UNIVERSITÀ
DEGLI STUDI
DI TRIESTE

University of Trieste
Department of Economics, Business, Mathematics and Statistics
Ph.D. Circular Economy – Statistical and Quantitative Methods for
Sustainability and Risk Assessment

Swarm Optimization Algorithms for Sustainable Investing

Ph.D. Candidate:

Dott. Filippo Piccotto

XXXVII Cycle

Supervisor

Prof. Massimiliano Kaucic

University of Trieste

Co-supervisor

Prof.ssa Paola Rossi

University of Trieste

Ph.D. Coordinator

Prof. Guido Bortoluzzi

University of Trieste

Academic Year 2023/2024

UNIVERSITÀ DEGLI STUDI DI TRIESTE

XXXVII CICLO DEL DOTTORATO DI RICERCA IN

CIRCULAR ECONOMY
STATISTICAL AND QUANTITATIVE METHODS FOR SUSTAINABILITY AND RISK ASSESSMENT

Swarm Optimization Algorithms for Sustainable Investing

Settore scientifico-disciplinare: STAT-04/A

DOTTORANDO / A
FILIPPO PICCOTTO

COORDINATORE
PROF. GUIDO BORTOLUZZI

SUPERVISORE DI TESI
PROF. MASSIMILIANO KAUCIC

CO-SUPERVISORE DI TESI
PROF. PAOLA ROSSI

ANNO ACCADEMICO 2023/2024

Are we Human?
Or are we Dancers?
My sign if vidal
My hands are cold

And I'm on my knees
Looking for the answers
Are we Human
Or are we Dancers?

Contents

Abstract

This thesis presents four research papers that were published during my PhD period.

The main theme that links these publications is portfolio optimization, which has been extensively investigated by analyzing models in the literature and exploring the most recent trends. The thesis is divided into two parts.

In the first part, I begin with an introduction about heuristic algorithms, and I illustrate some basics of the particle swarm optimizer (PSO). Then, in the presented chapters, I propose several single-objective portfolio selection models considering different functions to maximize that extend the Sharpe ratio formulation to include the information obtained from the third- and fourth-order moments of the distribution of financial returns. Given the great interest of institutions, managers, and investors in the practical application of these models, I consider several real-world constraints that increase the complexity of the optimization problems. Consequently, I use meta-heuristic swarm algorithms, starting from the PSO and getting to the latest state-of-the-art proposals, developing novel ad-hoc algorithms to find feasible solutions for our portfolio problems. Subsequently, I conduct experimental analyses using real-world data sets from European and World equity market indexes. On the one hand, I analyze the solving capacity of the proposed algorithms, comparing them with some state-of-the-art competitors. On the other hand, I evaluate the profitability of the proposed portfolio strategies.

After a short introduction about the concept of sustainable investment, in the second part of the thesis, my focus moves towards the topic of including the information related to ESG (environmental, social, and governance) criteria in the asset allocation models. To extend the classical mean-risk framework, I consider a tri-objective optimization model that maximizes portfolio expected return and ESG score while minimizing a risk measure. In addition, I expand my interest in the topic of behavioral finance, specifically trying to understand how to integrate the concept of sustainable investing within the behavioral finance paradigms. To do this, I consider a portfolio selection model based on the cumulative prospect theory introduced by Kahneman and Tversky in 1979. Moreover, I propose a bi-objective optimization problem that considers two s-shaped utility functions as objectives. In addition, instead of considering the ESG information using a deterministic score, I introduce the concept of sustainability gross returns as the counterpart of financial returns on the sustainability side. In the computational analyses, I investigate how adding the ESG information changes the ex-post performance of the considered portfolio strategies, particularly focusing on the possibility of having better risk control during periods of higher market volatility.

Sommario

Lo scopo della presente tesi è presentare i risultati di quattro lavori che sono stati pubblicati durante il tre anni di dottorato.

Il tema principale che lega le pubblicazioni è l'ottimizzazione di portafoglio, argomento che è stato ampiamente approfondito analizzando i modelli presenti nella letteratura ed esplorando le tendenze più recenti. La tesi è divisa in due parti.

Nella prima parte, dopo un'introduzione generale sugli algoritmi euristici ed una descrizione dell'algoritmo PSO (particle swarm optimizer), ho proposto diversi modelli di selezione di portafoglio ad obiettivo singolo che utilizzano diverse funzioni da massimizzare che estendono la formulazione dello Sharpe ratio al fine di incorporare al suo interno anche l'informazione ottenuta dai momenti di ordine terzo e quarto della distribuzione dei rendimenti finanziari. Visto il grande interesse da parte delle istituzioni, dei manager, e degli investitori nell'applicazione pratica di questi modelli, ho considerato diversi vincoli che hanno inevitabilmente reso la trattazione di tali modelli più complessa. Di conseguenza, per trovare soluzioni ammissibili ai problemi proposti, ho utilizzato degli algoritmi meta-euristici di tipo swarm, partendo dal PSO (particle swarm optimizer) fino ad arrivare alle innovazioni più recenti, sviluppando dei nuovi algoritmi costruiti ad-hoc per risolvere problemi pratici di portafoglio vincolati. Successivamente, ho condotto delle analisi sperimentali utilizzando dati reali relativi al mercato azionario a livello sia europeo che mondiale. Da una parte, è stata analizzata la capacità risolutoria degli algoritmi proposti, confrontandola con quella di alcuni competitor già presenti e ben noti in letteratura. Dall'altra, è stata comprovata la bontà applicativa dei modelli proposti ed analizzata la profittabilità dei modelli di portafoglio trattati.

Nella seconda parte della tesi, ho rivolto la mia attenzione verso l'incorporazione dell'informazione legata ai criteri ESG (environmental, social, and governance) nei modelli di allocazione del capitale. In un primo lavoro, nell'ottica di estendere la classica architettura media-rischio, ho considerato un modello di ottimizzazione tri-obiettivo che mira a massimizzare gli score ESG di portafoglio ed il rendimento atteso, al contempo minimizzando il rischio. Successivamente, ho cercato di capire come integrare il concetto di investimento sostenibile all'interno dei paradigmi della finanza comportamentale. Per fare ciò, ho considerato un modello di selezione di portafoglio basato sulla cumulative prospect theory introdotta da Kahneman e Tversky nel 1979. Nello specifico, ho analizzato un problema di ottimizzazione bi-obiettivo dove le funzioni obiettivo da massimizzare sono due utilità di tipo s-shaped. Inoltre, invece che considerare l'informazione ESG utilizzando uno score deterministico, ho considerato il concetto di rendimenti di sostenibilità come controparte dei rendimenti finanziari sul piano ESG. Nelle analisi sperimentali, ho investigato come l'aggiunta dell'informazione ESG cambia le performanca finanziaria degli investimenti, concentrandomi in particolare sulla possibilità di avere un miglior controllo del rischio nei periodi di maggior volatilità del mercato.

Per finire, segnalo che la tesi non include una sezione finale, poiché ogni capitolo, essendo un articolo indipendente, si chiude con una propria sezione conclusiva.

Part I

Swarm optimization algorithms for large-scale portfolio selection

In the first part of the PhD thesis, we are going to introduce two works that represent the core of the research activity. The first chapter, entitled*"A hybrid level-based learning swarm algorithm with mutation operator for solving large-scale cardinality-constrained portfolio optimization problems"*, has been published in the journal Information Sciences (Elsevier) in April 2023. The main point of novelty of this paper is the application of a recently proposed optimization algorithm, called LLSO, for solving a large-scale instance of a portfolio optimization problem that uses the modified Sharpe ratio as objective function. We recall that this performance measure is a modification of the Sharpe ratio that is coherent with agent preferences when the risk premia are negative.

The second chapter, *"A constrained swarm optimization algorithm for large-scale long-run investments using Sharpe ratio-based performance measures"*, has been published in Computational Management Sciences (Springer) in June 2024, as part of the collection *"New challenges for optimization methods, models and applications for decision making in complex environments - Stochastic, robust, distributionally robust optimization, artificial intelligence and data analytics"*. In this paper, we propose two asset allocation models that use as objective functions to optimize two modifications of the Sharpe ratio, namely the adjusted for skewness Sharpe ratio and the adjusted for skewness and kurtosis Sharpe ratio, that incorporate information from the third and the fourth moment of returns' distribution in the portfolio selection process. Then, we use an ad-hoc version of the LLSO to solve several instances of the proposed portfolio problems.

Before delving into the details of the first two chapters of the thesis, we will provide a brief overview of the PSO algorithm, since it represents the foundation upon which our proposed algorithms are built and thus it will be mentioned multiple times in the following sections.

Basics of the particle swarm optimizer (PSO)

Ezugwu et al. (2021) distinguish between heuristic and meta-heuristic algorithms in terms of their application range. The former type of solvers is very specific in the search for solution and problem-dependent, while the latter is high-level problem-independent and can be used in a broad range of optimization problems.

In the last decade, meta-heuristics have emerged as a powerful optimization tool in many different disciplines, from engineering design to operations research, due to their flexibility and extensive search capability within a limited computational budget. There is a plethora of such algorithms (see the recent survey by Rajwar et al. (2023), in which a set of 540 different meta-heuristic methods are analysed). In particular, we focus on the so-called population-based stochastic algorithms that use a search operator to gradually minimize (or maximize) the objective function value of the acceptable candidate solutions while maintaining population diversity. The operator employs the information present in the current set of candidate solutions to balance exploration and exploitation. A very general form of this offspring generation process is the following

(Kumar et al. (2023)):

$$\mathbf{y}' = \mathbf{y}_0 + \sum_{k=1}^{n} \boldsymbol{\varepsilon}_k \odot \mathbf{f}_k(\mathbf{y}_{a,k}, \mathbf{y}_{b,k})$$

where $\mathbf{y}' \in \mathbb{R}^d$ is the new generated solution, $\mathbf{y}_0,\ \mathbf{y}_{a,k},\ \mathbf{y}_{b,k} \in \mathbb{R}^d$ are different solutions corresponding to the current population, $\boldsymbol{\varepsilon}_k \in \mathbb{R}^d$ is a vector of parameters, with d representing the number of decision variables, $\mathbf{f}_k$ is the k-th mutation function of the search operator, with $k = 1, \ldots, n$, and $\odot$ denotes the element-wise product of two vectors.

The rules we use to select the so-called base-point $\mathbf{y}_0$, parents $\mathbf{y}_{a,k},\ \mathbf{y}_{b,k}$, scale vectors $\boldsymbol{\varepsilon}_k$, and mutation functions $\mathbf{f}_k$ define a given meta-heuristic algorithm. For instance, the particle swarm optimization algorithm by Clerc and Kennedy (2002), which is one of the most used meta-heuristics in portfolio optimization problems (see Ertenlice and Kalayci (2018)), mimicks the movements of a bird flock, or a fish schooling, that searches for food. The PSO mechanism is based on the communication of information about good solutions through the swarm. In this manner, the particles will tend to move toward better areas in the search space. Let $\mathcal{P}_g = \{\mathbf{y}_1(g), \ldots, \mathbf{y}_{N_{pop}}(g)\}$ be the set of candidate solutions at a given generation g, with $\mathbf{y}_j(g) \in \mathbb{R}^d$, $j = 1, \ldots, N_{pop}$, then the PSO calculates the j-th trial solution of the next generation by employing the current solution $\mathbf{y}_j(g)$ as base-point and considering the following three mutation functions

$$\begin{aligned}
\mathbf{f}_1(\mathbf{y}_j(g), \mathbf{y}_j(g-1)) &= \mathbf{y}_j(g) - \mathbf{y}_j(g-1) \\
\mathbf{f}_2(\mathbf{y}_j(g), \mathbf{pbest}_j(g)) &= \mathbf{pbest}_j(g) - \mathbf{y}_j(g) \\
\mathbf{f}_3(\mathbf{y}_j(g), \mathbf{gbest}(g)) &= \mathbf{gbest}(g) - \mathbf{y}_j(g)
\end{aligned}$$

where $\mathbf{pbest}_j(g)$ represents the best solution corresponding to candidate j and $\mathbf{gbest}(g)$ is the best solution among the $\mathbf{pbest}_j(g)$, with $j = 1, \ldots, N_{pop}$. The corresponding parameters vectors can be written as

$$\begin{aligned}
\boldsymbol{\varepsilon}_1 &= w\, \mathbf{e}_d \\
\boldsymbol{\varepsilon}_2 &= c_1\ (r_{1,1}, \ldots, r_{1,d})^\top \\
\boldsymbol{\varepsilon}_3 &= c_2\ (r_{2,1}, \ldots, r_{2,d})^\top
\end{aligned}$$

where $w \in \mathbb{R}$ denotes the inertia weight, $\mathbf{e}_d \in \mathbb{R}^d$ is a vector of ones, $c_1,\ c_2 \in \mathbb{R}$ are the so-called acceleration factors, $r_{1,i}$ and $r_{2,i}$ are i.i.d. random numbers drawn from a uniform distribution between 0 and 1, with $i = 1, \ldots, d$. In particular:

- the inertia weight w controls the exploration and the exploitation steps of the algorithm by scaling the contribution of the current vector of weight changes;
- $\mathbf{pbest}_j(g) - \mathbf{y}_j(g)$ represents the so-called cognitive component and quantifies how much displacement, starting at its current composition $\mathbf{y}_j(g)$, the candidate solution will need to reach its own best composition $\mathbf{pbest}_j(g)$;

- $\mathbf{gbest}(g) - \mathbf{y}_j(g)$, called social component, measures the distance of the particle $\mathbf{y}_j(g)$ to the best composition $\mathbf{gbest}(g)$ of the entire group of candidate solutions;
- c_1 and c_2 are positive acceleration coefficients used to weight the contribution of the cognitive and social components, respectively.

Algorithm 1: Pseudocode of the PSO (for minimization).

Input : w, c_1, c_2, N_{pop}, G
Output: $\mathbf{y}_{gbest}$

Set $g = 0$
Initialize the swarm $\mathcal{P}_g = \{\mathbf{y}_j(g) \in \mathbb{R}^d \colon j = 1, \ldots, N_{pop}\}$
Calculate the objective function $\varphi(\cdot)$ at the trial solutions in $\mathcal{P}_g$
Set $\mathbf{pbest}_j(g) = \mathbf{y}_j(0)$, with $j = 1, \ldots, N_{pop}$
Find $\mathbf{gbest}(g)$ as the best solution among the $\mathbf{pbest}_j(g)$
while $g < G$ **do**
    for $j = 1$ *to* N_{pop} **do**
        Generate the mutated solution $\mathbf{y}'_j(g)$ using the mutation procedure described above
        Calculate the objective function φ at $\mathbf{y}'_j(g)$
        if $\varphi(\mathbf{y}'_j(g)) < \varphi(\mathbf{y}_j(g))$ **then**
            Set $\mathbf{y}_j(g+1) = \mathbf{y}'_j(g)$
        else
            Set $\mathbf{y}_j(g+1) = \mathbf{y}_j(g)$
        end
    end
    Set $\mathcal{P}_{g+1} = \{\mathbf{y}_j(g+1) \colon j = 1, \ldots, N_{pop}\}$
    Calculate the objective function $\varphi(\cdot)$ at the solutions in $\mathcal{P}_{g+1}$
    Update $\mathbf{pbest}_j(g)$ and $\mathbf{gbest}(g)$
    Set $g = g + 1$
end
Denote the solution with the lowest φ-value as $\mathbf{y}_{gbest}$

References

i. M. Clerc, J. Kennedy: *The particle swarm-explosion, stability, and convergence in a multidimensional complex space*. IEEE Trans. Evol. Comput., **6** (2002), n. 1, 58–73, https://doi.org/10.1109/4235.985692.

ii. O. Ertenlice, C. B. Kalayci: *A survey of swarm intelligence for portfolio optimization: algorithms and applications*. Swarm Evol. Comput., **39** (2018), 36-52, https://doi.org/10.1016/j.swevo.2018.01.009.

iii. A. E. Ezugwu, A.K. Shukla, R. Nath, A.A. Akinyelu, J.O. Agushaka, H. Chiroma, P.K. Muhuri: *Meta-heuristics: a comprehensive overview and classification along with bibliometric analysis*. Artif. Intell. Rev.,**54** (2021), 4237–4316, https://doi.org/10.1007/s10462-020-09952-0.

iv. A. Kumar, S. Das, L. Kong, V. Snášel: *Self-adaptive spherical search with a low-precision projection matrix for real-world optimization*. IEEE Trans. Cyber., **53** (2023), n. 7, 4107–4121, https://doi.org/10.1109/TCYB.2021.3119386.

v. K. Rajwar, K. Deep, S. Das: *An exhaustive review of the metaheuristic algorithms for search and optimization: taxonomy, applications, and open challenges*. Artif. Intell. Rev., **56** (2023), 13187–13257, https://doi.org/10.1007/s10462-023-10470-y.

1 A hybrid level-based learning swarm algorithm with mutation operator for solving large-scale cardinality-constrained portfolio optimization problems

This chapter presents a novel hybrid variant of the level-based learning swarm optimizer (LLSO) for solving large-scale portfolio optimization problems. Since particle swarm optimization (PSO) does not work efficiently as the problem size increases, several authors have designed many variants to improve its performance for solving large-scale optimization problems. In particular, the LLSO algorithm is inspired by the teaching concept that teachers should treat students differently according to their cognitive and learning abilities.

Regarding the model formulation, we extend the classical mean-variance framework by maximizing a modified version of the Sharpe ratio subject to cardinality, box, and budget constraints. These three constraints are simultaneously handled using a projection operator. Further, we implicitly control transaction costs through a rebalancing constraint controlled by a suitable exact penalty function. In addition, we develop an ad hoc mutation operator to modify candidate exemplars in the highest level of the swarm.

We conduct some experiments on three public large-scale data sets. Firstly, we empirically prove that including the mutation procedure improves the accuracy of the solutions. Then, we compare our proposed solver with other variants of the LLSO algorithm and two state-of-the-art swarm optimizers, pointing out the outstanding performance of our hybrid LLSO both in terms of exploration capabilities and solution quality. Finally, we assess the profitability of the portfolio allocation strategy in the last five years using an investable pool of 1119 constituents from the MSCI World Index.

1.1 Introduction

The mean-variance portfolio selection problem developed by Markowitz underpins the modern portfolio theory. According to this approach, the portfolio returns distribution's mean and variance represents the investment profit and the risk, respectively. A portfolio is efficient if it provides the maximum return for a given level of risk or, equivalently, if it has the minimum risk for a given level of return. Therefore, the efficient frontier is provided by the set of optimal mean-variance trade-offs in the risk-return space.
In recent years, many developments of the basic model have been investigated (see, for instance, Guerard (2010)). On the one hand, researchers have defined new risk measures to describe better investor attitudes Kolm et al. (2014). On the other hand, the issues linked to portfolio management have been addressed by introducing several constraints to handle portfolio weights (a detailed list is given in Liagkouras and Metaxiotis (2017)). In this context, some authors have studied the multi-objective formulation of the mean-variance portfolio optimization problem, in which the expected portfolio return is maximized and, at the same time, its variance is minimized (see, for instance, Kaucic (2019)). The scope is to provide heuristics able to generate accurate dotted representations of the set of efficient portfolios in a few iterations.
Instead, in the single-objective counterpart, one usually aims to approximate the efficient frontier by minimizing the risk for different values of the desired mean return of the portfolio. Another standard practice is to optimize a sum of the portfolio mean and variance, weighted by a risk aversion parameter. In this case, one varies the sensitivity of the investor to build up the set of optimal portfolios Crama and Schyns (2003).
A third formulation maximizes the so-called Sharpe ratio, defined the ratio between the investment's excess return with respect to a risk-free rate and its standard deviation, in order to directly identify the best portfolio on the efficient frontier Zhu et al. (2011). From a rational investor's point of view, using the Sharpe indicator in periods of market downturns is questionable because it leads to preferring riskier portfolios. A solution to this issue is to multiply, instead of dividing, this quantity by the standard deviation when the excess rate of return of the portfolio is negative (see Israelsen (2005)). In this way, the performance measure is coherent with agent preferences even if the risk premia are negative. A first application of this modified Sharpe ratio for a passive portfolio optimization problem with an investable universe of almost fifty assets has been proposed in Kaucic et al. (2020).

In this paper, we focus on using the modified Sharpe ratio for active investments with an upper threshold on the portfolio size. We analyze the asset allocation problem from the perspective of an institutional investor who operates in large equity markets composed of hundreds or thousands of constituents and selects a restricted pool of stocks to build up a portfolio with a suitable performance with respect to the benchmark. In our portfolio design, we also consider the following standard real-world constraints.

First, a budget constraint ensures that all the available capital is invested; next, bound constraints prescribe lower and upper bounds on the fraction of capital invested in each asset; finally, a turnover constraint implicitly controls the effect of the transaction costs on the rebalancing phases. The resulting mixed-integer optimization problem belongs to the general family of cardinality-constrained portfolio optimization problems which are NP-hard, and finding possible optimal solutions becomes computationally challenging Moral-Escudero et al. (2006). For this reason, on the one hand, exact methods were proposed to supply optimal solutions, but they demand a significant amount of computation time when the problem size increases (Shaw et al. (2008) and Bertsimas and Shioda (2009)). On the other hand, heuristic approaches can identify approximate and sometimes optimal solutions within reasonable computation time even when the problem size is huge Woodside-Oriakhi et al. (2011).

Many other real-world and engineering applications can also be modeled as highly constrained optimization problems (Das and Prasad (2015) and Singh and Das (2016)). The lack of a universal optimizer Ho and Pepyne (2002) suggests the development of efficient algorithmic designs which exploit the problem-specific features. For this purpose, hyper-heuristic, memetic, and ensemble algorithms have been recently introduced (see Caraffini et al. (2019) and Xue et al. (2022)). It is worth noting that these meta-heuristics are all linked and address the same topic from different perspectives. More specifically, the first type of solvers consists of different search algorithms equipped by a coordination engine that selects and activates the various algorithms for the problem at hand. The second class combines the benefits of ad hoc local search heuristics and multi-agent systems in the algorithmic framework. Finally, an ensemble of strategies involves multiple and complementary search techniques to handle the optimization problem issues.

Within the context of memetic optimizers, swarm optimization algorithms, inspired by the self-organizing interaction among agents, have become popular in portfolio selection theory Ertenlice and Kalayci (2018). In particular, the particle swarm optimization (PSO) algorithm has been widely employed to solve real-world financial problems since its first proposal (Cura (2009), Zhu et al. (2011) and Corazza et al. (2021)) due to its effectiveness in reaching optimal solutions. However, the algorithm above does not work efficiently when the problem size is large, leading to population stagnation and premature convergence Oldewage et al. (2020). To improve PSO performance for large-scale optimization problems, several authors have designed many variants, such as competitive swarm optimizer Cheng and Jin (2015), social learning particle swarm optimizer Cheng and Jin (2015) and level-based learning swarm optimizer (LLSO) Yang et al. (2018). In particular, the latter one has shown better exploitation ability in different environments. For this reason, we propose its use to solve our portfolio optimization problem. LLSO is inspired by the teaching concept that teachers should treat students differently according to their cognitive and learning abilities. Based on that, the general idea of the LLSO is to sort the swarm individuals in ascending order with respect

to their fitness and then separate them into distinct levels. The best individuals are stored at a higher level and are not updated, preserving the most valuable information conveyed in the swarm. Unlike PSO, which uses the historically best positions to update the particles, LLSO employs predominant particles in the current swarm to guide the learning of the worst particles and to enhance the swarm diversity. Thus, particles in lower levels have more individuals in the upper levels to learn from and are focused on exploring the search space; those in higher levels mainly concentrate on the exploitation task. Even though LLSO shows promising capabilities in dealing with large-scale optimization problems, it is overly sensitive to its parameters. To mitigate this influence, an adaptive variant, henceforth ALLSO, has been introduced in Song et al. (2021), which takes advantage of a swarm aggregation indicator to estimate the evolution state of the swarm. Two adaptive adjustment strategies are then applied to identify the best configuration setting for each generation.

Due to the fact that the swarm optimization algorithms are usually blind to the constraints, they have to be equipped with constraint-handling techniques Mezura-Montes and Coello Coello (2011) to be effective in real-world applications. A class of constraint-handling methods widely used in literature is represented by the penalty function methods, where a penalty term reduces the fitness value of the infeasible candidates. However, despite its simplicity, this method usually requires the definition of problem-dependent parameters that significantly impact algorithm performance. To overcome this issue, adaptive penalty techniques have been developed, in which the parameters are automatically set by using information gathered from the violated constraints at the current generation. We refer the reader to Barbosa et al. (2015) for a more exhaustive overview of adaptive penalty techniques.

Therefore, to tackle the presented large-scale cardinality-constrained portfolio optimization problem, our memetic algorithm combines an improved ALLSO with a novel hybrid constraint-handling technique, in which we integrate a projection operator into the self-adaptive penalty scheme developed by Costa et al. (2017). Moreover, to further improve the algorithm's exploitation power and the solutions' quality, we introduce a novel mutation procedure, applied to the best individuals in the first level, which generalizes the one inspected in Kaucic et al. (2020).

Even though similar asset allocation models with real-world constraints have been already studied in the literature, our work represents the first application of the LLSO algorithm for solving a large-scale instance of the problem with the modified Sharpe ratio as the objective function. In addition, to the best of our knowledge, this is the first time such a hybrid constraint-handling technique and the mutation operator are presented and involved in the level-based learning paradigm.

The contributions of the paper can be summarised as follows:

- we develop a memetic algorithm integrating the ALLSO with an ad hoc mutation operator and a novel hybrid constraint-handling technique to solve large-scale cardinality-constrained portfolio optimization problems;

- we use a modified version of the Sharpe ratio as the objective function in active portfolio optimization;
- we introduce a novel initialization procedure for population-based heuristics to effectively address the low degree of feasibility in our portfolio rebalancing problem;
- we validate the profitability of the investment strategy in the last five years using a pool of 1119 assets from a global equity market index.

Let us now give a more precise overview of the paper's contents. The following section describes the investment framework, focusing on the portfolio optimization problem. In Section 1.3, we introduce the developed solver. More precisely, we first explain the adaptive LLSO and then detail the proposed methods, namely the novel mutation operator and the hybrid constraint-handling technique. In the last part of the section, we summarise the entire procedure. In Section 1.4, we show the experimental results; in the last section, we depict the conclusions and future perspectives.

1.2 Portfolio design

1.2.a Investment framework

We consider a frictionless market in which no short selling is allowed, and all investors act as price takers. Assuming that n assets represent the investable universe, a portfolio is identified with the vector of assets weights $\boldsymbol{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$, where $x_i \in \mathbb{R}$ denotes the proportion of capital invested in asset i, with $i = 1, \dots, n$. Let R_i be the random variable which stands for the rate of return of asset i, with expected value μ_i. Hence, the random variable $R_p(\boldsymbol{x}) = \sum_{i=1}^{n} R_i x_i$ indicates the rate of return of portfolio $\boldsymbol{x}$. The expected rate of return of portfolio $\boldsymbol{x}$ is then defined as

$$\mu_p(\boldsymbol{x}) = \sum_{i=1}^{n} x_i \mu_i \tag{1.2.1}$$

and its standard deviation, also called volatility, is given by

$$\sigma_p(\boldsymbol{x}) = \sqrt{\sum_{i=1}^{n} \sum_{j=1}^{n} c_{ij} x_i x_j} \tag{1.2.2}$$

where $(C)_{ij} = c_{ij}$ is the covariance between stocks i and j, with $i, j = 1, \dots, n$.
Since investors perceive large deviations from the portfolio mean value as damaging, (1.2.2) represents the so-called portfolio risk.

In such a setting, the portfolio choice is made only with respect to the expected portfolio rate of return and the portfolio risk, as stated in the following definition.

Definition 1. *Given two portfolios $\boldsymbol{x}, \boldsymbol{y}$, we say that $\boldsymbol{x}$ is preferred to $\boldsymbol{y}$ if and only if $\mu_p(\boldsymbol{x}) \geq \mu_p(\boldsymbol{y})$ and $\sigma_p(\boldsymbol{x}) \leq \sigma_p(\boldsymbol{y})$, with at least one strict inequality.*

In other words, an investor prefers one portfolio to another if it has a higher expected rate of return and lower risk. This decision making approach is known as mean-variance analysis.

1.2.b Objective function

According to the mean-variance analysis, we can use the so-called Sharpe ratio to identify the best investment among efficient portfolios. As stated before, it is defined as the ratio between the excess return of an investment with respect to a risk-free rate, r_f, and its standard deviation

$$SR(\boldsymbol{x}) = \frac{\mu_p(\boldsymbol{x}) - r_f}{\sigma_p(\boldsymbol{x})}. \tag{1.2.3}$$

This performance measure evaluates the compensation earned by the investor per unit of both systematic and idiosyncratic risks (Caporin et al. (2014)). Thus, higher values of SR indicate more promising portfolios.

From a theoretical point of view, this choice is justified by the fact that several widely used performance measures are increasing functions of the Sharpe ratio (Schumacher and Eling (2012) and Schumacher and Eling (2012)). Moreover, when the numerator in (1.2.3) is positive, this indicator is coherent with the risk-return profile of a rational investor. From a practical point of view, it can be easily calculated and its interpretation is simpler than most of recently proposed complex performance measures Auer and Schumacher (2013).

However, as pointed out in Israelsen (2005), the reliability of this performance measure decreases when the excess rate of return is negative. In that case, one would prefer higher-risk portfolios using the Sharpe ratio. To overcome this issue, in our portfolio selection problem we adopt as objective function the following modification of (1.2.3), the so-called modified Sharpe ratio

$$MSR(\boldsymbol{x}) = \frac{\mu_p(\boldsymbol{x}) - r_f}{\sigma_p(\boldsymbol{x})^{\operatorname{sign}(\mu_p(\boldsymbol{x}) - r_f)}} \tag{1.2.4}$$

where $\operatorname{sign}(z)$ is the sign function of $z \in \mathbb{R}$. Observe that if the portfolio excess return is non-negative, the modified Sharpe ratio is equal to the Sharpe ratio. Otherwise, it multiplies the portfolio excess return by the standard deviation. In this manner, even in adverse conditions, portfolios with lower risk and higher excess return will be preferred.

1.2.c Constraints

In our portfolio model, we consider the following constraints.

- *Budget*. All the available capital needs to be invested. In terms of portfolio weights, this translates to
$$\sum_{i=1}^{n} x_i = 1. \tag{1.2.5}$$

- *Cardinality*. We assume that the portfolio includes up to k assets, where $k \leq n$. To model the inclusion or the exclusion of the i-th asset in the portfolio, a binary variable δ_i is introduced as
$$\delta_i = \begin{cases} 0, \text{ if asset } i \text{ is excluded} \\ 1, \text{ if asset } i \text{ is included} \end{cases} \tag{1.2.6}$$
for $i = 1, \ldots, n$. The resulting vector of selected assets is $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_n) \in \{0,1\}^n$, and the cardinality constraint can be written as
$$\sum_{i=1}^{n} \delta_i \leq k. \tag{1.2.7}$$

- *Box*. A balanced portfolio should avoid extreme positions and foster diversification. Hence, we impose a maximum and a minimum limit for portfolio weights, that is
$$\delta_i l_i \leq x_i \leq \delta_i u_i, \quad i = 1, \ldots, n \tag{1.2.8}$$
where l_i and u_i are the lower and the upper bounds for the weight of the i-th asset, respectively, with $0 < l_i < u_i \leq 1$ to exclude short sales.

- *Turnover*. To control the effect of the transaction costs in the portfolio rebalancing phases, we consider a portfolio turnover constraint. Let $\boldsymbol{x}_0$ be a vector containing the current portfolio positions Shen et al. (2014). Then, the portfolio turnover constraint is
$$\sum_{i=1}^{n} |x_i - x_{0,i}| \leq TR \tag{1.2.9}$$
where TR denotes the maximum turnover rate, which lies between 0 and 1. Note that if $TR = 0$ rebalancing is not allowed, and more trades are allowed when TR increases.

The pairs $(\boldsymbol{\delta}, \boldsymbol{x}) \in \{0,1\}^n \times \mathbb{R}^n$ that satisfy (1.2.5), (1.2.7), (1.2.8) and (1.2.9) form the feasible set $\mathcal{F}$. Then, our portfolio optimization problem can be written as

$$\begin{aligned} \max_{\boldsymbol{\delta}, \boldsymbol{x}} \quad & MSR(\boldsymbol{x}) \\ \text{s.t.} \quad & (\boldsymbol{\delta}, \boldsymbol{x}) \in \mathcal{F}. \end{aligned} \tag{1.2.10}$$

Remark 2. *To simplify the following treatment, we reformulate our maximization problem into the equivalent minimization problem*

$$\begin{aligned} \min_{\boldsymbol{\delta},\boldsymbol{x}} \quad & f(\boldsymbol{x}) \\ s.t. \quad & (\boldsymbol{\delta},\boldsymbol{x}) \in \mathcal{F} \end{aligned} \tag{1.2.11}$$

where $f(\boldsymbol{x}) = -MSR(\boldsymbol{x})$.

1.3 Optimization algorithm

1.3.a Adaptive level-based learning swarm optimizer

The algorithm evolves a swarm of NP candidate solutions using the so-called level-based population structure Yang et al. (2018), according to which the evolution process is defined as follows.

1. At each iteration g, the individuals in the swarm are first sorted ascending based on their fitness and grouped into NL_g levels, each one containing $LP_g = \lfloor {}^{NP}/_{NL_g} \rfloor$ particles. In the last level, there are $\lfloor {}^{NP}/_{NL_g} \rfloor + NP\%NL_g$ particles.[1] Better individuals belong to higher levels, and a higher level corresponds to a smaller level index. Thus, L_1 represents the best level, while L_{NL_g} is the worst one.

2. To preserve the most valuable information conveyed in the current swarm, individuals belonging to L_1 are not updated and enter directly in the next generation. The p-th particle in level L_l, denoted by $\boldsymbol{x}^{l,p}(g)$, where $l = 3, \ldots, NL_g$ and $p = 1, \ldots, LP_g$, is allowed to learn from two particles $\boldsymbol{x}^{l_1,p_1}(g)$, $\boldsymbol{x}^{l_2,p_2}(g)$ randomly extracted from two different higher levels L_{l_1} and L_{l_2} with $l_1 < l_2$, and p_1 and p_2 are randomly chosen from $\{1, \ldots, LP_g\}$. For $l = 2$, we sample two particles from L_1 in such a way that $\boldsymbol{x}^{l_1,p_1}(g)$ is better than $\boldsymbol{x}^{l_1,p_2}(g)$ in terms of fitness function. Thus, the update rule for particle $\boldsymbol{x}^{l,p}(g)$ is given component-wise by

$$v_i^{l,p}(g+1) = r_1 v_i^{l,p}(g) + r_2 \left(x_i^{l_1,p_1}(g) - x_i^{l,p}(g) \right) + \phi_g r_3 \left(x_i^{l_2,p_2}(g) - x_i^{l,p}(g) \right) \tag{1.3.12}$$

$$x_i^{l,p}(g+1) = x_i^{l,p}(g) + v_i^{l,p}(g+1) \tag{1.3.13}$$

for $i = 1, \ldots, n$, where $v_i^{l,p}(g)$ denotes the i-th component of the velocity of particle p in level L_l at generation g, and r_1, r_2, r_3 are real numbers randomly generated within $[0, 1]$. The parameter $\phi_g \in [0, 1]$ controls the influence of the less performing exemplar $\boldsymbol{x}^{l_2,p_2}(g)$ on $\boldsymbol{v}^{l,p}(g)$.

Based on Song et al. (2021), both the parameters involved in the learning process at generation g, namely NL_g and ϕ_g, are adaptively adjusted based on the evolution state

[1] We denote by $\lfloor x \rfloor$ the floor of x and by $x\%y$ the rest of the division of x by y.

of the swarm by an aggregation indicator, which is defined as

$$s(g) = \frac{\overline{f}_g - f\left(\boldsymbol{x}_{gbest}(g)\right)}{f\left(\boldsymbol{x}_{gbest}(g)\right) + \xi} \tag{1.3.14}$$

where $\overline{f}_g$ is the average fitness of the population at generation g, $f\left(\boldsymbol{x}_{gbest}(g)\right)$ denotes the historically global best fitness up to iteration g, and ξ is a small positive value to avoid zero denominators.

Remark 3. *When $s(g)$ is high, particles are far from the current global best solution. Thus, the swarm is in an exploration phase. On the contrary, when $s(g)$ is low, particles are close to the global best solution $\boldsymbol{x}_{gbest}(g)$ and the swarm is in an exploitation phase.*

To guarantee a control on the number of levels, NL_g takes values in the set $\{NL_{min}, \ldots, NL_{max}\}$, where $NL_{min}, NL_{max} \in \mathbb{N}$ are predefined lower and upper bounds. Moreover, to balance the level selection diversity and the exemplar diversity, NL_g can be modified only when the relative improvement of the global fitness between generation g and generation $g-1$, given by

$$t(g) = \frac{f\left(\boldsymbol{x}_{gbest}(g-1)\right) - f\left(\boldsymbol{x}_{gbest}(g)\right)}{f\left(\boldsymbol{x}_{gbest}(g)\right) + \xi}, \tag{1.3.15}$$

slows down or stops, which corresponds to the cases $t(g) < t(g-1)$ or $t(g) = 0$ respectively. The update of NL_g then follows the rule

$$NL_g = \begin{cases} 2 \cdot NL_{g-1} & \text{if} \quad s(g) < \bar{\delta} \\ \frac{1}{2} \cdot NL_{g-1} & \text{if} \quad s(g) \geq \bar{\delta} \end{cases} \tag{1.3.16}$$

where $\bar{\delta}$ is a threshold in terms of the aggregation indicator to control the adjustment of NL_g.

When NL_g is out of the range, it is adjusted as follows

$$NL_g = \begin{cases} NL_{rand} & \text{if} \quad r < px \\ NL_{max} & \text{if} \quad r \geq px \text{ and } NL_g > NL_{max} \\ NL_{min} & \text{if} \quad r \geq px \text{ and } NL_g < NL_{min} \end{cases} \tag{1.3.17}$$

where NL_{rand} is uniformly sampled from $\{NL_{min}, \ldots, NL_{max}\}$, r is a real number randomly generated within $[0, 1]$, and px is a fixed probability employed to reset NL_g.

The update for ϕ_g is designed in the following way

$$\phi_g = 0.35 + 0.1 \cdot \frac{1}{1 + 10 \cdot s(g)} \tag{1.3.18}$$

where $s(g)$ is the value of aggregation indicator given in (1.3.14). For a theoretical

justification of the formula in Eqn. (1.3.18) and the rationale behind choosing these specific values, we refer the reader to the paper of Song et al. (2021).

A preliminary numerical analysis reveals that a clamping procedure, limiting the magnitude of the velocity $\boldsymbol{v}^{l,p}(g)$, provides a better exploration of the search space compared to the traditional particle swarm optimization method, as evidenced by Oldewage et al. (2017). This function can be written component-wise as

$$v_i^{l,p}(g) = \min\{\max\{v_i^{l,p}(g), v_i^{min}\}, v_i^{max}\} \tag{1.3.19}$$

where v_i^{min} and v_i^{max} are the minimum and the maximum velocity allowed for component i, with $i = 1, \ldots, n$. In the experiments, recalling equation (2.3.6), we set $v_i^{max} = u_i$ and $v_i^{min} = -v_i^{max}$.

1.3.b Mutation operator

Instead of directly moving the individuals of the first level to the next generation, we propose to mutate them using an operator that combines two perturbation strategies properly developed for our portfolio optimization problem.

More specifically, one technique is inspired by the swap operator proposed in Krink et al. (2009) and works as follows. First, we fix the maximum allowed number of non-null positions that could become zero, namely k_{max}^{swap}. Then, for each particle $\boldsymbol{x}^{1,p}(g)$ in level L_1 subject to swapping, we randomly sample from $\{1, \ldots, k_{max}^{swap}\}$ the number k^{swap} of non-null positions that will be set to zero. At this point, for $j = 1, \ldots, k^{swap}$, let a_j and b_j be two randomly chosen positions in $\boldsymbol{x}^{1,p}(g)$, such that $x_{a_j}^{1,p}(g) = 0$ and $l_{b_j} \leq x_{b_j}^{1,p}(g) \leq u_{b_j}$. Thus, the modified individual, $\hat{\boldsymbol{x}}^{1,p}(g)$, is defined component-wise as

$$\hat{x}_i^{1,p}(g) = \begin{cases} x_i^{1,p}(g), & \text{if } i \neq a_j \text{ and } i \neq b_j \\ l_{a_j} + \dfrac{x_{b_j}^{1,p}(g) - l_{b_j}}{u_{b_j} - l_{b_j}}(u_{a_j} - l_{a_j}), & \text{if } i = a_j \\ 0, & \text{if } i = b_j. \end{cases} \tag{1.3.20}$$

In this paper, based on the preliminary experiments, $k_{max}^{swap} = \lfloor 0.05 \cdot k \rfloor$, where k represents the maximum number of assets included in the portfolio.

Remark 4. *This generalisation, allowing multiple swaps at the same time, improves the search capabilities of the original swap operator.*

The other perturbation scheme focuses solely on the non-null components. For each $\boldsymbol{x}^{1,p}(g)$ to be mutated, let $I_+^{1,p}(g) = \{i : x_i^{1,p}(g) > 0\}$ then, for all $i \in I_+^{1,p}(g)$, we define the interval

$$W_i^{1,p}(g) = \left[x_i^{1,p}(g) - \Delta_i(g),\, x_i^{1,p}(g) + \Delta_i(g)\right] \tag{1.3.21}$$

where $\Delta(g) = \left(1 - \frac{g}{g_{max}+1}\right)(\mathbf{u} - \mathbf{l})$, with g_{max} be the maximum allowed number of iterations. The mutated component $\hat{x}_i^{1,p}(g)$ is randomly generated from the interval

$W_i^{1,p}(g) \cap [l_i, u_i]$. For $i \notin I_+^{1,p}(g)$, we set $\hat{x}_i^{1,p}(g) = 0$. By narrowing the range over time, this procedure increases the exploration around the particles in L_1.

