

The Economic Value of Volatility Timing

JEFF FLEMING, CHRIS KIRBY, and BARBARA OSTDIEK*

ABSTRACT

Numerous studies report that standard volatility models have low explanatory power, leading some researchers to question whether these models have economic value. We examine this question by using conditional mean-variance analysis to assess the value of volatility timing to short-horizon investors. We find that the volatility timing strategies outperform the unconditionally efficient static portfolios that have the same target expected return and volatility. This finding is robust to estimation risk and transaction costs.

VOLATILITY PLAYS A CENTRAL ROLE in derivatives pricing, optimal portfolio selection, and risk management. These applications motivate an extensive literature on volatility modeling. Starting with Engle (1982), researchers have fit a variety of autoregressive conditional heteroskedasticity (ARCH), generalized ARCH (Bollerslev (1986)), exponential ARCH (Nelson (1991)), and stochastic volatility models to asset returns. This literature, however, has centered on evaluating the statistical performance of volatility models rather than the economic significance of time-varying, predictable volatility. In contrast, we focus on the latter. Specifically, we examine the economic value of volatility timing to risk-averse investors.

Several review articles summarize the empirical findings on volatility (see, e.g., Bollerslev, Chou, and Kroner (1992), Bollerslev, Engle, and Nelson (1994), Diebold and Lopez (1995), and Palm (1996)). The evidence is generally consistent across a broad range of assets and econometric specifications, and overwhelmingly suggests that volatility is to some extent predictable. However, standard volatility models typically explain only a small fraction of the variation in squared returns. This has led some researchers to question the relevance of these models. Andersen and Bollerslev (1998) argue that the

* Fleming and Ostdiek are at Rice University and Kirby is at the Australian Graduate School of Management. This paper was completed while Kirby was at the University of Texas at Dallas. We have benefitted from many helpful comments by René Stulz (the editor), an anonymous referee, Amir Barnea, Jay Coughenour, Rob Engle, Wayne Ferson, John Graham, David Hsieh, David Ikenberry, Ravi Jagannathan, Andrew Karolyi, Uri Loewenstein, Jack Mosevich, Ehud Ronn, Paul Seguin, Guofu Zhou, and seminar participants at the University of Minnesota, University of Utah, Texas A&M University, the 1998 Chicago Board of Trade Spring Research Seminar, the 1998 Conference on Financial Economics and Accounting, the 1999 Computational Finance Conference, the 1999 Computational and Quantitative Finance Conference, the 1999 AGSM Accounting and Finance Research Camp, and the 2000 American Finance Association meetings.

low explanatory power is an inevitable consequence of the noise inherent in the return-generating process. They propose a more precise measure of ex post volatility (cumulative squared intradaily returns) and find that GARCH models explain about 50 percent of the variation in this measure. This suggests that standard volatility models deliver reasonably accurate forecasts, but it leaves unanswered the question of whether volatility timing has economic value.

There has been little research that specifically considers this issue. A few studies, such as Graham and Harvey (1996) and Copeland and Copeland (1999), examine trading rules designed to exploit predictable changes in volatility. But these studies typically limit their analysis to simple switching strategies. Busse (1999), on the other hand, examines the trading behavior of active portfolio managers. He finds that a significant percentage of mutual fund managers tend to reduce their market exposure during periods of high expected volatility. Although this suggests that many fund managers behave like volatility timers, their trading decisions may be driven by factors other than volatility modeling.

In this paper, we systematically examine the value of volatility timing for short-horizon asset-allocation strategies. The framework for our analysis is straightforward. We consider an investor who uses a mean-variance optimization rule to allocate funds across four asset classes: stocks, bonds, gold, and cash. The investor's objective is to maximize expected return (or minimize volatility) while matching a target volatility (or expected return). Allowing for daily rebalancing, the solution to the investor's portfolio problem is a dynamic trading strategy that specifies the optimal asset weights as a function of time. Implementing this strategy, in general, requires estimates of both the conditional expected returns and the conditional covariance matrix. The variances and covariances, however, can typically be estimated with far greater precision than the expected returns (Merton (1980)). Therefore, we treat expected returns as constant and let the variation in the portfolio weights be driven purely by changes in the conditional covariance matrix.

To estimate the conditional covariance matrix, we employ a general non-parametric approach developed by Foster and Nelson (1996). The estimator is a weighted rolling average of the squares and cross products of past return innovations that nests most ARCH, GARCH, and stochastic volatility models as special cases. We determine the weights by minimizing the asymptotic mean squared error (MSE) of the estimator. After constructing the covariance matrix estimates, we form the dynamic portfolios and evaluate their performance. Our measure of the value of volatility timing is the estimated fee that a risk-averse investor would be willing to pay to switch from the ex ante optimal static portfolio to the dynamic portfolio.

The data for our analysis consist of daily returns for stock, bond, and gold futures. We use futures to avoid short sale constraints and microstructure effects, but our analysis generalizes to the underlying spot assets via standard no-arbitrage arguments. Our results indicate that when estimation risk regarding expected returns is negligible, the volatility timing strategies plot

above the ex post efficient frontier for fixed-weight portfolios. We obtain similar results after using a bootstrap procedure to control for estimation risk. For example, when the risk is comparable to that over our sample, the maximum return strategy has a higher Sharpe ratio than the ex ante optimal static portfolio in 92 percent of our simulations. Moreover, the effectiveness of volatility timing increases if we use smoother volatility estimates than those obtained using our MSE criterion. When the volatility persistence is comparable to that implied by GARCH models, a risk-averse investor would pay a positive fee to switch from the optimal static to the dynamic portfolio in nearly 100 percent of the simulations. On average, the estimated fee exceeds 170 basis points per year.

The remainder of the paper is organized as follows. Section I develops our methodology for measuring the value of volatility timing. Section II describes the data used in our analysis. Section III reports the empirical results. Section IV summarizes our conclusions and outlines the implications for future research.

I. Methodology

Our methodology for measuring the value of volatility timing is to evaluate the impact of predictable changes in volatility on the performance of short-horizon asset-allocation strategies.¹ Both theoretical and empirical considerations motivate our focus on short-horizon strategies. Many theoretical models of the trading process imply that daily returns are characterized by stochastic volatility (see, e.g., Tauchen and Pitts (1983) and Andersen (1996)). And, empirically, the persistence in volatility is stronger for daily returns than for returns measured over longer horizons (see, e.g., Glosten, Jagannathan, and Runkle (1993)).

We use mean-variance analysis to implement the asset-allocation strategies. Consequently, they are optimal only if investors have logarithmic utility and the first two moments completely characterize the joint distribution of returns. This is not problematic, however, given the nature of our investigation. Specifically, if volatility timing has value using a suboptimal strategy, then more sophisticated strategies are likely to yield even greater value. More importantly, the mean-variance approach facilitates several aspects of our analysis. First, it underlies most of the common measures of portfolio performance. Second, the relation between mean-variance optimization and quadratic utility allows us to quantify how risk aversion affects the value of volatility timing. Finally, this framework accommodates a straightforward simulation approach to assess the significance and robustness of our results.

¹ This approach is similar in spirit to Kandel and Stambaugh (1996). They use a Bayesian framework to study how stock return predictability influences the asset-allocation decisions of an investor with power utility. Our analysis also focuses on the economic significance of predictable variation in the inputs to a portfolio problem, but we use a less complex approach that is more consistent with traditional performance measurement methods.

A. Volatility Timing in a Mean-variance Framework

Consider an investor with a one-day horizon who wants to minimize portfolio variance subject to achieving a particular expected return. In general, constructing the portfolio weights requires one-step-ahead estimates of both the vector of conditional means and the conditional covariance matrix. There is little empirical evidence, however, that we can detect variation in expected returns at the daily level. Moreover, Merton (1980) shows that a very long sample period would be needed to produce reliable coefficient estimates in a predictive regression. We assume, therefore, that our investor models expected returns as constant. Given our mean-variance framework, this is equivalent to following a volatility-timing strategy.

