JOURNAL OF INVESTMENT MANAGEMENT, Vol. 9, No. 4, (2011), pp. 52–74

www.joim.com

EFFICIENT INDEXATION: AN ALTERNATIVE TO CAP-WEIGHTED INDICES*

Noël Amenc, Felix Goltz, Lionel Martellini and Patrice Retkowsky

This paper introduces a novel method for the construction of equity indices that, unlike their cap-weighted counterparts, offer an efficient risk/return trade-off. The index construction method goes back to the roots of modern portfolio theory and focuses on the tangency portfolio, the portfolio that weights index constituents so as to obtain the highest possible Sharpe ratio. The major challenge is to generate the required input parameters in a robust manner. The expected excess return of each stock is estimated from portfolio sorts according to the stock's total downside risk. This estimate uses the economic insight that stocks with higher risk should compensate their holders with higher expected returns. To estimate the covariance matrix, we use principal component analysis to extract the common factors driving stock returns. Moreover, we introduce a procedure to control turnover in order to implement the method with low transaction costs. Our empirical results show that portfolio optimization with our robust parameter estimates generates out-of-sample Sharpe ratios significantly higher than those of the corresponding cap-weighted indices. In addition, the higher risk-return efficiency is achieved consistently and across varying economic and market conditions.

0 Introduction

The capital asset pricing model (CAPM), introduced by Sharpe (1964), has had a profound influence on the management of institutional portfolios. The CAPM starts with a series of assumptions and theories that the market portfolio of all assets is risk-return efficient in the sense that it provides the highest possible expected return above the risk-free rate per unit of volatility, i.e., the highest Sharpe ratio. Since the CAPM is taught in business schools around the world,

*The authors can be contacted at research@edhec-risk.com.
We thank Erika Richter and Dev Sahoo for excellent research assistance.

there is a widespread belief that all investors should hold the market portfolio, leveraged or de-leveraged with positions in the risk-free asset depending on their risk aversion. In practice, "holding the market" becomes virtually impossible, but its approximate implementation in terms of some market-capitalization-weighted equity indices has become the standard practice for most investors and asset managers.

Capitalization weighting in equity index construction has, however, come in for harsh criticism. Early papers by Haugen and Baker (1991) or Grinold (1992) provide empirical evidence that market-cap-weighted indices provide an inefficient risk-return trade-off. From the theoretical standpoint, the poor risk-adjusted performance of such indices should not come as a surprise, as market-cap-weighting schemes are risk-return efficient only at the cost of heroic assumptions.

- The theoretical basis for holding the market portfolio is the CAPM. An extensive body of literature has shown that the theoretical prediction of an efficient market portfolio breaks down when some of the highly unrealistic assumptions of the CAPM do not bear out. In particular, financial theory does not predict that the market portfolio is efficient if investors have different time horizons, if they derive wealth from nontraded assets such as housing, social security benefits, or human capital, if short sales are constrained or if frictions in the form of taxes exist. Unsurprisingly, when testing the CAPM on securities data, the model is commonly rejected.[1]
- In addition, even if the CAPM were the true asset pricing model, holding a market-cap-weighted equity index would be a rather poor proxy for holding the market portfolio, which in principle is a combination of all assets, traded or nontraded, financial or nonfinancial, including human capital.

In the wake of criticism of market-cap-weighted indices, alternative weighting schemes have been introduced. In pursuit of a more representative weighting scheme, researchers have proposed to weight stocks by firm characteristics such as earnings, dividends, or book value (Arnott *et al.*, 2005; Siegel *et al.*, 2007). Other research has focused on constructing minimum variance benchmarks (Chan *et al.*, 1999; Clarke *et al.*, 2006), maximum diversification benchmarks (Choueifaty and Coignard, 2008), equal-risk contribution benchmarks (Maillard *et al.*, 2008) and risk factor benchmarks (Lee, 2003; Eggins and Hill, 2008; Wagner and Stocker, 2009).

This paper focuses on the portfolio that achieves the highest risk-adjusted performance. In the end, if investors care about a portfolio's risk-adjusted performance, one should focus on designing a portfolio with the highest reward-to-risk ratio, i.e., with the highest Sharpe ratio. This portfolio is known as the tangency portfolio. Following Markowitz (1952), Tobin (1958) notes that any investor can separate his investment decisions into two steps. First, find the tangency portfolio and then use an investment in the risk-free asset to obtain an overall portfolio that corresponds to the investor's risk aversion. Our approach is to focus on the design of this tangency portfolio. We thus return to the roots of modern portfolio theory to provide an alternative to the current indexation methods of constructing equity indices. The aim of this efficient indexation approach is to provide investors with benchmarks that reflect the possible risk-reward ratio from a broadly diversified stock market portfolio, and that are thus a proxy for the normal returns of an exposure to equity risk.

To generate the tangency portfolio, we resort to standard mean–variance optimization. Although our aim to maximize risk-return efficiency is fully consistent with financial theory, successful implementation of the theory depends not only on its conceptual grounds but also on the reliability of the input to the model. In our case, the results depend greatly on the quality of the parameter estimate (the covariance matrix and the expected returns of all stocks in the index).

The CAPM, as it happens, is a poor guide to the input parameters. For the CAPM, expected returns should be proportional to the stock's beta, though it has in fact been shown that such a relationship does not hold (Fama and French, 1992). Likewise, the single-factor nature of the CAPM would mean that there is a single (market) factor driving the correlation of stocks, whereas the consensus in both academe and business is that multifactor models do a better job capturing the common drivers behind stock comovements.

Extending preliminary tests reported in Martellini (2008), we generate proxies for tangency portfolios that rely on robust input parameters for both the covariance matrix and expected returns. One challenge is the estimation of expected return parameters. Instead of relying purely on statistics, which is known to generate poor expected return estimates, we use a common sense estimate of expected returns that relies on a risk-reward trade-off. We use the insight that the return on a given stock in excess of the risk-free rate is proportional to the riskiness of the stock. Investors are often underdiversified and averse not only to systematic risk but also the specific risk of a stock. Investors shun the volatility, negative skewness, and kurtosis of a stock's returns. We use a suitably designed risk measure that integrates these aspects and estimate expected returns by sorting stocks into high risk and low risk categories. The second central ingredient in the tangency portfolio is an estimate of the covariance of stock returns. We use a robust estimation procedure that first extracts the common factors of stock returns and then uses these factors to model the comovement of individual stocks. This efficient indexation procedure allows us to construct proxies for the tangency portfolios whose risk/reward ratio is significantly better than that of cap-weighted indices.

We use constituent data for the S&P 500 index to construct tangency portfolio proxies based on the same set of stocks as the cap-weighted index. Overall, our efficient indices obtain both higher average returns and lower volatility than do cap-weighted indices. However, portfolios rebalanced every quarter are subject to high turnover. We reduce turnover by limiting rebalancing; only when significant new information arrives do we rebalance our optimal weights.

This approach leads to significantly less turnover yet maintains high Sharpe ratios. Annual turnover in excess of the cap-weighted index is less than 20%. Over the long term, our indices increase the Sharpe ratio of the S&P 500 cap-weighted index by more than 70%. Interestingly, this improved risk/reward trade-off does not come at the cost of an increase in extreme risks, and it holds when conditioning on business cycles or implied volatility. When performance over several 10-year periods is analyzed, the efficient indexation strategy had lower Sharpe ratios only during the bull markets of the 1990s, although volatility was still lower than that of the cap-weighted indices.

Cap-weighted indices weight stocks by the footprint they leave on the stock market. Characteristics-based indices weight stocks by their footprint in the economy. Investors probably care little about these aspects, unless they want portfolios representative of the stock market or the economy. Our approach weights stocks by their "risk/return footprint" on the investor's portfolio. Investors, of course, would prefer high weights in stocks that contribute positively to the portfolio's Sharpe ratio and low weights in stocks that contribute less to increasing the Sharpe ratio.

The contribution of this paper is to provide an index construction method that explicitly takes into account this investor objective.

It is sometimes argued that financial theory provides arguments in favor of the efficiency of cap-weighted equity indices, and that alternative weighting schemes are necessarily inconsistent with general equilibrium. However, it should be noted, that while the CAPM predicts a mean–variance efficient market portfolio, (i) this prediction is based on highly unrealistic assumptions (as emphasized in Markowitz (2005)) and (ii) existing stock market indices can only be very poor proxies of the theoretical market portfolio which is in fact unobservable in practice (see Roll 1977). While it is beyond the scope of this paper to analyze general equilibrium implications of our approach, the index construction method analyzed in this paper is fully consistent with financial theory as we consider a partial equilibrium by constructing the tangency portfolio for a mean variance investor given a set of robust parameter estimates.

