delta}^* \hat{\sigma}_{jt,\Delta}^* \left[\frac{\hat{\sigma}_{ijt,\Delta}}{\hat{\sigma}_{it,\Delta}^* \hat{\sigma}_{jt,\Delta}^*} + \frac{1}{q} \sum_{l=1}^q \left(\frac{\hat{\sigma}_{ij,t-l}}{\hat{\sigma}_{i,t-l,\Delta} \hat{\sigma}_{j,t-l,\Delta}} - \frac{\hat{\sigma}_{ij,t-l,\Delta}}{\hat{\sigma}_{i,t-l,\Delta} \hat{\sigma}_{j,t-l,\Delta}} \right) \right]. \quad (13)$$

Fleming *et al.* (2003) use this additive correction for the covariances because the bias measure can be positive, negative, or zero. Also, the term inside brackets (‘correlation’) is constrained by plus and minus one. Note that by correcting the rolling covariance estimator directly rather than the realized covariance matrix ($V_{t,\Delta}$), the bias correction becomes a function of the decay rates α and α_Δ . Thus, first α is estimated, and subsequently the bias corrections are computed in the course of estimating α_Δ . In (12) and (13) the choice of q is found not to be very important, as long as it is small enough to capture the time variation in the biases. Fleming *et al.* (2003) settle on $q = 22$ trading days and we do so here as well.

Equations (8) through (13) fully describe the approach of Fleming *et al.* (2001, 2003) that we employ here to investigate the benefits of intraday data for 78 S&P 100 stocks. As mentioned before, we examine different sampling frequencies to construct the realized covariance matrix V_t . In particular, we divide the 390-minute NYSE trading session in 1, 2, 3, 5, 10, 15, 30, 65 or 130 minute intervals, all of which cover the full length of the trading day. In addition for each of these sampling frequencies we explore the benefits of subsampling and of bias-correction by means of including a single lead and lag.

4 Results

4.1 Positive definite problems

The results discussed in this section do not concern portfolios constructed with the bias-adjusted covariance matrix estimates using the scaling approach advocated by Fleming *et al.* (2003). The key problem with the bias adjustments as given in (12) and (13) is that these adjust each individual element in the covariance matrix separately, with a possibly different correction factor. Hence, whereas the unadjusted covariance matrix $\hat{\Sigma}_{t,\Delta}$ obtained from (11) is guaranteed to be positive definite, this does not hold for the bias-adjusted matrix $\hat{\Sigma}_{t,\Delta}^*$. In fact, we repeatedly ran into trouble trying to estimate the optimal decay rates that maximize the likelihood of the model in equations (8) and (10) due to the determinant of the covariance matrix

being negative (whereas the likelihood requires the log of this determinant) and $\hat{\Sigma}_{t,\Delta}^*$ not being invertible (also needed to compute the likelihood). We tried to address this problem in several different ways but to no avail. First, following Ledoit and Wolf (2003) we considered shrinking the adjusted matrix $\hat{\Sigma}_{t,\Delta}^*$ back towards the unadjusted matrix $\hat{\Sigma}_{t,\Delta}$ that is known to be positive definite. Although this solved the problems in computing the log-likelihood, the log-likelihood surface became discontinuous with discomfoting spikes rather than a smooth run towards an optimal decay parameter. The use of ‘optimal’ decay parameters resulted in a dramatic out-of-sample performance of all portfolios. Second, we adjusted all stocks with the same average required bias correction but this also gave unsatisfactory results. One reason for this finding is that some of the required bias corrections to the realized covariance, as given in (13), are enormous for individual pairs of stocks, simply because the realized covariance for that stock pair for the last month is close to zero. For these reasons we only consider non-adjusted realized covariance matrices, leaving the search for satisfactory bias adjustment procedures for further research. We do consider the realized covariance matrices obtained with subsampling, and the matrices obtained with the lead-lag correction in (2). For the latter, we only report results obtained with Bartlett weights, $d_l = 1 - l/(q + 1)$. As mentioned before, the popular scheme with unit weights, $d_l = 1$, does not guarantee a positive definite covariance matrix. For the 78 S&P 100 stocks this weighting scheme indeed ran into this problem quite frequently. Again shrinkage led to discomfoting spikes in the log-likelihood surface. In contrast, the Bartlett weighting scheme does guarantee a positive definite covariance matrix.

