JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 47, No. 2, Apr. 2012, pp. 437–467

doi:10.1017/S0022109012000117

It's All in the Timing: Simple Active Portfolio Strategies that Outperform Naïve Diversification

Chris Kirby and Barbara Ostdiek*

Abstract

DeMiguel, Garlappi, and Uppal (2009) report that naïve diversification dominates mean-variance optimization in out-of-sample asset allocation tests. Our analysis suggests that this is largely due to their research design, which focuses on portfolios that are subject to high estimation risk and extreme turnover. We find that mean-variance optimization often outperforms naïve diversification, but turnover can erode its advantage in the presence of transaction costs. To address this issue, we develop 2 new methods of mean-variance portfolio selection (volatility timing and reward-to-risk timing) that deliver portfolios characterized by low turnover. These timing strategies outperform naïve diversification even in the presence of high transaction costs.

I. Introduction

Mean-variance optimization is a cornerstone of modern portfolio theory. However, a recent study by DeMiguel, Garlappi, and Uppal (2009) questions the value of mean-variance optimization relative to naïve diversification (i.e., relative to a strategy that places a weight of $1/N$ on each of the N assets under consideration). The authors of the study implement 14 variants of the standard mean-variance model for a number of data sets and find that "there is no single model that consistently delivers a Sharpe ratio or a CEQ return that is higher than that of the $1/N$ portfolio." This finding presents researchers with 2 clear challenges. The first is to understand why the mean-variance approach to portfolio selection performs so poorly in the DeMiguel et al. study. The second is to develop more effective procedures for using sample information about means and variances in portfolio problems.

*Kirby, christopher.kirby@uncc.edu, Belk College of Business, University of North Carolina at Charlotte, 9201 University City Blvd., Charlotte, NC 28223; and Ostdiek, ostdiek@rice.edu, Jones Graduate School of Business, Rice University, PO Box 2932, Houston, TX 77252. We are thankful for comments from Nick Bollen, Stephen Brown (the editor), Jun Tu (the referee), and seminar participants at Rice University, Universidad Carlos III de Madrid, Georgia State University, the University of North Carolina at Charlotte, and the 2010 Conference on Financial Economics and Accounting.

With respect to the 1st challenge, we show that the DeMiguel et al. (2009) research design places the mean-variance model at an inherent disadvantage relative to naïve diversification. Specifically, it delivers mean-variance efficient (MVE) portfolios that are very aggressive, with target conditional expected excess returns that often exceed 100% per year. Targeting conditional expected excess returns of this magnitude leads to poor out-of-sample performance because it greatly magnifies both estimation risk and portfolio turnover. If the mean-variance model is implemented by targeting the conditional expected return of the $1/N$ portfolio, the resulting MVE portfolios outperform naïve diversification for most of the DeMiguel et al. data sets. However, it is not clear that this finding is robust to the presence of high transaction costs.

In response to the 2nd challenge, we develop simple active portfolio strategies that retain the most appealing features of the $1/N$ strategy (no optimization, no covariance matrix inversion, and no short sales) while exploiting sample information about the reward and risk characteristics of the assets under consideration. In particular, we specify the portfolio weights in terms of conditional expected returns, conditional return volatilities, and a tuning parameter that allows some control over portfolio turnover. The empirical evidence shows that the proposed strategies outperform naïve diversification by statistically and economically significant margins. This is true even though we implement the strategies using estimators of the conditional expected returns and conditional return volatilities that are likely to be relatively inefficient. Moreover, the advantage of the proposed strategies persists even in the presence of high transaction costs.

Although our strategies are most naturally interpreted as mean-variance timing rules, they are rooted in the literature on asset allocation in the presence of estimation error and constraints on portfolio holdings. There have been several notable additions to this literature in recent years. Pastor (2000) and Pastor and Stambaugh (2000) use Bayesian methods to address parameter uncertainty. Jagannathan and Ma (2003) consider short-sale constraints and show that these restrictions improve performance by reducing the effect of estimation error. Kan and Zhou (2007) develop a 3-fund strategy that optimally diversifies across both factor and estimation risk. Tu and Zhou (2010) incorporate economic objectives into Bayesian priors and show that the resulting portfolio rules outperform some of the best rules that have been proposed in the classical framework.

The recent contribution of Tu and Zhou (2011) is particularly interesting. The authors develop a theory of portfolio choice under estimation risk by assuming that asset returns are independently and identically distributed (i.i.d.) over time. Building on the idea that the $1/N$ portfolio is a reasonable shrinkage target, they propose a new strategy that optimally combines the $1/N$ portfolio and the Kan and Zhou (2007) 3-fund portfolio, with the rate of shrinkage toward the $1/N$ portfolio determined by the level of estimation risk. Simulations demonstrate that this "4-fund strategy" outperforms the $1/N$ strategy in an i.i.d. setting under a range of assumptions about the data generating process.

This paper addresses many of the issues studied by Tu and Zhou (2011), but it does so from the perspective of an investor who assumes that the conditional means and variances of asset returns change through time. We start by proposing a new class of active portfolio strategies that are designed to exploit sample information about volatility dynamics in a way that mitigates the effect of estimation risk. Under our approach, which we refer to as volatility timing (VT), the portfolios are rebalanced monthly based solely on changes in the estimated conditional volatilities of asset returns. We control the sensitivity of the portfolio weights to these changes via a tuning parameter that can be interpreted as a measure of timing aggressiveness. This allows us to keep the turnover of the proposed strategies to a level competitive with that of naïve diversification.

We also propose a more general class of timing strategies that incorporates sample information about the dynamics of conditional expected returns. Under our approach, which we refer to as reward-to-risk timing (RRT), the portfolios are rebalanced monthly based solely on changes in estimated reward-to-risk ratios. We implement these strategies using 2 estimators of conditional expected returns: a simple rolling estimator that imposes no parametric assumptions and an estimator that is designed to reduce estimation risk by exploiting the predictions of asset pricing theory.

The empirical section of the paper compares the out-of-sample performance of the VT and RRT strategies to that of the $1/N$ strategy using 4 different data sets. All of the strategies are implemented using the same rolling sample approach used by DeMiguel et al. (2009). We find that both types of timing strategies can significantly outperform the $1/N$ strategy for a range of data sets, and that this performance advantage persists after accounting for the impact of transaction costs. Thus, in contrast to DeMiguel et al., we conclude that there can be substantial value in using sample information to guide portfolio selection decisions.

For example, with proportional transaction costs of 50 basis points (bp), the $1/N$ strategy has an estimated Sharpe ratio of 0.46 for a data set comprising 25 portfolios formed on size and book-to-market characteristics. In comparison, the VT strategies have estimated Sharpe ratios that range from 0.47 to 0.49. These differences translate into significant performance gains because the returns for the VT and the $1/N$ strategies are highly correlated, resulting in small standard errors for relative performance measures. We estimate that risk-averse investors would be willing to pay from 81 bp to 108 bp per year to switch from the $1/N$ strategy to our VT strategies. These estimated fees are statistically different from 0 at the 10% significance level.

The performance gains are more pronounced for the RRT strategies. If we implement these strategies using a simple rolling estimator of conditional expected returns, then the estimated Sharpe ratios range from 0.52 to 0.54. We estimate that risk-averse investors would be willing to pay from 127 bp to 167 bp per year to switch from the $1/N$ strategy to these strategies. RRT appears to be a particularly promising strategy when it is implemented using the estimator of expected returns implied by the Carhart (1997) 4-factor risk model. In this case, the estimated Sharpe ratios range from 0.52 to 0.57 and the estimated performance fees range from 118 bp to 220 bp per year. These gains are all statistically significant at the 1% level.

The evidence also suggests that the performance gains increase with the cross-sectional dispersion in the means and variances of returns. The data set that poses the biggest challenge to the timing strategies contains 10 industry portfolios. Sorting firms into industries produces only a modest spread in the estimated means. In contrast, some of the most compelling results are obtained using a data set created by sorting firms into 10 portfolios based on momentum, which produces a pronounced spread in the estimated means. The $1/N$ strategy has an estimated annualized Sharpe ratio of 0.28 for this data set, while the RRT strategies have estimated annualized Sharpe ratios that range from 0.43 to 0.47 when they are implemented using a simple rolling estimator of conditional expected returns. We observe similar performance gains for the VT strategies using a data set constructed by sorting firms into portfolios based on historical volatility.

The remainder of the paper is organized as follows. Section II considers the portfolio choice problem of an investor with quadratic risk preferences and shows how the resulting framework can be used to motivate VT and RRT strategies. Section III describes our estimators of the conditional mean vector and conditional covariance matrix of excess returns. Section IV discusses our approach to performance evaluation and statistical inference. Section V describes the data and presents the empirical results. Section VI provides concluding remarks.

II. Portfolio Strategies

This section describes the portfolio strategies we investigate. The 1st strategy (naïve diversification) consists of making an equal allocation to each asset. The remaining strategies fall into 1 of 2 basic categories: strategies that are ex ante optimal under quadratic loss, that is, they invest in conditionally MVE portfolios, and strategies that do not entail formal optimization, but invest in portfolios that exploit sample information about conditional means and variances in a manner that mitigates estimation risk. We assume throughout that there are N risky assets and a single risk-free asset, and we refer to naïve diversification across the risky assets as the $1/N$ strategy.

A. Optimal Portfolios under Quadratic Loss

Let $r_t = R_t - \iota R_{ft}$, where R_t is an $N \times 1$ vector of risky-asset returns, R_{ft} is the risk-free rate, and ι denotes an $N \times 1$ vector of 1s. Under the standard approach to conditional mean-variance optimization, the investor's objective in period t is to choose the $N \times 1$ vector of risky-asset weights ω_{pt} that maximizes the quadratic objective function

$$(1) \qquad Q(\omega_{pt}) = \omega_{pt}'\mu_t - \frac{\gamma}{2}\omega_{pt}'\Sigma_t\omega_{pt},$$

where $\mu_t = \mathrm{E}_t(r_{t+1})$ is the conditional mean vector of the excess risky-asset returns, $\Sigma_t = \mathrm{E}_t(r_{t+1}r_{t+1}') - \mathrm{E}_t(r_{t+1})\mathrm{E}_t(r_{t+1})'$ is the conditional covariance matrix of the excess risky-asset returns, and γ denotes the investor's coefficient of relative risk aversion. The weight in the risk-free asset is determined implicitly by $1-\omega_{pt}'\iota$.

This problem has a straightforward and well-known solution: $\omega_{pt} = \Sigma_t^{-1}\mu_t/\gamma$. The solution implies that, in general, the investor divides his wealth between the risk-free asset and a tangency portfolio (TP) of risky assets with weights $\omega_{\mathrm{TP},t} = \Sigma_t^{-1}\mu_t/\iota'\Sigma_t^{-1}\mu_t$. That is, he holds a conditionally MVE portfolio. The fraction of wealth allocated to the TP is $x_{\mathrm{TP},t} = \iota'\Sigma_t^{-1}\mu_t/\gamma$. Because there is a

1-to-1 correspondence between γ and $\mu_{pt} = \omega'_{pt}\mu_t$ for each t, we can express the vector of optimal weights as

$$(2) \qquad \omega_{pt} = \mu_{pt}\left(\frac{\Sigma_t^{-1}\mu_t}{\mu_t'\Sigma_t^{-1}\mu_t}\right)$$

and view the investor as choosing the period t portfolio by minimizing the conditional risk of the portfolio for a specified value of the conditional expected excess return.

We refer to the portfolio in equation (2) as the optimal unconstrained (OU) portfolio because the sum of the risky-asset weights is unconstrained. DeMiguel et al. (2009) focus on portfolios that constrain these weights to sum to 1 to ensure that performance differences across portfolios are not driven by different allocations to the risk-free and risky assets. This constraint is imposed by rescaling the weights of the OU portfolio to obtain the TP.[1] This may seem innocuous given that the tangency and OU portfolios have the same conditional Sharpe ratio. However, the TP differs from the OU portfolio in 2 important respects: estimation risk and turnover. These differences can have a substantial impact on the relative performance of the corresponding investment strategies.

1. Estimation Risk, Turnover, and the Tangency Portfolio

First consider the estimation risk issue. Intuitively, estimation risk stems from uncertainty about the parameters of the data generating process. This leads to errors in estimating the portfolio weights, which drives up portfolio risk. Suppose, for example, that $\mu_t = \mu$ and $\Sigma_t = \Sigma$ for all t. If μ and Σ are known to the investor, then the OU portfolio and the TP have known, time-invariant weights given by $\omega_p = \Sigma^{-1}\mu/\gamma$ and $\omega_{\text{TP}} = \Sigma^{-1}\mu/\iota'\Sigma^{-1}\mu$. Because the weights are proportional to $\Sigma^{-1}\mu$ in each case, the excess return on the OU portfolio is perfectly correlated with the excess return on the TP, and the 2 portfolios have the same *unconditional* Sharpe ratio.

This is not true, however, if we replace μ and Σ with the sample mean vector $\hat{\mu}$ and sample covariance matrix $\hat{\Sigma}$. Sampling variation in $\hat{\mu}$ and $\hat{\Sigma}$ translates into sampling variation in the portfolio weights, which increases the variance of the returns and lowers the unconditional Sharpe ratio. Although both portfolios are affected, the deterioration in the Sharpe ratio is likely to be more severe for the TP. Consider the expression for the estimated weights: $\hat{\omega}_{\text{TP}} = \hat{\Sigma}^{-1}\hat{\mu}/\iota'\hat{\Sigma}^{-1}\hat{\mu}$. If $|\iota'\hat{\Sigma}^{-1}\hat{\mu}|$ is small, then the TP can display extreme weights and hence extreme returns. The OU portfolio does not suffer from this problem because the vector $\hat{\Sigma}^{-1}\hat{\mu}$ is scaled by $1/\gamma$ rather than by $1/\iota'\hat{\Sigma}^{-1}\hat{\mu}$. The investor chooses γ, while $\iota'\hat{\Sigma}^{-1}\hat{\mu}$ is a random variable that can take on values close to 0 if there is sufficient sampling variation in $\hat{\mu}$ and $\hat{\Sigma}$.

