



WORKING PAPER SERIES

Cardinality versus q -Norm Constraints for Index Tracking

Bjorn Fastrich, Sandra Paterlini and Peter Winker

Working Paper 56

January 2011

www.recent.unimore.it

Cardinality versus q -Norm Constraints for Index Tracking

Björn Fastrich* Sandra Paterlini† Peter Winker‡

July 22, 2010

Abstract

Index tracking aims at replicating a given benchmark with a smaller number of its constituents. Different quantitative models can be set up to determine the optimal index replicating portfolio. In this paper, we propose an alternative based on imposing a constraint on the q -norm, $0 < q < 1$, of the replicating portfolios' asset weights: the q -norm constraint regularises the problem and identifies a sparse model. Both approaches are challenging from an optimisation viewpoint due to either the presence of the cardinality constraint or a non-convex constraint on the q -norm. The problem can become even more complex when non-convex distance measures or other real-world constraints are considered. We employ a hybrid heuristic as a flexible tool to tackle both optimisation problems. The empirical analysis on real-world financial data allows to compare the two index tracking approaches. Moreover, we propose a strategy to determine the optimal number of constituents and the corresponding optimal portfolio asset weights.

Keywords: Index tracking, Cardinality constraint, q -norm, Regularization methods, Heuristic algorithms.

1 Introduction

Index tracking (or benchmark replication) is a passive financial strategy that aims at replicating the performance and risk-profile of a given index (or benchmark). One of the most common approaches to tackle the index tracking problem consists of minimizing a given tracking error measure while limiting the maximum number of assets held in the portfolio. Having few active positions reduces the administrative and transaction costs and avoids detaining very small and illiquid positions, especially when the index has a large number of constituents. However, imposing an upper bound on the number of constituents

*Dept. of Economics, University of Giessen, Germany.

†Dept. of Economics, CEFIN and RECent, University of Modena and Reggio E., Italy.

‡Dept. of Economics, University of Giessen, Germany, and Center for European Economic Research (ZEW), Mannheim, Germany.

of the tracking portfolio makes the optimisation problem NP-hard (see, e.g., Coleman *et al.*, 2006). Different quantitative approaches have been proposed to tackle such an optimisation problem. Two comprehensive literature reviews on the main quantitative methods can be found in Beasley *et al.* (2003) and Canagkoz and Beasley (2008). Most approaches rely on search heuristics, which have proven successful in high dimensional contexts (see, e.g., Gilli and Kellezi, 2002*b*; Beasley *et al.*, 2003; Derigs and Nickel, 2003; Maringer and Oyewumi, 2007; Krink *et al.*, 2009; Gilli and Winker, 2009).

While the optimisation challenge has attracted large interest among researchers and practitioners, so far not much attention has been devoted to develop strategies that would provide the tracking portfolio with some other ideal characteristics. This includes the development of less expensive strategies by controlling the turnover and the maintenance of a good tracking performance both in-sample and out-of-sample. Recently, statistical regularisation methods have found application in mean-variance portfolio settings (DeMiguel *et al.*, 2008; Brodie *et al.*, 2008) in order to promote the identification of sparse (with a small number of constituents) portfolios with good out-of-sample properties and low turnover. The proposed approaches rely on imposing upper bounds on the 2-norm or on the 1-norm of the vector of the portfolio weights as suggested by the Ridge regression (Hoerl and Kennard, 1970) and the LASSO (Tibshirani, 1996) approach, respectively. Both methods can be considered as a special version of the (B)ridge regression approach (Frank and Friendman, 1993), where an upper bound is imposed on the q -norm ($0 < q < \infty$).¹ Empirical results in a mean-variance framework (DeMiguel *et al.*, 2008; Brodie *et al.*, 2008) support the use of the LASSO method when short selling is allowed. However, the LASSO approach is ineffective in promoting sparsity in presence of the budget and no-short selling (typical in index tracking) constraints, since the 1-norm of the asset weights will have a constant value of one.

One valid alternative could then be to consider a constraint on the q -norm with $0 < q < 1$. The lower the upper bound on the q -norm is, the more sparse and less diversified (with larger weights) are the portfolios. In fact, as suggested by Fernholz *et al.* (1998), the q -norm of the asset weights can also be considered as a measure of diversity of the portfolio, which, when the no-short selling constraint is imposed, has maximum value for the equally weighted portfolio and minimum value for a portfolio totally invested in a single asset. Hence, by simply imposing an upper bound on the q -norm, we could identify the tracking portfolio with the desirable maximum number of assets. Furthermore, this new formulation of the index tracking problem allows the development of strategies to identify the optimal maximum number of assets to be held in the tracking portfolio. From an optimisation viewpoint, the problem is still NP-hard due to the presence of the non-convex constraint on the q -norm and consequently heuristics could be preferred to classic optimisation techniques (Canagkoz and Beasley, 2008).

¹We follow the common practice in the literature to refer to a norm despite the fact that for cases with $0 < q < 1$ it does not define a norm but a quasi-norm, because it violates the triangle inequality.

In this paper, we introduce the q -norm formulation of the index-tracking problem, provide a flexible search heuristic that can effectively deal with the two NP-Hard optimisation problems and present some strategies to determine the optimal number of constituents and the corresponding optimal portfolio asset weights.

Section 2 introduces the formulation of the index tracking problem based on the q -norm constraint as an alternative to the one which relies on the cardinality constraint. Section 3 describes the main challenges of the two optimisation problems as well as the heuristic we propose to tackle the two problems. Section 4 presents the experimental set-up to compare the two approaches, while section 5 reports the results of an empirical comparison on real-world data. Section 6 discusses some possible strategies to determine the optimal number of active weights before we conclude in section 7.

2 Cardinality and q -Norm Constraints for Index Tracking

Index tracking aims at replicating a given benchmark with a smaller number of benchmark constituents. One of the most common quantitative approaches is to tackle the problem as an optimisation problem with a cardinality constraint. Hence, when no short selling is allowed, the optimisation problem can be formulated as:

$$\arg \min_{\mathbf{w}} TE(\mathbf{w}), \quad (1)$$

$$\sum_{i \in \mathcal{J}} w_i = 1, \quad (2)$$

$$0 \leq w_i \leq 1, \quad (3)$$

$$\#\mathcal{J} \leq K_{max}, \quad (4)$$

where the target is to determine the $K \times 1$ asset weight vector $\mathbf{w} = [w_1, \dots, w_K]'$ that minimizes a given tracking error measure $TE(\mathbf{w})$ with number of active positions $\#\mathcal{J}$ ($\mathcal{J} = \{i \in \{1, \dots, K\} | w_i > 0\}$) at most equal to $K_{max} \in \{1, \dots, K\}$. Two of the most commonly used tracking error measures are the root mean squared error (5) and the squared weighted difference between benchmark and portfolio weights (6), i.e.

