Demystifying Equity Risk–Based Strategies: *A Simple Alpha plus Beta Description*

RAUL LEOTE DE CARVALHO, XIAO LU, AND PIERRE MOULIN

RAUL LEOTE DE CARVALHO
is head of quantitative strategies and research in financial engineering at BNP Paribas Investment Partners in Paris, France.
raul.leotedecarvalho@bnpparibas.com

XIAO LU
is a quantitative analyst in financial engineering at BNP Paribas Investment Partners in Paris, France.
xiao.lu@bnpparibas.com

PIERRE MOULIN
is head of financial engineering at BNP Paribas Investment Partners in Paris, France.
pierre.moulin@bnpparibas.com

Equity risk–based strategies are systematic quantitative approaches to stock allocation that rely only on risk views to manage risk and increase diversification. These strategies do not require any explicit stock return forecasts. The portfolios are periodically rebalanced to take into account drift and changes in risk views.

The simplest of these strategies is based on the equally weighed (EW) portfolio that simply follows the principle of not putting all your eggs in one basket. The portfolio invests the same amount in each stock. The strategy makes sense if we believe that neither stock returns nor risk can be forecast.

The equal-risk budget (ERB) strategy invests in portfolios with the same risk budget for each stock, which is defined as the product of the stock's weight and its volatility. Risk is equally distributed among the stocks and hence riskier stocks get smaller weights. This can be viewed as an extension of EW, if we trust volatility forecasts.

If correlations are also taken into account, then we can think in terms of equal-risk contribution (ERC), where the contribution to risk from each stock is the same. Unlike the risk budget, the contribution to risk (defined as the product of the stock's weight and its marginal risk)[1] also takes into account the impact of correlations. The contribution to portfolio risk from two stocks with the same volatility but different correlations is higher for the stock with the higher correlation, and hence it gets a smaller weight in the ERC portfolio. The ERC strategy was discussed recently by Maillard, Roncalli, and Teiletche [2010].

These three strategies assume diversification can be achieved by equally allocating wealth or risk across the investment universe. The two other risk-based strategies we analyze are different.

Minimum variance (MV) invests in the portfolio with the lowest ex ante volatility. MV is the least risky approach to investing in equities and is expected to deliver the lowest volatility over time. It uses volatilities and correlations as inputs and should invest in stocks with the lowest volatility and low correlations.

The maximum diversification (MD) strategy, introduced by Choueifaty and Coignard [2008], invests in the portfolio that maximizes a diversification ratio. The ratio is the sum of the risk budget allocated to each stock in the portfolio divided by the portfolio volatility. This strategy should invest in stocks that are less correlated to other stocks.

Haugen and Baker [1991] and Clarke, de Silva, and Thorley [2006, 2011] investigated the MV strategy; Choueifaty and Coignard [2008] investigated the MD, MV, and EW strategies; and Demey, Maillard, and Roncalli [2010] investigated the EW, ERC, MV, and MD strategies. All reported that these strategies have been outperforming the market-cap index, and that all, bar EW, have done so with lower volatility. In the capital asset pricing model (CAPM), stock returns are proportional to the stock beta, and the market-cap portfolio is already the most diversified with the highest risk-adjusted return. Thus, the empirical results reported in the extant literature contravene the CAPM. These authors also noted that the variation of excess returns of the MV, MD, and ERC strategies could not be fully explained by exposure to the value and size Fama–French [1992] factors.

Scherer [2010] showed that the MV strategy is essentially exposed to two risk-based pricing anomalies: 1) low-beta stocks delivering higher returns than high-beta stocks, even after adjusting for beta, and 2) low residual volatility stocks delivering higher returns than high residual volatility stocks, also after adjusting for beta.

The low-beta pricing anomaly has been documented for some time. Fama and French [2004] and Baker, Bradley, and Wurgler [2011] gave empirical evidence that low-beta stocks do outperform the CAPM prediction with the converse being true for high-beta stocks. Messikh and Oderda [2010] presented a theoretical proof of the beta anomaly in the case of stock prices that are modeled using Brownian processes.

There is mixed evidence of a pricing anomaly for stocks with low residual risk. Scherer [2010] and Ang et al. [2006] found that stocks with high residual volatility generate lower returns than those predicted by CAPM, with the converse being true for low residual volatility stocks. But as shown by Tinic and West [1986] and Malkiel and Xu [1997, 2002], if stocks are ranked by average residual volatility in portfolios of stocks previously sorted by beta or market cap, then there is a negative relationship between residual volatility and average returns.

OVERVIEW

In this article, we look at the five risk-based strategies EW, ERB, ECR, MV, and MD in detail. We apply them to a universe of global stocks from developed countries going back to 1997 and free of survivorship bias. We also consider them in the case of the U.S., Europe, and Japan. In all cases, we show that the variation of excess returns over the market-cap index of all five strategies can be largely explained by a five-factor model inspired by that of Sherer [2010]. We find that all five strategies can be almost fully explained by exposure to the market and to the four additional factors: value, small caps, low beta, and low residual volatility. The regressions show no additional alpha beyond that generated from these factors, with the intercept of regressions equal to zero.

In that sense, we have translated the portfolios behind these five risk-based portfolio construction algorithms for which investors may have some aversion due to the poor visibility of what goes on behind the scenes into a language that is well understood by most investors these days: the market-cap index as the benchmark plus stock active weights determined by each stock's exposure to some simple factors.

We show that EW, ERB, and ERC are not too different from each other. The three have comparable turnover and invest in all the stocks included in the investment universe with EW more exposed to small-cap stocks than ERB, which is in turn more exposed than ERC. ERB and ERC are increasingly exposed to low-beta stocks and are defensive strategies with ERC being the most defensive.

MV and MD are different. Both tend to have short positions. Therefore, we considered an unconstrained and a long-only constrained version of each. Their long-only versions invest in a small number of stocks. For the global universe of stocks in the developed world, that number is around 120 stocks for both MV and MD. Their unconstrained versions invest in about 60% of the stocks in the global universe and sell short the other 40% of stocks. Both MV and MD require large portfolio turnover. The unconstrained versions require an even larger turnover than the long-only versions. Both MV and MD—both the long-only and unconstrained versions—are defensive strategies exposed to low-beta stocks.

The difference between MV and MD arises from a larger exposure to low residual volatility stocks in MV. We discover that MV reaches levels of diversification not too far from those of MD, when using the diversification ratio that MD maximizes as its measure of diversification. The diversification ratio of both MV and MD is much larger than that of EW, ERB, ERC, or the market-cap index. The unconstrained versions are even more diversified than the long-only versions. MV and MD, both the long-only and unconstrained versions, had little or no exposure to small-cap stocks, which were the key driver behind EW and also important in ERB and ERC.

We also looked at the historic average overlap of the portfolios behind each strategy and the correlation of the excess returns of each strategy to the market-cap index returns in order to show how remarkably high they are for EW, ERB, and ERC. We also show that the commonality between MV and MD is much higher than would perhaps be expected. This commonality prevails when we compare their respective long-only versions or their unconstrained versions. But between the first three strategies and MV and MD, long-only or unconstrained, there is less in common.

