BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2007-20

Multivariate Realized Stock Market Volatility

by Gregory H. Bauer and Keith Vorkink

www.bankofcanada.ca

Bank of Canada Working Paper 2007-20

March 2007

Multivariate Realized Stock Market Volatility

by

Gregory H. Bauer[1] and Keith Vorkink[2]

[1]Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
gbauer@bankofcanada.ca

[2]Sloan School of Management
MIT
Cambridge, MA 02142 USA
kvorkink@mit.edu

ISSN 1701-9397

Acknowledgements

We thank Robert Engle, Scott Hendry, Christopher Jones, Nour Meddahi, Miwako Nitani, Denis Pelletier and seminar participants at CIRANO-CIREQ Workshop on Financial Econometrics, the CIREQ Conference on Realized Volatility, and the Northern Finance Association meeting for helpful suggestions. Greg Adams, Scott Gusba, Stephan Mercier, Elena Nemykina and Tiago Rubin provided excellent research assistance.

Abstract

We present a new matrix-logarithm model of the realized covariance matrix of stock returns. The model uses latent factors which are functions of both lagged volatility and returns. The model has several advantages: it is parsimonious; it does not require imposing parameter restrictions; and, it results in a positive-definite covariance matrix. We apply the model to the covariance matrix of size-sorted stock returns and find that two factors are sufficient to capture most of the dynamics. We also introduce a new method to track an index using our model of the realized volatility covariance matrix.

JEL classification: G14, C53, C32
Bank classification: Econometric and statistical methods; Financial markets

Résumé

Les auteurs présentent un nouveau modèle de la matrice des covariances réalisées des rendements boursiers dans lequel la matrice est exprimée sous forme logarithmique et les facteurs latents sont fonction à la fois de la volatilité passée et des rendements historiques. Le modèle offre plusieurs avantages : il est parcimonieux, il ne nécessite pas l'imposition de restrictions sur les paramètres et il produit une matrice des covariances définie positive. L'application du modèle à la prévision de la matrice des covariances des rendements classés selon la taille de l'entreprise fait ressortir que deux facteurs suffisent pour rendre compte de l'essentiel de la dynamique. Les auteurs proposent aussi une méthode permettant de reproduire l'évolution d'un indice à l'aide de leur modèle de la matrice des covariances réalisées.

Classification JEL : G14, C53, C32
Classification de la Banque : Méthodes économétriques et statistiques; Marchés financiers

1 Introduction

The covariance matrix of stock returns resides at the center of a number of important concepts in financial economics. The covariance of a stock's return with other stocks forms the basis of the CAPM and other asset pricing models. The covariance matrix is used in designing tracking strategies, where the portfolio manager attempts to closely follow the return on a benchmark portfolio. The matrix is also used for risk management measures such as Value at Risk. An assessment of market stability and contagion depends on measuring time-varying volatilities and correlations. Finally, accurate measures of covariances are required for corporate hedging strategies.

One key stylized fact in empirical finance is that the variances and covariances of stock returns vary over time.[1] As a result, many important financial applications require a model of the conditional covariance matrix. Three distinct categories of methods for estimating a conditional covariance matrix have evolved in the literature. In the first, and probably best known category, are the various forms of the multivariate GARCH model where forecasts of future volatility depend on past volatility and shocks.[2] In the second category, authors have modeled asset return variances and covariances as functions of a number of predetermined variables.[3] The third category includes multivariate stochastic volatility models.[4]

Existing approaches to modeling the conditional covariance matrix suffer from three primary limitations. First, in the extant models, the covariances of asset returns are treated as an unobserved or latent process. However, the availability of high-frequency data has spawned the use of 'realized volatility' modeling, allowing more precise estimates of the volatility process to be constructed.[5] While the majority of the existing papers have focused on estimates of an individual realized volatility series, we investigate methods for analyzing the realized conditional covariance matrix. Second, the existing approaches to modeling multivariate latent volatility require a number of restrictions on the model's parameters to ensure that the estimated covariance matrix is positive definite. These

[1] For a very comprehensive survey of the literature on volatility modeling and forecasting, see Andersen et al. (2005b).

[2] See Bauwens et al. (2004) for a recent survey of this literature.

[3] Examples include Campbell (1987) and Harvey (1989). See Ferson (1995) for a survey of this literature.

[4] Asai et al. (2005) survey the multivariate stochastic volatility literature.

[5] Using daily data to construct estimates of monthly volatility originated with French, Schwert and Stambaugh (1989), and Schwert (1989, 1990). Andersen and Bollerslev (1998) introduced the idea of using microstructure data to construct estimates of daily 'realized volatility'. Andersen et al. (2003) formalized the definition which was applied to equity markets in Andersen et al. (2001) and exchange rates in Andersen et al. (2001).

restrictions may confound accurate estimates of the drivers of conditional covariances. Third, the existing approaches model volatility either as a function of past volatility or as a function of a number of predetermined variables. Yet, the conditional covariance matrix may be a function of both.

In this paper, we introduce a new model of the *realized* covariance matrix that overcomes these three limitations. Regarding problem one, unobservability, we use high-frequency data to construct estimates of the daily variances and covariances of five size-sorted stock portfolios. By using high-frequency data we obtain an estimate of the matrix of 'quadratic variations and covariations' that differs from the true conditional covariance matrix by mean zero errors (e.g. Andersen et al. (2001), (2003)). Thus, we can treat our conditional covariance matrix not as latent, but observed. This implies that very accurate estimates of the factors driving the conditional covariances can be found.

We overcome problem two, cumbersome parameter restrictions to ensure a positive definite covariance matrix, by transforming the realized covariance matrix using the matrix logarithm function to yield a series of transformed volatilities which we term the *log-space volatilities*. The matrix logarithm is a non-linear function of all of the elements of the covariance matrix and thus, the log-space volatilities do not correspond one-to-one with their counterparts in the realized covariance matrix.[6] However, modelling the time variation of the log-space volatilities is straightforward and avoids the problems that plague existing estimators of the latent volatility matrix.

We eliminate problem three by modelling the dynamics of the log-space volatility matrix using functions of both past volatilities and other factors that can help forecast future volatility. The model is estimated by Generalized Method of Moments (GMM) yielding a series of fitted values. We then transform these fitted values, using the matrix exponential function, back into forecasts of the realized covariance matrix. Our estimated matrix is positive definite by construction and does not require any parameter restrictions to be imposed. The approach can thus be viewed as a multivariate version of standard stochastic variance models, where the variance is an exponential function of the factors and the associated parameters.

In addition to introducing our new realized covariance matrix we also test the forecasting ability of alternative variables for time-varying equity market covariances. Our motivation is that researchers have examined a number of variables for forecasting returns but there is much less evidence that the variables forecast risks. The cross-section of small- and large-firm volatility has been examined in a number of earlier papers (Con-

[6] The matrix logarithm has been used for estimators of latent volatility by Chiu, Leonard and Tsui (1996) and Kawakatsu (2003) and was also suggested in Asai et al. (2005). However, to the best of our knowledge, we are the first to use it for realized covariance modeling.

rad, Gultekin and Kaul (1991), Kroner and Ng (1998), Chan, Karceski and Lakonishok (1999), and Moskowitz (2003)). However, these papers used models of latent volatility to capture the variation in the covariances.[7] In contrast, we use quote-by-quote data to construct daily measures of the realized covariance matrix of small and large firms over the 1988 to 2002 period. Our measures of volatility are more precise than those in previous work and allow a detailed examination of the drivers of conditional covariances.

We use three sets of forecasting variables from the existing literature to capture movements in the covariance matrix. The first set contains lagged realized variances and covariances only. Following Corsi (2004) and Andersen, Bollerslev and Diebold (2003), we use past weekly and monthly realized volatilities in a Heterogeneous Autoregressive model of realized volatility (HAR-RV). We implement a multivariate version of their model by using our matrix-logarithm transformation and by estimating latent volatility factors. Our factor model can thus be viewed as a multivariate HAR model of realized volatility (MHAR-RV).

The second set of forecasting variables adds lagged stock returns to the first set. The asymmetric relationship between realized returns and subsequent equity market volatility is well documented (Black (1976), Pagan and Schwert (1990), Engle and Ng (1993), Bollerslev, Litvinova and Tauchen (2005)) and is easily handled in our model.[8] The third set of forecasting variables uses the 'usual suspects' that have been shown to have predictive power for equity market returns (e.g., dividend yields, term structure slopes, etc.). We add to that literature by showing that these variables have predictive power for conditional volatilities.[9]

We evaluate our models of the conditional covariance matrix in two ways. The first is a set of standard econometric tests of the fit of the log-space and actual volatilities.

---

[7] Schwert and Seguin (1990) use daily return data to construct measures of monthly stock volatility. They find that a single index model describes the cross section of volatility in contrast to the multi-factor model found here.

[8] A few papers have analyzed the asymmetric relationship in a multivariate setting. Conrad, Gultekin and Kaul (1991) present a new GARCH model and show that surprises to large firm returns have an effect on small firm volatility but the reverse is not true. Kroner and Ng (1998) also introduce a new GARCH model of latent volatility and find that negative large firm returns affect future small-firm volatility. Chan, Karceski and Lakonishok (1999) and Moskowitz (2003) examine whether the Fama and French (1992) size factor can capture the cross section of stock volatility.

[9] A number of papers have examined market volatility and its relationship to various forecasting variables. Schwert (1989) examines the ability of several financial and macroeconomic factors to forecast monthly volatility over the 1859 to 1987 period. Attanasio (1991) examines the relationship between dividend yield and market volatility. Glosten, Jagannathan, and Runkle (1993) use the one-month Treasury bill yield in an augmented GARCH model. Whitelaw (1994) examines how the yield spread (Baa-Aaa rated), commercial paper-Treasury yield spread, one year Treasury yield and S&P dividend yield forecast monthly measures of volatility based on daily returns. Shanken and Tamayo (2004) present a Bayesian analysis of the effects of dividend yields on latent volatility.

The second is to analyze the dynamics of the global minimum variance portfolio and a tracking error portfolio. The portfolio analysis provides an economic comparison across alternative specifications and reveals interesting differences among the variables used. In fact, our tests suggest that including variables that forecast stock returns, along with lagged volatility factors, produce portfolios with superior performance.

Our method of estimating the conditional realized covariance matrix has a number of advantages when compared to existing approaches. First, standard GARCH type models of latent volatility can yield volatility estimates that are quite noisy (Andersen and Bollerslev (1998)). Our estimator is based on realized volatility constructed using high-frequency data. Thus, our measures of equity market variances and covariances are more precise. This allows us greater power in determining the effects of alternative forecasting variables on equity market volatility when compared to earlier efforts based on latent volatility models.

Second, existing models require numerous constraints on the parameters to yield a positive-definite covariance matrix. However, the constrained values of the parameters may not reveal the true influence of the forecasting variables. In contrast, we do not impose any constraints on our estimates of the log-space volatilities. Using the matrix exponential function results in estimates of the covariance matrix that are positive semi-definite by construction.

Third, most existing models require a large number of parameters to be estimated, making practical implementation difficult. By applying a factor approach to the log-space volatilities, we are able to model the conditional covariance matrix using a relatively small number of variables. Below, we show that two factors can capture the volatility dynamics of the size-sorted portfolios.[10]

Fourth, existing studies have usually taken two extreme approaches to modeling volatility. The first approach focuses on the autoregressive nature of volatility via ARCH and GARCH type effects and usually ignores the effect of other forecasting variables. The second ignores the autoregressive nature of volatility and uses variables from the expected return forecasting literature to capture time variation in volatilities. In our approach, we can easily incorporate both sources of dynamics into our estimates.

The final advantage of our model is its ability to characterize the volatility of volatility. A number of authors have noted that volatility process itself depends on the level of volatility in the market. For example, both Das and Sundaram (1999) and Jones (2003) show that the levels of skewness and kurtosis displayed by equity index data are both functions of the level of volatility in the stock market. However, to the best of our

[10] In addition, the factor structure allows us to capture the long-memory property of realized volatility measure. We return to the issue of modeling long memory in volatility below.

knowledge, no one has yet determined the economic variables that cause volatility to be volatile. Using our factor model estimates, we can obtain the derivatives of the realized covariance matrix with respect to the variables in the model.[11] We are able to calculate the derivatives at each point in our sample, yielding a series of conditional volatility elasticities that are functions of both the level of the volatility and the factors driving the volatility. This time series allows us to characterize the time-varying elasticities of expected volatility with respect to its underlying factors.

Naturally all of these advantages come at a cost. The main cost is that by performing our analysis on the log-space volatilities and then using the (non-linear) matrix exponential function, the estimated volatilities will not be unbiased.[12] However, as we show below, a simple bias correction is available that greatly reduces the problem. Another cost is that direct interpretation of the effects of an instrument on expected volatility is difficult due to the non-linear nature of the model. However, the volatility elasticities show the ultimate impact of any variable of interest.

Our paper is part of the growing literature examining the dynamics of realized volatility. Most of these papers look at the volatility of a single series. Andersen et al. (2001) construct multivariate measures of realized volatility in two foreign exchange rate markets as well as in the "cross market." They are able to model a conditional covariance matrix by relying on the triangular arbitrage between three spot exchange rates. We note that this approach would not work for general assets whose spot market prices are not constrained by arbitrage relationships such as the stock market portfolios that we examine. Andersen et al. (2004) construct quarterly estimates of realized variances, covariances and stock-market betas using daily data. They focus on the time-series predictability of betas for individual firms.

The paper is organized as follows. In section 2, we present our model of matrix logarithmic realized volatility. In section 3, we outline our method for constructing the realized volatility matrices and give the sources of the data. In section 4, we give our results. In section 5, we evaluate the alternative models using the global minimum variance and tracking error portfolios criteria. In section 6, we conclude.

[11] The derivatives are based on calculations presented in Najfeld and Havel (1995) and Mathias (1996).

[12] However, any univariate estimator that models log volatility will produce biased estimates of volatility in the real space.

2 Model

2.1 The matrix log transformation

In this paper, we use the matrix exponential and matrix logarithm functions to model the time-varying covariance matrix. The matrix exponential function performs a power series expansion on a square matrix A

$$V = \mathrm{expm}(A) \equiv \sum_{n=0}^{\infty} \frac{1}{n!} A^n. \tag{1}$$

The matrix exponential function has a number of useful properties (Chiu, Leonard and Tsui (1996)).[13] The most important of these is that if A is a real, symmetric matrix, then V is a real, positive definite matrix. The converse is also true and describes one of the useful properties of the matrix logarithm function. The matrix logarithm function is the inverse of the matrix exponential function. Taking the matrix logarithm of a real, positive definite matrix V results in a real, symmetric matrix A:

$$A = \mathrm{logm}(V).$$

The matrix logarithm and matrix exponential functions are used in our three-step procedure to obtain forecasts of the conditional covariance matrix of stock returns. In the first step, for each day t, we use high-frequency (quote-by-quote) data to construct the $P \times P$ realized conditional covariance matrix V_t.[14] The V_t matrix is positive definite by construction. Applying the matrix logarithm function,

$$A_t = \mathrm{logm}(V_t), \tag{2}$$

yields a real, symmetric $P \times P$ matrix A_t. We term the elements of A_t the "log-space volatilities" and note that the elements of A_t do not correspond one-to-one with the elements of V_t. Thus, for example, the $(1,1)$ element of A_t is not the log volatility of the first portfolio.

In the second step, we model the dynamics of the A_t matrix. To do this, we follow Chiu, Leonard and Tsui (1996) and apply the vech operator to the matrix A_t

$$a_t = vech(A_t),$$

which stacks the elements on and below the diagonal of A_t to obtain the $p \times 1$ vector a_t, where $p = \frac{1}{2}P(P+1)$. We note a_t contains all of the unique elements of the symmetric A_t matrix.

[13] The matrix exponential function has a long history in physics (e.g., Thompson (1965)) and has recently been used in modeling spatially dependent data by LeSage and Pace (2003).

[14] The details of how the matrix is constructed are presented below.

The a_t vector forms the basis for all subsequent models. For example, to model time-varying volatilities based on their past values, we could estimate a simple first-order vector auto-regression of the a_t processes:

$$a_t = \gamma_0 + \gamma_1 a_{t-1} + \varepsilon_t, \tag{3}$$

where γ_0 is a $p \times 1$ vector of intercepts, γ_1 is a $p \times p$ matrix of coefficients, and ε_t is a vector of residuals. This model would yield a series of fitted values

$$\widehat{a}_t \equiv E_t(a_t) = \widehat{\gamma}_0 + \widehat{\gamma}_1 a_{t-1}.$$

Below, we present a factor model for the a_t processes which has a much smaller number of parameters to be estimated and allows other variables to forecast volatility.

In the third step, we transform the fitted values from the log-volatility space into fitted values in the actual volatility space. We use the inverse of the *vech* function to form a $P \times P$ symmetric matrix $\widehat{A}_t$ of the fitted values at each time t from the vector $\widehat{a}_t$. Applying the matrix exponential function

$$\widehat{V}_t = \text{expm}(\widehat{A}_t), \tag{4}$$

yields the matrix $\widehat{V}_t$, which is our estimate of the conditional covariance matrix for day t. We note that as long as the elements of $\widehat{A}_t$ are real, then $\widehat{V}_t$ is positive definite by construction (Chiu, Leonard and Tsui (1996)).

2.2 A small example

A small example adapted from Williams (1999) can help clarify how our matrix logarithm model works. Suppose that the time-varying volatility matrix for two assets has the form

$$V_t = \begin{bmatrix} 1 & \rho_t \\ \rho_t & 1 \end{bmatrix}, \tag{5}$$

so that the volatilities are constant but the correlation is time varying. As V_t is a positive-definite, symmetric matrix then there exists matrices B_t and D_t such that $V_t = B_t D_t B_t^T$. B_t is an orthonormal 2×2 matrix with the eigenvectors of V_t in its columns, while D_t is a diagonal matrix with the eigenvalues on its diagonal. Then, by the properties of the matrix logarithm function, we have

$$A_t = \text{logm}(V_t) = B_t \log(D_t) B_t'.$$

The matrix logarithm function takes the log of the eigenvalues of the matrix while leaving the eigenvectors intact.[15]

[15] Chiu, Leonard and Tsui (1996) discuss a number of properties of matrix logarithm function.

