Toward Better Risk Forecasts

Using data with higher frequency than the forecasting horizon.

Kenneth Gosier, Ananth Madhavan, Vitaly Serbin, and Jian Yang

Volatility forecasts are crucial to financial applications ranging from portfolio construction to option pricing. They are typically constructed under the implicit premise that the sampling interval should match the forecast horizon. For example, most portfolio managers rely on monthly volatility estimates computed using returns sampled at monthly intervals. (See, e.g., Fama and French [1993] and Chan, Karceski, and Lakonishok [1999].)

Although this practice is intuitively appealing, it has some drawbacks. First, the use of low-frequency data such as monthly or quarterly information can induce staleness into forward-looking volatility estimates. Practitioners who recognize this problem sometimes place a greater weight on recent returns, but this practice does not eliminate dependence on events in the distant past. Risk models estimated with monthly data, for example, typically use 60 data periods, implying that estimates reflect events as long as five years ago.

Second, and more important, return events occurring in long sampling intervals are obscured, thus confounding estimation. A month with a gradual fall in asset prices is very different from a month of rising prices and a sudden fall on the last day, even if overall monthly returns are identical.

These concerns are important for portfolio managers, motivating our examination of the relation between data frequency and forecast horizon. We propose using daily data to construct monthly volatility estimates and forecasts.[1]

Kenneth Gosier is a research analyst at Putnam Investments in Boston.
kenneth_gosier@putnam.com

Ananth Madhavan is a managing director at Barclays Global Investors in San Francisco.
ananth.madhavan@barclaysglobal.com

Vitaly Serbin is a research analyst at ITG, Inc., in Boston.
vserbin@itginc.com

Jian Yang is co-head of financial engineering at ITG, Inc., in Boston.
jyang@itginc.com

The use of daily data would seem to enhance monthly forecasts by providing more timely information that accurately captures all return events in a month. Sampling at a daily frequency yields more information.

Holding the number of observations in each forecast constant, the point estimates of variance using different data frequencies are affected by return autocorrelation, heteroscedasticity, and microstructure effects such as bid-ask bounce or non-trading effects. Accordingly, the question of choice of data frequency is ultimately an empirical question. Our analysis shows that additional information in higher-frequency data permits the use of simpler and more intuitive forecasting models.

Most important, we propose a new approach to constructing monthly volatility forecasts using daily data. We show empirically that monthly risk measures using daily data are more robust than those using lower-frequency monthly data. This finding is consistent with other research predicting the volatility of a single asset or index using intraday returns. Microstructure effects explain the difference in actual estimates.[2]

MONTHLY VERSUS DAILY RETURNS

The rationale behind using daily returns to measure monthly volatilities has its roots in Merton [1980]. The basic idea is that monthly returns provide an unbiased but noisy estimate of monthly volatility. Theoretically, sampling returns with infinite frequency provides a perfect volatility measure, but in practice microstructure effects limit its usefulness (see Andersen, Bollerslev, and Lange [1999]).

Intuition

First, we need to define a measure of monthly volatility using different sampling intervals. Formally, we define a monthly risk measure for period t (computed with monthly returns) as $\sigma^2_{t,m} = r^2_{t,m}$ and define the corresponding measure using daily returns as:

$$\sigma^2_{t,d} = \sum_{\tau \in t} r^2_{\tau,d}$$

where $r_{t,m}$ is the monthly return of an asset in month t and $r_{\tau,d}$ the daily return on day τ in month t. The latter measure was first used by French, Schwert, and Stambaugh [1987] and then developed further in Schwert [1998, 2002].

Assume an AR(1) process described as follows:

$$\begin{aligned} r_{\tau,d} &= \rho r_{\tau-1,d} + \varepsilon_{\tau,d} \\ Var(\varepsilon_{\tau,d}) &= \sigma^2 \end{aligned} \tag{1}$$

where ρ is the autocorrelation parameter in the interval (−1, +1), and $\varepsilon_{\tau,d}$ is the disturbance term.

We assume the unconditional mean of $r_{\tau,d}$ to be zero. This is a common assumption in working with daily returns; allowing for a non-zero mean does not change the basic result. We also assume a constant volatility process for ε to simplify the exposition.

