Portfolio allocation under market uncertainty: an application of multi-stage stochastic mixed-integer models

Lars Beute

Groningen, January 1, 2024

Master's thesis Econometrics, Operations Research and Actuarial Studies

Supervisor: Dr. W. Romeijnders
Second Assessor: Dr. K.W. Chau

Portfolio allocation under market uncertainty: an application of multi-stage stochastic mixed-integer models

Lars Beute

Fall 2023

Abstract

By representing future uncertain asset prices through scenario trees, we formulate the portfolio selection problem as a two-stage stochastic model. The objective is to minimize the conditional value at risk of the portfolio loss while simultaneously adhering to a set of practical constraints. Particularly, we incorporate cardinality constraints, determining the fixed number of distinct assets to include in the portfolio, and minimum position size constraints. Introducing these constraints involves utilizing integer variables, thereby transforming the two-stage stochastic model into a two-stage stochastic mixed-integer model. We propose a novel decomposition-based algorithm to solve this model. Furthermore, we introduce a straightforward parametric copula-based method for generating asset return scenarios and demonstrate that this approach effectively captures the fat-tail property of asset return distributions. Applying this method to generate scenarios for the proposed two-stage model, we find that the portfolios produced by our model not only track the return levels of a benchmark index but also exhibit significantly greater risk-efficiency. Moreover, we demonstrate that by employing a dynamic portfolio rebalancing strategy, risk-efficiency is even further enhanced. We demonstrate that these results are consistent under our proposed scenario-generation method and are robust to parameter changes. Lastly, a three-stage extension is proposed and we show that portfolios produced under this three-stage extension exhibit the greatest risk-efficiency.

Contents

1 Introduction & Literature Review **5**

2 PSP Model **7**
2.1 Problem Setting 7
2.2 General Framework 8
2.3 Risk Function 10
2.4 Return- & Loss Function 12
2.5 Two-Stage SMIM 13
2.6 Three-Stage SMIM 15

3 Solution Methodology **17**
3.1 MILP 17
3.2 Decomposition Algorithm 19

4 Scenario Modelling **22**
4.1 Second-Stage Scenarios 22
4.2 Third-Stage Scenarios 25

5 Data **26**

6 Scenarios **27**
6.1 Copula Selection 27
6.2 Number of Scenarios 27
6.3 In-Sample Scenario Analaysis 29

7 Portfolio Results **30**
7.1 Benchmark Experiments 31
7.2 Sensitivity Analysis 32
7.3 Out-of-Sample Scenario Analysis 35
7.4 Rolling Horizon 35
7.5 Three-Stage SMIM 37

8 Conclusion & Discussion **38**

1 Introduction & Literature Review

The foundational work by Markowitz (1952) on modern portfolio theory laid the groundwork for the portfolio selection problem (PSP). In modern portfolio theory, investors are assumed to be risk-averse, implying they would consistently opt for the portfolio with the lowest risk when presented with two financial asset portfolios offering the same level of return. The PSP is amongst the most widely studied problems in financial literature derived from this theory and revolves around the creation of an optimal allocation of financial assets. This allocation aims to either minimize expected portfolio risk for a given expected level of return or maximize expected portfolio return for a given level of risk. In the paper by Markowitz, the PSP is formulated as a mean-variance (MV) model, which takes on the form of a quadratic optimization model. In this model, given its tractability, the covariance of individual asset returns in a portfolio is used as a risk measure.

Despite its successes, the MV model has faced criticism. Notably, one of the assumptions made in the model is that the asset returns from which the covariance is derived follow a normal distribution. However, as demonstrated in e.g. McNeil, Frey, and Embrechts (2015) and Fama (1970), financial return series typically exhibit a fat-tailed leptokurtic distribution, characterized by excess kurtosis. Furthermore, as argued in He and Qu (2014), market uncertainty is a significant factor that investors commonly weigh in their decision-making. However, the conventional MV model traditionally computes the expected return and covariance between assets by directly utilizing historical data. This approach has been demonstrated to inadequately address future market fluctuations, emphasized by the substantial estimation errors in both the vector of expected returns and the covariance matrix found in Jobson and Korkie (1980) and Merton (1980). Moreover, employing the covariance of individual asset returns as the risk measure in a portfolio has faced criticism. In the context of the PSP, where the primary emphasis lies in minimizing downside portfolio risk exclusively, utilizing risk metrics like Value at Risk (VaR) or Conditional Value at Risk (CVaR) may be more suitable. Finally, when including practical constraints in the MV model, the PSP transforms in an NP-complete problem for feasibility and an NP-hard problem for optimization (Bienstock (1996); Mansini and Speranza (1999)). Examples of such practical constraints involve measures designed to mitigate operational costs in portfolio management, including limitations on the number of unique stocks in a feasible allocation (cardinality constraints) (see e.g. Gao and Li (2013)) and minimum position size constraints for each unique stock (see e.g. Bonami and Lejeune (2009)).

Evidently, we can argue that utilizing the classical MV model is a suboptimal approach for addressing the PSP. In existing literature, different directions have been proposed to deal with the model limitations. One approach is to modify the risk measure in the MV model to a coherent risk measure such as CVaR as developed in Rockafellar and Uryasev (2000). Utilizing the CVaR as a risk measure in the MV model, the PSP transforms to a linear optimization problem, enhancing solvability. Other suggested approaches vary from improved estimates of correlation coefficients (Elton, Gruber, and Spitzer, 2006) to adding an ex-post time dimension for obtaining the optimal portfolio duration time (Fahmy, 2020) to effective risk diversification approaches (Cesarone and Tardella, 2017).

The absent yet crucial factor that many of these proposed approaches do not take into account is market uncertainty. The MV model makes simplifies the complex reality of financial markets, and its assumptions may not adequately capture the uncertainties and dynamics that investors face. It is imperative to make decisions on investment strategies while considering market conditions and the inherent uncertainty that comes with it (Li and Xu, 2013).

In this study, we prioritize addressing this specific factor by modeling asset returns as a scenario tree. The type of asset we focus on are stocks. We generate scenarios for future stock returns utilizing a parametric copula-based method. We illustrate that this approach provides a straightforward method for generating scenarios and, importantly, proves to be effective in addressing distributional concerns that were previously raised in the classic MV model. We introduce a general structure for the PSP and we show that, leveraging the scenario tree for future asset prices, this framework naturally evolves into the structure of a two-stage stochastic model (SM). This two-stage SM is tailored to the CVaR risk measure, resulting in a linear two-stage SM. With the incorporation of additional practical constraints such as cardinality and position size constraints, we will find that the resulting model adopts the structure of a linear two-stage stochastic mixed-integer model (SMIM).

SMIMs, which merge mixed-integer models (MIMs) with SMs, present significant conceptual and computational difficulties due to the discrete and non-convex aspects of MIMs and the presence of uncertainty through SMs. Notwithstanding the hurdles, numerous applications have emerged across various domains, including supply chain design (Werner, Uggen, Fodstad, Lium, and Egging, 2011), process engineering (Li, Armagan, and Barton, 2011) and many more. In such models, first-stage decisions are made here-and-now while some information is uncertain, whereas later-stage decisions, so-called recourse decisions, are made after some of the uncertainties are revealed. In the context of the PSP, this could be portfolio rebalancing decisions based on stock price fluctuations. The goal of a SMIM is to optimize an objective function that accounts for both the costs of the first-stage decisions as well as the expected costs of the recourse decisions.

In Gülten and Ruszczyński (2015), a two-stage SM with conditional measures of risk is constructed. They find that two-stage models outperform single-stage models in terms of long-term portfolio performance. In Topaloglou, Vladimirou, and Zenios (2008), a multi-stage SM for international portfolio management is developed. In line with Gülten and Ruszczyński (2015), they find that multi-stage models consistently outperform single-stage models. One of the first approaches to solving a PSP with a two-stage SMIM is reported by Stoyan and Kwon (2011). Inspired by this approach, He and Qu (2014) propose a two-stage SMIM similar to the SM in Topaloglou et al. (2008). They show that the portfolios derived from the two-stage SMIM can match the market index while holding less assets. Cui, Bai, Ding, Parkers, Qu, He, and Li (2020) extend on this work by considering a similar two-stage SMIM. Results are in line with the aforementioned studies.

Compared to linear SMs with continuous recourse decision variables, for which efficient solution methods exist, two main difficulties arise in solving SMIMs. Firstly, one has to solve many different integer programs, which generally are NP-hard. Secondly, unlike the situation with continuous recourse decision variables where the expected value function of the future-stage costs is typically convex, the imposition of integer restrictions on recourse decision variables often leads to the non-convexity of the expected value function (Schutlz, Stougie, and van der Vlerk, 1998). Efficient solution methods for SMIMs that follow a certain standard structure, typically referred to as the class of mixed-integer recourse models, exist and are listed in e.g. Li and Grossmann (2019). However, evident from the existing literature, SMIMs that address the PSP typically do not adhere to this standard mixed-integer recourse structure. This acknowledgement requires for innovative, novel solution methods. In Stoyan and Kwon (2011), the developed SMIM is decomposed into sub-problems and these sub-problems are further broken-down by relaxing difficult constraints. He and Qu (2014) deal with the computational difficulties by utilizing a simplification and hybrid solution method to integrate a

standard branch-and-bound. Cui et al. (2020) extend on this approach by integrating a genetic algorithm and a linear programming solver to solve the SMIM.

This study contributes to the existing literature through several dimensions. Firstly, it introduces a two-stage SMIM for the PSP similar to those explored in prior research, but the stock return scenarios we use in the SMIM are generated utilizing a parametric copula-based method. We demonstrate that, in contrast to the approaches employed in existing literature, this method offers a simpler and well-defined mathematical framework for generating scenarios. Nonetheless, it effectively captures the distinctive statistical properties of stock returns. Secondly, where the existing literature on SMIMs that address the PSP utilize solving algorithms that yield approximate solutions, we propose a novel decomposition-based algorithm that yields exact solutions. Thirdly, while the existing literature typically focus solely on the return rates and composition of the derived portfolios, our study goes a step further. In addition to these aspects, we incorporate statistics on the observed CVaR. This inclusion is crucial in demonstrating that SMIMs not only have the capability to align with benchmark indices in terms of return rates but also possess the potential to generate portfolios that exhibit significantly greater risk efficiency. This expanded perspective contributes to a more comprehensive understanding of the implications and benefits of utilizing SMIMs in the context of portfolio selection. Lastly, we are first to propose a three-stage SMIM for the PSP based on the two-stage SMIM. The impact of including an additional stage is thoroughly analyzed both through considering the derived portfolio performance as well as the derived portfolio compositions.

The remainder of this study is organized as follows. In Section 2, we briefly describe the specific problem setting we consider and we provide an extensive discussion on the methodology for both the two-stage- and three-stage SMIMs. In Section 3, we discuss the solution methodology utilized to solve the two- and three-stage SMIM. Section 4 covers the methodology used for scenario modelling. In Section 5 a preliminary analysis on the data is presented. Following this analysis, in Section 6, the methodology on scenario-generation is employed on the data to generate scenarios and the in-sample quality of these scenarios is analyzed. In Section 7, the empirical portfolio experiments performed are described and the results are discussed. Finally, in Section 8, we summarize and discuss the findings of this study. Additionally, we propose directions for further research.

2 PSP Model

The main focus of this section is to develop a two-stage SMIM for the PSP. We initiate this section with a brief description of the problem setting. We continue by introducing a notation for the general PSP framework. By representing asset prices through scenario trees, we discover that this general PSP framework can be adapted to have the structure of a two-stage SMIM. Expanding this model with practical constraints, we observe that the resulting formulation aligns with the structure of a two-stage SMIM. Furthermore, we propose an extension to transform this two-stage SMIM into a three-stage SMIM.

2.1 Problem Setting

In this paper, we consider the problem of a company concerned with the allocation and active management of a set of one type of financial assets: stocks. The company wants to generate profit from market returns simultaneously while maintaining control of their potential downside risk exposure. We model the problem in such a way that it presents itself in stages.

In the first stage, the company decides on an optimal portfolio allocation while simultaneously taking into consideration potential market fluctuations and resulting future stock prices. Consequently, the company ends up with an initial portfolio consisting of a mix of stocks. The stock prices in the initial phase are assumed to be known, that is, we can accurately value the portfolio of the company in the first stage. The initial allocation must be feasible within a set of real world constraints and circumstances that investors face. These constraints include cardinality constraints and minimum position size constraints. In the second stage, the company has the opportunity to reallocate their portfolio. This can be of interest for the company, as it is likely that new market information on stock prices (returns) affect the optimal portfolio allocation. Trivially, the reallocation should be a feasible selection and the goal of the company remains to generate profit from market returns simultaneously while maintaining control of the potential downside risk exposure. This portfolio is used to assess the ultimate portfolio return and risk in the third stage. We operate under the assumption that the time periods between the stages are one week each. This choice of a weekly period is deemed appropriate for the problem at hand. Smaller time periods might incur high portfolio management costs, whereas longer time periods could potentially lead to underperforming portfolios, as their performance is assessed less frequently.

2.2 General Framework

Consider portfolio weight vector $w = (w_1, ..., w_S)$ corresponding to a portfolio comprising S stocks with $w_i \in [0,1]$ the portfolio weight of stock i such that $\sum_{i=1}^{S} w_i = 1$. Assume that portfolio returns are generated by the function $r(w,\xi) : [0,1]^S \times \mathbb{R}^S \to \mathbb{R}$, with ξ a random vector representing uncertain market returns. The most well-known example of such a function is given by $r(w,\xi) = w^\top \xi$, which yields the weighted average random return. Similarly, assume that portfolio risk is generated by the function $\rho(w,\xi) : [0,1]^S \times \mathbb{R}^S \to \mathbb{R}$, with $\rho(\cdot)$ a risk measure. We introduce a general notation for the PSP framework, which will also serve as the framework we use to establish a model for the PSP with a two-stage SMIM structure:

$$\begin{aligned} \min_{w} \; & E_\xi[\rho(w,\xi)] \\ \text{s.t. } & E_\xi[r(w,\xi)] \geq \mu \\ & w \in F. \end{aligned} \tag{1}$$

We aim to construct an optimal portfolio weight vector w^* with the dual purpose of minimizing the expected portfolio risk and guaranteeing the achievement of a predetermined return rate $\mu \in \mathbb{R}$. The set F represents the set of feasible portfolio weight vectors. To illustrate how this framework can be applied, consider the MV model outlined in Markowitz (1952). In the Markowitz model, $r(w,\xi)$ is defined as $r(w,\xi) = w^\top \xi$ and the risk measure $\rho(w,\xi)$ is defined as $\rho(w,\xi) = w^\top \Sigma_\xi w$, where Σ_ξ is the covariance matrix of the uncertain market returns encapsulated in ξ.