Finally, we compared results obtained using two different risk models—principal components analysis, and Bayesian shrinkage—to show that the behavior of ERC, MV, and MD is relatively insensitive to the choice of risk model. We also looked at the impact of changing the period of time used in the estimation of the risk model and the frequency of data. We found that these parameters had little impact on the results.

RISK-BASED PORTFOLIOS

The market-cap index simply allocates stock weights according to their market cap, $M_i = n_i p_i$, which is the product of the number of shares outstanding n_i of company i and its current market price, p_i,

$$w_i = \frac{M_i}{\sum_j M_j} \tag{1}$$

The market-cap index is mean–variance efficient if the CAPM holds (i.e., if stock returns are equal to the stock beta times the index excess returns over the risk-free rate). This is clearly the strategy with the lowest turnover, reflecting occasional changes in the composition of the indices and reinvestment of dividends.

For a universe of N stocks, the EW portfolio weight allocates the same dollar amount to each stock,

$$w_i = \frac{1}{N} \tag{2}$$

This portfolio is mean–variance efficient, maximizing the Sharpe ratio, if the returns and volatility are the same for all stocks and all pair-wise correlations are equal. It is clear that, when compared to the market-cap index, the EW portfolio will overweight small-cap stocks and underweight large-cap stocks. The larger the dispersion of the stocks' capitalization, the larger is the difference between the EW and MC portfolios.

If sigma, σ_i, is the volatility of stock I, then the risk budget, $w_i \times \sigma_i$, allocated to each stock is the same for each stock in the ERB portfolio, and the weight of stock i is

$$w_i = \frac{1/\sigma_i}{\sum_j 1/\sigma_j} \tag{3}$$

If the Sharpe ratio for each stock is the same and all pair-wise correlations are equal, then the ERB portfolio is mean–variance efficient with the highest possible Sharpe ratio. We can think of the ERB as an EW portfolio tilted in favor of low-risk stocks and away from high-risk stocks. Therefore, the ERB portfolio is not only overweight small-cap stocks relative to the MC index, it is also overweight low-volatility stocks, which tend to have a low beta.[2]

The ERC portfolio allocates the same risk contribution to each stock. Thus, the stock weights in the ERB portfolio are such that the difference of the contribution to risk of any two stocks is equal to zero. The stock weights cannot be written in a closed form, but as Maillard, Roncalli, and Teilietche [2008] proposed, they can be found when minimizing the sum of the squares of the difference between the contributions to risk of any two stocks,

$$w^* = \arg\min\left(\sum_i \sum_j \left(w_i \left(\Sigma w\right)_i - w_j \left(\Sigma w\right)_j\right)^2\right) \quad \text{u.c.}$$
$$\times \sum_i w_i = 1 \text{ and } w_i \geq 0 \tag{4}$$

The ERC portfolio can be viewed as an ERB portfolio that is tilted toward the stocks less correlated with other stocks. It is not difficult to show that there are multiple solutions to Equation (4) unless we constrain the solution to be positive. Therefore, we consider the long-only solution, which matters only to long-only investors. Maillard, Roncalli, and Teilietche [2008] proved that the ex ante volatility of the ERC portfolio is always between those of the EW and MV portfolios. Another property of the ERC portfolio is that $\tilde{\beta}_i w_i = 1/N$ where $\tilde{\beta}_i$ is the beta of stock i estimated against the ERC portfolio, and w_i is the weight of stock i in the ERC portfolio.

Using the Kuhn–Tucker conditions, Maillard, Roncalli and Teilietche [2008] proved that the ERC portfolio must also be a solution of the following optimization problem:

$$w^* = \arg\min \sqrt{w'\Sigma w} \quad \text{u.c.} \quad \begin{cases} \sum_i \ln w_i \geq c \\ w_i \geq 0 \end{cases} \tag{5}$$

The solution is exactly the same if we replace the constraint $\Sigma_i \ln w_i \geq c$ in the optimization problem with $\Sigma_i \ln w_i = c$. Replacing w^* with $w^*/\Sigma_i w_i^*$, we can show that this optimization problem is exactly equivalent to

$$w^* = \arg\min \frac{\sqrt{w'\Sigma w}}{(\prod_i w_i)^{1/N}} \quad \text{u.c.} \quad \begin{cases} \sum_i w_i = 1 \\ w_i \geq 0 \end{cases} \tag{6}$$

N is the total number of stocks. This minimization problem under long-only and fully invested constraints falls somewhere between the minimization of $\sqrt{w'\Sigma w}$, which is the MV portfolio under a long-only constraint, and the maximization of $(\prod_i w_i)^{1/N}$, which is satisfied by the EW portfolio. The ERC portfolio is a trade-off between minimizing variance and equally weighting all stocks. Thus, we expect to find some small-cap bias, which characterizes the EW portfolio, as well as bias toward low-beta stocks and low residual volatility stocks, which is a feature of the MV portfolio.

If the Sharpe ratio for each stock is exactly the same and all pair-wise correlations are equal, then the ERC and ERB portfolios have the same stock allocation. In addition, if the volatility of all stocks is the same, then the ERC and ERB portfolios turn into the EW portfolio.

For a small number of stocks, the ERC portfolio can be found by numerically minimizing the problem stated in Equation (4). This approach fails, however, for a large universe of stocks, such as the MSCI World Index global universe of about 1,700 stocks. Instead, we use an iterative procedure to find the ERC solution when the number of stocks is large. We start with a solution that is not too far from the final ERC portfolio, then we increase the weights of stocks for which $\tilde{\beta}_i w_i < 1/N$ and decrease the weights of stocks for which $\tilde{\beta}_i w_i > 1/N$, where $\tilde{\beta}_i$ is the beta of stock i estimated against the current estimated solution. This procedure is repeated until the solution converges toward the ERC. A good starting portfolio is a minimum-variance portfolio optimized with a penalty on its tracking error against the EW portfolio. This starting point is inspired by Equation (6). Our target is to find, on average, 99% of the stocks in the final ERC solution with a risk contribution between 97% and 103% of the average risk contribution of all stocks invested.

The MV portfolio is the portfolio allocation with the lowest possible ex ante variance,

$$w^* = \arg\min\left(w'\Sigma w\right) \quad \text{u.c.} \quad \sum_i w_i = 1 \tag{7}$$

The portfolio is likely to contain short positions. Therefore, it is also appropriate to add a long-only constraint. The MV portfolio is always optimal in a mean–variance sense. In the case when all stock returns are equal, the MV portfolio is also the mean–variance portfolio, which maximizes the Sharpe ratio. Clarke, de Silva, and Thorley [2011] used a one-factor model (CAPM) to show that the unconstrained MV portfolio solution can be written as

$$w_i \approx \frac{\kappa}{\sigma_{\varepsilon,i}^2}(1 - \beta_i) \tag{8}$$

where β_i and $\sigma_{\varepsilon,i}^2$ represent the beta and residual variance, respectively, of the stock i and κ represents a normalization constant. This shows that the MV portfolio will be tilted toward the stocks that have the lowest betas and lowest residual volatilities.

