Toward the Design of Better Equity Benchmarks:

Rehabilitating the Tangency Portfolio from Modern Portfolio Theory

LIONEL MARTELLINI

LIONEL MARTELLINI
is a professor of finance at EDHEC Business School and the scientific director of EDHEC Risk and Asset Management Center in Nice, France.
lionel.martellini@edhec.edu

With an ever increasing volume of assets managed under passive indexing strategies, whether packaged under a mutual fund format or under an exchanged-traded fund format, the standard practice of constructing stock market indices based on a capitalization weighting scheme has recently been subject to heightened scrutiny and renewed criticism. More than fifteen years ago, a number of articles (in particular, Haugen and Baker [1991] and Grinold [1992]) had already offered convincing empirical evidence that market cap–weighted indices provide an inefficient risk–return trade-off. From a theoretical standpoint, the poor risk-adjusted performance of such indices should not come as a surprise given that the efficiency of market cap–weighting schemes can only be warranted at the cost of heroic assumptions.

More recently, equity indices with different weighting schemes have emerged as alternatives. A particular type of alternative weighting scheme is currently implemented by indices that use accounting characteristics to weight the component stocks (Arnott, Hsu, and Moore [2005]). In a different spirit, Fernholz, Garvy, and Hannon [1998] proposed the use of stock market diversity—a measure of the distribution of capital in an equity market, further discussed by Fernholz [2000]—as a weighting factor. Under a set of simplifying assumptions, Fernholz showed that the return of a diversity weighting scheme relative to the market cap–weighting scheme is, among other things, a function of the difference between the weighted-average variance of the individual stocks and the portfolio variance. Closely related is a recent paper by Choueifaty and Coignard [2006] which introduced, in a portfolio optimization context, a so-called diversification index given by the ratio of the average volatility of stocks in a portfolio divided by the portfolio volatility.[1] One last example is the emergence of equal-weighted counterparts to well-known indices such as the S&P 500 Index.

From a theoretical standpoint, however, it is not clear how any of these attempts are grounded in modern portfolio theory, which states that a mean–variance utility-maximizing agent should hold the portfolio with the highest reward-to-risk ratio, also known as a tangency portfolio, or maximum Sharpe ratio (MSR) portfolio. In fact, the equal-weighted portfolio can be regarded as an MSR portfolio, albeit under the rather specific assumption that all stocks' expected returns, volatilities, and pairwise correlations are identical. While discarding all possible information about the moment and co-moment structure of stock returns may appear to be an extreme stance, most attempts at implementing the prescriptions of scientific diversification on the basis of sample-based parameter estimates have generated rather disappointing results in thorough out-of-sample empirical tests.

In a recent article, DeMiguel, Garlappi, and Uppal [2007] argued that the presence of estimation error in input parameters almost entirely invalidates the performance of formal optimization models, even when improved estimators are used, except for an unreasonably large sample size. That article, following most existing research on the topic including, for example, Chan, Karceski, and Lakonishok [1999] and Jagannathan and Ma [2003], focused on testing the out-of-sample performance of global minimum variance (GMV) portfolios based on various improved estimators of the variance–covariance matrix and found they were dominated by the simple equal-weighted rule in terms of risk-adjusted performance. The reason these authors have chosen to focus on the out-of-sample performance of GMV portfolios, as opposed to MSR portfolios, is related to the widespread consensus that expected returns are difficult to obtain with a reasonable estimation error. Making the problem worse is that optimization techniques are well known to be very sensitive to differences in expected returns, so that portfolio optimizers typically allocate the largest fraction of capital to the asset class for which estimation error in the expected returns is the largest (e.g., Britten-Jones [1999] or Michaud [1998]).

This article attempts to move the focus back to the only truly optimal weighting scheme consistent with modern portfolio theory, that is, the one that is designed to achieve the highest Sharpe ratio. In a nutshell, I claim that using improved estimators for the variance–covariance matrix of stock returns, as well as for the expected returns, allows the design of proxies for the tangency portfolio that outperform their market cap–weighted counterparts on a risk-adjusted basis in out-of-sample exercises. I also find evidence that such benchmarks may outperform equal-weighted schemes, a rare occurrence in similar studies that have been previously conducted.