Assuming N daily observations in a month, the unconditional mean monthly volatility estimate is:

$$E(\sigma^2_{t,d}) = E\left(\sum_{\tau \in t} r^2_{\tau,d}\right) = N\frac{\sigma^2}{1-\rho^2} \tag{2}$$

measured using daily returns. The corresponding estimate of volatility using monthly returns is given by:

$$\begin{aligned} E(\sigma^2_{t,m}) &= E\left(\left[\sum_{\tau \in t} r_{\tau,d}\right]^2\right) \\ &= N\frac{\sigma^2}{1-\rho^2} + 2\sum_{i \neq j} E(r_{i,d} r_{j,d}) \end{aligned} \tag{3}$$

where the last term captures the covariance between successive returns. This latter term will reflect serial correlation, if present, and microstructure frictions that might confound the estimation process (see Madhavan, Richardson, and Roomans [1997]).

This simple decomposition shows that the two estimates of volatility differ only in terms of high-frequency microstructure effects. Sampling at a daily frequency yields more information than a monthly sample. The point estimates of variance using different data frequencies are a function of the extent of microstructure frictions and return autocorrelation.

Exhibit 1 reports statistics of interest for a variety of portfolios: factor-mimicking portfolios for the market (proxied by the S&P 500), size, and value/growth, and the six Fama and French size/book-to-market portfolios. We also report average statistics for the individual securities constituting the DJIA and the Russell 2000.[3]

The estimated monthly volatility computed using daily returns is lower than the volatility computed using monthly returns for portfolios, but it is higher for indi-

EXHIBIT 1

Sample Statistics for Monthly Volatility—1963–2002

	Mean monthly (*100)	Mean daily (*100)	St. dev. Monthly (*100)	St. dev. daily (*100)	abs (t-stat) on null monthly = daily
Bootstrapped S&P 500	0.2065	0.1949	0.4083	0.3309	0.7519
Bootstrapped SMB	0.0950	0.0479	0.2000	0.0768	6.061***
Bootstrapped HML	0.0957	0.0532	0.2249	0.1030	4.808***
Other Portfolios					
Small – Low BM	0.5095	0.2394	0.9041	0.3781	7.817***
Small – Medium BM	0.3058	0.1204	0.5598	0.1858	8.480***
Small – High BM	0.3286	0.1140	0.6566	0.1823	8.269***
Big – Low BM	0.2596	0.2265	0.5034	0.3186	0.1613
Big – Medium BM	0.2006	0.1603	0.3979	0.2973	2.726*
Big – High BM	0.2196	0.1576	0.3996	0.2485	3.917**
Individual Stocks					
DJIA components	0.7352	0.7814	1.4066	0.8692	1.065
Russell 2000 components	3.1036	3.3472	27.895	27.674	0.724

*** *Significant at 1% level.*

** *Significant at 5% level.*

* *Significant at 10% level.*

vidual stocks. Statistically, however, the difference between the two measures is not significant for the market portfolio, for the big/low BM portfolio, or for individual stock averages.

The economic significance of the difference in the mean estimates is generally quite low, although there are significant differences for some portfolios. For the value/growth portfolio, for instance, the difference between annualized standard deviations is 2.73%.

Of greater interest, the difference in the *precision* of the two measures is consistently in favor of using daily data. The standard deviations of the measure using daily data range from 28% (for the small/high BM portfolio) to 81% (for the market portfolio) of the standard deviations of the measure using monthly data. The average precision of the two measures for the Russell 2000 stocks is roughly the same, most likely because of noise inherent in estimates for individual stocks.

The daily return autocorrelation is a critical determinant of the difference in the mean values of the two measures. Specifically, the autocorrelation functions (ACFs) for factor portfolios in Exhibit 1 are positive, reflecting non-trading (Lo and MacKinlay [1990]), but for small stocks the majority of autocorrelation coefficients up to lag 20 are negative, consistent with French and Roll [1986] and Jegadeesh and Titman [1995]. A mixture of information effects from the underreaction of traders to the new information provided (inducing positive autocorrelation) and bid-ask bounce (inducing negative autocorrelation) can explain these findings, as in Madhavan, Richardson, and Roomans [1997].