The MD portfolio is the solution to the problem of maximization of a diversification ratio, D_r, which is defined as the ratio of the weighted average of stock volatility for the stocks in the portfolio to the actual portfolio volatility,

$$w^* = \arg\max(D_r) \quad \text{with} \quad D_r = \frac{\sigma' w}{\sqrt{w'\Sigma w}} \tag{9}$$

We may expect the solution to contain short positions. Therefore, a long-only constrained version may be of more interest to investors. If all stocks have the same volatility, then the MD and MV portfolios are equal. And if all stocks have the same Sharpe ratio, then the MD portfolio is mean–variance efficient and is the portfolio with the maximum Sharpe ratio. Using a one-factor model (CAPM), the unconstrained MD portfolio can be written as

$$w_i = \frac{\lambda}{\sigma_{\varepsilon,i}^2}(\sigma_i - \bar{\sigma}\beta_i) \quad \text{with} \quad \bar{\sigma} = \sum_i \frac{\beta_i \sigma_i}{\sigma_{\varepsilon,i}^2} \Big/ \sum_i \frac{\beta_i}{\sigma_{\varepsilon,i}^2} \qquad (10)$$

where $\bar{\sigma}$ is a weighted average of the stock volatilities to all stocks, and λ represents a normalization constant. A Taylor expansion will lead to

$$w_i \approx \lambda\left(\frac{1}{2\beta_i \sigma_{\text{MKT}}} + \frac{\beta_i}{\sigma_{\varepsilon,i}^2}(\sigma_{\text{MKT}} - \bar{\sigma})\right) \qquad (11)$$

Because stock pair-wise correlations are smaller than 1, $\bar{\sigma}$ is always larger than σ_{MKT}, the volatility of the market-cap index. Numerically, we find that $\bar{\sigma}$ falls between $1.5 \times \sigma_{\text{MKT}}$ and $2.5 \times \sigma_{\text{MKT}}$ using data for global stocks in the MSCI World Index from January 1997 through December 2010. Because $\sigma_{\text{MKT}} - \bar{\sigma}$ is always negative, the stock weight will decrease with the stock beta, and the MD portfolio will invest in the stocks with the lowest betas. For two stocks with the same beta, the MD will give a higher weight to the one with higher residual volatility. But because the first term, $1/(2\beta_i\sigma_{\text{MKT}})$, dominates the second term, $\beta_i(\sigma_{\text{MKT}} - \bar{\sigma})/\sigma_{\varepsilon,i}^2$, particularly for low beta stocks, the residual risk effect is much less important.

DATA AND METHODOLOGY

We run our simulations using weekly total returns for the stocks in the MSCI World Index of developed countries from January 1997 to December 2010.[3] Our data source is the Exshare database. We use log-returns and calculated all returns in U.S. dollars. The one-month Treasury bill rate is the proxy for the risk-free rate. We apply the five risk-based strategies to a global universe of stocks and also to the sub-universes of U.S., European, and Japanese stocks.

We simulate the performance of the strategies with quarterly rebalancing on the third Friday of March, June, September, and December of each year. The portfolios are allowed to drift between rebalancing dates. No constraints are imposed on the EW and ERB portfolios, which are long-only by construction. There are multiple solutions for the ERC strategy, so we retain it with positive weights. For the MV and MD strategies we consider two versions, the first is constrained by capping stock weights at a maximum of 5% on each rebalancing date, and the second is fully unconstrained, allowing for short positions.

For the ERC, MV, and MD strategies, we use a principal component analysis (PCA) risk model following the methodology proposed by Plerou et al. [2002] that considers results from random matrix theory. Plerou et al. showed that the eigenvalues λ of a $T \times N$ random matrix with variance σ^2 are capped asymptotically at

$$\lambda_{\max} = \sigma^2\left(1 + N/T + 2\sqrt{N/T}\right) \qquad (12)$$

with T the number of periods and N the number of stocks. Thus, we discard all eigenvalues smaller than $\lambda_{\max}$, which we consider to be statistical noise.

At each rebalancing, we update the risk model from a two-year rolling window of the most recent historical data. We keep only stocks for which there are at least 94 weeks of returns available.

In our analysis of results, we report the average performance, volatility, tracking error risk, and average excess returns against the market-cap index, average beta, and turnover. We also look at the correlation of the excess returns of the strategies to the market-cap index returns and at the average overlap of the portfolios of each strategy.

Finally, we use an extended Fama–French regression model to explain the variation of excess returns of each strategy over the market-cap index inspired by the model introduced by Scherer [2010]. He proposed an extension of the Fama–French model, which includes two additional factors designed to capture risk-based pricing anomalies. He called these the low-beta anomaly and the low residual volatility anomaly. They are based on the empirical observation that the stocks with the lowest beta and the stocks with the lowest residual volatility seem to have delivered higher returns than predicted by CAPM. He used two cash-neutral and beta-neutral long–short factor portfolios to estimate the returns of the risk-based factors: one is long low-beta stocks and short large-beta stocks, and the other is long low residual volatility stocks and short large residual volatility stocks. We find, however, that the returns of these two portfolios are correlated by more than 60%. In order to reduce this correlation, we use an orthogonalization approach similar to that of the Fama–French HML and SMB factors.[4]

Next, we describe how we build the factor portfolios. All factor portfolios are rebalanced quarterly at the same time that the strategies are rebalanced.

Market Rate over Risk-Free Rate Factor (MKT)

We use the return of the market-cap index minus the U.S. one-month T-Bill rate as a proxy for the risk-free rate.

Value and Size Factors (HML and SMB)

The Fama–French approach to value and size factors is well known and described on the Kenneth French website.[5] Ken French uses a large universe of U.S. stocks, but we re-estimate the factors, keeping only the stocks in our investment universe.[6] For the HML (high-minus-low) factor (value anomaly), we first rank stocks into quintiles by market cap, and then in each quintile, we rank stocks by quintiles of book-to-market value in order to form 25 portfolios. The returns of the HML factor are the returns to a portfolio that is long the stocks with the highest book-to-market value in each quintile of market cap and short the stocks with lowest book-to-market value.

The returns to the SMB (small-minus-big) factor (small-cap stock anomaly) are built using a similar procedure in which we first rank stocks by quintiles of book-to-market value and then by quintiles of market cap.

Beta and Residual Volatility Factors (LBMHB and LRVMHRV)

After estimating the beta and residual volatility of each stock from the two years of historical data, we then use a similar approach to that described for the HML and SMB factors, but in addition we neutralize beta.[7] For LBMHB, we first rank stocks by residual volatility into quintiles, and then by beta in each volatility quintile. The returns of the LBMHB factor are given by

$$R_{SBMLB} = (R_{SB} - R_f) - \kappa(R_{LB} - R_f) \tag{13}$$

where $\kappa = \beta_{SB}/\beta_{LB}$ · β_{SB} is the ex ante beta and R_{SB} is the return of the portfolio with the quintile of lowest-beta stocks in each quintile of residual volatility. Similar definitions exist for β_{LB} and R_{LB}, but considering the largest-beta stocks.

