1 Does Volatility Timing Matter?

Jeff Fleming, Chris Kirby, and Barbara Ostdiek

This chapter examines the performance of volatility-timing strategies. We consider a short-horizon investor who uses mean-variance optimization to allocate funds between stocks, bonds, gold, and cash. Specifically, the investor rebalances his portfolio daily based on the current estimate of the conditional covariance matrix of returns. Our results indicate that volatility timing can yield substantial benefits. Moreover, the benefits are robust to practical considerations such as estimation risk and transaction costs.

A great deal of evidence suggests that volatility in financial markets is predictable.[1] This is reflected in products such as Barra's *Short Term Risk Model* and J.P. Morgan's *RiskMetrics* that promise to use volatility modeling to enhance the performance of standard portfolio optimization and risk management techniques. In addition, Busse (1998) finds that many portfolio managers behave like volatility timers, reducing their market exposure during periods of high expected volatility. Despite this anecdotal evidence that volatility timing matters, researchers have yet to establish whether these strategies yield any real economic benefits. We examine this issue by measuring the value of volatility forecasts to investors who engage in short-horizon asset allocation strategies.

We consider a short-horizon investor who uses mean-variance optimization to allocate funds between stocks, bonds, gold, and cash. The investor's objective is to maximize expected return (or minimize volatility) while matching the volatility (or expected return) of a fixed-weight benchmark portfolio. To solve the portfolio problem, we need inputs for the conditional expected asset returns and the conditional covariance matrix. There is little evidence, however, that we can detect short-term variation in expected returns. Therefore, we treat the expected returns as constant, and let the variation in the optimal portfolio weights be driven purely by changes in the conditional covariance matrix.

To estimate this matrix, we use a general nonparametric approach developed by Foster and Nelson (1996). The estimator is a rolling weighted average of the squares and cross products of past return innovations, constructed to be asymptotically optimal. This approach accommodates a variety of return generating processes and nests most ARCH, GARCH, and stochastic volatility models as special cases. We use the time-series of these covariance matrix estimates to construct the investor's optimal portfolio weights. This yields a dynamic trading strategy that specifies the proportion of funds invested in each asset class as a function of time. The

1. See Bollerslev, Chou, and Kroner (1992) for a review of this literature.

performance differential between the dynamic strategy and an appropriate fixed-weight benchmark portfolio reveals the economic value of the volatility forecasts.

Our analysis is based on a 15-year sample of daily returns for stock, bond, and gold futures.[2] The results indicate that volatility timing can have substantial benefits. Ignoring the uncertainty in estimating expected returns, the dynamic strategy that minimizes conditional volatility reduces the portfolio's realized volatility from 11.3% to 8.8% relative to a traditional benchmark portfolio (a 60/40 split between stocks and bonds for 95% of the portfolio and 5% invested in gold). The maximum expected return strategy increases the realized return from 8.1% to 11.6%. The benefits are smaller, but still positive, after we control for estimation risk and the uncertainty in selecting an appropriate benchmark portfolio. Moreover, although these benefits are achieved by active trading, they are still apparent after we account for transaction costs.

1.1 Data

Our data consist of daily returns for stock, bond, and gold futures for January 3, 1983 to December 31, 1997. 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. The gold futures contract closes at 1:30 CST each day while the bond and stock contracts close at 2:00 CST and 3:15 CST, respectively. To align the price observations across contracts, we use daily closing prices for gold futures and the last transaction prices before 1:30 CST for the bond and stock contracts. 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's intraday transactions data. To maintain a uniform measurement interval across contracts, we exclude all days when any of the three markets is closed.

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 a new contract, timing 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

2. We use futures data to avoid problems induced by infrequent trading, but our analysis generalizes to the underlying spot assets via the no-arbitrage relation between futures and spot prices.

3,763 daily returns for each market. The summary statistics indicate that the average return is highest for stock index futures (10.82%) followed by bonds (6.53%) and then gold (−7.76%). Stocks also have a greater volatility (16.2%) than bonds (10.5%) and gold (14.7%). Finally, the correlation between stock and bond returns is positive (0.397), while the correlations between stock and gold returns (−0.105) and bond and gold returns (−0.157) are negative.

