Realized Covariance Estimation in Dynamic Portfolio Optimization

Lada Kyj*
School of Business and Economics
Humboldt-Universität zu Berlin
and Quantitative Products Laboratory

Barbara Ostdiek
Jones Graduate School of Management
Rice University

Katherine Ensor
Department of Statistics
Rice University

March 19, 2009

Abstract

Mean-variance portfolio optimization requires both invertible and well-conditioned covariance matrices. This paper compares the performance of covariance conditioning techniques applied to the realized covariance matrices of the portfolio of constituents of the Dow Jones Industrial Average. We use the volatility of portfolio returns derived from volatility-timing investment strategies employing different conditioning techniques as the criterion of assessment. As portfolio dimensions increase there is increasing need for matrix conditioning to maintain the precision improvement offered by intraday data. We find that the relative performance of the single factor model provides a computationally tractable alternative to fully estimated realized covariance matrices in a global minimum variance dynamic portfolio setting.

1 Introduction

Markowitz mean-variance (MV) optimization is the standard framework for optimal portfolio construction (See Chan, Karceski, and Lakonishok (1999), Jagannathan and Ma

*Corresponding author: Lada Kyj, Quantitative Products Laboratory, Alexanderstrasse 5, D-10178 Berlin, Germany, email: lada.kyj@wiwi.hu-berlin.de

(2003), and references therein). MV optimization requires covariance matrices to be not only invertible, but also well-conditioned. Michaud (1989) points out that matrix inversion maximizes the effects of errors in the input assumptions and, as a result, practical implementation is problematic. Consistent with this, Britten-Jones (1999) find the sampling error of the weights of mean-variance efficient portfolios to be very large.

Realized covariance estimation has emerged as a viable candidate for covariance estimation. This class of estimators employs high-frequency data and provides more precise estimates. In a low dimensional setting, Fleming, Kirby, and Ostdiek (2003) have shown that realized covariance estimates provide utility gains over implied covariance estimates for risk averse investor following a MV optimization strategy. More recently, Liu (2008) determined that a manager tracking the S&P 500 index with the DJIA stocks will switch from covariance estimates based on daily data to estimates using intraday data when there are less than 6 months of historical data available, or if she rebalances her portfolio daily.

The realized covariance literature has focused on improving these estimators by using techniques such as cross-market tick-matching, optimal sampling frequencies, and sub-sampling. Cross-market tick matching was introduced independently by Corsi (2006), and Hayashi and Yoshida (2005), and implementation concerns have been examined by Voev and Lunde (2007), and Griffin and Oomen (2006). Optimal sampling techniques attempt to identify the sampling frequency that balances the introduction of market microstructure effects by too frequent sampling against the loss of information from sampling too sparsely. (See Bandi and Russell (2006), Bandi, Russell, and Zhu (2008), Oomen (2006), and de Pooter, Martens, and van Dijk (2008).) Sub-sampling, introduced by Zhang, Mykland, and Ait-Sahalia (2005), is a technique used to reduce the variance of realized covariance estimators. Realized kernel estimation using refresh time sampling introduced by Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008b) synthesizes these refinements into one estimator.

In general, even without sampling refinements, estimation of high dimensional realized covariance matrices is both computationally expensive and plagued by sampling error. The required computational time affects implementation feasibility and sampling error can result in numerically ill-conditioned matrices, yielding noisy estimate and possibly making inversion problematic. Under the assumption of normally distributed data, Ledoit and Wolf (2003) suggest that to minimize ill-conditioning the number of observations, n, needs to be at least ten times the number of dimensions, p.

Previous literature has addressed imprecise covariance matrix estimates by imposing more structure on the covariance matrix. Variants of shrinkage are employed to mitigate ill-conditioned matrices. Fleming, Kirby, and Ostdiek (2003) and de Pooter, Martens, and van Dijk (2008) use rolling estimators, Jagannathan and Ma (2003) use non-negative constraints on the portfolio weights, and Ledoit and Wolf (2003) use shrinkage toward the market estimate. Fan, Fan, and Lv (2008) find that the major advantage of factor models is in the estimation of the inverse of the covariance matrix and demonstrate that the factor model provides a better conditioned alternative to the fully estimated covariance matrix. Using daily level data, Chan, Karceski, and Lakonishok (1999) and Jagannathan and Ma (2003) show that factor models can reduce the variance of optimal mean-variance portfolios and offer utility gains over strategies employing full sample covariance matrices. Specifically, Jagannathan and Ma (2003) compare the performance of portfolios determined using sample covariance matrix estimates (based on daily data). They conclude that an ad-hoc non-negativity on portfolio weights provides performance similar to the performance of portfolios based on factor and shrinkage covariance estimators. They show that the single factor model performs well when the number of observation is not much greater than the number of dimensions.

To date, the realized covariance literature has focused on evaluating estimators for a very small number of assets. The exceptions to this are Voev (2008) and Liu (2008) who focus on forecasting the covariance matrix for portfolios of 15 assets and 30 assets, respectively. Both studies conclude that some form of smoothed realized covariance estimates provides improvements over daily level estimates. Barndorff-Nielsen and Shephard (2004) provides the framework for estimating realized betas. Bollerslev and Zhang (2003), Andersen, Bollerslev, Diebold, and Wu (2006), and Morana (2007) consider realized betas using daily returns. They conclude that betas are time-varying and persistent. Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008b) exploit this persistence and forecast realized betas using intra-day returns.

Our paper contributes to the discussion by examining three features of interest to both academics and practitioners. First, we examine both the in-sample and out-of-sample performance of the estimators to disentangle estimation precision from the efficacy of the forecasting model. Second, we examine the performance of realized covariance estimators as the dimension increases from 3 to 30 and show the importance of conditioning the estimates at high dimensions. We find that at high dimensions, computationally sparse estimators offer portfolio volatilities that are similar to the best performing fully estimated covariance matrix. Third, we compare the performance of covariance conditioning techniques in a volatility timing setting and with traditional forecast evaluation regressions. We compare the performance of the single factor model to that of a fully estimated realized covariance matrix for portfolios of the 30 Dow Jones Industrial component stocks. We conclude that the single factor model offers a compelling, computationally tractable alternative.

The remainder of the paper is organized as follows. Section 2 provides the realized covariance estimator and the conditioning approaches. Section 3 provides the performance assessment criteria including the global minimum variance volatility timing strategy and the portfolio-level Mincer-Zarnowitz forecast evaluation. In Section 4 we discuss the data, the necessary parameter estimates, and the empirical results. Section 5 concludes.

2 Methods

2.1 Realized Covariance Estimators

The discretely observed price process $p(t_i)$ is a function of both the latent price process $x(t_i)$ and the market microstructure effects $u(t_i)$, which are treated as "observation error" such that the price of asset A is observed as:

$$p_A(t_i) = x_A(t_i) + u_A(t_i), \quad i = 1, 2, \ldots, n. \tag{1}$$

Returns are written as

$$r_A(t_i) = \Delta p_A(t_i). \tag{2}$$

Andersen, Bollerslev, Diebold, and Labys (2001) first proposed realized variance estimation using ad-hoc calendar-time sampling. Synchronous observations across markets are achieved by interpolating previous tick prices onto an ad-hoc common sampling grid (e.g., every 5 minutes) yielding m equally spaced intraday observations. The calendar time realized covariance estimator for asset A and B can be written as:

$$\widehat{\Sigma}_{AB}(m) = \sum_{i=1}^{m} r_A(t_i) \times r_B(t_i). \tag{3}$$

For variance reduction, Zhang, Mykland, and Ait-Sahalia (2005) have advocated sub-sampling as a technique that exploits the richness of high frequency data by sparsely sampling the observations into non-overlapping grids and then smoothing the estimator by averaging over subgrids. The resulting sub-sampled estimator of K subgrids is:

$$\widehat{\Sigma}^{*}_{AB}(m) = \sum_{k=1}^{K} \widehat{\Sigma}^{(k)}_{AB}(m)$$

where $\widehat{\Sigma}^{(k)}_{AB}(m)$ is the realized covariance for the kth grid. As outlined in Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008a), sub-sampling has the same asymptotic distribution as the Bartlett kernel within the realized kernel framework. Alternative methods for estimating realized covariance in the presence of market micro structure noise are numerous, but are beyond the scope of this paper. (See McAleer and Medeiros (2008) for a comprehensive review.) In this study, we limit the analysis to actively traded stocks which reduces market microstructure biases on both the variance and covariance estimates. Following Andersen, Bollerslev, Diebold, and Labys (2003) we using a five-minute sampling frequency.