The LRVMHRV factor is built using a similar approach, but the stocks are first ranked by beta and then by residual volatility. Beta is also neutralized in a similar way.

In Exhibit 1, we show the performance and volatility of each factor for global stocks estimated for the period January 1997–December 2010. The LBMHB factor has the highest Sharpe ratio followed closely by the HML and then by the SMB, which is also positive.

The Sharpe ratio for the LRVMHRV factor is small and negative. This result appears initially to contradict Scherer [2010] and Ang et al. [2006], who found that low residual volatility stocks outperformed high residual volatility stocks. The difference is that we have orthogonalized the low-beta factor and the low residual volatility factor. We believe that the results found by Scherer and Ang et al. are due to the high correlation between the two factors. After the orthogonalization is performed, the pricing anomaly for low-beta stocks seems to prevail whereas that for low residual volatility stocks does not, at least in a global universe of stocks for the sample period considered. We discuss the results for the U.S., Europe, and Japan later in the article.

The factor correlations shown in Exhibit 2 are low. The correlations between the SMB and HML factors and the LBMHB and LRVMHRV factors are indeed small,

EXHIBIT 1
Factor Annualized Returns and Volatility: World Universe (January 1997–December 2010)

	Mkt-Rf	HML	SMB	LBMHB	LRVMHRV
Return	2.1%	3.7%	2.1%	4.1%	–1.9%
Volatility	18.1%	13.2%	9.8%	10.2%	11.6%
Sharpe Ratio	0.12	0.32	0.20	0.42	–0.14

EXHIBIT 2
Correlations of Factor Returns: World Universe (January 1997–December 2010)

	Mkt-Rf	HML	SMB	LBMHB	LRVMHRV
Mkt-Rf		–23%	–7%	36%	–8%
HML			3%	25%	3%
SMB				20%	–48%
LBMHB					–26%
LRVMHRV					

as expected, due to the orthogonalization process used in their construction. SMB and LRVMHRV tend to be somewhat negatively correlated at –48%, which could explain the small negative return in the low residual volatility factor. This correlation arises from the fact that stocks with high residual volatility tend to be small-cap stocks, and stocks with low residual volatility tend to be large-cap stocks.[8]

NUMERICAL RESULTS FOR GLOBAL STOCKS

In Exhibit 3, we show the results from our historical backtests of the risk-based strategies EW, ERB, ERC, MV and MD long-only constrained with stock weights capped at 5%, and MV and MD fully unconstrained (allowing for short positions). We use the PCA risk model for the global universe of stocks from January 1997 to December 2010. All results are based on weekly returns. Transaction costs are not considered, a detailed analysis of transaction costs will be considered elsewhere.[9]

All strategies have outperformed the market-cap index since 1997, and all, bar EW, have lower volatility. The Sharpe ratio of the strategies is higher than that of the market-cap index. The MV and MD strategies have the smallest ex post volatility and the highest tracking error against the market-cap index. We find that the unconstrained versions of MV and MD have even lower volatility and higher excess returns over the market-cap index than their long-only versions, resulting in a dramatic improvement in the Sharpe ratio. The tracking error against the market-cap index is also larger, but the improvement in excess returns is enough to make the information ratio higher.

The long-only MV and MD versions invest, on average, in only about 120 stocks of 1,700. The unconstrained MV invest in about 1,000 stocks and sell short, on average, around 700 stocks. The unconstrained MD invests in about 1,100 stocks and shorts about 600 on average. The other three strategies (EW, ERB, and ERC) invest in the entire investment universe.

The turnover of MV and MD is comparable and much higher than that of the other strategies. The turnover in the unconstrained versions is the highest. We believe that the high turnover is due to the higher exposure of MV and MD to noise in the risk model. The turnover can be reduced substantially without a major impact on performance or risk by imposing a sensible turnover constraint.[10]

As expected, MD has the highest diversification ratio. The average diversification ratios of MD and MV are much higher than those for EW, ERB, or ERC, which are actually close to those of the market-cap index. It is, however, interesting to observe that the diversification ratio of MV is not much lower than that of MD. The MV portfolios are already almost as diversified as the MD portfolios when the diversification ratio defined in Equation (9) is used as the measure of diversification.

MV and MD are defensive strategies with a very low beta, particularly in their unconstrained versions. ERB and ERC also show some defensive characteristics with a beta lower than one, which is more pronounced for ERC, but less pronounced than for MV and MD.

EXHIBIT 3

Simulation Results for Risk-Based Strategies: World Universe (January 1997–December 2010)

	Mkt	EW	ERB	ERC	MV Long Only	MD Long Only	MV	MD
Excess Return over RF	2.1%	5.7%	5.9%	5.6%	5.2%	4.8%	6.3%	6.2%
Volatility	18.1%	18.2%	16.5%	14.8%	9.9%	11.5%	9.1%	10.8%
Sharpe Ratio	0.12	0.31	0.36	0.38	0.52	0.41	0.70	0.58
Excess Return over BM		3.6%	3.8%	3.5%	3.1%	2.6%	4.2%	4.1%
Tracking Error		5.1%	5.4%	6.6%	13.1%	12.1%	15.2%	14.5%
Information Ratio		0.70	0.70	0.52	0.23	0.22	0.28	0.28
Beta		0.96	0.87	0.76	0.39	0.48	0.27	0.36
Maximum Drawdown	–56%	–58%	–55%	–53%	–29%	–39%	–22%	–31%
Annual Turnover	18%	39%	37%	58%	151%	162%	220%	296%
Diversification Ratio	2.2	2.5	2.5	2.8	5.6	6.3	8.6	11.7

MV and MD have the smallest drawdowns, which are smaller in the unconstrained versions. Even if they seem to outperform the market-cap index in the long term, they can underperform significantly at times. The long-only versions of MV and MD underperform the market-cap index by as much as 41% and 38%, respectively, from the mid-1990s to the peak of the market in March 2000. All risk-based strategies that we consider in our analysis generate positive returns in the bull market period but fail to outperform the market-cap index. The Sharpe ratio in this period is lower for all the risk-based strategies than for the market-cap index.

It is interesting to observe that the MV strategy has the highest Sharpe ratio. Blitz and van Vliet [2007] observed that when global stocks are ranked by historical volatility into deciles, the average returns between January 1986 and December 1995 of each decile portfolio were comparable, whereas volatility of the decile portfolio with high-volatility stocks was about twice that of the decile portfolio with low-volatility stocks. Fama and French [2004] formed 10 value-weighted portfolios (using the CRSP database for U.S. stocks) based on ranked betas and computed their returns for the next 12 months in the period 1928–2003. They also found comparable average returns for the 10 portfolios. This suggests that, in first approximation, we can assume that all stocks have the same expected return irrespective of their risk or beta. It happens that when all stocks have the same expected return, the portfolio that maximizes the ex ante Sharpe ratio is the minimum-variance portfolio. Not only does the MV strategy have the highest Sharpe ratio when compared to the other risk-based strategies, but it is also possible that the MV has been one of the most efficient strategies over the last few decades in terms of maximizing the return per unit of risk.