4.2 Optimal decay rates

Table 2 shows the optimal decay rates α and α_Δ that maximize the likelihood of the model in equations (10) with (9) for daily returns and with (11) 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 subsampling, 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 $V_{t-1,\Delta}$ in (11) 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 decay parameters when subsampling is applied to obtain the realized covariance matrices are in general somewhat larger, which can be attributed to an overall reduction in the variance (noise) of the updating realized (co-)variances. For the lead-lag correction the bias decreases whereas the variance increases for a particular sampling frequency. It appears that the latter is more important here, given that the decay parameters are lower for the corrected covariance matrices compared to the standard case. For the highest frequencies the log-likelihood is improved, however, when using the Bartlett correction. 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.

4.3 Portfolio performance

Table 3 shows the performance of the overall minimum variance portfolio, with weights defined in (5), and the minimum variance portfolio given an annualized target return of 10%,⁹ with weights given by (7). 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 almost 14% 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.

⁹We 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 to those reported below. Detailed results are therefore not shown here, but are available on request from the authors.

- insert Table 3 about here -

The results in Panel B of Table 3, using subsampling, show only a marginal improvement for the overall minimum variance portfolio with a standard deviation of 12.07% 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 certain frequencies subsampling leads to a higher Sharpe ratio but for other frequencies it declines. At the optimal frequency of 130 minutes the Sharpe ratio is lower at 0.709, compared to 0.786 for the 65-minute frequency in the standard case.

The Bartlett correction in (2) 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 12.07% at the optimal 65-minute frequency in the standard and subsampling cases, respectively. In fact, the the Bartlett 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, both due to the use of autocovariances and due to the influence of microstructure effects, 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 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. Again subsampling provides a marginal improvement with the minimum tracking error equal to 4.17% for the 65-minute frequency. Finally, using one lead and lag with Bartlett weights results in a higher optimal sampling frequency of fifteen minutes as for the minimum variance portfolio, with a lower tracking error at 4.35%. Hence here we also observe that bias-correction using leads and lags helps.

In sum, the general conclusion from Tables 3 and 4 when computing the minimum variance portfolio or minimum tracking error portfolio is that subsampling and

the lead-lag bias correction marginally improve the out-of-sample performance for both the minimum variance portfolios and the minimum tracking error portfolios. 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 12.69% 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%).

- 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. 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. Subsampling results in an optimal frequency of three minutes with an information ratio of 0.413. Although this is below the optimum in the standard case, the benefits of subsampling are clear when comparing the information ratios at other frequencies with the corresponding results in the standard case. Apart from the two-minute frequency subsampling always improves the information ratio. Also subsampling leads to a lower annualized tracking error in all cases, including the two-minute frequency.

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.

4.4 Genuine out-of-sample forecasting

FKO (2001; 2003) suggest that determining the decay parameters α and α_Δ 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 proceeded as follows. First we find the decay parameters that maximize the performance of the various stock portfolios over the first 250 days following the initial burn-in period, i.e. the values of α and α_Δ that minimize the (relative) variance or maximizes the Sharpe (or Information) ratio. These decay parameters are then used to estimate the conditional covariance matrices $\hat{\Sigma}_t$ and $\hat{\Sigma}_{t,\Delta}$ for the first day following the in-sample period, for which optimal portfolio weights are then constructed using (5) and (7). This procedure is repeated using a moving in-sample estimation window of 250 days, where every day a new observation is added and the oldest one deleted. 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 5 about here -

The results are presented in Table 5. First of all, for both the minimum variance (MV) and minimum tracking error (MTE) portfolios the results are re-assuring. The optimal sampling frequency is still 65 minutes and 30 minutes for MV and MTE, respectively. Also the performance itself is similar to that of the standard case. Second, 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 minutes, and the performance has deteriorated from 0.640 to 0.554. In the excess return case the optimal sampling frequency is now 1 minute instead of 2 minutes, but with a better information ratio at 0.457 versus 0.373 for the maximum likelihood case. Third, and 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 alphas over time, we also did a datasnooping exercise with a constant decay parameter, where we determine the decay parameter that maximizes performance (rather than the log-likelihood) over the entire out-of-sample period. These results (not reported here) confirm that performance-based alphas are much lower than the log-likelihood based alphas. 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 U.S. S&P 100 index the optimum is more likely to be in the neighbourhood of an hour rather than five or thirty minutes.