Because the portfolio weights in this strategy ignore any time variation in expected returns, our methodology for measuring the value of volatility timing should yield conservative results. To see why, note that theory implies a positive relation between expected returns and volatility. Ignoring this linkage causes our portfolio weights to decrease by more than is optimal when volatility rises and to increase by more than is optimal when volatility falls. This reduces the potential effectiveness of our volatility-timing strategy, but it also mitigates the concern that our results are actually driven by variation in expected returns.

To develop our methodology, let \mathbf{R}_{t+1} , $\boldsymbol{\mu} \equiv E[\mathbf{R}_{t+1}]$, and $\boldsymbol{\Sigma}_t \equiv E_t[(\mathbf{R}_{t+1} - \boldsymbol{\mu})(\mathbf{R}_{t+1} - \boldsymbol{\mu})']$ denote, respectively, an $N \times 1$ vector of risky asset returns, the expected value of \mathbf{R}_{t+1} , and the conditional covariance matrix of \mathbf{R}_{t+1} . For each date t , the investor solves the quadratic program

$$\begin{aligned} \min_{\mathbf{w}_t} \quad & \mathbf{w}_t' \boldsymbol{\Sigma}_t \mathbf{w}_t \\ \text{s.t.} \quad & \mathbf{w}_t' \boldsymbol{\mu} + (1 - \mathbf{w}_t' \mathbf{1}) R_f = \mu_p, \end{aligned} \tag{1}$$

where \mathbf{w}_t is an $N \times 1$ vector of portfolio weights on the risky assets, R_f is the return on the riskless asset, and μ_p is the target expected rate of return. The solution to this optimization problem,

$$\mathbf{w}_t = \frac{(\mu_p - R_f) \boldsymbol{\Sigma}_t^{-1} (\boldsymbol{\mu} - R_f \mathbf{1})}{(\boldsymbol{\mu} - R_f \mathbf{1})' \boldsymbol{\Sigma}_t^{-1} (\boldsymbol{\mu} - R_f \mathbf{1})}, \tag{2}$$

delivers the risky asset weights. The weight on the riskless asset is $1 - \mathbf{w}_t' \mathbf{1}$.

We can express equation (2) in terms of futures returns by applying standard no-arbitrage arguments. Under the cost-of-carry model, the return on a futures contract equals the total return on the underlying asset minus the riskless interest rate (because futures entail no initial investment). Subtract-

ing the riskless rate from each element of \mathbf{R}_{t+1} has no effect on the conditional covariance matrix, so we can use the cost-of-carry relation to express equation (2) as

$$\mathbf{w}_t = \frac{\mu_p \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\mu}}{\boldsymbol{\mu}' \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\mu}}, \quad (3)$$

with the vector $\boldsymbol{\mu} \equiv E[\mathbf{r}_{t+1}]$ and matrix $\boldsymbol{\Sigma}_t \equiv E_t[(\mathbf{r}_{t+1} - \boldsymbol{\mu})(\mathbf{r}_{t+1} - \boldsymbol{\mu})']$ redefined in terms of excess returns.

The trading strategy implicit in equation (3) identifies the dynamically rebalanced portfolio that has minimum conditional variance for any choice of expected return. We could conduct a similar analysis where the objective is to maximize the expected return subject to achieving a particular conditional variance. Thus, our mean-variance framework suggests two candidate volatility-timing strategies. First, we set the portfolio expected return equal to a fixed target and solve for the weights that minimize conditional variance (the minimum volatility strategy). Second, we set the portfolio variance equal to a fixed target and solve for the weights that maximize conditional expected return (the maximum return strategy).

B. Estimating the Conditional Covariance Matrix

To implement the volatility-timing strategies, we need to form one-step-ahead estimates of the conditional covariance matrix. A number of estimation methods have been developed in the literature.² We follow Foster and Nelson (1996) and use rolling estimators that are constructed in an asymptotically optimal manner. This approach has some distinct advantages in our application. Unlike multivariate ARCH and GARCH models, which are heavily parameterized and difficult to estimate, the computational demands of rolling estimators are modest. In addition, the nonparametric nature of the approach is consistent with our objective of providing baseline evidence—without searching for the best volatility model—on the economic significance of time-varying, predictable volatility.

The class of rolling estimators that we employ can be written as

$$\hat{\sigma}_{ij,t} = \sum_{l=-t+1}^{T-t} \omega_{ij,t+l} (r_{i,t+l} - \mu_i)(r_{j,t+l} - \mu_j), \quad (4)$$

² Officer (1973) and Fama and MacBeth (1973) employ ad hoc rolling estimators. Merton (1980) and French, Schwert, and Stambaugh (1987) divide the data into nonoverlapping blocks and treat the conditional variances and covariances as constant within each block. More recently, ARCH models (e.g., Engle (1982) and Bollerslev (1986)) have gained popularity. Our approach nests a broad range of ARCH and GARCH models as special cases.

where r_{it} and r_{jt} denote the returns on assets i and j , respectively, $\omega_{ij,t+l}$ is the weight placed on the product of the return innovations for date $t+l$, and T is the number of observations in the sample. To distinguish between variance estimators and covariance estimators, we will use the notation $\hat{\sigma}_{it}^2$ for the case where $i = j$. Although equation (4) admits a wide range of potential weighting schemes, Foster and Nelson (1996) demonstrate that the optimal strategy is to let the weights decline in an exponential fashion as the magnitude of l increases.

The precise form of the optimal weights depends on the characteristics of the volatility process. If volatility is stochastic, then the optimal weights for the two-sided rolling estimator are given by

$$\omega_{ij,t+l} = (\alpha_{ij,t}/2)e^{-\alpha_{ij,t}|l|}, \quad (5)$$

where $\alpha_{ij,t}$ is the decay rate. This estimator uses both leads and lags of returns to estimate $\sigma_{ij,t}$. To construct the corresponding one-sided estimator, we set $\omega_{ij,t+l} = 0$ for $l > 0$ and double each of the weights for $l \leq 0$.

Applying this methodology requires an estimate of the optimal decay rate. Foster and Nelson (1996) show how to estimate the $\alpha_{ij,t}$ that minimizes the asymptotic MSE of the estimator in equation (4). Their procedure, however, implies a different decay rate for each element of the conditional covariance matrix. Because this makes it difficult to ensure that the matrix is positive definite, we impose the restriction $\alpha_{ij,t} = \alpha_t$ for all i and j . Given this restriction, we can show that under empirically plausible assumptions the optimal decay rate is constant. We estimate this decay rate by minimizing the asymptotic MSE of our rolling estimator of $\sigma_{pt}^2 = \mathbf{w}'_t \boldsymbol{\Sigma}_t \mathbf{w}_t$. Using the data for our sample yields an estimate of 0.063.³

C. Measuring the Value of Volatility Timing

To measure the value of volatility timing, we compare the performance of the dynamic strategies to that of the unconditional mean-variance efficient static strategies that have the same target expected return and volatility. If volatility timing has no value, then the ex post performance of the static and dynamic strategies should be statistically indistinguishable. Making this comparison requires a performance measure that captures the trade-off between risk and return. We use a generalization of West, Edison, and Cho's (1993) criterion for ranking the performance of forecasting models. This measure is based on the close relation between mean-variance analysis and quadratic utility.

³ An appendix is available on the journal's web site (www.afajof.org) that describes our estimation procedure in detail. Note that the procedure relies on the actual data, so we potentially introduce a look-ahead bias into our results. However, the sensitivity analysis in Section III.D suggests this is not a significant concern. In particular, we find that volatility timing is more effective using a smaller decay rate than implied by the minimum MSE criterion. A smaller decay rate generates smoother volatility estimates than the ones we use and is consistent with the findings in the GARCH literature.