Taking the matrix logarithm of our simple example (5) yields

$$A_t = \text{logm}(V_t) = \frac{1}{2}\begin{bmatrix} \log(1-\rho_t^2) & \log(\frac{1+\rho_t}{1-\rho_t}) \\ \log(\frac{1+\rho_t}{1-\rho_t}) & \log(1-\rho_t^2) \end{bmatrix}.$$

Estimation of the time-series process of $a_t = \left\{0.5\log(1-\rho_t^2), 0.5\log(\frac{1+\rho_t}{1-\rho_t})\right\}$ proceeds with no need for parameter restrictions on the user specified model. Using the fitted values from the estimation, we can form the matrix

$$\widehat{A}_t = \frac{1}{2}\begin{bmatrix} \widehat{\log(1-\rho_t^2)} & \widehat{\log(\frac{1+\rho_t}{1-\rho_t})} \\ \widehat{\log(\frac{1+\rho_t}{1-\rho_t})} & \widehat{\log(1-\rho_t^2)} \end{bmatrix}. \tag{6}$$

Taking the matrix exponential of (6) yields the estimate of $\widehat{V}_t$. We note that restrictions on the parameters are not required except that the elements of $\widehat{A}_t$ must be real. Given this, the resulting $\widehat{V}_t$ matrix is positive definite.

2.3 Factor models of volatility

As mentioned above, we will use several different groups of variables to forecast the conditional covariances. Based on the existing literature, we can separate the variables into two groups. The first are matrix-log values of realized volatility $(a_t, a_{t-1}, a_{t-2}, ...)$ which are used to capture the autoregressive nature of the volatility. We note that the existing literature shows that capturing volatility dynamics will likely require a long lag structure. The second group are those variables that have been shown to forecast equity market returns, X_t. In equilibrium, expected returns should be related to risk, so it is natural to question whether these variables also forecast the components of market wide volatility. We denote the augmented matrix of the forecasting variables as $Z_t = (a_t, a_{t-1}, a_{t-2}, ...\ X_t)$ and note that Z_t will differ depending on the information set chosen.

The simplest approach to modeling variation in the log-space transformation of the conditional covariance matrix is then

$$a_t = \gamma_0 + \gamma_1 a_{t-1} + \gamma_2 a_{t-2} + ... + \gamma_k a_{t-L} + \gamma_X X_{t-1} + \varepsilon_t, \tag{7}$$

where log-space volatility is modeled as a function of L lags of volatility plus the predetermined forecast variables. In this specification, there are a large number of parameters in the $\gamma_1, ..., \gamma_L$ and γ_X matrixes: $L \cdot p^2$ in the former and $p \cdot M$, where M is the number of variables in X_{t-1}, in the latter.[16] It would be preferable to adopt a specification with

[16] For example, to model the 5×5 matrix V_t of the size sorted-stock portfolios would require 15 parameters in γ_0, $225 = 15^2$ parameters in each of the γ_L, and $15{\cdot}M$ parameters in γ_X.

fewer parameters, *ceterus paribus*. To do this, we use a sequence of three dimension reduction techniques.

The first technique is designed to reduce the number of lagged log-space volatility variables by adapting the Heterogeneous Autoregressive model of realized volatility (HAR-RV) of Corsi (2004) and Andersen, Bollerslev and Diebold (2003) to a multivariate setting. These authors show that aggregate market realized volatility is best forecast by a (linear) combination of lagged daily, weekly and monthly realized volatility. Other authors have indicated that lagged realized volatility may not be the best predictor, however. In particular, both Andersen, Bollerslev and Diebold (2003) and Ghysels, Santa-Clara and Valkanov (2004b) find that bi-power covariation – an estimate of the continuous part of the volatility diffusion that is defined below – is a good predictor of aggregate market realized volatility.[17]

Following the approaches in these papers, we construct daily, weekly and monthly multivariate bi-power covariation as summary measures of lagged volatility (details below). As in (2) above, we take the matrix logarithm of the bi-power covariation matrix over the past d days to yield $A^{BP}(d)_t$. Taking the *vech* of this matrix yields the unique elements $a^{BP}(d)_t$, which can act as forecasting variables. We can then replace the large number of lagged log-space volatilities $a_{t-1}, ..., a_{t-k}$ on the right hand side of (7) to yield:

$$a_t = \gamma_0 + \gamma_1 a^{BP}(1)_{t-1} + \gamma_5 a^{BP}(5)_{t-1} + \gamma_{20} a^{BP}(20)_{t-1} + \gamma_X X_{t-1} + \varepsilon_t, \qquad (8)$$

where $a^{BP}(1)_t$, $a^{BP}(5)_t$, and $a^{BP}(20)_t$ are the matrix-logarithms of daily, weekly and monthly multivariate bi-power covariation, respectively. This specification reduces the number of parameters from (7) by $(L-3) \cdot p^2$.[18] Our approach (8) can be viewed as a multivariate approach to the HAR-RV model.

While multivariate HAR-RV approach reduces the number of parameters considerably, there is still a large number of variables in the $p \times p$ matrixes $a^{BP}(d)_t$. While the V_t matrix is full rank, it is likely that the volatilities, and hence the bi-power covariation, are driven by a smaller number of components. Our second dimension reduction technique is thus based on the (testable) hypothesis that the $a^{BP}(d)_t$ series are driven by a smaller number of factors. We test this by estimating the principal components of $a^{BP}(d)_t$,

$$a^{BP}(d,i), \quad i = 1, ..., pc, \qquad (9)$$

[17] Andersen, Bollerslev and Diebold (2003) also find that jumps help predict future volatility. We have tried some preliminary analysis using multivariate jumps and find that while they are statistically significant predictors of future daily volatility, they have little explanatory power.

[18] The reduction in parameters is quite large considering $p^2 = 225$ and a large number of k potential lags.

where $a^{BP}(d,i)$ is the i^{th} principal component of the d-day log-space bi-power covariation matrix. We find that a small number of components capture the volatility of the daily, weekly and monthly log-space bi-power covariation series (results below). We replace these variables in (8) with a small number of their principal components (9) to yield:

$$\begin{aligned} a_t &= \gamma_0 + \gamma_{1,1} a^{BP}(1,1)_{t-1} + \gamma_{1,2} a^{BP}(1,2)_{t-1} + \gamma_{5,1} a^{BP}(5,1)_{t-1} \\ &\quad + \gamma_{5,2} a^{BP}(5,2)_{t-1} + \gamma_{20,1} a^{BP}(20,1)_{t-1} + \gamma_{20,2} a^{BP}(20,2)_{t-1} + ... + \gamma_X X_{t-1} + \varepsilon_t. \end{aligned} \tag{10}$$

Although this specification appears complicated, it reduces the number of parameters dramatically as we have replaced the $a^{BP}(d)_t$ matrixes (each of which is $p \times p$) with their $a^{BP}(d,i)_t$ counterparts (each of which is $p \times 1$).[19] In our models below, we find that including a small number of principal components ($pc < p$) is sufficient to model the realized covariance matrix.

The first two techniques considerably reduce the number of parameters required to estimate the effects of lagged volatility on future volatility. However, there will still be a large number of parameters associated with the combined $\left[\gamma_{1,1}, ..., \gamma_{5,1}, ..., \gamma_{20,1}, ..., \gamma_X\right]$ matrix. The third dimension reduction technique is to use a latent factor approach where the factors that drive the time-varying volatility are not specified directly. Rather, we assume that our set of forecasting variables in (10)

$$Z_t = (a^{BP}(1,1)_t, ..., a^{BP}(5,1)_t, ..., a^{BP}(20,1)_t, ..., X_t),$$

is related to the true, but unknown, volatility factors. We thus specify the k-th volatility factor $\upsilon_{k,t}$ as a linear combination of the set of N variables Z_t:

$$\upsilon_{k,t} = \theta_k Z_{t-1}, \tag{11}$$

where the $\theta_k = \{\theta_{k,(1)}, ..., \theta_{k,(N)}\}$ are coefficients that aggregate the forecasting variables in Z_t. Each of the log-space volatilities a_t^i is a function of the K volatility factors:

$$a_t^i = \gamma_0^i + \beta^i \theta Z_{t-1} + \varepsilon_t^i, \quad i = 1, ..., p,$$

where γ_0^i is the i^{th} element of the intercept vector γ_0, β^i is the $1 \times K$ vector of the loadings of log-space volatility i on the K factors, and the $K \times N\theta$ matrix contains the coefficients on the Z_{t-1} variables for the K factors. Assembling the model for all p log-space volatilities yields

$$a_t = \gamma_0 + \beta\theta Z_{t-1} + \varepsilon_t, \tag{12}$$

[19] The $\gamma_{d,i}$ parameters in (10) are each $p \times 1$, while the γ_d coefficinet matrixes in (8) are each $p \times p$.

where the $p \times K$ matrix β is the loading of the log-space volatilities on the time-varying factors.

We note that using latent factors to model covariance matrixes have a number of advantages over existing methods. First, it allows us to combine both lagged volatility measures (the principal components in (9)) as well as the X_t variables in a parsimonious manner. Previous models required each variable to be a separate factor. While the large number of variables may help forecast the covariance matrix, it is unlikely that each variable represents a specific volatility factor.[20] Our approach can be used to weigh (via the θ coefficients) all of the variables in a way that is optimal for forecasting the covariance matrix.

A second advantage to our approach is that it avoids using expected returns in modeling the volatility matrix. Aggregating squared return or bi-power covariation data over high frequencies means that expected return variation can be ignored. Thus, we do not need to rely on expected returns to obtain the loadings on the factors as in Chan, Karceski and Lakonishok (1999) or Moskowitz (2003). As the realized covariance matrix can be estimated more precisely than can expected returns, we should obtain more precise measures of the determinants of the covariance matrix.

The third advantage is parsimony. For example, assume that we require 20 lags of daily log-space bi-power covariation plus 5 forecasting variables in X_t to capture volatility dynamics in our 5×5 volatility matrix. The number of parameters in the base line model (7) would be $4,590$ while a $K = 2$ factor version of (12) using the first three principal components of the log-space bi-power covariation matrixes (for $d = 1, 5$ and 20 days) has only 69. The small number of parameters in the factor model helps in estimating and interpreting the model in-sample and should help in out-of-sample forecasting.

Our approach using latent factors to model volatility dynamics is related to previous work. Diebold and Nerlove (1989) propose a factor ARCH model of the cross section of exchange rate changes. In their model, a single (latent) factor captures the common variation in the three exchange rates.[21] In addition, a number of papers have suggested that two or more factors are necessary to capture the dynamics of a single latent volatility series. Ding and Granger (1996) propose an N component model to capture the long memory in stock and foreign exchange volatility. They need to let N get very large to capture the long-run properties. Engle and Lee (1999) suggest that two factors are necessary to capture the dynamics of stock return volatility. One factor is very long lived while the other is a more quickly mean-reverting short-term factor. Gallant, Hsu and

[20] For example, Chan, Karceski and Lakonishok (1999) find little difference between the forecasting ability of volatility models with 3 and10 factors.

[21] See also King, Sentana and Wadhwani (1994) and Harvey, Ruiz and Shephard (1994).

Tauchen (1999) estimate a two factor model of volatility using the daily range to capture volatility dynamics. They also find both a long-run and short-run volatility component. Andersen and Bollerslev (1997) show that foreign exchange volatility can be modeled by a mixture process, where the short-run volatility dynamics are driven by "news".

2.4 Estimation

Our multivariate factor model is derived from the latent factor models of expected return variation that originated with Hansen and Hodrick (1983) and Gibbons and Ferson (1985). As in these papers, we estimate our factor model of volatility in (12) by GMM with the Newey-West (1987) form of the optimal weighting matrix.[22] In its present form, (12) is unidentified due to the $\beta\theta$ combination. We thus impose the standard identification that the first K rows of the matrix β are equal to an identity matrix. The cross-equation restrictions imposed on (10):

$$H_0: \quad \left[\gamma_{1,1}, ..., \gamma_{5,1}, ..., \gamma_{20,1}, ..., \gamma_X\right] = \beta\theta \tag{13}$$

can then be tested using the standard χ^2 test statistic from a GMM system.

Our model has a potential errors-in-variables problem as the log-space bi-power covariation matrix $A^{BP}(d)_t$ is constructed with error. Using its principal components as regressors will result in biased estimates of the coefficients. Ghysels and Jacquier (2005) have noted a similar problem with estimates of time-varying beta coefficients for portfolio selection. They advocate using lagged values of the betas in an instrumental variables regression to overcome the biases. We follow that approach here and use the twice lagged values of the principal components in the GMM instrument set.

Once the coefficients have been estimated by GMM, the fitted values are reassembled into a square matrix $\widehat{A}_t$. Applying the matrix exponential function as in (4) yields the prediction for the covariance matrix in period t,

$$\widehat{V}_t = \text{expm}\left(\widehat{A}_t\right). \tag{14}$$

We can then apply standard forecasting evaluation techniques to compare $\widehat{V}_t$ to V_t.

Our $\widehat{V}_t$ estimator will be biased as the estimation is done in the log-volatility space. We correct for this bias with a simple fix that is discussed in the appendix.

[22] We use 5 lags in the Newey-West standard errors to account for any autocorrelation in the residuals. The results are robust to this choice.

2.5 Interpreting expected volatility

The matrix logarithmic volatility model has the disadvantage that the estimated coefficients cannot be interpreted directly as the effect of the variable on the specified element of the realized volatility. This results from the non-linear relationship between particular elements of $\widehat{V}_t$ and $\widehat{A}_t$. However, derivatives of the estimated covariance matrix $\widehat{V}_t$ with respect to the elements of the factor model can be easily obtained (Najfeld and Havel 1995 and Mathias 1996).

Let $\widehat{A}(z)$ be the $P \times P$ expected conditional covariance matrix from the third step (4) of our estimation procedure, where we consider the matrix to be a function of a particular forecasting variable, say $z \in Z$. The matrix of the element-by-element derivatives of $\widehat{A}(z)$ with respect to z, $\frac{d\widehat{A}(z)}{dz}$, is calculated using the estimated coefficients from our factor model (12). The $P \times P$ matrix of the derivatives of the actual volatilities with respect to z,

$$\frac{d}{dz}\widehat{V}(z) = \frac{d}{dz}\operatorname{expm}\left(\widehat{A}(z)\right),$$

can be extracted from the upper $P \times P$ right block of the following $2P \times 2P$ matrix:

$$\begin{bmatrix} \operatorname{expm}\left(\widehat{A}(z)\right) & \frac{d}{dz}\operatorname{expm}\left(\widehat{A}(z)\right) \\ 0 & \operatorname{expm}\left(\widehat{A}(z)\right) \end{bmatrix} = \operatorname{expm}\begin{bmatrix} \widehat{A}(z) & \frac{d\widehat{A}(z)}{dz} \\ 0 & \widehat{A}(z) \end{bmatrix}, \tag{15}$$

where 0 represents a $P \times P$ matrix of zeros. Equation (15) allows the observation and interpretation of the impact of the forecasting variables on the realized covariance matrix, even though the estimation occurs in the matrix-log space.

We can calculate either the average impact across the entire sample, or the conditional impact at a point in time. For example, let $\widehat{A}_t(z_t)$ be the estimated realized log-space volatilities for day t. To find the response of the expected covariance matrix to the forecasting variable z_t, we need to calculate the derivative $\frac{d}{d(z_t)}\operatorname{expm}\left(\widehat{A_t}(z_t)\right)$. Given our two-factor structure as defined in (12), we can represent $\widehat{A}_t$ as functions of both parameters as well as variables, and the matrix of element-by-element derivatives, $\frac{d\widehat{A}_t(z_t)}{d(z_t)}$, is easily obtained. Plugging $\frac{d\widehat{A}_t(z_t)}{d(z_t)}$ into (15) yields the matrix $\frac{d\widehat{V}_t(z_t)}{d(z_t)}$. We can then calculate the time-varying elasticity

$$\epsilon(i,j,z,t) \equiv \frac{d\widehat{V}_t^{(i,j)}(z_t)}{d(z_t)}\frac{\sigma(z_t)}{\widehat{V}_{t(i,j)}}, \tag{16}$$

which represents the percent increase in the $(i,j)^{th}$ element of $\widehat{V}_t$ due to a one standard deviation shock in the forecasting variable, $\sigma(z_t)$, at time t.[23] We can therefore examine how the elasticity of a particular equity market variance or covariance changes over time in response to changes in the forecasting variables. In our results below, we show that there is significant time variation in the elasticities.

3 Data

3.1 Constructing the realized covariance matrix

We construct our realized covariance matrixes from two data sets: the Institute for the Study of Securities Markets' (ISSM) database and the Trades and Quotes (TAQ) database. Both data sets contain continuously recorded information on stock quotes and trades for securities listed on the New York Stock Exchange (NYSE). The ISSM database provides quotes from January 1988 through December 1992 while the TAQ database provides quotes from January 1993 through December 2002.[24]

Following Anderson, Bollerslev, Diebold and Wu (2004), realized covariances for a given day are constructed by summing high-frequency returns as follows:

$$V_t = \sum_{j=1,\ldots,m} r_{t,j\cdot\triangle} r'_{t,j\cdot\triangle}, \tag{17}$$

where V_t is the realized covariance matrix of a set of stocks/portfolios for day t, $r_{t,j\cdot\triangle}$ is the $P \times 1$ vector of high-frequency returns of length $\triangle$ for interval j of day t, and m is the number of intervals in one day. We use our high-frequency portfolio returns to calculate a total of $3,781$ daily realized covariance matrixes.

Value-weighted portfolio returns are created by assigning stocks to one of five size-sorted portfolios based on the prior month's ending price and shares outstanding. Our choice of portfolios is partially motivated by an interest to see if the systematic components of conditional volatility are common across the size portfolios. We use the CRSP database to obtain shares outstanding and prior month ending prices. Only stocks that are found in both the quotes databases (ISSM and TAQ) and CRSP are included in the sample.[25]

---

[23] Measuring the elasticity for the covariance elements ($i \neq j$) is problematic as the covariances can become arbitrarily small. For these elements, we therefore use $\left(\widehat{V}_{t(i,i)}\widehat{V}_{t(j,j)}\right)^{1/2}$ in the denominator of (16).

