

Classical Decomposition of Markowitz Portfolio Selection

by Christopher Demone,¹ Olivia Di Matteo^{2, 3, 4} and
Barbara Collignon^{1, 4}

¹Quantum Lab
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9
cdemone@bankofcanada.ca

²TRIUMF

³Department of Physics and Astronomy, University of Waterloo

⁴Institute for Quantum Computing, University of Waterloo



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Acknowledgements

We would like to thank the Bank of Canada's Narayan Bulusu for financial data and useful discussions, Ivan Komarov for additional comments and suggestions, and Vladimir Skavysh for his guidance and support. Funding for Olivia Di Matteo was provided by the National Research Council of Canada. The Institute for Quantum Computing is supported in part by the Government of Canada and the Province of Ontario.

Abstract

In this study, we enhance Markowitz portfolio selection with graph theory for the analysis of two portfolios composed of either EU or US assets. Using a threshold-based decomposition of their respective covariance matrices, we perturb the level of risk in each portfolio and build the corresponding sets of graphs. We show that the “superimposition” of all graphs in a set allows for the (re)construction of the efficient frontiers. We also identify a relationship between the Sharpe ratio (SR) of a given portfolio and the topology of the corresponding network of assets. More specifically, we suggest $SR = f(\text{topology}) \approx f(\text{ECC}/\text{BC})$, where ECC is the eccentricity and BC is the betweenness centrality averaged over all nodes in the network. At each threshold, the structural analysis of the correlated networks provides unique insights into the relationships between assets, agencies, risks, returns and cash flows. We observe that the best threshold or best graph representation corresponds to the portfolio with the highest Sharpe ratio. We also show that simulated annealing performs better than a gradient-based solver.

Topic: Central bank research

JEL code: C02

Résumé

Dans la présente étude, nous renforçons la méthode de sélection des actifs de Markowitz au moyen d’une théorie à représentation graphique afin d’analyser deux portefeuilles d’actifs libellés soit en euros, soit en dollars américains. Nous procédons à une décomposition fondée sur des variations ou seuils de la matrice de covariance de chacun de ces portefeuilles. Chaque seuil équivaut à un différent niveau de risque, et la structure de portefeuille correspondante est représentée graphiquement. Nous montrons que la « superposition » de tous les graphiques d’un ensemble permet de (re)construire les frontières efficientes. Nous notons également une relation entre le ratio de Sharpe (SR) d’un portefeuille donné et la topologie du réseau d’actifs correspondant. Plus précisément, nous suggérons la relation suivante : $SR = f(\text{topologie}) \approx f(\text{ECC}/\text{BC})$, où ECC correspond à l’excentricité et BC, à la centralité d’intermédiation moyenne de tous les nœuds du réseau. À chaque seuil, l’analyse structurelle des réseaux corrélés fournit de l’information unique sur les relations entre les actifs, les organismes, les risques, les rendements et les flux de trésorerie. Nous constatons que le seuil optimal ou la meilleure représentation graphique cadrent avec le portefeuille au ratio de Sharpe le plus élevé. Nous montrons aussi que le recuit simulé fonctionne mieux qu’un solveur basé sur un gradient.

Sujet : Recherches menées par les banques centrales

Code JEL : C02

Non-technical Summary

Portfolio optimization and diversification have played a pivotal role in the understanding of financial markets and the evolution of financial decision-making. In his seminal work published in 1952, Markowitz [1] laid the foundation for Modern Portfolio Theory, where investors decide on how to optimally allocate their budgets across a set of assets in a manner that maximizes their expected returns and diminishes their exposure to risk.

In this study, we enhance Markowitz portfolio selection with graph theory for the analysis of two portfolios composed of either EU or US assets. Using a threshold-based decomposition of their respective covariance matrices, we perturb the level of risk in each portfolio and build the corresponding sets of graphs in an attempt to find reduced, yet optimal, representations of our portfolio selection problem. Our procedure induces artificial diversification, impacting network topologies, asset selection and the contribution of each asset to the portfolio’s risk and return. The structural analysis of the resulting networks allows the identification of two types of nodes, bulk and surface ones, displaying opposite geometric properties. In general, we find that surface nodes correspond with low-risk assets while the bulk nodes, which have a greater influence over the topology of the network, are the higher risk assets. To quantify an asset’s influence within a portfolio, we utilize two graph properties: (1) eccentricity and (2) betweenness centrality. Both of these measures show how the variation in returns of a given asset influences the expected returns of the other assets in the portfolio. As in Markowitz mean-variance optimization, we show that assets that have the least influence over the correlated network are highly targeted for investments. The risks and returns in the EU portfolio are primarily driven by bond issuers based in Germany. On the other hand, US assets influence each other more equally. The “superimposition” of each threshold-based topology allows for the reconstruction of the efficient frontiers of the two portfolios, while we identify a relationship between the Sharpe ratio (SR) of a given portfolio and the topology of the corresponding network of assets. More specifically, we suggest $SR = f(\text{topology}) \approx f(\text{ECC}/\text{BC})$, where ECC is the eccentricity and BC is the betweenness centrality averaged over all nodes in the network. We find that the best threshold or best graph representation corresponds to the portfolio with the highest Sharpe ratio.

In future work, we want to use our decomposition method to decrease the problem size for applications on quantum hardware.

1 Introduction

Portfolio optimization and diversification have played a pivotal role in the understanding of financial markets and the evolution of financial decision-making. In his seminal work published in 1952, Markowitz [1] laid the foundation for Modern Portfolio Theory (MPT), where investors decide on how to optimally allocate their budgets across a set of assets in a manner that maximizes their expected returns and diminishes their exposure to risk. Presently, many quantitative techniques are utilized in the financial industry to inform decision-making with respect to investment in equity markets [4, 5]. Especially, market forecasting [6] and portfolio selection [7] have been streamlined by the reduced cost of computing power and the increased availability of sophisticated computational tools that allow investors to include their analysis of future market movements in their decision-making processes. However, there is no evidence that such methods can consistently outperform the simple but effective equally weighted portfolio rule [8–10].

Many recent publications have employed graph theory to represent financial markets as correlated networks [11–17] and grasp the long-term behaviour of a collection of assets. Following this framework, individual assets correspond to nodes while the edges connecting the nodes together are related to the covariance in the expected returns, encoding the complex interplay of financial equities towards maximizing the wealth and/or minimizing the risk associated with a given portfolio [11, 18–20]. Nevertheless, most of these studies remain fundamentally descriptive and lack direct applications for the portfolio selection process. Recently, Peralta and Zareei [21] showed a negative relationship between an asset’s centrality within a graph and its optimal weight within the Markowitz framework, validating the network-based approach to portfolio selection and promising an improvement over the Markowitz scheme.

Here, we leverage MPT with graph theory to probe the risk profile of a given collection of assets. We identify the best topology as well as securities that play a central role in driving the Sharpe ratio [22] of a portfolio. As we decompose our problem into sub-problems or sub-graphs, we show that the connectivity of the correlated networks used to seed the portfolio optimization at each level of risk plays an important role in determining the risk-return behaviour of the respective portfolios.

While in the present work we focus on a small collection of bonds, modern portfolio selection can become quite complex. Relative to classical information processing, quantum computation holds the promise of providing exponential speedups for some technologically important problems [23]. The idea of applying quantum mechanics to finance is not a new one, where many of the problems can be formulated as optimization problems. While these are tasks that can be hard for classical computers, they find a natural expression using quantum optimization methods [24].

In recent years, this area of research has seen tremendous growth, which is in part due to the increasing commercial availability of quantum annealers. These are a specialized type of quantum computer that employ quantum effects towards solving unconstrained optimization problems [25]. The portfolio selection problem is particularly well suited to run on the quantum annealer developed by the Canadian technology company D-wave Systems [26–28]. While a complete discussion of the methodologies behind this technology is outside the scope of this paper, the literature cited here provides more in-depth discussions of these concepts.

In light of this growing paradigm shift in computation, the motivation for this work is to provide a high-level yet reliable selection method for portfolio managers, while opening doors to the development of new algorithms for better and faster sampling on advanced hardware. Specifically, in future work, we are interested in implementing adaptive sampling of the portfolio optimization problem on a D-wave quantum annealer using well-defined boundaries for shocks and a good starting point. However, this technology is still in its infancy and, therefore, is limited by the size of problem that can be mapped onto the hardware. Hence, by using graph theory we aim to reduce the degrees of freedom in the portfolio

selection problem.

For the portfolio optimization carried out in this work we focus on results obtained using traditional computational methods. Two sampling methods are applied: simulated annealing (SA) and the interior/direct algorithm, a gradient-based method included with the Artelys KNITRO library. The SA approach is a heuristic optimization algorithm that is known to outperform gradient-based methods in sampling the solution space and is analogous to the process that occurs on a quantum annealer [29, 30]. Furthermore, previous work has shown it is well suited to solve the mean-variance portfolio selection model [31]. As such, the two methods should provide similar results, with the gradient-based method serving as our control.

The remainder of this paper is structured as follows. In section 2, we describe the methodology for decomposing two portfolios by including distinct assets in graphs and how it can be used to (re)construct the efficient frontiers. In section 3, we perform a structural analysis of the resulting correlated networks in an attempt to reinforce asset selection and help optimize cash flows. Results and discussion are presented in section 4. In section 5, we briefly compare the performance of simulated annealing with a gradient-based solver. The attached appendix provides supporting information for the methods used and results obtained within the work presented.

2 Methodology

2.1 The Portfolio Optimization Problem and Its Graph Representation

To formalize the problem of optimal portfolio selection, suppose that we wish to invest in M different stocks from a given portfolio, but we cannot purchase more than k of any given stock. Applying the mean-variance optimization introduced by Markowitz, the problem takes on the following mathematical representation:

$$\max \left\{ \sum_{i=1}^M X_i R_i - \gamma \sum_{i,j} Q_{ij} X_i X_j \right\} \quad (1)$$

where $X = [X_1, X_2, \dots, X_M]$ is a vector of weights describing the percentage allocation of the total budget to asset i , R_i represents the expected returns from asset i , Q is the covariance matrix of stock returns used to define risk, and γ determines how an investor trades off risk due to the covariances and expected return. The individual elements of the covariance matrix Q_{ij} represent the covariances of the variables R_i and R_j , which reduce to the variance of R_i for $i = j$.

In a correlated network, or graph $G = (V, E)$, the notation $(i, j) \in E$, signifies that vertex or node i shares an edge with vertex or node j , and we say that these two nodes are connected. In general, the connections of a graph are described using an *adjacency matrix*. The adjacency matrix, $\Omega = [\Omega_{ij}]$, is an $n \times n$ matrix that, in general, is constructed from a corresponding covariance or correlation matrix, where n corresponds to the number of nodes in the graph (see the detailed procedure in section 2.2). The element $\Omega_{ij} = 1$ in the adjacency matrix whenever $(i, j) \in E$, and $\Omega_{ij} = 0$ whenever $(i, j) \notin E$. The network G is said to be *undirected* if no relationships are assigned to the edges, which implies that $\Omega = \Omega^T$ since $(i, j) \in E \Leftrightarrow (j, i) \in E$. Undirected graphs are the foundation of many of the works referenced, as well as the work presented here.