$$TE_1(\mathbf{w}) = \sqrt{(\mathbf{y} - \mathbf{R}\mathbf{w})'(\mathbf{y} - \mathbf{R}\mathbf{w})/T} \quad (5)$$

or

$$TE_2(\mathbf{w}) = (\mathbf{w}\mathbf{b} - \mathbf{w})' \boldsymbol{\Sigma}(\mathbf{w}\mathbf{b} - \mathbf{w}), \quad (6)$$

where \mathbf{y} is the $T \times 1$ vector of the index returns, \mathbf{R} is the $T \times K$ matrix of the index constituents returns, $\boldsymbol{\Sigma}$ is the corresponding covariance matrix and $\mathbf{w}\mathbf{b} = [wb_1, \dots, wb_K]'$ is the vector of actual index portfolio weights. The tracking error measures (5) and (6) are convex, and classic optimisation tools (linear

or quadratic programming) can easily deal with them, but commonly fail when the cardinality constraint (4) is introduced (unless a linearisation of the problem is introduced, see, e.g., Canagkoz and Beasley (2008)). Search heuristics could then become a valid alternative. Furthermore, their application allows for the consideration of other real-world non-linear constraints such as the limited divisibility of asset shares, minimum transaction lots and transaction costs. An integer constraint ensures the minimum fraction of the investor's capital endowment V being invested in asset i with current price P_i is P_i/V , corresponding to one piece of asset i . This leads to discrete weights

$$w_i = \frac{n_i P_i}{V}, \quad n_i \in \mathbb{N}_0^+, \quad (7)$$

and as a result, it is likely that the investor's capital endowment cannot be entirely invested into assets; the remainder R will be held in cash

$$R = V - \sum_{i \in \mathcal{J}} n_i P_i. \quad (8)$$

Then the budget constraint (2) becomes

$$\sum_{i \in \mathcal{J}} w_i + \frac{R}{V} = 1. \quad (9)$$

We also notice that imposing only the no-short-selling (3) and budget (2 or 9) constraints determines lower and upper bounds on the weight vector's q -norm $\|\mathbf{w}\|_q = \left[\sum_{i=1}^N (w_i)^q \right]^{1/q}$ with $0 < q < 1$, which from now on we will simply refer to as q -norm:

$$1 < \|\mathbf{w}\|_q \leq K^{\frac{1}{q}-1}. \quad (10)$$

Then, the q -norm has maximum value for the most diversified portfolio which is the equally weighted one, while it has minimum value for the least diversified portfolio, which is the one totally invested in a single stock. The q -norm can then be interpreted, as suggested by Fernholz *et al.* (1998), as a diversity measure: the smaller the q -norm, the more sparse is the portfolio. Imposing the cardinality constraint (4) results in implicitly imposing tighter upper bounds on the q -norm, such that

$$1 < \|\mathbf{w}\|_q \leq K_{max}^{\frac{1}{q}-1}, \quad (11)$$

where clearly $K_{max}^{\frac{1}{q}-1} \leq K^{\frac{1}{q}-1}$. Hence, the tracking error problem could also be formulated by replacing the cardinality constraint (4) with a constraint that poses an upper bound on the q -norm, such as

$$\|\mathbf{w}\|_q \leq t, \quad (12)$$

where $t = K_{max}^{\frac{1}{q}-1}$. The optimisation problem is still NP-Hard given the presence of the non-convex constraint on the q -norm (12) and the integer constraint (7), and it can be re-written in terms of the Lagrangian:

$$\arg \min_{\mathbf{w}} TE(\mathbf{w}) + \lambda \|\mathbf{w}\|_q, \quad (13)$$

$$w_i = \frac{n_i P_i}{V}, \quad n_i \in \mathbb{N}_0^+ \quad (14)$$

$$R = V - \sum_{i=1}^K n_i P_i \quad (15)$$

$$\sum_{i=1}^K w_i + \frac{R}{V} = 1 \quad (16)$$

$$0 \leq w_i \leq 1 \quad (17)$$

where λ is inversely proportional to t .

Finally, when $q \rightarrow 0$, there exists a value of λ such that the optimal solution to the tracking error problem (13)-(17) is equal to the optimal solution of the optimisation problem (1), (3), (7)-(9) with the following cardinality equality constraint:

$$\#\mathcal{J} = K_{max}. \quad (18)$$

See Jansen and van Dijk (2002) for a formal proof.

3 Optimization

3.1 The main Challenges

From an optimisation viewpoint both the cardinality constrained problem (1), (3), (4), (7)-(9) and the q -norm constrained problem (13)-(17) are challenging due to the presence of local optima and discontinuities in the search space.

Figure 3.1 shows an example of the effect of a cardinality constraint on the search space. We consider a portfolio that consists of at most three assets out of a choice of four (i.e. the cardinality constraint is $\#\mathcal{J} \leq 3$). We draw all tracking errors for each of the four possible asset combinations in a three dimensional graph. In fact, the no-short-selling and the budget constraints assure that the weight combinations of three assets can be illustrated by triangular regions that map only two portfolio weights while the third (not shown) is determined by the budget constraint. These feasible regions build a common base area when arranged in the order shown in Figure 3.1. Point A corresponds to a portfolio that consists of assets 1, 3, and 4 with portfolio weights $w_1=0.2$, $w_3=0.3$, and $w_4 = 1 - w_1 - w_3 = 0.5$. Reducing w_1 to a value of zero, keeping w_3 constant, and consequently increasing portfolio weight w_4 to a value of 0.8, corresponds to a movement that ends in point B. There, the inactive asset 1 ($w_1=0$) can be replaced by the only asset left that is not already constituent of the portfolio, i.e., asset 2. Hence, the portfolio's cardinality (the number of active positions) does not exceed a value of three. Increasing w_2 by 0.4 and decreasing w_4 by the same amount gives point C (point D results analogously). Obviously, the exchange of an asset, which is represented by the transition of one triangular region to another, is discontinuous, since it requires a movement over an edge. Furthermore, each region has its own (local) minimum. For this example, the

identification of the portfolio that globally minimizes the tracking error subject to the cardinality constraint requires a four-fold application of a traditional optimisation technique.




Figure 1: Search space when integer constraints are neglected.

Figure 3.1 shows the tracking error of all portfolios that can be constructed with only three constituents of the Dow Jones Industrial Average 65 when integer constraints are not considered. All not depicted asset weights are determined by the budget constraint. Initial endowment: \$ 1 million.

Consequently, these techniques will quickly become unfeasible in real-world problems, which are usually characterised by high dimensionality.² Heuristics are then an attractive and viable solution for tackling the index tracking problem, as numerous studies have already shown (see, e.g., Gilli and Kellezi, 2002*b*; Beasley *et al.*, 2003; Maringer and Oyewumi, 2007; Krink *et al.*, 2009).

