Covariance Estimation in Dynamic Portfolio Optimization: A Realized Single Factor Model*

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

Barbara Ostdiek
Jones Graduate School of Business
Rice University

Katherine Ensor
Department of Statistics
Rice University

First Draft: October 2008
Second Draft: March 2009
Current Draft: June 2009

Abstract

Realized covariance estimation for large dimension problems is little explored and poses challenges in terms of computational burden and estimation error. In a global minimum volatility setting, we investigate the performance of covariance conditioning techniques applied to the realized covariance matrices of the 30 DJIA stocks. We find that not only is matrix conditioning necessary to deliver the benefits of high frequency data, but a single factor model, with a smoothed covariance estimate, outperforms the fully estimated realized covariance in one-step ahead forecasts. Furthermore, a mixed-frequency single-factor model - with factor coefficients estimated using low-frequency data and variances estimated using high-frequency data - performs better than the realized single-factor estimator. The mixed-frequency model is not only parsimonious but it also avoids estimation of high-frequency covariances, an attractive feature for less frequently traded assets. Volatility dimension curves reveal that it is difficult to distinguish among estimators at low portfolio dimensions, but for well-conditioned estimators the performance gain relative to the benchmark $\frac{1}{N}$ portfolio increases with N.

*This paper previously circulated under the title "Realized Covariance Estimation in Dynamic Portfolio Optimization." We are thankful for comments from Brad Paye and from seminar participants at the 2008 Advances in Portfolio Optimization Workshop at Imperial College, London and at HEC, University of Lausanne. We thank Alexander Chaudhry and Scott Baggett for excellent research assistance.

1 Introduction

Realized covariance estimation employs high-frequency asset price data to approximate the integrated covariance of asset returns. The Theory of Quadratic Variation developed in Andersen and Bollerslev (1998) implies that with more frequent observation of the price process the actual (co)variance over a given period of time can be precisely measured. In the decade following this research, the literature has focused on refinements that address practical estimation issues, including market microstructure effects and nonsynchroneity of price observations. Empirical analysis has been largely confined to variance estimation and covariance estimation at small dimensions.

We investigate realized covariance estimation in a 30 asset universe and offer three primary contributions. First, we examine the performance of a realized single factor estimator for a high-dimensional covariance matrix. Second, we consider a mixed-frequency single factor estimator with factor coefficients estimated using low-frequency data and variance elements estimated using high-frequency data. We compare these estimators to low- and high-frequency sample estimators and to alternative shrinkage estimators. We find that forecasts based on the parsimonious realized single-factor estimators match the best – and substantially outperform many – of the alternative forecasts. Finally, we examine the importance of matrix dimensions in the covariance estimation problem. We find that some realized covariance estimators exhibit signs of ill-conditioning at low dimensions and that with a well-conditioned matrix the relative gain to forecast quality increases with the number of dimensions.

In general, covariance estimation precision decreases as the dimensions of the matrix increases relative to the number of observations used in estimation. One solution to reducing estimation error for large matrices is to increase the number of observations. For covariance estimation this requires either extending the estimation window or sampling more intensely within a window. Neither approach is without cost. Extending the window, the standard approach with low-frequency data, risks loss in precision as the estimator incorporates less informative historical observations. Increasing the sampling intensity, made possible with high-frequency data, risks both increased bias and loss in precision resulting from market microstructure effects. Covariance estimates, in particular, are sensitive to the Epps (1979) effect, the downward bias induced by nonsynchroneity of observed prices. The bias increases with increased sampling intensity.

Even with the thousands of intraday transaction and quote observations available for many traded assets, sampling strategies designed to account for market-microstructure effects and nonsynchroneity effects in realized variance and covariance estimates can quickly limit the effective sample size of even ultra-high frequency data. For example, in the case of "classic" five-minute previous-tick realized covariance estimation only 78 observations are used each day. To decrease nonsynchroneity, researchers have considered less frequent intraday sampling. With the 30-minute sampling frequency there are only 13 observations, making ill-conditioning a likely problem for realized covariance estimates at even small dimensions.

Imposing additional structure on the sample covariance matrix to directly condition the estimate can reduce sampling error. Several variants of shrinkage - shrinking the most extreme estimates toward more central values - are used in the literature.[1] Fan, Fan, and Lv (2008) consider factor models as a conditioning approach and demonstrate that this model provides a better conditioned alternative to the fully-estimated covariance matrix of daily returns. We examine the performance of a single-factor estimator for a large-dimensional covariance matrix in a high-frequency setting. Using volatility timing strategies, we compare the realized single-factor covariance estimator based on five-minute returns to alternative high- and low-frequency sample and shrinkage estimators. Using forecasts generated with naive exponential smoothing, the computationally efficient single-factor high-frequency estimator provides better forecast performance than

[1] See Fleming, Kirby, and Ostdiek (2003), de Pooter, Martens, and van Dijk (2008), and Liu (2009) for examples in a high-frequency setting and Ledoit and Wolf (2003), Ledoit and Wolf (2004b), and Jagannathan and Ma (2003) for recent examples in a low-frequency setting.

both low-frequency estimators and alternative high-frequency shrinkage estimators, and matches the performance of a fully-estimated, sub-sampled matrix.

In addition, motivated by the lower persistence of the high-frequency beta estimates relative to the corresponding variance and covariance estimates, we estimate low-frequency factor coefficients in a mixed-frequency single-factor model. This approach reduces sampling error in the covariance elements while preserving the precision offered by the high-frequency variance estimates. The hybrid model provides forecast performance better than the realized single-factor model and comparable to the smoothed single-factor and smoothed fully-estimated models. In addition to the decreased computational burden offered by the factor structure, the mixed-frequency estimator, relying on high-frequency sampling for realized variance estimation only, avoids the more complex, and more restrictive, sampling techniques suggested for realized covariance estimation. This feature may be important for estimating covariances for assets with low liquidity and estimating covariances over non-trading periods.

To directly investigate the importance of matrix dimensions in the covariance estimation problem, we examine the performance of the alternative estimators for portfolios ranging from two to 30 assets. Sparsely sampled realized covariance estimators exhibit signs of ill-conditioning at small dimensions. Even the popular five-minute previous-tick realized covariance estimator exhibits increasing estimation error at 30 assets. Furthermore, we find that with a well-conditioned matrix, the relative gain to forecast quality increases with the number of dimensions. This effect is a function of the scope over which the forecast information can be used: the greater the investment opportunity set, the greater the benefit offered by high-quality forecasts.

The remainder of the paper is organized as follows. Section 2 provides the single factor realized covariance estimators and traditional shrinkage estimators. In Section 3 we present the data and the empirical analysis. Section 4 offers concluding remarks and suggestions for future research.

2 Estimators

Many applied problems in finance require a covariance matrix estimator that is accurate, well-conditioned and, in some cases, invertible. While the true covariance matrix is guaranteed to have all of these characteristics, estimated matrices may not due to measurement and sampling error. Stein (1956) proposed shrinkage estimators to reduce the most extreme estimation error by shrinking the sample matrix eigenvalues towards a more central value. The general linear shrinkage model can be written as

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

where $\widehat{\Sigma}_t$ is the sample estimate of Σ_t, the covariance matrix of returns, and G_t is an idealized covariance structure. To ensure positive definiteness, the target matrix G_t is chosen to be positive definite and the shrinkage factor α_t is chosen to optimize a criteria, such as MSE, which imposes positive definiteness on the resulting estimator $\widetilde{\Sigma}_t$.

If the extreme elements in the covariance matrix are the result of sampling error, shrinkage leads to more precise estimates. If, however, the true covariance element is extreme, shrinkage introduces specification error. Even a miss-specified shrinkage model can yield more accurate estimates when sampling error is large compared to specification error. We consider three shrinkage-related covariance estimators: the single-factor model, the rolling estimator, and the Ledoit-Wolf estimator.

2.1 Single Factor Models

A factor model can be thought of as an extreme form of shrinkage that puts all of the weight on the target matrix and none on the sample estimate; i.e. $\alpha_t = 1 \ \forall \ t$ in (1). Fan, Fan, and Lv (2008) show that when the number of factors is small relative to the matrix dimension, 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 yields a covariance matrix with less sampling error than the sample estimate. Factor models simplify estimation of the covariance matrix by greatly reducing the number of parameters estimated, a particularly attractive feature when estimating large realized covariance matrices. Finally, single-factor covariance estimates are strictly positive definite, guaranteeing invertible matrices.

In a low frequency setting, Chan, Karceski, and Lakonishok (1999) and Jagannathan and Ma (2003) show that factor models can reduce the variance of optimal mean-variance portfolios and, out-of-sample, can out-perform strategies employing full sample covariance matrices. Jagannathan and Ma (2003) find that the single-factor model performs well even when the number of observation is not much greater than the number of dimensions. Bollerslev and Zhang (2003) employ high-frequency data in a multi-factor model and find improved asset pricing predictions when compared with conventional monthly rolling estimates.