As alluded to, each correlated network can be described using graph properties like the eccentricity and the betweenness centrality of a node, as well as the degeneracy of a given topology. The concept of eccentricity of a given node i , ϵ_i , is fundamental in graph theory and is formally defined as the longest

hopcount between node i and any other node in G [32]. In essence, the eccentricity of a node provides a measure of how closely connected a given node is to all others in the network. The associated notion of centrality, which is intimately related to social network analysis, aims to quantify the number of connections a given node shares with all others in the graph [33, 34]. More specifically, the betweenness centrality is formally defined as the number of shortest paths that pass through a given node, which serves to quantify the influence a node has over the flow of information in a graph. Here, we choose to track the betweenness centrality and eccentricity of each node to assess its potential as a plausible new variable to consider in the portfolio selection procedure. We use the degeneracy as an indicator of the degree of connectivity between the nodes.

In the following sections, we present a graph-based decomposition of two fixed-income portfolio problems using US and EU sovereign and agency issuers. We generate multiple sub-graphs by perturbing the covariance matrix or by gradually eliminating connections between assets and then investigate the resulting correlated structures. These sub-graphs provide snapshots of different levels of variance or, in the context of portfolio selection, risk.

2.2 Decomposition of the Portfolio Optimization Problem

In the work presented, the covariance, ρ , is used as a means to study the relationship between 60 assets taken from European Union (EU) markets, in addition to 30 assets taken from the US market, where the bond prices were obtained from Bloomberg. Following the methodology proposed by Gürkaynak et al. [35], we fit the Svensson yield curve to the data. We then use the resulting estimated parameters to obtain zero-coupon yields at 1- to 10-year maturities. The covariance matrices used are then obtained from the zero-coupon yields.

We propose that the portfolio selection problem can be broken into a set of independent sub-problems by following a graph-based decomposition. We use the covariance thresholds listed in Table 2.1 to generate multiple graph representations (i.e., adjacency matrices) for each set of assets. Each threshold is a percentage, θ , of the maximum covariance found within the respective covariance matrices. For $\rho > \theta$, we insert 1 in the adjacency matrix and zero otherwise. The covariance threshold, θ , is thus a measure of the stringency of the stocks included in the adjacency matrix. Selecting the appropriate threshold is, therefore, an important factor to take into consideration.

At each threshold, the Hadamard (element-wise) product is taken between the adjacency matrix and the unperturbed covariance matrix, generating a sub-graph that conserves the original number of assets in the portfolio and removing only some of the risk contributions. Thus, each sub-graph represents a region around the efficient frontier [36, 37], a point with a given level of risk and a given level of returns associated with the investment. All generated sub-graphs are then used as starting configurations in two optimization procedures using either SA [29] or a gradient-based solver from the KNITRO library [38]. Each solution found is a portfolio, with optimal proportions of US or EU assets. The corresponding set of optimal portfolios will be later used to reconstruct the efficient frontier.

In Figure 2.1, the risk and return contributions for the optimal portfolios are plotted as functions of the covariance threshold. Both risk and return tend to increase as the covariance threshold increases, with a noticeable slope change at 50% of the maximum covariance. The EU portfolio shows the steadiest increase along θ , until reaching a plateau around 50–70% of the maximum covariance. Such a plateau occurs at the highest thresholds, when most of the nodes in the corresponding sub-graph are disconnected (i.e., most covariances are set to zero). Therefore, only a few assets in the portfolio are likely to drive the expected returns until those “major” nodes become detached from the correlated network and lose their influence.

In Figure 2.2, we plot the degeneracy of each threshold-based graph calculated with the following

formula:

$$\Gamma(\theta) = \left(\frac{\max(K(\theta))}{\max(K(0))} \right) \cdot 100 \quad (2)$$

where $\max(K(\theta))$ is the maximum k -core for the sub-graph generated at the threshold θ , and $\max(K(0))$ is the maximum k -core for the sub-graph obtained when no threshold is applied (when all nodes are connected). A k -core of a graph G is a maximal connected sub-graph of G in which all nodes are connected to at least k other nodes. Hence, the degeneracy provides a means of quantifying the change in connectivity as the covariance threshold is varied.

Table 2.1: Thresholds used to generate graphs

Graph	US Assets Threshold ϕ_{US}	EU Assets Threshold ϕ_{EU}
Maximum Covariance Q_{max}	6.1425×10^{-3}	6.5128×10^{-3}
$G_1(\theta = 10)$	6.1425×10^{-4}	6.5128×10^{-4}
$G_2(\theta = 20)$	1.2285×10^{-3}	1.3026×10^{-3}
$G_3(\theta = 30)$	1.8428×10^{-3}	1.9538×10^{-3}
$G_4(\theta = 40)$	2.4570×10^{-3}	2.6051×10^{-3}
$G_5(\theta = 50)$	3.0713×10^{-3}	3.2564×10^{-3}
$G_6(\theta = 70)$	4.2998×10^{-3}	4.5590×10^{-3}
$G_7(\theta = 90)$	5.5283×10^{-3}	5.8612×10^{-3}




Figure 2.1: Risks and returns for the optimal portfolios obtained using (a) simulated annealing and (b) gradient-based optimization, with starting configurations generated at different covariance thresholds.

The degeneracy for the US portfolio decreases linearly with respect to the threshold, while the degeneracy of the EU portfolio decreases in a more exponential manner. The latter suggests that the initial EU portfolio of 60 assets was composed of a cluster of loosely correlated bonds. This cluster rapidly reduces in size before stabilizing around $\theta = 50$, when we are left with a more tightly knit cluster of highly correlated assets. In contrast, the linear decrease in degeneracy for the US portfolio implies that there is a higher level of correlation amongst all the assets included in the US portfolio and, therefore, they are

less sparse than the EU assets. Hence, increasing the covariance threshold in this instance results in the connected components decreasing in a consistent manner.




Figure 2.2: Change in degeneracy with respect to increasing covariance thresholds, given as a percent of the maximum degeneracy when all assets are included in either portfolio.

To further understand the existence of a plateau in the degeneracy, as well as the risk and return plots, we show, in Figure 2.3, the purchase frequency of each EU asset as the covariance threshold increases. The frequency converges towards 1 as it reaches a threshold of 50%. Beyond, each asset has the same probability to be picked, as shown in the purchase frequency plots, which correspond to the plateau regions in Figure 2.1 and Figure 2.2. The corresponding distribution of weights is presented in the Appendix in Figure A.2. It is worth noting that asset 40 is always purchased across thresholds. In section 3, we show this security to be the greatest contributor to both the risk and returns.

Figure 2.4 shows that, when considering the purchase frequencies for the US portfolios across all thresholds, we see that a large proportion of the assets is always purchased. Especially, asset 20, which we later show to have the strongest influence over the topology of the network (section 3). Interestingly, there is a subset of assets that is consistently purchased from each group of bond distributors in both the EU and US portfolios. For the US portfolio these securities correspond to assets 1–5 and 10–15, while the EU portfolio bonds issued by 1–3, 10–13, 20–23, 30–33, 40–43 and 50–53 are purchased nearly 100% of the time across all thresholds. This is corroborated by the weight distributions in both Figure A.2 and Figure A.3 (see Appendix).

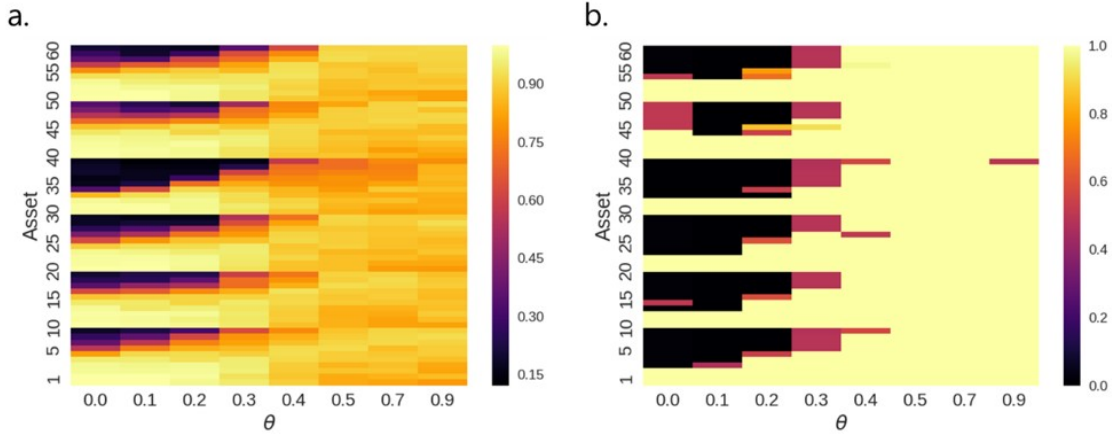


Figure 2.3: Purchase frequencies over 300 independent runs of the optimization at each starting configuration defined by the covariance threshold using (a) simulated annealing and (b) gradient-based optimization algorithms to determine optimal EU portfolios.




Figure 2.4: Purchase frequencies over 300 independent runs of the optimization at each starting configuration defined by the covariance threshold using (a) simulated annealing and (b) gradient-based optimization algorithms to determine optimal US portfolios.

2.3 Validation of Sampling

Figure 2.5 presents the return distributions of the optimal portfolios found with either SA or our gradient-based solver. For both the EU and US portfolios, SA yields to symmetric distributions over all thresholds considered. To confirm the normality of these distributions, we perform a Shapiro-Wilks test for normality at an alpha level of 0.01. The corresponding p-values can be found in Table A.1 and Table A.2 in the Appendix. For the two distributions displayed in Figure 2.5a and Figure 2.5c, the p-values are found to exceed 0.01 in all cases, providing evidence that the resulting distributions are normal within a confidence interval of 99%. In contrast, the resulting returns identified using a gradient, shown in Figure 2.5b and Figure 2.5d, are very narrow for both portfolios. Along with the purchase frequencies presented in Figure 2.3 and Figure 2.4, these results suggest that the gradient-based optimization converges to the same portfolio with a very high probability. The corresponding p-values remain far below 0.01, suggesting the formation of a Dirac peak. We should note that increasing the sampling further (up to 1000 samples) does not change the distribution of expected returns.