Remark 5. *By construction, the solutions modified by both the perturbation operators have at most k non-null positions and satisfy the box constraints.*

For each particle in L_1, the probability of applying the generalised swap operator decreases as the iteration counter increases according to the following rule

$$p_{swap}(g) = \frac{1}{1 + \exp(-0.005 \cdot g)}. \tag{1.3.22}$$

In the initial stages, the proposed mutation favours the global search, using the generalised swap operator to identify the most promising subset of non-null decision variables. With the progress of the generations, the role of the refinement operator increases and, in the late stages, the algorithm focuses primarily on the local search.

The pseudo-code of the developed mutation procedure is reported in Algorithm 2.

Algorithm 2: Mutation procedure

Input : $\boldsymbol{x}^{1,p}(g)$, $\boldsymbol{l}$, $\boldsymbol{u}$, k_{max}^{swap}, g, $\Delta(g)$
Output: $\hat{\boldsymbol{x}}^{1,p}(g)$

Set $\hat{\boldsymbol{x}}^{1,p}(g) = \boldsymbol{x}^{1,p}(g)$
Set $I^0 = \left\{i \colon x_i^{1,p}(g) = 0\right\}$
Set $I^+ = \left\{i \colon l_i \leq x_i^{1,p}(g) \leq u_i\right\}$
Calculate $p_{swap}(g)$ according to (1.3.22)
if $rand() \leq p_{swap}(g)$ **then**
    $k^{swap} \rightarrow \{1, \ldots, k_{max}^{swap}\}$
    for $j = 1$ *to* k^{swap} **do**
        $a_j \rightarrow I^0$
        $b_j \rightarrow I^+$
        $\hat{x}_{a_j}^{1,p}(g) = l_{a_j} + \dfrac{x_{b_j}^{1,p}(g) - l_{b_j}}{u_{b_j} - l_{b_j}}(u_{a_j} - l_{a_j})$
        $\hat{x}_{b_j}^{1,p}(g) = 0$
    end
else
    for i *in* I^+ **do**
        $lb = \max\left(x_i^{1,p}(g) - \Delta_i(g), l_i\right)$
        $ub = \min\left(x_i^{1,p}(g) + \Delta_i(g), u_i\right)$
        $\hat{x}_i^{1,p}(g) \rightarrow [lb, ub]$
    end
end

1.3.c Solution coding and hybrid constraint-handling procedure

Let us introduce some notation. Let $\mathcal{C}_i$ denote a closed convex subset of $\mathbb{R}_+$, with $i = 1, \dots, n$, and $K = \{i_1, \dots, i_k\}$ be any subset of indices of $I = \{1, \dots, n\}$ with cardinality $k \in \mathbb{N}$, so that $I \setminus K$ is its complement in I. For all $\boldsymbol{x} \in \mathbb{R}^n$, let $\boldsymbol{x}_K$ be defined component-wise as

$$\boldsymbol{x}_{K,i} = \begin{cases} x_i, & \text{if } i \in K \\ 0, & \text{if } i \in I \setminus K \end{cases} \tag{1.3.23}$$

and let $\pi_K \colon \mathbb{R}^n \to \mathbb{R}^k$ be the projection such that $\pi_K(\boldsymbol{x}) = (x_{i_1}, \dots, x_{i_k})$.

We start by presenting the following proposition (in this regard, see also Zhang et al. (2019)).

Proposition 6. *Let $\boldsymbol{y} \in \mathbb{R}^n$, with $n \geq 2$. Then, the optimal K for the problem*

$$\min_{\boldsymbol{x}_K \colon x_i \in C_i} \frac{1}{2} \|\boldsymbol{x}_K - \boldsymbol{y}\|^2 \tag{1.3.24}$$

is the set K^* of indices corresponding to the k largest components of $\boldsymbol{y}$.

The proof of this result is reported in Appendix A.

In other words, the proposition states that $\boldsymbol{x}_{K^*}$ is the vector with at most k non-null components which has minimum Euclidean distance from $\boldsymbol{y}$ among all $\boldsymbol{x}_K$, with $K \subset I$ of cardinality k.

Thanks to this projection, which implicitly enforces cardinality fulfillment, we can remove the vector of binary variables $\boldsymbol{\delta}$ from the coding scheme of the solutions and we reformulate the portfolio optimization problem (2.3.8) only in terms of $\boldsymbol{x}_{K^*}$. To this end, we introduce the set

$$\mathcal{B} = \left\{ \boldsymbol{x} \in \mathbb{R}^n \colon x_i = 0 \text{ or } x_i \in [l_i,\, u_i] \text{ for } i \in K^*,\, x_i = 0 \text{ for } i \in I \setminus K^*,\, \sum_{i=1}^{n} x_i = 1 \right\}, \tag{1.3.25}$$

that is the set of the points satisfying all the constraints apart from the turnover condition. Further, let $\psi(\boldsymbol{x})$ represent the value of the turnover function at $\boldsymbol{x}$, which is given by

$$\psi(\boldsymbol{x}) = \sum_{i=1}^{n} |x_i - x_{0,i}| - TR. \tag{1.3.26}$$

Then, the constrained optimization problem can be rewritten as

$$\begin{aligned} \min_{\boldsymbol{x} \in \mathcal{B}} \quad & f(\boldsymbol{x}) \\ \text{s.t.} \quad & \psi(\boldsymbol{x}) \leq 0. \end{aligned} \tag{1.3.27}$$

The following proposition, whose proof is given in Appendix A, establishes the equivalence between problem (1.3.27) and the mixed-integer optimization problem (2.3.8).

Proposition 7. *We assume that* $(\boldsymbol{\delta}^*, \boldsymbol{x}^*)$ *is a global solution to problem* (2.3.8)*, then* $\boldsymbol{x}^*_{K^*}$ *is a global solution to problem* (1.3.27)*. Conversely, if* $\boldsymbol{x}^*$ *is a global solution to* (1.3.27)*, then* $(\boldsymbol{\delta}^*, \boldsymbol{x}^*)$ *is a global solution to* (2.3.8)*, with*

$$\delta_i^* = \begin{cases} 1 & \text{if} \quad i \in K^* \\ 0 & \text{otherwise.} \end{cases}$$

As previously observed, the standard ALLSO algorithm can only deal with unconstrained problems; thus, we propose to incorporate a hybrid constraint-handling technique in order to solve problem (1.3.27).

The building block of our procedure is based on the following lemma.

Lemma 8. *Let* $\boldsymbol{l} = (l_1, \dots, l_n)$ *and* $\boldsymbol{u} = (u_1, \dots, u_n)$ *be such that* $l_i \leq u_i$ *for* $i = 1, \dots, n$*. Let* $\boldsymbol{y} \in \mathbb{R}^n$ *and define* $[\boldsymbol{l}, \boldsymbol{u}] = \{\boldsymbol{x} \in \mathbb{R}^n \colon l_i \leq x_i \leq u_i\}$*. Then, the orthogonal projection of* $\boldsymbol{y}$ *onto* $[\boldsymbol{l}, \boldsymbol{u}]$ *is given component-wise by*

$$P_{[\boldsymbol{l}, \boldsymbol{u}], i}(\boldsymbol{y}) = \min\{\max\{y_i, l_i\}, u_i\} \tag{1.3.28}$$

with $i = 1, \dots, n$*.*

The derivation of the orthogonal projection $P_{[\boldsymbol{l}, \boldsymbol{u}]}$ can be found in Beck (2017). We now provide the main result concerning the projection phase. We refer the reader to Appendix A for the proof.

Proposition 9. *Let* $\boldsymbol{y} \in \mathbb{R}^n$*, with* $n \geq 2$*, and* $K^{**} = \{i \in K^* \colon y_i > 0\}$*, with* K^* *being the optimal set in Proposition 6. Assume that* $\mathcal{B}$ *in* (1.3.25) *is non-empty. Then, the orthogonal projection of* $\boldsymbol{y}$ *onto* $\mathcal{B}$ *is*

$$P_{\mathcal{B}}(\boldsymbol{y}) = \pi^{-1}_{K^{**}} \left(P_{[\pi_{K^{**}}(\boldsymbol{l}), \pi_{K^{**}}(\boldsymbol{u})]}(\pi_{K^{**}}(\boldsymbol{y} - \eta^* \mathbf{1}) \right) \tag{1.3.29}$$

where $\pi^{-1}_K(\boldsymbol{z})$ *is the pre-image of* $\boldsymbol{z} \in \mathbb{R}^{|K^{**}|}$ *under* π_K *and* $\eta^* \in \mathbb{R}$ *is a solution of*

$$\sum_{i=1}^{k} P_{[\pi_{K^{**}}(\boldsymbol{l}), \pi_{K^{**}}(\boldsymbol{u})]}(\boldsymbol{y} - \eta \mathbf{1}) = 1. \tag{1.3.30}$$

Let $\mathcal{P}_g = \{\boldsymbol{x}^p(g) \in \mathbb{R}^n \colon p = 1, \dots, NP\}$ be the swarm at generation g, with $g = 1, \dots, g_{max}$. Then, the proposed ALLSO variant maps the individuals in $\mathcal{P}_g$, which are updated using (1.3.12) and (1.3.13), onto the set $\mathcal{B}$ by means of the projector defined in (1.3.29). The resulting mutated swarm is denoted by $\check{\mathcal{P}}_g$. Successively, we apply the self-adaptive penalty approach by Costa et al. (2017) to handle the turnover constraint and to guarantee the global optimality of solutions. More precisely, the objective function value at each projected individual in $\check{\mathcal{P}}_g$, namely $f(\check{\boldsymbol{x}}^p)$, is normalized according to the formula

$$\hat{f}(\check{\boldsymbol{x}}^p) = \frac{f(\check{\boldsymbol{x}}^p) - f^{min}}{f^{max} - f^{min}}$$

where $f^{min} = \min_{\check{\boldsymbol{x}}^p \in \check{\mathcal{P}}_g} f(\check{\boldsymbol{x}}^p)$ and $f^{max} = \max_{\check{\boldsymbol{x}}^p \in \check{\mathcal{P}}_g} f(\check{\boldsymbol{x}}^p)$. Similarly, the corresponding normalized constraint violation is given by

$$\Psi\left(\check{\boldsymbol{x}}^p\right) = \begin{cases} \dfrac{\max\{\psi(\check{\boldsymbol{x}}^p), 0\}}{\psi^{max}}, & \text{if } \psi^{max} > 0 \\ 0, & \text{otherwise} \end{cases}$$

where ψ^{max} denotes the maximum of $\psi(\check{\boldsymbol{x}}^p)$ over all the mutated solutions in $\check{\mathcal{P}}_g$ which do not satisfy the turnover constraint.

Finally, the penalty function is

$$F\left(\check{\boldsymbol{x}}^p\right) = \begin{cases} \hat{f}\left(\check{\boldsymbol{x}}^p\right) & \text{if} \quad \psi(\check{\boldsymbol{x}}^p) \leq 0 \\ \hat{f}(\check{\boldsymbol{z}}) + R_f \Psi\left(\check{\boldsymbol{x}}^p\right) & \text{if} \quad \psi(\check{\boldsymbol{x}}^p) > 0 \text{ and } f(\check{\boldsymbol{x}}^p) \leq f(\check{\boldsymbol{z}}) \\ \hat{f}(\check{\boldsymbol{x}}^p) + R_f \Psi\left(\check{\boldsymbol{x}}^p\right) & \text{if} \quad \psi(\check{\boldsymbol{x}}^p) > 0 \text{ and } f(\check{\boldsymbol{x}}^p) > f(\check{\boldsymbol{z}}), \end{cases} \tag{1.3.31}$$

where R_f represents the feasibility ratio for $\check{\mathcal{P}}_g$, that is the percentage of individuals in $\check{\mathcal{P}}_g$ satisfying the turnover constraint. In (1.3.31), the reference point $\check{\boldsymbol{z}}$ is a point belonging to $\check{\mathcal{P}}_g$ that satisfies the turnover constraint and has the lowest objective function value found so far. As in Costa et al. (2017), if the population has no feasible points, $f\left(\check{\boldsymbol{z}}\right)$ is initially and temporarily set to f^{max}, so that $f(\check{\boldsymbol{x}}^p) \leq f(\check{\boldsymbol{z}})$ for all $\check{\boldsymbol{x}}^p \in \check{\mathcal{P}}_g$ and $\hat{f}(\check{\boldsymbol{z}}) = 1$. The value of $f\left(\check{\boldsymbol{z}}\right)$ is updated only when the first feasible point is encountered.

We conclude this subsection by stating the following theorem, whose proof is omitted since it is similar to the one presented in Costa et al. (2017).

Proposition 10. *The problem*

$$\min_{\boldsymbol{x} \in \mathcal{B}} F(\boldsymbol{x})$$

with F as in (1.3.31)*, is equivalent to the problem* (1.3.27)*.*

1.3.d Initialisation strategy and complete algorithm

Following the financial literature Bertsimas et al. (2022), we consider a portfolio optimization problem of high dimensions if it involves more than 400 securities. For this kind of problems, the common strategies of seeking a search space coverage by initializing the particles uniformly throughout the space as well as by increasing the size of the swarm are inefficient, because the search space grows exponentially with the dimension Van Zyl and Engelbrecht (2015). Moreover, the presence of highly constrained feasible regions in the search space exacerbates even more the initialization issue Datta and Deb (2015).

To effectively address the low degree of feasibility in our portfolio rebalancing problem due to the complexity of the turnover constraint, we propose a direct initialization of the candidate solutions in a neighbourhood of $\boldsymbol{x}_0$.

Let d_i^{min} and d_i^{max} be the minimum and the maximum allowed weight changes for $\boldsymbol{x}_{0,i}$

respectively, with $i = 1, \ldots, n$. Let D^p denote the total portfolio weight allowed to be re-allocated in $\boldsymbol{x}_0$ for defining the p-th candidate solution $\boldsymbol{x}^p(0)$, with $p = 1, \ldots, NP$. Then, for each p,

1. we randomly select D^p within $[0,\ {}^{TR}/_2]$;

2. we select a subset J^- of k' assets from the k assets with positive weight in $\boldsymbol{x}_0$, so that
$$x_j^p(0) = x_{0,j} - d_j, \text{ for } j \in J^-$$
where d_j is randomly sampled in $\left[d_j^{min},\ d_j^{max}\right]$ in such a way that $\sum_{j \in J^-} d_j = D^p$, and $x_j^p(0) = 0$ or $l_j \leq x_j^p(0) \leq u_j$;

3. we select a subset J^+ of k'' assets from the $n - k$ assets with zero weight in $\boldsymbol{x}_0$, with $k'' \leq k'$, so that
$$x_j^p(0) = x_{0,j} + d_j, \text{ for } j \in J^+$$
where d_j is randomly sampled in $[d^{min},\ d^{max}]$ in such a way that $\sum_{j \in J^+} d_j = D^p$, and $l_j \leq x_j \leq u_j$;

4. for $j \in I \setminus (J^- \cup J^+)$, we set $x_j^p(0) = x_{0,j}$.

The portfolios assembled using this scheme satisfy cardinality, box and turnover constraints. In this way, the initialization strategy encourages the swarm to focus on exploitation rather than exploration, thereby allowing it to identify promising solutions, even in problems with high dimension and small feasible regions.
Regarding the initial velocities, we set them all equal to the zero vector, that is $\boldsymbol{v}^p(0) = \boldsymbol{0}$, for $p = 1, \ldots, NP$.

The pseudocode of the proposed LLSO variant with adaptive parameters update, mutation of the particles in the first level and hybrid constraint-handling technique, shortly ALLSO-MUT-H, is reported in Algorithm 3. It can be noticed that, setting $\boldsymbol{x}_0 = \boldsymbol{0}$ and $TR = 1$, ALLSO-MUT-H can also tackle portfolio optimization problems with no rebalancing. In this case, only the orthogonal projector is needed to move the unfeasible solutions to the feasible region.

Algorithm 3: ALLSO-MUT-H

Input : $\boldsymbol{x}_0$, $\boldsymbol{l}$, $\boldsymbol{u}$, k, TR, NL_{min}, NL_{max}, NP, $\bar{\delta}$, px, ξ
Output: $\boldsymbol{x}_{gbest}$

Set $g = 0$ and $NL_g = 20$
Initialize the swarm $\mathcal{P}_g = \{\boldsymbol{x}^p(g) \colon p = 1, \ldots, NP\}$ and the velocities $\boldsymbol{v}^p(g)$
for $i = 1$ *to* NP **do**
    Project $\boldsymbol{x}^p(g)$ onto $\mathcal{B}$ using (1.3.29)
    Calculate the turnover violation using (1.3.26)
end
Calculate the penalty F for particles in $\mathcal{P}_g$
Sort $\mathcal{P}_g$ by F value and divide it in NL_g levels
Set $\boldsymbol{x}_{gbest}(g) = \boldsymbol{x}^1(g)$
while $g < g_{max}$ **do**
    $g = g + 1$
    Set $\mathcal{P}^{L1} = \{\boldsymbol{x}^p(g) \colon p \in L_1\}$
    for $p = 1$ *to* LP_g **do**
        Use Algorithm 2 to generate the mutated particle $\hat{\boldsymbol{x}}^{1,p}(g)$ from $\boldsymbol{x}^{1,p}(g)$
        Project $\hat{\boldsymbol{x}}^{1,p}(g)$ onto $\mathcal{B}$ using (1.3.29)
        Calculate the turnover violation using (1.3.26)
    end
    Set $\mathcal{P}^{L1}_{mut} = \{\hat{\boldsymbol{x}}^p(g) \colon p \in L_1\}$
    Calculate the penalty F for $\mathcal{P}^{L1} \cup \mathcal{P}^{L1}_{mut}$ and update $\mathcal{P}^{L1}$ based on $\mathcal{P}^{L1}_{mut}$ using F
    Sort $\mathcal{P}^{L1}$ by the F value
    Calculate the swarm aggregation indicator using (1.3.14) and update ϕ_g using (1.3.18)
    for $p = LP_{g+1}$ *to* NP **do**
        Update $\boldsymbol{v}^p(g)$ using (1.3.12) and clamp it using (2.4.13)
        Update $\boldsymbol{x}^p(g)$ using (1.3.13)
        Project $\boldsymbol{x}^p(g)$ onto $\mathcal{B}$
        Calculate the turnover violation using (1.3.26)
    end
    Set $\check{\mathcal{P}}_g$ be the set of updated particles
    Calculate ϕ for $\mathcal{P}_g \cup \check{\mathcal{P}}_g$ and update $\mathcal{P}_g$ based on $\check{\mathcal{P}}_g$ using F
    Sort $\mathcal{P}_g$ by F value
    Calculate F for the set $\{\boldsymbol{x}_{gbest}(g), \boldsymbol{x}^1(g)\}$ and update $\boldsymbol{x}_{gbest}(g)$
    Calculate the relative improvement $t(g)$ of $\boldsymbol{x}_{gbest}(g)$ using (1.3.15)
    if $t(g) < t(g-1)$ *or* $t(g) = 0$ **then**
        Update NL_g using (1.3.16) and (1.3.17)
end

1.4 Experimental analysis

This section is divided into two parts. On the one hand, we point out the strengths and weaknesses of using the proposed algorithm to tackle large-scale cardinality-constrained portfolio optimization problems. On the other hand, we assess the profitability of the investment strategy in a real-world case study by varying the size of portfolios.

As a first task, we study the impact of the developed mutation operator on the LLSO-

type algorithms, all equipped with our hybrid constraint-handling technique. For this purpose, we compare the dynamic LLSO (DLLSO) Yang et al. (2018), the adaptive LLSO (ALLSO) Song et al. (2021), and the reinforcement learning level-based particle swarm optimization (RLLPSO) algorithm Wang et al. (2022). These are the most recent solvers proposed in the literature, that employ the level-based learning paradigm into the PSO framework. Their superiority in solution accuracy with respect to several state-of-the-art algorithms has already been proved theoretically and empirically for unconstrained optimization problems. Hence, we inspect the best variant of the LLSO-type algorithm for our constrained optimization problems. To highlight the benefits of our hybrid self-adaptive penalty approach, we compare the ALLSO-MUT-H with an ALLSO equipped by our mutation operator and the exact ℓ_1-penalty proposed in Corazza et al. (2021). At the same time, we consider two other state-of-the-art swarm optimizers, namely PSO and FA Yang (2010), both endowed with the adaptive penalty by Corazza et al. (2021). It is worth noting that the self-adaptive method Costa et al. (2017) and the adaptive penalty Corazza et al. (2021) guarantee the convergence to optimal solutions.
In the numerical experiments, we follow the suggestions in the reference papers to set up the parameters of each algorithm, and the complete list of parameter values is reported in Table 1.1.

Table 1.1: Parameter settings of the algorithms used in the comparisons.

Algorithm	Parameter settings	Reference
DLLSO	$S = \{4, 6, 8, 10, 20, 50\}$, $NP = 500$, $\phi = 0.4$	Yang et al. (2018)
ALLSO	$NL_{min} = 2$, $NL_{max} = 50$, $NP = 500$, $\bar{\delta} = 0.01$, $px = 0.01$, $\xi = 10^{-6}$	Song et al. (2021)
RLLPSO	$S = \{4, 6, 8, 10, 20, 50\}$, $NP = 500$, $\phi = 0.4$, $\alpha = 0.4$, $\gamma = 0.8$, $\varepsilon = 0.9$, $\xi = 10^{-6}$	Wang et al. (2022)
PSO	$\omega_{min} = 0.4$, $\omega_{max} = 0.9$, $c_{1,min} = c_{2,min} = 0.5$, $c_{1,max} = c_{2,max} = 2.5$,	Ratnaweera et al. (2004), Corazza et al. (2021)
FA	$\alpha = 0.5$, $\beta_{min} = 0.2$, $\gamma = 1$	Yang (2010)

1.4.a Algorithmic comparisons

For the algorithmic comparisons, we use three data sets from the OR-Library Canakgoz and Beasley (2009), namely S&P 500, Russell 2000, and Russell 3000, which represent large capital market indices. Table 1.2 summarizes the data sets employed in the experimental analysis. The second and third columns report the number of assets involved and the number of observations for each data set, respectively. The last columns show the cardinality thresholds used in the experiments. In particular, we explore the cases where k is 30%, 15%, and 5% of the size n of the corresponding data set. Since the results are similar, in the following analysis we describe only the outcomes for the cardinality threshold equal to 30% of the size of the corresponding data set. The cases regarding the last two columns of Table 1.2 are addressed in the

supplementary file. For calculating the expected rates of return, we adopt a historical approach based on all the information available, consisting of 290 weekly prices for each asset. Since our investable universes have more decision variables than observations, the estimator of the associated covariance matrices is biased. We employ the shrinkage estimator proposed in Ledoit and Wolf (2004) to reduce this effect. One can appreciate the difference between these estimators for the three data sets in the supplementary file.

For a fair comparison of the solvers, we adopt the following assumptions.

- All the algorithms have the same initial population of 500 individuals for each test set.
- We set $d_i^{min} = 0.0005$ and $d_i^{max} = 0.0050$ in the initialization strategy.
- For each algorithm, we perform 30 independent runs with 2000 generations, which represents the terminal condition.

Moreover, all the portfolios from a given test set employ the following parameter setting.

- The risk-free value in (1.2.4) is set to zero.
- The box thresholds in (2.3.6) are $l_i = 0.001$ and $u_i = 0.05$ for each asset.
- Regarding the turnover constraint (2.3.7), TR is set equal to 0.20.
- The vector of current positions $\boldsymbol{x}_0$ in (2.3.7) is fixed for all the compared algorithms and in all simulations by randomly sampling once from each set of feasible portfolios.

Table 1.2: Data sets from Canakgoz and Beasley (2009) with the corresponding number of weeks, the number of market constituents (n) used in the estimation of parameters, and the cardinality thresholds.

Data set name	Weekly prices	Assets (n)	Cardinality threshold (k)		
			$\lfloor 30\% \cdot n \rfloor$	$\lfloor 15\% \cdot n \rfloor$	$\lfloor 5\% \cdot n \rfloor$
S&P 500	290	457	137	68	22
Russell 2000	290	1318	395	197	65
Russell 3000	290	2151	645	322	107

Mutation effects on LLSO variants

We recall that in the following analysis k is set always equal to 30% of the size of the corresponding data set, and we refer the interested reader to the supplementary file for the other cases.

Table 1.3 shows four performance metrics linked to the best objective function value

over 30 runs, on the three public data sets. The best results are highlighted in bold font. We note that the ALLSO-MUT-H outperforms the competitors in all the case studies, presenting the lowest mean objective function value. Further, we remark that all the inspected LLSO variants are able to find feasible solutions. Focusing on the mutation benefits, we observe from Table 1.4 that the mutation has a significant impact on the performance of solvers. Specifically, although mutation-based optimizers present higher volatility than their counterparts, they always show lower results in terms of minimum-maximum range of the best solutions.

Table 1.3: Statistics regarding the best values of the objective function over 30 runs.

Data set	Statistics	DLLSO-H	DLLSO-MUT-H	ALLSO-H	ALLSO-MUT-H	RLLPSO-H	RLLPSO-MUT-H
S&P 500	mean	-0.1526	-0.1670	-0.1524	**-0.1672**	-0.1523	-0.1622
	std	0.0004	0.0015	**0.0003**	0.0014	0.0004	0.0023
	min	-0.1532	**-0.1718**	-0.1529	-0.1698	-0.1532	-0.1662
	max	-0.1517	**-0.1643**	-0.1517	-0.1638	-0.1509	-0.1546
Russell 2000	mean	-0.1920	-0.2104	-0.1925	**-0.2124**	-0.1929	-0.2049
	std	**0.0010**	0.0042	0.0011	0.0037	0.0013	0.0043
	min	-0.1932	-0.2193	-0.1946	**-0.2196**	-0.1963	-0.2146
	max	-0.1895	-0.2003	-0.1895	**-0.2039**	-0.1903	-0.1959
Russell 3000	mean	-0.2059	-0.2247	-0.2091	**-0.2301**	-0.2082	-0.2271
	std	0.0011	0.0031	0.0015	0.0025	**0.0010**	0.0043
	min	-0.2085	-0.2328	-0.2120	-0.2352	-0.2111	**-0.2358**
	max	-0.2042	-0.2192	-0.2067	**-0.2260**	-0.2065	-0.2196

Table 1.4: Relative change of the mutated algorithms versus non-mutated counterparts. The p-values for the paired t-tests are displayed in brackets. Note that in all cases the p-values are under the significance level $\alpha = 0.05$, indicating the rejection of the null hypothesis of equality of the means, against the alternative left-sided hypothesis.

Data set	DLLSO-MUT-H vs. DLLSO-H (%)	ALLSO-MUT-H vs. ALLSO-H (%)	RLLPSO-MUT-H vs. RLLPSO-H (%)
S&P 500	9.4432	9.7153	6.5109
	$(2.8638 \cdot 10^{-30})$	$(8.6384 \cdot 10^{-31})$	$(7.5328 \cdot 10^{-20})$
Russell 2000	9.5859	10.3564	6.2501
	$(6.2697 \cdot 10^{-21})$	$(2.7589 \cdot 10^{-23})$	$(3.2532 \cdot 10^{-15})$
Russell 3000	9.1309	10.0823	9.1097
	$(1.2435 \cdot 10^{-23})$	$(7.4039 \cdot 10^{-29})$	$(5.6022 \cdot 10^{-20})$

Figure 1.1 shows the convergence and the diversity analyses of the compared solvers on the three data sets. From the first set of graphs, we note that the three mutated algorithms are able to reach significantly lower objective function values, and the ALLSO-MUT-H performs better than the others. Moreover, the algorithms without mutation show population stagnation around 100 generations, meaning that they converge to a

local minimum and are not able to further explore the search space. This is confirmed by the results showed in the logarithmic scale plots of the diversity measures. We can observe that the ALLSO-MUT-H and the DLLSO-MUT-H are able to escape from the local minima, due to the oscillatory behaviour of the swarm diversity.

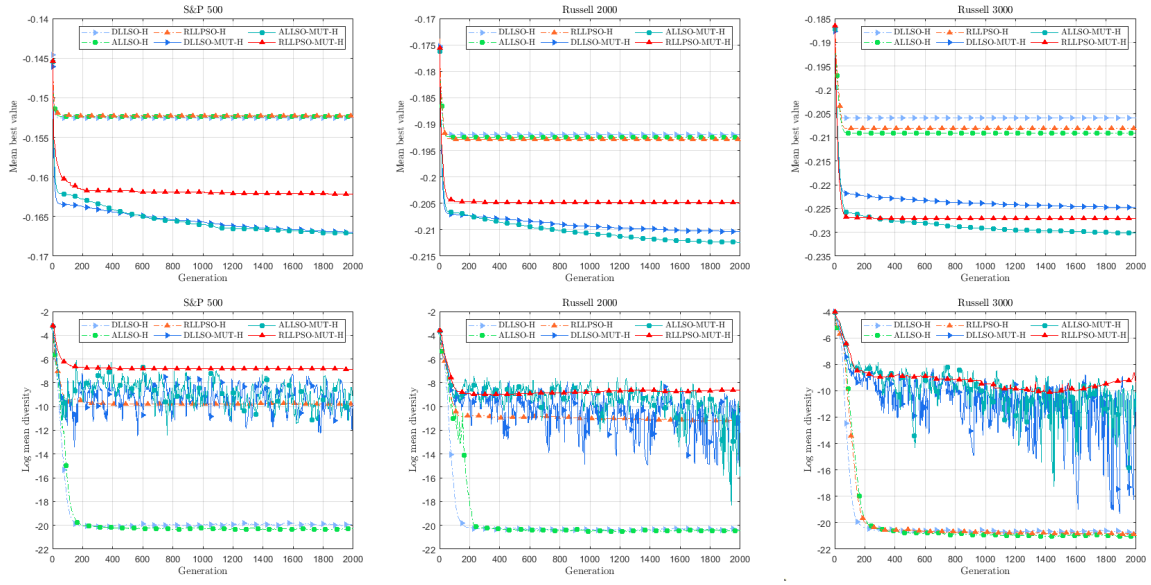

Figure 1.1: Convergence and diversity analyses on the three data sets. Graphs in the first row show the behaviour of algorithms in terms of mean best value of the objective function, while in the row below are displayed the logarithmic scale plots of the diversity scores.

Comparison with state-of-the-art swarm optimization algorithms

In the previous subsection, we have analysed the impact of the mutation on the capabilities of LLSO-based algorithms, finding that the ALLSO-MUT-H is the more efficient choice in terms of convergence and quality of solutions. We recall also that we adopt the same parameter setup presented above in Table 1.1. We exhibit the statistics of the comparison in Table 1.5, where are displayed the percentage of feasible solutions provided by the different solvers over the 30 runs; the mean of the constraint violation function CV for the non-feasible solutions; the average value of the penalty function F_{ℓ_1} over 30 runs. Notice that, the penalty function corresponds to the objective function when the solutions are feasible. Looking at the results, we can argue that the ALLSO-MUT-H reaches the best mean value of the penalty function in all the data sets, and it always provides feasible solutions. This insight is confirmed by the convergence analysis plots in Figure 1.2, which show the benefits of our constraint-handling technique in terms of accuracy of solutions. Moreover, the diversity graphs suggest that our solver is the sole algorithm able to exhibit exploration and exploitation phases alternatively.

Table 1.5: Comparison with state-of-the-art swarm optimization algorithms implementing the exact ℓ_1-penalty framework proposed in Corazza et al. (2021).

Data set	Statistics	PSO-ℓ_1	FA-ℓ_1	ALLSO-MUT-ℓ_1	ALLSO-MUT-H
S&P 500	feasible sol. (%)	0	0	100	100
	mean CV	$3.1572 \cdot 10^{-12}$	$3.1572 \cdot 10^{-12}$	0	0
	mean F_{ℓ_1}	-0.1325	-0.1325	-0.1325	-0.1672
Russell 2000	feasible sol. (%)	63	100	100	100
	mean CV	$1.1661 \cdot 10^{-5}$	0	0	0
	mean F_{ℓ_1}	-0.1717	-0.1639	-0.1639	-0.2124
Russell 3000	feasible sol. (%)	83	100	0	100
	mean CV	$7.4015 \cdot 10^{-12}$	0	0.0326	0
	mean F_{ℓ_1}	-0.1862	-0.1291	-0.1487	-0.2301

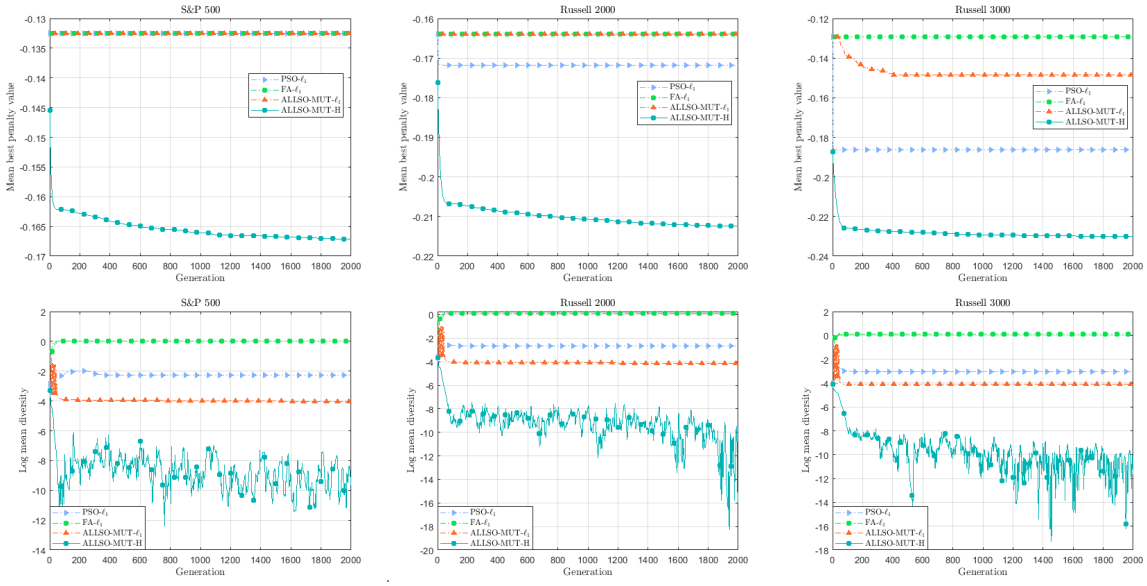

Figure 1.2: Plots in the first row show the behaviour of the algorithms in terms of mean best value of the penalty function, while in the second row are presented the logarithmic scale graphs of the diversity.

1.4.b Real-world application

Data description and investment setting

The constituents of the MSCI World index, on 31st January 2022, form our investable pool. The data set has been downloaded from DataStream and consists of monthly

prices covering the period from January 2012 to January 2022 for a total of 121 months. Stocks with missing observations were disregarded, and thus the final data set includes 1119 stocks. For the performance comparisons, we introduce a value-weighted benchmark index with the same constituents and, as in the previous analysis, we set $r_f = 0$.

The portfolio design employs the following parameter setting. For the cardinality constraint (2.3.5), we consider $k \in \{335, 167, 55, 22\}$, corresponding to $k_{\%} = 30\%$, 15%, 5% and 2% of the pool size, respectively. As stated in the introductory part of Section 1.4.a, we recall that the box thresholds in (2.3.6) are $l_i = 0.001$ and $u_i = 0.05$ for each asset i, with $i = 1, \ldots, 1119$, and the turnover rate in (2.3.7) is set equal to 0.20, as in Kaucic et al. (2020).

We use a rolling time window procedure to rebalance optimal portfolios every month, from January 2017 to January 2022, to point out the effects of the market changes on the behaviour of the investments and, as a consequence, the total number of ex-post dates is 61. We solve the corresponding problem instances by employing overlapping 60-months windows, which are updated every month by removing the oldest data and including the latest information.

In each quoted window, as already pointed out above, we adopt a historical approach to calculate expected rates of return, and to reduce the bias in the estimation of the covariance matrix C, we take advantage of the shrinkage estimator proposed in Ledoit and Wolf (2004). Let us denote by $\boldsymbol{x}_t$ the optimal portfolio at the ex-post month t, with $t = 1, \ldots, 61$. Due to the time dependence of the considered investment plan, we rewrite the turnover constraint (2.3.7) as follows

$$\sum_{i=1}^{n} |x_{t,i} - x_{t^-,i}| \leq TR\,. \tag{1.4.32}$$

In the previous equation, $\boldsymbol{x}_{t^-} = (x_{t^-,1}, \ldots, x_{t^-,n})$ represents the portfolio to be rebalanced Shen et al. (2014), which is defined for $t = 2, \ldots, 61$ as

$$x_{t^-,i} = \frac{x_{t-1,i} R^g_{t-1,i}}{\sum_{j=1}^{n} x_{t-1,j} R^g_{t-1,j}} \tag{1.4.33}$$

with the denominator being the gross portfolio return at month $t-1$.[2] At time $t = 1$, we set $\boldsymbol{x}_{t^-} = \mathbf{0}$ and $TR = 1$.

Let us assume a self-financing strategy with an initial wealth $W_0 = 10\,000\,000\,\$$. Then, we explicitly evaluate the magnitude of the trading through the cost function $\lambda(\boldsymbol{x}_t, \boldsymbol{x}_{t^-})$ introduced in Beraldi et al. (2021). As reported in Table 1.6, we consider the transaction cost structure characterized by decreasing cost rates as the traded value increases.

[2]The gross return of asset i at month t is defined as $R^g_{i,t} = \frac{S_{i,t}}{S_{i,t-1}}$, where $S_{i,t}$ is the price of the i-th asset at the end of month t.

Table 1.6: Structure of transaction costs.

Trading segment ($)	Fixed fee ($)	Proportional cost (%)
0 – 7 999	40	0
8 000 – 49 999	0	0.5
50 000 – 99 999	0	0.4
100 000 – 199 999	0	0.25
$\geq$ 200 000	400	0

Ex-post performance measures

The following measures (see Kaucic (2019) and Kaucic et al. (2020)) are considered to evaluate the profitability of the investment strategies. Let $r_{p,t}^{out}$ be the ex-post portfolio rate of return realized at time t, with $t = 1, \ldots, 61$. First, we consider the so-called ex-post Sharpe ratio, defined as

$$SR^{out} = \frac{\mu^{out}}{\sigma^{out}} \tag{1.4.34}$$

where μ^{out} and σ^{out} are the mean and the standard deviation of the ex-post portfolio rates of return, respectively.
The second measure employed in the analysis is the so-called Omega ratio, defined as the ratio between the gains over a threshold level and the losses under a threshold level. In this study, we set both thresholds equal to zero, that is

$$Omega = \frac{\sum_{t=1}^{61} r_{p,t}^{out} \mathbb{1}_{\{r_{p,t}^{out}>0\}}}{-\sum_{t=1}^{61} r_{p,t}^{out} \mathbb{1}_{\{r_{p,t}^{out}<0\}}} \tag{1.4.35}$$

where $\mathbb{1}_A$ is the indicator function on A.
The information gathered from these performance measures draws a complete picture of the ex-post portfolio return distribution. In particular, the ex-post Sharpe ratio describes the central part of the portfolio return distribution, while the Omega ratio considers the behaviour of profits and losses.

Further, to measure the profitability of the investment at time t, we compute the net wealth as

$$W_t = W_{t-1}\left(1 + r_{p,t}^{out}\right) - \lambda(\boldsymbol{x}_t, \boldsymbol{x}_{t^-}). \tag{1.4.36}$$

Then, we compare the profitability of the investments using the so-called compound annual growth rate, which in our case is calculated as

$$CAGR = \left(\frac{W_{61}}{W_0}\right)^{\frac{12}{61}} - 1 \tag{1.4.37}$$

where W_0 represents the initial wealth and W_{61} is the final wealth.
To evaluate the capacity of a strategy to avoid high losses, we introduce the drawdown

measure, which can be written

$$DD_t = \min\left\{0, \frac{W_t - W_{peak}}{W_{peak}}\right\} \tag{1.4.38}$$

where W_{peak} is the maximum amount of wealth reached by the strategy until time t. In particular, we consider the mean and the standard deviation of the drawdown measure over time.

Finally, we propose to measure the effect of the costs on the available capital in the out-of-sample period by

$$\Lambda_{\%} = \frac{1}{61}\sum_{t=1}^{61}\frac{\lambda(\boldsymbol{x}_t, \boldsymbol{x}_{t^-})}{W_{t-1}} \cdot 100. \tag{1.4.39}$$

Ex-post performance analysis

In the ex-post analysis, we investigate how the performance of the proposed asset allocation model changes by varying the k parameter in (2.3.5).

First, we remark that, for any ex-post dates, the proposed hybrid LLSO variant identifies feasible solutions for all the portfolio sizes. The empirical results are summarised in Table 1.7, where the number of assets of the considered strategies is also displayed. Note that for each value of k, the proposed investments provide better performances than the value-weighted benchmark. This result implies that introducing a cardinality constraint in the portfolio model allows to choose a subset of the most profitable assets in the investable pool.

In terms of the return-risk profile, strategies with $k_{\%} = 30\%$, 15%, and 5% show comparable performances, while the strategy with $k_{\%} = 2\%$ has a lower Sharpe ratio, which is due to its large volatility. Similar conclusions can be made about the Omega ratio, which expresses the gain-loss profile of the strategies. Despite better performance with respect to Sharpe and Omega ratios, strategies involving portfolios with a larger number of assets generate less wealth. Moreover, we observe that reducing cardinality leads to more profitable portfolio strategies.

Concerning the drawdown measures, the 5% asset allocation model is the most conservative, while the one with $k_{\%} = 2\%$ is the worst. Thus, the performance deteriorates by reducing portfolio size below a critical threshold.

As highlighted in the last two rows of Table 1.7 and in Figure 1.3, the impact of transaction costs for strategies with small k is negligible. On the contrary, portfolios with many assets have more fluctuations in the rebalancing phases, leading to higher trading commissions with a significant impact on the wealth generated.

Summing up, we can infer that the strategy with $k_{\%} = 5\%$ shows the best balance between risk-adjusted performance measures and capability to generate net profits.

Table 1.7: Performance of the proposed cardinality-constrained portfolio allocation model for different cardinalities in comparison to the benchmark.