[24] The ISSM data actually begins in January 1983; however, the first four years of the data have many missing days and the necessity of a contiguous data set for our time-series analysis precludes our use of these years.

[25] We use a variety of other filters that reduces the set of securities included in our data base. For

We use 20 minutes as our high-frequency return interval ($\triangle$) for a number of reasons. First, a well-known trade-off between interval length and microstructure effects exists in high-frequency stock returns (see Campbell, Lo and MacKinlay (1997)). The quadratic variation theory suggest that the finest return interval possible will lead to the most precise measure of covariances; however, microstructure effects such as bid-ask bounce and clustering are magnified by looking at finer return intervals. Our choice of 20 minute return intervals is based on rule-of-thumb suggestions by Anderson, Bollerslev, Diebold and Ebens (2001) for mitigating this trade-off for highly liquid securities. Second, while many of our stocks are illiquid, the fact that we are constructing portfolio realized covariances should diversify away the impact of the bid-ask bounce and other microstructure effects of individual securities. Last, we have investigated a number of return intervals varying from two minutes up to one hour and found the results to be quite robust to the length of $\triangle$.

Many of the stocks in our database trade less frequently than 20 minute intervals and so we follow Anderson and Bollerslev (1997) by constructing artificial equally-spaced returns for stocks by obtaining the closest quotes surrounding a given interval break and using linear interpolation to construct an artificial price at the interval break point.

Once we have our time series of high-frequency portfolio returns, we construct our measure of realized covariance matrixes using a slightly modified version of equation (17) from Hansen and Lunde (2004).[26] Hansen and Lunde (2004) suggest an extension to the usual construction of realized volatility whose intuition is based on the Newey-West (1987) variance estimator. They note that the serially autocorrelated nature of the data is ignored using simple variance calculations such as (17). To incorporate the serial-correlation effects, equation (17) can be extended to

$$V_t = \sum_{j=1}^{m} r_{t,j\cdot\triangle} r'_{t,j\cdot\triangle} + 2\sum_{k=1}^{n}\left(1-\frac{k}{n+1}\right)\sum_{j=1}^{m-k} r_{t,j\cdot\triangle} r'_{t,j+k\cdot\triangle}, \tag{18}$$

where n corresponds to number of lagged intervals where serial correlation may exist. We find that estimates of equation (18) can vary quite dramatically from equation (17) and for good reasons. Our portfolios of smaller stocks will include securities that are more illiquid than stocks in the larger quintiles. The illiquidity of small stocks suggests that prices and volatility responses to information shocks may take more time to be incorporated into prices, leading to time series autocorrelation in the high-frequency

example, we exclude securities with CRSP share codes that are not 10 or 11, leading to the exclusion of preferred stocks, warrants, etc.

[26] Jagannathan and Ma (2003) present a similar idea for constructing monthly volatility measures using daily data.

returns. The method of Hansen and Lunde (2004) corrects the realized covariance matrix for these high-frequency autocorrelations. In robustness checks, we find that our estimation results become quite robust to the choice of $\triangle$, if we estimated realized covariances following equation (18).

We also follow Hansen and Lunde (2005) by including an estimate of the volatility that occurs when the markets are closed. During the close-to-open hours, no data are recorded but the price of the asset will still be responding to news. Hansen and Lunde (2005) provide a way of combining the squared close-to-open return and the high-frequency open-to-close data to yield an estimate of the daily volatility. It is important in our portfolio tests below to have estimates of volatility over the entire day as the returns are over this period.

Table 1a provides summary statistics of the realized covariance matrix V_t. The smallest stocks are labelled portfolio 1 while the largest are labelled 5. The elements are labelled by their position in the matrix. Thus the $(1,3)$ element is the covariance of the returns on the smallest quintile with those on the mid-quintile. The diagonal elements show that daily volatility increases as the sizes of the firms increase. This result runs counter to volatility estimates that are constructed using lower frequency returns and is driven by larger firms generally trading more frequently and consequently experiencing more intra-day price movements than their small stock counterparts. All of the variance and covariance measures are skewed to the right as the means are above the medians. Volatility is quite volatile: the standard deviation of the realized variances and covariances are much larger than their mean values.

The data are persistent as the auto-regressive coefficients are above 0.3 although not as high as one might expect. As noted by Barndorff-Nielsen and Shephard (2004), failing to remove the jump component of quadratic volatility is likely to diminish the serial correlation coefficient of volatility.[27] Indeed, the average serial correlation of the bi-power covariation measures (detailed below), essentially quadratic volatility measures purged of return shocks, have much higher serial correlation coefficients with the average across all variances and covariances increasing to 0.55.

The Geweke and Porter-Hudak (1983, hereafter GPH) statistics estimate the degree of fractional integration. The statistics are clustered around 0.4 indicating that the series are stationary but have 'long memory': shocks to volatility persist for a long time. Ding and Granger (1996) show that standard GARCH type volatility models cannot capture the long memory that is usually present in financial time series. They propose a multi-factor volatility model to capture this effect.[28] Baillie, Bollerslev and Mikkelsen (1996) and

[27] See also Forsberg and Ghysels (2004).
[28] See also Ding, Granger, and (1993) and Granger, and Ding (1996).

Bollerslev and Mikkelsen (1996) propose using fractionally integrated GARCH models. Engle and Lee (1999) suggest that a two factor model of volatility — with one persistent and one short-term factor — would be able to match the volatility dynamics.

The volatility series are not near normally distributed. The final three columns of Table 1a show the skewness and kurtosis statistics as well as the asymptotic marginal significance levels (P-values) of the Jarque-Bera tests for normality. The data are quite skewed and the volatility of volatility causes a great deal of kurtosis. Normality is rejected for all elements of the realized covariance matrix.[29]

The summary statistics of the log-space volatility matrix A_t are shown in Table 1b. Taking the matrix logarithm of the data changes its properties along several dimensions. First, while the mean and median values of the series change, the skewness of the volatility series is decreased. Indeed, many of the skewness coefficients are now close to 0. While the series are still volatile, the kurtosis statistics are close to 3.00 indicating that there excess kurtosis levels relative to the normal distribution are not obvious. Although, all of the Jarque-Bera statistics reject the null of normally distributed data, the test statistic values (not reported) have decreased quite dramatically. Thus, taking the matrix logarithm of multivariate realized volatility results in series that are much closer to being normally distributed. This parallels the univariate finding of Andersen et al. (2001).

The log-space volatilities are still persistent. However, the GPH statistics show a great deal of heterogeneity when compared to their counterparts for the realized volatility matrix. Some of the elements are above 0.5 indicating a non-stationary series. As the factor model combines all of the elements, we can observe if the model can capture the long-memory property of the actual data.

3.2 Bi-power covariation

While realized volatility provides an estimate of the quadratic variation of the covariance matrix, it is constructed assuming continuous sample paths for the stock prices. However, researchers have found that equity markets are characterized by the presence of jumps in volatility, which results in discontinuous sample paths that can reduce the predictability of quadratic variation estimates. Barndorff-Nielsen and Shephard (2004b and 2006) develop the theory of bi-power variation, a measure that, in essence, is the quadratic variation with the jump components removed. Ghysels, Santa-Clara and Valkanov (2004b) find that bi-power variation is a good predictor of aggregate market volatil-

[29] Thomakos and Wang (2003) discuss the influence of long-memory on tests of Normality. They find that the standard Jarque-Bera test is oversized while the QQ and Kolmogorov-Smirnov tests are approximately correctly sized. These tests also reject normality.

ity.[30] Barndorff-Nielsen and Shephard (2005) extend univariate bi-power variation to the multivariate case, aptly named bi-power covariation.

Given that our exercise is multivariate, we construct bi-power covariation measures for our portfolios using Definition 3 of Barndorff-Nielsen and Shephard (2005). Let (BP_t, q) represent the matrix of bi-power variation for portfolios on day t and lag q:

$$\{\mathbf{BP}_t, q\} = \left\{ \begin{array}{cccc} \{BP_t^1, BP_t^1; q\} & \{BP_t^1, BP_t^2; q\} & ... & \{BP_t^1, BP_t^P; q\} \\ \{BP_t^2, BP_t^1; q\} & \{BP_t^2, BP_t^2; q\} & ... & \{BP_t^2, BP_t^P; q\} \\ \vdots & \vdots & \ddots & \vdots \\ \{BP_t^P, BP_t^1; q\} & \{BP_t^P, BP_t^2; q\} & ... & \{BP_t^P, BP_t^P; q\} \end{array} \right\}, \tag{19}$$

where

$$\begin{aligned} \{BP_t^k, BP_t^l; q\} &= \frac{\delta_{q\triangle}}{4} \sum_{j=q+1}^{m} \left[\left| r_{t,j\cdot\triangle}^k + r_{t,j\cdot\triangle}^l \right| \left| r_{t,j\cdot\triangle-q}^k + r_{t,j\cdot\triangle-q}^l \right| \right. \\ &\quad \left. - \left| r_{t,j\cdot\triangle}^k - r_{t,j\cdot\triangle}^l \right| \left| r_{t,j\cdot\triangle-q}^k - r_{t,j\cdot\triangle-q}^l \right| \right] \end{aligned} \tag{20}$$

and $\delta_{q\triangle} = \frac{m}{m-q}$. The intuition for equation (20) follows from the result that $cov(x, y) = \frac{1}{4}[Var(x+y) - Var(x-y)]$. Equation (20) looks similar to a traditional realized volatility construction with the exception that contemporaneous returns are not squared but absolute returns in window j are multiplied by absolute returns in window $j - q$. This removes the jumps which are present in contemporaneous squared returns, but have little impact on the product of contemporaneous and lagged returns. For the case of $k = l$, equation (20) simplifies to the univariate bi-power variation measures of Barndorff-Nielsen and Shephard (2004b and 2006).

Motivated by our earlier discussion of MHAR models and the desire to reduce the parameter space, we construct log-space bi-power covariation matrixes, $A^{BP}(d)$, aggregated over the last $d = 1, 5$ and 20 days.[31] We then estimate the principal components (9) of the matrixes, $a^{BP}(d, i)$. Table 1c presents the first three principal components for each of the holding periods. As the holding period lengthens, the first principal component explains more of the variation of the matrix. The component places a large positive weight on the diagonal terms of the $A(d)^{BP}$ matrix while the weights on the off-diagonal elements are small.

[30] Anderson, Bollerslev, and Diebold (2003) construct volatility measures less influenced by jumps and also find these measures predict future volatility better than quadratic variation measures

[31] The bi-power covariation measures used in our study are constructed setting $q = 3$ (60 minutes) following Barndorff-Nielsen and Shephard (2005), although preliminary sensitivity analysis found little variation in the predictive ability of alternative bi-power covariation measures constructed using neighboring values of q.

3.3 Alternative forecasting variables

Our goal in this paper is to compare alternative models of the conditional covariance matrix. While all of the models use the latent factor form given in (12), they differ by the forecasting variables Z_{t-1} used. We construct four alternative models using forecasting variables that correspond to existing approaches in the literature. Table 1d gives summary statistics of all of the forecasting variables, while Table 1e shows the correlations among them.

Our first model, labelled "MHAR-RV-BP," is a multivariate HAR model of daily realized volatility using the principal components of bi-power covariation as predictors. In our estimations, we find multicollinearity among all of the principal components. This is not surprising as the bi-power covariations are measuring the volatility of the equity market over different holding periods. Because of the multicollinearity, we use only the first principal component of the 5-day series $(a^{BP}(5,1)_t)$ and the first three principal components of the 20-day series $(a^{BP}(20,1)_t, a^{BP}(20,2)_t, a^{BP}(20,3)_t)$ as regressors. Intuitively, by including the 5-day measure we hope to capture the higher-frequency volatility variation. We found the 5-day measures were more stable than the 1-day measures, likely due to reduced estimation error. Inclusion of the 20-day measures follows from HAR models where lagged longer horizon volatility measures are included as predictors.

Thus, the set of forecasting variables for this model is

$$Z_t^{MHAR-RV-BP} = (a^{BP}(5,1)_t, a^{BP}(20,1)_t, a^{BP}(20,2)_t, a^{BP}(20,3)_t).$$

We show below that these four series do a good job in capturing the forecastability of the 5×5 matrix. Incorporating shorter-run (e.g., $d = 1$ day) volatility measures do not improve the estimates.

The second model, "MHAR-RV-BPA", adds the asymmetric response of volatility to past return shocks to the first model. A number of authors have shown that past negative returns causes higher future equity market volatility. We therefore include the returns on the small and large stock portfolios when they are negative ($R_{1,t-1} < 0$ and $R_{5,t-1} < 0$, respectively).

Our third model, "MHAR-RV-X," uses variables that have been shown to forecast stock returns. These include: a risk-free interest rate, the dividend yield, the credit spread and the slope of the term structure. We use the return on 30-day Treasury Bill as our measure of the risk-free rate, TB. Our source is the Ibbotson SBBI database. We use the trailing 12 months' dividends divided by the price at the end of month t for our measure of the dividend yield, DY. The data are created from the CRSP value-weighted total returns and the monthly CRSP value-weighted capital gains returns. We use the difference between the end of month yield-to-maturity on BAA bonds and AAA bonds as provided by Moody's Corporate Bond Indices, for our measure of the credit spread, CS. Our source is the FRED® database. We use the difference between the yield-to-maturity on Ibbotson's Long-Term Government bond portfolio and the yield-to-maturity on 3-month Treasury Bills (from FRED®) as our measure of the term spread, TS.

The above variables are standard in models of rational asset pricing. We also include a variable that originates in the behavioral finance literature. McQueen and Vorkink (2004) show that the market's sensitivity to news can be captured by a measure called the scorecard, which is a function of prior market returns. They show that the scorecard adds incremental value to standard asymmetric GARCH models. We construct our measure of the scorecard, SC, following McQueen and Vorkink (2004). In particular, we iterate on equations (5) and (6) found in their paper using CRSP value-weighted daily market returns.[32]

The fourth model, "MHAR-RV-BPAX," uses all of the variables from the other models.

4 Results

4.1 Model fit

We estimated our four versions of (12) with $K = 1$ factors. However, the model was rejected for each of the variable sets used. We then estimated the models with 2 factors. The fits of the two factor models are much better and in the subsequent analysis we report results for $K = 2$ factors.

Table 2 presents summary statistics for the overall model fit. We present the statistics for the non-bias corrected results to show how the model does by itself. Panel A shows the J statistics that test the over-identifying restrictions (13) from the four models. All of the models are rejected indicating that the restrictions of a two-factor model are too onerous. However, we have been unable to estimate a three-factor version using any of the data sets. We interpret the rejection of the two-factor versions of the models as the result of Lindley's paradox (i.e., having a large number of data points implies that any precise null can be rejected). As three-factor versions of the models are not estimable, we use the two factor versions in the subsequent analysis.

[32] Initial evidence of the predictive power of dividend yields on returns is found in Shiller (1981). Initial evidence for the credit and term spreads on returns come from Fama and French (1989). Initial evidence for the relative short rate on returns is from Campbell (1991). Initial evidence on the scorecard's ability to predict market volatility comes from McQueen and Vorkink (2004).

The next panel of Table 2 present statistics that evaluate the overall model fit. The first three statistics are from Moskowitz (2003). The eigenvalue statistic is

$$eig_t = \frac{\sqrt{trace(\widehat{V}_t'\widehat{V}_t)}}{\sqrt{trace(V_t'V_t)}}.$$

The trace of the matrix equals the sum of its eigenvalues. As noted by Moskowitz (2003), the sum of the eigenvalues of the volatility matrix is a summary measure of time-varying volatility in the five stock returns. Thus, this ratio shows how the model is able to capture the time-varying volatility.

The magnitude statistic is

$$mag_t = \frac{i'\left(\left|V_t - \widehat{V}_t\right|\right)i}{i'\left(|V_t|\right)i},$$

where i is a conformable vector of ones. The numerator is the sum of the absolute values of the difference between the fitted and true values of the volatility matrix. The denominator is the sum of the absolute values of the true volatility matrix. The statistic measures the percent error of the fitted values of the model.

The direction statistic is the ratio of the sum of the signed fitted values of the model to the squared rank of the true volatility matrix,

$$dir_t = \frac{i'sign\left(V_t\widehat{V}_t\right)i}{[rank(V_t)]^2},$$

and is designed to capture how well the model predicts the correct direction of volatility.

The average values of the three statistics shown in Table 2 reveal that all four models do a reasonable job of capturing time-varying volatility.[33] Indeed, according to these metrics, there appears to be little difference between the models. In moving from the first to second model, the eigenvalue statistic decreases slightly as the return variables are added as regressors. The third model (MHAR-RV-X) that contains the return forecasting variables is closest to the theoretical value of 1.00. However, the statistics from all four models are close to this value, indicating that they are able to capture the time-varying volatility matrix.

The magnitude statistics show a slight improvement as the return variables are added to the MHAR-RV-BP model. The model using just the return forecasting variables (MHAR-RV-X) has the highest (i.e., poorest) magnitude statistic. The four values of the

[33] Following Moskowitz (2003), we calculate the standard errors of the statistics using a Newey-West (1987) estimator to account for autocorrelation and heteroskedasticity of unknown form.

direction statistic are the same and close to the optimal value of 1.000. The last part of the table shows the multivariate versions of the root-mean squared error and mean absolute deviation statistics. Once again we observe little variation across the models, although the return forecasting variable model has the largest errors.

The inability of our tests to distinguish among the models is similar to the results of Green and Hollified (1992) and Chan et al. (1999). These papers show that there is a dominant factor in multivariate models of latent volatility. Model comparison tests will have difficulty differentiating among models that contain this dominant factor. Chan et al. (1999) suggest using a tracking portfolio metric to better distinguish among the models. We adopt this approach in section 5.