Evidently, the framework presented in (1) is adaptable in numerous ways to develop a PSP model that aligns with the particular context of a given problem. Our attention will be directed towards an adaptation wherein the PSP is structured as a two-stage SMIM. To enact this modification, we utilize a modeling approach where future stock prices are conceptualized as part of a large three-stage scenario tree. Let $P = (P_1, ..., P_S)$ denote a vector comprising the current price levels of S distinct stocks, representing the first-stage price levels. Similarly, let $P^n = (P_1^n, ..., P_S^n)$ be a vector encompassing the future price levels in scenario n in the second

stage ($1 \leq n \leq N$). Additionally, let $P^{(n,j)} = (P_1^{(n,j)}, ..., P_S^{(n,j)})$ be a vector comprising the future price levels in scenario (n,j) in the third stage stage, conditional on P^n ($1 \leq j \leq J_n$). We make the modelling decision to have an equal amount of third-stage scenarios (n,j) for every second-stage scenario n, such that we can write $J_n = J$. In Figure 1a, an example scenario tree in this setting is displayed. This scenario tree can be directly linked to our problem setting: in the first stage, the company allocates its capital to a portfolio at prices P, given the potential second- and third-stage prices. Subsequently, in the second stage, the company observes the price vector P^n and has the opportunity to adjust its portfolio through rebalancing. Rebalancing decisions are driven by an assessment of potential third-stage prices $P^{(n,j)}$ and the risk and return rates achieved up to that point. In the context of two-stage SMIMs, the initial portfolio decisions can be referred to as first-stage decisions and the rebalance decisions can be referred to as recourse actions.

To overcome potential challenges highlighted in e.g. Li and Grossmann (2019) regarding the size of the scenario tree and its potential impact on solving an associated SMIM, we will initially make a slight adjustment to the problem setting as outlined above. This adjustment aligns with the modeling methodology presented in e.g. He and Qu (2014) and Cui et al. (2020). For each scenario n in the second stage, we aggregate third-stage scenarios (n,j) and compute the average price vector $\bar{P}^n = (\bar{P}_1^n, ..., \bar{P}_S^n)$. This consolidation allows us to transform the original three-stage tree into a more manageable two-stage structure, as displayed in Figure 1b. Observe that now, all the information regarding future price levels is encapsulated in the second-stage scenarios.

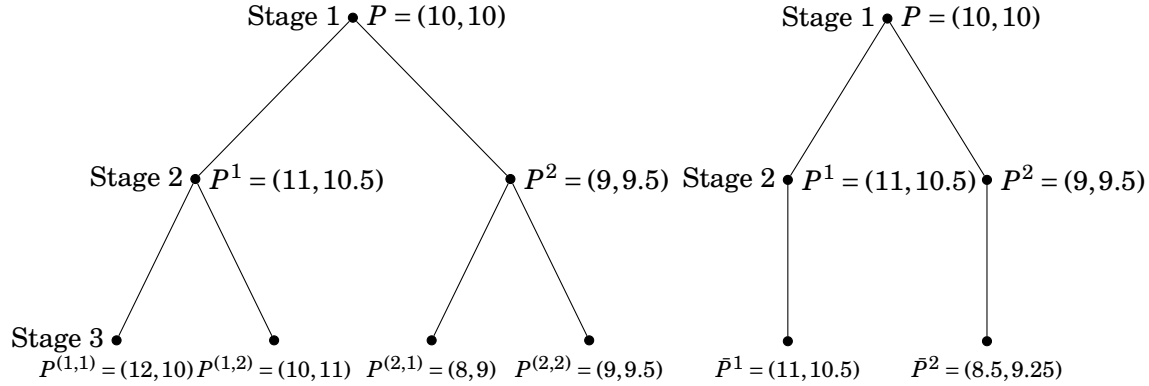

(a) Three-Stage Scenario Tree

(b) Two-Stage Scenario Tree

Figure 1: Example scenario trees, $S = N = J = 2$.

Consider again the framework introduced in (1). Assume that ξ has a finite number of scenarios $\xi_1, ..., \xi_N$ with corresponding probability masses $p_1, ..., p_N$. Observe that, starting from the initial price level P, the scenarios of ξ can be directly utilized to construct the future price scenarios. That is, P^n, and consequently $\bar{P}^n$, can be written as a function of ξ_n, as for any arbitrary $n = 1, ..., N$, we can write $P^n(\xi_n) = P \odot \xi_n$, where $\odot$ is the Hadamard product (element-wise multiplication). This implies that p_n is also the probability of reaching scenario n in the scenario tree. We can use this result to work out the expectations in (1) to obtain the following

model formulation corresponding to the two-stage scenario tree:

$$\begin{aligned} \min_{w} \quad & \sum_{n=1}^{N} p_n \rho(w, \xi_n) \\ \text{s.t.} \quad & \sum_{n=1}^{N} p_n r(w, \xi_n) \geq \mu \\ & w \in F. \end{aligned} \tag{2}$$

As we will discover, this modified framework allows for the integration of the conceptualization of future stock prices within a two-stage scenario tree into the general PSP framework introduced in (1). The next step is to develop the risk- and return functions, as well as to introduce additional constraints to adapt the model to the specifics of our problem, such as minimum position size- and cardinality constraints.

2.3 Risk Function

We will start by developing the risk function $\rho(w, \xi)$. As touched on in the introduction, the risk measure we consider is CVaR. In words, the α% CVaR denotes the expected portfolio return, provided that the portfolio return falls within the most unfavorable $(1-\alpha)$% quantile of the overall portfolio return distribution $(0 < \alpha < 1)$. Figure 2 visually depicts VaR and CVaR. In practical situations, portfolio managers typically strive to minimize the probability of encountering substantial losses. As evident from Figure 2, VaR falls short in addressing this concern, as it focuses solely on the worst-case scenario at a specified confidence level. In contrast, CVaR takes into account the entire distribution of losses beyond the VaR threshold. This characteristic establishes CVaR as a more robust and reliable risk metric for portfolio optimization (Stoyanov, Rachev, and Fabozzi, 2013).

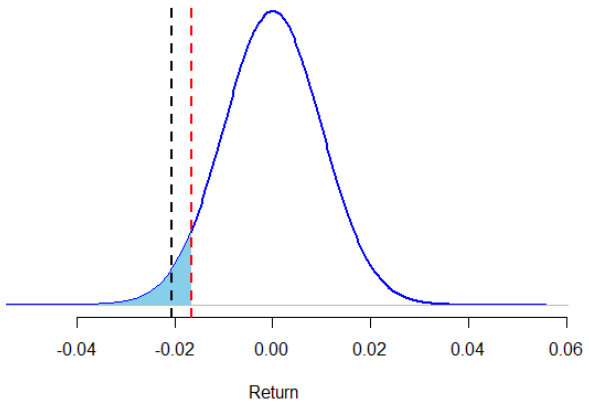

Figure 2: Example portfolio return distribution, 95 % VaR (red) & 95% CVaR (black).

In the context of portfolio optimization, adopting CVaR is not just a practical enhancement over VaR; it also addresses optimization-related shortcomings associated with VaR. VaR is not a coherent risk measure in the definition of Artzner, Delbaen, Eber, and Heath (1999),

implying it fails to adequately reward the benefits of portfolio diversification. Moreover, of particular significance to this study is that when considering discrete portfolio scenarios, VaR behaves as a non-smooth and non-convex function of the portfolio weights (Rockafellar and Uryasev, 2002). This attribute renders the application of VaR impractical within our modeling approach. In contrast to these deficiencies observed with the VaR risk measure, Rockafellar and Uryasev (2002) introduce a definition for CVaR applicable to general distributions, that is, both continuous and discrete distributions. In this definition, they demonstrate that CVaR is a continuous and convex function of the portfolio weights. Moreover, they establish that we can formulate a CVaR optimization model as a linear program.

For simplicity, let us assume that the portfolio return function $r(w,\xi)$ can be separated in a function for profits and a function for losses. Assume that portfolio losses are generated by the function $l(w,\xi):[0,1]^S \times \mathbb{R}^S \to \mathbb{R}$. write $p(\xi)$ for the complete underlying probability distribution of ξ. Rockafellar and Uryasev (2000) formally define the α%-CVaR as

$$\text{CVaR}_\alpha(w) := \frac{1}{1-\alpha}\int_{l(w,\xi)\geq \text{VaR}_\alpha(w,\xi)} l(w,\xi)p(\xi)d\xi. \tag{3}$$

To deal with the α%-VaR in the integral, a new function is introduced:

$$G_\alpha(w,\gamma) = \gamma + \frac{1}{1-\alpha}\int_{l(w,\xi)} (l(w,\xi)-\gamma)p(\xi)d\xi. \tag{4}$$

Rockafellar and Uryasev (2000) show that (4) is a convex function with respect to γ. Moreover, they find that if the solution of the minimization of (4) is given by (w^*,γ^*), the optimal α%-CVaR is given by $G_\alpha(w^*,\gamma^*)$ and the corresponding α%-VaR is given by γ^*.

To demonstrate the usefulness of this result, recall the framework developed in (2). Suppose we have generated N scenarios $\xi_1,...,\xi_N$. We can rewrite (4) as

$$G_\alpha(w,\gamma) = \gamma + \frac{1}{1-\alpha}\sum_{n=1}^{N} p_n[l(w,\xi_n)-\gamma]^+. \tag{5}$$

To deal with the maximum function in (5), introduce the auxiliary variable z_n such that $z_n \geq l(w,\xi_n)-\gamma$ and $z_n \geq 0$, $n = 1,...,N$. Then, to utilize CVaR as a risk measure in (2), we obtain

$$\begin{aligned} \min_{\substack{w,\ \gamma,\\ z_1,...,z_N}} \quad & \gamma + \frac{1}{1-\alpha}\sum_{n=1}^{N} p_n z_n \\ \text{s.t.} \quad & \sum_{n=1}^{N} p_n r(w,\xi_n) \geq \mu \\ & z_n \geq l(w,\xi_n)-\gamma,\ n = 1,...,N \\ & z_n \geq 0,\ n = 1,...,N \\ & \gamma \in \mathbb{R} \\ & w \in F. \end{aligned} \tag{6}$$

Observe that we have effectively integrated a linear expression for $\rho(w,\xi)$ in (2), such that $\rho(\cdot)$ now denotes the portfolio CVaR.

2.4 Return- & Loss Function

The next step involves deriving expressions for the return function $r(w,\xi)$ and the loss function $l(w,\xi)$ in (6). Given our specific problem context and the inherent two-stage structure of the resulting model, we must first introduce additional notation. Consider again the notation introduced in Section 2.2. Similar to the first-stage portfolio weight decision vector w, let $w^0 = (w_1^0,...,w_S^0) \in [0,1]^S$ represent the initial portfolio weight vector. This implies that we allow for the possibility that the company may already manage a portfolio. For $n = 1,...,N$, let $w^n = (w_1^n,...,w_S^n) \in [0,1]^S$ the portfolio weight vector after the described rebalancing procedure in scenario n. Let $\tilde{w}^0$, $\tilde{w}$, $\tilde{w}^n \in \mathbb{R}_+^S$ represent the portfolio holding vectors associated with weight vectors w^0, w, and w^n, respectively. For instance, for any arbitrary stock $i = 1,...,S$, $w_i^0 = \frac{\tilde{w}_i^0}{\sum_{i=1}^S \tilde{w}_i^0}$. Write $C \in \mathbb{R}_+$ for the initial available capital intended for investment in the portfolio. It is straightforward to see that $V^0 = C + \sum_{i=1}^S \tilde{w}_i^0 P_i$, where V^0 is the initial portfolio wealth. Similarly, given that we assume that all available capital is invested in the portfolio in the first stage, we have $V^n = \sum_{i=1}^S \tilde{w}_i^n \bar{P}_i^n(\xi_n)$, where V^n is the expected portfolio wealth in scenario n. The expected portfolio return in scenario n is then given by $R^n = \frac{V^n}{V^0} - 1$. To model the dynamics between $\tilde{w}^0$, $\tilde{w}$ and $\tilde{w}^n$, let $b = (b_1,...,b_S) \in \mathbb{R}_+^S$ and $s = (s_1,...,s_S) \in \mathbb{R}_+^S$ vectors where each element signifies the quantity of the respective stock bought and sold, respectively. Likewise, define b^n, $s^n \in \mathbb{R}_+^S$, representing the number of stocks purchased and sold, respectively, in scenario n. Let $\Phi = (\tilde{w},\ \gamma,\ b,\ s,\ \tilde{w}^1,...,\tilde{w}^n,\ b^1,...,b^N,\ s^1,...,s^N,\ z_1,...,z_N,\ V^1,...,V^N,\ R^1,...,R^N)$ be the decision vector with the newly introduced decision variables. Note that the portfolio weight vectors are dropped from the decision vector, as they can be directly computed from the portfolio holding vectors. We can extend (6) to obtain

$$
\begin{aligned}
\min_{\Phi}\ & \gamma + \frac{1}{1-\alpha} \sum_{n=1}^{N} p_n z_n && (7)\\
\text{s.t. } & \tilde{w}_i = \tilde{w}_i^0 + b_i - s_i,\ i = 1,...,S\\
& \tilde{w}_i^n = \tilde{w}_i + b_i^n - s_i^n,\ i = 1,...,S,\ n = 1,...,N\\
& C + \sum_{i=1}^{S} s_i P_i = \sum_{i=1}^{S} b_i P_i\\
& \sum_{i=1}^{S} s_i^n P_i^n(\xi_n) = \sum_{i=1}^{S} b_i^n P_i^n(\xi_n),\ n = 1,...,N\\
& V^n = \sum_{i=1}^{S} \tilde{w}_i^n \bar{P}_i^n(\xi_n),\ n = 1,...,N\\
& R^n = \frac{V^n}{V^0} - 1,\ n = 1,...,N\\
& \sum_{n=1}^{N} p_n R^n \geq \mu\\
& z_n \geq -R^n - \gamma,\ n = 1,...,N\\
& z_n \geq 0,\ n = 1,...,N\\
& \gamma \in \mathbb{R}\\
& \tilde{w} \in F\\
& \tilde{w}^n \in F,\ n = 1,...,N.
\end{aligned}
$$

We find that, in our model, the return function is described by R^n and its corresponding dynamics and the loss function is described by $-R^n$ and its corresponding dynamics. Compared to (6), the first six constraints are new in this model formulation. The first- and second constraint are referred to as the first- and second-stage stock balance conditions, respectively. These conditions ensure that the position size in a stock at the end of a stage equals the position size of that stock at the initiation of a stage, adjusted for stocks sold and purchased during the stage. The third- and fourth constraint are the first- and second-stage cash flow constraints, respectively. In the first-stage, all initially available capital and proceeds from stock sales is allocated to the portfolio. In the second stage, the income generated from selling assets is directly utilized to purchase new stocks. It is worth noting that by introducing additional variables to represent the capital position in the second stage, the assumption that all available capital must be invested can be easily modified to accommodate a capital position in the second stage as well. However, for the specific empirical experiments conducted in this study, we choose not to make this adjustment. The fifth- and sixth constraints have been detailed above.

The model in (7) adopts the structure of a two-stage SM. This implies that we have successfully illustrated the adaptation of the general PSP framework outlined in (1) to a two-stage SM under the CVaR risk measure. We can proceed by further enhancing the two-stage SM through the inclusion of additional practical constraints. Through this augmentation, we will discover that the resulting model takes on the structure of a two-stage SMIM.

2.5 Two-Stage SMIM

We initiate this section by formally defining all notation used.
Sets:

S	The set of available stocks, indexed by $i = 1, ..., S$.