Estimation risk is particularly important in the presence of transaction costs, because anything that increases turnover (the fraction of invested wealth traded

[1]To be precise, DeMiguel et al. (2009) consider a portfolio with weights $\omega_t^* = \Sigma_t^{-1}\mu_t/|\iota'\Sigma_t^{-1}\mu_t|$. This portfolio invests 100% in the TP if $\iota'\Sigma_t^{-1}\mu_t > 0$. However, it invests -100% in the TP and 200% in the risk-free asset if $\iota'\Sigma_t^{-1}\mu_t < 0$ (i.e., if the TP is conditionally inefficient).

in a given period) can cause performance to deteriorate. To see how we compute turnover for the OU portfolio, note that if \$1 is invested in the portfolio at time $t-1$, there will be $\omega_{i,t-1}(1+R_{it})$ dollars invested in the ith risky asset at time t. Hence, the weight in asset i before the portfolio is rebalanced at time t is

$$\text{(3)} \qquad \tilde{\omega}_{it} \quad = \quad \frac{\omega_{i,t-1}(1+R_{it})}{\sum_{i=1}^{N}\omega_{i,t-1}(1+R_{it}) + (1-\sum_{i=1}^{N}\omega_{i,t-1})(1+R_{ft})},$$

and the turnover at time t is given by

$$\text{(4)} \qquad \tau_{pt} \quad = \quad \sum_{i=1}^{N}|\omega_{it}-\tilde{\omega}_{it}| + \left|\sum_{i=1}^{N}(\omega_{it}-\tilde{\omega}_{it})\right|,$$

where ω_{it} is the desired weight in asset i at time t.

Equations (3) and (4) imply that turnover is largely determined by the value of γ. Consider our earlier example with $\mu_t=\mu$ and $\Sigma_t=\Sigma$ for all t. In this case, the estimated weights of the OU portfolio are given by $\hat{\omega}_p=\hat{\Sigma}^{-1}\hat{\mu}/\gamma$. Changing γ has no effect on the before-transaction-costs Sharpe ratio because both the mean and standard deviation of the portfolio return are proportional to $1/\gamma$. However, it can have a dramatic impact on turnover because $|\omega_{it}-\tilde{\omega}_{it}|$ is approximately proportional to $1/\gamma$.[2] Reducing γ causes rebalancing costs to rise, driving down the mean return and, correspondingly, the after-transaction-costs Sharpe ratio. Hence, the choice of γ is an important consideration when developing a research design to evaluate the effectiveness of mean-variance optimization.

Turnover is an even greater concern for the TP strategy. If there is more than a small chance that $\iota'\hat{\Sigma}^{-1}\hat{\mu}$ is less than 0, then realizations of this quantity that are close to 0 can push turnover to very high levels. The effect is the same as setting γ close to 0 for the OU strategy. Unfortunately, we can do little to mitigate this problem. In contrast to the OU strategy, we cannot reduce turnover by specifying a higher γ. This is a major drawback in the presence of transaction costs. By focusing on the TP strategy, DeMiguel et al. (2009) place the mean-variance model at an inherent disadvantage with respect to turnover and estimation risk. Thus, their results could produce an overly pessimistic picture of the usefulness of mean-variance optimization.

2. Optimization Over the Risky Assets Only

If the objective is to study MVE portfolios that exclude the risk-free asset, an alternative to considering the TP is to solve the investor's portfolio problem subject to the constraint $\omega_{pt}'\iota=1$. The 1st-order condition for this problem is

$$\text{(5)} \qquad \mu_t+\delta_t\iota-\gamma\Sigma_t\omega_{pt} \quad = \quad 0,$$

[2]Note that $\tilde{\omega}_{it}=\omega_{i,t-1}(1+R_{it})$ for the special case in which the portfolio has a zero return in period t. Because $\omega_{i,t-1}$ and ω_{it} are proportional to $1/\gamma$, it follows that $|\omega_{it}-\tilde{\omega}_{it}|$ is proportional to $1/\gamma$ in this case. More generally, approximate proportionality holds.

where δ_t is the Lagrange multiplier associated with the constraint. Hence, the optimal vector of constrained portfolio weights is

$$\omega_{pt} = \frac{1}{\gamma}\Sigma_t^{-1}\mu_t + \frac{\delta_t}{\gamma}\Sigma_t^{-1}\iota. \tag{6}$$

The 1st term on the right-hand side of equation (6) is proportional to $\omega_{\text{TP},t}$. The 2nd term is proportional to $\omega_{\text{MV},t} = \Sigma_t^{-1}\iota/\iota'\Sigma_t^{-1}\iota$, which is the vector of weights for the minimum-variance (MV) portfolio; that is, the portfolio obtained by minimizing $\omega_{pt}'\Sigma_t\omega_{pt}$ subject to $\omega_{pt}'\iota = 1$. Thus, the solution to the constrained problem takes the same general form as that for the unconstrained problem with the MV portfolio replacing the risk-free asset.

If we solve for δ_t and substitute the resulting expression into equation (6), we obtain

$$\omega_{pt} = x_{\text{TP},t}\left(\frac{\Sigma_t^{-1}\mu_t}{\iota'\Sigma_t^{-1}\mu_t}\right) + (1 - x_{\text{TP},t})\left(\frac{\Sigma_t^{-1}\iota}{\iota'\Sigma_t^{-1}\iota}\right), \tag{7}$$

which implies that $\mu_{pt} = x_{\text{TP},t}\mu_{\text{TP},t} + (1 - x_{\text{TP},t})\mu_{\text{MV},t}$, where $\mu_{\text{TP},t}$ and $\mu_{\text{MV},t}$ denote the conditional expected excess returns for the tangency and MV portfolios. Accordingly, we can express equation (7) as

$$\begin{aligned}\omega_{pt} = {} & \left(\frac{\mu_{pt} - \mu_{\text{MV},t}}{\mu_{\text{TP},t} - \mu_{\text{MV},t}}\right)\left(\frac{\Sigma_t^{-1}\mu_t}{\iota'\Sigma_t^{-1}\mu_t}\right) \\ & + \left(1 - \frac{\mu_{pt} - \mu_{\text{MV},t}}{\mu_{\text{TP},t} - \mu_{\text{MV},t}}\right)\left(\frac{\Sigma_t^{-1}\iota}{\iota'\Sigma_t^{-1}\iota}\right),\end{aligned} \tag{8}$$

and view the investor as choosing ω_{pt} by minimizing conditional risk for a specified μ_{pt}. We refer to the portfolio in equation (8) as the optimal constrained (OC) portfolio.

The OC portfolio is identical to the OU portfolio except that the weight in the risk-free asset has been transferred to the MV portfolio. It follows, therefore, that the increase in estimation risk from imposing the constraint is due solely to errors in estimating $\omega_{\text{MV},t}$. Two observations suggest this increase should be considerably less than that incurred by rescaling the OU portfolio weights (provided the OU portfolio is a convex combination of the risk-free asset and TP). First, $\omega_{\text{MV},t}$ does not depend on μ_t. The return variances and covariances can typically be estimated more precisely than the mean returns (see, e.g., Merton (1980)), so the errors in estimating $\omega_{\text{MV},t}$ should be smaller than those in estimating the weights of other MVE portfolios. Second, the denominator of $\omega_{\text{MV},t} = \Sigma_t^{-1}\iota/\iota'\Sigma_t^{-1}\iota$ is the reciprocal of the variance of the MV portfolio. This suggests that the potential for generating extreme weights is much lower for the MV portfolio than for the TP because the value of $\iota'\hat{\Sigma}_t^{-1}\iota$ is, by construction, both positive and maximized.

B. The DeMiguel et al. (2009) Results Revisited

To provide direct evidence on the role of the issues identified in Section II.A in the DeMiguel et al. (2009) results, we replicate their analysis for the tangency, MV, and $1/N$ portfolios, and compare the results to those for the OC portfolio. In particular, we document the performance of tangency, MV, $1/N$, and OC strategies under circumstances in which the OC portfolio targets the estimated expected excess return of the $1/N$ portfolio (i.e., we set $\hat{\mu}_{pt}$ for the OC portfolio equal to $\hat{\mu}_t'\iota/N$).[3] This ensures that the aggressiveness of the OC portfolio is comparable to that of $1/N$ portfolio in each period. The analysis is conducted for 5 of the 6 DeMiguel et al. data sets.[4] Three are constructed by sorting U.S. firms into portfolios based on size and book-to-market values (the FF 1-factor, FF 4-factor, and Mkt/SMB/HML data sets), one is constructed by sorting U.S. firms into industries using Standard Industrial Classification (SIC) codes (the FF 10 Industry data set), and one contains international equity market indexes (the international data set). The sample size is 497 monthly observations except for the international data set, which contains 379 observations.

We report the results of the comparison in Table 1. Panel A reports the annualized mean, standard deviation, and Sharpe ratio for the time series of monthly excess returns generated by each of the strategies. We use rolling estimators with a 120-month window length and assume that transaction costs are 0 when computing these statistics. Panel B reports the minimum, median, and maximum values of the estimated conditional expected return for each strategy over the months in the out-of-sample period. The estimated Sharpe ratios for the tangency, MV, and $1/N$ strategies match those reported by DeMiguel et al. (2009) in their Table 3 (after scaling by $1/\sqrt{12}$ to obtain monthly statistics).

It is clear that the reward and risk characteristics of the TP strategy are markedly different from those of the other strategies. The estimated mean and estimated standard deviation of the TP excess return are greater than 100% per year for 2 of the data sets. This contrasts sharply with the results for the OC strategy, which has reward and risk characteristics similar to those of the $1/N$ and MV strategies. Moreover, the TP strategy typically has an estimated expected monthly turnover that is orders of magnitude higher than that of the OC strategy. The only exception is for the Mkt/SMB/HML data set. The TP and OC strategies have the same estimated expected turnover for this data set.

Panel B of Table 1 points to why the Mkt/SMB/HML data set produces atypical results. The median value of $\hat{\mu}_{\mathrm{TP},t}$ for this data set is 6.7% per year, which is relatively low compared to median values ranging from 30.9% to 60.8% per year for the remaining data sets. This probably reflects the impact of predominately negative estimated return correlations.[5] More importantly, the maximum value of $\hat{\mu}_{\mathrm{TP},t}$ is only 15.5% for Mkt/SMB/HML, but it ranges from 2,486% to 12,216% per year for the remaining data sets. It is not surprising to find that target estimated expected excess returns of this magnitude produce extreme turnover. Furthermore, the weights that deliver these targets are not feasible in practice.

[3]Occasionally, targeting the estimated expected excess return of the $1/N$ portfolio delivers a conditionally inefficient portfolio. In these cases we replace $\hat{\mu}_{pt} - \hat{\mu}_{\mathrm{MV},t}$ in the sample analog of equation (8) with $|\hat{\mu}_{pt} - \hat{\mu}_{\mathrm{MV},t}|$. This delivers a conditionally efficient portfolio with the same conditional volatility as the identified inefficient portfolio.

[4]We thank Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal for sharing these data. The Standard & Poor's (S&P) sector data set is proprietary and thus not included in analysis.

[5]The estimates of $\mathrm{corr}(r_{\mathrm{Mkt},t}, r_{\mathrm{SMB},t})$ and $\mathrm{corr}(r_{\mathrm{Mkt},t}, r_{\mathrm{HML},t})$ are -0.29 and -0.47.

TABLE 1

Characteristics of the 1/*N* and MVE Strategies for the DeMiguel et al. (2009) Data Sets

Table 1 documents key sample characteristics of the 1/*N*, MV, OC, and TP strategies for 5 of the 6 DeMiguel et al. (2009) data sets: the 10 Fama-French (FF) industry portfolios plus the market portfolio (FF 10 Industry); the 8 Morgan Stanley Capital International (MSCI) developed market portfolios plus the MSCI world market portfolio (International); the FF market, size, and value factor portfolios (Mkt/SMB/HML); the 20 FF size and book-to-market portfolios plus the market portfolio (FF 1-Factor); and the 20 FF size and book-to-market portfolios plus the market, size, value, and momentum factor portfolios (FF 4-Factor). Panel A reports the annualized mean excess return ($\widehat{\mu}_p$), annualized excess return standard deviation ($\widehat{\sigma}_p$), annualized Sharpe ratio ($\widehat{\lambda}_p$), and average monthly turnover expressed as a fraction of wealth invested ($\widehat{\tau}_p$) for each strategy under the assumption that transaction costs are 0. Panel B reports the minimum, median, and maximum of the annualized time-series estimates of the conditional expected excess return ($\widehat{\mu}_{pt}$) for each strategy. The MV, OC, and TP strategies are implemented using a 120-month rolling estimator of the conditional mean vector and conditional covariance matrix of excess returns, and the OC strategy targets the estimated conditional expected excess return of the 1/*N* portfolio each period as described in Section II.B. The sample period is July 1963–November 2004 for the FF data sets and January 1970–July 2001 for the international data set, with 497 and 379 monthly observations, respectively. In each case, the first 120 observations are held out to initialize the rolling estimators. See the text for a detailed description of each strategy.

