

POLYTECHNIC OF TURIN



Master's Degree in Mathematical Engineering

Robust optimization for enhanced index tracking problem

Supervisors

Prof. Edoardo FADDA

Prof. Daniele MANERBA

Candidate

Andrea PRINCIPE

October 2025

*To my family,
for their constant love
and support*

Abstract

This thesis addresses the Enhanced Index Tracking Problem (EITP), which aims to construct portfolios that outperform a market benchmark while containing the additional risk introduced by active deviations. The models implemented in this work follow existing formulations from the literature; the thesis focuses on adapting and evaluating these approaches in a large-scale setting.

One approach applies robust optimization within factor models based on Fama–French specifications. The study implements both the three- and five-factor variants on the S&P-500 universe, introduces uncertainty sets on expected returns and on the factor loading matrix to account for estimation errors, and enforces cardinality constraints to control portfolio sparsity. Under the chosen formulation, the factor-based model leads to a Mixed-Integer Second-Order Cone Program (MISOCP).

A second approach implements a mixture-based model within the Lower Partial Moments (LPM) framework: asset returns are represented by a Gaussian mixture and uncertainty in mixture weights is modeled through ϕ -divergences. Lagrange duality is used to obtain a tractable formulation for numerical solution.

Empirical evaluation comprises static out-of-sample tests for both implementations and a rolling-window (dynamic) out-of-sample analysis for the factor-based implementation; dynamic evaluation was not feasible for the mixture model due to its computational profile. Results indicate that robust specifications can enhance stability and risk-adjusted performance relative to standard EITP benchmarks. Overall, the thesis provides an empirical assessment of robust EITP methodologies applied to a large-scale index and offers practical insights for portfolio construction.

Keywords: Enhanced Index Tracking, Lower Partial Model, Mixture Model, Robust Optimization, ϕ -divergence, Factor Model, Uncertainty

Table of Contents

List of Tables	VI
List of Figures	VII
1 Introduction	1
2 Literature Review	4
2.1 Linear and regression-based formulations	4
2.2 Stochastic and dominance-based models	5
2.3 Robust optimization	6
2.4 Multi-objective and reward–risk models	6
2.5 Summary	7
3 Models	8
3.1 Problem 1: Factor-based	8
3.1.1 Nominal model	8
3.1.2 Robust factor model	9
3.2 Problem 2: Mixture of Distribution	13
3.2.1 Nominal mixture model	13
3.2.2 Robust mixture model	14
4 Computational Experiments - Problem 1	16
4.1 Performance metrics	16
4.1.1 Return	16
4.1.2 Variance	17
4.1.3 Sharpe ratio	17
4.1.4 Tracking Ratio	17
4.1.5 Tracking Error Variance (TEV)	18
4.2 Dataset and Data Pre-Processing	18
4.3 In-sample Analysis	20
4.3.1 Objective Function	20

4.3.2	Norm of Weight Differences	22
4.3.3	Portfolio Return	22
4.3.4	Portfolio Variance	23
4.3.5	Portfolio Sharpe Ratio	23
4.4	Out-Of-Sample Analysis	25
4.4.1	Return, Variance and Sharpe Ratio	26
4.4.2	Tracking Ratio	29
4.4.3	Tracking Error Variance	29
4.5	Comparative Analysis: In-Sample vs. Out-of-Sample	31
4.6	Rolling Window Analysis	33
4.6.1	Rolling Window Setup	33
4.6.2	Return Out-of-Sample	34
4.6.3	Variance Out-of-Sample	34
4.6.4	Sharpe ratio Out-of-Sample	35
4.6.5	Tracking ratio Out-of-Sample	36
4.6.6	Tracking error Out-of-Sample	37
5	Computational Experiments - Problem 2	40
5.1	Parameter Estimation	41
5.2	Out-of-sample Analysis	42
5.2.1	Return	42
5.2.2	Variance	42
5.2.3	Sharpe ratio	44
5.2.4	Tracking Ratio	44
5.2.5	Tracking Error	46
6	Conclusion	48
	Bibliography	50

List of Tables

4.1	Return, variance and Sharpe ratio in-sample vs out-of-sample for selected portfolio sizes	32
4.2	Examples of training and test windows in the rolling window analysis.	34

List of Figures

4.1	Objective function in-sample	21
4.2	Norm of weight differences in-sample	22
4.3	Return in-sample	23
4.4	Variance in-sample	24
4.5	Sharpe ratio in-sample	25
4.6	Return out-of-sample	27
4.7	Variance out-of-sample	27
4.8	Sharpe ratio out-of-sample	28
4.9	Tracking ratio out-of-sample	30
4.10	Tracking error out-of-sample	30
4.11	Comparison in-sample/out-of-sample return	31
4.12	Comparison in-sample/out-of-sample variance	31
4.13	Comparison in-sample/out-of-sample Sharpe ratio	32
4.14	Return rolling window out-of-sample	35
4.15	Variance rolling window out-of-sample	36
4.16	Sharpe ratio rolling window out-of-sample	37
4.17	Tracking ratio rolling window out-of-sample	38
4.18	Tracking error rolling window out-of-sample	39
5.1	Return out-of-sample	43
5.2	Variance out-of-sample	44
5.3	Sharpe ratio out-of-sample	45
5.4	Tracking ratio out-of-sample	46
5.5	Tracking error out-of-sample	47

Chapter 1

Introduction

In the field of quantitative finance, index funds refer to investment strategies designed to replicate the performance of a specific market index, commonly referred to as the benchmark. This practice, known as *index tracking*, constitutes a classical form of passive portfolio management, whereby the portfolio is constructed to replicate the benchmark's returns as closely as possible. The core objective is the minimization of the tracking error, a statistical measure that quantifies how closely the portfolio follows the benchmark's performance. Over the years, numerous optimization models have been proposed to tackle the index tracking problem, differing primarily in the mathematical formulation of the tracking error and in the algorithms employed for its minimization.

Within this context, *enhanced index tracking* has emerged as a more flexible alternative. While still oriented toward replicating the benchmark, enhanced strategies allow for moderate deviations from the index composition, with the goal of achieving excess return while maintaining limited tracking error. The resulting optimization problem—known as EITP—is typically formulated as a bi-objective model, balancing tracking error minimization with the pursuit of active returns.

Empirical evidence suggests that enhanced index funds have gained considerable traction over the last three decades. For instance, Jorion (2002) reports that the volume of assets under enhanced index management in the United States grew more than tenfold between 1994 and 2000, outpacing the growth of traditional passive funds over the same period. Moreover, the adoption of enhanced strategies is not limited to developed markets: Weng and Wang (2017) document a substantial increase in the popularity of enhanced index funds in China between 2008 and 2015. Despite this growing practical interest, the academic literature on EITP remains relatively limited compared to the extensive body of work on classical index tracking, with most contributions emerging only in recent years.

This work examines two recent approaches to enhanced index tracking, specifically those introduced by Kwon and Wu (2017) and Kang et al. (2022). Both are

grounded in *robust optimization*, an area that remains relatively underexplored in the context of the EITP but offers considerable potential. The remainder of this section outlines the key ideas underlying these models.

The model developed by Kwon and Wu (2017) adopts a robust optimization approach that aims to maximize the expected return of a tracking portfolio while maintaining a controlled level of tracking error and portfolio risk. Their formulation, based on a *factor model* and incorporating a cardinality constraint, addresses the sensitivity of mean-variance optimization to estimation errors in asset returns and factor exposures. *Uncertainty* is captured through data-driven uncertainty sets within a robust optimization framework, and the resulting problem is formulated as a MISOCP.

In the present study, their methodology is adapted in two main directions. First, the implementation targets the broader S&P 500 index rather than the S&P 100, increasing both the dimensionality and the practical relevance of the model. Second, in addition to the Fama–French three-factor model employed in the original study, the five-factor extension is also considered, allowing for a more comprehensive analysis of how factor structures influence portfolio performance. The detailed mathematical formulation of both the nominal and robust versions of the model is provided in the subsequent chapters.

In contrast to the factor-based robust formulation, the second approach examined in this thesis is based on the enhanced index tracking framework proposed by Kang et al. (2022), which aims to capture the heavy-tailed nature of asset returns more accurately. To this end, the authors model asset returns using a mixture distribution framework, specifically *Gaussian Mixture Models*. This choice offers greater flexibility in approximating real-world return distributions and facilitates more accurate estimation of risk and reward.

Their contribution lies in formulating a stochastic EIT model under the assumption that asset returns follow a mixture of normal distributions, from which a deterministic equivalent can be derived. The tracking error is measured asymmetrically using LPM, penalizing only downside deviations from the benchmark-plus-alpha return — an approach aligned with the objectives of enhanced tracking.

The model is further extended to a robust formulation by incorporating uncertainty in the mixture proportions using ϕ -divergence. This leads to a distributionally robust optimization problem that, while more complex, remains theoretically tractable. In practice, however, the computational burden of solving the full robust model is substantial. To address this challenge, the implementation presented here applies a fix-and-relax strategy to obtain approximate but feasible solutions within reasonable computational time. While the original study evaluates the methodology on both the FTSE 100 and S&P 500 indices, the implementation developed in this work focuses exclusively on the S&P 500.

The subsequent chapters present the detailed formulation, implementation

strategies, and empirical evaluation of both models, highlighting their respective strengths, limitations, and potential for real-world application.