Even when the q -norm constraint is introduced to promote sparsity (instead of a cardinality constraint), traditional optimization techniques are still likely to fail due to the non-convexity of the resulting problem. Figure 2 shows the effect on the objective function of introducing a q -norm penalty. The top-left graph shows the value of the q -norm, with $q = 0.5$, of all feasible weight combinations: as expected, the maximum value is determined for the equally weighted portfolio (i.e.: $w_1 = w_2 = w_3 = 1/3$). The remaining three graphs show values of the objective function (13) for different values of λ , namely 0, 0.0025, and 0.005 respectively. When $\lambda = 0$ the objective function is convex (top right

²If we consider a portfolio that is made of at most 20 assets out of 100, we have approximately $7.07 \cdot 10^{20}$ applications of a traditional optimisation technique.

graph). The (global) minimum corresponds to projection A' in the solution space ($w_1 = 0.18, w_2 = 0.5, w_3 = 0.32$). When $\lambda > 0$ the objective function is non-convex. The bottom-left graph shows that increasing

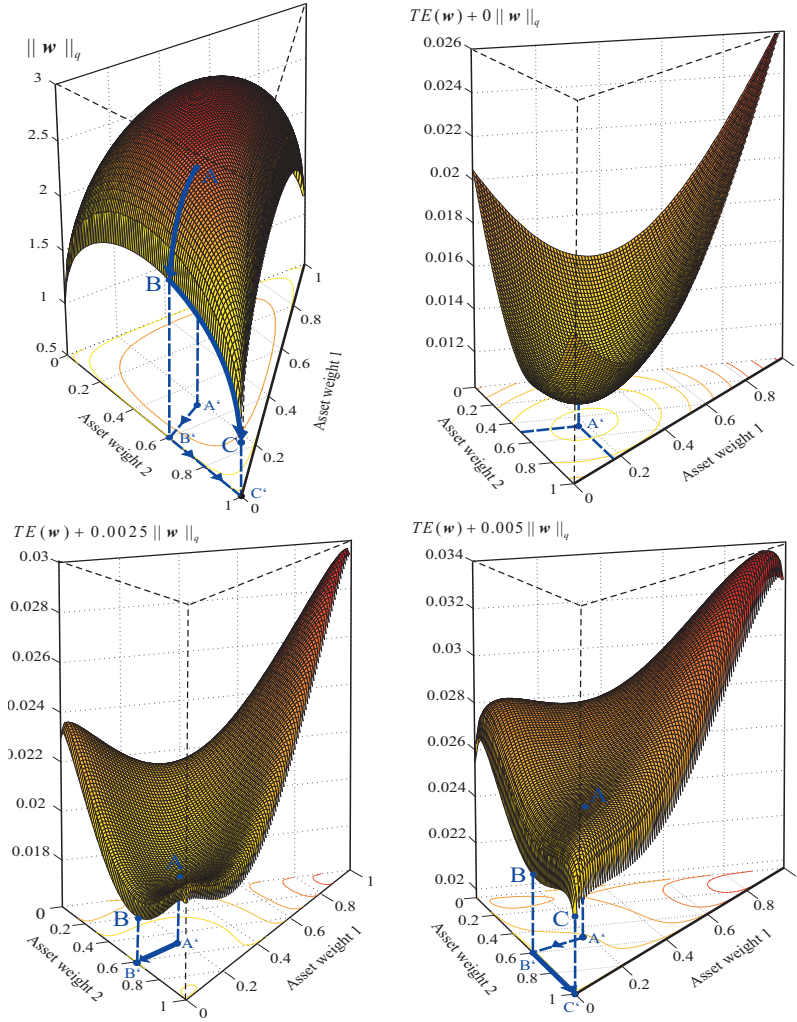


Figure 2: Mechanism of the q -norm penalty.

Figure 2 shows the q -norm (top-left) and objective function (13), where $q = 0.5$ and $\lambda = 0$ (top-right), $\lambda = 0.0025$ (bottom-left), and $\lambda = 0.005$ (bottom-right). The non-convexity of the q -norm constraint problem formulation as well as its promotion of sparsity when increasing λ is shown.

λ to a value of 0.0025 changes the objective function shape and its minimum is now B, which corresponds to projection B' in the solution space (w_1, w_2, w_3).

This is a sparse portfolio, since it consists of only two active positions ($w_1=0$). A is now only a local minimum. If λ is then set equal to 0.005, the tracking error is penalized more by the q -norm than in the previous cases and therefore sparsity increases such that only a single asset weight is active. This is shown in the bottom-right graph where minimum C corresponds to a portfolio that is fully invested in a single asset ($w_1=0$, $w_2=1$, $w_3=0$).

3.2 An Hybrid Heuristic Algorithm for Index Tracking

In this paper, we use the hybrid heuristic algorithm (HHA) proposed by Fastrich and Winker (2010), which is based on the hybrid local search algorithm introduced by Maringer and Kellerer (2003). The main differences between these two similar algorithms are related to, firstly, the embedded local search strategy, for which the HHA does not use Simulated Annealing (Kirkpatrick *et al.*, 1983) but its deterministic analogue Threshold Accepting (TA) (Dueck and Scheuer, 1990). Secondly, the HHA applies the TA-acceptance rule also within the evolutionary procedures in order to lessen the selective pressure on the population of solutions.

Algorithm 1 reports the pseudo code. In generation $g = 0$, the algorithm generates and evaluates a population of Pop random feasible solutions I_p^0 , $p = 1, \dots, Pop$. The solutions $\{I_p^0\}$ are referred to as search agents or individuals, and each individual encodes a feasible asset allocation. The population is then evolved by means of evolutionary operators. For each step size U_t , with $t = 1, \dots, thresh$, the population undergoes evolutionary operators for a number of $iter$ generations. The step size U_t , which linearly decreases from U_{max} to U_{min} by ΔU when t increases, can be interpreted as a fraction of the total capital endowment V that is subject to portfolio adjustments between two evolutionary steps. More precisely, these modifications, which are conducted within the

Algorithm 1 Hybrid Heuristic Algorithm.

- 1: Initialize Pop , $thresh$, $iter$, U_{min} , U_{max} and $\Delta U = (U_{max} - U_{min})/thresh$
 - 2: Generate a valid initial ($g = 0$) population of random solutions $\{I_p^0\}$, $p = 1, \dots, Pop$
 - 3: Evaluate the objective function for $I_p^0 \forall p$
 - 4: **for** $t = 1$ to $thresh$ **do**
 - 5: Determine the step size $U_t = U_{max} - \Delta U \cdot (t - 1)$
 - 6: **for** $l = 1$ to $iter$ **do**
 - 7: $g = (t - 1) \cdot iter + l$
 - 8: **Modification Phase**
 - 9: **Evaluation Phase**
 - 10: **Replacement Phase**
 - 11: **end for**
 - 12: **end for**
 - 13: terminate algorithm and report the best solution found
-

Modification Phase, independently adjust the individuals in a componentwise manner (i.e.: one asset's position is reduced while that of another is increased, see for example (Dueck and Winker, 1992; Gilli and Kellezi, 2002a)). After the *Modification Phase*, the population is then evaluated (*Evaluation Phase*) by

ranking the individuals according to their objective function values. Finally, in the *Replacement Phase* the worst individuals are replaced either by so-called *Clones*, i.e., exact copies of the current population’s best individuals, or by so-called *Averaged Idols*, i.e., individuals that combine assets which have been proven to be successful in other portfolios. More precisely, an *Averaged Idol* draws assets not from the whole asset universe, but of an (regularly updated) subset, which contains only those assets that were held by individuals with good objective function values. This arrangement of the phases helps to, firstly, find a successful combination of assets, i.e., a core structure, in earlier generations, before it, secondly, contributes to assigning proper weights to this core structure’s assets. The reader is referred to Fastrich and Winker (2010) for a detailed description and a parameter tuning analysis of the HHA.