Panel A. Summary Statistics

	FF 10 Industry				International				Mkt/SMB/HML				FF 1-Factor				FF 4-Factor			
Strategy	$\widehat{\mu}_p$	$\widehat{\sigma}_p$	$\widehat{\lambda}_p$	$\widehat{\tau}_p$	$\widehat{\mu}_p$	$\widehat{\sigma}_p$	$\widehat{\lambda}_p$	$\widehat{\tau}_p$	$\widehat{\mu}_p$	$\widehat{\sigma}_p$	$\widehat{\lambda}_p$	$\widehat{\tau}_p$	$\widehat{\mu}_p$	$\widehat{\sigma}_p$	$\widehat{\lambda}_p$	$\widehat{\tau}_p$	$\widehat{\mu}_p$	$\widehat{\sigma}_p$	$\widehat{\lambda}_p$	$\widehat{\tau}_p$
TP	108	457	0.24	473	−19.5	170	−0.11	1,077	5.66	7.47	0.76	0.06	35.92	812	0.04	212	139	219	0.64	69.5
1/*N*	7.14	15.23	0.47	0.02	6.66	15.07	0.44	0.03	4.93	6.36	0.78	0.02	10.47	18.62	0.56	0.02	9.93	16.35	0.61	0.02
MV	7.07	13.13	0.54	0.46	7.60	14.72	0.52	0.21	4.80	5.56	0.86	0.02	12.86	13.37	0.96	0.74	−0.20	3.09	−0.06	0.13
OC	6.55	13.30	0.49	0.65	7.10	14.75	0.48	0.29	4.77	6.26	0.76	0.06	16.32	14.18	1.15	0.95	5.69	4.94	1.15	0.46

Panel B. Estimated Conditional Expected Returns ($\widehat{\mu}_{pt}$)

	FF 10 Industry			International			Mkt/SMB/HML			FF 1-Factor			FF 4-Factor		
Strategy	Min.	Med.	Max.	Min.	Med.	Max.	Min.	Med.	Max.	Min.	Med.	Max.	Min.	Med.	Max.
TP	22.7	47.4	12,216	17.2	30.9	2,486	3.4	6.7	15.5	38.1	60.8	9,191	3.6	34.6	7,635
1/*N*	−3.2	7.9	14.1	−2.4	7.6	12.6	1.3	3.7	8.5	−0.5	8.6	16.5	0.4	8.0	15.5
MV	−1.7	4.2	12.2	−3.4	6.4	12.0	1.8	4.3	8.6	−3.1	14.2	22.5	−0.7	−0.1	0.8
OC	−0.3	8.4	15.8	−2.4	8.9	16.2	2.6	5.0	8.6	−0.5	16.8	36.7	0.4	8.0	15.5

In comparison, the maximum value of $\hat{\mu}_{pt}$ for the OC strategy ranges from 8.6% to 36.7% per year, a much more reasonable range. If we discount the results for the TP strategy, the picture that emerges from Table 1 is far more supportive of mean-variance optimization than that suggested by DeMiguel et al. (2009). The estimated Sharpe ratio of the OC strategy exceeds that of the $1/N$ strategy for 4 of the data sets. Indeed, for the FF 1-factor and FF 4-factor data sets, the estimated Sharpe ratio for the OC strategy is greater than 1 and about twice that of the $1/N$ strategy. This supports our view that the DeMiguel et al. research design, which makes no attempt to match the risk characteristics of the MVE portfolios under consideration to those of the naïve diversification benchmark, is skewed in favor of naïve diversification, especially with respect to turnover and after-transaction-costs performance.[6]

For example, DeMiguel et al. (2009) implement a version of the "3-fund" strategy proposed by Kan and Zhou (2007), but it contains only 2 funds (the tangency and MV portfolios) because they rescale the weights for the risky assets to sum to 1. The estimated expected turnover of the resulting strategy exceeds that of the $1/N$ strategy by a factor of more than 1,000 for several of the data sets. However, we find that the OC strategy, which is a combination of the same 2 funds, has a vastly lower turnover when we target the conditional expected excess return of the $1/N$ portfolio. We should not interpret the turnover and other performance figures reported by DeMiguel et al. as representative of the 3-fund strategy.

Our findings with respect to the DeMiguel et al. (2009) research design are consistent with the analysis of Brown (1979). He argues that the $1/N$ benchmark is difficult to beat in the standard mean-variance framework because this framework fails to account for the impact of estimation risk. Intuitively, optimization based on sample means, variances, and covariances tends to produce portfolio choices that are too extreme.[7] To make it less likely that estimation risk will overwhelm the potential benefits of incorporating sample information, we target the estimated expected excess return of the $1/N$ portfolio in the optimization. This reduces the chances (relative to the scenario considered in DeMiguel et al.) that optimization produces extreme weights due to estimation error.

C. Mean-Variance Timing Strategies

Although the OC strategy performs much better than the TP strategy, it is not clear that the OC strategy would consistently outperform the $1/N$ strategy under plausible assumptions about transaction costs. For instance, if we assume that establishing or liquidating a portfolio position costs 50 bp, then the estimates of expected turnover for the OC strategy reported in Table 1 would entail transaction costs of between 0.4% and 5.7% per year. In view of the potential impact of transaction costs, we regard turnover as the primary barrier to capitalizing on the gains promised by mean-variance optimization.[8]

[6] Tu and Zhou (2011) discuss additional reasons for the unusually strong relative performance of the $1/N$ portfolio in DeMiguel et al. (2009).

[7] See Jorion (1985), (1991) for further discussion and possible solutions.

[8] Transaction costs might be less of an issue for large institutional investors. Establishing or liquidating a portfolio position could plausibly cost as little as 5 bp for such investors.

We could consider reducing turnover by using various techniques proposed to improve the performance of mean-variance optimization.[9] However, our interest lies in a different direction. Instead of focusing strictly on portfolio optimization, we expand the scope of the investigation to include alternative methods of exploiting sample information. Our objective is to develop methods of portfolio selection that retain the features that make naïve diversification appealing (nonnegative weights, low turnover, and wide applicability) while improving on its performance.

1. Volatility Timing

Fleming, Kirby, and Ostdiek (2001), (2003) study a class of active portfolio strategies in which the portfolio weights are rebalanced based on changes in the estimated conditional covariance matrix of returns. They find that these "volatility-timing" strategies outperform unconditionally MVE portfolio strategies by statistically significant margins. This points to the potential for a VT approach to outperform naïve diversification. The question is how to implement VT in the present setting. Unlike Fleming et al. (2001), (2003), who use futures contracts for their analysis, we want to avoid short sales and keep turnover as low as possible. Accordingly, we propose a new class of VT strategies characterized by 4 notable features: They do not require optimization, they do not require covariance matrix inversion, they do not generate negative weights, and they allow the sensitivity of the weights to volatility changes to be adjusted via a tuning parameter. The last feature facilitates control over turnover and transaction costs.

To motivate our approach, consider a scenario in which all of the estimated pair-wise correlations between the excess risky-asset returns are 0 (i.e., $\hat{\Sigma}_t$ is a diagonal matrix). In this case, the weights for the sample MV portfolio are given by

$$\hat{\omega}_{it} = \frac{(1/\hat{\sigma}_{it}^2)}{\sum_{i=1}^{N}(1/\hat{\sigma}_{it}^2)}, \qquad i = 1, 2, \ldots, N, \tag{9}$$

where $\hat{\sigma}_{it}$ is the estimated conditional volatility of the excess return on the ith risky asset. Thus, if $\hat{\Sigma}_t$ is restricted to be diagonal for all t, the investor will follow a very simple VT strategy. Obviously, we do not expect $\hat{\Sigma}_t$ to actually be diagonal. However, the sample MV portfolio obtained by setting the off-diagonal elements of $\hat{\Sigma}_t$ to 0 might perform better than that obtained using the usual estimator of Σ_t.

To see why, note that weights in equation (9) are strictly nonnegative, while the weights obtained using a nondiagonal estimator of Σ_t will typically involve short positions in some assets. In general, strategies that permit short sales are more likely to generate extreme weights. We view setting the off-diagonal elements of $\hat{\Sigma}_t$ to 0 as an aggressive form of shrinkage. Because this results in $N(N-1)/2$ fewer parameters to estimate, the reduction in estimation risk could outweigh the information loss. We can also reduce the impact of the information

[9] Some recent examples of work in this area include Pastor (2000), Pastor and Stambaugh (2000), MacKinlay and Pastor (2000), Jagannathan and Ma (2003), Ledoit and Wolf (2003), (2004), Garlappi, Uppal, and Wang (2007), and Kan and Zhou (2007).

loss by modifying the way in which the portfolio weights respond to volatility changes. Consider the $N=2$ case. The estimated weights of the MV portfolio are in general given by

(10) $$\hat{\omega}_{1t} = \frac{\hat{\sigma}_{2t}^2 - \hat{\sigma}_{1t}\hat{\sigma}_{2t}\hat{\rho}_t}{\hat{\sigma}_{1t}^2 + \hat{\sigma}_{2t}^2 - 2\hat{\sigma}_{1t}\hat{\sigma}_{2t}\hat{\rho}_t}$$

and $\hat{\omega}_{2t} = 1 - \hat{\omega}_{1t}$, where $\hat{\rho}_t$ is the estimated conditional correlation between the excess returns. Now suppose $\hat{\sigma}_{1t} = \hat{\sigma}_{2t}$ so that $\hat{\omega}_{pt} = (1/2, 1/2)'$. If asset 1's estimated conditional volatility doubles in period $t+1$, we adjust the portfolio weights to $\hat{\omega}_{t+1} = (0,1)'$ for $\hat{\rho}_{t+1} = 1/2$ and to $\hat{\omega}_{t+1} = (1/5, 4/5)'$ for $\hat{\rho}_{t+1} = 0$. Thus, the weights are more responsive to volatility changes when the estimated correlation between the returns is positive.

Although the strategy in equation (9) provides no flexibility in determining how the portfolio weights respond to volatility changes, it belongs to a more general class of VT strategies with weights of the form

(11) $$\hat{\omega}_{it} = \frac{(1/\hat{\sigma}_{it}^2)^{\eta}}{\sum_{i=1}^{N}(1/\hat{\sigma}_{it}^2)^{\eta}}, \qquad i = 1, 2, \ldots, N,$$

where $\eta \geq 0$. The idea behind this generalization is straightforward. The tuning parameter η is a measure of timing aggressiveness (i.e., it determines how aggressively we adjust the portfolio weights in response to volatility changes). As $\eta \to 0$ we recover the naïve diversification portfolio, and as $\eta \to \infty$ the weight on the asset with the lowest volatility approaches 1. Setting $\eta > 1$ should help compensate for the information loss caused by ignoring the correlations. We refer to the portfolio in equation (11) as the VT(η) portfolio.

2. Reward-to-Risk Timing

The VT strategies of Section II.C.1 ignore information about conditional expected returns. It is natural to ask, therefore, whether we can improve upon their performance by incorporating such information. Suppose we again consider a scenario in which all of the estimated pair-wise correlations between the excess risky-asset returns are 0. The weights for the sample TP in this case are given by

(12) $$\hat{\omega}_{it} = \frac{(\hat{\mu}_{it}/\hat{\sigma}_{it}^2)}{\sum_{i=1}^{N}(\hat{\mu}_{it}/\hat{\sigma}_{it}^2)}, \qquad i = 1, 2, \ldots, N,$$

where $\hat{\mu}_{it}$ is the estimated conditional mean for the ith asset. Thus, if $\hat{\Sigma}_t$ is restricted to be diagonal for all t, the investor will follow a simple RRT strategy.

Because expected returns are typically estimated with less precision than variances, the strategy in equation (12) is likely to entail significantly higher levels of estimation risk than the VT strategies. Setting the off-diagonal elements of $\hat{\Sigma}_t$ to 0 reduces the tendency for the sample TP to be characterized by extreme long and short weights, but we could still see extreme weights if $\hat{\mu}_{it}$ is negative for some assets because this could cause the denominator of the fraction on the right-hand side of equation (12) to be close to 0. We address this possibility by assuming that the investor has a strong prior belief that $\mu_{it} \geq 0$ for all i and therefore constructs the RRT weights as

$$(13) \qquad \hat{\omega}_{it} = \frac{(\hat{\mu}_{it}^{+}/\hat{\sigma}_{it}^{2})}{\sum_{i=1}^{N}(\hat{\mu}_{it}^{+}/\hat{\sigma}_{it}^{2})}, \qquad i = 1,2,\ldots,N,$$

where $\hat{\mu}_{it}^{+} = \max(\hat{\mu}_{it}, 0)$. This is equivalent to assuming that the investor eliminates any asset with $\hat{\mu}_{it} \leq 0$ from consideration in period t.

Using the same approach as in Section II.C.1, we can view equation (13) as an example of a more general class of RRT strategies that have weights of the form

$$(14) \qquad \hat{\omega}_{it} = \frac{(\hat{\mu}_{it}^{+}/\hat{\sigma}_{it}^{2})^{\eta}}{\sum_{i=1}^{N}(\hat{\mu}_{it}^{+}/\hat{\sigma}_{it}^{2})^{\eta}}, \qquad i = 1,2,\ldots,N,$$

where $\eta \geq 0$. These strategies approach naïve diversification across the assets with positive estimated expected excess returns as $\eta \to 0$ and put a weight that approaches 1 on the asset with the maximum estimated reward-to-risk ratio as $\eta \to \infty$. We refer to the portfolio in equation (14) as the RRT(μ_t^{+}, η) portfolio.

III. Estimating the Conditional Moments of Returns

To implement the portfolio strategies, we must estimate μ_t and Σ_t for each portfolio rebalancing date t. For our baseline analysis we use fixed-window rolling estimators of μ_t and Σ_t. This allows us to directly compare our results to those of DeMiguel et al. (2009). In the case of the RRT strategies, we also consider an alternative estimator of μ_t that is designed to reduce estimation risk by exploiting the predictions of asset pricing theory.

A. Rolling Estimators

Using a rolling data window to estimate μ_t and Σ_t is designed to balance the tradeoff between the efficiency gains from using more observations and the loss in forecast precision from including less timely observations that are less likely to reflect current market conditions. To implement this approach, we define our estimators to be $\hat{\mu}_t = (1/L)\sum_{l=0}^{L-1} r_{t-l}$ and $\hat{\Sigma}_t = (1/L)\sum_{l=0}^{L-1}(r_{t-l}-\hat{\mu}_t)(r_{t-l}-\hat{\mu}_t)'$ for some window length L. Common choices of the window length for monthly data are $L = 60$ and $L = 120$ (i.e., 5- and 10-year rolling windows). We follow DeMiguel et al. (2009) and set $L = 120$.