1.2 Conditional covariance matrix estimation

Numerous techniques for estimating conditional covariance matrices have been developed in the literature.[3] We rely on a simple nonparametric approach that uses rolling estimators constructed in an asymptotically optimal manner. This approach nests a broad range of ARCH and GARCH models as special cases and has some distinct advantages in our application. Unlike multivariate ARCH and GARCH models, the computational demands of rolling estimators are modest and it is easy to ensure that the covariance matrix estimate is invertible. In addition, the general nature of the approach allows us to provide baseline evidence — without searching for the "best" volatility model — on the economic significance of volatility timing.

To develop the rolling estimator, let $\boldsymbol{r}_{t+1}$, $\boldsymbol{\mu}_t \equiv \mathrm{E}_t[\boldsymbol{r}_{t+1}]$, and $\boldsymbol{\Sigma}_t \equiv \mathrm{E}_t[(\boldsymbol{r}_{t+1} - \boldsymbol{\mu}_t)(\boldsymbol{r}_{t+1} - \boldsymbol{\mu}_t)']$ denote, respectively, a 3×1 vector of returns on stock, bond, and gold futures, the conditional expected value of $\boldsymbol{r}_{t+1}$, and the conditional covariance matrix of $\boldsymbol{r}_{t+1}$. The rolling estimator can be written as

$$\hat{\boldsymbol{\Sigma}}_t = \sum_{l=-t+1}^{T-t} \omega_{t+l}(\boldsymbol{r}_{t+l} - \boldsymbol{\mu}_{t+l})(\boldsymbol{r}_{t+l} - \boldsymbol{\mu}_{t+l})' \tag{1.1}$$

where ω_{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. Foster and Nelson (1996) demonstrate that the optimal weights depend on the characteristics of the process. If volatility is stochastic, the optimal weights for the two-sided rolling estimator are given by

$$\omega_{t+l} = (\alpha/2)e^{-\alpha|l|}, \tag{1.2}$$

where α is the decay rate. This is called a two-sided estimator because it uses both

3. Officer (1973) and Fama and MacBeth (1973) employ *ad hoc* rolling estimators. Merton (1980), Poterba and Summers (1986), and French, Schwert, and Stambaugh (1987) divide the data into nonoverlapping blocks and treat the variances and covariances as constant within each block. More recently, ARCH models [e.g., Engle (1982) and Bollerslev (1986)] have gained popularity.

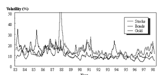

Figure 1.1
Daily volatility estimates for S&P 500, T-bond, and gold futures returns.

leads and lags of returns to estimate $\mathbf{\Sigma}_t$. To construct the corresponding one-sided estimator, we set $\omega_{t+l} = 0$ for $l > 0$ and double each of the weights for $l \leq 0$.

We estimate the optimal decay rate using the procedure developed by Fleming, Kirby, and Ostdiek (1999). This yields $\alpha = 0.0679$. The resulting two-sided estimates of the conditional volatilities are plotted in Figure 1.1.[4] The average estimates are consistent with the unconditional volatilities reported earlier with stocks the most volatile, followed by gold and then bonds. Moreover, the variability in these estimates suggests that volatility changes over time.

Figure 1.2 shows the two-sided estimates of the conditional return correlations. As with the volatilities, the average estimates are generally consistent with the

---

4. As is common practice, we use seasonally-adjusted data to generate these estimates. We begin by regressing the raw returns on a set of six variables: a dummy variable for each weekday and a variable (NTDYS) that counts the number of nontrading days covered by each return. The residual in this regression, r^*_{it}, is the unexpected component of the day t return in market i. To remove the daily seasonal in volatility, we estimate the regression,

$$\left[(r^*_{it})^2/var(r^*_{it})\right] - 1 = \beta_{1i}\text{MON}_t + \beta_{2i}\text{TUE}_t + \cdots + \beta_{5i}\text{FRI}_t + \beta_{6i}\text{NTDYS}_t + e_t$$

where $\text{var}(r^*_{it})$ is the sample variance of r^*_{it}. The resulting coefficient estimates are used to construct the seasonally-adjusted returns,

$$r_{it} = r^*_{it}/\sqrt{1 + \beta_{1i}\text{MON}_t + \beta_{2i}\text{TUE}_t + \cdots + \beta_{5i}\text{FRI}_t + \beta_{6i}\text{NTDYS}_t},$$

that we use to construct the covariance matrix estimates.

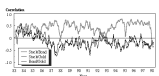

Figure 1.2
Daily correlation estimates between S&P 500, T-bond, and gold futures returns.

unconditional correlations and there is considerable evidence that the correlations change over time. In particular, the estimated stock/gold and bond/gold correlations sharply decrease after 1986 and the stock/bond estimates widely fluctuate throughout the sample.