These results are obtained from an approach that consists of linking expected return estimation to volatility estimation, building upon the recent literature on the explanatory power of idiosyncratic volatility for the cross section of expected returns. In particular, Malkiel and Xu [2006], extending an insight from Merton [1987], showed that the inability to hold the market portfolio, for whatever the cause, forces investors to care about total risk to some degree in addition to market risk. The result is that firms with larger firm-specific variances require higher average returns to compensate investors for holding imperfectly diversified portfolios (see also Barberis and Huang [2001] for a similar conclusion from a behavioral perspective). That stocks with high idiosyncratic volatility earn higher returns has also been confirmed in a number of recent empirical studies, including Tinic and West [1986] and Malkiel and Xu [1997, 2002].

These findings, in view of the fact that asset pricing theory implies a positive premium for systematic risk, suggest that expected returns should be positively related to total volatility, a model-free quantity given by the sum of systematic and idiosyncratic risk components as measured against the prevailing asset pricing model. In this context, I propose the use of stock volatility as a robust proxy for expected excess return and investigate the implication in terms of portfolio construction and performance.[2]

ON THE RELATIONSHIP BETWEEN TOTAL VOLATILITY AND EXPECTED RETURN

In this section, I report empirical evidence that supports the relationship between total volatility and expected return. For this analysis, I use monthly excess stock return data from the Center for Research in Security Prices (CRSP) for 682 stocks that survived over the entire 1975–2004 sample period.

I consider the time-series evidence on whether stocks with high volatility have higher average returns than stocks with lower volatility. For this, I follow the procedure of Ang et al. [2006] and sort individual stocks into equal-weighted quintiles based on their total volatility calculated on the basis of the returns of the previous 120 months (10 years). I hold these portfolios for 1 month and then repeat the process. I thus obtain a time series of returns for the quintiles starting in January 1985 and compute the difference that I call the 1–5 premium, which is the average return between the top quintile (20% of the highest-volatility stocks) and the bottom quintile (20% of the lowest-volatility stocks). The results are reported in Exhibit 1 and show that the 1–5 annualized premium over the 1985–2004 period reaches a spectacular annualized value of 8.47%, with a *p*-value of 2.38%, indicating that stocks with higher volatility earn significantly higher returns than stocks with lower volatility. While the premium is not significantly positive for quintiles 2, 3, and 4, a monotonic increase is found for annualized performance as a function of volatility. Exhibit 2 shows the

EXHIBIT 1
Annualized Performance of Five Quintiles and Quintile Premia

	Annualized Return	Annualized Premium	*P*-values (>0)
Quintile 1	22.99%	8.47%	2.38%
Quintile 2	15.94%	1.42%	30.39%
Quintile 3	15.31%	0.79%	35.62%
Quintile 4	15.24%	0.72%	32.68%
Quintile 5	14.52%	0.00%	NA

Note: This exhibit shows the out-of-sample performance of volatility quintiles over the 1985–2004 period with respect to the high-volatility quintile. Returns are given in percent per year. P-values indicate if the respective annualized premium is statistically significant.

cumulative returns on a strategy that goes long the top quintile and short the bottom quintile, and confirms the economic significance of the effect.

I further analyze whether total volatility has explanatory power for the cross section of expected returns by following the procedure of Fama and Macbeth [1973]. At each point in time, I run a cross-sectional regression of realized return on total estimated volatility,

$$R_{i,t} = \alpha_{i,t} + \gamma_t \sigma_{i,t} + \varepsilon_{i,t}$$

I then average the slope coefficient in this regression across time ($\bar{\gamma}$) to obtain an estimate for the unconditional slope coefficient. The *t*-statistic for the slope coefficient is given by $t_\gamma = \frac{\sqrt{T_\gamma}}{\sigma_\gamma}$, where $\sigma_\gamma = \frac{1}{T-1}\Sigma(\gamma_t - \bar{\gamma})^2$. Total volatility is computed using the previous 120 months of data so that the regression begins in 1985. I find that the