The remainder of this paper is organized as follows. In Section 1, we describe the parameter estimates used in the portfolio optimization, namely, the covariance matrix and the expected returns. Section 2 is an overview of the implementation of the method, addressing issues such as data, timing, weight constraints, and turnover control. Section 3 analyzes the performance of the resulting portfolios both over the long run and in terms of consistency over time and across different market conditions. A final section concludes.

1 Robust estimation of return comovements and expected returns

A key to providing truly efficient equity indices is, first, to recognize this objective explicitly in the index construction process. However, improvement of the objective function is possible only if input parameters are reliable. We now turn to describing the derivation of input parameters, first for the covariance matrix and then for the expected returns.

1.1 Improved estimation of the comovements of stock returns

Several improved estimates for the covariance matrix have been proposed, including the factor-based approach (Sharpe, 1964), the constant correlation approach (Elton and Gruber, 1973), and the statistical shrinkage approach (Ledoit and Wolf, 2004). In addition, Jagannathan and Ma (2003) find that imposing (short selling) constraints on the weights in the optimization program improves the risk-adjusted out-of-sample performance in a manner similar to some of the improved covariance matrix estimators.

In these papers, the authors focus on testing the out-of-sample performance of global minimum variance (GMV) portfolios, as opposed to the MSR portfolios (also known as tangency portfolios), as there is a consensus that available estimates of expected returns are not robust enough to be used (see Section 3 for a new approach to expected return estimation).

The key problem in covariance matrix estimation is the curse of dimensionality; when a large number of stocks is considered, the number of parameters to estimate grows exponentially. Furthermore, the sample covariance matrix will be noninvertible if the number of assets N exceeds the number of available observations T; this is particularly disturbing since the minimum variance (MV) investor's optimal portfolio depends on the inverse of the covariance matrix.

Therefore, at the estimation stage, the challenge is to reduce the number of factors. In general, a multifactor model decomposes the (excess) return (in excess of the risk-free asset) on an asset into its expected rewards for exposure to the "true" risk factors. The use of a multifactor model originates in Ross's (1976) arbitrage pricing theory (APT). Formally, the returns on an asset are governed by the following linear factor model:

$$r_{it} = \alpha_{it} + \sum_{j=1}^{K} \beta_{i,jt} \cdot F_{jt} + \varepsilon_{it}$$

or in matrix form for all N assets

$$r_t = \alpha_t + \beta_t F_t + \varepsilon_t, \qquad (1.1)$$

where β_t is an $N \times K$ matrix containing the sensitivities of each asset i with respect to the corresponding jth factor movements; r_t is the vector of the N assets' (excess) returns, F_t a vector containing the K risk factors' (excess) returns, and ε_t the $N \times 1$ vector containing the zero mean uncorrelated residuals ε_{it}. The covariance matrix for the asset returns, implied by a factor model, is given by

$$\Omega = \beta \cdot \Sigma_F \cdot \beta^{\mathrm{T}} + \Sigma_\varepsilon, \qquad (1.2)$$

where Σ_F is the $K \times K$ covariance matrix of the risk factors and Σ_ε an $N \times N$ covariance matrix of the residuals corresponding to each asset.

1.1.1 Choosing the appropriate factors

Although the factor-based estimator is expected to allow a reasonable trade-off between sample risk and model risk, the problem of choosing the "right" factor model remains. We take a somewhat agnostic approach to this question, and aim to rely as little as possible on strong theoretical assumptions by using principal component analysis (PCA) to determine the underlying risk factors from the data. The PCA method is based on a spectral decomposition of the sample covariance matrix and its goal is to explain covariance structures using only a few linear combinations of the original stochastic variables that will constitute the set of (unobservable) factors.

We can use PCA in the context of a factor model, making the assumption that all stock returns depend on a number of underlying and unobservable stochastic factors $F_1, F_2, \ldots, F_K$, as well as on the variable specific errors/variations $\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_N$. Consider the N-dimensional stochastic (demeaned) vector $\boldsymbol{r}$ to be any of the stochastic variables r_t for $t = 1, 2, \ldots, T$ with sample covariance matrix S. The factor model in matrix form would be:

$$r = LF + \varepsilon, \qquad (1.3)$$

where the coefficients l_{ij} of $\boldsymbol{L}$ correspond to the loading on variable i by factor j and $\boldsymbol{F}$ is a vector with the unobservable underlying factors. Equation (1.3) corresponds to Equation (1.1) assuming zero intercept (from a pricing theory standpoint, this should be valid if we have a correct factor model). We also assume that:

$$E[\varepsilon] = 0 \quad \mathrm{Var}[\varepsilon] = E[\varepsilon\varepsilon^{\mathrm{T}}] = \Psi,$$

where Ψ is a diagonal matrix of specific variances in which the factors and the specific variances are meant to be uncorrelated. Letting the covariance matrix of the factors be Λ and taking the variance of (1.3) gives:

$$\mathrm{Var}[r] = \bar{L}\dot{\Lambda}\bar{L}^{\mathrm{T}} + \dot{\Psi} \qquad (1.4)$$

The principal components are those linear combinations that give the direction of maximum variance in the sample such that they are uncorrelated with each other (orthogonal). The ith principal component is given by:

$$f_i = l_i^{\mathrm{T}} r = \sum l_{ni} r_n \quad i = 1, 2, \ldots, N \qquad (1.5)$$

for which variances and covariances are:

$$\mathrm{Var}[f_i] = l_i^{\mathrm{T}} S l_i = \lambda_i \quad i = 1, 2, \ldots, N \qquad (1.6)$$

$$\mathrm{Cov}[f_i, f_j] = l_i^{\mathrm{T}} S l_j \quad i \neq j \qquad (1.7)$$

and $l_i^{\mathrm{T}} l_i = 1, i = 1, 2, \ldots, N$. The loadings are determined by the eigenvectors of S in Equation (1.5) and their variances in (1.6) equal the corresponding eigenvalues λ_i.[2] The application of this procedure using standardized returns in $\boldsymbol{r}$ make S (in this notation) the correlation matrix. For clarity we use P to denote the correlation matrix.

Taking the eigenvalues–eigenvector pairs (e_1, λ_1), $(e_2, \lambda_2), \ldots, (e_N, \lambda_N)$, where $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N \geq 0$ and $e_j = \left[e_{1,j}, e_{2,j}, \ldots, e_{N,j}\right]^{\mathrm{T}}$ of the matrix P, we can re-write it as:

$$\begin{aligned} \mathrm{P} &= \lambda_1 e_1 e_1^{\mathrm{T}} + \lambda_2 e_2 e_2^{\mathrm{T}} + \cdots + \lambda_N e_N e_N^{\mathrm{T}} \\ &= [e_1 \quad e_2 \quad \ldots \quad e_N] \\ &\quad \times \begin{bmatrix} \lambda_1 & 0 & 0 & 0 \\ 0 & \lambda_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & \lambda_N \end{bmatrix} \begin{bmatrix} e_1^{\mathrm{T}} \\ e_2^{\mathrm{T}} \\ \vdots \\ e_N^{\mathrm{T}} \end{bmatrix} \\ &= \bar{L}\dot{\Lambda}\bar{L}^{\mathrm{T}}. \end{aligned} \tag{1.8}$$

The decomposition in (1.8) fits exactly into Equation (1.4), taking $\Psi = 0$ and noting that $\mathrm{Cov}[f_i, f_j] = 0$. This form yields an exact representation of the covariance structure; however, a great deal of the variability can be often explained by only a few of the principal components without losing much information.

The advantage of this procedure is that it can lead to a very significant reduction of the number of parameters to estimate. This can be implemented by neglecting the effect of the smallest eigenvalues; hence, we can write Equation (1.8) as:

$$\begin{aligned} \mathrm{P} &= \lambda_1 e_1 e_1^{\mathrm{T}} + \lambda_2 e_2 e_2^{\mathrm{T}} + \cdots + \lambda_{K_1} e_{K_1} e_{K_1}^{\mathrm{T}} \\ &= \left[e_1 \quad e_2 \quad \ldots \quad e_{K_1}\right] \\ &\quad \times \begin{bmatrix} \lambda_1 & 0 & 0 & 0 \\ 0 & \lambda_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & \lambda_{K_1} \end{bmatrix} \begin{bmatrix} e_1^{\mathrm{T}} \\ e_2^{\mathrm{T}} \\ \vdots \\ e_{K_1}^{\mathrm{T}} \end{bmatrix} \\ &= L\Lambda L^{\mathrm{T}}, \end{aligned} \tag{1.9}$$

where L is an $N \times K_1$ matrix with $K_1 < N$. If we now take into account the effect of the error embedded in the approximation we get:

$$\mathrm{P} \approx L\Lambda L^{\mathrm{T}} + \Psi \tag{1.10}$$

This is equivalent to obliging the diagonal elements in the correlation matrix to be equal to one. Note the correspondence of Equations (1.10) and (1.2); both use a factor model to decompose the matrix, but Equation (1.10) corresponds to a correlation matrix given that we take $\boldsymbol{r}$ as standardized returns (zero mean and unit variance).