$k_{\%}$	30%	15%	5%	2%	Benchmark
num. assets	335	167	55	22	1119
Sharpe ratio	0.4917	0.4794	0.4467	0.2878	0.1731
Omega ratio	3.3269	3.2428	3.0360	2.0980	1.5388
CAGR	1.1148	1.6899	2.1118	2.2502	1.3460
std	0.0052	0.0052	0.0054	0.0077	0.0066
mean DD	-0.0066	-0.0057	-0.0050	-0.0072	-0.0070
std DD	0.0070	0.0064	0.0060	0.0106	0.0089
mean λ (\$)	16,046	10,713	6,571.7	3,121.9	–
$\Lambda_{\%}$	0.1536	0.1011	0.0610	0.0290	–

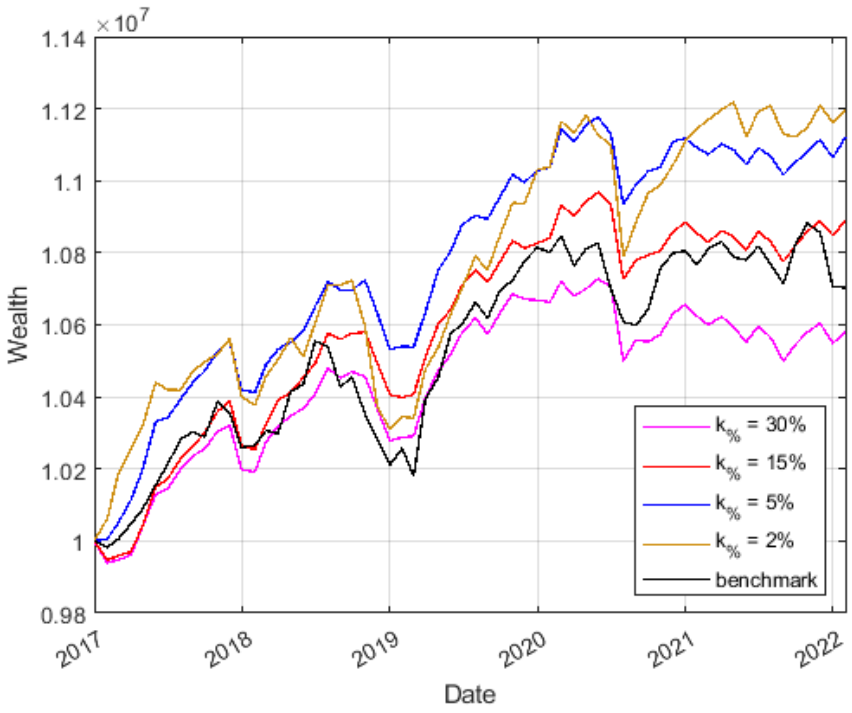

Figure 1.3: Ex-post evolution of net wealth of the benchmark and of the proposed cardinality-constrained portfolio allocation model with different cardinalities.

1.5 Conclusions and future works

In this paper, we have developed a swarm optimization algorithm for solving a large-scale cardinality-constrained portfolio optimization problem, where a modified Sharpe ratio performance measure represents the objective function. We have considered

four real-world constraints: cardinality, box, budget, and turnover constraints. Due to the properties of the model inspected, we have proposed a variant of the LLSO equipped with a hybrid procedure to manage the constraints efficiently. Moreover, a novel mutation operator has been introduced to improve the accuracy of solutions. Our solver capabilities have been compared with those of two variants of the LLSO as well as other state-of-the-art swarm optimization algorithms endowed with an ℓ_1-penalty function. Numerical experiments on three publicly available large-scale data sets showed the outperformance of our hybrid procedure. From the financial point of view, we have analyzed the sensitivity of the portfolio model to the cardinality constraint with data from the last five years of the MSCI World index. We have found that portfolios of small size are more competitive with respect to the value-weighted benchmark index, also in periods of market downturns. Specifically, the losses are reduced, and the cost impact on the available capital is marginal compared to the profits.

The developed algorithm has a direct application to solving financial management problems. However, it could be easily modified for large-scale parameter optimization, such as for training feed-forward neural networks. Moreover, since the projection paradigm has already been used profitably to tackle real-world problems in the engineering field like sensor networks, radiation therapy treatment planning, computerized tomography, magnetic resonance imaging, and optics to name a few, in our future research we plan to extend our portfolio optimization model by including other constraints. On the one hand, we will consider the so-called risk-budgeting constraints to control the portfolio risk exposition explicitly. On the other hand, based on the European Green Deal and the ESG Disclosure requirements for funds and investments, we will add sustainable-policy constraints to guarantee a minimum level of ESG rating to the investment. The inclusion of these new constraints in the optimization model requires an adequate adaptation of the hybrid-constraint handling technique, by considering novel projection schemes. Another possible research direction to fill the gap in handling the tail-risk is represented by the introduction of alternative performance measures in the optimization framework, like the conditional value-at-risk and its related measures.

Appendix A. Proofs of the main results

Proof. 6 We rewrite problem (1.3.24) in the following way

$$\min_{\boldsymbol{x}_K:\, x_i \in C_i} \left\{ \frac{1}{2} \sum_{j \in K} (x_j - y_j)^2 \right\} + \frac{1}{2} \sum_{j \in I \setminus K} y_j^2 \,,$$

that is equivalent to

$$\min_{\boldsymbol{x}_K:\, x_i \in C_i} \left\{ \frac{1}{2} \sum_{j \in K} (x_j - y_j)^2 \right\} - \frac{1}{2} \sum_{j \in K} y_j^2 + \frac{1}{2} \sum_{j \in I} y_j^2 \,. \tag{A.1}$$

We note that the last term in (A.1) does not depend on x_j, $j \in K$, so we can focus our attention on the first two terms, i.e.

$$\min_{\pi_K(\boldsymbol{x}):\, x_i \in C_i} \frac{1}{2} \|\pi_K(\boldsymbol{x} - \boldsymbol{y})\|^2 - \frac{1}{2} \|\pi_K(\boldsymbol{y})\|^2 \,.$$

By contradiction, we suppose that there is a K' different from K^*, where we recall that K^* is the set of indices corresponding to the k largest components of $\boldsymbol{y}$. At this point, we define

$$f(\pi_K(\boldsymbol{y})) := -\frac{1}{2} \|\pi_K(\boldsymbol{y})\|^2 + \min_{\pi_K(\boldsymbol{x}):\, x_i \in C_i} \frac{1}{2} \|\pi_K(\boldsymbol{x} - \boldsymbol{y})\|^2$$

and

$$g(t) := f((1-t)\pi_{K^*}(\boldsymbol{y}) + t\pi_{K'}(\boldsymbol{y})) \quad \text{with} \quad t \in [0,1] \,.$$

Then we have,

$$f(\pi_{K'}(\boldsymbol{y})) - f(\pi_{K^*}(\boldsymbol{y})) = g(1) - g(0) = \int_0^1 g'(t)\, \mathrm{d}t$$

and

$$g'(t) = \nabla f((1-t)\pi_{K^*}(\boldsymbol{y}) + t\pi_{K'}(\boldsymbol{y})) \cdot (-\pi_{K^*}(\boldsymbol{y}) + \pi_{K'}(\boldsymbol{y})) \,,$$

where $\nabla f(\pi_K(\boldsymbol{y})) = -\pi_K(\boldsymbol{y}) + \pi_K(\boldsymbol{y}) - \pi_K(\boldsymbol{x}^*) = -\pi_K(\boldsymbol{x}^*)$ and $\pi_K(\boldsymbol{x}^*) = \underset{\pi_K(\boldsymbol{x}):\, x_i \in C_i}{\operatorname{argmin}} \frac{1}{2} \|\pi_K(\boldsymbol{x} - \boldsymbol{y})\|^2$. Now, since $\mathcal{C}_i \subset \mathbb{R}_+$, we have that $\nabla f(\pi_K(\boldsymbol{y}))$ is non-positive in all components. Moreover, $-\pi_{K^*}(\boldsymbol{y}) + \pi_{K'}(\boldsymbol{y}) \leq 0$ due to the fact that $\pi_{K^*}(\boldsymbol{y})$ is the projection of $\boldsymbol{y}$ onto the set of its largest components. As a result, we obtain $g'(t) \geq 0$, which implies $f(\pi_{K'}(\boldsymbol{y})) \geq f(\pi_{K^*}(\boldsymbol{y}))$. This means that K^* must be the optimal choice. □ □

Proof. 7 The proof of the first part of the proposition follows by defining $\boldsymbol{x}^*_{K^*}$ such that $\boldsymbol{x}^*_{K^*} = \boldsymbol{\delta}^* \otimes \boldsymbol{x}^*$, where $\otimes$ stands for the Hadamard product.

On the contrary, if $\boldsymbol{x}^*$ solves (1.3.27), then $\boldsymbol{x}^*_{K^*} = \boldsymbol{x}^*$. By taking

$$\delta_i^* = \begin{cases} 1 & \text{if} \quad i \in K^* \\ 0 & \text{otherwise} \end{cases}$$

we deduce that $(\boldsymbol{\delta}^*,\ \boldsymbol{x}^*)$ solves (2.3.8). □ □

Proof. 9 The result follows from Theorem 6.27 in Beck (2017), where the hyperplane is represented by the budget constraint (2.3.4) and the box is $[\pi_{K^{**}}(\boldsymbol{l}),\ \pi_{K^{**}}(\boldsymbol{u})]$. We recall also that $\pi_{K^{**}}(\boldsymbol{z}) = (z_{i_1}, \ldots, z_{i_k})$ with $z_{i_j} > 0$ and $i_j \in K^{**}$. □

Appendix B. An approach based on an exact ℓ_1-penalty function

In this appendix we introduce a procedure based on the exact ℓ_1-penalty function for solving cardinality-constrained portfolio optimization problems. We adapt the approach discussed in Corazza et al. (2021) to the algorithms used in the comparison analysis of Subsection 1.4.a and thus we define the constraint violations as follows

$$\begin{aligned}
CV_1 &= \left| \sum_{i=1}^{n} x_i - 1 \right| \\
CV_2 &= \max \left\{ \sum_{i=1}^{n} \delta_i - k, 0 \right\} \\
CV_3 &= \sum_{i=1}^{n} \max \left\{ \delta_i l_i - x_i, 0 \right\} \\
CV_4 &= \sum_{i=1}^{n} \max \left\{ x_i - \delta_i u_i, 0 \right\} \\
CV_5 &= \sum_{i=1}^{n} |\delta_i (1 - \delta_i)| \\
CV_6 &= \max \left\{ \sum_{i=1}^{n} |x_i - x_{0,i}| - TR,\ 0 \right\} .
\end{aligned}$$

In this manner, we introduce the exact ℓ_1-penalty function

$$F_{\ell_1}(\boldsymbol{x}, \boldsymbol{\delta}; \boldsymbol{\varepsilon}) = f(\boldsymbol{x}) + \frac{1}{\varepsilon_0} \left[\varepsilon_1 CV_1 + \varepsilon_2 CV_2 + \varepsilon_3 CV_3 + \varepsilon_4 CV_4 + \varepsilon_5 CV_5 + \varepsilon_6 CV_6 \right] \tag{B.1}$$

where $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_6)$, with $\varepsilon > 0$ for all i.

The initial parameters vector $\boldsymbol{\varepsilon}^0$ is set to $\boldsymbol{\varepsilon}^0 = (\varepsilon_0^0, \varepsilon_1^0, \ldots, \varepsilon_6^0) = (10^{-4}, 1, \ldots, 1) \in \mathbb{R}^7$, where ε_0^0 is chosen in order to privilege feasible solutions, and the other parameters are equally penalized for all constraint violations.

The vector $\boldsymbol{\varepsilon}$ is updated by checking the decrease of the function $f(\boldsymbol{x})$ and the violation of the constraints. More precisely, on the one hand, every 5 iterations the entry $\varepsilon_0(g)$ is updated according to the rule

$$\varepsilon_0(g+1) = \begin{cases} \min\{3 \cdot \varepsilon_0(g),\, 1\} & \text{if} \quad f(\boldsymbol{x}(g)) \geq f(\boldsymbol{x}(g-1)) \\ \max\{0.6 \cdot \varepsilon_0(g),\, 10^{-15}\} & \text{if} \quad f(\boldsymbol{x}(g)) < 0.9 \cdot f(\boldsymbol{x}(g-1)) \\ \varepsilon_0(g) & \text{otherwise.} \end{cases} \tag{B.2}$$

On the other hand, every 10 iterations the entries $\varepsilon_i(g)$, $i = 1, \ldots 6$, are updated following the scheme

$$\varepsilon_i(g+1) = \begin{cases} \min\{2 \cdot \varepsilon_i(g),\, 10^4\} & \text{if} \quad CV_i(g) > 0.95 \cdot CV_i(g-1) \\ \max\{0.5 \cdot \varepsilon_i^g,\, 10^{-4}\} & \text{if} \quad CV_i(g) < 0.9 \cdot CV_i(g-1) \\ \varepsilon_i(g) & \text{otherwise,} \end{cases} \tag{B.3}$$

with CV_i the respective constraint violation linked to the ε_i parameter.

The above quoted strategy privileges optimality of solutions possibly at the expenses of their feasibility, due to the fact that $\varepsilon_0(g+1)$ in (2.4.19) is increasing in $F_{\ell_1}(\boldsymbol{x}, \boldsymbol{\delta}; \boldsymbol{\varepsilon}(g+1))$ when the function value $f(\boldsymbol{x}(g))$ increases. Moreover, to favour feasibility of solutions possibly at the expenses of their optimality, the penalty parameter $\varepsilon_i(g+1)$ in (2.4.20) is increased when the relative constraint violation in the g-th generation increases with respect to the previous one.

The procedure is also equipped by a splitting and refining technique for the positions of the particles. In particular, at each iteration, a particle p is split in its components $\boldsymbol{x}^p(g)$ and $\boldsymbol{\delta}^p(g)$ that are updated separately. For the vector $\boldsymbol{\delta}^p(g)$ we employ the following updating rule

$$\delta_i^p(g+1) = \begin{cases} 1 & \text{if} \quad x_i^p(g) \in [l_i,\, u_i] \\ 0 & \text{otherwise} \end{cases}$$

for $i = 1, \ldots, n$. Then, $\boldsymbol{\delta}^p(g+1)$ is kept fixed and $F_{\ell_1}(\boldsymbol{x}, \boldsymbol{\delta}^p(g+1); \boldsymbol{\varepsilon}(g+1))$ is minimized with respect to $\boldsymbol{x}$, obtaining $\widetilde{\boldsymbol{x}}^p(g+1)$. Finally, $\widetilde{\boldsymbol{x}}^p(g+1)$ is refined getting

$$x_i^p(g+1) = \frac{\widetilde{x}_i^p(g+1)\delta_i^p(g+1)}{\sum_{i=1}^n \widetilde{x}_i^p(g+1)\delta_i^p(g+1)}\,, \tag{B.4}$$

for $j = 1, \ldots, NP$.

Appendix C. Nomenclature

Symbol	Description
n	Number of assets in the investable universe
$\boldsymbol{x}$	$n \times 1$ vector of asset weights
R_i	Random variable representing the rate of return of asset i
μ_i	Expected value of R_i
$R_p(\boldsymbol{x})$	Random variable representing the rate of return of portfolio $\boldsymbol{x}$
$\mu_p(\boldsymbol{x})$	Expected rate of return of portfolio $\boldsymbol{x}$
C	$n \times n$ covariance matrix of the n stocks
$\sigma_p(\boldsymbol{x})$	Volatility (standard deviation) of portfolio $\boldsymbol{x}$
$SR(\boldsymbol{x})$	Sharpe ratio of portfolio $\boldsymbol{x}$
$MSR(\boldsymbol{x})$	Modified Sharpe ratio of portfolio $\boldsymbol{x}$
k	Maximum number ($\leq n$) of assets included in the portfolio
$k_{\%}$	Fraction of assets making up the portfolio
$\boldsymbol{\delta}$	$n \times 1$ vector of binary variables denoting inclusion/exclusion from a portfolio
l_i	Lower bound for the i-th portfolio weight, $i = 1, \ldots, n$
u_i	Upper bound for the i-th portfolio weight, $i = 1, \ldots, n$
$\boldsymbol{x}_0$	$n \times 1$ vector containing the current portfolio positions to rebalance
TR	Maximum turnover rate
NP	Number of candidate solutions in the swarm
MAX_{GEN}	Maximum number of generations
NL_g	Number of levels in generation g
LP_g	Number of particles in each level at generation g
$\boldsymbol{x}^{l,p}(g)$	$n \times 1$ vector denoting the p-th particle in level L_l at generation g
$\boldsymbol{v}^{l,p}(g)$	$n \times 1$ vector denoting the p-th velocity in level L_l at generation g
ϕ_g	Adaptive control parameter at generation g
$s(g)$	Swarm aggregation indicator at generation g
$t(g)$	Relative improvement between two consecutive generations
$\hat{\boldsymbol{x}}^{l,\,p}(g)$	$n \times 1$ vector denoting the p-th mutated individual in level L_l at generation g
$\phi(\boldsymbol{x})$	Turnover constraint violation of candidate solution $\boldsymbol{x}$
$F(\boldsymbol{x})$	Penalty function value for candidate solution $\boldsymbol{x}$
CV_j	Constraint violation for the j-th constraint in the exact ℓ_1-penalty function approach
$\boldsymbol{\varepsilon}$	6×1 vector of the exact ℓ_1-penalty function parameters
$F_{\ell_1}(\boldsymbol{x}, \boldsymbol{\delta}; \boldsymbol{\varepsilon})$	Exact ℓ_1-penalty function value for candidate solution $\boldsymbol{x}$

Table 1.8: Table of notation.

Bibliography

B. R. Auer, F. Schuhmacher: *Performance hypothesis testing with the sharpe ratio: the case of hedge funds*. Finance Res. Lett., **10** (2013), n. 4, 196-208, https://doi.org/10.1016/j.frl.2013.08.001.

H. J. C. Barbosa, A. C. C. Lemonge, H. S. Bernardino: *A critical review of adaptive penalty techniques in evolutionary computation*. In: *"Evolutionary Constrained Optimization"*. Infosys Science Foundation Series, Springer, New Delhi (2015).

A. Beck: *"First-order methods in optimization"*. Society for Industrial and Applied Mathematics and Mathematical Optimization Society, Philadelphia (2017).

P. Beraldi, A. Violi, M. Ferrara, C. Ciancio, B. A. Pansera: *Dealing with complex transaction costs in portfolio management*. Ann. Oper. Res., **299** (2021), n.1, 7-22, https://doi.org/10.1007/s10479-019-03210-5.

D. Bertsimas, R. Cory-Wright: *A scalable algorithm for sparse portfolio selection*. INFORMS J. Comput., **34** (2022), n. 3, 1489-1511, https://doi.org/10.1287/ijoc.2021.1127.

D. Bertsimas, R. Shioda: *Algorithm for cardinality-constrained quadratic optimization*. Comput. Optim. Appl., **43** (2009), n. 1, 1-22, https://doi.org/10.1007/s10589-007-9126-9.

F. Caraffini, F. Neri, M. Epitropakis: *HyperSPAM: a study on hyper-heuristic coordination strategies in the continuous domain*. Inf. Sci., **477** (2019), 186-202, https://doi.org/10.1016/j.ins.2018.10.033.

N. A. Canakgoz, J. E. Beasley: *Mixed-integer programming approaches for index tracking and enhanced indexation*. Eur. J. Oper. Res., **196** (2009), n. 1, 384-399, https://doi.org/10.1016/j.ejor.2008.03.015.

M. Caporin, G. Jannin, F. Lisi, B. Maillet: *A survey on the four families of performance measures*. J. Econ. Surv., **28** (2014), n. 5, 917-942, https://doi.org/10.1111/joes.12041.

R. Cheng, Y. Jin: *A competitive swarm optimizer for large scale optimization*. IEEE Trans. Cybern., **45** (2015), n. 2, 191-204, https://doi.org/10.1109/TCYB.2014.2322602.

R. Cheng, Y. Jin: *A social learning particle swarm optimization algorithm for scalable optimization*. Inf. Sci., **291** (2015), 43-60, https://doi.org/10.1016/j.ins.2014.08.039.

M. Corazza, G. di Tollo, G. Fasano, R. Pesenti: *A novel hybrid PSO-based metaheuristic for costly portfolio selection problems*. Ann. Oper. Res., **304** (2021), 109-137, https://doi.org/10.1007/s10479-021-04075-3.

M. F. P. Costa, R. B. Francisco, A. M. A. C. Rocha, E. M. G. P. Fernandes: *Theoretical and practical convergence of a self-adaptive penalty algorithm for constrained global optimization*. J. Optim. Theory Appl., **174** (2017), 875-893, https://doi.org/10.1007/s10957-016-1042-7.

Y. Crama, M. Schyns: *Simulated annealing for complex portfolio selection problems*. Eur. J. Oper. Res., **150** (2003), n. 3, 546-571, https://doi.org/10.1016/S0377-2217(02)00784-1.

T. Cura: *Particle swarm optimization approach to portfolio optimization*. Nonlinear Anal. Real World Appl., **10** (2009), n. 4, 2396-2406, https://doi.org/10.1016/j.nonrwa.2008.04.023.

R. Das, D. K. Prasad: *Prediction of porosity and thermal diffusivity in a porous fin using differential evolution algorithm*. Swarm Evol. Comput., **23** (2015), 27-39, https://doi.org/10.1016/j.swevo.2015.03.001.

R. Datta, K. Deb: *"Evolutionary Constrained Optimization"*. Infosys Science Foundation Series, Springer, New Delhi (2015).

O. Ertenlice, C. B. Kalayci: *A survey of swarm intelligence for portfolio optimization: algorithms and applications*. Swarm Evol. Comput., **39** (2018), 36-52, https://doi.org/10.1016/j.swevo.2018.01.009.

J. B. Guerard: *"Handbook of portfolio construction. Contemporary applications of Markowitz techniques"*. Springer, New York (2010).

Y. C. Ho, D. L. Pepyne: *Simple explanation of the no free lunch theorem of optimization*. Cybern. Syst. Anal., **38** (2002), n. 2, 292-298, https://doi.org/10.1023/A:1021251113462.

C. Israelsen: *A refinement to the Sharpe ratio and information ratio*. J. Asset Manag., **5** (2005), 423-427, https://doi.org/10.1057/palgrave.jam.2240158.

M. Kaucic: *Equity portfolio management with cardinality constraints and risk parity control using multi-objective particle swarm optimization*. Comput. Oper. Res., **109** (2019), 300-316, https://doi.org/10.1016/j.cor.2019.05.014.

M. Kaucic, F. Barbini, F. J. Camerota Verdù: *Polynomial goal programming and particle swarm optimization for enhanced indexation*. Soft Comput., **24** (2020), 8535-8551, https://doi.org/10.1007/s00500-019-04378-5.

P. N. Kolm, R. Tütüncü, F. J. Fabozzi: *60 years of portfolio optimization: practical challenges and current trends*. European J. Oper. Res., **234** (2014), 356-371, https://doi.org/10.1016/j.ejor.2013.10.060.

T. Krink, S. Mittnik, S. Paterlini: *Differential evolution and combinatorial search for constrained index-tracking*. Ann. Oper. Res. **172** (2009), 153-176, https://doi.org/10.1007/s10479-009-0552-1.

O. Ledoit, M. Wolf: *Honey, I shrunk the sample covariance matrix*. J. Portf. Manag., **30** (2004), n. 4, 110-119, https://doi.org/10.3905/jpm.2004.110.

K. Liagkouras, K. Metaxiotis: *Examining the effect of different configuration issues of the multiobjective evolutionary algorithms on the efficient frontier formulation for the constrained portfolio optimization problem*. J. Oper. Res. Soc. (2017), https://doi.org/10.1057/jors.2016.38.

E. Mezura-Montes, C. A. Coello Coello: *Constraint-handling in nature-inspired numerical optimization: past, present and future*. Swarm Evol. Comput., **1** (2011), n. 4, 173-194, https://doi.org/10.1016/j.swevo.2011.10.001.

R. Moral-Escudero, R. Ruiz-Torrubiano, A. Suarez: *Selection of optimal investment portfolios with cardinality constraints*. IEEE Trans. Evol. Comput. (2006), 2382-2388, https://doi.org/10.1109/CEC.2006.1688603.

E. T. Oldewage, A. P. Engelbrecht, C. W. Cleghorn: *The merits of velocity clamping particle swarm optimisation in high dimensional spaces*. IEEE Symposium Series on Computational Intelligence (2017), 1-8, https://doi.org/10.1109/SSCI.2017.8280887.

E. T. Oldewage, A. P. Engelbrecht, C. W. Cleghorn: *Movement patterns of a particle swarm in high dimensional spaces*. Inf. Sci., **512** (2020), 1043-1062, https://doi.org/10.1016/j.ins.2019.09.057.

K. E. Parsopoulos, M. N. Vrahatis: *On the computation of all global minimizers through particle swarm optimization*. IEEE Trans. Evol. Comput., **8** (2004), n. 3, 211-224, https://doi.org/10.1109/TEVC.2004.826076.

A. Ratnaweera, S. K. Halgamuge, H. C. Watson: *Self-organizing hierarchical particle swarm optimizer with time-varying acceleration coefficients*. IEEE Trans. Evol. Comput., **8** (2004), n. 3, 240-255, https://doi.org/10.1109/TEVC.2004.826071.

F. Schuhmacher, M. Eling: *Sufficient conditions for expected utility to imply drawdown-based performance rankings*. J. Bank Financ., **35** (2011), n. 9, 2311-2318, https://doi.org/10.1016/j.jbankfin.2011.01.031.

F. Schuhmacher, M. Eling: *A decision-theoretic foundation for reward-to-risk performance measures*. J. Bank Financ., **36** (2012), n. 7, 2077-2082, https://doi.org/10.1016/j.jbankfin.2012.03.013.

D. X. Shaw, S. Liu, L. Kopman: *Lagrangian relaxation procedure for cardinality-constrained portfolio optimization*. Optim. Methods Softw., **23** (2008), n. 3, 411-420, https://doi.org/10.1080/10556780701722542.

W. Shen, J. Wang, S. Ma: *Doubly regularized portfolio with risk minimization*. AAAI Conference on Artificial Intelligence, **28** (2014), n. 1, 1286-1292, https://doi.org/10.1609/aaai.v28i1.8906.

K. Singh, R. Das: *An experimental and multi-objective optimization study of a forced draft cooling tower with different fills*. Energy Conv. Manag., **111** (2016), 417-430, https://doi.org/10.1016/j.enconman.2015.12.080.

G.-W. Song, Q. Yang, X.-D. Gao, Y.-Y. Ma, Z.-Y. Lu, J. Zhang: *An adaptive level-based learning swarm optimizer for large-scale optimization*. IEEE Trans. Syst., Man, Cybern. (2021), 152-159, https://doi.org/10.1109/SMC52423.2021.9658644.

E. T. van Zyl, A. P. Engelbrecht: *A subspace-based method for PSO initialization*. IEEE Symposium Series on Computational Intelligence (2015), 226-233, https://doi.org/10.1109/SSCI.2015.42.

F. Wang, X. Wang, S. Sun: *A reinforcement learning level-based particle swarm optimization algorithm for large-scale optimization*. Inf. Sci., **602** (2022), 298-312, https://doi.org/10.1016/j.ins.2022.04.053.

M. Woodside-Oriakhi, C. Lucas, J. E. Beasley: *Heuristic algorithms for the cardinality constrained efficient frontier*. European J. Oper. Res. **213** (2011), n. 3, 538-550, https://doi.org/10.1016/j.ejor.2011.03.030.

Y. Xue, Y. Tong, F. Neri: *An ensemble of differential evolution and Adam for training feed-forward neural networks*. Inf. Sci., **608** (2022), 453-471, https://doi.org/10.1016/j.ins.2022.06.036.

X.-S. Yang: *Firefly algorithm, stochastic test functions and design optimisation*. Int. J. Bio-Inspired Comput., **2** (2010), n. 2, 78-84, https://doi.org/10.1504/IJBIC.2010.032124.

Q. Yang, W.-N. Chen, J. Da Deng, Y. Li, T. Gu, J. Zhang: *A level-based learning swarm optimizer for large-scale optimization*. IEEE Trans. Evol. Comput., **22** (2018), n. 4, 578-594, https://doi.org/10.1109/TEVC.2017.2743016.

J. Zhang, T. Leung, A. Aravkin: *A relaxed optimization approach for cardinality-constrained portfolios*. European Control Conference, Italy (2019), https://doi.org/10.23919/ECC.2019.8796164.

H. Zhu, Y. Wang, K. Wang, Y. Chen: *Particle Swarm Optimization (PSO) for the constrained portfolio optimization problem*. Expert Syst. Appl., **38** (2011), n. 8, 10161-10169, https://doi.org/10.1016/j.eswa.2011.02.075.

2 A constrained swarm optimization algorithm for large-scale long-run investments using Sharpe ratio-based performance measures

The objective of this chapter is to study large-scale portfolio optimization problems in which the objective function to maximize consists in a multi-moment performance measure that extends the Sharpe ratio. More specifically, we consider the adjusted for skewness Sharpe ratio, which incorporates the third moment of the returns distribution, and the adjusted for skewness and kurtosis Sharpe ratio, which exploits in addition the fourth moment.

In the model formulation, we account for two types of real-world trading constraints. On the one hand, we impose stock market restrictions through cardinality, buy-in thresholds, and budget constraints. On the other hand, a turnover threshold restricts the total allowed amount of trades in the rebalancing phases.

To deal with the complexity of the proposed asset allocation models, we embed a novel hybrid constraint-handling procedure into an improved dynamic level-based learning swarm optimizer (shortly, LLSO). A repair operator maps candidate solutions onto the set characterized by the first type of constraints. Unlike the LLSO presented in the previous chapter, this version employs an adaptive ℓ_1-exact penalty function to manage turnover violations.

The focus of the experimental part is to highlight the importance of including higher-order moments in the performance measures for long-run investments, to capitalize especially when the market is turbulent. We carry out empirical tests on two worldwide sets of assets, namely the MSCI World and the MSCI Pacific. On the one hand, we aim to pointing out the capabilities of the proposed dynamic LLSO algorithm in comparison to a state-of-the-art solver, which has been ad-hoc developed to solve cardinality-constrained portfolio optimization problems. On the other hand, we illustrate the scalability and effectiveness of the proposed portfolio strategies. The computational analysis reveals that the inclusion of higher-order moments in the performance measures produces superior results in terms of net wealth with respect to the benchmark and to the portfolio optimized through the standard Sharpe ratio. This is more evident after the pandemic outbreak of 2020, where more market fluctuations are present.

2.1 Introduction

In the portfolio selection problem, two phases are usually involved. The first one is devoted to the selection of promising assets, the second one focuses on the allocation of capital among them. In the modern portfolio theory, the so-called mean-variance analysis developed by Markowitz (1952) represents a milestone paper, and it has gained widespread acceptance as a practical tool for portfolio optimization among researchers and practitioners (Guerard (2010)). In this model the mean return of a portfolio represents the profit measure, while the portfolio variance is the risk. Accordingly, a portfolio is efficient if it provides the maximum return for a given level of risk or, equivalently, if it has the minimum risk for a given level of return. The set of optimal mean-variance tradeoffs in the risk-return space forms the efficient frontier. To guide the choice of an investor among these efficient portfolios, Sharpe (1994) has introduced a performance measure defined as the ratio between the excess return of an investment with respect to a risk-free asset and its standard deviation. Sharpe ratios with higher values correspond to more promising alternatives. In this manner, only the first two moments of the returns distribution are involved in the portfolio selection problem. However, the issue of whether higher moments should be considered to properly represent the investors' behaviour has been widely debated in literature and it is still open. The pioneering studies by Arditti (1967) and Samuelson (1970) have pointed out the importance of the third-order central moment of the returns distribution in the portfolio allocation process. In the same direction, Scott and Horvath (1980) have observed that a positive preference for skewness and a negative preference for kurtosis, known as prudence and temperance respectively (Kimball (1990)), properly explain the behaviour of investors. Thus, to take advantage of the potential upside return represented by positive skewness and the possible benefit given by small kurtosis, in the mean-variance framework, two novel performance measures have been recently proposed in literature, namely the adjusted for skewness Sharpe ratio (Zakamuline and Koekebakker (2009)) and the adjusted for skewness and kurtosis Sharpe ratio (Pézier and White (2008)). The former multiplies the classical Sharpe index by a factor linked to the portfolio skewness, while the latter extends the Sharpe ratio by including a factor that incorporates both skewness and kurtosis.

These two Sharpe ratio-based performance measures have been used for stock performance evaluation (Nagy and Benedek (2021)) but, to the best of our knowledge, they have not been applied as objective functions into the portfolio optimization process. Therefore, we propose portfolio selection strategies which maximize the aforementioned performance measures, considering four types of real-world constraints. A cardinality constraint is used to manage the portfolio size, buy-in threshold constraints ensure that all the available capital is invested. To characterize the investment profile, we introduce a set of bound constraints and a turnover threshold. In this manner, thanks to the former constraints, we avoid both the concentration of money in a few

assets and its splitting into too many assets. Further, due to the turnover bound, we limit the possibility of portfolio changes over time. The proposed asset allocation problem is analyzed from the perspective of an institutional investor who operates in equity markets with hundreds or thousands of constituents and selects a restricted pool of stocks to build up an active portfolio.

The introduction of cardinality constraints in the portfolio design leads to optimization problems for which finding optimal solutions becomes computationally challenging (Moral-Escudero et al. (2006)). For this reason, in recent years, swarm optimization algorithms, inspired by the self-organizing interaction among agents, have become popular for this topic. In particular, the particle swarm optimization (PSO) algorithm has shown a good capability in solving small and mid-size portfolio allocation models (Cura (2009); Zhu et al. (2011); Kaucic et al. (2020); Corazza et al. (2021)). This algorithm, mimicking the swarm behaviour of social animals such as bird flocking, gathers the information about good solutions through the swarm, and floats in the whole search space to find the global solution of the problem (Wang et al. (2018)). However, PSO does not work efficiently in solving high-dimensional optimization problems, due to the so-called curse of dimensionality (Gilli and Schumann (2012); Oldewage et al. (2020)). To overcome this issue, Yang et al. (2018) have developed a variant of the level-based learning swarm optimizer (LLSO), which exhibits superiority in achieving higher quality solutions with respect to other competitors in the literature for large scale optimization problems. The algorithm is based on the teaching paradigm where individuals are divided into levels according to their fitness and they are treated differently. The most performing candidates are stored in higher levels and guide the learning of the other particles in the swarm.

In this paper, we adopt a dynamic variant of the LLSO algorithm to solve our portfolio optimization problems, with a specific clamping and reversing procedure for the particles update rule, to improve the exploration efficiency. Moreover, since the LLSO is blind to the constraints, we equip it with a novel hybrid constraint-handling technique which works as follows. In order to deal with the cardinality constraint we use a projection operator, that selects the largest components of the candidate solutions and sets equal to zero the remaining ones. With this technique, we can relax the cardinality equality condition proposed in Kaucic and Piccotto (2022) as an inequality, assuming that the number of stocks included in the portfolio is lower than or equal to a fixed threshold. Furthermore, through this process, we transform the original mixed-integer optimization problem into a problem involving only real variables. Then, buy-in thresholds and budget constraints are handled using the repair operator proposed in Meghwani and Thakur (2017). Finally, to control the turnover constraint, a ℓ_1-exact penalty function method is adopted, as in Corazza et al. (2021).

Summing up, the contribution of this work to the current literature is threefold. On the one hand, it is the first time that the adjusted for skewness Sharpe ratio and the adjusted for skewness and kurtosis Sharpe ratio measures are employed in the portfolio

optimization problem. On the other hand, regarding the algorithmic novelties, we propose an improved variant of the LLSO equipped with a novel ad-hoc constraint-handling procedure which involves a repair operator as well as an ℓ_1-exact penalty function strategy. Moreover, from a practical point of view, we study the robustness of the proposed multi-moment strategies by comparing their profitability in a long-run setting, with almost 14 years of observations, involving also the recent phases of large market fluctuation due to the COVID-19 pandemic and the Ukrainian crisis.

The remainder of the paper is organized as follows. In the next section we review some literature related to our work. In Section 2.3 we introduce the investment framework and the objective functions involved. Section 2.4 presents the improved LLSO algorithm with the novel hybrid constraint-handling technique. Section 2.5 is devoted to the experimental analysis while the conclusions and future works are reported in Section 2.6.

2.2 Related works

In this section, we first survey the foremost contributions appeared in literature on the multi-moment formulations of the portfolio optimization problem. Since a complete review of the most recent population-based heuristics for large-scale optimization problem can be found in Omidvar et al. (2022a) and Omidvar et al. (2022b), we focus solely on the papers related to our study and concerning the PSO improvements.

2.2.a Multi-moment portfolio optimization models

In order to highlight the critical effects of prudence and temperance on investment decisions, and to provide a more complete characterization of investor preferences, many authors have revised the mean-variance framework by incorporating the third and fourth moments when constructing a portfolio (see Jurczenko and Maillet (2006) and references therein). Lai (1991) has introduced a new term in the objective function for including skewness into the asset allocation problem. Konno et al. (1993) have developed a model in which they maximize the third-order moment given a threshold for portfolio expected return and for the variance. In addition, Liu et al. (2003) have employed also a transaction costs constraint. Lai et al. (2006) have proposed a multi-objective portfolio optimization problem which involves the first four central moments of the portfolio return distributions.

The direct optimization of the third and the fourth moment terms in the problem formulation has been shown to be computationally demanding, due to the difficulty in obtaining reliable estimators for the co-skewness and co-kurtosis matrices, as highlighted in Kim et al. (2014). In some recent contributions, Chaigneau and Eeckhoudt (2020) and Gao et al. (2022) have extended the mean-variance framework by considering alternative measures which incorporate prudence and temperance in risk exposure.

Finally, several experimental investigations in the field of behavioural finance have been carried out in the last years, pointing out the importance of prudence, temperance, and higher-order preferences on the investors' behaviour (see, among others, Colasante and Riccetti (2020, 2021)).

2.2.b PSO enhancements for large-scale optimization

In literature, there are two major algorithmic approaches to solve the curse of dimensionality for PSO, namely decomposition-based and non decomposition-based approaches. The first type of procedures separate a high-dimensional problem into several small-dimensional instances using a divide-and-conquer strategy to reduce the dimensionality. In this direction, Van der Bergh and Engelbrecht (2004) and Li and Yao (2012) have proposed cooperative co-evolutionary particle swarm optimization algorithm, which randomly divides the decision variables into subgroups and then uses PSO to optimize each subgroup separately. The second type of algorithms directly optimizes all the variables at the same time, employing a learning mechanism to properly balance diversity and convergence. In this context, several learning strategies have been proposed recently. For instance, the competitive swarm optimizer (CSO, Cheng and Jin (2015a)) compares two randomly chosen particles, and then the superior particle guides the update of the inferior one. Inspired by social animal behaviors, Cheng and Jin (2015b) have proposed the social learning particle swarm optimizer (SL-PSO) which first sorts particles by fitness and then worse individuals learn from the better ones. Following the teaching concept that teachers should treat students in accordance with their abilities, Yang et al. (2018) have developed the so-called level-based learning swarm optimizer (LLSO). The optimal compromise between exploration and exploitation of this learning technique guarantees more accurate solutions than the above cited CSO and SL-PSO. For this reason, several extensions of LLSO have been developed. The dynamic LLSO (Yang et al. (2018)) dynamically adjusts the number of groups in which the particle swarm is divided based on the performance of the algorithm over time. Due to the oversensitivity of the standard LLSO to the parameter setting, Song et al. (2021) have proposed an adaptive variant in which the evolution state of the swarm is adjusted on the information given by the swarm aggregation indicator. Similarly, the reinforcement learning level-based particle swarm optimization algorithm by Wang et al. (2022) introduces a reinforcement learning strategy to control the number of levels and to improve the search efficiency.

2.3 Investment framework

In this study, we consider a frictionless market where short selling is not allowed and all investors act as price takers. The investable universe is represented by n risky assets. A portfolio is denoted by the vector of its assets weights $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n$. In

our dynamic setting, portfolio weights are periodically rebalanced, with an investment horizon of length h. We observe the market over a time window $\mathcal{T} = \{0, 1, \dots, T\}$ and we adopt the following two-steps scheme for the investment strategy:

1. at time T the optimal portfolio composition is determined using a scenario-based approach;
2. the same portfolio composition is retained until time $T + h$, assuming that the stocks selected are still available at $T + h$.

The prices of the n risky assets are available for the time window $\mathcal{T}$. Then, we define the observed price of asset i at time t, $t \in \mathcal{T}$ and $i = 1, \dots n$ with $p_{i,t}$, and the realized rate of return at time t, with $t \geq 1$, as $r_{i,t} = {}^{p_{i,t}}/_{p_{i,t-1}} - 1$. Based on this information, we aim to identify the optimal allocation at time T that guarantees the best performance at time $T + h$.

Let now $(\Omega, \mathcal{F}, P)$ be the probability space on which we assume the random variables are defined. We denote by:

- $R_i^{(h)}$ the random variable representing the rate of return of asset i at future time $T + h$, with expected value μ_i;
- $R_{\boldsymbol{x}}^{(h)} = \sum_{i=1}^{n} x_i R_i^{(h)}$ the random variable that expresses the rate of return of portfolio $\boldsymbol{x}$ at future time $T + h$.

We repeat this investment procedure over time, updating the observation window $\mathcal{T}$ by eliminating the h oldest observations and adding the most recent ones. For semplicity of notation, we will indicate the random rate of return of portfolio $\boldsymbol{x}$ at $T+h$, $R_{\boldsymbol{x}}^{(h)}$, by $R_{\boldsymbol{x}}$.

2.3.a Sharpe ratio-based performance measures

In this section, we will introduce the three performance measures that we will consider in our portfolio optimization problems.

Sharpe ratio

In the asset allocation problem proposed by Markowitz (1952), portfolio risk is represented by the volatility, given by $\sigma(R_{\boldsymbol{x}}) = \sqrt{\sum_{i=1}^{n} \sum_{j=1}^{n} c_{ij} x_i x_j}$, where $(C)_{ij} = c_{ij}$ is the covariance between stocks i and j, with $i, j = 1, \dots, n$. In this model, the portfolio choice is made solely with respect to the expected rate of return μ of the portfolio $\boldsymbol{x}$ and its risk, where a large portfolio volatility is perceived as damaging by the investors. As mentioned in the introduction, the Sharpe ratio is then defined as

$$SR(\boldsymbol{x}) = \frac{\mu - r_f}{\sigma(R_{\boldsymbol{x}})} \tag{2.3.1}$$

where r_f is a risk-free rate. This performance measure can be interpreted as the compensation earned by the investor per unit of risk. Thus, higher values of SR indicate more promising portfolios and are preferred by rational investors.