In general, we can view quadratic utility as a second-order approximation to the investor's true utility function. Under this approximation, the investor's realized utility in period $t + 1$ can be written as

$$U(W_{t+1}) = W_t R_{p,t+1} - \frac{aW_t^2}{2} R_{p,t+1}^2, \tag{6}$$

where W_{t+1} is the investor's wealth at $t + 1$, a is his absolute risk aversion, and

$$R_{p,t+1} = R_f + \mathbf{w}'_t \mathbf{r}_{t+1}$$

is the period $t + 1$ return on his portfolio. To facilitate comparisons across portfolios, we hold aW_t constant. This is equivalent to setting the investor's relative risk aversion, $\gamma_t = aW_t/(1 - aW_t)$, equal to some fixed value γ . With relative risk aversion held constant, we can use the average realized utility, $\bar{U}(\cdot)$, to consistently estimate the expected utility generated by a given level of initial wealth. In particular, we have

$$\bar{U}(\cdot) = W_0 \left(\sum_{t=0}^{T-1} R_{p,t+1} - \frac{\gamma}{2(1 + \gamma)} R_{p,t+1}^2 \right), \tag{7}$$

where W_0 is the investor's initial wealth.

We estimate the value of volatility timing by equating the average utilities for two alternative portfolios. Suppose, for example, that holding a static portfolio yields the same average utility as holding a dynamic portfolio that is subject to daily expenses of Δ , expressed as a fraction of wealth invested. Because the investor would be indifferent between these two alternatives, we interpret Δ as the maximum performance fee that he would be willing to pay to switch from the static to the dynamic strategy. To estimate this fee, we find the value of Δ that satisfies

$$\sum_{t=0}^{T-1} (R_{d,t+1} - \Delta) - \frac{\gamma}{2(1 + \gamma)} (R_{d,t+1} - \Delta)^2 = \sum_{t=0}^{T-1} R_{s,t+1} - \frac{\gamma}{2(1 + \gamma)} R_{s,t+1}^2, \tag{8}$$

where $R_{d,t+1}$ and $R_{s,t+1}$ denote the returns for the dynamic and static strategies, respectively. We report our estimates of Δ as annualized fees in basis points using two different values of γ , 1 and 10.⁴

⁴ Because utility depends on total returns rather than excess returns, we also have to specify a value for the riskless rate. We set $R_f = 6$ percent in the empirical analysis.

II. Data and Preliminary Analysis

Our empirical analysis focuses on four broadly defined asset classes: stocks, bonds, gold, and cash equivalents. As explained earlier, we use futures contracts for the analysis. The specific contracts are the S&P 500 index futures traded at the Chicago Mercantile Exchange, the Treasury bond futures traded at the Chicago Board of Trade, and the gold futures traded at the New York Mercantile Exchange. Because futures returns are approximately equivalent to excess spot returns, we can eliminate cash equivalents from explicit consideration. The weight placed in cash is implicit in the solution to the portfolio optimization problem using only the risky assets.

The source for the gold futures data is Datastream International and the source for the bond and stock futures data is the Futures Industry Institute. The gold futures contract closes at 1:30 CST each day whereas the bond and stock contracts close at 2:00 CST and 3:15 CST, respectively. We align the price observations across contracts by using daily closing prices for gold futures and the last transaction prices before 1:30 CST for the bond and stock contracts. In addition, we exclude all days when any of the three markets is closed in order to maintain a uniform measurement interval across contracts. The sample period is January 3, 1983 to December 31, 1997.

A. The Returns

We compute the daily returns using the day-to-day price relatives for the nearest-to-maturity contract. As the nearby contract approaches maturity, we switch to the second nearby contract. We time the switch to capture the contract month with the greatest trading volume. This results in switching contracts for S&P 500 futures once the nearby contract enters its final week and for bond and gold futures once the nearby contract enters the delivery month. This procedure yields a continuous series of 3,763 daily returns for each market.

Table I provides descriptive statistics for the returns. Panel A shows that the average returns, μ , are highest for stock index futures followed by bonds and then gold. The standard deviations, σ , indicate that stocks are most volatile and that gold is more volatile than bonds. Panel B reports the sample return correlations, $\rho(R)$. The correlation between stock and bond returns is positive (0.397), whereas the correlations between stock and gold returns (-0.105) and bond and gold returns (-0.157) are negative. These findings seem reasonable given historic spot market returns and the implications of the cost-of-carry relation.

B. The Conditional Covariance Matrix Estimates

We use equation (4) to estimate the conditional covariance matrix. We subtract the sample mean from the raw returns for each asset and then form the two-sided rolling weighted average of the squares and cross-products of the return innovations. The weights are given by equation (5) with $\alpha_{ij,t}$ equal

Table I
Summary Statistics for Daily S&P 500, T-Bond, and Gold Futures Returns

The table provides summary statistics for daily returns on S&P 500 stock index futures, T-bond futures, and gold futures. The returns are computed as daily price relatives minus one. We estimate the conditional covariance matrix using the Foster and Nelson (1996) two-sided procedure described in the text. Panel A reports the mean returns (μ), standard deviations (σ), and mean conditional volatilities ($\bar{\sigma}_t$). These values are annualized using 252 trading days per year. Panel B reports the cross-market correlations of returns ($\rho(R)$), the average conditional correlation based on our covariance matrix estimates ($\rho(r)$), and the correlations of the changes in conditional volatilities ($\rho(\sigma_t)$). Panel C reports the estimated autocorrelations of the conditional volatility estimates (ρ_l) for $l = 1, \dots, 10$ lags. The sample period is January 3, 1983 through December 31, 1997 (3,763 observations), and is divided into a pre-crash period ending September 30, 1987 (1,194 observations) and a post-crash period beginning November 2, 1987 (2,547 observations).

Period	S&P 500 Futures			T-Bond Futures			Gold Futures		
	μ	σ	$\bar{\sigma}_t$	μ	σ	$\bar{\sigma}_t$	μ	σ	$\bar{\sigma}_t$
Entire sample	10.82	16.16	14.37	6.53	10.46	10.07	-7.76	14.67	13.51
Pre-crash	13.99	14.82	15.12	6.09	12.57	12.31	-5.98	19.07	18.47
Post-crash	11.18	14.05	13.52	6.06	9.03	8.92	-8.94	12.00	11.16

Period	S&P 500 & T-Bonds			S&P 500 & Gold			T-Bonds & Gold		
	$\rho(R)$	$\rho(r)$	$\rho(\sigma_t)$	$\rho(R)$	$\rho(r)$	$\rho(\sigma_t)$	$\rho(R)$	$\rho(r)$	$\rho(\sigma_t)$
Entire sample	0.397	0.454	0.527	-0.105	-0.120	0.192	-0.157	-0.169	0.195
Pre-crash	0.475	0.447	0.374	0.050	0.073	0.125	-0.080	-0.066	0.188
Post-crash	0.426	0.459	0.506	-0.237	-0.210	0.204	-0.228	-0.215	0.160

Contract	Panel C: Autocorrelations of Volatility Changes									
	ρ_1	ρ_2	ρ_3	ρ_4	ρ_5	ρ_6	ρ_7	ρ_8	ρ_9	ρ_{10}
S&P 500	0.960	0.918	0.864	0.809	0.757	0.703	0.649	0.597	0.546	0.495
T-Bond	0.923	0.853	0.791	0.736	0.682	0.625	0.570	0.520	0.472	0.426
Gold	0.925	0.851	0.783	0.719	0.658	0.601	0.549	0.502	0.459	0.414

to our estimate of 0.063. Figure 1 plots the resulting estimates of the conditional volatilities (Panel A) and correlations (Panel B). Table I reports the associated summary statistics.

The average volatility estimates are reported as $\bar{\sigma}_t$ in Table I, Panel A, and are generally consistent with the standard deviations of raw returns. The autocorrelation structure, shown in Panel C of Table I, reveals a strong degree of persistence for each series. Figure 1 shows that the volatility estimates vary considerably over the course of the sample, and the correlations, reported as $\rho(\sigma_t)$ in Panel B of Table I, indicate that the volatility changes are positively correlated across markets. Figure 1 also shows substantial time-series variation in the rolling return correlations. The average estimates, $\rho(r)$ in Panel B of Table I, are consistent with the sample correlations, $\rho(R)$, but the stock/gold and bond/gold correlations sharply decrease after the crash and the stock/bond correlations widely fluctuate across the sample.