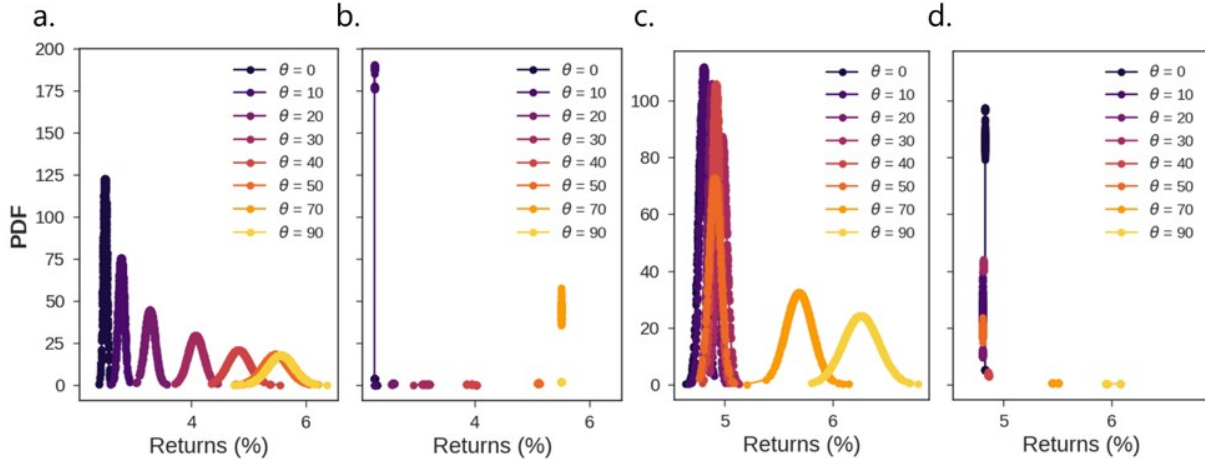


Figure 2.5: Distribution of expected returns over increasing covariance thresholds obtained for the EU portfolio using (a) simulated annealing and (b) gradient-based optimization and for the US portfolio using (c) simulated annealing and (d) gradient-based optimization. As the threshold increases the mean return and spread of the distributions increase, with the latter indicating a higher associated risk. In all cases shown, 500 samples are used to generate the probability density functions (PDF).

2.4 Superimposition of Independent Sub-problems

Figure 2.6 shows, for a different risk-aversion coefficient, efficient frontiers constructed from the optimal portfolios that are identified using threshold-based configurations as inputs to our optimization routines. Therefore, at each covariance threshold, an optimal portfolio offers the highest expected return for a specific level of risk. Here, we define the risk, σ , as:

$$\sigma = \sqrt{X^T Q_\theta X} \quad (3)$$

We show that a slight separation appears between efficient frontiers across increasing risk aversion, suggesting a low level of risk in the government bond asset class evaluated in this work. The dark blue point in each of the plots corresponds with investing in all assets equally. Under this scheme, the US portfolio has a Sharpe ratio of 0.97, while the equally weighted EU portfolio has a Sharpe ratio of 1.16. As seen in Figure 2.6, these equally weighted portfolios fall below the respective efficient frontiers. This indicates that the equally weighted portfolios generate a smaller return relative to a portfolio that lies on the efficient frontier for the same level of risk. We should note that results from the SA optimization show a greater amount of differentiation between the efficient frontiers. These findings will be revisited in more detail in section 5, where the two solvers are more formally compared.

3 Structural Analysis

To investigate how the topology of the networks generated over increasing covariance limits affects the resulting portfolios, we consider the two graph metrics that we defined in section 2.1. Specifically, we consider the betweenness centrality and eccentricities of the connected nodes at each covariance threshold. Recall that the betweenness centrality and eccentricity provide a means of quantifying how much influence a given asset has over all others in the network. This indicates whether a given asset plays an important role in determining the variation in returns of all other assets. In other words, this provides a means of identifying securities that contribute the most to risk. The eccentricity quantifies




Figure 2.6: Efficient frontiers for the EU portfolio obtained using (a) simulated annealing and (b) gradient-based optimization, and for the US portfolio using (c) simulated annealing and (d) gradient-based optimization. The point indicated on the graphs corresponds to the portfolio obtained when all assets are equally weighted.

how closely connected a given asset is to all others in the network. Here, when we refer to a sub-graph, we mean the connected nodes of a (larger) graph, while we name each disconnected node of the same (larger) graph a singleton.

Figure 3.1 shows the betweenness centrality (a) and eccentricity (b) of the individual nodes in the US portfolio with respect to θ . At low thresholds, the centrality is more equally distributed across all nodes in the network. The reverse is true when considering the eccentricity of the nodes across varying thresholds where, as the covariance limit increases, the eccentricity becomes more evenly distributed across nodes, and the assets form a tighter network. However, the eccentricities of the US nodes never exceed 2, where only nodes 2, 12 and 22 have eccentricities greater than 1. Figure 3.2 shows that nodes 2, 12 and 22 form singleton graphs as θ increases from 10% to 20%. The centrality of asset 20 is also particularly evident and remains significant for thresholds greater than 20%, as confirmed in Figure 3.1a. Nodes 2, 12 and 22 are typical “surface nodes” or nodes with low centralities and high eccentricities, while the opposite is true for “bulk nodes”. In general, as the threshold is increased, surface nodes are consistently seen forming singleton graphs with no connections to any other nodes. In contrast, bulk nodes act as connecting interfaces between neighbouring assets, strongly influencing the overall network topology. In this regard, the changing eccentricity and betweenness centrality of a given node indicates a change in the configuration of the network, where higher eccentricity vertices move towards the edge

of the resulting sub-graph, pushing lower eccentricity nodes inwards. In other words, highly influential nodes in the network are consistently found to be bulk nodes that act as the bridges linking together many of the other assets in the portfolio.

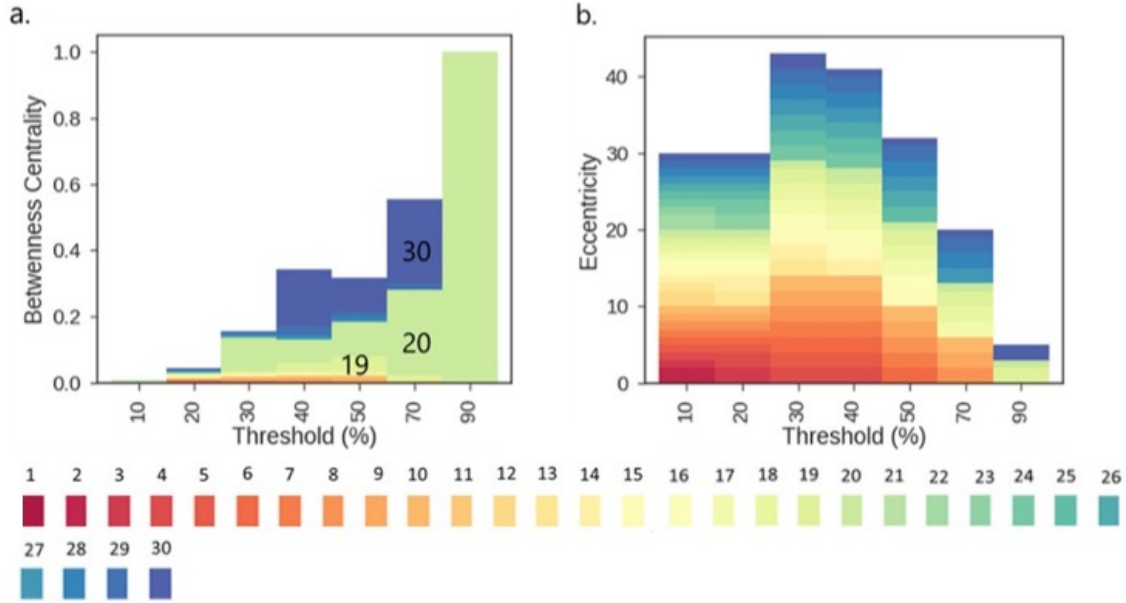


Figure 3.1: (a) Betweenness centrality of US portfolio, (b) Eccentricity for US portfolio

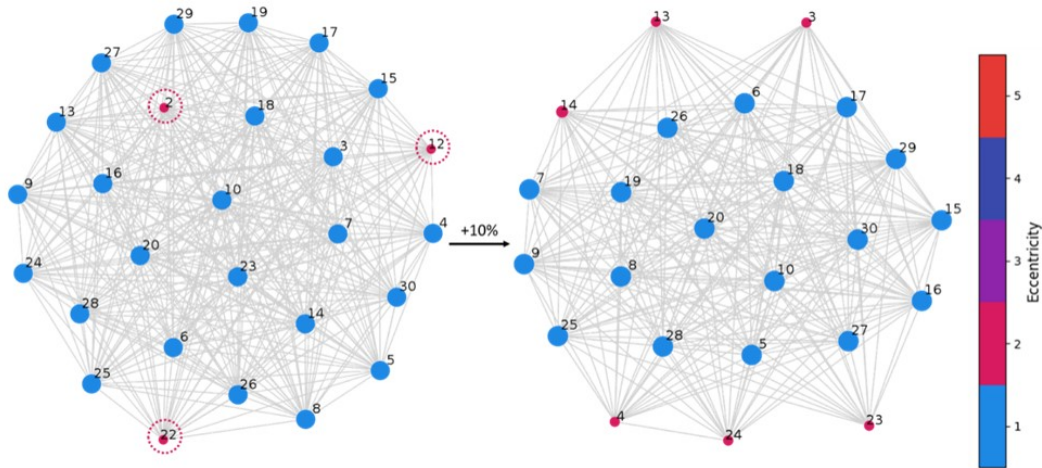


Figure 3.2: US sub-graphs at thresholds of 10% (left) and 20% (right) of the maximum covariance. The encircled nodes form singleton graphs upon increasing the threshold by 10%.

Similarly, in the collection of EU assets, the surface nodes tend to be disconnected at higher thresholds. In contrast with the US portfolio, however, the betweenness centrality of the nodes in the EU portfolio (Figure 3.3a) are less evenly distributed, even at lower thresholds, suggesting that there are fewer highly correlated assets in the EU portfolio. This is further illustrated by the nodes with the highest centralities remaining relatively constant across all covariance limits investigated. The most notable of these are assets 40, 50 and 60, which dominate the distribution of betweenness centralities at a threshold of 50% of the maximum covariance. Hence, the consistency with which these nodes are found to have the highest centralities suggests that they will have a significant influence on the returns and corresponding risks associated with the EU portfolio. This is further corroborated by the eccentricities presented in Figure

3.3b, where these nodes consistently act as the shortest path between abetweenll other assets in the resulting networks.

