CARISMA

A multi-disciplinary research centre focussed on understanding, modelling, quantification, management and control of RISK

- School of Information Systems, Computing and Mathematics
- Brunel Business School
- School of Engineering
- School of Social Sciences

TECHNICAL REPORT

CTR/90/11

Enhanced Index Based on Second-Order Stochastic Dominance

Diana Roman, Gautam Mitra, Victor Zviarovich

CARISMA

Centre for the Analysis of Risk and Optimisation Modelling Applications

Brunel UNIVERSITY WEST LONDON

Enhanced Indexation Based on Second-Order Stochastic Dominance

Diana Roman[*] Gautam Mitra[†] Victor Zviarovich[‡]

---

[*] CARISMA: The Centre for the Analysis of Risk and Optimisation Modelling Applications, School of Information Systems, Computing and Mathematics, Brunel University, UK. E-mail: diana.roman@brunel.ac.uk

[†] CARISMA: The Centre for the Analysis of Risk and Optimisation Modelling Applications, School of Information Systems, Computing and Mathematics, Brunel University, UK and Optirisk Systems, Uxbridge, West London, UK.

[‡] CARISMA: The Centre for the Analysis of Risk and Optimisation Modelling Applications, School of Information Systems, Computing and Mathematics, Brunel University, UK and Optirisk Systems, Uxbridge, West London, UK.

Contents

Abstract

Second order Stochastic Dominance (SSD) has a well recognised importance in portfolio selection, since it provides a natural interpretation of the theory of risk-averse investor behaviour. Recently, SSD-based models of portfolio choice have been proposed; these assume that a reference distribution is available and a portfolio is constructed, whose return distribution dominates the reference distribution with respect to SSD. We present an empirical study which analyses the effectiveness of such strategies in the context of enhanced indexation. Several datasets, drawn from FTSE 100, SP 500 and Nikkei 225 are investigated through portfolio rebalancing and backtesting. Three main conclusions are drawn. First, the portfolios chosen by the SSD based models consistently outperformed the indices and the traditional index trackers. Secondly, the SSD based models do not require imposition of cardinality constraints since naturally a small number of stocks are selected. Thus, they do not present the computational difficulty normally associated with index tracking models. Finally, the SSD based models are robust with respect to small changes in the scenario set and little or no rebalancing is necessary.

In this paper we present a unified framework which incorporates (a) SSD, (b) downside risk (Conditional Value-at-Risk) minimisation and (c) enhanced indexation.

1 Introduction

Second order Stochastic Dominance (SSD) has a well recognised importance in portfolio selection, due to its connection to the theory of risk-averse investor behaviour and tail risk minimisation. Until recently, stochastic dominance models were considered intractable or at least very demanding from a computational point of view. Computationally tractable and scalable portfolio optimization models which apply the concept of SSD were proposed recently (Dentcheva and Ruszczyński, 2006; Roman, Darby-Dowman, and Mitra, 2006; Fabian, Mitra and Roman, 2009). These portfolio optimisation models assume that a benchmark, that is, a desirable "reference" distribution is available and a portfolio is constructed, whose return distribution dominates the reference distribution with respect to SSD.
Index tracking models also assume that a reference distribution (that of a financial index) is available. A portfolio is then constructed, with the aim of replicating, or tracking, the financial index. Traditionally, this is done by minimising the tracking error: the standard deviation of the differences between the portfolio and index returns. Other methods have been proposed (for a review of these methods, see for example Beasley, Meade and Chang, 2003; Canakgoz and Beasley, 2008). The passive portfolio strategy of index tracking is based on the well established "Efficient Market Hypothesis" (Fama, 1970) which implies that financial indices achieve the best returns over time.

Enhanced indexation models are related to index tracking, in the sense that they also consider the return distribution of an index as a reference. They however aim to outperform the index by generating "excess" return (diBartolomeo, 2000; Scowcroft and Sefton, 2003). Enhanced indexation is a very new area of research and there is no generally accepted portfolio construction method in this field (Canakgoz and Beasley, 2008). Although the idea of enhanced indexation was formulated as early as 2000, the (few) enhanced indexation methods were proposed later in the research community (a review in Canakgoz and Beasley, 2008). Moreover, these methods are mostly concentrated on overcoming the computational difficulty raised by restricting the cardinality of the portfolios - not on answering the question if they do attain their stated purpose, i.e. obtain return in excess of the index.

From a theoretical perspective, enhanced indexation calls for further justification. The Efficient Market Hypothesis (EMH) is based on the key assumption that security prices fully reflect all available information (see Elton and Gruber 1995, also Lo 2004 for an insightful review of this topic). However, this hypothesis has been continuously challenged; the mere fact that academicians and practitioners commonly use "active" (i.e. non-index tracking) strategies is an indication for this. An attempt to reconcile the advocates and opponents of the EMH is the "adaptive market hypothesis" (Lo, 2004). Here, the idea is that the market "adapts" to the information received and is generally efficient but there are periods of time when it is not (and thus, these periods can be used by investors to make profit in excess of the market index.) This would justify, from a theoretical point of view, the quest for techniques that seek to obtain excess return as compared to financial indices. Enhanced indexation aims to discover and exploit market inefficiencies.

There are few empirical studies comparing performance of enhanced index funds with that of their proxy indices; a review is given in Krause (2009), see also Ahmed and Nanda (2005). They mostly come to the conclusion that, although overall the universe of enhanced funds does not seem to outperform the market, there are situations when outperformance does occur and persists for some periods of time. This seems to be in line with the adaptive market hypothesis. Another conclusion is that there may be specific types of funds that do add value; however, to the best of our knowledge, research has not been done in this direction. Indeed, although there are some empirical studies comparing overall performance of enhanced index funds with their proxies, no studies have been done regarding the effectiveness of a specific enhanced indexation *method*.

A common problem with the index tracking and enhanced indexation models is raised by the computational difficulty. This difficulty is posed by cardinality constraints that limit the number of stocks in the chosen portfolio. It is known that most index tracking models naturally select a very large number of stocks in the composition of the portfolio. Cardinality constraints overcome this problem, but they require introduction of binary variables and thus the resulting model becomes much more difficult to solve. Most of the literature in the field is concerned with overcoming this computational difficulty.

The purpose of this paper is to analyse the *effectiveness* of the SSD based models proposed by Roman, Darby-Dowman, and Mitra (2006) and Fabian et al. (2010) as enhanced indexation strategies, in a rebalancing approach. The good in-sample properties of the return distribution of the chosen portfolios have been underlined in previous papers (Roman, Darby-Dowman and Mitra, 2006, using historical data; Fabian et al., 2010 using scenarios generated via

Geometric Brownian Motion). However, it is the actual / historical performance of the chosen portfolios (measured over time and compared with the historical performance of the index) that provides empirical validation of whether the models achieved their stated purpose (to generate excess return).

We also investigate aspects related to the practical application of portfolio models in which the asset universe is very large - which is usually the case in index tracking and enhanced indexation models.

It has been recently shown that very large SSD-based models can be solved in seconds, using solution methods which apply the cutting plane approach, as proposed by Fabian, Mitra and Roman (2009). However, imposing additional constraints that add realism (for example cardinality constraints, normally required in index tracking) could increase dramatically the computational time.

Another aspect that require investigation is the amount of rebalancing needed; how does the current solution change when new information comes into place. This aspect is in connection with the stability of the model to changes in the input data. In the stochastic programming approach, where the distribution of random returns is considered as discrete and described by scenarios, new information is represented by new scenarios.