2.2 Ill-Conditioned Covariance Matrices

Many applied problems in finance require a covariance matrix estimator that is not only invertible, but also well-conditioned. While the true covariance matrix is guaranteed to be well-conditioned, covariance estimates may not be due to sampling error. The sample covariance matrix is a consistent estimate of the true covariance matrix as $\frac{p}{n} \to 0$, but when $\frac{p}{n} \to constant$, it may be ill-conditioned.

A strictly positive definite matrix is necessary for matrix inversion, a necessary step in the mean-variance framework. A symmetric matrix A, with eigenvalues λ_i, is positive definite when all the eigenvalues of A satisfy $\lambda_i > 0$. The relationship between positive definiteness and invertibility is understood via the eigenvalues, where the determinant is defined as: $det(A) = \prod_{i=1}^{p} \lambda_i$. A matrix is invertible when the $det(A) \neq 0$, therefore non-negative eigenvalues ensure that a positive definite matrix is invertible.

"Well-conditioned" is a numerical property of an operator and states that all small perturbations of x lead to only small changes in $f(x)$. The condition number of a matrix A is defined as: $\kappa(A) = \|A\|\|A^{-1}\|$, where $\| \cdot \|$ is the Frobenius norm and can be expressed as: $\|A\| = \sqrt{tr(AA^T)}$. We note that $\|A\|^2 = \sum_{i=1}^{p} \sum_{j=1}^{p} A_{ij}^2$, and in our case A is symmetric so $\|A\|^2 = \sum_{i=1}^{p} \lambda_{Ai}^2$ where λ_{Ai} are the eigenvalues of A. In this setting, the condition number can be interpreted as the eccentricity of the ratio of eigenvalues $\kappa(A) = \frac{\lambda_{max}}{\lambda_{min}}$. An ill-conditioned matrix has a very large condition number and is close to being numerically non-invertible.

The relative magnitude of the eigenvalues of the realized covariance matrices play a prominent role in mean-variance asset allocation. Positive definiteness requires the eigenvalues be positive and the well-conditioned property imposes an additional requirement of non-vanishing eigenvalues. Intuitively, as an eigenvalue shrinks toward 0, a principal component vanishes and the rank is reduced and the system is collinear. Numerically, as the smallest eigenvalue vanishes toward zero, the condition number explodes to infinity. This problem highlights the importance of preserving the smallest eigenvalue which is accomplished by imposing more structure to mitigate the imprecision of covariance estimates. In the following subsections we present two approaches to generating well-conditioned, consistent estimates of the true covariance matrix: shrinkage and single-factor models.

2.3 Shrinkage Estimators

As the number of dimensions increases relative to the number of observations, the resulting covariance matrices become ill-conditioned and, in particular, are characterized by the smallest estimated eigenvalues being too small and the largest being too big relative to the true eigenvalues. Shrinking these estimated covariance matrices towards some idealized structure yields more stable estimates. The resulting eigenvalues are more compressed and the covariance estimates are better conditioned.

Generally, linear shrinkage can be written as:

$$\widetilde{\Sigma}_t(\alpha) = \alpha G + (1-\alpha)\widehat{\Sigma}_t, \quad \alpha \in [0,1] \tag{4}$$

where $\widehat{\Sigma}_t$ is the estimate of Σ_t, the covariance matrix of returns, and G is the idealized covariance structure. The positive definiteness of the shrinkage estimator can be considered by examining

$$v^T\widetilde{\Sigma}_t v \;=\; \alpha v^T G v + (1-\alpha) v^T \widehat{\Sigma}_t v,$$

where positive definiteness requires $v^T\widetilde{\Sigma}_t v > 0$ for all nonzero vectors v. To ensure positive definiteness, the target matrix G is chosen to be positive definite and the shrinkage factor α is chosen to optimize a criteria, such as MSE, which requires maintaining positive definiteness of the resulting estimator $\widetilde{\Sigma}_t$. We consider two forms of linear shrinkage: the rolling estimator and the Ledoit-Wolf estimator.

2.3.1 Rolling Estimators

Rolling estimation is a common feature in covariance applications (See Fleming, Kirby, and Ostdiek (2003), Bandi, Russell, and Zhu (2008), de Pooter, Martens, and van Dijk (2008), Bandi and Russell (2006) for examples in the realized covariance literature). This approach is motivated by the conditional heteroskedasticity of financial time series. In the presence of conditional heteroskedasticity, estimation of the covariance matrix involves a trade-off between considering a sufficiently large number of observations to obtain an unbiased and consistent estimate and considering a short enough history to accommodate changes in the covariance structure. Rolling realized covariance estimation attempts to balance the statistical power obtained using a large sample against the loss of precision from including stale information. It is easy to see that

$$\widetilde{\Sigma}_t = \alpha\widetilde{\Sigma}_{t-1} + (1-\alpha)\widehat{\Sigma}_t, \quad \alpha \in [0,1] \tag{5}$$

is a variant of shrinkage estimation where the current realized covariance estimate is shrunk toward a function of past estimates. $\widetilde{\Sigma}_t$ then, is an exponentially weighted rolling covariance estimator for Σ_t. Based on the work of Foster and Nelson (1996) and Andreou and Ghysels (2002) we use an exponentially weighted rolling scheme which provides

MSE efficiency gains for realized covariance estimators. We avoid estimating α and instead, following the Risk Metrics approach, we set $\alpha = 0.94$, which Morgan (1996) found produces the best backtesting results. This exponentially weighted moving average (EWMA) method, applied to daily data, is perhaps the most widely used volatility model in industry.

Determining positive definiteness of the rolling estimator draws upon Foster and Nelson (1996) and assumes that $\widetilde{\Sigma}_{t-1}$ is a consistent estimator of Σ_{t-1}. The target matrix possesses the desirable properties of the true covariance matrix: positive definite and well-conditioned. As a result the rolling estimation shrinks the eigenvalues of the realized covariance estimator towards the more consistent eigenvalues of the target matrix resulting in better conditioned covariance matrices.

2.3.2 Ledoit-Wolf Estimators

Ledoit and Wolf (2003, 2004b) introduced an estimator that is an optimal linear combination of a target matrix and the sample covariance matrix under squared error loss:

$$E[\|\widetilde{\Sigma}_t - \Sigma_t\|^2], \tag{6}$$

where $\widetilde{\Sigma}_t(\alpha) = \alpha G + (1-\alpha)\widehat{\Sigma}_t, \quad \alpha \in [0,1]$, and G is the target matrix. Intuitively, this estimator seeks to minimize the quadratic distance between the true and the estimated covariance matrices. We use the equicorrelated matrix, suggested by Ledoit and Wolf (2004a) and used by Voev (2008), as our target matrix.

The equicorrelated matrix is defined as having all the off-diagonal elements of the covariance matrix having the average sample correlation. Φ, with elements ϕ_{ij}, denotes the unobserved true equicorrelated covariance matrix and F, with elements f_{ij}, as the corresponding estimate. The diagonal elements are the variance elements of the sample covariance matrix. The optimal shrinkage parameter α is estimated according to the method outlined in Ledoit and Wolf (2003), where $s_{i,j}$ corresponds to elements of matrix S:

$$\alpha = \frac{\sum_{i=1}^{p}\sum_{j=1}^{p} Var(s_{ij}) - Cov(f_{ij}, s_{ij})}{\sum_{i=1}^{p}\sum Var(f_{ij} - s_{ij}) + (\phi_{ij} - \sigma_{ij})^2} \quad s.t. \quad \alpha \in [0,1]. \tag{7}$$

This procedure shrinks the estimated eigenvalues towards the eigenvalues of the target matrix. The resulting eigenvalues have a smaller maximum and a larger minimum, resulting in a better conditioned estimator.

For implementation within the high-frequency data setting, we draw upon the well know long memory property of realized covariance and assume that the covariance process is locally constant. For a window of length l, we assume that $E[s_{ij,\tau-l}] = \sigma_{ij,t}$ where $\tau \in [t-l, t]$and estimate α. In this formulation S and $s_{i,j}$ represent the realized covariance matrix estimate and its individual elements. The quantities $Var(s_{ij}), Cov(f_{ij}, s_{ij}), Var(f_{ij} - s_{ij})$, and $(\phi_{ij} - \sigma_{ij})^2$ are calculated using the daily time series of realized covariance estimates.