We also investigated the use of bias- and variance-reduction methods for computing the realized covariances. Both the bias correction procedure proposed by Fleming *et al.* (2003) and the use of one lead and one lag following Scholes and Williams (1977) fail to produce better results. Subsampling, however, does help in reducing the variance and results in a marginal improvement over the use of a single

sample at the same frequency.

For further research it would be interesting to test other ways to correct for biases in realized covariances due to non-synchronous trading. In addition 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. Bollerslev and Zhang (2003) is an example where the Fama and French 3-factor model coefficients are estimated using five-minute returns, albeit not with the purpose of estimating or predicting the covariance matrix. Another interesting topic for further research is putting restrictions on the portfolio weights, which we do 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.

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.
- 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.
- 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.
- 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.
- 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), Semi-Parametric Modelling of Correlation Dynamics, Econometric Institute Research Report No. 2005-26.
- 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.
- Jagannathan, R. and T. Ma (2003), Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps, *Journal of Finance* **58**, 1651–1683.
- 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.
- Ledoit, O., and M. Wolf (2003), Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection, *Journal of Empirical Finance* **10**, 603–621.
- 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.
- 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., 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
BMJ*	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		Realized Variance		Realized Covariance	
	Mean	Variance	Mean	Variance	Mean	Variance	Mean	Variance
Daily	7.386	1763	1.568	93.58				
	Panel A: Standard				Panel B: Subsampling			
130 minutes	7.369	689.7	1.394	31.94	7.849	1083.8	1.458	36.63
65 minutes	7.324	624.3	1.357	22.46	7.329	755.2	1.409	22.52
30 minutes	7.311	563.5	1.316	16.48	7.228	601.7	1.394	16.30
15 minutes	7.463	545.8	1.307	14.46	7.384	565.2	1.381	14.26
10 minutes	7.614	547.2	1.305	13.25	7.533	557.5	1.349	13.05
5 minutes	7.912	531.7	1.239	11.49	7.856	550.3	1.254	11.18
3 minutes	8.193	527.6	1.136	10.34	8.165	539.0	1.140	10.12
2 minutes	8.525	537.1	1.025	9.60	8.509	541.8	1.029	9.47
1 minute	9.494	597.0	0.826	8.73	9.494	597.0	0.826	8.73
	Panel C: lead-lag, $q = 1$, $d_l = 1$				Panel C: lead-lag, $q = 1$, $d_l = 1 - l/(q + 1)$			
130 minutes	7.476	883.3	1.442	48.93	7.422	758.6	1.418	36.73
65 minutes	7.412	747.1	1.422	35.60	7.368	667.0	1.389	26.27
30 minutes	7.344	654.7	1.386	25.32	7.329	595.6	1.351	18.97
15 minutes	7.221	577.7	1.357	18.95	7.342	552.7	1.332	15.51
10 minutes	7.227	565.7	1.363	16.98	7.420	545.6	1.334	14.23
5 minutes	7.310	550.5	1.362	14.94	7.611	532.3	1.300	12.67
3 minutes	7.431	562.5	1.377	13.62	7.812	538.3	1.257	11.56
2 minutes	7.514	550.7	1.360	12.50	8.020	536.3	1.193	10.70
1 minute	7.556	501.1	1.223	10.70	8.525	533.4	1.025	9.47