In Exhibit 4, we show the correlation of weekly excess returns over market cap for the risk-based strategies. The excess returns of EW are 89% correlated with those of ERB, which in turn are 93% correlated with ERC. The correlation between the excess returns of the long-only MV and long-only MD is 96%. The correlation between the excess returns of the unconstrained MV and unconstrained MD is 93%. This clearly confirms that much is shared among the EW, ERB, and ERC strategies and also between the MV and MD strategies, which are poorly correlated with the EW strategy. Their correlation is higher with the ERB strategy, however, and above 80% with the ERC strategy.

The correlation between excess returns of the long-only MV and the unconstrained MV is 94%. For MD, it is 92%. The long-only versions, despite investing in for fewer stocks, seem to retain much of the underlying features of their unconstrained versions.

In Exhibit 5, we compare the average overlap of the underlying portfolios (i.e., the historical average of the sum of the smallest weight of each stock present in the portfolios of the two strategies). The overlap among the EW, ERB, and ERC portfolios is extremely high at over 80% and is consistent with the correlation of excess returns over the market-cap index. These three strategies exhibit an average overlap with the market-cap index of nearly 50%. The overlap between the long-only MV and long-only MD is 61%. This is actually remarkably high considering that the two strategies invest in only about 120 stocks of the 1,700 stocks available. The overlap of the long-only MV and long-only MD with the EW, ERB, ERC portfolios and the market-cap index is small.

On average, the unconstrained MV and MD strategies invest in 850 stocks in common of the 1,000 stocks for MV and the 1,100 stocks for MD and sell short 450 stocks in common of the 700 stocks for MV and 600 stocks for MD.

The results from the regression of the log-excess returns over market-cap index log-returns against the five factors previously described can be found in Exhibit 6. The regressions are based on 732 weekly returns from

EXHIBIT 4

Correlation of Excess Returns of Risk-Based Strategies over the Market-Cap Index Returns: World Universe (January 1997–December 2010)

	EW	ERB	ERC	MV Long Only	MD Long Only	MV	MD
EW		89%	72%	25%	33%	17%	30%
ERB			93%	57%	61%	52%	56%
ERC				81%	83%	75%	79%
MV Long Only					96%	94%	92%
MD Long Only						89%	92%
MV							93%
MD							

Exhibit 5

Historical Average Overlap between Risk-Based Strategy Portfolios: World Universe (January 1997–December 2010)

	Mkt	EW	ERB	ERC	MV Long Only	MD Long Only
Market Cap		49%	49%	45%	5%	5%
EW			88%	77%	6%	7%
ERB				82%	8%	7%
ERC					17%	18%
MV Long Only						61%
MD Long Only						

January 1997 to December 2010. The regression coefficients, R^2, and Durbin–Watson test statistics are shown.

The factor regressions explain the variation of the excess returns of each strategy rather well with remarkably high R^2s that range from 75% for the EW strategy to 87% for the unconstrained MV strategy and with regression intercepts virtually equal to zero. The estimated slope of the regression line is very close to one in all cases. The fact that the factors are built using only stocks in the available investment universe adds to the quality of the regression results. That is important because our aim is to understand how the different risk-based strategies invest when applied to a given universe of stocks.

The EW approach is essentially explained by exposure to small-cap stocks. ERB and ERC put increasing emphasis on low-beta stocks and decreasing emphasis on small-cap stocks, which is more pronounced for ERC than for ERB. Both ERB and ERC exhibit a defensive beta, but with that of ERC being lower. We also observe a small exposure to value stocks in EW, ERB, and ERC, but this seems much less important.

MV and MD are explained by a strong positive exposure to low-beta stocks and a very low beta. Additionally, MV is also exposed to low residual volatility stocks, but MD is not, which is consistent with our expectations. We find little exposure in MV and MD to small-cap stocks or to value stocks. It is reassuring that the factor exposures are in line with the theoretical expectations.

The unconstrained MV and MD versions have similar factor exposures to those of their long-only versions, with marginally higher R^2s. The defensive character of MV and MD is accentuated by even lower exposures to the market than the long-only versions. The exposure to low-beta stocks is also slightly increased. For MV, the exposure to low residual volatility stocks also increases, whereas for MD the exposure turns marginally negative. The exposure to small-cap stocks remains very small and is even lower, turning marginally negative for MV. The exposure to value stocks is also very small and remains marginally negative.

We also consider a six-factor model in which we include a momentum factor (Jegadeesh and Titman [1993]), using the same approach described by Kenneth French.[11] We find that momentum does not increase the explanatory power of the regressions, although we had no compelling reason to expect these strategies to be exposed to momentum. Clarke, de Silva, and Thorley [2006] already found that momentum does not add explanatory power to the particular case of the MV strategy for U.S. stocks.

Exhibit 6

Factor Regression Coefficients for Risk-Based Strategies: World Universe (January 1997–December 2010)

	EW	ERB	ERC	MV Long Only	MD Long Only	MV	MD
Intercept	0.000[b]	0.000[b]	0.000	0.000	0.000	0.000	0.000
Mkt-Rf	–0.013[c]	–0.119[a]	–0.258[a]	–0.708[a]	–0.629[a]	–0.847[a]	–0.763[a]
HML	0.135[a]	0.155[a]	0.122[a]	–0.020	–0.040[c]	–0.084[a]	–0.062[b]
SMB	0.373[a]	0.315[a]	0.241[a]	0.025	0.059[c]	–0.059[c]	0.012
LBMHB	0.060[a]	0.160[a]	0.287[a]	0.543[a]	0.567[a]	0.637[a]	0.585[a]
LRVMHRV	–0.019[c]	0.090[a]	0.060[a]	0.124[a]	0.044[c]	0.196[a]	–0.087[a]
R^2	0.75	0.79	0.85	0.84	0.78	0.87	0.79
DW	2.08	1.99	2.05	2.20	2.15	2.22	2.20

Note: Significance levels at 0.1%, 1% and 5% are marked by a, b, and c, respectively.