Although rolling estimators of conditional expected excess returns have the advantage of simplicity, using these estimators in portfolio optimization is likely to entail a high level of estimation risk. It is well known that we need a long time series of returns to estimate μ_t accurately (Merton (1980)). This is true even if μ_t is time invariant. We therefore consider an alternative estimator of μ_t for implementing the RRT strategies that should reduce estimation risk under the circumstances described later.

B. Alternative Estimator of Conditional Expected Returns

Many asset pricing models imply a direct relationship between the 1st and 2nd moments of excess returns. To see how we can exploit this relationship in the context of RRT, suppose that a conditional version of the capital asset pricing model (CAPM) holds. The conditional CAPM implies that the cross-sectional variation in conditional expected excess returns is due to cross-sectional variation in conditional betas. Since the market risk premium is just a scaling factor that multiplies each of the conditional betas, we can express the weights for the RRT portfolio as

$$\text{(15)} \qquad \omega_{it} = \frac{(\beta_{it}^{+}/\sigma_{it}^{2})^{\eta}}{\sum_{i=1}^{N}(\beta_{it}^{+}/\sigma_{it}^{2})^{\eta}}, \qquad i = 1, 2, \ldots, N,$$

where $\beta_{it}^{+} = \max(\beta_{it}, 0)$ and β_{it} is the period t conditional market beta of asset i.

Replacing μ_{it}^{+} with β_{it}^{+} can potentially lower the sampling variation of the weights. Consider, for illustration purposes, a scenario in which $r_{t+1} \sim$ i.i.d. $N(\mu, \Sigma)$. Upon further simplification, equation (15) reduces to

$$\text{(16)} \qquad \omega_{it} = \frac{(\rho_{i}^{+}/\sigma_{i})^{\eta}}{\sum_{i=1}^{N}(\rho_{i}^{+}/\sigma_{i})^{\eta}}, \qquad i = 1, 2, \ldots, N,$$

where $\rho_{i}^{+} = \max(\rho_{i}, 0)$ and ρ_{i} is the correlation between the excess return on asset i and the excess return on the market. Hence, we have replaced $\hat{\mu}_{i}$ with $\hat{\sigma}_{i}\hat{\rho}_{i}$ in the sample version of the strategy. With a window length of L, the asymptotic variances of $\hat{\mu}_{i}$ and $\hat{\sigma}_{i}\hat{\rho}_{i}$ are given by σ_{i}^{2}/L and $(\sigma_{i}^{2}/L)(1-\rho_{i}^{2}/2)$.[10]

Now suppose in the alternative that the conditional CAPM does not hold, but conditional betas capture at least some of the cross-sectional variation in conditional expected returns. In this case, replacing $\hat{\mu}_{it}$ with $\hat{\sigma}_{it}\hat{\rho}_{it}$ will introduce bias. The substitution may still prove beneficial, however, if we are replacing an unbiased but high variance estimator with a biased but lower variance estimator. Consider an estimator of the form $\hat{\mu}_{t} = \mu_{0}\iota$, where $\mu_{0} > 0$ is a scalar. This estimator is undoubtedly biased. Nonetheless, an investor who uses it, and imposes the constraint $\omega_{t}'\iota = 1$, holds the sample analog of the MV portfolio. This portfolio often performs better than other sample efficient portfolios because its weights do not depend on $\hat{\mu}_{t}$, which reduces estimation risk (see, e.g., Jagannathan and Ma (2003)).

This methodology can be extended to allow for multiple risk factors. Consider a K-factor model, and let $\beta_{ij,t}$ denote the period t conditional beta of the ith asset with respect to the jth factor. With a single factor the portfolio weights do not depend on the factor risk premium provided that $\omega_{t}'\iota = 1$. This is not the case for $K > 1$. Estimating the factor risk premiums would introduce additional errors

[10]To see this, let σ_{m}^{2} and σ_{im} denote the variance of the excess market return and the covariance between the excess return on asset i and the excess market return. It is easy to show that

$$\sqrt{L}\begin{pmatrix} \hat{\sigma}_{m}^{2} - \sigma_{m}^{2} \\ \hat{\sigma}_{im} - \sigma_{im} \end{pmatrix} \overset{d}{\to} N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 2\sigma_{m}^{4} & 2\sigma_{m}^{2}\sigma_{im} \\ 2\sigma_{m}^{2}\sigma_{im} & \sigma_{m}^{2}\sigma_{i}^{2}(1+\rho_{i}^{2}) \end{bmatrix} \right).$$

See, for example, Hamilton ((1994), p. 301). The asymptotic variance of $\hat{\sigma}_{i}\hat{\rho}_{i} = \hat{\sigma}_{im}/\hat{\sigma}_{m}$ follows immediately via the delta method.

that could easily overwhelm any benefits from employing the model. Instead, we mimic the approach used to obtain the MV portfolio (i.e., we assume for the purpose of computing the weights that the factors have identical risk premiums). The resulting weights for the RRT strategy are

(17) $$\omega_{it} = \frac{(\bar{\beta}_{it}^{+}/\sigma_{it}^{2})^{\eta}}{\sum_{i=1}^{N}(\bar{\beta}_{it}^{+}/\sigma_{it}^{2})^{\eta}}, \qquad i = 1, 2, \ldots, N,$$

where $\bar{\beta}_{it}^{+} = \max(\bar{\beta}_{it}, 0)$ and $\bar{\beta}_{it} = (1/K)\sum_{j=1}^{K}\beta_{ij,t}$ is the average conditional beta of asset i with respect to the K factors. We refer to this portfolio as the RRT($\bar{\beta}_{t}^{+}, \eta$) portfolio.

We implement the RRT($\bar{\beta}_{t}^{+}, \eta$) strategy using the Carhart (1997) 4-factor extension of the Fama and French (FF) (1993) 3-factor model. The conditional CAPM is not used for the empirical analysis because the literature finds little evidence of a relationship between market betas and average returns after accounting for size and book-to-market characteristics (FF). To estimate the conditional factor betas, we use a rolling estimator of the joint conditional covariance matrix of the excess risky asset and factor returns with the same L used to estimate μ_t.

IV. Evaluating Portfolio Performance

Suppose a given data set contains $T + L$ observations, where T is the length of the out-of-sample period. After we compute the sequence $\{r_{pt}\}_{t=L+1}^{T+L}$ of out-of-sample excess returns for each strategy considered, we evaluate the performance of the strategies using 2 criteria. The first is the Sharpe ratio for the strategy, that is, $\lambda_p = \mu_p/\sigma_p$, where $\mu_p = \mathrm{E}(r_{pt})$ and $\sigma_p^2 = \mathrm{Var}(r_{pt})$. We estimate this ratio using the sample mean and variance of the excess returns for the out-of-sample period: $\hat{\mu}_p = (1/T)\sum_{t=L+1}^{T+L} r_{pt}$ and $\hat{\sigma}_p^2 = (1/T)\sum_{t=L+1}^{T+L}(r_{pt} - \hat{\mu}_p)^2$. The difference in the estimated Sharpe ratios for any 2 strategies is one measure of relative performance in the out-of-sample tests. We base the reported values of $\hat{\lambda}_p$ on the annualized values of $\hat{\mu}_p$ and $\hat{\sigma}_p$ for each strategy.

Following Fleming et al. (2001), (2003), we also report a relative performance measure based on quadratic utility. The assumption is that quadratic utility is a 2nd-order approximation of the investor's true utility function. Under this approximation, the investor's realized utility in period $t + 1$ can be expressed as

(18) $$U(R_{p,t+1}) = W_t(1 + R_{p,t+1}) - \frac{1}{2}\alpha W_t^2(1 + R_{p,t+1})^2,$$

where $R_{p,t+1} = R_{f,t+1} + \omega_{pt}' r_{t+1}$ is the portfolio return in period $t + 1$, W_t is wealth in period t, and α is the coefficient of absolute risk aversion. To facilitate comparisons across strategies, we hold αW_t constant. This is equivalent to setting the investor's coefficient of relative risk, $\gamma_t = \alpha W_t/(1 - \alpha W_t)$, equal to some fixed value γ. Our performance measure is the fee (expressed as a fraction of wealth invested) that would equate the expected utilities generated by 2 alternative strategies.

Suppose, for example, that with a Δ_γ fee imposed on strategy j each period, the strategies i and j yield the same expected utility (i.e., $\mathrm{E}[U(R_{p_i t})] = \mathrm{E}[U(R_{p_j t} - \Delta_\gamma)]$). The investor would be indifferent between these 2 alternatives, so we interpret Δ_γ as the maximum per period fee the investor would pay to switch from strategy i to strategy j. Recognizing that $\mathrm{E}[U(R_{p_i t})] = \mathrm{E}[U(R_{p_j t})] \iff \Delta_\gamma = 0$, it follows from the quadratic formula that

$$(19) \quad \Delta_\gamma = -\gamma^{-1}(1 - \gamma \mathrm{E}[R_{p_j t}]) + \gamma^{-1}((1 - \gamma \mathrm{E}[R_{p_j t}])^2 - 2\gamma \mathrm{E}[U(R_{p_i t}) - U(R_{p_j t})])^{1/2}.$$

We consider 2 levels of relative risk aversion ($\gamma = 1$ and $\gamma = 5$), specify naïve diversification as strategy i, and report $\hat{\Delta}_\gamma$, the sample analog of Δ_γ, as an annualized basis point value for each active strategy j considered.

A. Portfolio Turnover and Adjustments for Transaction Costs

Because naïve diversification generates low turnover, active strategies that generate high turnover are disproportionately affected by the imposition of transaction costs. We document the impact of turnover by reporting a 2nd set of results using returns measured net of transaction costs. The cost of rebalancing to the desired period $t+1$ weights is subtracted from the excess portfolio return for period t. Our analysis assumes that the level of transaction costs is constant across assets and over the sample period.

To illustrate, let $\tilde{R}_{pt}$ denote the portfolio return net of transaction costs for period t. Under our assumptions, this return is given by $\tilde{R}_{pt} = (1+R_{pt})(1-\tau_{pt}c)-1$, where c is the level of proportional costs per transaction. To a close approximation, the impact of imposing transaction costs can be deduced by subtracting $\hat{\tau}_p c$ from the sample mean of R_{pt}, where $\hat{\tau}_p = (1/T)\sum_{t=L+1}^{T+L} \tau_{pt}$ is our estimate of the expected turnover. Because we report $\hat{\tau}_p$ for each strategy, the choice of c used to compute the net returns is not critical. DeMiguel et al. (2009) follow Balduzzi and Lynch (1999) and set c equal to 50 bp. We do the same to facilitate comparisons with their results.

B. Statistical Inference

We conduct inferences about the relative performance of different strategies using large-sample t-statistics. To illustrate, let $\hat{\lambda}_{p_j}$ and $\hat{\lambda}_{p_i}$ denote the estimated Sharpe ratios for strategies i and j. If the 2 strategies have the same population Sharpe ratio, then we have the large-sample approximation

$$(20) \quad \sqrt{T}\left(\frac{\hat{\lambda}_{p_j} - \hat{\lambda}_{p_i}}{\hat{V}_\lambda^{1/2}}\right) \overset{a}{\sim} N(0,1),$$

where $\hat{V}_\lambda$ denotes a consistent estimator of the asymptotic variance of $\sqrt{T}$ $(\hat{\lambda}_{p_j} - \hat{\lambda}_{p_i})$.[11] To identify strategies that outperform naïve diversification, we specify

[11]We use the generalized method of moments to construct this estimator. Let

$$e_t(\hat{\theta}) = \begin{pmatrix} r_{p_i t} - \hat{\sigma}_{p_i}\hat{\lambda}_{p_i} \\ r_{p_j t} - \hat{\sigma}_{p_j}\hat{\lambda}_{p_j} \\ (r_{p_i t} - \hat{\sigma}_{p_i}\hat{\lambda}_{p_i})^2 - \hat{\sigma}_{p_i}^2 \\ (r_{p_j t} - \hat{\sigma}_{p_j}\hat{\lambda}_{p_j})^2 - \hat{\sigma}_{p_j}^2 \end{pmatrix},$$

naïve diversification as strategy i and report p-values for H_0: $\lambda_{p_j} - \lambda_{p_i} \leq 0$ based on the previous t-statistic.

Because we have no evidence on the quality of the approximation in equation (20), we compute the p-values using a block bootstrap approach. Each bootstrap trial consists of 2 steps. Let $y = (y_{L+1}, y_{L+2}, \ldots, y_{T+L})$, where $y_t = (r_{p_i t}, r_{p_j t})$, denote the set of out-of-sample excess returns for strategies i and j. First, we construct a resample $y^* = (y^*_{L+1}, y^*_{L+2}, \ldots, y^*_{T+L})$ using the stationary bootstrap of Politis and Romano (1994). The resample is such that, in general, if $y^*_s = y_t$, then $y^*_{s+1} = y_{t+1}$ with probability π and y^*_{s+1} is drawn randomly from y with probability $1 - \pi$. This delivers an expected block length of $1/(1-\pi)$. Second, we calculate

$$(21) \qquad \hat{\vartheta}^* = \sqrt{T}\left(\frac{(\hat{\lambda}^*_{p_j} - \hat{\lambda}^*_{p_i}) - (\hat{\lambda}_{p_j} - \hat{\lambda}_{p_i})}{\hat{V}^{*1/2}_{\lambda}}\right),$$

where $\hat{\lambda}^*_{p_i}$, $\hat{\lambda}^*_{p_j}$, and $\hat{V}^*_{\lambda}$ denote the estimates for the resample. After carrying out M bootstrap trials in total, we compute the p-values for the t-statistic in equation (20) using the observed percentiles of $\hat{\vartheta}^*$. We set $M = 10,000$ and $\pi = 0.9$ for an expected block length of 10. We use a similar approach to assess the statistical significance of the estimated performance fees.

