Summer 2023

Optimal Portfolio Construction for Oil-Based Sovereign Wealth Funds

Mohammed Saeed Alshowaikhat
Claremont Graduate University

Optimal Portfolio Construction for Oil-Based Sovereign Wealth Funds

By
Mohammed Saeed Alshowaikhat

Claremont Graduate University
2023

Approval of the Dissertation Committee

This dissertation has been duly read, reviewed, and critiqued by the Committee listed below, which hereby approves the manuscript of Mohammed Saeed Alshowaikhat as fulfilling the scope and quality requirements for meriting the degree of Doctor of Philosophy in Economics.

Thomas Willett, Co-Chair
Claremont Graduate University
Horton Professor of Economics Emeritus

Pierangelo De Pace, Co-Chair
Pomona College
Associate Dean of the College; Professor of Economics

Benjamin Gillen
Claremont McKenna College
Associate Professor of Economics

Abstract

Optimal Portfolio Construction for Oil-Based Sovereign Wealth Funds

By

Mohammed Saeed Alshowaikhat

Claremont Graduate University: 2023

Chapter 1 of this dissertation delves into the economic challenges faced by oil-exporting countries that rely heavily on a single income source, with a particular focus on Saudi Arabia as a case study. The primary objective is to examine the efforts of Saudi Arabia's sovereign wealth fund in diversifying revenue streams and mitigating risks associated with an excessive dependence on oil. To achieve this, the study proposes an adaptation of the subset-optimization algorithm within the mean-variance model, aiming to enhance portfolio construction in sovereign wealth funds.

Chapter 2 of the dissertation conducts a comparative analysis between portfolios constructed using the subset-optimization algorithm and a benchmark portfolio that does not employ the algorithm. The findings show that the subset-optimized portfolios outperform the benchmark across various performance metrics. Notably, these portfolios exhibit higher Sharpe ratios, greater investor utility, and lower volatility compared to the benchmark. Additionally, the use of the algorithm leads to reduced exposure to oil beta across different subset sizes. Notably, as the subset size decreases, the portfolio's volatility also decreases, suggesting the algorithm's effectiveness in diversification.

In Chapter 3, the research explores advanced estimation strategies for portfolio construction, considering two distinct cases. The first case incorporates a four-factor model, including the Carhart four-factor model, along with an additional factor specifically related to oil.

The second case uses the Bayesian-shrinkage estimator for estimating the variance–covariance matrix and incorporates an informative prior within a Bayesian framework to estimate expected returns. Comparisons among the different models and inputs demonstrate that these advanced estimation techniques lead to improved portfolio performance. Specifically, the Sharpe ratio and investor utility are enhanced, indicating the contribution of these cutting-edge techniques to the creation of more effective and efficient portfolios.

Overall, the findings highlight the algorithm's potential to enhance risk-adjusted returns, reduce exposure to specific market factors such as oil, and ultimately contribute to the overall enhancement of portfolio management.

Acknowledgements

I would like to express my deepest gratitude to the following individuals who have been instrumental in the successful completion of my dissertation: First and foremost, I extend my heartfelt appreciation to Professor Tom Willett for his unwavering guidance, invaluable insights, and constant encouragement throughout my research journey. His expertise and mentorship have been pivotal in shaping my academic growth. I am also immensely grateful to Professor Ben Gillen, whose thoughtful feedback and constructive criticism have significantly enriched the quality of my dissertation. His commitment to excellence has inspired me to push the boundaries of my research. I extend my sincere thanks to Professor Pierangelo De Pace for his valuable contributions and support during the development of my dissertation. His expertise in the subject matter has been a source of inspiration for my work. My heartfelt thanks also go to my parents, whose unwavering love, encouragement, and sacrifices have been the cornerstone of my academic pursuits. Their continuous support has been the driving force behind my accomplishments. I would like to acknowledge my brother and two sisters for their unwavering belief in my abilities and their constant encouragement. Their presence in my life has been a source of strength and motivation. This dissertation would not have been possible without the unwavering support and encouragement from all these incredible individuals. Their belief in me and their contributions have been a guiding light throughout this endeavor. Thank you all for being an integral part of my academic journey and for shaping me into the researcher I am today.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction and Overview

Countries heavily reliant on a single source of income face significant economic risks, similar to investors who concentrate their portfolio in a single stock. Recognizing this, countries with oil-based economies have acknowledged the perils of overdependence on oil revenue and established sovereign wealth funds (SWFs). The Kuwaiti SWF, established in 1953,

was the pioneer in this regard; established in 1990, the Norwegian SWF has emerged as the largest fund globally, boasting a remarkable current asset value of $1.37 trillion (Sovereign Wealth Fund Institute, 2023). Currently, there are over 100 SWFs globally, collectively managing $11.49 trillion in assets. SWFs, if sufficiently large, can play a crucial role in reducing reliance on revenue derived from natural resources, such as oil. They have the capacity to diversify their investments across various asset classes, including equities, long-term debt, foreign direct investments, real estate, and even cryptocurrencies. Moreover, SWFs can contribute significantly to financing the development of new economic sectors within a country, encompassing areas such as energy transformation, manufacturing, tourism, entertainment, banking, trade, and industry. The concept revolves around utilizing the substantial size of the SWF to finance mega projects. Under this approach, a portion of the fund's annual profits would be allocated towards funding these projects. In instances where the fund experiences losses, it would liquidate its non-losing investments to generate the necessary funds for project financing. As these funds are not derived from the fund's profits, they can be perceived as zero-interest loans, which the state can reimburse to the fund during periods of economic prosperity. SWFs hold considerable importance for oil-dependent economies. It's worth mentioning that according to Yu and Yang (2017), SWFs have the potential to transform oil revenues into investments in various industries using carefully planned strategies. This means diversifying the funds generated from oil sales into different sectors to reduce dependence on oil as the sole source of income. Additionally, the authors suggest that

SWFs can introduce valuable resources and expertise from other countries through cross-border direct investment. By doing so, they bring in scarce factors of production, such as specialized skills and technologies, to the domestic market. This approach helps stimulate economic growth and development by expanding the range of industries and resources available in the country.

One key aspect of this dissertation revolves around the construction of a purposeful portfolio, aimed at ensuring that the returns of the portfolio are uncorrelated with a specific commodity, which aligns with the overarching goal of establishing a SWF that generates returns independent of oil prices. One part of our research centers on the potential of SWF returns to serve as a crucial tool for financing and promoting the development of new economic sectors, with a particular emphasis on Saudi Arabia and its Vision 2030 initiative. The Vision 2030 initiative of Saudi Arabia aims to diversify its economy and foster long-term economic growth. However, to successfully finance these new economic sectors, a stable and dependable stream of income becomes imperative. This becomes even more significant for countries heavily reliant on oil-based economies, as they are susceptible to business-cycle fluctuations that can hinder project funding and result in substantial delays.

Sovereign Wealth Funds offer a viable solution by providing a means to finance these ambitious projects, thereby mitigating the likelihood of delays. By decoupling the portfolio returns of the SWF from the fluctuations in oil prices, countries can ensure that their economic plans remain unaffected by the volatility of the oil market. This enables the funding and support necessary for the development of new economic sectors, contributing to economic diversification and long-term growth.

Our research underscores the significance of constructing purposeful portfolios that exhibit no correlation with a specific commodity. This supports the broader objective of establishing SWFs that generate independent returns from oil prices.

In testing the viability of our technical approach, we make certain assumptions for the purpose of this study. Firstly, we assume that a SWF allocates all its resources exclusively to U.S. equities. While this restriction can be relaxed in future studies to include international asset holdings, we focus on U.S. equities for the sake of feasibility. Additionally, we consider a scenario where the SWF allows both short selling and long positions.

Our primary objective in this study is to create a portfolio for the SWF that generates returns independent from oil returns, given the aforementioned reasons. If found to be successful this approach can be adopted by other countries that are heavily dependent on other types of exports. Furthermore, we aim to maximize the investor's utility and enhance the portfolio's Sharpe ratio. Essentially, our focus lies in developing an optimal portfolio management strategy for a commodity-exporting country's SWF.

To achieve a portfolio that is independent from any form of correlation, we introduce a constraint to the optimization problem. Throughout this dissertation, we employ the Subset Optimization approach by Gillen (2016) to guide our portfolio construction process. By utilizing this methodology, we strive to achieve our objectives of portfolio independence, utility maximization, and enhanced Sharpe ratio.

The chapter is structured as follows:1.1 Literature Overview and Research Utilization: Provides an overview of existing literature and its relevance to the research.1.2 Dissertation Objectives: Outlines the specific objectives of the dissertation.1.3 Case Study: Focuses on Saudi Arabia as a case study, highlighting how the research can benefit its sovereign wealth fund.1.4 Data Sources: Explains the sources of data used in the research. 1.5 Methodologies: Describes the methodologies employed in Chapters 2 and 3. 1.6 Conclusion: Provides a concise summary and serves as a transition to the subsequent chapters.

1.1. Overview

Harry Markowitz's mean-variance optimization model, developed in 1952, is a highly effective tool for determining optimal portfolio weights. However, estimating the inputs for this model presents significant challenges, particularly in cases where portfolios involve a large number of securities, resulting in higher-dimensional problems. The presence of a substantial number of securities in a portfolio introduces complexities in the optimization process, often referred to as the Curse of Dimensionality (Bellman, 1957). This phenomenon can make it challenging to impose specific constraints that align with the investor's objectives and construct purposeful portfolios. Additionally, the complexity of estimating inputs is compounded by the fact that parameters need to be estimated based on samples, leading to increased estimation error since ex post behavior may not be the same as in the sample. As the number of parameters to be estimated grows, the potential for estimation error also increases. It is important to note that even minor estimation errors in the input values can lead to misleading outcomes when deriving Markowitz's mean-variance optimal weights (Kritzman, 2006). It's worth mentioning that during times of financial stability, the estimation of returns and covariance matrices in financial models can be subject to small errors. These errors may become more significant during a crisis due to increased co-movement of returns among stocks, as highlighted by Baur (2012). Baur's findings suggest that financial crises lead to a change in correlations among some stocks; which in turn can contribute to an increase in estimation errors. Therefore, the magnitude of estimation error tends to be amplified during financial crises due to the altered correlation structure among stocks. Kan and Zhou (2007) also pointed out this issue, suggesting "using sample means and the covariance matrix of returns to estimate the optimal portfolio…would lead to very poor out-of-sample performance" (p. 621). The estimation error issue is a big one, as stated by Gillen (2016): "This sensitivity is so severe as to cause some researchers and many professional investment managers to question the relevance of the mean-variance paradigm and the value of incorporating any optimization into the asset allocation decision at all" (p. 1). The accuracy and reliability of a model's outputs are strongly influenced by the accuracy and reliability of its inputs. If the inputs of a model are flawed or inaccurate, the resulting outputs will also be compromised.

In this dissertation one of our objective was to construct a portfolio that exhibits minimal correlation with oil returns. To achieve this objective, we employed Markowitz's mean-variance model and implemented a constraint that ensures the correlations between the assets' returns and oil returns are set to zero. By incorporating this constraint, we aimed to create a portfolio that is less influenced by the fluctuations and movements of the oil market. Therefore, it is crucial to obtain more accurate estimates that are less susceptible to noise; in this paper we use Gillen (2016) subset optimization to get more accurate estimates. This is essential for the constraint in the optimization problem to hold effectively, resulting in a portfolio that maintains its independence from oil price fluctuations.

While we found Markowitz's mean-variance model to be valuable, we encountered challenges in accurately estimating expected returns and the variance-covariance matrix. These estimation issues can have significant implications for the model's outputs. Firstly, there is a risk that the optimal weights may not truly be optimal if the input estimates are inaccurate. Secondly, our constraint aimed to avoid any correlation between asset returns and oil returns, but inaccurate estimates of correlations from the variance-covariance matrix could render the constraint ineffective. In the absence of robust estimations of risk and return, the performance of Markowitz's mean-variance model can be compromised.

As the number of assets in the portfolio increases, the number of estimates required grows rapidly, dealing with a large number of estimates increases the risk of estimation error. Gillen's (2016) "The Subset Optimization for Asset Allocation" offers a solution to the challenge of higher dimensionality in portfolio optimization. The subset optimization algorithm divides the overall optimization task into smaller components. “Rather than optimizing weights for all securities jointly, subset optimization constructs Complete Subset Portfolios” (Gillen, 2016, p. 1). The approach involves breaking down the estimation process into smaller subsets to reduce the overall estimation error. For example, if we consider a subset of 10 stocks, estimating their returns and variance-covariance involves fewer estimations compared to a portfolio of 1000 stocks. By estimating returns for a smaller subset, we can potentially reduce the estimation mistakes. To calculate the number of estimations required for the variance-covariance matrix, we can use the formula n + (n x (n-1) / 2). For instance, with 10 assets in the portfolio, we would need to estimate 55 pairs, whereas for a portfolio of 1000 stocks, it would require 500,500 estimations. It is important to note that these estimations are not perfect from the beginning, so at an aggregate level we end up making more mistakes when we optimize 1000 stocks at once.

To address this, the algorithm takes multiple subsets, performs optimization within each subset, and then averages the weights at the end. This averaging process helps to approach the true optimal weights and mitigates the impact of estimation errors, ultimately improving the overall estimation accuracy. So, the subset optimization algorithm operates by selecting random subsets of assets, calculating their optimal weights for a given constraint, and recording the weights. This process is repeated multiple times with different subsets. Once the optimization is complete, the algorithm calculates the average weight for each security by summing up the weights and dividing by the number of iterations. For example, when an investor decides to have a portfolio comprising 1000 assets, the optimization algorithm runs, records the weights, then aggregates the weights for each asset and then divides them by the number of times the optimization process was executed. As a result, the investor obtains the optimal weights for all 1000 assets at the end of the optimization procedure. This approach ensures that each asset's contribution to the portfolio is adequately considered and balanced, leading to an optimized and well-diversified investment strategy.

This algorithm offers two advantages. Firstly, it reduces the dimensionality of the optimization task, resulting in lower estimation error. Secondly, by performing multiple optimizations and averaging the weights, the algorithm approaches the expected weights more closely, benefiting from the law of large numbers. By performing multiple optimizations and averaging the weights, we improve the estimation of true values by reducing the impact of random variations or noise in the optimization process. As Gillen stated, “As long as the subset size is smaller than the sample size, the variance of individual portfolio weights will be bounded, preventing the curse of dimensionality” (p. 3). In the subset optimization algorithm, the zero correlation with oil constraint is applied N times as the algorithm performs N optimizations. Due to estimation error, it is unlikely that any individual subset’s constraint would achieve a perfect zero correlation with oil. However, by imposing the constraint N times , then aggregating the weights obtained from multiple optimizations and taking their average, the resulting constraint average tends to have a lower correlation with oil than the individual constraints. Since all individual constraints strive for zero correlation with oil, their average is expected to exhibit a correlation closer to zero than any individual constraint. By taking the average we try to reduce the impact of random variations or noise, hence, the average of the constraint would be more accurate. The subset optimization technique played a crucial role in improving the robustness of our risk and return estimates, resulting in reduced estimation error. This reduction in error allows us to effectively enforce desired constraints within the optimization problem. With accurate estimates of asset correlations and the desired commodity or factor, we are able to impose robust constraints that effectively minimize the correlation between the returns of the SWF portfolio and commodity returns.

1.2. Objectives

Our contribution to the existing literature lies in the application of Gillen's subset optimization (2006) for constructing purposeful portfolios by imposing specific constraints. In this dissertation, our main focus was to explore the optimal utilization of subset optimization in portfolio weight optimization. To achieve this, we formulated two key research questions:

i. What is the optimal subset size (number of randomly selected assets in each optimization round) for the model?

ii. Considering the effectiveness of subset optimization, what is the most suitable input source for the optimization problem?

To answer these questions, we developed various models with different subset sizes and risk aversion parameters, evaluating their performance based on key metrics such as average returns, standard deviation, Sharpe ratio, and utility. Our objective was to identify the "sweet spot" for subset optimization that aligns with our goals.

Once we determined the optimal way to employ the algorithm (i.e., identifying the ideal subset size), we proceeded to assess the most effective input source for the subset optimization. We compared different input options, including sample moments, factor models (such as the Carhart four-factor model and the Oil Factor), and a Bayesian model with a shrinkage estimator. This investigation aimed to identify the best input source for our model, shedding light on how subset optimization should be utilized in conjunction with the appropriate input source for the optimization problem.

1.3. The Case of Saudi Arabia:

Saudi Arabia's Vision 2030 aims to transform the country's economy into a more diversified and sustainable one (Vision 2030). Prior to this vision, in 2011, oil revenues accounted for a significant 92% of Saudi Arabia's total government revenue (Saudi Central Bank, 2021). This high percentage clearly indicates a strong positive correlation between Saudi Arabia's budget and oil prices. Figure 1.1 illustrates the correlation between oil prices and Saudi Arabia's annual revenue, revealing a calculated correlation of 97.6%. Over the years, there has been a gradual decline in the proportion of oil revenue to total revenue. For instance, in 2019, the oil revenue to total revenue percentage decreased to 64% (Saudi Central Bank, 2021). Analyzing the breakdown of Saudi Arabia's GDP into main activities, it becomes apparent that the Oil and Gas sector constituted a significant 79.4% of the total GDP in 1973, as depicted in Figure 1.1. However, this percentage has been gradually decreasing over time (Saudi Arabian Government, 2023). While the reduction to 64% in 2019 signifies a substantial improvement, it is imperative for this number to continue decreasing significantly to achieve substantial diversification away from oil.

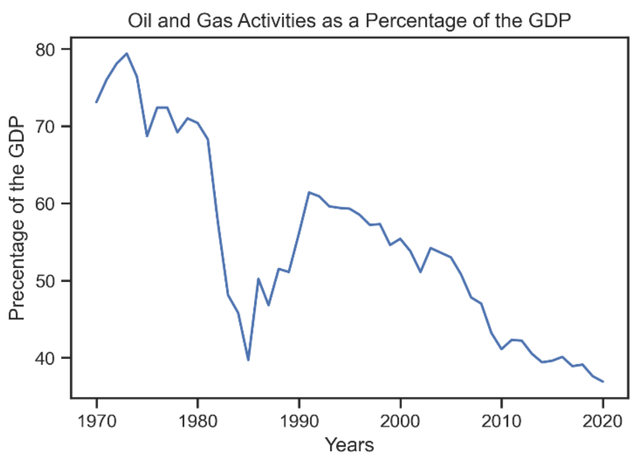

Figure 1.1: Saudi Arabia's Oil and Gas Activities as a Percentage of GDP

In addition to examining the correlation between revenue and oil prices depicted in Figure 1.2, it is essential to consider the concept of elasticity of demand and its impact on total revenue for Saudi Arabia. Different scenarios based on the elasticity of demand should be taken into account, including elastic demand, inelastic demand, and unit elastic demand. Each scenario presents varying effects on Saudi Arabia's revenue. In the context of elastic demand for oil, a reduction in oil prices leads to a higher percentage increase in the quantity demanded comparing with the percentage decrease in price, which would result in an increase in revenue for the oil exporter. Conversely, when oil prices rise, revenue declines due to the percentage decrease in quantity demanded is higher than the percentage increase in price. But, when it comes to inelastic demand for oil, the percentage change in quantity demanded is lower than the change in price in terms of percentage. As a result, a reduction in oil prices would lead to a lower revenue for the oil exporter. However, an increase in oil prices would result in higher revenue, given that the change in percentage when it comes to price is more than the percentage change in quantity demanded. For the unit elastic demand case, where the percentage change in quantity demanded would lead to the same percentage change in price, adjustments in oil prices would not impact the overall revenue.

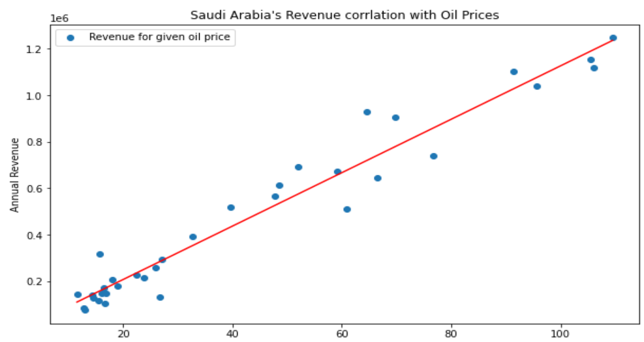

Figure 1.2: Saudi Arabia's Revenue correlation with oil prices

Cooper (2003) highlights that short-run estimations indicate a high level of price inelasticity in oil demand. This implies that fluctuations in oil prices have a limited impact on the quantity demanded of oil in the short term. As discussed earlier, based on this understanding, a decrease in oil prices would result in a decrease in revenue, while an increase in oil prices would lead to an increase in revenue for oil-based economies. This observation aligns with the fact that viable alternative energy sources are limited in the short run. However, it is important to note that the elasticity of demand for oil may change in the long run, particularly as alternative energy sources become more available. In the long run, the elasticity of oil demand could be influenced by the availability and adoption of alternative energy options. Additionally, the price of oil in the short run is influenced by the nature of oil supply shocks. With low elasticities oil prices give a good indicator of revenues. For commodities with elastic demand using price alone would not be a good proxy for earnings from that source. Fortunately in the case of oil exports demand tends to be inelastic for total exports of oil in the short run and this is what is assumed in this dissertation. An example would be responses to shift in world oil demand. However, in the case of supply shocks in particular countries more complicated analysis would be required.

Overall, if the price of oil decreases for an oil-exporting country, it will likely lead to a decline in revenue in the short term and probably in the medium term as well. This dependence on oil revenue exposes oil-based economies to fluctuations in the business cycle, which can further result in delays in ongoing projects. To mitigate such risks and provide a stable source of income for financing projects, SWFs can play a crucial role. SWFs offer a valuable mechanism for oil-exporting countries to establish an independent source of income, reducing their reliance on oil revenue and promoting economic stability.

To facilitate the diversification of its economy away from oil dependency, the government of Saudi Arabia established the Public Investment Fund (PIF). The PIF is regarded as a key driver for achieving economic diversity within the country, as outlined in the Saudi Arabia Vision 2030 objectives. The structure of the PIF's investments encompasses various areas, including Saudi equity holdings, sector development within Saudi Arabia, real estate and infrastructure development, Giga projects, as well as international strategic investments and a diversified international pool (Saudi Central Bank, 2021). Under the international strategic investments and diversified international pool, the PIF focuses on investing in both public and private equity. These investments serve as valuable tools to generate returns that can be utilized to finance the development of new economic sectors. By fostering the growth of these sectors, the PIF plays a crucial role in promoting the diversification of the Saudi Arabian economy, thereby reducing its reliance on oil (Public Investment Fund Program, 2017).

This dissertation focused on a specific case involving a SWF from an oil-exporting economy that made investments in U.S. equities. The primary objective was to achieve a higher level of independence between the SWF's returns and the fluctuations in oil prices. To accomplish this, a zero correlation constraint was imposed between the returns of the SWF and oil returns. The aim was to ensure that the SWF generated a stream of returns that was less influenced by oil price movements. Throughout the course of this study, several key questions were addressed, primarily related to the technical aspects of SWFs. Chapter 2 delved into the exploration of the most effective means of imposing the desired constraint, while Chapter 3 focused on determining the optimal source of input for the optimization problem. By addressing these questions, valuable insights were obtained, shedding light on how to enhance the independence of SWF returns from oil returns within the given context.

1.4. Data

In Chapter 1 of this dissertation, data was collected from the Saudi General Authority of Statistics. For Chapter 2, data was sourced from various databases, including Wharton Research Data Services, the Federal Reserve Economic Data, the Saudi Arabia Central Bank, and DataHub. Additionally, Kenneth French's website was utilized to obtain factor loading data, specifically related to market, high minus low, small over big, and momentum risk premium factors. The data spanned from January 1985 to December 2020.

The data obtained from Wharton Research Data Services consisted of price information for the largest 500 companies in the United States based on market capitalization. To estimate the inputs required for the mean-variance model, the prices of these companies were collected for the 11 years preceding a given year. The first 10 years of data served as the in-sample data, while the prices from the 11th year were used as the out-of-sample data. This process was repeated for each year from 1995 to 2019, with the identification of the largest 500 companies and the collection of their in-sample and out-of-sample prices. The prices were then converted into returns.

Data on oil returns was sourced from the Federal Reserve Economic Data, while information on Saudi Arabia's total revenue and oil revenue was obtained from the Saudi Arabia Central Bank. These datasets formed an integral part of the analysis conducted in this dissertation.

1.5. Methodology

In this dissertation, Markowitz's mean-variance model will be employed to construct optimal portfolios by determining the optimal weights for the assets. The portfolios will be optimized while imposing a zero correlation constraint between the returns of the assets and oil returns, gold returns, and imposing the constraint on both returns simultaneously. The portfolio is designed to generate returns independent of the performance of the oil, and gold market. It can be viewed as a hedge, with the level of hedging determined by the choice of the fund. In this particular study, we assumed that the portfolio returns, and oil, and gold returns are independent, reflecting the preferences of the fund manager. However, if the fund manager desires a negative correlation, the constraint can be adjusted accordingly, or the fund manager can choose to incorporate financial derivatives into the fund's holdings to achieve the desired correlation. We get the commodities returns through the formula log(new price /old price). Gold returns are utilized as a robustness check in this analysis. The selection of gold returns is not driven by a specific reason, but rather as an illustrative example since gold is considered a commodity. It serves as a comparative benchmark to assess the stability and reliability of the results obtained. Chapter 2 will utilize subset optimization to construct a few portfolios, including a benchmark portfolio constructed without using subset optimization. The portfolios will be constructed using historical data of oil, gold and asset prices, specifically utilizing sample data.

The process for obtaining the results will involve the following steps. For each year, a dataset spanning 11 years (10 previous years and 1 current year) will be used as in-sample data. The average returns of the assets and the covariance matrix will be calculated from this data, and they will be input into the mean-variance model to determine the optimal weights. These optimal weights will then be multiplied by the returns of the following year (the 11th year), resulting in realized out-of-sample returns. This process will be repeated for the period from 1995 to 2019, generating monthly time-series data spanning 20 years. The procedure will be applied to all portfolios, leading to time-series returns for each portfolio.

The returns of each portfolio will be used to calculate various performance statistics, including average return, Sharpe ratio, correlations with oil returns, gold returns, and both returns, standard deviation, utility, alpha, and beta.

Finally, the performance statistics of all portfolios will be compared. Chapter 3 will build upon the results obtained in Chapter 2, utilizing the identified “sweet spot” and varying the source of inputs for the model. Specifically, the analysis in Chapter 3 will involve examining the performance when using inputs from factor models and inputs from a Bayesian model with a shrinkage estimator.

1.6. Conclusion

Overall, in this chapter of the dissertation we explore the concept of sovereign wealth funds (SWFs) and their role in reducing the dependence of oil-based economies on oil revenue. We emphasize the importance of diversifying income resources and exports to mitigate economic risks and promote long-term growth. To address this, SWFs have been established as a means to manage and diversify countries’ assets. We introduce the construction of a purposeful portfolio for SWFs that generates returns independent of oil prices. We provided a concise overview of the construction process of the purposeful portfolio, highlighting the crucial role played by subset optimization in its formation. Our contribution to the existing literature lies in the application of subset optimization to construct purposeful portfolios that adhere to specific constraints. We aim to determine the most suitable source of estimation for the model input that aligns with our purpose.

Chapter 2 Unveiling the Benefits of Subset Optimization in Portfolio Management

2.1. Introduction

In this chapter, the dissertation aims to investigate the effectiveness of subset optimization in constructing purposeful portfolios. The purpose of this chapter was to investigate the utility of the subset optimization algorithm in constructing effective portfolios. Specifically, we aimed to answer two main questions:

I. Can the subset optimization algorithm be a valuable tool in constructing purposeful portfolios (portfolios with specific hedge)? If so, does it outperform the benchmark?

