Minimum-Variance Portfolios in the U.S. Equity Market

Reducing volatility without sacrificing returns.

Roger Clarke, Harindra de Silva, and Steven Thorley

Roger Clarke
is the chairman of Analytic Investors in Los Angeles, CA.
rclarke@aninvestor.com

Harindra de Silva
is the president and portfolio manager at Analytic Investors in Los Angeles, CA.
hdesilva@aninvestor.com

Steven Thorley
is the H. Taylor Peery professor of finasnce at the Marriott School of Management at Brigham Young University in Provo, UT.
steven.thorley@byu.edu

The concept of an efficient frontier has been a mainstay of financial economics and to some extent portfolio management practice since the 1960s when modern portfolio theory was first articulated. The basic Markowitz [1952] prescription is to estimate expected returns and a covariance matrix for individual securities, and then to minimize the portfolio's ex ante risk for any given expected return by adjusting security weights.

The practical implementation of this prescription for large security sets requires computing power and econometric techniques that were hard to come by until recently, but the essential algorithm has been well-documented and understood for decades. Economists with a goal of describing financial market equilibrium have since specified a number of simplifying assumptions, most notably the information efficiency of security prices, which together with the Markowitz algorithm lead to the traditional capital asset pricing model.[1]

One hurdle to practical implementation of the Markowitz prescription is the estimation of expected security returns. Financial economists usually assume information efficiency and employ some equilibrium model of risk premiums, while active quantitative portfolio managers assume inefficiently priced securities and search for data or signals that predict returns. Both find that numerical optimizers provide security weights that seem overly sensitive to small perturbations in the forecasted security returns.

The minimum-variance portfolio at the left-most tip of the mean-variance efficient frontier has the unique property that security weights are independent of the forecasted or expected returns on the individual securities. While all portfolios on the efficient frontier are designed to minimize risk for a given return, the minimum-variance portfolio minimizes risk without an expected return input.

We construct minimum-variance portfolios using a large set of U.S. equity securities, and examine the realized risk and return statistics over several decades. Our primary motivation is to exploit the unique opportunity to conduct empirical research that is independent of any particular assumption about expected returns.

Given the enormous amount of collective searching for predictive factors by both academics and practitioners over the last half century, any study that involves expected returns is subject to the pervasive data-mining criticism that we know how market history has turned out only after the fact. Our collective examination of historical databases never available before taints much of how we think about markets and where we look for predictive signals. While data mining is a useful tool in many contexts, the concern in financial markets research is that we tend to focus on what might in the end just be historical quirks that have little predictive power in forecasting the future.

Our research is similar to a study by Haugen and Baker [1991], an examination of large-scale minimum-variance portfolios from 1972 through 1989. We extend that research both through 2005 and back to 1968, as well as use more current covariance matrix structuring methodologies, specifically principal components and Bayesian shrinkage. These covariance matrix methodologies are of relatively recent origin, but avoid prespecified factor models for covariance matrix estimation. That is, we do not assume that the markets follow a three-factor, four-factor, or any model that includes risk factors such as market capitalization or book-to-market ratios that have been discovered over the years.

The empirical importance of these factors in explaining the covariance structure of U.S. equity security returns is now clear, although it was not understood by market participants when the Markowitz prescription was first introduced. Our hope is that empirical results that are less dependent on our collective data mining will be more robust going forward. We also examine the relation of the realized average returns and risks of minimum-variance portfolios to the now commonly accepted risk and return factors of size, value, and momentum.

We first review the data and methodologies (providing a technical appendix that summarizes the matrix algebra of both principal components and Bayesian shrinkage). We next report on our basic empirical result that minimum-variance portfolios have about three-fourths the realized risk of the general market, and that this lower risk does not come at the expense of lower realized returns. We explore the imposition of ex ante factor neutrality constraints that match minimum-variance and market portfolio exposures along the now commonly accepted factors of size, value, and momentum.

DATA AND METHODOLOGY

At the beginning of each month from January 1968 through December 2005 (456 months), we estimate a covariance matrix for the 1,000 largest market capitalization U.S. stocks with 60 months of historical return data. The 1968 start date is the first calendar year that 1,000 stocks with sufficient historical data became available in the Center for Research in Security Prices database.[2]

Because of the large-cap selection process and the quality of the CRSP database, companies that have been listed for five years have very few missing returns, and we require all 60 monthly historical observations with little concern for selection bias. The monthly Treasury bill return from Ibbotson Associates is subtracted from security returns for both covariance estimation and performance tracking.

Thus, all results reported are for equity returns in excess of the contemporaneous risk-free rate. The analysis of excess portfolio returns facilitates a simple Sharpe ratio-based analysis of the risk-return trade-offs.

We refer to the *sample* covariance matrix as the simple multiplication of the $1,000 \times 60$ matrix of historical excess returns by its transpose, resulting in a $1,000 \times 1,000$ (1 million element) matrix. Following French, Schwert, and Stambaugh [1987], the average realized excess security return in each security's time series is not subtracted, contrary to the more common definition of statistical variance.

Restricting covariance matrix estimation to historical security returns is required as a matter of practicality for long-history market analysis, because the stock characteristic data required for specified-factor risk models are often not available. Further, the use of factor-based risk models is subject to the data-mining criticism that such models reflect what we collectively know now about the covariance structure of security returns, but could not have known when the Markowitz [1952] prescription for mean-variance portfolio optimization was first popularized. The analysis of historical return covariance estimation procedures may also address whether the purchase of proprietary factor-based risk models is warranted.

A well-known problem in portfolio optimization is that the sample covariance matrix is singular (non-invertable) because there are fewer historical periods (in our case $T = 60$)

than the number of securities in the investable set ($N = 1{,}000$). Besides non-invertability, which causes problems in the search routines of most optimizers, large-sample covariance matrixes include many separate security volatility estimates (i.e., 1,000) and pairwise covariance estimates (i.e., $1{,}000 \times 999/2 = 499{,}500$), leading to estimation outliers that can dominate the optimized portfolio, a problem sometimes referred to as *error maximization*. Further discussion of these issues, in the context of minimum variance portfolios, is found in Chan, Karceski, and Lakonishok [1999].

One resolution to this problem is to provide structure to the sample covariance matrix using the Connor and Korajczyk [1988] asymptotic principal components procedure. We extract returns on the five highest eigenvalued principal components from the historical return cross-product matrix using SAS Procedure PRINCOMP. The five PC return vectors are then used to estimate a 5×5 factor covariance matrix, factor exposures for each security, and residual security return volatilities. The invertability of the final security covariance matrix is ensured by the invertability of the factor covariance matrix and the security residual return covariance matrix, which is diagonal.