Notes: The table shows mean and variance of the realized (co-)variance 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 variance the mean reflects the average of all realized variances taken over 78 stocks and 1788 trading days. For the variance of the variance the average is taken over the 78 sample variances of the (realized) variances. For the covariance the mean reflects the average of all realized covariances taken over all 3003 combinations and 1788 trading days. The variance of the covariance is the average taken over 3003 sample variances of the (realized) covariances. Panel A contains results for the ‘standard’ realized covariance matrix as given in (1). In Panel B subsampling is used, implying that for example the 10-minute window is laid over the data in 10 different ways rather than just one as in Panel A. Subsequently each day the average is taken over the 10 resulting estimates. In panels C and D a single lead and lag are added to the variance and covariances according to (2) with $q = 1$, with weights $d_l = 1$ and $d_l = 1 - l/(q + 1)$, respectively.

Table 2: Optimal decay parameters

Frequency	Standard		Subsampling		1 lead, 1 lag	
	α_Δ	LogL	α_Δ	LogL	α_Δ	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	0.0131	-275,282	0.0111	-276,939
65 minutes	0.0149	-274,580	0.0162	-273,335	0.0137	-274,782
30 minutes	0.0204	-273,747	0.0216	-272,419	0.0179	-273,519
15 minutes	0.0273	-273,802	0.0297	-272,587	0.0231	-273,186
10 minutes	0.0329	-274,121	0.0365	-273,068	0.0273	-273,261
5 minutes	0.0481	-275,004	0.0556	-274,322	0.0375	-273,774
3 minutes	0.0678	-275,975	0.0840	-275,551	0.0493	-274,407
2 minutes	0.1025	-276,985	0.1232	-276,631	0.0643	-275,164
1 minute	0.2106	-278,971	0.2106	-278,972	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	0.0133	-273,255	0.0112	-275,033
65 minutes	0.0151	-272,712	0.0165	-271,296	0.0138	-272,907
30 minutes	0.0208	-271,926	0.0222	-270,457	0.0183	-271,669
15 minutes	0.0282	-272,100	0.0308	-270,784	0.0238	-271,381
10 minutes	0.0342	-272,516	0.0384	-271,384	0.0282	-271,522
5 minutes	0.0514	-273,562	0.0607	-272,841	0.0393	-272,157
3 minutes	0.0757	-274,663	0.0966	-274,183	0.0529	-272,925
2 minutes	0.1178	-275,675	0.1415	-275,299	0.0713	-273,751
1 minute	0.2468	-277,537	0.2468	-277,537	0.1440	-275,393

Notes: The table shows the decay rates (α_Δ) that maximize the log-likelihood of the model in (10) and (9) for daily returns and (10) and (11) for intraday returns. 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 index returns). The second and third columns show the optimal decay rates and accompanying log-likelihood values when the covariance updates are based on the standard realized (co-)variances obtained from (1), the fourth and fifth column when the updates are based on subsampling, and the final two columns when a single lead and lag of the covariance is added to the contemporaneous (realized) covariance, according to (2) with $q = 1$ and Bartlett weights $d_l = 1 - l/(q + 1)$.

Table 3: Out-of-sample total performance

Frequency	Target return portfolio			MVP
	μ	σ	SR	σ
Daily	8.842	14.828	0.596	13.997
<u>Panel A: Standard</u>				
130 minutes	10.238	13.134	0.780	12.463
65 minutes	10.116	12.871	0.786	12.157
30 minutes	8.159	12.993	0.628	12.165
15 minutes	8.823	13.120	0.673	12.211
10 minutes	8.288	13.299	0.623	12.381
5 minutes	8.564	13.685	0.626	12.683
3 minutes	8.484	13.916	0.610	12.836
2 minutes	7.671	14.176	0.541	13.062
1 minute	7.889	14.439	0.546	13.325
<u>Panel B: Subsampling</u>				
130 minutes	9.231	13.026	0.709	12.241
65 minutes	8.747	12.869	0.680	12.073
30 minutes	8.000	12.878	0.621	12.098
15 minutes	8.145	13.055	0.624	12.214
10 minutes	8.393	13.223	0.635	12.325
5 minutes	8.456	13.611	0.621	12.606
3 minutes	8.213	13.898	0.591	12.820
2 minutes	8.121	14.090	0.576	12.986
1 minute	7.889	14.439	0.546	13.325
<u>Panel C: 1 lead and 1 lag</u>				
130 minutes	10.230	13.114	0.780	12.435
65 minutes	10.250	12.867	0.797	12.156
30 minutes	9.071	12.807	0.708	12.013
15 minutes	8.410	12.888	0.653	12.054
10 minutes	7.808	13.014	0.600	12.151
5 minutes	8.536	13.283	0.643	12.354
3 minutes	8.502	13.551	0.627	12.567
2 minutes	8.049	13.797	0.583	12.754
1 minute	8.055	14.098	0.571	12.994