V. Data and Empirical Results

The data for the empirical analysis consist of monthly excess returns on broadly based U.S. equity portfolios. The sample period is July 1963–December 2008 ($T + L = 546$ monthly observations with $L = 120$). We consider 4 data sets in total. Three are drawn from the data library maintained by Ken French, and one is constructed using the Center for Research in Security Prices (CRSP) daily stock file. The data library is also the source of the T-bill rate and the factor returns that are used to estimate the beta coefficients for the 4-factor risk model.[12] The risk factors are the excess return on the market index and the returns on a set of 3 zero-investment portfolios. These portfolios are designed to mimic the unobserved factors that lead to systematic differences in expected excess returns between small and large capitalization stocks (SMB), low and high book-to-market equity stocks (HML), and low and high momentum stocks (UMD).

We begin the analysis with a data set formed by sorting firms into 10 industry portfolios (10 Industry). With similar data, DeMiguel et al. (2009) find that none of the 14 portfolio selection methods considered performs better than naïve diversification by a statistically significant margin. Next, we consider a data set formed by using market capitalization and book-to-market value to sort firms into 25 portfolios (25 Size/BTM). Sorting firms on these criteria is known to produce a

---

where $\hat{\theta} = (\hat{\lambda}_{p_i}, \hat{\lambda}_{p_j}, \hat{\sigma}^2_{p_i}, \hat{\sigma}^2_{p_j})'$. Under suitable regularity conditions (see Hansen (1982)), $\sqrt{T}(\hat{\theta} - \theta) \overset{a}{\sim} N(0, \hat{D}^{-1}\hat{S}\hat{D}^{-1\prime})$, where $\hat{D} = (1/T)\sum_{t=L+1}^{T+L} \partial e_t(\hat{\theta})/\partial\hat{\theta}'$ and $\hat{S} = \hat{\Gamma}_0 + \sum_{l=1}^{m}(1 - l/(m + 1))(\hat{\Gamma}_l + \hat{\Gamma}_l')$ with $\hat{\Gamma}_l = (1/T)\sum_{t=L+l+1}^{T+L} e_t(\hat{\theta})e_{t-l}(\hat{\theta})'$. For the empirical analysis, we set $m = 5$ and $\hat{V}_{\lambda} = \hat{V}_{22} - 2\hat{V}_{21} + \hat{V}_{11}$, where $\hat{V} \equiv \hat{D}^{-1}\hat{S}\hat{D}^{-1\prime}$.

[12] See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

large cross-sectional dispersion in average returns. Finally, we consider 2 data sets that are chosen based on their potential to provide improved timing opportunities. The 1st is obtained by using a momentum measure to sort firms into 10 portfolios (10 Momentum). The 2nd is obtained by using estimated return standard deviations to sort firms into 10 portfolios (10 Volatility).[13] All of the portfolios except those in the 10 Volatility data set are value weighted. The 10 Volatility portfolios are formed in January of each year and rebalanced monthly to equal weights.

A. Settings for the Empirical Analysis

Most of the settings for the empirical analysis are described earlier. The one remaining task is to specify η for the timing strategies. Setting $\eta = 1$ is a natural choice for the baseline analysis because it delivers the VT and RRT portfolios that are implied by mean-variance optimization with a diagonal covariance matrix. We also consider 2 other values of the timing aggressiveness parameter: $\eta = 2$ and $\eta = 4$. These choices are motivated by our analysis of the effect of ignoring the estimated return correlations, which suggests setting $\eta > 1$ to mitigate information loss. To document how the timing strategies perform relative to mean-variance optimization, we report results for 4 MVE strategies: the MV, OC, and TP strategies, plus a version of the OC strategy that prohibits short sales (OC^+). We implement the OC and OC^+ strategies by targeting the estimated conditional expected excess return of the $1/N$ portfolio for each date t in the sample.[14]

B. Results for the 10 Industry Data Set

Table 2 documents the out-of-sample performance of the $1/N$, timing, and MVE strategies for the 10 Industry data set. The results for the $1/N$ strategy are reported in the leading row of the table. In the absence of transaction costs, the values of $\hat{\mu}_p$ and $\hat{\sigma}_p$ are 5.80% and 15.04%, respectively, which translates into a $\hat{\lambda}_p$ of 0.386.[15] As expected, the estimated expected turnover is quite low: 2.3% per month. Thus, imposing transaction costs of 50 bp has a minor impact on performance. Specifically, $\hat{\lambda}_p$ falls from 0.386 to 0.376.

The results for the VT strategies are given in Panel A of Table 2. In the absence of transaction costs, $\hat{\lambda}_p$ is 0.426 for $\eta = 1$, 0.454 for $\eta = 2$, and 0.463 for $\eta = 4$. These values exceed the estimated Sharpe ratio for the $1/N$ strategy, and the differences for $\eta = 1$ and $\eta = 2$ are statistically significant at the 5% level. Moreover, the evidence suggests that risk-averse investors could reap substantial benefits from VT. The value of $\hat{\sigma}_p$ for the VT strategies is about 0.8 to 2.2 percentage points lower than for the $1/N$ strategy. As a consequence, $\hat{\Delta}_\gamma$ ranges from

[13] We are grateful to Richard Price for making these data available to us. The volatility portfolios are constructed using the methodology developed in Crawford, Hansen, and Price (2009).

[14] See footnote 3 for additional details.

[15] Note that the results reported in Table 2 for this data set differ from those reported in Table 1 because the sample period is different and because the SIC codes included in each industry were changed in 2004. See the data library Web page (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html) for details.

TABLE 2

Results for the 10 Industry Data Set

Table 2 summarizes the out-of-sample performance of the 1/*N* strategy (row 1), 3 volatility timing (VT) strategies (Panel A), 6 reward-to-risk timing (RRT) strategies (Panel B), and 4 mean-variance efficient (MVE) strategies (Panel C) for the 10 Industry portfolios. It reports the following sample statistics for the time series of monthly excess returns generated by each strategy: the annualized mean ($\widehat{\mu}_p$), the annualized standard deviation ($\widehat{\sigma}_p$), the annualized Sharpe ratio ($\widehat{\lambda}_p$), the average monthly turnover expressed as a fraction of wealth invested ($\widehat{\tau}_p$), the annualized basis point fee that an investor with quadratic utility and constant relative risk aversion of γ = 1 and γ = 5 would be willing to pay to switch from the 1/*N* strategy to the timing or MVE strategy ($\widehat{\Delta}_\gamma$), the *p*-value for the difference between the annualized Sharpe ratio produced by the timing or MVE strategy and the 1/*N* strategy (vs. 1/*N* *p*-val), and the *p*-values (*p*-val) of the basis point fees. The timing and MVE strategies are implemented using a 120-month rolling estimator of the conditional mean vector and conditional covariance matrix of the excess returns, and the OC and OC$^+$ strategies target the estimated conditional expected excess return of the 1/*N* portfolio each period as described in Section II.B. The performance measures are reported assuming no transaction costs and assuming proportional transaction costs of 50 bp, and the *p*-values are determined from 10,000 trials of a stationary block bootstrap with expected block length of 10. The values of $\widehat{\mu}_p$ and $\widehat{\sigma}_p$ are not reported for the results with transaction costs. An entry of "—" for the TP strategy indicates that the corresponding sample statistic cannot be computed. This occurs if there is no real value of the performance fee that makes the investor indifferent between the TP strategy and 1/*N* strategy, or if the turnover for the TP strategy exceeds 20,000% in 1 or more months, which drives wealth to 0 under the assumed level of transaction costs. The sample period is July 1963–December 2008 (546 monthly observations). The first 120 observations are held out to initialize the rolling estimators. See the text for a detailed description of each strategy.

	No Transaction Costs									Transaction Costs = 50 bp					
				vs. 1/*N*							vs. 1/*N*				
Strategy	$\widehat{\mu}_p$	$\widehat{\sigma}_p$	$\widehat{\lambda}_p$	*p*-val	$\widehat{\Delta}_1$	*p*-val	$\widehat{\Delta}_5$	*p*-val	$\widehat{\tau}_p$	$\widehat{\lambda}_p$	*p*-val	$\widehat{\Delta}_1$	*p*-val	$\widehat{\Delta}_5$	*p*-val
1/*N*	5.80	15.04	0.386						0.023	0.376					
Panel A. Volatility Timing Strategies															
VT(1)	6.04	14.18	0.426	0.015	37	0.119	89	0.004	0.024	0.416	0.016	36	0.124	90	0.004
VT(2)	6.12	13.48	0.454	0.040	54	0.196	147	0.012	0.028	0.442	0.046	51	0.209	145	0.013
VT(4)	5.93	12.81	0.463	0.163	44	0.355	172	0.083	0.036	0.446	0.187	36	0.378	166	0.090
Panel B. Reward-to-Risk Timing Strategies															
RRT(μ_t^+,1)	5.01	15.30	0.328	0.908	−82	0.890	−99	0.883	0.077	0.298	0.960	−114	0.949	−131	0.934
RRT(μ_t^+,2)	5.00	15.65	0.320	0.865	−89	0.823	−129	0.880	0.089	0.285	0.929	−128	0.909	−169	0.932
RRT(μ_t^+,4)	5.06	16.35	0.310	0.830	−94	0.767	−181	0.902	0.119	0.266	0.914	−151	0.877	−241	0.957
RRT(β_t^+,1)	6.12	14.63	0.418	0.115	38	0.169	64	0.064	0.030	0.406	0.140	34	0.198	61	0.076
RRT(β_t^+,2)	6.20	14.41	0.430	0.130	49	0.190	88	0.065	0.040	0.413	0.174	39	0.247	78	0.091
RRT(β_t^+,4)	6.17	14.13	0.437	0.181	51	0.260	106	0.096	0.062	0.410	0.272	27	0.366	83	0.154
Panel C. Mean-Variance Efficient Strategies															
MV	6.48	12.81	0.506	0.129	99	0.269	228	0.084	0.161	0.431	0.303	16	0.461	147	0.189
OC	5.96	13.35	0.447	0.290	40	0.407	139	0.211	0.285	0.317	0.683	−118	0.758	−20	0.542
OC$^+$	5.39	13.04	0.414	0.364	−13	0.539	104	0.212	0.109	0.363	0.561	−64	0.698	52	0.347
TP	−14.86	59.02	−0.252	0.995	−3,717	0.985	−13,423	0.754	27.59	—	—	—	—	—	—

89 bp to 172 bp for $\gamma = 5$, and all of the estimates are statistically significant at the 10% level. Of course the gains to VT are smaller for less risk-averse investors. The values of $\hat{\Delta}_\gamma$ range from 37 bp to 54 bp for $\gamma = 1$, and none of the estimates is statistically significant at the 10% level.

Importantly, all of the VT strategies have low estimated expected turnover. It ranges from 2.4% with $\eta = 1$ to 3.6% with $\eta = 4$. These values are close to the estimated expected turnover for the $1/N$ strategy, so imposing transaction costs has little impact on relative performance. With transaction costs of 50 bp, $\hat{\lambda}_p$ falls to 0.416 for $\eta = 1$, 0.442 for $\eta = 2$, and 0.446 for $\eta = 4$, and the range of $\hat{\Delta}_\gamma$ becomes 90 bp to 166 bp for $\gamma = 5$ and 36 bp to 51 bp for $\gamma = 1$. Thus, turnover is not a concern for the VT strategies. The difference in the estimated Sharpe ratios of the VT and $1/N$ strategies remains statistically significant at the 5% level for $\eta = 1$ and $\eta = 2$, and all of the estimated performance fees for $\gamma = 5$ remain statistically significant at the 10% level.

The results for the RRT strategies are given in Panel B of Table 2. First consider the case in which we implement the strategies using the standard rolling estimator of μ_t. The results are clearly less favorable than for the VT strategies: $\hat{\lambda}_p$ is 0.328 for $\eta = 1$, 0.320 for $\eta = 2$, and 0.310 for $\eta = 4$, and all of the estimated performance fees are negative, ranging from -181 bp to -82 bp. Moreover, the estimated expected turnover is higher than for the VT strategies. It ranges from 7.7% per month for $\eta = 1$ to 11.9% per month for $\eta = 4$. Hence, the performance of the RRT strategies deteriorates when we impose transaction costs. The value of $\hat{\lambda}_p$ falls to 0.298 for $\eta = 1$, 0.285 for $\eta = 2$, and 0.266 for $\eta = 4$, and the values of $\hat{\Delta}_\gamma$ range from -241 bp to -114 bp.

Although we find no support for RRT in these results, this may be because the 10 Industry data set poses an especially difficult challenge for the RRT strategies. Even if sorting firms according to SIC codes is a reasonable way to identify different industries, there is no guarantee that the industries will display significant cross-sectional variation in conditional expected excess returns. If the variation in conditional expected excess returns across assets is relatively low, then the estimates of μ_t may convey little useful information. Our analysis suggests that this plays a role in the unimpressive performance of the RRT strategies for the 10 Industry data set.

Consider, for instance, the evidence in Graph A of Figure 1. The plot on the left side of the figure shows the cross-sectional dispersion in the annualized sample mean of the excess industry portfolio returns. The sample means lie in a relatively narrow range: 9%–14%. Moreover, 7 of the values are between 10% and 13%. We would expect to find greater variation in these values if there were substantial cross-sectional dispersion in conditional expected excess returns. In contrast, the plot on the right side of the figure reveals that the dispersion in the sample volatilities is considerably larger than the dispersion in the sample means. The range is from 14.6% to 24.5%. This is consistent with VT outperforming RRT for this data set.