One question of interest is the reduction in the explanatory power of a variable set resulting from the imposition of the two-factor structure. While a parsimonious structure is likely to be preferred in the out-of-sample tests, the in-sample fit may be reduced by the over-identfying restrictions (13). To test this, we calculate variance ratio statistics following Campbell (1987). In Panel A of Table 3, we measure the restrictions of the factor model on the log-space volatilities, A_t. In these ratios, the numerator is the variance of the fitted values $\widehat{A}_t$ of the log-space volatilities from the factor model (12). The denominator is the variance of the fitted values from an ordinary least squares regression of the log-space volatilities on the variables of the model as in (10). The ratio thus shows how much imposing the factor structure in (13) reduces the in-sample predictive power.

In Panel B of Table 3, we measure the restrictions of the factor model on the realized volatility matrix, V_t. In addition, the calculations for these ratios vary by whether they are for diagonal or off-diagonal elements of V_t. For the diagonal elements, the numerator is the variance of the (bias-corrected) fitted values of the realized volatility from (14). The denominator is the variance of the fitted values from an ordinary least squares regression of the realized volatilities on the variables of the model. For the off-diagonal elements, we repeat the same exercise using the Fisher transforms of the estimated correlations in the numerator and the Fisher transforms of the realized correlations in the dominator. This analysis allows us to measure the effects of the factor model on both the volatilities and the correlations contained in V_t.

While some of the ratios in Panel A are occasionally quite small (e.g. 0.014 for the $(1, 4)$ element of the A_t matrix for the MHAR-RV-BPA model), most of the ratios are above 0.7. The factor structure does not appear to impose much restrictions on the dynamics of the log-space volatilities. The results are just as strong for the realized volatilities in Panel B. Here the average ratios are above 1.00 for the volatilities in all four models suggesting that the factor model combined with the non-linear transformation in the matrix exponential function allows our model to capture more of the variation than a simple regression can. The results in panel B do show that the models with more variables have greater difficulty in capturing the dynamics of the time-varying correlations as the average ratio for the fourth model is 0.528. Overall, the results in Table 3 suggest that the two factor representation is not too stringent and that much of the variation in expected variances and correlations is captured by our models.

4.2 Estimated coefficients

Table 4 presents the estimates of the β coefficients from the model that uses all of the forecasting variables (MHAR-RV-BPAX). The coefficients are arranged according to the elements of the A_t matrix that they correspond to. The elements that are normalized for identification are in the upper left corners of both factors. Each cell presents the estimated coefficient from the GMM estimation procedure, the Newey-West (1987) standard error, and the corresponding t-statistic.[34]

In the first factor, the loadings for the first four diagonal elements are significant and positive while those for the off-diagonal elements are much smaller. In the second factor, the estimated coefficients are large and significant for the diagonal elements, while the off-diagonal elements are smaller, but mostly significant. With the exception of the $A(4,1)$ element, all of the log-space volatilities have a significant loading on at least one of the two factors. The overall significance of the coefficients show that the linear combinations of the forecasting variables (θZ_{t-1}) are able to forecast elements of the A_t matrix.

The estimated θ coefficients for the two factors in all of the models are shown in Table 5. Panel A presents the coefficients for the basic multivariate HAR model (MHAR-RV-BP). The coefficients on the first principal component of the lagged 5-day log-space bi-power covariation are significant in both factors. The coefficient on the first principal component of the lagged 20-day bi-power covariation is both positive and significant in the first factor. The coefficient on the second principal component is positive in the first factor, with a larger magnitude than the coefficient on the first principal component, and negative in the second factor, suggesting that the factors are picking up different elements of long-run volatility. The coefficients on the third principal component are both positive and significant, but smaller in size.

It is difficult to say much more about the influence of the variables on realized volatility due to the highly non-linear nature of the model: the θ coefficients in the two factors interact with each other as well as the β coefficients in (12). For example, it is not necessarily true that the second principal component of the 20-day log-space bi-power

[34] Only the results for the MHAR-RV-BPAX model are shown. The coefficients for the other models are similar and are available by request.

covariation (with a coefficient of 0.499 in the first factor) has a larger degree of explanatory power for realized volatilities than does the first principal component of the 5-day volatility (with a coefficient of 0.153 on the first factor). The elasticities presented below show the ultimate impact of any particular forecasting variable on the volatilities.

Panel B presents the coefficients for the model that includes the (asymmetric) effects of past stock returns (MHAR-RV-BPA). The coefficients on the lagged principal components of bi-power covariation are similar in size and significance to their values in the first model. In addition, an interesting pattern in the asymmetric response of volatility to past negative returns emerges. The coefficient on lagged negative returns on small stocks is negative in the first factor and not significant in the second factor. The coefficient on negative returns on large stocks has an opposite sign in the two factors. The elasticities presented below will show how the effects net out.

The coefficients on the variables usually used to forecast stock returns, model 'MHAR-RV-X', are presented in Panel C. The coefficients on the lagged short-term interest rate and the lagged credit spread are large and significant in both factors. The coefficients on the lagged term spread are insignificant. The coefficients on the scorecard and dividend yield variables are negative and significant.

Panel D presents the coefficients for the model that includes both the lagged volatility variables as well as the standard return forecasting variables (MHAR-RV-BPAX). The result for this model illustrates an important point in determining the effects of various economic factors on volatility: including lagged volatility as a variable to determine the true influence of all variables on the volatility proves to be important. This can be seen by examining the coefficient on the credit spread which shrinks in size and becomes mostly insignificant in the first factor. In contrast, the term spread appears as significant in the first factor while the scorecard becomes insignificant in the second factor. The autoregressive nature of the realized volatility implies that care must be taken when regressing volatility on lagged predictive variables which are themselves persistent.

In sum, a number of the coefficients on the variables that forecast stock returns remain significant in the multivariate model of stock market volatility, even when accounting for the effect of autoregressive volatility. The ability of these variables to forecast both means and volatilities suggests that better conditional portfolio allocation polices could be obtained than those based on return predictability alone. We return to this issue below.

4.3 Estimated elasticities

Table 6 presents the elasticities $\epsilon(i,j,z,t)$ calculated using (16). We calculate the elasticities for each day t in our sample and for each forecasting variable z. While we calculate the elasticities for each element (i,j) of the estimated volatility matrix $\widehat{V}_t$, we present only the results for the small stock variance $(\widehat{V}_{(1,1)})$, the covariance between small and large stocks $(\widehat{V}_{(1,5)})$, and the variance of the large stock portfolio $(\widehat{V}_{(5,5)})$ for brevity. The results for the other elements of $\widehat{V}$ are similar to those shown and are available on request. Table 6 reports the average value of the elasticities as well as their Newey-West (1987) standard errors.

Panel A of Table 6 gives the elasticities for the MHAR-RV-BP model. The first principal components of the lagged weekly and monthly volatilities have a positive effect on all of the realized volatilities. These components appear to capture the overall level of volatility in the market. In contrast, the second principal factor has an asymmetric effects on volatility. The elasticities change sign depending on the particular element of $\widehat{V}$. For example, changing the value of the second principal component by a one standard deviation shock would cause tomorrow's large stock volatility to decrease by 36.4 percent while small stock volatility would increase by 8.4 percent. We find monotonic decreases in the elasticities associated with this variable as we move from the small stock volatility towards the large stock portfolio, including those portfolio results not reported. The third principal component has a small influence on realized volatilities.

Panel B presents the results for the second model which includes the principal components and the lagged stock returns (MHAR-RV-BPA). The elasticities on the principal components do not change much from their values for the first model. The elasticities on the two negative return variables ($R_{1,t-1} < 0$ and $R_{5,t-1} < 0$) are quite small and negative showing that yesterday's return shocks do not explain a large part of today's volatility. This is in contrast to Kroner and Ng (1998), who find that different models give very different news impact surfaces showing the effects of past return shocks on current volatility. Thus, the asymmetric response of volatility to past returns is quite small in our data. We speculate that this may result from two reasons.

The first is that the response may only show up in longer (e.g., monthly) data so that we would not capture it in our model of daily volatility. The second is that our Newey-West type estimator of the realized volatility matrix V_t accounts for the interactions of the volatility between large and small stock portfolios. If new information is incorporated into the prices of large stocks first, then the returns and variances of large stocks would lead the returns and variances of small stocks. Indeed, Lo and MacKinlay (1990) show that the returns of large stocks lead those of small stocks, while Conrad, Gultekin and

Kaul (1991) show that volatilities of large stocks lead those of small stocks. However, this latter effect would be captured by our estimator (18) of the realized covariance matrix as lags of large stock volatility are used in constructing current small stock volatility. This implies that there would be little return asymmetries to be captured by any right-hand-side variable.

The elasticities on the variables usually used to model expected stock returns are shown in Panel C. An increase in the short-term interest rate or the credit spread causes the elements of the conditional covariance matrix to increase. An increase in the scorecard or dividend yield causes the elements to decrease. The slope of the term structure has a small effect.

Panel D presents the elasticities for the model that includes all of the variables. The one large change from previous smaller models is that the elasticities on the variables used in the MHAR-RV-X model greatly decrease in magnitude. For example, the impact of a change in the short-term interest rate almost completely disappears. This, once again, shows the importance of including lagged volatility in the models.

While the averages of the estimated elasticities reveal a large influence for certain variables on volatility, we also find it instructive to examine the time series properties of the elasticities. Figure 1a presents the estimated $\epsilon(i, j, z, t)$ for the first six forecasting variables in the MHAR-RV-BPAX model. Only the elasticities for $(i, j) = (1, 1)$, the volatility of the small stock portfolio, are shown. The graphs reveal considerable variation over time in the elasticities. For example, the elasticity of the first factor on the lagged 5-day bi-power covariation rises from 0.25 near the start of the sample to 0.26 near the end. The elasticity on the first factor of the 20-day bi-power covariation is around 0.40 for most of the sample but rises towards the end. The elasticity on the negative return on large stocks ($R_{5,t-1} < 0$) is hovers just below 0 at the start of the sample before turning more negative. Figure 1b shows the plots for the last five variables of the model. The graphs show large swings in the elasticities of all of the variables towards the end of the sample.

These plots show that our approach can help explain the volatility of volatility in the size-sorted portfolios.[35] Chernov et al. (2001) examine a number of volatility models. They note that the main difficulty is in capturing the persistence of volatility and the high degree of kurtosis found in daily stock index data. They find that a 'feedback' effect, where higher volatility causes the volatility of volatility to be higher, is important. Consequently, modeling the state dependence of volatility as our approach does is crucial to capturing stylized features of conditional variances.

[35] The plots in Figure 1b show time variation in *expected* volatility. In future versions of this paper, we hope to show the influence of shocks to the forecasting variables on *unexpected* volatility.

5 Global Minimum Variance and Tracking Portfolios

5.1 Motivation

The above analysis showed that it was difficult to differentiate among the models on the basis of unconditional statistics. In this section, we calculate global minimum variance portfolios and minimum tracking error portfolios to evaluate the conditional performance of the various models. These tests will provide an economic metric for the alternative information sets. They also provide a way to evaluate the importance of the factor structure of our model.

The problem of distinguishing among alternative models of the covariance matrix has been noted before. Green and Hollified (1992) show that portfolios can be poorly diversified — in the sense of having large negative holdings to offset large positive holdings — when the covariance matrix of stock returns is dominated by a single factor. They note that in real data determining whether the extreme positions are due to the poor sampling properties of the data or to the existence of a single factor in the population covariance matrix is difficult.

Chan et. al (1999) introduce the idea of using tracking error portfolios to evaluate alternative models. In their approach, the portfolio manager combines a number of tracking portfolios or assets to closely follow the return on a benchmark asset. If the volatilities of both the benchmark and tracking portfolios are exposed to the dominant factor, then minimizing the variance of the difference in returns between the two portfolios will allow most of the volatility driven by the dominant factor to be removed. Consequently, incremental differences between the models will be magnified.

Jagannathan and Ma (2003) raise another issue related to evaluating models of the covariance matrix. They note that practitioners often impose non-negativity constraints. This could help subsequent portfolio performance if the constraints had the effect of offsetting the sampling error in the estimates. Conversely, the constraints would diminish subsequent performance if the sampling error was small and the true covariance matrix was dominated by the single factor. Thus, whether imposing (possibly incorrect) constraints is better than unconstrained optimization depends on the degree to which the principal factor dominates the estimated covariance matrix.

The matrix logarithmic factor model presented above can shed some light on this trade-off. By using realized covariance matrixes, we can reduce the sampling errors that are present when latent covariance matrix estimates are derived from monthly return data. In addition, we can evaluate the importance of the factors along with the specific variables that associate strongly with the factor. We can also evaluate the constraints to see if they augment or diminish subsequent portfolio performance.

5.2 Global minimum variance and tracking error portfolios

We construct two sets of global minimum variance portfolios, one set with unconstrained portfolio weights and the other set with margin and shorting limits that are constrained to be lower than 30 percent. Each week, we solve the following minimization problem,

$$\begin{aligned} &\min_{w}\left(w_t'\hat{V}_t w_t\right) \\ &s.t. : -0.30 \le w_i \le 1.30, \end{aligned} \tag{21}$$

where the weight restrictions are ignored for the unconstrained optimization. We then use these weights to form portfolios and evaluate the time-series properties of the minimum variance portfolios.

For our tracking error exercise, we use the size-sorted portfolios to track the return on a value portfolio, constructed following Fama and French (1992 and 1993). Specifically, we construct book-to-market measures for each firm with Compustat and CRSP data and divide stocks into ten deciles monthly based on their relative book-to-market measures. Our tracking portfolio, high minus low (HML) is constructed by taking a long position in the 20% of stocks with high book-to-market (value stocks) and taking a short position in the 20% of stocks with the lowest book-to-market (growth stocks). We construct a value-weighted portfolio and rebalance the positions monthly.[36]

Let Υ_t denote the realized covariance matrix at time t that includes the volatility on the benchmark HML return and the five size-sorted portfolios:

$$\Upsilon_t = \begin{bmatrix} & \sigma_{1,HML,t} \\ V_t & \cdots \\ & \sigma_{5,HML,t} \\ \sigma_{1,HML,t} \ldots \sigma_{5,HML,t} & \sigma_{HML,HML,t} \end{bmatrix}.$$

Here, the upper-left quadrant is the realized volatility matrix of the five size sorted portfolios, V_t, as used above. The lower-right element is the realized volatility of the HML portfolio, $\sigma_{HML,HML,t}$.The lower-left and upper-right quadrants are the covariances of the size sorted portfolios and HML ($\sigma_{1,HML,t} \ldots \sigma_{5,HML,t}$).

We estimate the forecasted value of the augmented volatility matrix, $\widehat{\Upsilon}_t$, using the same factor models as above. We can then use the elements of this matrix to solve for

[36] Earlier versions of the paper used the S&P 500 SPDR as the tracking portfolio. However, we found that given our portfolios were sorted on size and the fact that the S&P SPDR is essentially a portfolio of large-cap stocks, our models performed very similar in tracking exercises.

the following tracking error portfolio. We wish to minimize the variance of the difference between the return on the HML portfolio, r_t^{HML}, and the return on the portfolio of size-sorted stock returns, $w_t' r_t$:

$$\min_{w} var(r_t^{HML} - w_t' r_t),$$

where w_t is a 5×1 vector of portfolio weights. Using the variance decomposition, the problem can be written as

$$\min_{w_t}\{var(r_t^{HML}) + var(w_t' r_t) - 2cov(r_t^{HML}, w_t' r_t)\}.$$

Using matrix notation, this minimization problem can be rewritten as:

$$\min_{w_t}\{e_6' \widehat{\Upsilon}_t e_6 + w_t' \widehat{V}_t w_t - 2\sum_{i=1}^{5} w_t^i \cdot e_6' \widehat{\Upsilon}_t e_i\}, \tag{22}$$

where e_i is a 6×1 vector of zeros with a 1 in the i^{th} position. We construct tracking error portfolios following (22) along with imposing the constraints that margin limits are 30 percent of the portfolio and that portfolios can be shorted up to 30 percent as well. To test overfitting of the data, we reestimate the models using only the data from 1988 - 2000 that allows us to construct out-of-sample statistics over the 2001-2002 period.

5.3 Results

Table 7 presents summary statistics of the global minimum variance portfolio and the tracking error portfolios. We report the return averages and standard deviations for the cases where the weights are constrained between -0.30 and 1.30 and unconstrained. We also report portfolio return averages and standard deviations for an in-sample (1998 - 2000) and an out-of-sample (2001-2002) period.

A number of patterns emerge from Table 7. First, we find that placing constrains on the portfolio weights in almost all cases reduces the standard deviations of the portfolios. This occurs for both in-sample and out-of-sample tests over both the minimum variance and tracking error portfolio tests. These results are consistent with Jagannathan and Ma's (2003) conjecture that constraints may help (or at least not hurt) a portfolio optimization exercise.

Second, we find, not unexpectedly, that portfolio standard deviations increase as we move from in-sample tests to out-of-sample tests. This again holds across both the minimum variance and tracking portfolio exercises.

Third, the average returns on both the minimum variance and tracking error portfolios are small and variable. Out-of-sample returns are, in general, smaller than in-sample returns. Imposing the constraints appears to have little effect. However, the returns are so small and we have not accounted for transactions costs in this exercise so it is hard to say anything definitive about them.

Fourth, performance improves when we use a model to calculate the minimum variance portfolio. The in and out-of-sample standard deviations on the equally-weighted portfolio are presented at the bottom of the table. We use this portfolio as it is the simplest passive portfolio that could be constructed. In sample, the minimum variance portfolios constructed using the two models containing the return forecasting variables (MHAR-RV-X and MHAR-RV-BPAX) and imposing the constraints on the portfolio weights have standard deviations that are lower than that on the equally weighted portfolio. Out of sample, only the portfolio constructed using the basic GARCH type model (MHAR-RV-BP) without any constraints yields a standard deviation that is greater than that of the equally-weighted portfolio. In other words, most of the models with or without constraints yield lower volatilities out of sample. Including the conditioning information improves the estimates of the covariance matrix out of sample.

Fifth, using the conditioning information in our two-factor model also leads to better performance using the tracking error portfolio metric. The in-sample standard deviations of the tracking error portfolios constructed using (22) are lower than when the equally-weighted portfolio is used to track the index. In the out-of-sample tests, only the standard deviation of the fourth model (MHAR-RV-BPAX) is higher than that on the equally weighted tracking portfolio. Once again, we conclude that including the conditioning information in our two factor model improves portfolio performance.