The results of the robustness analysis in Appendix A validate our approach of assuming locally constant realized covariance estimates. Specifically we see that the performance of the Ledoit Wolf shrinkage technique is not overly sensitive to the estimation of the α parameter. While our method may result in less precise α estimates, the sensitivity analysis suggests that in this application misspecification does not have an impact on the shrinkage estimator. Voev (2008) outlines an alternative adaptation of the Ledoit-Wolf estimator for realized covariance estimators, drawing upon the the asymptotic variance and asymptotic covariance results derived in Barndorff-Nielsen and Shephard (2004).

2.4 Factor Model

The market model is a one-factor model that relates individual equity returns to return on a market factor:

$$r_{A(t_i)} = \alpha_A + \beta_A r_{M(t_i)} + \varepsilon_{A(t_i)}, \tag{8}$$

where r_A is the return of the individual asset, r_M is the return of the market as represented by the index, and β_A represents the systemic risk of the asset with the market. We assume that $\varepsilon_{A,t} \sim N(0, \sigma_A^2)$ and that residuals ε_A are uncorrelated to market returns. The resulting covariance matrix is:

$$\Phi = \sigma_M^2 \beta \beta' + D. \tag{9}$$

Here β is the $p \times 1$ vector of betas for all the assets, σ_M^2 is the variance of the market and D is the diagonal matrix of residual variances. It is natural to estimate Φ as:

$$\widehat{\phi}_{i,j} = \begin{cases} s_i^2 & \text{if i=j} \\ s_M^2 \widehat{\beta_i} \widehat{\beta_j} & \text{if i} \neq \text{j}, \end{cases} \tag{10}$$

where, as above, s_i, j are the elements of the estimated realized covariance matix and

$$\widehat{\beta}_A = \frac{\sum_{i=1}^m r_A(t_i) r_M(t_i)}{\sum_{i=1}^m r_M^2(t_i)} = \frac{s_{A,M}}{s_M^2}.$$

In terms of conditioning the matrix, Fan, Fan, and Lv (2008) show that when $K = o(p)$, where K is the number of factors, the inverse of the factor model covariance matrix converges to the true inverse covariance faster than the inverse of the sample covariance matrix, implying that the factor model is a better conditioned alternative to the fully estimated covariance matrix. Moreover, they also show that the factor structure offers quicker convergence to the global minimum variance portfolio. As expected, when the number of factors is proportional, $K = O(p)$, to the number of dimensions, then the two estimators have the same performance. It is easy to show that the covariance matrix of the single-factor model is always positive definite:

$$v'\Phi v = v'\left(\sigma_M^2 \beta\beta' + D\right) v = \sigma_M^2 \underbrace{v'\beta\beta' v}_{\geq 0} + \underbrace{v' D v}_{>0}, \tag{11}$$

for all nonzero vectors v, $v'Dv = \sum_{i=1}^{p} v_i^2 d_{i,i} > 0$ as D is a diagonal matrix of positive values. Likewise, $\beta' v = z$, is a scalar, and as a result $v'\beta\beta' v = z^2 \geq 0$. Hence, single-factor covariance estimates are strictly positive definite. The faster convergence rate of the inverse matrix implies a better conditioned matrix.

3 Assessment Criteria

3.1 Volatility Timing

We consider portfolio allocation as our criterion for assessing estimator performance. Chan, Karceski, and Lakonishok (1999) and Jagannathan and Ma (2003) advocate the use of the Global Minimum Variance (GMV) portfolio to avoid the problematic estimation of μ_t, the vector of expected returns. The GMV portfolio is independent of μ_t and is the solution to the following optimization problem:

$$\begin{aligned} min_{w_t} \quad & w_t' \Sigma_t w_t \\ s.t. \quad & w_t' \iota = 1 \end{aligned} \tag{12}$$

where Σ is the covariance matrix and ι is a unitary vector of length p. The GMV weights are given as:

$$w_{t,GMV} = \frac{\Sigma_t^{-1}\iota}{\iota' \Sigma_t^{-1} \iota}. \tag{13}$$

At the end of day t we generate several covariance estimates using either historical open-to-close returns or day t intraday returns. We use these covariance estimates to determine GMV portfolio weights. We then compare the covariance estimators on the basis of the volatility of the resulting time series of portfolio returns. We consider both in-sample precision of the estimators and out-of-sample forecast quality.

3.2 Estimation and Forecasting

We present both in-sample volatility timing results and out-of sample forecasted volatility timing results. For the in-sample experiments, the time series of realized portfolio returns are defined based on the log difference of the open price to close price of day t. As such, the returns are known and only the covariance structure is left to estimation. This allows us to focus on the incremental benefits of covariance conditioning methods, independent of forecast models. In the out-of-sample experiments, one step ahead returns are defined as the log difference between the day $t + 1$ closing and opening prices. We distinguish between these two settings because the out-of-sample results rely on both the precision of the estimate and the robustness of the forecast model. Recognizing forecast model specification as an essential part of the problem, in the spirit of Briner and Connor (2008) we seek to separate estimation error from forecast model specification error.

We restrict our volatility timing positions to open-to-close returns. This is done with the intention of avoiding overnight returns. Realized covariance estimation is designed to exploit the richness of intraday data and overnight returns run counter to this idea. Works such as Hansen and Lunde (2005) and Gallo (2001) address overnight returns by comparing a number of additive and scaling models. Our study is focused on the estimation of realized covariance, and including overnight returns introduces another source of model specfication error.

3.2.1 Forecasting Realized Betas

Motivated by previous investigations (Andersen, Bollerslev, Diebold, and Wu (2006), Morana (2007), Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008b)) showing that realized beta is persistent, we consider forecasting realized betas. Andersen, Bollerslev, Diebold, and Wu (2006)assess the dynamics and predictability of realized betas and conclude that although displaying less persistence than realized variances or covariances, realized betas can be modeled well as stationary I(0) processes. Similar to Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008b), we consider the performance of a simple ARMA(1,1) forecast of the realized betas and realized variances. The motivations for this specification are three-fold. First, the factor structure is a parsimonious representation of the covariance structure. Second, the resulting covariance matrix forecasts are guaranteed to be positive definite. Finally, the ARMA framework allows us to capture the persistence of realized variances and realized betas.

3.3 Mincer-Zarnowitz Evaluation

Following Briner and Connor (2008), we employ the Mincer and Zarnowitz (1969) forecast evaluation framework to test the performance of the covariance estimators via the resulting portfolio volatility forecasts. (See) In order to give equal weight to the accuracy of each element of the covariance matrix, portfolio volatility is assessed for an equally-weighted portfolio. $\widehat{\sigma}_p^2 = w'\widehat{\Sigma}w$, where w is a vector of equally-weighted positions. With this approach, select assets - or elements of the estimated covariance matrix - do not dominate the analysis.

Specifically, for each covariance forecast we regress a proxy for ex post volatility of the equally-weighted portfolio on an intercept and the candidate portfolio volatility forecast:

$$\widehat{\sigma}_p(t) = b_0 + b_1\widetilde{\sigma}_p(t-1) + \varepsilon(t).$$

Under this specification, $\widehat{\sigma}_p(t)$ is the portfolio volatility proxy at time t, and $\widetilde{\sigma}_p(t-1)$ is the candidate forecast of portfolio volatility. The null hypothesis for this test is $H_0 : b_0 = 0$ and $H_0 : b_1 = 1$.

Following the recommendations in Hansen and Lunde (2006) we perform the Mincer-Zarnowitz regressions directly on the level instead of considering log-volatility. In simulation and empirical analysis, Hansen and Lunde (2006) find that this formulation is more robust and provides a consistent ranking of volatility forecasting models. The Mincer-Zarnowitz regression has heteroskedastic error terms when applied to variances, leading us to focus on volatility. We report robust Newey West standard errors for our parameter estimates to account for autocorrelated errors.

4 Empirical Analysis

We consider volatility timing strategies using daily covariance estimates based on low-frequency (open-to-close) returns and sub-sampled five-minute returns. For each class of estimator, in addition to the fully estimated covariance matrices we also consider shrinkage estimators, single index estimators, and ad hoc approaches to matrix conditioning. We assess the performance of the estimators by considering the volatility of the time series of open-to-close returns of the dynamic global minimum variance (GMV) portfolios based on each of the estimated covariances. We consider both in-sample precision and forecast quality. Finally, we consider the effect of portfolio dimensions on both precision and forecast quality of the different estimators.