1.3 Volatility timing in a mean-variance framework

To determine whether the variation in the conditional covariance matrix can be exploited to improve asset allocation decisions, we consider a hypothetical investor who uses conditional mean-variance analysis to allocate funds.[5] The investor wants to minimize his portfolio's conditional variance subject to achieving a particular conditional expected rate of return. For each date t, he solves the quadratic program

$$\begin{aligned} \min_{\mathbf{w}_t} \quad & \mathbf{w}_t' \boldsymbol{\Sigma}_t \mathbf{w}_t \\ \text{s.t.} \quad & \mathbf{w}_t' \boldsymbol{\mu}_t = \mu_{pt}, \end{aligned} \tag{1.3}$$

5. Sufficient conditions for an investor to demand a conditionally mean-variance efficient portfolio are (i) returns are conditionally multivariate normal, which is consistent with an important class of trading models [see, e.g., Clark (1973) and Tauchen and Pitts (1983)]; and (ii) the investor has lognormal utility of wealth.

where $\mathbf{w}_t$ is a 3×1 vector of portfolio weights on stock, bond, and gold futures. Note that we omit the riskless interest rate from this specification because we are dealing with futures returns. Under the cost-of-carry model, the return on a futures contract equals the total return on the spot asset minus the riskless rate. $1 - \mathbf{w}_t'\mathbf{1}$ represents the weight held in "cash equivalent" securities which earn an excess return equal to zero. The solution to our optimization problem,

$$\mathbf{w}_{pt} = \frac{\mu_{pt}\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t}{\boldsymbol{\mu}_t'\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t}, \tag{1.4}$$

delivers the weights on the risky assets.[6]

To construct the optimal dynamic portfolio weights, we use the one-step-ahead forecasts of the conditional covariance matrix from Section 1.2. We also, in general, need one-step-ahead forecasts of the expected returns. However, there is little evidence to suggest that we can detect daily variation in expected returns. Therefore, we treat the expected returns as constant and let the trading decisions depend only on changes in our estimates of the conditional covariance matrix.

We could conduct a similar analysis where the investor's objective is to maximize conditional expected return subject to achieving a particular conditional variance. Therefore, our optimal portfolio analysis suggests two candidate volatility-timing strategies. First, we solve for the weights that set the expected return equal to some fixed target and minimize volatility (the "minimum volatility strategy"). Second, we solve for the weights that set volatility equal to some fixed target and maximize expected return (the "maximum expected return strategy").

Implementing the dynamic strategies To implement our methodology, two remaining issues must be resolved. The first is that we must estimate the expected returns for the assets. One possible approach is to use data available at the beginning of our sample period. Unfortunately, this produces unreliable estimates for bonds and gold due to dramatic changes in the 1970s caused by the shift in Federal Reserve interest rate policy and the elimination of the gold standard. Alternatively, we could use the first few years of our sample to estimate the expected returns and use the remaining period to evaluate the dynamic strategies. This is also problematic, however, because the expected return estimates would be very imprecise.

We address this issue by adopting a bootstrap approach [Efron (1979)]. First, we randomly sample with replacement from the actual data to generate a series of

6. Fleming, Kirby, and Ostdiek (1999) show that under standard no-arbitrage arguments we would obtain the same weights by formulating the optimization in terms of spot assets.

artificial returns. We then 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 the weights to the actual returns and evaluate the performance of the dynamic strategies. This approach ensures that our analysis is based on a representative sample and allows us to directly assess the impact of estimation risk.

The second unresolved issue is how to measure the performance gains attributable to volatility timing. The most straightforward approach is to compare the performance of the dynamic strategies to that of a fixed-weight benchmark portfolio. But then we must choose an appropriate benchmark. To control for the uncertainty regarding this choice, we adopt the following approach. Initially, we use a traditional benchmark portfolio, defined as a 57/38/5 split between stocks, bonds, and gold. This reflects the general guideline among asset allocation managers of a 60/40 split between stocks and bonds, with five percent of the portfolio invested in gold to provide a hedge against inflation. After providing this evidence, we evaluate the performance of the dynamic strategies relative to the ex-post efficient frontier.