Reconciling Competing Measures of Monthly Volatility

The evidence from the ACFs suggests adding a temporal cross-product term when computing monthly volatility with daily returns to capture the autocorrelation of daily returns induced by microstructure effects. We add the cross-products through a Newey-West procedure in order to ensure positive semidefiniteness of the covariance matrix:

$$S(q) = \Gamma_0 + \sum_{j=1}^{q} \frac{q-j}{q} (\Gamma_j + \Gamma_j')$$

where

$$\Gamma_j = E[r_{\tau,d} r_{\tau+1-j,d}]$$

We use $q = 2$, $j = 1$ (only the first cross-product considered) for most portfolios and $q = 5$, $j = 4$ (autocorrela-

EXHIBIT 2

Sample Statistics for Monthly Volatility Taking into Account Daily Autocorrelation

	Mean monthly (*100)	Mean daily (*100)	St. dev. monthly (*100)	St. dev. daily (*100)	abs (t-stat) on null monthly = daily
Factors					
Bootstrapped S&P 500 (NW q = 2, j = 1)	0.2065	02105	0.3987	0.3357	0.3286
Bootstrapped SMB (NW q = 2, j = 1)	0.0950	0.0580	0.2000	0.0893	4.968***
Bootstrapped HML (NW q = 2, j = 1)	0.0957	0.0572	0.2249	0.1014	4.689***
Other Portfolios					
Small-Low BM (NW q = 2, j = 1)	0.5095	0.2995	0.9041	0.4627	6.349***
Small-Medium (NW q = 2, j = 1)	0.3058	0.1531	0.5598	0.2321	7.416***
Small-High (NW q = 2, j = 1)	0.3286	0.1470	0.6566	0.2319	7.389***
Big-Low (NW q = 2, j = 1)	0.2596	0.2522	0.5034	0.3475	0.3320
Big-Medium (NW q = 2, j = 1)	0.2006	0.1817	0.3979	0.3195	0.616
Big-High (NW q = 2, j = 1)	0.2196	0.1789	0.3996	0.2647	2.609*
Individual Stocks					
DJIA components (NW q = 5, j = 4)	0.7352	0.7757	1.4066	0.8878	1.081
Russell 2000 components (NW q = 2, j = 1)		3.2739		27.448	0.535

*** *Significant at 1% level.*

** *Significant at 5% level.*

* *Significant at 10% level.*

tions up to lag four considered) for DJIA stocks.

Exhibit 2 demonstrates that including covariance terms reduces the difference between the two measures of monthly volatility, and in some cases even makes the difference statistically insignificant (three Fama-French portfolios of large firms).

We conclude that, in accordance with theoretical models and earlier empirical results, volatility computed using higher-frequency data is less noisy than volatility computed using same-frequency returns. The two measures can differ in their means because of microstructure effects. Using daily returns without cross-products effectively removes these effects, and produces a better measure than using monthly returns. But is this superiority practically important in real-world financial applications?

FORECASTING MONTHLY VOLATILITY

To examine the forecasting capabilities of models that use daily versus monthly measures of the monthly variance, we apply the widely known Fama-French three-factor model, which includes market, size, and value/growth as factors. We focus first on forecasting individual factor volatility and then turn to risk models and portfolio volatility.

Forecasting Single-Asset Volatility

The factors considered here are the S&P 500 index, SMB, and HML for data from January 1963 through December 2002. For both the daily and monthly measures, we consider three forecasts, which are simple and feasible when estimating a factor risk model with a large number of stocks:

Forecast 1: $\sigma^2_{t+1} = \sigma^2_t$, which assumes a random walk (RW);

Forecast 2: An AR(1) process, $\sigma^2_{t+1} = \sigma^2_0 + \lambda\sigma^2_t + \varepsilon_{t+1}$, where λ is the persistence parameter; and

Forecast 3: Naive historical average.