The rest of the paper is organised as follows. In section two we introduce index tracking and enhanced indexation. In section three we discuss how second order stochastic dominance (SSD) is used as a choice criterion in portfolio selection. In section four we formulate the proposed models for enhanced indexation based on SSD. The numerical experiments are presented in section five. Three datasets, drawn from FTSE 100, Nikkei and SP500 are used for backtesting the proposed models in a rebalancing frame. Conclusions are presented in section six.

2 Index Tracking and Enhanced Indexation

Let n denote the number of the assets into which we may invest at the beginning of a fixed time period. A portfolio $x = (x_1,\dots,x_n)^T$ represents the proportions of initial capital invested in the different assets. Let the random vector $R = (R_1,\dots,R_n)^T$ denote the returns of the assets at the end of the investment period. The return of the portfolio x is denoted by $R_x = R^T x$, a random variable.

Let $X \subset \mathbb{R}^n$ denote the set of the feasible portfolios. We assume that X is a convex polyhedron; for example, in the simplest case,

$$X = \{(x_1,\dots,x_n)/\sum_{j=1}^{n} x_j = 1, x_j \geq 0, \forall j \in \{1,\dots,n\}\}$$

It is usual to assume that the future returns of the assets are discrete random variables with a finite number of outcomes, obtained by scenario generation or finite sampling of historical data (this is also the assumption used throughout this paper). Consider S scenarios and p_i the probability of scenario i, $i \in \{1,\dots,S\}$; $\sum_{i=1}^{S} p_i = 1$. Let r_{ij} be the return of asset j under scenario i, $i \in \{1,\dots,S\}$, $j \in \{1,...,n\}$. Thus, the random variable representing the return of asset j is finitely distributed over $\{r_{1j},\dots,r_{Sj}\}$ with probabilities $p_1,\dots,p_S$. The random variable R_x representing the return of portfolio $x = (x_1,...,x_n)$ is finitely distributed over $\{r_{x1},...,r_{xS}\}$, where $r_{xi} = x_1 r_{i1} + ... + x_n r_{in}$, $\forall i \in \{1,...,S\}$.

The primary problem in "active" portfolio selection is how to find a portfolio x such that its return R_x is "maximised". (Since R_x is a random variable, this requires further clarification. There are various models of choice under risk that specify a preference relation among random returns. A portfolio x is then chosen such that its return R_x is non-dominated with respect to the preference relation considered. We resume this discussion in Section 4.)

Index tracking models are a somewhat special category; they are a "passive" portfolio selection strategy. Their aim is to track a financial index's return as close as possible, thus, to "minimise" the difference between R_x and the (known) return distribution R_I of the financial index. The rationale behind it is the belief (advocated by many financial experts) that a financial index routinely beats any portfolio actively created from its component assets.

Traditionally, index tracking is done by minimising the volatility of the tracking error: the sum of the squared deviations of returns on the replicating portfolio from the index (Roll, 1992). Other approaches suggest the use of absolute deviations instead of the square deviations, which leads to a linear program (LP) instead of a QP (Clarke et al., 1994;

Rudolf et al., 1999). There are also approaches where only the downside deviations from the index' returns are considered (Rudolf et al., 1999).
One drawback of these methods is that the number of the stocks included in the portfolio is very large (as an example, in our numerical experiments, described in section 5, the number of stocks selected in a tracking portfolio is at least half of the total number of stocks available); this makes the solution impractical in many cases, particularly if transaction costs are taken into consideration. For this reason, cardinality constraints (limiting the number of stocks in the composition) are normally imposed. It is known that cardinality constraints require the introduction of binary variables (one binary variables for each asset in the basket) and this hugely increases the computational difficulty of these models.
A lot of the research in the area of passive portfolio selection has been thus concentrated on modelling or solution techniques meant to handle the computational difficulty associated to index tracking - for a review, see Canakgoz and Beasley (2008), Beasley et al. (2003).
Enhanced indexation is a very new area of research - for a review, see Canakgoz and Beasley (2008). As in index tracking, the return distribution of a financial index is available and has to be "tracked" but with the intention of seeking excess return. There is no generally accepted method in this area. Usually, the same computational problems as in index tracking are encountered: cardinality constraints (thus, introduction of binary variables) have to be introduced, resulting in computational difficulty.

3 Second Order Stochastic Dominance

As stated in section 2, the problem in "active" portfolio selection is how to find a portfolio x such that its return at the end of the investment period R_x is "maximised". Since portfolio returns are random variables, models that specify a preference relation among random returns are required. A portfolio x is then chosen such that its return R_x is non-dominated with respect to the preference relation considered - this is done via an optimisation model.
For portfolio selection, mean-risk models have been by far the most popular. They describe and compare random variables using two statistics: the expected value (mean) and a risk value. Various risk measures have been proposed in the literature, see for example Markowitz (1952), Fishburn (1977), Ogryczak and Ruszczynski (1999, 2001), Rockafellar and Uryasev (2000, 2002). Mean-risk models are convenient from a computational point of view and have an intuitive appeal, but their approach is somewhat oversimplified.
Expected utility theory (von Neumann and Morgenstern, 1944) compare random returns by comparing their expected utilities (larger value preferred). However, the expected utility values depend on the utility function that is used; the choice of a specific utility function is somewhat subjective.

Stochastic dominance (SD) has been recognized as a sounder model of choice, as it exploits "the three p's: price, probability and preference" (Lo, 1999). It is closely connected to the expected utility theory, but it eliminates the need to explicitly specify a utility function (see Whitmore and Findlay, 1978 for a detailed description of stochastic dominance relations, Kroll and Levy, 1980 for a review). With stochastic dominance, random variables are compared by pointwise comparison of functions constructed from their distribution functions. There are progressively stronger assumptions about the form of utility functions used in investment, which lead to first, second and higher orders of SD. For example, first order stochastic dominance (FSD) is connected to "non-satiation" behaviour. A random return is preferred to another with respect to FSD relation if its expected utility is higher, for any *increasing* utility function. This is a strong condition and thus many random returns cannot be ordered with respect to FSD.
In portfolio selection, second order stochastic dominance (SSD) is particularly important, due to its relation to risk-averse behaviour, as explained below.

Let R and R' denote two random returns. Second-order Stochastic Dominance is defined by the following equivalent criteria:

(a) $\mathrm{E}(U(R)) \geq \mathrm{E}(U(R'))$ holds for any *increasing and concave* (integrable) utility function U.

(b) $\mathrm{E}([t-R]_+) \leq \mathrm{E}([t-R']_+)$ holds for each $t \in \mathbb{R}$.

(c) $\mathrm{Tail}_\alpha(R) \geq \mathrm{Tail}_\alpha(R')$ holds for each $0 < \alpha \leq 1$, where $\mathrm{Tail}_\alpha(R)$ denotes the unconditional expectation of the least $\alpha * 100\%$ of the outcomes of R.

For the equivalence of *(a)* and *(b)* see for example Whitmore and Findlay (1978). The equivalence of *(b)* and *(c)* is shown in Ogryczak and Ruszczyński (2002).

If the relations above hold, we say that R dominates R' with respect to SSD; we use the notation $R \succeq_{SSD} R'$.

A portfolio x is said to dominate (or be preferred to) another portfolio y with respect to SSD if $R_x \succeq_{SSD} R_y$, where R_x and R_y are the (random) returns of portfolios x and y respectively. A similar notation is used: $x \succeq_{SSD} y$.

A portfolio x is said to be non-dominated (or efficient) with respect to SSD if there is no other feasible portfolio y such that $y \succ_{SSD} x$.