Notes: The table shows the out-of-sample performance of the overall minimum variance portfolio, with weights given in (5), and the minimum variance portfolio given a annualized target level of return of 10%, with weights given in (7), constructed using rolling covariance matrix forecasts based on various sampling frequencies and based on different ways to measure the realized covariance matrix (standard, subsampling and using a single lead and lag with Bartlett weights). Columns 2 to 4 show the return, standard deviation and Sharpe ratio of the target return portfolios, and column 5 shows the variance of the global minimum variance portfolio (MVP).

Table 4: Out-of-sample relative performance

Frequency	Target return portfolio			MTEP
	μ	TE	IR	TE
Daily	0.523	4.772	0.110	4.754
<u>Panel A: Standard</u>				
130 minutes	0.450	4.549	0.099	4.532
65 minutes	0.284	4.513	0.063	4.480
30 minutes	0.085	4.477	0.019	4.426
15 minutes	0.574	4.516	0.127	4.461
10 minutes	0.751	4.666	0.161	4.588
5 minutes	0.892	4.986	0.179	4.918
3 minutes	1.845	5.230	0.353	5.176
2 minutes	2.326	5.421	0.436	5.404
1 minute	1.393	5.790	0.241	5.777
<u>Panel B: Subsampling</u>				
130 minutes	0.562	4.258	0.132	4.230
65 minutes	0.497	4.200	0.118	4.169
30 minutes	0.618	4.214	0.147	4.181
15 minutes	0.492	4.306	0.211	4.266
10 minutes	0.932	4.419	0.339	4.362
5 minutes	1.610	4.753	0.413	4.694
3 minutes	2.103	5.090	0.417	5.048
2 minutes	2.214	5.306	0.241	5.287
1 minute				
<u>Panel C: 1 lead and 1 lag</u>				
130 minutes	0.175	4.581	0.038	4.564
65 minutes	0.320	4.503	0.071	4.476
30 minutes	0.409	4.439	0.092	4.400
15 minutes	0.621	4.400	0.141	4.353
10 minutes	0.627	4.475	0.140	4.415
5 minutes	0.716	4.584	0.156	4.505
3 minutes	1.570	4.772	0.329	4.706
2 minutes	1.954	4.965	0.394	4.911
1 minute	2.164	5.330	0.406	5.314

Notes: The table shows the out-of-sample performance of the over-all minimum tracking error portfolio, with weights given in (5), and the minimum variance portfolio given an annualized target level of return of 5%, with weights given in (7), constructed using rolling covariance matrix forecasts based on various sampling frequencies and based on different ways to measure the realized covariance matrix (standard, subsampling and using a single lead and lag with Bartlett weights). Columns 2 to 4 show the excess return, tracking error and Information Ratio of the target return portfolio, and column 5 shows the ex-post tracking error of the minimum tracking error portfolio (MTEP).

Table 5: Out-of-sample α 's

Frequency	Target return portfolio				Minimum risk portfolio			
	α		SR/IR	SR/IR	α		σ /TE	σ /TE
	Mean	St.Dev			Mean	St.Dev		
<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 portfolio, with weights given in (5), and the minimum variance portfolio given an annualized target level of return of 5%, with weights given in (7), 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 a moving window period of 250 days. Columns 2 and 3 report the mean and standard deviation of the resulting estimates of α_Δ . Column 4 shows the Sharpe ratio (panel A) or the Information ratio (panel B) for the resulting portfolios. Column 5 shows the SR/IR 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.