II. What are the properties and characteristics of the subset optimization approach that can benefit portfolio managers?

By addressing these questions, we sought to gain insights into the practical applications and advantages of the subset optimization algorithm in portfolio management. We Will compare the performance of portfolios constructed with the subset optimization constraint against portfolios constructed with the standard constraint (benchmark portfolio). Additionally, the study seeks to determine the optimal way to utilize the subset optimization algorithm to achieve the desired outcome. Key considerations include identifying the optimal subset size, assessing the correlation between portfolios generated by the algorithm and a specific commodity, and evaluating their correlation with the benchmark portfolio.

To address these objectives, a constraint will be imposed in which the correlation between oil returns and asset returns is set to zero. Performance statistics will be employed for comparing the portfolios. Robustness checks will be conducted by introducing another constraint, applying zero correlation between gold returns and asset returns, and examining the combined constraint of zero correlation for both oil and gold returns with asset returns.

The chapter aims to identify the "sweet spot" for subset optimization, which represents the subset size that best aligns with the purpose of the study. Moreover, the robustness of the constraint will be assessed to ensure its effectiveness in different scenarios. By addressing these objectives, the dissertation seeks to provide valuable insights into the optimal use of subset optimization in constructing purposeful portfolios while considering correlations with specific commodities and benchmark portfolios.

The chapter follows the following organization: 2.2 Methodologies: In this section, the methodologies utilized in this chapter are presented in detail. It includes an explanation of the specific methodologies employed and an overview of the performance metrics utilized in the analysis. 2.3 Results and Summary: This section discusses the findings of the analysis and provides a summary of the results. It presents the key outcomes and highlights the significant findings from the research. 2.4 Concluding Remarks: In this section, the chapter concludes by offering concluding remarks. It summarizes the main points discussed, provides insights into the implications of the results.

2.2. Methodology

In addition to the methodology discussed in Section 1.5, this chapter will follow the following structure. We will optimize multiple portfolios using Markowitz mean-variance for an investor with a high risk appetite, indicated by a risk aversion score of one. Several portfolios will be optimized using subset optimization, while one portfolio will be optimized using the standard approach (withing using the subset optimization). For the subset optimized portfolios, we will vary the number of subsets (subset sizes: 100, 50, 35, 25, 10) and compare their performance statistics to determine the optimal number of subsets.

To ensure robustness, we will conduct a robustness check by introducing a gold constraint and a combined gold and oil constraint. Following this, we will repeat the same process for an investor with a higher level of risk aversion and another investor with an even higher level of risk aversion. All portfolios will utilize the same source of input, specifically sample moments.

In this section, we will present the results in separate parts, focusing on various levels of investor risk appetite denoted by risk aversion scores of 1, 2, and 10. Throughout this dissertation, we will refer to risk aversion as “Gamma.” We will begin by showcasing the performance metrics for the oil constraint, followed by the metrics for the gold constraint. Furthermore, we will examine the correlations for both cases, presenting the correlations for Gamma 1, Gamma 2, and Gamma 10. To provide a comprehensive analysis, we will display a correlation matrix that represents the correlation between oil, gold returns and the generated portfolio, as well as the correlations among all portfolios. Finally, we will present the results obtained when simultaneously imposing both constraints and thoroughly discuss the performance in this particular scenario. It is important to note that when calculating the Sharpe ratio, we have considered a risk-free rate of zero. When calculating the utility, we employ a formula that combines the average portfolio return with a risk aversion factor, represented by gamma, and the variance of the portfolio. Specifically, the utility is computed as the average portfolio return minus gamma divided by 2, multiplied by the variance of the portfolio. The mean-variance utility function quantifies an investor's preference for balancing expected returns and risk in their portfolio. Investors aim to maximize utility by finding the optimal asset allocation that provides the desired trade-off between returns and risk. By using this function, investors can make more informed decisions about portfolio construction and asset allocation, considering their individual risk preferences and return expectations (Markowitz, 1952).

Utility = (Average Portfolio Return) − (gamma/2) * (Variance of the Portfolio)

Note: The formula assumes a linear relationship between risk and return, where a higher gamma indicates a higher risk aversion.

The betas are calculated using linear regression models. In the case of the oil constraint, the portfolio returns are considered the dependent variable, while the independent variable is the oil returns. Similarly, in the gold constraint case, the portfolio returns are the dependent variable, and the independent variable is the gold returns. When both constraints are imposed simultaneously, a multivariate regression model is used, where the portfolio returns are the dependent variable, and the independent variables are the gold returns and oil returns.

2.3. Empirical results

2.3.1. Oil Constraint

Figure 2.1 presents the performance metrics used to evaluate the portfolios in this section. The portfolios are categorized based on risk aversion levels, indicated by gamma values of 1, 2, and 10. The graph includes the standard portfolio as well as portfolios generated using subset optimization with subset sizes of 100, 50, 35, 25, and 10. The labels on the graphs serve as a reference to level of risk aversion, while the x-axis represents the number of subsets used in the algorithm.

Volatility: A notable pattern observed in Figure 2.1 is the decreasing volatility as the subset size decreases. Regardless of the risk aversion level, the standard portfolio consistently exhibits the highest volatility.

Average Returns: A noteworthy observation is that the average return tends to decrease as the subset optimization is employed. The average return keeps on decreasing as the subset size decreases, the continuous decrease in average return as the subset size decreases is a significant finding. It would be particularly intriguing to determine whether this pattern persists during the robustness check in Sections 2.3.2 and 2.3.3. If the results consistently demonstrate this trend, it would provide valuable insights into optimizing the utilization of subset optimization.

Sharpe Ratio: In the case of Gamma 2 and Gamma 10, we observe an increasing trend in the Sharpe ratio as the subset size decreases. However, for Gamma 1, implementing the subset optimization initially leads to a decrease in the Sharpe ratio, but it outperforms the standard approach when the subset size reaches 25 and 10, as shown in Figure 2.3. This observation is significant, and we need to verify if this trend persists during the robustness check in Sections 2.3.2 and 2.3.3. The reduction in volatility associated with smaller subset sizes drives the improvement in the Sharpe ratio, as we stated earlier that the average return tends to decrease as subset size decreases.

Utility: The utility derived from the standard portfolio is consistently negative across all levels of risk aversion. However, as the subset optimization algorithm is employed, the utility improves and eventually becomes positive with subset size 10.

Beta (Oil Exposure): The beta values in Figure 2.1 represent the exposure of portfolio returns to an oil factor. Upon examining the Beta graph in Figure 2.1, a significant decrease in oil exposure is evident when utilizing the subset optimization compared to the standard approach. However, with subset optimization, the exposure decreases as the subset size is reduced from 100 to 50, and further reduces from 50 to 35. The exposure remains low for subset sizes 35, 25, and 10, but the reduction is not consistently observed as the subset size continues to decrease.

To investigate this further, consider the Beta equation, $\beta = \frac{COV(Port,Oil)}{\sigma_{Oil}^2}$, where the denominator is the covariance of the portfolio and oil, and the denominator is the variance of oil. Rewriting this $\beta = (\sigma_{Port} \times \rho_{Oil,Port})/\sigma_{Oil}$. As observed earlier, the standard deviation of portfolios decrease as the subset size decreases, hence, the reduction in oil exposure can be attributed to the decrease in the portfolio's standard deviation as the subset size decreases, as per the beta equation. In the Gamma 1 case, we observe an increase in beta when the subset size is 10, which can be attributed to an uptick in correlation at that subset size. Similarly, in the Gamma 10 case, we do not see a consistent reduction in beta as the subset size decreases. Instead, we observe an increase in beta when the subset size is 25 and 10, which can be explained by an increase in correlations at those subset sizes.

Correlations: Figure 2.1 displays the correlation patterns, which tend to decrease as the subset size decreases, and then increase again when the subset size reaches 10. These observations indicate the potential effectiveness of the subset portfolio in constructing purposeful portfolios. However, further confirmation through a robustness check is necessary.

Figure 2.2 shows that the standard portfolio exhibits higher correlations with portfolios having larger subset sizes. As the subset size decreases in the algorithm, the correlations tend to decrease consistently across all risk aversion levels. Additional correlation matrices can be found in the appendix. Conducting a robustness check will help determine the optimal subset size and enhance our understanding of effectively utilizing subset optimization in portfolio management.

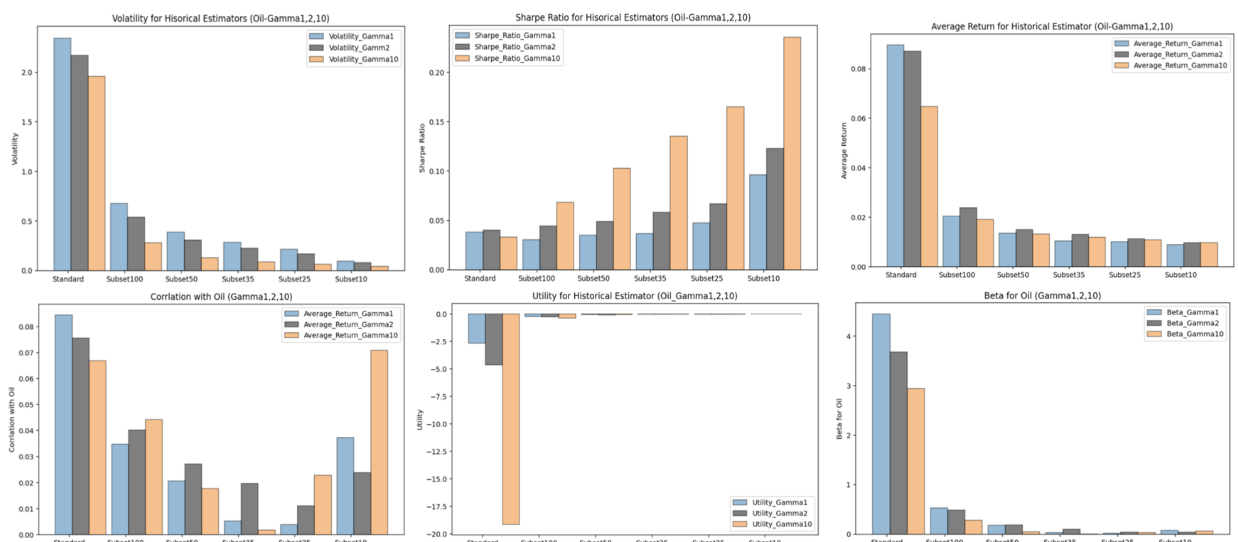

Figure 2.1: Performance Metrics for historical returns, Oil Constraint, all Gammas

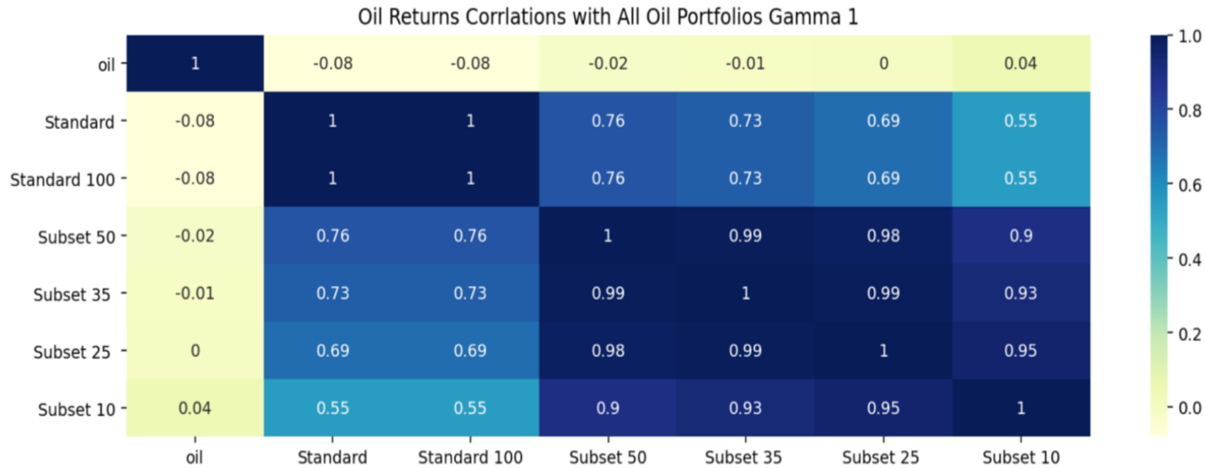

Figure 2.2: Oil Returns Correlations with All (Standard & Subset) Portfolios Gamma 1

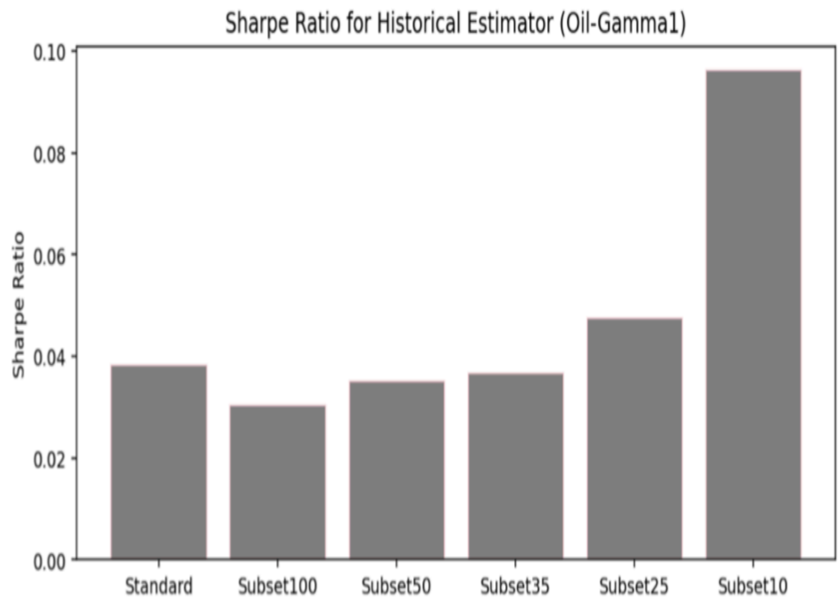

Figure 2.3: Sharpe Ratio for Historical Estimator (Oil Constraint -Gamma 1)

2.3.2. Robustness Check (Gold Constraint)

Figure 2.3 illustrates the performance metrics used to evaluate the portfolios in the Gold section. Similar to the observations in Section 2.3.1, we aim to examine whether the patterns observed in the performance metrics hold here as well. Upon analyzing Figure 2.4, we can clearly observe that the same patterns observed in Section 2.3.1 are evident in this section as well.

Volatility: The observed volatility pattern in Section 2.3.1 holds true, as the implementation of the subset algorithm leads to a decrease in volatility, which further decreases as the subset size decreases.

Sharpe ratio: The increasing trend in Sharpe ratio as the subset size decreases, observed in both Gamma 2 and Gamma 10 portfolios, follows a similar pattern witnessed in Section 2.3.1. However, for Gamma 1 portfolio, the increase in Sharpe ratio begins at subset sizes 35, 25, and 10, resembling the pattern observed in the oil constraint scenario. To validate the consistency of this pattern, an additional robustness check will be conducted in Section 2.3.3. The movement of the Sharpe ratio is depicted in Figure 2.6.

Average return: The average return is initially higher for the standard portfolio but decreases as the subset optimization is employed, with a further decrease as the subset size decreases, exhibiting a similar trend observed in the case of the oil constraint. To validate and further investigate this intriguing finding, additional robustness checks will be conducted in Section 2.3.3.

Utility: Utility shows a significant improvement when the subset optimization is utilized, reaching a positive value once the subset size is 10, the same pattern was observed in Section 2.3.1, and to validate this interesting result, an additional robustness check will be performed in Section 2.3.3.

Beta: A noticeable reduction in the gold beta is observed, with the beta consistently decreasing as the subset size decreases. The pattern of reduction in beta is more pronounced in the gold constraint scenario compared to the oil case. Unlike the oil case, where the beta showed substantial reduction, the relationship between subset size and beta was not linear. However, in the gold case, the relationship is more apparent, with the beta decreasing as the subset size decreases. The decrease in volatility in the gold case has a greater impact compared to the increase in correlation between gold returns and portfolio returns. This is why the beta in the gold case continues to decrease as the subset size decreases. In contrast, although there was a decrease in volatility in the oil case as the subset size decreased, there was also a notable increase in the correlation between oil returns and portfolio returns when subset size was 10. As a result, the beta in the oil case did not consistently decrease with decreasing subset size.

Correlations: Despite the low correlation between gold returns and portfolios' returns, we did not observe a decrease in correlation as the subset size decreased, followed by an increase in correlation with subset size 10, as observed in the oil hedge section. In the gold hedge scenario, we noticed an increase in correlation as the subset optimization was employed, followed by a decrease when the subset size reached 10, the pattern of correlations in this section almost inverted. This contrasting correlation pattern in the gold section compared to the oil section has significant implications, as it hinders us from drawing the conclusion that subset optimization outperforms the benchmark in terms of constructing a purposeful portfolio.

Figure 2.4 illustrates that the standard portfolio demonstrates higher correlations with portfolios possessing larger subset sizes, while the correlation between the standard portfolio and portfolios generated by subset optimization appears to decrease as the subset size decreases. This pattern closely resembles the findings in the oil case as depicted in Figure 2.2. In Section 2.3, we will conduct an additional robustness check to assess whether the observed pattern persists. Figure 2.5 presents the performance metrics used to evaluate the portfolios when both oil and gold constraints were simultaneously imposed in the optimization. This represents a second robustness check to validate the results obtained in Section 2.3.1. The findings in this section will serve as conclusive empirical evidence in this chapter of the dissertation.

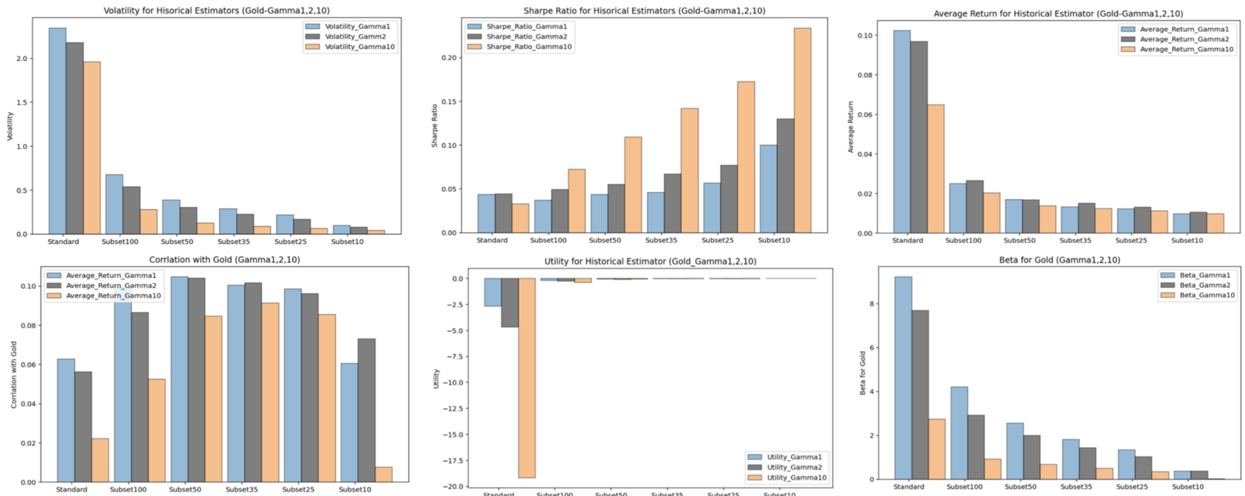

Figure 2.4: Performance Metrics for historical returns, Gold Constraint, all Gammas

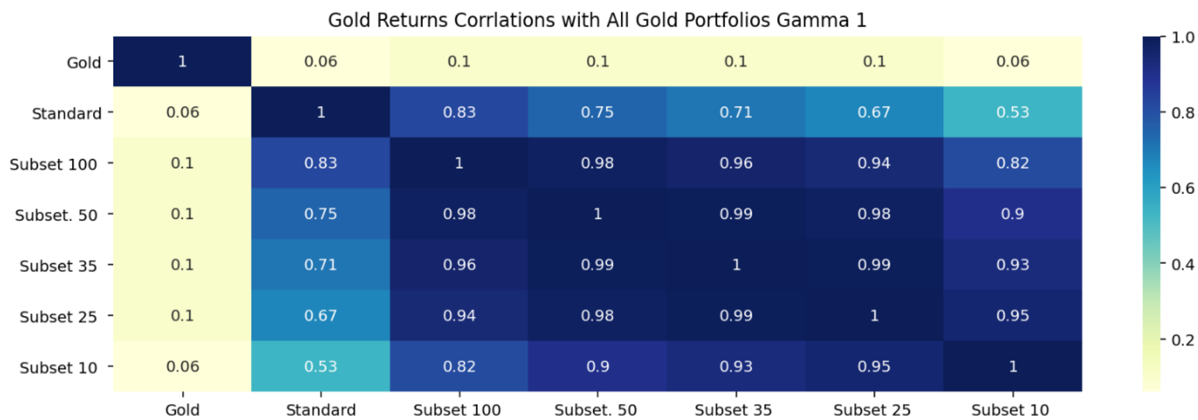

Figure 2.5: Gold Returns Correlations with All (Standard & Subset) Portfolios Gamma 1

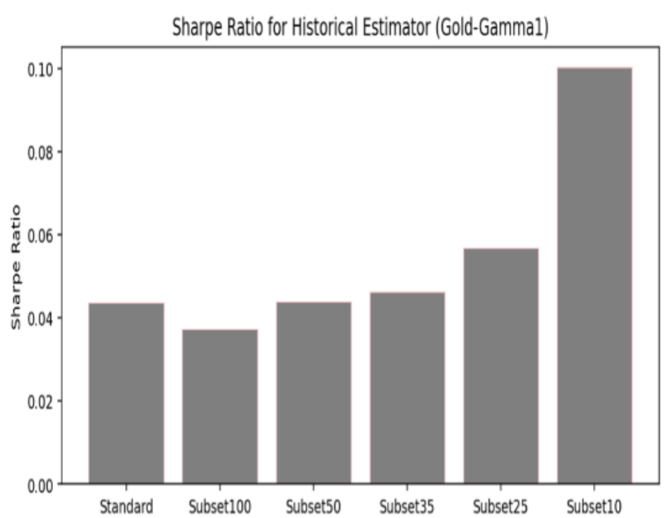

Figure 2.6: Sharpe Ratio for Historical Estimator (Gold Constraint-Gamma 1)

2.3.3. Robustness Check (Oil, and Gold Constraints)

Figure 2.7 presents the performance metrics used to evaluate the portfolios when both oil and gold constraints were simultaneously imposed in the optimization. This represents a second robustness check to validate the results obtained in Section 2.3.1. The findings in this section will serve as conclusive empirical evidence in this chapter of the dissertation.

Volatility: The volatility pattern observed in Sections 2.3.1 and 2.3.2 remains consistent. Implementing the subset algorithm leads to a decrease in portfolio volatility, and this reduction is further amplified as the subset size decreases. These findings provide compelling evidence that the subset optimization significantly contributes to reducing portfolio volatility. This result is noteworthy as it enhances our understanding of the practical utility of this algorithm.

Sharpe Ration: The consistent observation of an increasing trend in the Sharpe ratio as the subset size decreases for both Gamma 2 and Gamma 10 portfolios mirrors the pattern previously identified in Sections 2.3.1 and 2.3.2. For Gamma 1 portfolios, the implementation of the subset optimization initially leads to a decrease in Sharpe ratios when using subset size 100. However, as the subset size decreases, the Sharpe ratios gradually start to increase, until it outperforms the standard portfolio. In the case of the gold constraint, the subset optimization portfolio outperforms the benchmark portfolio when the subset size is 35, 25, and 10. Similarly, in the oil constraint case, the outperformance starts when the subset size is 25 and 10. When both constraints are imposed, the outperformance also occurs when the subset size is 25 and 10. These findings highlight that the subset optimization consistently provides a higher Sharpe ratio, regardless of the investor's risk aversion level, when the subset size in the algorithm is 25 and 10. The robustness check further confirms the significance of this empirical finding.

Average return: As the subset optimization is implemented, the average return of the portfolio decreases in comparison with the standard portfolio, and this reduction is further pronounced as the subset size decreases. This trend mirrors the observations made in the oil constraint and gold constraint scenarios. The findings in this section serve as a secondary robustness check, confirming the empirical result that utilizing the subset optimization leads to a decrease in the average return of the portfolio compared to the standard portfolio that is generated without the algorithm.

Utility: The utility metric demonstrates a substantial improvement when the subset optimization is employed, with a positive value achieved when the subset size reaches 10. This pattern aligns with the observations made in the presence of an oil constraint in Section 2.3.1 and a gold constraint in Section 2.3.2. The findings in this section further validate the empirical result that utilizing the subset algorithm leads to an enhancement in the investor’s utility, regardless of their level of risk aversion. Moreover, it confirms that as the subset size decreases in the subset optimization, the investor’s utility increases correspondingly.

Betas: When considering the gold constraint, we observe a decreasing trend in beta as the subset size decreases, indicating a reduction in exposure between the portfolio returns and gold. However, this trend is not consistent in the oil constraint scenario. For Gamma 1 and Gamma 2, the oil beta increases when the subset size is 10, and for Gamma 10, it increases when the subset size is 25 and 10. These increases in beta are driven by an upsurge in the correlation between oil returns and the subset portfolios in those specific cases. Based on the findings from Sections 2.3.1, 2.3.2, and the current section, we can infer that using the subset optimization generally leads to a decrease in exposure between the portfolio returns and a particular commodity, in comparison with the standard approach (when the algorithm is not used), when a constraint is imposed to ensure that the correlation between the two returns is equal to zero.

Correlation: When both constraints are imposed, we measure the correlation between the portfolio returns and the returns of the respective commodities (oil and gold). The observed correlations with oil exhibit a similar pattern to the correlations obtained when only the oil constraint is imposed. Likewise, the correlations with gold display a similar pattern to those obtained when only the gold constraint is imposed. However, overall, there is no clear and consistent trend in the correlations between the portfolio returns and commodity returns when utilizing the subset optimization. The correlations tend to be low, and they do not consistently outperform the correlations obtained when using the standard portfolio.

Figure 2.7 presents the correlations between the commodity returns and the portfolios generated by the subset optimization, and the standard portfolio returns. Notably, we observe a distinct pattern in the correlation between the standard portfolio and the subset optimization portfolios. Specifically, the correlation is stronger between the standard portfolio and the portfolio generated with a subset size of 100, and this correlation gradually decreases as the subset size decreases. This consistent pattern aligns with the observations made in Sections 2.3.1 and 2.2. The findings in this section confirm the relationship between the standard portfolio and the subset optimization portfolios, which states that the correlation between them decreases as the subset size decreases.