In an alternative to principal components, we also structure the sample covariance matrix using the Bayesian shrinkage procedure of Ledoit and Wolf [2004]. That is, we use as a Bayesian prior a two-parameter covariance matrix, assigning all securities the average historical volatility and all security pairs the average pairwise covariance. The final covariance matrix, input to the optimization routine, is a weighted average of the sample covariance matrix and the Bayesian prior, where the shrinkage parameter (i.e., weight placed on the prior) is estimated each month from the historical return data. (See the appendix for a brief discussion of the matrix algebra for both principal components and Bayesian shrinkage.)

Minimum-variance optimization is conducted using the numerical routines in SAS Procedure NLP.[3] Our portfolios are not true minimum-variance because of two types of constraints, in addition to the budget constraint that the sum of the security weights is 1.0. First, we focus on long-only optimizations with a lower bound of zero on security weights and specify an upper bound of 3%. The long-only and upper bound constraints on individual securities positions produce portfolios that are directly comparable to the market, which is by definition long-only and large enough that individual securities rarely exceed 3% of total portfolio capitalization. (We also discuss the impact of relaxing the 3% upper bound on security positions.)

Second, where specified, we apply market neutrality constraints to the optimized portfolio on three factors: size (market capitalization), value (book-to-market ratio), and momentum (prior year less prior month return). These Fama/French [1993] factors, excluding market beta, plus Jegadeesh and Titman [1993] momentum, have been canonized as the key marketwide determinants of realized portfolio returns in the U.S. equity market. Although return forecasts are not part of this study, the factor neutrality constraints help ensure that the minimum-variance portfolios have ex ante characteristics similar to those of the market.

We also study the use of high-frequency (daily) versus monthly historical returns data in estimating the security covariance matrix. The use of high-frequency data in covariance estimation is a topic of recent academic research, with promising results summarized in Gosier et al. [2005]. At each beginning-of-month optimization date, we take the past year of daily returns, which allows for over four times as many return observations ($T = 250$) as monthly data ($T = 60$) while relying on the implicit assumption of covariance matrix stationarity over only the prior year, as opposed to the prior five years. We do not address the market microstructure problems of high-frequency data noted in French, Schwert, and Stambaugh [1987].

At the beginning of each month from January 1968 through December 2005, we complete five steps and then roll forward to the next month. First, we select the largest 1,000 common stocks with sufficient historical information, and save relevant security data including returns and factor exposures. Second, we calculate the sample covariance matrix on the basis of historical returns for each of the 1,000 securities. Third, we structure the sample covariance matrix using Bayesian shrinkage or principal components to produce an estimated covariance matrix. Fourth, we feed the estimated covariance matrix into the optimizer and determine optimal security weights for the minimum-variance portfolio under various constraints. Finally, we use the realized security returns in the current month to track the performance of the optimized and market portfolios.

BASIC EMPIRICAL RESULTS

Exhibit 1 reports statistics on the realized portfolio return time series over 456 months for several portfolios. *Market* is the capitalization-weighted portfolio of all 1,000 stocks, a close approximation of the Russell 1,000 Large-Cap Index.[4] The realized return of the market

EXHIBIT 1
Portfolio Return Statistics 1968:01–2005:12
(returns in excess of one month T-bill)

Monthly Historical Data	Monthly returns (bp) Mean	Std Dev	Annualized Mean	Std Dev	Sharpe Ratio
Market	46	445	5.6%	15.4%	0.36
Min-Var: Bayesian Shrinkage					
Unconstrained	54	337	6.5%	11.7%	0.55
Factor Constrained	46	364	5.6%	12.6%	0.44
Sensitivity Constrained	47	344	5.6%	11.9%	0.47
Min-Var: Principal Components					
Unconstrained	56	333	6.7%	11.6%	0.58
Factor Constrained	45	363	5.5%	12.6%	0.43
Sensitivity Constrained	49	339	5.9%	11.8%	0.50
Long/short portfolios (Bayesian)					
Unconstrained	45	303	5.5%	10.5%	0.52
Factor Constrained	40	321	4.8%	11.1%	0.43
Sensitivity Constrained	45	300	5.4%	10.4%	0.52

Daily Historical Data	Monthly returns (bp) Mean	Std Dev	Annualized Mean	Std Dev	Sharpe Ratio
Market	45	452	5.3%	15.7%	0.34
Min-Var: Bayesian Shrinkage					
Unconstrained	51	326	6.1%	11.3%	0.54
Factor Constrained	41	346	4.9%	12.0%	0.41
Long/short portfolios (Bayesian)					
Unconstrained	51	302	6.1%	10.4%	0.59
Factor Constrained	55	292	6.6%	10.1%	0.65

EXHIBIT 2
Equilibrium Portfolio Theory
(return in excess of the risk-free rate)

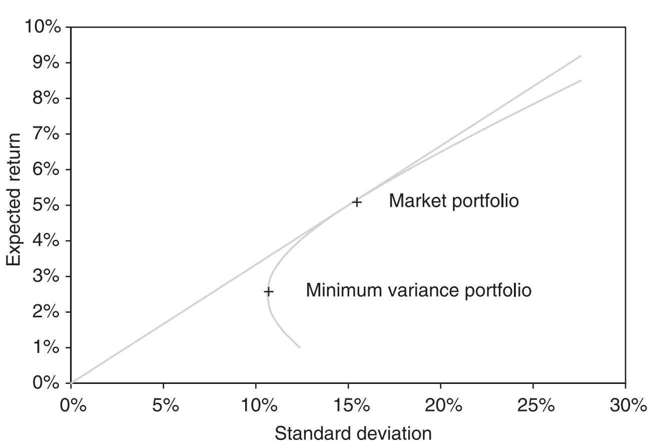

portfolio, in excess of the T-bill rate, averaged 46 basis points per month, for an annualized excess return of 5.6%. As an aside, the average T-bill return over this time period was 5.9%, so that the total average market return was 11.5%.

The realized risk of the market as measured by the standard deviation of monthly excess returns was 445 basis points, an annualized value (multiplication by the square root of 12) of 15.4%. Dividing the annualized excess return by annualized standard deviation gives a Sharpe ratio for the market of 0.36.

For much of the discussion following, we focus on the second line in Exhibit 1, which reports the realized returns on a base case minimum-variance portfolio. This base case uses the Bayesian shrinkage covariance matrix technique and is constrained by the long-only (positive security weights) and maximum security position size specifications.

The realized mean return of 54 basis points (6.5% annualized) in the base case minimum-variance portfolio is higher than the market return, while the realized risk of 337 basis points (11.7% annualized) is substantially lower than the market. In fact, the realized risk of the minimum-variance portfolio is only about three-fourths of the market, indicating material value-added from optimization. The Sharpe ratio of the minimum-variance portfolio is 0.55, much higher than the market's Sharpe ratio of 0.36 because of the higher mean and the lower standard deviation.