4 Experimental Set-Up

We consider the daily log-returns of three stock market indices and their constituents, namely the German DAX 100 (Period: 17.03.05 to 06.03.08), the Dow Jones Industrial Average 65 (DJ 65, Period: 12.04.2002 to 19.12.2003) and the Standard & Poor’s 500 (S&P 500, Period: 20.09.2006 to 01.10.2008). Table 1 reports the summary statistics of the daily log-returns of the indices. The three return times series exhibit the typical patterns of financial times series: mean values around zeros, light asymmetry but fat tails.

Table 1: Summary Statistics of the indices daily log-returns

	number of sample		standard				
	constituents	size	mean	deviation	skewness	min	max
DAX 100	98	750	0.0471	1.0372	-0.7120	-7.4636	5.6375
DJ 65	65	440	-0.0029	1.3402	0.1917	-4.6097	5.3455
S&P 500	491	530	-0.0249	1.2439	-0.8724	-9.2002	5.2785

In all the experiments, we assume an initial endowment of one million Euros for the DAX 100 or Dollars for the DJ 65 and the S&P 500 and determine index tracking investment strategies using a moving time window procedure. In particular, we assume that the optimal tracking portfolio is determined by using a window of 250 observations and held unchanged for the subsequent 21 out-of-sample trading days. Then, the (in-sample) window is moved forward by 21 days and the new optimal tracking portfolio is determined using a window of 250 observations and again held unchanged for the next 21 out-of-sample days and so forth. Consequently, tracking portfolios are revised once a month.

5 A comparison of the two approaches on real-world data

The first step of our empirical analysis aims to compare the cardinality and the q -norm constraint approach. Using a rolling window scheme and initial

endowment, as described in Section 4, we compare the performances of the two approaches over the whole period. In particular, Table 2 and 3 report the results when setting $q=0.5$ and λ equal to $1.5 \cdot 10^{-4}$, $0.5 \cdot 10^{-4}$, and $0.22 \cdot 10^{-4}$ in order to determine tracking portfolios with different sizes (with respectively 8-13, 17-25 and 28-42 active weights). The maximum number of assets K_{max} in the cardinality constraint is then set equal to the number of active weights (positive weights) found with the q -norm approach in the different windows. Then, given that $q=0.5$, the upper bound for the q -norm $K_{max}^{\frac{1-q}{q}}$ is equal to K_{max} . Clearly, this could imply a less tight constraint on the q -norm for the cardinality constraint approach. To explain this, let's consider a simple example with only three assets and set $q=0.5$. The q -norm surface for all feasible portfolios is represented in Figure 2 (top-left): the minimum values correspond to the cases when only one weight is active ($\|\mathbf{w}\|_q = 1$), while the maximum values correspond to the equally weighted portfolio ($\|\mathbf{w}\|_q = 3$). As Fernholz *et al.* (1998) report, the q -norm is also a measure of investment diversification. Then, since we know that there is an inverse relationship between λ and t , which cannot be easily derived, it could happen that we fix λ such that only two weights are active but which does not exactly corresponds to solve the optimisation problem with $t = 2$ but for a value of t such that $1 < t \leq 2$. Then, not all feasible portfolios with two active assets are included in the search space. As a consequence, the feasible search space for the q -norm formulation is smaller than for the cardinality constraint, when setting K_{max} equal to 2. This appears evident when looking at the q -norm values in Table 3: the minimum, mean and maximum q -norm values of the optimal portfolios are always larger for the cardinality approach, revealing a small comparative advantage in terms of diversification for the cardinality constraint approach.

Table 2 reports the in-sample and out-of-sample average tracking error, excess return, correlation and turnover, while Table 3 reports the skewness, kurtosis, value-at-risk of the out-of-sample return time series and the minimum, mean and maximum value of the q -norm for $q=0.5$. As Table 2 shows, the differences in the in-sample and out-of sample average tracking error volatilities among the two approaches are - with one exception - not statistically significant, while the average excess returns are statistically different in seven cases. Among them, five suggest that the q -norm approach provide tracking portfolios with better performances. In contrast, the two approaches lead to identify equivalent portfolios not only in terms of tracking error volatilities but also with respect to the average turnover and the correlation with the benchmark, although the correlation appears to be slightly higher for the q -norm approach for the DJ 65 and the S&P 500 data sets. Furthermore, looking at Table 3, we also notice that although the cardinality constraints portfolios are more diversified, as the q -norms show, the values of skewness, kurtosis and Value at Risk at 99% level are very similar. Finally, when comparing the portfolio compositions, as Figure 3 shows, we find further support of the fact that the two approaches lead to identify equivalent portfolios also in terms of asset composition. Differences between asset weights are, as Figure 3 shows, often negligible.