4.1 Data: Dow 30 Stocks

We estimate the time-varying covariance structure of the stocks in the Dow Jones Industrial Average (DJIA) over the period from January 1, 2002 to December 31, 2006. We consider several realized covariance estimators based on intraday returns and we consider estimators based on open-to-close returns. For both return frequencies, we consider a single factor model with the S&P 500 Index as the sole factor. We use the SPDR S&P 500 exchange traded fund (ETF) (SPY) as our proxy for the S&P 500 index. This ETF holds a market-value-weighted portfolio of the equities that comprise the S&P 500 Index. We use quote data from the TAQ database for intraday price observations. We collect quotes from only the primary exchange for each security and filter the observations before filling five-minute log-price grids.[1] To facilitate sub-sampling across the calendar-time returns, we generate five five-minute log-price sample grids with the starting time for each grid shifted forward one minute. Returns are calculated as the log-price-difference of the midpoints of the quotes that are closest to, but not past, the grid endpoints. Open-to-close returns are calculated from the first and last grid log prices.

In Table 1 we provide volatility summary statistics based on high-frequency (HF) and low-frequency (LF) returns for the stocks in the DJIA and for the SPY over our sample period. The estimates based on high-frequency data calculate daily realized variance as the sum of squared five-minute returns, averaged over the five sub-samples. The estimates based on low-frequency data are squared open-to-close returns. For both estimators, the table reports annualized average volatility (Vol.), annualized volatility of volatility (VV), and the lag one autocorrelation of the volatility estimates (Auto). While the realized volatility estimates are consistently larger than the estimates based on low-frequency returns, the volatility estimates using intraday returns are much less volatile, and display much greater first-order autocorrelation, than the estimates based on only open-to-close returns.

4.2 Model Parameter Estimates

To implement the shrinkage and single-factor covariance estimators we must determine the shrinkage parameter for the former and fit the single-index model to the data each

[1] We filter the quote data as follows: eliminate quotes 1) not from primary exchange; 2) with a time stamp outside the 9:30 a.m. and 4:00 p.m. window; 3) with bid or offer prices less than or equal to zero; 4) with TAQ-identified errors (Mode equals 4, 7, 9, 11, 13, 14, 15, 19, 20, 27, or 28); 5) not matched to a trade using the Lee and Ready (1991) algorithm with a 1 second lag on reported trades Henker and Wang (2006); 6)all duplicate quote records (same time stamp and same bid and ask prices); 6) reflecting a 10% move from the previous quote midpoint; 7) with a spread greater than 10% of the midpoint price; and 8)"redundant" quotes that reflect no revision to the bid or ask from the most recent quote.

day for the latter. Recall that we will use the Risk Metrics industry standard decay rate of .94 for the EWMA rolling estimator and for the Ledoit Wolf estimator we estimate the shrinkage parameter for each day in our sample based on a lagged one-year rolling window. Across the sample period we obtain a mean alpha of 0.4636, with a minimum estimate of 0.2521 and maximum of 0.8338. For simplicity, we estimate the Ledoit Wolf covariance matrix using the mean alpha for the entire sample.

In Table 2 we provide average coefficient estimates and the average R^2 for each stock for estimates based on high frequency data and those based on low frequency data. The high-frequency betas are estimated daily with sub-sampled five-minute returns and the low-frequency betas are estimated each day using the past 250 open-to-close returns. Using high-frequency data, the average coefficient estimate is 0.85 and the average R^2 is only 0.24. Estimated high-frequency β ranges from 0.61 for JNJ to 1.41 for INTC and the R^2 ranges from 0.08 for HP to 0.36 for C and INTC. Using open-to-close returns, the average beta is much higher at 0.96 and the average R^2 is marginally higher at 0.35. This is consistent with the estimates presented in Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008b).

4.3 Estimation Precision Results

In order to first focus on the precision of our estimators, we implement the GMV volatility timing strategies with the portfolio weights determined in sample. The results of this analysis are provided in Table 3. We present the standard deviation of the time series of open-to-close returns that result from rebalancing the GMV portfolio each day t based on the day t covariance matrix estimate. We include the 0.025 and 0.975 bootstrapped confidence interval, the maximum daily loss, as well as the minimum and maximum portfolio weights, and the median minimum and median maximum weights across all days in the sample. The maximum loss and the weight characteristics allow us to monitor the presence of ill-conditioned covariance estimates that result in extreme portfolio weights.

In addition to the results for the dynamic portfolios, we present the volatility of open-to-close returns over the estimation period for the equally-weighted portfolio (rebalanced daily).

In Panel B we present the results for dynamic GMV strategies based on covariance estimated using historic open-to-close returns up to and including the day t return. OC_{250} uses the past 250 days of open-to-close returns. This estimator captures conventional covariance estimation techniques and provides the benchmark for comparison for the high frequency estimators. OC_{250}^{W+} uses the OC_{250} covariance estimates, but adds the ad-hoc non-negativity constraint on the GMV portfolio weights. OC_{250}^{SF} is a single factor model applied to open-to-close returns. The results for both of these estimators indicate that little relevant information is lost to the constraints but, there is also no evidence that the conditioning of the sample estimate is needed.

OC_{78}, using only 78 days of historical returns, provides an indication of performance when the number of observations used for the low-frequency estimator equals the number of observations used for the realized covariance estimators. Capturing more local information than the OC_{250}, OC_{78} provides a substantially lower GMV portfolio volatility. OC_{78} does result in more extreme weights than OC_{250} but, at least in sample, these weights appear to be based on useful information about the structure of open-to-close returns. Finally, we include OC_{30}, using only the most recent 30 days of returns, to provide an indication of performance when the number of observations equals the number of assets. Note, OC_{30} does yield a positive definite covariance matrix, but the resulting minimum and maximum portfolio weights indicate that this estimator is ill-conditioned. These extreme weights lead to very high portfolio return volatility. In this case the value of more local information is overwhelmed by estimation error. This estimation error results computational multicollinearity.

In Panel C we present the results for dynamic GMV strategies based on covariance estimates using day t five-minute returns. All of the realized covariance estimates are the average of five sub-sampled estimates, each made using 78 five-minute returns. RC_{78s} is the fully estimated, sub-sampled realized covariance. In sample, RC_{78s} provides the dynamic GMV portfolio with the lowest volatility. Compared to the low-frequency estimators, the results for RC_{78s} look closest to those for OC_{78} but with somewhat less extreme portfolio weights perhaps due to the conditioning provided by sub-sampling. Note, in this case, all of the information is local as only day t intraday returns are used. The performance of the estimator, however, is based on the time series of squared open-to-close returns earned by the dynamic portfolio. It appears that the realized covariance provides better information regarding same day open-to-close returns when compared against an estimator using even a relatively short history of open-to-close returns.

RC_{78s}^{W+} uses RC_{78s} but imposes the non-negative weights constraint. While the RC_{78s}^{W+} is well conditioned, important information is not exploited in determining portfolio weights. RC_{78s}^{RM} is the rolling estimator with $\alpha = 0.94$. In sample, RC_{78s}^{RM} yields portfolio volatilities that are substantially higher than those for RC_{78s} and not significantly lower than the long-horizon low-frequency estimators. The Ledoit Wolf estimator (RC_{78s}^{LW}) offers comparatively low GMV volatility, much lower than all but the OC_{78} of the low-frequency estimators but, with extreme weights similar to OC_{250}, RC_{78s}^{LW} appears to be much better conditioned than OC_{78}. RC_{78s}^{LW} evidently captures day-to-day covariance changes but at the same time preserves the well-conditioned properties of the underlying covariance estimate.

In Panel D we provide the results for the realized covariance single-factor model. We consider both the day t estimate, RC_{78s}^{SF} and the rolling estimate variant, $\text{RC}_{78s}^{SF,RM}$. In sample, the single-factor model generates GMV portfolios that are somewhat more volatile than those of the full covariance matrix, RC_{78s}, but much less volatile than OC_{250} portfolios and in-line with OC_{78}, the best performing low-frequency model. RC_{78s}^{SF} appears to be better conditioned than OC_{78} and the less extreme weight characteristics indicate that it is also better conditioned than RC_{78s}. As with RC_{78s}, smoothing RC_{78s}^{SF}

with $\alpha = 0.94$ substantially decreases the in-sample precision of the estimator.