It is known that increasing and concave utility functions express the preference of risk-averse investors, which is the observed economic behavior. This underlines the importance of choosing SSD efficient solutions. Unfortunately, the SSD relation is expressed as a continuum of constraints in the form (b) above. This makes SSD-based portfolio models very difficult from a computational point of view. Only recently such models have been proposed in the literature (Dentcheva and Ruszczyński 2003, 2006; Roman et al., 2006; Fabian et al., 2010).

Dentcheva and Ruszczyński (2003, 2006) consider a benchmark return; a portfolio is then constructed, such that its return dominates the benchmark with respect to SSD (in addition, a functional of the portfolio's return is optimized).

Roman et al. (2006), Fabian et al. (2010) propose models whose solutions are SSD efficient portfolios. In addition, these portfolios have return distributions that comes uniformly close to given benchmark distributions (e.g. those of financial indices).

In the following section, we present the models proposed by Roman et al. (2006), Fabian et al. (2010) from an enhanced indexation perspective.

4 Enhanced Indexation based on SSD

Roman et al. (2006), Fabian et al. (2010) consider the case of S equally probable scenarios; under this assumption, the SD relations greatly simplify, as explained below.

Denote by R_I the return of the financial index considered as a benchmark; this is a random variable with a known distribution, with S equally probable outcomes (provided, for example, from historical observations). Its ordered outcomes are denoted by $R_I^{(1)} \leq \ldots \leq R_I^{(S)}$. Denote by R a random return and with ordered outcomes $R^{(1)} \leq \ldots \leq R^{(S)}$. The first and second order SD relations can be expressed as follows:

(a) R dominates R_I with respect to FSD (notation: $R \succeq_{FSD} R_I$) if and only if: $R^{(i)} \geq R_I^{(i)}$, $i = 1 \ldots S$.

(b) R dominates R_I with respect to SSD (notation: $R \succeq_{SSD} R_I$) if and only if: $\sum\limits_{j=1}^{i} R^{(i)} \geq \sum\limits_{j=1}^{i} R_I^{(i)}$, $i = 1 \ldots S$.

Following Fabian et al. (2009), given α $(0 < \alpha \leq 1)$, we denote by $\mathrm{Tail}_\alpha(R)$ the unconditional expectation of the least $\alpha * 100\%$ outcomes of the random variable R. Thus, $\mathrm{Tail}_{\frac{i}{S}}(R) = \sum\limits_{j=1}^{i} R^{(i)}$.

The SSD relation can be expressed in relation to Conditional Value-at-Risk (CVaR) at S different confidence levels. The CVaR of a random return R at confidence level $\alpha \in (0,1)$ is the mathematical transcription of the concept "mean of losses" in the worst $\alpha \cdot 100\%$ of cases (Acerbi and Tasche 2002), where the loss is relative to zero. A formal definition of CVAR is given for example in Rockafellar and Uryasev (2000, 2002).

It follows easily that, in the case of equi-probable scenarios, $\mathrm{CVaR}_{\frac{i}{S}}(R) = -\sum_{j=1}^{i} \frac{1}{i} R^{(j)}, \quad \forall i \in \{1,\ldots,S\}$, where $\mathrm{CVaR}_{\frac{i}{S}}$ denotes the Conditional Value-at-Risk at confidence level $\frac{i}{S}$. The above SSD equivalence can be further written as:

$$\mathrm{CVaR}_{\frac{i}{S}}(R) \;\leq\; \mathrm{CVaR}_{\frac{i}{S}}(R_I) \quad \text{holding for} \quad i = 1,\ldots,S$$

Thus, for finding the SSD efficient portfolios, a multi-objective approach is proposed, in which the S objective functions (to be minimized) can be written as CVaR at confidence levels $\frac{i}{S}$, for all i in $1\ldots S$.

$$\begin{gathered}\min \quad \Big(\mathrm{CVaR}_{\frac{1}{S}}\left(R^T x\right), \ldots, \mathrm{CVaR}_{\frac{i}{S}}\left(R^T x\right), \ldots, \mathrm{CVaR}_{\frac{S}{S}}\left(R^T x\right) \Big) \\ \text{subject to} \quad x \in X, \end{gathered} \tag{1}$$

Equivalently, the objective functions (to be maximized) can be written as $\mathrm{Tail}_{\frac{i}{S}}$, for all i in $1\ldots S$:

$$\begin{gathered}\max \quad \Big(\mathrm{Tail}_{\frac{1}{S}}\left(R^T x\right), \ldots, \mathrm{Tail}_{\frac{i}{S}}\left(R^T x\right), \ldots, \mathrm{Tail}_{\frac{S}{S}}\left(R^T x\right) \Big) \\ \text{such that} \quad x \in X, \end{gathered} \tag{2}$$

The specific SSD efficient solution that comes closest, in a uniform sense, to the reference distribution R_I, is chosen by using the *reference-point method* (Wierzbiki 1982); this transforms the multi-objective formulation into a single-objective optimisation problem.
We shortly describe this approach, applied to the multi-objective model (2).
We use the following notation:

$$\widehat{\tau} = (\widehat{\tau}_1,\ldots,\widehat{\tau}_S) \;:=\; \Big(\mathrm{Tail}_{\frac{1}{S}}(R_I), \ldots, \mathrm{Tail}_{\frac{S}{S}}(R_I) \Big).$$

The reference point method introduces a concave "achievement" function $\Gamma_{\widehat{\tau}}$ whose arguments are the components of the objective in (2). The simplest achievement function is

$$\Gamma_{\widehat{\tau}}(\tau_1,\ldots,\tau_S) \;:=\; \min_{1\leq i\leq S}(\tau_i - \widehat{\tau}_i) = \min_{1\leq i\leq S}(\mathrm{Tail}_{\frac{i}{S}}(R^T x) - \mathrm{Tail}_{\frac{i}{S}}(R_I)). \tag{3}$$

(Thus, the "achievement" function considers the worst difference between the tails of the resulting portfolio return and the tails of the index. A term $\varepsilon \sum_{i=1}^{S} (\tau_i - \widehat{\tau}_i)$ with a small positive ε is usually added to ensure Pareto-efficiency of the optimal solution, as described in Roman et al., 2006. The differences between the tails of the resulting portfolio return and the tails of the index at confidence levels $\frac{i}{S}$, $i = 1\ldots S$, are called "partial achievements").

The single objective optimization problem basically maximizes the worst "partial achievement" over $x \in X$:

$$\begin{gathered}\max \quad \Gamma_{\widehat{\tau}} \Big(\mathrm{Tail}_{\frac{1}{S}}\left(R^T x\right), \ldots, \mathrm{Tail}_{\frac{S}{S}}\left(R^T x\right) \Big) \\ \text{such that} \quad x \in X. \end{gathered} \tag{4}$$

Denoting by $\vartheta = \min_{1\leq i\leq S}(\mathrm{Tail}_{\frac{i}{S}}\left(R^T x\right) - \widehat{\tau}_i)$ the worst partial achievement, the above problem is written as:

$$\begin{gathered}\max \quad \vartheta \\ \text{such that} \quad \vartheta \in \mathbb{R}, \quad x \in X \\ \vartheta \leq \mathrm{Tail}_{\frac{i}{S}}\left(R^T x\right) - \widehat{\tau}_i \qquad (i = 1,\ldots,S). \end{gathered} \tag{5}$$

To compute the quantities $\mathrm{Tail}_{\frac{i}{S}}(\left(R^T x\right))$, Roman et al. (2006) used the CVaR-optimization formula of Rockafellar and Uryasev (2000, 2002). This approach requires introduction of a large number of additional variables. The result is a LP model of very large size, if the number of scenarios S is high. (There are more than S^2 variables and constraints. The number of assets n poses much less difficulty, since the number of constraints/ variables grows only linearly with n. No binary variables are required). Only models of relatively small sizes (up to 500 scenarios) could be solved with this approach.