IMPACT OF RISK MODELS

In order to assess the impact of a changing risk model we repeat our simulations using the Bayesian shrinkage approach proposed by Ledoit and Wolf [2003]. The final covariance matrix is a weighted average of the historical covariance matrix and the CAPM prior matrix with the weight and the prior (shrinkage parameter) estimated from historical returns. We find that all of our results are practically the same whether we use a PCA or a Bayesian approach. In Exhibit 7, we show as an example the cumulative returns

Exhibit 7

Cumulate Performance of a Long-Only MV Strategy Using PCA and Bayesian Risk Models: World Universe (January 1997–December 2010)

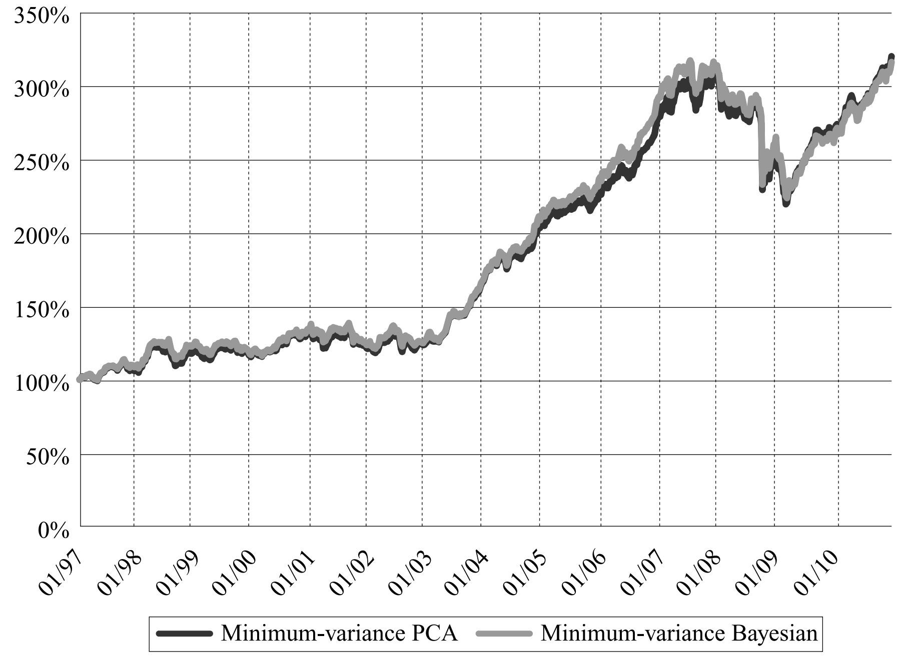

of the long-only MV strategy simulated using the two different risk models. The average annualized excess return of the long-only MV strategy using a Bayesian risk model is 5.3% higher than the risk-free rate and using PCA is 5.2% higher than the risk-free rate. The volatility is unchanged at 9.9%. We also look at the impact of changing the data frequency and increasing the window of estimation from two years to three, four, and five years. Again, we find little impact on the results of our analysis.

Choueifaty and Coignard [2008] also found little impact from changing the risk model and its parameters. They investigated the impact on long-only MV and MD portfolios, using a variety of approaches to estimate the covariance matrix, including simple rolling windows of different length, decayed weighting schemes, and GARCH and Bayesian approaches. They also found that changing the risk model, the frequency of data, and the estimation period had little impact on results and that even portfolios built on forward-looking covariance matrices (having perfect covariance foresight) are only slightly different. Clarke, de Silva, and Thorley [2006] also found little difference between the results of long-only and unconstrained MV strategies based on Bayesian and PCA risk models.

NUMERICAL RESULTS AT REGIONAL LEVEL

In order to check whether our results for the global universe of stocks are verified at the regional level, we apply the risk-based strategies to U.S., European, and Japanese stocks using their respective MSCI stock universes. In Exhibit 8, we show the results from the simulations for EW, long-only MV, and long-only MD with the stock weight limited to 5%.

Indeed, the results are in line with those found for global stocks. In all three regions, the risk-based strategies outperform the market-cap index and all, bar EW, have lower volatility. The MV strategy has the lowest volatility, largest Sharpe ratio, and largest tracking error compared to the market-cap index. EW has the lowest tracking error in all regions.

We build factor returns for each region using the same approach described for global stocks, with all returns in U.S. dollars. The performance of the factors is reported in Exhibit 9. During this period, the HML value factor shows good performance in Europe.

EXHIBIT 8

Simulation Results for Risk-Based Strategies: U.S., Europe, and Japan (January 1997–December 2010)

	U.S.				Europe				Japan			
	Mkt	EW	MV	MD	Mkt	EW	MV	MD	Mkt	EW	MV	MD
Excess Return over RF	2.8%	5.4%	4.1%	4.2%	4.2%	5.7%	5.4%	1.6%	–1.8%	–0.3%	4.9%	3.4%
Volatility	19.3%	20.8%	12.9%	15.3%	21.6%	21.8%	14.5%	17.8%	22.6%	22.0%	15.8%	18.9%
Sharpe Ratio	0.15	0.26	0.31	0.27	0.19	0.26	0.38	0.09	–0.08	–0.01	0.31	0.18
Excess Return over BM		2.5%	1.2%	1.3%		1.5%	1.2%	–2.6%		1.5%	6.6%	5.2%
Tracking Error		6.2%	11.9%	10.4%		5.7%	12.3%	10.8%		5.6%	13.3%	12.0%

EXHIBIT 9

Factor Annualized Returns and Volatility: U.S., Europe, and Japan (January 1997–December 2010)

		Mkt-Rf	HML	SMB	LBMHB	LRVMHRV
U.S.	Return	2.8%	–0.8%	2.0%	3.2%	–2.3%
	Volatility	19.2%	16.0%	13.3%	10.0%	13.1%
	Sharpe Ratio	0.15	–0.05	0.15	0.32	–0.18
Europe	Return	4.4%	6.6%	–0.2%	4.3%	–1.2%
	Volatility	21.5%	12.5%	11.3%	9.5%	9.4%
	Sharpe Ratio	0.20	0.53	–0.02	0.46	–0.13
Japan	Return	–1.8%	4.5%	–0.6%	0.8%	2.2%
	Volatility	22.6%	15.4%	10.9%	14.4%	11.8%
	Sharpe Ratio	–0.08	0.29	–0.06	0.05	0.18

It is also positive in Japan but is flat in the U.S. The size factor measured by SMB is only positive for U.S. stocks. The performance of the SMB factor in Europe and Japan is marginally negative for the period. The active returns of EW essentially depend only on the premium of small-cap stocks. In the U.S., where this premium is high, EW significantly outperforms the market-cap index. In Europe and Japan, where this premium has been absent, the outperformance of EW over the market-cap index is much smaller and due to less important exposure to value and its positive premium in Europe and Japan. Low-beta stocks outperform in Europe and in the U.S., but only marginally in Japan. Low residual volatility stocks underperform in the U.S. and Europe, but outperform in Japan. Finally, the performance of the market-cap index is negative in Japan. Of all these factors, only low-beta stocks have positive average returns in all three regions. Despite the difference in average returns across regions, the correlations of factor returns are similar to those for the global universe of stocks presented in Exhibit 3.