Finally, to provide an indication of the importance of estimation of the covariance elements as opposed to the variances, we consider RC_{78s}^{ZC}, with all diagonal elements set to zero, an extreme form of shrinkage. In sample, RC_{78s}^{ZC} volatility is 70% larger than RC_{78s}, out-performing only the very unstable OC_{30}. This indicates that the covariance estimates, even with the single-factor estimates, contain important information. RC_{78s}^{ZC} is well-conditioned by design but introduces severe estimation error.

4.4 Volatility Timing Forecast Quality Results

In Table 4 we provide the out-of-sample performance characteristics of dynamic GMV portfolios using a series of covariance matrix forecasts. We evaluate the one-step ahead performance of the forecasts based on the volatility of the open-to-close returns of the optimal portfolios. We report the out-of-sample portfolio volatility along with the 0.25 and 0.975 bootstrapped confidence intervals. The results for the forecasted covariance provide an interesting contrast to the estimation results. Consider first the results for forecasts using low-frequencies returns, presented in Panel A. The three long-horizon estimators perform quite similarly. Note that the out-of-sample volatility for the OC_{250} model is only 12% greater than the corresponding in-sample volatility reflecting the 99% overlap in the information set for the day $t+1$ covariance estimate and the day t forecast for day $t+1$. While both the nonnegativity constraint and the single factor approach to conditioning somewhat decreased the performance of the OC_{250} estimator, the one-step ahead forecasting performance is somewhat improved. OC_{78}, the best performing low-frequency estimator in sample, resulted in a portfolio volatility 60% higher than the in-sample volatility. The OC_{30} covariance, already shown to be ill-conditioned in sample, led to exploding portfolio volatility out of sample.

Turning to the results for forecasts using high-frequency returns, presented in Panel B, we see that RC_{Sub} performs slightly worse than OC_{250}, as the precision gains reflected in sample are lost one-step ahead. Based only on information from day t, RC_{Sub} provides a noisy forecast of the day $t+1$ covariance matrix compared to forecasts based on the long-horizon, low-frequency estimators. The conditioned RC_{Sub}^{RM}, however, exhibits much better forecasting performance, yielding a portfolio volatility that is only slightly higher in sample than out of sample. The ad-hoc non-negatively constraint does not have much affect on the forecast performance of RC_{Sub} and, in this application, the Ledoit Wolf conditioning approach is dominated by the exponentially-weighted rolling estimator.

The results for the parsimonious single factor model, reported in Panel C indicates the loss of precision relative to RC_{Sub} revealed in sample is offset out of sample by the advantages of better conditioning offered by the factor model. Apply exponential smoothing to a time series of these estimates and forecast performance is improved even further to the level of the smoothed full-estimated covariance forecast, RC_{78s}^{RM}.

We also consider a less naive forecast method, fitting an ARMA(1,1) model to forecast the realized variances and betas. This approach yields portfolio volatility lower than RC_{78s}^{SI} but somewhat higher than the simpler rolling estimate. To provide a check on the importance of the covariance terms in the performance of the forecasts, we again consider the zero-correlation matrix RC_{78s}^{ZC} . As was the case in sample, the out-of-sample zero-correlation forecast does not perform well, yielding portfolio volatility substantially greater than found for the single factor model.

4.5 Mincer-Zarnowitz Forecast Evaluation

In addition to the volatility timing experiment, we also evaluate the quality of the covariance forecasts using traditional Mincer-Zarnowitz regressions. Recall that we evaluate the forecasts for an equally-weighted portfolio to avoid distortions in the forecast quality assessment due to GMV asset weights. The coefficients and R^2 statistics are reported in Tables 5 and 6. Consider first Table 5 where the proxy for the unobservable ex post portfolio volatility is the sub-sampled realized volatility for day t. The statistical forecast quality results confirm the inferences from the volatility timing analysis. Consistent with existing literature, the smoothed high-frequency forecasts exhibit less bias and less noise than the low-frequency estimators and the raw realized covariance matrix. However, the analysis also indicates that the single factor model provides forecast quality equivalent to the fully-estimated realized covariance matrix. With insignificant bias, a slope coefficient insignificantly different from one, and an R^2 of 0.60, $\mathrm{RC}_{Sub}^{SF,RM}$ provides excellent forecast quality. In terms of goodness of fit, incorporating ARMA forcasts offers some improvement ($\mathrm{R}^2 = 0.66$) but this comes with increased bias.

We also evaluate the forecasts using the squared open-to-close returns as the proxy for ex post portfolio volatility. These results are reported in Table 6. Consistent with the literature, forecast quality is substantially diminished, across the board, for this target volatility. This is not surprising given the volatility of the volatility reported for the individual stocks in Table 1: the target itself is very noisy. Nonetheless, the smoothed high-frequency estimators again provide generally higher forecast quality with the lowest bias and the best fit. The margin of improvement over the low-frequency estimators, however, is greatly diminished evaluated for this target.

Taken together these results indicate that realized covariance-based forecasts provide greater predictability than low-frequency forecasts even when one is restricted to making low-frequency investment decisions. Moreover we have shown that the forecast improvements associated with realized covariance translate into forecast improvements for the resulting portfolios.

4.6 Portfolio Dimension Analysis

To further investigate the effects of ill-conditioning on our covariance estimations and forecasts, we consider the GMV volatility timing strategies across portfolios with an increasing number of assets. We generate the portfolios by averaging a large number of randomly sampled subsets, thereby avoiding stock selection bias at lower dimensions. The results, presented in Table 7 and Table 8, highlight the importance of dimensionality in the covariance estimation and forecasting problem. In Panel A of both tables, we report the average volatility for equally-weighted portfolios of the sampled assets. This provides a benchmark for the pure diversification effect with volatility falling as the number of assets increase, independent of any information on the structure of asset returns. For the set of assets we study, the diversification effect is strong from three to ten assets with a 17% reduction in volatility over this range. The effect tapers off with only a five percent additional improvement as assets increase from ten to the full portfolio of thirty assets.

Consider the in-sample estimation results presented in Table 7. With the exception of OC_{30}, the estimators all provide substantial improvement on the naive benchmark at all dimensions. Furthermore, with the same exception, the results indicate a steady reduction in the volatility of returns to the GMV strategies and a steadily increasing performance improvement relative to the benchmark. The intuition for the increasing performance gap is that in addition to the diversification effect captured in equally-weighted portfolios, the dynamic portfolios exploit information in the estimated covariance matrix. With an increasing number of assets, precise information is more valuable when there are more degrees of freedom over which to capitalize on the information.

Consider now, the OC_{30}. This short-horizon estimator provides large performance advantages over the longer horizon OC_{250} and OC_{78} estimators for dimensions 3 through 20. At fifteen assets, the volatility of the OC_{30} portfolio is 300 basis points or nearly 30% less than the volatility for OC_{250}. This highlights the importance of using a short horizon to capture relevant information that can be lost when using an excessively long horizon. This performance advantage, however, dissipates as the number of assets increases relative to the number of observations. Looking across the dimensions, we see that the volatility of the OC_{30} portfolio is 100 basis points higher at 25 than at 10 assets as poor matrix conditioning out-weighs the diversification effect. At 30 assets, the OC_{30} estimator, close to singular at this dimension, performs very poorly. In contrast, OC_{78} offers less performance gains at lower dimensions than OC_{30} but OC_{78} volatility continues to fall through 30 assets.

As expected, the realized covariance estimators dominate the long-horizon, low-frequency estimators at all horizons. RC_{78s}^{SF} performs comparable to OC_{78} across the dimensions and RC_{78s} reduces portfolio volatility by about 100 basis over OC_{78} at all dimensions.

Recognizing that the in-sample analysis is vulnerable to overfitting the data, we forecast one-step ahead to examine the robustness of our results. Estimators which overfit to our variance proxy will display limited persistence and poor out of sample results. Indeed we see that OC_{78} now fails to outperform OC_{250} at any dimension, an indication that OC_{78} is overfit in sample and is of limited practical use.