Chapter 2

Literature Review

The Enhanced Index Tracking Problem (EITP) has evolved over the last three decades into a rich research area at the intersection of finance and optimization. Early surveys such as Canakgoz and Beasley (2009) and more recent contributions like Guastaroba et al. (2020) highlight the variety of formulations and solution strategies proposed. In this chapter, we synthesize these contributions, focusing on how modeling approaches and methodologies have progressed.

The EITP was first formally introduced by Beasley et al. (2003), who proposed an evolutionary heuristic to balance two conflicting objectives: minimizing tracking error and maximizing excess return over the benchmark. By introducing a parameter into the objective function, they allowed the decision-maker to adjust the relative importance of these goals. The resulting model is a nonlinear mixed-integer program, solved through heuristic search rather than exact optimization.

2.1 Linear and regression-based formulations

A second important line of research has focused on linear and regression-based formulations of the EITP, aiming to improve tractability while retaining realistic features of portfolio-benchmark relationships. Canakgoz and Beasley (2009) introduced regression-based models where the slope and intercept of the regression line are incorporated through a multi-stage methodology. This approach provides a linear representation of the link between portfolio and benchmark and proved effective both for index tracking and for enhanced indexation.

Building on the idea of embedding statistical tools in linear optimization, Mezali and Beasley (2013) employed quantile regression to formulate a Mixed-Integer Linear Program (MILP) that minimizes a weighted sum of absolute positive and negative residuals, thereby capturing distributional aspects of returns within a linear structure and extending the regression-based perspective introduced earlier.

Similarly, Valle et al. (2014) analyzed absolute return portfolios and proposed a three-stage mixed-integer formulation applicable to EITP, particularly under volatile market conditions; their multi-stage design echoes the staged logic of regression-based models while emphasizing flexibility in handling instability. Another notable contribution is by Wu et al. (2017) that combined clustering with linear optimization, embedding cluster analysis into a mixed-integer framework; through Lagrangian and Semi-Lagrangian relaxations, they addressed scalability in large asset universes, illustrating how linear formulations integrate naturally with data-driven preprocessing.

2.2 Stochastic and dominance-based models

Parallel to linear modeling, a substantial body of work formulates the EITP in stochastic terms to account explicitly for randomness in returns. Dose and Cincotti (2005) introduced a two-stage stochastic procedure grounded in time-series cluster analysis, where representative assets are selected before weight optimization, foreshadowing later integrations of data-driven screening with exact models. Lejeune and Samatlı-Paç (2013) then modeled returns and covariances as random variables within a stochastic Mixed-Integer Non-Linear Programming (MINLP) and proposed an outer approximation strategy that improves tractability by relaxing selected binaries and reformulating cardinality constraints.

Within the same stochastic paradigm, stochastic-dominance methods encode distributional preferences directly in the optimization model. Bruni et al. (2012) cast the problem as a large-scale linear program. Their solution relied on a constraint generation technique based on a separation procedure. Among the various stochastic methods explored, the relaxed Cumulative Zero-Order Dominance stands out. While Sharma et al. (2017) developed a linear programming model maximizing mean portfolio return while controlling violations of relaxed Second-Order Stochastic Dominance (SSD) constraints. To handle the large number of variables, they employed a modified cutting plane algorithm. In the same vein, Roman et al. (2013) developed two models to ensure that the portfolio return distribution outperforms the benchmark under SSD conditions: the first contrasts cumulative sums of ordered portfolio returns with those of the index; the second uses averages of ordered outcomes, which is equivalent to evaluating the Conditional Value at Risk (CVaR) across multiple confidence levels. These studies highlight the potential of stochastic dominance in embedding investor preferences directly into the optimization model.

2.3 Robust optimization

Beyond stochastic programming, robust optimization offers an alternative route to handle parameter uncertainty by prescribing feasible performance over uncertainty sets. Kwon and Wu (2017) formulated a robust benchmark-tracking model as a mixed-integer second-order cone program that simultaneously controls portfolio risk—measured by return standard deviation—and tracking error—measured by the standard deviation of excess returns—while extending the Fama-French three-factor structure to include uncertainty in expected returns and factor loadings. In a related line, Kang et al. (2022) modeled returns as mixtures of Gaussian components to capture higher-order moments and non-normality; their robust formulation aims to mitigate misspecification and represent extreme market scenarios more faithfully. This stream complements the stochastic line by focusing on worst-case rather than probabilistic descriptions of uncertainty.

2.4 Multi-objective and reward–risk models

Given the bi-objective nature of the EITP, several studies have explicitly modeled it as a multi-objective optimization problem. Li et al. (2011) formulated a bi-objective MINLP minimizing downside tracking error—defined as the downside standard deviation of portfolio return from the benchmark—while maximizing excess return, solved via an immunity-based metaheuristic. Bruni et al. (2015) proposed a bi-objective MILP model aimed at maximizing excess return while minimizing tracking error, measured as the absolute deviation between portfolio and benchmark values. Due to the computational complexity of solving the model exactly, the authors designed a heuristic solution framework based on the Kernel Search method.

Multi-objective approaches to the EITP often yield a set of Pareto-optimal solutions, offering flexibility but requiring subjective choices by the decision-maker. To overcome this, some studies reformulate the problem as a single-objective model using reward-risk ratios, which compare expected returns to the risk undertaken. Among the most widely used metrics are the Sharpe ratio (Sharpe (1966)), based on total volatility, and the Sortino ratio (Sortino and Price (1994)), which focuses only on downside risk. These measures are extensively applied to assess and rank strategies by risk-adjusted performance.

Meade and Beasley (2011) were among the first to integrate a reward–risk measure in EITP, maximizing a modified Sortino ratio via a genetic algorithm. However, nonlinearities and practical constraints posed challenges. To address these limitations, Guastaroba et al. (2016) proposed two EITP formulations based on the Omega ratio (Keating and Shadwick (2002)). The first uses the standard Omega

ratio with a fixed target, while the Extended Omega Ratio model incorporates a random target. Both nonlinear models can be reformulated as LPs, with empirical results showing the extended version consistently outperforms the traditional one in out-of-sample tests. Building on the idea of incorporating reward-risk measures in the EITP, Guastaroba et al. (2020) have recently explored approaches based on Conditional Value-at-Risk (CVaR). In particular, they proposed a class of bi-criteria models where risk is measured using the weighted multiple CVaR (WCVaR), a combination of several CVaR measures that allows more refined modeling of risk aversion. Both theoretical and empirical analyses indicate that portfolios optimized with WCVaR outperform those obtained with other recent models, especially in out-of-sample evaluations.

2.5 Summary

The literature on the EITP has developed from heuristic approaches to linear, stochastic, robust, and reward-risk formulations. Among these, robust optimization has gained particular prominence, as it directly tackles parameter uncertainty and offers guarantees that are especially valuable in practice. For this reason, this thesis adopts robust optimization as the primary framework for addressing the EITP.

Chapter 3

Models

3.1 Problem 1: Factor-based

3.1.1 Nominal model

The first model proposed, developed by Kwon and Wu (2017), is the nominal enhanced index tracking model. The model maximizes the portfolio's expected return while enforcing limits on portfolio risk, tracking error and the number of selected assets. Additionally, lower and upper bounds on individual asset weights are imposed to avoid excessive concentration and prevent negligible holdings, thereby improving diversification and ensuring the solution is implementable in practice.

The formulation is as follows:

$$\max_x \mu^\top x \quad (3.1)$$

$$\text{s.t. } \left\| \Sigma^{1/2} x \right\|_2 \leq \sigma, \quad (3.2)$$

$$\left\| \Sigma^{1/2} (x - x_{BM}) \right\|_2 \leq TE, \quad (3.3)$$

$$e^\top x = 1, \quad (3.4)$$

$$e^\top y = q, \quad (3.5)$$

$$lb_i y_i \leq x_i \leq ub_i y_i, \quad i = 1, \dots, n, \quad (3.6)$$

$$y_i \in \{0,1\}, \quad x_i \geq 0 \quad i = 1, \dots, n. \quad (3.7)$$

Decision variables:

- $x \in \mathbb{R}^n$: vector of portfolio weights;

- $y \in \{0,1\}^n$: binary selection vector, where $y_i = 1$ if the i -th asset is included in the portfolio.

Parameters:

- $\mu \in \mathbb{R}^n$: vector of expected returns of assets;
- $\Sigma \in \mathbb{R}^{n \times n}$: covariance matrix of returns;
- $x_{BM} \in \mathbb{R}^n$: benchmark index weights;
- $\sigma \in \mathbb{R}$: allowable portfolio risk level;
- $TE \in \mathbb{R}$: tracking error limit;
- $\mathbf{e} \in \mathbb{R}^n$: vector of ones;
- $q \in \mathbb{R}$: size of the tracking portfolio;
- $lb, ub \in \mathbb{R}^n$: lower and upper bounds on the portfolio weights.

The resulting model is a MISOCP.

3.1.2 Robust factor model

In this section, we present a robust factor counterpart of the nominal model, following the framework of Kwon and Wu (2017), which in turn builds on the seminal work of Goldfarb and Iyengar (2003).

Asset returns are modeled as

$$r = \mu + V^T f + \epsilon,$$

where $\mu \in \mathbb{R}^n$ denotes the vector of expected returns, $f \sim N(0, F) \in \mathbb{R}^m$ represents the vector of market factors, $V \in \mathbb{R}^{m \times n}$ is the factor loading matrix, and $\epsilon \sim N(0, D) \in \mathbb{R}^n$ captures the residual returns, with $D = \text{diag}(d) \in \mathbb{R}^{n \times n}$.