User-specified parameters:

μ	Minimum expected portfolio return
α	Critical percentile for (C)VaR

Deterministic input data:

N	Number of second-stage scenarios, indexed by $n = 1, ..., N$
C	Initial available capital
$\tilde{w}_i^0$	Initial position in stock i
K	Number of unique stocks to be held in the portfolio
$\tilde{w}_{\min}$	Minimum position per unique stock to be held
P_i	Initial price per unit of stock i

Scenario dependent data:

p_n	Probability of observing scenario n
ξ_n	Stock returns in scenario n
$P_i^n(\xi_n)$	Price per unit of stock i in scenario n in the second-stage
$\bar{P}_i^n(\xi_n)$	Expected price per unit of stock i in scenario n in the third-stage

Auxiliary variables:

z_n Expected portfolio shortfall in excess of VaR in scenario n
γ VaR value
V^n Expected portfolio wealth in scenario n
R^n Expected portfolio return in scenario n

First-stage decision variables:

b_i Number of units of stock i purchased in stage 1
s_i Number of units of stock i sold in stage 1
$\tilde{w}_i$ Position size in stock i after transactions in stage 1
h_i Equals 1 if we hold stock i in stage 1, 0 otherwise

Second-stage / recourse decision variables:

b_i^n Number of units of stock i purchased in scenario n in stage 2
s_i^n Number of units of stock i sold in scenario n in stage 2
$\tilde{w}_i^n$ Position size in stock i after transactions in stage 2 in scenario n
h_i^n Equals 1 if we hold stock i in scenario n in stage 2, 0 otherwise

Extend the decision vector Φ introduced in the previous section with the newly introduced decision variables $h_i \in \{0,1\}$ and $h_i^1,...,h_i^n \in \{0,1\}$ for $i = 1,...,S$. Then, extending on (7), inspired by the model in Cui et al. (2020) and Topaloglou et al. (2008), we define the following two-stage SMIM:

Objective Function

$$\min_{\Phi} \gamma + \frac{1}{1-\alpha} \sum_{n \in N} p_n z_n \tag{8}$$

First-Stage Constraints

$$\text{s.t.} \quad \tilde{w}_i = \tilde{w}_i^0 + b_i - s_i, \ i = 1,...,S \tag{9}$$

$$C + \sum_{i=1}^{S} s_i P_i = \sum_{i=1}^{S} b_i P_i \tag{10}$$

$$\sum_{i=1}^{S} h_i = K \tag{11}$$

$$\tilde{w}_{\min} h_i \leq \tilde{w}_i, \ i = 1,...,S \tag{12}$$

$$\tilde{w}_i \leq M h_i, \ i = 1,...,S \tag{13}$$

$$h_i \in \{0,1\}, \ i = 1,...,S \tag{14}$$

$$\tilde{w}_i, b_i, s_i \in \mathbb{R}_+, \ i = 1,...,S \tag{15}$$

$$\gamma \in \mathbb{R} \tag{16}$$

Second-Stage Constraints

$$\tilde{w}_i^n = \tilde{w}_i + b_i^n - s_i^n, \ i = 1,...,S, \ n = 1,...,N \tag{17}$$

$$\sum_{i=1}^{S} s_i^n P_i^n(\xi_n) = \sum_{i=1}^{S} b_i^n P_i^n(\xi_n), \ n = 1,...,N \tag{18}$$

$$\sum_{i=1}^{S} h_i^n = K,\ n = 1,...,N \tag{19}$$

$$\tilde{w}_{\min} h_i^n \leq \tilde{w}_i^n,\ i = 1,...,S,\ n = 1,...,N \tag{20}$$

$$\tilde{w}_i^n \leq M h_i^n,\ i = 1,...,S,\ n = 1,...,N \tag{21}$$

$$V^n = \sum_{i=1}^{S} \bar{P}_i^n(\xi_n)\tilde{w}_i^n,\ n = 1,...,N \tag{22}$$

$$R^n = \frac{V^n}{V^0} - 1,\ n = 1,...,N \tag{23}$$

$$z_n \geq -R^n - \gamma,\ n = 1,...,N \tag{24}$$

$$z_n \geq 0,\ n = 1,...,N \tag{25}$$

$$\sum_{n \in N} p_n R^n \geq \mu \tag{26}$$

$$h_i^n \in \{0,1\},\ i = 1,...,S,\ n = 1,...,N \tag{27}$$

$$\tilde{w}_i^n, b_i^n, s_i^n \in \mathbb{R}_+,\ i = 1,...,S,\ n = 1,...,N \tag{28}$$

$$z_n \in \mathbb{R}_+,\ n = 1,...,N. \tag{29}$$

Let us focus on the constraints not covered in the previous sections. Constraints (11) and (19) are the first- and second-stage cardinality constraints, respectively. These constraints restrict the number of unique stock holdings in a stage to a fixed value $K \leq S$. In (12) and (20), a minimum position size per unique stock is ensured in the first- and second-stage, respectively. For M sufficiently large, constraints (13) and (21) are linear expressions that ensure that, for $i = 1,...,S$ and $n = 1,...,N$, h_i and h_i^n behave as defined in the notation in both the first- and second-stage, respectively. Even though it could be argued that the first- and second-stage holding, buying and selling variables in respectively (15) and (28) should take on strictly integer values, we have chosen to adopt a different approach in line with the perspective outlined by Woodside-Oriakhi, Lucas, and Beasley (2013). Specifically, we will treat these variables as continuous due to the assumption that the invested sums are of significant magnitude. Beyond this rational argument, one should recognize that without assuming continuity for these variables, we would face a substantial and challenging increase in computational complexity, as highlighted by He and Qu (2014). We believe this is beyond the scope of our research and warrants further research. We do restrict the decision variables in (14) and (27) to the binary set. This implies that by imposing cardinality constraints as well as minimum position size constraints, the two-stage SM with continuous recourse decision variables in (7) evolves into a two-stage SMIM.

2.6 Three-Stage SMIM

Literature exists on SMs that address the PSP in a multi-stage setting, see e.g. Topaloglou et al. (2008). However, to our knowledge, current research on modeling the PSP using a SMIM is primarily limited to two-stage formulations. The challenges associated with solving two-stage SMIMs are acknowledged in the literature, making the prospect of addressing similar problems in multi-stage SMIMs even more challenging. To illustrate the complexity, consider a two-stage SMIM with 5000 scenarios, corresponding to a scenario tree of size 5000. Introducing an additional third-stage, with only 2 scenarios for each second-stage scenario, causes the three-stage SMIM to grow to $5000 \cdot 2 = 10000$ scenarios. For continuous SMs, this would

not directly pose problems, as the convexity in the model can still be exploited to derive solutions efficiently. However, as one can imagine, for SMIMs, given the increase in the number of associated NP-hard integer programs when we include just one additional stage, we reach a point where the model becomes practically unsolvable with modern-day technology. Despite the challenges, we find it valuable to introduce a three-stage SMIM for the PSP. Our reasoning is straightforward: existing literature demonstrates favorable outcomes in portfolio performance derived from two-stage SMs and SMIMs addressing the PSP. The introduction of just a single additional stage already enhances the capacity of the model to quantify future uncertainty. Consequently, we anticipate that portfolios derived from the three-stage SMIM should demonstrate superior performance compared to those derived from the two-stage SMIM.

Consider the three-stage scenario tree in Figure 1a. For $n = 1,...,N$ and for $j = 1,...,J$, write $p_{(n,j)}$ for the probability of reaching scenario (n,j) in the third-stage. Similarly, write $R^{(n,j)}$, $V^{(n,j)}$ and $z_{(n,j)}$ for the return, wealth and portfolio shortfall in excess of VaR in scenario (n,j). Modify the decision vector Φ by dropping R^n, V^n and z_n and replace these for $R^{(n,j)}$, $V^{(n,j)}$ and $z_{(n,j)}$. Refer to the new decision vector as $\tilde{\Phi}$. We define the three-stage SMIM as:

Objective Function

$$\min_{\tilde{\Phi}} \gamma + \frac{1}{1-\alpha} \sum_{n=1}^{N} \sum_{j=1}^{J} p_{(n,j)} z_{(n,j)} \tag{30}$$

First-Stage Constraints

$$\text{s.t.} \quad \tilde{w}_i = \tilde{w}_i^0 + b_i - s_i,\ i = 1,...,S \tag{31}$$

$$C + \sum_{i=1}^{S} s_i P_i = \sum_{i=1}^{S} b_i P_i \tag{32}$$

$$\sum_{i=1}^{S} h_i = K \tag{33}$$

$$\tilde{w}_{\min} h_i \leq \tilde{w}_i,\ i = 1,...,S \tag{34}$$

$$\tilde{w}_i \leq M h_i,\ i = 1,...,S \tag{35}$$

$$h_i \in \{0,1\},\ i = 1,...,S \tag{36}$$

$$\tilde{w}_i, b_i, s_i \in \mathbb{R}_+,\ i = 1,...,S \tag{37}$$

$$\gamma \in \mathbb{R} \tag{38}$$

Second-Stage Constraints

$$\tilde{w}_i^n = \tilde{w}_i + b_i^n - s_i^n,\ i = 1,...,S,\ n = 1,...,N \tag{39}$$

$$\sum_{i=1}^{S} s_i^n P_i^n = \sum_{i=1}^{S} b_i^n P_i^n,\ n = 1,...,N \tag{40}$$

$$\sum_{i=1}^{S} h_i^n = K,\ n = 1,...,N \tag{41}$$

$$\tilde{w}_{\min} h_i^n \leq \tilde{w}_i^n,\ i = 1,...,S,\ n = 1,...,N \tag{42}$$

$$\tilde{w}_i^n \leq M h_i^n,\ i = 1,...,S,\ n = 1,...,N \tag{43}$$

$$h_i^n \in \{0,1\},\ i = 1,...,S,\ n = 1,...,N \tag{44}$$

$$\tilde{w}_i^n, b_i^n, s_i^n \in \mathbb{R}_+,\ i = 1,...,S,\ n = 1,...,N \tag{45}$$

Third-Stage Constraints

$$V^{(n.j)} = \sum_{i=1}^{S} P_i^{(n,j)} \tilde{w}_i^n,\ n = 1,...,N,\ j = 1,...,J \tag{46}$$

$$R^{(n,j)} = \frac{V^{(n,j)}}{V^0} - 1,\ n = 1,...,N,\ j = 1,...,J \tag{47}$$

$$z_{(n,j)} \geq -R^{(n,j)} - \gamma,\ n = 1,...,N,\ j = 1,...,J \tag{48}$$

$$z_{(n,j)} \geq 0,\ n = 1,...,N,\ j = 1,...,J \tag{49}$$

$$\sum_{n \in N} \sum_{j \in J} p_{(n,j)} R^{(n,j)} \geq \mu \tag{50}$$

$$z_{(n,j)} \in \mathbb{R}_+,\ n = 1,...,N,\ j = 1,...,J. \tag{51}$$

The overall structure and constraints of the three-stage SMIM closely resemble those of the two-stage SMIM discussed in the preceding section. The key distinction lies in the prices at which the portfolio wealth, return, and CVaR are evaluated. In the previous section, we referred to R^n, V^n, and z_n as the *expected* portfolio return, wealth, and shortfall in excess of VaR in scenario n, respectively. This nomenclature was employed because we utilized the average price vector $\bar{P}^n$, allowing us to formulate the PSP in the problem setting as a two-stage SMIM. If we shift our focus to the individual third-stage scenario prices $P^{(n,j)}$ instead of the average third-stage price vector $\bar{P}^n$, the portfolio wealth, return, and CVaR are evaluated at the true individual third-stage prices. Consequently, we no longer refer to them as their expected values. We hypothesize that, by considering each third-stage return scenario emanating from a given scenario n individually, as opposed to relying on the average third-stage return scenario, the three-stage SMIM should lead to an improved quality of derived portfolio solutions.

3 Solution Methodology

In this section, we delve into the methodology employed for solving the introduced two-stage- and three-stage SMIMs discussed in the previous section. We commence by exploring the most direct approach to solving such models and underscore the limitations associated with this method. Subsequently, we demonstrate that conventional techniques designed to address these limitations are not directly applicable to our specific SMIMs. To overcome these challenges, we propose a novel decomposition-based solution methodology.

3.1 MILP

Consider the two-stage- and three-stage SMIM introduced in the previous section. Completely formulating the objective function and constraints for all scenarios reveals that both SMIMs adopt the format of a large mixed-integer linear program (MILP). MILPs are known to be NP-complete, making them challenging to solve. Typically, efficient MILP-solving software is employed to tackle such problems, leveraging established branch-and-bound algorithms. Despite continuous advancements in solution times achieved by these software tools, depending on the problem setting, MILPs can become excessively large depending on the problem setting, making it impractical to obtain solutions from such software. MILPs deducted from two-stage SMIMs, let alone three-stage SMIMs, typically become very large, with the program size growing significantly as the number of considered scenarios increases. To address these challenges, existing literature traditionally leverages the observation that many SMIMs adhere to a specific standard structure, which we refer to as the class of linear mixed-integer recourse (MIR)

models. Appendix A provides a concise introduction to this class of models.

Assume that, for now, we relax the minimum expected return constraint (26). Then, the two-stage SMIM can also be formulated as a two-stage MIR model. That is, following the notation introduced in Appendix A, we can write the two-stage SMIM as

$$\min_{x\in\mathscr{X}}\{c^\top x + Q(x) : Ax \le b\}, \tag{52}$$

with

$$Q(x) = \sum_{n\in N} p_n v_n(x) \tag{53}$$

and

$$v_n(x) := \min_{y_n\in\mathscr{Y}}\{q_n^\top y_n : W_n y_n \ge h_n - T_n x\}. \tag{54}$$

To understand this, note that the variable x represents the first-stage decision vector, i.e. $\forall i = 1,...,S$, $x = (\gamma,\ h_i,\ w_i,\ b_i,\ s_i)$ and $\mathscr{X} = \mathbb{R}\times\{0,1\}^S\times\mathbb{R}_+^{3S}$. Examining the objective function (8), we observe that the first-stage cost vector c is defined as $c = (1,0,...,0)$. The matrix A and the right-hand side vector b in (52) can be straightforwardly constructed to represent the first-stage constraints. The expected-value function $Q(x)$ calculates the weighted sum of the second-stage value functions $v_n(x)$, $\forall n = 1,...,N$. These second-stage value functions correspond to the objective values of individual second-stage minimization problems associated with scenario n. To understand this, consider (54) and fix a scenario n. Then, y_n represents the second-stage decision vector, i.e. $\forall i \in S$, $y_n = (z_n,\ R^n,\ V^n,\ h_i^n,\ w_i^n,\ b_i^n,\ s_i^n)$ and $\mathscr{Y} = \mathbb{R}_+\times\{0,1\}^S\times\mathbb{R}^{3S}$. Examining the objective function (8), the second-stage cost vector q_n is given by $q_n = (\frac{1}{1-\alpha},0,...,0)$. Appropriately constructing the recourse matrix W_n and technology matrix T_n, along with the right-hand side vector h_n, they represent the second-stage constraints of the two-stage SMIM. The structure of the associated MILP is given by

$$\begin{aligned}
\min_{x\in\mathscr{X},y_1,..,y_N\in\mathscr{Y}}\quad & c^\top x + q_1^\top y_1 + q_2^\top y_2 + \cdots + q_N^\top y_N && && && (55)\\
& Ax && && && \le b\\
& T_1 x + W_1 y_1 && && && \ge h_1\\
& T_2 x && + W_2 y_2 && && \ge h_2\\
& \vdots && && \ddots && \vdots\\
& T_N x && && && + W_N y_N \ge h_N.
\end{aligned}$$

The block-diagonal structure of this MILP, commonly referred to as the L-shaped structure, naturally lends itself to a classical decomposition technique known as Benders decomposition (Benders, 1962). This approach involves dividing the problem into two subsets: a master problem, which focuses on the first-stage variables, and a set of N second-stage subproblems. The subproblems are solved for the second-stage variables for a given first-stage solution obtained from the master problem. Algorithms utilizing Benders decomposition often exhibit significantly reduced computation times, as typically the master problem and subproblems are considerably smaller compared to the non-partitioned original problem. Moreover, the subproblems can be parallelly solved, further enhancing computational efficiency. The L-shaped method (van Slyke and Wets, 1969) is one of the most well-known algorithms that exploits this decomposition.