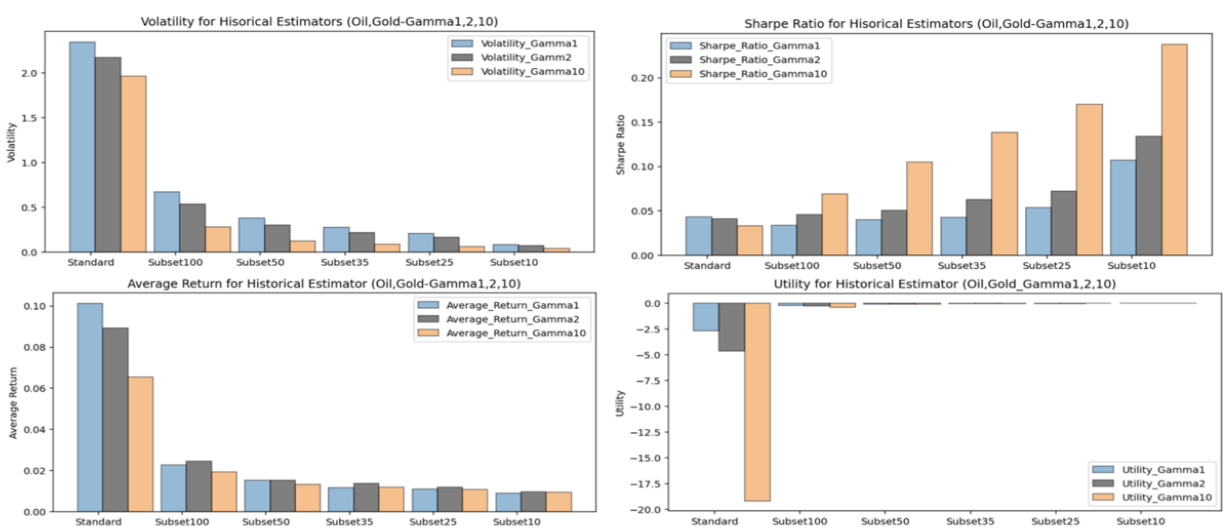

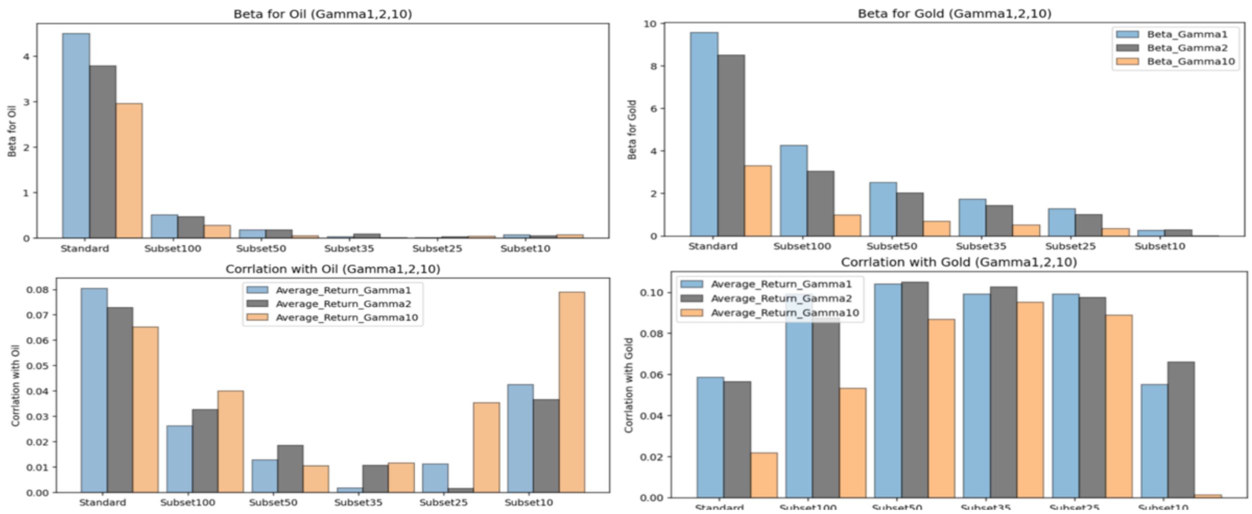

Figure 2.7: Performance Metrics for historical returns, both Constraint, all Gammas

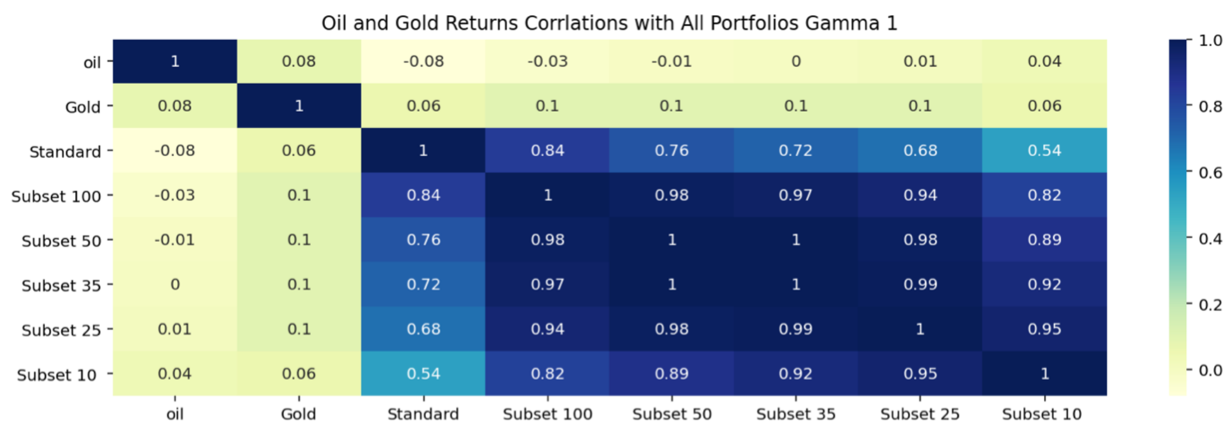

Figure 2.8: Oil and Gold Returns Correlations with All (Standard & Subset) Portfolios Gamma 1

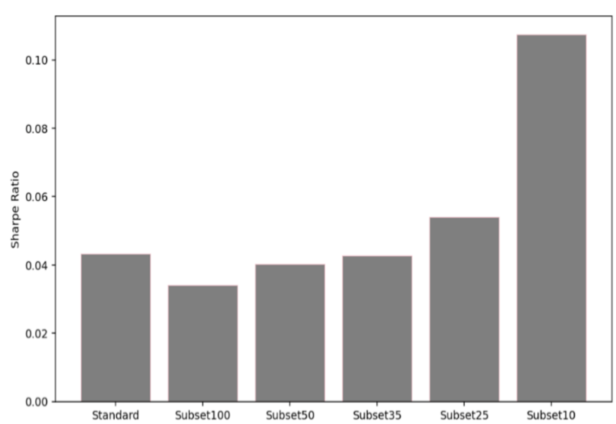

Figure 2.9: Sharpe Ratio for Historical Estimator (Both Constraints -Gamma 1)

2.3.4. Bottom Line

In this section, we will summarize our findings obtained from evaluating the performance metrics after conducting two robustness checks in Sections 2.3.2 and 2.3.3.

Correlations: When constructing portfolios with an oil constraint, we observed a decrease in correlation between the subset portfolios and oil returns as the subset size decreased. However, there was an increase in correlation when the subset size reached 10 for the Gamma 1 and Gamma 2 cases. In the Gamma 10 case, the correlation at subset size 10 was higher than the benchmark. In the gold hedge scenario, we observed an initial increase in correlation as the subset optimization was employed, followed by a decrease at subset size 10, showing an almost inverted pattern compared to the oil hedge correlations, Figures 2.10, 2.11, 2.12 show the correlations. When both constraints were imposed, the correlations with oil and gold displayed similar patterns as when only their respective constraints were imposed. In summary, the correlations between portfolio returns and commodity returns did not demonstrate a clear and consistent trend when employing the subset optimization. These correlations tended to be low, indicating that the subset optimization is useful in constructing purposeful portfolios with reduced exposure to commodity returns. However, it did not consistently outperform the standard portfolio in terms of correlation with commodity returns.

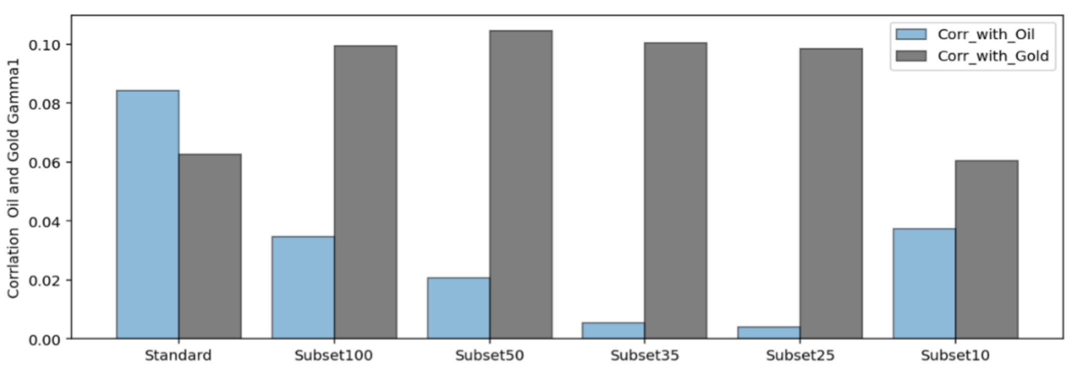

Figure 2.10: Portfolios Correlation with Oil and Gold for Historical Estimator Gamma 1

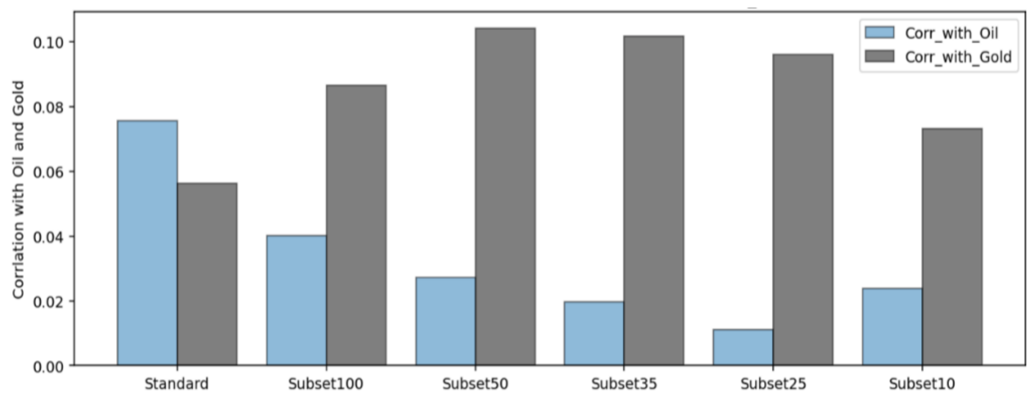

Figure 2.11: Portfolios Correlation with Oil and Gold for Historical Estimator Gamma 2

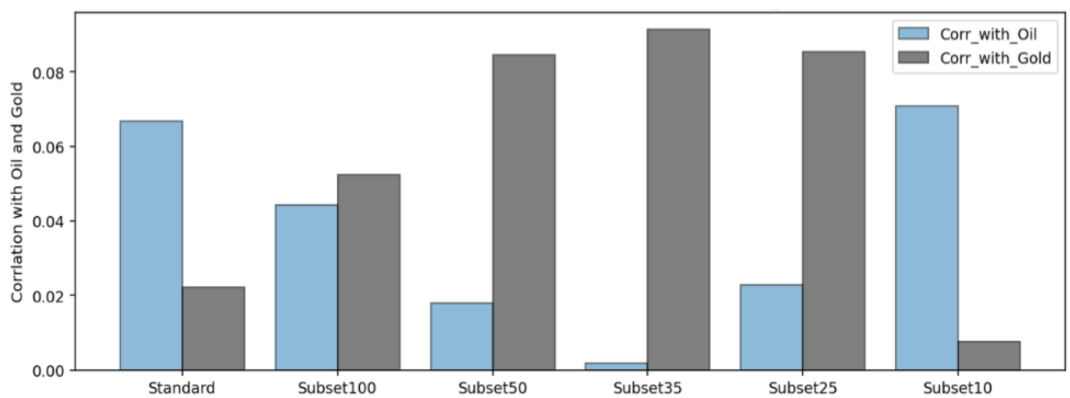

Figure 2.12: Portfolios Correlation with Oil and Gold for Historical Estimator Gamma 10

Sharpe ration: A distinct pattern emerges when examining the Sharpe ratio in relation to the subset optimization. In the cases of Gamma 2 and Gamma 10, it is evident that the Sharpe ratio obtained through the subset optimization consistently outperforms the Sharpe ratio obtained from the standard approach. With Gamma 1, the initial decrease in the Sharpe ratio when employing the subset optimization is followed by a gradual increase, ultimately surpassing the Sharpe ratio of the standard portfolio at subset sizes of 25 and 10, in both the oil constraint and double constraint scenarios. In the gold constraint scenario, the outperformance occurs at subset sizes of 35, 25, and 10. Overall, the Sharpe ratio obtained from the subset optimization is consistently higher than the Sharpe ratio of the standard portfolio at subset sizes of 25 and 10.

This finding is of great importance as it highlights a valuable property of the subset optimization that is beneficial to investors, regardless of their level of risk aversion. From a cost-benefit analysis perspective, despite the decrease in average return as the subset size decreases, the subset optimization consistently yields higher Sharpe ratios at subset sizes of 25 and 10 compared to the standard approach. This indicates that the reduction in volatility achieved through subset optimization outweighs the reduction in average returns observed as the subset size decreases. Therefore, the benefits derived from the volatility reduction outweigh the trade-off of lower average returns when employing smaller subset sizes.

Average Return: Consistently across various constraints, including the oil constraint, gold constraint, and double constraint, the average return of the subset-generated portfolios was found to be lower than that of the standard portfolio. Furthermore, as the subset size decreased, the average return also decreased, a pattern that was confirmed through robustness checks. This empirical result demonstrates that the utilization of subset optimization leads to a reduction in the average return of the portfolio compared to a standard portfolio generated without the algorithm, and this decrease in average return persists as the subset size decreases.

Utility: In contrast to the negative utility obtained with the standard approach across all Gammas, the utility significantly improves when utilizing the subset optimization. Moreover, as the subset size decreases, the utility continues to enhance, ultimately reaching a positive value when the subset size is 10. This empirical finding, supported by robustness checks, highlights the efficacy of subset optimization in increasing investors' utility. Additionally, our analysis reveals that decreasing the subset size while employing the optimization further enhances the utility, providing valuable insights for portfolio management.

Beta: In summary, The utilization of the subset optimization algorithm generally leads to reduced exposure between portfolio returns and a specific commodity, compared to the standard approach. This reduction in exposure is observed when a constraint is applied, requiring the correlation between portfolio returns and the commodity return to be zero. This reduction in exposure is consistent across all gamma values and has been confirmed through robustness checks. By imposing a constraint to achieve a correlation of zero between the portfolio returns and the commodity returns, we observe a consistent pattern of decreased exposure. These empirical findings provide valuable insights into the effectiveness of subset optimization in managing exposure to commodities within a portfolio, supporting its relevance and significance in portfolio management practices.

2.3.5. The Sweet Spot

Based on the findings presented in Section 2.3.4, it can be concluded that incorporating the subset optimization algorithm generally results in decreased exposure between portfolio returns and a specific commodity, as compared to the standard approach. This reduction in exposure is particularly evident when a constraint is implemented, mandating a zero correlation between portfolio returns and the commodity return. Furthermore, our analysis revealed that the subset optimization algorithm leads to an improvement in investor utility compared to the standard approach, with a greater enhancement observed as the subset size decreases.

Additionally, the subset optimization results in higher Sharpe ratios when the subset size is between 25 and 10, comparing with the standard approach. These findings provide valuable insights into the properties of the subset optimization algorithm, highlighting the significance of a subset size within the range of 25 to 10 as an optimal "sweet spot." This empirical discovery holds considerable importance in portfolio management.

Table 2.1 to 2.18 presented a comprehensive comparison of various metrics, including volatility, utility, Sharpe ratio, and beta, between the standard portfolio and the subset portfolio for subset sizes of 25 and 10. Across all the metrics, we consistently observed improvements when employing the subset optimization algorithm compared to the standard portfolio. This indicates the effectiveness and superiority of the subset optimization approach in enhancing portfolio performance and risk management.

Table 2.1: Difference Between Standard and Subset 10 Portfolio Gamma 1 Oil Constraints

	Percentage Change
Volatility	−96%
Sharpe Ratio	152%
Utility	−100%
Beta	−98%

Table 2.2: Difference Between Standard and Subset 10 Portfolio Gamma 2 Oil Constraints

	Percentage Change
Volatility	−96%
Sharpe Ratio	206%
Utility	−100%
Beta	−99%

Table 2.3: Difference Between Standard and Subset 10 Portfolio Gamma 10 Oil Constraints

	Percentage Change
Volatility	−98%
Sharpe Ratio	613%
Utility	−100%
Beta	−98%

Table 2.4: Difference Between Standard and Subset 25 Portfolio Gamma 1 Oil Constraints

	Percentage Change
Volatility	−91%
Sharpe Ratio	24%
Utility	−100%
Beta	−100%

Table 2.5: Difference Between Standard and Subset 25 Portfolio Gamma 2 Oil Constraints

	Percentage Change
Volatility	−92%
Sharpe Ratio	66%
Utility	−100%
Beta	−99%

Table 2.6: Difference Between Standard and Subset 25 Portfolio Gamma 10 Oil Constraints

	Percentage Change
Volatility	−97%
Sharpe Ratio	399%
Utility	−100%
Beta	−99%

Table 2.7: Difference Between Standard and Subset 10 Portfolio Gamma 1 Gold Constraints

	Percentage Change
Volatility	−96%
Sharpe Ratio	130%
Utility	−100%
Beta	−96%

Table 2.8: Difference Between Standard and Subset 10 Portfolio Gamma 2 Gold Constraints

	Percentage Change
Volatility	−96%
Sharpe Ratio	193%
Utility	−100%
Beta	−95%

Table 2.9: Difference Between Standard and Subset 10 Portfolio Gamma 10 Gold Constraints

	Percentage Change
Volatility	−98%
Sharpe Ratio	606%
Utility	−100%
Beta	−99%

Table 2.10: Difference Between Standard and Subset 25 Portfolio Gamma 1 Gold Constraints

	Percentage Change
Volatility	−91%
Sharpe Ratio	30%
Utility	−100%
Beta	−85%

Table 2.11: Difference Between Standard and Subset 25 Portfolio Gamma 2 Gold Constraints

	Percentage Change
Volatility	−92%
Sharpe Ratio	73%
Utility	−100%
Beta	−87%

Table 2.12: Difference Between Standard and Subset 25 Portfolio Gamma 10 Gold Constraints

	Percentage Change
Volatility	−97%
Sharpe Ratio	422%
Utility	−100%
Beta	−87%

Table 2.13: Difference Between Standard and Subset 10 Portfolio Gamma 1 Both Constraints

	Percentage Change
Volatility	-96%
Sharpe Ratio	149%
Utility	-100%
Beta Oil	-98%
Beta Gold	-97%

Table 2.14: Difference Between Standard and Subset 10 Portfolio Gamma 2 Both Constraints

	Percentage Change
Volatility	-97%
Sharpe Ratio	227%
Utility	-100%
Beta Oil	-99%
Beta Gold	-97%

Table 2.15: Difference Between Standard and Subset 10 Portfolio Gamma 10 Both Constraints

	Percentage Change
Volatility	-98%
Sharpe Ratio	613%
Utility	-100%
Beta Oil	-98%
Beta Gold	-100%

Table 2.16: Difference Between Standard and Subset 25 Portfolio Gamma 1 Both Constraints

	Percentage Change
Volatility	-91%
Sharpe Ratio	25%
Utility	-100%
Beta Oil	-100%
Beta Gold	-87%

Table 2.17: Difference Between Standard and Subset 25 Portfolio Gamma 2 Both Constraints

	Percentage Change
Volatility	-92%
Sharpe Ratio	75%
Utility	-100%
Beta Oil	-99%
Beta Gold	-88%

Table 2.18: Difference Between Standard and Subset 25 Portfolio Gamma 10 Both Constraints

	Percentage Change
Volatility	-97%
Sharpe Ratio	410%
Utility	-100%
Beta Oil	-99%
Beta Gold	-89%

2.3.6. Concluding Remarks

In this chapter of the dissertation, our main objective was to empirically examine the value of the subset optimization algorithm in constructing purposeful portfolios and its performance in hedging against exposure to specific commodities. Additionally, we aimed to gain insights into the properties and characteristics of this algorithm for effective portfolio management.

In response to the initial question, we examined the correlations between the returns of subset portfolios and commodity returns. Additionally, we analyzed the correlations between the returns of the standard portfolio and commodity returns. Overall, the correlations between portfolio returns and commodity returns did not exhibit a clear and consistent trend when utilizing the subset optimization. The correlations tended to be low, but they did not consistently outperform the correlations obtained from the standard portfolio. This leads us to conclude that while the subset optimization provides portfolios with low correlations to commodity returns, it does not consistently outperform the standard portfolio in terms of correlation. Therefore, our findings suggest that the subset optimization algorithm can be useful in constructing purposeful portfolios (to impose a specific hedge). However, it is important to note that the obtained correlations do not consistently outperform those of the standard portfolio.

This chapter of the dissertation has also provided valuable insights into the properties and benefits of utilizing the subset optimization algorithm in portfolio management. Through our analysis, we have learned how to effectively leverage the algorithm to enhance investor utility, improve the Sharpe ratio, and mitigate portfolio volatility. While it was observed that the average return tends to decrease when employing the subset optimization compared to the standard approach, we found that this decrease is outweighed by the significant improvements in the Sharpe ratio. The cost-benefit analysis demonstrates that the reduction in portfolio volatility achieved through the algorithm surpasses the impact on average returns. Furthermore, our analysis revealed that the subset optimization algorithm successfully reduces the exposure between portfolio returns and a specific commodity when constraints are imposed, requiring a zero correlation between the two. This finding highlights the effectiveness of the algorithm in managing and controlling risk. Overall, the results suggest that the subset optimization algorithm yields the best outcomes across various metrics when the subset size ranges between 25 and 10. These findings provide valuable guidance for investors seeking to optimize their portfolios in terms of risk management and returns. In summary, this chapter contributes essential empirical insights that can support investors in making informed decisions regarding portfolio construction and risk management.

Appendix Chapter 2

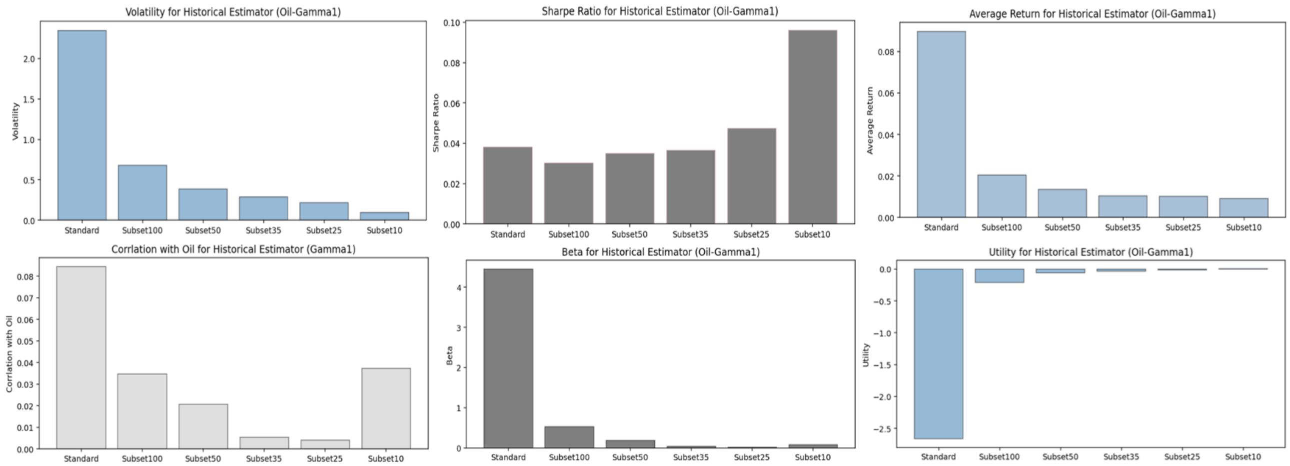

Figure 1A: Performance Metrics for historical returns inputs, Oil Constraint, Gammas 1

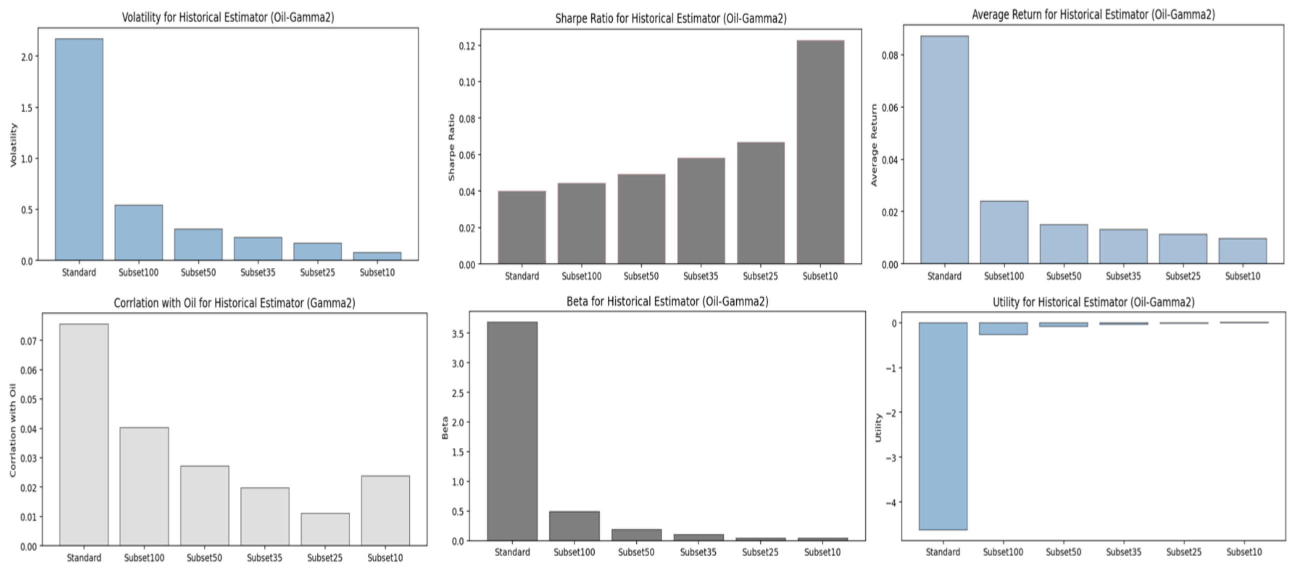

Figure 2B: Performance Metrics for historical returns inputs, Oil Constraint, Gammas 2

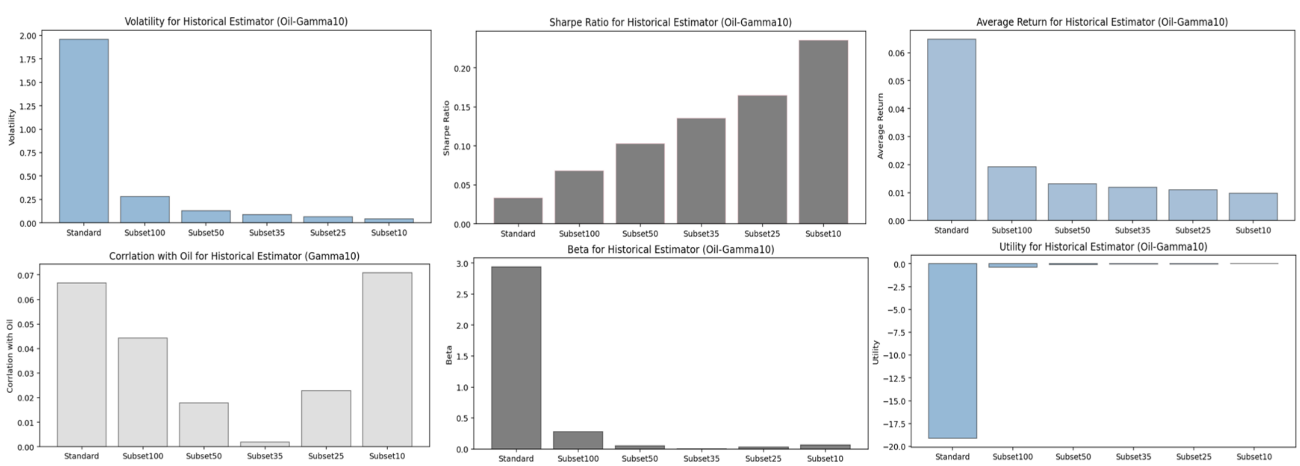

Figure 3C: Performance Metrics for historical returns inputs, Oil Constraint, Gammas 10