Considering the high frequency estimators, the sub-sampled variant, RC_{78s}, provides lower volatility of portfolio returns than the RC_{78} for the entire set of dimensions considered. This suggests the RC_{78} accumulates estimation error at dimensions 20 and 30, and that sub-sampling offers a more consistent estimator by averaging across a number of covariance estimates each day. The performance of RC_{78s}^{SF}, constructed using elements of RC_{78}, shows how conditioning strategies play a greater role as the dimension increases. For low dimensions, 3 to 10, RC_{78s}^{SF} produces results similar RC_{78}, and from 15 to 30 the RC_{78s}^{SF} provides much lower variance than the fully-estmated model.

We also consider the realized covariance sampled every 30 minutes, denoted as RC_{13}, as this results in 13 returns as well as the sub-sampled counterpart RC_{13s}. We introduce these estimators to show the conditioning benefits of sub-sampling. We can see that RC_{13} has comparatively large portfolio variance for low dimensions and is not positive definite over 10 dimensions. RC_{13s} displays large portfolio variance, but is positive definite for all the dimensions considered. In contrast, the RC_{13s}^{SF} estimator shows portfolio variance reduction as dimensions increase. This suggests that the single factor framework can correct for an surprisingly large amount of estimation error.

Finally, we consider Risk Metrics modeling of the high frequency estimators. As anticipated, the estimators with Risk Metrics show a minimum portfolio variance, and outperform the low frequency counterparts. This is takes advantage of the well known long memory properties of realized covariance estimators. Of particular interest, $RC_{78s}^{SF,RM}$ performs as well as RC_{78s}^{RM} once again demonstrating the appeal of the single factor model.

5 Future Work and Conclusion

Our paper contributes to the realized covariance literature by examining three features of interest to both academics and practitioners. We examine the performance of realized covariance estimators as the dimension increases from 3 to 30 and show the importance of conditioning the estimates at higher dimensions. We show that at high dimensions, computationally simple estimators offer portfolio volatilities that are similar to the best performing fully estimated covariance matrix. Finally, we examine in-sample performance to disentangle precision from forecasting. Both estimation and forecasting present sources of error, and we argue that by only looking at the final forecast results, the precision of an estimator can be overshadowed by the errors of a naive forecast.

As future work we will consider more sophisticated one-step ahead forecasts using the single factor estimators. The current results using the ARMA(1,1) model are encouraging and suggest that this is a sensible starting point for high dimensional covariance forecasting. An attractive feature of the single-factor model is that it is straightforward to apply, as there is no intermediate step requiring estimation of a smoothing factor. Furthermore, it circumvents the problem of ensuring positive definite forecasts. Our empirical analysis has focused primarily on open-to-close returns. In future work we will include overnight returns according to techniques suggested by Gallo (2001) and Hansen and Lunde (2005). This will allow for the replication of more practitioner oriented trading strategies and will allow us to examine allocation strategies in the presence of transaction costs.

We have compared a number of conditioning techniques for realized covariance and the characteristics of the GMV portfolios they generate. Using a straight-forward estimation approach we confirm that improved forecast quality of realized variance and covariances translates to better forecast qaulity in portfolios. The results are encouraging and suggest that single-factor models can offer similar results for even higher dimension portfolios.

References

Andersen, T., Bollerslev, T., Diebold, F., and Labys, P. (2001), "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association*, 96, 42–55.

— (2003), "Modeling and Forecasting Realized Volatility," *Econometrica*, 71, 579–625.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Wu, G. (2006), "Realized Beta: Persistence and Predictability," *Advances in Econometrics*, 20, 1–39.

Andreou, E. and Ghysels, E. (2002), "Rolling-Sample Volatility Estimators: Some New Theoretical, Simulation, and Empirical Results," *Journal of Business and Economic Statistics*, 363–376.

Bandi, F. and Russell, J. (2006), "Separating Microstructure Noise from Volatility," *Journal of Financial Economics*, 79, 655–692.

Bandi, F., Russell, J., and Zhu, Y. (2008), "Using High-Frequency Data in Dynamic Portfolio Choice," *Econometric Reviews*, 27, 163–198, forthcoming.

Barndorff-Nielsen, O., Hansen, P., Lunde, A., and Shephard, N. (2008a), "Designing Realized kernels to measure the ex-post variation of equity prices in the presence of noise," *Econometrica*, to appear.

— (2008b), "Multivariate Realized Kernels: Consistent Positive Semi-Definite Estimators of the Covariation of Equity Prices with noise and non-synchronous trading," Working Paper, University of Oxford.

Barndorff-Nielsen, O. and Shephard, N. (2004), "Econometric Analysis of Realized Covariation: High Frequency Based Covariance, Regression, and Correlation in Financial Economics," *Econometrica*, 72, 885–925.

Bollerslev, T. and Zhang, B. Y. B. (2003), "Measuring and modeling systematic risk in factor pricing models using high-frequency data," *Journal of Empirical Finance*, 10, 533–558.

Briner, B. and Connor, G. (2008), "How much structure is best? A comparison of market model, factor model and unstructured equity covariance matrces," *The Journal of Risk*, 10, 3–30.

Britten-Jones, M. (1999), "The Sampling Error in Estimates of Mean-Variance Efficient Portfolio Weights," *Journal of Finance*, 54, 655–671.

Chan, L., Karceski, J., and Lakonishok, J. (1999), "On Portfolio Optimization: Forecasting Covarainces and Choosing the Risk Model," *The Review of Financial Studies*, 12, 937–974.

Corsi, F. (2006), "Realized Correlation Tick-by-Tick," Working Paper University of Lugano.

de Pooter, M., Martens, M., and van Dijk, D. (2008), "Predicting the Daily Covariance Matrix of S&P100 Stocks Using Intraday Data - but which frequency to use?" *Econometric Reviews*, 27, 199–229, forthcoming.

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

Fleming, J., Kirby, C., and Ostdiek, B. (2003), "The Economic Value of Volatility Timing Using "Realized" Volatility," *Journal of Financial Economics*, 473–509.

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

Gallo, G. (2001), "Modelling the Impact of Overnight Surprises on Intra-Daily Volatilty," *Australian Ecnonomic Papers*, 40, 567–580.

Griffin, J. and Oomen, R. (2006), "Covariance Measurement in the Presence of Non-synchronous Trading and Market Microstructure Noise," Working Paper University of Warwick.

Hansen, P. and Lunde, A. (2006), "Consistent ranking of volatility models," *Journal of Econometrics*, 131, 97–121.

Hansen, P. R. and Lunde, A. (2005), "A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data," *Journal of Financial Econometrics*, 3, 525–554.

Hayashi, T. and Yoshida, N. (2005), "On Covariance Estimation of Non-synchronously Observed Diffusion Processes," *Bernoulli*, 11, 359–379.

Henker, T. and Wang, J.-X. (2006), "On the Importance of Timing Specifications in Market Microstructure Research," *Journal of Financial Markets*, 9, 162–179.

Jagannathan, R. and Ma, T. (2003), "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps," *Journal of Finance*, 58, 1651–1683.

Ledoit, O. and Wolf, M. (2003), "Improved estimation of the covariance matrix of stock returns with an application to portfolio selection," *Journal of Empirical Finance*, 10, 603–621.

— (2004a), "Honey, I Shrunk the Sample Covaraince Matrix," *The Journal of Portfolio Management*, 30, 110–119.

— (2004b), "A well-conditioned estimator for large-dimensional covariance matrices," *Journal of Multivariate Analysis*, 88, 365–411.

Lee, C. and Ready, M. (1991), "Inferring Trade Direction from Intraday Data," *Journal of Finance*, 46, 733–747.

Liu, Q. (2008), "On Portfolio Optimization: How Do We Benefit from High-Frequency Data?" University of Hawaii Working Paper.

McAleer, M. and Medeiros, M. (2008), "Realized Volatility: A Review," *Econometric Reviews*, 27, 10–45.

Michaud, R. O. (1989), "The Markowitz Optimization Enigma: Is Optimized Optimal," *Financial Analysts Journal*, 45, 31–42.

Mincer, J. and Zarnowitz, V. (1969), "The Evaluation of Economic Forecasts," in *Economic Forecasts and Expectations*, ed. Mincer, J., Columbia University Press, pp. 3–46.

Morana, C. (2007), "Estimating, filtering and forecasting realized betas," *Journal of Financial Forecasting*, 1, 83–111.

Morgan, J. (1996), *RiskMetrics Technical Document*, J.P.Morgan/Reuters.

Oomen, R. (2006), "Properties of Realized Variance Under Alternative Sampling Schemes," *Journal of Business and Economic Statistics*, 24, 219–237.