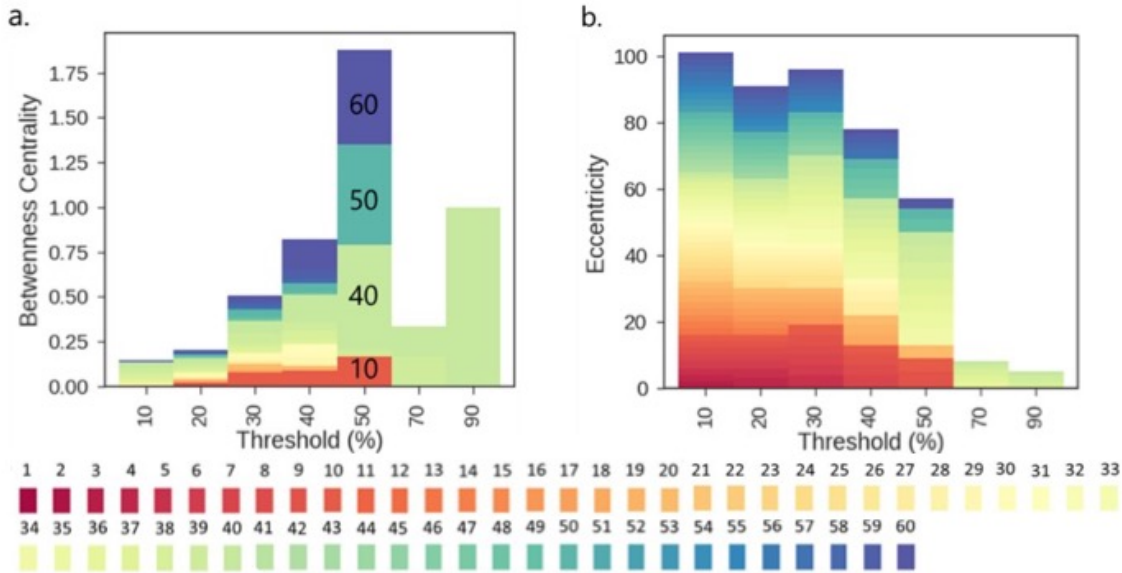


Figure 3.3: (a) Betweenness centrality for EU portfolio, (b) Eccentricity for EU portfolio.

Figure 3.4 shows the network of nodes corresponding to a covariance threshold of 50% of the maximum covariance. Here, we can see that the three EU assets – 40, 50 and 60 – act as the only connection between two tightly knit clusters on opposing sides of the sub-graph. Taking a closer look at the identity of these assets, which are principally bonds that require 1 to 10 years to reach maturity, one cluster is a set of bonds from Germany, while another is largely composed of bonds issued from Italy. The bridge between the bonds issued by the two regions, asset 60, is a supranational agency. Up to a covariance threshold of 50%, the bonds issued by the supranational agency are clearly seen as acting as the central connection between the bonds issued by different countries. Beyond 50%, the clusters of Italian and international bonds are no longer included in the subsequent sub-graphs, corresponding to a cut-off between the German holdings and the supranational agency. Therefore, the cluster of German assets present in the portfolio has the strongest influence over the topology of the network and, hence, acts as the principal driver for returns and risks. This is further reflected as the covariance threshold is increased and only bonds issued from Germany remain connected.




Figure 3.4: Sub-graphs for the EU portfolio at a covariance threshold of 50% (left) and 70% (right). The size of the nodes in the graphs corresponds to the betweenness centrality of the node, where larger nodes have higher centralities. Above the 50% threshold, the sub-graph is composed entirely of German assets, where node 40 is consistently found to have a high centrality and low eccentricity.

A similar trend is observed in the US portfolio (Figure 3.5), where a supranational agency, asset 30, acts as the central hub for connecting together the bonds distributed by different issuers from around the country. From the sub-graph shown in Figure 3.5, it can be concluded that expected returns on the bonds issued by a government-sponsored enterprise, asset 19, are strongly affected by the government-sponsored and supranational agencies represented by nodes 20 and 30. This indicates the likelihood of a non-negligible relationship between these three entities. Given the consistently high centrality of asset 30 over the different thresholds, this further indicates that much of the risk in the US portfolio is the result of variation in returns of this node. Along the same logic, then, asset 20 acts as the definitive bridge between the supranational agencies (nodes 21–30), the US Treasury (nodes 1–10) and all other government-sponsored bonds (nodes 11–20).

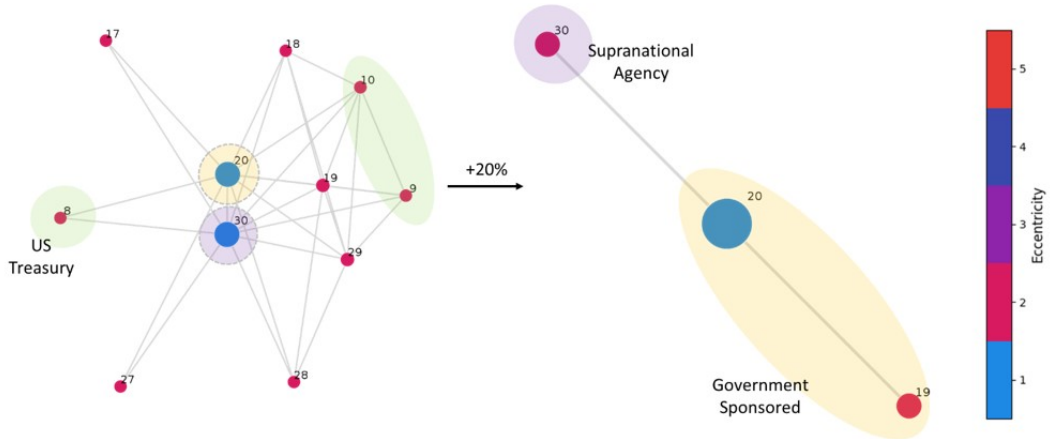


Figure 3.5: Sub-graphs for the US portfolio at a covariance threshold of 70% (left) and 90% (right) of the maximum covariance. At the lower threshold, the supranational agency represented by node 20 acts as the bridge between all other assets in the portfolio and is consistently found to have high centrality and low eccentricity.

4 Results and Discussion

4.1 Topology and Asset Selection

In Figure 2.2, we show a more marked decrease in degeneracy for the EU portfolio relative to the US portfolio across increasing thresholds. This suggests a greater level of correlation between the US assets, which aligns with all assets originating from same-country government agencies. In either portfolio, the maximum eccentricity for the EU assets is 5, while for the US assets the maximum eccentricity is 2, providing further indication that the EU assets are more sparse than the US securities. The purchase frequencies (Figure 2.4) suggest that, for the EU portfolio, random selection happened for high-threshold topologies, which corresponds to the plateau observed in Figure 2.2 or the sub-graphs primarily composed of German assets.

Interestingly, there is a subset of assets that are consistently purchased from each group of bond issuers in both portfolios. For the US portfolio, these securities correspond to assets 1–5 and 10–15, while for the EU portfolio bonds issued by 1–3, 10–13, 20–23, 30–33, 40–43 and 50–53 are purchased nearly 100% of the time across all thresholds. This finding aligns with the corresponding weight distributions shown in Figure A.1 and Figure A.2, where assets have more or less the same weight in their respective optimal portfolios.

Our detailed structural analysis also emphasizes asset 20 as the most influential node over the topology of the US network, while Figure 2.4 shows it is purchased 100% of the time across all threshold-based configurations. The same is true for asset 40 in the EU portfolio, which is the greatest contributor to both the risk and returns (see Figure A.7 in Appendix).

Finally, we propose to quantify each threshold-based topology using the ratio between the mean normalized eccentricity ($E\tilde{C}_{Avg}$) and the mean normalized betweenness centrality ($B\tilde{C}_{Avg}$). Higher values of this ratio indicate a greater number of surface nodes in the corresponding sub-graph. In Figure 4.1, we show a relationship between the Sharpe ratio and the topologies generated at thresholds 40, 50, 70 and 90, when all connected nodes in the corresponding graph are equally weighted. As the proportion of surface nodes increases, we may expect the trade-off between the expected return and the associated risk to improve.




Figure 4.1: Sharpe ratio as a function of topology, defined in terms of the ratio of the mean normalized eccentricity and the mean normalized betweenness centrality of the corresponding sub-graph for (a) EU topologies and (b) US topologies. The Sharpe ratio was calculated from sub-graphs with all connected nodes equally weighted.

4.2 Identifying the Optimal Risky Portfolio

The portfolios that exist along the efficient frontier present entities such as firms or individuals with many choices as to the best investment to make. Additionally, the economic theory of choice employs the concept of a utility function to describe the way entities make decisions when faced with a set of choices. Put simply, the utility function describes the preferences of entities over perceived risk and expected return combinations that comprise each optimal portfolio along the efficient frontier – namely, an efficient portfolio. These preferences can be plotted as a set of indifference curves, where all points on a given curve indicate combinations of risk and expected return that will give the same level of utility to a given investor. At each level of risk aversion, the utility, U , was calculated using equation (3) for the efficient portfolios presented in Figure 2.6. In equation (4), $E(R)$ is the expected returns of a given portfolio, γ is the risk-aversion coefficient, and σ is the associated risk. The indifference curves that were obtained are presented in Figure 4.2.

$$U = E(R) - 0.5\gamma\sigma^2 \quad (4)$$

In Figure 4.2, vertical lines are drawn at points that correspond to each threshold. For a risk-averse investor, the intersections of these lines with the indifference curve represent points at which the investor would be indifferent regarding the different portfolios that exist at each of these points. Hence, at a given threshold, the investor remains indifferent regarding the choices that remain within their accepted level of risk.

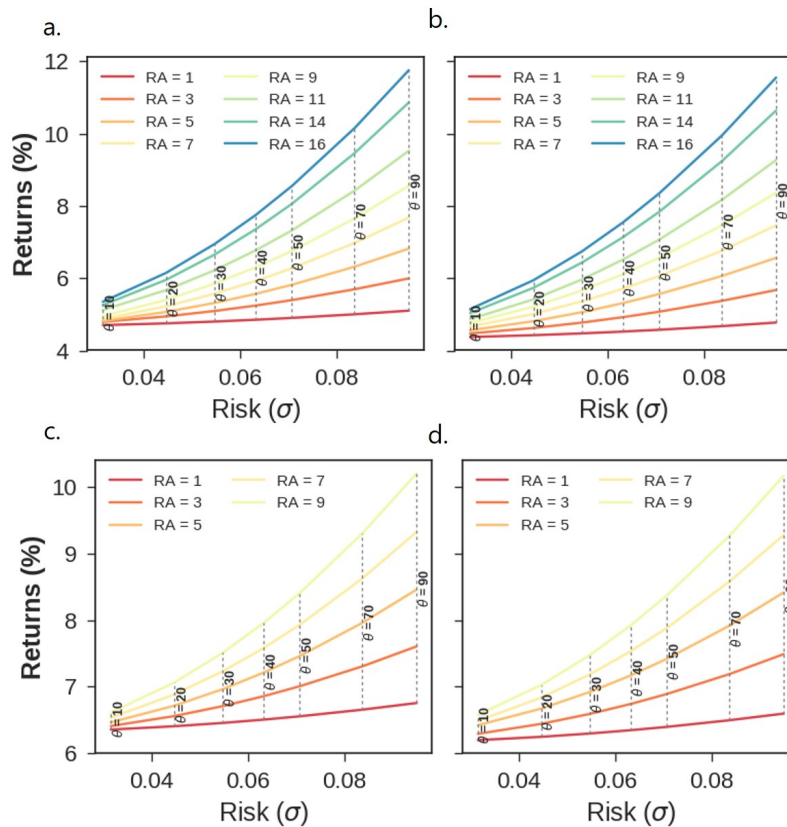


Figure 4.2: Indifference curves for the EU portfolio generated using results from the (a) simulated annealing and (b) gradient-based optimization algorithms, and indifference curves for the US portfolio generated using results from the (c) simulated annealing and (d) gradient-based optimization algorithms.