EXHIBIT 2
Equal-Weighted Top-Minus-Bottom Cumulative Returns, 1985–2004

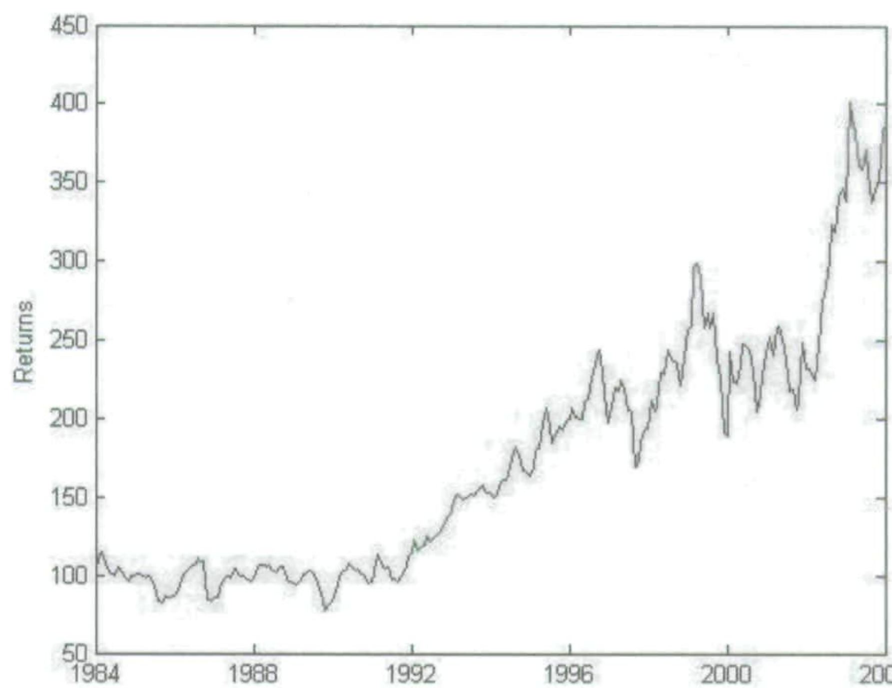

Note: For the 1985–2004 period, this exhibit shows the cumulative returns on a zero-cost portfolio strategy that is long the top equal-weighted quintile of stocks sorted on total volatility and short the bottom equal-weighted quintile sorted on total volatility.

average slope coefficient, when using total volatility as a regressor over the 1985–2004 period, is 0.1095 with a p-value of 0.01%, which shows that total volatility has explanatory power for the cross section of stock returns.

Given that total volatility is the sum of systematic volatility with respect to an asset pricing model and of idiosyncratic volatility with respect to the same asset pricing model, it is possible that these results are driven mostly by the systematic component of volatility. To test this possibility, I decompose total volatility into Fama–French systematic volatility and idiosyncratic volatility, using these as separate regressors. I find that over the 1985–2004 period the slope coefficient for Fama–French systematic volatility is 0.0257 and not significant, while the slope coefficient for idiosyncratic volatility is 0.1122 with a p-value of 0.03%, indicating that the explanatory power of total volatility is driven by idiosyncratic volatility with respect to the Fama–French model.

IMPLICATIONS FOR PORTFOLIO OPTIMIZATION AND DESIGN OF EFFICIENT EQUITY BENCHMARKS

The focus of this article is not to empirically assess the statistical significance and asset pricing implications of the relationship between total volatility and expected return at the individual stock level (a task performed in Basu and Martellini [2007]), but instead to try and assess the economic significance of such findings for portfolio selection and equity benchmark construction. To this end, I perform the following experiment. I randomly draw 100 portfolios of 100 stocks each from the base universe of stocks. For each of the 100 sub-universes, I estimate on a monthly basis the maximum Sharpe ratio (MSR) portfolio and the global minimum variance (GMV) portfolio based on the past 120 months of sample information and record their performance on an out-of-sample basis.