Finally, we find interesting variations across the models that are fairly consistent across the two exercises. Adding asymmetric return shocks to the base model, MHAR-RV-BP, decreases the portfolio standard deviations for the minimum variance portfolios and has a mixed impact on the tracking error portfolios. The third model, MHAR-RV-X, that includes only the economic variables, often has the lowest return standard deviation of the four models. The economic variables included in this model have long been used to predict returns. Here we show that they forecast portfolio standard deviations. Including both past volatilities and the economic variables in the fourth model (MHAR-RV-BPAX) rarely leads to the best result.

6 Conclusions

This paper has introduced a new model for the realized covariance matrix of returns. The model is parsimonious, guarantees a positive-definite covariance matrix, and does not require parameter constraints to be imposed. The model allows a number of variables to forecast the covariance matrix, yet restricts the number of factors in the estimation process. In addition, time-varying elasticities can be calculated that show the extent to which a percent shock to the forecasting variable influences any particular element of the realized covariance matrix.

The model is applied to the covariance matrix of realized stock returns over the 1988 to 2002 period. Four alternative sets of forecasting variables are tested. The alternative sets of forecasting variables produce results that are roughly similar according to standard unconditional tests. However, there are differences between the fitted values from the alternative models. We evaluate these differences using minimum variance and minimum tracking error portfolios. These results highlight the importance of including return forecasting variables in models of conditional volatility.

We hope to extend this analysis in other papers. In particular, we would like to include variables related to the volume of trading and order flow in the size-sorted portfolios given the links that have been found in past work.

A Construction of the Data

Since Epps (1979), it has been realized that measuring correlations using high-frequency data can be problematic. In particular, estimated correlations can be substantially lower when using high-frequency data (intra-day TAQ) as opposed to low-frequency data (daily and monthly CRSP). We were particularly sensitive to this issue given our finding that small stock return standard deviations being smaller than large stock return standard deviations using the TAQ data. We conducted data checks to ensure our methodology on the TAQ database was not creating data that was materially different in standard deviations and correlations than one would obtain is lower frequency data sets were used.

Table A1 provides an analysis of the portfolio return standard deviations using a number of different approaches. The first column reports the average portfolio return standard deviations (annualized) using in our analysis and corresponds to numbers found in Table 1A. The second column reports monthly portfolio return standard deviations (annualized) using daily returns following a similar methodology as our 20-minute high frequency approach where we only construct one price per stock each day based on the final quotes. The final two columns report portfolio return standard deviations constructed using daily and monthly CRSP returns representing a low frequency approach. Table A1 provides two interesting results. First, we find that our high frequency approach leads to portfolio return standard deviations that are lower than all of the alternatives, including daily TAQ returns. While our high frequency approach leads to lower standard deviations, the differences are quite small and likely due to our methods inability to capture information from the close of markets on one day to the open of markets on the following day. This likely leads to the daily TAQ results matching up so closely with the daily CRSP results.

Second, we note that small stock return variances are smaller than large stock counterparts for all of the approaches other than monthly CRSP, suggesting that our finding of this result in the high-frequency TAQ returns is not driven by the particular methodology we employ. We interpret the results of Table A1 as positive, that our high-frequency approach captures most of the return volatility as captures in lower-frequency returns but should do so with more precision relying on the sampling results of quadratic variation. One could account for overnight return information by adopting methods such as Hansen and Lunde (2005) that would likely increase the first column of Table A1 towards the results in the other three columns.

We also investigated the realized correlations to assess the impact of the Epps (1979) effect in our data. Table A2 provides results of these investigations. In this table, we compare the correlations that results from our high-frequency approach to those constructed using lower-frequency daily CRSP returns. The first three columns of Table A2 summarizes correlation information taken from daily CRSP returns, where we estimated monthly correlations each month from January 1988 through December 2002. The ten unique portfolio correlations correspond to the rows of the table and the time-series mean of the monthly estimated correlations is found in first column, the time series standard deviations in second, and the volatility ratio (standard deviation divided by mean) in the third column. Corresponding results for the high-frequency correlations are found in columns four through six. Comparing columns one and six we find that the low frequency approach leads to higher average correlations in all cases although the differences in correlations does not appear to be large. In fact, the final column reports a t-statistic on the differences of the two means where the correlations are transformed using the Fisher transformation so that the distribution of the resulting t-statistic is well-behaved. We see that at standard levels only two of the ten correlations are found to have statistically different means. Comparison of columns two and five allow one to see the time variation in correlations of the two approaches. We find that the low-frequency leads to greater time variation in monthly correlations for some pairs of portfolios but the high-frequency approach leads to greater time variation for others. In particular pairs of smaller stock portfolios have greater time-variation using the low-frequency approach while pairs of large stock portfolios have greater variation using the high-frequency approach. In any respect, our high-frequency approach does not seem to lose a substantial amount of time variation in correlations as compared to low-frequency approaches.

B Bias Correction

Our estimator $\widehat{V}_t$ will be biased as the estimation is done in the log-volatility space:

$$E\left(\widehat{V}_t\right) = \text{expm}\left(\widehat{A_t}\right) \neq E\left(V\right)$$

In this section we present two simple bias correction methods for the estimator, depending on the maintained assumption of the distribution and functional form of the Z_{t-1} and ε_t in (12).

B.1 Normally distributed variables

In the first correction, we assume that the $Z_{t-1}$and ε_t are jointly Normally distributed and that a standard linear regression model (10) has been estimated. We can thus also assume that the $P \times P$ matrix of fitted values $\widehat{A_t}$ and the $P \times P$ matrix of estimated residuals $\widehat{\varepsilon}_t$ are both distributed matrix Normal and are independent of each other (Gupta and Nagar (2000), p. 55). Also, by definition, $V = \text{expm}\left(\widehat{A} + \widehat{\varepsilon}\right)$.

Now consider using the definition of the matrix exponential function in (1) to expand the expected value of expm $\left(\widehat{A} + \widehat{\varepsilon}\right)$ up to the fourth power. Using Theorem 2.3.3 Gupta and Nagar (2000), a number of expectations of the products will vanish:

$$\begin{aligned} E(\widehat{A}\widehat{\varepsilon}) &= 0, \\ E(\widehat{\varepsilon}\widehat{\varepsilon}\widehat{\varepsilon}) &= 0, \\ E(\widehat{A}\widehat{A}\widehat{\varepsilon}) = E(\widehat{A}\widehat{\varepsilon}\widehat{A}) = E(\widehat{\varepsilon}\widehat{A}\widehat{A}) &= 0, \\ E(\widehat{A}\widehat{A}\widehat{A}\widehat{\varepsilon}) \text{ and permutations } &= 0, \\ E(\widehat{\varepsilon}\widehat{\varepsilon}\widehat{\varepsilon}\widehat{A}) \text{ and permutations } &= 0. \end{aligned}$$

Using these results and collecting like terms, the expansion becomes:

$$E(V) = E\left[\begin{array}{c} I + \left(\widehat{A} + \frac{\widehat{A}^2}{2!} + \frac{\widehat{A}^3}{3!} + \frac{\widehat{A}^4}{4!}\right) + \left(\frac{\widehat{\varepsilon}^2}{2!} + \frac{\widehat{\varepsilon}^4}{4!}\right) + \left(\frac{\widehat{\varepsilon}\widehat{\varepsilon}\widehat{A} + \widehat{\varepsilon}\widehat{A}\widehat{\varepsilon} + \widehat{A}\widehat{\varepsilon}\widehat{\varepsilon}}{3!}\right) \\ + \left(\frac{\widehat{A}\widehat{A}\widehat{\varepsilon}\widehat{\varepsilon} + \widehat{\varepsilon}\widehat{\varepsilon}\widehat{A}\widehat{A} + \widehat{A}\widehat{\varepsilon}\widehat{A}\widehat{\varepsilon} + \widehat{\varepsilon}\widehat{A}\widehat{\varepsilon}\widehat{A} + \widehat{A}\widehat{\varepsilon}\widehat{\varepsilon}\widehat{A} + \widehat{\varepsilon}\widehat{A}\widehat{A}\widehat{\varepsilon}}{4!}\right) + \ldots \end{array}\right] \tag{23}$$

Now we note from (1) that the first five terms on the right-hand-side of (23) are the leading terms of the matrix exponential of $\widehat{A}$. Thus, we have the approximation:

$$E(V) \approx E\left(\text{expm}\left(\widehat{A}\right)\right) + bc_1 \tag{24}$$

where the bias correction term is

$$bc_1 = E\left(\begin{array}{c} \left(\frac{\widehat{\varepsilon}^2}{2!} + \frac{\widehat{\varepsilon}^4}{4!}\right) + \left(\frac{\widehat{\varepsilon}\widehat{\varepsilon}\widehat{A} + \widehat{\varepsilon}\widehat{A}\widehat{\varepsilon} + \widehat{A}\widehat{\varepsilon}\widehat{\varepsilon}}{3!}\right) \\ + \left(\frac{\widehat{A}\widehat{A}\widehat{\varepsilon}\widehat{\varepsilon} + \widehat{\varepsilon}\widehat{\varepsilon}\widehat{A}\widehat{A} + \widehat{A}\widehat{\varepsilon}\widehat{A}\widehat{\varepsilon} + \widehat{\varepsilon}\widehat{A}\widehat{\varepsilon}\widehat{A} + \widehat{A}\widehat{\varepsilon}\widehat{\varepsilon}\widehat{A} + \widehat{\varepsilon}\widehat{A}\widehat{A}\widehat{\varepsilon}}{4!}\right) \end{array}\right). \tag{25}$$

In simulated data (not reported), this bias correction works well. We speculate that it would also work well in models of multivariate volatility such as ours if the original realized covariance matrix was estimated at a lower frequency (e.g. a monthly matrix could be constructed from daily data). Unfortunately, in our high-frequency application, the estimated $\widehat{A}$ matrix is too far from being Normally distributed with many outliers for this bias correction to work.

B.2 Non-Normally Distributed Data

With $\widehat{A}$ and $\widehat{\varepsilon}$ not being Normally distributed, there is no obvious closed-form bias correction. Thus, we do a simple numerical bias correction on the individual volatility series. The realized volatility matrix V_t can be decomposed into a matrix of standard deviations and correlations:

$$V_t = SD_t * C_t * SD_t',$$

where SD_t is a $P \times P$ diagonal matrix of the standard deviations and C_t is a $P \times P$ symmetric matrix of the correlations. A similar decomposition can be done for the fitted value $\widehat{V}_t$ to yield the $\widehat{SD}_t$ and $\widehat{C}_t$ matrixes. We then estimate a bias correction factor as the ratio of the median values of the two standard deviation series:

$$bc_2 = \frac{\text{med}\,(SD_t(i,i))}{\text{med}\left(\widehat{SD}_t(i,i)\right)}, \quad i = 1, ..., 5$$

We then bias correct the standard deviations while leaving the correlations intact. This simple method works well in that the fitted values are of the approximate magnitude of the actual realized volatility series and the statistical and economic tests presented in the paper support its use. We recognize that other, more sophisticated bias-correction methods could produce better results.

References

[1] Alizadeh, S., M. W. Brandt, and F. X. Diebold, 1999, Range-based estimation of stochastic volatility models. *Journal of Finance* 57, 1047-1091.

[2] Andersen, T. G. and T. Bollerslev, 1997, Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns. *Journal of Finance* 52, 975-1005.

[3] Andersen, T. G. and T. Bollerslev, 1998, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885-905.

[4] Andersen, T. G., T. Bollerslev, P. F. Christoffersen and F. X. Diebold, 2005a, Practical volatility and correlation modeling for financial market risk management. forthcoming in Mark Carey and Rene Stulz (eds.), *Risks of Financial Institutions*, University of Chicago Press.

[5] Andersen, T. G., T. Bollerslev, P. F. Christoffersen and F. X. Diebold, 2005b, Volatility forecasting. NBER working paper 11188.

[6] Andersen, T. G., T. Bollerslev, and F. X. Diebold, 2003, Some like it smooth and some like it rough: Untangling continuous and jump components in measuring, modeling, and forecasting asset return volatility. working paper, Northwestern University.

[7] Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens, 2001, The distribution of realized stock return volatility. *Journal of Financial Economics* 61, 43-76.

[8] Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys, 2001, The Distribution of Realized Exchange Rate Volatility. *Journal of the American Statistical Association* 96, 42-55.

[9] Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys, 2003, Modeling and Forecasting Realized Volatility. *Econometrica* 71, 579-625.

[10] Andersen, T. G., T. Bollerslev, F. X. Diebold, and J.Wu, 2004, Realized Beta: Persistence and Predictability. working paper, Northwestern University.

[11] Asai, M., M. McAleer, and J. Yu, 2005, Multivariate stochastic volatility. working paper, Tokyo Metropolitan University.

[12] Attanasio, O., 1991, Risk, time-varying second moments and market efficiency. *Review of Economic Studies* 58, 479-94.

[13] Baillie, R.T. and T. Bollerslev, 1994, Long memory in the forward premium. *Journal of International Money and Finance* 13, 565-71.

[14] Baillie, R.T., T. Bollerslev and H.O. Mikkelsen, 1996, Fractionally integrated generalized autoregressive heteroskedasticity. *Journal of Econometrics* 73, 3-33.

[15] Barndorff-Nielsen, O. E. and N. Shephard, 2004a, Econometric analysis of realised covariation: High frequency based covariance, regression, and correlation in financial economics. *Econometrica* 72, 885-925.

[16] Barndorff-Nielsen, O. E. and N. Shephard, 2004b, Power and bipower variation with stochastic volatility and jumps, with discussion. *Journal of Financial Econometrics* 2, 1-48.

[17] Barndorff-Nielsen, O. E. and N. Shephard, 2005, Measuring the impact of jumps in multivariate price processes using bipower variation. mimeo Oxford University.

[18] Barndorff-Nielsen, O. E. and N. Shephard, 2006, Econometrics of testing for jumps in financial economics using bipower variation. J*ournal of Financial Econometrics* 4, 1-30.

[19] Bauwens, L., S. Laurent and J.V.K. Rombouts, 2004, Multivariate GARCH models: A survey. CORE Discussion Paper 2003/31.

[20] Black, F., 1976, Studies of stock market volatility changes. *Proceedings of the American Statisitical Association, Business and Economics Statistics Section*, 177-181.

[21] Bollerslev, T. and H.O. Mikkelsen, 1996, Modeling and pricing long memory in stock market volatility. *Journal of Econometrics* 73, 151-184.

[22] Campbell, J.Y., 1987, Stock returns and the term structure. *Journal of Financial Economics* 18, 373-399.

[23] Chan, L. K.C., J. Karceski, and J. Lakonishok, 1999, On portfolio optimization: Forecasting covariances and choosing the risk model. *Review of Financial Studies* 12, 937-973.

[24] Chernov, M., A.R. Gallant, E. Ghysels and G. Tauchen, 2003, Alternative models for stock price dynamics. *Journal of Econometrics* 116, 329-64.

[25] Chib, S., F. Nardari and N. Shephard, 2004, Analysis of high dimensional multivariate stochastic volatility models. *Journal of Econometrics*, forthcoming.

[26] Chiu, T.Y.M., T. Leonrad and K.W. Tsui, 1996, The matrix-logarithmic covariance model. *Journal of the American Statistical Association* 91, 198-210.

[27] Conrad, J., M.N. Giltekin and G. Kaul, 1991, Asymmetric predictability of conditional variances. *Review of Financial Studies* 4, 597-622.

[28] Corsi, F., 2004, A simple long memory model of realized volatility. working paper, University of Lugano.

[29] Diebold, F. X. and M. Nerlove, 1989, The dynamics of exchange rate volatility: A multivariate latent factor ARCH model. *Journal of Applied Econometrics* 4, 1-21.

[30] Ding, Z. and C.W.J. Granger, 1996, Modeling volatiliy persistence of speculative returns: A new approach. *Journal of Econometrics* 73, 185-215.

[31] Ding, Z., C.W.J. Granger, and R.F. Engle, 1993, A long memory property of stock markets and a new model. *Journal of Empirical Finance* 1, 83-106.

[32] Doornik, J.A. and H. Hansen, 1994, An omnibus test for univariate and multivariate normality. working paper, Nuffield College, Oxford.

[33] Elton, E. J., J. M. Gruber, G. Comer and K. Li, 2002, Spiders: Where are the bugs?, *Journal of Business* 75, ppp-ppp.

[34] Engle, R.F. and G.G.J. Lee, 1999, A long-run and short-run component model of stock return volatility. in R.F. Engle and H. White, eds. *Cointegration, Causality and Forecasting – A Festchrift in Honour of Clive W.J. Grange*r, Oxford University Press.

[35] Engle, R.F. and V.K. Ng, 1993, Measuring and testing the impact of news on volatility. *Journal of Finance* 48, 1749-1778.

[36] Epps, T.W., 1979, Comovements in stock prices in the very short run. *Journal of the American Statistical Association* 74, 291-98.

[37] Fama, E., and K. French, 1992, The cross-section of expected returns. *Journal of Finance* 47, 427-465.

[38] Fama, E., and K. French, 1993, Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3-56.

[39] French, K., G.W. Schwert and R. F. Stambaugh, 1987, Expected stock returns and volatility. *Journal of Financial Economics* 19, 3-29.

[40] Gallant, A.R., C.T. Hsu, and G. Tauchen, 1999, Using daily range data to calibrate volatility diffusions and extract the forward integrated variance. *The Review of Economics and Statistics* 81, 617-631.

[41] Gallant, A. R., P. E. Rossi and G. Tauchen, 1993, Nonlinear dynamic structures. *Econometrica* 61, 871-907.

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

[43] Ghysels, E. and E. Jacquier, 2005, Market beta dynamics and portfolio efficiency. working paper, University of North Carolina.

[44] Ghysels, E., P. Santa-Clara and R. Valkanov, 2004a, The MIDAS touch: Mixed data sampling regression models. University of North Carolina working paper.

[45] Ghysels, E., P. Santa-Clara and R. Valkanov, 2004b, Predicting volatility: Getting the most out of return data sampled at different frequencies. University of North Carolina working paper.