In an ideal situation, our model exhibits a similar structure such that we can make use of these efficient solution algorithms. However, without relaxing the minimum expected return constraint (26), we observe a deviation from this structure. This is evident when considering that the constraint encompasses the portfolio returns across all N scenarios. Consequently, the structure of the associated MILP, with this constraint included, takes on a form of

$$\begin{aligned}
\min_{x\in\mathcal{X},y_1,..,y_N\in\mathcal{Y}} \quad & c^\top x + q_1^\top y_1 + q_2^\top y_2 + \cdots + q_N^\top y_N && (56)\\
& Ax && \geq b\\
& T_1 x + W_1 y_1 && \geq h_1\\
& T_2 x \qquad + W_2 y_2 && \geq h_2\\
& \vdots \qquad\qquad\qquad \ddots && \vdots\\
& T_N x \qquad\qquad\qquad + W_N y_N && \geq h_N\\
& \qquad e_2^\top y_1 + e_2^\top y_2 + \ldots \quad + e_2^\top y_N && \geq \mu,
\end{aligned}$$

where e_2 is the appropriately-sized standard unit vector, i.e. $e_2 = (0,1,0,...,0)$. This somewhat atypical structure does not allow for direct application of decomposition techniques such as Benders decomposition.

Studies examining SMs and SMIMs for the PSP similar to ours acknowledge this atypical structure and resort to alternative solving approaches. While some studies, like (Topaloglou et al., 2008), choose to solve the complete MILP using MILP-solving software, potentially incurring less efficient computation times, others opt for obtaining solutions through computational algorithms from computer science disciplines, as seen in works such as (He and Qu, 2014) and (Cui et al., 2020). Notably, a clear advantage of these algorithms is their significantly faster computation times. However, a disadvantage is that these algorithms provide approximate solutions, whereas solutions derived from solving the associated MILPs offer exact solutions. We contend that, particularly in the context of portfolio selection, where typically substantial amounts of capital are involved, the importance of exact solutions surpasses the advantages of faster computation times. Therefore, aligning with the approach in Topaloglou et al. (2008), all presented solutions of the SMIMs are obtained by solving the associated complete MILPs. The MILPs are implemented in the Python programming language and solved using the Gurobi optimization software (Gurobi Optimization, LLC, 2023).

We do want to address the gap in literature on obtaining exact solutions to SMIMs similar to ours in a more efficient manner. Hence, while not employed in this study, we now put forth a novel decomposition-based algorithm designed to efficiently solve the MILP of the two-stage SMIM as presented in (56). A notable advantage of our proposed algorithm over the computer science-related algorithms explored thus far is its capability to provide exact solutions rather than approximate ones. Numerical results obtained using this algorithm will be presented in a later work.

3.2 Decomposition Algorithm

Consider an arbitrary two-stage continuous stochastic program of the form

$$\min_{x\in\mathcal{X}}\{c^\top x : \sum_{n=1}^{N} p_n v_n(x) \leq \eta\}, \tag{57}$$

where

$$v_n(x) = \min_{y \in \mathcal{Y}} \{q_n^\top y : W_n y \geq h_n - T_n x\} \tag{58}$$

and η is a fixed parameter. Henceforth, we designate constraints like the one on the right-hand side in (57) as *global second-stage constraints*, as they articulate a connection between second-stage variables across multiple scenarios. As elaborated on in our examination of the MILP associated with the two-stage SMIM under consideration, the minimum expected return constraint exemplifies a global second-stage constraint. These global second-stage constraints introduce a violation of the standard MILP structure used in existing decomposition algorithms. Consequently, in this section, we propose a decomposition algorithm specific for such constraints. The idea is to apply this algorithm specifically to address the minimum expected return constraint. We can then eliminate the constraint from the MILP. Subsequently, we can leverage established decomposition algorithms to handle the remaining MILP, which, as demonstrated earlier, adheres to the standard MILP structure compatible with existing decomposition algorithms.

By linear programming duality theory, (58) can be equivalently written as

$$v_n(x) = \max_{\lambda \geq 0} \{\lambda_n^\top (h_n - T_n x) : \lambda_n^\top W_n \leq q_n^\top\}, \tag{59}$$

where λ_n is the vector of dual variables in scenario n. The right-hand side of (59) is a polyhedral set. Then, by the Fundamental Theorem of Linear Programming, which tells us that if an optimal solution exists, it is guaranteed to occur at an extreme point, we can rewrite (59) as

$$v_n(x) = \max_{k=1,\ldots,K_n} (\lambda_n^k)^\top (h_n - T_n x), \tag{60}$$

where λ_n^k is the kth extreme point in scenario n. Hence, the minimization problem (57) is equivalent to

$$\min_{\substack{x \in \mathcal{X} \\ \psi_n \in \mathbb{R} \ \forall n=1,\ldots,N}} \{c^\top x : \sum_{n=1}^{N} p_n \psi_n \leq \eta,\ \psi_n \geq \max_{k=1,\ldots,K_n} (\lambda_n^k)^\top (h_n - T_n x)\ \forall n = 1,\ldots,N\}. \tag{61}$$

To deal with the maximum function in (61), we can rewrite to

$$\min_{\substack{x \in \mathcal{X} \\ \psi_n \in \mathbb{R} \ \forall n=1,\ldots,N}} \{c^\top x : \sum_{n=1}^{N} p_n \psi_n \leq \eta,\ \psi_n \geq (\lambda_n^k)^\top (h_n - T_n x)\ \forall k = 1,\ldots,K_n\ \forall n = 1,\ldots,N\}. \tag{62}$$

The reformulation of the original two-stage stochastic program in (57) to (62) allows for a straightforward decomposition algorithm.

The algorithm is presented in Algorithm 1. In the initialization phase, the master problem is defined and the iteration counter is started. In iteration t in the iteration phase, we first solve the master problem, yielding solutions x^t and $\psi_n^t\ \forall n = 1,\ldots,N$. In line 5, note that the maximization term is equal to $v_n(x^t)$. We check if for the obtained solutions x^t and ψ_n^t (61) holds. If this is true, from above, we know that x^t and ψ_n^t are optimal: the algorithm stops. If this does not hold, $\forall n = 1,.,,,N$, we add the violating constraints $\psi_n \geq (\lambda_n^{k^t})^\top (h_n - T_n x)$, where $\lambda_n^{k^t}$ is the solution of $v_n(x^t)$ in line 5. We update the iteration counter to $t = t + 1$ and the iteration phase repeats. In the worst case, the algorithm as presented in Algorithm 1 will result in a master problem as large as the original two-stage stochastic program. Consequently,

from a computational standpoint, employing the proposed algorithm is consistently advantageous compared to solving the entire stochastic program as a single extensive linear program. Additionally, we may contemplate introducing an acceptable error parameter $\epsilon > 0$ in line 5, modifying the if-statement to verify whether $\psi_n^t + \epsilon \geq \max_{k=1,...,K_n^t}(\lambda_n^k)^\top(h_n - T_n x^t)\ \forall n = 1,...,N$. This would speed up the algorithm even more.

Algorithm 1 Global Second-Stage Constraints Decomposition Algorithm

Initialization Phase:

1: Define Master Problem (MP):

$$\min_{\substack{x\in\mathscr{X}\\ \psi_n\in\mathbb{R}\ \forall n=1,...,N}} \{c^\top x : \sum_{n=1}^{N} p_n\psi_n \leq \eta\}$$

2: Set $t = 1$

Iteration Phase:

3: **while** true **do**
4: Solve the MP, yielding solutions x^t and $\psi_n^t\ \forall n = 1,...,N$.
5: **if** $\psi_n^t \geq \max_{k=1,...,K_n^t}(\lambda_n^k)^\top(h_n - T_n x^t)\ \forall n = 1,...,N$ **then**
6: x^t is optimal. Break loop.
7: **else**
8: Add violating constraints $\psi_n \geq (\lambda_n^{k^t})^\top(h_n - T_n x)$ to the master problem $\forall n = 1,...,N$.
9: Set $t = t + 1$
10: **end if**
11: **end while**

Now, consider the minimum expected return constraint (26) in the two-stage SMIM. Observe that we can equivalently write

$$\sum_{n=1}^{N} -p_n R^n \leq -\mu. \tag{63}$$

Observe that the value of R^n can also be obtained by solving

$$\hat{v}_n(x) = \max_{R^n}\{R^n : W_n y \geq h_n - T_n x\} = -\min_{R^n}\{-R^n : W_n y \geq h_n - T_n x\}, \tag{64}$$

where W_n, T_n and h_n are defined as in Section 3.1, i.e. such that they correspond to the second-stage minimization problem in scenario n in the two-stage SMIM without minimum expected return constraint. We thus can write the minimum expected return constraint as

$$\sum_{n=1}^{N} -p_n \hat{v}_n \leq -\mu, \tag{65}$$

that is, the minimum expected return constraint can be written in the format used in the algorithm. We can now write the complete two-stage SMIM as

$$\min_{x\in\mathscr{X}}\{c^\top x + Q(x) : Ax \leq b,\ \sum_{n=1}^{N} -p_n \hat{v}_n(x) \leq \mu\}, \tag{66}$$

with the vectors, matrices and functions defined as above. Algorithm 1 can now be applied to deal with the minimum expected return constraint, such that the remaining standard-structured associated MILP can be solved using established (integer) decomposition techniques. This is something we will demonstrate in a later work.

4 Scenario Modelling

This section covers the methods utilized in modeling the uncertain market return scenarios introduced in Section 2. Given the general structure of scenario trees, where the third-stage scenarios depend on the second-stage scenarios they branched from,we distinctly outline the methods for the second-stage scenarios and those for the third-stage scenarios. We initiate this section by discussing the methodology for the second-stage scenarios, followed by a discussion on the methods used for the third-stage scenarios.

4.1 Second-Stage Scenarios

It is well known that modelling stock returns is challenging. Typically, returns display correlations, and, as depicted in Figure 3, their historical distribution often demonstrates asymmetries and fat tails. Consequently, modeling them in accordance with a conventional distribution, such as a normal distribution, may not be suitable. The complexity intensifies when turning attention to modeling portfolio returns. In this context, the objective is to establish a multivariate distribution for asset prices. Multivariate return series typically demonstrate dynamic non-linear cross-correlation for contemporaneous returns (McNeil et al., 2015). Moreover, as is evident from Figure 3, the marginal return series of the individual stocks differ, introducing an additional layer of complexity to the modeling process. However, particularly when concerned with controlling downside risk in the tail of the portfolio return distribution, effectively capturing observed multivariate return moments, such as skewness and excess kurtosis in the distribution of asset prices, becomes of vital importance (Topaloglou et al., 2008). Additionally, considering that periods characterized by high volatility and extreme returns are often directly associated with increased cross-correlation, significantly shaping the tail of the return distribution (Sandoval and De Paula Franca, 2012), it becomes essential to employ an adequate modeling approach.

Evidently, portfolio solutions generated by risk-controlling optimization processes are massively dependent on the scenarios (Quaranta and Zaffaroni, 2008). To overcome modelling difficulties, Topaloglou et al. (2008) and He and Qu (2014) utilize the moment-matching scenario generation procedure for stock returns as proposed in Høyland and Wallace (2001) and Høyland, Kaut, and Wallace (2003) in their two-stage SM and SMIM, respectively. The idea of this method is that we generate scenario sets such that key statistics of in this case the stock returns match specified target values closely. In this particular case, we would consider the first four moments of the stock return distributions as key statistics. The specified target values for these statistics are derived from historical data on the stock returns. Even though the moment-matching algorithm offers a straightforward method for generating a multivariate distribution for stock returns and, consequently, prices, we are a bit conservative towards this approach. The generated distribution not only heavily relies on the quality of the initial scenario moments but also lacks a provided convergence proof, as emphasized in Mehrotra and Papp (2013). In Cui et al. (2020), the copula method introduced by Kaut and Wallace (2011) is employed. Not only do they find that utilizing the scenarios generated from this copula method

in their model yield portfolio solutions that can adequately track a benchmark index, but also that the obtained portfolio solutions exhibit stability across different samples of generated scenarios. Encouraged by these findings, we also employ a copula-based method for generating second-stage stock return scenarios. Different from the method introduced in Kaut and Wallace (2011), this copula-based method is a parametric approach and makes use of the copula methodology covered in McNeil et al. (2015).

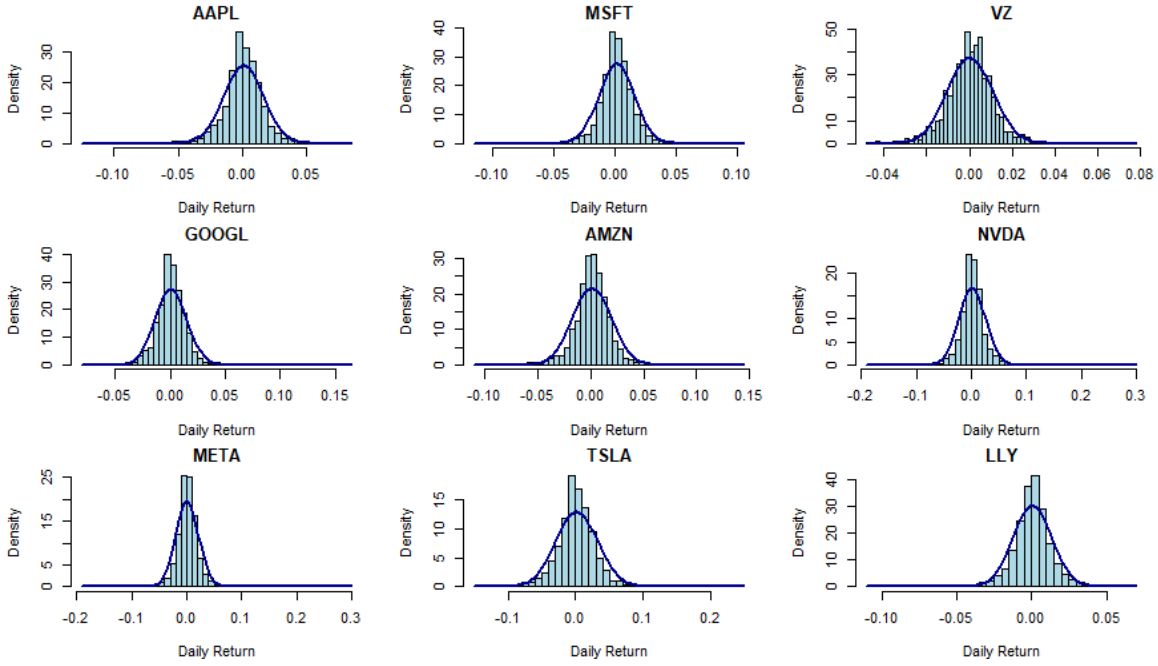

Figure 3: Daily returns of 9 large cap stocks with overlaid normal distribution curve, 2013-2019.