We consider a realized single-factor model:

$$r_{at_i} = \beta_{at}^R r_{Mt_i} + \varepsilon_{at_i}, \tag{2}$$

where r_{at_i} is the i^{th}-sampled intra-day return on day t for asset a, r_{Mt_i} is the corresponding intra-day return of the market proxy asset (the factor), and β_{at}^R is the realized market beta of the asset. We assume that the residuals ε_{at_i} are uncorrelated with factor returns and that $\varepsilon_{at_i} \sim N(0, \sigma_{at}^2)$. The corresponding realized covariance matrix for p assets is

$$\Sigma_t^{SF} = \sigma_{Mt}^2 \beta_t^R \beta_t^{R'} + D_t^R, \tag{3}$$

where β_t^R is the $p \times 1$ vector of realized factor loadings, σ_{Mt}^2 is the realized factor variance and D_t^R is the diagonal matrix of residual realized sample variances. The elements of

Σ_t^{SF} (with daily time-subscripts suppressed) are estimated as

$$\widehat{\sigma}_{ab}^{SF} = \begin{cases} (s_M^R)^2 \widehat{\beta}_a^R \widehat{\beta}_a^R + \widehat{d}_{aa}^R & \text{if a=b} \\ (s_M^R)^2 \widehat{\beta}_a^R \widehat{\beta}_b^R & \text{if a} \neq \text{b,} \end{cases} \tag{4}$$

where $\widehat{\beta}_{at}^R$ is the sample realized factor coefficient and the diagonal elements of $\widehat{D}_t^R$ are estimated as $\widehat{d}_{aat}^R = \sum_{i=1}^m (r_{at_i} - \widehat{\beta}_{at}^R r_{Mt_i})^2$.

Applying a single-factor model to high-frequency data is particularly appealing, since this form of shrinkage avoids over-fitting sample realized covariance estimates that may be noisy due to nonsynchroneity and other market microstructure effects. Estimating co-dependence through the market model has the advantage of matching each asset with the active index asset, potentially reducing the market microstructure noise in each factor-implied covariance estimate relative to the noise in each sample realized covariance.

Motivated by further reducing forecast error in the covariance estimates, we consider a mixed-frequency single-factor model with asset betas estimated using low-frequency returns but with market variance and residual variances estimated using high-frequency returns. This formulation is similar in spirit to the MIDAS (mixed data sampling regression models) developed in Ghysels, Santa-Clara, and Valkanov (2006). The elements of the mixed-frequency single-factor covariance matrix are estimated as in (3) and (4) above but with the vector of factor coefficients replaced by a low-frequency estimate, $\widehat{\beta}_t^{LF}$, and the residual variances estimated as $\widehat{d}_{aat}^{MF} = \sum_{i=1}^m (r_{at_i} - \widehat{\beta}_{at}^{LF} r_{Mt_i})^2$. This hybrid estimator combines the advantages of estimating market and residual risk with high-frequency data while avoiding the complexity of optimal realized covariance estimation in the presence of market microstructure noise.

This mixed-frequency single-factor estimator is a specialized case of the mixed-frequency factor model contemporaneously developed in Bannouh, Martens, Oomen, and van Dijk (2009). They present a multi-factor model relying on high-frequency observations to estimate the factor variances and low-frequency observations to estimate factor betas and the idiosyncratic component of asset variances for vast dimensional covariance matrices.[2] In this analysis, we estimate a modest-sized covariance matrix using a single-factor variation of this approach that relies on low-frequency data to estimate the factor coefficients, retaining high-frequency estimation for the factor variance and the idiosyncratic component of asset variances.

We also consider forecasting realized betas. Andersen, Bollerslev, Diebold, and Wu (2006) assess the dynamics and predictability of realized betas and conclude that although they display 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.

2.2 Rolling Estimators

Rolling realized covariance estimation, motivated by the conditional heteroskedasticity of financial time series, attempts to balance the statistical power of a large sample against loss of precision from including stale information.[3] It is easy to see that

$$\Sigma_t^{RM} = \alpha \Sigma_{t-1}^{RM} + (1-\alpha)\widehat{\Sigma}_t, \quad \alpha \in [0,1] \tag{5}$$

is a variant of shrinkage estimation where the current realized covariance estimate is shrunk toward a function of past estimates. The target matrix possesses the desirable properties of the true covariance matrix: positive definite and well-conditioned. 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

[2] In an extensive simulation exercise Bannouh et al. (2009) show that this estimator is more efficient than the sample realized covariance estimator. The estimator also exhibits superior empirical performance in applications in the S&P 500, S&P 400, and S&P 600 stock universes.

[3] See Fleming, Kirby, and Ostdiek (2003), Bandi, Russell, and Zhu (2008), de Pooter, Martens, and van Dijk (2008), and Bandi and Russell (2006) for examples in the realized covariance literature.

covariance estimators. We do not estimate an optimal α for our dataset but instead set $\alpha = 0.94$, which RiskMetrics (1996) found produces the best backtesting results for volatility estimates using daily stock returns. This is the RiskMetrics EWMA volatility model commonly used by practitioners for daily data.

2.3 Ledoit-Wolf Estimators

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

$$E[\|\Sigma_t^{LW} - \Sigma_t\|^2], \tag{6}$$

where $\Sigma_t^{LW}(\alpha) = \alpha G_t + (1-\alpha)\widehat{\Sigma}_t,\ \alpha \in [0,1]$, and G_t is the target matrix. This approach is equivalent to finding the optimal linear shrinkage of the sample eigenvalues.

To implement (6) we use the equicorrelated matrix, suggested by Ledoit and Wolf (2004a) and used by Voev (2008), as our target matrix. All the off-diagonal elements of the equicorrelated matrix are set to the covariance that, in combination with the estimated variances, results in all correlations equaling the average sample correlation. To estimate α within the high-frequency data setting, we exploit the long memory property of realized covariance (see Tables 1 and 2 below) and assume that the covariance process is locally constant. Appendix A provides the details of our approach and the results of parameter sensitivity analysis.

3 Empirical Analysis

We implement volatility timing strategies for a large dimension portfolio using daily covariance estimates based on high-frequency returns. We consider realized and mixed-frequency single-factor models in addition to traditional shrinkage estimators. We compare the performance of these approaches to high-frequency and low-frequency sample estimators. We assess the performance of the estimators based on the volatility of the dynamic minimum variance portfolios. We also examine the effect of portfolio dimensions on the forecast quality of the estimators. For robustness, we conduct sub-period analysis and Mincer and Zarnowitz (1969) forecast evaluations.

3.1 Data and Sample Realized Covariance Estimates

We estimate the time-varying covariance structure of the stocks in the Dow Jones Industrial Average (DJIA) from January 1, 2003 to December 31, 2006. For rolling estimators, we use returns proceeding January 1, 2003 to initialize the estimates. For the single factor models, we use the SPDR S&P 500 exchange traded fund (ETF) (SPY). Filtered price observations are sampled from TAQ quote records for the primary exchange for each security. Returns are calculated as the log-price-difference of the midpoints of quotes. See Appendix B for details.

The calendar-time realized covariance sample estimator for assets a and b, first proposed by Andersen, Bollerslev, Diebold, and Labys (2001), is

$$\Sigma_t^R(m) = \sum_{i=1}^{m} r_{at_i} \times r_{bt_i} \tag{7}$$

where m is the number of equally spaced intraday observations. m is chosen to maximize observations while minimizing market microstructure effects due to bid-ask bounce and price discreteness. For covariance estimation the Epps (1979) effect, a downward bias in the magnitude of the covariance estimates that increases as the sampling frequency increases, is of particular concern. The bias is the result of observed returns for one asset being paired with returns for the other that are measured as zero due to asynchroneity. As documented in Hansen and Lunde (2006b), the frequency of quote revisions for liquid stocks in the post-decimalization period reduces market microstructure biases for both realized variance and covariance estimates and reduces asynchroneity of observations across assets. We confine our analysis to a post-decimalization sample of liquid stocks and, for our primary analysis, estimate (7) with five-minute returns ($m = 78$).[4]

Zhang, Mykland, and Ait-Sahalia (2005) advocate sub-sampling as a technique to further reduce the influence of market microstructure effects and sampling error. Sub-sampling requires sparsely sampling the observations into multiple overlapping grids, estimating the realized covariance matrix for each sub-sample, then averaging the sub-estimates to generate the final estimate. The sub-sampled estimator over K sub-grids is

$$\widehat{\Sigma}_t^{R(K)}(m) = \sum_{k=1}^{K} \widehat{\Sigma}_t^{R(k)}(m), \tag{8}$$

where $\widehat{\Sigma}_t^{R(k)}(m)$ is the sample realized covariance for the kth grid. Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008a) point out that in the realized kernel framework $\widehat{\Sigma}_t^{R(K)}(m)$ has the same asymptotic distribution as the Bartlett kernel. We implement (8) with $K = 5$ to generate the realized covariance estimates for our primary analysis.[5]

Tables 1 and 2 provide summary variance and covariance statistics for our sample. The daily high-frequency estimates are calculated according to (8). The daily low-frequency variance estimates are squared open-to-close returns and the covariance estimates are the corresponding cross-products. For both estimators, the tables report the average, standard deviation (both $\times 1000$) and autocorrelation for 1, 15, and 30 lags of the daily variance and covariance estimates.

---

[4] Barndorff-Nielsen et al. (2008b) propose refresh-time tick-matching to align observations across assets and minimize the Epps effect. In this procedure, the most recent observation of each asset is used each time the price of all N assets in the chosen set have been "refreshed" with at least one new quote. 2x2 refresh-time tick-matching allows for maximal alignment of observations for each pair of assets, using a different set of price observations for a given asset for the $N - 1$ associated covariance elements. For a set of liquid assets similar to the stocks considered in this study, Barndorff-Nielsen et al. (2008b) report that the "classic" realized five-minute previous-tick covariance estimate behaves in a manner similar to the 2x2 refresh-time realized kernel covariance estimator.