As Exhibits 2 and 3 illustrate, the historical track record of the minimum-variance portfolio is at odds with conventional equilibrium portfolio theory. Our results both confirm the practicality of large-scale numerical portfolio optimization techniques and pose a significant puzzle with respect to the basic risk-return trade-off in financial economics.

Ignoring for a moment the differences in realized average returns, we focus first on realized risk. The economic significance of the minimum-variance portfolio risk reduction is substantial: almost 4 percentage points annually below the market. Given the approximate two-to-one ratio of risk to average return (i.e., Sharpe ratios of about one-half), this risk reduction can be translated by leverage into market outperformance of well over 2 percentage points, as visualized in the vertical distance between the lines in Exhibit 3 at the market risk level.

A direct measure of minimum-variance value-added, using ex post regression analysis, yields a more accurate alpha of 2.8%. The one-fourth reduction in risk of the minimum-variance portfolio compared to the market is also significant in a statistical sense, given the sample size of 456 months. Specifically, the null hypothesis of the equality of standard deviation estimates from two samples can be tested using the ratio of the variances (squared standard deviations), which has an F-distribution with denominator and numerator degrees of freedom equal to the sample size minus one. The observed 24% reduction in risk gives a test statistic of $(0.76 \times 0.76) = 0.578$, which is highly significant.

Exhibit 4 shows the cumulative excess return over the 456 months for both the minimum-variance portfolio and the market. While the ending cumulative returns for both portfolios are similar, the impact of the optimization is apparent: less volatility in the line that tracks the minimum-variance portfolio. For example, the sustained market drop of 1982 and the market crash of October 1987 are both muted in the minimum-variance portfolio. Perhaps the greatest impact is seen in the market portfolio's large rise and drop associated with the turn-of-the-century tech-stock bubble, which is almost invisible in the minimum-variance portfolio line.

Exhibit 5 shows the realized risk of the two portfolios using a rolling five-year (60-month) standard deviation of returns. The graphs indicate that the minimum-variance portfolio has lower realized risk than the market over all five-year periods. While the rolling standard deviations are not independent, the time series of 33 years (1968–2005, less 60 months) allows for at least six non-overlapping historical observations of the 60-month standard deviation, each with less realized risk than the market. Exhibit 5 also suggests that the risk reduction is greatest in the most recent period, ending 2005 at about 30% in risk reduction over the market (i.e., 10.0% compared to 14.4%).

The dashed lines in Exhibits 4 and 5 show the cumulative return and rolling standard deviation (i.e., tracking error) of relative portfolio returns, where the relative return is the simple difference between minimum-variance and market portfolio returns. (We discuss a more precise measure of the tracking error of the minimum-variance portfolio later in conjunction with Exhibit 10.)

Exhibit 6 plots the number of securities in the minimum-variance portfolio over time in comparison to the market. The actual number of securities in the minimum-variance portfolio ranges from a low of about 75 to over 250, as shown by the dashed line. The

EXHIBIT 3

Empirical Results 1968 to 2005 (return in excess of T-bill)

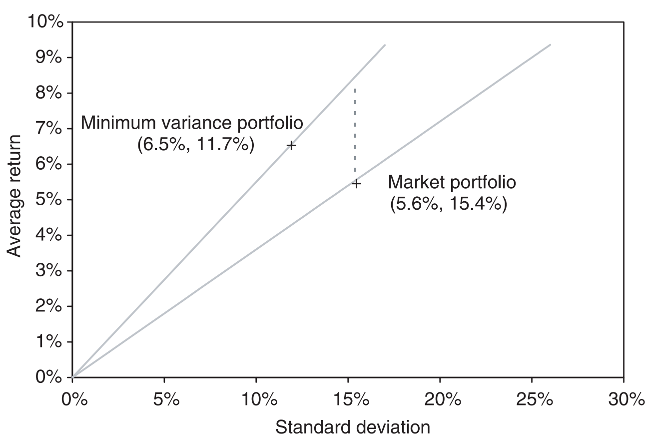

EXHIBIT 4

Cumulative Portfolio Returns

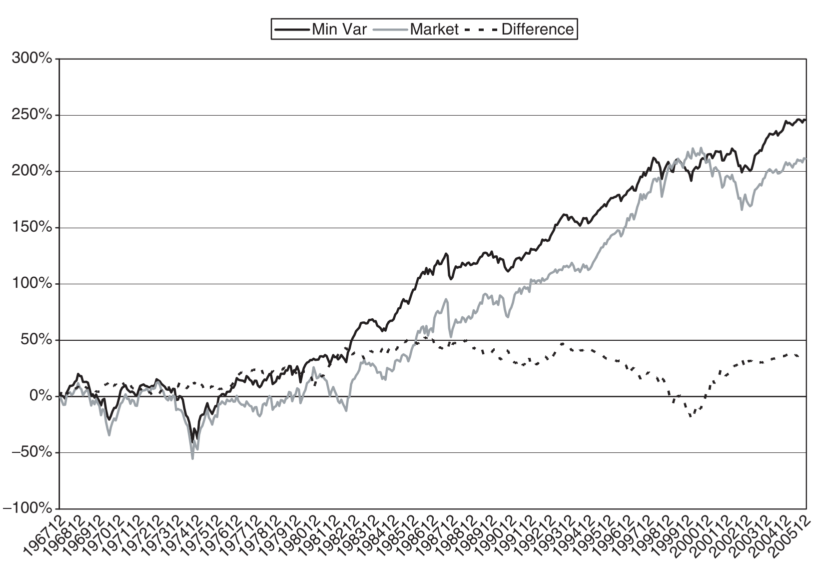

Exhibit 5
Realized Portfolio Risk (trailing 60 months)