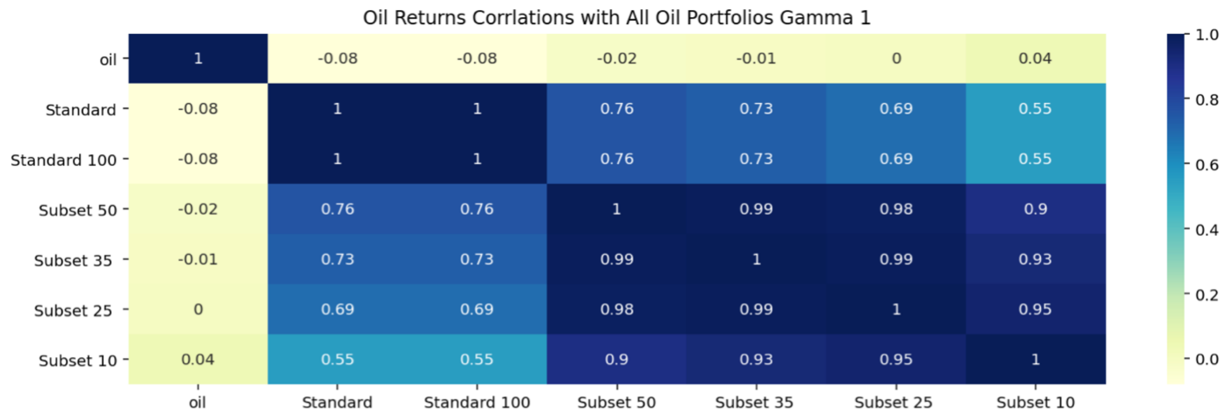

Figure 1D : Oil Returns Correlations with All Portfolios Gamma 1

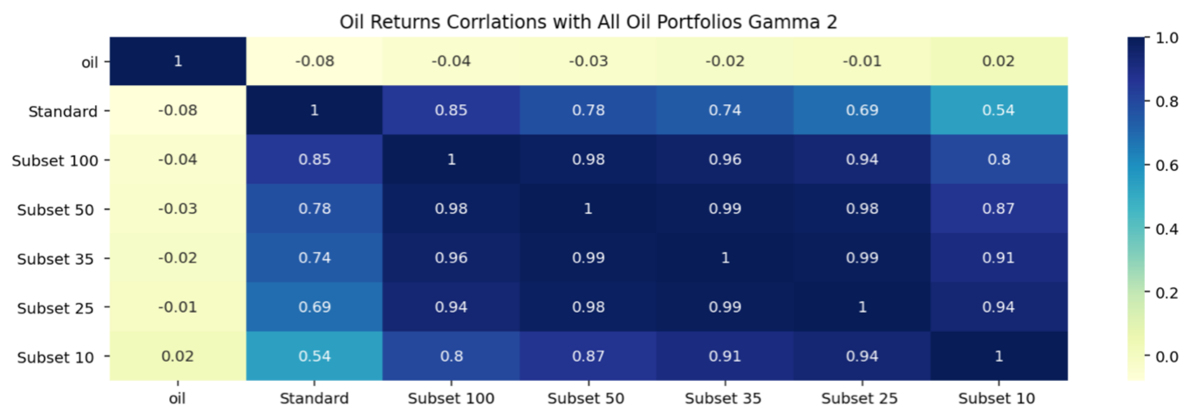

Figure 1E : Oil Returns Correlations with All Portfolios Gamma 2

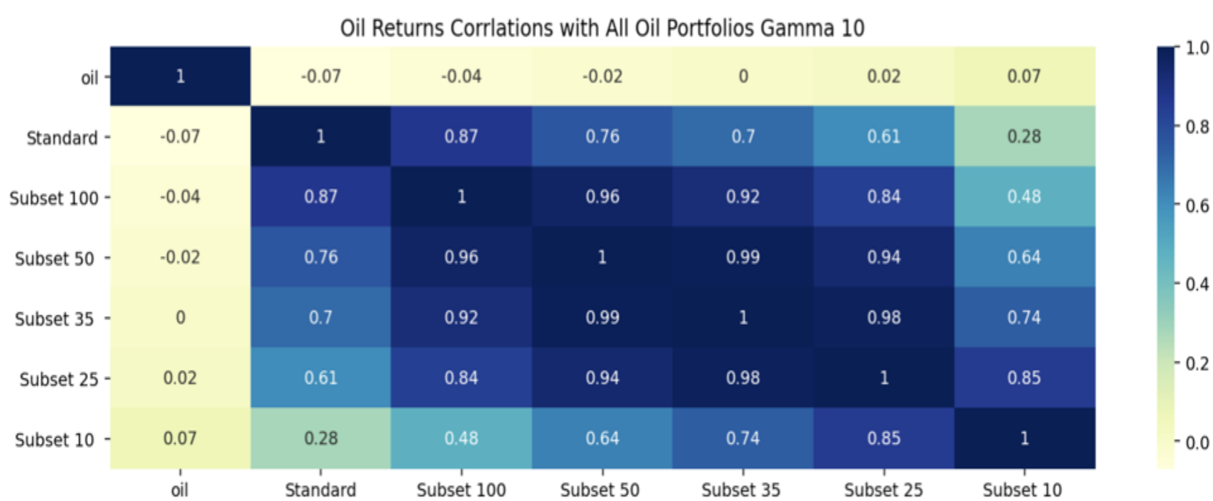

Figure 1F : Returns Correlations with All Portfolios Gamma 10

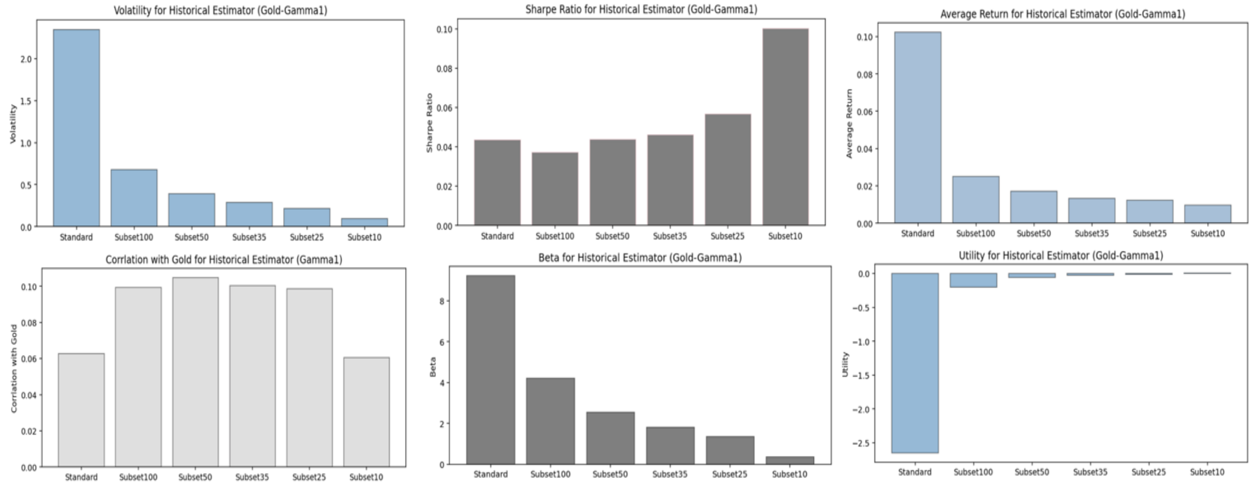

Figure 1G: Performance Metrics for historical returns inputs, Gold Constraint, Gammas 1

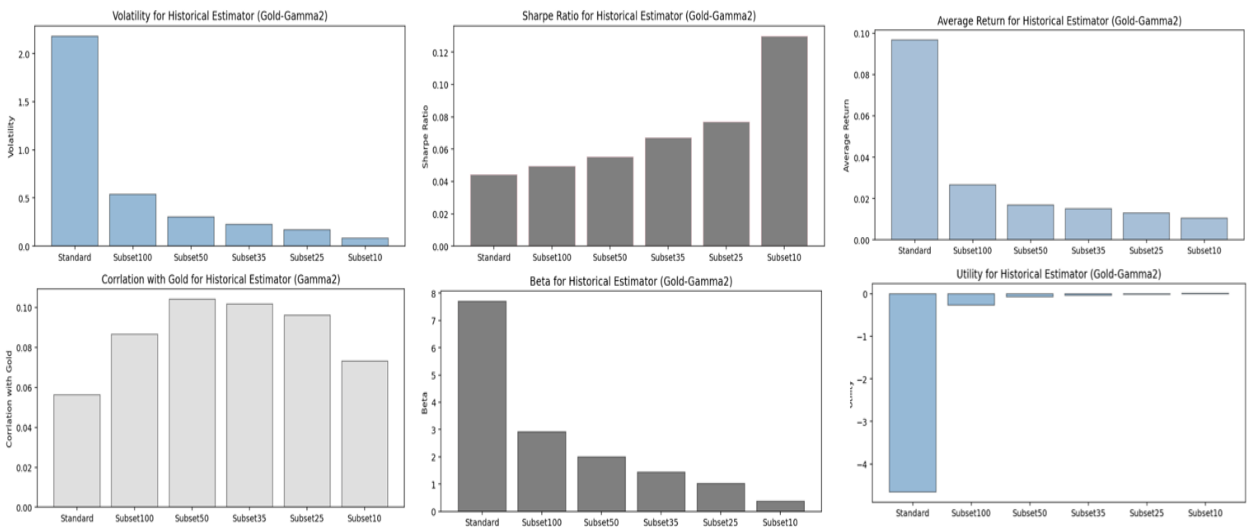

Figure 1H: Performance Metrics for historical returns inputs, Gold Constraint, Gammas 2

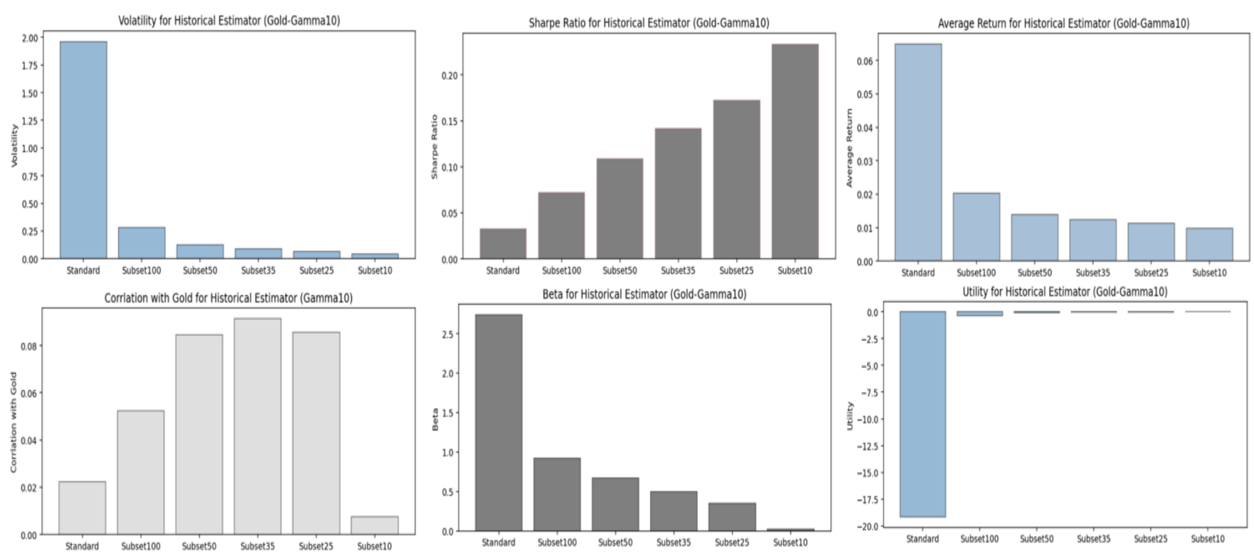

Figure 1I: Performance Metrics for historical returns inputs, Gold Constraint, Gammas 10

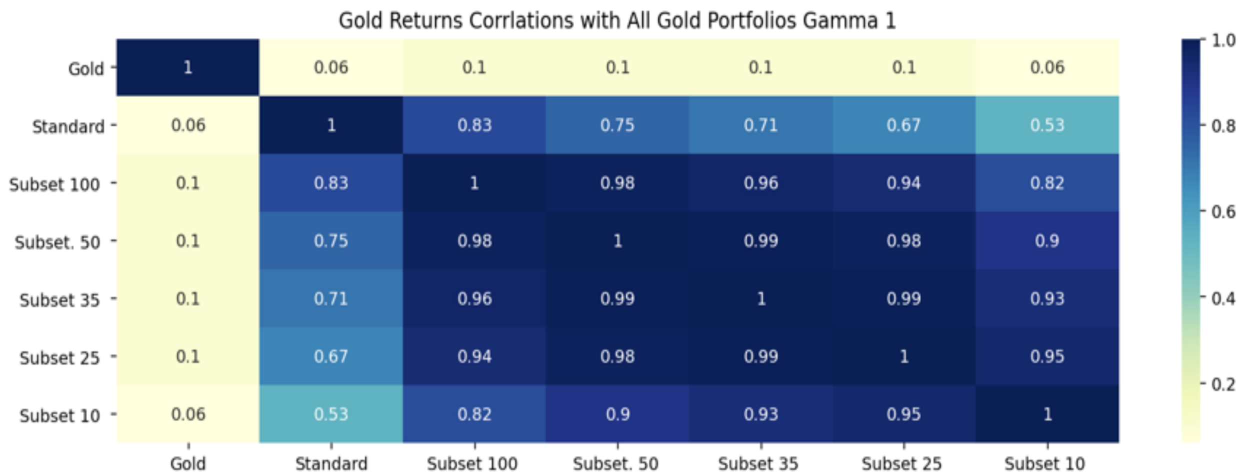

Figure 1J: Gold Returns Correlations with All Portfolios Gamma 1

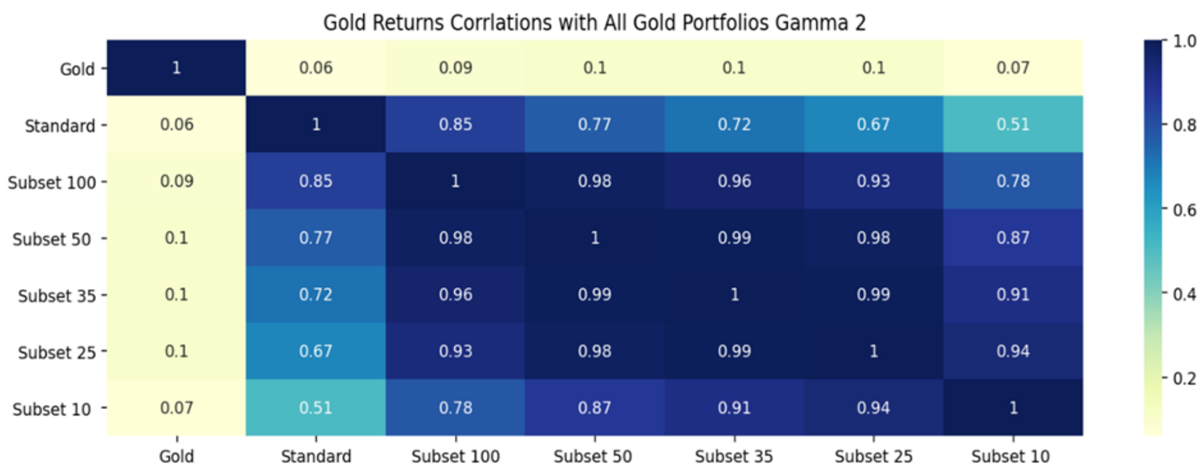

Figure 1K: Gold Returns Correlations with All Portfolios Gamma 2

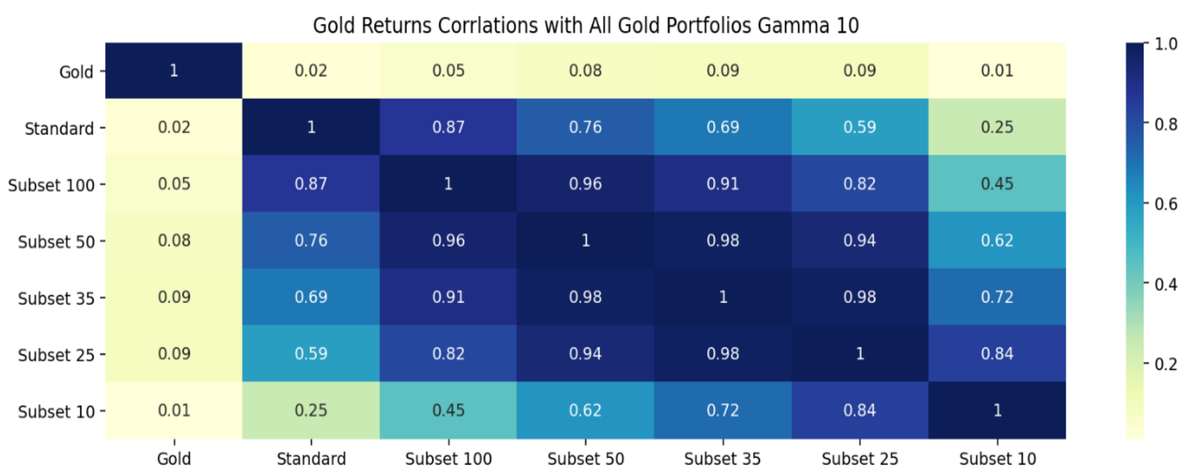

Figure 1L: Gold Returns Correlations with All Portfolios Gamma 10

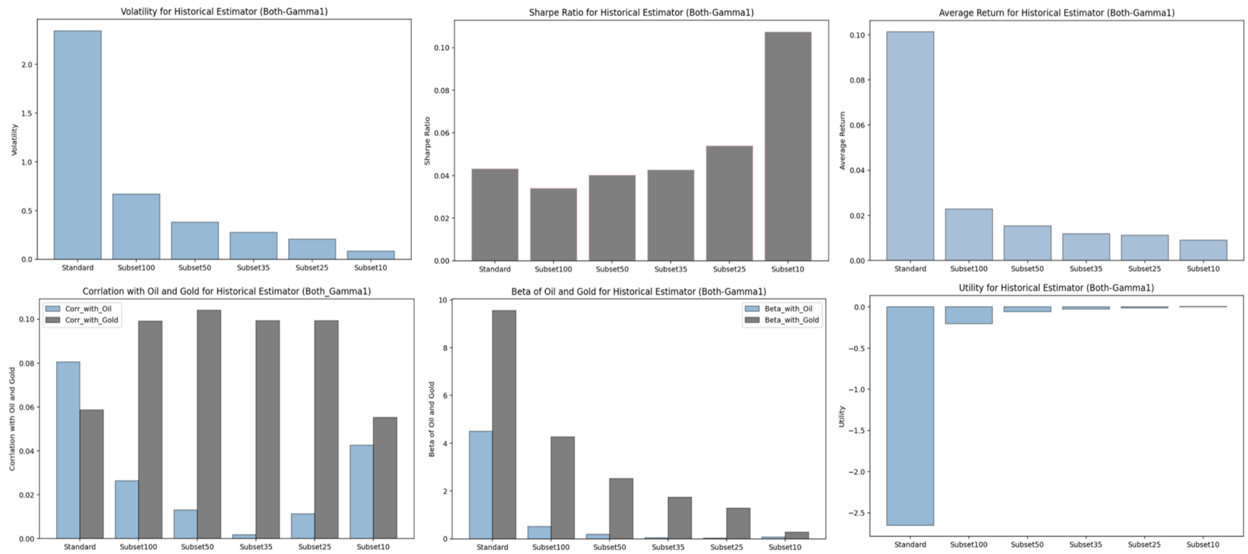

Figure 1M: Performance Metrics for historical returns inputs, both Constraint, Gammas 1

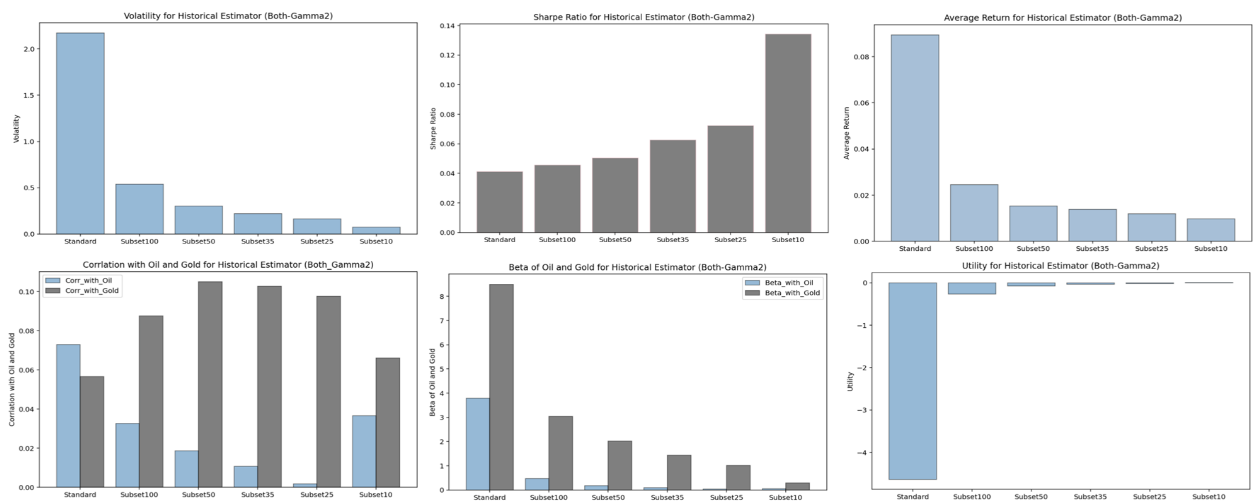

Figure 1N: Performance Metrics for historical returns inputs, both Constraint, Gammas 2

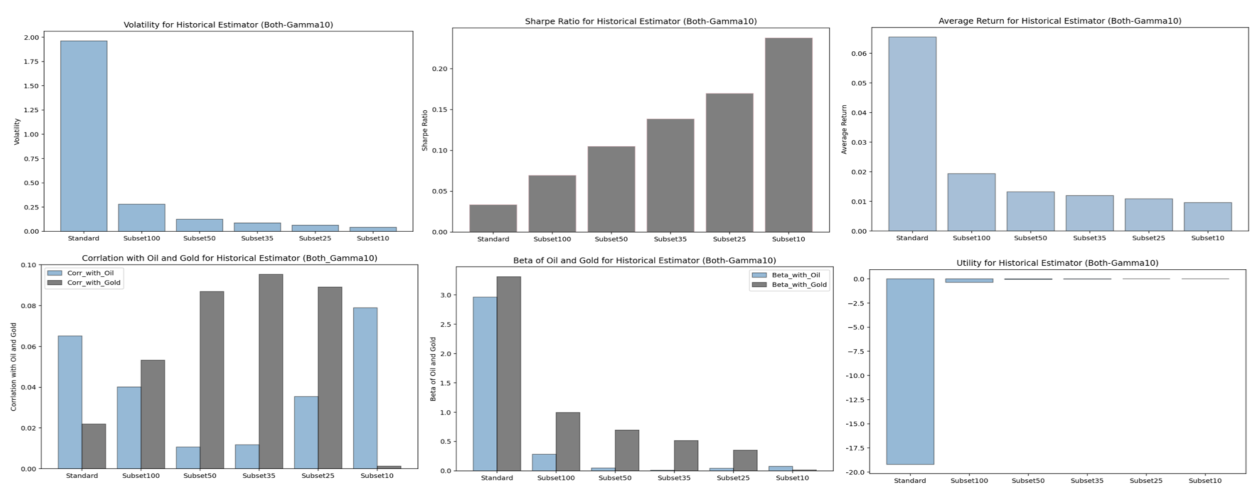

Figure 1O: Performance Metrics for historical returns inputs, both Constraint, Gammas 10

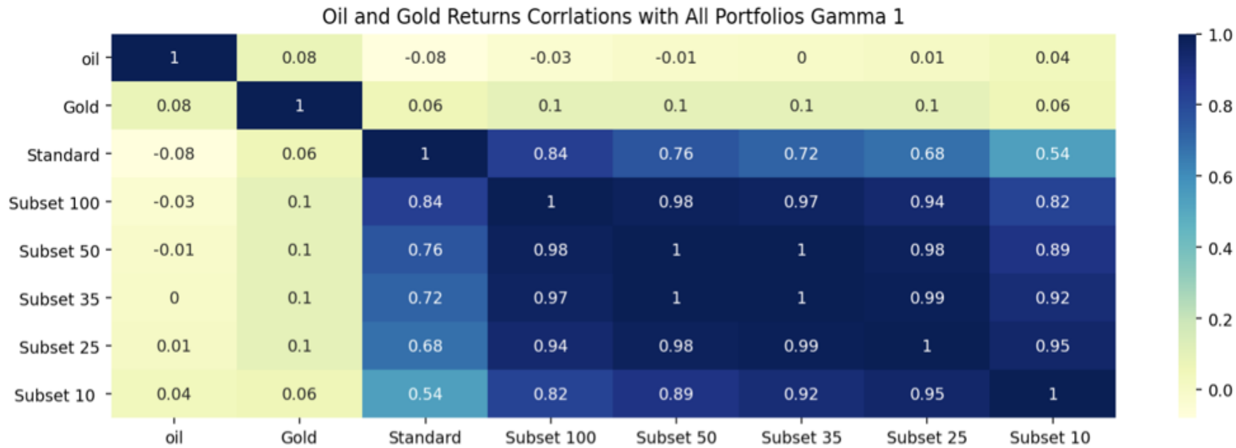

	oil	Gold	Standard	Subset 100	Subset 50	Subset 35	Subset 25	Subset 10
oil	1	0.08	-0.08	-0.03	-0.01	0	0.01	0.04
Gold	0.08	1	0.06	0.1	0.1	0.1	0.1	0.06
Standard	-0.08	0.06	1	0.84	0.76	0.72	0.68	0.54
Subset 100	-0.03	0.1	0.84	1	0.98	0.97	0.94	0.82
Subset 50	-0.01	0.1	0.76	0.98	1	1	0.98	0.89
Subset 35	0	0.1	0.72	0.97	1	1	0.99	0.92
Subset 25	0.01	0.1	0.68	0.94	0.98	0.99	1	0.95
Subset 10	0.04	0.06	0.54	0.82	0.89	0.92	0.95	1

Figure 1P: Oil and Gold Returns Correlations with All Portfolios Gamma 1

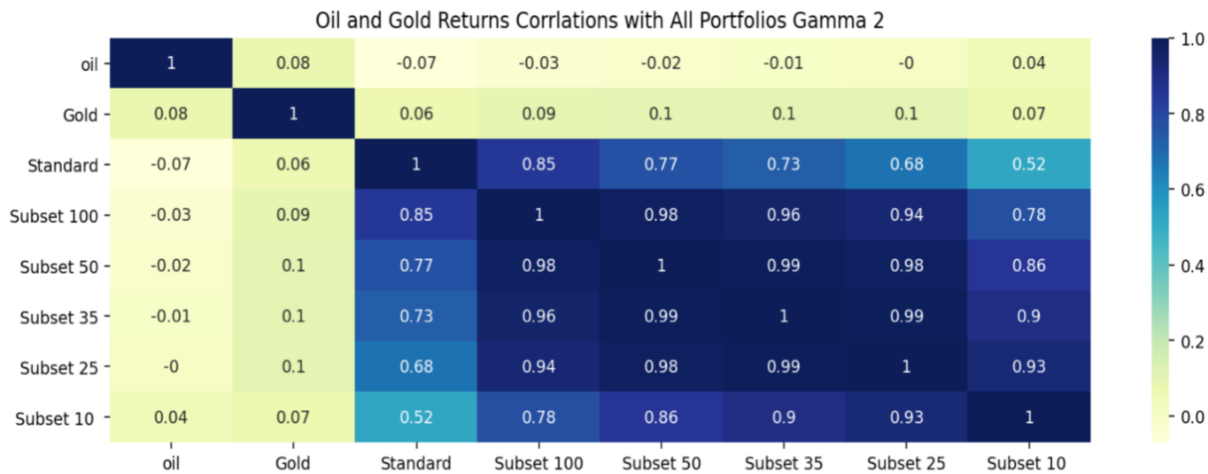

	oil	Gold	Standard	Subset 100	Subset 50	Subset 35	Subset 25	Subset 10
oil	1	0.08	-0.07	-0.03	-0.02	-0.01	-0	0.04
Gold	0.08	1	0.06	0.09	0.1	0.1	0.1	0.07
Standard	-0.07	0.06	1	0.85	0.77	0.73	0.68	0.52
Subset 100	-0.03	0.09	0.85	1	0.98	0.96	0.94	0.78
Subset 50	-0.02	0.1	0.77	0.98	1	0.99	0.98	0.86
Subset 35	-0.01	0.1	0.73	0.96	0.99	1	0.99	0.9
Subset 25	-0	0.1	0.68	0.94	0.98	0.99	1	0.93
Subset 10	0.04	0.07	0.52	0.78	0.86	0.9	0.93	1