Even if it has an easy interpretation, the Sharpe ratio presents some pitfalls. The first is the choice of the risk-free rate used to define the excess rate of return. The debate among scholars and practitioners is still open. For instance, Hitaj and Zambruno (2016) consider $r_f = 0$ as a reasonable value. Similarly, Amédée-Manesme and Barthélémy (2022) set exogenously $r_f = 2\%$. Alternatively, Deguest et al. (2022) suggest the use of 1 month or 3 month maturity US Treasury Bills. Since these sovereign bonds exhibited values close to zero during a large part of the investment period analyzed in this paper, in the empirical part we follow Hitaj and Zambruno (2016) and set $r_f = 0$.

Another problem related to this indicator is its incoherence with preference relations in periods of market downturns, when the expected excess return of the portfolio is negative. In these cases, a Sharpe ratio-oriented agent could select portfolios with higher volatility. To overcome this issue, we consider only the first moment, that means we prefer, among two portfolios with negative expected excess rate of return, the one with the less expected loss.

Adjusted for skewness Sharpe ratio

As observed in the introduction, a possible drawback of the Sharpe ratio is that it uses only the first two moments of the portfolio returns distribution, and does not consider the potential upside return represented by positive skewness. For this reason, Zakamuline and Koekebakker (2009) have derived an adjustment to the Sharpe ratio for skewness, which depends on the investor's utility function.

The proposed alternative performance measure is called adjusted for skewness Sharpe ratio (ASSR), and is given by

$$ASSR_b(\boldsymbol{x}) = SR(\boldsymbol{x})\sqrt{1 + b\frac{S_3(R_{\mathbf{x}})}{3}SR(\mathbf{x})} \tag{2.3.2}$$

where b expresses the individual's relative preference to the third moment of the returns distribution, and $S_3(R_{\mathbf{x}})$ is the skewness, defined by $E\left[\left(\frac{R_{\mathbf{x}}-\mu}{\sigma(R_{\boldsymbol{x}})}\right)^3\right]$.

The properties of this performance measure have been investigated by Cheridito and Kromer (2013). In particular, we can note that the $ASSR_b$ preserves the standard Sharpe ratio for zero skewness, while it is higher if the skewness is positive. However, in order to compute the $ASSR_b$ one needs to determinate the value of b, which depends on the choice of the utility function. Thus, this performance measure is not unique for all investors, but it is rather an individual performance measure. In this work, we will set $b = 1$, meaning that we consider an investor with exponential utility, and we will indicate the above quoted performance measure by $ASSR$.

Adjusted for skewness and kurtosis Sharpe ratio

In order to account for both skewness and kurtosis, Pézier and White (2008) have proposed the so-called adjusted for skewness and kurtosis Sharpe ratio ($AKSR$), that is

$$AKSR(\boldsymbol{x}) = SR(\boldsymbol{x}) \left[1 + \frac{S_3(R_{\mathbf{x}})}{3!} SR(\boldsymbol{x}) - \left(\frac{K_4(R_{\mathbf{x}}) - 3}{4!}\right) SR(\boldsymbol{x})^2\right] \tag{2.3.3}$$

where $K_4(R_{\mathbf{x}})$ is the kurtosis, defined by $E\left[\left(\frac{R_{\mathbf{x}}-\mu}{\sigma(R_{\boldsymbol{x}})}\right)^4\right]$.

According to this performance measure, an investor prefers portfolios with higher skewness and dislikes portfolios with kurtosis values higher than 3, meaning that leptokurtic portfolio distributions are penalized in order to avoid extreme events.

2.3.b Portfolio optimization model

Now, we introduce the family of constraints used to define the set of admissible portfolios.

1) *Budget constraint*. We require that all the available capital is invested. In terms of portfolio weights this translates into

$$\sum_{i=1}^{n} x_i = 1. \tag{2.3.4}$$

2) *Cardinality constraint*. We assume that the portfolio includes up to K stocks out of the n available, where $K < n$ is a predefined number. To model the inclusion or the exclusion of the i-th asset in the portfolio, an auxiliary variable δ_i is defined as follows

$$\delta_i = \begin{cases} 1, \text{ if asset } i \text{ is included} \\ 0, \text{ otherwise} \end{cases}$$

for $i = 1, \ldots, n$. The resulting vector of selected assets is $\delta = (\delta_1, \ldots, \delta_n)$. We can write the cardinality constraint as

$$\sum_{i=1}^{n} \delta_i \leq K. \tag{2.3.5}$$

3) *Box constraints*. To avoid extreme positions and foster diversification, we introduce a maximum and a minimum limit for the wealth allocation in the i-th stock included in the portfolio. Let l_i and u_i be respectively the lower and the upper bound for the weight of the i-th asset, with $0 < l_i < u_i \leq 1$, then we can write the box constraints as

$$\delta_i l_i \leq x_i \leq \delta_i u_i, \quad i = 1, \ldots, n\,. \tag{2.3.6}$$

Note that if an asset is not included in the portfolio, no capital is invested on it.

4) *Turnover constraint*. In every rebalancing phase, the portfolio composition used in the previous investment window, denoted by $\boldsymbol{x}_0$, is updated. Let $\boldsymbol{x}$ be the vector of weights in the rebalanced portfolio and $\widetilde{\boldsymbol{x}}_0$ be the vector of re-normalized weights associated to $\boldsymbol{x}_0$ (see Shen et al. (2014)), which is calculated component-wise as

$$\widetilde{x}_{0,i} = \frac{x_{0,i}(r_i + 1)}{\sum_{j=1}^{n} x_{0,j}(r_j + 1)}$$

and r_i, $i = 1, \ldots, n$ is the rate of return of the i-th stock in the portfolio at the moment of the rebalancing phase. Then, the portfolio turnover constraint is given by

$$\sum_{i=1}^{n} |x_i - \widetilde{x}_{0,i}| \leq TR \tag{2.3.7}$$

where TR denotes the maximum turnover rate, which lies between 0 and 1. It can be noted that if $TR = 0$ rebalancing is not allowed, and as TR increases more trades are allowed.

We indicate with $\mathcal{X}$ the feasible set comprising the pairs $(\delta, \boldsymbol{x})$ that satisfy (2.3.4), (2.3.5), (2.3.6) and (2.3.7).
Summing up, our asset allocation model can be written as

$$\begin{aligned} \max \quad & \Phi(\boldsymbol{x}) \\ \text{s.t.} \quad & (\delta, \boldsymbol{x}) \in \mathcal{X} \end{aligned}$$

where $\Phi(\boldsymbol{x})$ is one of the three Sharpe ratio-based performance measures introduced in Section 2.3.a. As it is customary in the programming literature, we transform this maximization problem into the equivalent minimization instance

$$\begin{aligned} \min \quad & f(\boldsymbol{x}) \\ \text{s.t.} \quad & (\delta, \boldsymbol{x}) \in \mathcal{X} \end{aligned} \tag{2.3.8}$$

where $f(\boldsymbol{x}) = -\Phi(\boldsymbol{x})$.

Scenario generation technique

To solve the optimization problem (2.3.8), we use the following scenario-based generation technique. To estimate the values of the considered performance measures for a given portfolio $\boldsymbol{x}$, we need to simulate the distribution of the h-step ahead rate of return $R_{\boldsymbol{x}}$. To this end, we consider the historical rates of return of the n risky assets realized on the time window $[0, T]$. We assume that historical observations are good proxies for the future rates of return. Then, we define a scenario as the set of the joint realizations of the rates of return for the n assets in a given time period. Due to the good performance of block bootstrapping techniques to preserve correlations between time series (see, for instance, Guastaroba et al. (2009)), we adopt the so-called stationary bootstrap

(Politis and Romano (1994)). This technique considers a random block size, that is a set of consecutive scenarios with variable length, in order to bring some robustness with respect to the standard block bootstrap, which uses fixed block size. The procedure works as follows. First, we select the optimal average block size B^* by the procedure developed in Politis and White (2004). Then, we extract randomly from the observed data frame a block of length B^*. We repeat this exercise until the extracted sample reaches the desired size h, adjusting the last block length if the procedure exceeds the desired number of periods.
After having a bootstrap sample of h rates of return for each asset i, denoted as $\widehat{R}_{i,t'}$, with $i = 1, \ldots, n$ and $t' = 1, \ldots, h$, we calculate the simulated h-step ahead rate of return of the i-th asset as $\widehat{R}_i = \prod_{t'=1}^{h} \left(1 + \widehat{R}_{i,t'}\right)$ and the simulated h-step ahead rate of return of portfolio $\boldsymbol{x}$ as $\widehat{R}_{\boldsymbol{x}} = \sum_{i=1}^{n} x_i \widehat{R}_i$. We repeat this procedure S times to have an estimate of the empirical distribution of $R_{\boldsymbol{x}}$. With an abuse of notation, let $\widehat{R}_{\boldsymbol{x}}(s)$ be the s-th simulation of the h-step ahead rate of return of portfolio $\boldsymbol{x}$. We can then calculate all the quantities used to evaluate the Sharpe ratio-based performance measures. Indeed, the sample mean of the empirical distribution is given by

$$\widehat{\mu} = \frac{1}{S} \sum_{s=1}^{S} \widehat{R}_{\boldsymbol{x}}(s)$$

and the sample standard deviation is

$$\widehat{\sigma} = \sqrt{\frac{1}{S-1} \sum_{s=1}^{S} (\widehat{R}_{\boldsymbol{x}}(s) - \widehat{\mu})^2} \,.$$

Similarly, we estimate the skewness and kurtosis as follows

$$\widehat{S}_3(R_{\boldsymbol{x}}) = \frac{1}{S-1} \sum_{s=1}^{S} \left(\frac{\widehat{R}_{\boldsymbol{x}}(s) - \widehat{\mu}}{\widehat{\sigma}} \right)^3$$

and

$$\widehat{K}_4(R_{\boldsymbol{x}}) = \frac{1}{S-1} \sum_{s=1}^{S} \left(\frac{\widehat{R}_{\boldsymbol{x}}(s) - \widehat{\mu}}{\widehat{\sigma}} \right)^4 .$$

2.4 Optimization algorithm

After introducing the LSSO paradigm, we present the dynamic LLSO with the proposed improvements as well as the hybrid constraint handling technique.

2.4.a Level-based learning swarm optimizer

The LLSO algorithm, developed by Yang et al. (2018), evolves a swarm $\mathcal{P}$ of NP candidate solutions using the so-called level-based population strategy. This process operates in accordance with the following two steps.

1. First, at each generation g the individuals in $\mathcal{P}$ are sorted ascending based on their fitness and grouped into NL_g levels. Each level contains $LP_g = \lfloor NP/NL_g \rfloor$ particles and in the last one, there are $\lfloor NP/NL_g \rfloor + NP\%NL_g$ candidate solutions[1]. Better individuals belong to higher levels, and a higher level corresponds to a smaller level index. Therefore, we denote with L_1 the best level and with L_{NL} the worst one.
2. Individuals belonging to the first level L_1 are not updated and directly enter into the next generation, because they represent the most valuable information conveyed in the swarm at the current generation. On the contrary, the p-th particle in level L_l, denoted by $\boldsymbol{x}^{l,p}(g)$, where $l = 3, \ldots, NL_g$ and $p = 1, \ldots, LP_g$, is allowed to learn from two particles $\boldsymbol{x}^{l_1,p_1}(g)$ and $\boldsymbol{x}^{l_2,p_2}(g)$. These two individuals are randomly extracted from two different higher levels L_{l_1} and L_{l_2}, with $l_1 < l_2$, and p_1, p_2 randomly chosen from $\{1, \ldots, LP_g\}$. For $l = 2$, we sample two particles from L_1 in such a way that $\boldsymbol{x}^{l_1,p_1}(g)$ is better than $\boldsymbol{x}^{l_1,p_2}(g)$ in terms of fitness function. Thus, the update rule for particle $\boldsymbol{x}^{l,p}(g)$ is given by

$$\mathbf{v}^{l,p}(g+1) = r_1\mathbf{v}^{l,p}(g) + r_2(\boldsymbol{x}^{l_1,p_1}(g) - \boldsymbol{x}^{l,p}(g)) + \psi r_3(\boldsymbol{x}^{l_2,p_2}(g) - \boldsymbol{x}^{l,p}(g)) \tag{2.4.9}$$

$$\boldsymbol{x}^{l,p}(g+1) = \boldsymbol{x}^{l,p}(g) + \mathbf{v}^{l,p}(g+1) \tag{2.4.10}$$

for $i = 1, \ldots, n$, where $\mathbf{v}^{l,p}(g)$ is the so-called velocity of particle p in level L_l at generation g, and r_1, r_2, r_3 are three real numbers randomly generated within $[0, 1]$. The initial velocities, at generation 0, are all set equal to the zero vector, that is $\mathbf{v}^{l,p}(0) = \mathbf{0}$. The parameter $\psi \in [0, 1]$ controls the influence of the less performing exemplar $\boldsymbol{x}^{l_2,p_2}(g)$ on $\mathbf{v}^{l,p}(g+1)$.

The algorithm repeats these two steps until a maximum number of generations, MAX_{GEN}, is reached.

Dynamic LLSO with clamping and reversion

Following the suggestions in Yang et al. (2018), we adopt a dynamic setting for the number of levels NL_g by designing a pool $S = \{l_1, \ldots, l_s\}$ containing s different candidate integers. Then, at each generation g, the algorithm selects one of the elements of S based on their probabilities. At the end of the generation, the performance of the algorithm with the current level number is recorded, in order to update the probability of this level number for the next generation selection. To compute the probabilities of the elements of S, a record list $\Upsilon_s = \{\gamma_1, \ldots \gamma_s\}$ is defined. Initially, each γ_i is set equal to 1. Then, the element γ_i corresponding to the level number l_i in the current generation is updated as follows

$$\gamma_i = \frac{|F - \tilde{F}|}{|F|} \tag{2.4.11}$$

[1]We denote by $\lfloor z \rfloor$ the floor of z and by $z\%y$ the rest of the division of z by y.

where F is the global best fitness of the last generation, and $\tilde{F}$ is the global best fitness of the current generation. Then, the i-th element of the probability vector $P_s = \{p_1, \ldots, p_s\}$ is computed as

$$p_i = \frac{e^{7 \cdot \gamma_i}}{\sum_{j=1}^{s} e^{7 \cdot \gamma_j}} \tag{2.4.12}$$

with $i = 1, \ldots, s$. At each generation, based on P_s, an integer from the pool S is selected as the level number following a roulette wheel scheme.

From now on, we will remove the dependency on the generation g if it will be clear from context.

Based on the results of a preliminary analysis, we introduce a "clamping and reversion" procedure to increase the exploration capability of the LLSO. The clamping mechanism is applied component-wise to the velocity vector in (2.4.9) as follows

$$v_i^{l,p} = \min\{\max\{v_i^{l,p}, v_i^{\min}\}, v_i^{\max}\} \tag{2.4.13}$$

where $v_i^{\min}$ and $v_i^{\max}$ are the minimum and the maximum velocity allowed for component i, with $i = 1, \ldots, n$. In the experimental part, recalling (2.3.6), we set the maximum velocity as $v_i^{\max} = u_i$ and the minimum velocity as $v_i^{\min} = -v_i^{\max}$. Moreover, when $v_i^{l,p}$ and $x_i^{l,p}$ are both negative, we reverse and scale $v_i^{l,p}$ as follows:

$$v_i^{l,p} = -r_4 \cdot v_i^{l,p} \tag{2.4.14}$$

where r_4 is a random number uniformly generated in $[0, 1]$.

2.4.b Hybrid constraint-handling procedure

In the so-called construction phase, admissible portfolios have to satisfy cardinality, buy-in threshold, budget and turnover constraints. However, the LLSO update procedure is blind to these constraints. To overcome this issue, we equip the solver with a hybrid constraint-handling procedure.

First, in order to assure the cardinality requirement, for each candidate solution $\boldsymbol{x}^p$, the K largest components enter the corresponding portfolio, while zero weights are assigned to the remaining $n - K$. In this manner, we implicitly remove the binary variables from problem (2.3.8).

To guarantee the feasibility with respect to the bound constraints (2.3.6), we consider the following projection

$$x_i^p = \min\left\{\max\left\{x_i^p, l_i\right\}, u_i\right\} \tag{2.4.15}$$

where $p = 1, \ldots, NP$ and $i \in I_K^p = \{i = 1, \ldots, n \colon x_i^p > 0\}$. Note that $|I_K^p| \leq K$. Then, we use the repair transformations developed in Meghwani and Thakur (2017) to also satisfy the budget constraint (2.3.4). More precisely, for each $p = 1, \ldots, NP$, assuming that l_i and u_i are such that $\sum_{i \in I_K^p} l_i < 1$ and $\sum_{i \in I_K^p} u_i > 1$, we adjust the candidate

solution $\boldsymbol{x}^p$ component-wise:

$$x_i^p = \begin{cases} l_i + \dfrac{(x_i^p - l_i)}{\sum_{j \in I_K^p}(x_j^p - l_j)} \left(1 - \sum_{j \in I_K^p} l_j\right) & \text{, if } \sum_{j \in I_K^p} x_j^p > 1 \\ x_i^p, \quad \text{if } \sum_{j \in I_K^p} x_j^p = 1 & \\ u_i - \dfrac{(u_i - x_i^p)}{\sum_{j \in I_K^p}(u_j - x_j^p)} \left(\sum_{j \in I_K^p} u_j - 1\right) & \text{, if } \sum_{j \in I_K^p} x_j^p < 1 \end{cases} \tag{2.4.16}$$

for all $i \in I_K^p$. As proved in Meghwani and Thakur (2017), solutions transformed through (2.4.16) fulfill at the same time budget and box constraints.

Finally, to handle the turnover constraint, we use the ℓ_1-exact penalty function approach as in Corazza et al. (2021). In particular, we define the constraint violation of (2.3.7) as

$$CV = \max\left\{\sum_{i=1}^{n} |x_i - \tilde{x}_{0,i}| - TR, 0\right\}. \tag{2.4.17}$$

Then, we introduce the ℓ_1-exact penalty function

$$F_{\ell_1}(\boldsymbol{x}, \varepsilon_0(g), \varepsilon_1(g)) = f(\boldsymbol{x}) + \frac{\varepsilon_1(g)}{\varepsilon_0(g)} CV \tag{2.4.18}$$

where ε_0 and ε_1 are two positive real numbers, defined adaptively at each generation g. Initially, $\varepsilon_0(0) = 10^{-4}$ and $\varepsilon_1(0) = 1$ in order to privilege feasible solutions. This parameters are then updated by checking the decrease of the objective function $f(\boldsymbol{x})$ and the violation of the constraints. More precisely, every 5 iterations $\varepsilon_0(g)$ is updated according to the rule

$$\varepsilon_0(g+1) = \begin{cases} \min\{3 \cdot \varepsilon_0(g),\, 1\} & \text{if} \quad f(\boldsymbol{x}(g)) \geq f(\boldsymbol{x}(g-1)) \\ \max\{0.6 \cdot \varepsilon_0(g),\, 10^{-15}\} & \text{if} \quad f(\boldsymbol{x}(g)) < 0.9 \cdot f(\boldsymbol{x}(g-1)) \\ \varepsilon_0(g) & \text{otherwise} \end{cases} \tag{2.4.19}$$

while, every 10 iterations, $\varepsilon_1(g)$ is updated following the scheme

$$\varepsilon_1(g+1) = \begin{cases} \min\{2 \cdot \varepsilon_1(g),\, 10^4\} & \text{if} \quad CV(g) > 0.95 \cdot CV(g-1) \\ \max\{0.5 \cdot \varepsilon_1(g),\, 10^{-4}\} & \text{if} \quad CV(g) < 0.9 \cdot CV(g-1) \\ \varepsilon_1(g) & \text{otherwise.} \end{cases} \tag{2.4.20}$$

With this strategy, we privilege optimality of solutions possibly at the expenses of their feasibility, due to the fact that $\varepsilon_0(g+1)$ in (2.4.19) is increasing in $F_{\ell_1}(\boldsymbol{x}, \varepsilon_0(g+1), \varepsilon_1(g+1))$ when the function value $f(\boldsymbol{x}(g))$ increases. Moreover, to favour feasibility of solutions

possibly at the expenses of their optimality, the penalty parameter $\varepsilon_1(g+1)$ in (2.4.20) increases when the relative constraint violation in the g-th generation increases with respect to the previous one.

Using the penalty approach, the constrained optimization problem (2.3.8) reduces to an unconstrained one, in which we minimize F_{ℓ_1} and thus it can be solved by the proposed LLSO variant.

2.4.c Initialisation strategy

Due to the complexity of the problem and to the fact that the search space grows exponentially with the dimension, the common strategies of seeking a search space coverage by initializing the particles uniformly throughout the space are inefficient (see Van Zyl and Engelbrecht (2015)). In particular, for our portfolio optimization problems, the presence of the turnover constraint exacerbates even more the initialization phase. To address this issue, we directly initialize the candidate solutions in a neighbourhood of $\boldsymbol{x}_0$ as proposed by Kaucic et al. (2022). A brief description of the procedure follows. Let d_i^{min} and d_i^{max} be the minimum and the maximum allowed weight changes for $\boldsymbol{x}_{0,i}$ respectively, with $i = 1, \dots, n$. Let D^p denote the total portfolio weight allowed to be re-allocated in $\boldsymbol{x}_0$ for defining the p-th candidate solution $\boldsymbol{x}^p(0)$ at generation 0, with $p = 1, \dots, NP$. Then, for each p,

1. we randomly select D^p within $[0,\ {}^{TR}/_2]$;

2. we select a subset J^- of K' assets from the assets with positive weight in $\boldsymbol{x}_0$, so that
$$x_j^p(0) = x_{0,j} - d_j, \ \text{ for } j \in J^-$$
where d_j is randomly sampled in $\left[d_j^{min},\ d_j^{max}\right]$ in such a way that $\sum_{j\in J^-} d_j = D^p$, and $x_j^p(0) = 0$ or $l_j \le x_j^p(0) \le u_j$;

3. we select a subset J^+ of K'' assets from the assets with zero weight in $\boldsymbol{x}_0$, with $K'' \le K'$, so that
$$x_j^p(0) = x_{0,j} + d_j, \ \text{ for } j \in J^+$$
where d_j is randomly sampled in $[d^{min},\ d^{max}]$ in such a way that $\sum_{j\in J^+} d_j = D^p$, and $l_j \le x_j^p(0) \le u_j$;

4. for $j \in I \setminus (J^- \cup J^+)$, we set $x_j^p(0) = x_{0,j}$.

The portfolios assembled using this scheme satisfy cardinality, buy-in threshold and turnover constraints. This initialization strategy encourages the swarm to focus on exploitation rather than exploration, and allows to identify promising solutions, even in problems with high dimension and small feasible regions.

2.5 Computational analysis

2.5.a Data description and portfolio parameters

We select two equity investment universes that differ for the number of constituents and for the geographic area, to highlight the scalability of the proposed portfolio strategies. The first data set, called Pacific, consists of 323 assets selected among the constituents of the MSCI Pacific Index at 28/07/2022. For the second data set, called shortly World, we consider 1229 assets listed in the MSCI World Index at 28/07/2022.
We obtain the daily rates of return and the market values for each asset from Bloomberg, covering the period 01/01/2008 to 28/07/2022 for a total of 3803 observations. For comparison purposes, we build up an auxiliary market value-weighted benchmark for each data set.

Table 2.1 reports a preliminary analysis concerning normality assumption for the time series of rates of return in the two data sets. In both cases, the number of assets exhibiting high skewness is around 20%, while those with large kurtosis is close to 90%. The Jarque-Bera (J-B) test rejects the null hypothesis of normality at the 5% significance level in almost half of both samples.

Table 2.1: Preliminary analysis on the normality assumption of the assets in the two data sets. The first column reports the name of the data set considered, while the corresponding percentage of assets with high skewness and large kurtosis is given in the second and third columns, respectively. The last column displays the rejections percentage of the null hypothesis of normality for the Jarque-Bera (J-B) test at 5% significance level.

Data set name	$\lvert S_3 \rvert > 0.5$	$K_4 > 3$	J-B rejections
Pacific	19%	88%	46%
World	22%	89%	53%

The goal of our computational analysis is twofold. On the one hand, we aim at pointing out the capabilities of the proposed dynamic LLSO algorithm in comparison to a state-of-the-art solver, which has been ad-hoc developed to solve cardinality-constrained portfolio optimization problems (see Corazza et al. (2021)). On the other hand, we study the profitability of the proposed investment strategies focusing on the impact of both portfolio size and amount of trades.

To this end, we consider an investment plan with monthly portfolio rebalancing. The out-of-sample window is given by 126 months, covering the period from 02/01/2012 to 28/07/2022. For each month in the out-of-sample window, we generate 1000 scenarios of monthly rates of return for the assets in each data set by using the stationary bootstrap technique introduced above. The procedure employs an in-sample window of 1000 days, which is updated monthly by including the daily rates of return of the last month and by removing the information about the oldest month.

For the analysis, we set the buy-in thresholds l_i and u_i equal to 0.001 and 0.2 respectively, according to Kaucic and Piccotto (2022). We express the cardinality parameter K as a fraction $K_\%$ of the number of assets in a given data set, that is $K = \lfloor K_\% \cdot n \rfloor$, and pick up $K_\%$ in the set $\{15\%, 30\%, 50\%\}$. The turnover rate TR is in $\{10\%, 20\%, 40\%\}$. For each data set and each out-of-sample month, an instance of Problem (2.3.8) is then obtained by fixing one of the three Sharpe-based performance measures, a value of $K_\%$ and TR, for a total of 27 alternative investment schemes.

Finally, the role of the trades is analyzed ex-post through the cost function introduced in Beraldi et al. (2021), with an initial wealth $W_0 = 10\,000\,000\,\$$.

2.5.b Algorithm performance evaluation

In this subsection, we compare the performance of the proposed dynamic LLSO algorithm with respect to the PSO developed by Corazza et al. (2021). The latter has been showed to tackle efficiently non-smooth portfolio optimization problems with real-world constraints.

For each data set, the test suite consists of the 27 instances of Problem (2.3.8) previously introduced and specified at three dates randomly drawn from the out-of-sample window. For each sampled date and each portfolio optimization problem, we compute the initial portfolio without including the turnover constraint. The next out-of-sample month, we optimize the portfolio weights, accounting for the rebalancing constraint. It is worth noticing that we select three different dates for the evaluations in order to avoid the possible time dependence of the results.

The parameter settings for the considered algorithms follow the suggestions in the reference papers. More specifically, for the dynamic LLSO we set $\psi = 0.4$ and the set of candidate level numbers $S = \{4, 6, 8, 10, 20, 50\}$, as in Yang et al. (2018). For the PSO variant by Corazza et al. (2021), we consider $\omega_{min} = 0.4$, $\omega_{max} = 0.9$, $c_{1,min} = c_{2,min} = 0.5$, and $c_{1,max} = c_{2,max} = 2.5$. To guarantee a fair comparison, we set for both solvers the maximum number of generations $MAX_{GEN} = 1000$, and the swarm size NP equal to 300 for the Pacific data set and 500 for the World one. To obtain more robust results, we run 30 times each test instance. The analysis have been implemented in MATLAB 2023a and carried out on a 3.3 GHz Intel Core i9-7900X workstation with 16 GB of RAM. To prove the efficiency of the proposed algorithm, comparisons are made in terms of run time, capability to identify optimal solutions which satisfy the constraints, and accuracy in solving the optimization problems.

Due to the negligible impact of the turnover rate levels on the results, in the following we show only the findings related to $TR = 20\%$. Tables 2.2 and 2.3 display the average computational time on 30 runs for the dynamic LLSO and the PSO. For both methods, we can observe that the results remain relatively stable when transitioning from one date to another, regardless of the cardinality threshold. However, the outcomes depend on the employed objective function and the number of involved decision variables. In our analysis the PSO exhibit a slightly lower average computational time. This difference

can be attributed to the proposed hybrid constraint-handling technique.

Table 2.2: Average computational time in seconds on 30 runs for the two compared algorithms, for the Pacific data set and the three ex-post dates, with $TR = 20\%$ and increasing values of $K_{\%}$.

Pacific data set - LLSO									
	Date 1			**Date 2**			**Date 3**		
$K_{\%}$	SR	ASSR	AKSR	SR	ASSR	AKSR	SR	ASSR	AKSR
15%	97.98	162.23	233.07	103.16	164.91	237.26	94.77	148.86	214.20
30%	100.62	165.15	236.22	103.96	168.93	239.78	91.26	151.94	217.56
50%	100.83	165.83	235.90	104.14	168.95	240.60	91.52	152.00	217.91
Pacific data set - PSO									
	Date 1			**Date 2**			**Date 3**		
$K_{\%}$	SR	ASSR	AKSR	SR	ASSR	AKSR	SR	ASSR	AKSR
15%	80.51	125.99	193.27	82.882	127.37	195.91	79.041	119.62	181.70
30%	80.54	126.89	195.39	81.827	128.40	197.57	78.70	119.83	182.73
50%	80.55	127.40	195.57	81.604	127.90	196.56	79.08	120.50	182.88

Tables 2.4 and 2.5 show the average percentage of feasible solutions provided by the two solvers over the 30 runs at the final generation. Analyzing the results of the Pacific data set, it is evident that our LLSO is able to identify feasible solutions in almost all the cases. Conversely, the PSO algorithm struggles to properly handle constraints in the first test date, performing accurately in the other ones. Furthermore, the PSO is heavily influenced by the complexity of the objective function and the increase of portfolio size. As the number of decision variables increases, moving to the World case, the difference between the algorithms becomes even more pronounced. Specifically, for the cardinality threshold of 50%, the dynamic LLSO consistently finds optimal solutions, while the PSO fails to identify feasible portfolios in any of the 30 runs.

Moreover, we validate the capabilities of the proposed LLSO over the PSO in solving our test problems through a non-parametric statistical test. In particular, we focus on the average of the best values of the ℓ_1-exact penalty function (2.4.18), and we conduct a Wilcoxon signed-rank test to determine if there is a significant difference in the distributions of the values obtained by the two algorithms (Derrac et al. (2011)). The results are given in Table 2.6, where R^+ is the sum of ranks for the problems in which LLSO outperformed PSO, and R^- denotes the sum of the ranks for the opposite. Based on the p-values reported on the last column of this table, we conclude that our dynamic LLSO outperforms its competitor at the 5% significance level in all the case studies.

Table 2.3: Average computational time in seconds on 30 runs for the two compared algorithms, for the World data set and the three ex-post dates, with $TR = 20\%$ and increasing values of $K_{\%}$.

World data set - LLSO									
	Date 1			**Date 2**			**Date 3**		
$K_{\%}$	SR	ASSR	AKSR	SR	ASSR	AKSR	SR	ASSR	AKSR
15%	307.78	357.05	418.34	305.47	361.52	421.76	284.03	334.24	387.68
30%	307.32	363.84	424.71	307.87	365.03	426.48	285.68	337.22	391.22
50%	307.52	365.38	424.22	308.85	365.99	426.91	286.77	337.59	392.12
World data set - PSO									
	Date 1			**Date 2**			**Date 3**		
$K_{\%}$	SR	ASSR	AKSR	SR	ASSR	AKSR	SR	ASSR	AKSR
15%	281.87	314.34	357.79	278.45	305.44	345.67	278.90	310.53	348.13
30%	281.39	315.74	358.73	273.24	304.48	345.08	278.90	310.73	355.27
50%	282.65	314.61	356.33	273.07	304.66	345.48	277.97	309.80	347.59

2.5.c Long-run sensitivity analysis

Ex-post performance metrics

In this subsection, we present the ex-post performance measures that we will use to assess the profitability of the proposed investment strategies. Let $r_{p,t}^{out}$ be the ex-post portfolio rate of return at the month of the out-of-sample window. We compute the net wealth at time t as

$$W_t = (W_{t-1} - c_t)\left(1 + r_{p,t}^{out}\right) \tag{2.5.21}$$

where c_t represents the transaction costs associated to the rebalancing at time t. Given the optimal portfolio at time t, $\boldsymbol{x}_t$, and the re-normalized portfolio at time $t-1$, $\widetilde{\boldsymbol{x}}_{t-1}$, we define the portfolio cost as the sum of the trading costs of each constituent, that is

$$c_t = \sum_{i=1}^{n} c(ts_{i,t})$$

where $ts_{i,t} = W_{t-1}|x_{t,i} - \widetilde{x}_{t-1,i}|$ is the so-called trade segment for asset i at time t, with $i = 1, \ldots, n$ and $t = 1, \ldots, 126$, and $c(\cdot)$ is given by

$$c(ts_{i,t}) = \begin{cases} 0 & ts_{i,t} = 0 \\ 40 & 0 < ts_{i,t} < 8\,000\$ \\ 0.05 \times ts_{i,t} & 8\,000\$ \leq ts_{i,t} < 50\,000\$ \\ 0.04 \times ts_{i,t} & 50\,000\$ \leq ts_{i,t} < 100\,000\$ \\ 0.025 \times ts_{i,t} & 100\,000\$ \leq ts_{i,t} < 200\,000\$ \\ 400 & ts_{i,t} \geq 200\,000\$\,. \end{cases} \tag{2.5.22}$$

Table 2.4: Percentage of optimal solutions that satisfy all the constraints, over the 30 runs for the two compared algorithms for the Pacific data set and for the three ex-post dates, with $TR = 20\%$ and increasing values of $K_{\%}$.

Percentage of feasible solutions found (%) - Pacific data set - LLSO									
	Date 1			**Date 2**			**Date 3**		
$K_{\%}$	SR	ASSR	AKSR	SR	ASSR	AKSR	SR	ASSR	AKSR
15%	100	100	100	100	98.89	98.89	98.89	97.778	100
30%	100	100	100	100	98.89	100	98.89	100	98.89
50%	100	100	100	100	100	100	100	97.78	98.89
Percentage of feasible solutions found (%) - Pacific data set - PSO									
	Date 1			**Date 2**			**Date 3**		
$K_{\%}$	SR	ASSR	AKSR	SR	ASSR	AKSR	SR	ASSR	AKSR
15%	33.33	33.33	0	100	100	100	100	100	100
30%	33.33	33.33	0	98.89	96.67	96.67	96.67	93.33	86.67
50%	0	0	0	96.67	96.67	95.56	61.11	64.44	55.56

According to (2.5.22), we divide the traded monetary amount into non-overlapping intervals and apply a different cost percentage on the interval in which the traded capital lies. This transaction costs structure reflects the main configurations proposed nowadays by financial brokers (Beraldi et al. (2021)).

Next, we evaluate the attractiveness of the proposed asset allocation strategies through the so-called compound annual growth rate (shortly $CAGR$), which is calculated as

$$CAGR = \left(\frac{W_{T_{out}}}{W_0}\right)^{\frac{12}{T_{out}}} - 1 \tag{2.5.23}$$

where T_{out} denotes the length of the out-of-sample window, while W_0 and $W_{T_{out}}$ represent the initial wealth and the wealth at the end of the investment period, respectively. We also compute the monthly ex-post average rate of return and the ex-post standard deviation

$$\mu^{out} = \frac{1}{T_{out}} \sum_{t=1}^{T_{out}} r_{p,t}^{out} \tag{2.5.24}$$

$$\sigma^{out} = \sqrt{\frac{1}{T_{out} - 1} \sum_{t=1}^{T_{out}} (r_{p,t}^{out} - \mu^{out})^2}\,. \tag{2.5.25}$$

Furthermore, to analyze more precisely the distribution of the ex-post returns, we

Table 2.5: Percentage of optimal solutions that satisfy all the constraints, over the 30 runs for the two compared algorithms for the World data set and for the three ex-post dates, with $TR = 20\%$ and increasing values of $K_{\%}$.

Percentage of feasible solutions found (%) - World data set - LLSO									
	Date 1			**Date 2**			**Date 3**		
$K_{\%}$	SR	ASSR	AKSR	SR	ASSR	AKSR	SR	ASSR	AKSR
15%	98.89	100	98.89	98.89	95.56	98.89	97.78	100	98.89
30%	100	100	100	97.78	96.67	93.33	98.89	96.67	96.67
50%	100	100	100	100	100	96.67	100	100	98.89
Percentage of feasible solutions found (%) - World data set - PSO									
	Date 1			**Date 2**			**Date 3**		
$K_{\%}$	SR	ASSR	AKSR	SR	ASSR	AKSR	SR	ASSR	AKSR
15%	64.44	33.33	32.22	66.67	33.33	62.22	65.56	61.11	56.67
30%	33.33	33.33	0	33.33	33.33	32.22	0	0	0
50%	0	0	0	0	0	0	0	0	0

calculate the ex-post skewness and kurtosis

$$S_3^{out} = \frac{1}{T_{out}-1}\sum_{t=1}^{T_{out}}\left(\frac{r_{p,t}^{out}-\mu^{out}}{\sigma^{out}}\right)^3 \tag{2.5.26}$$

$$K_4^{out} = \frac{1}{T_{out}-1}\sum_{t=1}^{T_{out}}\left(\frac{r_{p,t}^{out}-\mu^{out}}{\sigma^{out}}\right)^4 . \tag{2.5.27}$$

To evaluate the capability of a strategy to avoid high losses, we introduce the drawdown measure, which is defined as follows (Chekhlov et al. (2005))

$$DD_t = \min\left\{0, \frac{W_t - W_{peak}}{W_{peak}}\right\} \quad t = 1, \ldots, T_{out} \tag{2.5.28}$$

where W_{peak} is the maximum amount of wealth reached by the strategy until time t. Particularly, we focus on the mean, standard deviation, and maximum value of the drawdown measure over time.

Finally, we measure the effect of the costs on the available capital in the out-of-sample period as in Kaucic et al. (2020) by

$$\Lambda_{\%} = \frac{1}{T_{out}}\sum_{t=1}^{T_{out}}\frac{c_t}{W_{t-1}}\cdot 100\,. \tag{2.5.29}$$

Table 2.6: Wilcoxon signed-rank test results for the two data sets at different dates. The null hypothesis assumes that the distributions of the best values of the dynamic LLSO and the PSO by Corazza et al. (2021) are equal, meaning there is no significant difference between the solvers' performance. The alternative hypothesis suggests that the dynamic LLSO consistently outperforms the compared algorithm.

Data set	Ex-post date	R^+	R^-	p-value
Pacific	Date 1	378	0	$2.97 \cdot 10^{-6}$
Pacific	Date 2	378	0	$2.97 \cdot 10^{-6}$
Pacific	Date 3	273	105	0.0224
World	Date 1	357	21	$2.85 \cdot 10^{-5}$
World	Date 2	375	3	$4.16 \cdot 10^{-6}$
World	Date 3	376	2	$3.72 \cdot 10^{-6}$

Results for the two data sets

In this subsection we present the results of the ex-post analysis. We start by highlighting the relation between portfolio size, allowed trades in the rebalancing phases, and the ex-post performance. Tables 2.7 and 2.8 report the findings for the three proposed portfolio models employing the Pacific and the World data sets, respectively. In the last row, we also include the results for the associated market value-weighted benchmark. Overall, we can observe that the impact of costs, expressed as percentage of the total wealth, does not exceed 0.057% for the Pacific data set, whereas it attains 0.15% in the World case study. Therefore, the influence of transaction costs seems relatively marginal for the smaller data set and becomes more pronounced for the larger data set. Furthermore, it is worth noting that increasing the cardinality and turnover thresholds leads to an expected rise in costs. Instead, if we fix TR and vary $K_{\%}$, we can notice the quasi-linear growth of costs. On the contrary, with $K_{\%}$ fixed, the cost increase is not proportional to the twofold increase of the turnover levels. This is coherent with the transaction cost structure defined in (2.5.22).

Regarding the Pacific data set, the ASSR and AKSR-based investment plans show similar results in the ex-post analysis for all the choices of $K_{\%}$ and TR, with the AKSR-based strategy with $K_{\%} = 15\%$ and $TR = 10\%$ which outperforms the other tested alternatives.

For the World data set, maximizing the ASSR with $K_{\%} = 15\%$ and $TR = 10\%$ gives promising results in terms of ex-post returns, CAGR, and drawdown measures. However, this investment strategy lacks of robustness, as it underperforms significantly in all other parameter configurations compared to both the AKSR and SR strategies. As in the Pacific case, the AKSR-based model appears to be the most resilient and viable choice for the long-term investments.

Summing up, the optimal combination of the parameters is $K_{\%} = 15\%$ as cardinality

threshold and $TR = 10\%$ as turnover rate. The comparison with the buy-and-hold strategy with the auxiliary benchmarks points out that the proposed models, in both the investable universes, outperform the benchmark in terms of profits, generating a higher ex-post net wealth on the long-run. In addition, the multi-moment strategies show less volatility and lower drawdown than the benchmark for the Pacific data set. However, for the World data set, the ASSR and AKSR-based investments present higher standard deviations and drawdowns.

Table 2.7: Results of the long-run sensitivity analysis for the three Sharpe ratio-based optimization strategies with the Pacific data set. In the first two columns we report the value of the fraction $K_{\%}$ of assets making up the portfolio and the value of the turnover rate, TR, respectively. The other columns show the results of the ex-post metrics presented in Section 2.5.c.