Table 2: Tracking Error, Excess returns, Correlation and Turnover

	in-sample cardin q -norm	out-of-sample cardin q -norm	in-sample cardin q -norm	out-of-sample cardin q -norm	in-sample cardin q -norm	out-of-sample cardin q -norm	in-sample cardin q -norm	out-of-sample cardin q -norm	in-sample cardin q -norm	out-of-sample cardin q -norm	in-sample cardin q -norm	out-of-sample cardin q -norm
Panel A: DAX 100												
	9 to 12 stocks ($\emptyset=10,25$)				18 to 24 stocks ($\emptyset=21,42$)				31 to 37 stocks ($\emptyset=34,29$)			
λ	$1.5 \cdot 10^{-4}$				$0.5 \cdot 10^{-4}$				$0.22 \cdot 10^{-4}$			
<i>annualized tracking error volatility in percent</i>												
mean	2,43	2,46	3,24	3,27	1,15	1,17	1,68	1,68	0,66	0,67	1,02	1,07
std	0,20	0,20	0,75	0,77	0,12	0,12	0,45	0,46	0,08	0,09	0,23	0,23
t_{diff}	0,49		0,16		0,49		0,04		0,63		0,74	
<i>annualized excess return over the benchmark in percent</i>												
mean	-1,64	-1,28	-5,75	-5,75	0,39	0,15	-1,86	-1,54	0,19	0,12	-2,09	-2,44
std	0,52	0,49	2,55	2,91	0,33	0,36	1,48	1,68	0,18	0,16	0,79	0,83
t_{diff}	2,51**		0,01		2,43**		0,68		1,33		1,46	
<i>corr</i>	0,980	0,984	0,989	0,986	0,992	0,991	0,997	0,996	0,998	0,998	0,999	0,999
\emptyset - <i>t.o.</i>	0,38	0,35			0,26	0,27			0,18	0,20		
Panel B: DJ 65												
	8 to 13 stocks ($\emptyset=10,70$)				19 to 25 stocks ($\emptyset=22,30$)				30 to 40 stocks ($\emptyset=35,40$)			
λ	$1.5 \cdot 10^{-4}$				$0.5 \cdot 10^{-4}$				$0.22 \cdot 10^{-4}$			
<i>annualized tracking error volatility in percent</i>												
mean	3,76	3,80	3,98	3,71	2,34	2,37	2,91	3,09	1,74	1,73	2,29	2,43
std	0,08	0,09	0,21	0,11	0,06	0,05	0,19	0,26	0,03	0,03	0,21	0,26
t_{diff}	0,35		1,13		0,33		0,56		0,13		0,42	
<i>annualized excess return over the benchmark in percent</i>												
mean	2,34	2,33	-2,09	-2,17	1,74	2,07	3,55	4,99	1,47	1,69	-0,62	2,47
std	0,54	0,47	1,66	1,66	0,13	0,11	0,89	0,74	0,08	0,06	0,56	0,56
t_{diff}	0,02		0,04		1,93*		1,25		2,27**		3,90***	
<i>corr</i>	0,958	0,968	0,645	0,648	0,997	0,998	0,870	0,934	0,999	0,999	0,963	0,964
\emptyset - <i>t.o.</i>	0,82	0,72			0,64	0,63			0,34	0,36		
Panel C: S&P 500												
	9 to 13 stocks ($\emptyset=10,14$)				17 to 25 stocks ($\emptyset=20,71$)				28 to 42 stocks ($\emptyset=33,71$)			
λ	$1.5 \cdot 10^{-4}$				$0.5 \cdot 10^{-4}$				$0.22 \cdot 10^{-4}$			
<i>annualized tracking error volatility in percent</i>												
mean	3,14	3,14	5,90	5,66	1,88	1,87	4,31	4,67	1,60	1,30	3,86	3,82
std	0,09	0,08	0,55	0,40	0,06	0,05	0,29	0,39	0,05	0,04	0,23	0,31
t_{diff}	0,03		0,35		0,11		0,75		5,00***		0,12	
<i>annualized excess return over the benchmark in percent</i>												
mean	1,77	1,73	2,86	0,99	1,13	0,94	-3,80	-3,76	0,96	1,37	-0,63	-0,86
std	0,19	0,19	1,86	1,63	0,10	0,13	1,50	1,34	0,02	0,02	0,36	0,23
t_{diff}	0,17		0,76		1,14		0,02		13,04***		3,53***	
<i>corr</i>	0,973	0,975	0,886	0,892	0,993	0,987	0,891	0,913	0,994	0,994	0,961	0,969
\emptyset - <i>t.o.</i>	1,36	1,28			1,38	1,45			1,38	1,39		

Table 2 summarizes the results of both the cardinality constraint (abbreviated as cardin) and the q -norm constraint approach (where $q=0,5$) for all data sets. Each benchmark is replicated with three differently sized tracking portfolios headed with their number of constituents (average in parenthesis). A statistically significant difference with confidence level 10% (5%, 1%) between the two approaches is marked with one (two, three) asterisk(s) at the variable t_{diff} , which shows the t -statistic of the difference between the two approaches. The correlation coefficient between the returns of the tracking portfolio and the index is given by *corr* and the averaged turnover is given by \emptyset -*t.o.*

Table 3: Extreme statistics and q -norms

	cardinality	q -norm	cardinality	q -norm	cardinality	q -norm
Panel A: DAX 100						
	9 to 12 stocks ($\emptyset=10,25$)		18 to 24 stocks ($\emptyset=21,42$)		31 to 37 stocks ($\emptyset=34,29$)	
λ	$1,5 \cdot 10^{-4}$		$0,5 \cdot 10^{-4}$		$0,22 \cdot 10^{-4}$	
	<i>distribution of daily (out-of-sample) portfolio returns</i>					
skew	-0,6432	-0,6445	-0,7517	-0,7759	-0,7690	-0,7877
kurt	8,4661	8,4456	8,3175	8,4352	8,4150	8,3638
VaR_{99}	-0,0297	-0,0310	-0,0288	-0,0291	-0,0307	-0,0308
	<i>portfolio weight's q-norm</i>					
min	8,68	8,45	16,63	16,47	27,30	26,78
mean	9,91	9,71	19,80	19,36	29,58	28,81
max	11,35	11,23	22,07	21,81	31,36	30,89
Panel B: DJ 65						
	8 to 13 stocks ($\emptyset=10,70$)		19 to 25 stocks ($\emptyset=22,30$)		30 to 40 stocks ($\emptyset=35,40$)	
λ	$1,5 \cdot 10^{-4}$		$0,5 \cdot 10^{-4}$		$0,22 \cdot 10^{-4}$	
	<i>distribution of daily (out-of-sample) portfolio returns</i>					
skew	0,1446	0,1432	0,3124	0,2875	0,2141	0,2397
kurt	3,1407	3,0513	3,1106	3,0325	3,0978	3,1346
VaR_{99}	-0,0146	-0,0146	-0,0141	-0,0142	-0,0152	-0,0149
	<i>portfolio weight's q-norm</i>					
min	7,70	7,62	17,73	17,60	28,28	27,17
mean	10,32	9,99	21,20	20,55	33,14	32,14
max	12,34	12,00	23,78	23,56	37,59	36,33
Panel C: S&P 500						
	9 to 13 stocks ($\emptyset=10,25$)		17 to 25 stocks ($\emptyset=21,42$)		28 to 42 stocks ($\emptyset=33,71$)	
λ	$1,5 \cdot 10^{-4}$		$0,5 \cdot 10^{-4}$		$0,22 \cdot 10^{-4}$	
	<i>distribution of daily (out-of-sample) portfolio returns</i>					
skew	0,0445	-0,0476	0,0024	0,0009	-0,0844	-0,0309
kurt	4,2280	4,1603	3,9763	4,5136	3,2023	4,3820
VaR_{99}	-0,0349	-0,0349	-0,0343	-0,0353	-0,0283	-0,0283
	<i>portfolio weight's q-norm</i>					
min	8,60	8,35	15,86	15,52	25,67	25,04
mean	9,74	9,40	19,73	19,23	31,17	30,90
max	12,09	11,57	24,12	23,13	39,43	37,18

Table 3 shows the skewness, the kurtosis, and the 99-percent value at risk (VaR_{99}) of the out sample return-distribution of the portfolios (based on a daily frequency).




Figure 3: Differences in Portfolio compositions

Differences in weights of (small) portfolios; DAX 100. The scaling of the vertical axis is determined by the maximum portfolio weights (q -norm: 0.21 and cardinality: 0.19).