In Exhibit 10, we show the results of applying the extended Fama–French regression model to the excess returns of each strategy over the market-cap index. The results from the regression of the excess returns of each strategy over the respective market-cap index returns lead to a remarkably similar picture as the one we found

EXHIBIT 10

Factor Regression Coefficients for Risk-Based Strategies: U.S., Europe, and Japan (January 1997–December 2010)

	EW			MV Long Only			MD Long Only		
	U.S.	Europe	Japan	U.S.	Europe	Japan	U.S.	Europe	Japan
Intercept	0.000	0.000	0.000	0.000	0.000	0.000	0.000	–0.001[b]	0.000
Mkt	–0.058[a]	–0.001	–0.094[a]	–0.597[a]	–0.536[a]	–0.608[a]	–0.481[a]	–0.382[a]	–0.512[a]
HML	0.119[a]	0.078[a]	0.073[a]	0.093[a]	0.017	0.013	0.126[a]	0.051[b]	–0.104[a]
SMB	0.374[a]	0.400[a]	0.349[a]	0.139[a]	0.137[a]	0.029	0.166[a]	0.232[a]	0.052
LBMHB	0.185[a]	0.044[b]	0.103[a]	0.598[a]	0.663[a]	0.529[a]	0.591[a]	0.600[a]	0.503[a]
LRVMHRV	–0.012	–0.024	–0.02	0.138[a]	0.168[a]	0.282[a]	–0.049[c]	–0.072[c]	0.014
R^2	0.72	0.72	0.74	0.80	0.82	0.75	0.70	0.65	0.51
DW	2.15	2.07	2.06	2.12	2.17	2.21	2.07	2.09	2.12

Note: Significance levels at 0.1%, 1% and 5% are marked by a, b, and c, respectively.

for global stocks, with R^2s still quite high, in particular for EW and MV. Only MD in Japan has a lower R^2, but it is still reasonably high. The results confirm a small-cap bias in EW. The key exposure of MV and MD is to low-beta stocks. MV is also exposed to low residual volatility stocks. MV and MD are defensive strategies with an important negative load on the market excess returns over the risk-free rate.

CONCLUSIONS

We compare five risk-based strategies (EW, ERB, ERC, MV, and MD) and analyze the factors behind their risk and performance. We show that each of these five strategies, irrespective of their underlying complexity, can be viewed as simple active strategies in which the stock active weights compared to the market-cap index have consistent and well-defined exposures to just a few factors. All employ different algorithms for portfolio construction, which, in spite of not relying on any explicit stock return forecasts and focusing only on risk and diversification, generate portfolios with well-defined factor biases.

Five factors are sufficient to explain their behavior. The first factor is the market excess return over the risk-free rate. Two other factors are defined similarly to the well-known Fama–French small-cap and value factors. Two more factors are the low-beta factor and the low-residual volatility factor. Of these, the market excess return, small-cap, and low-beta factors are the three factors with the largest explanatory power.

Our findings turn what are, to some extent, complex algorithms for portfolio construction that rely on risk models and optimizers (in some cases with portfolio constraints) into quite simple approaches. They can be interpreted as the traditional benchmark and the market-cap index plus some active factor exposures, which tilt the final portfolio toward stocks with properties that are easy to grasp by investors: low beta, small cap, value, and low residual volatility.

In MV and MD, the key factor exposure is essentially to low-beta stocks, with MV also exposed to low residual volatility stocks. EW, ERB, and ERC are essentially exposed to small-cap stocks, with ERB and ERC also exposed to low-beta stocks. Value plays only a marginal role and momentum is not relevant. The exposure to low-beta stocks is responsible for the defensive characters of the MV, MD, ERC, and ERB strategies and for their lower volatilities when compared to the market-cap index. We show that risk and excess returns over the market-cap index of all the five strategies can be remarkably well explained by their respective factor exposures.

Moreover, we find that EW, ERB, and ERC share a lot in common, investing in all the stocks in the investment universe, which explains their exposure to small-cap stocks. The three require low turnover and generate low tracking risk against the market-cap index. The focus on low risk in ERB and on low risk and low correlation in ERC explains their additional tilt toward low-beta stocks and their consequent negative active exposure to the market and to lower volatility.

More remarkably, we find that MV and MD share much in common. Although this may not seem intuitive at first, further analysis uncovers a relatively similar structure with just one key difference: an additional tilt of MV toward low residual volatility stocks. This seems to play a second-order role as demonstrated by the numerical results. Other than that, both are defensive and invest essentially in low-beta stocks.

Numerical results show that all stocks seem to have delivered comparable returns in the past few decades independent of their (ex ante) risk (Fama and French [2004] for U.S. stocks between 1928 and 2003, and Blitz and van Vliet [2007] for global stocks between 1986 and 2006), which suggests that MV has historically been one of the most efficient strategies. The MV portfolio is the one that maximizes the (ex ante) Sharpe ratio when all stocks have the same expected returns. Our numerical simulations for the five risk-based strategies that we analyze show that the MV strategy does, indeed, have the highest Sharpe ratio on an ex post basis, and this was the case for global, U.S., European, and Japanese stocks.

Indeed, the negative load on the market will make the ERB, ERC, MV, and MD strategies resilient in bear markets and lag in bull markets. This behavior occurs because they are all tilted toward low-beta stocks. The additional premium from the exposure to low-beta stocks, which have delivered higher returns than predicted by the CAPM (the only premium found across all regions and at a global level), helps pay for this protection against bear markets, but there is certainly a trade-off between capturing the low-beta stock premium and lagging in bull markets.

The low-beta stock pricing anomaly has been attributed to the fact that equity investors typically seek higher returns and therefore prefer riskier stocks, which have a higher beta, and thus create a demand imbalance. The fact that many investors cannot leverage their portfolios is another reason for this demand imbalance. Even investors aware of the higher risk-adjusted returns of low-beta stocks may need to invest in higher-beta stocks in an attempt to meet target returns in the absence of leverage.

Haugen and Baker [1991] also attributed this demand imbalance to the fact that higher-volatility stocks, typically having a higher beta, are also more often in the news, creating more demand for them than for stocks that are rarely talked about. In some sense they become the "must have" stocks for fund managers because clients may not forgive them for lagging performance should these stocks be absent from the portfolio. Furthermore, high-beta stocks tend to be preferred by fund managers who have compensation tied to performance, which essentially can be attributed to lottery effects and to being a consequence of the long bull market that ended in 2001.

Investors may prefer a cleverer risk budget allocation to the factors behind risk-based strategies instead of being passively exposed to the risk budget that these allocate to each factor. In that sense, the strategies we present in this article could be replaced by a tailored portfolio strategy in which the investor adjusts his exposure to small-cap stocks, low-beta stocks, and value stocks according to his view of the premiums of these factors. Investors can also control the volatility of their equity investment and the level of defensiveness (beta) of the final portfolio. Clearly, that requires more advice from fund managers and interaction between managers and investors, which would result, in our view, in better portfolios. Scherer [2010] has already pointed this out for the MV strategy, and we have generalized the recommendation to the other four risk-based strategies that we consider.