Pacific data set										
$\mathbf{K_{\%}}$	**TR**	$\mu^{\mathbf{out}}$	$\sigma^{\mathbf{out}}$	$\mathbf{S_3^{out}}$	$\mathbf{K_4^{out}}$	**mean DD$_t$**	**std DD$_t$**	**max DD$_t$**	**CAGR (%)**	$\mathbf{\Lambda_{\%}}$
Sharpe ratio										
15%	10%	0.0155	0.0524	-0.3363	3.8968	-0.0401	0.0604	-0.3221	18.1213	0.0168
15%	20%	0.0130	0.0510	-0.4092	3.9483	-0.0464	0.0630	-0.3312	14.6364	0.0246
15%	40%	0.0129	0.0527	-0.2663	4.2228	-0.0469	0.0660	-0.3500	14.2350	0.0364
30%	10%	0.0144	0.0511	-0.3447	4.0604	-0.0413	0.0615	-0.3402	16.5397	0.0258
30%	20%	0.0126	0.0495	-0.6337	3.8347	-0.0456	0.0688	-0.3646	14.0217	0.0346
30%	40%	0.0123	0.0509	-0.3258	3.6120	-0.0520	0.0697	-0.3407	13.4481	0.0482
50%	10%	0.0141	0.0517	-0.3131	4.3767	-0.0401	0.0632	-0.3462	15.9818	0.0362
50%	20%	0.0132	0.0508	-0.3919	3.6058	-0.0470	0.0655	-0.3339	14.5913	0.0448
50%	40%	0.0140	0.0521	-0.3356	3.8593	-0.0498	0.0705	-0.3556	15.4638	0.0576
Adjusted for skewness Sharpe ratio										
15%	10%	0.0151	0.0530	-0.2048	4.1423	-0.0413	0.0605	-0.3186	17.4500	0.0173
15%	20%	0.0147	0.0526	-0.3414	3.6750	-0.0438	0.0602	-0.3200	16.8631	0.0227
15%	40%	0.0135	0.0526	-0.2439	4.2715	-0.0451	0.0660	-0.3476	15.0479	0.0350
30%	10%	0.0151	0.0523	-0.2658	3.8936	-0.0398	0.0596	-0.3309	17.5081	0.0245
30%	20%	0.0136	0.0505	-0.5488	3.6586	-0.0450	0.0648	-0.3387	15.4300	0.0324
30%	40%	0.0124	0.0513	-0.4045	3.6817	-0.0508	0.0697	-0.3527	13.4513	0.0466
50%	10%	0.0150	0.0532	-0.1995	4.3461	-0.0397	0.0618	-0.3391	17.1096	0.0351
50%	20%	0.0137	0.0522	-0.3742	3.7466	-0.0469	0.0673	-0.3503	15.2951	0.0422
50%	40%	0.0140	0.0532	-0.2455	4.0119	-0.0515	0.0709	-0.3538	15.3767	0.0569
Adjusted for skewness and kurtosis Sharpe ratio										
15%	10%	0.0159	0.0531	-0.1289	4.2466	-0.0391	0.0581	-0.3081	18.6239	0.0170
15%	20%	0.0143	0.0526	-0.3342	3.6512	-0.0437	0.0608	-0.3180	16.3558	0.0231
15%	40%	0.0141	0.0526	-0.3061	4.1253	-0.0440	0.0653	-0.3433	15.9187	0.0343
30%	10%	0.0153	0.0525	-0.2646	4.0030	-0.0406	0.0607	-0.3351	17.6604	0.0248
30%	20%	0.0139	0.0504	-0.5317	3.6388	-0.0442	0.0645	-0.3407	15.7871	0.0324
30%	40%	0.0131	0.0520	-0.4019	3.6642	-0.0524	0.0721	-0.3573	14.4739	0.0449
50%	10%	0.0151	0.0526	-0.2434	4.2324	-0.0403	0.0626	-0.3423	17.2423	0.0346
50%	20%	0.0140	0.0517	-0.4260	3.7441	-0.0467	0.0677	-0.3502	15.6687	0.0416
50%	40%	0.0148	0.0535	-0.2355	3.9622	-0.0508	0.0705	-0.3521	16.5384	0.0544
Benchmark										
–	–	0.0074	0.0541	-0.1912	8.7070	-0.0794	0.0817	-0.3261	7.3157	–

Table 2.8: Results of the long-run sensitivity analysis for the three Sharpe ratio-based optimization strategies with the World data set. In the first two columns we report the value of the fraction $K_{\%}$ of assets making up the portfolio and the value of the turnover rate, TR, respectively. The other columns show the results of the ex-post metrics presented in Section 2.5.c.

World data set										
$\mathbf{K_{\%}}$	**TR**	$\mu^{\mathbf{out}}$	$\sigma^{\mathbf{out}}$	$\mathbf{S_3^{out}}$	$\mathbf{K_4^{out}}$	**mean DD$_t$**	**std DD$_t$**	**max DD$_t$**	**CAGR (%)**	$\mathbf{\Lambda_{\%}}$
Sharpe ratio										
15%	10%	0.0161	0.0466	-0.2773	4.9703	-0.0291	0.0551	-0.2912	18.9178	0.0425
15%	20%	0.0140	0.0430	-0.7310	4.6233	-0.0285	0.0511	-0.2584	16.1381	0.0553
15%	40%	0.0148	0.0438	-0.3794	4.9213	-0.0291	0.0516	-0.2557	16.8147	0.0751
30%	10%	0.0150	0.0451	-0.4116	5.0116	-0.0285	0.0564	-0.3093	17.0248	0.0762
30%	20%	0.0142	0.0442	-0.5375	5.0115	-0.0280	0.0529	-0.2722	15.8571	0.0870
30%	40%	0.0153	0.0446	-0.2956	4.4928	-0.0264	0.0474	-0.2501	17.2135	0.1003
50%	10%	0.0156	0.0463	-0.2359	5.5967	-0.0272	0.0575	-0.3195	17.1385	0.1194
50%	20%	0.0155	0.0453	-0.1842	5.2368	-0.0258	0.0505	-0.2742	17.0534	0.1247
50%	40%	0.0145	0.0449	-0.1576	5.0796	-0.0281	0.0505	-0.2673	15.4072	0.1454
Adjusted for skewness Sharpe ratio										
15%	10%	0.0177	0.0497	-0.1205	4.7322	-0.0290	0.0534	-0.2791	21.0693	0.0402
15%	20%	0.0138	0.0439	-0.4892	4.3779	-0.0306	0.0509	-0.2521	15.7652	0.0556
15%	40%	0.0149	0.0460	-0.1817	4.9571	-0.0302	0.0524	-0.2579	16.9092	0.0712
30%	10%	0.0156	0.0467	-0.4320	4.8426	-0.0297	0.0575	-0.3079	17.7964	0.0746
30%	20%	0.0142	0.0452	-0.4714	4.7642	-0.0294	0.0539	-0.2745	15.8234	0.0828
30%	40%	0.0152	0.0453	-0.2490	4.5258	-0.0264	0.0468	-0.2472	16.9971	0.0959
50%	10%	0.0154	0.0469	-0.1386	5.3080	-0.0279	0.0566	-0.3129	16.7302	0.1260
50%	20%	0.0152	0.0467	-0.0542	5.3984	-0.0267	0.0524	-0.2905	16.5561	0.1274
50%	40%	0.0138	0.0464	-0.0221	5.2357	-0.0290	0.0513	-0.2788	14.3094	0.1490
Adjusted for skewness and kurtosis Sharpe ratio										
15%	10%	0.0167	0.0497	-0.2142	4.8412	-0.0305	0.0575	-0.3090	19.6955	0.0406
15%	20%	0.0146	0.0455	-0.5668	4.9556	-0.0318	0.0556	-0.2766	16.8653	0.0525
15%	40%	0.0151	0.0472	-0.1642	5.1791	-0.0316	0.0539	-0.2549	17.1044	0.0726
30%	10%	0.0159	0.0465	-0.3997	4.8263	-0.0297	0.0570	-0.3009	18.2756	0.0750
30%	20%	0.0154	0.0471	-0.2610	5.0832	-0.0290	0.0537	-0.2740	17.4193	0.0804
30%	40%	0.0160	0.0474	-0.0953	4.8810	-0.0269	0.0490	-0.2624	18.0243	0.0957
50%	10%	0.0151	0.0462	-0.2559	5.2358	-0.0276	0.0570	-0.3136	16.5143	0.1228
50%	20%	0.0154	0.0473	-0.0451	5.1902	-0.0276	0.0549	-0.3064	16.7104	0.1298
50%	40%	0.0149	0.0475	-0.0182	5.3073	-0.0287	0.0536	-0.2936	15.8154	0.1449
Benchmark										
–	–	0.0108	0.0381	-0.6071	3.8057	-0.0273	-0.2626	0.0541	12.8355	–

Pre- and post-COVID analysis

Figure 2.1 reports the evolution of the net wealth for the best three Sharpe ratio-based portfolio strategies in each data set. The plot confirms the insights previously evidenced. All these investment strategies perform better than the auxiliary benchmarks. In the Pacific case study, we can observe that the SR investment plan is the one with the best performance until February 2020, date of the pandemic outbreak, at which it realizes a very large drawdown. After this period, only the AKSR strategy seems to be able to gain a great advantage of the market fluctuations in the last years. Similarly, in the World data set, the three models have a comparable net wealth evolution until February 2020. Then, the two multi-moment strategies perform better than the SR one,

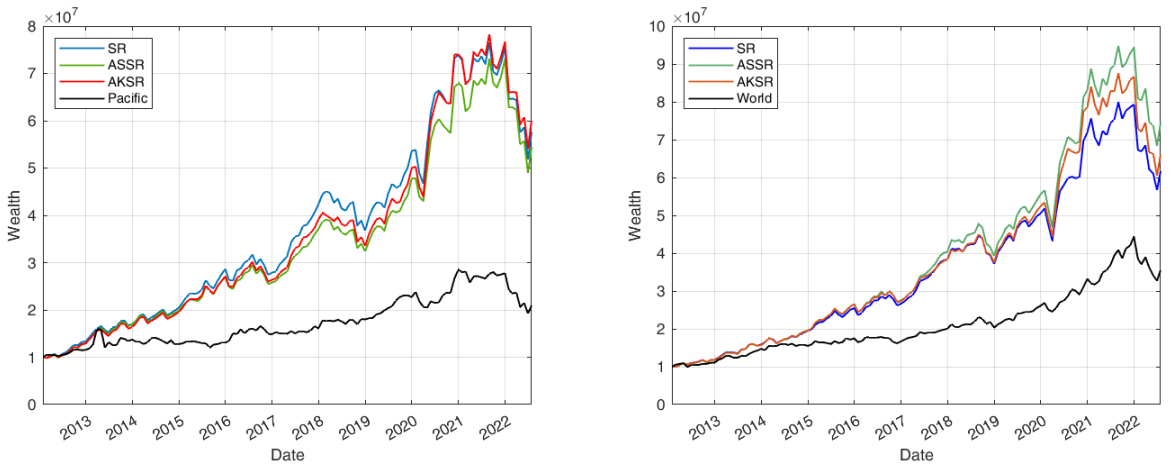

Figure 2.1: Evolution of the net wealth for the best three Sharpe ratio-based portfolio strategies ($K_{\%} = 15\%$ and $TR = 10\%$) in comparison to the market value weighted benchmark using the assets in the two data sets.

with the ASSR being the most profitable.

In Table 2.9 we highlight the different behaviours of the auxiliary markets and the proposed strategies before the above quoted date, namely pre-COVID period, and after that date, namely post-COVID period. For both data sets, the distributions of the rates of return in the pre-COVID period have a negative skewness while post-COVID rates of return show symmetric distributions and positive skewness. Conversely, regarding the benchmarks, the distribution of the rates of return presents a lower mean with negative skewness in both market phases. We summarize the results in Figures 2.2 and 2.3 in which, in both data sets, we can detect a positive value for the ex-post mean of the rates of return with a reduced dispersion in the pre-COVID period, while in the post-COVID epoch we can discover a lower ex-post mean of the rates of return with more dispersion and fat-tails.

These findings substantiate that extending the Sharpe ratio model by incorporating higher-order moments can yield financial performance benefits, particularly during periods characterized by market instability.

2.6 Conclusions and future works

In this study, we have proposed a comparison of three long-run investment strategies based on Sharpe ratio type performance measures on large-scale global market indices. In particular, we have considered the standard Sharpe ratio and two extensions which involve higher-order moments of the returns distribution. Furthermore, we have included four real-world constraints, namely the cardinality, buy-in threshold, budget and turnover constraints, in order to provide a complete control on the portfolio composition. To solve this family of optimization problems, we have developed a novel

Table 2.9: Out-of-sample statistics for the distribution of the ex-post rates of return of the three compared strategies with $K_{\%} = 15\%$ and $TR = 10\%$ and the benchmark for the two data sets. Results refer to the period before and after 28 February 2020, the watershed date for the COVID-19 pandemic outburst.

	Pre-COVID				**Post-COVID**			
Strategy	$\mu^{\mathbf{out}}$	$\sigma^{\mathbf{out}}$	$\mathbf{S_3^{out}}$	$\mathbf{K_4^{out}}$	$\mu^{\mathbf{out}}$	$\sigma^{\mathbf{out}}$	$\mathbf{S_3^{out}}$	$\mathbf{K_4^{out}}$
Pacific data set								
SR	0.0177	0.0439	-0.8178	3.2845	0.0083	0.0749	0.2446	3.3132
ASSR	0.0165	0.0437	-0.7641	3.1747	0.0103	0.0773	0.3176	3.4553
AKSR	0.0170	0.0435	-0.7728	3.2698	0.0122	0.0780	0.3933	3.4518
Benchmark	0.0095	0.0526	-0.0977	11.7410	0.0004	0.0602	-0.3545	3.5086
World data set								
SR	0.0174	0.0367	-0.7443	3.5505	0.0116	0.0711	0.1419	3.6941
ASSR	0.0183	0.0388	-0.6842	3.4463	0.0155	0.0766	0.2310	3.3632
AKSR	0.0177	0.0380	-0.6780	3.3421	0.0134	0.0780	0.1222	3.2992
Benchmark	0.0100	0.0324	-0.4454	3.6395	0.0135	0.0542	-0.7804	3.0798

swarm optimization algorithm equipped with an ad-hoc constraint handling technique combining the global convergence properties of the ℓ_1-exact penalty functions with a repair operator.

The empirical findings are obtained on two large-scale data sets of the Pacific and World areas, which include several hundreds of stocks, covering the last 14 years. We have performed a sensitivity analysis for the portfolio size and for the limit of the trades magnitude, in order to identify the best combination of these parameters in terms of ex-post performance and management cost. Results show that portfolios with a reduced number of constituents (15% of the investment pool) and with a two-sided turnover up to 10% provide, on both data sets, better profits which are stable over time.

A more detailed analysis reveals that the inclusion of higher-order moments in the performance measures produces superior results in terms of net wealth with respect to the benchmark and to the portfolio optimized through the standard Sharpe ratio. This is more evident after the pandemic outbreak of 2020, where more market fluctuations are present.

In the future works, on the one hand, we plan to extend the above experimental analysis to other markets and, on the other hand, we are interested in the possibility to introduce similar multi-moment performance measures in a passive investment framework.

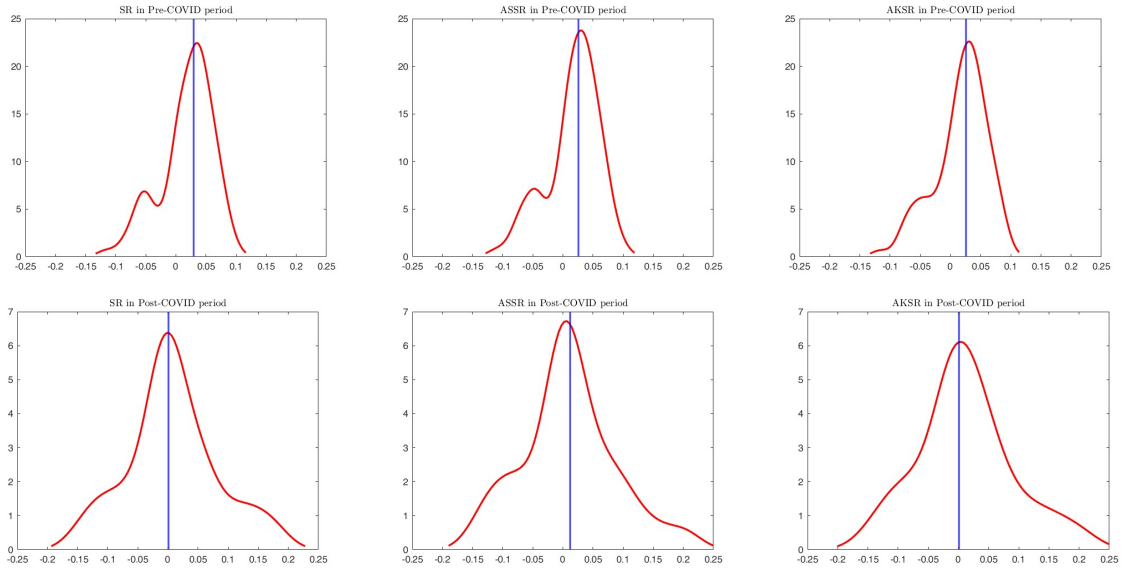

Figure 2.2: Density functions of the ex-post rates of return of the three strategies proposed in Section 2.3.a for the Pacific data set in the pre-COVID period (in the first row) and post-COVID epoch (in the second row). The vertical blue line represents the median of the distribution.

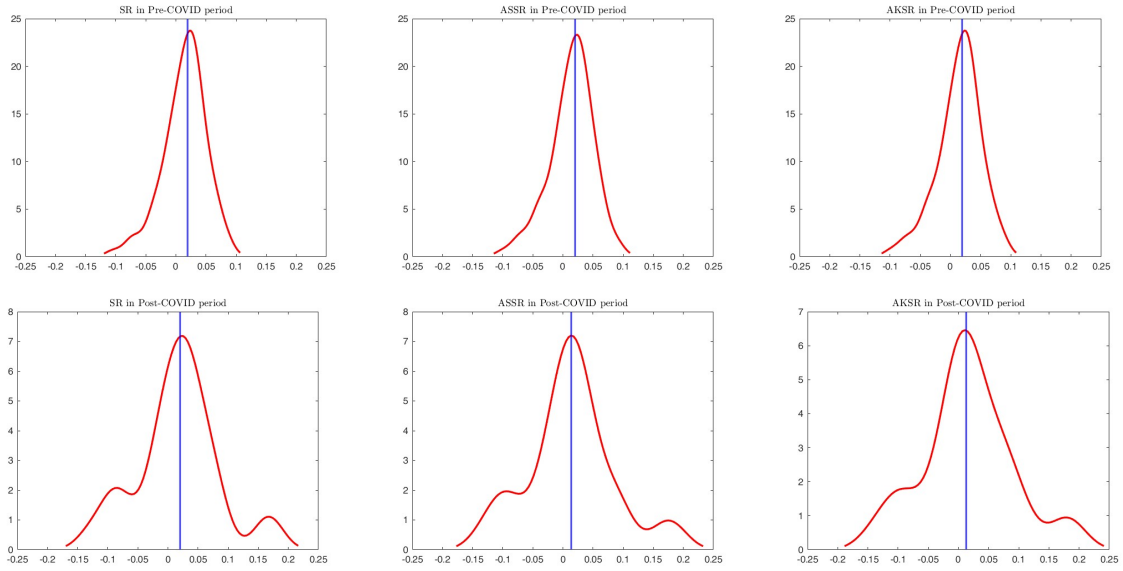

Figure 2.3: Density functions of the ex-post rates of return of the three strategies proposed in Section 2.3.a for the World data set in the pre-COVID period (in the first row) and post-COVID epoch (in the second row). The vertical blue line represents the median of the distribution.

Bibliography

C. Amédée-Manesme, F. Barthélémy: *Proper use of the modified Sharpe ratios in performance measurement: rearranging the Cornish Fisher expansion*. Ann. Oper. Res., **313** (2022), 691-712, https://doi.org/10.1007/s10479-020-03858-4.

F. D. Arditti: *Risk and the required return on equity*. J. Finance, **22** (1967), n. 1, 19–36, https://doi.org/10.1111/j.1540-6261.1967.tb01651.x.

P. Beraldi, A. Violi, M. Ferrara, C. Ciancio, B. A. Pansera: *Dealing with complex transaction costs in portfolio management*. Ann. Oper. Res., **299** (2021), n. 1, 7-22, https://doi.org/10.1007/s10479-019-03210-5.

P. Chaigneau, I. Eeckhoudt: *Downside risk-neutral probabilities*. Econ. Theory Bull., **8** (2020), 65-77, https://doi.org/10.1007/s40505-019-00165-5.

A. Chekhlov, S. Uryasev, M. Zabarankin: *Drawdown measure in portfolio optimization*. Int. J. Theor. Appl. Finance, **8** (2005), n. 1, 13-58, https://doi.org/10.1142/S0219024905002767.

R. Cheng, Y. Jin: *A competitive swarm optimizer for large scale optimization*. IEEE Trans. Cybern., **45** (2015), n. 2, 191-204, https://doi.org/10.1109/TCYB.2014.2322602.

R. Cheng, Y. Jin: *A social learning particle swarm optimization algorithm for scalable optimization*. Inf. Sciences, **291** (2015), 43-60, https://doi.org/10.1016/j.ins.2014.08.039.

P. Cheridito, E. Kromer: *Reward-risk ratios*. J. Invest. Strateg., **3** (2013), n. 1, 3-18, https://doi.org/10.21314/JOIS.2013.022.

A. Colasante, L. Riccetti: *Risk aversion, prudence and temperance: It is a matter of gap between moments*. J. Behav. Exp. Finance, **25** (2020), n. 100262, https://doi.org/10.1016/j.jbef.2019.100262.

A. Colasante, L. Riccetti: *Financial and non-financial risk attitudes: what does it matter?* J. Behav. Exp. Finance, **30** (2021), n. 100494, https://doi.org/10.1016/j.jbef.2021.100494.

M. Corazza, G. di Tollo, G. Fasano, R. Pesenti: *A novel hybrid PSO-based metaheuristic for costly portfolio selection problems*. Ann. Oper. Res., **304** (2021), 109-137, https://doi.org/10.1007/s10479-021-04075-3.

T. Cura: *Particle swarm optimization approach to portfolio optimization*. Nonlinear Anal. Real World Appl., **10** (2009), n. 4, 2396-2406, https://doi.org/10.1016/j.nonrwa.2008.04.023.

R. Deguest, L. Martellini, A. Meucci: *Risk parity and beyond - from asset allocation to risk allocation decisions*. J. Portf. Manag., **48** (2022), n. 4, 108-135, https://doi.org/10.3905/jpm.2022.1.340.

J. Derrac, S. Garcia, D. Molina, F. Herrera: *A practical tutorial on the use of nonparametric statistical tests as a methodology for comparing evolutionary and swarm intelligence algorithms*. Swarm Evol. Comput., **1** (2011), n. 1, 3-18, https://doi.org/10.1016/j.swevo.2011.02.002.

J. Gao, X. Gao, X. Liu, Z. Wang: *Detecting prudence and temperance in risk exposure: the hybrid variance framework*. J. Risk, **24** (2022), n. 5, 75-88, https://doi.org/10.21314/JOR.2022.034.

M. Gilli, E. Schumann: *Large-scale portfolio optimization with heuristics*. In: "Advanced Statistical Methods for the Analysis of Large Data-Sets". Springer, Berlin, Heidelberg (2012), 181-192, https://doi.org/10.1007/978-3-642-21037-2_17.

G. Guastaroba, R. Mansini, M. G. Speranza: *On the effectiveness of scenario generation techniques in single-period portfolio optimization*. Eur. J. Oper. Res., vol. **192** (2009), n. 2, 500-511, https://doi.org/10.1016/j.ejor.2007.09.042.

J. B. Guerard: *"Handbook of portfolio construction. Contemporary applications of Markowitz techniques"*. Springer, New York (2010).

A. Hitaj, G. Zambruno: *Are Smart Beta strategies suitable for hedge fund portfolios?* Rev. Financial Econ., **29** (2016), 37-51, https://dx.doi.org/10.1016/j.rfe.2016.03.001.

E. Jurczenko, B. Maillet: *"Multi-moment asset allocation and pricing moments"*. John Wiley & Sons, New Jersey, (2006).

M. Kaucic, F. Barbini, F. J. Camerota: *Polynomial goal programming and particle swarm optimization for enhanced indexation*. Soft Comp., **24** (2020), 8535-8551, https://doi.org/10.1007/s00500-019-04378-5.

M. Kaucic, F. Piccotto: *A level-based learning swarm optimizer with a hybrid constraint-handling technique for large-scale portfolio selection problems*. 2022 IEEE Congr. Ev. Comput., (2022), 1-8, https://doi.org/10.1109/CEC55065.2022.9870358.

M. Kaucic, F. Piccotto, G. Sbaiz, G. Valentinuz: *A hybrid level-based learning swarm algorithm with mutation operator for solving large-scale cardinality-constrained portfolio optimization problems*. Inf. Sci., **634** (2022), 321-339, https://doi.org/10.1016/j.ins.2023.03.115.

W. C. Kim, F. J. Fabozzi, P. Cheridito, C. Fox: *Controlling portfolio skewness and kurtosis without directly optimizing third and fourth moments*. Econ. Lett., **122** (2014), n. 2, 154-158, 2014, https://doi.org/10.1016/j.econlet.2013.11.024.

M. S. Kimball: *Precautionary saving in the small and in the large*. Econometrica, **58** (1990), n. 1, 53-73, 1990, https://doi.org/10.2307/2938334.

H. Konno, H. Shirakawa, H. Yamazaki: *A mean-absolute deviation-skewness portfolio optimization model*. Ann. Oper. Res., **45** (1993), 205–220, https://doi.org/10.1007/BF02282050.

T. Y. Lai: *Portfolio selection with skewness: A multiple-objective approach*. Rev. Quant. Finan. Acc., **1** (1991), 293–305, https://doi.org/10.1007/BF02408382.

K. K. Lai, L. Yu, S. Wang: *Mean-variance-skewness-kurtosis-based portfolio optimization*. 2006 Int. Symp. on Comp. and Comput. Science, (2006), 292-297, https://doi.org/10.1109/IMSCCS.2006.239.

X. Li, X. Yao: *Cooperatively coevolving particle swarms for large scale optimization*. IEEE Trans. Evol. Comput., **16** (2012), n. 2, 210-224, https://doi.org/10.1109/TEVC.2011.2112662.

S. Liu, S. Y. Wang, W. Qiu: *Mean-variance-skewness model for portfolio selection with transaction costs*. Int. J. Syst. Sci., **34** (2003), n. 4, 255-262, http://dx.doi.org/10.1080/0020772031000158492.

H. Markowitz: *Portfolio selection*. J. Finance, **7** (1952), n. 1, 77-91, https://doi.org/10.2307/2975974.

S. S. Meghwani, M. Thakur: *Multi-criteria algorithms for portfolio optimization under practical constraints*. Swarm Evol. Comput., **37** (2017), 104-125, https://doi.org/10.1016/j.swevo.2017.06.005.

R. Moral-Escudero, R. Ruiz-Torrubiano, A. Suarez: *Selection of optimal investment portfolios with cardinality constraints*. IEEE Conf. Evol. Comput. (2006), 2382-2388, https://doi.org/10.1109/CEC.2006.1688603.

B. Z. Nagy, B. Benedek: *Higher co-moments and adjusted Sharpe ratios for cryptocurrencies*. Finance Res. Lett., **39** (2021), 1-7, https://doi.org/10.1016/j.frl.2020.101543.

E. T. Oldewage, A. P. Engelbrecht, C. W. Cleghorn: *Movement patterns of a particle swarm in high dimensional spaces*. Inf. Sci., **512** (2020), 1043-1062, https://doi.org/10.1016/j.ins.2019.09.057.

M. N. Omidvar, X. Li, X. Yao: *A review of population-based metaheuristics for large-scale black-box global optimization—Part I*. IEEE Trans. Evol. Comput., **26** (2022), n. 5, 802-822, https://doi.org/10.1109/TEVC.2021.3130838.

M. N. Omidvar, X. Li, X. Yao: *A review of population-based metaheuristics for large-scale black-box global optimization—Part II*. IEEE Trans. Evol. Comput., **26** (2022), n. 5, 823-843, https://doi.org/10.1109/TEVC.2021.3130835.

J. Pézier, A. White: *The relative merits of alternative investments in passive portfolios*. J. Altern. Invest., **10** (2008), n. 4, 37-49, https://doi.org/10.3905/jai.2008.705531.

D. Politis, J. P. Romano: *The stationary bootstrap*. J. Am. Stat. Assoc., **89** (1994), n. 428, 1303-1313, https://doi.org/10.1080/01621459.1994.10476870.

D. Politis, H. White: *Automatic block-length selection for the dependent bootstrap*. Econom. Rev., **23** (2004), n. 1, 53-70, https://doi.org/10.1081/ETC-120028836.

P. A. Samuelson: *The fundamental approximation theorem of portfolio analysis in terms of means, variances and higher moments*. Rev. Econ. Stud., **37** (1970), n. 4, 537-542, https://doi.org/10.2307/2296483.

R. C. Scott, P. A. Horvath: *On the direction of preference for moments of higher order than the variance*. J. Finance, **35** (1980), n. 4, 915-919, https://doi.org/10.2307/2327209.

W. F. Sharpe: *The Sharpe ratio*. J. Portf. Manag., **21** (1994), n. 1, 45-58, https://doi.org/10.3905/jpm.1994.409501.

W. Shen, J. Wang, S. Ma: *Doubly regularized portfolio with risk minimization*. Proc. AAAI Conf. Artif. Int., **28** (2014), n. 1, 1286-1292, https://doi.org/10.1609/aaai.v28i1.8906.

G.-W. Song, Q. Yang, X.-D. Gao, Y.-Y. Ma, Z.-Y. Lu, J. Zhang: *An adaptive level-based learning swarm optimizer for large-scale optimization*. IEEE Trans. Syst., Man, Cybern., (2021), 152-159, https://doi.org/10.1109/SMC52423.2021.9658644.

F. Van der Bergh, A. P. Engelbrecht: *A cooperative approach to particle swarm optimization*. IEEE Trans. Evol. Comp., **8** (2004), n. 3, 225-239, https://doi.org/10.1109/TEVC.2004.826069.

E. T. van Zyl, A. P. Engelbrecht: *A subspace-based method for PSO initialization*. IEEE Symposium Series on Computational Intelligence (2015), 226-233, https://doi.org/10.1109/SSCI.2015.42.

D. Wang, D. Tan, L. Liu: *Particle swarm optimization algorithm: an overview*. Soft Comput., **22** (2018), n. 2, 387-408, https://doi.org/10.1007/s00500-016-2474-6.

F. Wang, X. Wang, S. Sun: *A reinforcement learning level-based particle swarm optimization algorithm for large-scale optimization*. Inf. Sci., **602** (2022), 298-312, https://doi.org/10.1016/j.ins.2022.04.053.

Q. Yang, W.-N. Chen, J. Da Deng, Y. Li, T. Gu, J. Zhang: *A level-based learning swarm optimizer for large-scale optimization*. IEEE Trans. Evol. Comput., **22** (2018), n. 4, 578-594, https://doi.org/10.1109/TEVC.2017.2743016.

V. Zakamuline, S. Koekebakker: *Portfolio performance evaluation with generalized Sharpe ratios: beyond the mean and variance*. J. Bank. Financ., **33** (2009), n. 7, 1242-1254, https://doi.org/10.1016/j.jbankfin.2009.01.005.

H. Zhu, Y. Wang, K. Wang, Y. Chen: *Particle Swarm Optimization (PSO) for the constrained portfolio optimization problem*. Expert Syst. Appl., **38** (2011), n. 8, 10161-10169, https://doi.org/10.1016/j.eswa.2011.02.075.

Part II

ESG integration in portfolio models

The second part of the PhD thesis contains the last two chapters. The third, entitled *"The role of ESG ratings in investment portfolio choices"* has been published as a chapter in the book *"Enhancing Sustainability Through Non-Financial Reporting"* edited by Albertina Monteiro, Ana Pinto Borges, Elvira Vieira and published in October 2023. In this chapter, we introduce a tri-objective asset allocation model that considers the portfolio ESG score as a third function to maximize along with return and risk.

The fourth chapter, *"Optimal portfolio with sustainable attitudes under cumulative prospect theory"*, is a research paper that has been published in the Journal of Applied Finance and Banking in 2023. This paper is devoted to the presentation of several cumulative prospect theory (shortly, CPT)-based portfolio strategies. In particular, employing the concept of sustainable returns obtained from firms' ESG scores, we propose a novel bi-objective portfolio selection model that involves CPT-value functions for the financial and sustainable returns as objectives to maximize simultaneously.

Before presenting the two works, we provide a general overview of the concept of socially responsible investment.

Basics of sustainable investment

The issue of socially responsible investment (SRI), born in a relatively circumscribed context and based on ethical or religious beliefs, has seen growing importance over time, assuming different and more widely shared connotations. In particular, SRIs were selected primarily for non-financial reasons, as opposed to traditional financial theory, which assumes that investors seek the best combination of risk and return. For example, Renneboog et al. (2008) emphasise that investors derive non-financial utility from investing in SRI practices. To identify investments made into companies and funds that produce a positive social and environmental impact while guaranteeing financial returns, the Rockefeller Foundation, in 2007, formally used for the first time the term "Impact investing" (we refer the interested reader to the book of Bugg-Levine and Emerson(2011)). However, in a traditional view, a company that pursues goals other than maximizing shareholder value, such as benefiting stakeholders not directly connected to the company, risks reducing its value and creating costly conflicts of interest. However, over time, the environmental issue and the consequent effects of global warming started to have an increasing weight on the agendas of international political authorities. The Paris Agreement on climate change and, among others, the UN Principles for Responsible Investment (PRI), launched in 2006, and the UN 2030 Agenda for Sustainable Development are the necessary premises for a different way of tackling economic development, questioning the models which had hitherto been dominant. Thanks to the large amount of resources they can mobilise, the institutional investors and financial intermediates have been seen as possible "agents of change" towards a more sustainable path for both the planet and the economic system, with inclusive growth. The European Commission, with the Action Plan in 2018, has attributed spe-

cific concerns to financial intermediaries to drive flows toward sustainable investment, explicitly requiring portfolio managers to integrate factors relating to sustainability into their process. The need to urgently allocate capital deriving from the financial services sector (alongside those made available by public policies) is identified in the preamble of Regulation (EU) 2019/2088 on sustainability-related disclosure in the financial services sector (SFDR). Through this new regulation, the European Union (EU) tries to change behaviour patterns in the financial industry, discouraging greenwashing and promoting responsible and sustainable investments. Albeit with significant differences, in an international context which shows a growing interest in sustainability issues, the EU is particularly active in transforming principles into behaviour, with the introduction of directives and regulations, both towards the world of investors and businesses. The EU has introduced some rules and initiatives to improve corporate disclosure and align it with the EU's priority objective of sustainable transition (reinforced with the adoption of the "Next Generation EU"). In particular, to raise transparency and comparability of sustainability-related information for investors, the EU introduced the Non-Financial Reporting Directive, the already mentioned Sustainable Finance Disclosure Regulation, and the new Corporate Sustainability Reporting Directive. Further, some ad-hoc directives (the Markets in Financial Instruments Directive and the Insurance Distribution Directive) require asset management firms and insurance companies to offer suitable products to meet their customers' needs and desires and to inform them of the sustainability of the financial instruments. As part of impact investments, it is possible to identify green, social, and sustainability bonds. In particular, a green bond can be defined as "any type of bond instrument where the proceeds will be exclusively applied to finance or re-finance, in part or in full, new or/and existing eligible green projects"[2], whereas social and sustainability bonds follow socially responsible goals. The difference between the yield on a conventional bond and a green bond with similar characteristics is defined as 'greenium' (or green premium) (Agliardi and Agliardi (2019)). More precisely, investors are disposed to accept a lower yield on a green asset compared to the traditional one, mainly on the secondary markets, putting in evidence strong investor demand in pro-environmental projects. This conclusion seems to be confirmed by most of the literature analysed by MacAskill et al. (2021) on green bonds.

While Corporate Social Responsibility (CSR), SRI, and ESG concepts differ, they are often used interchangeably. SRIs arise from excluding certain types of investments (and issuers) from the construction of portfolios, particularly those that do not respond to ethical considerations or do not follow socially responsible behaviour. CSR can instead be defined as a general framework of sustainability and responsible cultural influence of a company towards its shareholders. ESG is a measurable sustainability assessment. The ESG rating purpose is to measure compliance with the objectives and principles

[2]ICMA, 2018. Green Bond Principles 2018: Voluntary Process Guidelines for Issuing Green Bonds. Retrieved from. https://www.icmagroup.org/assets/documents/Regulatory/Green-Bonds/Green-Bonds-Principles-June-2018-270520.pdf.

of sustainability broadly and - consequently - the greater or lesser consistency of an investment with the idea of impact investing. Liang and Renneboog (2020) link the concepts of CSR and ESG, referring corporate social responsibility to the incorporation of environmental, social and governance considerations into corporate management, financial decision-making and investor portfolio decisions.

Nowadays, attention to ESG factors is relevant worldwide: Global Sustainable Investment Alliance[3] recently reported that 36% of total professionally managed assets (around 35,3 trillion dollars) incorporate ESG factors. Many investors use ESG scores (usually defined as ratings) to understand a firm's sustainability and – consequently – guide the allocation of their investments. These scores can be considered a framework used to evaluate a company's performance on specific environmental, social, and governance issues and transform the non-financial disclosure required, particularly by the EU rules, in numerical terms. So, the importance of ESG ratings is evident; but the lack of commonly unified standards for ESG measurement has led to considerable differences in how ESG is measured and evaluated by different data vendors (Zumente and Lāce (2021)). Berg et al. (2022), using a dataset of ESG ratings from six providers, identified that the correlations between the ratings are 0.54 on average, ranging from 0.38 to 0.71. As one of the consequences caused by this divergence in valuations, we suggest that the prices of stocks and bonds are less likely to reflect ESG performance, as investors do not have unique information. Depending on the rating agencies' preferences, weights of the constituting factors, and rating methodology, these rating inconsistencies represent a problem for the past and a challenge for future research (Liang and Renneboog (2020)). A consequence of such a broad and growing interest in ESG investments is the threat of a possible overvaluation of ESG products (Bofinger et al. (2022)), which could be amplified by the effect of a growing investment demand - driven both by the context regulatory and by the request of the new generation of investors - towards potentially overvalued firms (Becker et al. (2022)). Different studies have been conducted to prove the relationship between impact investment and the rate of return for bonds and stocks (or mutual funds investing in green assets). In particular, the researchers often are interested in examining whether green or socially responsible focused products outperform their benchmarks or the 'traditional' investment. Literature reviews have been conducted on the link between ESG indicators and financial performance or market trends. A considerable number of papers have highlighted either a positive relationship (Clark et al. (2015); Verheyden et al. (2016)) or at least a non-negative association (Friede et al. (2015)) between ESG measures and performance. Moreover, if yield is one of the most studied aspects in the various analyses, risk represents an equally important element in portfolio strategies. In the study by Fooladi and Hebb (2022), it is highlighted that the discussion of whether SRI in the portfolios of CSR firms can provide better risk-adjusted returns for investors is controversial (even if these firms

[3]GSIA. Global Sustainable Investment Review. (2020). Available online: https://www.gsi-alliance.org/wp-content/uploads/2021/08/GSIR-20201.pdf.

are all profitable).

References

i. E. Agliardi, R. Agliardi: *Financing environmentally-sustainable projects with green bonds*. Environ. Dev. Econ., **24** (2019), n. 6, 608-623, https://doi.org/10.1017/S1355770X19000020.

ii. M. G. Becker, F. Martin, A. Walter: *The power of ESG transparency: The effect of the new SFDR sustainability labels on mutual funds and individual investors*. Fin. Res. Lett., **47** (2022), 102708, https://doi.org/10.1016/j.frl.2022.102708.

iii. F. Berg, J. F. Koelbel, R. Rigobon: *Aggregate confusion: The divergence of ESG ratings*. Rev. Fin., **26** (2022), n. 6, 1315-1344, https://doi.org/10.1093/rof/rfac033.

iv. Y. Bofinger, K. J. Heyden, B. Rock: *Corporate social responsibility and market efficiency: Evidence from ESG and misvaluation measures*. J. Bank. Fin., **134** (2022), http://dx.doi.org/10.1016/j.jbankfin.2021.10632.

v. Bugg-Levine, A., Emerson, J: *"Impact investing: Transforming how we make money while making a difference"*. John Wiley & Sons (2011).

vi. G. Clark, A. Feiner, M. Viehs: *"From the stockholder to the stakeholder: How sustainability can drive financial outperformance"*. Oxford University & Arabesque Partners (2015).

vii. I. J. Fooladi, G. Hebb: *Drivers of differences in performance of ESG-focused funds relative to their underlying benchmarks*. Glob. Fin. J., (2022), 100745, https://doi.org/10.1016/j.gfj.2022.100745.

viii. G. Friede, T. Busch, A. Bassen: *ESG and financial performance: Aggregated evidence from more than 2,000 empirical studies*. J. Sustain. Fin. Invest., **5** (2015), n. 4, 210–233, https://doi.org/10.1080/20430795.2015.1118917.

ix. H. Liang, L. Renneboog: *Corporate social responsibility and sustainable finance: A review of the literature*. Eur. Corp. Govern. Inst.–Finance Working Pap., (2020), 701, www.ecgi.global/sites/default/files/working_papers/documents/liangrenneboogfinal.pdf.

x. S. MacAskill, E. Roca, B. Liu, R. A. Stewart, O. Sahin: *Is there a green premium in the green bond market? Systematic literature review revealing premium determinants*. J. Clean. Prod., **280** (2021), 124491, https://doi.org/10.1016/j.jclepro.2020.124491.

xi. L. Renneboog, J. Ter Horst, C. Zhang: *Socially responsible investments: institutional aspects, performance, and investor behavior*. J. Bank. Fin., **32** (2008), n. 9, 1723-1742, https://doi.org/10.1016/j.jbankfin.2007.12.039.

xii. T. Verheyden, R. G. Eccles, A. Feiner: *ESG for all? The impact of ESG screening on return, risk, and diversification*. J. Appl. Corp. Fin., **28** (2016), n. 2, 47–56, https://doi.org/10.1111/jacf.12174.

xiii. I. Zumente, N. Lāce: *ESG Rating: Necessity for the Investor or the Company?* Sustainability, **13** (2021), 8940, https://doi.org/10.3390/su13168940.

3 The role of ESG ratings in investment portfolio choices

This third chapter introduces a model for constructing investment portfolios that integrates environmental, social, and governance (ESG) ratings as an additional objective, alongside the traditional focus on risk and return. We specifically examine two models that utilize different risk measures: variance and conditional value at risk (CVaR). By using CVaR, we can take into account information about the tail of the returns' distribution when optimizing the portfolio weights.