In terms of the correlation matrix, I use the following three estimates: 1) sample correlation matrix, 2) constant correlation matrix (see Elton and Gruber [1973]), and 3) factor-based correlation matrix for which the Fama–French three-factor model is used.[2] The first estimator (the sample correlation matrix) is likely to be severely plagued by estimation error due to the high number of parameters which must be estimated and the second estimate (the constant correlation matrix) is likely to be impaired by the presence of significant specification risk and implies the inability to extract *any* useful information about the co-movement structure of stock returns from the sample. The factor-based estimator, however, is expected to allow for a reasonable trade-off between sample risk and model risk.

In terms of the volatility vector, I use the sample estimate. The volatility estimates are used to obtain the covariance matrix needed for the GMV portfolios. Consistent with the discussion in the previous section, I also use the volatility estimates as proxies for excess expected return estimates in the case of the MSR portfolios. Hence, the MSR portfolio, with a composition well known to be given by $\frac{\Sigma^{-1}\mu}{1_N \Sigma^{-1}\mu}$, can be rewritten as $\frac{\Sigma^{-1}\sigma}{1_N \Sigma^{-1}\sigma}$, where Σ is the N-dimensional stock return covaiance-matrix, μ is the N-dimensional vector of excess expected returns, σ is the N-dimensional vector of volatilities, and 1_N is an N-dimensional vector of ones. It should be noted that if the assumption is further made that volatility is identical for all stocks, then the MSR portfolio, which is based on volatility as a proxy for excess expected return, coincides with the GMV portfolio, $\frac{\Sigma^{-1}1_N}{1_N \Sigma^{-1}1_N}$.

In Exhibit 3, I report minimum, bottom-25th percentile, median, top-25th percentile, and maximum out-of-sample values (over the 100 randomly chosen portfolios) for expected returns, volatilities, and Sharpe ratios of the MSR and GMV portfolios. I compare these values to the risk and return indicators for value-weighted and equal-weighted counterparts. A number of conclusions can be drawn from this analysis. First, the value-weighted index is outperformed by its equal-weighted counterpart in terms of the Sharpe ratio, suggesting again that market cap–weighting may not be the most appropriate weighting scheme as far as portfolio efficiency is concerned. The GMV portfolios, however, appear to have lower out-of-sample volatility compared to both the value-weighted and equal-weighted benchmarks when an improved estimator of the correlation matrix is used, an estimator that offers a reasonable trade-off between sample and specification risks. As expected, the sample and correlation estimators do not yield very attractive out-of-sample results. The Sharpe ratios of GMV portfolios typically dominate the Sharpe ratios of the value-weighted benchmarks, but remain lower than the Sharpe ratios of the equal-weighted benchmarks because of poorer average performance. Taken together, these results are consistent with those of DeMiguel, Garlappi,

Exhibit 3

Performance Indicators Using Sample Estimate of Stock Volatilities with Total Volatility as Proxy for Excess Expected Return

Sample Volatility		Expected Return					Volatility					Sharpe Ratio				
		Min	25%	Median	75%	Max	Min	25%	Median	75%	Max	Min	25%	Median	75%	Max
Sample	MSR	-67.79%	4.31%	13.79%	25.23%	3233%	36.39%	55.05%	66.81%	111%	13082%	-0.233	-0.001	0.221	0.354	0.839
Correlation	GMV	-11.28%	0.40%	3.81%	7.23%	19.94%	22.12%	24.34%	25.28%	26.62%	49.08%	-0.665	-0.053	0.112	0.282	0.832
Constant	MSR	9.37%	10.97%	11.53%	12.07%	13.28%	13.75%	14.26%	14.40%	14.54%	14.90%	0.635	0.761	0.799	0.836	0.948
Correlation	GMV	2.71%	5.48%	6.54%	7.48%	11.25%	12.20%	12.85%	13.21%	13.50%	16.10%	0.177	0.397	0.493	0.567	0.852
Factor	MSR	5.76%	9.76%	11.27%	12.53%	15.85%	11.29%	12.65%	12.97%	13.43%	15.91%	0.449	0.747	0.874	0.955	1.192
Correlation	GMV	4.57%	6.98%	8.22%	9.14%	11.58%	9.84%	10.61%	10.85%	11.06%	11.97%	0.435	0.637	0.758	0.842	1.072
Equal-Weighted	EW	9.71%	11.67%	12.59%	13.03%	14.81%	14.97%	15.71%	15.95%	16.26%	17.28%	0.601	0.727	0.783	0.812	0.917
Value-Weighted	VW	7.22%	9.09%	9.83%	10.43%	12.12%	15.47%	15.83%	16.07%	16.27%	16.67%	0.455	0.574	0.613	0.651	0.770