Figure 1Q: Oil and Gold Returns Correlations with All Portfolios Gamma 2

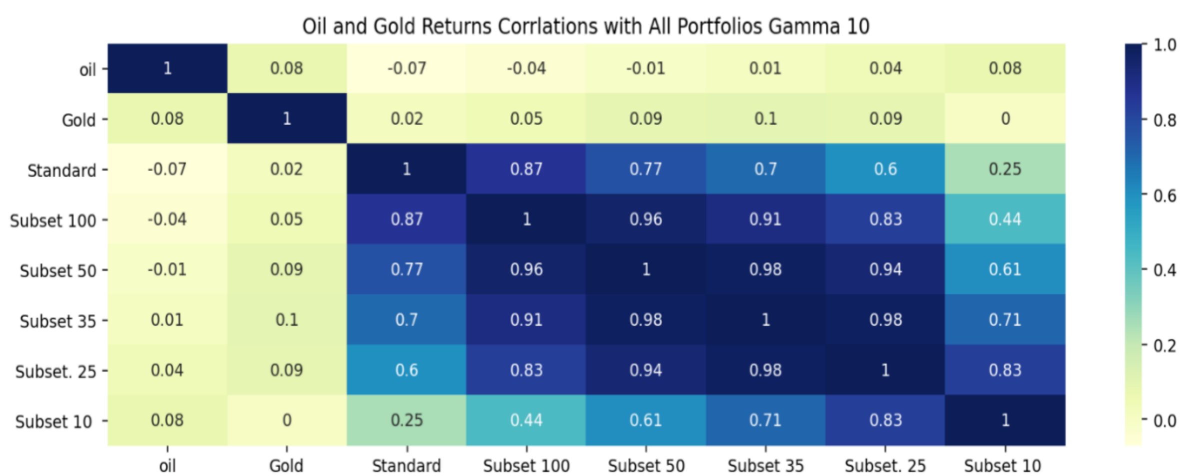

	oil	Gold	Standard	Subset 100	Subset 50	Subset 35	Subset. 25	Subset 10
oil	1	0.08	-0.07	-0.04	-0.01	0.01	0.04	0.08
Gold	0.08	1	0.02	0.05	0.09	0.1	0.09	0
Standard	-0.07	0.02	1	0.87	0.77	0.7	0.6	0.25
Subset 100	-0.04	0.05	0.87	1	0.96	0.91	0.83	0.44
Subset 50	-0.01	0.09	0.77	0.96	1	0.98	0.94	0.61
Subset 35	0.01	0.1	0.7	0.91	0.98	1	0.98	0.71
Subset. 25	0.04	0.09	0.6	0.83	0.94	0.98	1	0.83
Subset 10	0.08	0	0.25	0.44	0.61	0.71	0.83	1

Figure 1R: Oil and Gold Returns Correlations with All Portfolios Gamma 10

Table A1: Performance Statistics for Historical Inputs Portfolios Gamma 1, Oil Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0896	0.0205	0.0136	0.0104	0.0102	0.0091
Volatility	2.3468	0.6758	0.3874	0.2848	0.2151	0.0942
Sharpe Ratio	0.0382	0.0304	0.0351	0.0366	0.0475	0.0962
Utility	-2.6642	-0.2078	-0.0614	-0.0301	-0.0129	0.0046
Correlation	0.0844	0.0348	0.0207	0.0053	0.0039	0.0373
Beta	-4.4499	-0.5279	-0.1804	-0.0341	0.0191	0.0789
Alpha	0.0955	0.0212	0.0138	0.0105	0.0102	0.0090
Alpha *t*-Stat	0.7056	0.5432	0.6171	0.6359	0.8196	1.6452
Beta *t*-Stat	-1.4622	-0.6006	-0.3579	-0.0920	0.0682	0.6439

Table A2: Performance Statistics for Historical Inputs Portfolios Gamma 2, Oil Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0871	0.0239	0.0151	0.0132	0.0114	0.0097
Volatility	2.1707	0.5396	0.3067	0.2259	0.1703	0.0791
Sharpe Ratio	0.0401	0.0444	0.0492	0.0583	0.0667	0.1228
Utility	-4.6251	-0.2673	-0.0790	-0.0379	-0.0176	0.0035
Correlation	0.0755	0.0402	0.0272	0.0197	0.0111	0.0239
Beta	-3.6824	-0.4879	-0.1873	-0.1002	-0.0425	0.0424
Alpha	0.0919	0.0246	0.0153	0.0133	0.0114	0.0097
Alpha *t*-Stat	0.7340	0.7878	0.8639	1.0181	1.1594	2.1110
Beta *t*-Stat	-1.3072	-0.6953	-0.4696	-0.3407	-0.1919	0.4122

Table A3: Performance Statistics for Historical Inputs Portfolios Gamma 10, Oil Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0648	0.0192	0.0132	0.0120	0.0110	0.0098
Volatility	1.9590	0.2812	0.1281	0.0886	0.0663	0.0414
Sharpe Ratio	0.0331	0.0682	0.1030	0.1353	0.1651	0.2358
Utility	-19.1228	-0.3763	-0.0689	-0.0272	-0.0110	0.0012
Correlation	0.0668	0.0442	0.0178	0.0019	0.0229	0.0709
Beta	-2.9391	-0.2796	-0.0513	0.0037	0.0341	0.0659
Alpha	0.0686	0.0196	0.0133	0.0120	0.0109	0.0097
Alpha *t*-Stat	0.6071	1.2032	1.7901	2.3379	2.8430	4.0489
Beta *t*-Stat	-1.1554	-0.7646	-0.3077	0.0320	0.3954	1.2269

Table A4: Performance Statistics for Historical Inputs Portfolios Gamma 1, Gold Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.1024	0.0251	0.0170	0.0132	0.0123	0.0097
Volatility	2.3467	0.6754	0.3896	0.2875	0.2177	0.0972
Sharpe Ratio	0.0436	0.0372	0.0437	0.0460	0.0567	0.1002
Utility	-2.6512	-0.2029	-0.0588	-0.0281	-0.0114	0.0050
Correlation	0.0628	0.0994	0.1047	0.1004	0.0986	0.0605
Beta	9.2307	4.2059	2.5558	1.8090	1.3449	0.3684
Alpha	0.0814	0.0156	0.0112	0.0091	0.0093	0.0089
Alpha *t*-Stat	0.5950	0.3968	0.4965	0.5466	0.7335	1.5706
Beta *t*-Stat	1.0860	1.7245	1.8178	1.7425	1.7106	1.0460

Table A5: Performance Statistics for Historical Inputs Portfolios Gamma 2, Gold Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0968	0.0266	0.0168	0.0151	0.0131	0.0105
Volatility	2.1808	0.5382	0.3047	0.2250	0.1702	0.0811
Sharpe Ratio	0.0444	0.0495	0.0551	0.0671	0.0770	0.1299
Utility	-4.6593	-0.2630	-0.0760	-0.0355	-0.0159	0.0040
Correlation	0.0563	0.0865	0.1041	0.1017	0.0962	0.0731
Beta	7.6954	2.9175	1.9878	1.4334	1.0255	0.3715
Alpha	0.0793	0.0200	0.0123	0.0118	0.0108	0.0097
Alpha *t*-Stat	0.6235	0.6393	0.6944	0.9061	1.0885	2.0509
Beta *t*-Stat	0.9738	1.4993	1.8074	1.7649	1.6675	1.2652

Table A6: Performance Statistics for Historical Inputs Portfolios Gamma 10, Gold Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0649	0.0203	0.0138	0.0124	0.0113	0.0098
Volatility	1.9615	0.2804	0.1268	0.0873	0.0652	0.0420
Sharpe Ratio	0.0331	0.0725	0.1091	0.1420	0.1727	0.2337
Utility	-19.1722	-0.3727	-0.0666	-0.0257	-0.0100	0.0010
Correlation	0.0223	0.0524	0.0846	0.0914	0.0855	0.0077
Beta	2.7370	0.9209	0.6722	0.4996	0.3493	0.0201
Alpha	0.0587	0.0182	0.0123	0.0113	0.0105	0.0098
Alpha *t*-Stat	0.5123	1.1158	1.6682	2.2174	2.7585	3.9827
Beta *t*-Stat	0.3846	0.9063	1.4659	1.5842	1.4821	0.1321

Table A7: Performance Statistics for Historical Inputs Portfolios Gamma 1, Both Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.1013	0.0228	0.0153	0.0118	0.0111	0.0091
Volatility	2.3459	0.6703	0.3811	0.2771	0.2068	0.0846
Sharpe Ratio	0.0432	0.0339	0.0402	0.0426	0.0538	0.1074
Utility	-2.6503	-0.2019	-0.0573	-0.0266	-0.0102	0.0055
Correlation Oil	-0.0805	-0.0263	-0.0130	0.0017	0.0113	0.0426
Correlation Gold	0.0587	0.0991	0.1040	0.0993	0.0992	0.0553
Beta Oil	-4.4981	-0.5117	-0.1789	-0.0358	0.0182	0.0735
Beta Gold	9.5641	4.2694	2.5214	1.7318	1.2816	0.2775
Alpha	0.0855	0.0137	0.0098	0.0079	0.0082	0.0084
Alpha t-Stat	0.6260	0.3520	0.4433	0.4904	0.6809	1.6925
Beta Oil t-Stat	-1.4746	-0.5871	-0.3612	-0.0993	0.0676	0.6659
Beta Gold t-Stat	1.1243	1.7568	1.8252	1.7230	1.7085	0.9018

Table A8: Performance Statistics for Historical Inputs Portfolios Gamma 2, Both Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0894	0.0245	0.0152	0.0138	0.0119	0.0097
Volatility	2.1737	0.5364	0.3019	0.2204	0.1645	0.0725
Sharpe Ratio	0.0411	0.0457	0.0505	0.0627	0.0722	0.1343
Utility	-4.6356	-0.2632	-0.0759	-0.0348	-0.0152	0.0045
Correlation Oil	-0.0729	-0.0327	-0.0187	-0.0107	-0.0017	0.0366
Correlation Gold	0.0566	0.0876	0.1050	0.1028	0.0975	0.0661
Beta Oil	-3.7871	-0.4753	-0.1812	-0.0918	-0.0335	0.0518
Beta Gold	8.5009	3.0420	2.0231	1.4386	1.0117	0.2893
Alpha	0.0751	0.0182	0.0109	0.0107	0.0096	0.0090
Alpha t-Stat	0.5930	0.5833	0.6197	0.8317	1.0043	2.1303
Beta Oil t-Stat	-1.3389	-0.6810	-0.4619	-0.3203	-0.1563	0.5480
Beta Gold t-Stat	1.0777	1.5628	1.8492	1.8001	1.6957	1.0981

Table A9: Performance Statistics for Historical Inputs Portfolios Gamma 10, Both Constraint

	Standard	Subset Size				
		100	50	35	25	10
Arithmetic Mean	0.0655	0.0194	0.0132	0.0120	0.0109	0.0096
Volatility	1.9622	0.2797	0.1261	0.0864	0.0643	0.0402
Sharpe Ratio	0.0334	0.0693	0.1050	0.1384	0.1700	0.2380
Utility	-19.1862	-0.3718	-0.0662	-0.0254	-0.0097	0.0015
Correlation Oil	-0.0652	-0.0401	-0.0107	0.0117	0.0354	0.0790
Correlation Gold	0.0219	0.0533	0.0868	0.0953	0.0890	0.0013
Beta Oil	-2.9632	-0.2789	-0.0489	0.0089	0.0417	0.0716
Beta Gold	3.3154	0.9925	0.6962	0.5142	0.3497	-0.0117
Alpha	0.0618	0.0175	0.0117	0.0108	0.0101	0.0095
Alpha t-Stat	0.5396	1.0715	1.5958	2.1422	2.6907	4.0500
Beta Oil t-Stat	-1.1580	-0.7646	-0.2980	0.0789	0.4985	1.3671
Beta Gold t-Stat	0.4646	0.9756	1.5206	1.6399	1.4993	-0.0798

Chapter 3 Advanced Estimation Strategies for Portfolio Construction

3.1. Introduction

In Chapter 2, we empirically examined the value of the subset optimization algorithm in constructing purposeful portfolios and its performance in hedging against exposure to specific commodities empirically examine the value of the subset optimization algorithm in constructing purposeful portfolios and its performance in hedging against exposure to specific commodities, also, we explored the properties of subset optimization and its application in portfolio management, and how to better utilize it to max the investor utility. We employed historical returns as inputs for mean-variance optimization, a commonly used technique for portfolio construction. However, the accuracy of our optimization results heavily relies on the quality of these input estimates.

The application of mean-variance optimization in portfolio management has been widely studied and implemented by practitioners. Its objective is to find an allocation that maximizes the expected return while minimizing the portfolio's risk, as measured by the variance of returns. Traditionally, historical returns have been the primary source of information used in this process. However, the expected returns and variance-covariance estimates we obtain from historical returns suffer from several limitations, including data limitations, sensitivity to outliers, and inaccuracy in the correlations and expected returns. To address these limitations and enhance the estimation process, researchers have explored various alternative inputs that can provide better estimates for mean-variance optimization.

The estimation of covariance matrices in portfolio diversification and other economic applications is often challenging due to limited sample size compared to the dimensionality of the problem. In a study by Gillen (2016), the author focuses on addressing this issue by proposing a Bayesian regression model that incorporates prior beliefs and structures the estimation of covariance matrices. Gillen’s paper introduces a novel approach that utilizes a factor model and incorporates shrinkage estimation, resulting in more stable and accurate estimates of covariance matrices. The proposed methodology outperforms existing estimators in terms of mean-square error and demonstrates its effectiveness in portfolio optimization and various scenarios. The paper provides a robust and automated framework for estimating covariance matrices, particularly in large-scale settings. Given the significance of accurate covariance matrix estimation in portfolio optimization, we will leverage Gillen’s approach as a valuable input for one of our models. By utilizing his estimate of the covariance matrix, we aim to enhance the precision and reliability of our optimization process, which could be leading to improved portfolio construction and risk management. It is important to note that Gillen’s work offers valuable insights and addresses the limitations associated with traditional sample covariance matrix estimation. By incorporating his approach into our research, we can benefit from the advancements made in the estimation of covariance matrices and contribute to the ongoing efforts in portfolio diversification and optimization.

Frost and Savarino (1986) highlighted that traditional portfolio selection methods based on historical estimates of expected returns, variances-covariances matrix, could be enhanced by incorporating Bayesian framework. By specifying an informative prior that reduces estimation error, the paper suggests an Empirical Bayes method to reduce the estimation error. The empirical evaluation of the proposed rule demonstrates its superior performance compared to the Classical Estimated investment methodology. Overall, the paper highlights the benefits of considering estimation risk and utilizing informative priors in portfolio selection within a Bayesian framework. Incorporating the approach proposed by Frost and Savarino into our research allows us to leverage the advancements made in return estimation. Specifically, we will utilize the returns estimated by Frost and Savarino as crucial inputs for our second model.

It's important to note that if well-chosen priors can accurately capture the underlying characteristics of the data and provide meaningful insights, they can significantly improve the estimation process, leading to more robust and reliable results. On the other hand, poorly selected or inaccurate priors may introduce biases and negatively impact the estimation, potentially yielding misleading conclusions. Therefore, the careful consideration and incorporation of high-quality prior information are essential to enhance the effectiveness and validity of the estimation procedure.

In this chapter of the dissertation, we aim to improve the estimation process by incorporating alternative inputs, our objective in this chapter would be trying to find the best source of estimates for our optimization problem. We employ a four-factor model, specifically the Carhart four-factor model (Carhart, 1997), along with an additional factor related to oil. The Carhart four-factor model, used to explain variations in stock returns, incorporates four key factors: market risk, representing overall market-related risk; size, which considers company size and indicates that smaller companies tend to outperform larger ones; value, relating to stocks with low price-to-book ratios generating higher returns compared to high P/B stocks; and momentum, capturing the trend of stocks continuing their recent performance in the near future, influencing their returns (Carhart, 1997). In the second case, we utilize the Bayesian Shrinkage estimator (Gillen, 2016) for estimating the variance-covariance matrix, and we incorporate an informative prior within a Bayesian framework (Frost and Savarino, 1986) to estimate the expected returns. The question we aim to answer would be: Would using the Carhart four-factor model, or using Bayesian Shrinkage estimator as inputs for our model lead to an enhancement in the performance metrics?

The remainder of this chapter is organized as follows. Section 3.2 presents our proposed methodology for incorporating these inputs and outlines the steps involved in the estimation process. Finally, Section 3.3 presents and interprets the results, comparing the performance of our enhanced mean-variance optimization approaches with the traditional approach based solely on historical returns.

By addressing this research question, we aim to contribute to the existing literature on optimization methods and provide valuable insights into the practical implications of incorporating alternative inputs within the subset optimization framework.

The chapter follows the following organization: 3.2 Methodologies: This section provides a detailed explanation of the methodologies employed in this chapter. It elaborates on the specific methodologies used and offers further insights into their application. 3.3 Results and Summary: In this section, the findings of the analysis are discussed, and a summary of the results is provided. The key outcomes are presented, and a concise overview of the research findings is offered.

3.4 Comparison Across All Inputs: This section focuses on comparing the results obtained from various inputs used for optimization. It examines the outcomes across different input scenarios and highlights any significant variations or patterns observed. 3.5 Concluding Remarks: The chapter concludes with this section, which offers concluding remarks. It provides a concise summary of the main points discussed, emphasizes the implications of the results.

3.2. Methodology:

In this chapter, we will employ a methodology similar to that used in Sections 1.5 and 2.2. Our objective is to optimize multiple portfolios using the Markowitz mean-variance approach, considering investors with varying levels of risk aversion. Specifically, we will compare the performance of different portfolio optimization techniques.

Firstly, we will optimize one portfolio using the standard approach, without utilizing subset optimization. Subsequently, we will optimize several portfolios using subset optimization. We will adjust the level of risk aversion for an. We will explore subset sizes of 100, 50, 35, 25, and 10. Also adjust the risk aversion level, we start with risk neutral investor with Gamma 1, then switch to more risk averse investor with Gamma 2, then to a very risk averse investor with Gamma 10.

However, a significant departure from Chapter 2 lies in the input estimation process. In this chapter, we will introduce two models. The first model will utilize estimates obtained from a factor model, while the second model will employ estimates derived from a Bayesian Shrinkage approach. These alternative input estimation methods aim to enhance the accuracy and reliability of the optimization process.

To evaluate the performance of each model, we will assess various performance metrics, including volatility, Sharpe ratio, investor's utility, beta (which measures exposure to a specific commodity), average return, and correlations. To impose a constraint on the optimization problem, we will set the returns of the generated portfolios and a specific commodity equal to zero. The correlations analyzed will be between the generated portfolio returns and the return of the designated commodity.

In this chapter, our specific constraint will be that the correlation between the portfolio returns and oil returns is zero. Subsequently, we will conduct a robustness check by employing gold returns as an alternative constraint. This additional analysis will allow us to evaluate the sensitivity and robustness of the optimization results to different constraints, providing a more comprehensive assessment of the models' performance. By following this approach, we aim to gain insights into the effectiveness of subset optimization with alternative input estimation techniques.

3.3. Empirical results

3.3.1. Factors Model with Oil Constraint

Similar to Section 2.3, we will conduct a performance evaluation of our standard portfolio and compare it with the portfolios generated through subset optimization. To assess their effectiveness, we will utilize key performance metrics, including volatility, Sharpe ratio, average return, beta, correlation, and utility. These metrics will provide comprehensive insights into the risk-adjusted returns, and overall performance of the portfolios. By conducting this analysis, we aim to identify the strengths and weaknesses of each approach, helping us make informed decisions regarding portfolio construction strategies.

Figure 3.1 illustrates the performance metrics employed to evaluate the portfolios when optimizing with the inclusion of oil and constraints. In this scenario, our estimates are derived from the Carhart four-factor model along with an additional oil factor oil. To ensure the robustness of our findings, we will conduct a further analysis by imposing a constraint on gold. By doing so, we aim to verify the consistency and stability of the results obtained in this section, thereby strengthening the validity of our conclusions.

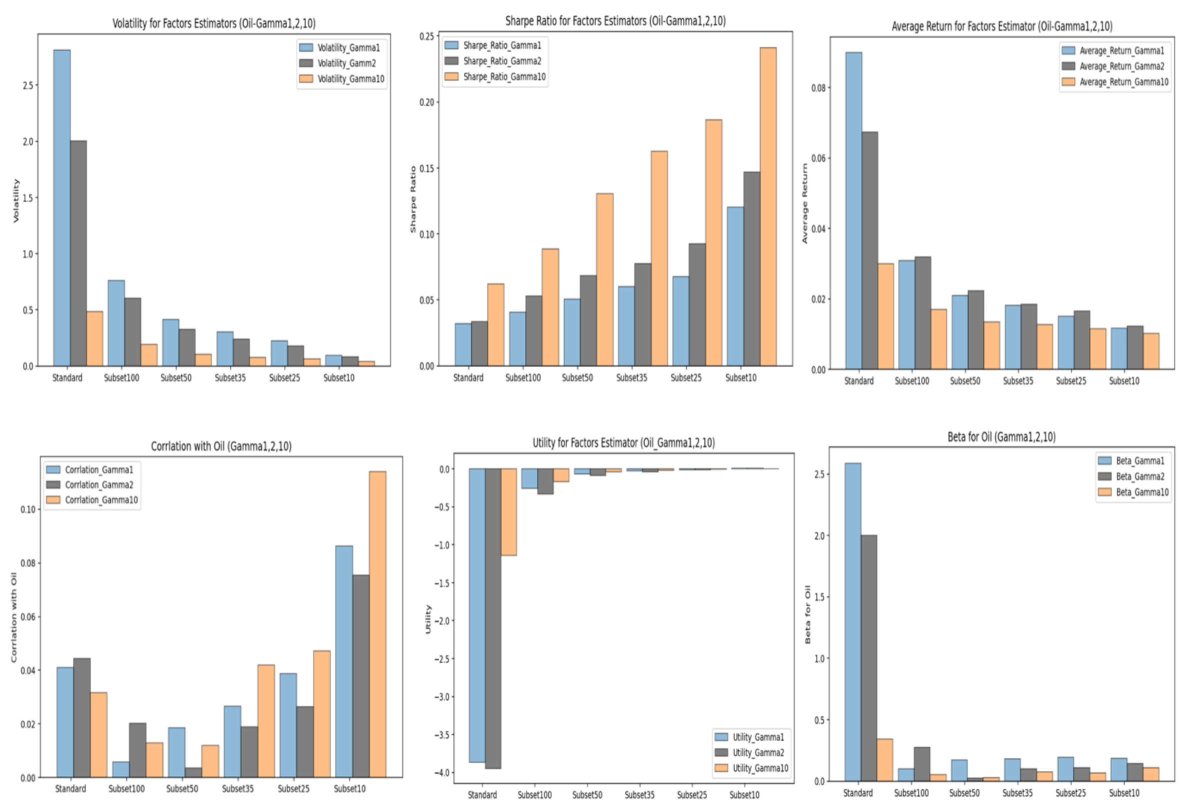

Figure 3.1: Performance Metrics for Factors inputs, Oil Constraint, all Gammas

Volatility: Upon implementing the subset optimization technique, we observe a consistent decrease in portfolios' volatilities. This outcome remains consistent across different levels of investor risk aversion. Furthermore, reducing the subset size in the optimization process leads to a further decrease in volatility. This pattern aligns with the findings from Chapter 2, where historical inputs were utilized in the optimization problem.

Average Returns: One noteworthy finding is that the implementation of subset optimization is associated with a decrease in the average return of the portfolio. This trend persists as the subset size decreases, revealing a consistent decrease in average return. This observation aligns with the pattern identified in Chapter 2 when historical estimates were employed. The continuous decline in average return with decreasing subset size underscores the significance of this finding.

Sharpe Ratio: A noteworthy finding is that the utilization of subset optimization leads to an increase in the Sharpe ratio of the portfolio. This trend persists consistently across different levels of risk aversion (Gamma), and becomes more pronounced as the subset size decreases. Interestingly, this pattern was not observed in Chapter 2 when historical inputs were employed for Gamma 1. However, a similar trend was noticed for Gamma 2 and Gamma 10. The significance of this finding warrants further investigation to determine if the trend persists during the robustness check in Section 3.3.2.

Utility: An interesting observation is that the utility derived from the standard portfolio remains consistently negative across all levels of risk aversion. However, as we implement the subset optimization algorithm, the utility shows improvement. Notably, as we decrease the number of subsets, the utility gradually increases and eventually becomes positive when the subset size reaches 10, this pattern in utility was observed in Chapter 2 of the dissertation as well.

Beta: The incorporation of subset optimization leads to a significant decrease in beta when compared to the standard portfolio. However, as the subset size decreases, we do not observe a further decrease in beta. Nonetheless, a consistent finding is the reduction in exposure to oil when employing the subset optimization approach. This implies that the subset optimization technique effectively mitigates the portfolio's sensitivity to oil returns, highlighting its potential benefits in constructing an effective hedge against a certain commodity exposure.

Correlations: When implementing subset optimization, we observe a decrease in correlations. However, a distinct pattern emerges as we vary the subset size. For Gamma 2 and Gamma 10, the correlations increase initially after reaching a minimum at subset size 35, and continue to rise as the subset size further decreases. In the case of Gamma 1, the correlations decrease initially at subset size 100, then exhibit an upward trend when the subset size is reduced to 50 and below.

Figure 3.2 illustrates that the standard portfolio demonstrates higher correlations with portfolios characterized by larger subset sizes. However, as we decrease the subset size within the optimization algorithm, the correlations consistently diminish. To further explore the correlations, additional correlation matrices are provided in the appendix.

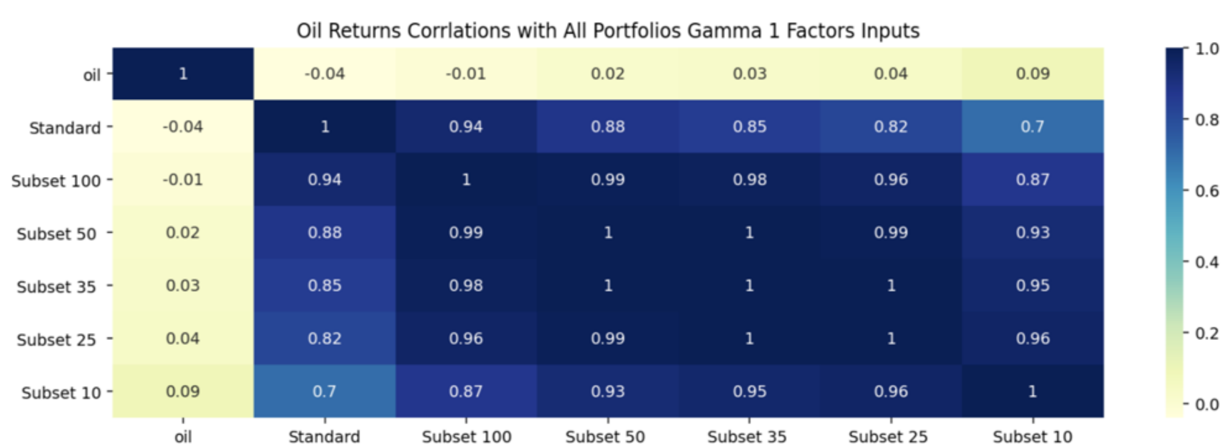

Figure 3.2: Oil Return Correlations with All Portfolios Gamma 1 Factors Inputs

3.3.2. Factors Model with Gold Constraint (Robustness check)

Figure 3.3 presents the performance metrics used to evaluate the portfolios in the Gold section, which serves as a robustness check for the results obtained in Section 3.3.1. The objective is to investigate whether the patterns observed in Section 3.3.1 also hold in the context of the Gold constraint.