The AR(1) process is estimated using 60 periods of data with non-negative constraints on the parameters. Other lengths of estimation window (36 and 100) and other AR processes were considered, with similar results. All forecasting models considered are identical except that they use data of different frequency. The statistics to evaluate the volatility forecasts are:

$$\text{HMAE: } E\{\,|\sigma^2_{true} - \sigma^2_{forecast}\,|/\sigma^2_{true}\} \quad (4)$$

$$\text{HRMSE: } E\{((\sigma^2_{true} - \sigma^2_{forecast})/\sigma^2_{true})^2\}^{\frac{1}{2}} \quad (5)$$

We also use a Mincer-Zarnowitz regression:

$$\sigma^2_{true,t} = b_0 + b_1\sigma^2_{forecast,t} + e_t$$

for further analysis; unbiased forecasts imply $b_0 = 0$ and $b_1 = 1$. We have three options to define the true volatility:

$$\sigma^2_{true,d} = \sum_{\tau\in t} r^2_{\tau,d} \tag{6}$$

$$\sigma^2_{true,m} = r^2_{t,m} \tag{7}$$

and

$$\sigma^2_{true,d} = \sum_{i\in t} r^2_{i,d} + \sum_{j=1}^{n} r_{i,d} r_{i-j,d} \tag{8}$$

That is, we use daily returns but including n cross-products in order to capture autocorrelation in daily returns.

Despite our view that autocorrelation is largely driven by microstructure noise, we still find it instructive to apply (8) with $n = 5$ as the true volatility measure. First, our measure of monthly volatility used during the estimation stage is different from the measure used for out-of-sample evaluation. Second, the typical autocorrelation function rarely has significant entries beyond lag 5. Therefore, the out-of sample risk measure falls between the daily measure and the monthly measure, and is closer to the monthly measure.

We attempt to proxy for the true volatility with squared monthly returns, but the R^2 from the Mincer-Zarnowitz regression is very low across all models, showing it to be a very noisy proxy of monthly volatility. Further, the monthly measure routinely comes close to zero, which makes HMAE and HRMSE statistics meaningless.

Finally, comparison of the two measures (6) and (7) with the widely accepted VIX values from January 1986 through December 2002 is extremely unfavorable to the monthly measure. The monthly measure clearly presents a noisy volatility proxy; the HMAE statistic for the monthly measure (with VIX as a benchmark) is more than twice that of the daily measure.[4]

Exhibit 3 reports HMAE and HRMSE test statistics for forecasting the monthly volatility of the S&P 500 together with the coefficients and R^2 from the Mincer-Zarnowitz regression. Several findings are apparent. First, using daily data in conjunction with a simple AR(1) model produces vastly superior forecasts. The HMAE is the best of all the models. Its R^2 is only slightly lower than the R^2 of the daily random walk model, and the null hypothesis of $b_0 = 0$ and $b_1 = 1$ could not be rejected.

Second, the random walk model of daily returns does well when no extreme events take place (as HMAE shows), but performs more poorly when volatility changes rapidly (HRMSE and regression results). Third, the AR(1) fitted to monthly data performs much more poorly than when fitted to daily data. Fourth, models that use historical volatility perform poorly by any measure.

Finally, in many cases the difference between corresponding daily and monthly models is much greater than between models that use the same-frequency data. For instance, the difference in R^2 between AR(1) models and random walk models of the same frequency is around 5% for SMB forecasts, but the difference between the same model that uses different frequencies is more than 10% (0.38 versus 0.23 for the AR(1) model and 0.34 versus 0.23 for the random walk model).

These results illustrate a major advantage of using higher-frequency returns, namely, that we can apply simple and robust forecasting models without sacrificing forecasting performance.

EXHIBIT 3

Forecasting Monthly Variance of Factor Portfolios

	HMAE	HRMSE	b_0	b_1	R^2
Daily_AR(1)	0.7890	1.5131	**0.0003**	**1.0405**	0.1361
Daily_Hist	1.2839	2.4118	0.0011	0.3278	0.0194
Daily_RW	1.0689	2.2104	0.0011	0.3857	0.1420
Monthly_AR(1)	0.8087	1.4864	0.0008	**0.7716**	0.0477
Monthly_Hist	1.3328	2.4662	0.0012	0.2985	0.0137
Monthly_RW	1.6641	3.9303	0.0016	0.1022	0.0270

Boldface indicates null hypothesis could not be rejected.

Other Risk Model Building Blocks

While using daily returns results in better forecasts of single-asset volatility, the majority of portfolio applications require forecasting a portfolio covariance matrix. Such forecasting requires estimates of the factor covariance matrix, the exposures or loadings of individual stocks for the factors, and the stock-specific risk.