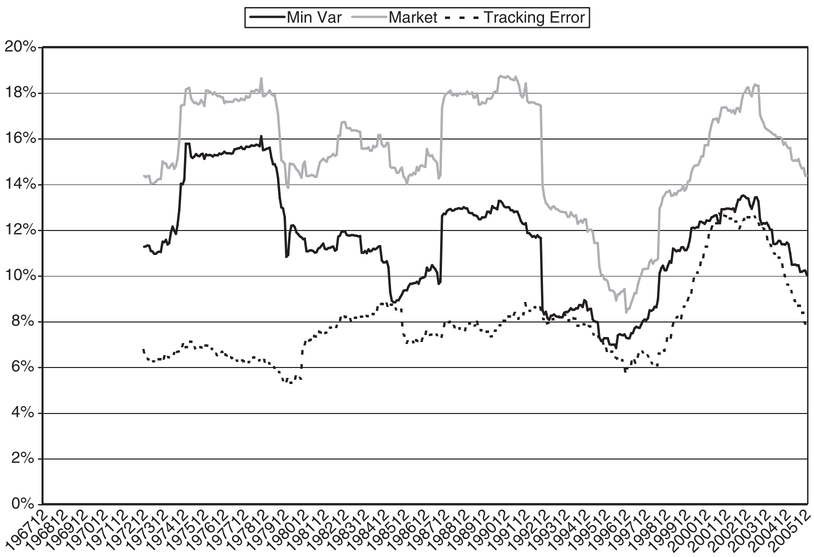

Exhibit 6
Securities in the Portfolio

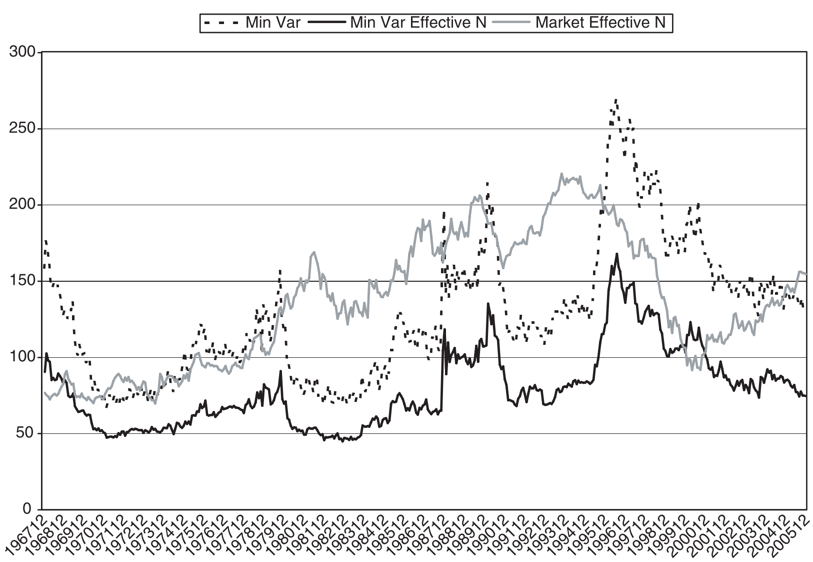

history ends in 2005 with 131 securities, ranging in weights of just greater than zero to the maximum position size of 3%.

Exhibit 6 also shows the *effective* number of securities in the minimum-variance portfolio, calculated as the inverse of the sum of the security weights squared. As explained in Strongin, Petsch, and Sharenow [2000], effective N is the number of equally weighted securities that would provide the same diversification benefit as the actual portfolio. For example, the effective N of the 1,000-stock market portfolio ranges between 75 to a little over 200 securities during the history shown in Exhibit 6. The effective N of the minimum-variance and market portfolios are in the same general range, further justifying the comparison of portfolio risk and return.

Under the Bayesian shrinkage covariance matrix procedure, very few securities are limited by the maximum security weight of 3%, indicating that the error maximization problem in numerical optimization is not much of a concern. As we note later, this concern may still apply to the principal components covariance matrix procedure, where the 3% security weight limit seems to be more binding.

The issue of turnover and transaction costs for minimum-variance portfolios is important in practical application, so we conduct some limited turnover analysis. The base case unconstrained minimum-variance portfolio reported in the second line of Exhibit 1 has security turnover of 11.9% per month, or 143% a year, a turnover rate that may be undesirable. That number compares with 2.0% per month for the 1,000 stock market portfolio in the first line of Exhibit 1, giving a net turnover increase over the passive benchmark of 9.9 percentage points per month, or 119 annualized.

In practice, turnover can be explicitly constrained by the optimizer or balanced with a consideration for transaction costs. We conduct sensitivity analysis by reducing the frequency of rebalancing. That is, we explore the same optimization as the base case unconstrained minimum-variance portfolio, but with calendar-year annual rebalancing instead of monthly rebalancing.

The annual rebalance produces average returns within a basis point of the base case, and with only a minor increase in realized risk from 337 basis points as reported in Exhibit 1, to 341 basis points. The average

EXHIBIT 7
Size Factor Exposure

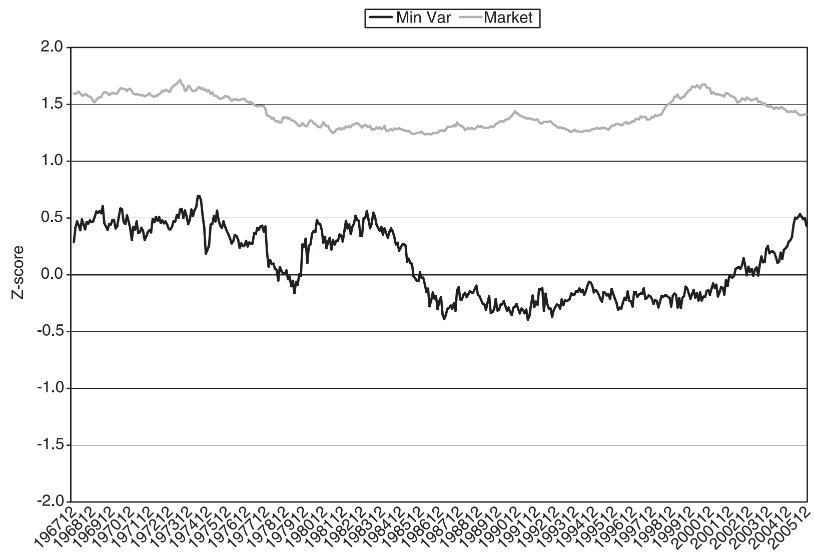

EXHIBIT 8
Value Factor Exposure

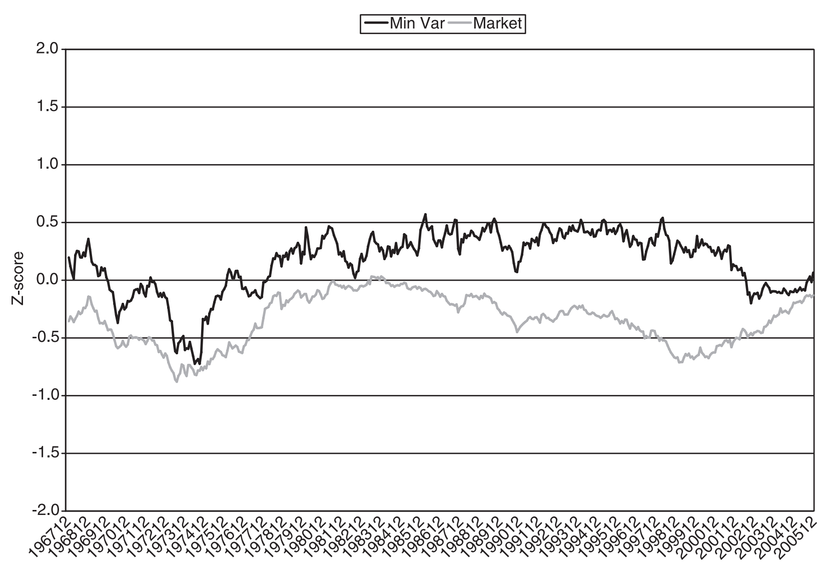

turnover is dramatically reduced to 4.7% per month (56% per year, primarily in January), or in terms of benchmark-relative turnover, 2.7% per month (32% annual.)