For consistency of the model we require the factor covariance $F \in \mathbb{R}^{m \times m}$ to be positive semidefinite. We further assume that the residuals are mutually independent and also independent of the factor realizations:

- ϵ_i and ϵ_j are independent whenever $i \neq j$;
- each ϵ_i is independent of the factor vector f .

We adopt a stability assumption whereby the factor covariance matrix F is fixed over time. Under this setting, uncertainty in (μ, Σ) originates from the parameters (μ, V, D) . Following Kwon and Wu (2017) and Goldfarb and Iyengar (2003), these are described through different uncertainty sets:

- For μ and D , interval bounds are adopted:

$$S_d = \{D : D = \text{diag}(d), d_i \in [d_i, \bar{d}_i], i = 1, \dots, n\},$$

$$S_m = \{\mu : \mu = \mu_0 + \xi, \xi_i \in [\underline{\gamma}_i, \bar{\gamma}_i], i = 1, \dots, n\}.$$

- For V , an ellipsoidal uncertainty set is considered:

$$S_v = \{V : V = V_0 + W, \|W_i\|_g \leq \rho_i, i = 1, \dots, n\},$$

where W_i is the i -th column of the strength matrix W and $\|\cdot\|_g$ is the elliptical norm induced by a matrix $G \succeq 0$. Choosing $G \succ 0$ ensures strict convexity of the robust problem.

Building on this framework, Goldfarb and Iyengar (2003), and later Kwon and Wu (2017), showed that the robust factor model can be reformulated using only linear and second-order cone constraints. To make this connection precise, we recall a lemma that provides the key equivalence used in the transformation of the risk bound (an analogous argument applies to the tracking-error constraint).

Lemma 3.1.1 *Let $r, v > 0, y_0, y \in \mathbb{R}^m$ and $F, G \in \mathbb{R}^{m \times m}$ be positive definite matrices. Then*

$$\max_{\{y : \|y\|_g \leq r\}} \|y_0 + y\|_f^2 \leq v$$

is equivalent to either of the following:

1. *there exist $\tau, \sigma \geq 0$, and $t \in \mathbb{R}_+^m$ such that*

$$\begin{aligned} v &\geq \tau + e^T t \\ \sigma &\leq \frac{1}{\lambda_{\max}(H)} \\ r^2 &\leq \sigma \tau \\ w_i^2 &\leq (1 - \sigma \lambda_i) t_i, \quad i = 1, \dots, m \end{aligned}$$

where $Q\Lambda Q^T$ is the spectral decomposition of $H = G^{-1/2}FG^{-1/2}$, $\Lambda = \text{diag}(\lambda_i)$ and $w = Q^T H^{1/2} G^{1/2} y_0$;

2. *there exist $\tau \geq 0$ and $s \in \mathbb{R}_+^m$ such that*

$$\begin{aligned} r^2 &\leq \tau(v - e^T s) \\ u_i^2 &\leq (1 - \tau \theta_i) s_i, \quad i = 1, \dots, m \\ \tau &\leq \frac{1}{\lambda_{\max}(K)} \end{aligned}$$

where $P\Theta P^T$ is the spectral decomposition of $K = F^{1/2}G^{-1}F^{1/2}$, $\Theta = \text{diag}(\theta_i)$ and $u = P^T F^{1/2} y_0$.

Applying the factor model to the nominal risk condition yields

$$\|\Sigma^{1/2}x\|_2^2 \leq \sigma^2 \iff x^T(V^T FV + D)x \leq \sigma^2 \iff x^T V^T FVx + x^T Dx \leq \sigma^2$$

Hence, in the robust case,

$$\max_{V \in S_v} x^T V^T FVx + \max_{D \in S_d} x^T Dx \leq \sigma^2.$$

Introducing slack variables $v, \delta \geq 0$ allows the two contributions to be treated separately:

$$\begin{cases} \max_{V \in S_v} x^T V^T FVx \leq v \\ \max_{D \in S_d} x^T Dx \leq \delta \\ v + \delta \leq \sigma^2 \end{cases} \iff \begin{cases} \max_{V \in S_v} x^T V^T FVx \leq v \\ \left\| \begin{bmatrix} 2\bar{D}^{1/2}x \\ 1 - \delta \end{bmatrix} \right\| \leq 1 + \delta \\ v + \delta \leq \sigma^2 \end{cases}$$

The first inequality is handled through Lemma 3.1.1. Substituting $y = Vx$ and carrying out the change of basis leads to the equivalent SOCP representation:

$$\begin{aligned} \max_{V \in S_v} x^T V^T FVx \leq v &\iff \max_{V \in S_v} \|Vx\|_f^2 \\ &\iff \begin{cases} u = P^T F^{1/2} V_0 x \\ \left\| \begin{bmatrix} 2\rho^T x \\ \tau - \nu + e^T s \end{bmatrix} \right\| \leq \tau + \nu - e^T s \\ \left\| \begin{bmatrix} 2u_i \\ \nu - \tau\theta_i - s_i \end{bmatrix} \right\| \leq \nu - \tau\theta_i + s_i, \forall i = 1, \dots, m \\ \nu - \tau\lambda_{\max}(K) \geq 0 \\ \tau \geq 0 \end{cases} \end{aligned}$$

This reformulation makes it possible to express the worst-case quadratic terms solely through linear and second-order cone conditions. As a result, the robust tracking formulation remains a MISOCP.

The complete problem can be written as:

$$\max (\mu_0 + \underline{\gamma})^\top x \tag{3.8}$$

$$\text{s.t. } u = P^T F^{1/2} V_0 x \tag{3.9}$$

$$w = P^T F^{1/2} V_0 (z^+ - z^-) \tag{3.10}$$

$$z^+ - z^- = x - x_{BM} \tag{3.11}$$

$$\left\| \begin{bmatrix} 2\rho^\top x \\ \tau - \nu + e^\top s \end{bmatrix} \right\| \leq \tau + \nu - e^\top s \quad (3.12)$$

$$\left\| \begin{bmatrix} 2u_i \\ \nu - \tau\theta_i - s_i \end{bmatrix} \right\| \leq \nu - \tau\theta_i + s_i, \quad \forall i = 1, \dots, m \quad (3.13)$$

$$\left\| \begin{bmatrix} 2\rho^\top(z^+ + z^-) \\ \tau - l + e^\top s \end{bmatrix} \right\| \leq \tau + l - e^\top s \quad (3.14)$$

$$\left\| \begin{bmatrix} 2w_i \\ l - \tau\theta_i - s_i \end{bmatrix} \right\| \leq l - \tau\theta_i + s_i, \quad \forall i = 1, \dots, m \quad (3.15)$$

$$\left\| \begin{bmatrix} 2\bar{D}^{1/2}x \\ 1 - \delta \end{bmatrix} \right\| \leq 1 + \delta \quad (3.16)$$

$$\nu + \delta \leq \sigma^2 \quad (3.17)$$

$$\left\| \begin{bmatrix} 2\bar{D}^{1/2}(z^+ - z^-) \\ 1 - \zeta \end{bmatrix} \right\| \leq 1 + \zeta \quad (3.18)$$

$$l + \zeta \leq TE^2 \quad (3.19)$$

$$\nu - \tau\lambda_{\max}(K) \geq 0 \quad (3.20)$$

$$l - \tau\lambda_{\max}(K) \geq 0 \quad (3.21)$$

$$e^\top x = 1 \quad (3.22)$$

$$e^\top y = q \quad (3.23)$$

$$lb_i y_i \leq x_i \leq ub_i y_i, \quad \forall i = 1, \dots, n \quad (3.24)$$

$$x \geq 0, \tau \geq 0, z^+ \geq 0, z^- \geq 0 \quad (3.25)$$

$$y \in \{0,1\}^n \quad (3.26)$$

where $x, z^+, z^- \in \mathbb{R}^n$, $u, w, s \in \mathbb{R}^m$ and $\nu, \delta, l, \zeta, \tau \in \mathbb{R}$.

Following Kwon and Wu (2017), we implement the Fama–French three-factor specification as the underlying factor structure. This extension of the classical CAPM accounts not only for the market excess return, but also for the value versus growth effect (book-to-market) and the size premium (small-minus-big). Compared with the single-factor CAPM, this richer structure captures a larger portion of the cross-sectional variation in excess returns and forms the basis of the uncertainty

sets for both expected returns and the covariance of the factor loadings.

Moreover, in order to broaden the empirical foundation, we also consider the Fama–French five-factor model. This variant supplements the three-factor version with two further risk drivers: the profitability factor (RMW), defined as the return spread between firms with robust and weak operating profitability, and the investment factor (CMA), capturing the return differential between firms pursuing conservative versus aggressive investment strategies.

3.2 Problem 2: Mixture of Distribution

3.2.1 Nominal mixture model

We now introduce the nominal formulation of the Enhanced Index Tracking (EIT) model with mixture distributions, following Kang et al. (2022).

Let $\tilde{\xi} = (\xi_1, \dots, \xi_n, \xi_b)^\top \in \mathbb{R}^{n+1}$ be the vector of random returns of the n assets and the benchmark, and $\tilde{x} = (x_1, \dots, x_n, -1)^\top \in \mathbb{R}^{n+1}$ the extended portfolio vector where $x_1, \dots, x_n$ represents the tracking portfolio weights. The relative excess return can be expressed as

$$Y = \tilde{\xi}^\top \tilde{x} = \sum_{i=1}^n x_i \xi_i - \xi_b,$$

so that the tracking error is defined in terms of Y .