Bengtsson and Holst (2002) and Fujiwara *et al.* (2006) also provide justification for the use of PCA in a similar way, extracting principal components to estimate expected correlation within MV portfolio optimization. Fujiwara *et al.* (2006) find that the realized reward-to-risk ratio of portfolios based on the PCA method is higher than that of the single-index and that the optimization gives a practically reasonable asset allocation. Overall, the main strength of the PCA approach at this stage is that it enables the data to talk and to show the underlying risk factors that govern most of the variability of the assets at each point in time. The PCA approach strongly contrasts with forced reliance on the assumption that a particular factor model is the true pricing model and reduces the specification risk embedded in the factor-based approach while keeping the sample risk reduction.

Furthermore, to reduce the specification risk to the minimum, we use an objective criterion to determine the number of factors in our estimation.

1.1.2 Choosing the appropriate number of factors

Determining the appropriate number of factors to structure the correlation matrix is critical to the risk estimation when using PCA as a factor model. Several options, some with more theoretical justification than others, have been proposed to make this determination.

Financial applications such as those of Laloux *et al*. (2001), Bengtsson and Holst (2002), Amenc and Martellini (2002), and Fujiwara *et al*. (2006) provide justification for the use of a rule derived from some explicit results from random matrix theory (RMT) (Plerou *et al*., 2001; Plerou *et al*., 1999; Laloux *et al*., 1998; Guhr, 2001; Marchenko and Pastur, 1967). Within the correlation matrix structure context, Fujiwara *et al*. (2006) find that the error in the estimation of the correlation matrix via RMT is more stable and smaller than that of the sample, single-index, or constant-correlation model.

The idea is to try to separate the real correlation from the estimation error by comparing the properties of the empirical correlation matrix with known results for a completely random correlation matrix. It has been shown that the asymptotic density of eigenvalues of the correlation matrix of strictly independent asset reads[3]:

$$f(\lambda) = \frac{T}{2N\pi}\frac{\sqrt{(\lambda - \lambda_{\max})(\lambda - \lambda_{\min})}}{\lambda}, \quad (1.11)$$

where

$$\lambda_{\min} \leq \lambda \leq \lambda_{\max}$$

and the minimum and the maximum eigenvalue bounds are given by:

$$\lambda_{\min}^{\max} = 1 + \frac{N}{T} \pm 2\sqrt{\frac{N}{T}}. \quad (1.12)$$

A conservative interpretation of this result to design a systematic decision rule is to regard as statistical noise all factors associated with an eigenvalue lower than $\lambda_{\max}$.

1.2 Improved estimators of expected returns

Although we rely on statistics to extract meaningful factor models for covariance estimation, they are nearly useless in estimating expected returns, since the data are extremely noisy (Britten-Jones, 1999). Recognizing the difficulty of using sample-based expected return estimates in portfolio optimization, we follow Martellini (2008) in using an economic relation to estimate expected returns. In particular, we use an estimate of the stock's risk to proxy for a stock's expected returns. This approach is based on the principle that investors expect an additional return for taking on more risk.

Although it seems reasonable to assume that there is a risk-return trade-off, how risk should be measured must be addressed. Standard asset pricing theories such as the capital asset pricing model (Sharpe, 1964) and the arbitrage pricing theory (Ross, 1976) imply that expected returns should be positively related to systematic volatility, as measured through a factor model that summarizes individual stock return exposure with respect to one or more rewarded risk factor(s). More recently, a series of papers has focused on the explanatory power of idiosyncratic, as opposed to systematic, risk for the cross-section of expected returns. Malkiel and Xu (2006), developing an insight from Merton (1987), show that an inability to hold the market portfolio, whatever the cause, will force investors to look, to some degree, at both total risk and market risk, so firms with larger firm-specific risk must deliver higher average returns to compensate investors for holding imperfectly diversified portfolios. That stocks with high idiosyncratic risk earn higher returns has also been confirmed in a number of recent empirical studies (Tinic and West, 1986; Malkiel and Xu, 1997, 2002).

Taken together, these findings suggest that total risk should be positively related to expected return. Most commonly, total risk is the volatility of a stock's returns. Martellini (2008) has investigated the portfolio implications of these findings, and they have found that tangency portfolios constructed on the assumption that the cross-section of excess expected returns could be approximated by the cross-section of volatility posted better out-of-sample risk-adjusted performance than their market-cap-weighted counterparts.

In this paper, we extend the results in Martellini (2008) to risk measures that take into account higher-order moments. Theoretical models have shown that, in exchange for higher skewness and lower kurtosis of returns, investors are willing to accept expected returns lower (and volatility higher) than those of the mean–variance benchmark (Rubinstein, 1973; Kraus and Litzenberger, 1976). More specifically, skewness and kurtosis in individual stock returns (as opposed to the skewness and kurtosis of aggregate portfolios) have been shown to matter in several papers. High skewness is associated with lower expected returns in Barberis and Huang (2004), Brunnermeier *et al.* (2005), and Mitton and Vorkink (2007). The intuition behind this result is that investors like to hold positively skewed portfolios. The highest skewness is achieved by concentrating portfolios in a small number of stocks that themselves have positively skewed returns. Thus investors tend to be underdiversified and drive up the price of stocks with high positive skewness, which in turn reduces their future expected returns. Stocks with negative skewness are relatively unattractive and thus have low prices and high returns. The preference for kurtosis is in the sense that investors like low kurtosis and thus expected returns should be positively related to kurtosis. Boyer *et al.* (2009) and Conrad *et al.* (2008) provide empirical evidence that individual stocks' skewness and kurtosis is indeed related to future returns.

An alternative to direct consideration of the higher moments of returns is to use a risk measure that aggregates the different dimensions of risk. In this line, Bali and Cakici (2004) show that future returns on stocks are positively related to their Value-at-Risk and Estrada (2000) and Chen *et al.* (2009) show that there is a relationship between downside risk and expected returns.

Our estimate of expected returns to construct the tangency portfolio proxy uses such a downside risk measure, and, in particular, the stock's semi-deviation. The semi-deviation is a more meaningful definition of risk than volatility, since it takes into account only deviations below the mean. We compute the semi-deviation of the returns of each constituent SEM_i with respect to the average return μ_i of the ith stock as

$$SEM_i = \sqrt{E\{\min[r_{i,t} - \mu_i], 0\}^2},$$

where $E(\cdot)$ is the expectation operator computed as the arithmetic average, $\min(x, y)$ the minimum of x and y, and $r_{i,t}$ the return of stock i in week t.

To estimate expected returns, we follow the portfolio sorting approach of Fama and French (1992). That is, rather than attribute an expected return to each stock, we sort stocks by their total risk and form decile portfolios. We then attribute the median total risk of stocks in that decile portfolio to all stocks in the portfolio and use this risk measure as an estimate of expected return. The relationship between risk and returns derived from these portfolio sorts provides an estimate of expected returns.

It should be noted that, while our working assumption is a positive risk-return trade-off in the cross-section of expected stock returns, there exists conflicting empirical evidence on the nature of the risk-return relation depending on how risk is defined. In particular, Ang *et al.* (2006, AHXZ hereafter) find empirically that there is a negative link between stock-specific volatility and expected returns. However, more recent studies have examined these results and have found that they do not withstand a range of robustness tests.

For example, Bali and Cakici (2008) show that the results in AHXZ break down when changing the weighting scheme applied to low volatility and high volatility stocks in portfolio sorting exercises, when changing the breakpoints to classify stocks into low and high idiosyncratic volatility, when changing the data frequency to estimate idiosyncratic volatility, or when avoiding micro caps stocks and highly illiquid stocks. Likewise, Huang *et al.* (2010) show that the negative relation put forward by AHXZ does not hold when one corrects for the effects of short term return reversals. Also, AHXZ measure volatility over the past month to examine how this is related to returns over the next month. Their results therefore, are not applicable to more long term holding periods or calibration periods for measuring risk. Recent papers have shown that when measuring stock-specific volatility over longer horizons (3 years) and using a time series model (EGARCH) to measure volatility, there is a positive link between expected stock returns and stock-specific volatility in both the united states (Fu, 2009) and in international stock markets (Brockmann *et al.*, 2009).