This suggests that some latitude in the security selection and weighting composition of minimum-variance portfolios makes the results robust to turnover and other constraints inherent in practical implementation.

FACTOR NEUTRALITY CONSTRAINTS

As we have noted, the realized average excess return of the minimum-variance portfolio in Exhibit 1 of 6.5% is actually higher than the market return of 5.6%, in contrast to the typical risk-return trade-off in financial markets. We next explore the extent to which the higher return is due to minimum-variance portfolio exposures to now well-known determinants of expected return that differ from the market.

Exhibits 7, 8, and 9 chart the exposures of the market and base case minimum-variance portfolios to three critical marketwide factors: size, value, and momentum. At the start of each month, we collect data on market capitalization, book-to-market ratio, and return momentum (prior year less prior month) for all 1,000 securities in the investable set. For a uniform comparison across factors and over time, these exposures are translated into standard normal (i.e., zero mean, unit variance, normal distribution) cross-sectional Z-scores.

Exhibit 7 shows the market and minimum-variance portfolio exposures to the size factor, where the portfolio exposure is the weighted average exposures of the securities in the portfolio. The minimum-variance portfolio fluctuates around zero for most of the history, where zero is by definition the exposure of an equally weighted portfolio. The size factor exposure of the market is naturally much higher than zero because of market-cap weighting. Thus, the higher realized return of the minimum-variance portfolio may be due to the small-cap return premium noted by Fama and French [1993] and others.

Exhibit 8 indicates that, for most of the history, the minimum-variance portfolio also had a substantially higher value (book-to-market) factor exposure than the market, with a narrowing of the spread in 1974 and the end of the history in 2005. Again, the higher average

EXHIBIT 9
Momentum Factor Exposure

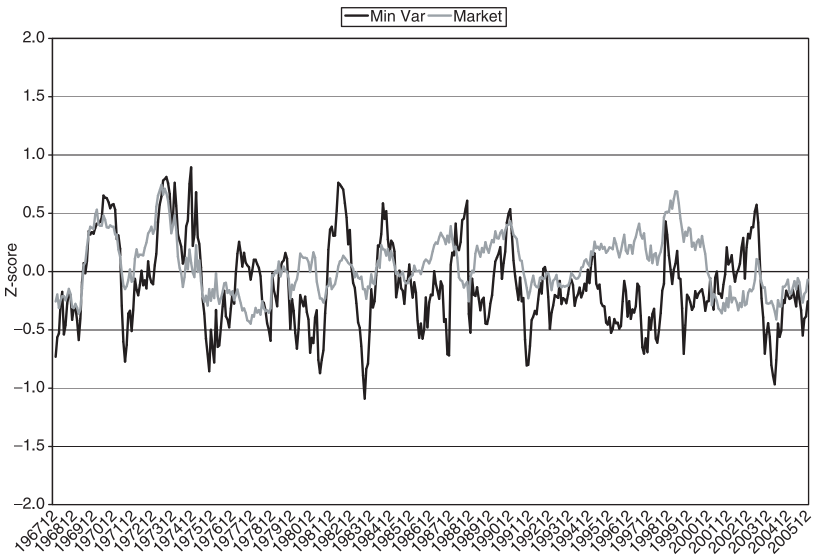

realized return on the minimum-variance portfolio might be associated with the well-documented return premium to value stocks over the last few decades.

Finally, Exhibit 9 shows the momentum exposure of the minimum-variance portfolio compared to the market. While there is no obvious difference in the mean momentum exposure over time, there appears to be more variation in the momentum factor exposure of the minimum-variance portfolio than the market portfolio.

To account for the differential factor exposures between the minimum-variance and market portfolios, we run optimizations with factor-neutrality constraints, programming the optimizer to find the minimum-variance portfolio under the constraint that the size, value, and momentum factors equal the market portfolio exposures at the beginning of each month. The factor constraints insure that the ex ante characteristics of the minimum-variance and market portfolios are equal, at least with respect to the factors now accepted as the major determinants of U.S. market returns over the last few decades.

Exhibit 1, third line, reports the realized return statistics for the factor-neutral-constrained minimum-variance portfolio. As it turns out, the actual realized return of the constrained minimum-variance portfolio is equal to the market (within a basis point) at 46 basis points per month, or 5.6% annualized. At the same time, the factor neutrality constraints limit the risk reduction, with an annualized standard deviation of 12.6% compared to the unconstrained value of 11.7%.

Daniel and Titman [1998] suggest that a more appropriate way to measure a security's factor exposure is the historical sensitivity to factor returns rather than the stock characteristic. For example, some large-cap stocks may covary more with the small-cap market (and vice versa). Using the Fama/French factor returns available on the French web site, we estimate factor sensitivities for all 1,000 stocks each month using a four-factor multivariate regression on the prior 60 monthly returns. We then apply factor-neutrality constraints to the minimum-variance portfolio using the estimated factor sensitivities, rather than the characteristics of the individual stocks.[5]

The results, reported in the fourth line of Exhibit 1, are similar to the direct factor-constrained portfolio, but with slightly higher realized average return and lower realized risk. The combined effect leads to a Sharpe ratio of 0.47 compared to 0.44 for the factor-constrained portfolio, and 0.55 for the unconstrained minimum-variance portfolio.

Although they are not shown here, the time series of factor return sensitivities have similar patterns to the characteristic exposures shown in Exhibits 7, 8, and 9. That is, the unconstrained minimum-variance portfolios have a small-cap and value bias compared to the capitalization-weighted market portfolio in most time periods, but no consistent bias on momentum.

Exhibit 1 next reports on the unconstrained and factor-constrained minimum-variance portfolios using the principal components procedure for structuring the security covariance matrix. The results are fairly similar to the results for Bayesian shrinkage, but a closer examination of the number of securities in the solution, similar to Exhibit 6, indicates that the security weight limit of 3% is now critical. In our principal components implementation, the minimum-variance solution has as few as 40 securities in some time periods, with most of these stocks constrained at the 3% weight limit. Further examination of the estimated covariance matrix explains this phenomenon.

The principal components technique for structuring the sample covariance matrix results in fairly reasonable correlation estimates between securities. Recall that the sample covariance matrix has about half a million independent pairwise correlation estimates (499,500). Given this high number, it's not surprising to find extreme correlation coefficient values in the sample matrix as high as 0.99 or as low as -0.80 in any given month. Structuring the matrix using principal components only allows for correlations associated with the five PC factors, so individual correlation coefficients are reduced to a more reasonable max to min range of about 0.80 to -0.20 in the estimated matrix.