The risk of underperforming the benchmark is measured through the τ -th order LPM is

$$\min_{x \in \mathcal{X}} \mathbb{E}_P \left[(\kappa - Y)_+^\tau \right], \quad (3.27)$$

where $\mathcal{X}$ denotes the feasible set of portfolios, κ is the desired excess return, and τ controls the order of downside risk considered. This formulation penalizes scenarios in which the portfolio return underperforms the benchmark by more than κ , thereby encouraging the portfolio to outperform the index of the parameter κ .

Since financial returns often display non-normality and multimodality, we model the distribution of $\tilde{\xi}$ as a Gaussian mixture with d components, i.e.,

$$\tilde{\xi} \sim \sum_{i=1}^d \lambda_i \mathcal{N}(\mu_i, \Sigma_i),$$

where λ_i are the component weights, and (μ_i, Σ_i) are the mean and covariance of the i -th Gaussian component. For a given portfolio vector $\tilde{x}$, the excess return $Y = \tilde{\xi}^\top \tilde{x}$ is then normally distributed under component i , with mean and variance

$$\nu_i(x) = \mu_i^\top \tilde{x}, \quad \sigma_i^2(x) = \tilde{x}^\top \Sigma_i \tilde{x}.$$

These quantities represent, respectively, the expected relative return and its volatility if the market were in regime i of the mixture.

Under this assumption, and following the derivation in Kang et al. (2022), the LPM admits closed-form expressions for $\tau = 1, 2$:

$$R_i^1(x) = (\kappa - \nu_i(x)) \Phi\left(\frac{\kappa - \nu_i(x)}{\sigma_i(x)}\right) + \sigma_i(x) \varphi\left(\frac{\kappa - \nu_i(x)}{\sigma_i(x)}\right), \quad i = 1, \dots, d, \quad (3.28)$$

$$\begin{aligned} R_i^2(x) = & \left(\sigma_i^2(x) + (\kappa - \nu_i(x))^2\right) \Phi\left(\frac{\kappa - \nu_i(x)}{\sigma_i(x)}\right) \\ & + \sigma_i(x)(\kappa - \nu_i(x)) \varphi\left(\frac{\kappa - \nu_i(x)}{\sigma_i(x)}\right), \quad i = 1, \dots, d, \end{aligned} \quad (3.29)$$

where $\Phi(\cdot)$ and $\varphi(\cdot)$ denote the standard normal cumulative distribution and density function, respectively.

Finally, aggregating across all mixture components with weights λ_i (the mixture proportions), the nominal EIT problem can be written as

$$\min_{x \in \mathcal{X}} \sum_{i=1}^d \lambda_i R_i^\tau(x), \quad (3.30)$$

where the λ_i 's are estimated through an Expectation-Maximization (EM) process, as described later in Section 5.1.

3.2.2 Robust mixture model

To incorporate distributional uncertainty, we reformulate the model by introducing an ambiguity set for the mixture proportions, following Kang et al. (2022).

In practice, the true vector of proportions λ is not assumed to coincide exactly with the estimate $\hat{\lambda}$ obtained via EM, but is instead allowed to vary within a neighborhood of $\hat{\lambda}$. This neighborhood is defined through the concept of ϕ -divergence, which provides a statistical measure of distance between probability distributions.

The resulting ambiguity set is defined as

$$\Lambda^\phi = \left\{ \lambda \in \mathbb{R}_+^d \mid \sum_{i=1}^d \lambda_i = 1, I_\phi(\lambda, \hat{\lambda}) \leq \rho \right\},$$

where ρ determines the size of the uncertainty region around $\hat{\lambda}$.

The robust counterpart of the EIT problem is then expressed as

$$\min_{x \in \mathcal{X}} \max_{\lambda \in \Lambda^\phi} \sum_{i=1}^d \lambda_i R_i^\tau(x),$$

which corresponds to optimizing the portfolio against the worst-case choice of mixture proportions lying in the divergence ball Λ^ϕ .

Specializing to the Kullback–Leibler (KL) divergence, the inner maximization problem admits a tractable dual representation. As shown in Kang et al. (2022), the robust EIT problem can then be written in the equivalent form:

$$\begin{aligned} \min_{x, \theta, \zeta} \quad & \theta + \zeta \rho + \zeta \sum_{i=1}^d \hat{\lambda}_i \left[\exp\left(\frac{R_i^r(x) - \theta}{\zeta}\right) - 1 \right] \\ \text{s.t.} \quad & \zeta \geq 0, \ x \in \mathcal{X}, \end{aligned} \tag{3.31}$$

where $\hat{\lambda}_i$ are the estimated mixture proportions, and θ and ζ are the dual variables corresponding to the constraints defining the ambiguity set Λ^ϕ .

This formulation, referred to as Distributionally robust Enhanced Index Tracking (DEIT), ensures robustness against estimation errors in the mixture distribution.

Chapter 4

Computational Experiments - Problem 1

4.1 Performance metrics

In order to evaluate the effectiveness of the tracking portfolios, we employ a set of performance metrics that capture different aspects of portfolio outcomes. These indicators include measures of profitability, risk, and replication accuracy. Specifically, we consider average return, variance, Sharpe ratio, as well as two replication-focused statistics — tracking ratio and tracking error variance. Together, these metrics provide a comprehensive view of portfolio performance, allowing us to compare alternative strategies, identify periods of divergence from the benchmark, and assess whether a portfolio achieves a satisfactory balance between risk and fidelity to the index.

4.1.1 Return

Let $r_{i,t} = \frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}}$ denote the daily return of asset i at time t , where $P_{i,t}$ is the closing price. The mean return of asset i over T days is

$$\bar{r}_i = \frac{1}{T} \sum_{t=1}^T r_{i,t}.$$

For a tracking portfolio holding q assets with weights $w = (w_1, \dots, w_q)^\top$ and a benchmark with weights w^b , the portfolio and index returns are

$$r_p = \sum_{i=1}^q w_i \bar{r}_i, \quad r_I = \sum_{i=1}^n w_i^b \bar{r}_i.$$

The mean return captures the portfolio's profitability over the evaluation horizon, providing a direct measure of its growth relative to the benchmark.

4.1.2 Variance

Variance quantifies the dispersion of returns around their mean. Let the sample covariance between assets i and j over T days be

$$\text{cov}_{i,j} = \frac{1}{T-1} \sum_{t=1}^T (r_{i,t} - \bar{r}_i)(r_{j,t} - \bar{r}_j),$$

and let Σ be the corresponding $n \times n$ covariance matrix of all index assets, with Σ_q the submatrix restricted to the q selected portfolio assets. Then the portfolio and benchmark variances are

$$\sigma_p^2 = w^\top \Sigma_q w, \quad \sigma_I^2 = (w^b)^\top \Sigma w^b,$$

where $w = (w_1, \dots, w_q)^\top$ and w^b are the portfolio and index weight vectors, respectively. The standard deviations σ_p and σ_I represent portfolio and benchmark volatility.

The variance thus provides a measure of the risk inherent in the portfolio by quantifying return fluctuations, thereby indicating the degree of volatility relative to the index.

4.1.3 Sharpe ratio

The Sharpe ratio, introduced by Sharpe (1966), compresses risk and return into a single metric. Using a risk-free rate r_f (here proxied by the 3-month U.S. Treasury Bill), portfolio and benchmark Sharpe ratios are

$$S_p = \frac{r_p - r_f}{\sigma_p}, \quad S_I = \frac{r_I - r_f}{\sigma_I},$$

where r_p, r_I are the portfolio and index returns, and σ_p, σ_I their standard deviations. Comparing S_p and S_I allows assessing whether the tracking portfolio not only replicates the benchmark but also achieves an efficient balance between risk and return.

4.1.4 Tracking Ratio

We define the Tracking Ratio as a relative measure that compares the cumulative performance of the index against that of the portfolio over the period from the

initial date 0 to a later date t . Formally, it is given by the equation:

$$TR_{0,t} = \frac{\frac{\sum_{i=1}^n V_{i,t}}{\sum_{i=1}^n V_{i,0}}}{\frac{\sum_{j=1}^q w_j V_{j,t}}{\sum_{j=1}^q w_j V_{j,0}}}. \quad (4.1)$$

The components of the formula can be interpreted as follows:

- **Numerator:** represents the index's cumulative growth factor, computed as the ratio of the total market capitalization at time t to the total market capitalization at the initial date 0.
- **Denominator:** represents the portfolio's cumulative growth factor, obtained as the ratio of the portfolio value at time t to its value at the initial date 0, keeping the portfolio composition fixed as in 0.

By construction, $TR_{0,t} = 1$ denotes perfect cumulative replication, values above 1 signal relative underperformance of the portfolio, while values below 1 indicate relative outperformance.

4.1.5 Tracking Error Variance (TEV)

Tracking error variance quantifies the variability of the return difference between the portfolio and the benchmark. Writing the full covariance matrix over the evaluation window as Σ and aligning the weight vectors w and w^b to the same asset ordering (padding zeros for non-held assets), the TEV is

$$TEV_p = (w - w^b)^\top \Sigma (w - w^b) = \sum_{i=1}^n \sum_{j=1}^n (w_i - w_i^b) \sigma_{i,j} (w_j - w_j^b). \quad (4.2)$$

A low TEV indicates that portfolio returns closely track the benchmark in variance terms, while higher values highlight greater deviations in risk exposure.