Factor Correlation. We estimate a so-called constant correlation model advocated in Bollerslev [1990], which forecasts the covariance matrix assuming multivariate con-

EXHIBIT 4
Forecasting Factor Correlation

	SP 500	SMB	HML
SP 500	--	0.2077	0.2193
		0.2779	0.6243
SMB	0.1888	--	0.2239
	0.2891		0.3918
HML	0.2129	0.2151	--
	0.5952	0.4557	

Comparing correlation estimated from 12 (36) months of past data to realized correlation over the next two months. Realized correlation is measured with daily data; one cross-product is included to reflect possible non-synchronicity. The first number is the MAE when in-sample correlation is measured with daily returns; the number below is the MAE when in-sample correlation is measured with monthly returns. The upper triangle corresponds to an estimation period of 36 months, while the lower triangle corresponds to an estimation period of 12 months.

ditional heteroscedasticity. Exhibit 4 reports the mean of the absolute difference (MAE) between forecasted and realized correlation among the three factors considered: S&P 500, SMB, and HML.[6] The correlations computed with the last three years of returns are a reasonable approximation of the correlations over the next two months. Moreover, with daily returns, shortening the estimation period to one year results in a slight improvement in the correlation forecast.

Optimal Length of Estimation Window. In selecting the number of observations for each estimation window, there is a trade-off between the statistical estimation accuracy and the relevance of the historical information. Many monthly risk models that use monthly data use five years of past returns in order to estimate factor loadings, raising a question as to whether the information is relevant for current market conditions. We run simple tests that vary the lengths of estimation and holding periods and measure the deviation of estimated exposure from the realized exposure.

The sample period extends from January 1973 through December 2002 with daily returns on Dow Jones stocks and on the factors. The estimated exposures from the last T_1 observations at the beginning of each month are compared to the realized exposures within the next T_2 months.

We use MSE as the statistical measure of deviation, but given that there are three factors in the model, we also compute the weighted MSE across three exposures. The weights are proportional to the sample variance of each factor in order to reflect its importance. Weighted least squares regression is used to estimate historical and realized exposures in order to eliminate the influence of influential observations.

Daily betas for small stocks may be biased due to non-trading, as noted by Scholes and Williams [1977]. While single-factor beta adjustment is a fairly standard procedure, adjusting loadings in a multivariate setting is a bit more complicated. One approach is to assume that factors are orthogonal and adjust each factor in turn (see Jagannathan and Ma [2003]). We use a slightly more complicated adjustment that does not require an assumption of factor orthogonality.

The results are reported in Exhibit 5. We vary T_1 from two months to three and a half years and T_2 from one month to six years. Setting T_2 equal to one month reflects the risk model updating frequency, while the other two values of T_2 match the investment cycles of monthly risk model users. The minimum values of weighted MSE across each row indicate the optimal length of the estimation period, given the length of the holding period.

Specifically, the optimal estimation periods are 375 days (one and a half years), 625 days (two and a half years), and 750 days (three years) for holding periods of one, three, and six months, respectively. These findings are consistent with the mean-reversion features of factor exposures documented by Vasicek [1973]. In other words, exposures for longer holding periods converge to their

EXHIBIT 5
Window Size and Forecasting Errors of Exposures

T1 / T2	60 days	125 days	250 days	375 days	500 days	625 days	750 days	875 days
1 month	1.4987	1.3770	1.2967	1.2924	1.2954	1.2970	1.2986	1.2979
3 month	0.6074	0.4449	0.3891	0.3815	0.3807	0.3800	0.3800	0.3824
6 month	0.4655	0.3006	0.2408	0.2325	0.2292	0.2270	0.2262	0.2289

EXHIBIT 6
Forecasting Monthly Covariance Matrix

	Factor model daily returns	Factor model monthly returns	t-statistic in absolute value for equality	Sample covar daily returns	Sample covar monthly returns	t-statistic in absolute value for equality
Test 1	0.6100	0.6376	10.38**	0.6469	0.6888	8.86**
Test 2	1.0022	1.1822	27.62**	1.0452	1.6261	43.12**

***Significant at the 5% level.*

long-run mean, so using longer periods to estimate them proves beneficial.