6 Selecting the Optimal Number of Active Weights

One of the most challenging questions a passive asset manager has to face is how to determine the optimal number of active weights to track a given index or benchmark. In fact, although having a large number of active weights can lead to a better tracking performance, the relationship between the number of active positions and the in-sample tracking error is not linearly decreasing but tends to have a hyperbolic shape: a sharpe decrease at the beginning and then a slight improvement for each newly added asset after enough assets are included in the tracking portfolio. Empirical results also show that out-of-sample tracking errors could be larger for a larger number of assets (Krink *et al.*, 2009). A possibility of choosing the best combination of q and λ and therefore to determine the optimal size of the tracking portfolio in the q -norm approach, is to use cross-validation. Section 6.1 describes in an ideal world with no transaction costs one approach to select the optimal tracking portfolio size and composition. Clearly, the assumption of no transaction costs is unrealistic. In fact, the choice of how many active positions to have is strongly influenced by the cost of implementing the trading strategy, which ideally should be as cheap as possible. Then, section 6.2 describes an approach to determine the number of active weights when taking into account transaction costs.

6.1 Ideal world: No Transaction Costs

As a first step for the cross-validation, we define a grid of 400 equally spaced nodes (q - λ -combinations), with $q \in [0.3, 0.7]$ and $\lambda \in [1 \cdot 10^{-8}, 2 \cdot 10^{-8}]$. Then, we estimate the optimal tracking portfolio on a window of 250 observations (in-

sample) and evaluate its performance on a window of 21 observations (out-of sample). In addition to the real sample we block bootstrap (block length 21) 49 further data samples (Efron and Tibshirani, 1993): 250 observations corresponds to the in-sample and 21 to the out-of-sample interval. For each sample and each q - λ -combination we estimate the optimal tracking portfolio. Figure 4 shows the

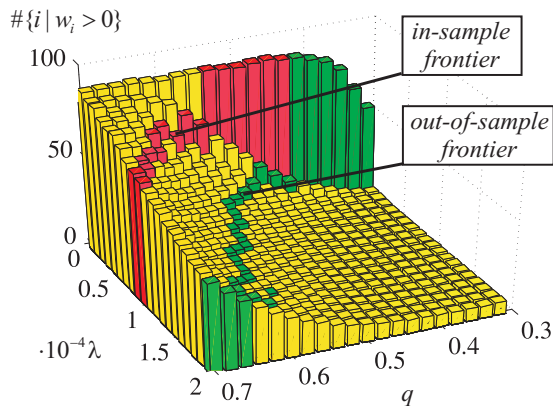


Figure 4: Number of active positions.

Figure 4 shows the average number of active positions in correspondence to the q - λ -combinations. Combinations that on average do not differ statistically significantly from the best solution found, in terms of the (in-sample vs. out-of-sample) tracking error lie on or are in the subset to the left of the depicted frontiers.

average number of active weights for each grid node. As expected, the portfolio size decreases for small values of q and large values of λ . Small values of q promote sparsity and large values of λ increase the size of the q -norm penalty in the objective function. The largest portfolios are identified for the smallest value of λ , independent of the value of q . In fact, in such cases, the penalty is basically equal to zero and there is no promotion of sparsity, since it is alike a situation with no q -norm constraint. Given λ , the optimal portfolios shrink in size when q decreases. Figure 5 shows the mean values of the in-sample (top-left) and out-of-sample (bottom-left) tracking errors and the mean values of the in-sample (top-right) and out-of-sample (bottom-right) excess returns. We also draw two frontiers, one for the in-sample and the other for the out-of-sample results, to identify the portfolios that are statistically equivalent in mean values to the optimal in-sample and out-of-sample portfolios (marked with stem and filled dot), respectively. The term “statistically equivalent” is used in cases in which the t -value of the difference between the means of the 50 samples is greater than two. The two frontiers are also plotted in in Figure 4. As Figure 5 shows, the in-sample and out-of-sample tracking errors increase when q decreases and λ increases, that is when the portfolio sizes shrink. The out-of-sample frontier is




Figure 5: Bootstrapped surfaces without costs.

Figure 5 shows the mean values of the in-sample (top-left) and out-of-sample (bottom-left) tracking errors as well as the mean values of the in-sample (top-right) and out-of-sample (bottom-right) excess returns. In-sample and out-of sample best solutions are marked with stem and filled dot. Combinations that on average do not differ statistically significantly from the best solution lie on or are in the subset to the left of the frontiers. In the excess return graphics those combinations are marked with a point, for which the null hypothesis of an expected value of zero could not be rejected.

quite far from the in-sample one. In fact, the number of in-sample statistically equivalent combinations is equal to 50 out of 400 and the average number of active positions in this subset varies between a minimum of 64.86 to a maximum of 90.843 with mean and median values equal to 80.78 and 81.88, respectively. The situation changes drastically in terms of portfolios sizes when we look at the out-of-sample statistically equivalent portfolios: the larger variability in the results lead to have 151 out of 400 equivalent combinations and the average portfolio size varies from a minimum value of 24.51 to a maximum value of 90.84 with mean and median values equal to 52.77 and 56.57, respectively. The best in-sample q - λ -combination, which leads to identify portfolios with on average 84.49 active weights, ranks 97th if we sort the optimal portfolios with respect to the out-of-sample tracking error, while the best out-of sample combination, which lead to identify portfolios with on average 35.63 active weights, ranks 124th when sorting with respect to the in-sample tracking error. A nice analysis regarding the in- and out-of-sample ranks of optimised portfolios is provided by Gilli and Schumann (forthcoming). Hence, the best in-sample portfolio can still be considered as statistically equivalent to the best out-of-sample one, but not vice-versa. This could result in discarding smaller portfolios, although out-of-sample they show good tracking performances and the average size of out-of-sample statistically equivalent portfolios is much smaller than the in-sample equivalent ones. Furthermore, the empirical results show that also excess returns exhibit a large variability when considering smaller portfolios, which could result in small but also in large excess returns.

6.2 Real world: Transaction Costs

As the previous section has shown, the out-of-sample optimal portfolios tend to be smaller in size than the in-sample optimally chosen ones, which makes them a more attractive investment opportunity. Clearly, this is even more important when considering a real-world set-up, where managers have to take into account transaction, administrative and monitoring costs. In this section, we show that introducing a simple but yet realistic way of modeling costs, could allow to even better tackle the problem of selecting the optimal size and composition of a tracking portfolio. The transaction, administrative and monitoring costs are modeled by a fixed payment $c_{f,i}$ that is due whenever stock i is purchased or sold.³ In addition, financial institutions charge investors for the execution of orders. We assume these fees to grow proportionally with the transaction volume by the factor $c_{p,i}$. Thus, we introduce the following cost function:

$$C = \sum_{j=\{i:i=1,\dots,K|0<|\Delta n_i|\}} c_{f,j} + c_{p,j}|\Delta n_j|P_j. \quad (19)$$

Costs occur with a monthly frequency whenever a stock is subject to portfolio rebalancing. We assume that the index tracking strategy must not exceed a cost

³Of course, further cost categories, such as regular payments for the custody account, can be seen as included in this fixed cost term.

limit of C_{max} . If the rebalancing costs exceed this limit, we add a penalty to the optimised tracking error. The penalty is not taken into account during the optimisation but added ex-post to not impose a further constraint which could at this stage narrow too much the feasible search space. The introduction of a cost constraint during the optimisation is left for further research.