[5] Alternative methods for estimating realized covariance in the presence of market microstructure effects and asychronous observations are numerous, but are beyond the scope of this paper. In particular, the literature investigates sampling based on dynamically optimized rules (Bandi and Russell, 2006; Bandi, Russell, and Zhu, 2008), cross-market tick-matching (Corsi, 2006; Hayashi and Yoshida, 2005), and refresh-time sampling (Barndorff-Nielsen, Hansen, Lunde, and Shephard, 2008b). See McAleer and Medeiros (2008) for a comprehensive review.

The cross-sectional mean and median of the high-frequency and low-frequency variance estimates are similar but the realized volatility estimates are less volatile and display greater autocorrelation than the estimates based on open-to-close returns. The 0.05 variance (x1000) for SPY corresponds to an 11% annualized open-to-close volatility and the average annualized open-to-close volatility for the stocks is 20%. For comparison, the corresponding close-to-close volatilities are 12% and 26% respectively.

For the covariance estimates, the table provides the average of the time-series statistics for all covariance pairs for the indicated stock. As with the variance estimates, the average realized covariance estimates have lower standard deviations and higher autocorrelations than the low-frequency estimates. The SPY autocovariance structure exhibits strong persistence, supporting consideration of a realized factor model for the covariance forecasts.

3.2 Model Parameter Estimates

To implement the single-factor and shrinkage covariance estimators we fit the market model to the data each day for the former and estimate the optimal shrinkage parameters for the latter. We use the RiskMetrics industry standard decay rate of 0.94 for the EWMA rolling estimator. For the Ledoit-Wolf estimator we estimate the shrinkage parameter daily based on a lagged one-year rolling window and, for simplicity, use the mean of these estimates, 0.4636.

In Table 3 we provide summary statistics for the coefficient estimates for the factor model applied to each stock using high-frequency and low-frequency returns. The realized betas (β_t^R) are estimated from sub-sampled realized variance and realized covariance estimates. The low-frequency betas (β_t^{LF}) are estimated using the variance and covariances calculated for a rolling 250-day window of open-to-close returns. The cross-sectional average of the realized betas is 0.84, ranging from a time-series average of 0.61 for JNJ to 1.35 for INTC. The corresponding average R^2 is 0.23 with values ranging from

0.08 for HP to 0.35 for C. The average standard deviation of the betas is 0.27 and the average autocorrelation is 0.33 at lag one, dropping to 0.19 at lag 15 and to 0.15 at lag 30. The realized betas, then, exhibit less serial correlation than do the realized variances and covariance.[6] Compared to the realized beta estimates, the average β_t^{LF} is higher at 0.95 with less time-series variation, and the average R^2 is 50% higher at 0.35, ranging from a low of 0.14 for HP and MO to a high of 0.55 for JPM.[7]

3.3 Volatility Timing

3.3.1 GMV Portfolio Optimization Problem

We compare the covariance estimators in terms of the ex post volatility of the dynamic global minimum variance portfolios (GMV) determined by optimizing portfolio weights with respect to each covariance estimate. Chan, Karceski, and Lakonishok (1999) and Jagannathan and Ma (2003) advocate focusing on this particular optimal allocation to avoid estimating the vector of expected returns. Patton and Sheppard (2008) show that the GMV portfolio constructed using the true covariance matrix has lower volatility than the corresponding portfolio for any other forecast.[8] The GMV portfolio solves the following constrained optimization problem:

$$\begin{aligned} min_{w_t} \quad & w_t^{'}\Sigma_t w_t \\ s.t. \quad & w_t^{'}\iota = 1 \end{aligned} \tag{9}$$

where Σ is the $p \times p$ covariance matrix and ι is the unit vector. The GMV weights are

$$w_{t,GMV} = \frac{\Sigma_t^{-1}\iota}{\iota'\Sigma_t^{-1}\iota}. \tag{10}$$

[6]These statistics are consistent with those reported in Barndorff-Nielsen et al. (2008b).

[7]We do not report the autocorrelation estimates for β_t^{LF} since the 250-day moving average approach results in predictably high statistics.

[8]In terms of general mean-variance portfolio optimization, Engle and Colacito (2006) show that for any assumed vector of constant expected returns the true covariance matrix results in a lower ex post portfolio volatility than any other covariance forecast.

At the end of day t we generate several covariance estimates using either day t intraday returns or historical open-to-close 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 open-to-close returns earned by the portfolio on day $t+1$. By confining our analysis to open-to-close returns, we avoid the additional model specification error introduced by overnight returns.[9]

3.3.2 Performance of the Estimators

In Table 4 we report the performance characteristics of dynamic GMV portfolios constructed based on the realized single factor covariance estimators, sample high-frequency and low-frequency estimators, and traditional shrinkage estimators. We evaluate the one-step ahead forecast performance of the estimators based on the volatility of the open-to-close returns of the optimal portfolios. We report the out-of-sample portfolio volatility and the 0.025 and 0.975 bootstrapped confidence intervals. The volatility of the equally-weighted portfolio, reported in the first row of the table, serves as a benchmark reflecting a naive asset allocation strategy.[10] We provide the in-sample results with portfolio returns evaluated at time t as an indication of estimation precision. For both sets of results we report the ratio of the dynamic portfolio volatility relative to the volatility of the equally-weighted portfolio and the maximum open-to-close loss experienced by the strategy. We report portfolio weight characteristics including the median minimum and maximum weights across the time series, the cross-sectional median of the standard deviation of the time-series of weights in each stock, and the median number of short portions.

[9] Hansen and Lunde (2005), Gallo (2001) and others address overnight returns by comparing a number of additive and scaling models.

[10] DeMiguel, Garlappi, and Uppal (2009) study the naive "$\frac{1}{N}$" diversification strategy in a low-frequency setting. They find that parameter estimation error makes it difficult for dynamic asset allocation strategies to generate higher Sharpe ratios than the $\frac{1}{N}$ portfolio over reasonable sample periods. This equally-weighted portfolio is equivalent to an extreme form of shrinkage that, under the assumption that expected returns are equal across the assets, places all of the weight on an identity matrix.

Before considering the single factor and shrinkage approaches, we present in Panel A the results for several sample covariance estimates using both high-frequency and low-frequency returns. RC_{78}, the estimate based on one sample of five-minute intra-day returns, yields volatility that is nearly equivalent to the benchmark. The sub-sampled realized covariance estimator, however, yields a volatility 120 basis points lower than the benchmark. With somewhat less extreme weights, less variable weights, and fewer short positions, the RC_{78s} dynamic portfolio reflects a less noisy estimate than RC_{78}. Both sample realized estimators deliver good performance in-sample, with GMV portfolio volatility at about 50% of the benchmark level. The contrast between the in-sample and out-of-sample performance indicates that, consistent with the literature, smoothing the daily estimates should be beneficial.

For comparison, we also consider rolling-window low-frequency estimators. OC_{78}, the low-frequency estimator that uses the same number of observations in the daily estimation as RC_{78}, outperforms the naive strategy and performs slightly better than the high-frequency estimator. The portfolio characteristics for OC_{78} are quite similar to those for RC_{78}, with the exception of somewhat smoother weights. OC_{78} does not outperform RC_{78s}, the sub-sampled high-frequency estimator. OC_{250}, based on a calendar-year of open-to-close returns, performs substantially worse than the other sample covariance estimators in sample, but provides the best forecast performance. Note that the out-of-sample volatility for OC_{250} is only 12% greater than the corresponding in-sample volatility, reflecting the extent of the smoothing in this estimator. The OC_{250} portfolio weights are 50% to 60% less variable than the weights associated with the other sample estimators, again indicating that smoothing the high-frequency estimates should be beneficial.

In Panel B we report the volatility timing results for several single factor models. RC_{78s}^{SF}, which applies the single-factor model to the sub-sampled high-frequency estimator, shows some loss of precision in-sample, as expected, but out-of-sample the estimator performs better than RC_{78s}, reducing volatility by 55 basis points relative to the subsampled estimator. The portfolio weights associated with RC_{78s}^{SF} are 30% less variable than the RC_{78s} weights but 35% more variable than the weights for OC_{250}, the sample estimate with the best forecast performance. This indicates that additional shrinkage, in particular smoothing, may further improve forecast quality of RC_{78s}^{SF}. We consider three approaches to generating smoother single-factor forecasts.

First we apply the EWMA RiskMetrics estimator ($\alpha = 0.94$) to the RC_{78s}^{SF} estimates. The resulting estimator, $\mathrm{RC}_{78s}^{SF,RM}$, yields a GMV volatility of 8.50%, over a 70 basis point drop from the volatility delivered by RC_{78s}^{SF}. While the extreme characteristics of the portfolio weights are quite similar for the two estimators, the weights associated with the smoothed single-factor model are about 40% less variable than those associated with RC_{78s}^{SF}. Next we consider a less naive forecast method, fitting an ARMA(1,1) model to the time series of RC_{78s}^{SF} to forecast the realized variances and realized betas. This forecast, $\mathrm{RC}_{78s}^{SF,ARMA}$, yields portfolio volatility of 8.72%, 50 basis points lower than for RC_{78s}^{SF} but 22 basis points higher than the simpler rolling estimate reflected in $\mathrm{RC}_{78s}^{SF,RM}$.

Our third approach is motivated by the time series characteristics of the realized variance, covariance, and beta estimates provided in Tables 1 through 3. These summary statistics indicate that the realized beta estimates are much less persistent than either the realized variances or covariances. Hence, we consider smoothing the single-factor estimates by directly introducing low-frequency beta estimates in conjunction with high-frequency variance estimates in a mixed-frequency single-factor model. Specifically, RC_{78s}^{MF} combines betas estimated daily using a 250-day window of open-to-close returns with high-frequency factor variance estimates and high-frequency residuals. Both in-sample and out-of-sample, this estimator performs well, with an out-of-sample volatility 23 basis points higher than $\mathrm{RC}_{78s}^{SF,RM}$, the best performing model. Variability of the RC_{78s}^{MF} weights is about 33% less than it is for RC_{78s}^{SF}, reflecting the effect of smoothing the covariance elements via the low-frequency beta.