[46] Glosten, L.R., R. Jagannathan and D.E. Runkle, 1993, On the relationship between the expected value and the variance of the nominal excess return on stocks. *Journal of Finance* 48, 1779-1801.

[47] Granger, C.W.J. and Z. Ding, 1996, Varieties of long memory models. *Journal of Econometrics* 73, 61-77.

[48] Green, R.C. and B. Hollifield, 1992, When will mean-variance efficient portfolios be well diversified?. *Journal of Finance* 47, 1785-1809.

[49] Gupta, A.K. and D.K. Nagar, 2000, *Matrix Variate Distributions*. New York, Chapman and Hall.

[50] Hansen, P., and A. Lunde, 2004, An unbiased measure of realized variance, working paper Stanford University.

[51] Hansen, P., and A. Lunde, 2005, A realized variance for the whole day based on intermittent high-frequency data, *Journal of Financial Econometrics,* 3, 525-554.

[52] Harvey, A.C., E. Ruiz and N. Shephard, 1994, Multivariate stochastic variance models. *Review of Economic Studies* 61, 247-264.

[53] Harvey, C.R., 1989, Time-varying conditional covariances in tests of asset pricing models. *Journal of Financial Economics* 24, 289-317.

[54] Jagannathan, R. and T. Ma, 2003, Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance* 58, 1651-1683.

[55] Jones, C.S., 2003, The dynamics of stochastic volatility: Evidence from underlying and options markets. *Journal of Econometrics* 116, 181-224.

[56] Kawakatsu, H., 2003, Matrix exponential GARCH. working paper, Queen's University, Belfast.

[57] King, M., E. Sentana, and S. Wadhwani, 1994, Volatility and links between national stock markets. *Econometrica* 62, 901-933.

[58] Kroner, K. and V.K. Ng, 1998, Modeling asymmetric comovements of asset returns. *Review of Financial Studies* 11, 817-844.

[59] Ledoit, O., P. Santa-Clara and M. Wolf, 2003, Flexible multivariate GARCH modeling with an application to international stock markets. *The Review of Economics and Statistics* 85, 735-747.

[60] LeSage, J.. and R. K. Pace, 2003, A matrix exponential spatial specification. working paper University of Toledo.

[61] Lo, A.W. and A.C. MacKinlay, 1990, When are contrarian profits due to stock market overreaction?. *Review of Financial Studies* 3, 175-205.

[62] Martens, M., 2003, Estimating unbiased and precise realized covariances. working paper, Erasmus University.

[63] Mathias, R.,, 1996, A chain rule for matrix functions and applications. *SIAM Journal on Matrix Analysis and Applications* 17, 610-620.

[64] McQueen, G.R. and K. Vorkink, 2004, Whence GARCH? A preference-based explanation for conditional volatility. *Review of Financial Studies* 17, 915-949.

[65] Moler, C. and C. van Loan, 2003, Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review* 45, 3-49.

[66] Moskowitz, T.J., 2003, An analysis of covariance risk and pricing anomalies. *The Review of Financial Studies* 16, 301-333.

[67] Najfeld, I. and T.F. Havel, 1995, Derivatives of the matrix exponential and their computation. *Advances in Applied Mathematics* 16, 321-375.

[68] Pagan, A.and G.W. Schwert, 1990, Alternative models for conditional stock volatility. *Journal of Econometrics* 45, 267-290.

[69] Poon, S. and C. Granger, 2003, Forecasting volatility in financial markets: A Review. *Journal of Economic Literature* 41, 478-539.

[70] Schwert, G.W., 1989a, Business cycles, financial crises and stock volatility. *Carnegie-Rochester Conference Series in Public Policy* 31, 83-126.

[71] Schwert, G.W., 1989b, Why does stock market volatility change over time?. *Journal of Finance* 64, 1115-1153.

[72] Schwert, G.W., 1990, Stock volatility and the crash of '87. *Review of Financial Studies* 3, 77-102.

[73] Schwert, G.W. and Paul J. Seguin, 1990, Heteroskedasticity in stock returns. *Journal of Finance* 45, 1129-1155.

[74] Scruggs, J. T., 1998, Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two factor approach. *Journal of Finance* 53, 575-603.

[75] Shanken, J. and A. Tamayo, 2004, Dividend yield, risk and mispricing: A bayesian analysis. working paper, Emory University.

[76] Thompson, C. J., 1965, Inequality with applications in statistical mechanics. *Journal of Mathematical Physics* 6, 1812-1813.

[77] Whitelaw, R.F., 1994, Time variations and covariations in the expectations and volatility of stock market returns. *The Journal of Finance* 49, 515-541.

[78] Williams, P.M., 1999, Matrix logarithm parameterizations for regularized neural network covariance models. *Neural Network*s 12, 299-308.

Table 1a
Summary Statistics of the Realized Covariance Matrix of Size-Sorted Stock Returns

The table shows summary statistics of the realized one-day covariance matrix of stock returns on the five size-sorted NYSE portfolios. All stocks on the exchange are sorted by market value at the beginning of each calendar year. The data are spaced into 30 minute observations and the price at the end of the observation is recorded. The log change in the prices are squared and summed over each day resulting in a total of 3,778 daily observations. The covariances are labeled $V(1)_{i,j}$, where i=1 contains the smallest and i=5 contains the largest stocks. The summary statistics of the upper triangular elements of the resulting matrix are shown here. The table shows the: mean; median; standard deviation; the first-order autoregressive coefficient; the Geweke-Porter-Hudack (GPH) statistic of long memory; and the skewness and kurtosis statistics. Also shown is the asymptotic marginal significance level (P-value) for the Jarque-Bera test of Normality. The bottom of the table presents the QQ test statistic of multivariate normality, multivariate skewness and multivariate kurtosis test statistics as well as their marginal significance levels. The sources of the data are the Institute for the Study of Securities Markets' (ISSM) database (January, 1988 to December, 1992) and the Trades and Quotes database (January, 1993 to December, 2002).

						Normality Test		
$V(1)_{i,j}$	Mean (%)	Median (%)	Std. Dev. (%)	AR(1) statistic	GPH statistic	Skewness statistic	Kurtosis statistic	Jarque-Bera P-value
$V(1)_{1,1}$	0.0061	0.0245	0.0129	0.390	0.464	8.779	132.325	<0.001
$V(1)_{1,2}$	0.0048	0.0143	0.0127	0.406	0.395	9.697	143.401	<0.001
$V(1)_{1,3}$	0.0048	0.0141	0.0132	0.390	0.376	10.048	154.365	<0.001
$V(1)_{1,4}$	0.0047	0.0144	0.0133	0.378	0.370	10.540	167.278	<0.001
$V(1)_{1,5}$	0.0047	0.0148	0.0145	0.301	0.361	12.153	214.203	<0.001
$V(1)_{2,2}$	0.0059	0.0196	0.0144	0.457	0.437	9.418	134.210	<0.001
$V(1)_{2,3}$	0.0057	0.0168	0.0150	0.442	0.423	10.189	160.920	<0.001
$V(1)_{2,4}$	0.0056	0.0172	0.0154	0.425	0.402	11.215	195.515	<0.001
$V(1)_{2,5}$	0.0057	0.0175	0.0165	0.352	0.372	12.583	231.825	<0.001
$V(1)_{3,3}$	0.0067	0.0225	0.0168	0.429	0.424	10.674	182.335	<0.001
$V(1)_{3,4}$	0.0065	0.0209	0.0174	0.409	0.398	11.719	220.234	<0.001
$V(1)_{3,5}$	0.0068	0.0218	0.0188	0.354	0.355	12.903	249.432	<0.001
$V(1)_{4,4}$	0.0074	0.0272	0.0189	0.395	0.384	12.029	234.751	<0.001
$V(1)_{4,5}$	0.0078	0.0283	0.0206	0.343	0.336	12.907	252.537	<0.001
$V(1)_{5,5}$	0.0107	0.0450	0.0251	0.324	0.438	12.198	219.257	<0.001
Multivariate Normality Test			*QQ*	Skewness	Kurtosis			
			791	1,011,336	5,220			
			(<0.001)	(<0.001)	(<0.001)			

Table 1b
Summary Statistics of the Log-Space Realized Covariance Matrix of Size-Sorted Stock Returns

The table shows summary statistics of the log-space one-day realized covariance of stock returns on the five size-sorted NYSE portfolios. For each day in the series, the matrix logarithm of the realized volatility matrix is calculated. Because of the non-linear nature of the matrix logarithm function, the element $A(1)_{i,j}$ is not the logarithm of $V(1)_{i,j}$ from Table 1a. The summary statistics of the upper triangular elements of the resulting matrix are shown here. The table shows the: mean; median; standard deviation; the first-order autoregressive coefficient; the Geweke-Porter-Hudack (GPH) statistic of long memory; and the skewness and kurtosis statistics. Also shown is the asymptotic marginal significance level (*P*-value) for the Jarque-Bera test of Normality. The bottom of the table presents the QQ test statistic of multivariate normality, multivariate skewness and multivariate kurtosis test statistics as well as their marginal significance levels. The sources of the data are given in Table 1a.

						Normality Tests		
$A(1)_{i,j}$	Mean (%)	Median (%)	Std. Dev. (%)	AR(1) statistic	GPH statistic	Skewness statistic	Kurtosis statistic	Jarque-Bera *P*-value
$A(1)_{1,1}$	-11.880	-11.907	1.228	0.452	0.722	0.093	2.746	<0.001
$A(1)_{1,2}$	0.715	0.721	0.701	0.123	0.417	-0.064	2.894	0.110
$A(1)_{1,3}$	0.536	0.550	0.648	0.085	0.309	-0.051	2.712	0.001
$A(1)_{1,4}$	0.429	0.441	0.631	0.069	0.309	-0.125	3.047	0.0061
$A(1)_{1,5}$	0.301	0.313	0.608	-0.006	0.120	-0.163	2.835	<0.001
$A(1)_{2,2}$	-12.793	-12.817	1.133	0.399	0.767	0.107	2.812	0.002
$A(1)_{2,3}$	0.987	0.994	0.737	0.172	0.569	-0.066	2.914	0.137
$A(1)_{2,4}$	0.708	0.724	0.638	0.079	0.236	-0.117	2.862	0.0029
$A(1)_{2,5}$	0.454	0.477	0.598	-0.007	0.211	-0.147	2.857	<0.001
$A(1)_{3,3}$	-13.099	-13.139	1.157	0.450	0.766	0.111	2.678	<0.001
$A(1)_{3,4}$	1.160	1.191	0.684	0.111	0.332	-0.182	2.952	<0.001
$A(1)_{3,5}$	0.752	0.778	0.615	0.029	0.384	-0.185	2.802	<0.001
$A(1)_{4,4}$	-13.068	-13.114	1.109	0.415	0.665	0.245	3.062	<0.001
$A(1)_{4,5}$	1.406	1.441	0.652	0.074	0.408	-0.198	2.800	<0.001
$A(1)_{5,5}$	-11.450	-11.511	1.187	0.420	0.620	0.212	2.924	<0.001
Multivariate Normality Test			*QQ*	Skewness	Kurtosis			
			158	1,748	-6			
			<0.001	<0.001	0.001			

Table 1c
Principal Components of the Log-Space Bi-power Covariation Matrixes

The table shows the loadings for the first three principal components $a^{BP}(d,i)$, $i = 1,\ldots,3$ of the log-space bi-power covariation matrixes $A^{BP}(d)$ for $d = 1$, 5 and 20 days. The proportion of volatility explained by the component is given at the bottom of the table.

	d = 1 day			d = 5 days			d = 20 days		
$A^{BP}(d)_{i,j}$	$a^{BP}(1,1)$	$a^{BP}(1,2)$	$a^{BP}(1,3)$	$a^{BP}(5,1)$	$a^{BP}(5,2)$	$a^{BP}(5,3)$	$a^{BP}(20,1)$	$a^{BP}(20,2)$	$a^{BP}(20,3)$
$A^{BP}(d)_{1,1}$	0.420	0.218	0.234	0.447	0.446	0.413	0.454	0.499	0.542
$A^{BP}(d)_{1,2}$	-0.001	-0.033	-0.121	0.027	-0.181	0.033	0.024	-0.190	-0.174
$A^{BP}(d)_{1,3}$	-0.016	-0.054	-0.062	0.012	-0.122	0.033	0.010	-0.119	-0.146
$A^{BP}(d)_{1,4}$	-0.011	0.028	0.023	0.008	-0.040	-0.029	0.002	-0.031	-0.122
$A^{BP}(d)_{1,5}$	-0.010	-0.008	-0.048	-0.004	-0.019	0.021	-0.004	-0.018	-0.014
$A^{BP}(d)_{2,2}$	0.409	0.217	0.219	0.457	0.147	0.205	0.464	0.166	-0.270
$A^{BP}(d)_{2,3}$	0.046	-0.095	-0.223	0.066	-0.253	0.080	0.064	-0.239	-0.180
$A^{BP}(d)_{2,4}$	0.021	0.053	-0.078	0.034	-0.065	0.005	0.033	-0.056	-0.189
$A^{BP}(d)_{2,5}$	-0.028	-0.033	0.007	0.012	-0.037	0.006	0.011	-0.041	-0.018
$A^{BP}(d)_{3,3}$	0.465	0.573	-0.103	0.488	0.104	-0.021	0.486	0.075	-0.142
$A^{BP}(d)_{3,4}$	0.070	-0.019	-0.553	0.079	-0.227	0.403	0.073	-0.159	-0.096
$A^{BP}(d)_{3,5}$	-0.046	-0.154	0.353	0.011	-0.015	-0.097	0.017	-0.021	-0.126
$A^{BP}(d)_{4,4}$	0.567	-0.627	0.176	0.468	-0.019	-0.711	0.457	-0.094	-0.433
$A^{BP}(d)_{4,5}$	-0.137	0.341	-0.015	-0.002	-0.142	0.305	0.007	-0.098	-0.033
$A^{BP}(d)_{5,5}$	0.299	-0.161	-0.590	0.348	-0.760	0.082	0.349	-0.749	0.514
variance (%)	25.518	13.245	9.811	67.731	12.519	5.08	79.769	12.807	1.623

Table 1d
Summary Statistics of the Forecasting Variables

The table shows the summary statistics of the variables used to forecast the realized log-space volatilities. The principal component $a^{BP}(5,1)$ is the first principal component of the 5-day log-space bi-power covariation matrix. The principal components ($a^{BP}(20,i)$ $i=1,\ldots,3$) are the first three principal components of the 20 day log-space bi-power covariation matrixes. The daily returns when the returns are negative on the small and large stock portfolios are represented by $R_1 < 0$ and $R_5 < 0$, respectively. The return on 30-day Treasury Bill is the measure of the risk-free rate, TB_t. The difference between the end of month yield-to-maturity on BAA bonds and AAA bonds as provided by Moody's Corporate Bond Indices are the measure of the credit spread, CS_t. The term spread, TS_t, is the difference between the yield-to-maturity on Ibbotson's Long-Term Government bond portfolio and the yield-to-maturity on 3-month Treasury Bills. The dividend yield, DY_t, is the trailing 12 months' dividends divided by the price at the end of month *t* from the CRSP value-weighted total returns and the monthly CRSP value-weighted capital gains returns. The scorecard, SC_t, is a measure of the sensitivity of the market news and is constructed using CRSP value-weighted daily market returns following the procedure in McQueen and Vorkink (2004). The table shows the: mean; median; standard deviation; the first-order autoregressive coefficient; the Geweke-Porter-Hudack (GPH) statistic of long memory; and the skewness and kurtosis statistics. Also shown is the asymptotic marginal significance level (*P*-value) for the Jarque-Bera test of Normality.

						Normality Tests		
	Mean (%)	Median (%)	Std. Dev. (%)	AR(1) statistic	GPH statistic	Skewness statistic	Kurtosis statistic	Jarque-Bera *P*-value
$a^{BP}(5,1)$	0.0005	0.4288	1.6470	0.9791	0.8622	-0.8079	3.0087	<0.001
$a^{BP}(20,1)$	0.0007	0.4551	1.5836	0.9986	0.8856	-0.8564	2.9607	<0.001
$a^{BP}(20,2)$	0.0001	-0.0950	0.6347	0.9965	0.7401	0.5485	2.4721	<0.001
$a^{BP}(20,3)$	0.0001	0.0033	0.2259	0.9871	0.5586	0.3066	3.5090	<0.001
$R_1 < 0$	-0.0057	0.0000	0.0111	0.8772	0.3158	-3.4799	21.5785	<0.001
$R_5 < 0$	-0.0058	0.0000	0.0099	0.7025	0.1840	-3.1471	21.2562	<0.001
TB	0.0494	0.0502	0.0172	0.9996	0.9918	0.0165	2.7101	0.001
CS	0.0084	0.0080	0.0022	0.9960	0.9111	0.8794	2.9791	<0.001
TS	-3.7216	-3.6562	0.3434	0.9964	0.9805	0.0149	1.6225	<0.001
SC	-0.0140	0.0097	0.3272	0.9985	1.1437	-1.1776	5.1466	<0.001
DY	0.0170	0.0154	0.0112	0.9995	0.9901	0.1567	1.8919	<0.001

Table 1e
Correlation Coefficients of the Forecasting Variables

The table shows the correlation coefficients of the variables used to forecast the realized log-space volatilities. The principal component $a^{BP}(5,1)$ is the first principal component of the 5-day log-space bi-power covariation matrix. The principal components ($a^{BP}(20,i)$ $i=1,\ldots,3$) are the first three principal components of the 20 day log-space bi-power covariation matrixes. The daily returns when the returns are negative on the small and large stock portfolios are represented by $R_1 < 0$ and $R_5 < 0$, respectively. The return on 30-day Treasury Bill is the measure of the risk-free rate, TB_t. The difference between the end of month yield-to-maturity on BAA bonds and AAA bonds as provided by Moody's Corporate Bond Indices are the measure of the credit spread, CS_t. The term spread, TS_t, is the difference between the yield-to-maturity on Ibbotson's Long-Term Government bond portfolio and the yield-to-maturity on 3-month Treasury Bills. The dividend yield, DY_t, is the trailing 12 months' dividends divided by the price at the end of month *t* from the CRSP value-weighted total returns and the monthly CRSP value-weighted capital gains returns. The scorecard, SC_t, is a measure of the sensitivity of the market news and is constructed using CRSP value-weighted daily market returns following the procedure in McQueen and Vorkink (2004).