We should underline that the expected return proxy used in our analysis is based on total downside risk, not stock-specific volatility and is thus not comparable in nature to the AHXZ measure of risk. It should also be noted that to decompose stock-specific and systematic risk components the above-mentioned papers need to specify an asset pricing model. In contrast, our approach of using total risk is entirely model free. Once we have computed the total risk measure, we simply work with the assumption of a positive relation between total risk and expected return without specifying a particular functional form or model for this relation. In addition, our expected return estimates are robust in the sense that they rely on a relation between return and risk that is a fundamental principle in financial theory. We also refrain from estimating individual expected returns, as we sort stocks into groups with high and low expected returns, consistent with cross-sectional asset pricing tests in the empirical finance literature.

2 Implementing efficient indexation

We now turn to the implementation of portfolio optimization with our robust input parameters, with the objective of deriving a reliable proxy for the tangency portfolio. This section describes the set of data used in our tests. In addition, practical implementation of the approach imposes further constraints, which we consider here. For example, our objective is to weight index constituents more efficiently, so we aim to match the actual constitution of the cap-weighted index as closely as possible. In addition, we introduce weight constraints and a method to control portfolio turnover.

To test our approach to constructing proxies for the tangency portfolio, we consider long-term US stock market data from CRSP. We consider the S&P 500 index and test whether we can improve its risk-return efficiency by weighting constituents differently than by their market cap.

We obtain the S&P 500 constituent lists directly from CRSP, where one can see for each day which stocks belong to the index. For the risk-free rate, we use the ML US T-Bill 3M index from Datastream, and we compute the corresponding weekly returns. All equity returns time series are weekly total returns (including reinvestment of periodic payments such as dividends), as computed by CRSP. The constitution of the S&P 500 is available from 1959.

We assess portfolios, rebalanced every quarter, of all index constituents. The rebalancing is done after the close of the first Friday of January, April, July, and October. To estimate optimal weights, we use returns for the 2 years before rebalancing. We select all constituents that are constituents of the underlying index at the rebalancing. To construct our index, we use the same constitution as that of the cap-weighted index.

We find the efficient weights as the set of weights that allow an investor to obtain the highest Sharpe ratio, given the risk and return inputs and the weight constraints. The constituent weights that solve this optimization are the efficient weights w^* that will be used in the efficient index, obtained with the following formula:

$$w^* = \arg\max_{w} \frac{w'\mu}{\sqrt{w'\Sigma w}},$$

where μ is the vector of expected returns in excess of the risk-free rate and Σ is the covariance matrix for returns of these constituents. The efficient weights lead to the highest expected returns per unit of risk, given the expected excess returns and given the covariance matrix for the index constituents in question.

As described above, the covariance matrix is estimated from a statistical factor model using principal component analysis, whereas the expected returns are estimated through a risk-return relation, in which we sort stocks by total downside risk to group them into deciles according to their expected returns. Each quarter, we use the updated input parameters to derive optimal weights, implement these optimal weights at rebalancing and then hold the stocks until the next quarter.

We impose the usual portfolio constraint that weights have to sum to one. In addition, we impose weight constraints that depend on the number of constituents (N) in the index. We impose an upper bound of λ/N and a lower bound of $1/(\lambda N)$, where λ is a flexibility parameter we set to two. Changing this parameter has no qualitative effect on the results. These constraints ensure that we include all index constituent stocks and that we do not obtain any negative weights that would lead to short sales. An appealing side effect of imposing weight constraints is that, not unlike statistical shrinkage techniques (Jagannathan and Ma, 2003), it makes possible a better trade-off between specification error and sampling error.

Although we wish to rebalance every quarter to be able to update information when necessary, it seems reasonable to rebalance the portfolio not at fixed intervals but only when weights have undergone significant shifts. This approach is consistent with insights from control techniques applied to portfolio optimization to lower transaction costs (Leland, 1999; El Bied *et al.*, 2002). To achieve lower turnover, we refrain from updating the optimal weights if the average absolute change in weights is less than 50% of the overall portfolio value. To ensure that we match the constitution of the underlying cap-weighted index, we continue to update the constitution in quarters in which we do not update the optimal weights. So exiting constituents will be deleted and new constituents will be included with the minimum weight $(1/\lambda N)$ at rebalancing. The following table shows the resulting turnover statistics and provides an analysis of indifference transaction costs. The table also shows results for another practical issue, portfolio concentration.

In addition to absolute turnover and to excess turnover relative to the cap-weighted index, Table 1 shows the impact this turnover would have on the performance of the efficient portfolios. Since transaction costs vary from one investor to another, it is unreasonable to assume fixed transaction costs. Instead, we compute the indifference level of transaction costs an investor would have to pay if these costs were to offset completely the difference in average returns compared to cap-weighting. This indifference level is 13% for the

Table 1 Turnover and concentration.

Index	Annual one-way turnover (%)	Excess turnover vs. cap-weighted (%)	Ann. return difference over cap-weighting (%)	Indifference level of transaction costs (%)	Average effective constituents (%)	Effective constituents to nominal constituents (%)
Efficient index	23.10	18.41	2.40	13.06	382	76
Cap-weighted	4.69	0.00	—	—	94	19

Note: The table shows the resulting turnover measures for portfolios that have been implemented using controlled reoptimization with a threshold value of 50%. The table also shows the indifference transaction costs, the difference between annualized return and cap-weighting divided by portfolio turnover. This measure indicates at which level (round-trip) transaction costs would neutralize the return difference with cap-weighting. The table also indicates the effective number of constituents in the efficient index and in the cap-weighted index, computed as the inverse of the sum of squared constituent weights. This measure is computed at the start of each quarter and averaged over the entire period. The results are based on weekly return data from 01/1959 to 12/2008 for S&P 500 constituents.

efficiently weighted portfolio. In practice, it is unlikely that any investor would pay costs of such magnitude.

Also of interest to investors is portfolio concentration. Indeed, it has been argued that one of the main drawbacks of capitalization weighting is excessive concentration in a few stocks with high market capitalization. The argument is that since few stocks will account for most of the weight in the index the effective number of stocks held in a cap-weighted index will be well below the actual number of constituents. We follow Strongin *et al.* (2000) in computing the effective number of stocks as the inverse of the sum of squared portfolio weights. For the S&P 500 universe, the efficient weighting method leads to portfolios with an average of 382 effective constituents, whereas the cap-weighted index has only 94 stocks effectively by this measure. Thus, with efficiently weighted portfolios concentration is considerably reduced.

3 Performance of efficient indexation

Now that we have described a method that controls turnover and shown the feasibility of the approach in terms of portfolio turnover and concentration, we turn to an analysis of the risk and return properties of the strategy. As our approach is an alternative to cap-weighted stock market indices that is based on the exact same constituents and changes only the weighting scheme, risk and return statistics for the cap-weighted index are shown for comparison. As the efficient index and the cap-weighted index have exactly the same constituents, the resulting portfolios will show commonalities in risk and return. At an annualized 5%, the tracking error of efficient indexation is lower than that of the cap-weighted index. This section looks first at long-term performance and then at the consistency of performance across market conditions.

3.1 Long-term risk and return

Both the absolute and the relative performance of the strategy must be analyzed. In addition, it is necessary to test whether the efficient weighting method's outperformance of capitalization weighting is statistically significant and to assess exposure to extreme risks and volatility. This section does these analyzes for the full historical time period, whereas the next section will focus on performance in different market conditions.

Table 2 shows summary performance statistics. For the average return, volatility and the Sharpe

Table 2 Risk and return.

Index	Ann. average return (compounded) (%)	Ann. standard deviation (%)	Sharpe ratio (compounded)	Information ratio	Tracking error (%)
Efficient index	11.63	14.65	0.41	0.52	4.65
Cap-weighted	9.23	15.20	0.24	0.00	0.00
Difference (efficient – cap-weighted)	2.40	−0.55	0.17	—	—
p-value for difference	0.14	6.04	0.04%	—	—

Note: The table shows risk and return statistics portfolios constructed with the same set of constituents as the cap-weighted index. Rebalancing is quarterly subject to an optimal control of portfolio turnover (by setting the reoptimization threshold to 50%). Portfolios are constructed by maximizing the Sharpe ratio given an expected return estimate and a covariance estimate. The expected return estimate is set to the median total risk of stocks in the same decile when sorting by total risk. The covariance matrix is estimated using an implicit factor model for stock returns. Weight constraints are set so that each stock's weight is between 1/2N and 2/N, where N is the number of index constituents. P-values for differences are computed using the paired t-test for the average, the F-test for volatility, and a Jobson–Korkie test for the Sharpe ratio. The results are based on weekly return data from 01/1959 to 12/2008.

ratio, we report differences with respect to cap-weighting and assess whether this difference is statistically significant. It is important to assess significance, as we base our conclusions on a limited amount of data, and any differences could, in principle, be the result of random effects.