In Fabian et al. (2009), a cutting-plane approach is used for computing the quantities $\mathrm{Tail}_{\frac{i}{S}}(R^T x)$; this is based on the cutting plane representation of CVaR proposed by Künzi-Bay and Mayer (2006).
In the case of equally probable scenarios, a very intuitive cutting-plane representation for $\mathrm{Tail}_{\frac{i}{S}}$ $(i = 1, \ldots, S)$ follows:

$$\begin{aligned} \mathrm{Tail}_{\frac{i}{S}}\left(R^T x\right) = \quad & \tfrac{1}{S} \min \sum_{j \in \mathscr{J}} r^{(j)T} x \\ & \text{such that} \quad \mathscr{J} \subset \{1, \ldots, S\}, \ |\mathscr{J}| = i. \end{aligned} \tag{6}$$

Using (6), the achievement-maximization problem (5) can be re-formulated to:

$$\begin{array}{lll} \max & \vartheta & \\ \text{such that} & \vartheta \in \mathbb{R}, \quad x \in X, & \\ & \vartheta + \widehat{\tau}_i \le \frac{1}{S} \sum\limits_{j \in \mathscr{J}_i} r^{(j)T} x & \text{for each} \quad \mathscr{J}_i \subset \{1, \ldots, S\}, \ |\mathscr{J}_i| = i, \\ & & \text{where} \quad i = 1, \ldots, S. \end{array} \tag{7}$$

No additional variables are introduced in the above formulation. Theoretically an astronomical number of cuts are required, but in practice only a few of them are needed. Fabian et al. (2009) propose a cutting-plane solution method for solving the above problem and show that the computational time is dramatically decreased: problems with tens of thousands of scenarios are solved within seconds.

The above model (7) is based on comparison of "unscaled" tails $\mathrm{Tail}_{\frac{i}{S}}$, $i = 1 \ldots S$ (of the chosen portfolio and of the index return distributions). A similar model, based on the comparison of "scaled" tails (or, equivalently, CVaR's at confidence levels $\frac{i}{S}$, $i = 1 \ldots S$) is proposed by Fabian et al. (2010). Applying similarly the reference point method to (1) and using cutting-plabe resprentations of CVaRs, the following model results:

$$\begin{array}{lll} \max & \vartheta & \\ \text{such that} & \vartheta \in \mathbb{R}, \quad x \in X, & \\ & \frac{i}{S}\vartheta + \widehat{\tau}_i \le \frac{1}{S} \sum\limits_{j \in \mathscr{J}_i} r^{(j)T} x & \text{for each} \quad \mathscr{J}_i \subset \{1, \ldots, S\}, \ |\mathscr{J}_i| = i, \\ & & \text{where} \quad i = 1, \ldots, S. \end{array} \tag{8}$$

Remark 1 *Models (8) and (7) provide different solutions (both of them non-dominated with respect to SSD). The "unscaled" model (7) usually provides the portfolio that improves most on the worst outcome of the reference distribution (i.e. that maximises the difference between the worst outcome of the chosen portfolio and the worst outcome of the reference distribution).*
The "scaled" model (8) provides the portfolio that improves the most on a CVaR at a confidence level i/S, $i = 1 \ldots S$ of the reference distribution.
A natural question is whether one provides better solutions than the other. Fabian et al. (2010) show that the model (8) presents advantages over the "unscaled" model (7) from a theoretical point of view: it can be formulated as a risk minimisation model, considering a convex risk measure. In addition, it may present advantages from a practical point of view: the (in-sample) return distributions of the portfolios chosen with (8) are somewhat "shifted to the right" as compared to those obtained with (7), indicating overall higher returns - except for a small portion in the left tail. That

is, under the most unfavourable scenarios, the unscaled model (7) could provide better solutions, i.e. leading to a less dramatic loss.

The models (8) and (7) described above are never infeasible, but always provide solutions that are SSD efficient - irrespective of the benchmark chosen by the user.
The benchmark reference distribution however plays an important role in choosing the solution. Our interest is the case when the benchmark distribution is that of a financial index. As far as we are aware, all numerical experiments checking the SSD efficiency of financial indexes' distributions (Roman et al. 2006, Fabian et al. 2009, Fabian et al. 2010, Dentcheva and Ruszczynski 2003, 2006) led to the same conclusion: that the indices were dominated with respect to SSD.
(In our models (8) and (7), it is easy to check this: a positive optimum indicates a reference distribution that is SSD dominated - please see Roman et al. 2006, Fabian et al. 2009 for more details.)

Thus, a portfolio is chosen that *improves* on the index's distribution, in the sense that it dominates it with respect to SSD. More precisely, the tails of the index are increased (or, equivalently, the CVaR's are decreased) until SSD efficiency is obtained.
For this reason, the models (8) and (7) can be viewed as enhanced indexation models.

Remark 2 *Both models (8) and (7) are based on CVaR minimisation (at different confidence levels). Empirical studies show that portfolios obtained by CVaR minimisation have less stocks in the composition than portfolios based on variance minimisation. This aspect is resumed in Section 5.2, where the cardinality of portfolios chosen by SSD-models is investigated.*

Roman et al. (2006), Fabian et al. (2010) conducted numerical tests in order to compare the return distributions of the chosen portfolios with those of the reference distributions. The chosen portfolios had clearly better return distributions than those of the corresponding financial indices (higher mean, lower variance, higher skewness, better left tail, etc.)
All these results were obtained in-sample, in a single period framework. Of more practical importance is however the actual, realized performance of these portfolios measured *over time*, as well as the easiness in applying these models. The realised performance of the portfolios is measured by the amount of excess return over the corresponding index. The easiness in applying the models is assessed by the number of stocks in the composition of the chosen portfolios and the amount of rebalancing needed.
(We earlier explained that models (8) and (7) do not pose computational problems and are fast to solve even for very large datasets. This is based on the assumption that no cardinality constraints are introduced. However, if the number of stocks selected in the portfolios is large, cardinality constraints have to be introduced and the same computational problems as in tracking error minimisation would be encountered.)

These aspects are investigated in the next section.

5 Computational study

5.1 Scope of the study, dataset and implementation issues

We consider the two SSD-based models (8) and (7) in a rebalancing frame; for the purpose of comparison, we also consider the classical tracking error minimisation model (which minimises the mean of the squared deviations from the index). We investigate the following aspects:

- the effectiveness in obtaining excess return on the financial index considered;
- the number of stocks in the composition of the chosen portfolios;
- the amount of rebalancing needed.

The data used in this analysis is drawn from 3 financial indices: FTSE 100, Nikkei 225 and SP 500. We considered a (rolling) 1 week investment period. We used weekly past historical returns of the component stocks and of the corresponding index as scenarios for the returns of the week following the decision.

The FTSE 100 dataset has 101 component stocks (in addition to the FTSE100 index) that are monitored weekly over the period 12 Dec 2006 to 5 March 2009 [1]. The period 1 January - 5 March 2009 (10 weeks) is used for backtesting analyis. (To be more precise, we use the period 12 Dec 2006 to 25 Dec 2008 as scenarios for the next week, i.e. 1 Jan 2009. We ran the models, obtain the solution portfolios and compute their *actual* returns for 1 Jan 2009. For the following problem, the in-sample data set has been updated to include 1 Jan 2009. We ran the models, obtain the solution portfolios and compute their *actual* returns for 8 Jan 2009, so forth.)