Results of the optimal portfolio analysis We now examine the optimization results for the limiting case in which the true expected returns are known. These results are obtained by solving for the optimal portfolio weights in Equation (1.4) where the inputs are (i) the sample mean returns for each asset, and (ii) our estimates of the conditional covariance matrix. We refer to this as the case of no estimation risk.[7]

Figure 1.3 plots the optimal weights for the minimum volatility strategy. The sign and magnitude of each weight depends on the estimated expected returns and the forecasted volatilities and correlations. For example, the average return on gold futures is negative in our sample and, in general, the weight in gold is negative. But the size of this short position decreases when the forecasted gold volatility increases, as in 1985 and 1993 (Figure 1.1), and also when gold's forecasted correlation with stocks and bonds becomes more negative, as in early- to mid-1987 and the beginning of 1991 (Figure 1.2). Similarly, the split between stocks and bonds is sensitive to their relative volatilities. Stock volatility decreases steadily from 1991 to 1994 while bond volatility remains relatively constant. As a result, the weight in stocks steadily increases over this period while the weight in bonds decreases. The opposite occurs from 1996 to 1998 as stock volatility rises and bond volatility falls.

---

7. We use this terminology loosely because we still face the risk associated with estimating the conditional covariance matrix. However, Merton (1980) implies that this risk is likely small in comparison to that associated with estimating expected returns.

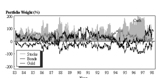

Figure 1.3
Optimal dynamic portfolio weights with minimum volatility.

Figure 1.4 shows the anticipated reduction in volatility according to our optimization results. The unconditional return on the benchmark portfolio is 8.26%, but its forecasted volatility changes as a function of our daily covariance matrix. The average volatilities for the benchmark and dynamic portfolios are, respectively, 10.6% and 7.7%. Not surprisingly, the anticipated volatility reduction is greatest during periods of high stock market volatility such as 1986, the 1987 crash, the 1989 mini-crash, and near the end of the sample. During these periods, the minimum volatility strategy has a much lower weight in stocks than the benchmark portfolio.

Turning to the maximum expected return strategy, the optimal weights (not shown) are similar to those for the minimum volatility strategy. The main difference is that the changes tend to be more pronounced because a greater risk exposure is generally needed to match the benchmark's volatility. Figure 1.5 shows the anticipated improvement using this strategy, measured with the estimated expected returns. The strategy's average expected return is 13.42%, a gain of more than five percentage points over the benchmark portfolio (8.26%), although gains of more than ten points are apparent during much of the second half of the sample.

1.4 Evaluating the performance of the dynamic strategies

Our portfolio optimization results suggest that the anticipated benefits to volatility timing are substantial. In practice, however, the realized benefits may be smaller

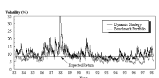

Figure 1.4
Ex-ante comparison of the minimum volatility dynamic strategy to the static benchmark portfolio.

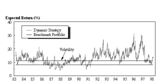

Figure 1.5
Ex-ante comparison of the maximum return dynamic strategy to the static benchmark portfolio.

Table 1.1
Ex-post performance of the static and dynamic asset allocation strategies.

	Static: Benchmark Portfolio			Dynamic: Minimum Volatility			Dynamic: Maximum Return		
Period	μ	σ	SR	μ	σ	SR	μ	σ	SR
Entire Sample	8.13	11.30	0.719	8.01	8.84	0.906	11.62	12.65	0.919
1983−1985	6.35	10.03	0.633	6.77	7.46	0.907	10.47	12.38	0.846
1986−1988	7.64	16.73	0.457	8.46	12.59	0.672	6.54	13.27	0.493
1989−1991	8.15	10.05	0.811	8.31	9.33	0.891	9.69	12.85	0.754
1992−1994	2.77	7.71	0.358	3.37	6.33	0.533	5.07	12.07	0.420
1995−1997	15.67	9.70	1.615	13.09	6.86	1.909	26.50	12.61	2.102
Excluding:									
Oct 19-30, 1987	8.12	10.43	0.778	7.49	8.66	0.865	11.15	12.56	0.888
1986−1988	8.25	9.41	0.877	7.89	7.59	1.040	12.93	12.49	1.035

than anticipated (Figures 1.4 and 1.5) because the portfolio weights are based on estimates of the expected returns and conditional covariance matrix. To assess the realized benefits, we need to form the portfolios implied by the dynamic weights and compute their realized returns and volatilities.