The term *copula* is introduced in Sklar (1959) and is described as a function that links a multivariate distribution to its one-dimensional marginal distributions. Sklar's theorem asserts that the joint distribution of multiple variables can be expressed by combining their individual marginal distribution functions with a copula that characterizes the dependence structure among these variables. Consider the random vector $X = (X_1, ..., X_d)$ and assume the marginal cumulative distribution functions (CDFs) $F_i(x) = Pr(X_i \leq x)$ are continuous ($1 \leq i \leq d$). By the probability integral transform, the marginals of the random vector

$$(U_1, ..., U_d) = (F_1(X_1), ..., F_d(X_d))$$

are distributed uniformly on $[0,1]$. The copula $C : [0,1]^d \rightarrow [0,1]$ of $(X_1, ..., X_d)$ is defined as the joint CDF of $U = (U_1, ..., U_d)$:

$$C(u_1, ..., u_d) = Pr[U_1 \leq u_1, ..., U_d \leq u_d].$$

In this formulation, C contains all information on the dependence structure between the separate variables in X whereas the information on the marginal distributions of X_i is in F_i ($1 \leq i \leq d$). We can reverse these steps to generate pseudo-random scenarios from a given multivariate probability distribution. That is, given samples $\hat{U} = (\hat{u}_1, ..., \hat{u}_d)$ generated from the

copula function, we can construct a scenario $\hat{X} = (\hat{x}_1, ..., \hat{x}_d)$ as

$$(\hat{x}_1, ..., \hat{x}_d) = (F_1^{-1}(\hat{u}_1), ..., F_d^{-1}(\hat{u}_d)).$$

By utilizing copulas to simulate joint stock return scenarios, we can adeptly represent the interdependence among individual stock returns without altering their individual marginal distributions. This approach is particularly advantageous as it facilitates adequately modeling the dynamic, non-linear cross-correlations within the multivariate series using copulas, while simultaneously accommodating variations in the marginal distributions of individual stocks. This approach enhances the realism of scenarios compared to modeling the joint distribution as, for instance, a multivariate normal distribution.

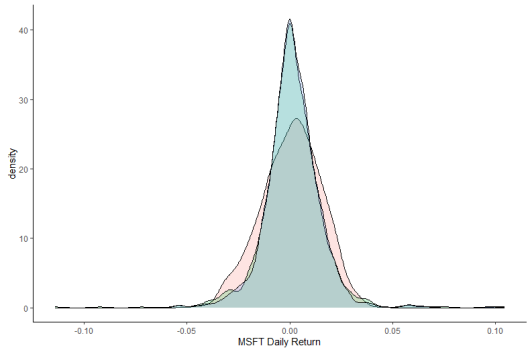

(a) Marginal Densities Microsoft Returns

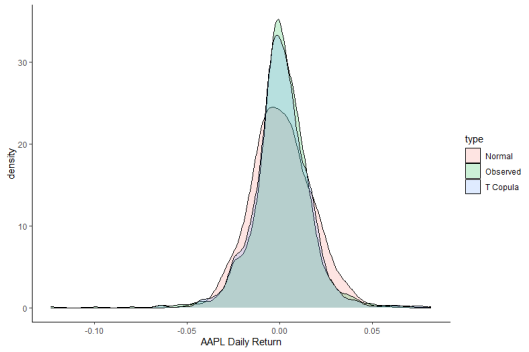

(b) Marginal Densities Apple Returns

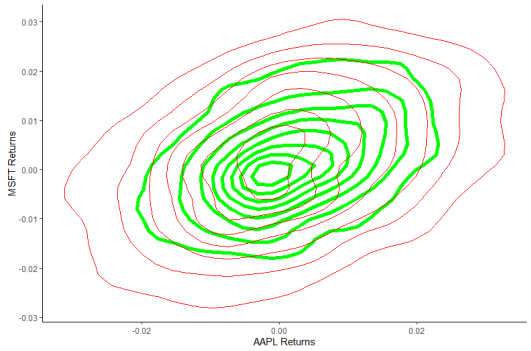

(c) Bivariate Density, observed (green) & bivariate normal (red)

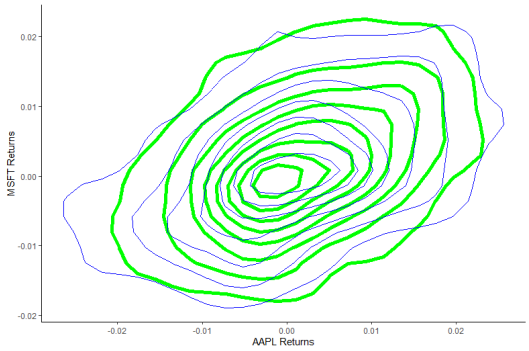

(d) Bivariate density, observed (green) & T copula (blue)

Figure 4: Marginal and bivariate densities Apple stock & Microsoft stock returns: observed & 1000 scenarios generated from multivariate normal distribution and T copula (2013-2019).

The copula-based method we employ is based on a bottom-up approach. We fit multivariate return data to four widely recognized parametric copula models: Gaussian, T, Clayton, and Gumbel (see Nelsen (1998) for an overview and fitting procedure). Consequently, the goodness-of-fit is assessed based on criteria such as (log) likelihood and Bayesian Information Criterion (BIC). Additionally, we employ the two-sided Kolmogorov-Smirnov (KS) test, which examines the null hypothesis that the data is drawn from the same distribution. This methodology is in line with the recommendations provided in Patton (2012). In Figure 4, we showcase the effectiveness of our approach through an illustrative example, focusing on modeling the bivariate distribution of stock returns for Apple and Microsoft. The historical returns, as depicted,

suggest a somewhat leptokurtic distribution characterized by heavy tails and positive kurtosis—typical traits of stock returns (McNeil et al., 2015). Through the bottom-up approach, the T Copula emerges as the best-fitting model for the return data. This selection is substantiated by both marginal and bivariate density plots. Notably, the T Copula significantly improves the distribution fit compared to a simple normal distribution approach, as evident in both the marginal and bivariate density plots.

This method has two major advantages over the approach presented in Kaut and Wallace (2011). Firstly, the scenario generation process is remarkably straightforward, eliminating the need for users to write algorithmic code. This simplicity allows for easy implementation in various statistical software applications. Secondly, the use of parametric copulas establishes a precise and well-defined mathematical framework, enhancing practicality across various applications. This is facilitated by the straightforward process of parameter estimation and the ease of interpreting results associated with this approach.

4.2 Third-Stage Scenarios

To model the distribution of the third-stage scenarios, a different approach is utilized. In an ideal situation, we would like the third-stage return scenarios to be predictions conditional on the second-stage return scenarios. However, as emphasized in the well-known paper on the efficient-market theory by Fama (1970), stock markets are efficient and predicting short-term stock price movements is impossible. The prevailing literature on SMs for the PSP also acknowledges this result, leading to the adoption of more time-efficient and simplified methods for modeling third-stage scenarios. Furthermore, given that the existing literature commonly employs a unified approach of computing the expected third-stage prices rather than working with the individual third-stage prices, one could argue that the quality of third-stage scenarios is of less significance compared to that of the second-stage scenarios.

In Najafi and Mushakhian (2015), third-stage scenarios are modeled as equiprobable bull and bear cases, where returns either shift up or down by multiplying a random parameter generated from a standard uniform distribution with the observed stock return variance and adding this to the average return in the second-stage scenario. In Topaloglou et al. (2008), the third-stage return scenarios are modeled by iteratively incorporating previously moment-matched generated scenarios at each stage. Cui et al. (2020) adopt a method where the third-stage return scenarios are determined by multiplying the preceding second-stage scenarios with a random parameter drawn from a uniform distribution within the interval $[0.9, 1.1]$. Our approach is similar to this method.

For each scenario $n \in N$, we generate J parameters $\lambda_{(n,j)} \sim \mathscr{U}(-0.01, 0.01)$ $(1 \leq j \leq J)$, with $\mathscr{U}$ the uniform distribution. These parameters are then used to construct third-stage prices $P_i^{(n,j)}(\xi_n) = P_i^n(\xi_n)(1 + \lambda_{(n,j)})$, $i = 1, ..., S$. This method is similar to the methods in Najafi and Mushakhian (2015) and Cui et al. (2020). Formally, for all $n = 1, ..., N$, write $\lambda_n = (\lambda_{(n,1)}, ..., \lambda_{(n,J)})$. The Law of Large Numbers informs us that $E[\lambda_n] = 0$ as $J \to \infty$. Consequently, for the two-stage SMIM in Section 2.5, we have that $\bar{P}^n \to P^n$ as $J \to \infty$. This observation underscores the earlier point in this section that, within the unified framework of computing expected third-stage scenarios in the existing literature, the quality of these scenarios may be of less significance than that of second-stage scenarios. This methodology introduces an intriguing dimension to the study, particularly in light of the three-stage SMIM presented in Section 2.6, which evaluates portfolio returns at the individual third-stage scenarios. The potential divergence in portfolio solutions between the two models under identical third-stage

scenarios could raise the argument for exploring enhanced methods for third-stage scenario generation in future research.

5 Data

The methodology developed for both the models and scenarios can now be merged to derive portfolio solutions. We will conduct an analysis on the portfolio solutions and scenario quality by deriving portfolio solutions for the Dow Jones Industrial Average (DJI) index. For this, we require return data for both the DJI index and the individual stocks constituting this index. Data is gathered from Yahoo Finance[1] and imported for analysis utilizing the *quantmod*-R package (Ryan, Ulrich, Smith, Thielen, Teetor, and Bronder, 2023).

The DJI, or Dow Jones Industrial Average, is a stock market index comprising 30 prominent large-cap companies listed on the American stock market. The index employs a price-weighted methodology, which implies that, in a simplified scenario, a stock with a price of $100 would carry ten times more weight in the DJI than a stock priced at $10. We opted for the DJI index due to its extensive data availability and the consistency in its holdings. By consistency, we refer to the fact that the composition of the DJI remains relatively constant over time, experiencing infrequent changes. This is in contrast to other indices, such as the S&P 500, which undergoes more frequent alterations in its constituent stocks. The stability in the composition of the DJI makes it a more convenient benchmark index for conducting assessments and analyses, providing a steady and consistent foundation for evaluation compared to indices with more dynamic compositions.

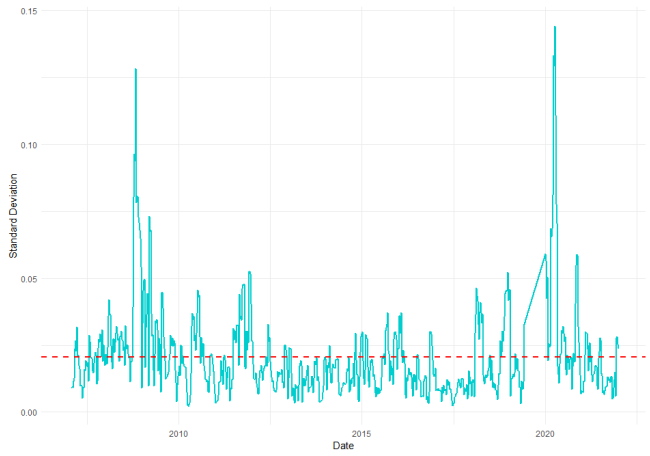

Figure 5: DJI weekly return volatility, 4-weeks rolling window, 2007-2021.

To justify a selection for the time-period considered in this study, we refer to the DJI return volatility depicted in Figure 5. We aim for a period that is both relatively stable and sufficiently large to provide ample information on the distribution moments of the returns. We decide to select the period from 2010 to 2018 as data used for generating scenarios and the period from 2018 to 2019 as data used for evaluation. This selection places our analysis between the extreme volatility peaks caused by the global financial crisis in 2008 and the emergence of

[1] `https://finance.yahoo.com/`

the COVID-19 pandemic in 2020. By excluding these extreme periods, we aim to mitigate the potential impact of outliers on the shape of the return distributions. It is worth noting that the chosen period does include some instances of higher volatility. However, this is not necessarily detrimental; in fact, data from these periods can effectively capture extreme returns in the tail of the return distributions, the importance of which is emphasized in Section 4.1.

In Table 1, we present the first four moments of the DJI returns Consistent with the analysis of stock return distributions outlined in Section 4.1, it is evident that the DJI return distribution exhibits excess kurtosis. Furthermore, given the negative skewness observed, it suggests that the excess kurtosis in the return distribution is mainly caused by a thicker left tail associated with the losses. This underscores the significance of employing an appropriate scenario generation procedure, as well as conducting an in-sample scenario analysis.

	DJI
Mean	0
Standard deviation	0.009
Kurtosis	3.940
Skewness	-0.380

Table 1: First four moments of DJI Returns, 2010-2018.

6 Scenarios

In this section, we implement the bottom-up copula selection approach to select the best-fitting copula. The resulting copula is then employed to generate second-stage scenarios of varying sizes. These scenario sets are subjected to both univariate and multivariate analyses to ascertain an empirically supported scenario sample size. Finally, the selected copula and scenario sample size are utilized to perform an in-sample analysis on the scenarios.

6.1 Copula Selection

We utilize the individual stock return data to fit the four proposed parametric copula models. The Log-Likelihood, AIC and BIC of the four fits is displayed in Table 2. For each fitted parametric model, we generate 1000 random weekly return second-stage scenarios, such that the two-side Kolmogorov-Smirnov (KS) test can be performed. Turning our attention to the p-values of the KS test in the last column, we observe that only the Gumbel Copula rejects the null hypothesis that the generated scenarios are from the same distribution as the original data. Subsequently, we focus on the Gaussian-, T-, and Clayton copulas. By examining the log-likelihood, AIC, and BIC, it becomes evident that the T copula emerges as the best-fitting parametric copula model. Consequently, the T-copula is selected for generating second-stage scenarios in the upcoming analyses.

6.2 Number of Scenarios

To justify a selection for the number of scenarios used, we consider the properties of differently sized scenario sets both from an univariate- and a multivariate point of view. We focus primarily on the properties of the second-stage scenarios, given that the subsequent scenarios are completely dependent on these scenarios, as shown in Section 4.2. For all second-stage

Copula	Log-Likelihood	AIC	BIC	KS
Gaussian Copula	2467.17	-4932.35	-4928.20	0.64
T Copula	2745.09	-5486.17	-5477.87	0.42
Gumbel Copula	1863.49	-3724.99	-3720.84	0.01
Clayton Copula	1907.46	-3812.91	-3808.76	0.53

Table 2: Parametric opula goodness-of-fit statistics.

scenarios $n = 1,...,N$, we set $J = 10$. We generate second-stage scenario sets from the fitted T copula with sizes of 1000, 2000, 3000, 4000, and 5000. Consequently, for each scenario set, the third-stage scenarios are generated in line with the first approach outlined in Section 4.2. To ensure robustness, this entire process is iterated 100 times.