To determine the most optimal portfolio or optimal risky portfolio for the investor using these indifference curves, we must identify the point where an indifference curve is tangent to the efficient frontier. In Figure 4.3, this point is highlighted in orange and corresponds to an optimal portfolio generated at a 70% covariance threshold, a sub-graph made of a cluster of German bonds and disconnected nodes. This is the point at which the assets are randomly selected, including all EU bonds, according to Figure 2.3, while it is also the point where both the expected rates of return and the associated risks begin to plateau, as shown in Figure 2.1. In contrast, results from Figure 2.4 suggest that, in the US portfolio, assets are not all included at this threshold and, according to Figure 3.1, assets 9, 18, 19, 29 and 30 will likely be excluded. In Figure 3.3, we see that at the highest covariance limit, asset 20 is the sole highly correlated bond included in the optimal portfolio. This result confirms the important role this security plays in driving the returns and the risk for the US portfolio. The SA results presented in Figure 2.3a show that only a few of the same assets are purchased at a lower frequency, but still point to a strong influence of asset 20 on the overall returns and risk of the US portfolio.

In Figure 4.4, we show the Sharpe ratio of the different efficient portfolios across varying covariance thresholds and risk-aversion coefficients. From the Sharpe ratio plots shown, the maximum corresponds directly to the optimal risky portfolio previously identified for either set of securities, as shown in Figure 4.3. For both collections of assets, this is the efficient portfolio generated from the sub-graph at a 70% covariance threshold used as the starting configuration of our optimization procedures. As for our previous results, SA shows a greater degree of variability in the Sharpe ratios found for the same portfolio at varying levels of risk aversion.

The optimal EU portfolio coincides with a Sharpe ratio of 1.26, while the optimal US portfolio has a Sharpe ratio of 1.16. When we compare the Sharpe ratios of the equally weighted EU and US portfolios, we find the method presented in this work offers a non-negligible improvement (EU Sharpe ratio = 1.16, US Sharpe ratio = 0.97). Given the low risk of either portfolio, it is not a surprise that the equally weighted portfolio provides a reasonable solution in comparison to the more involved method presented in this work. This is particularly apparent for the EU portfolio, where the assets are less correlated with each other. Hence, we see a much greater improvement over the equal weighting scheme for the US portfolio, where the overall riskiness of the portfolio is greater.

4.3 Analysis of the Optimal Risky Portfolio

The work of Peralta and Zareei [21] shows that the asset centrality is negatively correlated to the risk component of that same asset. In an attempt to evaluate the riskiness of the assets included in the optimal portfolio, we plot, in Figure 4.5, the individual Sharpe ratios of each security with respect to their betweenness centrality. From this analysis, we identify three groups of assets (low-, medium- and high-risk), while the corresponding data points are sized by their relative weight within the respective portfolio. Similarly to Peralta and Zareei [21], we find that the low-risk group contains assets with low centralities, and that these securities are purchased the most. In both portfolios, the assets that are found to have the greatest centralities across all covariance thresholds (asset 40 for the EU portfolio and asset 20 for the US portfolio) belong to the high-risk group. Results from the gradient-based optimization confirm our findings, as shown in Figure A.7 in the Appendix.




Figure 4.3: Selection of the optimal EU portfolio from (a) simulated annealing and (b) gradient-based optimization results, and selection of the optimal US portfolio from (c) simulated annealing and (d) gradient-based optimization results.

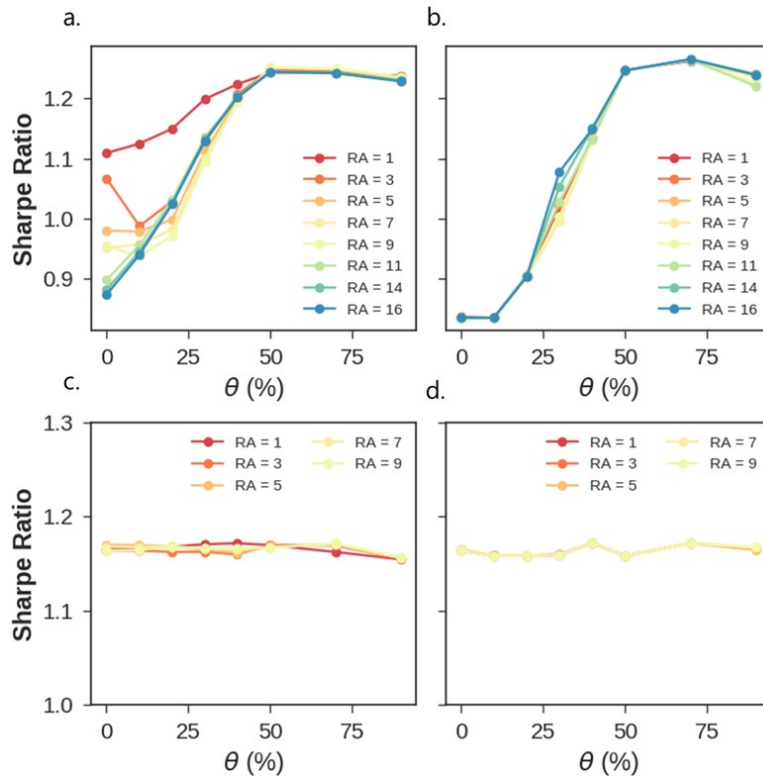


Figure 4.4: Sharpe ratio across a varying covariance threshold and for different risk-aversion coefficients. EU results with (a) simulated annealing and (b) gradient-based optimization. US results with (c) simulated annealing and (d) gradient-based optimization.




Figure 4.5: Relationship between the Sharpe ratio and betweenness centrality of individual assets included in the optimal risky (a) EU and (b) US portfolios identified by SA. The betweenness centralities plotted here were taken as the mean relative centrality of a given asset over all thresholds considered, such that assets with consistently high betweenness centralities will have the highest centralities.

In Figure A.8a, we investigate the returns and associated risks per individual asset. The pie chart given in Figure A.8b shows the allocation of the investor’s budget towards each of the assets included in the optimal portfolio. Of the collection of assets present in the EU portfolio, those that consistently carry the least weight in the optimal portfolio across both algorithms used are assets 38, 39 and 40. These assets, which originate from German bond issuers, coincide with nodes that were found to have high betweenness centralities and low eccentricities relative to the others in this group of assets (Figure 3.3). These also correspond to assets that have the highest correlation with each other and are the only remaining nodes that are connected at the highest covariance threshold. In this regard, the betweenness centrality and eccentricities of the nodes in the resulting sub-graphs provide insight into the riskiness of the associated asset. It should be noted that although these German bonds are not included within the optimal portfolio identified by the gradient-based solver from KNITRO, they are allocated a small portion (less than 1%) of the budget in the optimal portfolio identified by SA. Furthermore, the other German assets (31–37) are included in the optimal portfolio, with very similar weights. Looking at the returns and risks carried by each individual asset, it is also interesting to note that the relative proportion of risk to return for each asset is approximately equal.

Turning our attention to the optimal US portfolio, Figure A.9 presents a similar analysis. For the US portfolio, the assets that are consistently excluded from the optimal portfolio at a threshold of 70% include assets 20 and 30. Again, these assets are seen in Figure 3.2 as having both higher betweenness centralities and lower eccentricities relative to the other assets. Hence, these assets contribute significantly to the risk of the US securities considered in this work. Additionally, it was previously noted during the structural analysis presented in section 3 that asset 20, a government-sponsored bond, represented a bridging point between all other assets in the US portfolio and that much of the cash flow between the different agencies would pass through this point. Therefore, due to the high correlation with the other bond issuers in the United States, it is reasonable to assume that investors would be incurring added risk by including this security in their portfolios. Asset 30, which is issued by a supranational agency, can also be seen as acting as a central bridging point between the other assets in the US portfolio in Figure 3.5. A similar conclusion can be drawn for this asset as was drawn for asset 20, where together the two play a pivotal role in determining the risk of the respective portfolio.

4.4 Confirming the Best Covariance Threshold

In statistics and information theory, the relative entropy of two distributions serves as a measure of the change in amount of information contained in one distribution relative to the other. In this context, as the entropy increases there is a corresponding decrease in the information contained within the given distribution. Following this approach, we propose that the best reduced representation of a given portfolio can be identified as the minimal sub-graph (lowest degeneracy) that is most representative of the entire asset universe (smallest relative entropy). In a formal manner, we propose that the optimal threshold can be identified by simultaneously minimizing the graph degeneracy and relative entropy (i.e., change in entropy between the distributions of Sharpe ratios calculated at thresholds θ_i and θ_0).

In Figure 4.6 we plot the degeneracy of the sub-graph obtained at threshold θ_i , $\Gamma(\theta_i)$, with respect to the normalized relative entropy, $\tilde{\phi}(\theta_i - \theta_0)$, of the resulting distribution of Sharpe ratios obtained at θ_i compared with the distribution obtained when no threshold is applied (θ_0). The entropies are normalized so that both the degeneracy and relative entropy are (1) of the same scale numerically and (2) unitless. In Figure 4.6, we consider only the configurations from the SA sampling. It follows that if the expected returns are not normally distributed, the Sharpe ratios will also not be normally distributed. Given our proposed methodology, the optimal threshold should be found as the minimum in the plot given in Figure 4.6, which coincides with a covariance threshold of 70% for both portfolios investigated. Figure 4.7 summarizes the recommended procedure for decomposing a given asset portfolio.




Figure 4.6: Identifying the optimal reduced representation of the EU (red) and US (blue) portfolios by simultaneously minimizing the graph degeneracy and the entropy of the Sharpe ratio distribution relative to the distribution obtained when no threshold is applied.

5 Simulated Annealing vs Gradient-Based Optimization

Both SA and a gradient-based optimization are utilized in this work. The intrinsic search capability of each method is captured by the distribution of returns (Figure 2.5), where the KNITRO method produces a single pick of solutions, rather than near-normal distributions produced while using SA.