To investigate further, we turn to the case in which we implement the RRT strategies using an estimator of μ_t derived from our 4-factor risk model. Since this estimator should display less sampling variation than the standard rolling estimator of μ_t, we anticipate better performance from the RRT strategies. The

FIGURE 1

Reward and Risk Characteristics of the Data Sets

Figure 1 summarizes the sample reward and risk characteristics for the 10 Industry data set (Graph A), 25 Size/BTM data set (Graph B), 10 Momentum data set (Graph C), and 10 Volatility data set (Graph D). The 1st graph in each panel shows the cross section of annualized mean returns, and the 2nd graph shows the cross section of annualized return standard deviations. The sample period is July 1963–December 2008 (546 monthly observations). However, the reported statistics correspond to the subperiod used to evaluate the out-of-sample performance of the portfolio strategies (i.e., observations 121–546).

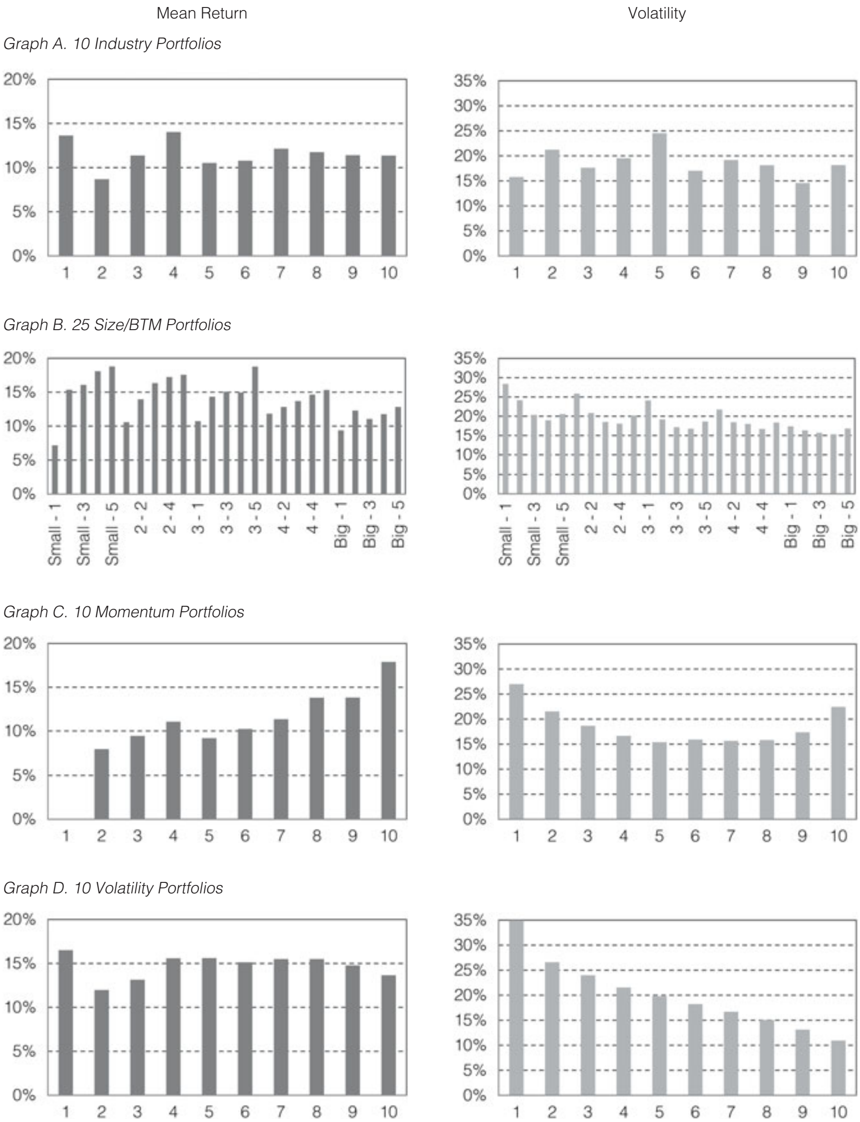

improvement in performance is substantial. In the absence of transaction costs, $\hat{\lambda}_p$ is 0.418 for $\eta = 1$, 0.430 for $\eta = 2$, and 0.437 for $\eta = 4$. The values of $\hat{\Delta}_\gamma$ range from 38 bp to 51 bp for $\gamma = 1$ and from 64 bp to 106 bp for $\gamma = 5$. These results are comparable to those for the VT strategies. Although the differences in the estimated Sharpe ratios of the RRT and $1/N$ strategies are not statistically significant, the estimated performance fees for $\gamma = 5$ are statistically significant at the 10% level.

These findings point to a sizeable reduction in estimation risk from using the 4-factor risk model. Additional evidence of this reduction can be found in the estimates of expected turnover, which are roughly ½ as large as those generated by using the standard rolling estimator of μ_t. The range is from 3.0% per month for $\eta = 1$ to 6.2% per month for $\eta = 4$. As a consequence, imposing transaction costs has a minor impact on performance. The value of $\hat{\lambda}_p$ falls to 0.406 for $\eta = 1$, 0.413 for $\eta = 2$, and 0.410 for $\eta = 4$ and the range of $\hat{\Delta}_\gamma$ becomes 27 bp to 39 bp for $\gamma = 1$ and 61 bp to 83 bp for $\gamma = 5$. Two of the 3 values of $\hat{\Delta}_\gamma$ for $\gamma = 5$ remain statistically significant at the 10% level. Hence, we find that the RRT strategies are capable of outperforming naïve diversification, provided they are implemented using our alternative estimator of μ_t.

The evidence in Panel C of Table 2, which documents the performance of the MVE strategies, lends additional context to these findings. In the absence of transaction costs, the MVE strategies deliver mixed results. The TP strategy displays the worst performance by far, with a negative $\hat{\mu}_p$ and a value of $\hat{\sigma}_p$ that exceeds the estimated volatility of every other MVE strategy by a factor of 4. Its estimated Sharpe ratio is exceedingly low: -0.252. In comparison, the MV and OC strategies have estimated Sharpe ratios of 0.506 and 0.447, respectively. Both outperform the $1/N$ strategy, although the differences in the estimated Sharpe ratios are not statistically significant. Prohibiting short sales causes the performance of the OC strategy to deteriorate. Its estimated Sharpe ratio falls to 0.414.

It may seem odd that the performance gains for the MV and OC strategies, which are larger than those for the VT strategies, are statistically insignificant. This is explained by the correlation between the excess returns for the different strategies. For example, the estimated correlation between the excess returns for the $1/N$ and VT(1) strategies is 0.99, while that for the $1/N$ and OC strategies is 0.74. Since the standard error of $\hat{\lambda}_{p_j} - \hat{\lambda}_{p_i}$ is decreasing in this correlation, it takes a larger difference in the estimated Sharpe ratios to conclude that the OC strategy outperforms the $1/N$ strategy. This is also the case for the estimated performance fees. The values of $\hat{\Delta}_\gamma$ for the MV and OC strategies range from 40 bp to 228 bp, but only one is statistically significant at the 10% level.

High turnover is the most troublesome aspect of the performance of the MVE strategies. The TP strategy is clearly an outlier in this respect, with an estimated expected turnover of 2,759% per month. This value exceeds the estimated expected turnover for the other MVE strategies by 2 orders of magnitude. Indeed, there are a number of months in which the turnover of the TP strategy exceeds 20,000%. Since this level of turnover implies that rebalancing costs would consume all of the investor's wealth, it is impossible to compute meaningful after-transaction-costs values of the performance measures. This is indicated by a "—" entry in Panel C of Table 2.

The remaining estimates of expected turnover are much more reasonable: 16.1% for the MV strategy, 28.5% for the OC strategy, and 10.9% for the OC^+ strategy. Nonetheless, turnover is sufficiently high that there is substantial deterioration in the performance of the strategies in the presence of transaction costs.

The estimated Sharpe ratios for the MV and OC strategies fall to 0.431 and 0.317, respectively, and the estimated performance fees are -118 bp and 16 bp for $\gamma = 1$ and -20 bp and 147 bp with $\gamma = 5$. None of the performance gains is statistically significant at the 10% level. Note also that once we account for transaction costs, the OC$^+$ strategy has a higher Sharpe ratio than the OC strategy because prohibiting short sales leads to a substantial reduction in turnover.

C. Results for the 25 Size/BTM Data Set

The results for the 10 Industry data set foreshadow the crucial role of turnover in our investigation. With the exception of the TP strategy, all of the MVE strategies perform better than naïve diversification in the absence of transaction costs. Although the analysis is inconclusive because most of the performance gains are statistically insignificant, the picture that emerges from the initial evidence is generally supportive of mean-variance optimization. This is no longer true in the presence of transaction costs because the advantage of the MVE strategies is eroded by high turnover. The results indicate, therefore, that controlling turnover is key to improving mean-variance methods of portfolio selection. The timing strategies are largely successful in this regard for the 10 Industry data set, but additional evidence is needed to draw firm conclusions.

To develop this evidence, we turn to the 25 Size/BTM data set. If our hypothesis regarding the relationship between the cross-sectional dispersion in conditional expected returns and the performance of the RRT strategies is correct, then we should find stronger support for RRT with this data set. As Graph B of Figure 1 shows, the annualized sample mean of the excess returns for the size and book-to-market portfolios ranges from 7.2% to 18.8%, which is more than twice the range for the industry portfolios. The estimates of μ_t should therefore have more value in this case than for the 10 Industry data set. We anticipate little change in the performance of VT strategies, since the range of the annualized sample volatilities is 15.3% to 28.4%, which is only a few percentage points larger than that for the industry portfolios.

Table 3 documents the out-of-sample performance of the $1/N$, timing, and MVE strategies for the 25 Size/BTM data set. The layout of the table is identical to that of Table 2. In the absence of transaction costs, the values of $\hat{\mu}_p$ and $\hat{\sigma}_p$ for the $1/N$ strategy are 8.25% and 17.64%, respectively, which translates into a $\hat{\lambda}_p$ of 0.468. Once again this strategy has low estimated expected turnover (1.7% per month), so the impact of imposing transaction costs is quite small. The value of $\hat{\lambda}_p$ falls to 0.462.

The results for VT strategies, given in Panel A of Table 3, are similar to those for the 10 Industry data set. In the absence of transaction costs, VT delivers higher values of $\hat{\lambda}_p$ than does naïve diversification: 0.492 for $\eta = 1$, 0.496 for $\eta = 2$, and 0.484 for $\eta = 4$. The difference in the estimated Sharpe ratios is statistically significant at the 10% level for $\eta = 1$. The estimated performance fees range from -39 bp to 15 bp for $\gamma = 1$ and from 81 bp to 111 bp for $\gamma = 5$. In the latter case, all of the fees are statistically significant at the 10% level. Because the VT strategies have low turnover (1.8%–3.1% per month), imposing transaction costs does not alter the nature of these findings.

TABLE 3

Results for the 25 Size/BTM Data Set

Table 3 summarizes the out-of-sample performance of the 1/*N* strategy (row 1), 3 volatility timing (VT) strategies (Panel A), 6 reward-to-risk timing (RRT) strategies (Panel B), and 4 mean-variance efficient (MVE) strategies (Panel C) for the 25 Size/BTM portfolios. It reports the following sample statistics for the time series of monthly excess returns generated by each strategy: the annualized mean ($\widehat{\mu}_p$), the annualized standard deviation ($\widehat{\sigma}_p$), the annualized Sharpe ratio ($\widehat{\lambda}_p$), the average monthly turnover expressed as a fraction of wealth invested ($\widehat{\tau}_p$), the annualized basis point fee that an investor with quadratic utility and constant relative risk aversion of $\gamma = 1$ and $\gamma = 5$ would be willing to pay to switch from the 1/*N* strategy to the timing or MVE strategy ($\widehat{\Delta}_\gamma$), the *p*-value for the difference between the annualized Sharpe ratio produced by the timing or MVE strategy and the 1/*N* strategy (vs. 1/*N* *p*-val), and the *p*-values (*p*-val) of the basis point fees. The timing and MVE strategies are implemented using a 120-month rolling estimator of the conditional mean vector and conditional covariance matrix of the excess returns, and the OC and OC$^+$ strategies target the estimated conditional expected excess return of the 1/*N* portfolio each period as described in Section II.B. The performance measures are reported assuming no transaction costs and assuming proportional transaction costs of 50 bp, and the *p*-values are determined from 10,000 trials of a stationary block bootstrap with expected block length of 10. The values of $\widehat{\mu}_p$ and $\widehat{\sigma}_p$ are not reported for the results with transaction costs. An entry of "—" for the TP strategy indicates that the corresponding sample statistic cannot be computed. This occurs if there is no real value of the performance fee that makes the investor indifferent between the TP strategy and 1/*N* strategy, or if the turnover for the TP strategy exceeds 20,000% in 1 or more months, which drives wealth to 0 under the assumed level of transaction costs. The sample period is July 1963–December 2008 (546 monthly observations). The first 120 observations are held out to initialize the rolling estimators. See the text for a detailed description of each strategy.