Yet because idiosyncratic security risk is estimated by regression of the security returns on PC factors, the total risk (i.e., diagonal elements) of the sample security covariance matrix is preserved in the estimated covariance matrix. With 1,000 securities, it's not uncommon that several individual securities will have estimated standard deviations in the lower range of 5% to 15%. In the absence of the security weight limit constraint, the optimizer tends to load up on these very low estimated risk securities without regard to correlation estimates. The two-parameter Bayesian shrinkage algorithm based on the statistical theory in Ledoit and Wolf [2003] resolves extreme estimates of both correlations and variances.

ADDITIONAL EMPIRICAL RESULTS

Exhibit 1 next reports return statistics for minimum-variance portfolios when we eliminate the short-sell constraint. These long/short portfolios conform to the budget constraint (security weights sum to 100%) but typically have about 70% of the notional value of the portfolio in short positions (i.e., a 170/70 long/short portfolio).

Not surprisingly, the long/short minimum-variance portfolios have even less risk than the long-only constrained minimum-variance portfolios, although the realized average returns tend to be lower over the time period of this study. The long/short portfolio statistics are included in Exhibit 1 for completeness, but we continue to focus on the long-only portfolios that are arguably more comparable to the market.

We also estimate the security covariance matrix using daily returns from the prior year (i.e., about 250 observations) instead of monthly returns from the prior five years. The added observations and shorter time period should improve the accuracy of the estimated covariance matrix as Gosier et al. [2005] suggest.

The last section in Exhibit 1 reports on optimizations using daily instead of monthly historical data. The statistics on the market portfolio differ slightly from the first line in Exhibit 1 because the one-year versus five-year historical data requirement allows for a slightly different set of 1,000 securities. For daily returns, we allow 2% of missing observations (i.e., 5 out of about 250 for any security), which are replaced by the cross-sectional market average return for the day. We make no attempt to subtract T-bill returns from the daily security returns in the covariance matrix estimation process, but continue to track realized performance using monthly returns in excess of T-bills.

The unconstrained and factor-constrained minimum-variance portfolios using daily historical data reported show some risk-return improvement over the same optimizations using monthly historical data.

EXHIBIT 10

Relative Return Statistics 1968:01–2005:12 (returns in excess of market benchmark)

Monthly historical data	Relative Return Mean	Tracking Error	Information Ratio
Min-var: Bayesian Shrinkage			
Unconstrained	0.9%	8.0%	0.11
Factor Constrained	0.0%	5.8%	0.00
Sensitivity Constrained	0.0%	7.0%	0.01
Min-var: Principal Components			
Unconstrained	1.2%	9.5%	0.12
Factor Constrained	-0.1%	7.0%	-0.01
Sensitivity Constrained	0.3%	8.3%	0.04
Long/short portfolios (Bayesian)			
Unconstrained	-0.1%	10.2%	-0.01
Factor Constrained	-0.7%	9.7%	-0.07
Sensitivity Constrained	-0.1%	10.1%	-0.01
Daily historical data	**Relative Return Mean**	**Tracking Error**	**Information Ratio**
Min-var: Bayesian Shrinkage			
Unconstrained	0.8%	9.5%	0.08
Factor Constrained	-0.4%	6.8%	-0.06
Long/short portfolios (Bayesian)			
Unconstrained	0.8%	11.5%	0.07
Factor Constrained	1.2%	11.3%	0.11

Factor sensitivity constraints for daily historical data are not included because the Fama/French factor returns are not generally available as daily observations. Optimizations using principal components to structure the daily security covariance matrix are not reported in Exhibit 1, but are similar to the Bayesian shrinkage results if the 3% upper bound on security weights is maintained.

The daily long/short results provide some improvement in Sharpe ratio over the long/short results for monthly historical data, validating the argument for using higher-frequency data when they are available, especially for portfolio optimizations that are not limited by the long-only constraint.

Exhibit 10 reports benchmark-relative return statistics. Here, the relative return refers to the simple difference between monthly minimum-variance and market portfolio returns. The absolute return statistics in Exhibit 1 are arguably more relevant for our analysis because the objective function used in the minimum-variance optimizations refers to total, not benchmark-relative, risk. In addition, the calculation of simple return differences ignores the fact that the market betas of the minimum-variance portfolios are well below unity.

For managers who are measured against market benchmarks, the annualized means of the relative returns reported in Exhibit 10 are not impressive, and the tracking error to the market may be uncomfortably high. For example, the base case unconstrained minimum-variance portfolio has an annualized tracking error of 8.0%. The tracking error to the market is reduced by the imposition of ex ante factor neutrality, but is still 5.8% for characteristic-based constraints and 7.0% for sensitivity-based constraints.[6] Consequently, the information ratio, defined as the mean relative return divided by tracking error, is quite low for most minimum-variance portfolios.

Exhibit 11 reports on ex post regression analysis of the base case minimum-variance portfolio returns on the market, as well as perturbations to the base case. A single regression over the entire 1968 to 2005 history (456 times series observations) assumes that the underlying risk parameters are stable over time, although closer examination (i.e., tabulated factor exposures and rolling regressions) indicates that the market risk and factor structure has *not* been stable over the last few decades.

The static regression analysis does produce a directly measured alpha of 2.8% for the base case minimum-variance portfolio and a market beta coefficient of 0.65. As expected, the imposition of neutrality constraints increases the market beta and reduces, but does not eliminate, the ex post alpha.

Note that the base case regression R-squared of 0.74 reported in the last column of Exhibit 11 increases with the imposition of factor neutrality constraints, indicating that the constraints are effective in moving the minimum-variance portfolio closer to the market. Like Daniel and Titman [1998], we find that more of the alpha is preserved using return sensitivity constraints than factor characteristic constraints. Finally, as in Exhibit 1, the use of daily historical data in estimating the covariance matrix improves the optimization value-added, particularly for long/short minimum-variance portfolios.