We compute the ex-post return for each strategy by multiplying the portfolio weights on a given day by the observed next-day returns on stock, bond, and gold futures. Table 1.1 summarizes the results. For the minimum volatility strategy, the average return is comparable to that of the benchmark portfolio, but its sample volatility (8.8%) is considerably lower than the benchmark's (11.3%). As a result, the Sharpe ratio for this strategy (0.91) is substantially higher than the Sharpe ratio for the benchmark (0.72), indicating greater ex-post efficiency. The maximum expected return strategy achieves similar results. Ex-post, the strategy has a slightly higher sample volatility (12.7%) than the benchmark, but a substantially greater average return (11.6%). Therefore, its Sharpe ratio (0.92) also indicates improvement.

Table 1.1 also breaks down the ex-post returns by three-year subperiods. These subperiod results indicate that our general findings are robust across the sample. Although there is substantial variation in the average returns and sample volatilities, both of the dynamic strategies outperform the benchmark portfolio in every three-year subperiod except 1989−1991. Even in this subperiod, the minimum volatility strategy outperforms the benchmark and the maximum expected return strategy achieves a Sharpe ratio (0.75) that is comparable to that of the benchmark (0.81).

The final two lines of Table 1.1 show the impact of the 1987 crash. Excluding the crash has a greater impact on the benchmark than on either of the dynamic strategies. At the time of the crash, the conditional volatilities for both stocks and bonds are relatively high (Figure 1.1). Therefore, the dynamic strategies are predominately invested in cash, with relatively low weights in stocks and bonds (Figure 1.3), causing them to outperform the benchmark. It is unclear, however, how much we should credit this to "timing" ability. The large cash positions persist before and after the crash, making it difficult to evaluate the strategies over a noncrash sample. If we exclude the period when the strategies did well because they were holding cash, we should also exclude the surrounding period when, for the same reason, they did poorly. One alternative is to simply exclude the entire 1986−1988 subperiod. As shown in Table 1.1, the relative performance of the static and dynamic strategies over this sample is similar to that over the full sample.

Significance tests We assess the statistical significance of the performance gains for the dynamic strategies by comparing the results in Table 1.1 to those obtained when the asset returns are generated independently of the portfolio weights. This is accomplished using a simple randomization scheme. First, we form a permutation of the actual data series by randomly sampling without replacement from the joint distribution of returns. Next, we apply the actual weights for the dynamic strategies to the randomized returns to get a time-series of daily portfolio returns.[8] We repeat this process 10,000 times. Note that the mean return and volatility of the benchmark portfolio are the same as in the actual data because we sample without replacement.

If the performance gains of the dynamic strategies are significant, then these strategies should perform substantially worse in the artificial samples than they do using the actual data. We find this to be the case. For the minimum volatility strategy, the mean return for the 10,000 trials is 8.0% with a standard error of 1.3%. The mean volatility is 10.6% with a standard error of 0.3%. None of the 10,000 trials produce an ex-post volatility as low as that observed using the actual return series (8.8%), and only 11.2% of the trials yield a higher Sharpe ratio. For the maximum expected return strategy, the mean return is 13.0% with a standard error of 2.6%, and the mean volatility is 18.3% with a standard error of 0.6%. Only 7.3% of the trials produce a Sharpe ratio that is higher than that observed for the actual returns. These findings indicate it is unlikely that the superior performance of the dynamic strategies is due to chance.

8. This procedure is asymptotically equivalent to using the actual returns series and randomizing the portfolio weights. Either way, the portfolio weights are independent of the realized returns on the assets.

The impact of estimation risk We evaluate the impact of estimation risk on our results using the bootstrap approach described earlier. First, we randomly sample with replacement from the actual returns to generate a series of 3,763 artificial returns. We then compute the average return for each asset in this artificial sample and use these averages (instead of the true sample means) to determine the optimal portfolio weights. Finally, we apply these weights to the actual returns and evaluate the performance of the dynamic strategies as before. Relying on the artificial sample to estimate the expected returns mimics the uncertainty about expected returns that an investor would face in practice.[9] Therefore, if the dynamic strategies still outperform the benchmark portfolio, we can conclude that estimation risk does not offset their superior performance.

In this set of experiments, the mean return for the minimum volatility strategy is 7.2% and the mean volatility is 8.5%, with an average Sharpe ratio of 0.84. This is lower than the value obtained using the true sample means (0.91), but higher than that for the benchmark (0.72). In 94.4% of the trials, this strategy produces a higher Sharpe ratio than the benchmark portfolio. This indicates that the benefits of the dynamic strategy are apparent even after we account for the uncertainty about expected returns. The results for the maximum expected return strategy are similar — the mean return is 10.8% and the mean volatility is 12.6%. The average Sharpe ratio for this strategy is 0.85 and 96.6% of the trials produce a higher Sharpe ratio than the benchmark portfolio. Based on these results, it does not appear that the superior performance of the dynamic strategies is eliminated when we account for estimation risk.