	No Transaction Costs								Transaction Costs = 50 bp						
				vs. 1/*N*							vs. 1/*N*				
Strategy	$\widehat{\mu}_p$	$\widehat{\sigma}_p$	$\widehat{\lambda}_p$	*p*-val	$\widehat{\Delta}_1$	*p*-val	$\widehat{\Delta}_5$	*p*-val	$\widehat{\tau}_p$	$\widehat{\lambda}_p$	*p*-val	$\widehat{\Delta}_1$	*p*-val	$\widehat{\Delta}_5$	*p*-val
1/*N*	8.25	17.64	0.468						0.017	0.462					
Panel A. Volatility Timing Strategies															
VT(1)	8.24	16.74	0.492	0.074	15	0.312	81	0.002	0.018	0.486	0.078	14	0.317	81	0.002
VT(2)	8.01	16.15	0.496	0.178	2	0.498	109	0.012	0.021	0.488	0.192	−1	0.518	108	0.013
VT(4)	7.51	15.52	0.484	0.382	−39	0.698	111	0.081	0.031	0.472	0.430	−47	0.731	104	0.093
Panel B. Reward-to-Risk Timing Strategies															
RRT(μ_t^+,1)	9.17	17.15	0.535	0.001	101	0.002	137	0.000	0.036	0.522	0.001	90	0.004	127	0.000
RRT(μ_t^+,2)	9.41	16.96	0.555	0.001	128	0.004	178	0.000	0.050	0.537	0.004	108	0.015	159	0.001
RRT(μ_t^+,4)	9.58	16.84	0.569	0.003	147	0.013	206	0.001	0.080	0.540	0.019	109	0.050	167	0.004
RRT($\bar{\beta}_t^+$,1)	9.00	17.17	0.524	0.001	83	0.004	118	0.000	0.018	0.518	0.001	83	0.004	118	0.000
RRT($\bar{\beta}_t^+$,2)	9.39	16.93	0.555	0.001	126	0.004	179	0.000	0.020	0.547	0.001	124	0.005	178	0.000
RRT($\bar{\beta}_t^+$,4)	9.68	16.72	0.579	0.002	158	0.008	226	0.000	0.029	0.568	0.004	152	0.011	220	0.000
Panel C. Mean-Variance Efficient Strategies															
MV	11.45	13.43	0.853	0.021	386	0.068	664	0.003	0.789	0.496	0.443	−81	0.641	193	0.224
OC	14.19	14.11	1.006	0.006	651	0.018	892	0.001	0.996	0.569	0.300	56	0.435	283	0.175
OC$^+$	7.18	15.12	0.475	0.455	−65	0.752	110	0.143	0.164	0.409	0.789	−154	0.941	20	0.432
TP	−58.80	777.8	−0.076	0.982	—	—	—	—	233.8	—	—	—	—	—	—

https://doi.org/10.1017/S0022109012000117 Published online by Cambridge University Press

As anticipated, the evidence on RRT, given in Panel B of Table 3, is more compelling for the 25 Size/BTM data set than for the 10 Industry data set. In the absence of transaction costs, using the standard rolling estimator of μ_t produces a $\hat{\lambda}_p$ of 0.535 for $\eta = 1$, 0.555 for $\eta = 2$, and 0.569 for $\eta = 4$. The increase in the estimated Sharpe ratio relative to naïve diversification is statistically significant at 1% in each case. Moreover, all the estimated performance fees, which range from 101 bp to 147 bp for $\gamma = 1$ and from 137 bp to 206 bp for $\gamma = 5$, are statistically significant at 5%. Because the gains are achieved with little increase in turnover relative to the $1/N$ strategy, they remain statistically significant in the presence of transaction costs, with estimated performance fees that range from 90 bp to 109 bp for $\gamma = 1$ and from 127 bp to 167 bp for $\gamma = 5$.

Interestingly, switching to the estimator of μ_t derived from our 4-factor risk model has little effect on the estimated Sharpe ratios for the RRT strategies. In the absence of transaction costs, the value of $\hat{\lambda}_p$ is 0.524 for $\eta = 1$, 0.555 for $\eta = 2$, and 0.579 for $\eta = 4$. Because these values are nearly identical to those obtained using the rolling estimator of μ_t, it may appear that there is no benefit from employing the risk model. This is not the case, however. The estimated expected turnover is 1.8% for $\eta = 1$, 2.0% for $\eta = 2$, and 2.9% for $\eta = 4$. These values are 2 to 3 times lower than those obtained using the rolling estimator of μ_t. Thus, the risk model achieves substantial reductions in turnover while delivering comparable performance gains.

With such low levels of turnover, imposing transaction costs has little impact on the performance of the RRT strategies. The value of $\hat{\lambda}_p$ falls to 0.518 for $\eta = 1$, 0.547 for $\eta = 2$, and 0.568 for $\eta = 4$. In each case, the increase in the estimated Sharpe ratio relative to naïve diversification is statistically significant at the 1% level. All of the estimated performance fees, which range from 83 bp to 152 bp for $\gamma = 1$ and from 118 bp to 220 bp for $\gamma = 5$, are statistically significant at the 5% level, and all but one are statistically significant at the 1% level.

The evidence for the MVE strategies, given in Panel C of Table 3, provides additional perspective on these results. The performance of the TP strategy is again exceedingly poor. In the absence of transaction costs, it has an estimated annualized mean return of −58.8% and an estimated annualized volatility of 777.8%. This translates into an estimated Sharpe ratio of −0.076. Note that we report "—" in Table 3 for the estimated performance fees. This is because the results for the TP strategy are so poor that the quadratic equation used to find the performance fees has no real roots. Thus, there is no fixed fee that equates the estimated expected utilities generated by the TP and $1/N$ strategies. The TP strategy is also characterized by extreme estimated expected turnover: 23,380% per month.

In contrast, the MV and OC strategies perform well in the absence of transaction costs. Their estimated Sharpe ratios (0.853 and 1.006) are much higher than the estimated Sharpe ratio of the $1/N$ strategy, and the differences are statistically significant at the 5% level. Moreover, the estimated performance fees for the MV and OC strategies, which are 386 bp and 651 bp for $\gamma = 1$ and 664 bp and 892 bp for $\gamma = 5$, respectively, are statistically significant at the 10% level. Thus, the evidence suggests that mean-variance optimization is superior to naïve diversification if transaction costs are 0.

Notice, however, that imposing transaction costs leads to a marked reduction in the performance of the MVE strategies. The estimated expected turnover for the MV and OC strategies is 79% and 100%, respectively, per month. As a consequence, $\hat{\lambda}_p$ falls to 0.496 for the MV strategy and to 0.569 for the OC strategy. These values are still higher than the estimated Sharpe ratio for the $1/N$ strategy, but the differences are no longer statistically significant. This is also true for the estimated performance fees. Unlike the timing strategies, the MVE strategies simply generate too much turnover to be competitive with naïve diversification.

D. Results for the 10 Momentum Data Set

In view of the results for the 25 Size/BTM portfolios, it seems clear that performance of the timing strategies is influenced by data set characteristics. This is not a surprise. Intuitively, we would expect the effectiveness of VT and RRT to depend on both the cross-sectional and time-series variation in the conditional means and volatilities. These characteristics are undoubtedly affected by the scheme used to sort firms into portfolios. In the case of the 25 Size/BTM data set, the sorting scheme is explicitly designed to increase the cross-sectional dispersion in conditional expected returns. To the extent that it does so, we should see an improvement in the signal-to-noise ratio of our rolling estimator of μ_t, which should benefit both the timing and MVE strategies.

To see if these findings hold more generally, we consider a data set obtained by sorting firms into portfolios using a momentum measure. This is another case in which the sorting scheme is explicitly designed to spread conditional expected returns (see, e.g., Jegadeesh and Titman (1993)). The evidence in Graph C of Figure 1 suggests that it succeeds in this respect. The annualized sample mean for the momentum portfolios ranges from 0% to 17.9%. Since this is larger than the range for the size and book-to-market portfolios, we again anticipate that the RRT strategies will perform well relative to naïve diversification. The range of the annualized sample volatilities (15.4% to 27.0%) is comparable to that for both the 25 Size/BTM and 10 Industry data sets.

Table 4 documents the out-of-sample performance of the $1/N$, timing, and MVE strategies for the 10 Momentum data set. In the absence of transaction costs, the values of $\hat{\mu}_p$, $\hat{\sigma}_p$, and $\hat{\lambda}_p$ for the $1/N$ strategy are 4.70%, 16.68%, and 0.282, respectively. Imposing transaction costs reduces the value of $\hat{\lambda}_p$ to 0.276. In comparison, the value of $\hat{\lambda}_p$ for the VT strategies, given in Panel A of Table 4, ranges from 0.330 to 0.375 with no transaction costs, and the increase in the estimated Sharpe ratio relative to naïve diversification is statistically significant at the 1% level in each case. These gains translate into estimated performance fees of 68 bp to 128 bp for $\gamma = 1$ and 117 bp to 215 bp for $\gamma = 5$, all of which are statistically significant at the 5% level. Because the estimated expected turnover is only 1.7% to 2.6% per month, imposing transaction costs has little effect on the results.

The performance of the RRT strategies, given in Panel B of Table 4, is even more compelling. In the absence of transaction costs, the standard rolling estimator of μ_t produces a $\hat{\lambda}_p$ of 0.455 for $\eta = 1$, 0.476 for $\eta = 2$, and 0.497 for $\eta = 4$. The differences relative to naïve diversification are statistically significant at the 1% level. This is also the case for the estimated performance fees, which range

TABLE 4

Results for the 10 Momentum Data Set

Table 4 summarizes the out-of-sample performance of the $1/N$ strategy (row 1), 3 volatility timing (VT) strategies (Panel A), 6 reward-to-risk timing (RRT) strategies (Panel B), and 4 mean-variance efficient (MVE) strategies (Panel C) for the 10 Momentum portfolios. It reports the following sample statistics for the time series of monthly excess returns generated by each strategy: the annualized mean ($\widehat{\mu}_p$), the annualized standard deviation ($\widehat{\sigma}_p$), the annualized Sharpe ratio ($\widehat{\lambda}_p$), the average monthly turnover expressed as a fraction of wealth invested ($\widehat{\tau}_p$), the annualized basis point fee that an investor with quadratic utility and constant relative risk aversion of $\gamma = 1$ and $\gamma = 5$ would be willing to pay to switch from the $1/N$ strategy to the timing or MVE strategy ($\widehat{\Delta}_\gamma$), the p-value for the difference between the annualized Sharpe ratio produced by the timing or MVE strategy and the $1/N$ strategy (vs. $1/N$ p-val), and the p-values (p-val) of the basis point fees. The timing and MVE strategies are implemented using a 120-month rolling estimator of the conditional mean vector and conditional covariance matrix of the excess returns, and the OC and OC$^+$ strategies target the estimated conditional expected excess return of the $1/N$ portfolio each period as described in Section II.B. The performance measures are reported assuming no transaction costs and assuming proportional transaction costs of 50 bp, and the p-values are determined from 10,000 trials of a stationary block bootstrap with expected block length of 10. The values of $\widehat{\mu}_p$ and $\widehat{\sigma}_p$ are not reported for the results with transaction costs. An entry of "—" for the TP strategy indicates that the corresponding sample statistic cannot be computed. This occurs if there is no real value of the performance fee that makes the investor indifferent between the TP strategy and $1/N$ strategy, or if the turnover for the TP strategy exceeds 20,000% in 1 or more months, which drives wealth to 0 under the assumed level of transaction costs. The sample period is July 1963–December 2008 (546 monthly observations). The first 120 observations are held out to initialize the rolling estimators. See the text for a detailed description of each strategy.

	No Transaction Costs									Transaction Costs = 50 bp					
				vs. $1/N$							vs. $1/N$				
Strategy	$\widehat{\mu}_p$	$\widehat{\sigma}_p$	$\widehat{\lambda}_p$	p-val	$\widehat{\Delta}_1$	p-val	$\widehat{\Delta}_5$	p-val	$\widehat{\tau}_p$	$\widehat{\lambda}_p$	p-val	$\widehat{\Delta}_1$	p-val	$\widehat{\Delta}_5$	p-val
$1/N$	4.70	16.68	0.282						0.018	0.276					
Panel A. Volatility Timing Strategies															
VT(1)	5.27	15.96	0.330	0.004	68	0.017	117	0.001	0.017	0.324	0.004	69	0.017	119	0.001
VT(2)	5.55	15.63	0.355	0.004	102	0.018	173	0.001	0.018	0.349	0.004	102	0.019	174	0.001
VT(4)	5.77	15.38	0.375	0.005	128	0.022	215	0.001	0.026	0.365	0.006	123	0.026	211	0.002
Panel B. Reward-to-Risk Timing Strategies															
RRT(μ_t^+,1)	7.62	16.75	0.455	0.000	290	0.000	285	0.000	0.058	0.434	0.000	266	0.000	260	0.000
RRT(μ_t^+,2)	8.14	17.09	0.476	0.000	336	0.000	306	0.001	0.061	0.454	0.000	310	0.000	279	0.002
RRT(μ_t^+,4)	8.64	17.41	0.497	0.000	381	0.000	328	0.002	0.068	0.473	0.001	351	0.001	298	0.005
RRT($\bar{\beta}_t^+$,1)	6.40	16.08	0.398	0.000	179	0.000	220	0.000	0.019	0.391	0.000	178	0.000	220	0.000
RRT($\bar{\beta}_t^+$,2)	7.23	16.19	0.447	0.000	261	0.000	295	0.000	0.024	0.438	0.000	257	0.000	292	0.000
RRT($\bar{\beta}_t^+$,4)	7.87	16.47	0.478	0.000	320	0.000	335	0.001	0.034	0.466	0.000	310	0.000	325	0.001
Panel C. Mean-Variance Efficient Strategies															
MV	7.47	15.05	0.496	0.051	302	0.068	409	0.036	0.281	0.385	0.191	144	0.242	253	0.132
OC	9.18	15.59	0.589	0.015	465	0.018	536	0.013	0.348	0.454	0.096	266	0.112	337	0.081
OC$^+$	6.49	15.71	0.413	0.007	195	0.014	261	0.004	0.129	0.363	0.044	127	0.072	193	0.031
TP	8,321	5,680	0.146	0.680	—	—	—	—	184.9	—	—	—	—	—	—

from 285 bp to 381 bp. The estimates of expected turnover (5.8%, 6.1%, and 6.8%) are higher than for the $1/N$ strategy, but the differences are not dramatic. Hence, imposing transaction costs does not alter our inferences.

The results using the estimator of μ_t derived from our 4-factor risk model are similar. In the absence of transaction costs, the value of $\hat{\lambda}_p$ is 0.398 for $\eta = 1$, 0.447 for $\eta = 2$, and 0.478 for $\eta = 4$, and the estimated performance fees range from 179 bp to 335 bp. All of the gains are statistically significant at the 1% level. As with the other data sets, the estimates of expected turnover are 2 to 3 times lower than those produced by the rolling estimator of μ_t. And, again, the impact of imposing transaction costs is minimal.