Note: Risk and return performance indicators of various portfolios constructed from 100 random draws of 100 stocks, based on a sample estimate of stock volatilities and using total volatility as a proxy for a stock's excess expected return. This exhibit shows the out-of-sample (geometric) annualized performance, annualized volatility, and Sharpe ratio of various optimized portfolios over the 1985–2004 period. The abbreviations MSR and GMV represent the maximum Sharpe ratio portfolio and the global minimum variance portfolio, respectively. In all cases, the sample estimate for volatility has been used in the exhibit. The expected return estimate is assumed to be given by the sample volatility estimate for the MSR portfolio. Various estimates of the correlation matrix have been used, including the sample correlation estimate, the constant correlation estimate, and the factor correlation estimate (see details in the text).

Exhibit 4

Performance Indicators Using Sample Estimate of Stock Volatilities with Fama–French Model to Forecast Excess Expected Return

Sample Volatility		Expected Return					Volatility					Sharpe Ratio				
		Min	25%	Median	75%	Max	Min	25%	Median	75%	Max	Min	25%	Median	75%	Max
Sample	MSR	-430.7%	2.35%	11.77%	22.08%	166.9%	33.24%	44.62%	56.92%	98.74%	1982%	-0.2375	0.0349	0.1960	0.3785	0.8162
Correlation	GMV	-9.04%	3.60%	6.93%	9.89%	21.48%	21.70%	23.71%	24.80%	25.71%	34.24%	-0.3820	0.1399	0.2800	0.3820	0.8559
Constant	MSR	-94.02%	7.67%	11.58%	14.30%	62.18%	23.09%	26.69%	28.31%	32.56%	449%	-0.2093	0.2558	0.3802	0.4949	0.7395
Correlation	GMV	2.71%	5.48%	6.54%	7.48%	11.25%	12.20%	12.85%	13.21%	13.50%	16.10%	0.1768	0.3973	0.4931	0.5670	0.8517
Factor	MSR	8.05%	10.20%	11.33%	12.61%	16.23%	13.60%	14.17%	14.46%	14.81%	15.80%	0.5496	0.6951	0.7819	0.8652	1.0652
Correlation	GMV	4.50%	6.96%	8.20%	9.10%	11.56%	9.84%	10.61%	10.85%	11.06%	11.94%	0.4278	0.6335	0.7574	0.8401	1.0688
Equal-Weighted	EW	9.71%	11.67%	12.59%	13.03%	14.81%	14.97%	15.71%	15.95%	16.26%	17.28%	0.601	0.727	0.783	0.812	0.917
Value-Weighted	VW	7.22%	9.09%	9.83%	10.43%	12.12%	15.47%	15.83%	16.07%	16.27%	16.67%	0.455	0.574	0.613	0.651	0.770

Note: Risk and return performance indicators of various portfolios constructed from 100 random draws of 100 stocks, based on a sample estimate of stock volatilities and using the Fama–French model to forecast a stock's excess expected return. This exhibit shows the out-of-sample (geometric) annualized performance, annualized volatility, and Sharpe ratio of various optimized portfolios over the 1985–2004 period. The abbreviations MSR and GMV represent the maximum Sharpe ratio portfolio and the global minimum variance portfolio, respectively. In all cases, the sample estimate for volatility has been used in this exhibit. The expected return estimate is assumed to be given by the Fama–French factor model for the MSR portfolio. Various estimates of the correlation matrix have been used, including the sample correlation estimate, the constant correlation estimate, and the factor correlation estimate (see details in the text).