Then, the ε-constraint method is employed to find a set of optimal solutions for the resulting tri-objective optimization model and obtain a dotted representation of the corresponding efficient frontiers. Moreover, we present a technique for selecting the optimal portfolio according to the investor's profile and preference toward the three objectives.

To assess the effectiveness of the proposed portfolio strategy, an empirical analysis is carried out, considering the Refinitiv's ESG ratings for 538 companies that are part of the STOXX Europe 600 Index from January 2016 to December 2021. Results show that the tri-objective portfolio selection model shows a reduction in the profit opportunities but simultaneously a higher control on the risk, and in particular the CVaR-based strategies yield superior ex-post results.

3.1 Introduction and literature review

Several methods have been proposed in the literature to incorporate ESG information into the portfolio optimization process.

Firstly, the introduction of a more formal definition of environmental, social, and governance principles and the development of a standardized metric that assesses the extent to which a company aligns with the ESG objectives help investors in selecting the most promising assets in terms of sustainable and ethical practices, mimicking the stock-picking strategies used by portfolio managers. Following this approach, Liagkouras et al. (2022) used a screening procedure to identify a subset of ESG-compliant stocks as constituents of a mean-variance portfolio. They showed that concerning the UK stock market, selecting only firms that comply with the ESG standards leads to portfolio strategies that provide subordinate risk-return combinations compared to the available investment opportunities. Similarly, Pacelli et al. (2023) employed the ESG score as an ex-ante criterion for selecting assets in the mean-CVaR optimization problem. However, they obtained mixed empirical results, suggesting that sustainable investment performance is still heterogeneous across the Global Industrial Classification Standard (GICS) sectors. In a merging perspective between behavioral finance and ESG investing, Kaucic et al. (2023b) studied several preselection techniques based on stocks' ESG scores in a CPT-based portfolio model. Using an investment pool from the STOXX Europe 600 index, they assessed the profitability of the proposed strategies and showed that investments' financial performance is improved using ESG-based stock-picking techniques.

Secondly, by defining portfolio ESG score as the weighted average of its constituents' scores, an investor can impose a constraint for the minimum acceptable level of portfolio ESG. Pedersen et al. (2021) have computed the ESG-efficient frontier, displaying the highest possible Sharpe ratio for each ESG score. Then, they observed that, on a theoretical basis, ethical investors should be willing to accept lower portfolio returns in exchange for a more sustainable portfolio choice. De Spiegeleer et al. (2021) extended the mean-variance model following this approach; their analysis, with assets in either the STOXX Europe 600 and the Russell 1000 index, pointed out the impact of the choice of the rating agency as well as the specific market universe and investment period on the risk-return results. In the same framework, Morelli (2023), incorporated a constraint on the selected parameter in the mean-CVaR model exploiting only the environmental scores of the firms in the STOXX Europe 50 index. In this context, results show that environmentally committed portfolios result in being more efficient over the entire sample period under the CVaR-return profile, providing rewarding investments combined with mitigation of the tail market risk.

The third approach consists in the extension of the bi-objective mean-risk framework by directly including sustainability as a third criterion. In this direction, Garcia-Bernabeu et al. (2019) formalised the preference relation of an ESG-aware mean-variance

investor and introduced a multiobjective genetic algorithm to identify the surface of optimal portfolios. Then, they analysed the range of values attainable and the trade-off between return, risk, and sustainability for a set of institutional SRI open-end funds from Morningstar. Then, Hilario-Caballero et al. (2020) developed a method that combines two well-known multi-criteria decision tools, namely the stochastic multi-criteria acceptability analysis (SMAA) and the technique for order preferences by similarity solutions (TOPSIS), to provide robust solutions for this tri-objective framework. Xidonas and Essner (2022) proposed a multi-objective minimax-based optimisation model to build optimal ESG portfolios, maximising at the same time the risk performance of the three ESG pillars. Cesarone et al. (2022) adapted the standard ε-constrained method from mathematical programming to solve the tri-objective mean-variance-ESG optimization problem, and investigated the possible effects of SRI regulatory developments and ESG rating on the portfolio selection process over the past 15 years using data involving five major equity market indices. Similarly to these works, Lindquist et al. (2022) combined ESG scores with financial returns to generate an ESG-valued return and applied this mixed measure in a general mean-risk optimisation framework. Then, they considered an application to the optimisation of real estate investment trust (REIT) portfolios, employing variance and CVaR to quantify the investment risk.

This work is structured as follows. Section 3.2 describes the proposed portfolio optimization model, while Section 3.3 the method used for the resolution of the optimization model. Finally, the experimental part is in Section 3.4. Some concluding remarks are given in Section 3.5.

3.2 Model description

In the optimal wealth allocation process among a finite number of risky assets, investors have to recognise the sources of risk, then quantify and control them. This section provides an overview of the general mean-variance portfolio selection model (Markowitz (1952)) and a further generalization of this approach that considers alternative portfolio risk estimations. Then, the proposed portfolio selection model introduces, in the mean-risk framework, a third dimension linked to sustainability.

3.2.a The mean-variance asset allocation framework

A simple and effective way to measure risk is to consider the deviations of the observations of a random variable from a reference point. Thus, one can measure portfolio uncertainty using mean deviations, and quantify the risk with the standard deviation. Formally, given a random variable Z, its standard deviation is defined as the square root of the variance:

$$\sigma(Z) = \sqrt{E((Z - E(Z))^2)} \tag{3.2.1}$$

This quantity measures the uncertainty associated with Z. Moreover, the standard deviation respects the properties given in the following axiomatic definition (Rockafellar et al. (2006)):

- $\sigma(Z) = \sigma(Z + c)$, for all Z and $c \geq 0$ (shift invariance)
- $\sigma(0) = 0$, $\sigma(\lambda Z) = \lambda\sigma(Z)$ for all Z and $\lambda > 0$ (positive homogeneity)
- $\sigma(Z) \geq 0$ for all Z, with $\sigma(Z) > 0$ if Z is nonconstant (positivity)
- $\sigma(Z + Y) \leq \sigma(Z) + \sigma(Y)$ for all Z and Y (subadditivity).

Subadditivity property ensures that the standard deviation of a portfolio composed of two positions is not higher than the two components' volatility taken separately. Thus, it means that diversification does not increase uncertainty. Since the presence of the square root can lead to complex expressions when differentiated, the use of $\sigma(Z)$ in solving minimisation problems is not popular. However, if the square of the standard deviation, called variance, is minimised, the resulting quadratic minimisation problem is more easily solvable, and it is worth noting that minimising the standard deviation is equivalent to minimising the variance. For this reason, the choice of variance as a proxy for risk is more popular in asset allocation practice. In this framework, Markowitz (1952) developed the so-called M-V analysis which represents the cornerstone of modern portfolio theory. This theory suggests that portfolio choice should be based on the information given by two quantities: the expected portfolio rate of return and the portfolio variance. The main principle behind this paradigm can be formulated in two ways.

1. Among all portfolios with a given lower bound on the expected rate of return, find the ones with the lower variance.
2. Among all portfolios with the same upper bound on volatility, find the ones with the highest expected rate of return.

Now, suppose that agents act their investment decisions over a one-period horizon, and assume that the investment universe consists of n risky assets, with their rates of return expressed by the random variables $R_1, \ldots, R_n$. The random variable $R_p(\boldsymbol{x}) = \sum_{i=1}^{n} x_i R_i$ represents the rate of return of the portfolio $\boldsymbol{x} = (x_1, \ldots, x_n)^\top$, where $^\top$ is the transpose operator. The portfolio expected rate of return is denoted by $\mu_p(\boldsymbol{x}) = \sum_{i=1}^{n} x_i \mu_i$, where $\mu = (\mu_1, \ldots, \mu_n)^\top$ is the vector of assets' expected rates of return. Finally, $\sigma_p^2(\boldsymbol{x}) = \sum_{j=1}^{n}\sum_{i=1}^{n} x_i x_j \sigma_{i,j}$ expresses portfolio variance, where $\sigma_{i,j}$ is the covariance between assets i and j. Note that, given the $n \times n$ covariance matrix $\boldsymbol{\Sigma}$, portfolio variance can also be expressed compactly as $\sigma_p^2(\boldsymbol{x}) = \boldsymbol{x}^\top \boldsymbol{\Sigma} \boldsymbol{x}$. In the proposed portfolio design, short-selling is not allowed, and the set of feasible portfolios can be expressed as follows:

$$\Delta = \left\{ \boldsymbol{x} \in \mathbb{R}^n : x_i \geq 0, i = 1, \ldots, n, \sum_{i=1}^{n} x_i = 1 \right\}. \tag{3.2.2}$$

Hence, the M-V optimisation problem behind the first paradigm formulation is the following quadratic programming problem:

$$\begin{aligned} \min_{\boldsymbol{x}\in\Delta} \quad & \boldsymbol{x}^\top \Sigma \boldsymbol{x} \\ \text{s.t.} \quad & \boldsymbol{x}^\top \mu \geq \mu_* \,, \end{aligned} \tag{3.2.3}$$

where μ_* is the lower bound on the portfolio expected rate of return. Further, the optimisation problem behind the second formulation of the M-V principle can be expressed as a maximisation problem:

$$\begin{aligned} \max_{\boldsymbol{x}\in\Delta} \quad & \boldsymbol{x}^\top \mu \\ \text{s.t.} \quad & \boldsymbol{x}^\top \Sigma \boldsymbol{x} \leq \sigma^* \,, \end{aligned} \tag{3.2.4}$$

where σ^* denotes the upper bound on the portfolio volatility.

A portfolio that is a solution of Problem (3.2.3) or Problem (3.2.4) is called efficient. By varying μ_* in (3.2.3) or σ^* in (3.2.4) one can obtain the so-called efficient frontier in the variance-mean plane. Analysing the shape of the efficient frontier, investors can assess the trade-off between the portfolio expected rate of return and its variance, to select a portfolio aligned with their risk tolerance and investment objectives.

Criticism of mean-variance optimization

The standard Markowitz's formulation uses the variance as a proxy for the portfolio risk. Even if it is considered a milestone in the field, the M-V approach presents two well-known drawbacks.

- Variance is not a good choice for a risk measure since it penalises both the negative and the positive deviation from the mean symmetrically. Hence, it does not account for the asymmetric nature of risk, which concerns losses only. When asset returns are heavily skewed, or portfolio returns distribution presents fat tails, risks could be underestimated, and the M-V approach could lead to sub-optimal portfolio choices.
- The M-V framework often identifies efficient portfolios concentrated in a small subset of the investable universe. This feature is particularly evident when there is a large number of assets. In particular, during times of financial crisis, a concentrated and unbalanced portfolio could lead to substantial losses.

A solution is to replace the variance with an alternative measure to estimate the uncertainty of a random variable, the downside semi-variance Markowitz (1959). Given a random variable Z, this measure only accounts for the negative deviations from the mean, and it is defined as follows:

$$\sigma(Z) = \sqrt{E((Z - E(Z))_+^2)} \tag{3.2.5}$$

where $(Z - E(Z))_+ = \max(Z - E(Z), 0)$. The use of this measure in the asset allocation practice has been studied extensively, and the interested reader may refer to the works of Konno et al. (2002), and Markowitz et al. (1993).

3.2.b Extending the M-V model: the mean-risk approach

The mean-variance analysis can be significantly extended by replacing the variance with a suitable risk measure to evaluate portfolio risk. In other words, a risk measure quantifies the risk by associating a real number with a random payoff's loss distribution. Artzner et al. (1999) introduced the axiomatic definition of coherent risk measure, considering $\mathcal{R}(Z)$ which assigns a real value to a random variable Z that should satisfy the following properties:

1. $\mathcal{R}(Y) \leq \mathcal{R}(Z)$ if $Y \geq Z$ almost surely (monotonicity)
2. $\mathcal{R}(0) = 0; \mathcal{R}(\lambda Z) = \lambda \mathcal{R}(Z)$, for all Z and $\lambda > 0$ (positive homogeneity)
3. $\mathcal{R}(Z + Y) \leq \mathcal{R}(Z) + \mathcal{R}(Y)$, for all Z and Y (subadditivity)
4. $\mathcal{R}(Z + c) = \mathcal{R}(Z) - c$, for all Z and $c \geq 0$ (invariance).

In this context, the M-V principle can be generalised using another risk measure instead of variance. This approach is called mean-risk (M-R) analysis. Now, the choice paradigm can be stated in the following two ways.

1. Among all portfolios with a given lower bound on the expected return, find the ones that minimise the risk.
2. Among all portfolios with a given upper limit on the risk, find the ones with the highest expected performance.

With an abuse of notation, we will indicate the risk measure of the portfolio rate of return $R_p(\boldsymbol{x})$ by using $\mathcal{R}_p(\boldsymbol{x})$. This notation evidences the dependence from the vector of the portfolio weights $\boldsymbol{x}$ and simplifies the treatment. Hence, the M-R optimisation problems can be formulated as follows:

$$\begin{aligned} \min_{\boldsymbol{x} \in \Delta} \quad & \mathcal{R}_p(\boldsymbol{x}) \\ \text{s.t.} \quad & \boldsymbol{x}^\top \mu \geq \mu_* , \end{aligned} \tag{3.2.6}$$

where μ_* is the lower bound on the expected performance. An equivalent formulation is the following:

$$\begin{aligned} \min_{\boldsymbol{x} \in \Delta} \quad & \boldsymbol{x}^\top \mu \\ \text{s.t.} \quad & \mathcal{R}_p(\boldsymbol{x}) \leq \mathcal{R}^* , \end{aligned} \tag{3.2.7}$$

where $\mathcal{R}^*$ is the upper bound on portfolio risk. Similarly to what described above, a representation of the mean-risk efficient frontier can be obtained by solving Problems (3.2.6) and (3.2.7) with different lower and upper bounds μ_* and $\mathcal{R}^*$.

Quantile-based risk measures

In the latest years, quantifying significant losses and dealing with tail events has become a priority in banking and insurance due to the importance of the recently introduced regulations for solvency capital requirement. For this reason, a class of risk measures called quantile-based risk measures occupies a leading position in the risk management sector. Measures such as value-at-risk (VaR) and conditional value-at-risk (CVaR) have become very popular in portfolio optimisation problems. If one considers a random variable Z that expresses a random payoff, and given the significance level α, with $\alpha \in [0, 1]$, VaR can be defined as the threshold below which the payoff value falls with probability α (the so-called tail probability):

$$VaR_\alpha(Z) = -\inf_z \{z : P(Z \leq z) \geq \alpha\} . \tag{3.2.8}$$

Although VaR is easy to use and intuitive, it presents some critical pitfalls. First, it does not account for losses that exceed the VaR threshold. Second, it does not respect subadditivity property, which is essential to guarantee that diversification does not enhance the risk. To deal with these shortcomings, the conditional value-at-risk (CVaR) has been introduced by Rockafellar et al. (2002). The CVaR of a random variable Z at confidence level α is Mausser and Romanko (2018):

$$CVaR_\alpha(Z) = -E\left[Z|Z \leq -VaR_\alpha(Z)\right]\delta_\alpha - VaR_\alpha(Z)(1-\delta_\alpha) , \tag{3.2.9}$$

where $\delta_\alpha = \frac{P(Z \leq -VaR_\alpha(Z))}{\alpha}$. Note that if Z has a continuous and strictly increasing distribution function, then $CVaR_\alpha(Z) = -E(Z|Z \leq -VaR_\alpha(Z))$. In this case, the CVaR is a conditional expectation, namely the average of all losses greater or equal to VaR_α. It is worth noting that the CVaR risk measure satisfies all the above-presented axioms and is thus a coherent risk measure. For this reason, it has become popular in practical applications.

CVaR-based portfolio optimization

The mean-CVaR portfolio selection problem is a specification of the M-R model where an agent seeks to maximise return while minimising CVaR. Conceptually, the minimum CVaR portfolio with confidence level α is formulated as follows:

$$\begin{aligned} \min_{\boldsymbol{x} \in \Delta} \quad & CVaR_\alpha(\boldsymbol{x}) \\ \text{s.t.} \quad & \boldsymbol{x}^\top \mu \geq \mu_* . \end{aligned} \tag{3.2.10}$$

In this formulation, we indicate the conditional value-at-risk of the portfolio returns distribution $R_p(\boldsymbol{x})$ with $CVaR_\alpha(\boldsymbol{x})$ to underline the dependence of this measure on the portfolio weights and to simplify the treatment.

3.2.c Inclusion of ESG issues in the mean-risk framework

The models introduced in the previous sections focus solely on financial features and do not account for the increasing sensibility of political institutions, portfolio managers, and investors toward socially responsible investing. For this reason, we propose a tri-objective model, which extends the M-R analysis by including an objective of sustainability. The proposed model maximises both the expected return and the ESG assessment and - in the meantime - minimises the risk profile. We define the sustainable grade of a given portfolio as the weighted value of portfolio constituents' ESG scores. More formally it can be calculated as:

$$ESG_p(\boldsymbol{x}) = \sum_{i=1}^{n} x_i ESG_i \tag{3.2.11}$$

where ESG_i is the value of the ESG score for the i-th asset. Thus, $ESG_p(\boldsymbol{x})$ represents the sustainability propensity of portfolio $\boldsymbol{x}$ expressed in terms of ESG score. The preference relation of the M-R principle can be extended through the following formulation.

Definition 11. *According to the mean-risk-ESG (M-R-ESG) portfolio selection model, a portfolio $\boldsymbol{x}$ is preferred to a portfolio $\mathbf{y}$ if and only if $\mu_p(\boldsymbol{x}) \geq \mu_p(\mathbf{y})$, $\mathcal{R}_p(\boldsymbol{x}) \leq \mathcal{R}_p(\mathbf{y})$, and $ESG_p(\boldsymbol{x}) \geq ESG_p(\mathbf{y})$, with at least one strict inequality.*

It follows that the efficient frontier of the proposed mean-risk-ESG portfolio selection model can be obtained by solving the following tri-objective optimisation problem:

$$\min_{\boldsymbol{x}\in\Delta} \quad (-\mu_p(\boldsymbol{x}), \mathcal{R}_p(\boldsymbol{x}), -ESG_p(\boldsymbol{x})) \ . \tag{3.2.12}$$

The choice of an appropriate portfolio risk measure $\mathcal{R}_p(\boldsymbol{x})$ in the model specification is crucial since different measures may capture completely different characteristics of the portfolio return distribution. For instance, the asset allocation problem when the standard deviation is considered as a proxy for portfolio risk declines as follows:

$$\min_{\boldsymbol{x}\in\Delta} \quad \left(-\mu_p(\boldsymbol{x}), \sigma_p^2(\boldsymbol{x}), -ESG_p(\boldsymbol{x})\right) \ . \tag{3.2.13}$$

while the respective model that minimises $CVaR_\alpha(\boldsymbol{x})$ is

$$\min_{\boldsymbol{x}\in\Delta} \quad (-\mu_p(\boldsymbol{x}), CVaR_\alpha(\boldsymbol{x}), -ESG_p(\boldsymbol{x})) \ . \tag{3.2.14}$$

3.3 Optimization method

3.3.a ε-constrained technique

In this section, we face the issue of how to practically find the efficient surface of the mean-variance-ESG and mean-CVaR-ESG models. The standard ε-constrained method (Ehrgott (2005)) allows us to reformulate Problem (3.2.12) into a single-objective optimisation problem, which is more easily manageable. Let η and λ be the required target levels of the portfolio expected return and sustainability score, respectively. The resulting optimisation problem is the following:

$$\begin{aligned} \min_{\boldsymbol{x}\in\Delta} \quad & \mathcal{R}_p(\boldsymbol{x}) \\ \text{s.t.} \quad & \boldsymbol{x}^\top \mu \geq \eta \\ & ESG_p(\boldsymbol{x}) \geq \lambda\,. \end{aligned} \tag{3.3.15}$$

This single-objective asset allocation model minimises the risk measure with parametric lower bounds on the portfolio return and ESG. Computing the solutions of Problem (3.3.15) by appropriately varying η and λ leads to the determination of the mean-risk-ESG optimal portfolios. Following a procedure similar to Cesarone et al. (2022) and Lindquist et al. (2022), with this approach, one can obtain the M-R-ESG efficient portfolios and represent them in the risk-expected rate of return for several fixed levels of expected ESG thresholds.
Practically speaking, a suitable interval $[\lambda_{min}, \lambda_{max}]$ for λ is set, where $\lambda_{min} = ESG_p\left(\boldsymbol{x}_{min\mathcal{R}}\right)$, and $\lambda_{max} = ESG_p(\boldsymbol{x}_{maxESG})$. More specifically, λ_{min} is the sustainability level of the portfolio that minimises the risk, $\boldsymbol{x}_{min\mathcal{R}} = argmin_{\boldsymbol{x}\in\Delta}\mathcal{R}_p\left(\boldsymbol{x}\right)$, and λ_{max} is the maximum level of ESG obtainable from constructing a portfolio with the n available assets, namely $\boldsymbol{x}_{maxESG} = argmax_{\boldsymbol{x}\in\Delta}ESG_p\left(\boldsymbol{x}\right)$. Then, given a fixed ESG level λ, an appropriate interval $[\eta_{min}(\lambda), \eta_{max}(\lambda)]$ is determined. Notice that this interval depends on the choice of the threshold λ. On the one hand, $\eta_{min}\left(\lambda\right) = \mu_p\left(\boldsymbol{x}_{min\mathcal{R}}\left(\lambda\right)\right)$ is the rate of return of the portfolio that minimises the risk with a given sustainability requirement, namely the optimal solution of the following problem:

$$\begin{aligned} \min_{\boldsymbol{x}\in\Delta} \quad & \mathcal{R}_p(\boldsymbol{x}) \\ \text{s.t.} \quad & ESG_p(\boldsymbol{x}) \geq \lambda,\,. \end{aligned} \tag{3.3.16}$$

On the other hand, $\eta_{max}\left(\lambda\right) = \mu_p(\boldsymbol{x}_{\max\mu}(\lambda))$ is the rate of return of the portfolio that maximises the performance with a lower bound on the ESG score, which is the solution of:

$$\begin{aligned} \max_{\boldsymbol{x}\in\Delta} \quad & \mu_p(\boldsymbol{x}) \\ \text{s.t.} \quad & ESG_p(\boldsymbol{x}) \geq \lambda,\,. \end{aligned} \tag{3.3.17}$$

Note that by solving Problem (3.3.15) with $\lambda \in [\lambda_{min}, \lambda_{max}]$ and with $\eta = \eta_{min}(\lambda)$, one can obtain the ESG-risk efficient frontier. More precisely, setting $\eta = \eta_{min}$ and $\lambda = \lambda_{min}$, the optimal solution is the global minimum risk portfolio. Furthermore, setting $\lambda = \lambda_{max}$, the frontier collapses in a single portfolio composed of the asset with the highest ESG value. Finally, solving Problem (3.3.15) with $\lambda \in [\lambda_{min}, \lambda_{max}]$ and with $\eta = \eta_{max}(\lambda)$, the optimal solutions form the mean-ESG frontier.

3.3.b Scenario-based framework for portfolio optimisation

Let h indicate the investment horizon for the considered asset allocation problem on a daily basis. In this work, we use a scenario-based generation technique to obtain simulations for the random h-step ahead portfolio return $R_p(\boldsymbol{x})$ to estimate the values of the quantities involved in Problems (3.2.13) and (3.2.14). To this end, we consider the historical daily rates of return of the n risky assets observed during the time window $[0, T]$. We further assume that historical observations are good proxies for future rates of return.

More specifically, a scenario is defined as the set of h consecutive joint realisations of the rates of return for the n assets in a given time horizon. In this context, block bootstrapping techniques have shown good performance in preserving time series correlations (Guastaroba et al. (2009)). Thus, we adopt the so-called stationary bootstrap technique (Politis and Romano (1994)). This procedure considers a random block size, that is a set of consecutive scenarios with variable length, in order to bring some robustness with respect to the standard block bootstrap, which uses a fixed block size. The procedure works as follows. First, an optimal average block size B^* is selected following the technique introduced by Politis and White (2004). Then, a block of B^* consecutive daily observations is extracted randomly from the observed data frame. Finally, we repeat this exercise until the extracted sample reaches the desired size h, adjusting the last block length if the procedure exceeds the desired number of periods. Through this technique, for each asset i there is a bootstrap sample $\hat{R}_{i,k}$ of h rates of return, with $i = 1, \ldots, n$ and $k = 1, \ldots, h$. Then, the bootstrap-simulated h step ahead rate of return of the i-th asset is calculated as $\hat{R}_i = \prod_{k=1}^{h} \left(1 + \hat{R}_{i,k}\right)$. Given a portfolio $\boldsymbol{x}$, its h-step ahead rate of return is computed as $\hat{R}_p(\boldsymbol{x}) = \sum_{i=1}^{n} x_i \hat{R}_i$. To obtain an estimation of the empirical distribution of $R_p(\boldsymbol{x})$, this bootstrap procedure is repeated S times. Let $\hat{R}_p^s(\boldsymbol{x})$ be the s-th simulation of the portfolio return $R_p(\boldsymbol{x})$. Then, concerning Problem (11), the bootstrap estimators for $\mu_p(\boldsymbol{x})$ and $\sigma_p^2(\boldsymbol{x})$ are:

$$\hat{\mu}_p(x) = \frac{1}{S} \sum_{s=1}^{S} \hat{R}_p^s(\boldsymbol{x}) \tag{3.3.18}$$

and

$$\hat{\sigma}^2{}_p(x) = \sqrt{\frac{1}{S}\sum_{s=1}^{S}\left(\hat{R}_p^s(\boldsymbol{x}) - \hat{\mu}_p(x)\right)^2} \tag{3.3.19}$$

Regarding Problem (3.2.14), the CVaR expression in (3.2.9) can be estimated as an arithmetic mean, following the procedure suggested by Kaucic et al. (2019) and Rachev et al. (2008). Let $\hat{R}_p^{(1)}(\boldsymbol{x}) \leq \hat{R}_p^{(2)}(\boldsymbol{x}) \leq \ldots \leq \hat{R}_p^{(S)}(\boldsymbol{x})$ be the sorted sequence of the bootstrap simulations of the portfolio rates of return. Thus, the value-at-risk at confidence level α can be defined as follows:

$$\widehat{VaR}_\alpha(\boldsymbol{x}) = -\hat{R}_p^{(\lfloor \alpha S \rfloor)}(\boldsymbol{x}), \tag{3.3.20}$$

where $\lfloor \cdot \rfloor$ denotes the floor operator. Then, after some algebra, one can obtain the bootstrap simulated conditional value-at-risk, which is:

$$\widehat{CVaR}_\alpha(\boldsymbol{x}) = -\frac{1}{\alpha S}\left(\sum_{s=1}^{\lfloor \alpha S \rfloor}\hat{R}_p^s(\boldsymbol{x}) + (\alpha S - \lfloor \alpha S \rfloor)\hat{R}_p^{(\lfloor \alpha S \rfloor)}(\boldsymbol{x})\right) \tag{3.3.21}$$

3.3.c Portfolio selection based on agent's preferences

The procedure described in sections 3.3.a and 3.3.b allow us to compute a set of efficient solutions, called efficient frontier, denoted by $\mathcal{P}$. The portfolio more compatible with the investor's attitude toward reward, risk, and sustainability is selected on $\mathcal{P}$ as follows. First, the objective functions for each efficient portfolio $\boldsymbol{x} \in \mathcal{P}$ are normalised using the formulas:

$$\widetilde{\mu}_p(\boldsymbol{x}) = \frac{\mu_p(\boldsymbol{x}) - \mu_p^{min}}{\mu_p^{max} - \mu_p^{min}}, \tag{3.3.22}$$

$$\widetilde{\mathcal{R}}_p(\boldsymbol{x}) = \frac{\mathcal{R}_p^{max} - \mathcal{R}_p(\boldsymbol{x})}{\mathcal{R}_p^{max} - \mathcal{R}_p^{min}}, \tag{3.3.23}$$

and

$$\widetilde{ESG}_p(\boldsymbol{x}) = \frac{ESG_p(\boldsymbol{x}) - ESG_p^{min}}{ESG_p^{max} - ESG_p^{min}}. \tag{3.3.24}$$

In the previous formulas, $\mu_p^{max} = max_{\boldsymbol{x}\in\mathcal{P}}\mu_p(\boldsymbol{x})$, and $\mu_p^{min} = min_{x\in\mathcal{P}}\mu_p(\boldsymbol{x})$; $\mathcal{R}_p^{max} = max_{\boldsymbol{x}\in\mathcal{P}}\mathcal{R}_p(\boldsymbol{x})$, and $\mathcal{R}_p^{min} = min_{\boldsymbol{x}\in\mathcal{P}}\mathcal{R}_p(\boldsymbol{x})$; $ESG_p^{max} = max_{\boldsymbol{x}\in\mathcal{P}}\ ESG_p(\boldsymbol{x})$, and $ESG_p^{min} = min_{\boldsymbol{x}\in\mathcal{P}}\ ESG_p(\boldsymbol{x})$. Notice that, with the normalization of Eqn. (3.3.23), a higher value of $\widetilde{\mathcal{R}}_p$ is preferable, and $\widetilde{\mathcal{R}}_p$ is well-suited in a maximization problem. Next, we introduce a vector that expresses the agent's preferences $\mathbf{w}_{pref} = (w_\mu, w_\mathcal{R}, w_{ESG}) \in [0,1]^3$ such that $w_\mu + w_\mathcal{R} + w_{ESG} = 1$. Finally, the efficient portfolio tailored to the investor's financial reward/risk/ESG profile $\mathbf{w}_{pref}$ can be obtained by solving the following single-objective problem that maximises the weighted sum of the normalised

objective functions:

$$\max_{\boldsymbol{x}\in\mathcal{P}} \left\{ w_\mu \widetilde{\mu}_p(\boldsymbol{x}) + w_{\mathcal{R}} \widetilde{\mathcal{R}}_p(\boldsymbol{x}) + w_{ESG} \widetilde{ESG}_p(\boldsymbol{x}) \right\}. \tag{3.3.25}$$

3.4 Experimental part

The experimental analysis focuses on the European market because from 2015 the European Union, through Directive 2014/95/EU, has imposed the disclosure of companies' non-financial information. In this context, the investable universe consists of a subset of the STOXX Europe 600 Index constituents. We select a pool of stocks for which an ESG score from the same provider is available during the whole period. Among all the possible providers available in the market, we have decided to utilise Refinitiv's ESG scores since they are applied widely in the financial industry. These scores are based on a company's performance on over 400 ESG indicators across ten categories: climate change, emissions, environmental product innovation, water resources, human rights, labor practices, community development, supply chain, corporate governance, and anti-corruption. The score provider uses company-reported data and third-party data sources to assess ESG performance (including ESG controversies impacting the firms). Each indicator's weighting is determined through expert judgment, data analysis, and consultation with industry experts and stakeholders. Refinitiv's ESG scores are a percentile ranking, with 100 indicating the highest score and 0 indicating the lowest. The scores are based on the relative performance of ESG factors within the company's sector (for environmental and social) and country of incorporation (for governance) and are updated monthly. As a result, the dataset consists of the monthly rates of return and the monthly ESG ratings from 01/01/2016 to 31/12/2021 for 538 firms, downloaded from Refinitiv Datastream.

The empirical analysis consists of two stages. First, we investigate the relation among reward, risk, and ESG scores for the istances (3.2.13) and (3.2.14) of the tri-objective optimisation problem (3.2.12). To achieve this goal, they consider a dotted representation of the efficient frontiers for each ex-post month, employing the procedure detailed in Section 3.3.a with 10 values for parameter λ and, for each value of λ, 10 suitable levels of $\eta(\lambda)$, for a total of 100 portfolios. Then, we analyse the ex-post performance of a set of portfolios selected on the efficient frontiers, based on six typical attitudes toward the reward, risk, and ESG criteria.

3.4.a Mean-risk-ESG efficient frontiers

We consider an investment plan with one-month horizon and employ a rolling window scheme with 72 end-of-month investments from 29/01/2016 to 31/12/2021. For each choice of the risk measure in Problem (3.3.15), and for each month, we exploit an in-sample window of 1000 days to set up the financial parameters, that are μ_p, σ_p^2 and

CVaR_α. Then, we calculate the portfolio ESG score in (3.2.11) by using the stocks' ESG scores at the beginning of the corresponding ex-post month.

Case 1: mean-variance-ESG model

Regardless of the selected month, we note that, from a qualitative perspective, the relationships between expected return, variance, and ESG score remain consistent over time. Consequently, we discuss and provide a detailed commentary on the findings pertaining to a particular month. Specifically, the 100 optimal portfolios at December 31, 2021 are displayed in Figure 3.1 on the left in the mean-variance-ESG tridimensional space. On the one hand, this dotted representation reveals that there exists a clear trade-off between portfolio reward, measured by the expected rate of return, and portfolio's sustainability attitude, indicated by the ESG score. The graph illustrates that portfolios with higher rates of return tend to exhibit lower ESG scores, implying that pursuing greater financial gains may involve compromising sustainability considerations within the portfolio. On the other hand, one can observe the relation between portfolio rate of return and portfolio variance. Indeed, portfolios with higher rates of return are typically associated with higher risks, in line with the modern portfolio theory. This is because the demand of higher financial reward may involve investing in riskier assets that prioritise financial gains that could be more sensible to market fluctuations. To better relate these findings to the sustainability principles, Figure 1 on the right reports the projection of the 3D efficient frontier on the variance-mean plane for increasing levels of the ESG threshold. It can be noted that, at the same level of variance, increasing the ESG threshold leads to a reduction in portfolio profitability.

Case 2: mean-CVaR-ESG model

To analyse the robustness of the results in terms of increasing tail-risk severity, we consider two different tail levels, $\alpha = 0.05$ and $\alpha = 0.01$. The first choice is the most commonly used in practice, and it is also the suggested value by the solvency regulamentation, while the second choice represents a more prudential approach. Building upon the methodology employed in Case 1, we investigate the interplay among portfolio rate of return, CVaR, and ESG criteria during the final ex-post month. These findings can be extended to the other months involved in the analysis, suggesting consistent conclusions.

More specifically, Figure 3.2 and Figure 3.3 report on the left the dotted approximations of the mean-CVaR_α-ESG frontiers with $\alpha = 0.05$ and $\alpha = 0.01$ respectively, and on the right the corresponding projections on the CVaR_α-mean plane for increasing ESG thresholds. A first comparison reveals that the severity of tail-risk does not have a significant effect on the frontier shape. As presented in Case 1, increasing the ESG threshold, for fixed levels of CVaR_α, leads to a decrease in the portfolio expected rate of return. These results align with the existing literature, confirming that a higher sus-

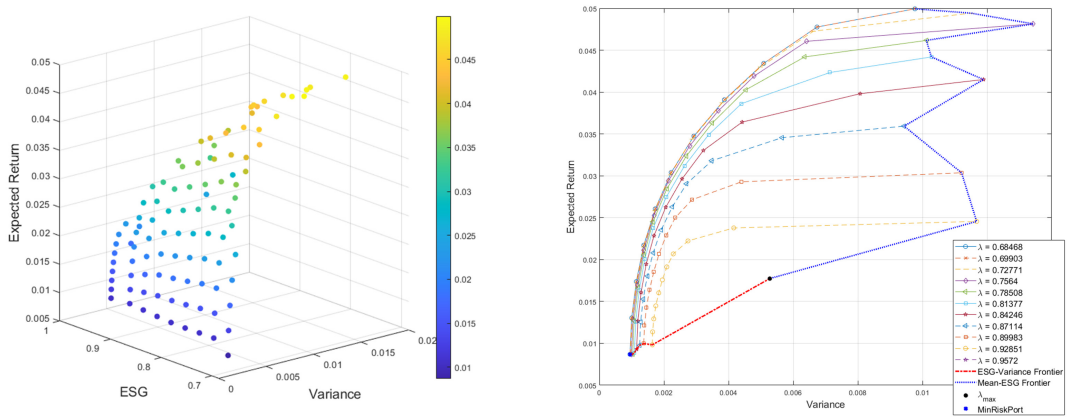

Figure 3.1: Dotted representation at 31/12/2021 of the mean-variance-ESG efficient frontier on the left, and corresponding projection on the variance-mean plane on the right for different values of the ESG threshold λ. The blue dotted line represents the mean-ESG frontier, while the red one is the ESG-variance frontier.

tainability demand within the portfolio entails a trade-off in terms of financial reward.

3.4.b Ex-post investment analysis

In this section, we focus on the profitability of portfolios selected from the proposed mean-risk-ESG efficient frontiers for different types of investors. Each agent described in Table 3.1 exhibits a specific attitude toward the three criteria. In particular, these preferences are encoded as a vector of non-negative weights that sum up to one, namely $\mathbf{w}_{pref} = (w_\mu, w_\mathcal{R}, w_{ESG})$. Higher values of a weight mean higher interest toward the corresponding criterion. The first three agents display a strong preference towards one criterion while equally weighting the other two. The fourth preference combination represents an investor who allocates equal weights to the three objectives. The fifth agent makes decisions based on the M-R principle with a small interest in sustainable criterion. Finally, Agent 6 pays great attention to both risk control and sustainability, with a marginal focus on financial reward.

Ex-post performance metrics

We evaluate the attractiveness of the investment opportunities from different points of view. In general terms, the random variable expressing the ex-post portfolio rate of return is denoted by r_p^{out}. Let $r_{p,t}^{out}$, $t = 1, \ldots, 72$ indicate the realised portfolio rates of return for each strategy in the ex-post (also called out-of-sample) window from 29/01/2016 to 31/12/2021. The STOXX Europe 600 Index from the european stock

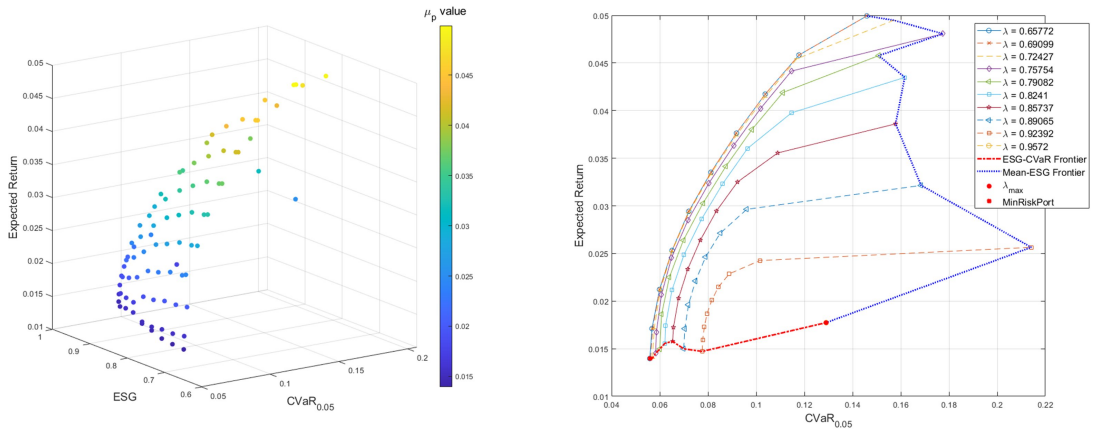

Figure 3.2: Dotted representation at 31/12/2021 of the mean-CVaR$_{0.05}$-ESG efficient frontier on the left, and corresponding projection on the CVaR$_{0.05}$-mean plane on the right for different values of the ESG threshold λ. The blue dotted line represents the mean-ESG frontier, while the red one is the ESG-CVaR$_{0.05}$ frontier.

Agent	$\mathbf{w}_\mu$	$\mathbf{w}_\mathcal{R}$	$\mathbf{w_{ESG}}$
1	0.25	0.25	0.50
2	0.25	0.50	0.25
3	0.50	0.25	0.25
4	0.33	0.33	0.33
5	0.45	0.45	0.10
6	0.10	0.45	0.45

Table 3.1: Types of agents considered in the ex-post analysis based on different attitudes toward the expected rate of return $\mathbf{w}_\mu$, the risk measure $\mathbf{w}_\mathcal{R}$, and the ESG score $\mathbf{w_{ESG}}$.

market is employed as benchmark, and its rate of return in the out of sample period is indicated with r_b^{out}. Similarly, the realised ex-post benchmark returns are specified as $r_{b,t}^{out}$, $t = 1, \ldots, 72$. At first, to measure the profitability of the analysed strategies, we consider the annualised mean return (AMR) during the ex-post window. Then, we examine the strategies' performance with respect to the benchmark by assessing the annualised excess mean return (AMER). Further, we conduct a t-test on the differences between observed portfolio and benchmark returns within the ex-post window. To evaluate the compensation earned per unit of portfolio variance during the investment period, we consider the so-called Sharpe ratio (Sharpe (1966)). This measure is defined as the ratio between the average of the realised rates of return, μ^{out}, and their standard deviation, σ^{out}:

$$SR = \frac{\mu^{out}}{\sigma^{out}}\,. \tag{3.4.26}$$

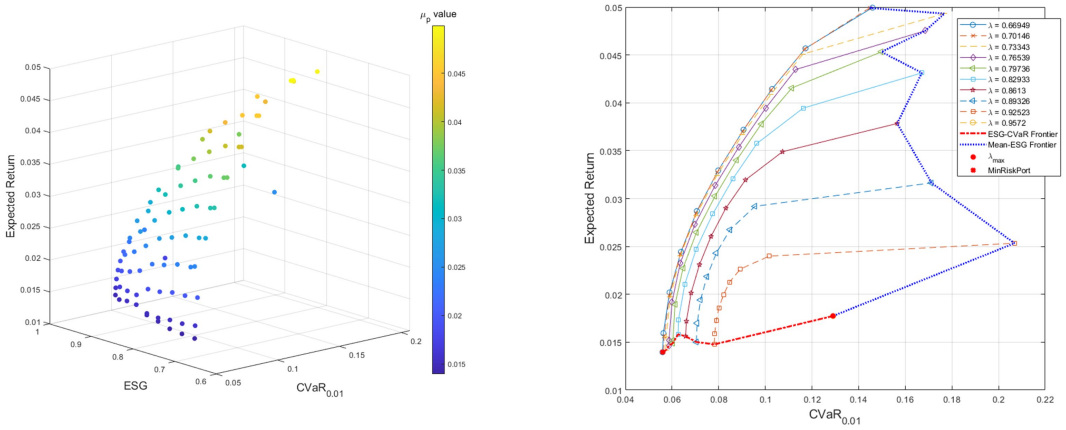

Figure 3.3: Dotted representation at 31/12/2021 of the mean-CVaR$_{0.01}$-ESG efficient frontier on the left, and corresponding projection on the CVaR$_{0.01}$-mean plane on the right for different values of the ESG threshold λ. The blue dotted line represents the mean-ESG frontier, while the red one is the ESG-CVaR$_{0.01}$ frontier.