In Panel C we provide the results for the two alternative shrinkage approaches applied to the sub-sampled realized covariance: the RiskMetrics EWMA model (RC_{78s}^{RM}) and the Ledoit-Wolf optimal shrinkage estimator (RC_{78s}^{LW}). Both estimators improve out-of-sample performance relative to the sample estimators. The reduction of estimation noise is evident in the portfolio characteristics.

We are interested in the performance of these estimators relative to the parsimonious single-factor models. RC_{78s}^{LW} performs in line with RC_{78s}^{SF} but substantially worse than $\mathrm{RC}_{78s}^{SF,RM}$. RC_{78s}^{RM}, on the other hand, with a volatility nearly 140 basis points lower than the RC_{78s} volatility, matches the performance of $\mathrm{RC}_{78s}^{SF,RM}$, the best performing single-factor model. The 70 basis point performance advantage of RC_{78s}^{RM} over RC_{78s}^{SF} highlights the importance of smoothing the daily estimates. While the extreme weights and the short positions are nearly matched for the two estimators, the RC_{78s}^{RM} portfolio weights are 45% less variable.[11]

The performance of the smoothed, realized single-factor models indicates that the reduction in estimation noise as the covariance estimates are shrunk to the factor covariances is greater than the loss in precision due to model misspecification. Our findings indicate that this trade-off in favor of structure holds even for a covariance matrix of modest dimension and for a portfolio of highly active stocks where sampling error in the covariance elements may be expected to be minimized. Furthermore, the precision offered by high-frequency variance estimates is maintained after introducing low-frequency factor coefficients in a mixed-frequency model. An advantage of the mixed-frequency estimator, which relies on high-frequency estimates for the variances only, is avoiding the need to estimate the realized covariances via computationally complex and data inten-

[11] We also consider the ad hoc approach of applying a nonnegative weight constraint on the portfolio optimization. Jagannathan and Ma (2003) demonstrate that this common practice can be viewed as a form of shrinkage on the covariance matrix estimate. The effect is the equivalent of shrinking the larger covariance matrix elements towards zero. In constrast to the results reported by Jagannathan and Ma (2003) in a low-frequency application, applying this constraint to dynamic portfolios based on the sub-sampled realized covariance estimate offers little benefit over the unconstrained optimization. The out-of-sample GMV volatility for the constrained portfolio is 0.0980.

sive procedures. A second advantage of the estimator is that low frequency betas can be estimated even for less liquid assets where non-synchronous observations and, hence, the Epps effect may be a more severe problem. Finally, given the structure imposed by the factor model, a third advantage is that the estimator is scalable to large dimension matrices.[12]

3.3.3 Portfolio Dimension Analysis

To further investigate the interplay of estimation error and portfolio dimensions, we consider the GMV volatility timing strategies across portfolios ranging from 2 to 30 assets. For each portfolio dimension N, we draw 500 random portfolios from the 30 DJIA stocks. For each of these portfolios, we execute the volatility timing strategies over the entire sample period and we track the volatility of the benchmark equally-weighted portfolio. We execute the portfolio bootstrap for $N = 2, 3, 5, 7, 10, 15, 20, 25$, and 27 stocks and we include the previously reported volatilities for the full 30-stock portfolio. The results of this analysis are summarized in the volatility dimension curves in Figure 1. The solid line in each panel is the mean volatility, at each dimension, for the equally-weighted portfolios of the sampled assets. This curve provides a benchmark for volatility reduction resulting from a pure diversification effect.

The curves in Panels A and B show little difference in the volatilities for the various dynamic strategies when the number of assets is small. For two, three, and five assets, they perform quite similarly. This indicates that the volatility timing assessment criteria may have little power to distinguish among estimators when the ratio of observations to dimensions is high for all candidates. Correspondingly, with a few exceptions, across the volatility dimension curves the performance differentials relative to the benchmark portfolio are greater as the number of assets increases. A larger N offers a larger investment opportunity set over which to use the information. If an estimator remains informative

[12]See Bannouh et al. (2009) for estimation of a mixed-frequency multi-factor model for vast dimensional covariance matrices for low-liquidity stocks.

as N increases, this greater scope can translate into relatively greater volatility reduction as the set of possible portfolios increases.

The results for the sample covariance estimators are provided in Panel A. We have added a 30-day rolling window low-frequency estimator, OC_{30}, to highlight the effect of the number of observations at increasing matrix dimensions. At seven assets the OC_{30} volatility dimension curve begins to diverge sharply upward. OC_{30} yields positive definite matrices across all dimensions but at 25 (30) assets the average volatility has reached 93% (147%), reflecting an ill-conditioned matrix. OC_{78}, on the other hand, provides performance consistent with the better performing sample estimators through 20 assets before estimation error begins to overwhelm the diversification effect. OC_{250} performs well through-out, with volatility falling at each increment of portfolio size. The RC_{78} volatility dimension curve reflects worse performance than OC_{78} and OC_{250} throughout. Estimation error for the single-sample five minute estimator begins to offset the diversification effect by 20 assets. The sub-sampled variant, RC_{78s}, provides lower volatility than RC_{78} throughout and shows no sign of increasing noise but performs the same or worse than OC_{250} at all dimensions.

Panel B provides the volatility dimension curves for the single-factor models, the equally-weighted portfolio, and RC_{78s}^{RM}, the best-performing shrinkage estimator. The realized single-factor model performs well throughout, even at smaller N where the factor model is not expected to fit the sampled portfolios as well as it does at higher dimensions. The mixed-frequency model dominates at all dimensions and performs closely in-line with the smoothed, fully-estimated model, RC_{78s}^{RM}.

To investigate the conditioning benefits of sub-sampling the high-frequency estimators, we also consider realized covariance estimates based on returns sampled every 30 minutes, a common sampling frequency in the realized volatility literature. The estimator, RC_{13}, uses 13 intra-day return observations and RC_{13s}, the sub-sampled counterpart, is the average of five estimated matrices with observation grids shifted forward by five minutes for each subsequent sample. As a robustness check on the high-frequency single-factor model, we consider a single factor model based on 30-minute returns, RC_{13s}^{SF}. Panel C provides the volatility dimension curves for these estimators.

With volatilities greater than the benchmark, RC_{13} is revealed as a noisy estimator at only two and three assets and the estimator is not positive definite above 10 dimensions. The curve for the sub-sampled variant, RC_{13s}, is consistently high throughout, reflecting substantial estimation error. It is, however, positive definite for all dimensions considered, revealing the conditioning power of sub-sampling. For less active assets where intra-day sampling frequency is limited, these results indicate that better-conditioned and invertible covariance matrices can still be obtained through sub-sampled estimates. Furthermore, applying the single-factor structure to this estimator sharply reduces estimation error. The volatility curve for RC_{13s}^{SF} falls below the benchmark curve by 15 assets and continues to fall through 30 assets.[13]

Finally, the volatility dimension curves allow us to quantify forecast quality comparisons in terms of the corresponding pure diversification benefit captured in the $\frac{1}{N}$ portfolio. Switching from the single-sample five-minute high-frequency estimator, RC_{78}, to the corresponding sub-sampled estimator, RC_{78s}, for example, reduces GMV portfolio volatility for the 30-asset portfolio by 100 basis points. This is equivalent to the pure diversification benefit of increasing from holding seven to holding 25 of this universe of 30 stocks. Switching from RC_{78s} to the smoothed single-factor model, $\mathrm{RC}_{78s}^{SF,RM}$, reduces GMV volatility a further 140 basis points, equivalent to the diversification benefit of increasing from six to 25 stocks.

[13] A comparison of the results for RC_{13s} to those for the high frequency estimators based on five-minute returns, indicates that for active stocks in our sample period the 30-minute sampling frequency is not optimal. This contrasts with the results reported by Liu (2009) and de Pooter et al. (2008) for similar samples of stocks but for pre-decimalization sample periods where trade and quote activity were much less than in the post-decimalization period.

3.3.4 Subperiod Robustness Analysis

Table 5 provides a summary of the robustness of the performance of these estimators over different sample periods and for different market conditions. We report the volatility for the dynamic GMV and the EW benchmark portfolios for the four calendar years in our sample and for trading days marked by high and low volatility. We define low (high) volatility days as those days with an absolute return on the equally-weighted benchmark portfolio that is less (greater) than the median absolute return on the portfolio over the sample period. The relative performance of these estimators is robust to all sub-periods (with benchmark volatility ranging from 9% to over 15%) and to high and low volatility periods specifically. In particular, in each sub-sample, the smoothed realized covariance estimator and the realized single-factor models perform well relative to the alternative estimators.

3.4 Mincer-Zarnowitz Forecast 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. To give equal weight to the accuracy of each element of the covariance matrix, portfolio volatility is assessed for an equally-weighted portfolio: $\widehat{\sigma}^2_{EW} = w'\widehat{\Sigma}w$, 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.

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:

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

$\widehat{\sigma}_{EW}(t)$ is the portfolio volatility proxy at time t and $\widetilde{\sigma}_{EW}(t-1)$ is the candidate forecast of portfolio volatility. The joint null hypothesis is $H_0 : b_0 = 0$ and $b_1 = 1$.