EXHIBIT 11
Ex Post Regression Analysis 1968:01–2005:12

Monthly historical data	Alpha (Annualized)	Market Beta	Regression R-squared
Min-var: Bayesian Shrinkage			
Unconstrained	2.8%	0.65	0.74
Factor Constrained	1.3%	0.76	0.87
Sensitivity Constrained	1.7%	0.69	0.81
Min-var: Principal Components			
Unconstrained	3.4%	0.59	0.62
Factor Constrained	1.4%	0.73	0.80
Sensitivity Constrained	2.3%	0.65	0.72
Long/short portfolios (Bayesian)			
Unconstrained	2.6%	0.51	0.57
Factor Constrained	1.7%	0.56	0.60
Sensitivity Constrained	2.6%	0.51	0.58

Daily historical data	Alpha (Annualized)	Market Beta	Regression R-squared
Min-var: Bayesian Shrinkage			
Unconstrained	3.0%	0.58	0.64
Factor Constrained	1.2%	0.70	0.83
Long/short portfolios (Bayesian)			
Unconstrained	3.7%	0.45	0.46
Factor Constrained	4.2%	0.45	0.48

A final caveat on the ex ante factor neutrality constraints is in order. Broadly speaking, one never knows if factor neutrality is achieved ex post because history provides us with only one portfolio return observation each month, not the spectrum of possible histories that could have occurred. In addition, each month has a slightly different covariance matrix, which evolves over time in an attempt to capture the changing risk characteristics of the market. We can only try to measure neutrality ex post by stringing together realized portfolio returns, and regressing them on realized factor returns, similar to the univariate regression analysis reported in Exhibit 11.

We find in this ex post analysis (results not tabulated for the sake of space) that factor-constrained long-only minimum-variance portfolios appear to retain part of their value bias, which may explain some of the value added over the market.[7]

CONCLUSION

So what have we learned? First, large-scale optimization procedures for variance reduction do seem to work. Security variance and covariance risk are persistent, and thus reasonably predictable. Large-sample covariance matrixes can be constructed that are practical and operational as inputs to numerical optimization routines, without the need for Monte Carlo procedures. The problems of error maximization associated with extreme individual security variance and covariance estimates can be ameliorated by structured covariance matrix procedures, particularly the more recently developed Bayesian methods.

Second, minimum-variance portfolios that do not rely on any specific expected return theory or return forecasting signal show promise in terms of adding value over the market-capitalization weighted benchmark. We can confirm Haugen and Baker's [1991] basic empirical result over the years 1972 through 1989 using a longer and more recent dataset from 1968 through 2005. Our results are also consistent with recent research by Ang et al. [2006], who conclude that stocks with higher historical idiosyncratic volatility have lower realized returns. In general, we find that realized standard deviation is lowered by about one-fourth, and risk as measured by market beta is lowered by about one-third, compared to the capitalization-weighted market benchmark.

Third, minimum-variance portfolios tend to have both a value and small-size bias. These biases can be controlled ex ante by the imposition of factor neutrality constraints using either stock characteristics or factor return sensitivities. Under either methodology, the imposition of factor neutrality constraints does little harm to the high realized Sharpe ratios of minimum-variance portfolios. We confirm that the use of high-frequency daily returns improves the risk estimation process, although only modestly, and that long/short minimum-variance portfolios provide greater value-added than long-only constrained optimizations using daily data.

We also cite ex post analysis that assigns some of the value-added of minimum-variance portfolios to the value factor, even after the imposition of ex ante neutrality constraints. On the other hand, we note that Markowitz did not base his portfolio optimization prescription on any awareness of the value or any other factor's performance. The value factor as we now know it was not an element in the general academic theory or practitioner practice of the 1960s, and may or may not yield favorable performance statistics in the future.

All empirical research on the U.S. stock market suffers from the data-mining criticism that we collectively rely on knowledge of market characteristics that we have extracted from an after-the-fact observation of our history. Conscientious empiricists are careful to test their specific hypothesis out of sample, but researchers with any financial markets education inevitably rely on perspectives tainted by what we know about markets now, as opposed to what one could have known in the past.

While the econometric methodologies and computing power are vastly improved, the analysis of minimum-variance portfolios in this study rests squarely on the Markowitz [1952] mean-variance portfolio management objective function popularized in the 1960s. We do not rely on prespecified risk factors, let alone any perspective or model of expected average returns. Results based on methodologies that are less dependent on researchers' collective data mining of the historical databases are arguably more reliable in the only out-of-sample empirical test that truly matters, the future.

Additional empirical research could be done on minimum-variance portfolio strategies applied to international equity and other financial markets. We look forward to exploring the apparent persistence of a value bias in minimum-variance portfolio returns even after controlling for the ex ante exposure. We also anticipate more academic research on differences in cross-sectional historical volatility as a factor in explaining the out-of-sample covariance and average returns on U.S. stocks.

APPENDIX

Principal Components and Bayesian Shrinkage

We estimate security covariance matrixes that rely solely on historical security return data. Let R be an $N \times T$ matrix of historical excess returns on N securities for T periods. In the base case of this study, N is 1,000 and T is 60 for monthly historical data. The *sample* covariance matrix, Ω, is calculated directly from the historical returns by the simple matrix equation

$$\Omega = RR' \tag{A-1}$$

where the symbol $'$ indicates the matrix transpose function.

The elements of this sample covariance matrix do not fully conform to the standard definition of historical pairwise security covariances in that the security returns are not de-meaned over time, following French, Schwert, and Stambaugh [1987], and Equation (A-1) does not have the requisite scaling factor (i.e., division by T, and multiplication by 12 for annualized returns). The key issue is that the $N \times N$ sample covariance matrix, Ω, is not invertable because N is much larger than T. We employ two common techniques for creating an invertable covariance matrix that are suitable for optimization: principal components and Bayesian shrinkage.

The asymptotic principal components procedure of Connor and Korajczyk [1988] is based on an eigenvalue decomposition of the $T \times T$ security return cross-product matrix, $R'R$, in contrast to the sample security covariance matrix, RR'.[8] The key outputs of the principal components algorithm are the eigenvectors associated with the K highest eigenvalues of the return cross-product matrix, which combine to form the $K \times T$ factor return matrix F. The principal components procedure also outputs a $K \times N$ matrix B of factor exposures for the securities, based on the regression equation:

$$B = (FF')^{-1}(FR') \tag{A-2}$$

The factor return and exposure matrixes are then used to calculate the $N \times T$ residual returns matrix:

$$E = R - B'F \tag{A-3}$$

and the K-factor risk model equation is used to calculate the *PC* version of the security covariance matrix:

$$\Omega_{PC} = B'(FF')B + diag(EE') \tag{A-4}$$

where diag(#) is the diagonalization function. We include the factor return covariance matrix, FF', in Equations (A-2) and (A-4) so that the forms are recognizable, although in this principal components application the factor return covariance matrix is an identity matrix and is thus not required.

While the principal components methodology structures the off-diagonal elements of the security covariance matrix to guarantee invertability, the procedure does not modify the diagonal elements (i.e., security variance estimates) of the sample covariance matrix. Given 1,000 securities, there are fairly extreme values for some estimated security variances in any given month. This motivates our use of the two-parameter Bayesian shrinkage model of Ledoit and Wolf [2003].