Table 2 shows that the efficient weighting of index constituents leads to higher average returns, lower volatility, and a higher Sharpe ratio. All these differences are statistically significant at the 10% level, whereas the difference in Sharpe ratios is significant even at the 0.1% level. Given the data, it is highly unlikely that the unobservable true performance of efficient weighting was not different from that of capitalization weighting. Economically, the performance difference is pronounced, as the Sharpe ratio increases by about 70%.

The performance measures used above adjust portfolio returns for absolute risk, i.e., for the variability in portfolio wealth without reference to an external benchmark. Since the efficient weighting procedure is an alternative to cap-weighted indexing for investors seeking exposure to the risk premium in equity markets, the standard cap-weighted index is a useful benchmark. We measure the performance of our index relative to the cap-weighted benchmark by computing alpha and beta from a single-factor analysis. This corresponds to a CAPM framework, in which the cap-weighted index is taken as a proxy for the market portfolio. Table 3 shows the performance of the efficient indexation method once we adjust for its exposure to market risk in the sense of its beta with the cap-weighted index. To account for other systematic factors which might explain returns of the efficient index, we used the Fama–French 3-Factor model. In addition to the market beta, we thus assess exposures to the value factor (i.e., to the return difference between high and low book-to-market ratio stocks) and to the small cap factor (i.e., to the return difference between high and low capitalization stock).

The results in the table show that for the S&P 500 universe the efficient indexation method significantly outperforms the cap-weighted benchmarks, since the intercept of the regression is significant. The annualized alpha is on the same order of magnitude as the annualized return difference in Table 2, suggesting that the higher returns of the efficient indexation strategy are not caused by greater exposure to market risk. The results in

Table 3 Factor model analysis.

		Ann. alpha	Market-beta	Size-beta	Value-beta	*R*-squared
CAPM	Coefficient	2.92%	0.9177			91%
	t-statistic	3.7	68.3			
	p-value	0.02%	0.00%			
Fama–French 3-Factor	Coefficient	0.72%	0.9989	0.0781	0.2916	94%
	t-statistic	1.4	185.1	8.7	29.2	
	p-value	16.77%	0.00%	0.00%	0.00%	

Note: The first panel of the table shows coefficient estimates from the regression of weekly returns of the efficient indexation strategy on weekly returns of the cap-weighted index. The data are for the period from 01/1959 to 12/2008. The second panel shows coefficient estimates from the regression of weekly returns of the efficient indexation strategy on and weekly Fama–French factors. The data are for the period from 07/1963 to 12/2008. P-values are obtained using Newey–West robust standards that are consistent in the case of heteroscedasticity and autocorrelation.

the second panel of Table 3 show that the alpha becomes insignificant in the 3 factor model. At the same time, the *R*-squared of the regression only increases slightly, showing that the size and value betas do not suffice to fully explain the return variations of the efficient indexation strategy.

In spite of the favorable absolute and relative performance of the efficient indexation method, it is interesting to analyze if the strategy exposes investors to other forms of risk. In particular, our analysis has focused on measures that do not take into account the presence of extreme risks. We ask whether the greater risk-reward efficiency in terms of the volatility of these indices comes at the cost of a higher risk of extreme losses. We first compute aggregate measures of extreme or downside risk, notably Value-at-Risk and semi-deviation. We compute Value-at-Risk to estimate the worst loss an investor can expect to incur over a weekly horizon with 95% confidence. Our Value-at-Risk estimate follows Zangari (1996) and takes into account not only the volatility but also the skewness and kurtosis of the return distribution. Portfolio semi-deviation is computed much as is the individual stock's semi-deviation.

Table 4 shows standard downside risk measures such as VaR and semideviation. In addition, the table shows first, fifth, and tenth percentiles of 3 and 12 month trailing returns, i.e., the rolling return that is exceeded in 99%, 95%, and 99% of the sample. We can see that 12-month trailing losses are considerably lower for the efficient index than for the cap-weighted index. Three-month trailing losses, VaR and semi-deviation for the efficient index are broadly similar to those of the cap-weighted index.

From the results in Table 4, we conclude that the improvement in the volatility-adjusted return (the higher Sharpe ratio) does not come at the cost of higher downside risk. This result suggests that, though we tend to overweight stocks with high downside risk through the expected return estimate, this risk is diversified away on the portfolio level.

It should be noted that, we do not see the results as evidence that efficient indexation reduces extreme risk. In fact, reducing extreme risk is not the objective of such an approach which merely focuses on obtaining the best long-term risk/reward trade-off. Our objective here was simply to check whether improving risk/reward efficiency in the sense of the Sharpe ratio comes at the expense of increasing extreme risk, and the evidence suggests that this is not the case.

Table 4 Extreme risk.

Index	95% value-at-risk over one week (%)	Ann. semi-deviation (%)	3-month trailing return: 1st percentile (%)	3-month trailing return: 5th percentile (%)	3-month trailing return: 10th percentile (%)	12-month trailing return: 1st percentile (%)	12-month trailing return: 5th percentile (%)	12-month trailing return: 10th percentile (%)
Efficient index	3.20	10.93	−21.89	−10.72	−6.57	−25.70	−14.21	−9.19
Cap-weighted index	3.28	11.13	−20.99	−10.12	−6.45	−28.18	−16.86	−11.07

Note: The table shows different measures of downside risk for efficient indexation and cap-weighting. The 95% Value-at-Risk is computed using a Cornish–Fisher expansion. Semi-deviation is the square-root of the second lower partial moment with respect to the mean and annualized by multiplying with the square-root of 52. Three-month trailing returns are computed by compounding the past 13 weeks of returns for each weekly observation, and 12-month trailing returns by compounding the past 52 weeks of returns. The table shows percentiles for the distribution of the available sample of trailing returns. The results are based on weekly return data from 01/1959 to 12/2008.

3.2 Efficient indexation versus Naïve de-concentration

Table 1 shows that the efficient indexation strategy leads to a portfolio that is considerably less concentrated than its cap-weighted counterpart. A different way to limit concentration would simply be to weight each stock equally. This naïve form of de-concentration ignores any possibility of portfolio optimization. It seems useful to compare the performance of the efficient indexation strategy and that of this naïve alternative. In fact, if the performance of efficient indexation could be attributed to a mere de-concentration effect, we would expect the performance of the equal-weighted strategy to be even stronger than that of the efficient indexation strategy, as concentration is, by definition, lower for an equal-weighted portfolio.

The equal-weighted portfolio is an appropriate benchmark for comparison, as several papers show that many alternative weighting mechanisms do not outperform simple equal-weighted portfolios. For example, DeMiguel *et al.* (2009) find that, across a wide range of datasets, the equal-weighting strategy is not consistently outperformed by mean–variance optimized portfolios including global minimum variance portfolios. Similarly, Amenc *et al.* (2009) report that equity indices that weight stocks by firm characteristics do not outperform equally-weighted indices. Therefore, using the equal-weighted portfolio as reference provides a parsimonious comparison of the performance of efficient indexation. Amenc *et al.* (2010) provide a performance comparison of different alternative weighting schemes, including efficient indexation.

Table 5 shows that efficient indexation based on robust portfolio construction seems preferable to a simple equal-weighting scheme. This suggests that portfolio optimization with robust parameter estimates, as introduced in Section 1, adds more useful information than an equal-weighted benchmark. The table below shows, in particular, that efficient indexation leads to higher expected returns and lower volatility than its equal-weighted counterpart. The tracking error and turnover of the efficient indexation strategy are also slightly lower than those of the equal-weighted strategy.

The bottom line from Table 5 is that efficient indexation leads to Sharpe and information ratios considerably higher than does equal weighting.

Table 5 Summary statistics: efficient indexation versus equal-weighting.

	Ann. average return (%)	Ann. standard deviation (%)	Sharpe ratio	Info ratio	Tracking error (%)	Annual one-way turnover (%)	Effective constituents
Equal weighting	11.1	15.8	0.35	0.39	4.8	24.2	500
Efficient indexation	11.6	14.6	0.41	0.52	4.7	23.1	382

Note: The table shows risk and return statistics portfolios constructed with the same set of constituents as the cap-weighted index. The efficient indexation method from above is compared to the equal-weighted portfolio with quarterly rebalancing that is based on the same set of constituents. The results are based on weekly return data from 01/1959 to 12/2008.

That efficient indexation makes possible performance superior to equal weighting, and with a lower effective number of constituents, also suggests that the efficient indexation method is suitable for constituent universes that include stocks that have low liquidity. As the effective number of stocks of the efficient indexation strategy is relatively low, it is possible to avoid holding the least liquid stocks. In practice, then, transaction and liquidity costs may be lower for efficient indexation.