Volatility: The implementation of subset optimization demonstrates a consistent reduction in portfolio volatility. This effect is observed across various levels of investor risk aversion. Moreover, as the subset size is decreased during the optimization process, the volatility further decrease. This consistent pattern is consistent with the findings reported in Chapter 2, where historical inputs were used in the optimization problem. Additionally, the findings in this section, where an oil constraint was imposed, align with this observed trend.

Average Returns: A significant discovery in this study is the consistent decrease in the average return of the portfolio when implementing subset optimization. This trend remains evident as the subset size decreases, reflecting a continuous decline in average return. This finding is consistent with the observed pattern in Chapter 2 when historical estimates were used and is further supported by the results obtained in this section when imposing the oil constraint. The persistent decrease in average return with decreasing subset size is one of the properties we discovered about the subset optimization algorithm.

Sharpe Ratio: Figure 3.3 reveals a distinct trend in the Sharpe ratio. When employing the subset optimization technique, we observe an improvement in the Sharpe ratio, which is further amplified as the subset size decreases. This consistent pattern was also observed in this chapter when the oil constraint was imposed. The persistent increase in the Sharpe ratio as the number of subsets decreases represents a significant empirical finding.

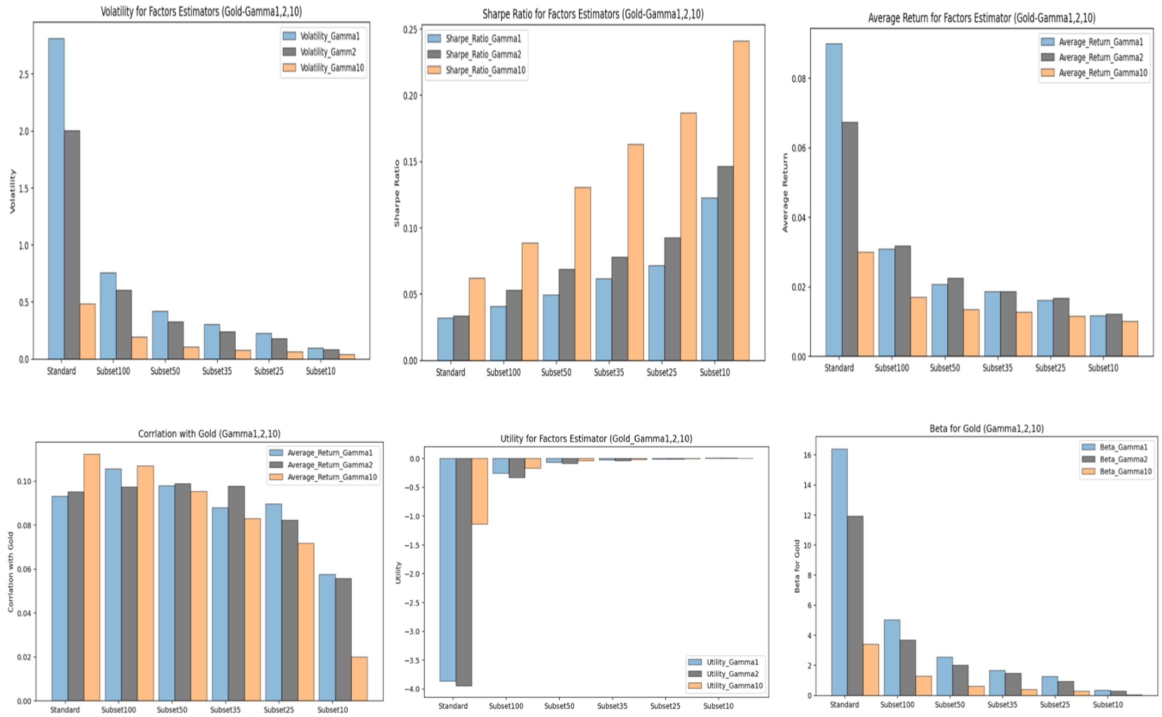

Figure 3.3: Performance Metrics for Factors inputs, Gold Constraint, all Gammas

Utility: A compelling observation arises from the analysis of utility derived from the portfolios. The utility of the standard portfolio consistently exhibits negative values across all levels of risk aversion. However, as the subset optimization algorithm is employed, there is a notable improvement in utility. Remarkably, as the number of subsets decreases, the utility demonstrates a gradual increase and eventually becomes positive when the subset size reaches 10. This observed pattern in utility aligns with the findings reported in Chapter 2 of the dissertation and is further supported by the results obtained in this section when the oil constraint was imposed.

Beta: Upon implementing the subset optimization, we consistently observe a decrease in beta compared to the standard portfolio. This finding is in line with what was observed in Section 3.3.1 when the oil constraint was imposed. Additionally, as the subset size decreases, the beta continues to decrease, a pattern that remains consistent across different levels of risk aversion (Gamma). However, it is important to note that this specific trend was not observed in Section 3.3.1 when the oil constraint was applied.

Correlations: The correlation between portfolio returns and gold returns exhibits inconsistency. Specifically, for Gamma 10, we observe a decrease in correlation as the subset optimization is implemented, and this correlation further decreases as the subset size decreases. However, for Gamma 1 and Gamma 2, we do not observe the same pattern. It is important to note that this particular pattern was not observed in Section 3.3.1 when the oil constraint was imposed.

Figure 3.4 presents a visual representation of the correlation dynamics between the standard portfolio and portfolios generated through subset optimization. The Figure reveals an interesting pattern: the standard portfolio shows higher correlations with portfolios characterized by larger subset sizes. However, as the subset size decreases within the optimization algorithm, the correlations consistently decrease as well. This finding aligns with our earlier observations regarding the imposition of the oil constraint and is in line with the correlation trends observed in Chapter 2. To provide a more detailed exploration of correlations, we have included additional correlation matrices in the appendix section.

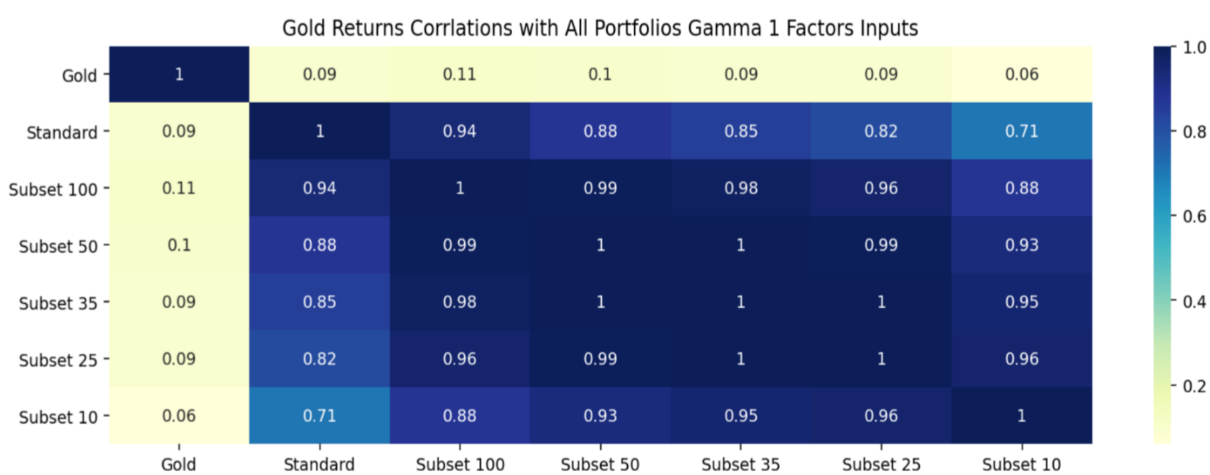

Figure 3.4: Gold Return Correlations with All Portfolios Gamma 1 Factors Inputs

3.3.3. Shrinkage Model with Oil Constraint

Figure 3.5 provides a visual representation of the performance metrics used to evaluate portfolios when incorporating oil and constraints. In this scenario, the variance-covariance matrix is estimated using the Bayesian Shrinkage estimator, while the expected returns are estimated using a Bayesian framework with informative priors. To enhance the reliability of our findings, we will perform an additional analysis by imposing a constraint on gold. This step is crucial in ensuring the consistency and stability of the results obtained in this section, thus reinforcing the validity of our conclusions.

Volatility: The application of subset optimization leads to a notable reduction in portfolio volatility compared to the standard approach. Additionally, as the subset size decreases, the volatility of the portfolio continues to decline. This consistent pattern of decreasing volatility is in line with the findings observed in Chapter 2, where historical inputs were used, as well as in the earlier section of Chapter 3, where inputs from a factor model were incorporated.

Average Returns: The utilization of subset optimization results in a decrease in average returns compared to the standard portfolio. Furthermore, as we decrease the subset size, the average returns continue to decrease. This consistent pattern of decreasing average returns aligns with the observations made in Chapter 2, where historical inputs were employed, as well as in the initial part of Chapter 3, where inputs from a factor model were utilized.

Sharpe Ratio: Figure 3.5 reveals interesting patterns in the Sharpe ratio for different levels of risk aversion (Gamma). In the cases of Gamma 1 and Gamma 2, implementing the subset optimization initially leads to a decrease in the Sharpe ratio compared to the standard portfolio. However, as the number of subsets decreases, the Sharpe ratio starts to improve. Notably, for Gamma 1, the subset optimization begins to outperform the standard portfolio when the subset size is 35 or smaller. Similarly, for Gamma 2, the subset optimization surpasses the standard portfolio when the subset size is 50 or smaller. In contrast, for Gamma 10, the Sharpe ratio exhibits consistent improvement when the subset optimization is utilized, compared to the standard portfolio. This improvement continues as the number of subsets decreases.

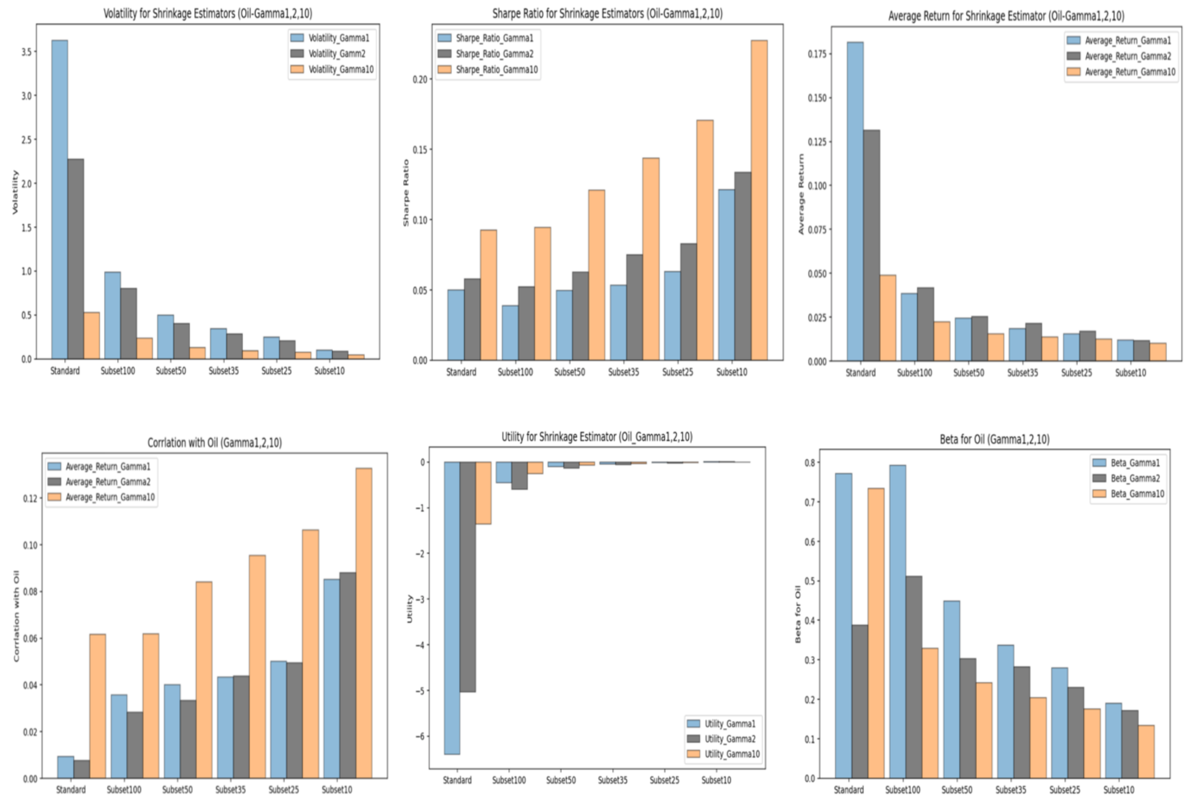

Figure 3.5: Performance Metrics for Shrinkage inputs, Oil Constraint, all Gammas

Utility: A pattern emerges when examining the utility derived from the portfolios. The standard portfolio consistently yields negative utility values across various levels of risk aversion. However, with the implementation of the subset optimization algorithm, there is a noticeable improvement in utility. Notably, as the number of subsets decreases, the utility gradually increases and eventually becomes positive when the subset size reaches 10. This observed pattern in utility is consistent with the findings reported in Chapter 2 of the dissertation and in the earlier section of Chapter 3, where inputs from a factor model were utilized.

Beta: The beta coefficient in this context represents the sensitivity of portfolios to changes in oil returns. Notably, when implementing the subset optimization technique, we observe an initial increase in beta compared to the standard portfolio for both Gamma 1 and Gamma 2. However, as we decrease the subset size to 50, the beta obtained from subset optimization becomes lower than the beta derived from the standard portfolio. This decreasing trend in beta continues as the number of subsets decreases further. In the case of Gamma 10, the beta obtained from subset optimization remains consistently lower than the beta from the standard portfolio, and this trend persists as we reduce the number of subsets.

The initial increase in beta for Gamma 1 and Gamma 2 when the subset size is 100 can be attributed to the higher correlation between oil returns and portfolio returns. Although the correlations continue to increase as we decrease the subset size, the beta decreases for subsets smaller than 100. This decrease in beta is driven by the reduction in portfolio volatility as the subset size decreases.

Correlations: The analysis of Figure 3.5 reveals a notable trend: the correlation between portfolio returns and oil returns increases as we incorporate the subset optimization technique. Additionally, this correlation further intensifies as we decrease the number of subsets within the optimization algorithm. To ensure the robustness of our findings, we will conduct a comprehensive robustness check by imposing a gold constraint on the optimization problem. Then we will observe the correlation between portfolio returns and gold returns and assess if the observed trend persists.

Figure 3.6 provides a visual depiction of the correlation dynamics between the standard portfolio and portfolios generated using the subset optimization technique. Notably, the figure highlights a pattern: the standard portfolio exhibits higher correlations with portfolios characterized by larger subset sizes. However, as we decrease the subset size within the optimization algorithm, the correlations consistently diminish. This observed pattern aligns with previous findings regarding portfolio correlations, although in this section, the correlations tend to be relatively higher compared to previous sections.

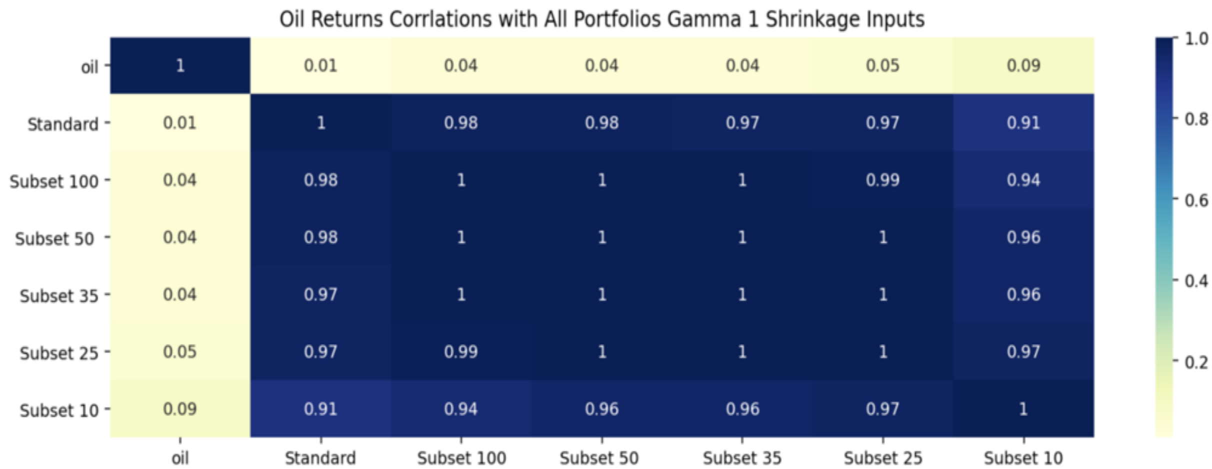

Figure 3.6: Oil Return Correlations with All Portfolios Gamma 1 Shrinkage Inputs

3.3.4. Shrinkage Model with Gold Constraint (Robustness Check)

In Figure 3.7, we examine the performance metrics used to evaluate the portfolios in the Gold section, serving as a robustness check for the findings obtained in Section 3.3.3. The purpose is to determine whether the patterns observed in Section 3.3.3 remain consistent when considering the Gold constraint. By conducting this analysis, we aim to validate the stability and reliability of our previous results, thus reinforcing the credibility of our conclusions.

Volatility: The implementation of subset optimization consistently results in a reduction in portfolio volatility compared to the standard approach. Moreover, as the subset size decreases, the volatility of the portfolio continues to decrease. This consistent pattern of decreasing volatility aligns with the findings observed in Chapter 2, where historical inputs were employed, as well as in the earlier section of Chapter 3, where inputs from a factor model were utilized, also aligns with the findings in this chapter, when the oil constraint was imposed. The reduction in volatility further demonstrates the effectiveness of subset optimization in managing portfolio risk. These findings contribute to our understanding of the properties and benefits of subset optimization, emphasizing its role in reducing portfolio volatility.

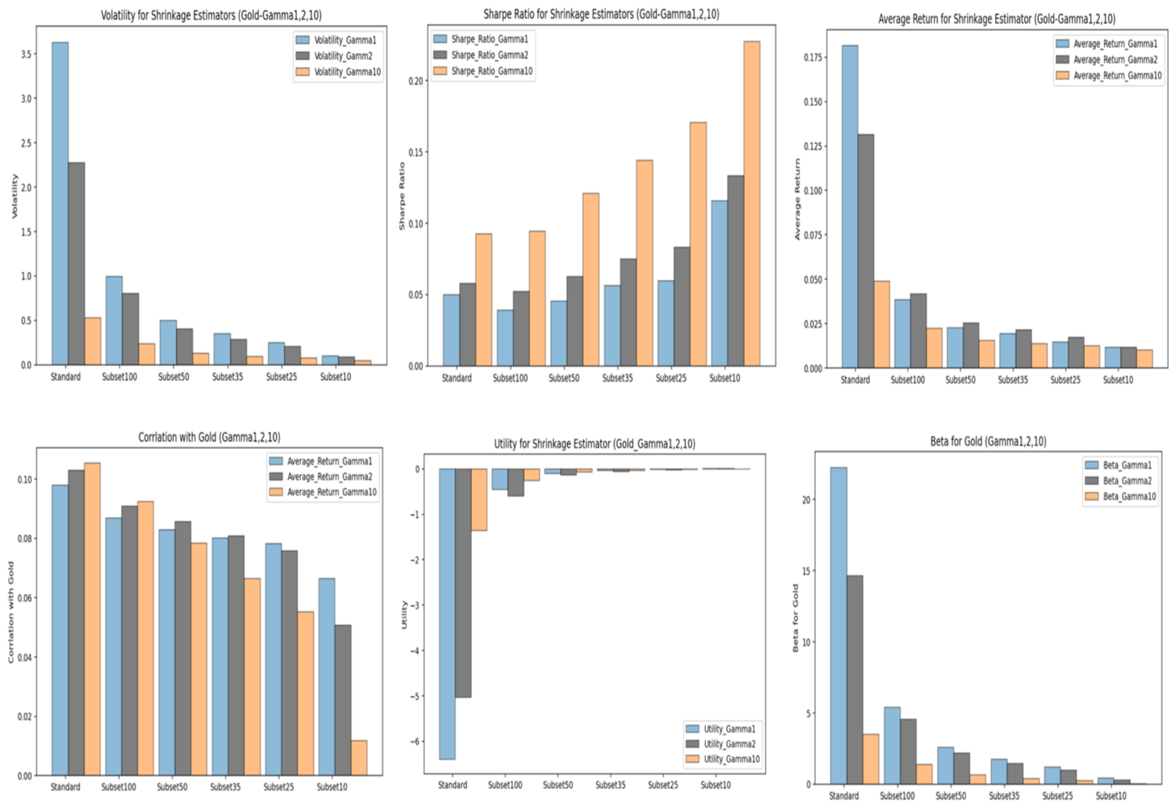

Figure 3.7: Performance Metrics for Shrinkage inputs, Gold Constraint, all Gammas

Average Returns: The application of subset optimization consistently leads to a decrease in average returns compared to the standard portfolio. Additionally, as the subset size decreases, the average returns further decline. This pattern aligns with the findings observed in Chapter 2, where historical inputs were used, and in the earlier part of Chapter 3, where inputs from a factor model were utilized. Moreover, this pattern persists in this section when the oil constraint is imposed. The consistent decrease in average returns with decreasing subset size highlights a notable property of subset optimization.

Sharpe Ratio: The findings in this section closely mirror those observed when imposing the oil constraint. For both Gamma 1 and Gamma 2, the initial implementation of the subset optimization leads to a decrease in the Sharpe ratio compared to the standard portfolio. However, as the number of subsets decreases, the Sharpe ratio shows signs of improvement. Specifically, for Gamma 1, the subset optimization begins to outperform the standard portfolio when the subset size is 35 or smaller.

Similarly, for Gamma 2, the subset optimization surpasses the standard portfolio when the subset size is 50 or smaller. On the other hand, for Gamma 10, the Sharpe ratio consistently improves when the subset optimization is employed, outperforming the standard portfolio.

Utility: A consistent pattern emerges when examining the utility derived from portfolios. The standard portfolio consistently generates negative utility values across various levels of risk aversion. However, with the implementation of the subset optimization algorithm, there is a notable improvement in utility. Particularly, as the number of subsets decreases, the utility gradually increases and eventually becomes positive when the subset size reaches 10. This observed pattern in utility is consistent with the findings reported in Chapter 2 of the dissertation and in the earlier section of Chapter 3, where inputs from a factor model were utilized. Additionally, in this chapter, when the oil constraint was imposed, the utility improvement further validates the positive impact of subset optimization on investor utility. These findings contribute to our understanding of the properties and benefits of subset optimization.

Beta: In this section, we introduced a constraint in the optimization problem to ensure that the portfolios generated have no correlation with gold returns. The beta coefficient, which represents the sensitivity of portfolios to changes in gold returns, exhibits a consistent decrease when applying the subset optimization technique, across all Gamma values. Furthermore, as the subset size is reduced, the beta decreases further. It is worth noting that this decreasing trend differs from the observations made in the previous section where an oil constraint was imposed. However, the beta obtained from subset optimization is consistently lower than the beta of the standard portfolio, particularly when the subset size is 50 or smaller in both, the gold case and oil case. This finding emphasizes the effectiveness of the subset optimization technique in reducing the portfolio's exposure to a specific commodity.

Correlations: The examination of Figure 3.7 reveals a notable decrease in the correlation between gold returns and portfolio returns when the subset optimization technique is applied. This decrease is observed in comparison to the case where the algorithm is not used. Additionally, as we decrease the subset size within the optimization algorithm, the correlation between gold returns and portfolio returns continues to decrease. Interestingly, this trend differs from the observations made earlier in this section, where an oil constraint was imposed and a positive correlation trend was identified. The inconsistent direction of the correlation trend hinders us from drawing a conclusive finding regarding the correlation pattern. However, it is important to note that the overall magnitude of the correlations remains relatively small.

Figure 3.8 visually illustrates the correlation patterns between the standard portfolio and portfolios generated through subset optimization. The figure reveals a consistent trend: larger subset sizes are associated with higher correlations with the standard portfolio, while decreasing the subset size leads to a consistent decrease in correlations. This observed pattern aligns with previous findings on portfolio correlations in Chapter 2, and in the first part of Chapter 3, and is consistent with the correlation patterns observed when applying the oil constraint in this section.

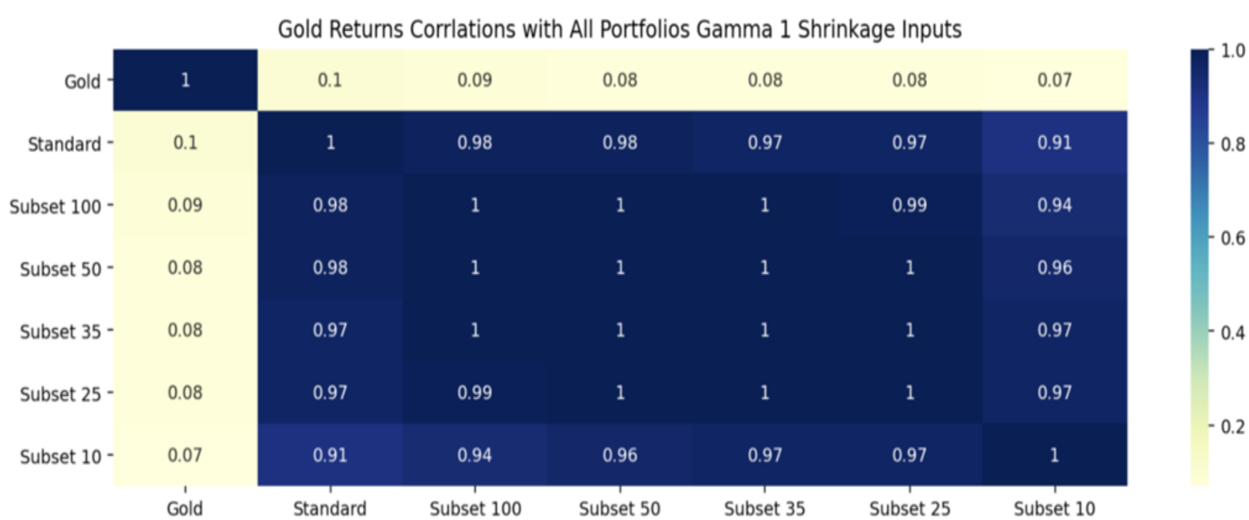

Figure 3.8: Gold Return Correlations with All Portfolios Gamma 1 Shrinkage Inputs

3.3.4.1. Findings Summary of Factors Model Inputs Approach

In this section, we will provide a summary of our findings in 3.3.1 after we conducted robustness checks. The implementation of subset optimization leads to a decrease in both average returns and volatility of the portfolios compared to the standard approach. This trend persists as the subset size decreases, which is consistent with the findings reported in Chapter 2 of this dissertation. The utilization of subset optimization results in an improvement in the investor's utility compared to the standard approach. This improvement continues as the subset size decreases, eventually reaching a positive utility level when the subset size reaches 10. A noteworthy empirical finding in this chapter is that the utilization of subset optimization, along with inputs from a factor model, consistently yields a superior Sharpe ratio compared to the standard portfolio. Moreover, as the subset size decreases, the Sharpe ratio further improves. These findings provide valuable insights into the properties and benefits of subset optimization in portfolio management. Consistently, the implementation of subset optimization results in a reduction in beta when compared to the standard portfolio. This signifies that the subset optimization technique effectively mitigates the portfolio's sensitivity to a specific factor return. This finding underscores the potential advantages of subset optimization in constructing portfolios that serve as effective hedges against exposure to particular commodities or sectors. During the robustness check conducted on our findings, we observed a different pattern in the correlation comparing with oil constraint. Also, the correlations obtained through the subset optimization technique were not consistently lower than those derived from the standard portfolio approach. However, it is important to note that overall, the correlations between portfolio returns and a specific commodity return remained low.