Voev, V. (2008), "Dynamic Modelling of Large Dimensional Covariance Matrices," in *High Frequency Financial Economics*, eds. Bauwens, L., Pohlmeier, W., and Veredas, D., Springer/ Pysica-Verlag, Studies in Empirical Economics, pp. 293–312.

Voev, V. and Lunde, A. (2007), "Integrated Covariance Estimation Using High-Frequency Data in the Presence of Noise," *Journal of Financial Econometrics*, 5, 68–104.

Zhang, L., Mykland, P., and Ait-Sahalia, Y. (2005), "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data," *Journal of the American Statistical Association*, 1394–1411.

APPENDIX

A Robustness to Shrinkage Parameter Estimation

We suspect that the sampling errors with respect to the estimated shrinkage parameter are rather large, and warranting a robustness analysis. In this section, we show the robustness of the shrinkage methods for different values of of the shrinkage parameter α. Table 9 presents the resulting volatility of portfolio returns using GMV as the investment strategy. We conclude that the variation in results is quite modest with respect to a large range of shrinkage parameters. Recall that our target matrix is the the constant correlation matrix which is somewhat similar to the single-factor estimator. In particular, both of these estimators have the same diagonal elements. We can easily see that the size and frequency of extreme weights are reduced as we smooth. Again, we see that out of sample forecasts call for more smoothing than do in sample estimates.

Table 1: Components of DJIA. Symbols denote the following time periods for inclusion in the index: † April 8, 2004 - Dec 31, 2006, § Jan 1, 2002 - April 7, 2004 and Nov 22, 2005 - Dec 31, 2006, * Jan 1, 2002 - April 7, 2004, and ‡ Jan 1, 2002 - Dec 21, 2005.

	HF Data			LF Data		
Ticker	**Vol.**	**VV**	**Auto**	**Vol.**	**VV**	**Auto**
AA	0.2605	0.5807	0.7529	0.1994	1.8030	0.1221
AIG†	0.2060	0.6920	0.7991	0.1613	1.8456	0.2738
AXP	0.1905	0.8773	0.8668	0.1545	1.6787	0.2278
BA	0.2306	0.6376	0.8127	0.1785	1.5042	0.1279
C	0.1968	0.9986	0.8859	0.1468	1.7066	0.3590
CAT	0.2218	0.3928	0.6628	0.1805	1.5067	0.0370
DD	0.1978	0.4459	0.8046	0.1469	1.1139	0.1413
DIS	0.2350	0.9157	0.8439	0.1808	1.7930	0.2432
EK*	0.2505	0.6440	0.5354	0.1906	2.1283	0.0930
GE	0.1881	0.7550	0.8533	0.1483	1.5106	0.3168
GM	0.2568	0.8672	0.6651	0.2197	2.8880	0.2274
HD	0.2262	0.6948	0.8221	0.1758	1.7319	0.1860
HON	0.2495	0.8259	0.7232	0.1823	2.0110	0.2054
HP	0.2942	0.4791	0.6239	0.2457	2.4342	0.0167
IBM	0.1771	0.4103	0.8404	0.1412	1.1686	0.1923
INTC	0.2793	1.0215	0.8268	0.2336	2.8705	0.2407
IP*	0.2081	0.3346	0.7237	0.1662	1.2301	0.1194
JNJ	0.1624	0.5263	0.8177	0.1183	0.8482	0.2149
JPM	0.2236	1.4882	0.8649	0.1758	3.0401	0.3855
KO	0.1638	0.3374	0.8566	0.1153	0.7880	0.2158
MCD	0.2200	0.6079	0.6733	0.1658	1.5477	0.1776
MMM	0.1701	0.2685	0.7087	0.1301	0.8697	0.0887
MO	0.1859	0.6933	0.5194	0.1401	1.7517	0.2076
MRK	0.2072	0.6574	0.5899	0.1582	1.6064	0.1733
MSFT	0.2036	0.6686	0.7625	0.1593	1.6832	0.3524
PFE†	0.2042	0.5405	0.6941	0.1565	1.3892	0.1546
PG	0.1507	0.1916	0.7559	0.1086	0.6026	0.1115
SBC‡	0.2534	1.4199	0.8685	0.1879	2.2862	0.2306
T§	0.2338	0.9575	0.7708	0.1933	2.5216	0.3266
UTX	0.1967	0.3999	0.7713	0.1526	1.3021	0.2027
VZ†	0.2147	0.8139	0.8431	0.1661	1.6047	0.3361
WMT	0.1829	0.4177	0.8257	0.1374	0.9706	0.1995
XOM	0.1878	0.4020	0.8240	0.1477	1.0665	0.1387
Mean	0.2130	0.6656	0.7633	0.1656	1.6607	0.2014
Median	0.2077	0.6507	0.7880	0.1625	1.6056	0.2020
SPY	0.1221	0.2523	0.8675	0.1036	0.6314	0.2048

Table 2: S&P 500 Factor Model Estimates with Sub-Sampled 5-minute Returns and 250-Day Rolling OC Returns

	HF Data		LF Data	
Ticker	$\boldsymbol{\beta}$	R^2	$\boldsymbol{\beta}$	R^2
AA	0.98	0.19	1.22	0.37
AIG	0.87	0.26	1.08	0.40
AXP	0.83	0.30	1.06	0.50
BA	0.91	0.22	0.98	0.31
C	0.95	0.36	1.03	0.54
CAT	1.00	0.27	1.22	0.43
DD	0.89	0.28	0.95	0.43
DIS	0.87	0.20	1.01	0.34
EK	0.72	0.12	0.98	0.22
GE	0.91	0.35	0.99	0.52
GM	0.80	0.16	1.12	0.28
HD	0.96	0.25	1.14	0.42
HON	1.04	0.25	1.12	0.39
HP	0.77	0.08	0.97	0.14
IBM	0.85	0.33	0.88	0.42
INTC	1.41	0.36	1.47	0.44
IP	0.83	0.22	1.06	0.39
JNJ	0.61	0.22	0.64	0.27
JPM	0.97	0.30	1.23	0.55
KO	0.66	0.23	0.64	0.30
MCD	0.79	0.18	0.80	0.19
MMM	0.79	0.29	0.83	0.37
MO	0.64	0.17	0.64	0.14
MRK	0.73	0.18	0.79	0.22
MSFT	0.98	0.33	0.98	0.43
PFE	0.79	0.22	0.86	0.29
PG	0.65	0.25	0.63	0.30
SBC	0.90	0.23	0.91	0.33
T	0.70	0.14	0.83	0.21
UTX	0.87	0.27	0.96	0.41
VZ	0.82	0.22	0.84	0.31
WMT	0.83	0.30	0.81	0.36
XOM	0.86	0.29	0.94	0.40
Mean	0.85	0.24	0.96	0.35
Median	0.85	0.25	0.97	0.37

Table 3: Volatility of Open-to-Close Returns of Global Minimum Variance (GMV) Portfolios 2003-2006, in sample

	Portfolio Volatility			Max	Min. Weight		Max. Weight	
Model	$\boldsymbol{\sigma_P}$	**0.025**	**0.975**	**Loss**	**Min**	**Median**	**Max**	**Median**
Panel A: Benchmark Portfolios								
EW	0.1107	0.1042	0.1177	-0.0266				
Panel B: LF Data								
OC_{250}	0.0855	0.0816	0.0906	-0.0185	-0.1837	-0.0963	0.4913	0.2365
OC_{250}^{W+}	0.0887	0.0848	0.0938	-0.0185	0.0000	0.0000	0.5698	0.2269
OC_{250}^{SF}	0.0871	0.0831	0.0922	-0.0187	-0.1293	-0.0858	0.4609	0.2409
OC_{78}	0.0666	0.0637	0.0700	-0.0130	-0.4096	-0.1660	0.8639	0.3400
OC_{30}	0.3233	0.2271	0.4373	-0.3292	-130.6358	-1.2014	98.4752	1.3985
Panel C: HF Data - Full Covariance Matrix								
RC_{78s}	0.0576	0.0534	0.0605	-0.0137	-0.3770	-0.1131	0.7869	0.2910
RC_{78s}^{W+}	0.0730	0.0692	0.0771	-0.0177	0.0000	0.0000	0.7319	0.2621
RC_{78s}^{RM}	0.0812	0.0771	0.0855	-0.0153	-0.1210	-0.0797	0.4584	0.2199
RC_{78s}^{LW}	0.0629	0.0598	0.0662	-0.0173	-0.1439	-0.0590	0.7663	0.2745
Panel D: HF Data - Single Factor Model								
RC_{78s}^{SF}	0.0687	0.0651	0.0722	-0.0147	-0.2003	-0.0673	0.6687	0.2312
$RC_{78s}^{SF,RM}$	0.0825	0.0782	0.0860	-0.0180	-0.1229	-0.0811	0.4217	0.2090
RC_{78s}^{ZC}	0.0977	0.0912	0.1046	-0.0239	0.0004	0.0077	0.2762	0.0837