Finally, we must add a word on transaction costs. The strategies that we have analyzed are not equal. The strategies that rely more heavily on small-cap stocks, such as EW, ERB, and ERC, will be more severely impacted by costs. For sufficiently large investors, costs and market impact may completely remove the excess returns that we find in our simulations because these strategies will have a smaller capacity. MV and MD seem to have greater capacity because they do not expose the investor to small-cap stocks. Nevertheless, MV and MD can be victims of their own success if most investors tilt their portfolios in favor of low-beta stocks, thus reducing the demand imbalance that has been around for a very long time. This would probably be the case if most investors were one day to benchmark themselves against the MV or MD portfolios. Because MV and MD are defensive strategies, the result from not having a premium on low-beta stocks would be a return in line with their portfolio beta, which is much lower than one for both strategies.

Therefore, our view is that the market-cap indices will remain the benchmarks for investing in equities, with the largest capacity and lower turnover. Alternative strategies will always involve higher turnover, bringing additional costs and market impact with it. A good compromise between the impact of turnover through transaction costs and market impact and the expected excess returns needs to be found and depends on the size of the investment, particularly when small-cap stocks are involved.

Major index providers have been proposing several indices based on risk-based strategies; see Amenc, Martellini, and Goltz [2010] for a recent review. Most of the indices tend to be constrained versions of the strategies that we discuss in this article. Having looked at the underlying approaches of several indices, we believe that they are well described by a simple factor model like the one we present. We expect constraints to have an impact on the factor exposures and the final beta of the index. Investors planning to add these indices to their portfolios are advised to carefully check the correlation of their excess returns over the market-cap index and their overlap and to compare their factor exposures. As we show, most of the risk-based strategies have more in common than was perhaps expected. Similarly, indices based on such strategies are likely to have more in common than anticipated.

ENDNOTES

[1]The marginal stock risk is defined as $\partial_{w_i}\sigma(w) = \partial\sigma(w)/\partial w_i$ with w_i the weight of stock *i*.

[2]We found that most stocks with low beta also have low volatility. We checked that this is valid not only for the universe of global stocks but also for the U.S., Europe, and Japan.

[3]Due to licensing constrains, for data prior to 2002, we use the global universe of stocks of developed countries

in the Exshare database for which the market-cap allocation minimizes the tracking risk against the total returns of the MSCI World Index in U.S. dollars. Therefore, the universe for the period prior to 2002 may not be the exact same universe that underlies the MSCI World Index. We believe that our universe is likely to contain more stocks than those in the MSCI Index in the period 1995–2001.

[4]http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/index.html

[5]http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/index.html

[6]We know that the strategies will never invest in stocks outside the available universe. That is why, in order to understand which exposures the portfolios develop, it does not make sense to consider factors built with stocks outside the investment universe. This means, however, that what we call small-cap stocks are really the stocks with the smallest market capitalization in the investment universe considered.

[7]Unlike the SMB and HML factors, the LBMHB and LRVMHRV factors have a very large negative beta before orthogonalization. When orthogonalized against each other, the large negative beta exposure is mainly concentrated in LBMHB. We have nevertheless chosen to neutralize the beta of both using the approach we describe in the text.

[8]Few large-cap stocks seem to have large residual volatility, and although there are a few small-cap stocks with low residual volatility, most tend to be large- and mid-cap stocks. We observe this not only in the case of a universe of global stocks but also for the U.S., Europe, and Japan.

[9]We do not take into account the impact of transaction costs in this analysis because our goal is to provide a better understanding of the risk-based strategies. Nevertheless, we provide a short discussion on their impact later in the article. We will consider in detail the problem of transaction costs in a future publication.

[10]We perform a simulation for the MV and MD strategies that imposes turnover constraints and find that turnover can be reduced substantially without sacrificing results. This suggests that part of the turnover arises from the larger exposure of these strategies to noise in the risk model and can therefore be minimized.

[11]http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/index.html

REFERENCES

Amenc, N., L. Martellini, and F. Goltz. "Improved Beta? A Comparison of Index Weighting Schemes." EDHEC Risk Institute, 2010.

Ang, A., R. Hodrick, Y. Xing, and X. Zhang. "The Cross-Section of Volatility and Expected Returns." *Journal of Finance*, Vol. 61, No. 1 (2006), pp. 259-299.

Baker, M., B. Bradley, and J. Wurgler. "Benchmarks as Limits to Arbitrage: Understanding the Low-Volatility Anomaly." *Financial Analysts Journal*, Vol. 67, No. 1 (2011), pp. 40-54.

Blitz, D., and P. van Vliet. "The Volatility Effect." *The Journal of Portfolio Management*, Vol. 34, No. 1 (2007), pp. 102-113.

Choueifaty, Y., and Y. Coignard. "Towards Maximum Diversification." *The Journal of Portfolio Management*, Vol. 34, No. 4 (2008), pp. 40-51.

Clarke, R., H. de Silva, and S. Thorley. "Minimum-Variance Portfolios in the U.S. Equity Market." *The Journal of Portfolio Management*, Vol. 33, No. 1 (2006), pp. 10-24.

——. "Minimum-Variance Portfolio Composition." *The Journal of Portfolio Management*, Vol. 37, No. 2 (2011), pp. 31-45.

Demey, P., S. Maillard, and T. Roncalli. "Risk-Based Indexation." Working paper, 2010. Available at http://ssrn.com/abstract=1582998.

Fama, E., and K. French. "The Cross-Section of Expected Stock Returns." *Journal of Finance*, Vol. 47, No. 2 (1992), pp. 427-465.

——. "The Capital Asset Pricing Model: Theory and Evidence." *Journal of Economic Perspectives*, Vol. 18, No. 3 (2004), pp. 25-46.

Haugen, R., and N. Baker. "The Efficient Market Inefficiency of Capitalization-Weighted Stock Portfolios." *The Journal of Portfolio Management*, Vol. 17, No. 3 (1991), pp. 35-40.

Jegadeesh, N., and S. Titman. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *Journal of Finance*, Vol. 48, No. 1 (1993), pp. 65-91.

Ledoit, O., and M. Wolf. "Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection." *Journal of Empirical Finance*, Vol. 10, No. 5 (2003), pp. 603-621.

Maillard, S., T. Roncalli, and J. Teiletche. "The Properties of Equally Weighted Risk Contribution Portfolios." *The Journal of Portfolio Management*, Vol. 36, No. 4 (2010), pp. 60-70.

Malkiel, B., and Y. Xu. "Risk and Return Revisited." *The Journal of Portfolio Management*, Vol. 23, No. 3 (1997), pp. 9-14.

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

Messikh, R., and G. Oderda. "A Proof of the Outperformance of Beta Arbitrage Strategies." Working paper, 2010. Available at http://ssrn.com/abstract=1534183.

Plerou, V., P. Gopikrishnan, B. Rosenow, L. Amaral, T. Guhr, and H. Stanley. "A Random Matrix Approach to Cross-Correlations in Financial Data." *Physical Review E*, Vol. 65, No. 6 (2002), pp. 1-18.

Scherer, B. "A Note on the Returns from Minimum Variance Investing." *Journal of Empirical Finance*, Vol. 18, No. 4 (2011), pp. 652-660.

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