1.5 Robustness Tests

The evidence reported in Section 1.4 is consistent with the view that volatility timing can improve the performance of asset allocation decisions. However, our analysis to this point assumes a specific benchmark portfolio, a given target expected return or volatility, no transaction costs, and infinitely divisible contract sizes. In this section, we assess the impact of these issues on our results.

---

9. Most investors would also 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 returns should be highest for stocks, 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 these conditions.

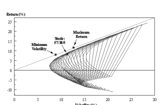

Figure 1.6
Ex-post performance of the static portfolio using different combinations of weights.

Alternative benchmark portfolios To control for the uncertainty about the appropriate benchmark portfolio, we consider *all* possible combinations of static weights in stocks, bonds, and gold. For each combination, we compute the realized return and volatility during our sample. This set of benchmarks, plotted in Figure 1.6, defines a region in expected return/standard deviation space whose boundary is the ex-post minimum variance frontier. To allow a positive weight in cash, we would draw a line from the origin through the tangency portfolio. The slope of this line represents the highest Sharpe ratio we could have attained using the most efficient (ex-post) set of static portfolio weights. Note that the Sharpe ratios for the two dynamic strategies are even greater.

The implication of this finding is clear. It is unlikely that we would have chosen ex-ante the portfolio that turns out to be most efficient; but, even if we had, the dynamic strategies still would have outperformed it. This indicates that our earlier findings are robust to alternative benchmarks.

Varying the target expected returns and volatilities Table 1.2 shows how the performance of the dynamic strategies varies with the target expected return or volatility used in solving our daily portfolio optimizations. For comparison, the mid-

Table 1.2
Ex-post performance of the dynamic strategies under different target returns and volatilities.

Minimum Volatility Strategy				Maximum Return Strategy			
Target	Realized returns			Target	Realized returns		
Return	μ	σ	SR	Volatility	μ	σ	SR
6.00	5.77	6.43	0.898	9.00	9.31	10.10	0.921
6.50	6.27	6.96	0.900	9.50	9.84	10.66	0.923
7.00	6.77	7.50	0.902	10.00	10.35	11.21	0.923
7.50	7.26	8.03	0.904	10.50	10.85	11.76	0.923
8.00	7.75	8.56	0.905	11.00	11.33	12.31	0.921
8.26	8.01	8.84	0.906	11.31	11.62	12.65	0.919
8.50	8.24	9.09	0.906	11.50	11.80	12.86	0.918
9.00	8.74	9.62	0.908	12.00	12.27	13.41	0.915
9.50	9.23	10.15	0.909	12.50	12.75	13.95	0.914
10.00	9.72	10.68	0.910	13.00	13.22	14.50	0.912
10.50	10.21	11.21	0.910	13.50	13.69	15.05	0.910

dle line of the table contains the results from Section 1.4 where the target is based on the benchmark portfolio's mean return or volatility. As expected, lowering the target return or volatility reduces both the realized return and volatility, however, the effects on the Sharpe ratios are ambiguous. Moreover, across the entire range of target parameter values, there is little variation in the realized performance. This indicates that the dynamic strategies are relatively insensitive to variation in the target return or target volatility.

Transaction costs We evaluate the effect of transaction costs by first estimating the transaction costs for S&P 500 futures. We estimate the bid/ask spread using the Smith and Whaley (1994) approach based on intraday transaction prices. This yields an average spread of 0.0593, or \$29.65 per contract.[10] In addition, roundtrip commissions and fees for large institutions are about \$6.00 per contract.[11] Combining these estimates with the average index level during our sample period (384.51) indicates that the average one-way transaction cost is \$17.825 on a contract size of \$384.51×500, or 2.34% annualized. In other words, if we traded one contract every day for a year (buy or sell), transaction costs would reduce the realized return by 234 basis points. We assume that the transaction costs for T-bond and gold futures are comparable to those for S&P 500 futures.

10. This slightly exceeds the estimate of 0.0508 reported in Smith and Whaley (1994) because their sample ends in 1987 while ours extends through 1997.

11. This estimate is the same as that used in Fleming, Ostdiek, and Whaley (1996).

Table 1.3
Performance of the static and dynamic asset allocation strategies after imposing transaction costs.