and Uppal [2007] and suggest that focusing solely on risk minimization leads to an opportunity cost in terms of performance with respect to a naïvely diversified benchmark. I find that for the MSR portfolios, which were not considered by DeMiguel, Garlappi, and Uppal, higher out-of-sample Sharpe ratios are consistently obtained when volatility is used as a proxy for a stock's expected return, provided that a reasonable (in this case, factor-based) estimator of the correlation structure of stock returns is used. In fact, when using a constant correlation estimator, which is equivalent to throwing away all information from the correlation matrix, Sharpe ratios for the MSR portfolios are already roughly equivalent to those of the equal-weighted benchmarks.[3]

As a robustness check, I also look at the maximum Sharpe ratio portfolios using expected return estimates implied by the Fama–French model that are based on the same competing covariance matrix estimates. The results, reported in Exhibit 4, suggest that using a factor model, not only for the covariance estimates but also for expected return forecasts, leads to lower out-of-sample Sharpe ratios than when total volatility is used as a proxy for expected returns. Furthermore, the out-of-sample Sharpe ratio of the MSR portfolio is no longer greater than the out-of-sample Sharpe ratio of the corresponding GMV portfolio. This result, which stands in sharp contrast to the case in which total volatility is used as an expected return estimate (see Exhibit 3), is consistent with the aforementioned finding that the explanatory power of total volatility in the cross section of expected returns is not driven by systematic volatility with respect to the Fama–French model.

IMPROVING VOLATILITY ESTIMATES

The analysis conducted in the previous sections puts a strong emphasis on volatility inputs which are used as proxies for both the risk and return parameters on individual stock returns. This suggests that particular attention should be devoted to the estimation of these parameters of interest. In particular, a vast amount of empirical evidence has been accumulated documenting that stock market volatility randomly evolves in time. To account for the presence of volatility risk, I extend the

EXHIBIT 5

Performance Indicators Using GARCH Forecasts of Stock Return Volatilities

GARCH Volatility		Expected Return					Volatility					Sharpe Ratio				
		Min	25%	Median	75%	Max	Min	25%	Median	75%	Max	Min	25%	Median	75%	Max
Sample	MSR	-2479%	-1.67%	11.88%	24.80%	3343%	39.56%	52.07%	61.13%	106%	14818%	-0.299	-0.032	0.212	0.354	0.663
Correlation	GMV	-6.52%	3.99%	6.95%	10.35%	21.66%	18.75%	20.72%	21.67%	22.54%	25.95%	-0.321	0.187	0.321	0.456	1.007
Constant	MSR	9.36%	11.02%	11.50%	11.98%	13.56%	13.66%	14.02%	14.16%	14.29%	14.73%	0.644	0.778	0.817	0.846	0.930
Correlation	GMV	5.16%	7.34%	8.12%	9.15%	12.32%	10.90%	11.68%	11.93%	12.15%	13.91%	0.408	0.601	0.692	0.755	1.023
Factor	MSR	6.46%	9.77%	11.26%	12.44%	15.61%	11.15%	12.23%	12.53%	12.90%	14.17%	0.502	0.782	0.902	0.988	1.267
Correlation	GMV	6.21%	8.56%	9.32%	10.52%	12.48%	9.60%	10.38%	10.60%	10.79%	11.93%	0.605	0.797	0.892	0.987	1.203
Equal-Weighted	EW	9.71%	11.67%	12.59%	13.03%	14.81%	14.97%	15.71%	15.95%	16.26%	17.28%	0.601	0.727	0.783	0.812	0.917
Value-Weighted	VW	7.22%	9.09%	9.83%	10.43%	12.12%	15.47%	15.83%	16.07%	16.27%	16.67%	0.455	0.574	0.613	0.651	0.770