We also consider encompassing forecast quality regressions to test the hypothesis that a given forecast provides incremental value over another. We estimate the following regression:

$$\widehat{\sigma}_{EW}(t) = b_0 + b_1 \widetilde{\sigma}^{(1)}_{EW}(t-1) + b_2 \widetilde{\sigma}^{(2)}_{EW}(t-1) + \varepsilon(t)$$

where $\widetilde{\sigma}^{(1)}_{EW}(t-1)$ and $\widetilde{\sigma}^{(2)}_{EW}(t-1)$ are candidate forecasts of portfolio volatility. The joint null hypothesis is $H_0 : b_0 = 0$ and $b_1 + b_2 = 1$. If either b_1 or b_2 equal zero, the corresponding forecast is encompassed by the other.

We perform the Mincer-Zarnowitz regressions directly on the volatility levels. In simulation and empirical analysis, Hansen and Lunde (2006a) find this specification provides consistent ranking of forecast models and is more robust than the log-volatility formulation. In addition, they find that while the Mincer-Zarnowitz regressions have heteroskedastic error terms when applied to variances this is not the case for volatility regressions. We report robust Newey-West standard errors for our parameter estimates to account for autocorrelated errors. The proxy for the unobservable ex post portfolio volatility is the sub-sampled realized portfolio volatility for day t. The regression statistics are reported in Table 6.

The statistical forecast quality results confirm the inference from the volatility timing analysis: the smoothed high-frequency forecasts, RC^{RM}_{78s} and $\mathrm{RC}^{SF,RM}_{78s}$, exhibit less bias and are more efficient than the high- and low-frequency sample estimators. Of the sample estimators, RC_{78s} has the highest R^2 at 0.57, a considerable margin over the R^2 of 0.45 for the long-horizon low-frequency estimator, OC_{250}. The Ledoit-Wolf optimally conditioned estimate, RC^{LW}_{78s}, performs quite similar to RC_{78s} in terms of bias, efficiency, and explanatory power. The single-factor models, RC^{SF}_{78s} and RC^{MF}_{78s}, provide forecast quality equivalent to the fully-estimated realized covariance matrix, RC_{78s}. With the EWMA shrinkage model applied to RC_{78s} and RC^{SF}_{78s}, their performance is equivalent with nearly the same R^2 and neither smoothed estimator exhibiting significant bias or significant inefficiency. In terms of goodness of fit, incorporating ARMA forecasts offers improved explanatory power ($R^2 = 0.66$) but this comes with increased bias and inefficiency.

In the final column of Table 6, we report the R^2 of the Mincer-Zarnwitz encompassing regressions with RC_{78s} as the candidate encompassing forecast. The results indicate that the addition of the sub-sampled high-frequency estimator substantially improves the explanatory power of the forecast regressions for the low-frequency sample estimators and somewhat increases the R^2 for the single-sample high-frequency estimator. The R^2 increases for the smoothed estimators, RC_{78s}^{RM} and $RC_{78s}^{SF,RM}$, but there is no evidence that RC_{78s} provides additional forecast information to the other single-factor models or to the Ledoit-Wolf estimator.

The Mincer-Zarnowitz forecast quality results indicate that high-frequency shrinkage estimates provide greater forecast predictability than forecasts based on either low- or high-frequency sample estimates. The results also confirm that the parsimonious high-frequency and mixed-frequency single-factor models provide good forecast quality and are not encompassed by the high-frequency sample estimator, RC_{78s}. The statistical assessment indicates that the $RC_{78s}^{SF,ARMA}$ is a more promising forecast than was indicated by the volatility timing assessment.

4 Future Work and Conclusion

Based on the volatility timing and Mincer-Zarnowitz forecast quality evaluation results for a post-decimalization sample of 30 DJIA stocks we conclude that sub-sampling is very effective in reducing estimation error in large realized covariance matrices; that further matrix conditioning is important for forecast quality; and, for fully-estimated matrices, a naive EWMA shrinkage approach works well. Our analysis of single factor models indicates that the structure provided by this model substantially reduces estimation noise and that smoothing these estimates yields performance equivalent to the smoothed fully-estimated realized covariance. Furthermore, using low-frequency betas as a smoothing technique in a mixed-frequency estimate provides performance close to the smoothed high-frequency single factor model. The parsimony of the factor model, in general, is attractive for large-dimension applications and the additional estimation simplification offered by the mixed-frequency model may be particularly important when considering less-liquid assets where nonsynchroneity becomes a greater problem.

Volatility timing bootstrap experiments at dimensions ranging from 2 to 27 provide additional insight on assessing covariance forecast quality. We find, not surprisingly, that the ability to discriminate across estimators in a portfolio optimization setting depends on portfolio dimensions. This dependence, however, is not simply a function of estimation noise but is also driven by a scope effect; i.e., with more assets there is greater opportunity to take advantage of a forecast signal.

There are many avenues for additional research. Of particular importance for general applicability of realized covariance estimation is developing additional techniques for the inclusion of overnight returns into these estimates. Estimating the covariance matrix for non-trading periods presents additional challenges but is necessary to accommodate varying forecast horizons and to consider more practitioner-oriented dynamic strategies. The success of the mixed-frequency factor model in estimating trading-period covariances suggests this approach for reducing noise in estimates based on observed overnight returns. In addition, the performance of the mixed-frequency model should be investigated using low-frequency betas estimated over shorter windows than the 250 days used in this analysis. Where more sophisticated forecasting techniques are appropriate, the results for the ARMA(1,1) forecast model in conjunction with the single factor model suggest that such techniques can add value at high dimensions. Finally, in terms of forecast assessment methodology, our dimensional analysis indicates the need for a more robust understanding of the choice of "N" on volatility timing experiments as well as on portfolio optimization experiments more generally.

References

Andersen, T. and Bollerslev, T. (1998), "Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts," *International Economic Review*, 39, 885–905.

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.

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.

Bannouh, K., Martens, M., Oomen, R., and van Dijk, D. (2009), "Realized Factor Models for Vast Dimensional Covariance Estimation," Working paper.

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*, 76.

— (2008b), "Multivariate Realized Kernels: Consistent Positive Semi-Definite Estimators of the Covariation of Equity Prices with Noise and Non-Synchronous Trading," Working paper.

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.

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.

DeMiguel, V., Garlappi, L., and Uppal, R. (2009), “Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?” *Review of Financial Studies*, 22, 1915–1953.

Engle, R. and Colacito, R. (2006), “Testing and Valuing Dynamic Correlations for Asset Allocation,” *Journal of Business and Economic Statistics*, 238–253.

Epps, T. (1979), “Comovement in Stock Prices in the Very Short Run,” *Journal of the American Statistical Association*, 291–298.

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.

Ghysels, E., Santa-Clara, P., and Valkanov, R. (2006), "Predicting Volatility: Getting the Most out of Return Data Sampled at Different Frequencies," *Journal of Econometrics*, 59–95.

Hansen, P. and Lunde, A. (2006a), "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.

— (2006b), "Realized Variance and Market Microstructure Noise," *Journal of Business and Economic Statistics*, 24, 127–161.

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 Covariance 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. (2009), "On Portfolio Optimization: How Do We Benefit from High-Frequency Data?" *Journal of Applied Econometrics*, forthcoming.

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

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.

Patton, A. and Sheppard, K. (2008), "Evaluating Volatility and Correlation Forecasts," in *Handbook of Financial Time series*, eds. Andersen, T. G., A.Davis, R., Kreiss, J.-P., and Mikosch, T., Springer, pp. 801–828.

RiskMetrics (1996), *RiskMetrics Technical Document - 4th Edition*, J.P.Morgan/Reuters.

Stein, C. (1956), "Inadmissibility of the Usual Estimator for the Mean of a Multivariate Normal Distribution," in *Proceedings of the Third Berkeley Symposium on Mathematical and Statistical Probability*, ed. Neyman, J., University of California, Berkeley, pp. 197–206.

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.

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 Implementation of Ledoit-Wolf Estimator

Φ, with elements ϕ_{ij}, denotes the unobserved true equicorrelated covariance matrix and F, with elements f_{ij}, is 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_{ij} are the elements of $\widehat{\Sigma}$:

$$\alpha = \frac{\sum_{i=1}^{p}\sum_{j=1}^{p} Var(s_{ij}) - Cov(f_{ij}, s_{ij})}{\sum_{i=1}^{p}\sum Var(f_{ij} - s_{ij}) + (\phi_{ij} - \sigma_{ij})^2} \quad s.t. \quad \alpha \in [0,1]. \tag{11}$$

For implementation within the high-frequency data setting, we exploit the long memory property of realized covariance and assume that the covariance process is locally constant. For a window of length l, we estimate α by assuming that $E[s_{ij,\tau-l}] = \sigma_{ij,t}$ where $\tau \in [t-l, t]$. The quantities $Var(s_{ij}), Cov(f_{ij}, s_{ij}), Var(f_{ij} - s_{ij})$, and $(\phi_{ij} - \sigma_{ij})^2$ are estimated using the daily time series of realized covariance estimates. σ_{ij} is estimated as the sample average of $s_{ij,t}$ over the l-day window and ϕ_{ij} is the corresponding equicorrelated matrix.[14]

Sensitivity analysis validates the approach of assuming a locally constant covariance structure. Table 7 reports the GMV volatility associated with the Ledoit-Wolf shrinkage approach over a large range of values for the shrinkage parameter, α. The variation in volatility is quite modest. Therefore we conclude that misspecification does not have a substantial impact on the performance of the shrinkage estimator.