Forecasting Portfolio Covariance

Beginning in December 1972 and ending in December 2002, we randomly pick 100 stocks from the eligible universe for each month. The stock must be a domestic common issue with price above $5, with no missing returns during the estimation period, and with market capitalization available at the time of the portfolio formation. Second, we form market capitalization quintiles of stocks and then pick 100 random eligible stocks from each quintile. We compute the realized covariance matrix as a product of sample variance, measured by the sum of the squared daily returns during one month immediately following the estimation date V_t and sample correlation R_t as $\Sigma_t = V_t^{1/2} R_t V_t^{1/2}$.

We use several test statistics to judge the model's performance. First, we measure model precision in predicting variance only:

$$\frac{\sum |diag(F) - diag(\Sigma)|}{\sum diag(\Sigma)}$$

where F is forecasted covariance matrix. Similarly, we measure model precision in forecasting covariance only, with:

$$\frac{\sum |tril(F) - tril(\Sigma)|}{\sum |tril(\Sigma)|}$$

as a test statistic, where *tril (X)* is the lower triangular matrix of X.

Results for random portfolios are presented in Exhibit 6. They show that using daily data (whether in conjunction with a factor risk model or not) produces a comfortable advantage over using monthly data in forecasting variance and covariance. We also find that using higher-frequency data is better than using a more complicated model. Overall, a factor model using daily data dominates all other models in forecasting covariances.

PORTFOLIO OPTIMIZATION

Portfolio optimization is used for a variety of tasks, including hedging risk, maximizing the information ratio, or tracking a benchmark. We are interested in whether documented improvements in estimating the covariance matrix translate to superior portfolios. We focus our attention on constructing global minimum-variance portfolios (GMV). Every month, starting in December 1972 and ending in December 2001, we run optimization routines on a group of stocks with the aim of minimizing portfolio variance over the ensuing holding period.

Optimization results for portfolios of 100 random stocks are summarized in Exhibit 7. As before, we measure monthly volatility as the sum of squared daily returns with correction for autocorrelation (up to lag 5). We assume a one-month holding period, consistent with the nature of the monthly risk model.

The first row gives standard deviations of optimized portfolios across all time periods. The rest of Exhibit 7 reports standard deviations of optimized portfolios under four different volatility regimes: lowest, low, high, and highest. Month t is considered the lowest volatility month if the standard deviation of the equally weighted portfolio during this month falls into the bottom 25% of the time series distribution of monthly volatility. A low volatility regime corresponds to an equally weighted portfolio falling between 25% and 50% of the time series distribution of monthly volatility; high volatility is 50%-75%, and the highest volatility exceeds 75% of its time series distribution.

Note first that the daily factor model has an edge over the monthly factor model (9.51% versus 10.18%

EXHIBIT 7
Portfolio Optimization—Time Series Averages of Standard Deviations

Volatility	Factor model daily returns (1)	Factor model monthly returns (2)	Sample cov daily returns (3)	Sample cov monthly returns (4)	EW portfolio (5)	(1) – (2)	(1) – (3)
100_random	0.0951	0.1018	0.0964	0.1062	0.1448	-0.0067	-0.0013
			By volatility regime				
Lowest	0.0522	0.0628	0.0521	0.0643	0.0569	-0.0107	0.0001
Low	0.0697	0.0821	0.0699	0.0848	0.0940	-0.0124	-0.0002
High	0.1098	0.1135	0.1115	0.1192	0.1785	-0.0037	-0.0018
Highest	0.1590	0.1660	0.1623	0.1738	0.2689	-0.0070	-0.0033

EXHIBIT 8
Tracking Error Portfolio Optimization

Volatility	Factor model daily returns (1)	Factor model monthly returns (2)	Sample cov daily returns (3)	Sample cov monthly returns (4)	EW portfolio (5)	(1) (2)	(1) (3)
100_random	0.0630	0.0664	0.0642	0.0761	0.0812	-0.0034	-0.0012
			By volatility regime				
Lowest	0.0409	0.0456	0.0417	0.0552	0.0402	-0.0047	-0.0008
Low	0.0550	0.0564	0.0562	0.0659	0.0647	-0.0014	-0.0012
High	0.0674	0.0682	0.0688	0.0771	0.0853	-0.0008	-0.0014
Highest	0.0888	0.0908	0.0901	0.0998	0.1347	-0.0020	-0.0013

annualized standard deviation). Second, the sample covariances matrix computed using daily data beats the monthly factor model (9.64% versus 10.18% annualized standard deviation). Third, the daily model beats the monthly model in all volatility regimes.