Note: Risk and return performance indicators of various portfolios constructed from 100 random draws of 100 stocks, based on GARCH forecasts of stock return volatilities. This exhibit shows the out-of-sample (geometric) annualized performance, annualized volatility, and Sharpe ratio of various optimized portfolios over the 1985–2004 period. The abbreviations MSR and GMV represent the maximum Sharpe ratio portfolio and the global minimum variance portfolio, respectively. In all cases, volatility estimates have been taken to be the quintile estimates. The expected return estimate is assumed to be given by the volatility estimate for the MSR portfolio. Various estimates of the correlation matrix have been used, including the sample correlation estimate, the constant correlation estimate, and the factor-based correlation estimate (see details in the text).

previous analysis to a setup in which variances are forecasted with a simple GARCH (1,1) model with the parameter estimates obtained by a maximum likelihood procedure for a model with normal innovations. Given the date t parameter estimates, estimates of the conditional variance for month $t + 1$ are obtained and used as one-step forecasts each month during the out-of-sample 1985–2004 period.

The results, which are reported in Exhibit 5, show that a simple dynamic volatility model allows for further improvement in the risk-adjusted performance of the MSR portfolio, as well as the GMV portfolio, and strongly dominates its equal-weighted and value-weighted counterparts.

CONCLUSION

In this article, I report empirical evidence on the relationship between a stock's total volatility estimate and its expected return, and examine the implications in terms of portfolio selection. I find that tangency portfolios with attractive out-of-sample risk-adjusted performance properties can be obtained when using volatility as a model-free proxy for a stock's excess return and when a reasonable amount of structure is imposed in the covariance matrix estimation process. These portfolios, which strongly dominate their simple value-weighted counterparts, can be used as convenient benchmarks for normal performance in equity markets.

It should be acknowledged that total volatility is probably not the best possible proxy for excess expected returns. The interesting fact, however, is that using even such a crude and simplistic measure allows for promising results with the potential for outperforming not only market cap–weighted indices, but also equal-weighted benchmarks, in terms of the out-of-sample Sharpe ratio.

This work can be extended in several directions. First, it is tempting to further improve upon the covariance matrix estimation process. Although I use the Fama–French model for simplicity, improved covariance forecasts can instead be generated from an implicit factor model in an attempt to alleviate the concern over factor-specification risk. In addition, a conditional version of the factor model would allow the capture of time variations in correlation and/or volatility parameters. The analysis can and should be extended beyond the mean–variance setting. In particular, recent research has shown that in the presence of non-normally distributed asset returns, investor preferences regarding skewness and kurtosis have an impact on expected returns (see Harvey and Siddique [2000], Dittmar [2002], or Mitton and Vorkink [2007]). In this context, robust estimators for the coskewness and cokurtosis parameters can be introduced, as in Martellini and Ziemann [2007], in an attempt to further enhance the risk-adjusted performance of equity benchmarks. These, as well as other related questions, are left for further research.

ENDNOTES

I would like to thank Daniel Mantilla for superb research assistance, as well as Noël Amenc and Devraj Basu for very useful discussions.

[1]Maximizing the diversification index is formally equivalent to maximizing the portfolio's Sharpe ratio under the assumption that expected excess return parameters are proportional to volatility parameters.

[2]See Basu and Martellini [2007] for the asset pricing implications of the relationship between total volatility and expected return.

REFERENCES

Ang, A., R. Hodrick, Y. Xing, and X. Zhang. "The Cross Section of Volatility and Expected Returns." *Journal of Finance*, Vol. 51, No. 1 (2006), pp. 259–299.

Arnott, R., J. Hsu, and P. Moore. "Fundamental Indexation." *Financial Analysts Journal*, Vol. 61, No. 2 (2005), pp. 83–99.

Barberis, N., and M. Huang. "Mental Accounting, Loss Aversion and Individual Stock Returns." *Journal of Finance*, 56 (2001), pp. 1247–1292.

Basu, D., and L. Martellini. "Total Volatility and the Cross-Section of Expected Stock Returns." Working paper, EDHEC Risk and Asset Management Research Center, 2007.

Britten-Jones, M. "The Sampling Error in Estimates of Mean-Variance Efficient Portfolio Weights." *Journal of Finance*, Vol. 54, No. 2 (1999), pp. 665–671.