Each month we calculate the average value of the N diagonal elements (security variance estimates) in the sample covariance matrix given in Equation (A-1), as well as the average value of the $N(N-1)/2$ independent off-diagonal elements (security covariance estimates). These two scalar parameters are used to populate an $N \times N$ Bayesian prior matrix Ω_{prior}. The Bayesian shrinkage or *BS* security covariance matrix is the weighted average of the two-parameter prior matrix Ω_{prior} and the sample covariance matrix Ω:

$$\Omega_{BS} = \lambda\,\Omega_{prior} + (1-\lambda)\Omega \tag{A-5}$$

where λ is a scalar shrinkage parameter bounded between zero and one.

The value of the unbounded shrinkage parameter is a complex function of the spread of variance and covariance values around the averages. Specifically:

$$\lambda = \frac{\text{SUM}(\text{SQ}(R)\text{SQ}(R)') - \text{SUM}[\text{SQ}(\Omega)]/T}{\text{SUM}[\text{SQ}(\Omega - \Omega_{prior})]} \tag{A-6}$$

where SQ( ) is the element-by-element squaring (power of two) function for a matrix argument, and SUM( ) is the sum of the argument matrix elements.

In this research, the shrinkage parameter 1for monthly historical returns (1,000 securities, 60 months) generally ranges from 40% to 80% over the 1968–2005 period, with an average value of 54.4%. The shrinkage back to the variance and covariance means for the daily historical returns (1,000 securities, 250 days) ranges from about 10% to 50%, with an average value of 31.3%. The lower shrinkage for daily data reflects the higher number of historical observations, which leads the Bayesian process to rely more on the variations around parameter means. The amount of shrinkage spikes up (but does not hit 100%) for a period after the October 1987 crash using both monthly and daily data.

ENDNOTES

[1]The mean-variance Markowitz algorithm is a fairly straightforward normative statement about what investors *should* do. The traditional CAPM by Sharpe [1964] and others is a positive statement of what *would* prevail in a marketplace of mean-variance investors under a number of simplifying and controversial assumptions. The CAPM has fallen on hard times empirically, and most researchers have adopted ad hoc multifactor models, specifically the three-factor model of Fama/French [1993], with the addition of momentum by Jegadeesh and Titman [1993] and others.

[2]The CRSP database has comprehensive monthly and daily returns data on several hundred NYSE shares dating back to 1926. Our specification of 1,000 U.S. domiciled equity securities, excluding closed-end funds and REITS, requires the addition of the Amex market, which is not included for a full calendar year in CRSP until 1963. In our base case portfolio, we use 60 months of historical return data on each security for covariance matrix estimation, so the optimized portfolio results start in 1968. The CRSP database and thus our study also include Nasdaq securities starting in 1974.

[3]One base case solution was checked using the commercial optimizer Axioma©, which produced the same security weights as SAS Procedure NLP for any given estimated covariance matrix.

[4]Monthly returns for the Russell 1000 large-capitalization index are available starting in January 1979, 11 years after the January 1968 start date of our study. For the overlapping 324 months from 1979:01 through 2005:12, the correlation coefficient between our CRSP-based market return (first line of Exhibit 1) and the Russell 1000 index is 99.7%, indicating a very close match.

[5]We constrain portfolio sensitivity to be equal to the sensitivity of the 1,000-stock market portfolio on the SMB, HML, and UMD factors as measured by a multivariate regression of each security return on the four Fama/French factor returns (Mkt-Rf, SMB, HML, and UMD) over the prior 60 months. We also constrain portfolio sensitivities to be *zero* (rather than matching the market) on the SMB, HML, and UMD factors, and find immaterial differences in the final results.

[6]The tracking error of the base case minimum-variance portfolio relative to the Russell 1000 index return is 8.9% over the overlapping 1979 to 2005 time period. The tracking error of the minimum-variance portfolio relative to the Russell 1000 Value is lower at 6.7%, suggesting a value bias in the unconstrained portfolio.

[7]Specifically, we conducted ex post multivariate regression analysis of minimum-variance portfolio returns on four factor returns, Mkt-Rf, SMB, HML, and UMD. In these regressions, the HML coefficient is reduced by the imposition of the factor neutrality constraints but is still statistically significant. On the other hand, the SMB coefficient indicates that the small-cap bias of minimum-variance portfolios is eliminated, and perhaps overcorrected, by the ex-ante neutrality constraints.

[8]The basic numerical algorithm in the asymptotic principal components procedure is an eigenvalue-based decomposition of the security return cross-product matrix, $R'R = F'(BB')F$, where F is an orthonormal eigenvector matrix (i.e., $FF' = I$), and BB' is a diagonal matrix of eigenvalues, multiplied by N.

REFERENCES

Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang. "The Cross-Section of Volatility and Expected Returns." *Journal of Finance*, 61 (2006), pp. 259-299.

Chan, Louis K.C., Jason Karceski, and Josef Lakonishok. "On Portfolio Optimization: Forecasting Covariance and Choosing the Risk Model." *The Review of Financial Studies,* Winter 1999, pp. 937-974.

Connor, Gregory, and Robert Korajczyk. "Risk and Return in Equilibrium APT: Application of a New Test Methodology." *Journal of Financial Economics*, 21 (1988), pp. 255-289.

Daniel, Kent, and Sheridan Titman. "Characteristics or Covariances?" *The Journal of Portfolio Management*, Summer 1998, pp. 24-33.

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

French, Kenneth, G. William Schwert, and Robert Stambaugh. "Expected Stock Returns and Volatility." *Journal of Financial Economics*, 19 (1987), pp. 3-25.

Gosier, Kenneth, Ananth Madhavan, Vitaly Serbin, and Jian Yang. "Toward Better Risk Forecasts." *The Journal of Portfolio Management*, Spring 2005, pp. 82-91.

Haugen, Robert, and Nardin Baker. "The Efficient Market Inefficiency of Capitalization-Weighted Stock Portfolios." *The Journal of Portfolio Management*, Spring 1991, pp. 35-40.

Jegadeesh, Narasimhan, and Sheridan Titman. "Returns to Buying Winners and Selling Losers; Implications for Stock Market Efficiency." *Journal of Finance*, 48 (1993), pp. 65-91.

Ledoit, Olivier, and Michael Wolf. "Honey, I Shrunk the Sample Covariance Matrix." *The Journal of Portfolio Management*, Summer 2004, pp. 110-119.

—— "Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection." *Journal of Empirical Finance*, 10 (2003), pp. 603-621.

Markowitz, Harry. "Portfolio Selection." *Journal of Finance*, 7 (1952), pp. 77-91.

Sharpe, William. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *Journal of Finance*, 19 (1964), pp. 425-442.

Strongin, Steven, Melanie Petsch, and Greg Sharenow. "Beating Benchmarks." *The Journal of Portfolio Management*, Summer 2000, pp. 11-27.