3.3 A closer look at the performance of efficient indexation

The evidence provided above suggests that the efficient indexation method greatly improves the risk/return efficiency of cap-weighted indices. In fact, Sharpe ratios are considerably higher than those of cap-weighted indices, even though the underlying constituents are identical. The analysis above is based on long-term historical data. For the investor, it is important that the improvement in risk-reward efficiency be consistent. To determine whether it is, we provide an overview of how efficient indexation fares in different time periods, stock market regimes, and economic conditions.

The upper graph in Figure 1 shows the growth over time of investments in the efficient index and in the cap-weighted index. The plots for the two constituent universes show that the return difference leads to spectacular differences in wealth over long time periods, the simple result of compounding.

For an idea about the consistency of the increase in returns through efficient indexation, we also plot the ratio of the portfolio wealth obtained with efficient indexation to the wealth obtained by cap-weighting the same stocks. The lower plot shows the ratio of the wealth of an investor in the efficient index to the wealth of an investor in the value-weighted index at each point of time, assuming that both investors start investing at the same date and with the same amount. Thus the plot shows how many dollars an investor in the efficient index has for every dollar he would have had when investing in the value-weighted index.

The graph shows that, over time, the efficient index's cumulative outperformance of the cap-weighted index is considerable. Efficient indexation does, however, underperform value-weighting in the years from January 1996 to December 1999. Wealth ratios for both indices fall over this period, the time of the bull markets that led to the "tech bubble." Except for this period, the wealth ratio either increases or is stable, suggesting that the method provides a consistent return enhancement other than in the period of the extremely bullish markets of the late 1990s.

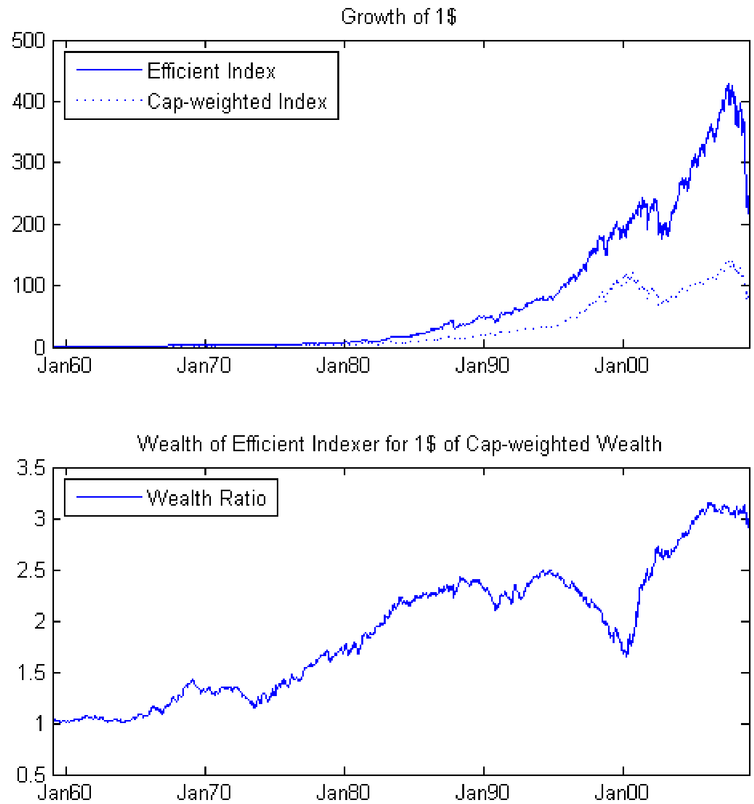

Figure 1 Growth of portfolio wealth.

The upper graph shows cumulative returns normalized to a starting value of one for efficient indexation and for cap-weighting. The lower graph indicates the ratio of the solid line to the dotted line in the upper plot. The results are for the efficient indexation portfolios with controlled reoptimization. The data have a weekly frequency and range from 01/1959 to 12/2008.

The long-term evolution of wealth highlights average returns rather than risk and risk-adjusted performance. We analyze performance statistics over sub-samples for a more systematic picture of the variations in performance by time period. Table 6 shows the annualized return, volatility, and Sharpe ratio for periods of a decade. We divided the sample into nonoverlapping periods of 10 years, going backwards from December 2008. We thus obtain five sub-periods of 10 years.

Table 6 shows that the Sharpe ratio of efficient indexation is higher in every 10-year period but that from 1989 to 1998. This confirms the underperformance in bull markets observed in the graphs on the growth of wealth. Interestingly, the underperformance in bull markets suggests that the performance of efficient indexation is, in general, more stable than that of cap-weighting. In addition, though the Sharpe ratio of efficient indexation is lower than that of the cap-weighted index in the bull markets of the 1990s, efficient indexation is less volatile over this period.

It is useful to look directly at the dependence of performance on cap-weighted market returns. We sort all weekly observations by the returns of the cap-weighted index and analyze the performance in five groups of the data that correspond to different ranges of return for the cap-weighted index. The first group contains the weeks in the sample during which the cap-weighted index has returns that are below a negative 4%. The fifth group contains the weeks in the sample that correspond to

Table 6 Risk and return in different decades.

Decade	Ann. average return: Cap-weighting (%)	Ann. average return: Efficient indexation (%)	Ann. volatility: Cap-weighting (%)	Ann. volatility: Efficient indexation (%)	Sharpe ratio: Cap-weighting	Sharpe ratio: Efficient indexation
1999–2008	−1.22	3.47	18.98	18.04	−0.23	0.01
1989–1998	19.16	16.43	12.84	12.45	1.07	0.89
1979–1988	16.32	20.82	16.02	15.82	0.42	0.71
1959–1978	2.96	4.24	16.02	15.47	−0.20	−0.13
1959–1968	10.33	14.29	10.65	10.05	0.62	1.05

Note: The table shows risk and return statistics when dividing the sample into periods of 10 years. The results are based on weekly return data from 01/1959 to 12/2008.

more than 4% returns for the cap-weighted index. Computing the average weekly returns for each group shows how the strategy depends on the returns of the cap-weighted index.

Table 7 shows that efficient indexation has higher average returns than cap-weighting in all the ranges except the top two sub-samples, which correspond to roughly 15% of observations with the most bullish market conditions. The results in Table 7 confirm, unsurprisingly, that efficient indexation has lower returns than cap-weighting bull markets. An intuitive explanation is that it is extremely difficult to outperform the trend-following strategy when markets continue to follow the trend and the stocks with price increases continue to go up. However, the dispersion of efficient-weighted portfolio returns across quintiles is also lower, again suggesting more stability.

Conditioning on the cap-weighted return does not provide a complete characterization of varying market conditions. For a look at economic conditions in a broader sense, we repeat the analysis of Table 7, in which we divided the sample into sub-samples, and computed performance statistics. This time, we sort the sample into sub-samples according to the prevailing economic conditions. To characterize economic conditions, we use two variables. The first is a recession indicator for the US economy, which we obtain from the NBER. The second is implied volatility, computed by the CBOE based on option prices for index options written on the S&P index.

Table 7 Dependence of returns on cap-weighted returns

	Average weekly return (%)					
Range of cap-weighted return	below −4 (%)	−4 to −2(%)	−2 to 0(%)	0 to 2(%)	2 to 4(%)	above 4(%)
Cap-weighted	−5.85	−2.73	−0.89	0.93	2.74	5.52
Efficient indexation	−5.48	−2.48	−0.74	0.94	2.54	4.78
Percentage of observations	2.64	9.16	31.16	42.51	11.65	2.87

Note: The table shows average returns computed for six sub-samples. The sub-samples are obtained by sorting the weekly observations based on the weekly return of the cap-weighted index. The samples ranges were chosen to be one standard deviation of the cap-weighted weekly data ($\sim$2%). The results are based on weekly return data from 01/1959 to 12/2008.

Table 8 Risk and return in recessions and expansions.

Business cycle	Ann. average return		Ann. volatility		Sharpe ratio	
	Cap-weighting (%)	Efficient indexation (%)	Cap-weighting (%)	Efficient indexation (%)	Cap-weighting	Efficient indexation
Recessions	−1.64	2.26	22.85	22.29	−0.37	−0.20
Expansions	11.19	13.30	13.47	12.92	0.43	0.61

Note: The table shows risk and return statistics computed for two sub-samples. The sub-samples are obtained by sorting the weekly observations based on a recession indicator for that week. The recession indicator is obtained from NBER dates for peaks and troughs of the business cycle. The results are based on weekly return data from 01/1959 to 12/2008.