Following the Markowitz portfolio selection scheme, the associated risks are simultaneously minimized. At the start of the optimization, the vector of weights is randomly initialized. For the optimization car-

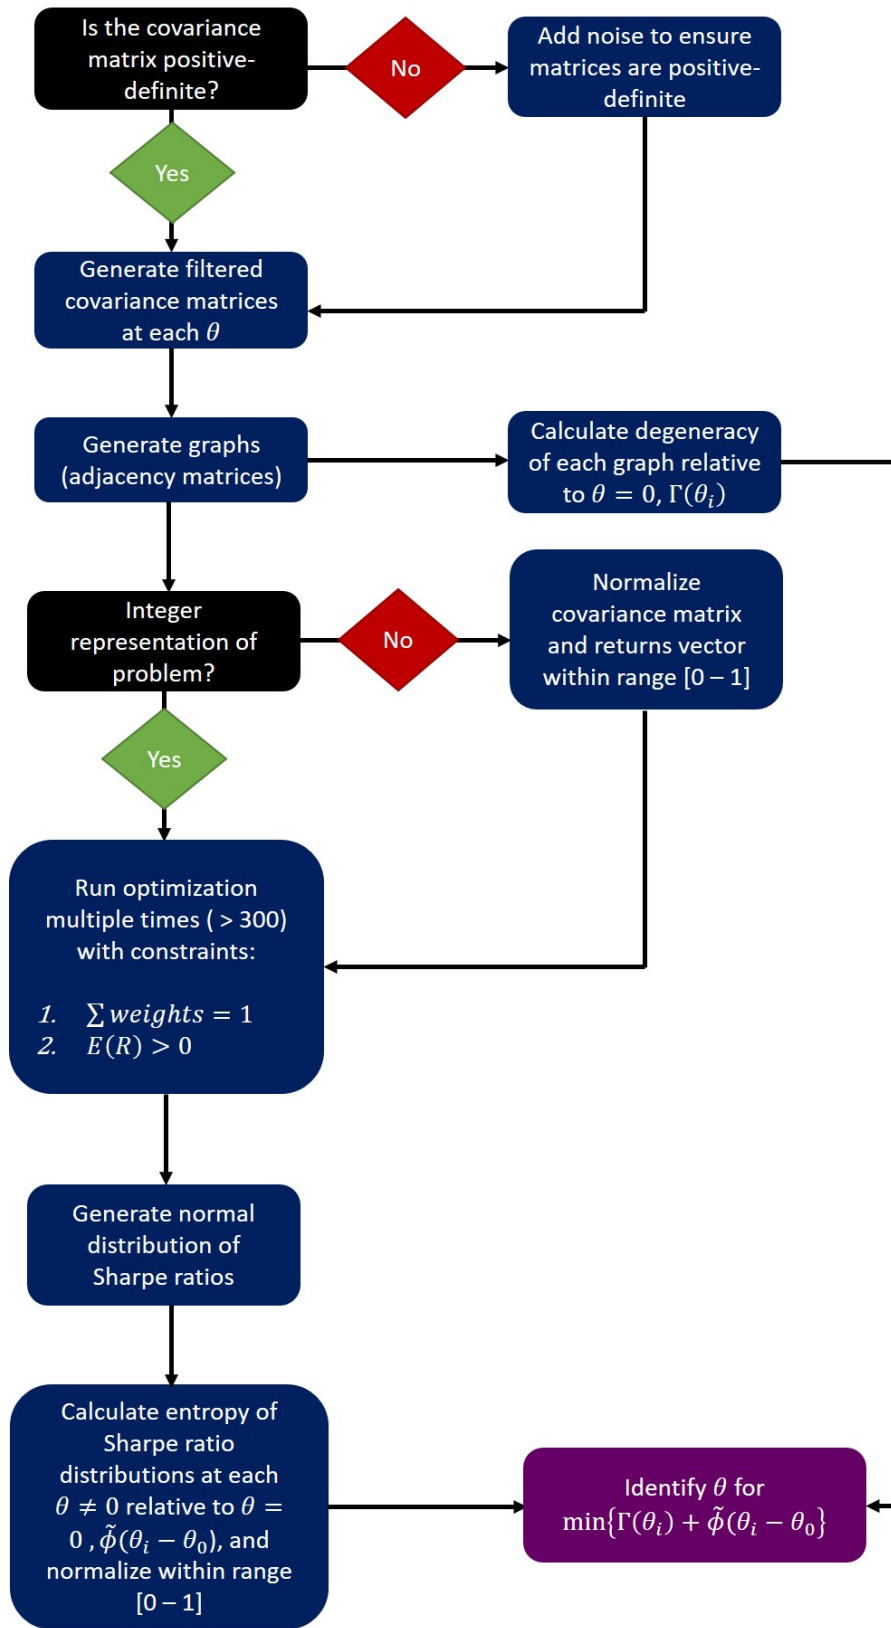


Figure 4.7: Proposed procedure for decomposing an asset portfolio. It should be noted, per the results of this work, that an appropriate optimization algorithm should be selected for this procedure. We suggest using a heuristic search algorithm such as SA.

ried out with the interior/direct algorithm, the expected return vector and the covariance matrix are normalized, producing values within a range of 0–1. In contrast, SA is carried out using an integer representation of the problem, and no such normalization is required. The choice for an integer representation is motivated by the necessity to eventually map the problem to a more advanced mathematical form in order to implement simulated quantum annealing (future work).

During the optimization procedure, the budget constraint is applied to ensure the sum of weights does not exceed the total possible investment of 100%. Additionally, the weights are constrained to remain within 5% of the total investable budget per asset, while another constraint forces the sum of returns to be greater than zero. The latter constraint is added to guarantee that portfolios that incur a loss are ignored, since the goal of making an investment is ultimately to increase the entities’ wealth.

A comparison of heatmaps built with either SA (Figure 2.3) or the KNITRO interior/direct algorithm (Figure 2.4) shows that the former results in more noticeable colour bands in purchase frequencies, suggesting that SA allows us to explore a larger space of solutions. This is further reflected in the distribution of weights in Figure A.2 and Figure A.3 in the Appendix, where, overall, the asset weights have a greater variance when using the SA algorithm.

Additionally, the budget allocation over all assets identified by the KNITRO algorithm varies little between assets, with only one asset having a weight that is slightly higher than the others at 1.82% (see Figure A.8 in the Appendix). This corresponds to asset 28, which coincides with a bond issued by a European supranational agency. In contrast, there is a much more significant variation in the budget allocation found with SA. Furthermore, the asset with the largest budget allocation is asset 49, rather than 28. Looking at the returns and risks of the individual assets, asset 49 coincides with the asset that yields the largest return. In this regard, SA appears to identify a more beneficial budget allocation. It is likely that the small increase in budget allocated to asset 28 is more the result of the constraints used in the optimization procedure.

As seen with the EU portfolio, the budget allocation in the optimal US portfolio identified by the KNITRO algorithm varies little between assets. Only asset 28 receives an allocation below 4.00%, at 2.53%. Looking at Figure A.9a (in the Appendix), this asset has the smallest relative return but also a large associated risk with respect to the other securities. This is further confirmed in Figure A.9c, where the same asset is seen with the smallest relative returns and a larger associated risk. Figure A.9d also shows this asset carrying the smallest weight in the optimal risky portfolio. Another feature similar to the EU portfolio analysis is that the assets that are allocated the largest portion of the budget differ, where SA assigns asset 16 the largest amount of the investor’s budget. In this instance, the KNITRO algorithm also assigns this asset a maximal amount of the budget, but since the weighting varies little between assets, this is not apparent. Looking at both Figure A.9a and Figure A.9c, we see that this security has notably the highest expected return with a smaller associated risk. Interestingly, in Figure A.9d, asset 17 is assigned the next largest portion of the investor’s budget, at 3.93%, and in Figure A.9c this asset shows a very similar expected return, but also has a slightly larger associated risk. Hence, the greater sensitivity of SA again captures the risks and returns of the assets in either portfolio with a greater level of reliability.

A final interesting characteristic of SA is that the optimal risk portfolios for both collections of assets are more diversified. For instance, in the optimal EU portfolio identified by the KNITRO algorithm, a greater number of assets are excluded, and the same is true for the optimal US portfolio. To give this result context, we compare the expected returns and risks of the optimal portfolios identified by either algorithm. For the EU portfolio, the expected rates of return are 5.25% and 5.40% for the KNITRO algorithm and SA, respectively. The risks for the same portfolio are 0.033 and 0.034, as identified by the gradient-based optimization and SA, respectively. Hence, for a marginal increase in risk, the investor would receive a larger expected return following the budget allocation outlined by the SA results. A

similar trend is observed for the optimal US portfolios identified by either algorithm, where SA results in a more diversified portfolio. In this instance, the return and risk identified for the optimal portfolio by the SA algorithm are 6.60% and 0.039, respectively. On the other hand, the optimal portfolio identified by the gradient-based optimization yielded a rate of return of 6.25% with an associated risk of 0.037.

6 Conclusion

In this study, we decompose two distinct covariance matrices representing either EU or US asset dependencies into multiple financial graphs in an attempt to find reduced yet optimal representations of our portfolio selection problem. Our threshold-based procedure induces artificial diversification, impacting network topologies, asset selection and the contribution of each asset to the risk and return of its respective portfolio. The structural analysis of the resulting networks allows the identification of two types of nodes, bulk and surface ones, displaying opposite geometric properties. In general, we find that surface nodes have high eccentricity and low betweenness centrality and correspond to low-risk assets. As in Markowitz mean-variance optimization, we show that assets that have the least influence over the correlated network are highly targeted for investments. The risks and returns in the EU portfolio are primarily driven by bond issuers based in Germany. On the other hand, US assets influence each other more equally. The “superimposition” of each threshold-based topology allows for the reconstruction of the efficient frontiers of the two portfolios, while we identify the optimal risky portfolio as the point with the highest Sharpe ratio. Therefore, we suggest the portfolio optimization problem can be simplified by ignoring assets’ contribution to risk or by cancelling covariance components for the assets that have little influence on the network. In future work, we want to investigate additional sampling methods, both classical and quantum-based, to refine the optimal risky portfolio.