Table 4: Volatility of Open-to-Close Returns of Global Minimum Variance (GMV) Portfolios 2003-2006, One-Step-Ahead Forecasts

	Portfolio Volatility			**Max**
Model	$\boldsymbol{\sigma_P}$	**0.025**	**0.975**	**Loss**
Panel A: LF Data				
OC_{250}	0.0961	0.0913	0.1013	-0.0210
OC_{250}^{W+}	0.0933	0.0882	0.0979	-0.0189
OC_{250}^{SF}	0.0909	0.0871	0.0953	-0.0192
OC_{78}	0.1074	0.1011	0.1131	-0.0285
OC_{30}	1.4709	0.8648	2.2022	-2.0368
Panel B: HF Data - Full Matrix				
RC_{78s}	0.0987	0.0930	0.1048	-0.0227
RC_{78s}^{RM}	0.0849	0.0809	0.0889	-0.0166
RC_{78s}^{W+}	0.0980	0.0927	0.1032	-0.0224
RC_{78s}^{LW}	0.0933	0.0895	0.0990	-0.0199
Panel C: HF Data - Single Factor				
RC_{78s}^{SF}	0.0922	0.0882	0.0967	-0.0197
$\text{RC}_{78s}^{SF,RM}$	0.0850	0.0801	0.0880	-0.0182
$\text{RC}_{78s}^{SF,ARMA}$	0.0872	0.0820	0.0933	-0.0366
RC_{78s}^{ZC}	0.1031	0.0980	0.1111	-0.0249

Table 5: Mincer-Zarnowitz Forecast Quality Regressions: Realized Volatility Proxy

Model	$\boldsymbol{b_0}$	**s.e.**	$\boldsymbol{b_1}$	**s.e.**	R^2
Panel A: LF Data					
OC_{250}	0.002708	0.000387	0.430082	0.052002	0.45
OC_{250}^{SF}	0.002662	0.000392	0.436030	0.052655	0.45
OC_{78}	0.001663	0.000379	0.643869	0.055473	0.52
OC_{30}	0.001494	0.000350	0.699039	0.056515	0.54
Panel B: HF Data - Full Matrix					
RC_{78s}	0.001522	0.000220	0.754240	0.039952	0.57
RC_{78s}^{RM}	0.000180	0.000323	0.933463	0.052788	0.60
RC_{78s}^{LW}	0.001460	0.000224	0.754383	0.040152	0.57
Panel C: HF Data - Single Factor					
RC_{78s}^{SF}	0.001633	0.000225	0.767534	0.042758	0.57
$RC_{78s}^{SF,RM}$	0.000261	0.000320	0.958125	0.054590	0.60
$RC_{78s}^{SF,ARMA}$	0.000663	0.000172	0.883619	0.028845	0.66

Table 6: Mincer-Zarnowitz Forecast Quality Regressions: Squared Open-to-Close Returns Volatility Proxy

Model	$\boldsymbol{b_0}$	**s.e.**	$\boldsymbol{b_1}$	**s.e.**	$\mathrm{R}^{\mathbf{2}}$
Panel A: LF Data					
OC_{250}	0.001949	0.000470	0.415540	0.062653	0.10
OC_{250}^{SF}	0.001904	0.000476	0.421354	0.063434	0.10
OC_{78}	0.000866	0.000488	0.632566	0.072831	0.12
OC_{30}	0.000903	0.000511	0.656462	0.081593	0.12
Panel B: HF Data - Full Matrix					
RC_{78s}	0.001361	0.000402	0.638799	0.068719	0.10
RC_{78s}^{RM}	-0.000428	0.000589	0.891820	0.097228	0.13
RC_{78s}^{LW}	0.001305	0.000408	0.639411	0.068809	0.10
Panel C: HF Data - Single Factor					
RC_{78s}^{SF}	0.001480	0.000395	0.645825	0.070831	0.10
$\mathrm{RC}_{78s}^{SF,RM}$	-0.000319	0.000583	0.910173	0.100653	0.13
$\mathrm{RC}_{78s}^{SF,ARMA}$	0.000427	0.000468	0.781323	0.078604	0.13

Table 7: Volatility of Open-to-Close Returns of Global Minimum Variance (GMV) Portfolios 2003-2006, at Different Dimensions, in-sample

	Number of Assets						
Model	3	5	10	15	20	25	30
Panel A: Benchmark Portfolio							
EW	0.1407	0.1291	0.1171	0.1149	0.1122	0.1114	0.1107
Panel B: LF Data							
OC_{250}	0.1299	0.1169	0.1023	0.0962	0.0916	0.0877	0.0855
OC_{250}^{SF}	0.1302	0.1173	0.1031	0.0972	0.0925	0.0887	0.0871
OC_{78}	0.1268	0.1122	0.0952	0.0861	0.0784	0.0713	0.0667
OC_{30}	0.1224	0.1055	0.0827	0.0691	0.0798	0.0848	0.3233
Panel C: HF Data							
RC_{78s}	0.1205	0.1058	0.0871	0.0765	0.0687	0.0617	0.0576
RC_{78s}^{SF}	0.1215	0.1079	0.0911	0.0826	0.0766	0.0713	0.0687

Table 8: Volatility of Open-to-Close Returns of Global Minimum Variance (GMV) Portfolios 2003-2006, at Different Dimensions, Naive Forecasts

	Number of Assets						
Model	3	5	10	15	20	25	30
Panel A: Benchmark Portfolio							
EW	0.1407	0.1291	0.1171	0.1149	0.1122	0.1114	0.1107
Panel B: LF Data							
OC_{250}	0.1315	0.1188	0.1071	0.1005	0.0986	0.0960	0.0961
OC_{250}^{SF}	0.1316	0.1187	0.1063	0.0989	0.0958	0.0922	0.0909
OC_{78}	0.1314	0.1197	0.1093	0.1047	0.1046	0.1043	0.1074
OC_{30}	0.1334	0.1242	0.1218	0.1293	0.4889	0.3750	1.4709
Panel C: HF Data							
RC_{78}	0.1325	0.1219	0.1118	0.1100	0.1086	0.1079	0.1091
RC_{78s}	0.1317	0.1205	0.1101	0.1040	0.1026	0.0996	0.0987
RC_{78s}^{RM}	0.1295	0.1160	0.1027	0.0943	0.0908	0.0865	0.0850
RC_{78s}^{SF}	0.1312	0.1195	0.1069	0.1015	0.0968	0.0938	0.0922
$RC_{78s}^{SF,RM}$	0.1293	0.1155	0.1018	0.0932	0.0899	0.0860	0.0847
RC_{13}	0.1463	0.1462	0.2005	npd	npd	npd	npd
RC_{13s}	0.1447	0.1404	0.1497	0.1552	0.1525	0.1488	0.1495
RC_{13s}^{SF}	0.1396	0.1287	0.1173	0.1110	0.1051	0.1014	0.0997

Table 9: Robustness analysis of smoothing parameters; Ledoit Wolf

α	σ		$\min(w)$		$\max(w)$	
	in	out	min	median	max	median
.1	0.0807	0.0964	−0.1087	−0.0394	0.7261	0.2561
.2	0.0791	0.0958	−0.1113	−0.0423	0.7340	0.2600
.3	0.0781	0.0955	−0.1150	−0.0472	0.7398	0.2653
.4	0.0774	0.0956	−0.1277	−0.0541	0.7533	0.2689
.5	0.0772	0.0960	−0.1456	−0.0615	0.7662	0.2748
.6	0.0773	0.0968	−0.1672	−0.0705	0.7783	0.2772
.7	0.0779	0.0983	−0.2882	−0.0817	0.7893	0.2836
.8	0.0805	0.1020	−2.2709	−0.0967	3.3067	0.2924
.9	0.0849	0.1441	−9.3478	−0.1190	6.2236	0.3034