4.2 Dataset and Data Pre-Processing

The empirical analysis takes the S&P 500 index as benchmark, given its role as one of the most widely followed indicators of the U.S. equity market and its ability to capture the performance of large-cap American companies. The index is constructed on a capitalization-weighted basis, which implies that larger firms contribute more strongly to its movements. Its composition is reviewed periodically to reflect market dynamics, corporate actions, and eligibility rules, making it a relevant and realistic case study for index replication.

The list of constituents was retrieved from publicly available sources, while daily financial data were collected from Yahoo Finance. For each company, we obtained historical prices and information on market capitalization over the period 2020–2023. This window was selected as it provides a sufficiently long sample for the estimation of statistical moments while keeping the analysis focused on recent market conditions.

To ensure consistency across the entire period, only firms that remained continuously in the S&P 500 from 2020 to 2023 were retained. Companies that entered or exited the index during the window were excluded, which led to the removal of 13 assets from the initial universe. This filtering step prevents distortions arising from changes in index composition and guarantees comparability of results over time.

In addition to the historical price series, daily market capitalizations were calculated for each firm by multiplying the closing price by the number of outstanding shares. This information is fundamental, since it establishes the relative size of each company within the index and directly affects the computation of index weights and the design of tracking portfolios.

Finally, the dataset was cleaned to address missing values and potential inconsistencies. The result of this process is a coherent and stable dataset composed of a fixed group of S&P 500 constituents, each with reliable time series of prices and market capitalizations, ready to be used in the subsequent modeling and evaluation phase.

In this chapter, we present the computational experiments conducted on the factor-based problem. We develop the models described in equations (3.1)–(3.7) and (3.8)–(3.26), considering both the three-factor and five-factor specifications for the latter. We begin by analyzing the in-sample results and subsequently evaluate the models’ performance in the static and dynamic out-of-sample settings.

Prior to the analysis, several key parameters were fixed to ensure consistency and comparability across the study: $lb_i = 0.01$ and $ub_i = 0.7$ for all i , $\sigma = 5 \cdot p_{95}$, and $TE = 3 \cdot p_{95}$, where p_{95} denotes the 95th percentile of the standard deviations, computed as $\text{std_dev}_i = \sqrt{\Sigma_{ii}}$ from the covariance matrix Σ . The remaining model parameters, namely μ_0 , V_0 , G , ρ_i , γ_i , d_i , and the corresponding uncertainty sets for μ_0 and V_0 , were chosen to be consistent with those reported in Kwon and Wu (2017). For a detailed account of their construction, readers are referred to Goldfarb and Iyengar (2003). Finally, the joint confidence level is set to $\omega = 0.95$.

All tracking models were solved using Gurobi 11.0.3, with Python 3.12.4 as the interface, on a PC equipped with 16 GB of RAM.

4.3 In-sample Analysis

In-sample evaluation serves as an initial benchmark for assessing how well the optimization models replicate the reference index. In this framework, the portfolio is constructed and assessed over the same historical period, with asset weights held constant throughout the evaluation horizon. While this approach offers a transparent way to test the internal consistency of the models, it should be emphasized that the results tend to be optimistic compared with out-of-sample analysis, as they do not account for changing market conditions or shifts in asset correlations.

The empirical exercise relies on daily data from the S&P 500 index between December 31, 2019 and December 31, 2020. The analysis focuses on the year 2020 as the primary window; however, to compute the first daily return in January 2020, the closing price of the previous trading day, December 31, 2019, was included as a reference point. This choice ensures the consistency of the return series without actually extending the analysis to the previous year. From this window, average returns and the covariance matrix were estimated. Market capitalizations for 2020 and sector classifications of the constituents were also incorporated in order to preserve the structural characteristics of the benchmark.

To examine the role of portfolio size, the optimization problems were solved for a range of cardinalities, specifically

$$q = \{11, 20, 30, 40, 50, 60, 70, 80, 90, 100\}.$$

For each case, the optimal solution was obtained and stored in a structured Python object that allows for subsequent comparisons across models. Alongside the portfolio composition and weights, the stored outputs include the value of the objective function, the mean returns of the selected assets, the deviation between portfolio and index weights, the overall portfolio return, the variance, and the Sharpe ratio.

These results provide the basis for evaluating the comparative performance of the different optimization approaches under in-sample conditions. Before turning to the analysis, it is useful to clarify the model labels used in the following figures: Model 1 denotes the nominal formulation, Model 2 corresponds to the robust specification based on the Fama–French three-factor model, and Model 3 represents the robust version relying on the Fama–French five-factor model. In all cases, outcomes are reported as a function of the portfolio cardinality q , that is, the number of assets included in the tracking portfolio.

4.3.1 Objective Function

Figure 4.1 reports the optimal values of the maximization problem as a function of the portfolio size q . Since the problem is one of maximization, higher values of the

objective function are preferable. However, the attained optima tend to decrease as q grows, reflecting how diversification and robustness interact.

The nominal formulation (Model 1) shows the steepest reduction: adding assets quickly lowers the maximum achievable value. In contrast, the robust counterparts (Models 2 and 3) display a more gradual decline, indicating that the conservativeness induced by parameter uncertainty partly offsets the diversification effect.

Interestingly, Model 3 exhibits a slight increase in the optimal value for very small portfolios before starting to decline. Although modest, this suggests that under limited diversification, the interaction between uncertainty and the five-factor specification may temporarily enhance performance. After this adjustment, the decline resumes, though more slowly than in Model 2.

Comparing the two robust models, the five-factor specification (Model 3) yields the flattest curve, with the smallest reduction in the optimal value as q increases. The broader uncertainty set induced by additional factors makes the objective almost insensitive to portfolio size, so enlarging the portfolio brings little gain—especially in Model 3—while significantly raising computational cost.




Figure 4.1: Objective function in-sample

4.3.2 Norm of Weight Differences

Figure 4.2 depicts the distance between portfolio weights and the S&P500 weights, measured in terms of the norm of their difference. The nominal model (Model 1) remains the farthest from the target, though the gap narrows as q grows. The three-factor robust model (Model 2) begins at an intermediate distance, diverges slightly for small portfolios, and then converges once q becomes larger. The five-factor version (Model 3) consistently stays closer to the benchmark, with deviations gradually shrinking as q increases. These patterns highlight how robustness and factor specification influence the alignment of the tracking portfolio with the benchmark, with Model 3 achieving the closest match, especially for larger portfolios.




Figure 4.2: Norm of weight differences in-sample

4.3.3 Portfolio Return

Figure 4.3 reports the in-sample returns of the tracking portfolios. The nominal model (Model 1) achieves the highest return, followed by Model 2 and Model 3. In every case, the constructed portfolios outperform the S&P 500 benchmark, often by a sizable margin. The higher return of the nominal model can be explained by its lack of robustness constraints: without accounting for parameter uncertainty, the optimization is free to exploit historical patterns fully, leading to more aggressive weight allocations that maximize in-sample performance. By contrast, the robust

models impose conservativeness to hedge against estimation errors, reducing potential gains while enhancing protection against uncertainty. This trade-off anticipates the variance and Sharpe ratio results discussed below.




Figure 4.3: Return in-sample

4.3.4 Portfolio Variance

Figure 4.4 illustrates the portfolio variances. As expected, the nominal model (Model 1) exhibits the highest variance, reflecting its aggressive in-sample optimization without accounting for uncertainty. Model 2 shows a reduced variance, while Model 3 achieves the lowest variance among the three. This highlights the stabilizing role of robustness: by incorporating parameters uncertainty, the robust models discourage extreme weight allocations, leading to more stable and less volatile portfolios, with the five-factor formulation being the most conservative.

4.3.5 Portfolio Sharpe Ratio

Figure 4.5 presents the Sharpe ratios of the tracking portfolios, which measure performance relative to risk. The robust models outperform the nominal one, with Model 2 achieving the highest ratio. This indicates that introducing robustness helps control variability while maintaining returns, leading to a more efficient




Figure 4.4: Variance in-sample

portfolio. Overall, robustness appears to improve the trade-off between return and volatility, resulting in a more stable and effective allocation.




Figure 4.5: Sharpe ratio in-sample

4.4 Out-Of-Sample Analysis

As a preliminary step beyond the in-sample experiments, the models are evaluated out-of-sample in order to test how well the portfolio allocations perform on data not used during training. This first validation exercise is carried out in a static framework: portfolios are constructed using in-sample information, and the resulting asset weights are then held constant throughout the evaluation horizon. No rebalancing or parameter updates are performed, so performance depends exclusively on how the initial allocation responds to market dynamics in the subsequent period. This setup reflects the behavior of an investor that avoids frequent adjustments—often to limit transaction costs—and provides an initial measure of how stable and robust the optimized portfolios are when confronted with new data.

Performance is assessed using standard indicators such as return, volatility, and the Sharpe ratio, always in comparison with the benchmark index. Concretely, portfolios built on the 2020 dataset are tested against market developments in 2021, thereby providing an initial out-of-sample assessment of tracking quality.

It is important to emphasize that this static assessment represents only a preliminary evaluation. While it provides useful insights into the robustness of the initial allocations, it does not reflect the dynamics of an actively monitored strategy. For this reason, the next section introduces a rolling-window framework with rebalancing, which allows for parameter updates and offers a more realistic representation of how the portfolios would perform in practice.