Interestingly, when volatility is very low, both monthly models fail to beat an equally weighted portfolio. The factor model's advantage over the sample covariance matrix increases as overall volatility increases.

Examining the standard deviation of optimized portfolios formed within each market cap quintile, we find the daily factor model performs better than the monthly factor model for all size quintiles (from 0.27% annualized standard deviation for the smallest quintile to 0.9% for the middle quintile). The daily factor model has a much smaller advantage over the daily model that uses the sample covariance matrix (ranging from zero to 0.14%). This reflects our finding that transition to higher-frequency data has more effect on forecasting covariances than using a more complicated model.

Risk models are often used to control tracking error. Exhibit 8 reports the results of a tracking error optimization for randomly chosen long-short portfolios. For each month starting in December 1972 and ending in December 2001, we randomly pick 100 stocks satisfying the eligibility criteria as described. Next, we pick 50 stocks from the preselected 100 and assign negative weights proportional to market capitalization and normalized to sum to negative one. These 50 stocks constitute the short side of the portfolio.

Then, we run portfolio optimizations to select positive weights for the remaining 50 preselected stocks. As before, an upper limit of 5% on the long side is imposed. The best model will yield the portfolio with the lowest volatility of the tracking error.

Exhibit 8 is consistent with our earlier results; using daily returns makes a superior tracking portfolio.

CONCLUSION

Volatility forecasts are crucial to many financial applications, and traditionally are constructed so that the sampling interval matches the forecast horizon. This practice has drawbacks. Using low-frequency data to create low-frequency risk forecasts induces staleness and might not accurately reflect return events within the sampling interval. So why should data frequency match the forecast horizon?

We demonstrate that using higher-frequency data can improve risk forecasts. The use of daily returns results in more precise estimates of monthly volatility. We also show that the difference between monthly volatility estimates and our approach can be bridged by factoring in the autocorrelation in daily returns (except this approach is subject to error arising from spurious estimates of autocorrelation). We maintain that the use of daily returns offers greater precision in practice.

The main advantage of higher-frequency returns for construction of volatility measures over a longer horizon is that it allows the use of simple and robust volatility forecasting models. Daily returns also have significant advantages for estimating exposures and factor correlations from a monthly risk model.

Whether these results can be extended to other areas of concern for portfolio managers and traders is an open question.

ENDNOTES

[1]Higher-frequency (daily or weekly) risk models are increasingly used in trading applications or by high-turnover managers and hedge funds as discussed in Madhavan and Yang [2003].

[2]Anderson and Bollerslev [1998] and Anderson, Bollerslev, and Lange [1999] explore the predictability of exchange rate return volatility using intraday returns. Blair, Poon, and Taylor [2001] examine the accuracy of volatility forecasts for a stock index.

[3]Data on S&P 500 returns are from CRSP. Data on all other portfolios are from the French data library. The Fama/French factors are constructed using the six value-weight portfolios formed on size and book-to-market. The universe includes all NYSE, AMEX, and Nasdaq firms. SMB and HML for July of year t through June of $t + 1$ include all NYSE, AMEX, and Nasdaq stocks for which market equity data for December of $t - 1$ and June of t, and (positive) book equity data for $t - 1$, are available. SMB (small minus big) is the average return on the three small portfolios minus the average return on the three big portfolios: *SMB* = 1/3 *(Small Value + Small Neutral + Small Growth)* – 1/3 *(Big Value + Big Neutral + Big Growth)*. HML (high minus low) is the average return on the two value portfolios minus the average return on the two growth portfolios: *HML* = 1/2 *(Small Value + Big Value)* – 1/2 *(Small Growth + Big Growth)*.

The Fama/French benchmark portfolios are rebalanced quarterly using two independent sorts, on size (market equity, ME) and book-to-market (the ratio of book equity to market equity, BE/ME). The size breakpoint (which determines the buy range for the small and big portfolios) is the median NYSE market equity. The BE/ME breakpoints (which determine the buy range for the growth, neutral, and value portfolios) are the 30th and 70th NYSE percentiles.