III. Empirical Results

The time-series variation in the covariance matrix estimates suggests a potential role for volatility timing in asset-allocation decisions. In this section, we assess the economic value of volatility timing to short-horizon investors. First, we operationalize the portfolio optimization procedure and examine the portfolio weights and the ex post returns for the dynamic strategies. We then compare the performance of the dynamic portfolios to that of the ex ante optimal static portfolios under various levels of estimation risk. Finally, we assess the sensitivity of our findings to the choice of decay rate used to generate the covariance matrix estimates.

Constructing the optimal portfolios requires estimates of the conditional expected returns, variances, and correlations. Estimates of the conditional variances and correlations are obtained using the procedure in Section I.B. We use the one-sided estimates that are based only on information available on a real-time basis. Although we treat the conditional expected returns as constant, it is not clear how we should estimate them. A natural choice would be to use spot data prior to the start of our sample. The 1970s, however, produced dramatic economic changes such as the oil crisis, the shift in Federal Reserve interest rate policy, and the elimination of the gold standard. Using returns data from this period would be appropriate only if we assume investors expected similar changes in the 1980s and 1990s. More generally, to the extent that the portfolio returns are sensitive to the expected return estimates, using any single set of estimates would make it difficult to assess the robustness of our results.

To avoid this problem, we consider a range of expected return estimates generated using a bootstrap approach (Efron (1979)). The bootstrap is a simple resampling technique that controls for the uncertainty in estimating population parameters from sample data. We begin by drawing randomly with replacement from the actual returns to generate a series of artificial re-

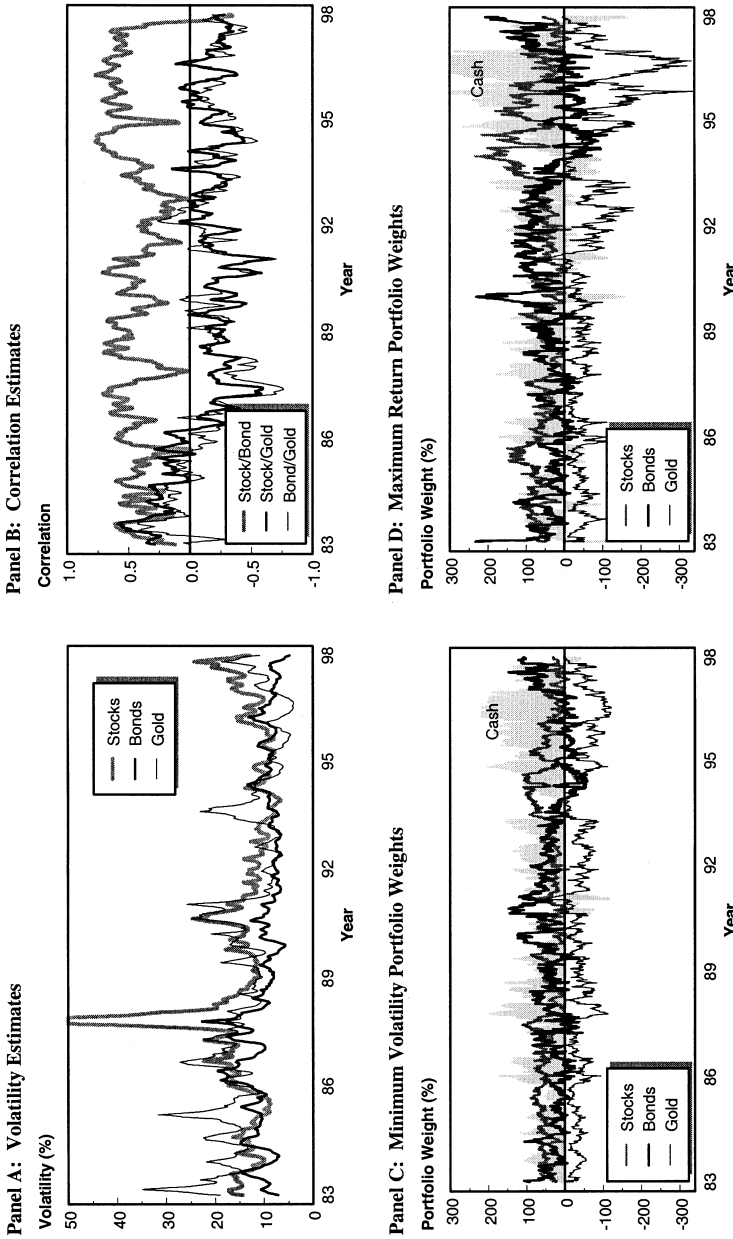


Figure 1. Daily volatility estimates, correlation estimates, and portfolio weights for the volatility-timing strategies. Panels A and B show the daily estimates of the conditional volatilities (annualized) and cross-market correlations, respectively, of S&P 500, T-bond, and gold futures returns. The estimates are obtained using the Foster and Nelson (1996) two-sided procedure described in the text. Panels C and D show the results of daily portfolio optimizations using stock, bond, and gold futures. We assume that the expected returns are equal to their in-sample means and we estimate the daily covariance matrix using the Foster and Nelson (1996) one-sided procedure. Panel C shows the weights that minimize conditional volatility while setting the expected return equal to 10 percent, and Panel D shows the weights that maximize expected return while setting the conditional volatility equal to 12 percent. The sample period is January 3, 1983 through December 31, 1997.

turns. Then, we compute the mean returns in this artificial sample and use them, along with our conditional covariance matrix estimates, to compute the optimal portfolio weights. Finally, we apply these weights to the actual returns and conduct our performance evaluations. This approach allows us to mimic the estimation risk that an investor would face when estimating expected returns using a sample of any given size.

Our approach also requires a benchmark portfolio in order to measure the value of volatility timing. The natural choice, given our mean-variance framework, is the unconditionally optimal static portfolio for the same target expected return or volatility used in the dynamic strategies. To construct this benchmark, we use our bootstrap estimates of the unconditional expected returns, volatilities, and covariances to solve for the required portfolio weights. This ensures that the static and dynamic portfolios are based on the same *ex ante* information.

A. The Portfolio Weights and Returns

Before implementing the bootstrap, it is useful to establish baseline results for the case where estimation risk is negligible. This corresponds to constructing the portfolio weights based on the mean returns over the entire sample along with the one-step-ahead estimates of the conditional covariance matrix. Figure 1 shows the resulting weights for the minimum volatility (Panel C) and maximum return (Panel D) strategies. The weights are based on a target expected return of 10 percent and a target volatility of 12 percent.

As expected, the sign and magnitude of each of the weights depends on the estimated expected returns and the conditional volatility and correlation estimates. For example, the weight in gold is generally negative because the average return on gold futures is negative. But the size of this short position decreases when gold volatility increases, as in 1985 and 1993 (Figure 1, Panel A), and when gold's correlation with stocks and bonds becomes more negative, as in 1987 and 1991 (Figure 1, Panel B). Similarly, the split between stocks and bonds is sensitive to their relative volatilities. Stock volatility decreases steadily from 1991 to 1994 whereas bond volatility remains relatively constant. As a result, the weight in stocks steadily increases over this period whereas the weight in bonds decreases. The swings in the weights are more pronounced in Panel D because a greater risk exposure is needed to match the target volatility of 12 percent. Panels C and D also show the implicit weights in cash. A negative cash weight means that the corresponding position in the underlying assets is levered; there are relatively few instances, however, when this occurs.

We compute the *ex post* daily returns for each strategy by multiplying the portfolio weights by the observed next-day returns on stock, bond, and gold futures. The minimum volatility strategy yields a mean return of 9.8 percent, a sample volatility of 10.7 percent, and an estimated Sharpe ratio of

0.92. For the maximum return strategy, the mean return is 12.4 percent, the sample volatility is 13.6 percent, and the estimated Sharpe ratio is 0.91. To put these results in perspective, consider the ex post minimum variance frontier for fixed-weight portfolios. The maximum realizable Sharpe ratio is 0.86, so both of the dynamic strategies plot above the efficient frontier. The implication of this finding is clear. It is unlikely that we would have chosen ex ante a fixed-weight portfolio that turned out to be ex post efficient; but, even if we had, we would not have outperformed either of the dynamic strategies.