Since transaction costs arise by changes from holding one portfolio to holding another (updated) portfolio, we optimise an initial portfolio with the first 250 observations of our data set. The performance of this portfolio is not part of the analysis; it only serves as a starting point to compute the transaction costs that are caused by the transition to all sample's portfolios.⁴ As in the previous section, we block bootstrap such that we obtain 50 samples with an in-sample and out-of-sample window of 250 and 21 observations, respectively.

Figure 6 shows the average number of active positions for each q - λ -combination for the DAX 100 data set and $c_{f,i}=20$, $c_{p,i}=0.0025$, and $C_{max}=2,500$. Darker shaded bars point out the optimal portfolios which are statistically equivalent




Figure 6: Number of active positions.

Figure 6 shows the average number of active positions dependent on the λ - q -combination. Combinations that on average do not differ statistically significantly from the best solution found, in terms of the (in-sample vs. out-of-sample) tracking error, are marked red and green. The combinations that are marked blue do not differ significantly in both cases, in-sample as well as out-of-sample.

in-sample, out-of-sample, or both. It appears evident that the optimal portfolios tend to have smaller sizes when considering transaction costs and only seven of the optimal out-of-sample portfolios which are statistically equivalent are not also considered statistically equivalent to the optimal one in-sample. The number of in-sample and out-of-sample statistically equivalent combinations is equal

⁴ One reason for this setup is to avoid the complications caused by the high transaction costs when constructing a portfolio at time 0.

to 61 and 40 out of 400, respectively. Compared to the analysis without cost, the number of out-of-sample equivalent combinations is much smaller (40 instead of 124). Furthermore, the average number of in-sample and out-of-sample active positions are now similar: in-sample the average portfolio size varies between a minimum of 25.56 to a maximum of 54.46 with mean and medians values equal to 36.74 and 36.83, out-of-sample between a minimum of 14.44 to a maximum of 43.84 with mean and medians values equal to 31.48 and 31.841. Furthermore, as Figure 7 shows, the tracking error surfaces are now very different than in the

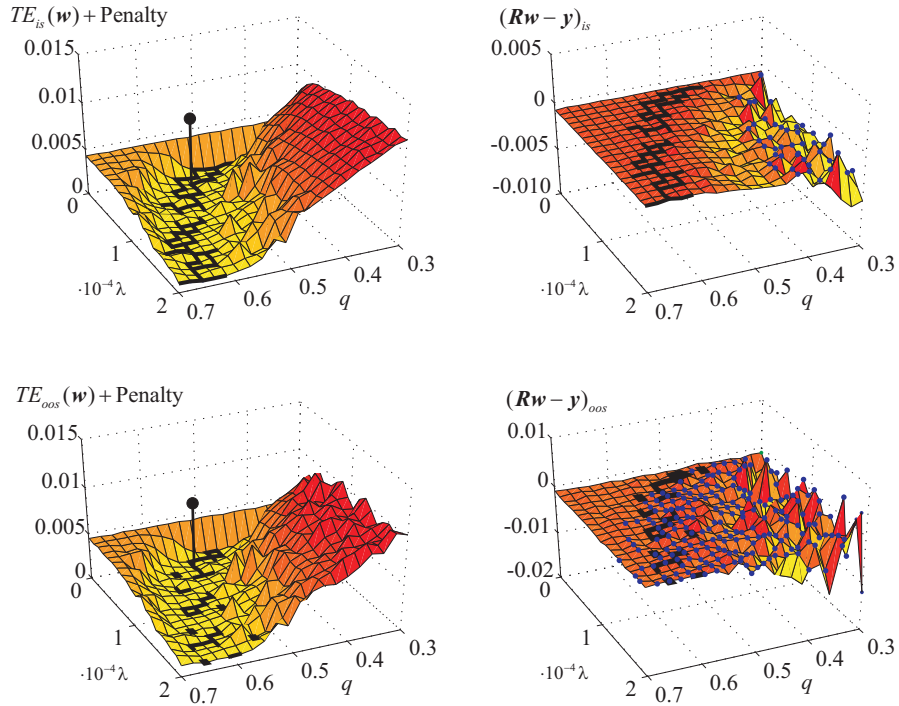


Figure 7: Bootstrapped surfaces with costs.

Figure 7 shows the mean values of the in-sample (top-left) and out-of-sample (bottom-left) tracking error, which is penalized with $+0.004$, if the rebalancing costs exceed $C_{max} = 2.500$. In the two tracking error graphics the best solution found is marked with stem and filled dot. Combinations that on average are not statistically significantly different from the best solution lie on the frontier or are enveloped by it. The figure also shows the mean values of the in-sample (top-right) and out-of-sample (bottom-right) excess return over the benchmark, where those combinations are marked with a point, for which the null hypothesis of an expected value of zero could not be rejected.

no-costs set-up. Introducing costs leads to strong preferences for more parsimonious portfolios: the portfolios identified in correspondence to small values of λ and large values of q are not any longer the optimal investment opportuni-

ties due to the excessive costs and again the portfolios with a too low number of active positions, identified for small values of q , do not allow to optimally track the index performance. The area of the optimal portfolios is identified in the valley. There is almost an one-to-one correspondence between in-sample and out-of-sample statistically equivalent portfolios. The presence of a large out-of-sample variability of the tracking error leads to identify a smaller set of statistically equivalent portfolios with smaller sizes than in the previous case without costs. Introducing a realistic but yet simple way of modelling costs can then allow to identify smaller optimal portfolios that exhibit nice properties both in-sample and out-of-sample.

7 Discussion and Further Research

Index tracking aims to replicate the performance of an index by using a small number of constituents. Ideally, the tracking investment strategy should be cheap to implement and should have not only good in-sample properties but also out-of-sample ones. The problem can be formulated as an optimisation problem: minimize a given distance measure between the tracking portfolio and the index such that at most K constituents have active weights. This approach is known in the literature as the cardinality constraint approach. The optimisation problem is NP-hard and it becomes extremely challenging when the problem size is large. In this paper, we introduce a formulation of the problem by imposing a constraint on the q -norm ($0 < q < 1$) of the portfolio weights. Such regularisation techniques are well-established in the statistical community and have recently gained much attention as a way for simultaneous model selection and estimation, especially in linear regression models. In fact, the introduction of the q -norm penalty allows to promote sparsity and therefore to select and estimate only few non-zero coefficients for a subset of explanatory variables. Furthermore, imposing a constraint on the q -norm of portfolio weights is a natural way for controlling diversification in the tracking portfolio: it can be easily shown that, given the budget constraint and the positivity constraint of the asset weights, the larger the q -norm the more diversified is the portfolio.