Data on individual stocks are from Datastream for 1973–2002 (360 months). DJIA stocks that did not stay in the index for the entire period are excluded (Citigroup, Home Depot, Honeywell, Microsoft, SBC Communications). Applying the same filter to Russell 2000 stocks results in 255 available for the test.

T-statistics are for a one-sided t-test of the monthly measure being lower than the daily measure of volatility. Bootstrapping proceeds in 20-day segments to preserve the return structure over a single month. The number of samples drawn depends on the amount of actual data. In each case the drawing procedure is repeated 100 times.

[4]See Blair, Poon, and Taylor [2001] for standard adjustment of the VIX.

REFERENCES

Anderson, Torben, and Tim Bollerslev. "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts." *International Economic Review*, Vol. 39 (1998), pp. 885-905.

Anderson, Torben, Tim Bollerslev, and Steve Lange. "Forecasting Financial Market Volatility: Sample Frequency vis-à-vis Forecast Horizon." *Journal of Empirical Finance*, Vol. 6 (1999), pp. 457-477.

Blair, Bevan, Ser-Huang Poon, and Stephen Taylor. "Forecasting S&P 100 Volatility: The Incremental Information Content of Implied Volatilities and High-Frequency Index Returns." *Journal of Econometrics*, 105(1) (2001), pp. 5-26.

Bollerslev, Tim. "Modeling the Coherence in Short-run Nominal Exchange Rates: A Multivariate Generalized ARCH Model." *Review of Economics and Statistics*, 72(3) (August 1990), pp. 498-505.

Chan, Louis, Jason Karceski, and Josef Lakonishok. "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model." *The Review of Financial Studies*, Vol. 12 (1999), pp. 937-974.

Fama, Eugene, and Kenneth French. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, Vol. 33 (1993), pp. 3-56.

French, Kenneth, and Richard Roll. "Stock Return Variances: The Arrival of Information and the Reaction of Traders." *Journal of Financial Economics*, Vol. 17 (1986), pp. 5-26.

French, Kenneth, William Schwert, and Robert Stambaugh. "Expected Stock Returns and Volatility." *Journal of Financial Economics*, Vol. 19 (1987), pp. 3-29.

Jagannathan, Ravi, and Tongshu Ma. "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps." *Journal of the American Finance Association*, Vol. 58 (2003).

Jegadeesh, Narasimhan, and Sheridan Titman. "Short-Horizon Return Reversals and the Bid-Ask Spread." *Journal of Financial Intermediation*, Vol. 4 (1995), pp. 116-132.

Lo, Andrew, and Craig MacKinlay. "An Econometric Analysis of Nonsynchronous Trading." *Journal of Econometrics*, Vol. 45 (1990), pp. 181-211.

Madhavan, Ananth, Matthew Richardson, and Mark Roomans. "Why Do Security Prices Fluctuate? A Transaction-Level Analysis of NYSE Stocks." *The Review of Financial Studies*, Vol. 10 (1997), pp. 1035-1064.

Madhavan, Ananth, and Jian Yang. "Practical Risk Analysis for Portfolio Managers and Traders." *The Journal of Portfolio Management*, Vol. 30 (2003), pp. 73-85.

Merton, Robert. "On Estimating the Expected Return on the Market: An Exploratory Investigation." *Journal of Financial Economics*, Vol. 8(4) (1980), pp. 323-361.

Scholes, Myron, and Joseph Williams. "Estimating Betas from Nonsynchronous Data." *Journal of Financial Economics*, 5 (1977), pp. 309-327.

Schwert, William. "Stock Market Volatility: Ten Years After the Crash." *Brookings-Wharton Papers on Financial Services*, Vol. 1 (1998), pp. 65-114.

——. "Stock Volatility in the New Millennium: How Wacky is Nasdaq?" *Journal of Monetary Economics*, Vol. 49 (2002), pp. 3-26.

Vasicek, Oldrich Alfons. "A Note on Using Cross-Sectional Information in Bayesian Estimation of Security Betas." *Journal of Finance*, Vol. 8, No. 5 (1973), pp. 1233-1239.