Chan, L., J. Karceski, and J. Lakonishok. "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model." *Review of Financial Studies*, 12 (1999), pp. 937–974.

Choueifaty, Y., and Y. Coignard. "Towards Maximum Diversification." Working paper, Lehman Brothers, 2006.

DeMiguel, V., L. Garlappi, and R. Uppal. "Optimal versus Naive Diversification: How Inefficient Is the 1/*N* Portfolio Strategy?" *Review of Financial Studies* Advanced Access, December 3, 2007.

Dittmar, R. "Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns." *The Journal of Finance*, Vol. 57, No. 1 (2002), pp. 369–403.

Elton, E., and M. Gruber. "Estimating the Dependence Structure of Share Prices—Implications for Portfolio Selection." *The Journal of Finance*, Vol. 28, No. 5 (1973), pp. 1203–1232.

Fama, E., and J. MacBeth. "Risk, Return and Equilibrium: Empirical Tests." *Journal of Political Economy*, 81 (1973), pp. 607–636.

Fernholz, R. *Stochastic Portfolio Theory*. New York: Springer-Verlag, 2000.

Fernholz, R., R. Garvy, and J. Hannon. "Diversity-Weighted Indexing." *Journal of Portfolio Management*, Vol. 4, No. 2 (1998), pp. 74–82.

Grinold, R.C. "Are Benchmark Portfolios Efficient?" *Journal of Portfolio Management*, 19 (1992), pp. 34–40.

Harvey, C., and A. Siddique. "Conditional Skewness in Asset Pricing Tests." *The Journal of Finance*, Vol. 55, No. 3 (2000), pp. 1263–1295.

Haugen, R.A., and N.L. Baker. "The Efficient Market Inefficiency of Capitalization-weighted Stock Portfolios." *Journal of Portfolio Management*, Spring 1991, pp. 35–40.

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

Johnstone, I. "On the Distribution of the Largest Eigenvalue in Principal Components Analysis." *Annals of Statistics*, Vol. 29, No. 2 (2001), pp. 295–327.

Jorion, P. "Bayes–Stein Estimation for Portfolio Analysis." *Journal of Financial and Quantitative Analysis*, Vol. 21, No. 3 (1986), pp. 279–292.

Laloux, L., P. Cizeau, J.P. Bouchard, and M. Potters. "Random Matrix Theory." *Risk Magazine*, 12 (1999), p. 69.

Ledoit, O., and M. Wolf. "Honey, I Shrunk the Sample Covariance Matrix." *Journal of Portfolio Management*, Vol. 30, No. 4 (2004), pp. 110–119.

Malkiel, B., and Y. Xu. "Risk and Return Revisited." *Journal of Portfolio Management*, 23 (1997), pp. 9–14.

———. "Idiosyncratic Risk and Security Returns." Working paper, University of Texas at Dallas, 2002.

———. "Idiosyncratic Risk and Security Returns." Working paper, University of Texas at Dallas, 2006.

Markowitz, H. "Portfolio Selection." *Journal of Finance*, Vol. 7, No. 1 (1952), pp. 77–91.

Martellini, L., and V. Ziemann. "Improved Forecasts of Higher-Order Moments and Implications for Portfolio Selection." Working paper, EDHEC Risk and Asset Management Research Center, 2007.

Merton, R. "A Simple Model of Capital Market Equilibrium with Incomplete Information." *Journal of Finance*, 42 (1987), pp. 483–510.

Michaud, R. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. Boston, MA: Harvard Business School Press, 1998.

Mitton, T., and K. Vorkink. "Equilibrium Underdiversification and the Preference for Skewness." *Review of Financial Studies*, Vol. 20, No. 4 (2007), p. 1255.

Sharpe, W. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *Journal of Finance*, Vol. 19, No. 3 (1964), pp. 425–442.

Tinic, S., and R. West. "Risk, Return, and Equilibrium: A Revisit." *Journal of Political Economy*, Vol. 94, No. 1 (1986), pp. 126–147.