To assess the statistical significance of the volatility timing results, we conduct simulations where the asset returns are generated independently of the portfolio weights. We first form a random permutation of the actual return series and then we apply the actual weights to the randomized returns to compute portfolio returns.⁵ If the volatility-timing gains are significant, the strategies should perform better using the actual data than in the simulations. We find this to be the case. For the minimum volatility strategy, the mean return and volatility across 10,000 trials are 9.7 percent and 12.9 percent. No trial produces a volatility as low as that observed using the actual returns, and only 9.2 percent of the trials yield a higher Sharpe ratio. For the maximum return strategy, the mean return and volatility are 14.0 percent and 23.6 percent, and only 0.4 percent of the trials yield a higher Sharpe ratio. These findings indicate it is unlikely that the gains to volatility timing are due to chance.

Table II breaks down the actual ex post portfolio returns for the dynamic strategies by three-year subperiods. The average returns and sample volatilities vary considerably across the subperiods, with the worst and best performance during the periods from 1992 to 1994 and 1995 to 1997, respectively. Of course, these results are based on fairly small samples, so much of the variation may be attributable to estimation error. The final two lines of the table show that the 1987 stock market crash has little impact on the results. Specifically, if we exclude either the two-week period following the crash or the entire 1986 to 1988 subperiod, the performance of the dynamic strategies is comparable to that observed over the entire sample.

B. The Impact of Estimation Risk

Although our previous results suggest volatility timing may have value, these results do not account for estimation risk. We evaluate the impact of this risk using the bootstrap approach described earlier. Suppose, for example, that we want to mimic the estimation risk that an investor would face using a sample size comparable to our sample. We generate an artificial sample of 4,000 returns by drawing randomly with replacement from the

⁵ Note that this is asymptotically equivalent to using the actual returns and randomizing the portfolio weights. Either way, the weights are independent of the asset returns.

Table II
Ex Post Performance of the Volatility-timing Strategies

The table summarizes the ex post performance of the volatility-timing strategies. The dynamic portfolio weights are determined by solving a daily portfolio optimization problem in which the expected returns are assumed to be constant and equal to the in-sample mean returns. The daily estimates of the conditional covariance matrix of returns are obtained using the Foster and Nelson (1996) one-sided procedure described in the text. We solve for two sets of dynamic weights: (1) those that minimize conditional volatility while setting the conditional expected return equal to 10 percent, and (2) those that maximize conditional expected return while setting the conditional volatility equal to 12 percent. For each set of weights, we report the annualized mean realized return (μ), the annualized realized volatility (σ), and the realized Sharpe ratio (SR). The sample period is January 3, 1983 through December 31, 1997. The first three months of data are withheld to initialize the volatility estimation procedure. We also report results for each three-year subsample and for two noncrash periods that exclude either October 19 to 30, 1987 or the entire 1986 to 1988 subperiod.

Period	Obs.	Minimum Volatility			Maximum Return		
		μ	σ	SR	μ	σ	SR
Entire sample	3,700	9.75	10.65	0.916	12.38	13.57	0.912
1983–1985	692	8.11	8.96	0.905	10.67	13.23	0.806
1986–1988	757	10.76	15.11	0.712	7.36	14.16	0.520
1989–1991	758	9.96	11.33	0.879	9.40	13.77	0.682
1992–1994	751	4.24	7.63	0.555	6.88	13.18	0.522
1995–1997	742	15.63	8.26	1.892	27.71	13.45	2.060
Noncrash samples							
ex. Oct. 19–30, 1987	3,690	9.18	10.46	0.878	11.91	13.48	0.883
ex. 1986–1988	2,943	9.50	9.17	1.036	13.67	13.42	1.019

actual returns. We then compute the portfolio weights using the artificial sample means (instead of the true sample means), apply these weights to the actual returns, and evaluate the performance of the dynamic strategies.⁶

We quantify the impact of estimation risk by comparing the simulation results to the case where estimation risk is negligible. Across 10,000 simulation trials, the mean Sharpe ratios for the minimum volatility and maximum return strategies are 0.84 and 0.85, compared to the 0.92 and 0.91 reported earlier. We can translate these differences into annualized basis point fees using our utility-based approach. This indicates that to eliminate the risk of estimating the expected returns from a sample size of 4,000, an

⁶ Most investors would weigh the expected returns obtained from sampling against their prior expectations. Specifically, asset pricing theory suggests that the unconditional expected returns for stock and bond futures should be positive, and that stocks should have the highest expected return, followed by bonds, and then gold. We incorporate these priors into our bootstrap experiment by requiring that the average returns in each of our artificial samples satisfy both conditions.

Table III
The Effect of Estimation Risk on the Performance
of the Volatility-timing Strategies

The table illustrates the effect of estimation risk on the performance of the volatility-timing strategies. We mimic this risk using a bootstrap approach. Specifically, we generate a random sample of k observations by sampling with replacement from the actual returns, and we use the mean returns for this sample to compute the dynamic portfolio weights. These weights are determined by solving a daily portfolio optimization problem in which the conditional covariance matrix of returns is estimated using the Foster and Nelson (1996) one-sided procedure described in the text. We repeat the bootstrap experiment 10,000 times for each value of k . The table reports the annualized mean realized returns (μ), annualized realized volatilities (σ), and the realized Sharpe ratios (SR), as well as the average annualized basis point fees (Δ_γ) that an investor with quadratic utility and constant relative risk aversion of $\gamma = 1$ or $\gamma = 10$ would be willing to pay to eliminate the risk associated with estimating expected returns.

k	Minimum Volatility Strategy					Maximum Return Strategy				
	μ	σ	SR	Δ_1	Δ_{10}	μ	σ	SR	Δ_1	Δ_{10}
500	5.28	7.22	0.736	423.6	207.3	10.38	13.55	0.765	199.7	197.3
1,000	6.57	8.65	0.766	307.4	204.7	10.73	13.56	0.791	164.3	162.5
1,500	7.32	9.44	0.782	239.9	204.1	10.92	13.56	0.805	145.7	143.9
2,000	7.78	9.83	0.799	196.4	189.4	11.12	13.56	0.820	125.5	124.3
2,500	8.15	10.13	0.810	162.4	179.9	11.24	13.56	0.829	113.5	112.3
3,000	8.42	10.29	0.824	136.1	159.3	11.40	13.57	0.841	97.4	96.4
3,500	8.59	10.38	0.833	119.1	141.5	11.50	13.57	0.848	87.8	87.0
4,000	8.73	10.46	0.840	105.0	130.9	11.58	13.57	0.853	80.3	79.7
4,500	8.90	10.58	0.845	89.7	126.7	11.64	13.57	0.858	74.0	73.3
5,000	8.96	10.59	0.850	83.1	115.0	11.69	13.57	0.861	68.8	68.4
10,000	9.34	10.65	0.879	43.0	60.9	11.99	13.57	0.884	38.5	38.3

investor with a relative risk aversion of one would be willing to pay 105 and 80 basis points, on average, when implementing the minimum volatility and maximum return strategies. The magnitude of these fees suggests that the impact of estimation risk can be substantial.

Table III shows how the fees vary with the level of estimation risk. As expected, the fees decline monotonically with the sample size, independent of whether the risk aversion is $\gamma = 1$ or $\gamma = 10$. The overall relation between the fees, risk aversion, and estimation risk, however, is more complex. In some instances, the fee for $\gamma = 1$ is higher than the fee for $\gamma = 10$. To see how this occurs, note that both the mean and volatility tend to fall as the number of observations in the sample decreases. Thus, when the number of observations is small, the resulting portfolio is likely to be less attractive to an investor with low risk aversion than to an investor with high risk aversion. Under these circumstances, eliminating the estimation risk can have a greater impact for investors with low risk aversion.

C. The Value of Volatility Timing

We use a similar approach to estimate the value of volatility timing under various levels of estimation risk. Specifically, we compare the performance of the dynamic portfolios to the performance of the unconditionally mean-variance efficient static portfolios that have the same target expected return and volatility. We construct the static portfolios using the unconditional means, variances, and covariances obtained from the artificial data.⁷ As before, we then apply the optimal fixed weights to the actual returns to evaluate ex post performance. Because the static and dynamic portfolios incorporate the same level of estimation risk, their relative performance should reliably indicate the value of volatility timing.