	Static: Benchmark Portfolio			Dynamic: Minimum Volatility			Dynamic: Maximum Return		
Pct. Costs	μ	σ	SR	μ	σ	SR	μ	σ	SR
0.0%	8.13	11.30	0.719	8.01	8.84	0.906	11.62	12.65	0.919
0.5%	8.11	11.30	0.717	7.95	8.84	0.899	11.52	12.65	0.911
1.0%	8.09	11.30	0.716	7.89	8.84	0.892	11.42	12.65	0.903
1.5%	8.07	11.30	0.714	7.83	8.84	0.886	11.32	12.65	0.895
2.0%	8.05	11.30	0.712	7.77	8.84	0.879	11.22	12.65	0.887
2.5%	8.03	11.30	0.711	7.71	8.84	0.872	11.12	12.65	0.879
3.0%	8.01	11.30	0.709	7.65	8.84	0.866	11.02	12.65	0.871
3.5%	7.99	11.30	0.707	7.59	8.84	0.859	10.92	12.65	0.863
4.0%	7.97	11.30	0.706	7.53	8.84	0.852	10.81	12.65	0.855
4.5%	7.95	11.30	0.704	7.47	8.84	0.846	10.71	12.65	0.847
5.0%	7.93	11.30	0.702	7.41	8.84	0.839	10.61	12.65	0.839

Table 1.3 shows the effect of various levels of transaction costs, centered around our estimate of 2.34%. We impose the transaction costs on every trade, including those required for establishing the initial position, daily rebalancing, rolling into each subsequent contract month, and liquidating the position at the end of the sample. For the dynamic strategies, the daily rebalancing represents the trading required to track time-variation in the optimal portfolio weights and, for the benchmark portfolio, it represents the trading needed to maintain a constant weight in each asset. The transaction costs are imposed each day by subtracting the percentage cost from that day's realized return.

As the table indicates, transaction costs have the greatest effect on the maximum expected return strategy. This is not surprising because, as noted earlier, the weights for this strategy exhibit the most time-variation. Imposing transaction costs of, say, 2.5% reduces the strategy's realized return from 11.6% to 11.1% with just a trivial effect on its volatility. As a result, the strategy still earns a greater Sharpe ratio (0.88) than the benchmark portfolio (0.71). In order to equate the Sharpe ratios for the two strategies, transaction costs would need to be 15.9% (not shown in the table) — an amount that is almost seven times our previous estimate.

A related issue is whether the strategies require so much trading that they impact market prices. To examine this, we need to consider a specific value for the underlying portfolio. Suppose we begin with $100 million, which by the end of our sample grows to over $300 million. The contract sizes for the stock, bond, and gold futures contracts are, respectively, 500, 1000, and 100 times price.[12] Given these parameters, the maximum return strategy requires, on average, daily trade sizes of 23, 67, and 116 contracts (in absolute value), respectively, for stock, bond, and gold futures. For comparison, the average daily volumes in these markets during our sample period are 59,000 for stocks, 267,000 for bonds, and 30,000 for gold. Based on this evidence, it seems unlikely that the dynamic strategies entail market-impact costs large enough to materially affect their performance.

Trading discrete quantities of contracts Our analysis to this point assumes that we can trade fractional contracts, i.e., if we have $100 to invest, and the optimal weight is 0.50 for a contract worth $100, we can buy only half a contract. If instead we require each trade to be a discrete number of contracts, then we incur rounding error in the sense that our portfolio weight is suboptimal. This error has the greatest effect when the portfolio value (i.e., the number of contracts traded) is small.

We examine the impact of this rounding error using the following procedure. For a given level of initial wealth, we compute the number of contracts implied by our optimal weights and we round to the nearest integer. Any residual funds created by the rounding procedure are held in cash.[13] After each day, we compute our new portfolio value, and we apply the same rounding procedure to the optimal weights for the following day to determine the new number of contracts. Table 1.4 reports the results.

Imposing discrete trade sizes has an unpredictable effect on the mean returns, but generally increases volatility. As expected, the effects are greatest for smaller initial portfolio values. Beginning with $100,000, for example, the volatilities for each trading strategy sharply increase relative to those using continuous trade sizes. The mean returns are higher for the benchmark and minimum volatility strategy but lower for the maximum return strategy. These effects cause the Sharpe ratios to increase for the benchmark portfolio and decrease for the dynamic strategies. This result seems random, however, when compared to the results using other levels of initial wealth. In any case, the effect of discrete trade sizes disappears quickly. For an initial wealth as low as $5,000,000, the performance of each of the strategies is comparable to that using continuous trade sizes. This suggests that rounding error does not have much of an effect for even moderately-sized portfolios.