References

- [1] Markowitz, H. Portfolio selection. *J. Financ.* 1952, 7, 77 – 91.
- [2] Merton, R. On Estimating the expected return on the Market. *J. Financ. Econ.* 1980, 8, 323 – 361.
- [3] Jobson, D. J., Korkie, B. M. Estimation for Markowitz efficient portfolios. *J. Am. Stat. Assoc.* 1980, 75, 544 – 554.
- [4] Fabozzi, J., Kolm, P. N., Pachamanova, A. D., Focardi, S. M. *Robust portfolio Optimization and Management*. 2007, John Wiley & Sons Inc., First Ed. New Jersey.
- [5] Dempster, M. A. H., Mitra, G., Pflug, G. *Quantitative Fund Management*. 2008, First Ed. Chapman and Hall/CRC, New York.
- [6] Hegazy, O., Soliman, O., S., Salam, M. A. A Machine Learning Model for Stock Market Prediction. *International Journal of Computer Science and Telecommunications*. 2013, 12, 16 – 23.
- [7] Parida, P. K., Sahoo, S. K. *Portfolio Optimization Theory in Computational Finance*. *Global Journal of Pure and Applied Mathematics*. 2015, 11, 941 – 948.
- [8] Bloomfield, T., Leftwich, R., Long, J. Portfolio strategies and performance. *J. Financ. Econ.* 1977, 5, 201 – 218.
- [9] Jorion, P. Bayesian and CAPM estimators of the means: implications for portfolio selection. *J. Bank. Financ.*, 1991, 15, 717 – 727.
- [10] DeMiguel, V., Garlappi, L., Uppal, R. Optimal versus naïve diversification: how inefficient is the 1/N portfolio strategy? *Rev. Financ. Stud.* 2009, 22, 1915 – 1953.
- [11] Mantegna, R. N. Hierarchical structure in financial markets. *Eur. Phys. J. B – Condens. Matter*. 1999, 11, 193 – 197.
- [12] Vandewalle, N., Brisbois, F., Tordoir, X. Non-random topology of stock markets. *Quant. Financ.* 2001, 1, 372 – 374.
- [13] Billio, M., Getmansky, M., Lo, A. W., Pelizzon, L. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *J. Financ. Econ.* 2012, 104, 535 – 559.
- [14] Diebold, F. X., Yilmaz, K. On the network topology of variance decompositions: measuring the connectedness of financial firms. *J. Econ.* 2014, 182, 119 – 134.
- [15] Hautsch, N., Schaumburg, J., Schienle, M. Financial network systemic risk contributions. *Eur. Finan. Rev.* 2015, 9, 685 – 738.
- [16] Peralta, G. Network-based measures as leading indicators of market instability: the case of the Spanish stock market. *J. Netw. Theory Financ.* 2015, 1, 91 – 122.
- [17] Zareei, A. Network centrality, failure prediction and systemic risk. *J. Netw. Theory Financ.* 2015, 1, 73 – 97.
- [18] Tumminello, M., Aste, T., Di Matteo, T., Mantegna, R. N. A tool for filtering information in complex systems. *Proceedings of the National Academy of Sciences*. 2005, 102, 10421 – 10426.
- [19] Aste, T., Di Matteo, T. Dynamical networks from correlations. *Physica A*. 2006, 370, 156 – 161.
- [20] Tumminello, M., Lillo, F., Mantegna, R. N. Correlation, hierarchies, and networks in financial markets. *J. Econ. Behav. Organ.* 2010, 75, 40 – 58.

- [21] Peralta, G., Zareei, A. A network approach to portfolio selection. *J. Emp. Financ.* 2016, 38, 157 -- 180.
- [22] Sharpe, W. F. Mutual Fund Performance. *J. Bus.* 1966, 39, 119 – 138.
- [23] Shor, P. W. Polynomial-Time Algorithms for Prime Factorization Logarithms on a Quantum Computer. *SIAM J. Sci. Statist. Comput.* 1999, 41, 303 – 332
- [24] Farhi, E., Goldstone, J., Gutmann, S., Sipser, M. Quantum Computation by Adiabatic Evolution, 2000, quant-ph/0001106
- [25] Bian, Z., Chudak, F., Macready, W. G., Rose, G. The Ising model: teaching an old problem new tricks. D-Wave Systems, 2010.
- [26] Elsokkary, N., Khan F. S., La Torre, D., Humble, T. S, Gottlieb, J. Financial Portfolio Management using D-Wave Quantum Optimizer: The Case of Abu Dhabi Securities Exchange. tech. rep., Oak Ridge National Lab, 2017.
- [27] Rosenberg, G., Haghnegahdar, P., Goddard, P., Carr, P., Wu, K., López de Prado, M. Solving the Optimal Trading Trajectory Problem using a Quantum Annealer. *IEEE*, 2016, 10, 1053 – 1060.
- [28] Venturelli, D., Kondratyev, A. Reverse Quantum Annealing Approach to Portfolio Optimization Problems. *Quantum Mach. Intell.*, 2018, 1, 17 – 30.
- [29] Bertsimas, D., Tsitsiklis, J. Simulated Annealing. *Stat. Sci.* 1993, 8, 10 – 15
- [30] de Falco, D., Tamascelli, D., An Introduction to Quantum Annealing. *RAIRO:ITA.* 2011, 45.
- [31] Crama, Y., Schyns, M. Simulated Annealing for Complex Portfolio Selection Problems. *Eur. J. Oper. Res.* 2003, 150, 546 -- 571.
- [32] P. Hage and F. Harary. Eccentricity and Centrality in Networks. *Social Networks.* 1995, 17, 57 -- 63.
- [33] Sabidussi, G. The Centrality index of a Graph. *Psychometrika.* 1966, 31, 581 -- 603.
- [34] Freeman, L. Centrality in Social Networks Conceptual Clarification. 1978, 1, 215 -- 239.
- [35] Gökaynak R. s., Sack B., Wright J.H. The U.S Treasury Yield Curve: 1961 to the Present. *J. Monet. Econ.* 2007, 54, 2291 – 2304
- [36] Jobson, D. J., Korkie, B. M. Estimation for Markowitz efficient portfolios. *J. Am. Stat. Assoc.* 1980, 75, 544 — 554.
- [37] Mantegna, R. N. Hierarchical Structure in Financial Markets. *Eur. Phys. J. B – Condens. Matter.* 1999, 11, 193 — 197.
- [38] Byrd R. H., Nocedal J., Waltz R. A. KNITRO: An Integrated Package for Nonlinear Optimization. In: Di Pillo G., Roma M. Large-Scale Nonlinear Optimization. Nonconvex Optimization and Its Applications, 2016, 83, Springer, Boston, MA

Appendix

A.1 Adding Noise to Covariance Matrices to Ensure They are Positive Semi-definite

The two fundamental ingredients of portfolio selection are the expected return for each stock, which represents the portfolio manager’s ability to forecast future price movements, and the covariance matrix of stock returns, which represents a means of controlling risk. The standard statistical method for obtaining the covariance matrix is to gather a history of past stock returns and compute the sample covariance over a set period. To ensure that the resulting covariance matrix of stock returns is positive semi-definite, noise has been added to the initial covariance matrices. To do so, the following procedure, adapted from Hardin et al.[1], was applied. Let:

1. Σ be a given $N \times N$ covariance matrix.
2. The maximum level of noise, ϵ , be constrained $0 < \epsilon < \lambda_N(\Sigma)$, such that the noise does not exceed the largest eigenvalue of Σ .
3. M be a positive integer that describes the dimension of the noise space.

Then, select N unit vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N$ from \mathbb{R}^M and form the $M \times N$ matrix $U = (\mathbf{u}_1 | \mathbf{u}_2 | \dots | \mathbf{u}_N)$ whose columns are the u_i . Then, the $N \times N$ matrix

$$S = \Sigma + \epsilon(U^T U - I) \quad (1)$$

is a covariance matrix whose entries satisfy $|S_{ij} - \Sigma_{ij}| \leq \epsilon$ for $1 \leq i, j \leq N$ and whose condition number, $\kappa(S)$ satisfies

$$\kappa(S) \leq \frac{\lambda_1(\Sigma) + (N - 1)\epsilon}{\lambda_N(\Sigma) - \epsilon} \quad (2)$$

A.2 Results of Shapiro-Wilks Test for Normality of Distribution of Expected Returns

Table A.1: P -values for distribution of expected returns from the EU portfolio across covariance thresholds.

θ	Simulated Annealing	Gradient-based Optimization
0	0.052	1.8×10^{-35}
10	0.189	1.1×10^{-31}
20	0.367	6.8×10^{-23}
30	0.594	7.8×10^{-19}
40	0.123	6.1×10^{-20}
50	0.057	1.1×10^{-23}
70	0.701	7.2×10^{-24}
90	0.439	1.2×10^{-24}

Table A.2: P -values for distribution of expected returns from the US portfolio across covariance thresholds.

θ	Simulated Annealing	Gradient-based Optimization
0	0.083	1.4×10^{-21}
10	0.055	1.3×10^{-20}
20	0.734	6.1×10^{-24}
30	0.054	1.4×10^{-24}
40	0.335	1.7×10^{-23}
50	0.130	4.2×10^{-24}
70	0.182	2.7×10^{-24}
90	0.954	1.5×10^{-24}

A.3 Distribution of Weights for the Optimized Portfolios




Figure A.1: Distribution of optimized weights for the EU portfolio obtained using different covariance thresholds as the seed for optimization carried out with (a) simulated annealing and (b) gradient-based optimization.

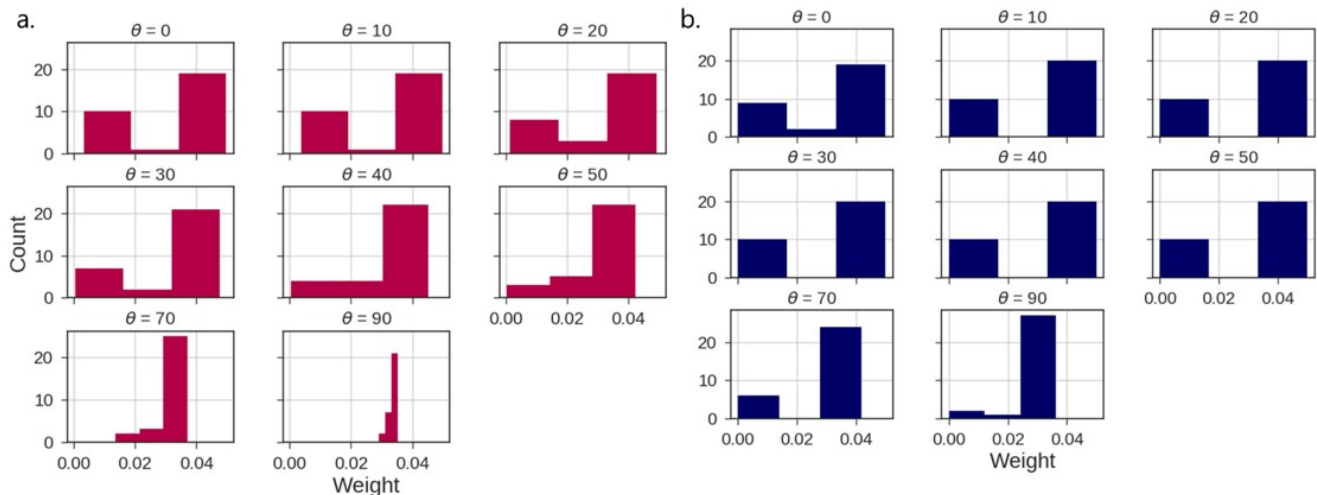


Figure A.2: Distribution of optimized weights for the US portfolio obtained using different covariance thresholds as the seed for optimization carried out with (a) simulated annealing and (b) gradient-based optimization.