To evaluate the scenario sets from a univariate perspective, we initially treat the generated multivariate scenario sets as combinations of univariate stock return scenarios. For each set, we calculate the first four moments of each of these univariate series. This process is also applied to the actual return data. Subsequently, we determine the average absolute distance between the moments of the scenario sets and the moments of the true data. The results of this process are presented in Table 3. The rationale behind this univariate analysis stems from the recognition of the importance of effectively capturing the moments of observed return distributions, as emphasized in Topaloglou et al. (2008). The objective is to analyze the differences in these moments across the various scenario sets. What we observe is that, starting from 4000, the differences start to become smaller and converge. We contend that the number of scenarios used should be set to either 4000 or 5000.

N	**Mean**	**Standard deviation**	**Skewness**	**Kurtosis**
1000	0.067%	0.081%	0.115	0.478
2000	0.051%	0.064%	0.085	0.435
3000	0.041%	0.060%	0.071	0.404
4000	0.035%	0.056%	0.065	0.402
5000	0.033%	0.055%	0.060	0.394

Table 3: Average absolute differences in moments of univariate stock return distributions, different scenario sizes; 100 replications.

To gain more insight, we also conduct an evaluation of the complete multivariate scenario sets. In this assessment, we solve the two-stage SMIM with $\alpha = 95\%$ under both the different scenario sets as well as the actual data for $\mu = 0.2\%$, $\mu = 0.3\%$, $\mu = 0.35\%$, $\mu = 0.38\%$ and $\mu = 0.4\%$. Linear interpolation on the obtained objective values is performed to plot the efficient frontiers generated from each scenario set. Without loss of generality, we will solely concentrate on the results and the associated solving times, omitting detailed examination of the portfolio solutions and other parameters. The efficient frontiers are illustrated in Figure 6 and the associated runtimes are presented in Table 4. Each point on the efficient frontier represents a specific portfolio allocation, and the curve itself depicts the trade-off between risk and return. Portfolios that lie on the efficient frontier are considered efficient because they provide the maximum return for a given level of risk or the minimum risk for a given level of return. The rationale for plotting different efficient frontiers against each other is to assess the similarity between the frontiers generated from the T copula scenario sets and those obtained using the

actual return data as second-stage scenarios. Ideally, we aim for the frontier of the selected T copula scenario set to lie relatively close to the frontier derived from the actual data. This proximity suggests that similar portfolio decisions are made, indicating that the generated multivariate scenario series likely stems from a distribution similar to that of the actual data. Based on this rationale, we choose to set $N = 5000$. By examining the information in Table 4, we contend that the runtimes and model size associated with $N = 5000$ are manageable. While the model size grows linearly with the number of scenarios, the runtime increases exponentially. In light of this, we argue against exploring scenario sets of any greater size. Therefore, in alliance with the insights from the univariate analysis, we opt to set $N = 5000$.

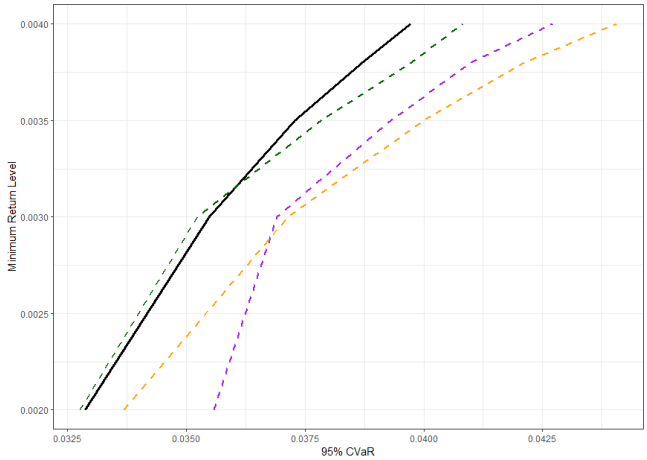

Figure 6: Efficient frontiers: actual data scenarios (black) &
$N = 3000$ (purple), $N = 4000$ (orange) and $N = 5000$ (green) T Copula scenarios.

	N **= 3000**	N **= 4000**	N **= 5000**
# Rows	249081	332081	415081
# Columns	321105	428105	535105
# Nonzeros	957260	1276260	1595260
Runtime (seconds)	139.82	184.74	327.55

Table 4: T Copula scenarios: problem sizes & runtimes.

6.3 In-Sample Scenario Analaysis

Now that we have determined a copula model and the number of scenarios, an in-sample analysis on the scenarios can be performed based on the tests proposed in Kaut, Wallace, Vladimirou, and Zenios (2007). In this analysis, we generate 10 scenario sets from the T copula with $N = 5000$. The objective is to demonstrate that, regardless of the specific scenario set chosen, the optimal objective value from the two-stage SMIM remains approximately the same. We argue that this implies the portfolio compositions under the different scenario sets are similar. The need for this additional in-sample analysis is underscored by Yamai and Yoshiba (2002), who identify a potential drawback of CVaR in that it can be notably unstable when considered with fat-tailed distributions. As evident, this is a characteristic often observed in

stock return distributions. If the in-sample analysis consistently yields stable optimal objective values across the different scenario sets, we contend that our method adequately captures the fat-tail property in stock return distributions.

In Table 5a, we present the complete results of the in-sample analysis for $\mu = 0.2\%$ and $\mu = 0.3\%$. Table 5b provides the relevant summary statistics. We argue that the objective values exhibit a desirable degree of stability for both values of μ. Notably, there is a slight uptick in instability for $\mu = 0.3\%$ compared to the results for $\mu = 0.2\%$. We attribute this increase in volatility to the fact that $\mu = 0.3\%$ represents an above-average minimum expected return rate. Specifically, given that the average weekly return rate for the DJI in the period 2010-2018 is approximately 0.2%. Consequently, we contend that this marginal rise in instability is not a cause for concern.

Table 5: In-sample scenario analysis.

Sample	$\mu = 0.2\%$	$\mu = 0.3\%$
1	3.24%	3.44%
2	3.19%	3.36%
3	3.12%	3.27%
4	3.47%	3.80%
5	3.17%	3.45%
6	3.25%	3.61%
7	3.27%	3.52%
8	3.21%	3.46%
9	3.22%	3.68%
10	3.03%	3.22%

(a) Model 95% CVaR values.

	$\mu = 0.2\%$	$\mu = 0.3\%$
Mean	3.22%	3.48%
Median	3.21%	3.45%
Standard deviation	0.11%	0.18%

(b) Average, median and standard deviation of 95% CVaR values.

7 Portfolio Results

The decisions made regarding the data and scenarios in the previous section will now be leveraged to generate portfolio solutions in various settings. We initiate this section with a direct comparison of the performance, composition, and characteristics of the obtained portfolio solutions and the actual market performance, composition, and characteristics of the DJI benchmark index. Subsequently, we perform a sensitivity analysis on these results for a set of relevant parameters. After this, we delve into an out-of-sample scenario analysis, wherein we compare the performance of the portfolio solutions generated under the proposed scenario generation method to the performance of the portfolio solutions generated under Monte Carlo scenarios. Following this, we investigate the impact of different rolling horizons on the performance of the portfolios. The section concludes by comparing the performance and compositions of the portfolios derived from the two-stage SMIM and the three-stage SMIM.

Throughout this analysis, if not mentioned differently, the user-specified parameters and deterministic input data is as follows: The critical percentage level for (C)VaR is set to $\alpha = 95\%$. We assume the initial portfolio only consists of cash, and we set the initial cash position to $C = 1,000,000$. The minimum position per stock is set to $w_{min} = 100$. M is set sufficiently large. Even though our model does allow for differing scenario probabilities for the scenarios, for simplicity, we assume that the probabilities for respectively the second-stage scenarios and

the third-stage scenarios are equal. That is, $p_n = \frac{1}{N} = \frac{1}{5000}$ and $p_{(n,j)} = \frac{1}{N}\frac{1}{J} = \frac{1}{5000}\frac{1}{10} = \frac{1}{50000}$. We set $K = 10$. All experiments have been carried out on a Windows machine with CPU AMD Ryzen 5 4500U @ 2.38GHz and 8 GB of RAM.

7.1 Benchmark Experiments

Figure 7 depicts the value of a hypothetical \$1 invested either directly in the DJI or in one out of four derived portfolio solutions under different μ values obtained from the two-stage SMIM. The obtained portfolio solutions are constant over time, that is, no rolling horizon is applied. Evident from the figure, in-line with the results in the existing literature, the portfolio solutions obtained from the two-stage SMIM either closely track or slightly outperform the DJI index. This is already a significant finding on its own, given that in the two-stage model, we can only select 10 out of 30 stocks to assign weights to. Moreover, we find that the return trajectories of the portfolios derived from the two-stage SMIM tend to be much more stable than the return trajectory of the DJI, suggesting that these portfolios likely exhibit significantly smaller risk levels.

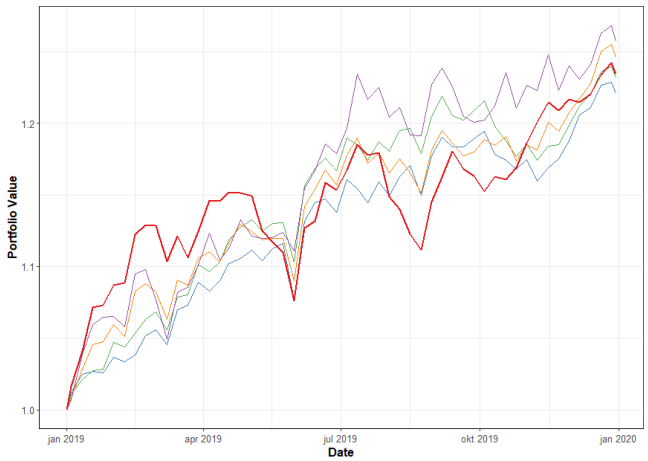

Figure 7: Value of \$1 invested in DJI (red) &
$\mu = 0.1\%$ (blue), $\mu = 0.2\%$ (green), $\mu = 0.3\%$ (orange), $\mu = 0.4\%$ (purple).

To validate this assertion, we present the summary statistics of the portfolios in Table 6. Extending on the existing literature, we augment these statistics by including metrics for the observed daily 95% CVaR. Clearly, all four derived portfolios consistently maintain a comparable weekly return level to the DJI index while demonstrating a markedly smaller 95% daily CVaR. This outcome is of considerable significance, suggesting that the two-stage SMIM we propose can allocate portfolio weights from the same set of stocks as the DJI in a more risk-efficient manner. Furthermore, we observe that the average weekly return of all four portfolios exceeds the pre-specified minimum expected portfolio return μ. Notably, the portfolios exhibit not only significantly reduced risk under the CVaR measure but also when considering standard deviation as a risk measure.

In a practical scenario, we argue that the company would have set $\mu = 0.2\%$, aligning with the approximate average weekly return observed in the historical return dataset. If the company were to construct its portfolio based on the recommendations of our proposed two-stage

SMIM, it would have achieved a weekly return rate close to that of the DJI, surpassing the pre-specified return rate μ. Moreover, the daily 95% CVaR for their portfolio would have been 0.5% lower than the daily 95% CVaR observed in the DJI returns. This reduction in CVaR, particularly on an initial portfolio value of e.g. $1,000,000 and larger, holds significant financial implications.

	$\mu = 0.1\%$	$\mu = 0.2\%$	$\mu = 0.3\%$	$\mu = 0.4\%$	**DJI**
Average weekly return	0.38%	0.39%	0.42%	0.44%	0.40%
Daily 95% CVaR	-1.48%	-1.47%	-1.53%	-1.83%	-1.97%
Daily S.D.	0.65%	0.65%	0.66%	0.77%	0.79%

Table 6: Benchmark portfolio performances: summary statistics.

7.2 Sensitivity Analysis

In this section, we perform sensitivity analysis on the values of α, K and $\tilde{w}_{\min}$. In this analysis, we place emphasis not only on the performance metrics of the derived portfolios but also on scrutinizing the composition of these portfolios. Performing sensitivity analysis on SMIMs is crucial, given the inherent uncertainties in such models. By analyzing to what extend the portfolio solutions are affected by changes in user-specified parameters, we can say something about the robustness of the SMIM.

In Table 7, portfolio statistics for different values of α are presented. Immediately, we observe that the portfolio statistics are stable for all considered values of α, indicating that the model tends to be robust to changes in α. Considering that in the two-stage SMIM, the α values determine the objective of the model, it seems fair to not only consider the observed daily 95% CVaR, but also the daily 90% CVaR and the daily 99% CVaR. It is noteworthy that the most favorable realized daily α%-CVaR consistently emerges from the SMIM solved under the corresponding α value. This observation emphasizes the capability of the proposed two-stage SMIM not only to optimize the CVaR level within the model but also to generate portfolios that are optimal for the observed α%-CVaR. Adding to this observation, even though differences are small, we would also like to mention that it makes sense that the average weekly return seems to be decreasing as the α level increases. This makes sense, as a higher critical CVaR level typically is associated with a lower expected return because the optimization strategy aims to control the tail risk, potentially sacrificing some average return to achieve a more robust portfolio against extreme events. Conversely, a lower critical level would imply a less conservative stance, allowing for higher expected returns but potentially exposing the portfolio to larger losses in tail scenarios. This is also what we see for the daily 99% CVaR values. The portfolio solution generated under $\alpha = 99\%$ has a significantly smaller observed daily CVaR level compared to the portfolios generated under $\alpha = 95\%$ and $\alpha = 90\%$.

In Table 8, we present portfolio statistics for different values of K. As anticipated, there is more variability in the portfolio statistics across the various K values. Specifically, for $K = 4$, it's notable that the two-stage SMIM fails to select four stocks in a manner that approximates the return level of the benchmark index. However, starting from $K = 6$ and extending up to $K = 16$, we observe a general consistency in the return levels. Significantly, the daily 95% CVaR levels for all values of K are notably smaller compared to the daily 95% CVaR of the benchmark index. This robust observation provides a high level of confidence that the two-stage SMIM, under our proposed scenario-generation methodology, consistently generates portfolios that significantly

enhance risk efficiency in comparison to the benchmark index.

Table 9 showcases portfolio statistics for various values of $\tilde{w}_{\min}$. In general, the findings align with what has been observed in the sensitivity analysis of α and K. It is intriguing to note that, in this specific setting, the highest level of risk-efficiency is associated with higher values of $\tilde{w}_{\min}$. We posit that this is attributed to the compelled higher level of diversification linked with an increased $\tilde{w}_{\min}$. This concept aligns with contemporary portfolio diversification theory, which posits that a well-diversified portfolio typically mitigates associated risk metrics.

	$\boldsymbol{\alpha = 90\%}$	$\boldsymbol{\alpha = 95\%}$	$\boldsymbol{\alpha = 99\%}$
Average Weekly Return	0.39%	0.39%	0.38%
Daily 90% CVaR	-1.15%	-1.17%	-1.15%
Daily 95% CVaR	-1.47%	-1.47%	-1.48%
Daily 99% CVaR	-1.98%	-2.00%	-1.89%
Daily S.D.	0.65%	0.65%	0.65%

Table 7: Portfolio statistics for different values of α.