[14]Voev (2008) outlines an alternative adaptation of the Ledoit-Wolf estimator for realized covariance estimators that exploits the asymptotic variance and covariance results derived in Barndorff-Nielsen and Shephard (2004).

B High Frequency Data Filtering and Sampling

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) that duplicate a previous record (same time stamp, same bid, same ask); 7) reflecting a 10% move from the previous quote midpoint; 8) with a spread greater than 10% of the midpoint price; and 9) that are "redundant," i.e., reflecting no revision to the bid or ask from the most recent quote.

To facilitate sub-sampling across 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 quotes recorded for the 9:30:00 and 16:00:00 grid points.

Figure1: Average GMV Volatility Across Bootstrapped Portfolios of N Assets

Panel A. Sample Estimators vs. EW Benchmark Portfolio

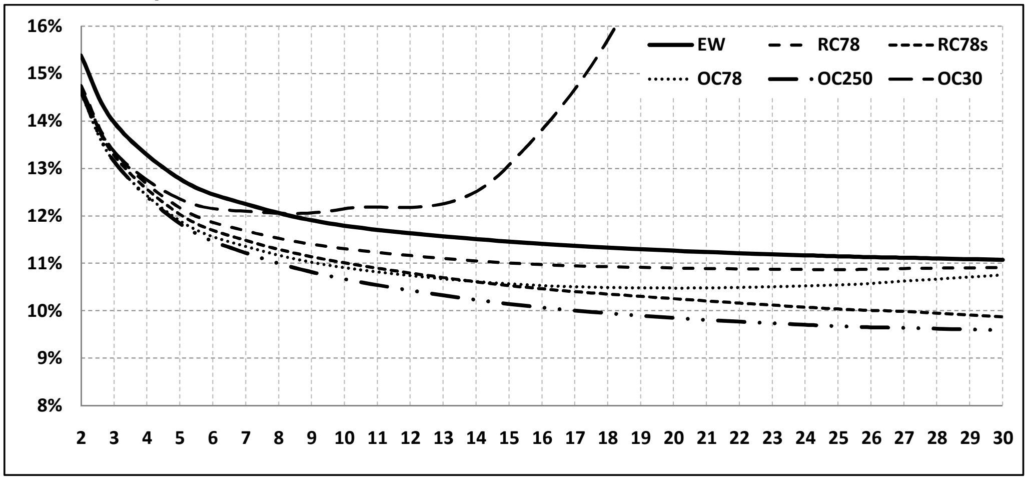

Panel B. HF Single Factor Models vs. EW and RC78sRM

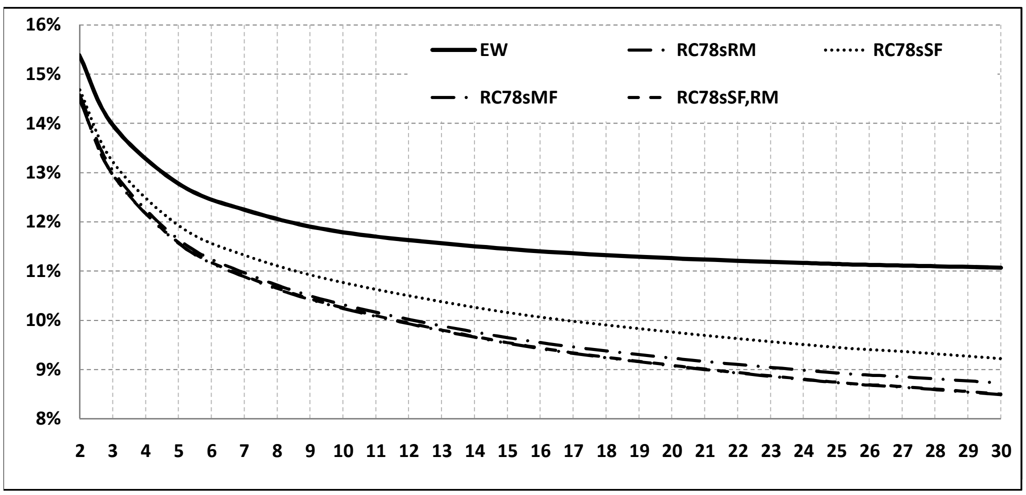

Panel C. RC13 Estimators vs. EW and RC78sRM

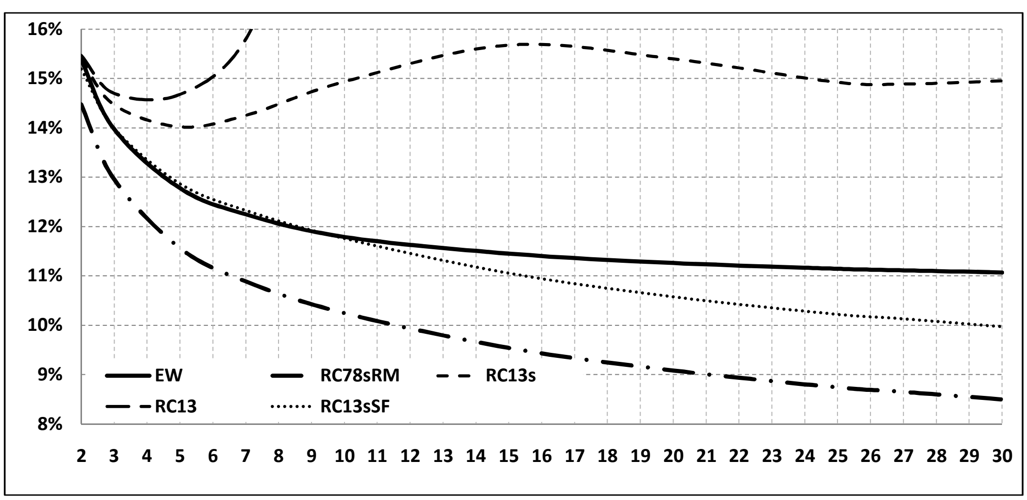

Table 1: **Summary Statistics for Daily Variance Estimates**

"Mean," "StDev," and "ρ_l" are the average ($\times$1000), the standard deviation ($\times$1000), and autocorrelation for $l = 1, 15$, and 30 lags, of daily variance estimates for the DJIA stocks and the S&P500 ETF (SPY). Using high(low)-frequency data, daily variance is the sum of squared five-minute returns averaged over five sub-samples (squared open-to-close returns). The sample period is January 1, 2003 to December 31, 2006. Superscripts denote sub-periods for inclusion of the stock in the index: †(4/8/2004–12/31/2006), §(1/1/2003–4/7/2004 and 11/22/2005–12/31/2006), *(1/1/2003–4/7/2004), and ‡(1/1/2003–12/21/2005).

	High Frequency Estimator					**Low Frequency Estimator**				
Ticker	**Mean** (×1000)	**StDev** (×1000)	ρ_1	ρ_{15}	ρ_{30}	**Mean** (×1000)	**StDev** (×1000)	ρ_1	ρ_{15}	ρ_{30}
AA	0.24	0.15	0.52	0.29	0.16	0.21	0.31	0.16	0.16	0.10
AIG†	0.15	0.18	0.54	0.21	0.17	0.15	0.47	0.24	0.13	0.08
AXP	0.10	0.10	0.68	0.53	0.55	0.11	0.23	0.17	0.15	0.13
BA	0.17	0.12	0.63	0.51	0.46	0.16	0.26	0.10	0.08	0.02
C	0.11	0.09	0.75	0.59	0.58	0.09	0.17	0.45	0.12	0.03
CAT	0.18	0.11	0.50	0.17	0.20	0.19	0.35	0.03	0.00	0.03
DD	0.13	0.07	0.52	0.25	0.20	0.10	0.16	0.06	0.07	0.08
DIS	0.17	0.15	0.70	0.51	0.52	0.16	0.32	0.21	0.19	0.22
EK*	0.24	0.25	0.27	-0.01	0.00	0.26	0.79	0.04	0.00	0.01
GE	0.10	0.08	0.71	0.57	0.56	0.09	0.18	0.25	0.18	0.15
GM	0.29	0.32	0.47	0.26	0.27	0.33	0.74	0.18	0.10	0.08
HD	0.17	0.11	0.50	0.30	0.24	0.16	0.34	0.15	0.21	0.10
HON	0.20	0.17	0.55	0.28	0.30	0.15	0.29	0.11	0.11	0.03
HP	0.34	0.20	0.53	0.26	0.16	0.35	0.56	0.04	0.03	0.02
IBM	0.10	0.06	0.56	0.28	0.22	0.09	0.19	0.12	0.22	0.16
INTC	0.24	0.19	0.46	0.24	0.28	0.24	0.38	0.30	0.30	0.17
IP*	0.16	0.10	0.53	0.30	0.31	0.15	0.29	0.10	0.04	0.05
JNJ	0.09	0.09	0.55	0.44	0.39	0.08	0.17	0.16	0.22	0.06
JPM	0.13	0.12	0.73	0.51	0.55	0.13	0.31	0.37	0.08	0.02
KO	0.09	0.06	0.69	0.47	0.48	0.07	0.14	0.15	0.16	0.08
MCD	0.18	0.17	0.40	0.26	0.32	0.17	0.35	0.17	0.04	0.10
MMM	0.10	0.06	0.39	0.04	0.00	0.09	0.19	0.02	0.03	0.02
MO	0.16	0.32	0.20	0.12	0.04	0.14	0.53	0.09	0.01	0.02
MRK	0.16	0.25	0.13	0.18	0.03	0.17	0.59	0.12	0.30	0.03
MSFT	0.13	0.18	0.20	0.16	0.14	0.12	0.24	0.45	0.20	0.15
PFE†	0.14	0.15	0.19	0.06	0.06	0.13	0.28	0.10	0.19	0.09
PG	0.08	0.05	0.47	0.08	0.11	0.06	0.10	0.14	0.10	0.05
SBC‡	0.19	0.20	0.69	0.55	0.58	0.17	0.38	0.15	0.24	0.10
T§	0.18	0.20	0.44	0.33	0.34	0.18	0.40	0.32	0.18	0.19
UTX	0.13	0.08	0.60	0.32	0.21	0.11	0.21	0.12	0.14	0.13
$VZ^{\dagger}$	0.15	0.13	0.63	0.55	0.44	0.12	0.25	0.26	0.32	0.14
WMT	0.11	0.07	0.47	0.27	0.24	0.09	0.16	0.16	0.27	0.10
XOM	0.12	0.08	0.61	0.29	0.14	0.12	0.18	0.07	0.08	0.04
Mean	0.16	0.14	0.51	0.31	0.28	0.15	0.32	0.17	0.14	0.08
Median	0.15	0.12	0.53	0.28	0.24	0.14	0.29	0.15	0.14	0.08
SPY	0.05	0.04	0.68	0.51	0.49	0.05	0.10	0.14	0.24	0.06