The q -norm approach could then provide a better way of tackling the index tracking problem for different reasons. First, we have shown, using three different data sets, that the two approaches lead to equivalent results in terms of tracking errors and portfolio size when considering almost equivalent set-ups, but the q -norm approach seems to provide on average portfolios with larger excess returns and out-of-sample correlations. We find then evidence that imposing the q -norm constraint allows to regularise the index tracking problem, determining in one single step the number of active weights and their optimal values, and providing models with good out-of-sample performances. Other recent studies (e.g., DeMiguel *et al.*, 2008; Brodie *et al.*, 2008) have found similar findings in a mean-variance context. Further research on different data sets and data frequency is currently high on our agenda.

Second, we have proposed a simple yet effective method to use the q -norm

approach to solve the problem of selecting the optimal portfolio size by using cross-validation not only in an ideal world, without transaction costs, but also in a real-world context, where costs are considered. The empirical results show that the q -norm approach combined with an easy but yet realistic way of modeling transaction costs could allow to identify a portfolio with relatively small size and attractive in-sample and out-of-sample tracking performances.

Third, the q -norm approach could provide an easier way of simplifying the optimisation problem or approximating the non-convex constraint on the q -norm (e.g., Coleman *et al.*, 2006) rather than the cardinality constraint and then develop more efficient optimisation methods. In this work, we have proposed to use a search heuristics for both optimisation problems but comparison with other optimisation algorithms is currently under investigation.

References

- Beasley, J.E., N. Meade and T.-J. Chang (2003). An evolutionary heuristic for the index tracking problem. *European Journal of Operational Research* **148**, 621–643.
- Brodie, J., I. Daubechies, C. DeMol, D. Giannone and D. Loris (2008). Sparse and stable markowitz portfolios. *ECB Working Paper Series*.
- Canagkoz, 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.
- Coleman, T.F., Y. Li and J. Henniger (2006). Minimizing tracking error while restricting the number of assets. *Journal of Risk* **8**, 33–56.
- DeMiguel, V., L. Garlappi, J. Nogales and R. Uppal (2008). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Working Paper*.
- Derigs, U. and N.-H. Nickel (2003). Meta-heuristic based decision support for portfolio optimization with a case study on tracking error minimization in passive portfolio management. *OR Spectrum* **25**, 345–378.
- Dueck, G. and P. Winker (1992). New concepts and algorithms for portfolio choice. *Applied Stochastic Models and Data Analysis* **8**, 159–178.
- Dueck, G. and T. Scheuer (1990). Threshold accepting: A general purpose optimization algorithm. *Journal of Computational Physics* **90**, 161–175.
- Efron, B. and R.J. Tibshirani (1993). *An Introduction to the Bootstrap*. Chapman & Hall.
- Fastrich, B. and P. Winker (2010). Robust portfolio optimization with a hybrid heuristic algorithm. *COMISEF Working Paper Series forthcoming* pp. 1–28. available from http://comisef.eu/?q=working_papers.

- Fernholz, R., R. Garvy and J. Hannon (1998). Diversity weighted indexing. *Journal of Portfolio Management* **24**, 83–92.
- Frank, I.E. and J.H. Friendman (1993). A statistical view of some chemometrics regression tools. *Technometrics* **35**, 109–135.
- Gilli, M. and E. Kellezi (2002a). A global optimization heuristic for portfolio choice with var and expected shortfall. *E. Kontogiorgos, B. Rustem & S. Siokos eds. "Computational Methods in Decision-making, Economics and Finance"*, Kluwer pp. 165–181.
- Gilli, M. and E. Kellezi (2002b). *The Threshold Accepting Heuristic for Index Tracking*. Financial Engineering, E-Commerce, and Supply Chain, Kluwer Applied Optimization Series.
- Gilli, M. and E. Schumann (forthcoming). Optimal Enough?. *Journal of Heuristics*.
- Gilli, M. and P. Winker (2009). Review of heuristic optimization methods in econometrics. *in: Handbook of Computational Econometrics, Ed. D. Besley and E. Kontogiorgos. Wiley. Chichester* pp. 81–120.
- Hoerl, A.E. and R.W. Kennard (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* **8**, 27–51.
- Jansen, R. and R. van Dijk (2002). Optimal benchmark tracking with small portfolios. *The Journal of Portfolio Management* **28**, 33–39.
- Kirkpatrick, S., C.D. Gelatt and M.P. Vecchi (1983). Optimization by simulated annealing. *Science* **220**, 671–680.
- Krink, T., S. Mittnik and S. Paterlini (2009). Differential evolution and combinatorial search for constrained index tracking. *Annals of Operations Research* **172**, 153–176.
- Maringer, D. and H. Kellerer (2003). Optimization of cardinality constrained portfolios with a hybrid local search algorithm. *OR Spectrum* **25**, 481–495.
- Maringer, D. and O. Oyewumi (2007). Index tracking with constrained portfolios. *Intelligent Systems in Accounting, Finance and Management* **15**, 57–71.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Royal Statistical Society* **58**, 267–288.

RECent Working Papers Series

The 10 most RECent releases are:

- No. 56 CARDINALITY VERSUS Q-NORM CONSTRAINTS FOR INDEX TRACKING (2011)
B. Fastrich, S. Paterlini and P. Winker
- No. 55 THE ATTRACTIVENESS OF COUNTRIES FOR FDI. A FUZZY APPROACH (2010)
M. Murat and T. Pirotti
- No. 54 IMMIGRANTS, SCHOOLING AND BACKGROUND. CROSS-COUNTRY EVIDENCE FROM PISA 2006 (2010)
M. Murat, D. Ferrari, P. Frederic and G. Pirani
- No. 53 AGRICULTURAL INSTITUTIONS, INDUSTRIALIZATION AND GROWTH: THE CASE OF NEW ZEALAND AND URUGUAY IN 1870-1940 (2010)
J. Álvarez, E. Bilancini, S. D'Alessandro and G. Porcile
- No. 52 ON THE RATIONALIZABILITY OF OBSERVED CONSUMERS' CHOICES WHEN PREFERENCES DEPEND ON BUDGET SETS AND (POTENTIALLY) ON ANYTHING ELSE (2010)
E. Bilancini
- No. 51 SLAVERY, EDUCATION, AND INEQUALITY (2010)
G. Bertocchi and A. Dimico
- No. 50 A DYNAMIC POLITICO-ECONOMIC MODEL OF INTERGENERATIONAL CONTRACTS (2010)
F. Lancia and A. Russo
- No. 49 FISCAL FORESIGHT AND THE EFFECTS OF GOVERNMENT SPENDING (2010)
M. Forni and L. Gambetti
- No. 48 THE HUB CONTINENT? IMMIGRANT NETWORKS, EMIGRANT DIASPORAS AND FDI (2010)
S. Flisi and M. Murat
- No. 47 TESTING FOR CONTAGION: A TIME-SCALE DECOMPOSITION (2010)
A. Cipollini and I. Lo Cascio

The full list of available working papers, together with their electronic versions, can be found on the RECent website: <http://www.recent.unimore.it/workingpapers.asp>