Figure 2 shows the results for 10,000 simulation trials where $k = 4,000$ observations. Each point in the figure represents a separate trial, plotting the realized Sharpe ratio for both the static (x-axis) and dynamic (y-axis) portfolios. For the minimum volatility strategy (Panel A), the points are clustered around a 45-degree line through the figure, suggesting that its advantage over the static portfolio is relatively modest. Nonetheless, in 84 percent of the trials, the dynamic portfolio achieves the higher Sharpe ratio. For the maximum return strategy (Panel B), the performance differential is even greater. The distribution clearly shifts above the 45-degree line, and the dynamic portfolio achieves the higher Sharpe ratio in 92 percent of the trials.


Averaging across the simulation trials, both of the static portfolios produce a mean Sharpe ratio of 0.80, compared to 0.84 and 0.85 for the dynamic strategies. Although these differences do not seem large, they can translate into sizable performance fees. With $\gamma = 1$, for example, our quadratic-utility investor would be willing to pay an estimated 173 basis points annually to switch from the ex ante optimal static portfolio to the maximum return strategy. In other cases, however, the estimated fees are much smaller. The fees to switch to the maximum return strategy when $\gamma = 10$ or to the minimum volatility strategy when $\gamma = 1$ are close to zero.

Figure 3 shows the distribution of the performance fees across the 10,000 simulation trials. For the minimum volatility strategy (Panel A), the distribution is close to symmetric around a mean of -1 basis point when $\gamma = 1$. For $\gamma = 10$, the distribution shifts to the right, with an estimated mean of 60 basis points. The results for the maximum return strategy (Panel B) look quite different. For $\gamma = 1$, the distribution is skewed to the left with an estimated mean of 173 basis points, whereas the distribution for $\gamma = 10$ is skewed to the right with an estimated mean of -3 basis points. These differences highlight the effect of risk aversion on the trade-off between risk

⁷ An investor implementing the static strategy would generally expect greater volatility for stocks than for bonds. Therefore, in addition to the conditions given in footnote 6, we require that each of our artificial samples has a higher volatility for stocks than for bonds.

Panel A: Minimum Volatility Strategies

Dynamic



Panel B: Maximum Return Strategies

Dynamic




Figure 2. Sharpe ratios for the volatility-timing and ex ante optimal static strategies.

The figure plots the realized Sharpe ratios for the volatility-timing and ex ante optimal static strategies. The input parameters used in the portfolio optimizations are determined by bootstrapping. Specifically, using actual returns, we generate an artificial sample of 4,000 observations from which we estimate the mean returns, volatilities, and covariances for each asset. For the volatility-timing strategies, we use the estimated mean returns and our daily covariance matrix estimate to determine the daily optimal portfolio. For the static strategies, we use the estimated mean returns, volatilities, and covariances from the artificial sample to determine the unconditional optimal portfolio. The figure plots the realized Sharpe ratios for each of 10,000 trials of this bootstrap experiment. The sample period for the realized returns is April 4, 1983 through December 31, 1997.

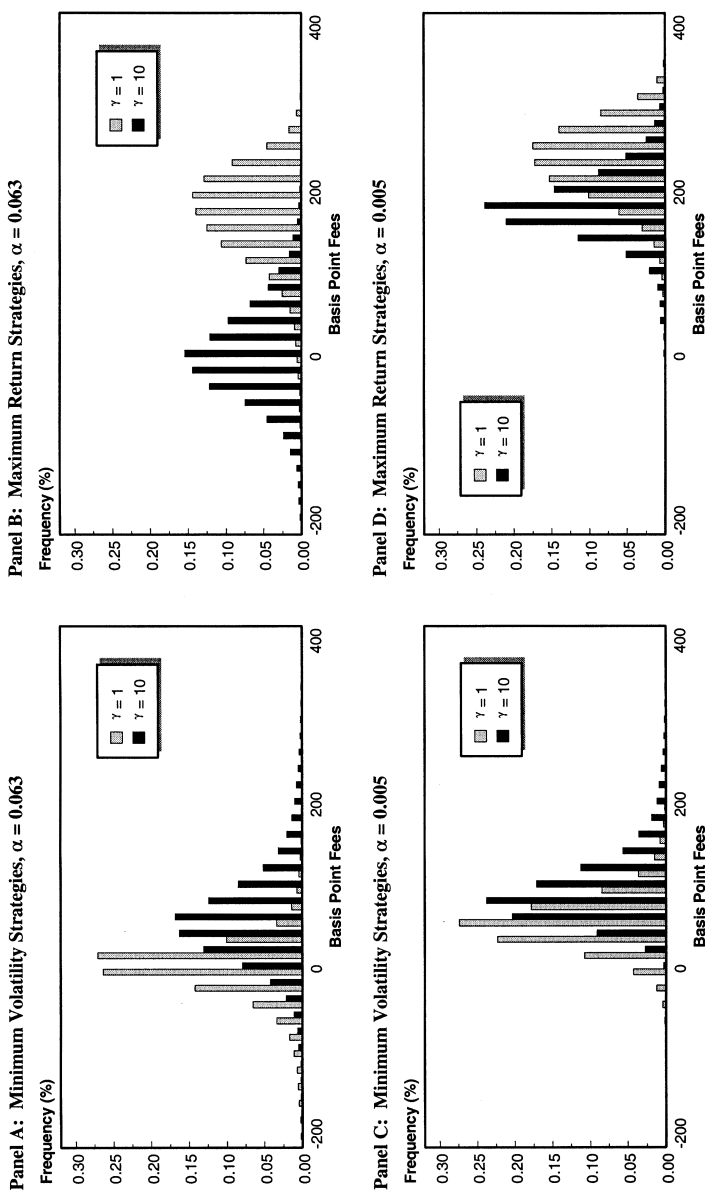


Figure 3. Distribution of the performance fees to switch from the optimal static strategies to the volatility-timing strategies. The figure shows the distribution of fees across 10,000 bootstrap experiments that an investor with quadratic utility and relative risk aversion of γ would pay to switch from the ex ante optimal static strategies to the volatility-timing strategies. In the bootstrap experiments, we generate a random sample of 4,000 observations by sampling from the actual returns to estimate the unconditional mean returns, volatilities, and covariances. We use these estimates to form the unconditional optimal portfolio. We implement the volatility-timing strategies using the estimated means and our daily covariance matrix estimates. Panels A and B show the results when $\alpha = 0.063$ is used to estimate the daily covariance matrix. Panels C and D show the results when $\alpha = 0.005$. The sample period for the realized returns is April 4, 1983 through December 31, 1997.

Table IV
Comparison of the Volatility-timing and Ex Ante Optimal Static Strategies

The table compares the performance of the volatility-timing strategies to that of the ex ante optimal static portfolios. We use a bootstrap procedure to simulate the ex ante information set. Specifically, we generate a random sample of k observations by sampling with replacement from the actual returns in order to estimate the unconditional mean returns, volatilities, and correlations. Using these estimates, we compute the weights that deliver the unconditionally efficient static portfolios. To obtain the weights for the volatility-timing strategies, we use the mean returns from the bootstrap sample and our daily estimates of the conditional covariance matrix to solve the daily optimal portfolio problem. We repeat the bootstrap experiment 10,000 times for each k using a target expected return and volatility of 10 percent and 12 percent, respectively. The table reports the annualized mean realized returns (μ), annualized realized volatilities (σ), and realized Sharpe ratios (SR) for each strategy, the proportion of trials in which the volatility-timing strategy has a higher Sharpe ratio than the static portfolio (p -value), and the average annualized basis point fees (Δ_1) that an investor with quadratic utility and constant relative risk aversion of $\gamma = 1$ or $\gamma = 10$ would be willing to pay to switch from the static portfolios to the volatility-timing strategies.