Table 8 shows results separately for recessionary and expansionary periods. The results show that both capitalization weighting and efficient indexation fare much better in expansions than in recessions. In recessions, average returns are lower and volatility of returns is higher. In both stages of the business cycle, efficient indexation provides higher average returns, lower volatility, and thus higher Sharpe ratios.

Another useful conditioning variable is implied volatility. Although the recession variable used above tells us something about the realization of economic variables, option-implied volatility directly captures investor uncertainty. The advantage of using option-implied volatility rather than realized volatility is that we can measure implied volatility precisely at a weekly frequency. In addition, implied volatility, which Whaley (2000) has described as a "fear gauge", directly reflects investor's instantaneous beliefs and preferences rather than past realizations. Table 9 repeats the analysis from the previous table. The difference is that the sub-samples are now formed according to volatility regimes. Data on implied volatility indices are available only from 1986 for the S&P index. We divide the data available since 1986 into one half that corresponds to high volatility weeks and another half to low volatility weeks.

As it does in both recessions and periods of growth, efficient indexation improves risk-return efficiency in both times of great uncertainty and times of low uncertainty. Its advantage over capitalization weighting in terms of reduced volatility is most pronounced in times of great uncertainty, suggesting that efficiently weighted portfolios provide risk reduction benefits precisely when they are most needed.

Table 9 Risk and return in times of high uncertainty and low uncertainty.

Implied volatility regime	Ann. average return		Ann. volatility		Sharpe ratio	
	Cap-weighting (%)	Efficient indexation (%)	Cap-weighting (%)	Efficient indexation (%)	Cap-weighting	Efficient indexation
High volatility	8.90	10.99	15.40	14.61	0.18	0.34
Low volatility	6.22	10.03	11.88	11.83	0.03	0.36

Note: The table shows risk and return statistics computed for two sub-samples of equal size. The sub-samples are obtained by sorting the weekly observations based on the value of the corresponding implied volatility index for that week. The median level of volatility is used to separate the two samples. The data for implied volatility indices start on 03/01/1986 (VXO index) and end on 26/12/2008.

In general, when the performance of our indexing method conditional on time, market conditions or economic conditions is analyzed, the improvements in risk-reward efficiency are confirmed. In fact, the performance of efficient indexation is extremely robust, regardless of the time period, point on the business cycle, or degree of uncertainty. When returns in rising and falling markets are analyzed, we find that efficient indexation lags capitalization weighting in pronounced bull markets, as in the late 1990s. From an investor's perspective, however, underperforming capitalization weighting when it returns 20% or more per year may be a risk worth taking in exchange for greater average efficiency.

4 International evidence

Although the results obtained for post-war US data suggest that the improvement in efficiency is highly significant both statistically and economically, it may be that these results are specific to US data. So it is important to gather evidence on how efficient indexation fares internationally. Since it is more challenging to obtain accurate data over long time periods for international markets, we analyze indices only for countries or regions with the largest stock market capitalizations and we concentrate on a recent time period for which data are available.

We apply the efficient indexation method with the same parameters as above to the constituents of the FTSE All World index from the following countries or regions: United States, Eurobloc, United Kingdom, developed Asia-Pacific ex Japan (including stocks from Australia, Hong Kong, New Zealand, and Singapore), and Japan. For these indices, we obtain daily constituent lists and constituent-level return data for a period of approximately 7 years (from 23/12/2002 to 18/09/2009). Table 10 shows the risk and return statistics obtained through efficient indexation based on these constituents and compares them to the corresponding statistics of the FTSE All World indices that weight constituent stocks by (free-float-adjusted) market capitalization.

The results in Table 10 show that risk/return efficiency in terms of the Sharpe ratio is improved considerably for all five indices. In addition, the improvement is actually quite similar across the five indices, with Sharpe ratios approximately 0.15 higher than those of the cap-weighted index. When the results for the four international indices and for the US index are compared, it is clear that the method works slightly better in the other

Table 10 Risk and return in different markets around the world.

	Ann. average return			Ann. std. dev.			Sharpe ratio		
	Efficient index (%)	Value weighted (%)	Diff. (%)	Efficient index (%)	Value weighted (%)	Diff. (%)	Efficient index	Value weighted	Diff.
USA	5.60	2.77	2.83	20.42	19.03	1.39	0.15	0.01	0.14
Eurobloc	7.48	4.19	3.30	18.61	21.40	−2.79	0.24	0.05	0.18
UK	9.66	5.44	4.23	19.65	19.43	0.22	0.27	0.06	0.21
Asia	17.19	15.80	1.40	21.36	23.83	−2.47	0.69	0.56	0.13
Japan	5.85	3.01	2.84	19.04	21.30	−2.26	0.30	0.13	0.16

Note: The table shows risk and return statistics computed for efficient indexation and capitalization weighting applied to stock market index constituents in five regions. The statistics are based on daily returns data from 23/12/2002 to 18/09/2009.

datasets, except perhaps in the Asia Pacific index. It is interesting to note that the Asia Pacific index had extremely high returns of more than 15% over the period, compared to returns of less than 6% for all other cap-weighted indices. Thus, the lesser improvement of the Sharpe ratio through efficient indexation in this dataset is actually coherent with the observation in the sub-sample analysis for the long-term US data, where it was found that, in strong bull markets, efficient indexation does not improve on capitalization weighting as much as it does in other market conditions.

In general, analysis of international data suggests that our results are not specific to US data, as the method yields similar results in stock markets around the world.

5 Conclusion

Evidence abounds of the inefficiency of cap-weighted indices. Currently available alternatives may well owe their success to that inefficiency, but, surprisingly, they do not explicitly address this problem. Characteristics-based indices, for example, attempt to be more representative of the economy by weighting stocks by each company's economic footprint. Their goal is not to weight stocks so as to improve risk-return efficiency. The approach described here, on the other hand, takes the investor's perspective into account and makes risk-return efficiency an explicit goal. Input parameters throw up a major conceptual obstacle to constructing efficient portfolios, as estimation error may weaken optimization results. From a practical standpoint, optimization methods may lead to high turnover and thus to transaction costs that wipe out favorable performance. In this paper, we draw on the academic literature to provide solutions to both the parameter estimation problem and the turnover problem. Our main contribution is to provide a novel approach, focusing on efficiency, to equity indexation; efficiency, after all, was arguably the motivation for the creation of index funds drawing on insights from the CAPM in the first place.

Our implementation of the fundamental insight of modern portfolio theory, that investors should hold the tangency portfolio, is based on robust estimates of risk and return parameters. To obtain robust parameter estimates for the comovements of stock returns, we use a multifactor model based on principal component analysis. For expected returns, we use the insight that there is a risk-return trade-off and generate estimates from a suitably designed risk measure that involves not only average risk but also higher moment risk, following the evidence of the link between downside risk and expected returns provided by Estrada (2000) and Chen *et al.* (2009), as well as the evidence of the importance of total risk for portfolio construction in Martellini (2008). Out of practical concerns, we also introduce a procedure, inspired by optimal control theory (Leland, 1999), to control turnover and transaction costs.

The empirical tests described in this paper show that this procedure allows us to generate efficient indices with out-of-sample Sharpe ratios considerably higher than those of their capitalization-weighted counterparts. In addition, performance is consistent across different business cycles, volatility regimes, and time periods. Lower risk-return efficiency occurs only in the extreme bull markets of the late 1990s. Even in this period, efficient indexation posted lower volatility than capitalization weighting, and expected returns were lower when the cap-weighted indices were returning in excess of 20% a year. It should also be kept in mind that, unlike that of other index construction methods that do not weight constituents by market capitalization, the performance of the method can be put down entirely to a different method of weighting constituents.

On the whole, when the evidence from post-war US data is taken into account, the differences in the efficiency of value-weighted indexation and efficient indexation (and efficient indexation is more efficient) are statistically significant. The increase in risk-return efficiency is similar when the method is applied to international stock market indices. In general, efficient indexation leads to an economically significant increase in efficiency for investors seeking exposure to the equity risk premium.

Notes

1 See Goltz and Le Sourd (2009) for a literature review.
2 For a proof see Johnson and Wichern (1992).
3 See Plerou *et al*. (2001).

References

Amenc, N. and Martellini, L. (2002). "Portfolio Optimization and Hedge Fund Style Allocation Decisions." *Journal of Alternative Investments* **5**(2), 7–20.

Amenc, N., Goltz, F., and Le Sourd, V. (2009). "The Performance of Characteristics-based Indices-super-1." *European Financial Management* **15**(2), 241–278.

Amenc, N., Goltz, F., and Martellini, L. (2010). "Improved Beta? A Comparison of Index-Weighting Schemes." Working Paper.