Moreover, we account for a performance measure which is constructed on the basis of conditional value-at-risk, the so-called Rachev ratio (Rachev et al. (2008)). It considers the extra rates of return of the portfolio on the ex-post window, and it is defined as follows:

$$\mathrm{RaR}^{\mathrm{out}}_{\alpha_1,\alpha_2} = \frac{\mathrm{CVaR}_{\alpha_1}\left(\mathrm{r}^{\mathrm{out}}_{\mathrm{b}} - \mathrm{r}^{\mathrm{out}}_{\mathrm{p}}\right)}{\mathrm{CVaR}_{\alpha_2}\left(\mathrm{r}^{\mathrm{out}}_{\mathrm{p}} - \mathrm{r}^{\mathrm{out}}_{\mathrm{b}}\right)}. \tag{3.4.27}$$

This ratio is computed by dividing the conditional value-at-risk at level α_1 of the ex-post portfolio returns' distribution on the right tail by the conditional value-at-risk at level α_1 on the left tail. These quantities are estimated using the formula (3.3.21) with $\hat{R}^s_p(\boldsymbol{x}) = r^{out}_{p,t} - r^{out}_{b,t}$, and $s = t$. The Rachev ratio represents the reward potential for positive returns compared to the risk potential for negative returns at a quantile level defined by the user. In this analysis, we set $\alpha_1 = \alpha_2 = 0.05$ and indicate the corresponding performance ratio with $\mathrm{RaR}^{out}_{0.05}$. Then, the diversification level of the optimal portfolios is measured using the normalised diversification index (DI), given by:

$$DI = \frac{1}{72}\sum_{t=1}^{72}\frac{1-\sum_{i=1}^{n}(x_{i,t})^2}{1-\frac{1}{n}}, \tag{3.4.28}$$

where $x_{i,t}$ is the weight of asset i at time t. This quantity is equal to 0 when all the capital is concentrated in one single asset and is 1 for the equally-weighted portfolio. Thus, highest diversified portfolios present an higher DI value. Moreover, to get an impression of the transaction costs involved, we calculate the average turnover over the

out-of-sample period:

$$TR = \frac{1}{72} \sum_{t=1}^{72} \sum_{i=1}^{n} |x_{i,t} - x_{i,t-}| \,, \tag{3.4.29}$$

where $x_{i,t-}$ is the weight of asset i at ex-post time t immediately before the rebalancing phase. A greater value of TR indicates a more expensive investment strategy. Finally, we analyse the risk exposure for the three proposed models (3.2.13) and (3.2.14) with $\alpha = 0.01$ and $\alpha = 0.05$ and the six considered agents using three measures: the ex-post standard deviation σ^{out}, the conditional value-at-risk at the alpha level 0.05, $\text{CVaR}^{out}_{0.05}$, and the skewness of the out-of-sample returns distribution, Skew^{out}.

Analysis of the best investment models

Table 3.2 and Table 3.3 present the results of the ex-post performance measures described in Section 4.2.1 for the three considered models and for the agents whose preferences are described in Table 3.1. The following points summarise the main findings for the ex-post analysis of the considered investment strategies.

1. Strategies employing conditional value-at-risk yield superior ex-post performance results in absolute terms (AMR) as well as in relative terms (AMER). For almost all the considered agents, the ex-post portfolio rates of return of the CVaR models exhibit statistically significant positive deviations from the benchmark. The only instance where the mean-variance-ESG model slightly outperforms in terms of AMR and AMER is for agent 5. In addition, all the considered strategies are superior to the benchmark in terms of ex-post Sharpe ratio, with values at least three times higher, meaning a better capability in terms of return-risk trade-off.

2. We analyse the relationships between the three objectives without introducing additional portfolio restrictions, such as diversification and turnover constraints, that could somehow affect the ex-post statistics. However, strategies that adopt CVaR as a risk measure show superior diversification levels while ensuring lower turnover. Additionally, regarding skewness, it is noteworthy to observe that only the realised rates of return of Agent 6 (the most conservative profile regarding risk-return-sustainability preferences) exhibit ex-post skewness of the same sign as the benchmark. In all other cases, the proposed strategies have positive skewness, suggesting a right-skewed distribution which implies potential profits in the right tails.

3. Incorporating ESG criteria into the standard mean-risk framework could lead to a reduction in tail-risk compared to the benchmark. As depicted in Table 3, all the considered asset allocation models exhibit a lower realised conditional value-at-risk than the benchmark.

   However, contrasting results emerge when examining portfolio standard deviation. Agents 3 and 5, representing profiles more inclined towards financial performance in addition to risk-ESG considerations, display higher exposure to standard deviation compared to both the benchmark and the other strategies considered.

4. Agents 3 and 5 exhibit the best ex-post performance in absolute and relative terms. Additionally, they demonstrate superior results in terms of the Rachev ratio. This

fact implies that an investor with a profile more inclined towards return rather than risk and sustainability considerations has the potential to achieve higher profits than others. However, as highlighted in the previous point, they also display higher ex-post standard deviations, being potentially more vulnerable to sudden market fluctuations.

5. A highly promising investment profile that demonstrates excellent performance across all considered measures is represented by Agent 4. This configuration consistently exhibits the lowest values of ex-post standard deviation and ex-post CVaR, regardless of the chosen risk measure in the multi-objective optimisation model. Additionally, it evidences good outcomes regarding diversification (ranked as the second-best strategy) and turnover (in line with the other agents but worse than Agent 6). The ability to contain variability and tail risk enables this configuration to achieve outstanding results in the Sharpe ratio while maintaining a Rachev ratio comparable to that of Agents 3 and 5, albeit with slightly lower average rates of return.

Agent	Risk	AMR	AMER	p-value	SR	$\mathbf{RaR}^{\mathbf{out}}_{\mathbf{0.05}}$
1	σ^2	0.1530	0.1035	0.1086	0.3228	1.8570
	$CVaR_{0.05}$	0.1854	0.1360	0.0431*	0.4322	1.9627
	$CVaR_{0.01}$	0.1798	0.1304	0.0479*	0.4252	1.9656
2	σ^2	0.1737	0.1243	0.0493	0.4616	1.9376
	$CVaR_{0.05}$	0.2045	0.1550	0.0251*	0.5009	2.0755
	$CVaR_{0.01}$	0.2009	0.1514	0.0271*	0.4959	2.0640
3	σ^2	0.3416	0.2922	0.0011**	0.5983	2.8373
	$CVaR_{0.05}$	0.4416	0.3922	0.0003**	0.5935	2.7423
	$CVaR_{0.01}$	0.4049	0.3555	0.0008**	0.5583	2.9789
4	σ^2	0.2061	0.1567	0.0204*	0.5192	2.2592
	$CVaR_{0.05}$	0.2803	0.2309	0.0022*	0.6064	2.1930
	$CVaR_{0.01}$	0.2777	0.2283	0.0021*	0.6155	2.1657
5	σ^2	0.3521	0.3027	0.0006**	0.6205	3.0140
	$CVaR_{0.05}$	0.3425	0.2931	0.0009**	0.5798	3.0260
	$CVaR_{0.01}$	0.3495	0.3001	0.0006**	0.6044	2.8550
6	σ^2	0.1114	0.0620	0.2231	0.2552	1.8384
	$CVaR_{0.05}$	0.1312	0.0818	0.1393	0.3371	1.6957
	$CVaR_{0.01}$	0.1305	0.0811	0.1432	0.3316	1.6997
Benchmark		0.0494	-	-	0.1040	-

Table 3.2: Ex-post performance measures for the three istances of Problem (3.2.12), corresponding to each preference vector outlined in Table 3.1. The third column represents the average rate of return, while the fourth column indicates the average excess rate of return over the STOXX Europe 600 Index. The fifth column reports the p-value of the t-test assessing the differences between ex-post portfolio returns and the benchmark. A single asterisk (*) denotes rejection of the null hypothesis of mean equality at a 5% confidence level, while two asterisks (**) signify rejection at a 1% confidence level. Lastly, the sixth and seventh columns display the ex-post Sharpe and Rachev ratios.

In conclusion, these findings point out the superiority of a well-balanced strategy (represented by Agent 4), which displays equal preferences for the three portfolio objectives. Moreover, while incorporating ESG criteria is crucial to reduce tail-risk, relying solely on them (as for configurations 2 and 6) leads to mediocre outcomes. Optimal financial results are achieved by Agents 3 and 5 who approach the standard mean-variance framework, confirming the strength of the M-V model in investment decision-making.

Agent	**Risk**	**DI**	**TR**	$\sigma^{\mathbf{out}}$	$\mathbf{CVaR}_{\mathbf{0.05}}^{\mathbf{out}}$	$\mathbf{Skew}^{\mathbf{out}}$
1	σ^2	0.7041	0.3564	0.0395	0.0626	0.1852
	$CVaR_{0.05}$	0.7500	0.2410	0.0358	0.0542	0.0903
	$CVaR_{0.01}$	0.7429	0.2177	0.0352	0.0545	0.1347
2	σ^2	0.8413	0.3255	0.0313	0.0504	-0.2293
	$CVaR_{0.05}$	0.8780	0.2920	0.0340	0.0526	0.1736
	$CVaR_{0.01}$	0.8728	0.3069	0.0337	0.0536	0.1381
3	σ^2	0.6576	0.6767	0.0476	0.0501	0.7494
	$CVaR_{0.05}$	0.6388	0.5440	0.0620	0.0748	0.7326
	$CVaR_{0.01}$	0.6352	0.5518	0.0604	0.0694	0.8596
4	σ^2	0.7975	0.4400	0.0331	0.0485	0.0483
	$CVaR_{0.05}$	0.8223	0.3195	0.0385	0.0482	0.2787
	$CVaR_{0.01}$	0.8203	0.3526	0.0376	0.0476	0.2748
5	σ^2	0.8397	0.3952	0.0473	0.0544	0.3290
	$CVaR_{0.05}$	0.8522	0.3331	0.0492	0.0553	0.4494
	$CVaR_{0.01}$	0.8481	0.3368	0.0482	0.0566	0.3862
6	σ^2	0.7849	0.1092	0.0364	0.0646	-0.0597
	$CVaR_{0.05}$	0.8228	0.0970	0.0324	0.0591	-0.4122
	$CVaR_{0.01}$	0.8191	0.0984	0.0328	0.0598	-0.4056
Benchmark		-	-	0.0396	0.0927	-0.5204

Table 3.3: Ex-post performance measures for the three istances of Problem (3.2.12), corresponding to each preference vector outlined in Table 3.1. The third and fourth columns provide the average ex-post diversification index and the mean portfolio turnover. Then, the ex-post standard deviation is displayed in column 5, while the sixth column focuses on the conditional value-at-risk of the distribution of ex-post rates of return at the 0.05 significance level. In the end, the last column presents the skewness of the portfolio returns' distribution during the out-of-sample window.

Ex-post net wealth evolution

To conclude this section, we assess the profitability of the different portfolio allocation strategies for each preference configuration. To achieve this goal, they consider an initial wealth allocation $W_0 = 1,000,000$€. Then, we explicitly evaluate the magnitude of the tradings through a transaction costs structure where the applied commissions depend on the range of traded monetary amount, as in Beraldi et al. (2021) and Kaucic et

al. (2023a). More precisely, the transaction cost structure is characterised by decreasing cost rates as the traded value increases.

Figure 3.4 highlights the behavior of the six investment strategies in terms of the produced net wealth for the two proposed investment models. The top graph displays the net wealth evolution of the six considered preference behaviours for the mean-variance-ESG model. The bottom left graph corresponds to the mean-CVaR$_{0.05}$-ESG model, while the bottom right graph shows the ex-post performance of the 6 agents using the CVaR$_{0.01}$ as a risk measure.

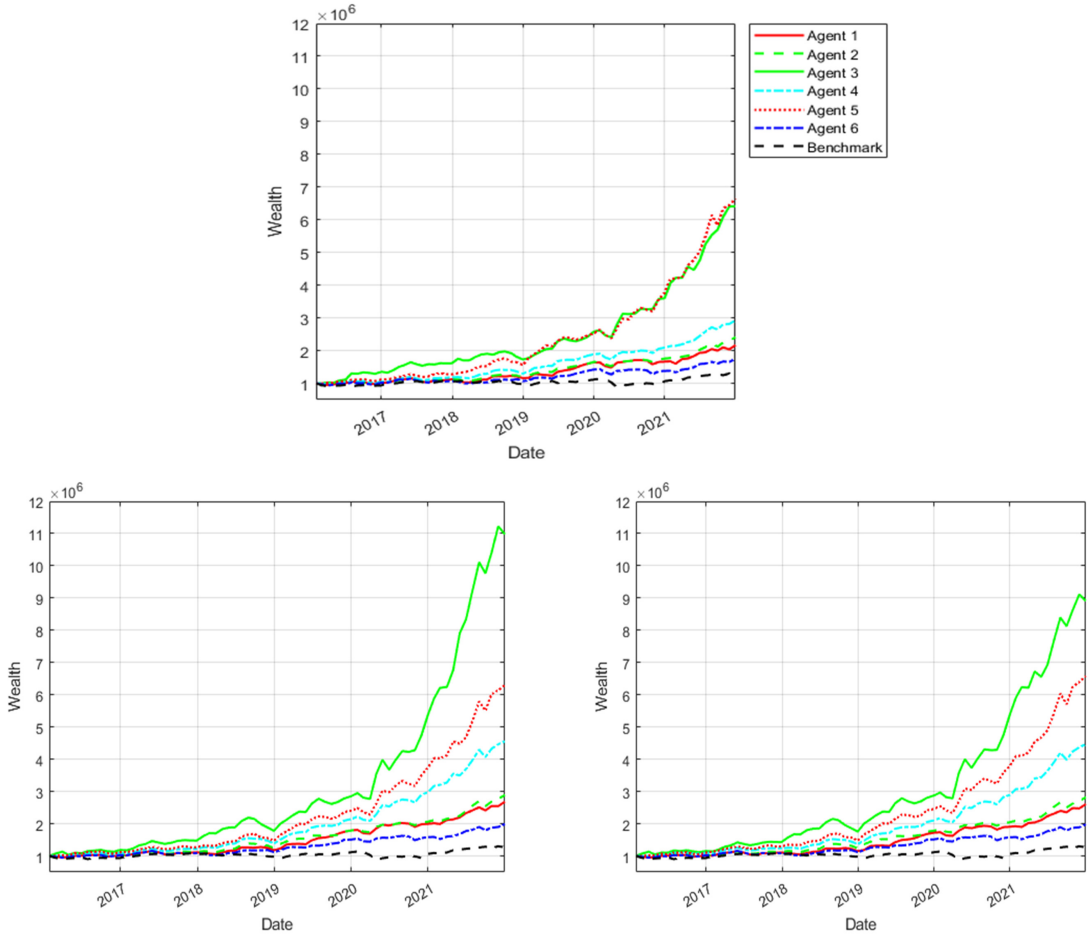

Figure 3.4: Ex-post net wealth evolution for the three models proposed and the six agents' configurations. Top figure shows the results for the mean-variance-ESG model. The figure on the bottom left displays the net wealth evolution for the mean-CVaR$_{0.05}$-ESG strategy. Finally, the on the bottom right shows results for the mean-CVaR$_{0.01}$-ESG model. The initial wealth investment is set equal to $1,000,000$€, and the transaction costs structure considered is characterized by decreasing cost rates as the traded amount increases.

As previously noted, configurations 3 and 5 demonstrate superior financial performance than the other ones. Concerning the mean-variance-ESG model, Agents 3 and 5 show a similar net wealth evolution. However, when employing the CVaR-based risk measure and specifically in the case of CVaR$_{0.01}$, Agent 3 exhibits a substantial

outperformance, attaining within a span of just five years a final wealth that is eleven times greater than the initial investment. Moreover, the three plots confirm that, as described formerly, CVaR-based models yield superior ex-post results. Overall, the proposed asset allocation models outperform the STOXX Europe 600 Index and show a remarkable trade-off between reward and risk. In particular, we highlight that configuration 3, which represents a balanced profile among risk, return, and sustainability, exhibits excellent results in terms of net wealth, thereby it could be considered a viable investment option.

3.5 Conclusions

The risk measures employed in the proposed tri-objective portfolio selection model show a reduction in the profit opportunities but simultaneously a higher control on the risk, following the literature. Consequently, the results show that ESG scoring can be used for better risk control, especially during periods of high volatility. This finding supports the broad debate among scholars and practitioners about ESG relevance and disclosure in investment decisions.

The discussion leads organisations to disclose information about their efforts to promote sustainability concerning the environment, society, and the economy. As noticed before, based on the increase in voluntary corporate non-financial disclosure, the European Commission (EC) has decided to implement mandatory disclosure across European countries. Introducing the European Union Directive 2014/95/EU, the EC requires organisations to disclose non-financial and diversity information. The Directive aims to increase transparency regarding social and environmental issues and ensure consistent and relevant information across companies. It requires large undertakings, public interest entities (PIEs), and organisations with over 500 employees to report information on environmental, social, employee, human rights, diversity, anti-corruption, and bribery matters.

Some studies (Clementino and Perkins (2021); Mervelskemper and Streit (2017); Santamaria et al. (2021)) showed the importance of assessing the combinations of quality and quantity disclosures for ESG score, demonstrating the effectiveness of disclosure measurements on ESG score. However, these studies are based on different disclosure frameworks used by companies that could influence the ESG rating developed by various agencies. Indeed, the critical characteristic of the Directive is the lack of a specific standard/guideline to prepare non-financial information. The Directive neither suggests the content nor the reporting method for measuring, monitoring and managing undertakings' performance and their impact on society. Many empirical studies recognise the failure of companies to include crucial information that investors and other stakeholders consider essential under the new mandatory regime.

Furthermore, comparing the reported information across different companies can be challenging, and data users often doubt their reliability. For these reasons, the

European Commission revised the non-financial Directive, issuing in December 2022 the new Corporate Social Responsibility Directive. It requires target companies to report ESG content applying a double materiality perspective in compliance with European Sustainability Reporting Standards (ESRS). These standards provide a common structure and content to disclose ESG factors, allowing stakeholders to better assess companies' ESG performance. We expect that applying the new ESRS standards could harmonise the ESG scores and impact the investment portfolio allocation strategies.

Bibliography

P. Artzner, F. Delbaen, M. Eber, D. Heath: *Coherent Measures of Risk*. Math. Fin., **9** (1999), n. 3, 203-228, https://doi.org/10.1111/1467-9965.00068.

B. R. Auer, F. Schuhmacher: *Do socially (ir)responsible investments pay? New evidence from international ESG data*. Quart. Rev. Econ. Fin., **59** (2016), 51–62, https://doi.org/10.1016/j.qref.2015.07.002.

P. Beraldi, A. Violi, M. Ferrara, C. Ciancio, B. A. Pansera: *Dealing with complex transaction costs in portfolio management*. Ann. Oper. Res., **299** (2021), n. 1, 7–22, https://doi.org/10.1007/s10479-019-03210-5.

X. Blasco, J. M. Herrero, J. Sanchis, M. Martínez: *A new graphical visualisation of n-dimensional Pareto front for decision-making in multiobjective optimisation*. Inf. Sci., **178** (2008), n. 20, 3908-3924, https://doi.org/10.1016/j.ins.2008.06.010.

F. Cesarone, M. L. Martino, A. Carleo: *Does ESG Impact Really Enhance Portfolio Profitability?* Sustainability, **14** (2022), n. 4, 2050, https://doi.org/10.3390/su14042050.

E. Clementino, R. Perkins: *How Do Companies Respond to Environmental, Social and Governance (ESG) ratings? Evidence from Italy*. J. Bus. Ethics, **171** (2021), 379–397, https://doi.org/10.1007/s10551-020-04441-4.

J. De Spiegeleer, S. Höcht, D. Jakubowski, S. Reyners, W. Schoutens: *ESG: A new dimension in portfolio allocation*. J. Sustain. Fin. Invest., (2021), 1-41, https://doi.org/10.1080/20430795.2021.1923336.

M. Ehrgott: *"Multicriteria Optimization"*. Springer Science & Business Media, Berlin (2005).

A. Garcia-Bernabeu, J. V. Salcedo, A. Hilario, D. Pla-Santamaria, J. M. Herrero: *Computing the mean-variance-sustainability nondominated surface by Ev-MOGA*. Complexity, (2019), 1-12, https://doi.org/10.1155/2019/6095712.

G. Guastaroba, R. Mansini, M. G. Speranza: *On the effectiveness of scenario generation techniques in single-period portfolio optimisation*. Eur. J. Oper. Res., **192** (2009), n. 2, 500-511, https://doi.org/10.1016/j.ejor.2007.09.042.

A. Hilario-Caballero, A. Garcia-Bernabeu, J. V. Salcedo, M. Vercher: *Tri-Criterion Model for Constructing Low-Carbon Mutual Fund Portfolios: A Preference-Based Multiobjective Genetic Algorithm Approach*. Int. J. Environ. Res. Public Health, **17** (2020), n. 17, 6324, https://doi.org/10.3390/ijerph17176324.

M. Kaucic, M. Moradi, M. Mirzazadeh: *Portfolio optimisation by improved NSGA-II and SPEA 2 based on different risk measures*. Fin. Innov., **5** (2019), n. 26, 1-28, https://doi.org/10.1186/s40854-019-0140-6.

M. Kaucic, F. Piccotto, G. Sbaiz, G. Valentinuz: *A hybrid level-based learning swarm algorithm with mutation operator for solving large-scale cardinality-constrained portfolio optimisation problems*. Inf. Sci., **634** (2023), 321-339, https://doi.org/10.1016/j.ins.2023.03.115.

M. Kaucic, F. Piccotto, G. Sbaiz, G. Valentinuz: *Optimal Portfolio with Sustainable Attitudes under Cumulative Prospect Theory*. J. Appl. Fin. Bank., (2023b), accepted for publication.

H. Konno, H. Waki, A. Yuuki: *Portfolio Optimisation under Lower Partial Risk Measures*. Asia-Pac. Fin. Mark., **9** (2002), 127–140, https://doi.org/10.1023/A:1022238119491.

K. Liagkouras, K. Metaxiotis, G. Tsihrintzis: *Incorporating environmental and social considerations into the portfolio optimisation process*. Ann. Oper. Res., **316** (2022), 1493–1518, https://doi.org/10.1007/s10479-020-03554-3.

W. B. Lindquist, S. T. Rachev, Y. Hu, A. Shirvani: *"Advanced REIT Portfolio Optimization"*. Springer, New York (2022).

H. M. Markowitz: *Portfolio Selection*. J. Fin., **7** (1952), n. 1, 77–91, https://doi.org/10.2307/2975974.

H. M. Markowitz: *"Portfolio Selection: Efficient Diversification of Investments"*. Yale University Press (1959).

H. Markowitz, P. Todd, G. Xu, Y. Yamane: *Computation of mean-semivariance efficient sets by the Critical Line Algorithm*. Ann. Oper. Res., **45** (1993), 307–317, https://doi.org/10.1007/BF02282055.

H. Mausser, O. Romanko: *Long-only equal risk contribution portfolios for CVaR under discrete distributions*. Quant. Fin., **18** (2018), n. 11, 1927-1945, https://doi.org/10.1080/14697688.2018.1434317.

L. Mervelskemper, D. Streit: *Enhancing market valuation of ESG performance: Is integrated reporting keeping its promise?* Bus. Strat. Environ., **26** (2017), n. 4, 536–549, https://doi.org/10.1002/bse.1935.

G. Morelli: *Responsible investing and portfolio selection: a shapley-CVaR approach*. Ann. Oper. Res., (2023), 1-29, https://doi.org/10.1007/s10479-022-05144-x.

V. Pacelli, F. Pampurini, A. G. Quaranta: *Environmental, Social and Governance Investing: Does rating matter?* Bus. Strat. Environ., **32** (2023), n. 1, 30-41, https://doi.org/10.1002/bse.3116.

D. Politis, J. P. Romano: *The stationary bootstrap*. J. Am. Stat. Assoc., **89** (1994), n. 428, 1303-1313, https://doi.org/10.1080/01621459.1994.10476870.

D. Politis, H. White: *Automatic block-length selection for the dependent bootstrap*. Econom. Rev., **23** (2004), n. 1, 53-70, https://doi.org/10.1081/ETC-120028836.

S. T. Rachev, S. V. Stoyanov, F. J. Fabozzi: *"Advanced stochastic models, risk assessment, and portfolio optimisation: The ideal risk, uncertainty, and performance measures"*. Wiley (2008).

R. T. Rockafellar, S. Uryasev, M. Zabarankin: *Optimisation of conditional value-at-risk*. J. Risk, **2** (2002), n. 3, 21-41, https://doi.org/10.21314/JOR.2000.038.

R. T. Rockafellar, S. Uryasev, M. Zabarankin: *Generalised deviations in risk analysis*. Fin. Stoch., **10** (2006), 51-74, https://doi.org/10.1007/s00780-005-0165-8.

R. Santamaria, F. Paolone, N. Cucari, L. Dezi: *Non-financial strategy disclosure and environmental, social and governance score: Insight from a configurational approach*. Bus. Strat. Environ., **30** (2021), n. 4, 1993-2007, https://doi.org/10.1002/bse.2728.

A. B. Schmidt: *Optimal ESG portfolios: an example for the Dow Jones Index*. J. Sustain. Fin. Invest., **12** (2022), n. 2, 529-535, https://doi.org/10.1080/20430795.2020.1783180.

W. F. Sharpe: *Mutual fund performance*. J. Business, **39** (1966), 119-138, http://dx.doi.org/10.1086/294846.

P. Xidonas, E. Essner: *On ESG Portfolio Construction: A Multiobjective Optimisation Approach*. Comput. Econ., (2022), 1-25, https://doi.org/10.1007/s10614-022-10327-6.

4 Optimal portfolio with sustainable attitudes under cumulative prospect theory

Since its introduction, the cumulative prospect theory (shortly, CPT) developed by Kahneman and Tversky in 1979 has emerged as a significant approach for understanding the behavioral patterns of decision-makers in financial contexts.

Additionally, extreme recent events such as the COVID-19 pandemic and the Russian invasion of Ukraine have highlighted the importance of corporate social responsibility and sustainable principles. Consequently, the investment process is changing toward more ethical choices.

In this chapter, we extend the classical optimization framework under the cumulative prospect theory (CPT) in two directions. On the one hand, we consider an agent who maximizes a financial CPT- based value function preselecting the assets to be included in the portfolio based on their environmental, social, and governance (ESG) scores. On the other hand, we first introduce the concept of sustainable returns obtained from firms' ESG scores. Then, assuming that investors present the same decision-making behavior in approaching financial and sustainable returns, we develop a bi-objective model that optimizes financial and sustainable CPT-value functions simultaneously.

Numerical results obtained on an investable universe from the constituents of the STOXX Europe 600 show that introducing ESG information improves the portfolio's financial performance.

4.1 Introduction

The portfolio selection process usually involves two stages. The first consists in analyzing the most promising assets from a pool. The second focuses on the most desirable combination of the selected securities, ending with a portfolio choice. The first theoretical framework developed to support investors' decisions under uncertainty is Von Neumann and Morgenstern's representation theorem (Von Neumann and Morgenstern (1947)), which represents the cornerstone of the so-called expected utility theory. According to this result, the choices of a rational investor can be described by a utility function, which is concave since agents are usually assumed to be risk-averse. Furthermore, when choosing among alternative investments, agents select the ones that maximize their expected utility. In the same direction, (Markowitz (1952)) introduced the so-called mean-variance (MV) analysis, which has become the foundation of the modern portfolio theory. Following this principle, investors should maximize the expected return while retaining a given level of volatility. The MV framework can be derived as a particular case of the expected utility maximization when a quadratic utility function approximates the agent's behavior in the risk-return trade-off. Over the years, the expected utility theory has become the normative model of rational choice, and the MV approach has been widely accepted as a practical tool for portfolio optimization among researchers and practitioners (Guerard (2010)). However, several financial paradoxes have evidenced that, beyond risk aversion, some behavioral aspects deviate from the implications of expected utility theory (see, for instance, Allais (1953) and Ellsberg (1961)). To solve these puzzles, Kahneman and Tversky (1979) have developed the so-called prospect theory (PT), which incorporates human psychology into financial decisions and asserts that investors consider the deviations of their terminal wealth from a subjective reference point as gains and losses. Agents react differently for negative and positive outcomes, following a non-smooth, asymmetrical S-shaped value function concave in the domain of gains (risk aversion) and convex in the domain of losses (risk propensity). Later, Tversky and Kahneman (1992) refined the original PT formulation and developed the so-called cumulative prospect theory (CPT). This extension adds a transformation for the probability of the realizations of gains and losses, which overweights small probabilities and underweights higher ones. Recently, researchers have shown more interest in behavioral finance, applying the CPT theoretical framework to the asset allocation practice. (De Giorgi and Hens (2006)) have built up a theoretical analysis of an asset allocation problem where they have maximized an S-shaped utility function that generalizes the one introduced by Tversky and Kahneman. The resulting optimization problem is non-smooth, and calculating its solutions can be challenging. Thus, to overcome this problem, De Giorgi et al. (2007) have developed a robust technique to compute optimal CPT-based portfolios. Pirvu and Schulze (2012) have considered how PT-investors allocate their wealth in a single period between a riskless bond and multiple risky assets. Hens and Mayer (2017) have compared the

financial performance of the CPT-based asset allocation strategy to the classical mean-variance portfolio. Consigli et al. (2019) have performed a sensitivity analysis of the impact of CPT parameters on the mean-risk efficient frontier. It is worth mentioning that the above-quoted papers rely on restrictive assumptions on the number of assets and on the distribution of their returns to compute the solutions of the CPT-based portfolio optimization problems. To avoid these issues, several authors have used stochastic optimization algorithms that do not require any assumptions on the objective function to optimize. In particular, genetic algorithms (GAs) have proved efficient for solving portfolio selection problems. Some significant examples are the papers of Chang et al. (2000), Yang (2006), Sefiane and Benbouziane (2012), and Rankovic et al. (2014), which have used GAs to compute the optimal solutions of MV models. Moreover, GAs have also been adopted for solving PT-based portfolio problems. Grishina et al. (2017) have proposed several intelligent algorithms to obtain an accurate solution in optimization problems involving the original PT. Gong et al. (2018) extended this research by considering the re-weighted probabilities and introduced a hybrid optimization method combining a bootstrap technique and a genetic algorithm. Environmental, social, and governance (ESG) issues and non-financial disclosure have recently gained considerable attention. In 2015, the European Union introduced Directive 2014/95/EU to improve the quality of corporate non-financial disclosure. Therefore, political institutions, portfolio managers, and new generations of investors show an increasing sensitivity toward investment choices' impact on society and the environment. Debate on green and sustainable transitions in finance and investments is growing. Several studies have extended the traditional MV model by including a third criterion related to sustainability (see Hirschberger et al. (2013) and Utz et al. (2014)). Jessen (2012) has incorporated investors' ESG tastes into the Markowitz model by adding the weighted sum of the constituents' ESG scores as a constraint of the portfolio optimization problem. Pedersen et al. (2021) have computed the ESG-efficient frontier, displaying the highest possible Sharpe ratio for each ESG score. Then, they observed that ethical investors should be willing to accept lower portfolio returns in exchange for a more sustainable portfolio choice. Mimicking the stock-picking strategies used by portfolio managers to select the most promising assets based on their financial performance, Liagkouras et al. (2022) have developed a screening procedure to identify a subset of ESG-compliant stocks as constituents of an MV portfolio. Then, they tested this approach through a real-world application on the European market. In parallel with the above study, Yu et al. (2021) have selected the most promising assets according to their ESG score from the Chinese stock market to construct an ethical CPT-based portfolio. While these literature contributions build on deterministic quantities of sustainability (e.g., the ESG scores), Dorfleitner and Utz (2012) have introduced the concept of sustainable returns obtained from firms' ESG scores. An ethical investor would prefer positive sustainable returns rather than negative ones. Then, they implemented this novelty in the MV portfolio selection framework. Later, Dorfleitner and Nguyen (2016) integrated

the concept of sustainability returns into the utility theory setting. More specifically, they have proposed an asset allocation model maximizing a convex combination of two utility functions linked to financial and sustainable returns, respectively. There is a gap in the current literature concerning the extension of CPT-based portfolio models to account for investors' ESG attitudes. This paper aims to fill this issue. On the one hand, we propose two optimization models that extend the CPT-based portfolio selection framework. In the first model, we adopt several preselection techniques that identify a subset of ESG-compliant stocks based on their ESG scores. Then, we find the optimal portfolio weights according to the CPT-based asset allocation strategy. This approach generalizes the studies of Yu et al. (2021) and Liagkouras et al. (2022). In the second model, we account for investors who could be sensitive to the ESG scores and the evolution of these values over time. To formalize this insight, we propose a novel portfolio design that involves CPT-value functions for the financial and sustainable returns as objectives to maximize simultaneously. On the other hand, we present a real-world case study that involves a set of securities from the constituents of the STOXX Europe 600 index, covering the period from 2015 to 2021. The choice to use a European market index and to start the analysis in 2015 is coherent with the introduction of disclosure rules in Europe and the increase of their importance, especially for listed companies that have to follow specific and strong rules on non-financial disclosure. We perform an ex-post analysis with a rolling window approach that covers the period from 2019 to 2021, involving the recent phases of market downturns due to the COVID-19 outbreak. Following the literature, we exploit evolutionary computation to find the optimal portfolio weights for our models. In particular, for the bi-objective case, we use a variant of the GA suited for multi-objective portfolio problems, the so-called improved NSGA-II (shortly, iNSGA-II), proposed in Kaucic et al. (2019).

We organize this paper as follows. Section 4.2 introduces the theoretical framework and presents the proposed CPT-based asset allocation models. Section 4.3 describes the GAs used to solve the optimization problems, whereas Section 4.4 presents the experimental results of the out-of-sample analysis implemented on the European stock market. Section 4.5 concludes with some remarks and future research directions.

4.2 Theoretical framework

4.2.a Financial and sustainable returns

This study considers a frictionless financial market where investors act as price takers, and short selling is not allowed. The investable universe is represented by n risky assets, and a portfolio is denoted by the vector $\boldsymbol{x} = (x_1, \ldots, x_n)^\top$ of its asset weights. We observe the market over a time window of $T + 1$ months $\{0, 1, \ldots, T\}$. We denote by $p_{(i,t)}$ the observed price at time t of asset i, where $i = 1, \ldots, n$ and $t = 0, 1, \ldots, T$. Then, the realized rate of return at time t for the i-th stock is defined as follows:

$$r_{i,t} = \frac{p_{i,t}}{p_{i,t-1}} - 1 \tag{4.2.1}$$

where $i = 1, \ldots, n$ and $t = 1, \ldots, T$. Next, we define the portfolio rate of return at time t as $R_{p,t} = \sum_{i=1}^{n} r_{i,t}\, x_{i,t}$. Following the papers of Dorfleitner and Utz (2012) and Arslan and Posch (2022), we introduce a performance ratio that uses ESG information to obtain the sustainable counterpart of the financial rate of return. To this end, we consider the so-called relative sustainability value, which is calculated as the ratio between the company ESG score and the average ESG score of all companies in the current investable universe, that is:

$$s_{i,t} = \frac{ESG_{i,t}}{\overline{ESG_t}} \tag{4.2.2}$$

with $i = 1, \ldots, n$, $t = 0, \ldots, T$ and $\overline{ESG}_t = \frac{1}{n}\sum_{i=1}^{n} ESG_{i,t}$. This ratio represents how a company performs compared to the average in terms of ESG score. It is worth noting that there is no uniformity among companies in publishing non-financial reports, so firms' ESG scores are not updated on the same date in a given year. To guarantee a fair interpretation, we define the sustainability rate of return of firm i at time t by comparing the sustainability performance score in t with the one of the previous year, as follows:

$$sr_{i,t} = \frac{s_{i,t}}{s_{i,t-12}} - 1\,. \tag{4.2.3}$$

A company with a positive sustainability rate of return at time t reflects its capability to improve its relative value (4.2.2) in a year. This fact is an indicator of a successful implementation of an ESG-compliant business. Conversely, a negative value of (4.2.3) could indicate the mismatching of a company with the market standards of ESG practices. Similar to the financial case, we also introduce the portfolio sustainability rate of return at time t as:

$$sr_{p,t} = \sum_{i=1}^{n} sr_{i,t}\, x_{i,t}. \tag{4.2.4}$$

4.2.b Cumulative prospect theory

According to the PT approach developed by Kahneman and Tversky (1979), individual preferences are modeled through a value function of the aggregate gains and losses from a reference point $z_{ref} \in \mathbb{R}$, which is usually fixed at the beginning of the investment period. Thus, the investor decides the values of these outcomes according to the following piecewise value function:

$$v(z;\alpha,\beta)=\begin{cases}(z-z_{ref})^{\alpha}, & z\geq z_{ref}\\ -\beta(z_{ref}-z)^{\alpha}, & z<z_{ref}\end{cases} \tag{4.2.5}$$

with $z\in\mathbb{R}$, $0<\alpha<1$, and $\beta>0$. Specifically, they argued for distinct attitudes among traders toward gains, represented by $z-z_{ref}\geq 0$, and losses, given by $z-z_{ref}<0$.

The concavity of (4.2.5) in the domain of gains and its convexity in the domain of losses imply that investor behavior exhibits risk aversion towards gains and risk seeking towards losses. Two weighting functions are then introduced to modify the outcome probabilities to underweight higher outcomes and overweight smaller ones. Formally, in the domain of gains, we have:

$$w_{+}(u;\delta^{+})=\frac{u^{\delta^{+}}}{(u^{\delta^{+}}+(1-u)^{\delta^{+}})^{\frac{1}{\delta^{+}}}} \tag{4.2.6}$$

while in the domain of losses:

$$w_{-}(u;\delta^{-})=\frac{u^{\delta^{-}}}{(u^{\delta^{-}}+(1-u)^{\delta^{-}})^{\frac{1}{\delta^{-}}}} \tag{4.2.7}$$

where $u\in[0,1]$ is the probability of an outcome and $\delta^{+},\delta^{-}>0$. The scaling of these probabilities is based on the following procedure. Given a random outcome Z, let $\mathbf{z}=(z_1,\ldots,z_n)^{\top}\in\mathbb{R}^n$ be the vector of its realizations and let denote by N_{+} and N_{-} the number of positive and negative values in $\mathbf{z}$. We sort the elements of $\mathbf{z}$ in ascending order:

$$z_{(1)}\leq\cdots\leq z_{(N_{-})}<0\leq z_{(N_{-}+1)}\leq\cdots\leq z_{(N)}\,.$$

Then, we define the positive and negative re-weighted probabilities as:

$$\pi_{+,j}=\begin{cases}w_{+}\left(\frac{N_{+}-j+1}{N};\delta_{+}\right)-w_{+}\left(\frac{N_{+}-j}{N};\delta_{+}\right), & \text{if } j\in\{1,\ldots,N_{+}-1\}\\ w_{+}\left(\frac{1}{N};\delta_{+}\right), & \text{if } j=N_{+}\end{cases} \tag{4.2.8}$$

and

$$\pi_{-,j}=\begin{cases}w_{-}\left(\frac{N_{+}-j+1}{N};\delta_{-}\right)-w_{-}\left(\frac{N_{+}-j}{N};\delta_{-}\right), & \text{if } j\in\{1,\ldots,N_{+}-1\}\\ w_{-}\left(\frac{1}{N};\delta_{+}\right), & \text{if } j=N_{+}\,.\end{cases} \tag{4.2.9}$$

Hence, according to the CPT, the utility value of Z is the weighted average of the value function (4.2.5) calculated in the positive and negative outcomes using the re-weighted probabilities $\pi_{+,j}$ and $\pi_{-,j}$:

$$V(\mathbf{z})=\sum_{j=1}^{N_{-}}\pi_{-,j}\,v(z_{(j)};\alpha,\beta)+\sum_{h=1}^{N_{+}}\pi_{+,h}\,v(z_{(N_{-}+h)};\alpha,\beta)$$

In the experimental section, we will adopt the same parameter setting as in Kahneman and Tversky (1979) and Tversky and Kahneman (1992), that is $\alpha = 0.88$, $\beta = 2.25$ and $z_{ref} = 0$ for the value function (4.2.5), $\delta_+ = 0.61$, and $\delta_- = 0.68$ for the weighting functions (4.2.6) and (4.2.7), respectively. Figure 4.1 displays the graphs of these functions.

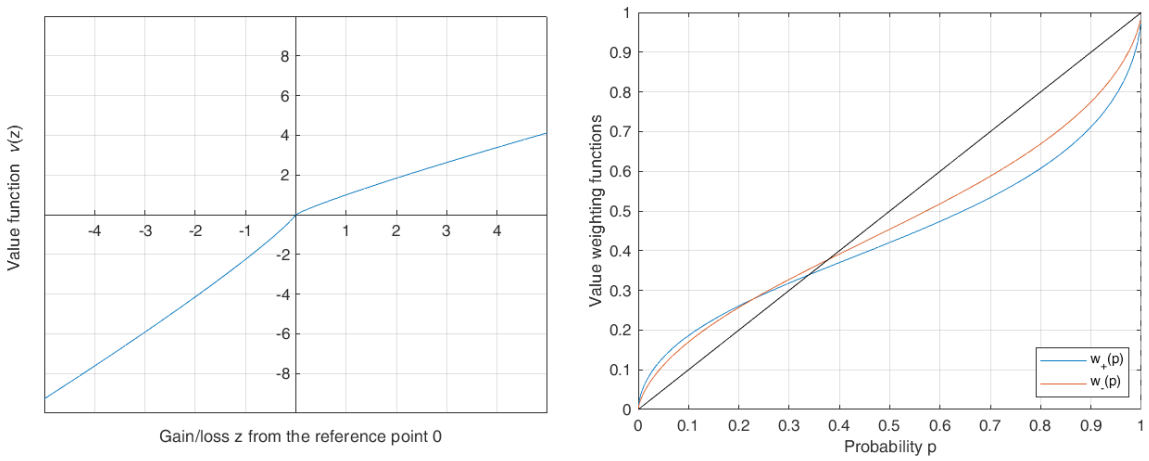

Figure 4.1: Cumulative prospect theory piecewise value function and weighting functions with parameters as in Kahneman and Tversky (1979) and Tversky and Kahneman (1992).