The Nikkei 225 dataset has 225 component stocks (in addition to the Nikkei 225 index), monitored over the period 28 Nov 2005 to 5 March 2009 (as with the FTSE 100 dataset, we chose the number of time periods in such a way that maximum number of stocks in the universe can be considered.) The same as with FTSE 100 dataset, the last ten periods (1 Jan - 5 March 2009) are used for backtesting.

The SP 500 dataset has 491 component stocks (in addition to the SP 500 index), monitored over the period 20 Sept 2006 to 5 March 2009. The same as with the other two datasets, the last ten periods (1 Jan - 5 March 2009) are used for backtesting. The returns for the in-sample time periods were bootstrapped in order to obtain a number of scenarios greater than the number of assets - we considered 600 (in-sample) scenarios. (For details on bootstrapping, please see Efron and Tibishirani, 1993).

For each dataset and each model, there are thus 10 chosen portfolios (weekly decisions for the period 1 Jan - 5 March 2009). The ex-post returns of the chosen portfolios are computed and compared to the historical weekly returns of the corresponding indices in the period 1 Jan - 5 March 2009.

Following Fabian et al. (2009), the models were implemented using the AMPL modelling system (Fourer, Gay and Kernighan 1989) and the AMPL COM Component Library (2005), integrated with C functions; the models were solved using the FortMP solver (Ellison et al., 1989). In all instances, the computational time was very small (a few seconds).

5.2 Test results

In all cases, the portfolios chosen by the SSD based models had an overall better performance than the corresponding indices and the tracking error minimisers. This is better underlined by computing the ex-post compounded returns (Figures 2, 4, 6).

The SSD-based portfolios have a small number of stocks in the composition (usually, much less than one tenth of the available stocks), which makes the imposition of cardinality constraints unnecessary. (This behaviour is due to the CVaR-minimisation nature of the SSD-based models, as explained in Section 5). In comparison, the tracking errors minimisers contain at least half of the available stocks in their composition.

In addition, the composition of the SSD-based portfolios is generally stable over the rebalancing period. This is true particularly for the case of the unscaled SSD model (7); in most cases, no rebalancing is necessary from one week to another.

In the case of the FTSE 100 dataset, both portfolios chosen by SSD models had a better ex-post performance than the FTSE 100 index and the ”tracker” (i.e. the tracking error minimiser). The difference between index tracking and enhanced indexation is well illustrated by Fig. 1. While the tracker follows very closely the index’s movements, the other two portfolios have generally higher returns than the index, particularly in the periods of large losses.

The good performance of portfolios chosen by the SSD-based models is even better underlined by computing the *compounded* ex-post returns (Fig. 2). Investment in the portfolios chosen by the SSD based models would not have led to loss, despite a backtesting period with falling prices. Particularly the portfolio chosen with the scaled model (8) had a very good performance, resulting in a 2% gain at the end of the 10-weeks backtesting period (the FTSE 100 index had a 16% loss, the tracker had a 15% loss while the portfolio obtained with the unscaled model (7) had a 0.3%

---

[1] We chose this period in order to have maximum number of stocks in the asset universe; had we considered a higher number of time periods, the number of stocks with full data would have been smaller, since stocks move in and out of index. Our asset universe is made up of stocks which are present in the index for the entire period of observation.

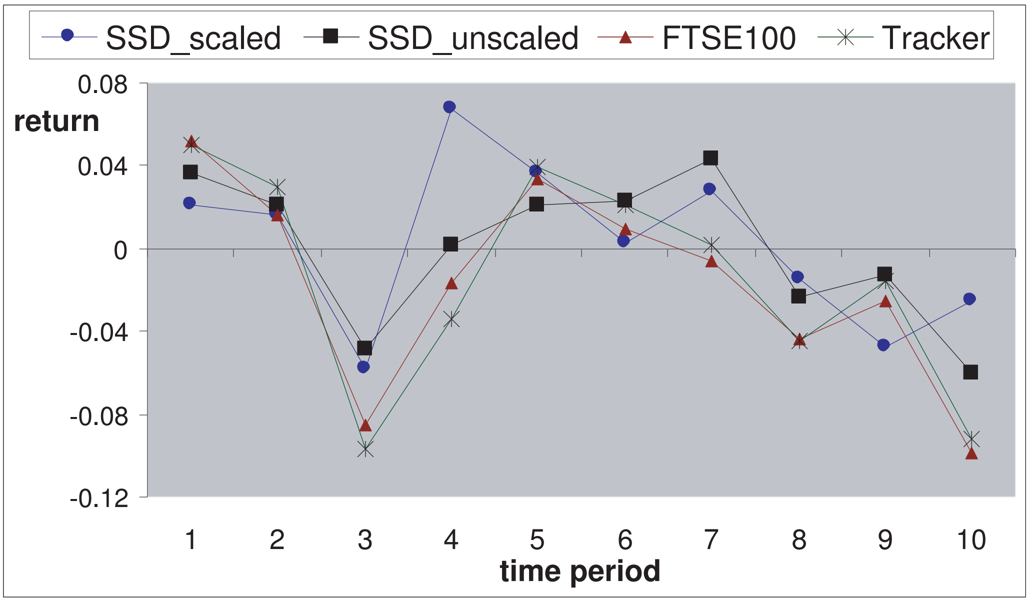

Figure 1: FTSE 100 dataset: Ex-post weekly returns 1 Jan - 5 Mar 2009; FTSE 100, index tracker, portfolios chosen by SSD models

loss. We notice that, for the previous backtesting periods the portfolio obtained with the unscaled model (7) did not incur losses.

The number of stocks in the composition of the SSD-based portfolios is about one tenth of the total number of stocks in the universe (which is 101).
The portfolios chosen by the scaled model (8) have on average 8 stocks in the composition (the highest number of stocks is 9, for the portfolio chosen on the 9-th backtesting week). The composition of the portfolios has small changes over the 10 weeks (see Table 5.2), which means that (only) little rebalancing is necessary from one period to another.
The case of the scaled model (8) is even more remarkable. The composition of the portfolios remains unchanged for the whole backtesting period. Thus, the same portfolio is chosen for all the 10 backtesting time periods; this portfolio has 11 stocks in the composition.
In comparison, the tracking error minimiser has on average 58 stocks in composition.