4.4.1 Return, Variance and Sharpe Ratio

Figures 4.6, 4.7, 4.8 report the out-of-sample performance of the tracking portfolios. In terms of returns, Models 2 and 3 achieve very similar outcomes, with Model 3 slightly outperforming Model 2. Both clearly dominate the nominal specification (Model 1), which in turn delivers higher returns than the S&P500 benchmark. This confirms that, even out-of-sample, the constructed portfolios consistently outperform the benchmark. Turning to variance, the same hierarchy observed in-sample persists: the nominal model exhibits the highest variance, followed by three-factor (Model 2) and five-factor models (Model 3). However, the differences among the three specifications are less pronounced than before, although all remain significantly above the benchmark. This alignment between higher returns and higher variance is expected, as superior profitability naturally entails greater risk exposure. Finally, the Sharpe ratio analysis highlights a distinct picture: the benchmark achieves the highest risk-adjusted performance, followed closely by the five-factor model and, at a slightly lower level, the three-factor model. The nominal specification performs markedly worse, underscoring its inefficiency when accounting for risk.

Overall, these findings highlight the advantages of adopting robust Fama–French specifications. By explicitly incorporating parameters uncertainty, the robust optimization framework avoids overfitting to historical patterns and instead encourages more stable allocations, which translates into improved efficiency out-of-sample. Among the robust formulations, the five-factor model emerges as the most effective: by capturing a richer structure of systematic risks, it delivers a better balance between return and volatility, ranking closest to the benchmark in terms of Sharpe ratio while still ensuring consistently higher raw returns. This result confirms the theoretical motivation behind robust portfolio construction: the additional conservativeness acts as a safeguard against estimation errors, enabling the portfolio to remain competitive even under adverse scenarios. It is also worth noting that further improvements could likely be obtained by lowering the portfolio risk and tracking error thresholds. Due to the computational resources available, such configurations could not be fully explored in this work, but they represent a promising direction in which the robust five-factor approach could show even more substantial out-of-sample benefits.




Figure 4.6: Return out-of-sample




Figure 4.7: Variance out-of-sample




Figure 4.8: Sharpe ratio out-of-sample

4.4.2 Tracking Ratio

Tracking ratios are computed as in Equation 4.1, with the benchmark growth factor placed in the numerator and the tracking portfolio growth factor in the denominator. The corresponding plot is shown in Figure 4.9. Under this convention, values below one indicate portfolio outperformance relative to the benchmark.

In the out-of-sample analysis, the nominal specification (Model 1) achieves the highest tracking ratios, while the robust formulations (Models 2 and 3) exhibit lower values. Taken at face value, this ranking could suggest weaker tracking performance of the robust portfolios. Yet, this interpretation would be misleading: the lower ratios primarily result from the stronger absolute outperformance of the robust models, which mechanically inflates the denominator of the ratio and reduces the indicator. In other words, a smaller ratio here does not necessarily correspond to poorer replication quality, but rather to higher cumulative returns relative to the benchmark.

It is also important to note—consistent with the discussion in the previous section—that these results are shaped by the constraints of the experimental setup. In particular, the limited computational resources restricted the range of feasible parameter configurations, especially for variance and tracking error bounds. With a broader set of configurations, the tracking ratios would likely provide a more accurate picture of the efficiency of the robust specifications. For this reason, the values reported here should be regarded as conservative and mainly reflecting design limitations, rather than as evidence of intrinsic shortcomings of the models.

4.4.3 Tracking Error Variance

Figure 4.10 presents the out-of-sample tracking error variance of the constructed portfolios. Five-factor model consistently achieves the lowest values, clearly outperforming both nominal and three-factor models across the entire range of cardinalities. Model 2 follows, generally improving upon Model 1, especially for lower portfolio sizes. Only at very high cardinalities does Model 1 occasionally achieve slightly lower tracking error than Model 2, but otherwise the three-factor specification dominates. Moreover, the nominal model appears highly sensitive to changes in cardinality, with strong fluctuations in its tracking error, while the five-factor model shows a smooth and consistent reduction in error as cardinality increases. In contrast, Model 2 does not display a clear relationship between error and portfolio size. Overall, these findings confirm the stabilizing effect of robustness, with Model 3 emerging as the most effective specification in minimizing deviations from the benchmark out-of-sample.




Figure 4.9: Tracking ratio out-of-sample




Figure 4.10: Tracking error out-of-sample

4.5 Comparative Analysis: In-Sample vs. Out-of-Sample

Figures 4.11, 4.12, 4.13 provide a direct comparison of the portfolio performance metrics across different cardinalities q , reported both in-sample and out-of-sample. While in-sample results sometimes show mixed trends, out-of-sample Model 3 consistently outperforms the others. Table 4.1 complements the plots by summarizing returns, variances, and Sharpe ratios for the benchmark (S&P 500) and the different models. For clarity of exposition, results are reported only for $q \in \{20, 50, 100\}$, which are representative of small, medium, and large portfolio sizes among the tested values.




Figure 4.11: Comparison in-sample/out-of-sample return




Figure 4.12: Comparison in-sample/out-of-sample variance




Figure 4.13: Comparison in-sample/out-of-sample Sharpe ratio

Model	In-sample			Out-of-sample		
	Return	Variance	Sharpe Ratio	Return	Variance	Sharpe Ratio
S&P500	0.140	0.00047	1.010	0.116	0.00007	2.159
q = 20						
Model 1	0.981	0.00254	3.087	0.186	0.00100	0.919
Model 2	0.854	0.00152	3.468	0.319	0.00078	1.804
Model 3	0.660	0.00124	2.966	0.337	0.00072	1.988
q = 50						
Model 1	0.960	0.00245	3.075	0.186	0.00097	0.942
Model 2	0.855	0.00153	3.464	0.316	0.00078	1.783
Model 3	0.666	0.00124	2.999	0.333	0.00070	1.990
q = 100						
Model 1	0.922	0.00231	3.041	0.187	0.00089	0.987
Model 2	0.851	0.00152	3.458	0.308	0.00078	1.745
Model 3	0.678	0.00120	3.102	0.323	0.00066	1.994

Table 4.1: Return, variance and Sharpe ratio in-sample vs out-of-sample for selected portfolio sizes

4.6 Rolling Window Analysis

In addition to the static out-of-sample setting, a dynamic evaluation is performed to better capture the evolution of portfolio performance over time. While the static framework relies on fixed portfolio weights throughout the test horizon, the dynamic approach incorporates periodic rebalancing and parameter re-estimation. This procedure aims to mimic a more realistic investment environment, where market conditions change and portfolio allocations must adapt accordingly.

The analysis is carried out over the extended horizon 2020–2023, with quarterly rebalancing. At each rebalancing date, the models are re-estimated using the most recent historical window, updating parameters such as factor exposures and correlations. This rolling window mechanism enables the strategies to adapt to time-varying market conditions and provides a more robust evaluation of their performance.

Performance is evaluated with the same set of indicators—return, volatility, and the Sharpe ratio—so that results can be directly compared between the static and dynamic settings, with the rolling window framework providing a natural extension of the static approach and enabling a more comprehensive evaluation of the strategies.

4.6.1 Rolling Window Setup

For this analysis, a rolling window of 12 months is adopted, with portfolio rebalancing taking place every three months. At each step of the evaluation, the preceding 12 months of data are used both for parameter estimation and for portfolio optimization. Specifically, the procedure begins with the initial training window from December 31, 2019 to December 31, 2020, on which the portfolio weights of all optimization models are computed. These weights are then applied to the immediately subsequent 12-month out-of-sample period in order to evaluate the standard performance metrics. The window is then shifted forward by three months (e.g., March 31, 2020 to March 31, 2021), at which point the covariance matrix and portfolio weights are re-estimated, and the resulting allocations are tested on the following 12-month investment horizon. This process continues until the final portfolio is tested over the period December 2022 to December 2023.

Such a rolling procedure is computationally demanding, since for each training window all parameters and all three optimization models must be re-estimated. To reduce the computational burden, the analysis is restricted to a limited set of portfolio sizes, namely $q = \{30, 40, 50\}$, and to eight rebalancing intervals in total.

Although this dynamic evaluation is more time-consuming than the static one, it provides deeper insights into the models' ability to continuously adjust to evolving market conditions, thereby improving their responsiveness to unforeseen economic

events. Table 4.2 reports a selection of the training and corresponding out-of-sample test windows used in the rolling window analysis. The same procedure was applied to all intermediate windows not shown here.

Training Window	Test Window
2020-03-31 → 2021-03-31	2021-04-01 → 2022-03-31
2020-06-30 → 2021-06-30	2021-07-01 → 2022-06-30
2020-09-30 → 2021-09-30	2021-10-01 → 2022-09-30
...	...
2021-12-31 → 2022-12-31	2023-01-01 → 2023-12-31

Table 4.2: Examples of training and test windows in the rolling window analysis.

4.6.2 Return Out-of-Sample

When comparing the ex-post returns of the three strategies, the contrast between the nominal and robust formulations is pronounced (see Figure 4.14). The nominal model (Model 1) behaves most aggressively: it frequently oscillates between sharp losses and swift gains. Although it occasionally outperforms the benchmark, this comes at the expense of pronounced drawdowns.