3.3.4.2. Findings Summary of Shrinkage Estimator Inputs Approach

In this section, we will summarize the key findings obtained 3.3.3 after we confirmed the finding with a robustness check.

The implementation of the subset optimization technique leads to a consistent decrease in average returns compared to the standard portfolio. Furthermore, as the subset size decreases, the average returns continue to decline. This pattern is mirrored in the volatility of the portfolios, which decreases with the implementation of the subset optimization and further decreases as the subset size decreases. On the other hand, the investor's utility improves when utilizing the subset optimization and shows further improvement as the subset size decreases. Notably, a positive utility is achieved when the subset size reaches 10. As for the Sharpe ratio, the utilization of the subset optimization consistently results in a higher Sharpe ratio compared to the standard portfolio with historical inputs when the subset size is 35 or smaller. The beta derived from subset optimization is lower than the beta of the standard portfolio, when the subset size is 50 or smaller. This indicates that the subset optimization technique effectively reduces the portfolio's sensitivity to a specific factor or sector. In terms of correlations, due to their inconsistent trend direction, it is challenging to draw a definitive conclusion regarding the pattern. However, it is worth noting that the correlations observed between the portfolio returns and the commodity returns, overall, exhibit relatively small magnitudes.

3.4. Comparison Across All Inputs

Throughout this dissertation, we have gained valuable insights into the properties of subset optimization. Moreover, we have examined how these properties are influenced by varying the inputs in the optimization problem. By exploring different inputs, we have deepened our understanding of the behavior and characteristics of the subset optimization technique. In this section of the dissertation, our focus is to compare the performance of portfolios under different input configurations in the optimization process.

Based on the findings presented in Chapter 2 and the current chapter, it can be concluded that the correlations observed between various portfolios and commodities are generally low, and the correlation trends lack consistency. Additionally, it has been established that the incorporation of the subset optimization technique results in lower beta values, particularly when the subset size is 50 or smaller, however the beta trends are not consistent as well. With these insights in mind, the focus of this section will be on analyzing the Sharpe ratio. By comparing the performance of different input estimates, we aim to identify which input estimation approach yields the highest Sharpe ratio and determine the most effective input estimate when combined with the subset optimization technique.

The analysis in this section will go as follow, we will select a specific gamma value, representing the risk aversion level, and focus on a particular commodity. By doing so, we can closely examine the Sharpe ratios of the portfolios generated through different inputs associated with that particular gamma and commodity. This structured approach allows us to observe which estimates work best with the subset optimization. we will examine the Sharpe ratio derived from models utilizing three different inputs. We will evaluate the performance of these models using both the standard approach and the subset optimization technique. Additionally, we will vary the subset size to observe its impact on the Sharpe ratio. By conducting this analysis, we aim to gain insights into the performance differences among the models and determine the optimal combination of inputs and subset size for maximizing the Sharpe ratio.

3.4.1. Gamma One Oil Constraint

Figure 3.9 demonstrates the trend in the Sharpe ratio for three different input models: factors model inputs, historical inputs, and shrinkage estimator inputs. When transitioning from the standard approach to utilizing the subset optimization, the factors model input model exhibits an increasing trend in the Sharpe ratio. This increase continues as the subset size decreases. In contrast, both the historical inputs and shrinkage estimator input models initially experience a decrease in the Sharpe ratio when implementing the subset algorithm. However, as the subset size decreases, the Sharpe ratio begins to improve and eventually surpasses the standard approach. Notably, the factors model input model and the shrinkage estimator input model consistently outperform the historical inputs model in terms of the Sharpe ratio when utilizing the subset optimization.

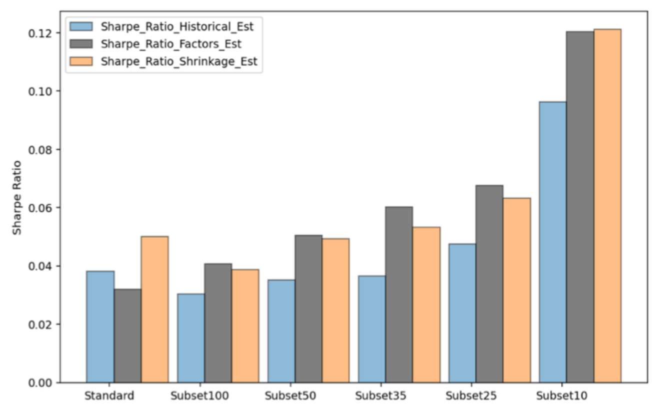

Figure 3.9: Sharpe Ratio for All Estimators (Oil Constraint -Gamma 1)

3.4.2. Gamma Two Oil Constraint

Figure 3.10 reveals that the factors model input model demonstrates a consistent increase in the Sharpe ratio when implementing the subset optimization, with further improvement as the subset size decreases. Both the shrinkage estimator input model and the factors model input model consistently outperform the historical inputs model when utilizing the subset optimization. Notably, all models outperform the portfolios generated with the standard approach when the subset size is 35 or smaller.

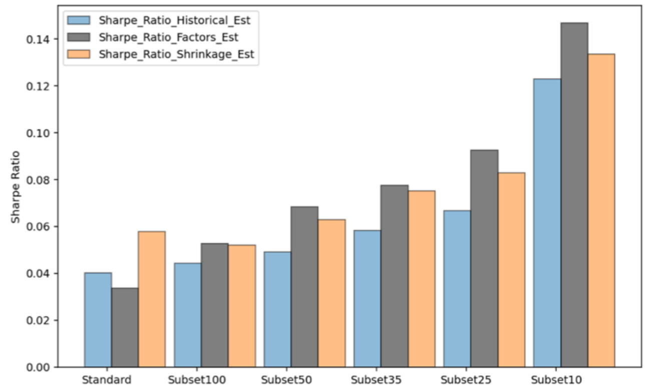

Figure 3.10: Sharpe Ratio for All Estimators (Oil Constraint -Gamma 2)

3.4.3. Gamma Ten Oil Constraint

Figure 3.11 illustrates a consistent increase in the Sharpe ratio when utilizing the subset optimization for both the factors model input model and the historical input model. As the subset size decreases, the Sharpe ratio further improves. Notably, all models outperform the portfolios generated with the standard approach when the subset size is 50 or smaller. The factors model input model consistently outperforms the historical inputs model when implementing the subset optimization.

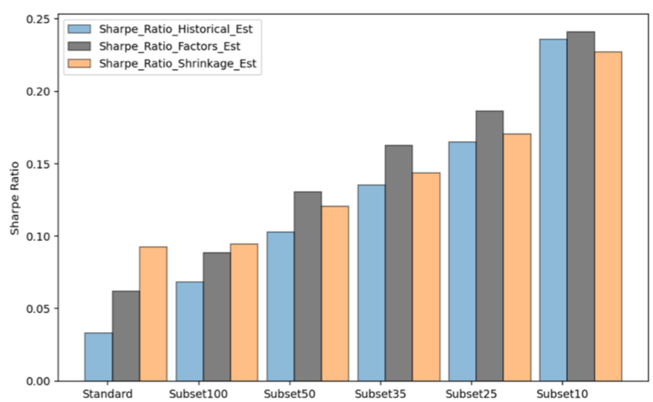

Figure 3.11: Sharpe Ratio for All Estimators (Oil Constraint -Gamma 10)

3.4.4. Gamma One Gold Constraint

In Figure 3.12, we observe a consistent upward trend in the Sharpe ratio for the factors model input model when applying the subset optimization. As the subset size decreases, the Sharpe ratio continues to improve. Importantly, all models outperform the portfolios generated using the standard approach when the subset size is 25 or smaller. Notably, both the factors model input model and the model with inputs from the Shrinkage estimator consistently outperform the historical inputs model when implementing the subset optimization.

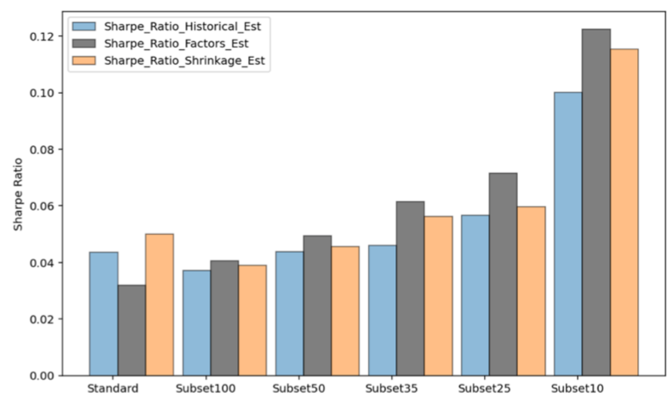

Figure 3.12: Sharpe Ratio for All Estimators (Gold Constraint -Gamma 1)

3.4.5. Gamma Two Gold Constraint

Figure 3.13 demonstrates a consistent increase in the Sharpe ratio for the factors model input model when implementing the subset optimization, the Sharpe ratio increases as the subset size decreases. All models outperform the portfolios generated with the standard approach when the subset size is 35 or smaller. Notably, the factors model input model consistently outperforms the historical inputs model when utilizing the subset optimization.

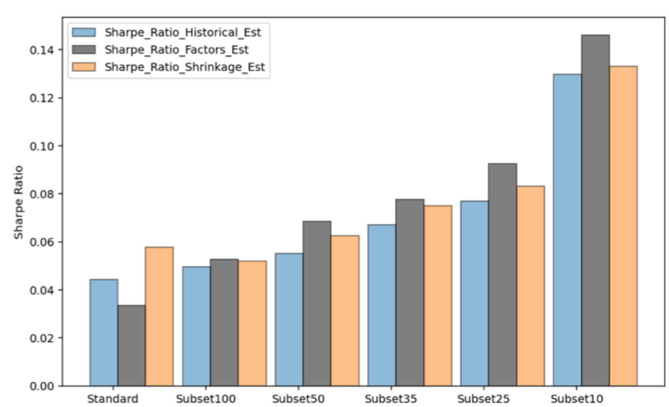

Figure 3.13: Sharpe Ratio for All Estimators (Gold Constraint -Gamma 2)

3.4.6. Gamma Ten Gold Constraint

In Figure 3.14, we can observe a consistent upward trend in the Sharpe ratio for both the factors model input model and the historical input model when utilizing the subset optimization.

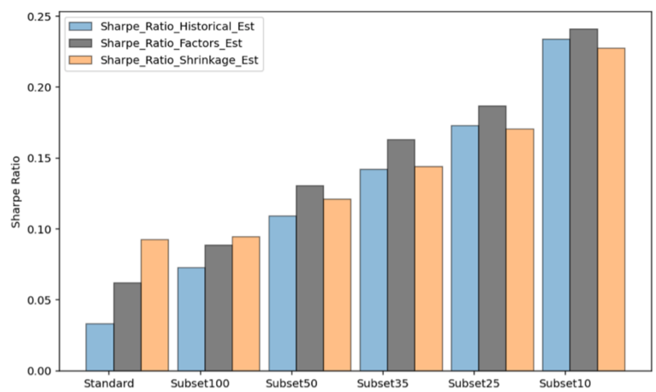

Figure 3.14: Sharpe Ratio for All Estimators (Gold Constraint -Gamma 10)

As the subset size decreases, the Sharpe ratio continues to improve. Importantly, all models outperform the portfolios generated with the standard approach when the subset size is 50 or smaller. Specifically, the factors model input model consistently outperforms the historical inputs model in the context of subset optimization.

In summary, when employing the subset optimization algorithm with inputs from the factors model, we consistently observe superior performance in terms of the Sharpe ratio compared to the standard approach. This improvement in performance is consistent across all levels of risk aversion (Gamma) and is observed when imposing both oil and gold constraints. Notably, as the subset size decreases when using factors models inputs, the Sharpe ratio continues to increase, indicating the effectiveness of the subset optimization approach. Additionally, the factors model inputs consistently outperform the historical inputs model, suggesting that incorporating factors-based estimates enhances the performance of the portfolio when used with the subset optimization. Overall, these findings highlight the usefulness of factors model inputs in conjunction with the subset optimization algorithm, as they consistently outperform the benchmark (historical inputs) in terms of Sharpe ratio across various scenarios.

3.5. Concluding Remarks

In this chapter, we delved into the properties of the subset optimization technique and its impact on portfolio performance when different inputs were employed. Our objective was to investigate the behavior of portfolios and identify the optimal inputs that work best with the subset optimization. We discovered several key insights throughout our analysis. Firstly, regardless of the inputs utilized, implementing the subset optimization led to a reduction in average returns and volatility compared to the standard approach. This trend persisted as the subset size decreased, indicating the effectiveness of the subset optimization in managing risk.

Furthermore, we observed a notable improvement in investor utility when employing the subset optimization, regardless of the input source. This improvement was further amplified as the subset size decreased, culminating in the highest utility level when the subset size reached 10. These findings underscored the advantages of the subset optimization technique in enhancing investor satisfaction. Notably, when considering the Sharpe ratio, we found that incorporating inputs from a factor model consistently outperformed the standard portfolio, demonstrating a superior risk-adjusted performance. Moreover, as the subset size decreased, the Sharpe ratio exhibited further improvement. However, it is important to mention that the same level of performance was not observed when utilizing inputs from the Shrinkage estimator. In summary, this chapter provided valuable insights into the behavior of portfolios when employing the subset optimization technique with different inputs. The findings highlighted the benefits of utilizing the subset optimization technique in managing risk, improving investor utility, and achieving superior risk-adjusted performance when inputs from a factor model were employed.

3.6. Overall Conclusion

Chapter 1 of this dissertation delved into the significance of sovereign wealth funds (SWFs) in mitigating the reliance of oil-based economies on oil revenue. The focus was on constructing purposeful portfolios for SWFs that generate returns independent of oil prices, thus promoting economic diversification. The chapter provided a concise overview of the portfolio construction process, with particular emphasis on the pivotal role of subset optimization in designing robust and diversified portfolios for SWFs.

In Chapter 2, our main objective was to empirically examine the value of the subset optimization algorithm in constructing purposeful portfolios and its performance in hedging against exposure to specific commodities. We also aimed to gain insights into the properties and benefits of this algorithm for effective portfolio management. Overall, the correlations between portfolio returns and commodity returns did not consistently outperform those of the standard portfolio when utilizing the subset optimization. However, we found that the algorithm is useful in constructing purposeful portfolios by imposing specific hedges. Additionally, the subset optimization algorithm proved effective in enhancing investor utility, improving the Sharpe ratio, and mitigating portfolio volatility. It successfully reduced exposure between portfolio returns and specific commodities when constraints were imposed. These findings provide valuable guidance for investors in optimizing their portfolios in terms of risk management and returns.

In Chapter 3, we examined the impact of the subset optimization technique on portfolio performance using different inputs. Regardless of the inputs used, the subset optimization consistently reduced average returns and volatility compared to the standard approach. Investor utility improved with the subset optimization, particularly as the subset size decreased. Incorporating inputs from a factor model into the subset optimization yielded superior risk-adjusted performance, as indicated by the Sharpe ratio. However, the performance was not as strong when utilizing inputs from the Shrinkage estimator. Overall, the findings emphasized the advantages of the subset optimization technique in managing risk and enhancing investor utility.

Chapter 4 Conclusion and Future Research

4.1 Concluding Remarks

Chapter 1 of this dissertation delved into the significance of SWFs in mitigating the reliance of oil-based economies on oil revenue. The focus was on constructing purposeful portfolios for SWFs that generate returns independent of oil prices, thus promoting economic diversification. The chapter provided a concise overview of the portfolio construction process, with particular emphasis on the pivotal role of subset optimization in designing robust and diversified portfolios for SWFs.

In Chapter 2, our main objective was to empirically examine the value of the subset optimization algorithm in constructing purposeful portfolios and its performance in hedging against exposure to specific commodities. We also aimed to gain insights into the properties and benefits of this algorithm for effective portfolio management. Overall, the correlations between portfolio returns and commodity returns we found, did not consistently outperform those of the standard portfolio when utilizing the subset optimization. However, we found that the algorithm is useful in constructing purposeful portfolios by imposing specific hedges as it successfully reduced exposure (commodity beta). between portfolio returns and specific commodities when constraints were imposed. Additionally, the subset optimization algorithm proved effective in enhancing investor's utility, improving the Sharpe ratio, and mitigating portfolio volatility. Through our analysis using the subset optimization approach, we discovered that portfolio subsets ranging from 25 to 10 securities consistently outperformed larger subsets. These smaller subsets exhibited lower volatility, higher Sharpe ratios, and greater investor utility, indicating their effectiveness in achieving favorable risk-adjusted returns. Additionally, portfolios constructed with subset sizes of 25 and 10 demonstrated lower correlation with the standard portfolio, suggesting their differences. These findings provide valuable guidance for investors in optimizing their portfolios in terms of risk management and returns.

In Chapter 3, we examined the impact of the subset optimization technique on portfolio performance using different inputs. Regardless of the inputs used, the subset optimization consistently reduced average returns and volatility compared to the standard approach. Also, regardless of the source of estimates, investor utility improved with the subset optimization, particularly as the subset size decreased. However, incorporating inputs from a factor model into the subset optimization yielded superior risk-adjusted performance, as indicated by the Sharpe ratio. It's worth mentioning that the performance was not as strong when utilizing inputs from the Shrinkage estimator. Overall, the findings emphasized the advantages incorporating cutting-edge inputs to the subset optimization technique in managing risk and enhancing investor's utility.

4.2 The Big Picture

This dissertation explores the construction of purposeful portfolios by incorporating desirable constraints and optimizing risk and return. The study specifically considers the challenges associated with managing portfolios that encompass a large number of securities. To address this issue, this dissertation proposes a solution by integrating the Subset Optimization approach into the optimization process. The subset optimization approach introduced by Gillen (2016) offers a solution to address the challenge of higher dimensionality in portfolio optimization. In this dissertation, we conducted an investigation to determine whether employing the subset optimization algorithm would result in improved hedging capabilities, specifically by reducing the correlation between portfolio returns and commodity returns. Our findings revealed that utilizing the subset optimization algorithm indeed led to portfolios with lower correlation, indicating a good hedge. However, interestingly, even without using the algorithm, we observed that imposing a constraint on the optimization problem still produced satisfactory results in terms of hedging.

Nonetheless, employing the subset optimization algorithm offered notable benefits. It enhanced the portfolio's Sharpe ratio, investor's utility, and exhibited lower volatility, as well as reducing exposure to Oil beta. Furthermore, we discovered that employing a lower subset size contributed to better portfolio performance. This included lower volatility, higher Sharpe ratio, and increased utility. Additionally, employing cutting-edge estimators in the portfolio optimization process resulted in further improvements in the Sharpe ratio and utility metrics. Overall, our research highlights the advantages of employing the subset optimization algorithm in portfolio optimization. While it offers enhanced hedging capabilities, it also improves risk-adjusted performance measures and mitigates exposure to specific factors such as Oil beta. Employing a smaller subset size and utilizing cutting-edge estimators further amplify these benefits, resulting in portfolios with improved volatility, Sharpe ratio, and utility. . By leveraging the Subset Optimization algorithm, investors can tackle the challenges associated with higher dimensionality and enhance their ability to construct purposeful portfolios that fulfill their desired constraints, while simultaneously optimizing risk and return.

4.3 Future Research

The current research focuses on the top 500 companies in the US based on market capitalization. This is a narrow focus, as it does not include companies of various sizes or from different sectors. Also, Investment funds allocate capital to a variety of asset classes, including stocks, bonds, private credit, private equity, cryptocurrencies, and more. To reflect this reality, our research should incorporate a wider range of asset classes in the portfolio holdings. Additionally, the research does not consider global opportunities. There is potential for significant expansion and enhancement in future research. For example, the research could be expanded to include companies of various sizes and from different sectors. This would allow for a more comprehensive analysis of investment strategies, as it would take into account a wider range of factors that can affect investment performance. The research could also be expanded to include global opportunities, and incorporate different asset classes. This would allow for a more diversified portfolio, as it would spread risk across different countries, assets, and markets. Aside from the methodologies used to estimate expected returns and the variance-covariance matrix in this research, a hybrid approach could be used to test the improvement and accuracy of the models. For example, the Frost and Savarino Bayesian framework could be used to estimate expected returns, while factor models could be used to estimate the variance-covariance matrix. Alternatively, Gillen's Bayesian Shrinkage estimator could be used to estimate the variance-covariance matrix and the factor models could be used to estimate expected returns. In addition to optimizing and determining the optimal weights for a portfolio, it is essential to account for potential structural shifts that could alter correlations and optimal allocations. To address this, conducting scenario analysis and stress testing after the optimization process is crucial. These measures serve to mitigate the element of surprise for portfolio managers, providing them with insights into how the portfolio might perform under various adverse scenarios. By proactively considering these possibilities, investors can better prepare for changing market conditions and potential risks, enhancing the overall resilience of their portfolios. By incorporating these adjustments and exploring alternative methodologies, future research could broaden its scope and provide a more comprehensive analysis of investment strategies. This would be beneficial for both researchers and investors, as it would provide a better understanding of the factors that affect investment performance and how to develop effective investment strategies.