In the case of the Nikkei 225 dataset, the portfolio chosen with the unscaled model (7) had a consistently better performance than the Nikkei 225 index and the index tracker, particularly in the periods of heavy loss (e.g. week 4); see Fig. 3. This is clearly underlined by computing the ex-post compounded returns (Fig. 4). The backtesting period was one with acute falling prices for the stocks in Nikkei 225: at the end of the 10 backtesting weeks, Nikkei 225 and the index tracker had a loss of 19%, while the portfolio chosen with the unscaled model (7) had a loss of ”only” 10%. The portfolio chosen with the scaled model (8) had a loss of 21%.
The better performance of the unscaled model (7) is possibly due to the unusually ”tough” backtesting period; as underlined in Remark 1, Section 4, this model performs better (than the scaled model) under the most unfavourable scenarios.
With respect to the composition of the portfolios, again the unscaled model (7) showed a remarkable behaviour. No rebalancing has been done over the ten weeks of 1 Jan - 5 Mar 2009 and the chosen portfolios has only 3 stocks in the composition.
The portfolios chosen with the scaled model (7) have on average 15 stocks in the composition (out of 225); 17 is the highest number of stocks in the composition of such portfolios (corresponding to the 9-th backtesting week). As in the

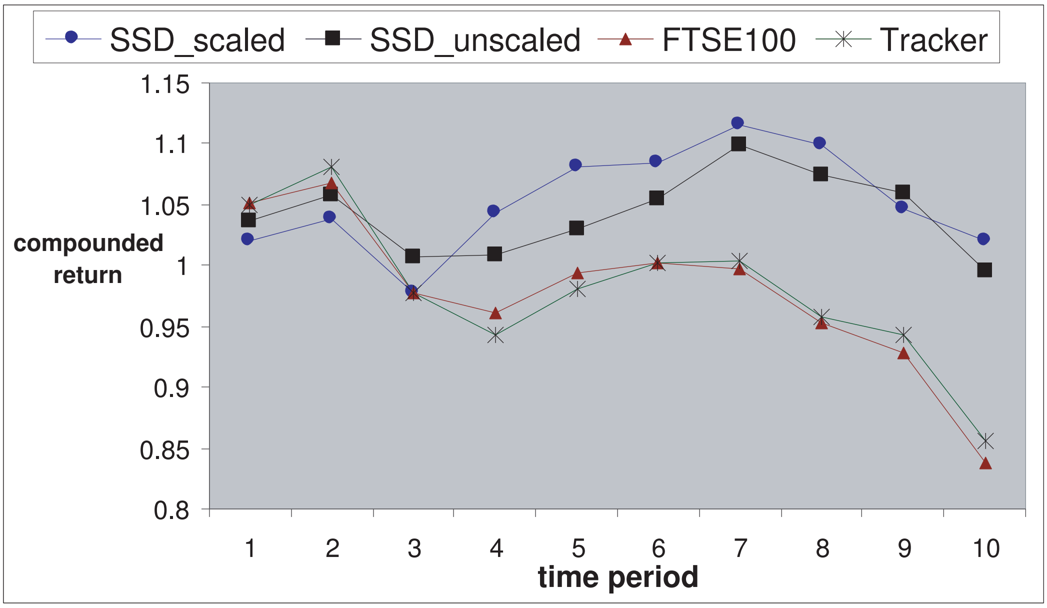

Figure 2: FTSE 100 dataset: Ex-post compounded weekly returns 1 Jan - 5 Mar 2009; FTSE 100, index tracker, portfolios chosen by SSD models

stock no	week1	week 2	week 3	week 4	week 5	week 6	week 7	week 8	week 9	week 10
9	19%	18%	25%	25%	21%	22%	22%	21%	21%	21%
18	28%	27%	25%	27%	31%	32%	30%	29%	29%	34%
43										2%
57	4%								1%	
64	22%	23%	18%	19%	22%	22%	21%	22%	21%	18%
65									1%	
80	3%	9%	2%		2%	3%	2%	4%	9%	2%
82	13%	12%	13%	13%	13%	13%	16%	15%	12%	9%
83							4%	4%	1%	6%
92	4%	5%	6%	6%	3%	3%	4%	5%	6%	8%
93	8%	5%	11%	10%	8%	5%				

Table 1: The FTSE 100 dataset: composition of the portfolios chosen with the SSD scaled model (8), over the 10 weeks 01 Jan - 05 Mar 2009

case of the FTSE 100 dataset, the composition of the portfolios is stable from one week to the next. In comparison, the tracking error minimisers have on average 118 stocks in composition.

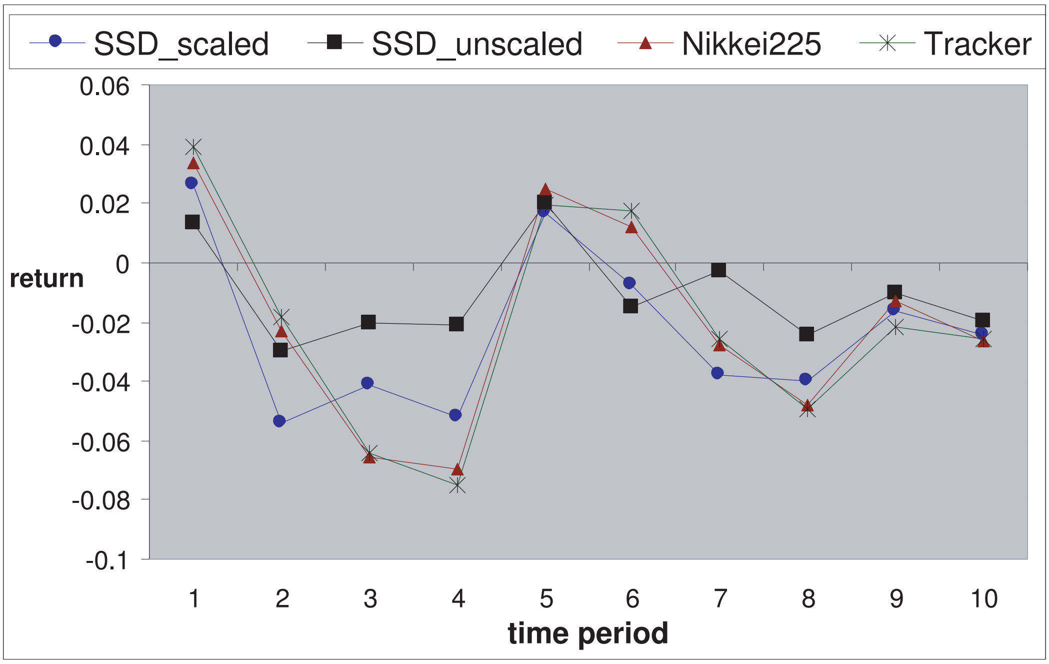

Figure 3: Nikkei 225 dataset: Ex-post weekly returns 1 Jan - 5 Mar 2009; Nikkei 225, index tracker, portfolios chosen by SSD models

In the case of the SP 500 dataset, both SSD-based portfolios, and especially those resulted from the scaled model (8), had a better performance than the index (Figures 5, 6). We notice that again, the SSD-based portfolios (particularly the "unscaled" ones) have a much "smoother" trajectory than the index (Fig. 5), in the sense that they did not incur such abrupt falls(like in weeks 2, 5, 7, 9).
The backtesting period was one with acute falling prices for the stocks in SP 500: at the end of the 10 backtesting weeks, the index had a loss of 20%, while the SSD-based portfolios incurred losses of "only" 17.6% (the scaled model) and 18% (the unscaled model); please see Fig. 6.

With respect to the composition of the portfolios, again the SSD-based portfolios contain a small number of stocks which makes the imposition of cardinality constraints unnecessary. The portfolios chosen with the scaled model (8) have on average 11 stocks in composition (out of 491). The highest number of stocks in the composition of such a portfolio is 16 (corresponding to the 2-nd backtesting period).
The unscaled model (8) shows again a remarkable behaviour in this respect. The initial portfolio is rebalanced only 3 times during the 10-week backtesting period. The four portfolios chosen over the backtesting period are composed of 5, 5, 7 and 8 stocks respectively (see Table 2 for the composition of these portfolios). The necessary rebalancing is however higher than in the case of the previous two datasets.