k	Static Portfolio			Dynamic Portfolio			p -value	Δ_1	Δ_{10}
	μ	σ	SR	μ	σ	SR			
Panel A: Minimum Volatility Strategies									
1,000	6.76	9.16	0.742	6.57	8.65	0.766	0.7137	-13.8	35.5
2,000	7.94	10.39	0.770	7.78	9.83	0.799	0.7653	-9.6	48.3
3,000	8.53	10.87	0.789	8.42	10.29	0.824	0.8075	-4.3	56.7
4,000	8.81	11.04	0.801	8.73	10.46	0.840	0.8379	-1.0	60.3
5,000	9.02	11.18	0.809	8.96	10.59	0.850	0.8603	0.5	63.3
10,000	9.36	11.26	0.832	9.34	10.65	0.879	0.9276	5.1	67.8
Panel B: Maximum Return Strategies									
1,000	9.10	12.27	0.742	10.73	13.56	0.791	0.8049	147.0	-0.5
2,000	9.34	12.14	0.770	11.12	13.56	0.820	0.8563	159.8	-3.9
3,000	9.53	12.08	0.789	11.40	13.57	0.841	0.8921	169.0	-2.3
4,000	9.65	12.05	0.801	11.58	13.57	0.853	0.9175	173.3	-2.7
5,000	9.74	12.04	0.809	11.69	13.57	0.861	0.9331	175.5	-1.9
10,000	9.98	12.00	0.832	11.99	13.57	0.884	0.9812	181.1	-1.2

and return. When risk aversion is low, the higher return of the dynamic strategy is worth the slight increase in volatility. But, when risk aversion is high, the increase in return is not enough to compensate for the greater volatility.

Table IV shows how the performance comparisons vary with the level of estimation risk. As we increase the sample size, the Sharpe ratios for both the static and dynamic strategies increase, as does the fraction of trials in which the dynamic strategies outperform the static strategies (p -value). However, varying sample size has little impact on the estimated performance fees. This is revealing given that estimation risk has a big impact on the

performance of the dynamic strategies (Table III). Apparently, estimation risk has a similar effect on both the static and dynamic strategies. Therefore, if the level of estimation risk is the same for both strategies, changing the level has little impact on the relative performance fees.

Finally, Table V illustrates the effect of changing the target expected return and volatility. Recall that our previous results are based on targets of $\mu_p = 10$ percent and $\sigma_p = 12$ percent. Changing the targets moves the portfolio to a different point along the efficient frontier. This changes the portfolio's expected return and volatility but not its Sharpe ratio. The minor differences in the Sharpe ratios in Table V are solely due to sampling variation associated with the bootstrap procedure. The performance fees, on the other hand, do vary with the targets because the trade-off between risk and return is nonlinear. Nonetheless, regardless of the choice of targets, volatility timing always has value for some reasonable level of risk aversion. Because this suggests that our overall results are robust to this issue, we continue to rely on targets of 10 percent and 12 percent for the remainder of the analysis.

D. Sensitivity Analysis

All of our previous results are based on the estimated decay rate that minimizes the asymptotic MSE of our covariance matrix estimator. Although this is the optimal choice with respect to estimator efficiency, it may be suboptimal in terms of implementing our volatility-timing strategies. Our estimated decay rate of 0.063 implies that the half-life of a volatility shock is about 11 trading days.⁸ This is much shorter than the two- to six-month half-lives commonly obtained using GARCH models for U.S. equity returns (see, e.g., Bollerslev et al. (1992)). Perhaps an investor would be better off using smoother covariance matrix estimates to reduce some of the noise inherent in the estimated portfolio weights. We investigate this by examining the sensitivity of our findings to the choice of decay rate.

Table VI reports the Sharpe ratios, probability values, and estimated performance fees for decay rates ranging from 0.001 to 0.1. Not surprisingly, the results vary substantially as we move from one end of this range to the other. The dynamic strategies have the greatest performance edge in the region $\alpha = 0.005$ to 0.02. For example, with $\gamma = 1$ and $\alpha = 0.005$, the maximum return strategy has a Sharpe ratio of 0.96, a probability value of 1.00, and an estimated performance fee of 241 basis points. This choice of decay rate implies that the half-life of a volatility shock is about 139 trading days, generating much smoother covariance matrix estimates than we used above. Thus, volatility timing appears to be more effective when we smooth the estimates more than is optimal under our asymptotic MSE criterion.

Figure 3 shows the distribution of the performance fees when $\alpha = 0.005$. For the minimum volatility strategy (Panel C), the distribution for $\gamma = 10$ still plots to the right of that for $\gamma = 1$, as in Panel A. Now, however, almost all of the fees are greater than zero, regardless of whether $\gamma = 1$ or $\gamma = 10$. A similar shift in

⁸ The half-life is $-\ln(0.5)/\alpha$.

Table V
Comparison of the Volatility-timing and Ex Ante Optimal Static Strategies Using Different Target Expected Returns and Volatilities

The table shows how the performance of the volatility-timing strategies varies with the target expected return and volatility. We form a random sample of 4,000 observations by sampling with replacement from the actual returns to estimate the unconditional mean returns, volatilities, and correlations. Using these estimates, we compute the weights that deliver the unconditionally efficient static portfolios. To obtain the weights for the volatility-timing strategies, we use the mean returns from the bootstrap sample and our daily estimates of the conditional covariance matrix to solve the daily optimal portfolio problem. We repeat this procedure 10,000 times for each target expected return and volatility. The table reports the annualized mean realized returns (μ), annualized realized volatilities (σ), and realized Sharpe ratios (SR) for each strategy, the proportion of trials in which the volatility-timing strategy has a higher Sharpe ratio than the static portfolio (p -value), and the average annualized basis point fees (Δ_y) that an investor with quadratic utility and constant relative risk aversion of $\gamma = 1$ or $\gamma = 10$ would be willing to pay to switch from the static portfolios to the volatility-timing strategies.

Panel A: Minimum Volatility Strategies									
Target Return	Static Portfolio			Dynamic Portfolio			p -value	Δ_1	Δ_{10}
	μ	σ	SR	μ	σ	SR			
6%	5.27	6.59	0.802	5.22	6.24	0.841	0.8401	-2.2	19.7
7%	6.18	7.75	0.801	6.13	7.34	0.840	0.8435	-1.5	28.7
8%	7.05	8.83	0.801	6.98	8.36	0.839	0.8366	-2.7	36.5
9%	7.97	9.99	0.800	7.89	9.46	0.839	0.8354	-2.0	48.1
10%	8.86	11.10	0.801	8.78	10.51	0.840	0.8401	-1.3	61.3
11%	9.67	12.13	0.800	9.59	11.49	0.839	0.8399	-0.2	73.8
12%	10.58	13.28	0.801	10.49	12.58	0.839	0.8398	0.4	89.0
13%	11.49	14.44	0.800	11.39	13.67	0.838	0.8411	1.5	107.6
14%	12.31	15.44	0.800	12.20	14.62	0.839	0.8379	2.3	122.8

Panel B: Maximum Return Strategies									
Target Volatility	Static Portfolio			Dynamic Portfolio			p -value	Δ_1	Δ_{10}
	μ	σ	SR	μ	σ	SR			
8%	6.44	8.03	0.802	7.73	9.04	0.854	0.9163	119.9	42.1
9%	7.24	9.04	0.801	8.69	10.18	0.854	0.9183	134.0	35.7
10%	8.04	10.04	0.801	9.64	11.31	0.853	0.9151	146.4	24.3
11%	8.84	11.05	0.800	10.60	12.44	0.852	0.9159	159.5	12.9
12%	9.66	12.06	0.801	11.58	13.57	0.853	0.9120	172.1	-1.8
13%	10.45	13.06	0.800	12.53	14.70	0.853	0.9132	185.9	-19.6
14%	11.26	14.07	0.801	13.49	15.83	0.852	0.9157	197.4	-40.3
15%	12.05	15.07	0.800	14.46	16.96	0.852	0.9162	210.7	-62.2
16%	12.86	16.07	0.800	15.43	18.09	0.853	0.9160	222.4	-89.1

the distribution is evident for the maximum return strategy (Panel D). For both levels of risk aversion, the fees almost always exceed 100 basis points, and roughly half are greater than 200 basis points. This provides the strongest evidence yet that volatility timing has significant value.