A.4 Distribution of Expected Returns over Covariance Thresholds

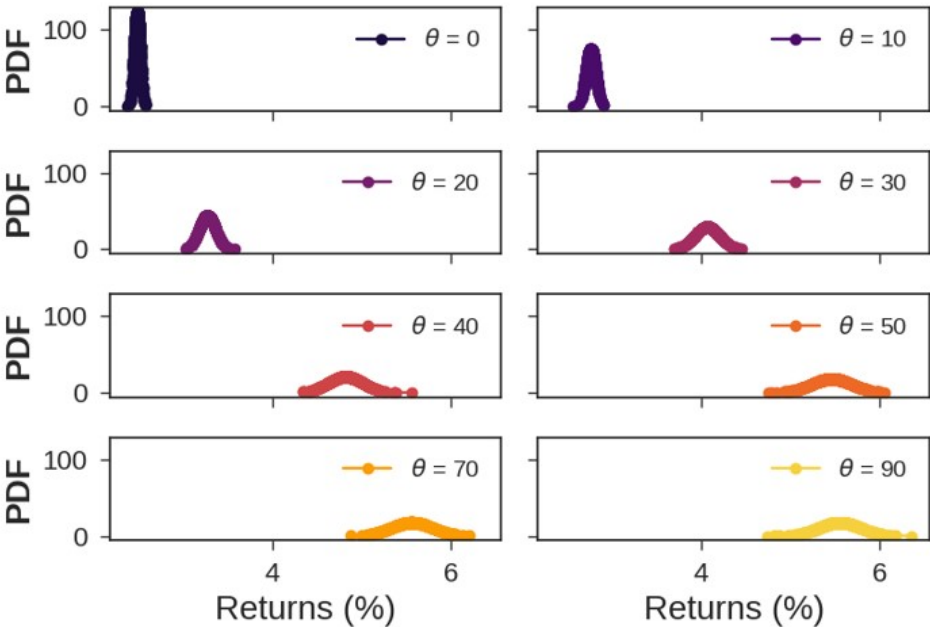


Figure A.3: Normal distributions of expected returns for the EU portfolio over varying covariance thresholds. These distributions are obtained for a risk-aversion coefficient of 11 using the SA algorithm. In all cases shown, 500 samples are used to generate the probability density functions (PDF).




Figure A.4: Normal distributions of expected returns for the US portfolio over varying covariance thresholds. These distributions are obtained for a risk-aversion coefficient of 5 using the SA algorithm. In all cases shown, 500 samples are used to generate the PDFs.

A.5 Distribution of Sharpe Ratios Obtained with SA




Figure A.5: Normal distributions of Sharpe ratios for the EU portfolio over varying covariance thresholds. These distributions are obtained for a risk-aversion coefficient of 11 using the SA algorithm. In all cases shown, 500 samples are used to generate the PDFs.

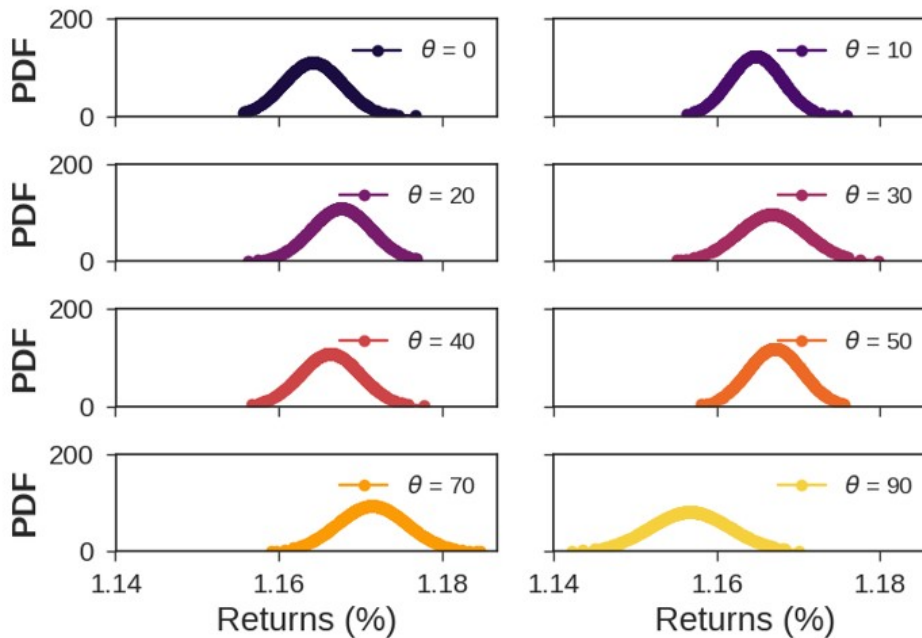


Figure A.6: Normal distributions of Sharpe ratios for the US portfolio over varying covariance thresholds. These distributions are obtained for a risk-aversion coefficient of 5 using the SA algorithm. In all cases shown, 500 samples are used to generate the PDFs.

A.6 Categorization of Assets in Optimal Risky Portfolio Obtained with Gradient-based Optimization

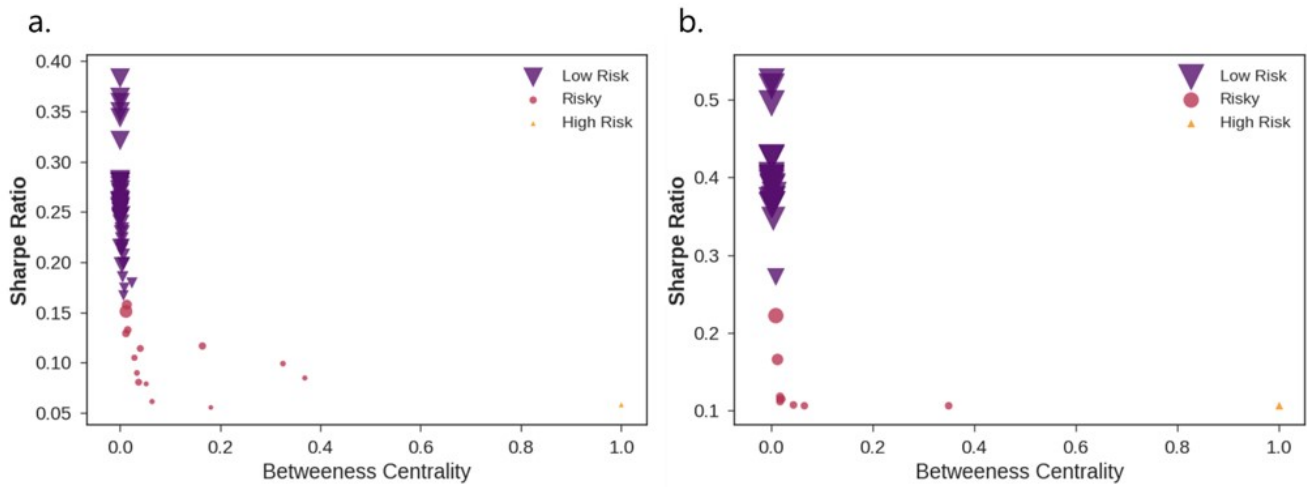


Figure A.7: Normal distributions of Sharpe ratios for the US portfolio over varying covariance thresholds. These distributions are obtained for a risk-aversion coefficient of 5.

A.7 Optimal Risky Portfolios

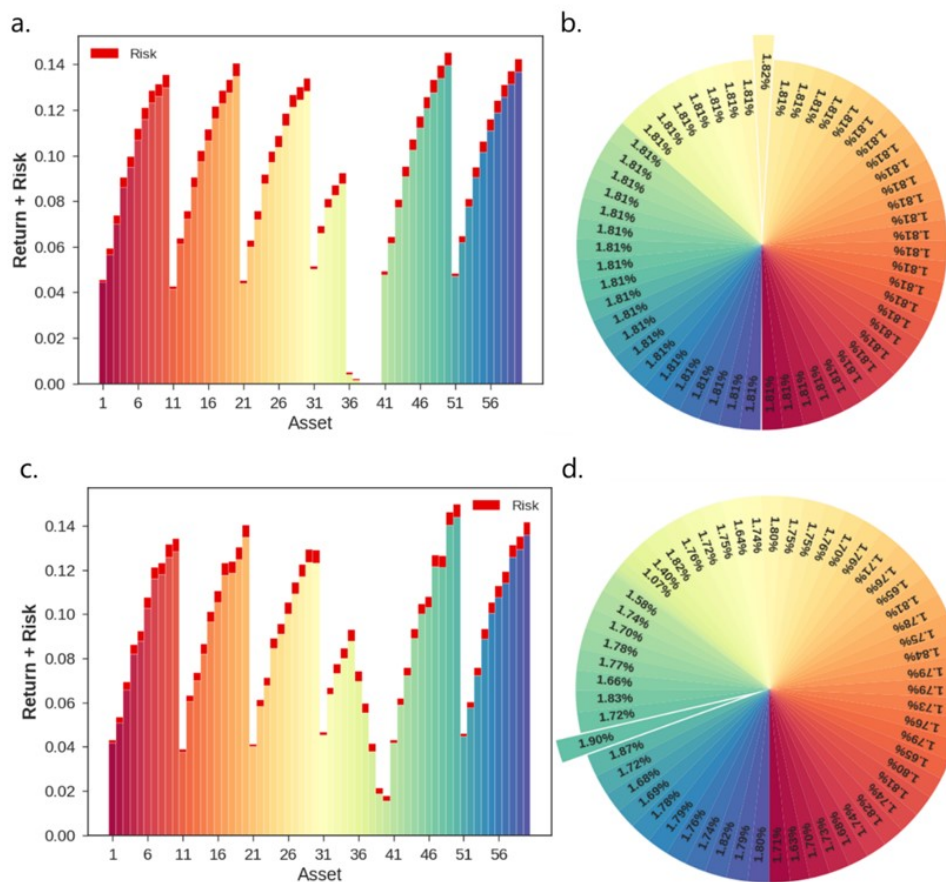


Figure A.8: Analysis of the (a) risks and returns by asset and (b) budget allocation per asset in optimal EU portfolio using gradient-based optimization, and (c) risks and returns by asset and (d) budget allocation per asset from SA results. Assets for which the budget allocation fell below 1.00% are not labelled in the pie charts.




Figure A.9: Analysis of the (a) risks and returns by asset and (b) budget allocation per asset in optimal US portfolio using gradient-based optimization, and (c) risks and returns by asset and (d) budget allocation per asset from SA results. Assets for which the budget allocation fell below 1.00% are not labelled in the pie charts.

References

[1] Hardin, J., Garcia, S. R., Golan, D. A method for generating realistic correlation matrices. *Annals of Applied Statistics*. 2013, 3, 1733-1762.