6 Summary and Conclusions

Two portfolio selection models that use SSD as a choice criterion were presented from an enhanced indexation perspective. Both models consider a financial index's return distribution; they produce portfolios whose return distributions improves on the index's until SSD efficiency is attained. One model is based on comparison of tails (i.e., sums of

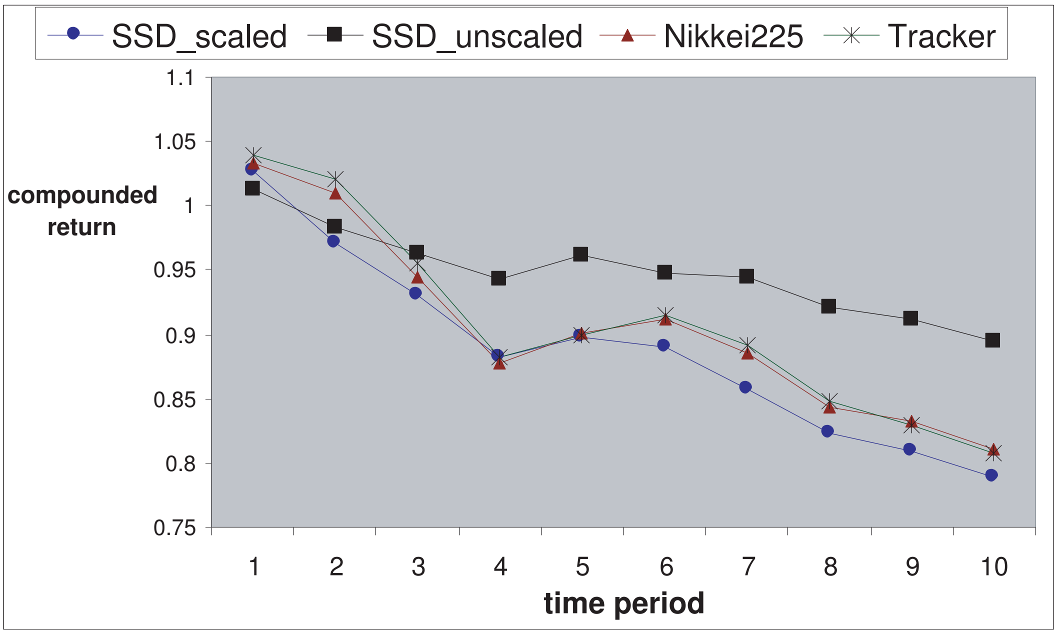

Figure 4: Nikkei 225 dataset: Ex-post compounded weekly returns 1 Jan - 5 Mar 2009; Nikkei 225, index tracker, portfolios chosen by SSD models

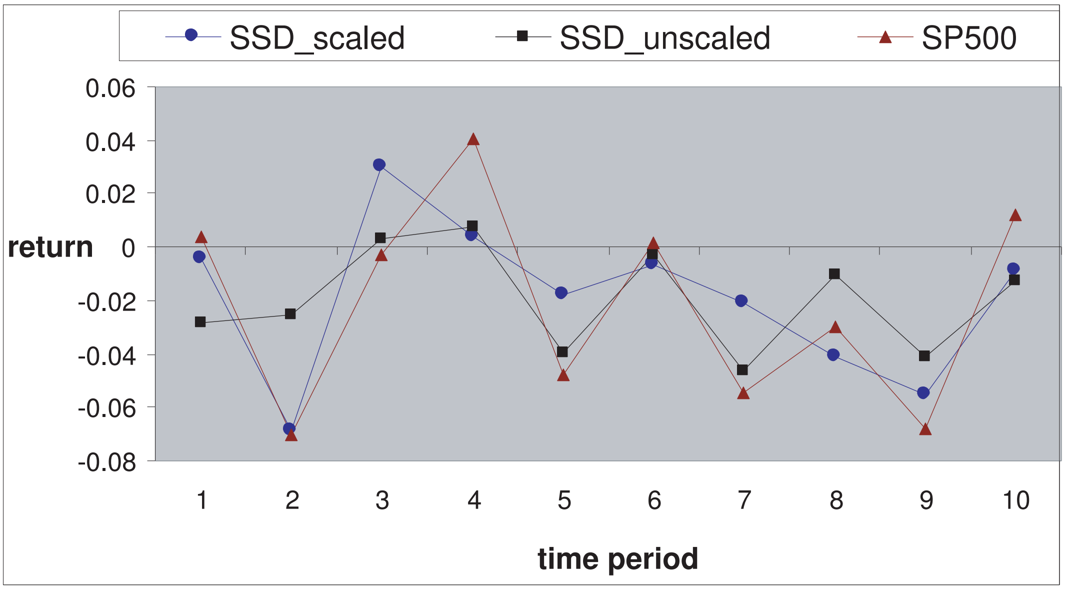

Figure 5: SP 500 dataset: Ex-post weekly returns 1 Jan - 5 Mar 2009; SP 500 and portfolios chosen by SSD models

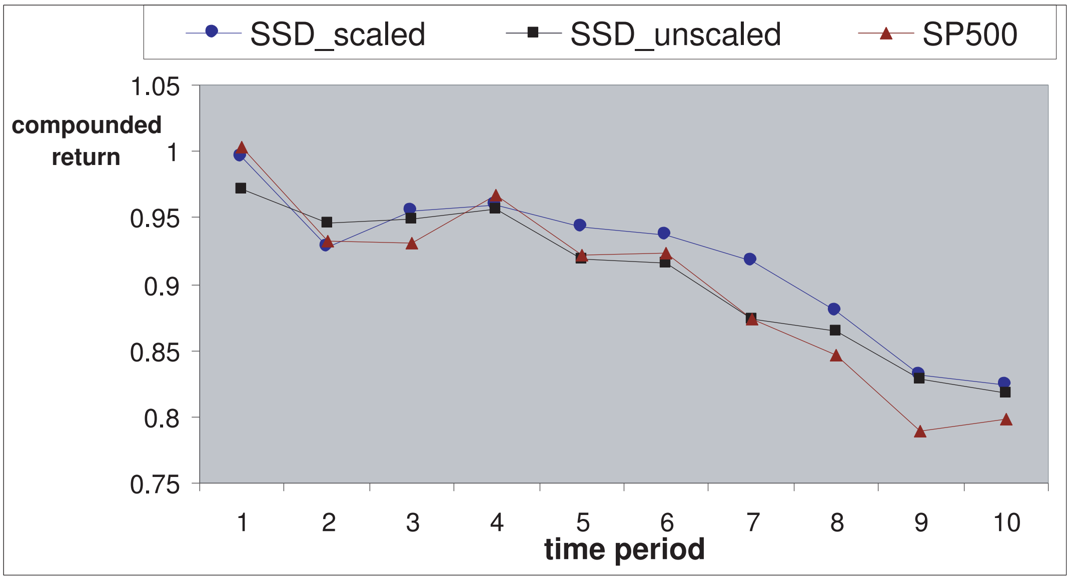

Figure 6: SP 500 dataset: Ex-post compounded weekly returns 1 Jan - 5 Mar 2009; SP 500 and portfolios chosen by SSD models

stock no	weeks 1-3	weeks 4-7	week 8	weeks 9-10
5		12%	8%	
33	31%		7%	22%
114				15%
118	8%	9%	8%	4%
172	1%			11%
194	53%	66%	58%	31%
218				6%
239		9%	7%	
383	7%	4%	2%	4%
400				6%
466			9%	

Table 2: The SP 500 dataset: composition of the portfolios chosen with the SSD unscaled model (7), over the 10 weeks 01 Jan - 05 Mar 2009

ordered outcomes) between the chosen portfolio's return distribution and the index's return distribution. (We referred to this model as the ”unscaled” model). The other model uses comparison of scaled tails (i.e., means of ordered outcomes), or, equivalently, comparison of CVaRs at different confidence levels. (We referred to this model as the ”scaled” model).