The robust models (Models 2 and 3) exhibit more conservative dynamics. Model 2 often falls short of the benchmark, particularly in the early and late stages of the evaluation period, and its recovery from drawdowns is slower. Model 3 delivers a more balanced outcome: its returns are less volatile than those of Model 1 and avoid the persistent underperformance of Model 2. In several intervals it tracks the index more closely, suggesting that the inclusion of the two additional factors enhances the explanatory power of the model and allows the robust set to represent the return structure more effectively.

These patterns remain qualitatively unchanged across all tested cardinalities ($q = 30, 40, 50$): changes in q affect the magnitude of returns but not the relative ranking of the three models.

4.6.3 Variance Out-of-Sample

The ex-post variance displayed in Figure 4.15 complements the return analysis by highlighting the risk differences among the formulations. The nominal model (Model 1), which generates the most extreme return swings, also exhibits the highest variance throughout the sample. Early-window peaks in variance underscore the instability produced by the lack of protection against estimation error: the aggressive return profile coincides with elevated volatility.




Figure 4.14: Return rolling window out-of-sample

Both robust models produce much smoother variance paths. Models 2 and 3 remain well below the nominal portfolio's variance for nearly all windows, demonstrating that robustness effectively dampens risk exposure. Across the sample, five factor model is consistently less volatile than three factor model, consistent with the more balanced return profile observed above. These conclusions remain essentially unchanged when varying the cardinality ($q = 30, 40, 50$): while small quantitative adjustments occur, the gap between the nominal and robust approaches persists throughout.

4.6.4 Sharpe ratio Out-of-Sample

Figure 4.16 shows rolling Sharpe ratios, which combine return and risk into a single metric. The S&P 500 typically outperforms the three strategies on a risk-adjusted basis, particularly in later periods where its Sharpe ratio rises above 1. This indicates that, despite capturing some market dynamics, none of the optimized portfolios consistently match the benchmark's efficiency.

Among the strategies, the nominal model (Model 1) is the most erratic: it occasionally produces Sharpe ratios comparable to the benchmark but also experiences




Figure 4.15: Variance rolling window out-of-sample

deep declines into negative territory, reflecting the instability induced by its higher variance. The robust formulations are steadier: Model 3 most often attains higher Sharpe ratios than Model 2 and occasionally narrows the gap with the benchmark, notably in the mid-sample period. Model 2 rarely achieves strong risk-adjusted performance. As with returns and variance, the observed ranking is not sensitive to the choice of q . Different cardinalities slightly affect levels but leave the overall Sharpe dynamics of the three models unaltered.

4.6.5 Tracking ratio Out-of-Sample

The tracking ratios in Figure 4.17 provide further insight into the models' ability to replicate the benchmark. The nominal formulation (Model 1) proves to be the least reliable: its values fluctuate widely, at times approaching one but also dropping well below 0.6 in some windows. This instability mirrors the volatility observed in returns and variance, confirming that the absence of robustness leads to poor and inconsistent index replication despite occasional strong episodes.

The robust models achieve more stable and accurate tracking. The three-factor specification (Model 2) delivers reasonably high ratios, though with notable dips




Figure 4.16: Sharpe ratio rolling window out-of-sample

that reveal some difficulty in capturing all market movements. The five-factor robust formulation (Model 3) performs more strongly: its tracking ratios are consistently higher than Model 2 and often close to one, particularly in the later periods. This suggests that the additional factors enhance the explanatory power of the model and allow the portfolio to follow the index more closely. As with the other performance measures, this ranking is practically unaffected by the choice of cardinality.

4.6.6 Tracking error Out-of-Sample

The tracking errors reported in Figure 4.18 further emphasize the gap between the nominal and the robust formulations. Model 1 exhibits the largest and most volatile errors, confirming that its instability translates into poor index replication. In contrast, both robust formulations maintain substantially lower errors throughout the sample.

Between the two, Model 3 consistently achieves the smallest deviations from the benchmark, particularly in the later windows, while Model 2 occasionally experiences modest spikes. This suggests that the five-factor specification not only






Figure 4.17: Tracking ratio rolling window out-of-sample


improves tracking ratios but also provides tighter benchmark replication in absolute terms. As in previous analyses, changes in q have only marginal effects and do not modify the relative performance of the three strategies.



(a) $q=30$



(b) $q=40$



(c) $q=50$

Figure 4.18: Tracking error rolling window out-of-sample

Chapter 5

Computational Experiments - Problem 2

In this chapter, we present the analysis of the results obtained using the models within the mixture-of-distributions framework. The same performance metrics introduced in the previous chapter are adopted here, along with the same dataset and preprocessing procedures.

The main difference lies in the analysis phase: in this case, we focus exclusively on the out-of-sample evaluation, without rebalancing and without parameter re-estimation. This choice is motivated by the implementation complexity of the model, which requires the use of fix-and-relax strategies, making it unsuitable for a rolling-window analysis with periodic rebalancing. Consequently, the out-of-sample period is extended from 31 December 2020 to 31 December 2023, while the in-sample period remains from 31 December 2019 to 31 December 2020. Since the in-sample results do not provide meaningful insights in this context, they are omitted, and the discussion focuses entirely on the out-of-sample performance.

In particular, for the model 3.30 and 3.31 the analysis are conducted only for $\tau = 1$. All other parameters follow the specifications in Kang et al. (2022), specifically: $k = 5\%$ on a yearly basis (corresponding to $k = 1.9841 \times 10^{-4}$ on a daily basis) and $\rho = 0.05$. Furthermore, the model accounts for transaction costs exactly as in Kang et al. (2022), defined by the function:

$$c(x) = \sum_{i=1}^n |x_i - x_i^0|,$$

where c_i denotes the unit transaction cost for the i -th asset, and x_i^0 represents the initial position held in the i -th stock at the beginning of the investment period. We set $x_i^0 = 0$, $\forall i = 1, \dots, n$, since no tracking portfolio is available at the starting point, and fix $c_i = 0.01$.

Accordingly, the feasible set of portfolios is defined as:

$$\mathcal{X} = \{x \in \mathbb{R}^n \mid e^\top x + c(x) = 1, x_i \geq 0 \ \forall i = 1, \dots, n\}.$$

5.1 Parameter Estimation

The goal of this section is to estimate the parameters of a Gaussian mixture model with d components using observed data. Due to the non-linear structure of the mixture, closed-form maximum likelihood solutions are not available, and iterative procedures are required. Among these, the Expectation-Maximization (EM) algorithm is widely used, as it provides a practical method to approximate the maximum likelihood estimates of the mixture parameters.

Let $\{\xi_1, \dots, \xi_N\}$ be an i.i.d. sample of observations. The EM algorithm estimates the means, covariances, and mixing weights of the Gaussian mixture through an iterative scheme alternating between two steps:

Initialization. The algorithm starts by assigning initial values to the parameters of each Gaussian component: mean vectors $\mu_i^{(0)}$, covariance matrices $\Sigma_i^{(0)}$, and mixture weights $\lambda_i^{(0)}$, for $i = 1, \dots, d$. An initial evaluation of the likelihood is also computed.

E-step. At iteration k , the algorithm evaluates the probability that each observation belongs to component i , given the current parameters. These probabilities, usually called *responsibilities*, are computed as

$$\gamma_{ji}^{(k)} = \frac{\lambda_i^{(k)} \varphi(\xi_j \mid \mu_i^{(k)}, \Sigma_i^{(k)})}{\sum_{\ell=1}^d \lambda_\ell^{(k)} \varphi(\xi_j \mid \mu_\ell^{(k)}, \Sigma_\ell^{(k)})}, \quad j = 1, \dots, N, \ i = 1, \dots, d,$$

where $\varphi(\cdot \mid \mu, \Sigma)$ denotes the multivariate Gaussian density.

M-step. Using the responsibilities, the parameters are updated as follows:

$$\begin{aligned} \mu_i^{(k+1)} &= \frac{1}{N_i} \sum_{j=1}^N \gamma_{ji}^{(k)} \xi_j, \quad \forall i = 1, \dots, d, \\ \Sigma_i^{(k+1)} &= \frac{1}{N_i} \sum_{j=1}^N \gamma_{ji}^{(k)} (\xi_j - \mu_i^{(k)}) (\xi_j - \mu_i^{(k)})^\top, \quad \forall i = 1, \dots, d, \\ \lambda_i^{(k+1)} &= \frac{N_i}{N}, \quad \forall i = 1, \dots, d, \end{aligned}$$

where $N_i = \sum_{j=1}^N \gamma_{ji}^{(k)}$.

Likelihood evaluation. At the end of each iteration, the log-likelihood function is recalculated:

$$\log L^{(k+1)} = \sum_{j=1}^N \ln \left(\sum_{i=1}^d \lambda_i^{(k+1)} \varphi(\xi_j | \mu_i^{(k+1)}, \Sigma_i^{(k+1)}) \right).$$

The process continues until the increase in likelihood between two consecutive iterations falls below a predefined tolerance threshold.

To reduce the sensitivity of the EM algorithm to the choice of initial parameters, as suggested in Kang et al. (2022), we apply the K-means algorithm to generate the starting values for the procedure.

Following Yan and Han (2019) and Kang et al. (2022), we assume that $d = 3$ components are sufficient to describe the distribution of daily stock returns. The three Gaussian components are interpreted as representing bear, neutral, and bull market regimes.

5.2 Out-of-sample Analysis


This section reports the out-of-sample performance of the benchmark (S&P500), the nominal Enhanced Index Tracking model (EIT), and its robust counterpart (DEIT). The evaluation is static, meaning that portfolios were constructed once and held without rebalancing.