	$a^{BP}(5,1)$	$a^{BP}(20,1)$	$a^{BP}(20,2)$	$a^{BP}(20,3)$	$R1 < 0$	$R5 < 0$	TB	CS	TS	SC	DY
$a^{BP}(5,1)$	1										
$a^{BP}(20,1)$	0.948	1									
$a^{BP}(20,2)$	0.035	0.000	1								
$a^{BP}(20,3)$	0.002	-0.001	0.000	1							
$R1 < 0$	0.196	0.121	-0.213	-0.017	1						
$R5 < 0$	0.212	0.167	-0.157	0.009	0.557	1					
TB	0.310	0.329	-0.006	0.332	0.000	0.038	1				
CS	-0.533	-0.549	-0.351	-0.184	-0.042	-0.078	-0.051	1			
TS	0.370	0.395	-0.442	-0.010	0.003	0.075	0.484	0.290	1		
SC	0.569	0.575	0.370	0.283	0.096	0.119	0.419	-0.570	-0.101	1	
DY	0.008	0.009	-0.338	-0.361	0.041	0.019	-0.621	0.263	0.273	-0.425	1

Table 2
Summary Statistics from and Tests of the Latent Factor Models

The table presents summary statistics from and tests of the two latent factor MHAR-RV models of the cross section of realized stock market volatility. Panel A shows the value of the *J*-statistics associated with the Wald tests of the over-identifying restrictions of the latent factor model. The statistics are distributed as χ^2 and are presented along with their degrees of freedom and asymptotic marginal significance levels (*P*-values). Panel B shows the eigenvalue, magnitude and direction statistics from Moskowitz (2003). The eigenvalue statistic is the ratio of the eigenvalue from the fitted value of the model to the eigenvalue of the realized covariance matrix. The statistic measures how well the fitted values of the model capture the total variation in the real volatility matrix. The numerator of the magnitude statistic is the sum of the absolute values of the difference between the fitted and true volatility matrix. The denominator is the sum of the absolute values of the true volatility matrix. The statistic measures the per cent error of the fitted values of the model. The direction statistic is the ratio of the sum of the signed fitted values of the model to the squared rank of the true volatility matrix. It is designed to capture how well the model predicts the correct direction of volatility. Also shown are multivariate measures of root means square error (RMSE) and mean absolute deviation (MAD), both multiplied by 1000. The forecasting variables used in each model are described in Table 5. The model is estimated separately for each set of forecasting variables by Generalized Method of Moments. The Newey-West standard errors (in parentheses) are asymptotically robust to general forms of heteroskedasticity and autocorrelation.

Model:	MHAR-RV-BP	MHAR-RV-BPA	MHAR-RV-X	MHAR-RV-BPAX
(A) *J* statistics from overidentifying restrictions				
χ^2 statistic	84.63	126.28	153.64	319.74
df	26	52	39	117
P-value	<0.001	<0.001	<0.001	<0.001
(B) Measures of fit				
Eigenvalue statistic				
average	0.934	0.911	1.007	0.929
(std. err.)	(0.016)	(0.016)	(0.021)	(0.017)
Magnitude statistic				
average	0.998	0.977	1.094	1.006
(std. err.)	(0.016)	(0.015)	(0.021)	(0.017)
Direction statistic				
average	0.958	0.958	0.958	0.959
(std. err.)	(0.002)	(0.002)	(0.002)	(0.002)
RMSE x 1000	2.586	2.529	2.634	2.546
MAD x 1000	12.623	12.477	13.020	12.672

Table 3
Variance Ratios

The table presents variance ratio measures of the statistical fit of the two latent factor MHAR-RV models of the cross section of realized stock market volatility. The ratio shows how the latent factor model captures the variation in the data. In panel (A), the numerator is the variance of the fitted values of the $(i,j)^{th}$ element of the log-space volatilities A (1) from the latent factor model, while the denominator is the variance of the fitted values of the log-space volatilities from an OLS regression of the volatility on the variables used in that model. In panel (B), the numerator is the variance of the fitted values of the $(i,j)^{th}$ element of the realized volatilities V from the latent factor model, while the denominator is the variance of the (bias-corrected) fitted values of the realized volatilities from an OLS regression of the volatility on the variables used in that model. For this panel, the off-diagonal elements are the Fisher transforms of the fitted correlation in the numerator and the actual correlation in the denominator. The average value of the ratios for the realized volatilities (avg. vol.) and realized correlations (avg. cor.) are shown at the bottom. The model is estimated separately for each set of forecasting variables by Generalized Method of Moments.

Model:	MHAR-RV-BP	MHAR-RV-BPA	MHAR-RV-X	MHAR-RV-BPAX
(A) Log-space volatilities				
$A\ (1)_{1,1}$	0.964	0.944	1.183	0.978
$A\ (1)_{1,2}$	0.902	0.552	0.648	0.536
$A\ (1)_{1,3}$	0.708	0.401	0.012	0.220
$A\ (1)_{1,4}$	0.126	0.014	0.504	0.004
$A\ (1)_{1,5}$	0.738	0.507	1.067	0.300
$A\ (1)_{2,2}$	1.028	1.010	1.162	1.066
$A\ (1)_{2,3}$	0.919	0.884	1.036	1.073
$A\ (1)_{2,4}$	0.953	0.475	0.441	0.366
$A\ (1)_{2,5}$	1.454	0.876	0.461	0.313
$A\ (1)_{3,3}$	1.039	1.005	1.137	0.994
$A\ (1)_{3,4}$	1.027	0.948	1.265	0.924
$A\ (1)_{3,5}$	1.987	1.010	0.846	0.529
$A\ (1)_{4,4}$	1.008	1.012	1.198	1.046
$A\ (1)_{4,5}$	0.186	0.101	0.929	0.236
$A\ (1)_{5,5}$	0.983	0.950	0.938	0.975
(B) Realized Volatilities				
$V\ (1)_{1,1}$	1.136	1.088	1.311	1.107
$V\ (1)_{1,2}$	1.177	0.837	0.818	0.685
$V\ (1)_{1,3}$	1.006	0.743	0.203	0.511
$V\ (1)_{1,4}$	0.750	0.520	0.045	0.328
$V\ (1)_{1,5}$	0.584	0.527	0.353	0.076
$V\ (1)_{2,2}$	1.246	1.184	1.357	1.246
$V\ (1)_{2,3}$	1.251	1.139	1.428	1.284
$V\ (1)_{2,4}$	1.212	0.911	0.920	0.736
$V\ (1)_{2,5}$	0.961	0.887	0.707	0.303
$V\ (1)_{3,3}$	1.209	1.162	1.320	1.195
$V\ (1)_{3,4}$	1.027	0.885	1.099	0.554
$V\ (1)_{3,5}$	3.008	1.493	1.094	0.438
$V\ (1)_{4,4}$	1.114	1.070	1.267	1.064
$V\ (1)_{4,5}$	0.555	0.325	1.131	0.363
$V\ (1)_{5,5}$	1.038	1.001	0.978	0.894
avg. vol.	1.149	1.101	1.247	1.101
avg. cor.	1.153	0.827	0.780	0.528

Table 4
Beta Coefficients on the Implied Volatility Factors

The table shows the β coefficients from the two latent factor MHAR-RV models of the cross section of realized stock market volatility. The latent volatility factors are the linear combination of the variables given in Table 5. The beta coefficients for the first two elements of the log-space volatility matrix have been normalized for identification. The model is estimated separately for each set of forecasting variables by Generalized Method of Moments. The coefficients for the MHAR-RV-BPAX model described in Table 5 are presented here. The coefficients for the models using other variables are similar to the values shown here. The Newey-West standard errors (in parentheses) are asymptotically robust to general forms of heteroskedasticity and autocorrelation. The *t* statistics are shown below the standard errors.

	$A(1)_{i.1}$	$A(1)_{i.2}$	$A(1)_{i.3}$	$A(1)_{i.4}$	$A(1)_{i.5}$
	First Factor				
$A(1)_{1,j}$	1.000				
$A(1)_{2,j}$	0.000	0.742			
		(0.045)			
		16.55			
$A(1)_{3,j}$	-0.043	-0.004	0.660		
	(0.015)	(0.032)	(0.062)		
	-2.99	-0.13	10.64		
$A(1)_{4,j}$	-0.009	0.068	0.074	0.515	
	(0.015)	(0.014)	(0.021)	(0.068)	
	-0.58	4.78	3.45	7.58	
$A(1)_{5,j}$	-0.016	0.034	0.080	0.059	0.066
	(0.013)	(0.013)	(0.017)	(0.020)	(0.091)
	-1.26	2.57	4.83	2.93	0.72
	Second Factor				
$A(1)_{1,j}$	0.000				
$A(1)_{2,j}$	1.000	2.310			
		(0.288)			
		8.01			
$A(1)_{3,j}$	0.451	2.040	3.492		
	(0.086)	(0.208)	(0.390)		
	5.22	9.83	8.94		
$A(1)_{4,j}$	0.021	0.179	1.021	3.980	
	(0.088)	(0.089)	(0.133)	(0.448)	
	0.24	2.01	7.66	8.89	
$A(1)_{5,j}$	-0.267	-0.120	-0.596	-0.753	5.754
	(0.075)	(0.076)	(0.097)	(0.121)	(0.593)
	-3.54	-1.58	-6.16	-6.23	9.70

Table 5
Coefficients on the Forecasting Variables

The table shows the θ coefficients on the forecasting variables in the two latent factor MHAR-RV models of the cross section of realized stock market volatility. The variables for the MHAR-RV-BP model (Panel A) are: the lagged value of the first principal component of the 5 day log-space bi-power variation matrix ($a^{BP}(5,1)$) and the lagged values of the first three principal components of the 20 day log-space bi-power variation matrixes ($a^{BP}(20,1)$, $a^{BP}(20,2)$, and $a^{BP}(20,3)$). In the MHAR-RV-BPA model (Panel B) the variables are those in the MHAR-RV-BP model plus the lagged values of the negative returns on the small and large stock portfolios ($R_1 < 0$ and $R_5 < 0$). The coefficients on the constant terms are not shown. The model is estimated separately for each set of forecasting variables by Generalized Method of Moments. The Newey-West standard errors (in parentheses) are asymptotically robust to general forms of heteroskedasticity and autocorrelation. The *t* statistics are reported below the standard errors.

	$a^{BP}(5,1)$	$a^{BP}(20,1)$	$a^{BP}(20,2)$	$a^{BP}(20,3)$	$R_1 < 0$	$R_5 < 0$
(A) MHAR-RV-BP						
θ	0.153	0.256	0.499	0.176		
(std err)	(0.028)	(0.028)	(0.032)	(0.053)		
t-stat	5.41	9.08	15.68	3.31		
θ	0.050	-0.003	-0.171	0.037		
(std err)	(0.009)	(0.008)	(0.021)	(0.015)		
t-stat	5.49	-0.31	-8.29	2.52		
(B) MHAR-RV-BPA						
θ	0.116	0.284	0537	0.170	-7.324	2.663
(std err)	(0.028)	(0.027)	(0.029)	(0.052)	(1.239)	(1.332)
t-stat	4.21	10.36	18.66	3.27	-5.91	2.00
θ	0.037	0.008	-0.133	0.033	0.400	-2.229
(std err)	(0.007)	(0.007)	(0.019)	(0.011)	(0.315)	(0.399)
t-stat	5.02	1.12	-6.92	2.97	1.27	-5.58

Table 5, continued
Coefficients on the Forecasting Variables

The table shows the θ coefficients on the forecasting variables in the two latent factor MHAR-RV models of the cross section of realized stock market volatility. The variables in the MHAR-RV-X model (Panel C) are: the lagged interest rate on a short-term Treasury bill (*TB*); the lagged credit spread (*CS*); the lagged 'scorecard' measure of the sensitivity of the market to news (*SC*); the lagged dividend yield (*DY*); and the lagged spread between the long Government bond and the 3 month Treasury bill (*TS*). The MHAR-RV-BPAX model (Panel D) includes all of the variables in the other models, The coefficients on the constant terms are not shown. The model is estimated separately for each set of forecasting variables by Generalized Method of Moments. The Newey-West standard errors (in parentheses) are asymptotically robust to general forms of heteroskedasticity and autocorrelation. The *t* statistics are reported below the standard errors.

	$a^{BP}(5,1)$	$a^{BP}(20,1)$	$a^{BP}(20,2)$	$a^{BP}(20,3)$	$R_1 < 0$	$R_5 < 0$	*TB*	*CS*	*TS*	*SC*	*DY*
(C) MHAR-RV-X											
θ							11.821	183.94	-0.269	-1.202	-0.856
(std err)							(2.742)	(10.857)	(3.586)	(0.078)	(0.126)
t-stat							4.31	16.94	-0.08	-15.49	-6.80
θ							0.868	16.329	-0.690	-0.069	-0.243
(std err)							(0.388)	(3.444)	(0.447)	(0.025)	(0.040)
t-stat							2.24	4.74	-1.54	-2.79	-6.06
(D) MHAR-RV-BPAX											
θ	0.154	0.257	0.413	0.243	-4.517	2.463	5.170	9.889	6.580	-0.161	0.133
(std err)	(0.027)	(0.029)	(0.026)	(0.059)	(1.304)	(1.323)	(1.886)	(8.454)	(2.345)	(0.062)	(0.088)
t-stat	5.65	8.77	15.66	4.16	-3.46	1.86	2.74	1.17	2.81	-2.58	1.50
θ	0.009	0.014	-0.061	0.023	-1.368	-1.895	0.205	9.729	-0.803	0.011	-0.213
(std err)	(0.006)	(0.007)	(0.011)	(0.012)	(0.281)	(0.344)	(0.370)	(1.858)	(0.491)	(0.013)	(0.028)
t-stat	1.65	2.07	-5.45	1.87	-4.87	-5.51	0.56	5.24	-1.64	0.83	-7.54

Table 6
Elasticities of the Forecasting Variables

The table shows the estimated means and standard errors of the time series of elasticities *ε(i,j,z,t)* associated with the forecasting variables in the two latent factors model of the cross section of realized stock market volatility. The elasticity for the $(i,j)^{th}$ element of the realized volatility matrix *V* are presented. The variables (*z*) for the four models are described in Table 5. The model is estimated separately by Generalized Method of Moments for each set of forecasting variables. The Newey-West standard errors (in parentheses) are asymptotically robust to general forms of heteroskedasticity and autocorrelation.

	$a^{BP}(5,1)$	$a^{BP}(20,1)$	$a^{BP}(20,2)$	$a^{BP}(20,3)$	$R_1 < 0$	$R_5 < 0$
(A) MHAR-RV-BP						
$V(1)_{1,1}$	0.369	0.350	0.084	0.050		
	(0.001)	(<0.001)	(0.002)	(<0.001)		
$V(1)_{1,5}$	0.298	0.156	-0.128	0.036		
	(0.001)	(<0.001)	(0.001)	(<0.001)		
$V(1)_{5,5}$	0.463	0.134	-0.364	0.052		
	(<0.001)	(0.001)	(0.001)	(<0.001)		
(B) MHAR-RV-BPA						
$V(1)_{1,1}$	0.292	0.421	0.128	0.049	-0.076	-0.026
	(0.001)	(<0.001)	(0.002)	(<0.001)	(<0.001)	(0.001)
$V(1)_{1,5}$	0.244	0.208	-0.108	0.036	-0.023	-0.061
	(0.001)	(<0.001)	(0.001)	(<0.001)	(<0.001)	(<0.001)
$V(1)_{5,5}$	0.379	0.199	-0.350	0.051	0.001	-0.126
	(<0.001)	(0.001)	(0.001)	(<0.001)	(<0.001)	(<0.001)

Table 6, continued
Elasticities of the Forecasting Variables

The table shows the estimated means and standard errors of the time series of elasticities *ε(i,j,z,t)* associated with the forecasting variables in the two latent factors model of the cross section of realized stock market volatility. The elasticity for the $(i,j)^{th}$ element of the realized volatility matrix *V* are presented. The variables (*z*) for the four models are described in Table 5. The model is estimated separately by Generalized Method of Moments for each set of forecasting variables. The Newey-West standard errors (in parentheses) are asymptotically robust to general forms of heteroskedasticity and autocorrelation.

	$a^{BP}(5,1)$	$a^{BP}(20,1)$	$a^{BP}(20,2)$	$a^{BP}(20,3)$	$R_1 < 0$	$R_5 < 0$	*TB*	*CS*	*TS*	*SC*	*DY*
(C) MHAR-RV-X											
$V(1)_{1,1}$							0.208	0.426	-0.011	-0.394	-0.367
							(<0.001)	(<0.001)	(<0.001)	(<0.001)	(0.001)
$V(1)_{1,5}$							0.121	0.253	-0.013	-0.224	-0.263
							(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
$V(1)_{5,5}$							0.140	0.307	-0.033	-0.244	-0.454
							(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
(D) MHAR-RV-BPAX											
$V(1)_{1,1}$	0.253	0.402	0.158	0.058	-0.091	-0.018	0.085	0.061	0.048	-0.039	-0.099
	(<0.001)	(<0.001)	(0.001)	(<0.001)	(<0.001)	(0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(0.001)
$V(1)_{1,5}$	0.140	0.219	0.014	0.035	-0.070	-0.040	0.045	0.061	0.008	-0.013	-0.157
	(<0.001)	(0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(0.001)
$V(1)_{5,5}$	0.157	0.242	-0.090	0.042	-0.107	-0.090	0.046	0.108	-0.017	-0.001	-0.323
	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)

Table 7
Global Minimum Variance and Tracking Error Portfolio Statistics

The table presents statistics on the performance of the global minimum variance and tracking error portfolios. The global minimum variance portfolio is the weighted sum of the five size-sorted stock portfolios with the smallest expected volatility. The table shows the return and standard deviation of the portfolio. The tracking error portfolio is designed to minimize the volatility of the difference between the weighted sum of the five size-sorted stock portfolios and the HML "value" portfolio (which takes a long position in value (high book-to-market) stocks and a short position in growth (low book-to-market) stocks). The table shows the average difference in returns between the two portfolios (Ex. Return), and the standard deviation of the difference. Also shown are the statistics for the equally-weighted portfolio and the equally-weighted less HML tracking error portfolio, both in and out-of-sample. The models of the conditional covariance matrix are estimated using the first 13 years of the data (1988–2000) and tested in-sample. In addition, the models are estimated and the portfolios evaluated out-of-sample (2001–2002) as noted in the first column. Weights in both the minimum-variance and tracking portfolio exercises are either constrained to be $-0.30 \leq w_i \leq 1.30$ or are left unconstrained as noted in the second column.