Both models can be solved within seconds even for very large datasets (thousands of assets and scenarios), using a cutting plane solution method and algorithm.

We tested the effectiveness of these two models as enhanced indexation strategies, using three datasets: FTSE 100 (101 stocks), Nikkei 225 (225 stocks) and SP 500 (491 stocks). We used the first 10 weeks of 2009 as a backtesting period, in a (weekly) rebalancing frame. Three conclusions were drawn.

First, the SSD-based models consistently outperformed the corresponding indices, in the sense that higher returns were obtained over most of the considered backtesting period. This aspect was better emphasised by computing the compounded returns. The backtesting period was one with severe falls in the stock market; the SSD-based models resulted in strategies that either managed to avoid loss (as in the case of FTSE 100) or obtained a less dramatic loss than the indices (as in the case of Nikkei 225 and SP 500).

Secondly, the imposition of cardinality constraints is unnecessary in the two SSD-based models. Due to their CVaR-minimisation nature, these models naturally select a much lower number of stocks (than the consecrated index tracking models). Usually, the SSD-based models select less than one tenth of the available stocks (in most cases, much less than this). (In comparison, the tracking errors minimisers select more than half of the available stocks and thus the cardinality has to be explicitly limited.) Not imposing cardinality constraints has big advantages from a computational point of view, since no additional binary variables are required.

Finally, the amount of necessary rebalancing in the SSD-based models is low, since the models are stable with the introduction of new scenarios, representing new information on the market. In particular, the unscaled model has a remarkable behaviour from this point of view. In the case of FTSE 100 and Nikkei 225 datasets, the composition of the initially chosen portfolios remained unchanged over the ten rebalancing weeks. In the SP 500 dataset, the composition of the initially chosen portfolio has changed three times over 10 rebalancing periods.

There are several future work directions in this area. First, we propose to evaluate the performance of the SSD-based models, compared to that of the indices, for a backtesting period with rising prices.

Secondly, we propose to formally investigate the stability of the SSD based models with respect to changes in the scenario set; the results of the empirical study presented in this paper are encouraging.

Finally, we propose to investigate the possibility of using the unscaled SSD-based model as a rebalancing trigger.

Acknowledgements

Dr Diana Roman's work has been supported by BRIEF (Brunel University Research Innovation and Enterprise Fund).
The PhD studies of Victor Zverovich have been supported by OptiRisk Systems.
These sources of support are gratefully acknowledged.

References

[1] AHMED, P. and S. NANDA (2005). Performance of Enhanced Index and Quantitative Equity Funds. *Financial Review* **40**(4), 459-479.

[2] AMPL COM COMPONENT LIBRARY (2005). User's Guide Version 1.6. *Internal T. Report.* http://www.optirik-systems.com/products/AMPLCOM.

[3] DIBARTOLOMEO, D. (2000). The Enhanced Index Fund as an Alternative to Indexed Equity Management. *Northfield Information Services documents* , available at http://www.northinfo.com/documents/70.pdf.

[4] BEASLEY, J. E. and MEADE, N. and T. J. CHANG (2003). An evolutionary heuristic for the index tracking problem. *European Journal of Operational Research* **148**, 621-643.

[5] CANAKGOZ, N. A. and J.E. BEASLEY (2008). Mixed-integer programming approaches for index tracking and enhanced indexation. *European Journal of Operational Research* **196**, 384-399.

[6] CLARKE, R.C. and KRASE, S. and M. STATMAN (1994). Tracking errors, regret, and tactical asset allocation. *The Journal of Portfolio Management* **20**, 16-24.

[7] DENTCHEVA, D. and A. RUSZCZYŃSKI (2006). Portfolio optimization with stochastic dominance constraints. *Journal of Banking & Finance* **30**, 433-451.

[8] ELTON, E. J. and M. J. GRUBER (1995). Modern Portfolio Theory and Investment Analysis. 5th edition, Wiley, New York.

[9] FABIAN, C.I., MITRA, G., and ROMAN, D. (2009). Processing second- order stochastic dominance models using cutting-plane representations. *Mathematical Programming*, DOI 10.1007/s10107-009-0326-1.

[10] FABIAN, C.I., MITRA, G., ROMAN, D.and ZVEROVICH, V. (2010). An enhanced model for portfolio choice with SSD criteria: a constructive approach. *Quantitative Finance*, to appear.

[11] FAMA, E. (1970). Efficient Capital Markets: A Review of Thoery and Empirical Work *Journal of Finance*, **25**(2), 383-417.

[12] FISHBURN, P.C. (1977). Mean-Risk Analysis with Risk Associated with Below-Target Returns, *American Economic Review* , **67** (2), 116-126.

[13] FOURER, R., D. M. GAY, and B. KERNIGHAN. (1989). *AMPL: A Mathematical Programming Language*.

[14] KRAUSE, S. (2009). *On the Performance of Quantitative Equity Mutual Funds and Enhanced Index Funds.* Master Thesis, Department of Finance, Maastricht University, The Netherlands.

[15] KROLL, Y. and H. LEVY (1980). Stochastic dominance: A review and some new evidence. *Research in Finance* **2**, 163-227.

[16] KÜNZI-BAY, A. and J. MAYER (2006). Computational aspects of minimizing conditional value-at-risk. *Computational Management Science* **3**, 3-27.

[17] LO, A. W. (1999).The Three Ps of Total Risk Management. *Financial Analysts Journal* **55**(1), 13-26.

[18] LO, A. W. (2005). Reconciling Efficient Markets with Behavioral Finance:The Adaptive Markets Hypothesis. *The Journal of Investement Consulting* **7**(2), 21-44. .

[19] MARKOWITZ, H. (1952). Portfolio Selection. *Journal of Finance* **7**, 77-91.

[20] OGRYCZAK, W., and A. RUSZCZYNSKI (1999). From Stochastic Dominance to Mean- Risk Models: Semideviations as Risk Measures. *European Journal of Operational Research* **116**, 33-50.

[21] OGRYCZAK, W., and A. RUSZCZYNSKI (2001). On Consistency of Stochastic Dominance and Mean-Semideviations Models. *Mathematical Programming* **89**, 217-232.

[22] OGRYCZAK, W. and A. RUSZCZYŃSKI (2002). Dual stochastic dominance and related mean-risk models. *SIAM Journal on Optimization* **13**, 60-78.

[23] ROCKAFELLAR, R.T. and S. URYASEV (2000). Optimization of conditional value-at-risk. *Journal of Risk* **2**, 21-41.

[24] ROCKAFELLAR, R.T. and S. URYASEV (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance* **26**, 1443-1471.

[25] ROLL, R. (1992). A mean/variance analysis of tracking error. *The Journal of Portfolio Management* **18**, 13-22.

[26] ROMAN, D., K. DARBY-DOWMAN, and G. MITRA (2006). Portfolio construction based on stochastic dominance and target return distributions. *Mathematical Programming,* Series B **108**, 541-569.

[27] RUDOLF, M. and WOLTER, H. J. and H. ZIMMERMANN (1999). A linear model for tracking error minimization. *Journal of Banking & Finance* **23**, 85-103.

[28] SCOWCROFT, A. and SEFTON, J. (2003). Enhanced Indexation, in *Advances in Portfolio Construction and Implementation*, Chapter 4 (S. Satchell and A. Scowcroft, eds), Oxford.

[29] WHITMORE, G.A. and M.C. FINDLAY (1978). Stochastic Dominance: An Approach to Decision- Making Under Risk. *D.C.Heath, Lexington, MA*.