Table 2: **Summary Statistics for Daily Covariance Estimates**

"Mean," "StDev," and "ρ_l" are the cross-sectional average of the time-series average ($\times 1000$), standard deviation ($\times 1000$), and autocorrelation for $l = 1, 15$, and 30 lags, of daily covariance estimates for each DJIA stock and for the S&P500 ETF (SPY) with DJIA stocks. Using high(low)-frequency data, daily covariance is the sum of the cross-product of five-minute returns averaged over five sub-samples (cross-product of open-to-close returns). The sample period is January 1, 2003 to December 31, 2006.

	High Frequency Estimator					**Low Frequency Estimator**				
	Cross-Sectional Mean					**Cross-Sectional Mean**				
Ticker	**Mean** (×1000)	**StDev** (×1000)	ρ_1	ρ_{15}	ρ_{30}	**Mean** (×1000)	**StDev** (×1000)	ρ_1	ρ_{15}	ρ_{30}
AA	0.05	0.05	0.48	0.33	0.27	0.06	0.20	0.02	0.08	0.02
AIG	0.04	0.05	0.62	0.48	0.48	0.05	0.20	0.02	0.12	0.11
AXP	0.04	0.05	0.66	0.53	0.53	0.05	0.18	0.00	0.09	0.11
BA	0.04	0.04	0.57	0.44	0.39	0.05	0.18	0.03	0.05	0.05
C	0.04	0.05	0.67	0.52	0.51	0.05	0.15	0.03	0.10	0.08
CAT	0.05	0.05	0.55	0.37	0.37	0.06	0.21	0.02	0.06	0.07
DD	0.04	0.04	0.57	0.40	0.39	0.05	0.15	0.05	0.05	0.08
DIS	0.04	0.05	0.68	0.51	0.54	0.05	0.20	0.03	0.09	0.11
EK	0.04	0.04	0.53	0.41	0.41	0.05	0.22	0.00	0.05	0.03
GE	0.04	0.05	0.66	0.51	0.51	0.05	0.16	0.03	0.10	0.09
GM	0.04	0.04	0.50	0.38	0.37	0.06	0.23	0.02	0.05	0.01
HD	0.04	0.04	0.55	0.38	0.36	0.06	0.21	0.00	0.09	0.05
HON	0.05	0.05	0.59	0.44	0.41	0.06	0.19	0.01	0.11	0.04
HP	0.03	0.04	0.30	0.20	0.13	0.04	0.23	0.01	0.01	0.01
IBM	0.04	0.04	0.61	0.46	0.44	0.04	0.16	0.01	0.09	0.05
INTC	0.06	0.05	0.62	0.46	0.47	0.07	0.23	0.00	0.08	0.04
IP	0.04	0.05	0.58	0.44	0.42	0.05	0.17	0.03	0.05	0.07
JNJ	0.03	0.03	0.57	0.49	0.44	0.03	0.14	0.04	0.10	0.10
JPM	0.04	0.05	0.62	0.45	0.45	0.06	0.20	0.00	0.10	0.10
KO	0.03	0.03	0.60	0.48	0.44	0.03	0.12	0.04	0.12	0.08
MCD	0.04	0.04	0.52	0.38	0.35	0.04	0.18	0.03	-0.01	-0.02
MMM	0.04	0.03	0.53	0.37	0.31	0.04	0.14	0.00	0.06	0.05
MO	0.03	0.04	0.51	0.39	0.35	0.03	0.16	-0.0	0.00	0.02
MRK	0.04	0.04	0.56	0.43	0.39	0.04	0.17	0.03	0.06	0.04
MSFT	0.04	0.05	0.68	0.53	0.48	0.05	0.18	0.04	0.15	0.05
PFE	0.04	0.04	0.55	0.40	0.35	0.04	0.16	0.02	0.07	0.03
PG	0.03	0.03	0.53	0.35	0.31	0.03	0.11	0.00	0.06	0.04
SBC	0.05	0.06	0.65	0.49	0.53	0.06	0.21	0.02	0.14	0.11
T	0.03	0.04	0.49	0.36	0.40	0.04	0.21	0.03	0.10	0.11
UTX	0.04	0.04	0.55	0.38	0.34	0.05	0.17	0.01	0.09	0.06
VZ	0.04	0.05	0.62	0.50	0.47	0.05	0.17	0.03	0.11	0.08
WMT	0.04	0.04	0.60	0.46	0.44	0.04	0.15	0.03	0.10	0.06
XOM	0.04	0.04	0.58	0.39	0.35	0.04	0.14	0.05	0.04	0.05
Mean	0.04	0.04	0.57	0.43	0.41	0.05	0.18	0.02	0.08	0.06
Median	0.04	0.04	0.57	0.44	0.41	0.05	0.18	0.02	0.09	0.05
SPY	0.04	0.04	0.68	0.51	0.48	0.05	0.12	0.03	0.12	0.10

Table 3: **S&P 500 Factor Model Estimates**

The table provides the means and standard deviations for β_t^R, the daily sub-sampled realized market factor coefficient estimate using five minute returns, and β_t^{LF}, the daily low-frequency factor coefficient estimate using the past 250 open-to-close returns. ρ_l, the autocorrelation for $l = 1, 15, 30$ lags, is reported for the realized beta estimates. The average of the daily R_t^2 are reported for each model. The sample period is January 2, 2003 to December 31, 2006.

	High Frequency Betas						**Low Frequency Betas**		
Ticker	**Mean** $\boldsymbol{\beta_t^R}$	**StDev** $\boldsymbol{\beta_t^R}$	$\boldsymbol{\rho_1}$	$\boldsymbol{\rho_{15}}$	$\boldsymbol{\rho_{30}}$	**Mean** $\boldsymbol{R_t^2}$	**Mean** $\boldsymbol{\beta_t^{LF}}$	**StDev** $\boldsymbol{\beta_t^{LF}}$	**Mean** $\boldsymbol{R_t^2}$
AA	1.01	0.35	0.39	0.17	0.15	0.18	1.22	0.11	0.37
AIG	0.89	0.30	0.37	0.25	0.29	0.26	1.08	0.12	0.40
AXP	0.77	0.24	0.45	0.29	0.26	0.28	1.06	0.15	0.50
BA	0.90	0.25	0.17	0.06	0.05	0.22	0.98	0.10	0.31
C	0.88	0.22	0.41	0.26	0.29	0.35	1.03	0.20	0.54
CAT	1.04	0.30	0.41	0.19	0.13	0.27	1.22	0.14	0.43
DD	0.89	0.23	0.33	0.09	0.05	0.28	0.95	0.09	0.43
DIS	0.81	0.30	0.50	0.34	0.29	0.19	1.01	0.18	0.34
EK	0.73	0.32	0.16	0.05	0.07	0.11	0.98	0.16	0.22
GE	0.84	0.23	0.50	0.39	0.35	0.33	0.99	0.16	0.52
GM	0.82	0.34	0.11	0.08	0.00	0.14	1.12	0.07	0.28
HD	0.95	0.27	0.27	0.06	0.09	0.24	1.14	0.08	0.42
HON	1.06	0.27	0.21	0.04	0.04	0.26	1.12	0.10	0.39
HP	0.83	0.54	0.55	0.47	0.39	0.08	0.97	0.47	0.14
IBM	0.83	0.19	0.17	0.06	0.00	0.32	0.88	0.12	0.42
INTC	1.35	0.37	0.46	0.26	0.11	0.33	1.47	0.15	0.44
IP	0.85	0.26	0.26	0.10	0.10	0.21	1.06	0.13	0.39
JNJ	0.61	0.21	0.43	0.33	0.26	0.22	0.64	0.10	0.27
JPM	0.91	0.24	0.33	0.17	0.14	0.30	1.23	0.23	0.55
KO	0.66	0.20	0.37	0.29	0.28	0.22	0.64	0.09	0.30
MCD	0.82	0.24	0.17	0.05	-0.0	0.19	0.80	0.16	0.19
MMM	0.80	0.20	0.28	0.14	0.03	0.28	0.83	0.14	0.37
MO	0.67	0.29	0.24	0.12	0.03	0.17	0.64	0.18	0.14
MRK	0.73	0.31	0.18	0.08	-0.0	0.17	0.79	0.10	0.23
MSFT	0.91	0.28	0.52	0.37	0.32	0.30	0.98	0.19	0.43
PFE	0.78	0.28	0.21	0.10	0.00	0.20	0.86	0.08	0.30
PG	0.66	0.20	0.31	0.17	0.13	0.24	0.63	0.10	0.30
SBC	0.83	0.30	0.47	0.31	0.34	0.21	0.91	0.23	0.33
T	0.67	0.28	0.29	0.16	0.16	0.13	0.83	0.16	0.21
UTX	0.90	0.23	0.30	0.06	0.02	0.28	0.96	0.07	0.41
VZ	0.78	0.27	0.41	0.29	0.26	0.20	0.85	0.18	0.31
WMT	0.81	0.22	0.26	0.12	0.12	0.27	0.81	0.09	0.36
XOM	0.86	0.27	0.40	0.28	0.23	0.27	0.94	0.20	0.40
Mean	0.84	0.27	0.33	0.19	0.15	0.23	0.96	0.15	0.35
Median	0.83	0.27	0.33	0.17	0.13	0.24	0.97	0.14	0.37