4.2.c Portfolio optimization problems under CPT and ethical principles

In this part, we first present the standard CPT-based portfolio optimization model. Then, we describe the two extensions we propose to include ESG criteria.

We consider the T observed portfolio rates of return $\{r_{p,t}\}_{(t=1,\dots,T)}$ as realizations of the random variable R_p, which expresses the stochastic rate of return of portfolio $\boldsymbol{x}$ with a monthly investment horizon. Following the steps described in the previous section, we calculate the vector of the re-weighted probabilities of the realized rates of return $\pi^{(r)} = \{\pi_1^{(r)}, \dots, \pi_T^{(r)}\}$. Then, the standard CPT-based portfolio optimization problem can be written as:

$$\begin{aligned} \max_{\boldsymbol{x}\in\mathbb{R}^n} \quad & \sum_{t=1}^{T} \pi_t^{(r)}\, v\left(\sum_{i=1}^{n} x_i r_{i,t}; \alpha, \beta\right) \\ \text{s.t.} \quad & x_i \geq 0, \quad i = 1, \dots, n \\ & \sum_{i=1}^{n} x_i = 1\,. \end{aligned} \tag{4.2.10}$$

Problem (4.2.10) does not explicitly account for the attitude of investors towards firms' sustainable targets. To overcome this issue, we present two novel asset allocation

strategies that directly exploit ESG scores. In the first model, we consider a CPT-based investor who preselects a subset of promising assets using ESG information to solve problem (4.2.10). The other model extends the previous one by including, besides the preselection based on ESG scores, a second CPT-value function for the sustainability rates of return, as an additional objective to maximize, in order to take advantage of the evolution of ESG scores over time.

CPT-based model with preselection using ESG scores

To include ethical goals in the asset allocation process, we adopt a stock-picking technique to preselect the subset of assets that go beyond a given ESG threshold $\overline{\theta}$. We denote by $U_T(\overline{\theta})$ the subset which includes k out of the n available assets with an ESG score higher than the threshold $\overline{\theta}$ at the end of the observed period T. Then, the portfolio selection model can be reformulated as follows:

$$\begin{aligned} \max_{\boldsymbol{x}\in\mathcal{X}} \quad & V_1(\boldsymbol{x}) = \sum_{t=1}^{T} \pi_t^{(r)}\, v\left(\sum_{i\in U_T(\overline{\theta})} x_i r_{i,t}; \alpha, \beta\right) \\ \text{s.t.} \quad & x_i \geq 0, \quad i \in U_T(\overline{\theta}) \\ & \sum_{i\in U_T(\overline{\theta})} x_i = 1\,. \end{aligned} \tag{4.2.11}$$

where $\mathcal{X} = \{\boldsymbol{x} \in \mathbb{R}^k : i \in U_T(\overline{\theta})\}$. If we do not adopt a preselection, problem (4.2.11) is equivalent to (4.2.10).

CPT-based model with financial and ethical objectives

Since the variations of the firms' ESG scores over time could suggest the company's capability to implement an ESG-compliant business, an ESG-aware investor may also be sensitive to this information. Assuming that investors present the same decision-making behavior in approaching financial and sustainable returns, we define the following CPT-type value function for the sustainability rates of return:

$$V_2(\boldsymbol{x}) = \sum_{t=1}^{T} \pi_t^{(s)}\, v\left(\sum_{i\in U_T(\overline{\theta})} x_i sr_{i,t}; \alpha, \beta\right) \tag{4.2.12}$$

where $sr_{i,t}$ is the sustainability rate of return of asset $i \in U_T(\overline{\theta})$ at time t and $\pi^{(s)} = \{\pi_1^{(s)}, \ldots, \pi_T^{(s)}\}$ is the vector of the re-weighted probabilities associated with the sustainability rates of return. In this extended formulation of the portfolio optimization problem, the agent simultaneously maximizes the value functions V_1 and V_2 with a preliminary ESG-based screening of the constituents. The resulting bi-objective optimization problem is:

$$\begin{aligned} \max_{\boldsymbol{x}\in\mathcal{X}} \quad & (V_1(\boldsymbol{x}), V_2(\boldsymbol{x})) \\ \text{s.t.} \quad & x_i \geq 0, \quad i \in U_T(\bar{\theta}) \\ & \sum_{i\in U_T(\bar{\theta})} x_i = 1\,. \end{aligned} \tag{4.2.13}$$

In such a setting, portfolio choices are made according to the following preference relation.

Definition 12. *Given two portfolios $\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}$, we say that $\boldsymbol{x}^{(1)}$ is preferred to (or non-dominated by) $\boldsymbol{x}^{(2)}$ if and only if $V_1(\boldsymbol{x}^{(1)}) \geq V_1(\boldsymbol{x}^{(2)})$ and $V_2(\boldsymbol{x}^{(1)}) \geq V_2(\boldsymbol{x}^{(2)})$, with at least one strict inequality.*

In this manner, an agent will prefer portfolios with higher financial and sustainable values. The set of non-dominated solutions to problem (4.2.13) forms the so-called Pareto front or efficient frontier, denoted by $\mathcal{P}$. At this point, we can choose on $\mathcal{P}$ the portfolio which best-fit investor's attitude toward financial and sustainable objectives as follows. First, the values of the two objective functions for each candidate portfolio $\boldsymbol{x}$ in $\mathcal{P}$ are normalized using the formula:

$$V_j^{(n)}(\boldsymbol{x}) = \frac{V_j(\boldsymbol{x}) - V_j^{min}}{V_j^{max} - V_j^{min}}$$

where $V_j^{min} = \min_{\boldsymbol{x}\in\mathcal{P}} V_j(\boldsymbol{x})$ and $V_j^{max} = \min_{\boldsymbol{x}\in\mathcal{P}} V_j(\boldsymbol{x})$. Next, we define the vector of preferences $\mathbf{w}_{pref} = (w_{ESG}, w_{fin})^\top \in [0,1]^2$ such that $w_{ESG} + w_{fin} = 1$. Finally, we calculate the weighted sum of the objective functions and maximize the result on the set of efficient portfolios:

$$\max_{\boldsymbol{x}\in\mathcal{P}} \left\{ w_{fin} V_1^{(n)}(\boldsymbol{x}) + w_{ESG} V_2^{(n)}(\boldsymbol{x}) \right\}. \tag{4.2.14}$$

The solution to this single-objective problem identifies an efficient portfolio tailored to the investor's financial/ESG profile.

4.3 Description of the optimizers

In this paper, we consider two evolutionary algorithms to solve the above optimization problems, which are nonlinear and non-smooth. More specifically, to tackle the single-objective problems (4.2.10) and (4.2.11), we employ a genetic algorithm (GA), while for the bi-objective one (4.2.13), we implement the iNSGA-II developed by Kaucic et al. (2019). In Figure 4.2, we report the flowcharts of the two solvers. In both cases, an initial set of candidate solutions, forming the so-called population, is randomly initialized and is evolved through the following three steps.

- Two sub-populations of parents are created by uniform selection. The first sub-population is involved in the recombination phase, while the second is subject to the mutation strategy. It is worth noting that the two sub-groups are not disjointed. In this way, the same individual may enter into both the crossover step and the mutation stage.
- The crossover is applied to the first sub-population of parents to give rise to a set of offspring. This operator produces two children $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{R}^n$ for each pair of parents $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{R}^n$:

$$\begin{aligned} \overline{x}_{1i} &= c_i x_{1i} - (1 - c_i) x_{2i} \\ \overline{x}_{2i} &= c_i x_{2i} - (1 - c_i) x_{1i} \end{aligned} \tag{4.3.15}$$

  where c_i is the scaling factor randomly chosen in $[-1, 2]$ with $i = 1, \ldots, n$. A Gaussian mutation modifies the second sub-group of parents.
- The sub-populations of children are pickled together to forge the offspring population.

Due to the requirement of the entire investment, we also introduce a constraint-handling procedure based on the repair mechanism developed by Liagkouras and Metaxiotis (2015). First, in order to satisfy the non-negativity constraint, each candidate solution $\boldsymbol{x} \in \mathbb{R}^n$ is clamped in $[0, 1]^n$ as:

$$\tilde{x}_i = \begin{cases} 0 & \text{if } x_i < 0 \\ 1 & \text{if } x_i > 1 \\ x_i & \text{otherwise} \end{cases} \tag{4.3.16}$$

where $i = 1, \ldots, n$. Then, the projected vector $\tilde{\boldsymbol{x}}$ is adjusted by normalization to satisfy the budget constraint in the following way:

$$\tilde{\tilde{\boldsymbol{x}}} = \frac{\tilde{x}_i}{\sum_{j=1}^{n} \tilde{x}_j} \text{ for } i = 1, \ldots, n\,. \tag{4.3.17}$$

The above-quoted procedure makes feasible the individuals in the search space. The main differences between the two algorithms are the following. In the GA, the best individual at each generation represents a sub-optimal solution to the problem under consideration. Instead, the iNSGA-II algorithm generates good approximations of the Pareto front. To this end, iNSGA-II exploits a two-step ranking scheme. The first ranking is based on the non-domination relation given in Definition 12. If the individuals have the same position in that stage, one applies the second-level ranking, which exploits a diversity-preserving mechanism, the so-called crowding distance. The elitist individuals are chosen from the current population, and the offspring set based on the ranking scheme. The candidates, saved in the following population, corresponding to the higher level of non-domination fronts and the higher values of crowding distance.

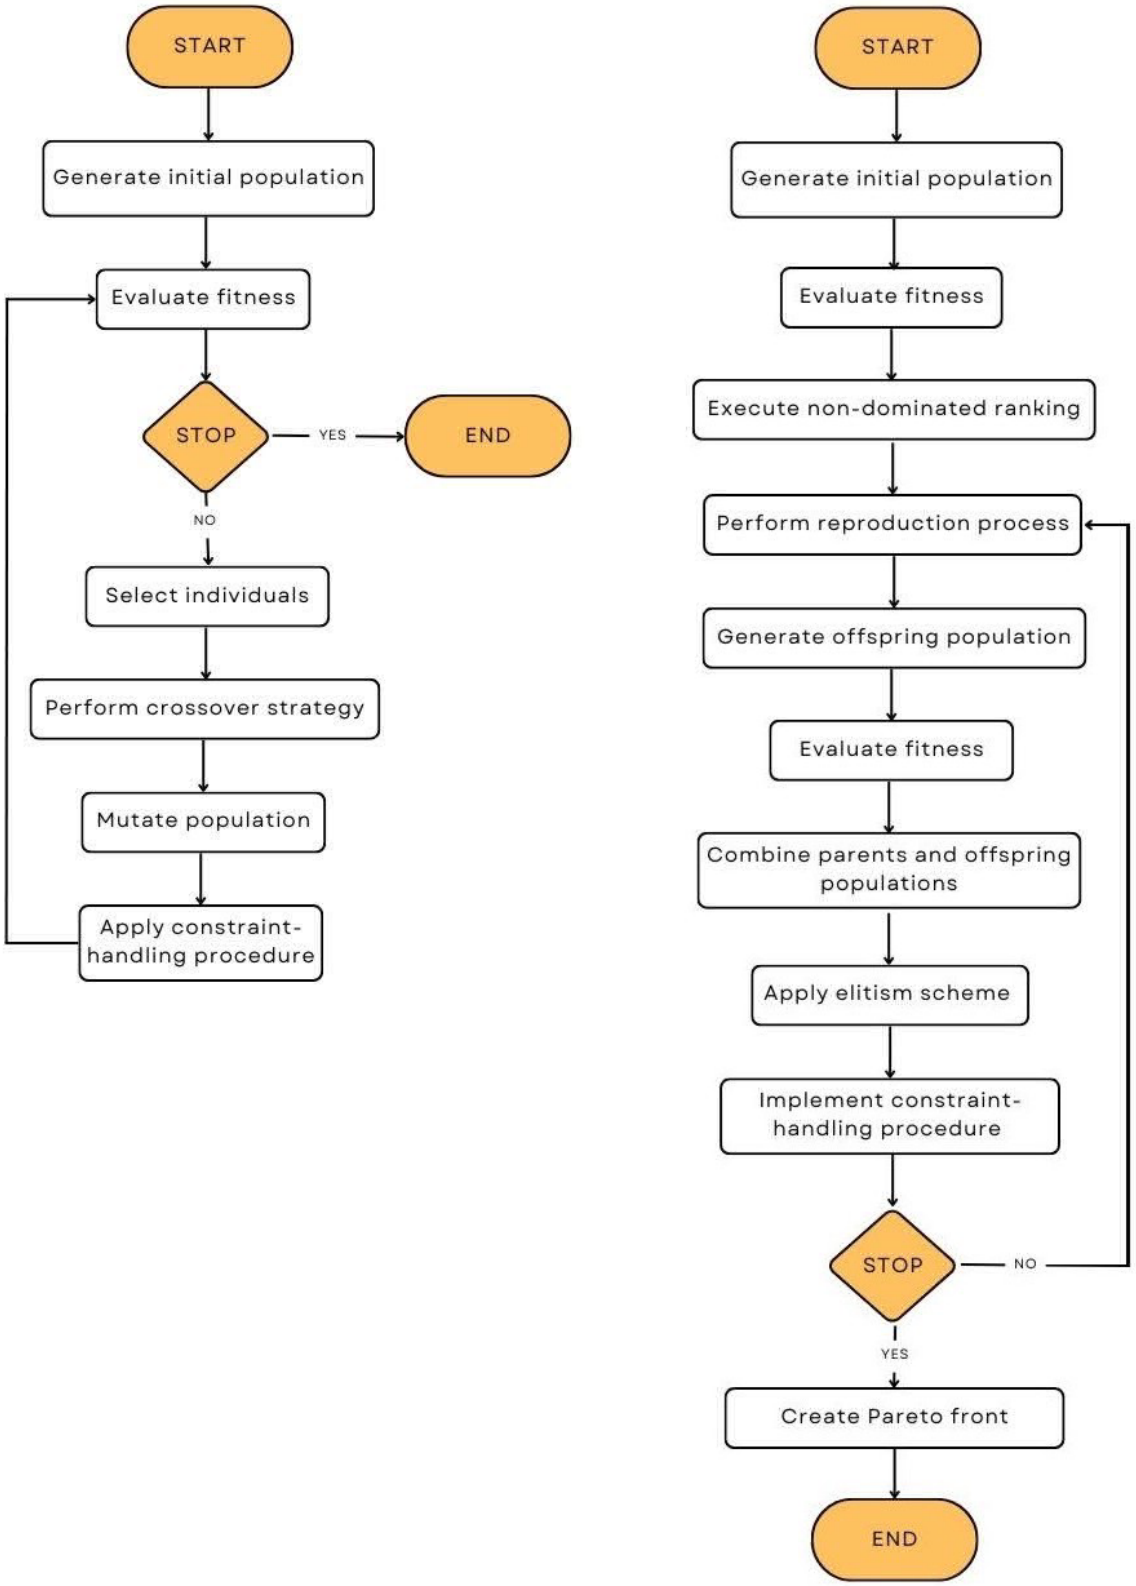

Figure 4.2: Flowcharts of the two GAs used in the paper. On the left, the GA for the single-objective problems, and on the right, the iNSGA-II for the bi-objective problem.

4.4 Experimental part

Our analysis focuses on the European market due to the advanced disclosure rules in this zone. More specifically, we consider a subset of the STOXX Europe 600 index constituents as the investable universe. We obtain the monthly rates of return and the monthly ESG ratings from 01/01/2015 to 31/12/2021 for 538 companies from Refinitiv Datastream. The ESG scores are scaled between 1 and 100 and represent an assessment of how the company operates on sustainability. The time window choice is motivated by Directive 2014/95/EU, which has imposed the disclosure of non-financial information for this type of company since 2015. In the first instance, our goal is to highlight the role in portfolio choices of a preselection technique based on the sustainable score of firms. In the second step, we analyze the benefits that derive from the inclusion in the investment process of a criterion involving the dynamics of the ESG scores. To this end, we consider an investment plan with a one-month horizon and employ a rolling window scheme based on historical data. The out-of-sample window covers the period from 31/11/2019 to 31/12/2021, for a total of 25 months. The procedure also exploits an in-sample window of 47 months to set the model parameters.

4.4.a Ex-post performance measures

In the sequel, we introduce the ex-post performance measures that will be used to assess the profitability of the proposed strategies. Let us denote by $\boldsymbol{x}_t$ the optimal portfolio at the ex-post month t, with $t = 1, \dots, 25$. Due to the time dependence of the considered investment plan, we calculate the value of turnover as follows:

$$\sum_{i=1}^{n} |\boldsymbol{x}_{t,i} - \boldsymbol{x}_{t-1,i}| \tag{4.4.18}$$

where $\boldsymbol{x}_{t-1,i} = (x_{t-1,1}, \dots, x_{t-1,n})$ represents the portfolio to be balanced. Now, let the $r_{p,t}^{out}$ be the ex-post portfolio rate of return realized at time t, with $t = 1, \dots, 25$. We consider the so-called ex-post Sharpe ratio, defined as:

$$SR^{out} = \frac{\mu^{out}}{\sigma^{out}} \tag{4.4.19}$$

where μ^{out} and σ^{out} are the mean and the standard deviation of the ex-post portfolio rates of return, respectively. Further, we compute the net wealth at time t as:

$$W_t = W_{t-1}(1 + r_{p,t}^{out}) - \lambda(\boldsymbol{x}_t, \boldsymbol{x}_{t-1}) \tag{4.4.20}$$

where $\lambda(\boldsymbol{x}_t, \boldsymbol{x}_{t-1})$ is the cost function, whose structure is presented in Table 4.1.

Moreover, we use the so-called compound annual growth rate, which in our case is calculated as:

Trading segment (€)	Fixed fee (€)	Proportional cost (%)
0 – 7999	40	0
8000 – 49999	0	0.50
50000 – 99999	0	0.40
100000 – 199999	0	0.25
$\geq$ 200000	400	0

Table 4.1: Structure of transaction costs.

$$CAGR = \left(\frac{W_{T_{end}}}{W_0}\right)^{\frac{12}{T_{end}}} - 1 \tag{4.4.21}$$

where W_0 represents the initial wealth and $W_{T_{end}}$ is the final wealth. Finally, we measure the downside risk by the so-called peak-to-valley drawdown:

$$DD_t = -\min\{0, \frac{W_t - W_{peak}}{W_{peak}}\} \tag{4.4.22}$$

where W_{peak} is the maximum amount of wealth reached by the strategy until time t. In particular, we consider the maximum of the drawdowns over time.

4.4.b Analysis of the proposed models

The preselection strategies involved at the time of investment consider six levels of sustainable-worthiness, as reported in Table 4.2. In particular, the NO ESG strategy, without requiring any ESG information, represents the entire investable universe. Conversely, the second strategy requires only the availability of the ESG score at the investment time for the assets to be included in the portfolio. Strategies B and B+ consist of a pool of companies with good ESG performance referring to their sector and an above-average degree of transparency in publicly reporting ESG material. Finally, the last two stock-picking levels translate excellent ESG performances of a firm compared to its sector and a maximum degree of transparency in disclosing sustainable materials.

Stock-picking	Description
NO ESG	All firms in the dataset
ANY ESG	Firms with an ESG score
B	Firms with ESG score ≥ 58.3
B+	Firms with ESG score ≥ 66.7
A−	Firms with ESG score ≥ 75
A	Firms with ESG score ≥ 83.3

Table 4.2: Preselection strategies description.

Financial-only CPT-based models

In models (4.2.10) and (4.2.11), agents first preselect the investment pool based on the strategies introduced above. Then, they maximize the financial CPT-value function V_1. In Figure 4.3, we can observe the evolution of this objective for the different stock-picking techniques. When we consider NO ESG and ANY ESG, we notice that the corresponding optimal portfolios generate similar CPT-values. The other screening procedures present significantly lower values. Overall, according to the literature, it is worth noting that when the ESG threshold increases, the financial utility progressively decreases.

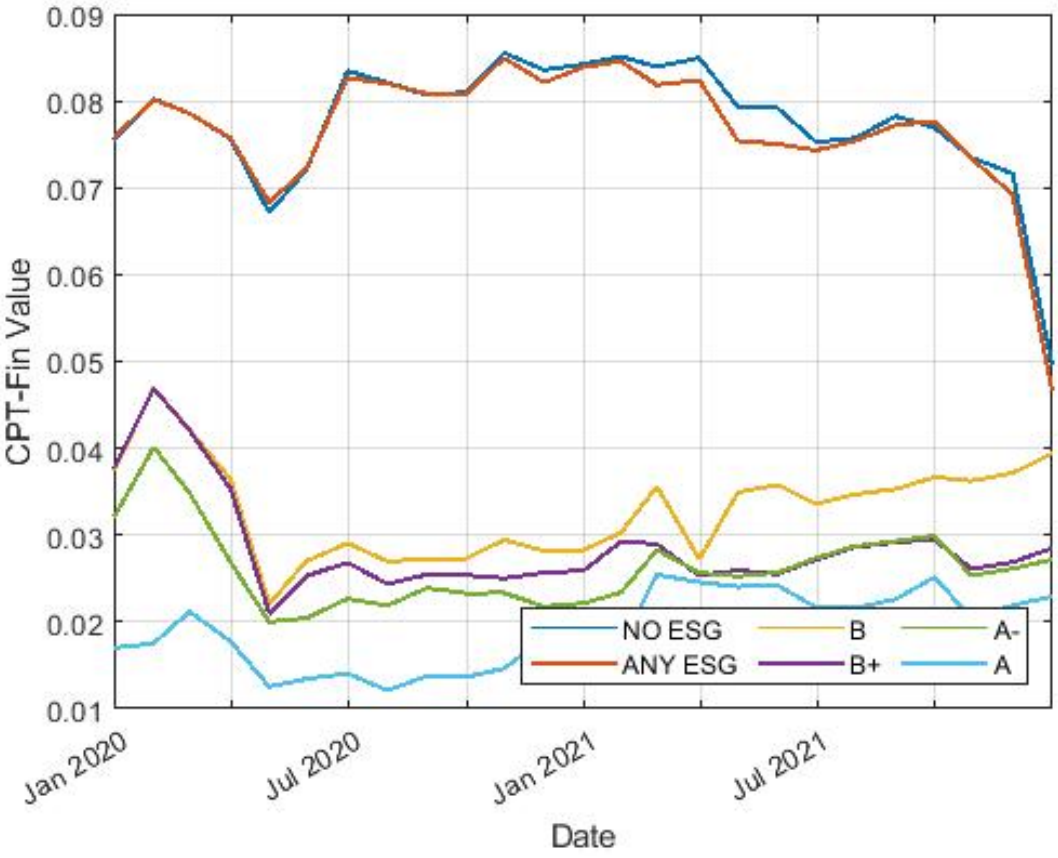

Figure 4.3: Evolution of the financial CPT-value function over time.

Next, we assess the profitability of the different portfolio allocation strategies with an initial wealth $W_0 = 1,000,000$€. Looking at Table 4.3, one can infer that including an ESG preference for the investable pool is a positive discriminant for a CPT-type agent investing in the European market. More specifically, the best strategy is associated with the ANY ESG screening, which presents a Sharpe ratio three times more than the benchmark with controlled maximum drawdown and turnover. The portfolios derived from NO ESG and A screenings are the second-best alternatives. The former shows a higher SR value, while the latter controls better risk and turnover.

	NO ESG	ANY ESG	B	B+	A-	A	Benchmark
SR	0.3130	0.4501	0.2914	0.1639	0.1317	0.2100	0.1638
CAGR (%)	6.5549	28.2080	-7.1879	-12.8220	-8.1606	4.1181	9.0256
Std	0.0978	0.0981	0.0651	0.0595	0.0553	0.0522	0.0539
Max DD	0.2922	0.2958	0.2910	0.2812	0.2480	0.2014	0.2303
Turnover	0.2891	0.2735	0.7386	0.7121	0.5976	0.5379	-

Table 4.3: Performance measures of portfolio allocation models (4.2.10) and (4.2.11) for the different stock-picking strategies.

Figure 4.4 highlights the behavior of the six investment strategies in terms of the produced wealth. The mere inclusion of the ESG criterion leads to an increase in investment, especially after 2021, if we do not consider transaction costs. However, focusing on the net wealth, only portfolios with NO and ANY ESG screenings still outperform the index. This fact is a consequence of the low turnover level of these portfolios.

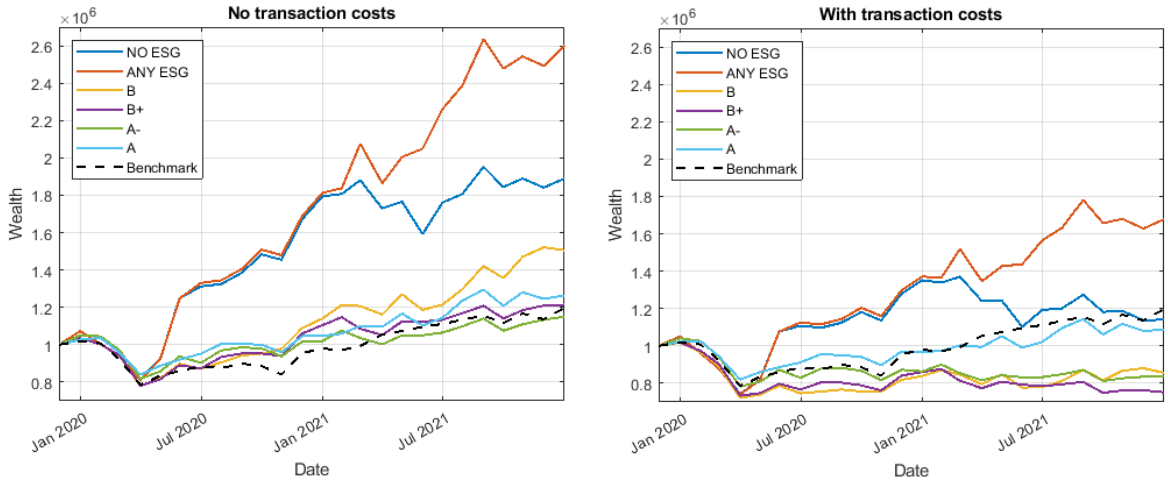

Figure 4.4: Ex-post evolution of gross wealth (on the left) and net wealth (on the right) of the proposed stock-picking strategies.

CPT-based models with financial and ethical objectives

Concerning the bi-objective model (4.2.13), Figure 4.5 shows for each preselection strategy the corresponding efficient frontier at the end of each year in the out-of-sample window. It is worth noting that the non-dominated fronts can be ranked in decreasing order with respect to the ESG threshold of the screening. In particular, we can observe the clustering of the efficient frontiers into three groups: ANY ESG, B and B+, A- and A.

Similar to the single-objective instances, we assess the ex-post profitability of the portfolio allocation strategies for three representative investor attitudes, whose vectors of preferences are listed in Table 4.4. The first agent gives more importance to the

financial criterion, while Agent 3 to the sustainable one. Finally, Agent 2 reflects a moderate profile and perfectly balances financial and ESG preferences.

	Weight of the financial preference	**Weight of the ESG preference**
Agent 1	0.75	0.25
Agent 2	0.50	0.50
Agent 3	0.25	0.75

Table 4.4: Structure of the weights for the three considered agents.

The ex-post results for these ESG-aware agents are reported in Table 4.5. As in the previous section, we consider an initial wealth $W_0 = 1,000,000$€, with the same transaction cost structure presented in Table 4.1. Each column refers to a specific ESG-based preselection. We observe that as the relevance of the CPT-type value function for the sustainability rates of return increases, the preselection threshold decreases to allow a broader range of ESG variations over time. Conversely, if there is a greater financial preference, the optimal ethical screening is attained at a higher ESG threshold.

We conclude this section by comparing the models presented above. In Figure 6, the best-performing model in terms of final wealth corresponds to the financial-only investor with the ANY ESG preselection criterion as it reaches $1,700,000$€. The bi-objective model for Agent 3 profile and preselection B+ is the second-best strategy having as final wealth around $1,300,000$€. We also stress the fact that for the same period, the buy-and-hold strategy finally shows $1,200,000$€as net wealth. Finally, the strategy that envisages no inclusion of ESG criteria underperforms the buy-and-hold strategy with a final wealth of $1,140,000$€.

4.5 Conclusions

In this paper, we have presented two extensions of the classical cumulative prospect theory framework. In the first model, we analyzed the impact of various ethical stock-picking strategies on the financial CPT-value function. In a second formulation, we have proposed a sustainable CPT-value function, which has been integrated into the objectives of the optimization problem. The aim has been to also include in the valuation framework of an ESG-aware investor the variations of the firms' ESG scores over time. We have introduced a scaling procedure based on the investor's preferences toward financial and ESG criteria to select the suitable portfolio on the associated efficient frontier. Finally, we have assessed the profitability of the developed portfolio allocation models using an investment pool from the STOXX Europe 600 covering 2019 to 2021. The results show that using a stock-picking procedure based on the ESG information improves the financial performance of the investment. Moreover, we have considered the portfolio construction for three representative investor attitudes. The ex-post analysis reveals that the bi-objective models with higher ESG preferences have

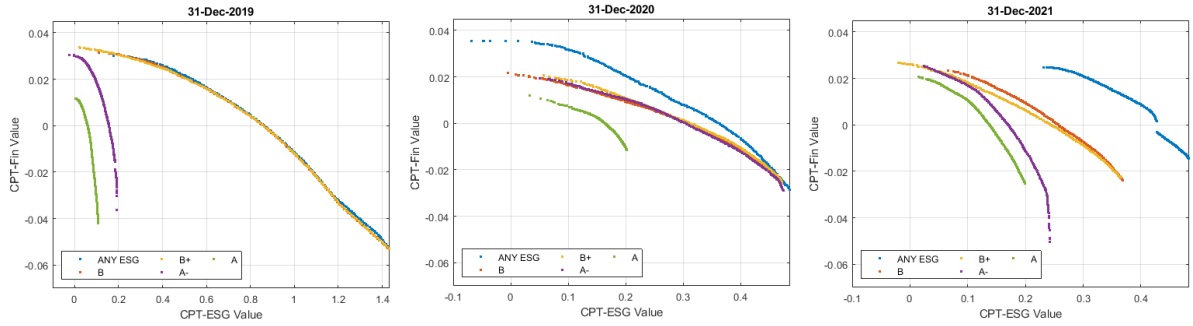

Figure 4.5: Efficient frontiers of model (4.2.13) at the end of each year.

	ANY ESG	B	B+	A-	A
Agent 1					
SR	-0.0399	0.0567	0.2218	0.1652	0.3592
CAGR (%)	-37.9350	-23.4090	-10.3210	-4.5579	10.5370
Std	0.0552	0.0444	0.0470	0.0573	0.0408
Max DD	0.6595	0.4263	0.2616	0.1882	0.1041
Turnover	0.7305	0.9666	0.8976	0.7693	
Agent 2					
SR	0.1073	0.2117	0.3737	0.2232	0.2090
CAGR (%)	-24.1850	-11.7920	3.6240	3.3902	4.3580
Std	0.0458	0.0471	0.0521	0.0672	0.0498
Max DD	0.4383	0.2330	0.1672	0.1815	0.1753
Turnover	0.7053	0.6048	0.5100	0.5462	0.3957
Agent 3					
SR	0.3323	0.3308	0.3368	0.2978	0.0944
CAGR (%)	5.6427	9.1423	12.8160	10.3200	-1.8447
Std	0.0703	0.0712	0.0711	0.0646	0.0560
Max DD	0.1645	0.1483	0.1471	0.1457	0.1810
Turnover	0.2102	0.2001	0.1632	0.3874	0.2648

Table 4.5: Performance measures of the proposed portfolio optimization for the three representative agents.

lower risk and better control of extreme losses with respect to the buy-and-hold strategy and the single-objective CPT-based models. However, these findings could depend on the selected provider's ESG rating evaluation methodology and the geographical area of the stocks employed in the investments. These factors offer an opportunity to develop further research by considering other ESG providers and different international markets.

Bibliography

M. Allais: *The behavior of the rational man before risk: criticism of postulates and axioms of the American school*. Econometrica, **21** (1953), 503-543, https://doi.org/10.2307/1907921.

A. Arslan, P. N. Posch: *How do investors value sustainability? A utility-based preference optimization*. Sustainability, **14** (2022), n. 23, https://doi.org/10.3390/su142315963.

T. J. Chang, N. Meade, J. E. Beasley, Y. M. Sharaiha: *Heuristics for cardinality constrained portfolio optimisation*. Comput. Oper. Res., **27** (2000), n. 13, 1271-1302, https://doi.org/10.1016/S0305-0548(99)00074-X.

G. Consigli, A. Hitaj, E. Mastrogiacomo: *Portfolio choice under cumulative prospect theory: sensitivity analysis and an empirical study*. Comput. Manag. Sci., **16** (2019), 129–154, https://doi.org/10.1007/s10287-018-0333-x.

G. Dorfleitner, M. Nguyen: *A new approach for optimizing responsible investments dependently on the initial wealth*. J. Asset Manag., **18** (2016), 81-98, https://doi.org/10.1057/s41260-016-0011-x.

G. Dorfleitner, S. Utz: *Safety first portfolio choice based on financial and sustainability returns*. Eur. J. Oper. Res., **221** (2012), n. 1, 155–164, https://doi.org/10.1016/j.ejor.2012.02.034.

E. De Giorgi, T. Hens: *Making prospect theory fit for finance*. Fin. Mark. Portf. Manag., **20** (2006), 339-360, https://doi.org/10.1007/s11408-006-0019-1.

E. De Giorgi, T. Hens, J. Mayer: *Computational aspects of prospect theory with asset pricing applications*. Comput. Econ., **29** (2007), 267–281, https://doi.org/10.1007/s10614-006-9062-2.

D. Ellsberg: *Risk, ambiguity, and the savage axioms*. Q. J. Econ., (1961), 643-669, https://doi.org/10.2307/1884324.

C. Gong, C. Xu, J. Wang: *An efficient adaptive real coded genetic algorithm to solve the portfolio choice problem under cumulative prospect theory*. Comput. Econ., **52** (2018), 227-252, https://doi.org/10.1007/s10614-017-9669-5.

N. Grishina, C. A. Lucas, P. Date: *Prospect theory-based portfolio optimization: an empirical study and analysis using intelligent algorithms*. Quant. Fin., **17** (2017), 1-15, https://doi.org/10.1080/14697688.2016.1149611.

J. B. Guerard: *"Handbook of portfolio construction. Contemporary applications of Markowitz techniques"*. Springer, New York (2010).

T. Hens, J. Mayer: *Cumulative prospect theory and mean-variance analysis: a rigorous comparison*. J. Comput. Fin., **21** (2017), n. 3, 47–73, http://dx.doi.org/10.2139/ssrn.2417191.

M. Hirschberger, R. Steuer, S. Utz, M. Wimmer, Y. Qi: *Computing the non-dominated surface in tri-criterion portfolio selection*. Oper. Res., **61** (2013), 169-183, https://doi.org/10.1287/opre.1120.1140.

P. Jessen: *Optimal responsible investment*. Appl. Fin. Econ., **22** (2012), n. 21, 1827-1840, https://doi.org/10.1080/09603107.2012.684786.

D. Kahneman, A. Tversky: *Prospect theory: an analysis of decision under risk*. Econometrica, **47** (1979), 263–291, https://doi.org/10.2307/1914185.

M. Kaucic, M. Moradi, M. Mirzazadeh: *Portfolio optimisation by improved NSGA-II and SPEA 2 based on different risk measures*. Fin. Innov., **5** (2019), n. 26, 1-28, https://doi.org/10.1186/s40854-019-0140-6.

K. Liagkouras, K. Metaxiotis: *Efficient portfolio construction with the use of multiobjective evolutionary algorithms: best practices and performance metrics*. Int. J. Inf. Technol. Decis. Making, **14** (2015), n. 3, 535-564, https://doi.org/10.1142/S0219622015300013.

K. Liagkouras, K. Metaxiotis, G. Tsihrintzis: *Incorporating environmental and social considerations into the portfolio optimisation process*. Ann. Oper. Res., **316** (2022), 1493–1518, https://doi.org/10.1007/s10479-020-03554-3.

H. M. Markowitz: *Portfolio Selection*. J. Fin., **7** (1952), n. 1, 77–91, https://doi.org/10.2307/2975974.

T. A. Pirvu, K. Schulze: *Multi-stock portfolio optimization under prospect theory*. Math. Fin. Econ., **6** (2012), 337–362, https://doi.org/10.1007/s11579-012-0079-0.

L. Pedersen, S. Fitzgibbons, L. Pomorski: *Responsible investing: the ESG-efficient frontier*. J. Fin. Econ., **142** (2021), 572–597, https://doi.org/10.1016/j.jfineco.2020.11.001.

V. Rankovic, M. Drenovak, B. Stojanovic, Z. Kalinic, Z. Arsovski: *The mean-value at risk static portfolio optimization using genetic algorithm*. Comput. Sci. Inf. Syst., **11** (2014), n. 1, 89-109, https://doi.org/10.2298/CSIS121024017R.

S. Sefiane, M. Benbouziane: *Portfolio selection using genetic algorithm*. J. Appl. Fin. Bank., **2** (2012), n. 4, 143-154, https://www.scienpress.com/Upload/JAFB/Vol%202_4_9.pdf.

A. Tversky, D. Kahneman: *Advances in prospect theory: cumulative representation of uncertainty*. J. Risk Uncertain., **5** (1992), n. 4, 297-323, https://doi.org/10.1007/BF00122574.

S. Utz, M. Wimmer, M. Hirschberger, R. Steuer: *Tri-criterion inverse portfolio optimization with application to socially responsible mutual funds*. Eur. J. Oper. Res., **234** (2014), 491-498, https://doi.org/10.1016/j.ejor.2013.07.024.

J. Von Neumann, O. Morgenstern: *Theory of games and economic behavior*. Princeton University Press Ltd, (1947).

X. Yang: *Improving portfolio efficiency: a genetic algorithm approach*. Comput. Econ., **28** (2006), n. 1, 1-14, https://doi.org/10.1007/s10614-006-9021-y.

W. Yu, S. Liu, L. Ding: *Efficiency evaluation and selection strategies for green portfolios under different risk appetites*. Sustainability, **13** (2021), n. 4, https://doi.org/10.3390/su13041933.

Conclusions

In this final section, I will briefly discuss the main findings of the works included in this PhD Thesis. In the first paper, the introduction of a hybrid Level-Based Learning Swarm Optimizer (LLSO) with mutation operators addresses challenges in a large-scale portfolio optimization model that maximizes a modified version of the Sharpe ratio. In the numerical experiments conducted on three publicly available data sets, on the one hand, I showed the superior performance of the proposed algorithm with respect to other state-of-the-art swarm optimization solvers in solving large-scale cardinality-constrained problems. On the other hand, an analysis of the portfolio model's sensitivity to the cardinality constraint using data from the MSCI World index reveals that smaller portfolios are more competitive than the value-weighted benchmark index and exhibit greater resilience during market downturns. The second paper is an analysis of three long-run investment strategies based on Sharpe-based performance measures on large-scale global market indices. Specifically, I considered some extensions of the classical Sharpe ratio that incorporate higher-order moments, in order to enhance the value of considering skewness and kurtosis in portfolio performance evaluation. Using the LLSO algorithm developed in the first work equipped with an ad-hoc constraint handling technique, I first conducted a sensitivity analysis for the portfolio size and for the turnover constraint, in order to identify the best combination of these parameters in terms of ex-post performance and management cost. Overall, the best portfolio combination turned out to be the one with the 15% of the investment pool and with a turnover of 10%. Then, I conducted a more detailed analysis of the portfolio performance before and after the COVID-19 pandemic outbreak, evidencing that including higher-order moments in the Sharpe ratio measure provides superior results in terms of net wealth and leads to a better control of risk. Finally, in the third and fourth paper, the incorporation of ESG criteria into portfolio optimization models reveals that sustainable investing can achieve competitive financial returns while enhancing risk control during market turbulence. Moreover, behavioral finance models further emphasize the alignment of ethical investment strategies with investor preferences.

I believe that the results proposed in this PhD thesis could bridge theoretical advancements with practical applications, being useful for both academics and practitioners. For the former, the proposed swarm algorithms represent a novel approach for solving large-scale cardinality-constrained portfolio optimization problems subject to real-world constraints. In particular, the hybridization of the standard swarm intelligence algorithms with different constraint-handling mechanisms represents a point of novelty in the current literature. For the latter, I think that the integration of ESG factors into quantitative investment frameworks could be useful for institutional investors seeking to align financial objectives with sustainability goals. Moreover, in the presented papers, I observed empirically that sustainable investment strategies not only exhibit risk mitigation during volatile periods, but also are able to create financial value with respect to their non-ESG counterparts.

However, despite showing very promising results, this PhD thesis presents also some limitations and open questions that can be tackled in future research works. Firstly, it is well known that the use of meta-heuristic algorithms is widely problem-dependent, and a solver that performs well on one real-world problem may not yield the same results on another. Therefore, there is a need in the literature to develop heuristics with strong theoretical foundations that can be adapted to various practical problems. To achieve this goal, a greater collaboration between practitioners, who utilize the outputs of these algorithms, and researchers, who develop and refine the solvers, is crucial, in order to ensure that the solutions found are practical and effective. Secondly, the proposed portfolio models have been tested primarily on developed markets, and especially Europe, since ESG criteria are more rigorously followed there due to the recent non-financial disclosure. Hence, the results might differ when applied to other markets, and to enhance the practical appeal of these models, it would be valuable to evaluate their performance in other contexts, in particular the American market. This could provide a broader understanding of the models' applicability and robustness across different regulatory environments and market conditions. Moreover, as ESG data becomes more granular and standardized in all the principal markets, future work could explore real-time, multi-source ESG indicators that could enhance decision-making precision.