12. The contract size for S&P 500 futures changed to 250 times price after October 31, 1997.

13. A more precise procedure for determining the number of contracts would be to include the quantity discreteness in our daily portfolio optimization. This approach, however, is more complex and seems unnecessary given the small effects we report below using the simpler approach.

Table 1.4
The effect of trading discrete contract sizes on the ex-post performance of the static and dynamic asset allocation strategies.

	Static: Benchmark Portfolio			Dynamic: Minimum Volatility			Dynamic: Maximum Return		
Initial Wealth	μ	σ	SR	μ	σ	SR	μ	σ	SR
$ 100,000	12.46	17.30	0.720	8.21	9.18	0.894	11.29	12.84	0.880
500,000	7.94	11.58	0.686	8.46	8.95	0.945	11.94	12.71	0.939
1,000,000	8.18	11.17	0.733	8.01	8.88	0.903	11.66	12.65	0.922
5,000,000	8.12	11.29	0.719	8.03	8.84	0.908	11.67	12.65	0.922
10,000,000	8.12	11.30	0.718	8.02	8.84	0.907	11.60	12.65	0.917
50,000,000	8.13	11.30	0.719	8.01	8.84	0.906	11.62	12.65	0.919
100,000,000	8.12	11.30	0.719	8.01	8.84	0.906	11.63	12.65	0.919
500,000,000	8.13	11.30	0.719	8.01	8.84	0.906	11.62	12.65	0.919

1.6 Conclusions

Our results indicate that volatility timing can improve the performance of short-horizon investment strategies. In particular, an investor trading stock, bond, and gold futures can use daily forecasts of the conditional covariance matrix to form a dynamic trading strategy that outperforms any fixed-weight benchmark portfolio. For the limiting case with no estimation risk, the dynamic strategies that minimize volatility and maximize return realize Sharpe ratios of 0.91 and 0.92, respectively. By comparison, the Sharpe ratio for a traditional benchmark portfolio is 0.72. The benefits of volatility timing are smaller, but still positive, after we account for estimation risk and the optimal choice of benchmark portfolio.

Additional tests indicate that the performance differential is robust to several practical considerations. First, accounting for the transaction costs of active trading and for discrete contract sizes does not eliminate the advantage of the dynamic strategies. Second, the results are insensitive to the target level of expected return or volatility. Finally, our results may be conservative in that they rely on a simple volatility specification rather than on a more complex parametric model.

References

Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–328.

Bollerslev, T., R. Y. Chou, and K. F. Kroner. 1992. ARCH modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics*, 52:5–60.

Busse, J. 1998. Volatility timing in mutual funds: Evidence from daily returns. Unpublished (Stern School of Business, New York University).

Clark, P. K. 1973. A subordinated stochastic process model with finite variance for speculative prices. *Econometrica*, 41:135–156.

Efron, B. 1979. Bootstrap methods: Another look at the jacknife. *Annals of Statistics*, 7:1–26.

Engle, R. F. 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50:987–1008.

Fama, E. F. and J. D. MacBeth. 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81:607–636.

Fleming, J., C. Kirby, and B. Ostdiek. 1999. The economic value of volatility timing. Unpublished (Jones Graduate School of Management, Rice University).

Fleming, J., B. Ostdiek, and R. E. Whaley. 1996. Trading costs and the relative rates of price discovery in stock, futures, and option markets. *Journal of Futures Markets*, 16:353–387.

Foster, D. P. and D. B. Nelson. 1996. Continuous record asymptotics for rolling sample variance estimators. *Econometrica*, 64:139–174.

French, K. R., G. W. Schwert, and R. F. Stambaugh. 1987. Expected stock returns and volatility. *Journal of Financial Economics*, 19:3–30.

Merton, R. C. 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics*, 8:323–361.

Officer, R. R. 1973. The variability of the market factor of the New York Stock Exchange. *Journal of Business*, 46:434–453.

Poterba, J. M. and L. H. Summers. 1986. The persistence of volatility and stock market fluctuations. *American Economic Review*, 76:1142–1151.

Smith, T. and R. E. Whaley. 1994. Estimating the effective bid/ask spread using time and sales data. *Journal of Futures Markets*, 14:437–455.

Tauchen, G. E. and M. Pitts. 1983. The price variability-volume relationship on speculative markets. *Econometrica*, 51:485–505.