Ang, A., Hodrick, R., Xing, Y., and Zhang, X. (2006). "The Cross-Section of Volatility and Expected Returns." *Journal of Finance* **61**, 259–299.

Arnott, R., Hsu, J., and Moore, J. (2005). "Fundamental Indexation." *Financial Analysts Journal* **60**(2), 83–99.

Bali, T. G. and Nusret, C. (2004). "Value at Risk and Expected Stock Returns." *Financial Analysts Journal* **60**(2), 57–73.

Barberis, N. and Huang, M. (2004). "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices." Working Paper.

Bengtsson, C. and Holst, J. (2002). "On Portfolio Selection: Improved Covariance Matrix Estimation for Swedish Asset Returns." Working Paper, Lund University and Lund Institute of Technology.

Boyer, B., Mitton, T., and Vorkink, K. (2009). "Expected Idiosyncratic Skewness." *Review of Financial Studies* (forthcoming).

Britten-Jones, M. (1999). "The Sampling Error in Estimates of Mean-Variance Efficient Portfolio Weights." *Journal of Finance* **54**, 655–671.

Brockman, P., Schutte, M., and Yu, W. "Is Idiosyncratic Risk Priced? The International Evidence." Working Paper.

Brunnermeier, M. K., Gollier, C., and Parker, J. (2007). "Optimal Beliefs, Asset Prices, and the Preference for Skewed Return." Working Paper.

Chan, L. K. C., Karceski, J., and Lakonishok, J. (1999). "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model." *Review of Financial Studies* **12**(5), 937–974.

Clarke, R., de Silva, H., and Thorley, S. (2006). "Minimum-Variance Portfolios in the U.S. Equity Market." *Journal of Portfolio Management* (fall), 1–14.

Chen, D. H., Chen, C. D., and Chen, J. (2009). "Downside Risk Measures and Equity Returns in the NYSE." *Applied Economics* **41**, 1055–1070.

Choueifaty, Y. and Coignard, Y. (2008). "Toward Maximum Diversification." *Journal of Portfolio Management* **35**(1): 40–51.

Conrad, J., Dittmar, R. F., and Ghysels, E. (2008). "Ex Ante Skewness and Expected Stock Returns." Working Paper.

DeMiguel, V., Garlappi, L., and Uppal, R. (2009). "Optimal Versus Naive Diversification: How Inefficient is the 1-N Portfolio Strategy?" *Review of Financial Studies* **22**(5), 1915–1953.

Eggins, J. and Hill, R. J. (2008). "Momentum and Contrarian Stock-Market Indices." Working Paper.

Elton, E. and Gruber, M. (1973). "Estimating the Dependence Structure of Share Prices: Implications for Portfolio Selection." *Journal of Finance* **28**:1203–1232.

Estrada, J. (2000). "The Cost of Equity in Emerging Markets: A Downside Risk Approach." *Emerging Markets Quarterly* 19–30.

Fama, E. and French, K. (1992). "The Cross-Section of Expected Stock Returns." *Journal of Finance* **47**(2), 427–465.

Fu, F. (2009). "Idiosyncratic Risk and the Cross-Section of Expected Stock Returns." *Journal of Financial Economics* **91**, 24–37.

Fujiwara, Y., Souma, W., Murasato, H., and Yoon, H. (2006). "Application of PCA and Random Matrix Theory to Passive Fund Management, *Practical Fruits of Econophysics*, Tokyo: Springer, pp. 226–230.

Goltz, F. and Le Sourd, V. (2009). "Does Finance Theory Make the Case for Cap-Weighted Indexing?" Working

Paper, EDHEC Risk and Asset Management Research Centre.

Grinold, R. (1992). "Are Benchmark Portfolios Efficient?" *Journal of Portfolio Management* (fall).

Guhr, T. (2001). *Econophysics. Lecture Notes*. Department of Mathematical Physics, Lund Institute of Technology.

Haugen, R. and Baker, N. (1991). "The Efficient Market Inefficiency of Capitalization-Weighted Stock Portfolios." *Journal of Portfolio Management* (spring).

Huang, W., Liu, Q., Rhee, S. G., and Zhang, L. (2010). "Return Reversals, Idiosyncratic Risk, and Expected Returns." *Review of Financial Studies* **23**, 147–168.

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

Johnson, R. and Wichern, D. (1992). *Applied Multivariate Statistical Analysis*. Englewood Cliffs, New Jersey: Prentice Hall.

Kraus A. and Litzenberger, R. H. (1976). "Skewness Preference and the Valuation of Risk Assets." *Journal of Finance*, 1085–1100.

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

Laloux, L., Cizeau, P., Bouchaud, J.-P., and Potters, M. (1998). Noise Dressing of Financial Correlation Matrices. arXiv:cond-mat/9810255 v1.

Ledoit, O. and Wolf, M. (2004). "Honey, I Shrunk the Sample Covariance Matrix." *Journal of Portfolio Management* **30**(4), 110–119.

Lee, Y. (2003). "Indexation of Momentum Effects." *Journal of Investment Management* **1**(4).

Leland, H. (1999). "Optimal Asset Rebalancing in the Presence of Transaction Costs." Working Paper, UC Berkeley.

Maillard, S., Roncalli, T., and Teiletche, J. (2008). "On the Properties of Equally-Weighted Risk Contributions Portfolios." Working Paper.

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

Malkiel, B. and Xu, Y. (2002). "Idiosyncratic Risk and Security Returns." Working Paper, University of Texas at Dallas.

Malkiel, B. and Xu, Y. (2006). "Idiosyncratic Risk and Security Returns." Working Paper, University of Texas at Dallas.

Markowitz, H. (1952). "Portfolio Selection." *Journal of Finance* **7**(1), 77–91.

Markowitz, H. (2005). "Market Efficiency: A Theoretical Distinction and So What?" *Financial Analysts Journal* **61**(5), 17–30.

Marchenko, V. and Pastur, L. (1967). "The Distribution of Eigenvalues in Certain Sets of Random Matrices." *Mat. Sb*. **72**, 507–536.

Martellini, L. (2008). "Towards the Design of Better Equity Benchmarks: Rehabilitating the Tangency Portfolio from Modern Portfolio Theory." *Journal of Portfolio Management* **34**(4), 34–41.

Merton, R. (1987). "A Simple Model of Capital Market Equilibrium with Incomplete Information." *Journal of Finance* **42**(3), 483–510.

Mitton, T. and Vorkink, K. (2007). "Equilibrium Underdiversification and the Preference for Skewness." *Review of Financial Studies* **20**(4), 1255–1288.

Nielsen, F. and Aylursubramanian, R. (2008). "Far from the Madding Crowd: Volatility Efficient Indices." Working Paper.

Plerou, V., Gopikrishnan, P. Rosenow, B., Amaral, L., Guhr, T., and Stanley, H. E. (1999). "Universal and Nonuniversal Properties of Cross Correlations in Financial Time Series." *Physical Review Letters* 83, American Physical Soc.

Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L., Guhr, T., and Stanley, H. E. (2001). "A Random Matrix Approach to Cross-Correlations in Financial Data." arXiv:condmat/0108023.

Roll, R. (1977). "A Critique of the Asset Pricing Theory's Tests Part I: On Past and Potential Testability of the Theory." *Journal of Financial Economics* **4**(2), 129–176.

Ross, S. (1976). "The Arbitrage Theory of Capital Asset Pricing." *Journal of Economic Theory* **13**, 341–360.

Rubinstein, M. E. (1973). "The Fundamental Theorem of Parameter-Preference Security Valuation." *Journal of Financial and Quantitative Analysis* 61–69.

Sharpe, W. (1964). "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance* **19**(3), 425–42.

Siegel, J. J., Schwartz, J. D., and Siracusano, L. (2007). "The Unique Risk and Return Characteristics of Dividend-Weighted Stock Indexes", Working Paper.

Strongin, S., Petsch, M., and Sharenow, G. (2000). "Beating Benchmarks." *Journal of Portfolio Management* **26**(4), 11–27.

Tinic, S. and West, R. (1986). "Risk, Return and Equilibrium: A Revisit." *Journal of Political Economy* **94**(1), 126–147.

Tobin, J. (1958). "Liquidity Preference as Behavior Towards Risk." *Review of Economic Studies* **67**, 65–86.

Wagner, N. and Stocker, E. (2009). "A New Family of Equity Style Indices." Working Paper.

Whaley, R. (2000). "The Investor Fear Gauge." *Journal of Portfolio Management* **26**, 12–17.

Zangari, P. (1996). "A VaR Methodology for Portfolios that Include Options." *RiskMetrics Monitor* (January), 4–12.

Keywords: Equity benchmarks; non-cap weighted equity indices; efficient equity indices