Table 4: **Dynamic GMV Portfolio Performance**

Volatility statistics are reported for the benchmark equally-weighted portfolio and for the dynamic GMV portfolios optimized each day t using the covariance estimator indicated in the first column. The out-of-sample (in-sample) portfolio volatility (σ_P) is the annualized standard deviation of the open-to-close day $t+1$ (t) returns for the portfolio. "0.025" and "0.975" report the empirical confidence intervals from 100 draws with replacement from the 993 daily portfolio returns, "Vol. vs. EW" is the ratio of the volatility of the dynamic strategy vs. the benchmark, and "Max. Loss" reports the largest negative open-to-close portfolio return. "Med. Min" ("Med. Max") is the median minimum (maximum) weight across the time series, "Med. StDev" is the cross-sectional median of the standard deviation the time-series of weights in each stock, and "Med. # Shorts" is the median of the time series of short positions in the portfolio. The sample period is January 2, 2003 - December 31, 2006.

	Out-of-Sample					Portfolio Weight Characteristics				In-Sample		
		Empirical C.I.		Vol. vs.	Max	Med.	Med.	Med.	Med. #		Vol. vs.	Max
Model	$\boldsymbol{\sigma_P}$	**0.025**	**0.975**	**EW**	**Loss**	**Min**	**Max**	**StDev**	**Shorts**	$\boldsymbol{\sigma_P}$	**EW**	**Loss**
EW	0.1107	0.1042	0.1177	-	-0.0266	0.03	0.03	-	-	0.1107	-	-0.0266
Panel A: Sample Estimators												
RC_{78}	0.1091	0.1029	0.1152	0.99	-0.0218	-0.16	0.33	0.0992	12	0.0583	0.53	-0.0151
RC_{78s}	0.0987	0.0941	0.1036	0.89	-0.0227	-0.11	0.29	0.0756	11	0.0576	0.52	-0.0137
OC_{78}	0.1075	0.1027	0.1126	0.97	-0.0285	-0.16	0.34	0.0871	12	0.0666	0.60	-0.0130
OC_{250}	0.0959	0.0919	0.0999	0.87	-0.0210	-0.10	0.23	0.0387	11	0.0858	0.78	-0.0185
Panel B: High Frequency Single Factor Models												
RC_{78s}^{SF}	0.0922	0.0881	0.0964	0.83	-0.0197	-0.07	0.23	0.0526	9	0.0687	0.62	-0.0147
$RC_{78s}^{SF,RM}$	0.0850	0.0811	0.0886	0.77	-0.0182	-0.08	0.21	0.0325	8	0.0825	0.75	-0.0180
$RC_{78s}^{SF,ARMA}$	0.0872	0.0825	0.0921	0.79	-0.0366	-0.08	0.19	0.0358	7	-	-	-
RC_{78s}^{MF}	0.0872	0.0835	0.0909	0.79	-0.0177	-0.08	0.25	0.0371	11	0.0762	0.69	-0.0173
Panel C: High Frequency Shrinkage Estimators												
RC_{78s}^{RM}	0.0849	0.0810	0.0884	0.77	-0.0166	-0.08	0.22	0.0340	8	0.0812	0.73	-0.0153
RC_{78s}^{LW}	0.0933	0.0888	0.0979	0.84	-0.0199	-0.06	0.27	0.0598	12	0.0629	0.57	-0.0173

Table 5: **Dynamic GMV Portfolio Sub-Period Performance**

The out-of-sample (day $t+1$) annualized portfolio volatility (σ_P) is reported for the benchmark equally-weighted portfolio and for the dynamic GMV portfolios optimized each day t using the covariance estimator indicated in the first column. The volatility is provided for the full sample period and for each of the four calendar years in the sample, as indicated, and for market volatility level sub-samples with "Low" ("High") volatility days defined as days with an absolute return less (greater) than the median absolute return on the equally-weighted benchmark portfolio. The full sample period is January 2, 2003 - December 31, 2006.

Model	**Full Sample**	**2003**	**2004**	**2005**	**2006**	**Volatility Level Low**	**Volatility Level High**
EW	0.1107	0.1534	0.0965	0.0912	0.0897	0.0365	0.1523
Panel A: Sample Estimators							
RC_{78}	0.1091	0.1376	0.1061	0.0980	0.0857	0.0806	0.1316
RC_{78s}	0.0987	0.1246	0.0989	0.0868	0.0766	0.0679	0.1220
OC_{78}	0.1075	0.1350	0.1051	0.0998	0.0835	0.0821	0.1278
OC_{250}	0.0959	0.1218	0.0932	0.0866	0.0764	0.0673	0.1180
Panel B: High Frequency Single Factor Models							
RC_{78s}^{SF}	0.0922	0.1140	0.0927	0.0843	0.0708	0.0580	0.1168
$RC_{78s}^{SF,RM}$	0.0850	0.1020	0.0857	0.0798	0.0680	0.0548	0.1071
$RC_{78s}^{SF,ARMA}$	0.0872	0.1096	0.0866	0.0792	0.0668	0.0551	0.1102
RC_{78s}^{MF}	0.0872	0.1041	0.0871	0.0826	0.0714	0.0631	0.1061
Panel C: High Frequency Shrinkage Estimators							
RC_{78s}^{RM}	0.0849	0.1033	0.0851	0.0791	0.0671	0.0525	0.1080
RC_{78s}^{LW}	0.0933	0.1170	0.0945	0.0834	0.0715	0.0601	0.1176

Table 6: **Mincer-Zarnowitz Forecast Quality Regressions**

Statistics are reported for the regression of the sub-sampled realized equally-weighted portfolio volatility for day t on the forecast generated on day $t-1$ for the same portfolio using the estimator indicated in the first column. For each single forecast regression the estimates and the robust Newey-West standard errors ("s.e.") are provided for the intercept (b_0) and the slope coefficient (b_1) along with the R^2. The R^2 for an encompassing Mincer-Zarnowitz regression with RC_{78s} as the encompassing forecast is reported in the final column. The sample period is January 2, 2003 - December 31, 2006. There are 997 observations.

	Single Forecast Regressions					$\mathbf{RC_{78s}}$ Encomp. Regressions
Model	$\boldsymbol{b_0}$	**s.e.**	$\boldsymbol{b_1}$	**s.e.**	$\mathbf{R^2}$	$\mathbf{R^2}$
Panel A: Sample Estimators						
RC_{78}	0.001647	0.000241	0.736537	0.043551	0.54	0.57
RC_{78s}	0.001522	0.000220	0.754240	0.039952	0.57	-
OC_{78}	0.001663	0.000379	0.643869	0.055473	0.52	0.63
OC_{250}	0.002708	0.000387	0.430082	0.052002	0.45	0.62
Panel B: High Frequency Single Factor Models						
RC_{78s}^{SF}	0.001633	0.000225	0.767534	0.042758	0.57	0.57
$RC_{78s}^{SF,ARMA}$	0.000663	0.000172	0.883619	0.028845	0.66	0.66
RC_{78s}^{MF}	0.001417	0.000223	0.734793	0.037532	0.57	0.58
$RC_{78s}^{SF,RM}$	0.000261	0.000320	0.958125	0.054590	0.60	0.65
Panel C: High Frequency Shrinkage Estimators						
RC_{78s}^{RM}	0.000180	0.000323	0.933463	0.052788	0.60	0.65
RC_{78s}^{LW}	0.001460	0.000224	0.754383	0.040152	0.57	0.57

Table 7: **Robustness Analysis of Ledoit-Wolf Smoothing Parameters**
Volatility and portfolio weight statistics are reported for the dynamic GMV portfolios based on Ledoit-Wolf covariance estimators using the smoothing parameter, α, indicated in the first column. The out-of-sample ("Out") and in-sample ("In") volatilities (σ_P) are the annualized standard deviation of the open-to-close day $t+1$ and day t returns, respectively, for the portfolio. "Med. Min" ("Med. Max") is the median minimum (maximum) weight across the time series. The sample period is January 2, 2003 - December 31, 2006.

			Weights	
	σ_P		**Med.**	**Med.**
α	**Out**	**In**	**Min**	**Max**
.1	0.0964	0.0807	-0.04	0.26
.2	0.0958	0.0791	-0.04	0.26
.3	0.0955	0.0781	-0.05	0.27
.4	0.0956	0.0774	-0.05	0.27
.5	0.0960	0.0772	-0.06	0.27
.6	0.0968	0.0773	-0.07	0.28
.7	0.0983	0.0779	-0.08	0.28
.8	0.1020	0.0805	-0.10	0.29
.9	0.1441	0.0849	-0.12	0.30