			Minimum Variance Portfolio		HML Tracking Error Portfolio	
Sample	Constraints	Model	Return (%)	Std Dev (%)	Ex. Return (%)	Std Dev (%)
In	Yes	MHAR-RV-BP	0.0461	0.8298	0.0562	1.3048
		MHAR-RV-BPA	0.0461	0.8244	0.0571	1.3067
		MHAR-RV-X	0.0547	0.8165	0.0642	1.3017
		MHAR-RV-BPAX	0.0527	0.8238	0.0636	1.3090
Out	Yes	MHAR-RV-BP	-0.0256	1.3548	-0.0993	2.0679
		MHAR-RV-BPA	-0.0423	1.3297	-0.1094	2.0602
		MHAR-RV-X	0.0180	1.3457	-0.0813	2.0403
		MHAR-RV-BPAX	-0.0139	1.3536	-0.0941	2.1281
In	No	MHAR-RV-BP	0.0447	0.8674	0.0552	1.3259
		MHAR-RV-BPA	0.0450	0.8538	0.0576	1.3168
		MHAR-RV-X	0.0544	0.8482	0.0644	1.3240
		MHAR-RV-BPAX	0.0534	0.8540	0.0600	1.3206
Out	No	MHAR-RV-BP	-0.0303	1.3769	-0.0993	2.0822
		MHAR-RV-BPA	-0.0489	1.3514	-0.1020	2.0842
		MHAR-RV-X	0.0235	1.3487	-0.0682	2.0456
		MHAR-RV-BPAX	-0.0165	1.3500	-0.0911	2.1562
			Equally-weighted portfolio		(Equally-weighted - HML) portfolio	
In			0.0515	0.8243	0.0555	1.3278
Out			-0.0250	1.3630	-0.0865	2.0749

Appendix Table A1
Summary Statistics of Volatility Estimates

The table presents estimates of annualized return standard deviations for the five size-sorted portfolios using both TAQ and CRSP data. The first column of data reports the annualized return standard deviation constructed from the average daily realized volatility values provided in Table 1a. The second column reports the annualized return standard deviation constructed from daily portfolio returns using end of day prices constructed following a similar procedure used to construct the 20-minute prices. The last two columns report annualized return standard deviations constructed from CRSP daily and monthly returns respectively.

	TAQ		CRSP	
	20-Min	Daily	Daily	Monthly
1 (Small)	0.0926	0.1228	0.1232	0.2098
2	0.0939	0.1265	0.1492	0.1933
3	0.0991	0.1332	0.1560	0.1767
4	0.1041	0.1376	0.1556	0.1662
5 (Large)	0.1240	0.1567	0.1642	0.1497

Appendix Table A2
Summary Statistics of Estimated Monthly Correlations

The table presents summary statistics of monthly correlations between all pairs of the five size-sorted stock portfolios. The correlations are labeled $C(m)_{i,j}$, where i=1 contains the smallest and i=5 contains the largest stocks and m is the method used to produce the correlations. The two methods for constructing the correlations are: (i) using daily CRSP data on the five size-sorted stock returns to estimate the correlations each month ("daily data"); and, (ii) using the realized volatility data from Table 1 to estimate a correlation each day and then average the correlations across the month ("realized volatility"). The table shows the: mean, standard deviation, and a volatility ratio (i.e., the ratio of the standard deviation to the mean of the Fisher transform of the monthly correlation). The final column shows an average of the monthly t-statistics on the equivalence between the two ways of measuring the correlations, applied to the Fisher transforms of the monthly correlations.

	m = daily data			m = realized volatility			
	mean	std. dev.	vol. ratio	mean	std. dev.	vol. ratio	mean t-stat.
$C(m)_{1,2}$	0.737	0.173	0.379	0.649	0.133	0.304	1.645
$C(m)_{1,3}$	0.685	0.194	0.410	0.610	0.148	0.329	1.562
$C(m)_{1,4}$	0.647	0.212	0.440	0.570	0.144	0.324	1.454
$C(m)_{1,5}$	0.560	0.216	0.479	0.479	0.138	0.340	1.271
$C(m)_{2,3}$	0.880	0.089	0.262	0.791	0.094	0.246	1.871
$C(m)_{2,4}$	0.837	0.119	0.283	0.735	0.101	0.242	1.685
$C(m)_{2,5}$	0.729	0.150	0.315	0.612	0.122	0.276	1.433
$C(m)_{3,4}$	0.923	0.056	0.211	0.840	0.076	0.215	2.083
$C(m)_{3,5}$	0.802	0.113	0.251	0.712	0.115	0.255	1.338
$C(m)_{4,5}$	0.882	0.085	0.229	0.819	0.094	0.238	1.466

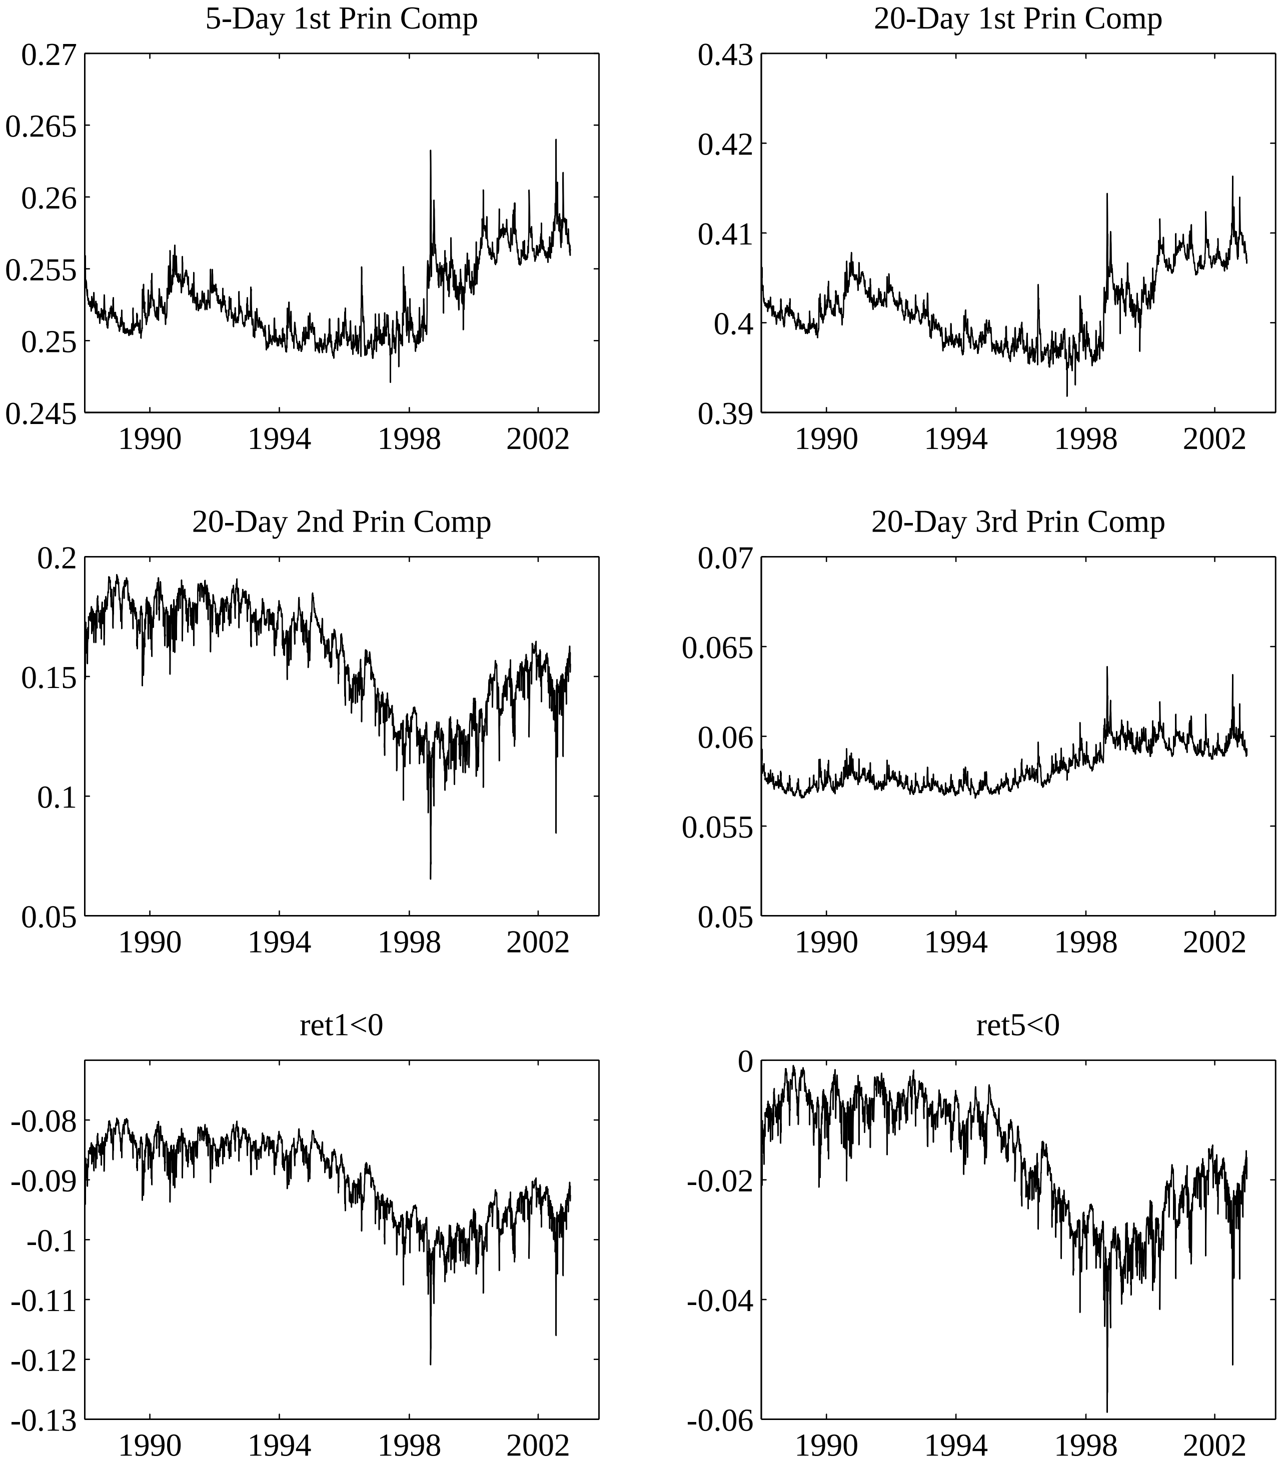
5-Day 1st Prin Comp
0.27
0.265
0.26
0.255
0.25
0.245
1990
1994
1998
2002
20-Day 1st Prin Comp
0.43
0.42
0.41
0.4
0.39
1990
1994
1998
2002
20-Day 2nd Prin Comp
0.2
0.15
0.1
0.05
1990
1994
1998
2002
20-Day 3rd Prin Comp
0.07
0.065
0.06
0.055
0.05
1990
1994
1998
2002
ret1<0
-0.08
-0.09
-0.1
-0.11
-0.12
-0.13
1990
1994
1998
2002
ret5<0
0
-0.02
-0.04
-0.06
1990
1994
1998
2002

Figure 1a

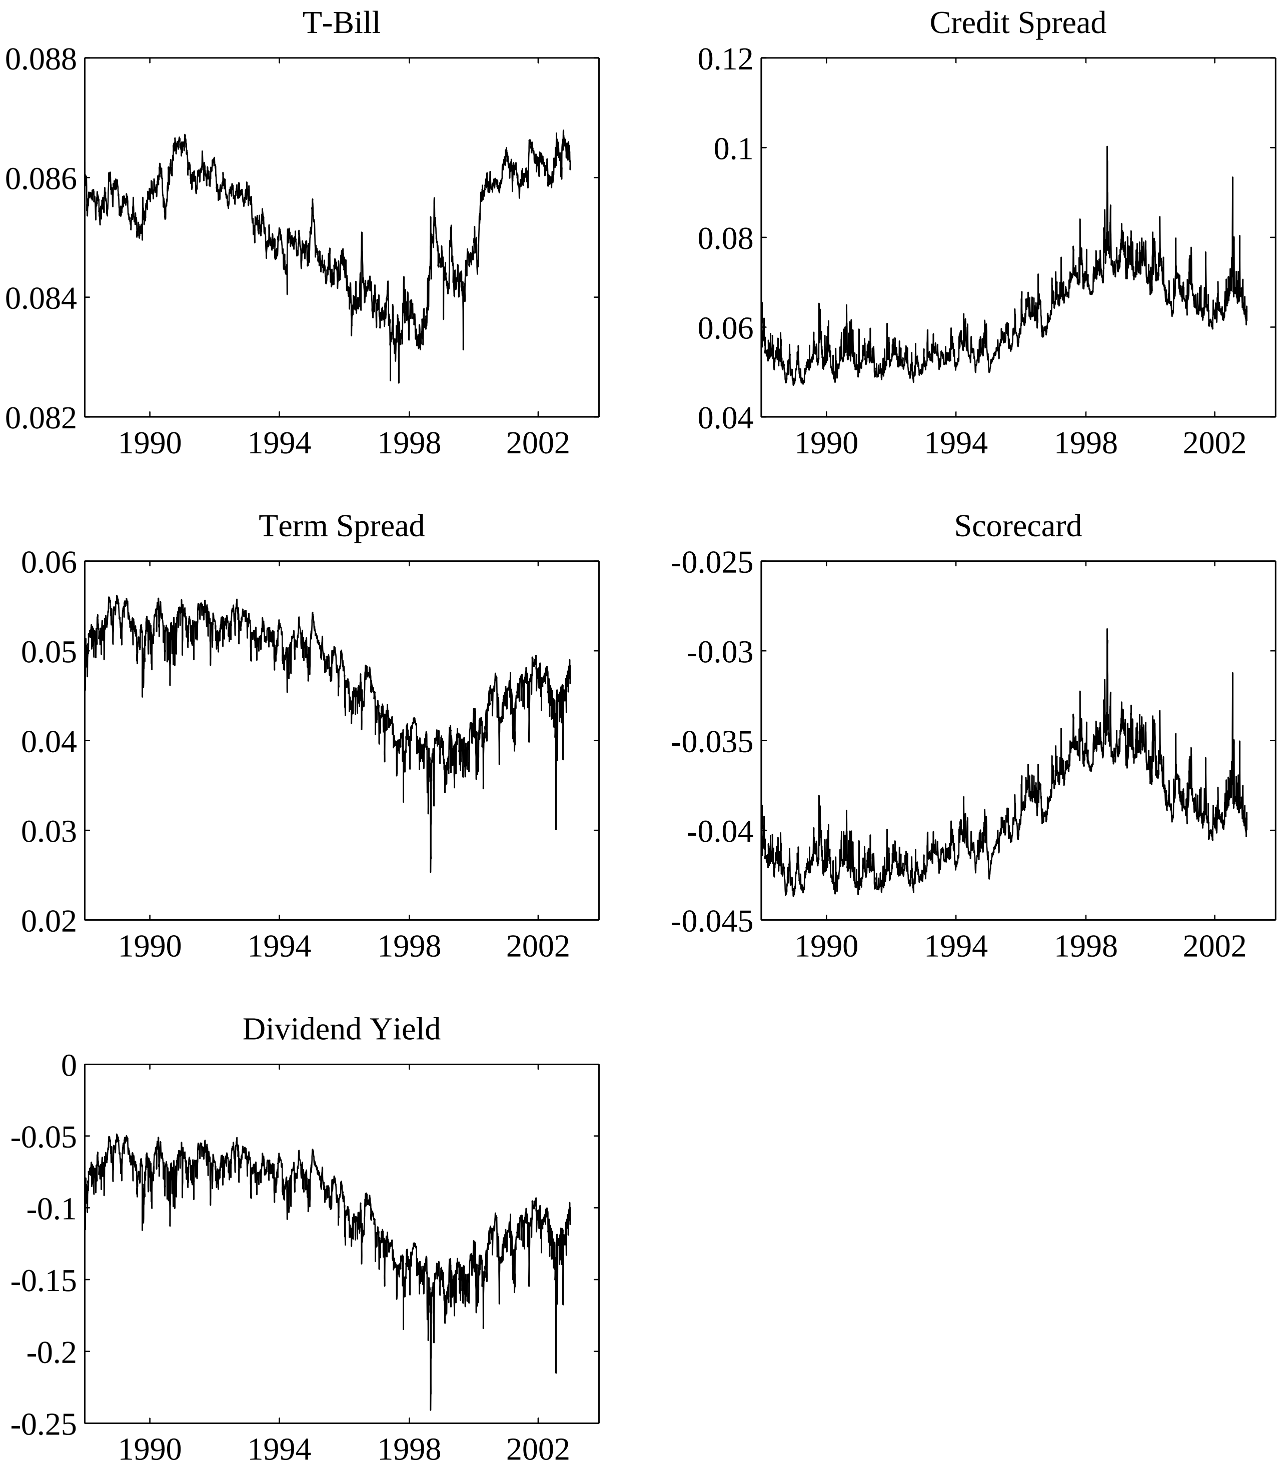
T-Bill
0.088
0.086
0.084
0.082
Credit Spread
0.12
0.1
0.08
0.06
0.04
Term Spread
0.06
0.05
0.04
0.03
0.02
Scorecard
-0.025
-0.03
-0.035
-0.04
-0.045
Dividend Yield
0
-0.05
-0.1
-0.15
-0.2
-0.25
1990
1994
1998
2002

Figure 1b