5.2.1 Return

Figure 5.1 shows the out-of-sample portfolio returns. The trajectories are closely aligned with the benchmark, but some distinctions can be observed. During the market downturn in mid-2022, the DEIT portfolio experiences slightly smaller losses compared to EIT, while in subsequent recoveries its upside is somewhat more contained. These shifts indicate a more conservative profile for the robust specification, which trades part of the potential gains for reduced exposure to adverse scenarios. In the final periods, both models converge very closely to the index, showing that the overall tracking ability is preserved.

5.2.2 Variance

Figure 5.2 reports the portfolio variance over time. All three lines co-move strongly, with spikes around mid-2022 and early 2023, consistent with episodes of market turbulence. The DEIT portfolio generally exhibits marginally lower variance than EIT, particularly during the most volatile quarters. Although the differences are not large, they are systematic and suggest that the robust design provides a stabilizing



effect, especially under higher uncertainty, without deviating from the benchmark in calmer phases.




Figure 5.2: Variance out-of-sample

5.2.3 Sharpe ratio

Figure 5.3 presents the Sharpe ratios. The three series display the expected pattern, with positive values during stable or bullish conditions and sharp drops in periods of stress, such as mid-2022. In several periods, most notably the beginning and the end of the evaluation, the DEIT portfolio attains slightly higher Sharpe ratios than EIT, indicating that its variance reduction is not achieved at the expense of returns. In downturns, differences remain limited, which shows that the model cannot fully avoid systemic risk. Overall, the three series remain closely aligned, suggesting that the DEIT strategy preserves the general dynamics of the benchmark while providing marginal improvements in selected periods.

5.2.4 Tracking Ratio

Figure 5.4 illustrates the out-of-sample portfolio tracking ratios relative to the S&P500 benchmark. Both the EIT and DEIT portfolios stay generally close to the ideal value of 1, confirming the effectiveness of the enhanced index tracking approach.




Figure 5.3: Sharpe ratio out-of-sample

Deviations emerge during periods of market stress or recovery. The EIT portfolio occasionally overshoots the benchmark, reflecting its more aggressive exposure to market dynamics. In contrast, the DEIT portfolio tends to remain more conservative, with ratios slightly below the benchmark in downturns and a more moderate response in upward phases. These patterns highlight the stabilizing effect of the robust specification. It should also be noted that, since no cardinality constraint is imposed, the nominal EIT portfolio contains more than 200 stocks, while the robust DEIT is restricted to only 50 assets. This structural difference suggests that, once other transaction costs and cardinality limits are introduced, the relative performance and computational burden may change considerably. The richer composition of the EIT portfolio explains its closer proximity to the ideal tracking ratio and anticipates the lower tracking error discussed in the next subsection.




Figure 5.4: Tracking ratio out-of-sample

5.2.5 Tracking Error

Figure 5.5 shows the out-of-sample portfolio tracking error with respect to the S&P500. The nominal model consistently achieves a lower error than its robust counterpart across the entire evaluation period. This result is consistent with the earlier observation regarding the larger number of assets included in the

EIT portfolio, which naturally allows for tighter replication of the benchmark. Nevertheless, in both cases the error remains of the order of 10^{-5} , which represents a very high level of accuracy. For the DEIT portfolio, two pronounced peaks can be observed in correspondence with the periods of highest variance, confirming its relatively more cautious allocation under turbulent conditions.




Figure 5.5: Tracking error out-of-sample

Chapter 6

Conclusion

This thesis has examined the problem of enhanced index tracking, with the S&P 500 serving as the benchmark index. Two alternative modeling frameworks were investigated: a factor-based approach relying on the Fama–French multi-factor models, and a mixture-of-distributions approach that explicitly accounts for regime changes in asset returns.

In the first framework, the analysis was conducted both in a static setting and through a rolling-window procedure. The results showed that, even though computational limitations prevented an extensive exploration of the parameter space—particularly with respect to the choice of the volatility parameter σ and the target tracking error—clear patterns emerged. Specifically, adopting the five-factor Fama–French specification, as opposed to the simpler three-factor model, led to more robust and stable solutions over time. Similarly, the nominal formulation consistently provided allocations closer to the benchmark than its robust counterpart, even though neither specification was able to fully reproduce the performance of the index. The performance gap is primarily due to the relaxations applied to obtain a tractable formulation, which induce sub-optimality relative to the original problem. Nevertheless, the results highlight the relevance of the factor-based framework and the potential benefits of extending the factor structure beyond the standard three-factor model.

The second framework relied on a mixture-of-distributions approach, where asset returns were modeled through a Gaussian mixture capturing different market regimes. In this case, the empirical evaluation was restricted to a static, multi-year setting due to the computational burden associated with fix-and-relax strategies, which made rolling-window rebalancing unfeasible. The results highlighted a clear distinction between the nominal and robust specifications. The nominal model, benefiting from a larger number of assets, naturally achieved lower tracking errors and a tracking ratio closer to one. Nevertheless, the robust counterpart tended to produce solutions with fewer assets, not due to explicit cardinality constraints,

but as a natural outcome of the model itself. Interestingly, despite this sparser composition, the robust portfolio delivered slightly better performance in terms of return, variance, and Sharpe ratio. This suggests that the robust approach, while less precise in tracking, may offer a more efficient trade-off between risk and return. Furthermore, one may argue that in settings where additional constraints—such as transaction costs or explicit cardinality limits—are introduced, the robust specification could prove even more advantageous, as its parsimonious structure would better accommodate the restrictions typically faced in practical portfolio management.

Taken together, the findings illustrate both the potential and the limitations of the proposed frameworks. On the one hand, they confirm the feasibility of constructing enhanced tracking portfolios that can preserve proximity to the benchmark while improving stability or risk-adjusted performance under certain conditions. On the other hand, they underline the significant computational challenges associated with robust and distributionally aware models, which restrict the scope of parameter exploration and prevent the implementation of dynamic rebalancing schemes. Despite these constraints, the thesis provides evidence that incorporating richer factor structures and distributional considerations can lead to meaningful improvements in tracking performance, while also opening directions for further refinement in more computationally tractable settings.

Bibliography

- Jorion, (2002). «Enhanced index funds and tracking error optimization». In: (cit. on p. 1).
- Weng, and Wang (2017). «Do enhanced index funds truly have enhanced performance? Evidence from the chinese market». In: (cit. on p. 1).
- Kwon, and Wu (2017). «Factor-based robust index tracking». In: (cit. on pp. 1, 2, 6, 8–10, 12, 19).
- Kang, Yao, Li, and Li (2022). «Robust enhanced index tracking problem with mixture of distributions». In: (cit. on pp. 1, 2, 6, 13–15, 40, 42).
- Canakgoz, and Beasley (2009). «Mixed-integer programming approaches for index tracking and enhanced indexation». In: (cit. on p. 4).
- Guastaroba, Speranza, Mansini, and Ogryczak (2020). «Enhanced index tracking with CVaR-based ratio measures». In: (cit. on pp. 4, 7).
- Beasley, Meade, and Chang (2003). «An evolutionary heuristic for the index tracking problem». In: (cit. on p. 4).
- Mezali, and Beasley (2013). «Quantile regression for index tracking and enhanced indexation». In: (cit. on p. 4).
- Valle, Meade, and Beasley (2014). «Absolute return portfolios». In: (cit. on p. 5).
- Wu, Kwon, and Costa (2017). «A constrained cluster-based approach for tracking the S&P500 index». In: (cit. on p. 5).
- Dose, and Cincotti (2005). «Clustering of financial time series with application to index and enhanced index tracking portfolio». In: (cit. on p. 5).
- Lejeune, and Samatlı-Paç (2013). «Construction of Risk-Averse Enhanced Index Funds». In: (cit. on p. 5).
- Bruni, Cesarone, Scozzari, and Tardarella (2012). «A new stochastic dominance approach to enhanced index tracking problems». In: (cit. on p. 5).
- Sharma, Agrawal, and Mehra (2017). «Enhanced indexing for risk averse investors using relaxed second order stochastic dominance». In: (cit. on p. 5).
- Roman, Mitra, and Zverovich (2013). «Enhanced indexation based on second-order stochastic dominance». In: (cit. on p. 5).
- Li, Sun, and Bao (2011). «Enhanced index tracking based on multi-objective immune algorithm». In: (cit. on p. 6).

- Bruni, Cesarone, Scozzari, and Tardella (2015). «A linear risk-return model for enhanced indexation in portfolio optimization». In: (cit. on p. 6).
- Sharpe, (1966). «Mutual fund performance». In: (cit. on pp. 6, 17).
- Sortino, and Price (1994). «Performance measurement in a downside risk framework». In: (cit. on p. 6).
- Meade, and Beasley (2011). «Detection of momentum effects using an index out-performance strategy». In: (cit. on p. 6).
- Guastaroba, Speranza, Mansini, and Ogryczak (2016). «Linear Programming Models based on Omega Ratio for the Enhanced Index Tracking Problem». In: (cit. on p. 6).
- Keating, and Shadwick (2002). «A universal performance measure». In: (cit. on p. 6).
- Goldfarb, and Iyengar (2003). «Robust Portfolio Selection Problems». In: (cit. on pp. 9, 10, 19).
- Yan, and Han (2019). «Empirical distributions of stock returns: Mixed normal or kernel density?» In: (cit. on p. 42).