	$\mathbf{K = 4}$	$\mathbf{K = 6}$	$\mathbf{K = 8}$	$\mathbf{K = 10}$	$\mathbf{K = 12}$	$\mathbf{K = 14}$	$\mathbf{K = 16}$
Average Weekly Return	0.29%	0.39%	0.37%	0.39%	0.31%	0.36%	0.37%
Daily 95% CVaR	-1.42%	-1.55%	-1.44%	-1.47%	-1.43%	-1.43%	-1.44%
Daily S.D.	0.65%	0.68%	0.65%	0.65%	0.63%	0.63%	0.65%

Table 8: Portfolio statistics for different values of K.

	$\boldsymbol{\tilde{w}_{\min} = 100}$	$\boldsymbol{\tilde{w}_{\min} = 200}$	$\boldsymbol{\tilde{w}_{\min} = 300}$	$\boldsymbol{\tilde{w}_{\min} = 400}$	$\boldsymbol{\tilde{w}_{\min} = 500}$
Average Weekly Return	0.39%	0.36%	0.38%	0.39%	0.41%
Daily 95% CVaR	-1.47%	-1.41%	-1.43%	-1.43%	-1.44%
Daily S.D.	0.65%	0.62%	0.63%	0.63%	0.64%

Table 9: Portfolio statistics for different values of $\tilde{w}_{\min}$

We now move attention to the changes in portfolio composition under different values of α, K and $\tilde{w}_{\min}$. The changes are visualized in Figure 8. Eyeballing the figures reveals some interesting trends. A noteworthy observation, consistent with the assertion made in Cui et al. (2020), is the prevalence of analogous structures in efficacious solutions. Irrespective of the parameter values considered, there tends to be a regularity wherein certain stocks consistently are assigned substantial weights in the two-stage SMIM solution. Examining Figure 8b, it becomes apparent that higher values of $\tilde{w}_{\min}$ correspond to a tendency for increased diversification. In other words, stocks, on average, tend to be assigned higher weights, supporting the earlier discussion on the benefits of diversification. Intriguingly, the anticipated increase in diversification associated with higher values of K does not materialize as one might expect, as illustrated in Figure 8c. Instead, a discernible pattern emerges where a specific group of 5 stocks consistently receives substantial weights, while others are allocated the minimum weight, $\tilde{w}_{\min}$.

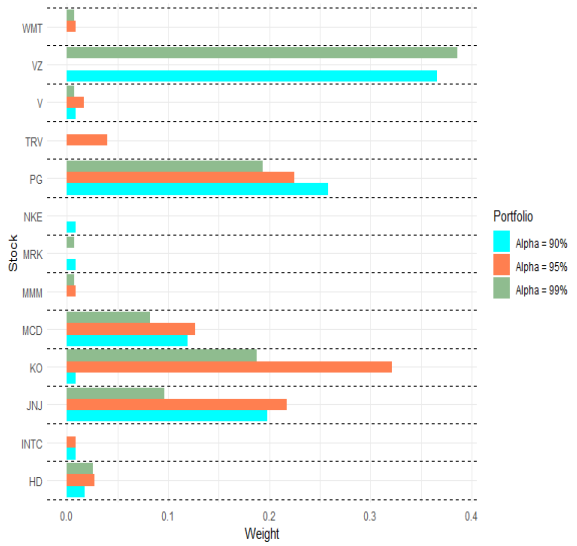

(a) Portfolio compositions for different values of α.

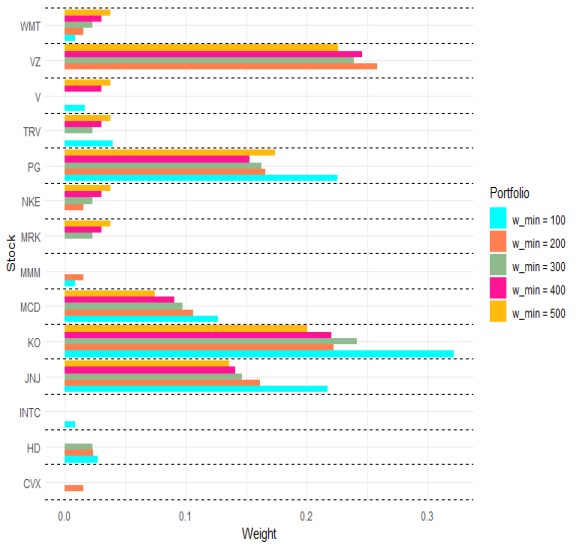

(b) Portfolio compositions for different values of $\tilde{w}_{\min}$.

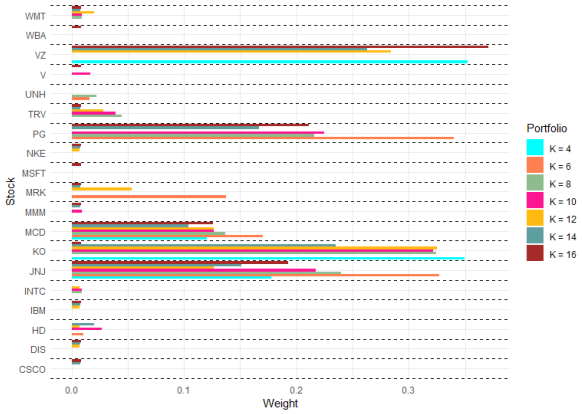

(c) Portfolio compositions for different values of K.

Figure 8: Portfolio compositions.

In concluding this section, we contend that, when K attains a sufficiently large value, the precise selection of K exerts minimal influence on both the resulting portfolio compositions and their performances. Similarly, regarding composition, an analogous assertion can be made for α. However, in the case of α, it is prudent to emphasize the significance of aligning α with the true critical level that the company aims to minimize CVaR for. Our analysis reveals that setting α to this critical level in the two-stage SMIM indeed yields portfolios with the most favorable empirical CVaR performance. In instances where a well-diversified portfolio is a paramount objective, attention to adjusting the parameter $\tilde{w}_{\min}$ to a desired level becomes essential.

7.3 Out-of-Sample Scenario Analysis

In Section 6.3, we performed an in-sample analysis on the scenarios, to assess the stability of the scenarios within the model. The out-of-sample scenario analysis in this section assesses the quality of the scenarios by comparing the performance of the portfolios generated under our proposed scenario-generation method with the performance of the portfolios generated under a scenario-generation method based on Monte Carlo (MC) sampling of the historical return data. In the existing literature, out-of-sample scenario analysis is typically performed by analyzing the portfolio performance similar to what we have done in the previous section. Then, it is argued that, given the portfolio solutions derived from their model are able to track some benchmark index, the scenarios are of high quality. From this rationale, the previous section suggests that our proposed scenario-generation methodology is able to generate scenarios of high quality. However, we believe that it is desirable to also provide results based on a different scenario generation process, such that the effect of the proposed scenario-generation method on the portfolio performance can be better isolated.

Utilizing MC sampling on the historical return data, we generated 1800 second-stage weekly return scenarios. Similar to the process for the copula second-stage scenarios, for each second-stage weekly return scenario, we generated 10 third-stage scenarios using the first method proposed in 4.2. In Table 10, the portfolio performances under these scenarios are juxtaposed with the portfolio performances under the copula scenarios for different values of μ. We find that the average weekly return rates are quite similar under the two scenario methods. However, portfolios generated under the copula scenarios exhibit a significantly smaller daily 95% CVaR in three out of the four cases. This reduction in CVaR underscores that, despite the copula method being less straightforward than a simpler method such as MC sampling historical data, it is advantageous due to its ability to produce more risk-efficient portfolio solutions.

	Copula	**MC**
Avg. Weekly Return	0.38%	0.36%
Daily 95% CVaR	-1.48%	-1.63%
Daily S.D.	0.65%	0.68%

(a) $\mu = 0.1\%$.

	Copula	**MC**
Avg. Weekly Return	0.39%	0.40%
Daily 95% CVaR	-1.47%	-1.53%
Daily S.D.	0.65%	0.65%

(b) $\mu = 0.2\%$.

	Copula	**MC**
Avg. Weekly Return	0.42%	0.42%
Daily 95% CVaR	-1.53%	-1.68%
Daily S.D.	0.66%	0.70%

(c) $\mu = 0.3\%$.

	Copula	**MC**
Avg. Weekly Return	0.44%	0.41%
Daily 95% CVaR	-1.83%	-1.73%
Daily S.D.	0.77%	0.75%

(d) $\mu = 0.4\%$.

Table 10: Portfolio statistics: copula scenarios & MC scenarios.

7.4 Rolling Horizon

Up to this point, we have only considered static portfolio solutions. However, given the nature of our problem setting, it is natural to analyze the impact of a rolling horizon on the portfolio performances. One might anticipate that, considering our model assumes the company has the flexibility to adjust its portfolio over time, employing a rolling horizon where portfolio compositions dynamically evolve throughout the evaluation period would result in more favorable

performance compared to portfolios that remain constant. This notion is well-acknowledged in the existing literature on the PSP in stochastic settings, where obtaining portfolio solutions within a rolling horizon is common. In fact, employing a rolling horizon is common in stochastic programming in general, given that it can essentially be considered an approximation of a multi-stage model. It allows for a nuanced exploration of the potential effects of additional stages. In our study, the rolling horizon solution approach involves solving the two-stage SMIM after every pre-specified time interval. During each iteration, the initial portfolio is updated with the portfolio from the previous period, and prices are adjusted to reflect the new time point. While existing literature often opts for a single rolling horizon period frequency, we aim to provide additional insights by exploring various rolling horizon frequencies. This investigation not only sheds light on the impact of a rolling horizon approach on portfolio performance but also allows us to gain insights into how the choice of the period affects performance.

Figure 9 showcases the evolution of portfolio composition within a monthly rolling horizon. Notably, while the overall composition of the portfolio generally remains consistent over time, there is discernible active rebalancing in every month. This implies that the composition established in the preceding month is likely sub-optimal in the current month. Consequently, it suggests a preference for a dynamic portfolio rebalancing strategy over a static approach.

To validate this assertion, we present portfolio statistics for $\mu = 0.2\%$ under various rolling horizons in Table 11. It should be noted that the yearly rolling horizon statistics are identical to those in Section 7.1, as both evaluate performance over a one-year period. A consistent observation is the improvement in both CVaR and weekly return levels when implementing a rolling horizon strategy compared to a static portfolio strategy. Consequently, the adoption of a rolling horizon strategy consistently contributes to the development of more risk-efficient portfolio solutions. Focusing on the differences in results under the selected rolling horizon frequencies, the differences are negligible. However, the rolling horizon results can be considered as a robustness check on our static results. That is, even if we solve the model 12 times in the monthly rolling horizon under different prices, the portfolios generated consistently yield stable returns and CVaR levels. Concluding, the results in this section suggest to employ a dynamic portfolio rebalancing strategy over a static portfolio strategy.

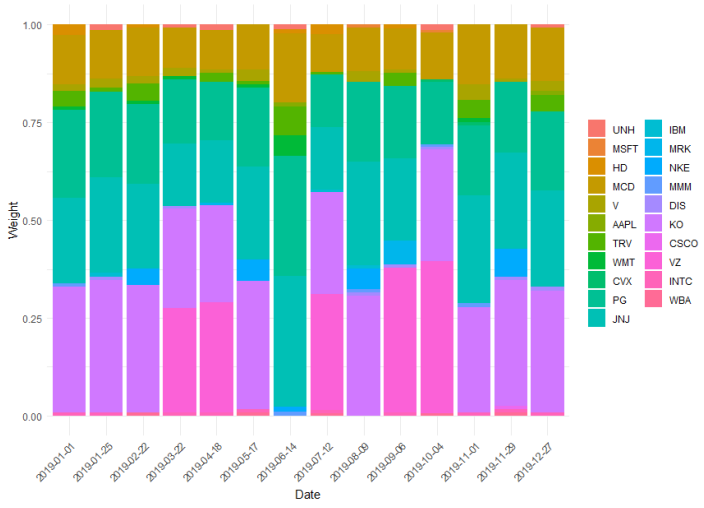

Figure 9: Portfolio composition over time, $\mu = 0.2\%$, monthly rolling horizon.

	Yearly	Half Yearly	Quarterly	Monthly
Average Weekly Return	0.39%	0.42%	0.41%	0.41%
Observed Daily 95% CVaR	-1.47%	-1.44%	-1.46%	-1.46%
Observed Daily S.D.	0.65%	0.63%	0.63%	0.65%

Table 11: Portfolio statistics for $\mu = 0.2\%$, different rolling horizons.

7.5 Three-Stage SMIM

We now move attention to portfolios generated from the three-stage model proposed in Section 2.6. Focus will be laid not only on the performance of the portfolios generated, but also on the impact on the portfolio composition compared to the portfolio composition generated with the two-stage model. Figure 10 displays the differences in portfolio compositions between the portfolios obtained from the two-stage model and the portfolios obtained from the three-stage model. In this comparison, the exact same second-stage- and third-stage scenarios are used such that the three-stage extension effect is isolated. It is clear that the three-stage extension changes the optimal portfolio composition. This suggests that it might be meaningful to consider third-stage scenarios of proven high quality, given that the optimal portfolio allocations clearly depend on the third-stage scenarios.

To evaluate the impact of the extension on portfolio performance, Table 12 compares the outcomes of portfolios obtained from the two-stage model with those derived from the three-stage model. While the weekly return levels show similarity, consistently smaller CVaR values are observed for portfolios derived from the three-stage model. This suggests that the three-stage model may be capable of producing portfolios with enhanced risk-efficiency compared to those generated from the two-stage model. Further exploration into the influence of improved third-stage scenario generation techniques on these results would be intriguing and is an avenue deserving of additional research.

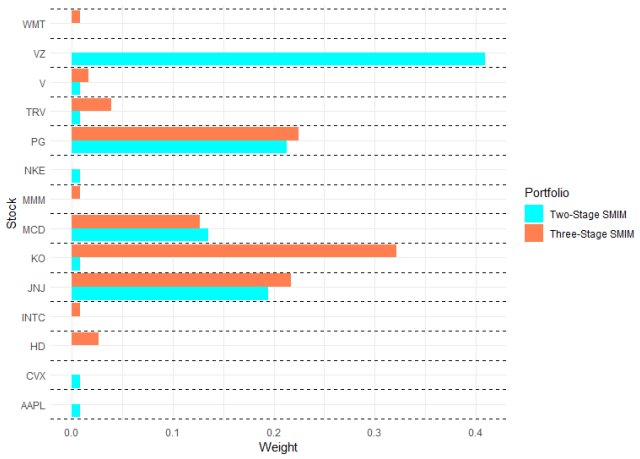

Figure 10: Portfolio compositions: two-stage SMIM- (blue) & three-stage SMIM (red), $\mu = 0.2\%$.

	$\mu = 0.1\%$	$\mu = 0.2\%$	$\mu = 0.3\%$	$\mu = 0.4\%$
Average Weekly Return Two-Stage SMIM	0.38%	0.39%	0.42%	0.44%
Average Weekly Return Three-Stage SMIM	0.37%	0.38%	0.41%	0.44%
Observed Daily 95% CVaR Two-Stage SMIM	-1.48%	-1.47%	-1.53%	-1.83%
Observed Daily 95% CVaR Three-Stage SMIM	-1.47%	-1.44%	-1.47%	-1.82%

Table 12: Portfolio performance for different μ values, yearly rolling horizon.