We also find that the performance of the MVE strategies, given in Panel C of Table 4, is more competitive for this data set. The TP strategy clearly delivers extreme results: an estimated annualized mean return of 8,321%, an estimated annualized volatility of 5,680%, and an estimated expected monthly turnover of 18,490%. However, the MV and OC strategies have estimated Sharpe ratios of 0.496 and 0.589, respectively, in the absence of transaction costs, and the estimated performance fees range from 302 bp to 536 bp. All of the gains are statistically significant at the 10% level. Turnover is still an issue, but the estimates are considerably lower than for the 25 Size/BTM data set. Thus, the majority of the performance gains remain statistically significant in the presence of transaction costs.

E. Results for the 10 Volatility Data Set

The evidence for the 25 Size/BTM and 10 Momentum data sets suggests that the performance of RRT strategies is related to the cross-sectional dispersion in conditional expected excess returns. Accordingly, we posit that the performance of the VT strategies is related to the cross-sectional dispersion in conditional return volatilities. To test this hypothesis, we consider a final data set that is obtained by sorting firms into portfolios based on estimates of historical volatility. Graph D of Figure 1 shows that the resulting cross-sectional dispersion in the annualized sample volatilities, which range from 10.9% to 34.9%, is approximately twice that for the other 3 data sets. The range for the annualized sample means, on the other hand, is relatively narrow at 12.0% to 16.5%.

Table 5 documents the out-of-sample performance of the $1/N$, timing, and MVE strategies for the 10 Volatility data set. As anticipated, the VT strategies perform particularly well. In the absence of transaction costs, the $1/N$ strategy has an estimated Sharpe ratio of 0.477 and an estimated expected turnover of 1.7%. The VT strategies have estimated Sharpe ratios that range from 0.576 to 0.712 and estimated expected turnovers that range from 1.5% to 1.6%, as given in Panel A of Table 5. The reductions in $\hat{\sigma}_p$ relative to naïve diversification, which are in the 3–7 percentage point range, would be quite valuable to risk-averse investors, even in the presence of transaction costs. The estimated performance fees range from 296 bp to 498 bp with $\gamma = 5$, and all of the estimates are statistically significant at the 1% level.

The RRT strategies also perform well, as indicated in Panel B of Table 5. In the absence of transaction costs, the value of $\hat{\lambda}_p$ ranges from 0.547 to 0.647 using

TABLE 5

Results for the 10 Volatility Data Set

Table 5 summarizes the out-of-sample performance of the $1/N$ strategy (row 1), 3 volatility timing (VT) strategies (Panel A), 6 reward-to-risk timing (RRT) strategies (Panel B), and 4 mean-variance efficient (MVE) strategies (Panel C) for the 10 Volatility portfolios. It reports the following sample statistics for the time series of monthly excess returns generated by each strategy: the annualized mean ($\widehat{\mu}_p$), the annualized standard deviation ($\widehat{\sigma}_p$), the annualized Sharpe ratio ($\widehat{\lambda}_p$), the average monthly turnover expressed as a fraction of wealth invested ($\widehat{\tau}_p$), the annualized basis point fee that an investor with quadratic utility and constant relative risk aversion of $\gamma = 1$ and $\gamma = 5$ would be willing to pay to switch from the $1/N$ strategy to the timing or MVE strategy ($\widehat{\Delta}_\gamma$), the p-value for the difference between the annualized Sharpe ratio produced by the timing or MVE strategy and the $1/N$ strategy (vs. $1/N$ p-val), and the p-values (p-val) of the basis point fees. The timing and MVE strategies are implemented using a 120-month rolling estimator of the conditional mean vector and conditional covariance matrix of the excess returns, and the OC and OC^+ strategies target the estimated conditional expected excess return of the $1/N$ portfolio each period as described in Section II.B. The performance measures are reported assuming no transaction costs and assuming proportional transaction costs of 50 bp, and the p-values are determined from 10,000 trials of a stationary block bootstrap with expected block length of 10. The values of $\widehat{\mu}_p$ and $\widehat{\sigma}_p$ are not reported for the results with transaction costs. An entry of "—" for the TP strategy indicates that the corresponding sample statistic cannot be computed. This occurs if there is no real value of the performance fee that makes the investor indifferent between the TP strategy and $1/N$ strategy, or if the turnover for the TP strategy exceeds 20,000% in 1 or more months, which drives wealth to 0 under the assumed level of transaction costs. The sample period is July 1963–December 2008 (546 monthly observations). The first 120 observations are held out to initialize the rolling estimators. See the text for a detailed description of each strategy.

	No Transaction Costs									Transaction Costs = 50 bp						
				vs. 1/N							vs. 1/N					
Strategy	$\widehat{\mu}_p$	$\widehat{\sigma}_p$	$\widehat{\lambda}_p$	p-val	$\widehat{\Delta}_1$	p-val	$\widehat{\Delta}_5$	p-val	$\widehat{\tau}_p$	$\widehat{\lambda}_p$	p-val	$\widehat{\Delta}_1$	p-val	$\widehat{\Delta}_5$	p-val
$1/N$	8.96	18.79	0.477						0.017	0.472					
Panel A. Volatility Timing Strategies															
VT(1)	8.83	15.34	0.576	0.006	46	0.298	296	0.000	0.015	0.570	0.006	47	0.293	297	0.000
VT(2)	8.58	13.17	0.652	0.010	52	0.357	432	0.002	0.015	0.645	0.010	53	0.354	433	0.001
VT(4)	8.17	11.47	0.712	0.023	32	0.438	498	0.007	0.016	0.703	0.024	32	0.437	498	0.007
Panel B. Reward-to-Risk Timing Strategies															
RRT(μ_t^+,1)	8.81	16.12	0.547	0.045	32	0.360	230	0.004	0.029	0.536	0.063	24	0.395	224	0.005
RRT(μ_t^+,2)	8.79	14.70	0.598	0.032	51	0.340	341	0.004	0.042	0.581	0.050	36	0.389	327	0.005
RRT(μ_t^+,4)	8.62	13.32	0.647	0.031	54	0.364	426	0.005	0.066	0.617	0.061	24	0.437	395	0.009
RRT(β_t^+,1)	8.88	17.07	0.520	0.017	22	0.329	154	0.001	0.017	0.514	0.017	22	0.326	154	0.001
RRT(β_t^+,2)	8.78	15.71	0.559	0.011	35	0.333	262	0.001	0.019	0.552	0.011	34	0.338	260	0.001
RRT(β_t^+,4)	8.51	13.67	0.623	0.012	38	0.380	390	0.002	0.027	0.611	0.015	32	0.399	384	0.002
Panel C. Mean-Variance Efficient Strategies															
MV	8.09	10.56	0.766	0.078	34	0.462	541	0.034	0.350	0.566	0.324	−168	0.725	340	0.129
OC	10.84	11.66	0.930	0.008	296	0.146	751	0.006	0.547	0.641	0.176	−25	0.538	425	0.072
OC^+	8.31	13.52	0.615	0.032	21	0.448	381	0.006	0.237	0.508	0.322	−113	0.786	247	0.054
TP	138.5	242.6	0.571	0.338	−21,691	0.904	—	—	191.1	—	—	—	—	—	—

the standard rolling estimator of μ_t and from 0.520 to 0.623 using the estimator of μ_t derived from the 4-factor risk model. In each case, however, the RRT strategy for a given η performs worse than the corresponding VT strategy. Apparently, the μ_t estimates convey little information for the volatility portfolios, which is not surprising given the low cross-sectional dispersion in the sample means. The estimates of expected turnover are similar to those for the other data sets: 2.9%–6.6% using the standard rolling estimator and 1.7%–2.7% using the 4-factor model.

The results for the MVE strategies, given in Panel C of Table 5, follow the same general pattern as for previous data sets. If we ignore the TP strategy, the evidence indicates that turnover is the primary barrier to outperforming naïve diversification. In the absence of transaction costs, the MV and OC strategies have estimated Sharpe ratios of 0.766 and 0.930, respectively, and the differences relative to naïve diversification are statistically significant at the 10% level. The estimated performance fees range from 34 bp to 751 bp. However, the estimates of expected turnover are 35% and 54.7% per month, so performance deteriorates sharply in the presence of transaction costs. The estimated Sharpe ratios fall to 0.566 and 0.641, and the differences relative to naïve diversification are no longer statistically significant.

VI. Closing Remarks

DeMiguel et al. (2009) raise serious questions about the value of mean-variance optimization. Using a range of data sets, they investigate the out-of-sample performance of the standard mean-variance model along with a large number of variants developed in the literature to mitigate the impact of estimation risk. They find that "there is no single model that consistently delivers a Sharpe ratio or a CEQ return that is higher than that of the $1/N$ portfolio" and conclude that "there are still many 'miles to go' before the gains promised by optimal portfolio choice can actually be realized out of sample."

We show that their analysis casts mean-variance optimization in such an unfavorable light largely because the research design implicitly targets a conditional expected return that greatly exceeds the conditional expected return of the $1/N$ strategy. This magnifies estimation risk and leads to excessive turnover. If the mean-variance model is instead implemented by targeting the conditional expected return of the $1/N$ strategy, it generally performs better than naïve diversification absent transaction costs. It is only after we consider the differences in transaction costs across strategies that the mean-variance model has difficulty outperforming the $1/N$ strategy by statistically significant margins.

Motivated by these findings, we propose 2 alternative methods of mean-variance portfolio selection (VT and RRT) that exploit sample information in a manner that mitigates the impact of estimation risk. Importantly, these methods allow us to exercise some control over turnover, and hence transaction costs, via a tuning parameter that can be interpreted as a measure of timing aggressiveness. We find that both types of timing strategies outperform naïve diversification for a range of data sets. This is true even after we incorporate high transaction costs. RRT appears to be a particularly promising strategy when it is implemented using estimates of conditional expected returns obtained from a 4-factor risk model.

References

Balduzzi, P., and A. W. Lynch. "Transaction Costs and Predictability: Some Utility Cost Calculations." *Journal of Financial Economics*, 52 (1999), 47–78.

Brown, S. "Optimal Portfolio Choice under Uncertainty: A Bayesian Approach." In *Estimation Risk and Optimal Portfolio Choice*, V. Bawa, S. Brown, and R. Klein, eds. Amsterdam: North Holland (1979).

Carhart, M. M. "On Persistence in Mutual Fund Performance." *Journal of Finance*, 52 (1997), 57–82.

Crawford, S.; J. Hansen; and R. Price. "CRSP Portfolio Methodology and the Effect on Benchmark Returns." Working Paper, Rice University (2009).

DeMiguel, V.; L. Garlappi; and R. Uppal. "Optimal versus Naive Diversification: How Inefficient Is the $1/N$ Portfolio Strategy." *Review of Financial Studies*, 22 (2009), 1915–1953.

Fama, E. F., and K. R. French. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, 33 (1993), 3–56.

Fleming, J.; C. Kirby; and B. Ostdiek. "The Economic Value of Volatility Timing." *Journal of Finance*, 56 (2001), 329–352.

Fleming, J.; C. Kirby; and B. Ostdiek. "The Economic Value of Volatility Timing Using 'Realized' Volatility." *Journal of Financial Economics*, 67 (2003), 473–509.

Garlappi, L.; R. Uppal; and T. Wang. "Portfolio Selection with Parameter and Model Uncertainty: A Multi-Prior Approach." *Review of Financial Studies*, 20 (2007), 41–81.

Hamilton, J. D. *Time Series Analysis*. Princeton, NJ: Princeton University Press (1994).

Hansen, L. P. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, 50 (1982), 1029–1054.

Jagannathan, R., and T. Ma. "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps." *Journal of Finance*, 58 (2003), 1651–1683.

Jegadeesh, N., and S. Titman. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *Journal of Finance*, 48 (1993), 65–91.

Jorion, P. "International Portfolio Diversification with Estimation Risk." *Journal of Business*, 58 (1985), 259–278.

Jorion, P. "Bayesian and CAPM Estimators of the Means: Implications for Portfolio Selection." *Journal of Banking and Finance*, 15 (1991), 717–727.

Kan, R., and G. Zhou. "Optimal Portfolio Choice with Parameter Uncertainty." *Journal of Financial and Quantitative Analysis*, 42 (2007), 621–656.

Ledoit, O., and M. Wolf. "Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection." *Journal of Empirical Finance*, 10 (2003), 603–621.

Ledoit, O., and M. Wolf. "Honey, I Shrunk the Sample Covariance Matrix: Problems in Mean-Variance Optimization." *Journal of Portfolio Management*, 30 (2004), 110–119.

MacKinlay, A. C., and L. Pastor. "Asset Pricing Models: Implications for Expected Returns and Portfolio Selection." *Review of Financial Studies*, 13 (2000), 883–916.

Merton, R. C. "On Estimating the Expected Return on the Market: An Exploratory Investigation." *Journal of Financial Economics*, 8 (1980), 323–361.

Pastor, L. "Portfolio Selection and Asset Pricing Models." *Journal of Finance*, 55 (2000), 179–223.

Pastor, L., and R. F. Stambaugh. "Comparing Asset Pricing Models: An Investment Perspective." *Journal of Financial Economics*, 56 (2000), 335–381.

Politis, D. N., and J. P. Romano. "The Stationary Bootstrap." *Journal of the American Statistical Association*, 89 (1994), 1303–1313.

Tu, J., and G. Zhou. "Incorporating Economic Objectives into Bayesian Priors: Portfolio Choice under Parameter Uncertainty." *Journal of Financial and Quantitative Analysis*, 45 (2010), 959–986.

Tu, J., and G. Zhou. "Markowitz Meets Talmud: A Combination of Sophisticated and Naive Diversification Strategies." *Journal of Financial Economics*, 99 (2011), 204–215.