Table VI
The Effect of the Decay Rate on the Performance
of the Volatility-timing Strategies

The table illustrates the effect of the decay rate (α) used to estimate the daily conditional covariance matrix on the performance of the volatility-timing strategies. We form a random sample of 4,000 observations by sampling with replacement from the actual returns to estimate the unconditional mean returns, volatilities, and correlations. Using these estimates, we compute the weights that deliver the unconditionally efficient static portfolios. To obtain the weights for the volatility-timing strategies, we use the mean returns from the bootstrap sample and our daily estimates of the conditional covariance matrix to solve the daily optimal portfolio problem. We repeat this procedure 10,000 times for each value of α using a target expected return and volatility of 10 percent and 12 percent, respectively. The table reports the annualized mean realized returns (μ), annualized realized volatilities (σ), and realized Sharpe ratios (SR) for each strategy, the proportion of trials in which the volatility-timing strategy has a higher Sharpe ratio than the static portfolio (p -value), and the average annualized basis point fees (Δ_γ) that an investor with quadratic utility and constant relative risk aversion of $\gamma = 1$ or $\gamma = 10$ would be willing to pay to switch from the static portfolios to the volatility-timing strategies.

α	Static Portfolio			Dynamic Portfolio			p -value	Δ_1	Δ_{10}
	μ	σ	SR	μ	σ	SR			
Panel A: Minimum Volatility Strategies									
0.001	8.78	10.99	0.802	9.01	10.99	0.823	0.9364	23.4	23.4
0.005	8.83	11.07	0.801	9.38	10.74	0.876	0.9925	59.1	94.5
0.010	8.81	11.04	0.801	9.15	10.56	0.870	0.9724	39.4	90.9
0.020	8.85	11.10	0.800	9.00	10.48	0.862	0.9440	22.0	87.9
0.040	8.79	11.03	0.800	8.84	10.37	0.857	0.9275	12.6	82.4
0.060	8.84	11.10	0.800	8.78	10.50	0.841	0.8565	1.1	65.3
0.080	8.78	11.02	0.800	8.58	10.53	0.820	0.7115	-13.9	37.3
0.100	8.85	11.07	0.802	8.54	10.70	0.802	0.5624	-26.9	12.5
Panel B: Maximum Return Strategies									
0.001	9.66	12.05	0.802	12.46	15.66	0.796	0.3740	230.2	-222.8
0.005	9.65	12.05	0.801	12.13	12.63	0.960	0.9994	240.9	177.0
0.010	9.65	12.05	0.801	11.95	12.51	0.955	0.9978	224.5	173.8
0.020	9.65	12.06	0.800	11.80	12.64	0.933	0.9921	207.7	143.5
0.040	9.64	12.05	0.800	11.68	13.05	0.895	0.9793	191.2	78.9
0.060	9.64	12.06	0.800	11.58	13.50	0.858	0.9334	176.0	9.9
0.080	9.64	12.05	0.800	11.45	13.96	0.820	0.6873	156.2	-68.0
0.100	9.66	12.05	0.802	11.38	14.44	0.788	0.3658	140.3	-145.9

A final issue is whether the value of volatility timing is offset by transaction costs. We assess this by running another set of simulations in which we impose various levels of proportional transaction costs. The results indicate that, over a wide range, the impact of transaction costs is approximately linear—when $\alpha = 0.005$, each percentage point increase in the one-way cost reduces the mean return for the dynamic strategies by 4 to 5 basis points. Thus, the transaction costs would need to be 19 (or 47) percent annualized to equate the Sharpe ratios for the static and minimum volatility (or maximum return) strategies. To give this some perspective, suppose conservatively that

the bid/ask spread and round-trip commission costs for S&P 500 futures total \$0.10 per index unit. With an average index level of \$384.51 over our sample period, this implies a one-way transaction cost of $0.05/384.51 \times 252 = 3.28$ percent annualized. Our results indicate that transaction costs need to be at least six times this estimate to offset the dynamic strategies' advantage.

IV. Conclusions

Researchers have long known that volatility is predictable. However, the low explanatory power of standard volatility models has led to questions about their economic relevance. Our analysis indicates that the predictability captured by volatility modeling is economically significant. In particular, we find that volatility-timing strategies based on one-step-ahead estimates of the conditional covariance matrix significantly outperform the unconditionally efficient static portfolios with the same target expected return and volatility. This finding is robust to both the level of estimation risk and transaction costs. Moreover, our results are probably conservative because we use a simple nonparametric volatility specification and ignore any linkage between volatility changes and changes in expected returns.

Our work suggests a number of possible directions for future research. First, because volatility timing requires active trading, hedge funds are a likely source for further empirical evidence. It would be interesting, for example, to develop a volatility-timing coefficient for the various classes of hedge funds and examine whether the cross-sectional variation in this coefficient explains differences in fund performance. Another possibility is to assess the importance of volatility modeling in applications unrelated to asset allocation. We are currently pursuing this issue in the context of derivatives risk management. Finally, recent work by Ferson and Siegel (1999) and Bekaert and Liu (1999) shows how to optimally incorporate conditioning information when the objective is to maximize an unconditional mean-variance criterion. This provides a natural framework for investigating the linkage between volatility timing, return predictability, and traditional methods of performance evaluation.

REFERENCES

- Andersen, Torben G., 1996, Return volatility and trading volume: An information flow interpretation of stochastic volatility, *Journal of Finance* 51, 169–204.
- Andersen, Torben G., and Tim Bollerslev, 1998, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* 39, 885–905.
- Bekaert, Geert, and Jun Liu, 1999, Conditioning information and variance bounds on pricing kernels, Working paper, Stanford University.
- Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307–328.
- Bollerslev, Tim, Ray Y. Chou, and Kenneth F. Kroner, 1992, ARCH modeling in finance: A review of the theory and empirical evidence, *Journal of Econometrics* 52, 5–60.
- Bollerslev, Tim, Robert F. Engle, and Daniel B. Nelson, 1994, ARCH models, in Robert F. Engle and Daniel McFadden, eds.: *Handbook of Econometrics*, Vol. IV (North-Holland, Amsterdam).
- Busse, Jeffrey A., 1999, Volatility timing in mutual funds: Evidence from daily returns, *Review of Financial Studies* 12, 1009–1041.

- Copeland, Maggie M., and Thomas E. Copeland, 1999, Market timing: Style and size rotation using VIX, *Financial Analysts Journal* 55, 73–81.
- Diebold, Francis X., and Jose A. Lopez, 1995, Modeling volatility dynamics, in Kevin Hoover, ed.: *Macroeconometrics: Developments, Tensions, and Prospects* (Kluwer Academic Press, Boston).
- Efron, Bradley, 1979, Bootstrap methods: Another look at the jackknife, *Annals of Statistics* 7, 1–26.
- Engle, Robert F., 1982, Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* 50, 987–1008.
- Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.
- Ferson, Wayne E., and Andrew F. Siegel, 1999, The efficient use of conditioning information in portfolios, Working paper, University of Washington.
- Foster, Dean P., and Dan B. Nelson, 1996, Continuous record asymptotics for rolling sample variance estimators, *Econometrica* 64, 139–174.
- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–30.
- Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.
- Graham, John R., and Campbell R. Harvey, 1996, Market timing ability and volatility implied by investment newsletters' asset allocation recommendations, *Journal of Financial Economics* 42, 397–421.
- Kandel, Shmuel, and Robert F. Stambaugh, 1996, On the predictability of stock returns: An asset-allocation perspective, *Journal of Finance* 51, 385–424.
- Merton, Robert C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.
- Nelson, Daniel B., 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347–370.
- Officer, Robert R., 1973, The variability of the market factor of the New York Stock Exchange, *Journal of Business* 46, 434–453.
- Palm, Franz C., 1996, GARCH models of volatility, in G.S. Maddala and C.R. Rao, eds.: *Handbook of Statistics*, Vol. 14. (North-Holland, Amsterdam).
- Tauchen, George E., and Mark Pitts, 1983, The price variability-volume relationship on speculative markets, *Econometrica* 51, 485–505.
- West, Kenneth D., Hali J. Edison, and Don Cho, 1993, A utility-based comparison of some models of exchange rate volatility, *Journal of International Economics* 35, 23–45.