8 Conclusion & Discussion

In this study, we underscore the significance of considering market uncertainty in portfolio allocation by employing a two-stage SMIM for the PSP. This model aims to minimize the CVaR of the portfolio while adhering to a set of practical constraints. The model integrates market uncertainty by representing future asset prices through a scenario tree. Our proposed approach utilizes a parametric copula-based method for generating scenarios within this tree. We illustrate that this method effectively captures the excess kurtosis evident in both univariate and multivariate asset return series. Moreover, the portfolio solutions derived under these scenarios are stable, establishing its suitability for use in a CVaR minimization model. The methodology for both the models and scenarios are merged to derive portfolio solutions for the stocks that contribute to the DJI index. Consistent with the existing literature, our findings demonstrate that the obtained portfolio solutions effectively track the DJI index, utilizing only a third of the stocks included in the DJI.

A key contribution of our study is that, despite the derived portfolio solutions demonstrating similar return levels to the DJI index, those generated from our two-stage model exhibit substantially smaller CVaR values. These results are robust to changes in parameters. This highlights the ability of our model to produce portfolios that are more risk-efficient, which has not been empirically demonstrated before in similar literature. Moreover, we demonstrate that by employing a rolling horizon strategy to solve our two-stage model, risk-efficiency is enhanced even more. This indicates that portfolios generated from our model exhibit optimal performance when implemented under a dynamic portfolio rebalancing strategy.

Another key contribution of our study is the inclusion of a three-stage extension to the two-stage SMIM. This extension shifts the objective from minimizing the CVaR associated with the average expected future returns in the two-stage model to minimizing the CVaR linked with the true expected future returns. Despite the substantial increase in the size of the problem introduced by this extension, our findings reveal that the three-stage model generates portfolios that further improve risk-efficiency.

The last significant contribution is the proposal of a decomposition algorithm for the two-stage SMIM. Two-stage SMIMs that address the PSP typically do not adhere to a standard structure utilized in many efficient solving algorithms, such that it requires for novel techniques. Where the existing literature employs methods that yield approximate solutions, the decomposition algorithm we propose is able to yield exact solutions.

Despite the significant findings, we do acknowledge certain limitations in this study. Firstly, while comparing portfolios derived from our model to the DJI index is a rational and well-justified approach given the inherent nature of the actual data, it also presents itself as the sole rational method. Comparing the derived portfolios to those of other two- and multi-stage PSP models is deemed impractical. This impracticality stems from the inherent stochasticity in

such models and the diverse modeling decisions made, rendering a direct parallel comparison with other methods in the literature exceedingly challenging. Therefore, it is challenging to make a direct comparison to determine if one model performs "better" than another model. Secondly, it is impossible to include all practical constraints investors face in the real-world market in our model. It could be that the company we consider in the problem setting faces other constraints that could influence the optimal portfolio allocation. Lastly, it is important to note that the solution algorithm presented in this study has not undergone numerical testing. Consequently, we are unable to provide a comprehensive assessment of its performance relative to existing algorithms employed in analogous models.

Building on the findings of our study, we propose several potential avenues for future research. Firstly, considering the effectiveness of a rolling horizon strategy in obtaining portfolio solutions, it prompts the exploration of incorporating additional stages in the original two-stage model. This suggestion arises from the notion that portfolio solutions under a rolling horizon strategy can be viewed as approximations of a multi-stage model. To our knowledge, no studies exist that solve a multi-stage SMIM for the PSP, so this is definitely an area that warrants further research. Secondly, considering the influence of the quality of third-stage scenarios on the derived portfolio solutions when extending the original two-stage model to a three-stage model, we propose further research and evaluation studies focused on asset return scenarios conditional on the last observed asset price, given that this could lead to even more risk-efficient portfolio solutions. Lastly, it is worth highlighting the significance of obtaining numerical results for the decomposition algorithm introduced in this study. This aspect stands as a compelling avenue for exploration, and we plan to delve into it in future research.

References

Artzner, P., F. Delbaen, J.M. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance 9*(3), 203–228.

Benders, J.F. (1962). Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik 4*(3), 238–252.

Bienstock, D. (1996). Computational study of a family of mixed-integer quadratic programming problems. *Mathematical Programming 74*(2), 121–140.

Bonami, P. and A. Lejeune (2009). An exact solution approach for portfolio optimization problems under stochastic and integer constraints. *INFORMS 57*(3), 650–670.

Cesarone, F. and F. Tardella (2017). Equal risk bounding is better than risk parity for portfolio selection. *Journal of Global Optimization 68*, 439–461.

Cui, T., R. Bai, S. Ding, A.J. Parkers, R. Qu, F. He, and J. Li (2020). A hybrid combinatorial approach to a two-stage stochastic portfolio optimization model with uncertain asset prices. *Soft Computing 24*, 2809–2831.

Elton, E.J., M.J. Gruber, and J. Spitzer (2006). Improved estimates of correlation coefficients and their impact on optimum portfolios. *European Financial Management 12*(3), 303–318.

Fahmy, H. (2020). Mean-variance-time: An extension of markowitz's mean-variance portfolio theory. *Journal of Economics and Business 109*, 105888.

Fama, E.F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance 25*(2), 383–417.

Gao, J. and D. Li (2013). Optimal cardinality constrained portfolio selection. *INFORMS 61*(3), 745–761.

Gurobi Optimization, LLC (2023). Gurobi Optimizer Reference Manual.

Gülten, S. and A. Ruszczyński (2015). Two-stage portfolio optimization with higher-order conditional measures of risk. *Annals of Operations Research 229*, 409–427.

Haneveld, K., M.H. van der Vlerk, and W. Romeijnders (2020). *Stochastic Programming*. Springer International Publishing.

He, F. and R. Qu (2014). A two-stage stochastic mixed-integer program modelling and hybrid solution approach to portfolio selection problems. *Information Sciences 289*, 190–205.

Høyland, K., M. Kaut, and S.W. Wallace (2003). A heuristic for moment-matching scenario generation. *Computational Optimization and Applications 24*(2), 169–185.

Høyland, K. and S.W. Wallace (2001). Generating scenario trees for multistage decision problems. *Management Science 47*(2), 295–307.

Jobson, J.D. and B. Korkie (1980). Estimation for markowitz efficient portfolios. *Journal of the American Statistical Association 75*(371), 544–554.

Kaut, M. and S. Wallace (2011). Shape-based scenario generation using copulas. *Computational Management Science 8*(1), 181–199.

Kaut, M., S. Wallace, H. Vladimirou, and S. Zenios (2007). Stability analysis of a portfolio management model based on the conditional value-at-risk measure. *Quantitative Finance 7*, 397–409.

Li, C. and E. Grossmann (2019). A generalized benders decomposition-based branch and cut algorithm for two-stage stochastic programs with nonconvex constraints and mixed-binary first and second stage variables. *Journal of Global Optimization 75*, 247–272.

Li, J. and J. Xu (2013). Multi-objective portfolio selection model with fuzzy random returns and a compromise approach-based genetic algorithm. *Information Sciences 220*, 507–521.

Li, X., E. Armagan, and P.I. Barton (2011). Stochastic pooling problem for natural gas production network design and operation under uncertainty. *AIChE Journal 57*(8), 2120–2135.

Mansini, R. and M.G. Speranza (1999). Heuristic algorithms for the portfolio selection problem with minimum transaction lots. *European Journal of Operational Research 114*(2), 219–233.

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

McNeil, A.J., R. Frey, and P. Embrechts (2015). *Quantitative Risk Management*. Princeton Series in Finance.

Mehrotra, S. and D. Papp (2013). Generating moment matching scenarios using optimization techniques. *SIAM Journal on Optimization 23*(2), 963–999.

Merton, R.C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics 8*(4), 323–361.

Najafi, A.A. and S. Mushakhian (2015). Multi-stage stochastic mean–semivariance–cvar portfolio optimization under transaction costs. *Applied Mathematics and Computation 256*, 445–458.

Nelsen, R.B. (1998). *An introduction to copulas*. Springer.

Patton, A.J. (2012). A review of copula models for economic time series. *Journal of Multivariate Analysis 110*, 4–18.

Quaranta, A.G. and A. Zaffaroni (2008). Robust optimization of conditional value at risk and portfolio selection. *Journal of Banking & Finance 32*(10), 2046–2056.

Rockafellar, R.T. and S. Uryasev (2000). Optimization of conditional value-at-risk. *Journal of Risk 2*, 21–41.

Rockafellar, R.T. and S. Uryasev (2002). Conditional value-at-risk for general distributions. *Journal of Banking & Finance 26*(7), 1443–1471.

Ryan, J.A., J.M. Ulrich, E.B. Smith, W. Thielen, P. Teetor, and S. Bronder (2023). Package 'quantmod'.

Sandoval, L. and I. De Paula Franca (2012). Correlation of financial markets in times of crisis. *Physica A: Statistical Mechanics and its Applications 391*(1), 187–208.

Schutlz, R., L. Stougie, and M.H. van der Vlerk (1998). Solving stochastic programs with integer recourse by enumeration: A framework using gröbner basis. *Mathematical Programming 83*, 229–252.

Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut Statistique de l'Université de Paris 8*, 229–231.

van Slyke, R.M. and R. Wets (1969). L-shaped linear programs with applications to optimal control and stochastic programming. *Applied Mathematics 17*(4), 638–663.

Stoyan, S.J. and R.H. Kwon (2011). A stochastic-goal mixed-integer programming approach for integrated stock and bond portfolio optimization. *Computers & Industrial Engineering 61*(4), 1285–1295.

Stoyanov, S.V., S.T. Rachev, and F.J. Fabozzi (2013). Sensitivity of portfolio var and cvar to portfolio return characteristics. *Annals of Operations Research 205*, 169–187.

Topaloglou, N., H. Vladimirou, and S.A. Zenios (2008). A dynamic stochastic programming model for international portfolio management. *European Journal of Operational Research 185*(3), 1501–1524.

Werner, A., K.T. Uggen, M. Fodstad, A-G. Lium, and R. Egging (2011). Stochastic mixed-integer programming for integrated portfolio planning in the lng supply chain. *The Energy Journal 35*(1), 2120–2135.

Woodside-Oriakhi, M., C. Lucas, and J.E. Beasley (2013). Portfolio rebalancing with an investment horizon and transaction costs. *Omega 41*(2), 406–420.

Yamai, Y. and T. Yoshiba (2002). Comparative analyses of expected shortfall and value-at-risk: Their estimation error, decomposition, and optimization. *Monetary and Economic Studies 20*, 87–121.

Appendix

A Introduction to two-stage recourse models

This section commences with a concise theoretical overview on the class of two-stage linear mixed-integer recourse (MIR) models. In upcoming sections, we will find that this class of models naturally arises in the context of a PSP with market uncertainty. The general formulation for a two-stage MIR model is given by the minimization problem

$$\min_{x \in \mathscr{X}} \{c^\top x + Q(x) \ : \ Ax = b\} \tag{67}$$

Here, $Q(x)$ is referred to as the expected value function (EVF)

$$Q(x) \ := \ E_\omega[v_\omega(x)] \tag{68}$$

, where $E[v_\omega(x)]$ is the mathematical expectation of the second-stage cost function

$$v_\omega(x) \ := \ \min_{y \in \mathscr{Y}} \{q_\omega^\top y \ : W_\omega y = h_\omega - T_\omega x\} \tag{69}$$

For n_1, n_2, p_1, p_2 nonnegative integers with $p_1 \leq n_1$ and $p_2 \leq n_2$, $x \in \mathscr{X} \subset \mathbb{R}_+^{n_1 - p_1} \times \mathbb{Z}_+^{p_1}$ is the first-stage decision variable vector and $y \in \mathscr{Y} \subset \mathbb{R}_+^{n_2 - p_2} \times \mathbb{Z}_+^{p2}$ is the second-stage decision variable vector. The first-stage cost vector is given by $c \in \mathbb{R}^{n_1}$ and the second-stage cost (recourse) vector is given by $q_\omega \in \mathbb{R}^{n_2}$ for all $\omega \in \Omega$, where Ω is the set of possible realizations of the random vector consisting of uncertain second-stage data ω with known probability distribution $\mathbb{P}$. In (67), the right-hand side represent the constraints related to the first-stage decisions. In (69), T_ω is the (random) technology matrix related to the first stage decisions x, W_ω is the (random) recourse matrix related to the second stage decisions y and h_ω the (random) right-hand side vector. Collectively, they formulate the constraints for the second-stage cost function $v_\omega(x)$. From (67), we observe that in the first stage, we minimize the costs associated with the first-stage decisions plus the expected costs of the optimal second-stage decisions. We can consider the solution to the second-stage problem (69) as a recourse action, where $W_\omega y$ compensates for violations in the random goal constraints $T_\omega x = h_\omega$ and $q_\omega^\top y$ are the associated costs of this recourse action.

Typically, in line with our approach in this paper, we assume that the random vector ω has a finite number of scenarios $\omega_1, .., \omega_N$ with corresponding probability masses $p_1, .., p_N$. With slight abuse of notation, let us write $v_{\omega_n} = v_n$ for the second-stage cost function in scenario n, $n \in 1, ..., N$. Then, we can reformulate the EVF in (68) as

$$Q(x) := \sum_{n=1}^{N} p_n v_n(x). \tag{70}$$

This notation allows us to formulate the minimization problem in (67) as one large linear mixed-integer programming (MIP) problem. This is referred to as the deterministic equivalent of the model, and, assuming N finite scenarios, is of the form

$$\begin{aligned} \min_{x \in \mathscr{X}, y_1, .., y_N \in \mathscr{Y}} \quad & c^\top x + q_1^\top y_1 + q_2^\top y_2 + \cdots + q_N^\top y_N \\ & Ax \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad = b \end{aligned} \tag{71}$$

$$
\begin{array}{lllll}
T_1 x + W_1 y_1 & & & & = h_1 \\
T_2 x & + W_2 y_2 & & & = h_2 \\
\vdots & & \ddots & & \vdots \\
T_N x & & & + W_N y_N & = h_N .
\end{array}
$$

The model presented in (71) takes the form of a conventional linear mixed-integer program (MIP), which can be effectively solved utilizing established solvers. However, it must be noted that as the number of scenarios increases significantly, these kind of models can become very large, posing heavy computational challenges that make it impractical to solve the model as a linear MIP. In the case where $p_2 = 0$ such that the set $\mathscr{Y}$ contains no integer restrictions, Q is convex (Haneveld, van der Vlerk, and Romeijnders, 2020). Numerous decomposition algorithms have been developed that exploit this convexity, with one of the most well-known being the L-shaped algorithm (van Slyke and Wets, 1969). However, for the case where $p_2 > 0$ such that the set $\mathscr{Y}$ can contain both continuous and integer restrictions, these conventional algorithms can not be utilized, as this often leads to non-convexity of Q (Schutlz et al., 1998). For an overview of algorithms and corresponding assumptions that can be utilized to handle the non-convexity of Q, we refer to Li and Grossmann (2019).