Appendix Chapter 3

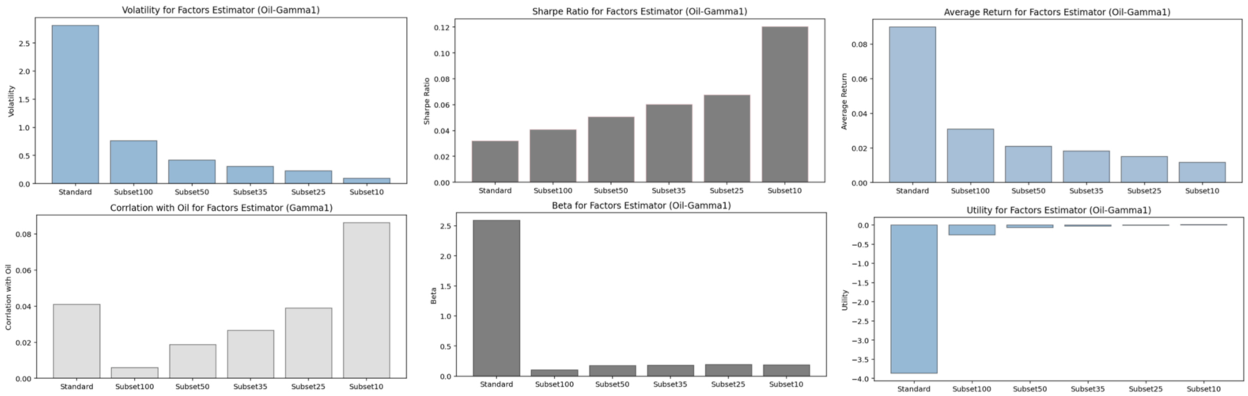

Figure 2A: Performance Metrics for Factors inputs, Oil Constraint, Gammas 1

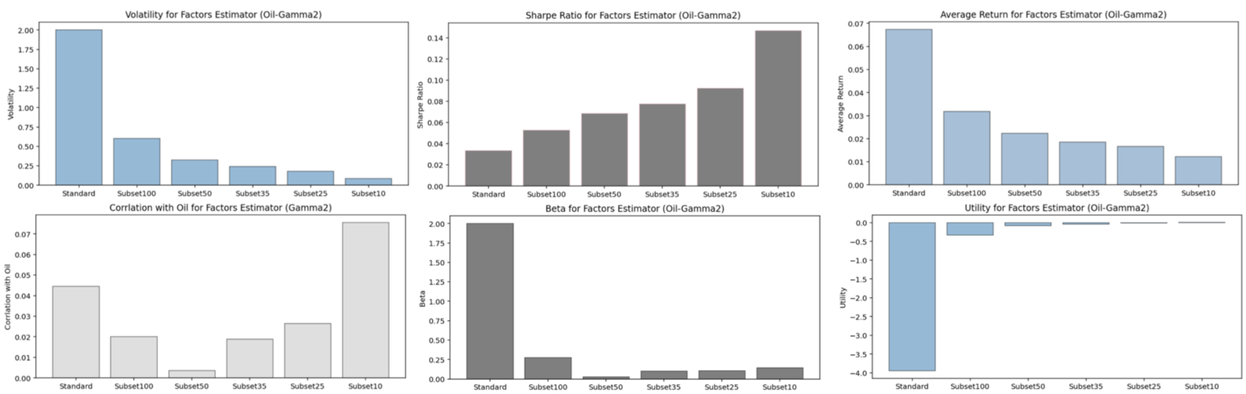

Figure 2B: Performance Metrics for Factors inputs, Oil Constraint, Gammas 2

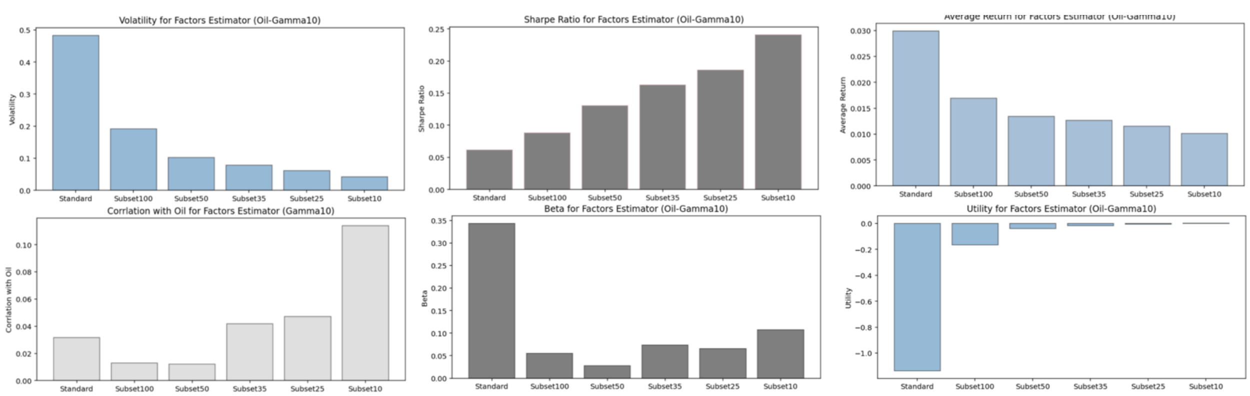

Figure 2C: Performance Metrics for Factors inputs, Oil Constraint, Gammas 10

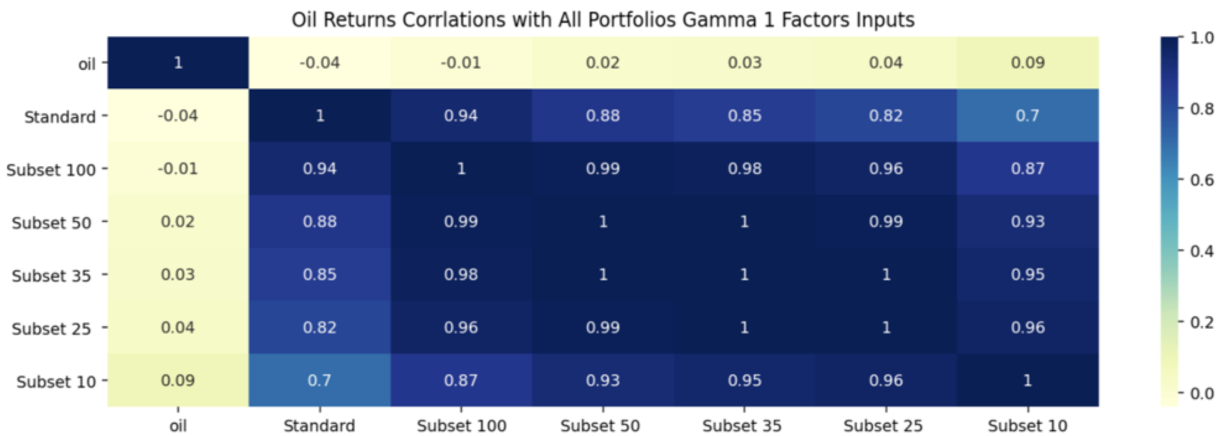

Figure 2D: Oil Return Correlations with All Portfolios Gamma 1 Factors Inputs

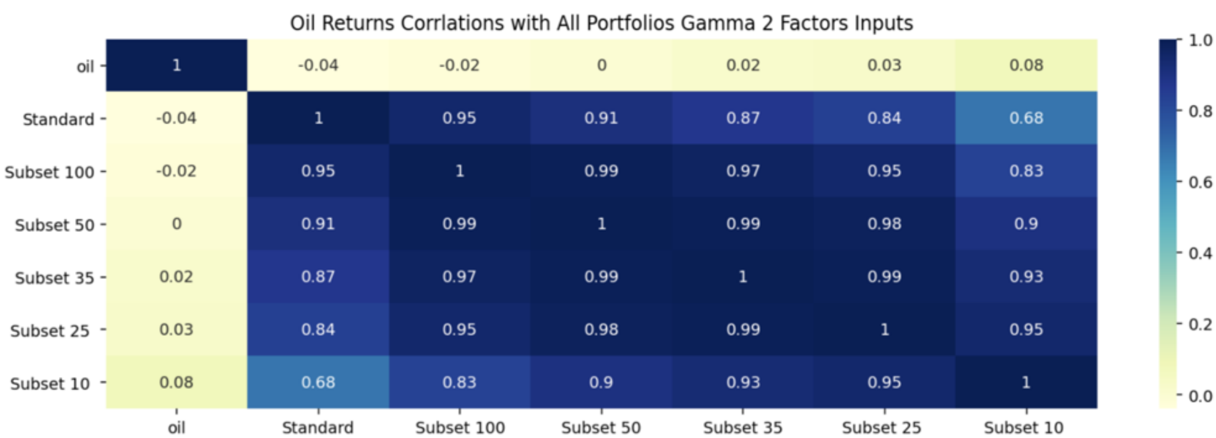

Figure 2E: Oil Return Correlations with All Portfolios Gamma 2 Factors Inputs

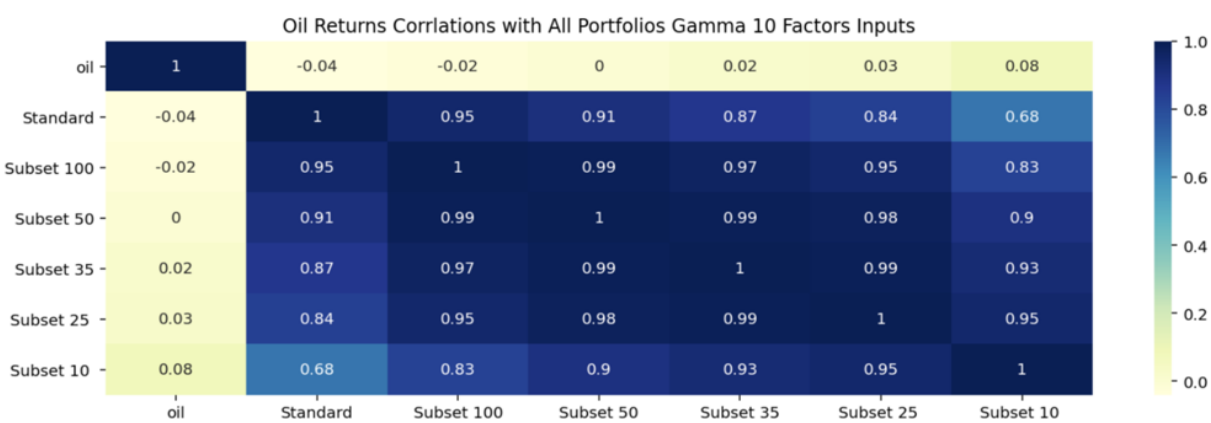

Figure 2F: Oil Return Correlations with All Portfolios Gamma 10 Factors Inputs

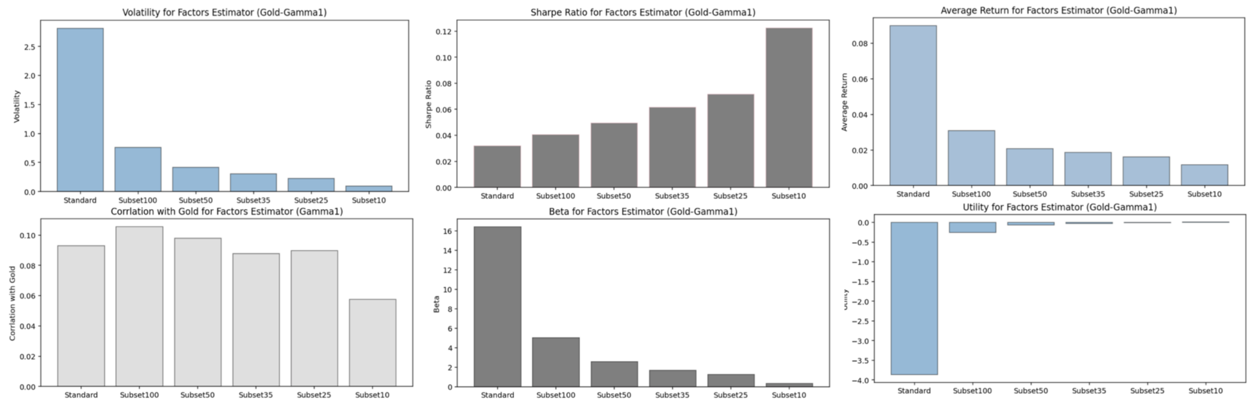

***Figure 2G:* Performance Metrics for Factors inputs, Gold Constraint, Gammas 1**

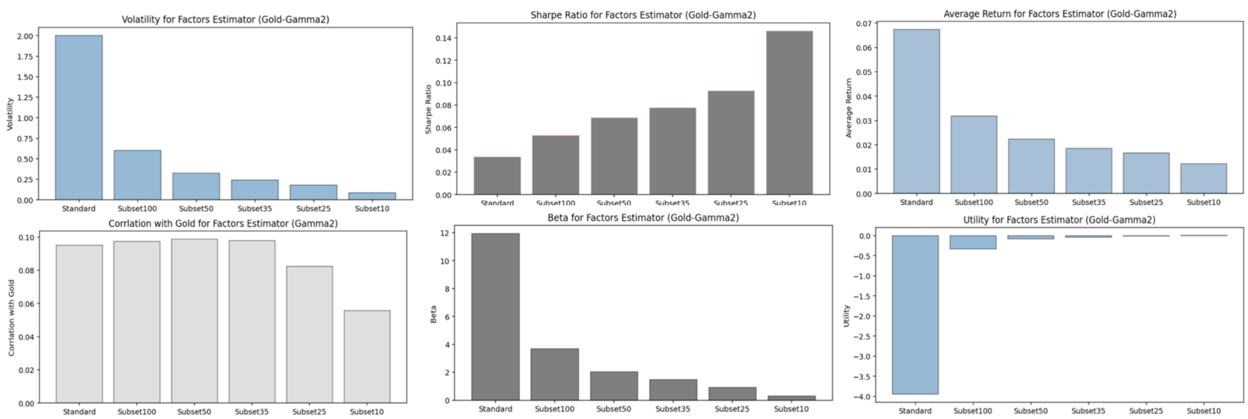

***Figure 2H:* Performance Metrics for Factors inputs, Gold Constraint, Gammas 2**

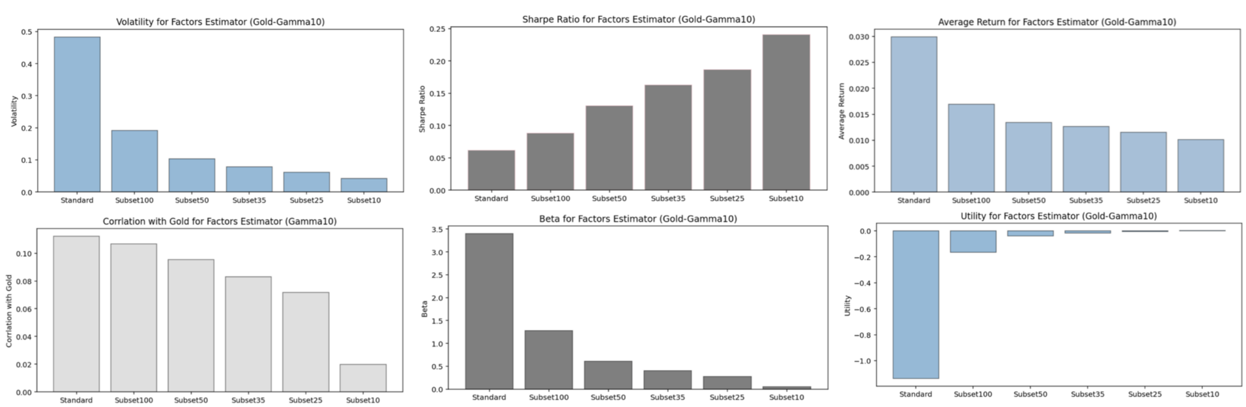

***Figure 2I:* Performance Metrics for Factors inputs, Gold Constraint, Gammas 10**

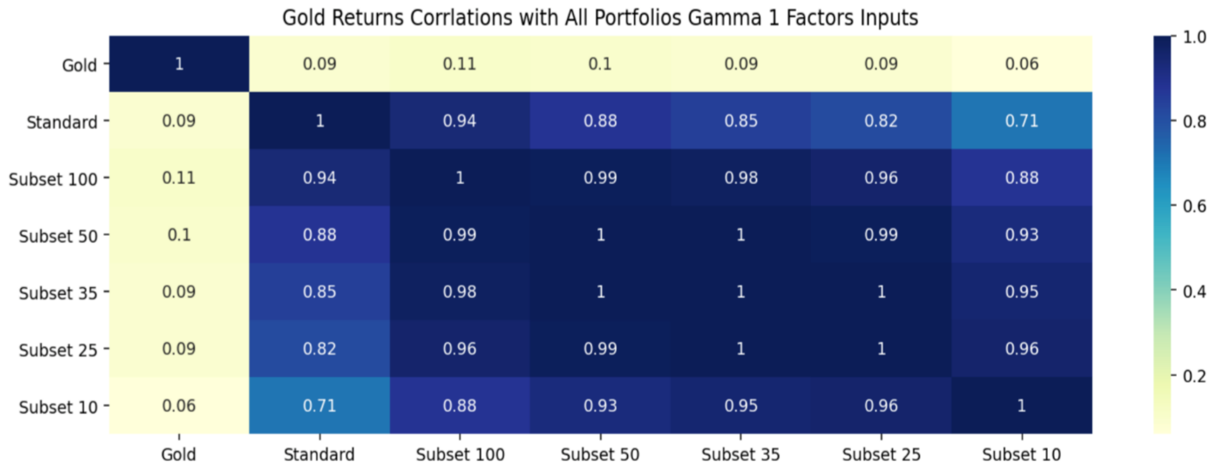

Figure 2J: ***Gold Return Correlations with All Portfolios Gamma 1 Factors Inputs***

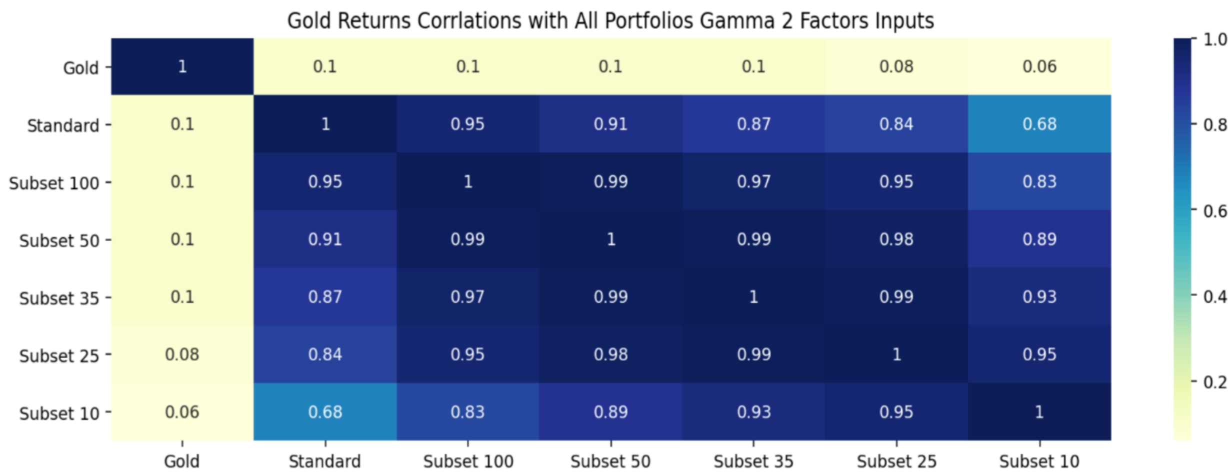

Figure 2K: ***Gold Return Correlations with All Portfolios Gamma 2 Factors Inputs***

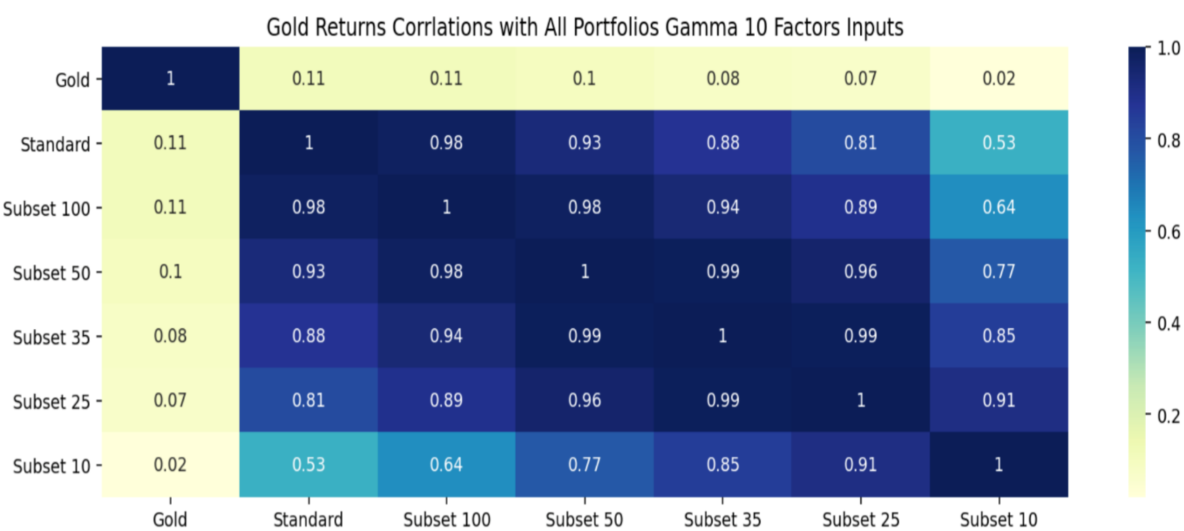

Figure 2L: ***Gold Return Correlations with All Portfolios Gamma 10 Factors Inputs***

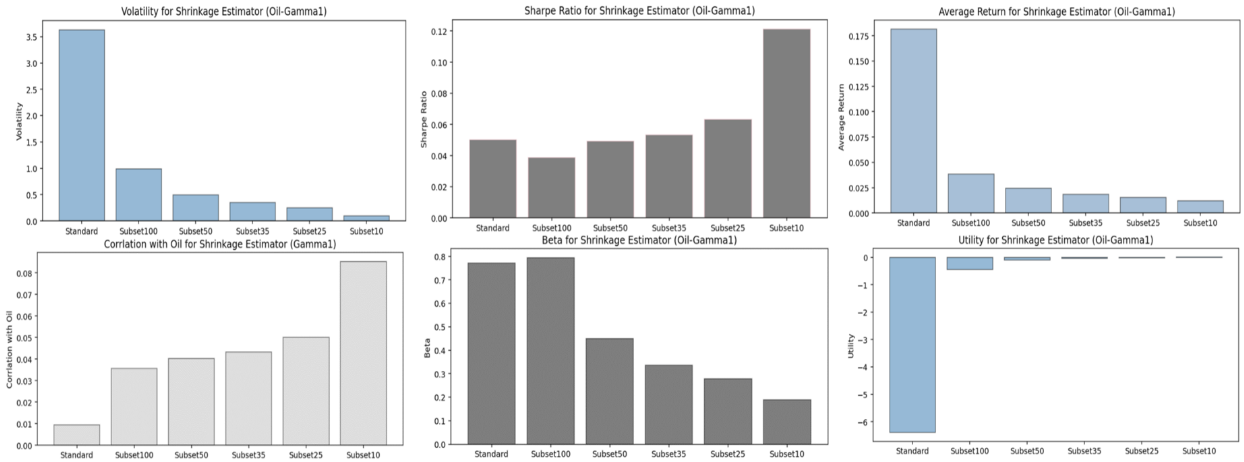

***Figure 2M:* Performance Metrics for Shrinkage inputs, Oil Constraint, Gammas 1**

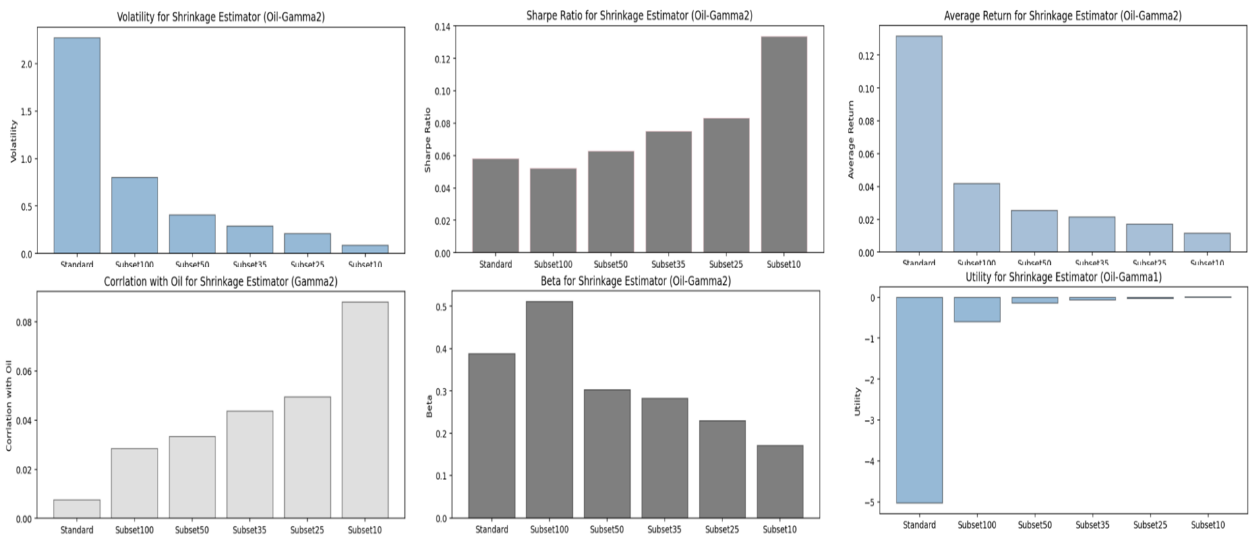

***Figure 2N:* Performance Metrics for Shrinkage inputs, Oil Constraint, Gammas 2**

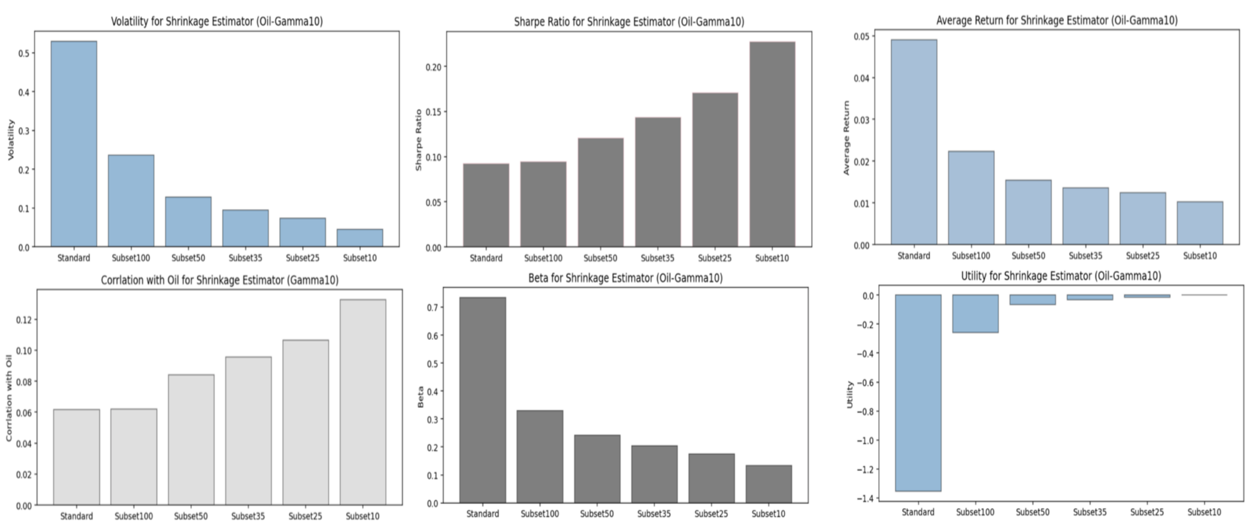

***Figure 2O:* Performance Metrics for Shrinkage inputs, Oil Constraint, Gammas 10**

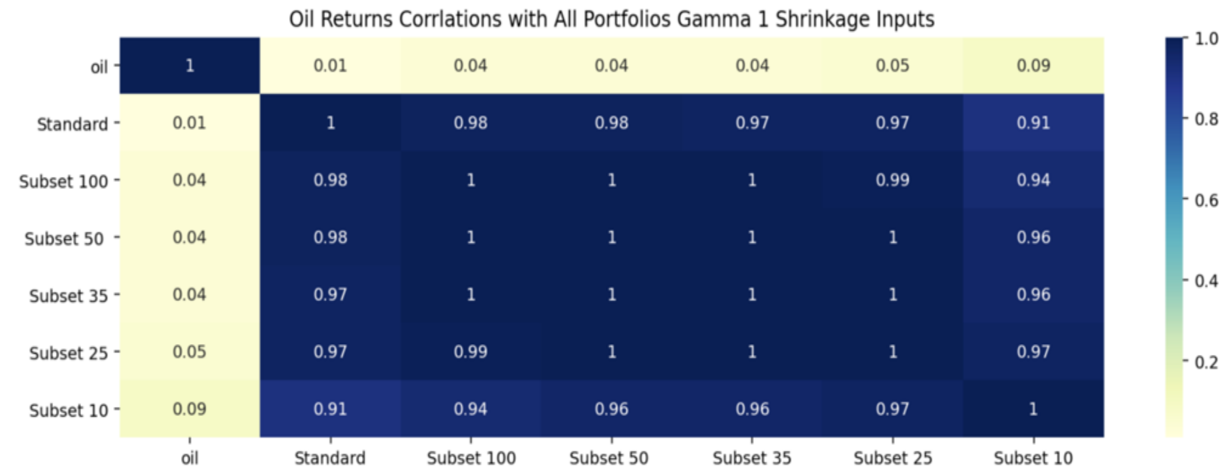

***Figure 2P:* Oil Return Correlations with All Portfolios Gamma 1 Shrinkage Inputs**

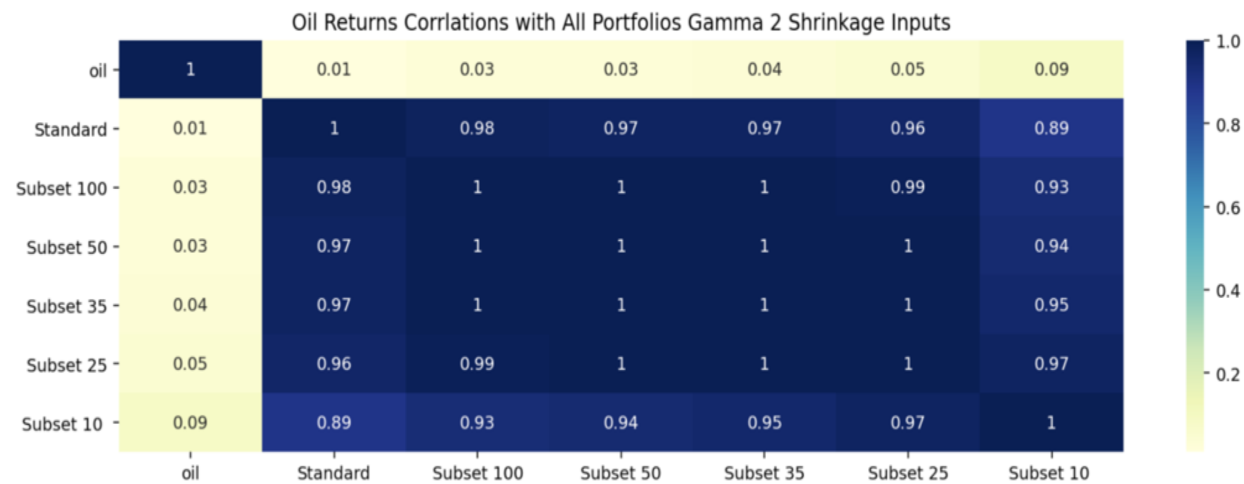

***Figure 2Q:* Oil Return Correlations with All Portfolios Gamma 2 Shrinkage Inputs**

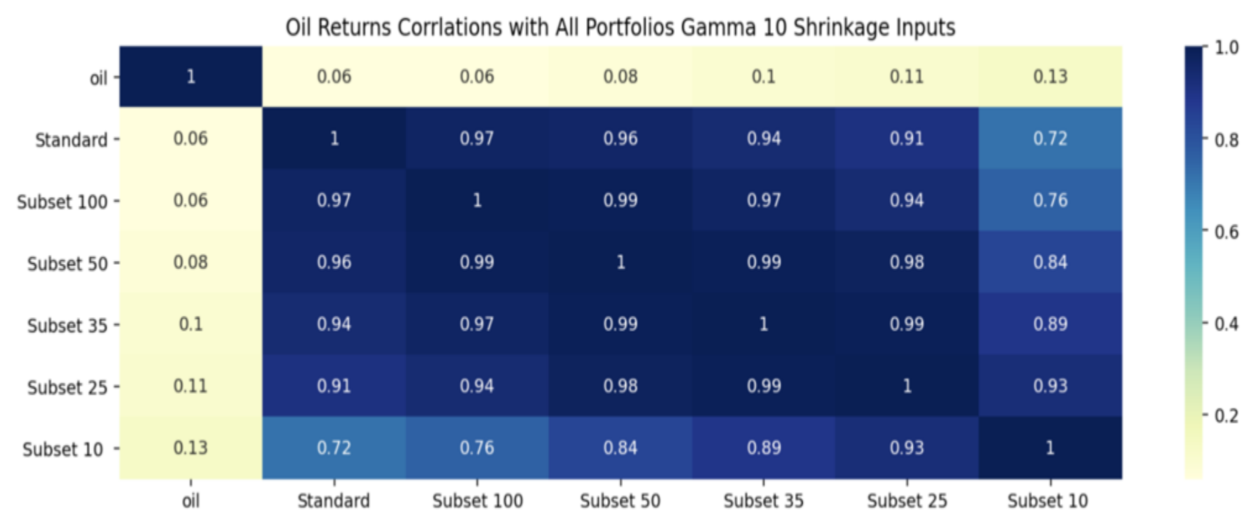

***Figure 2R:* Oil Return Correlations with All Portfolios Gamma 10 Shrinkage Inputs**

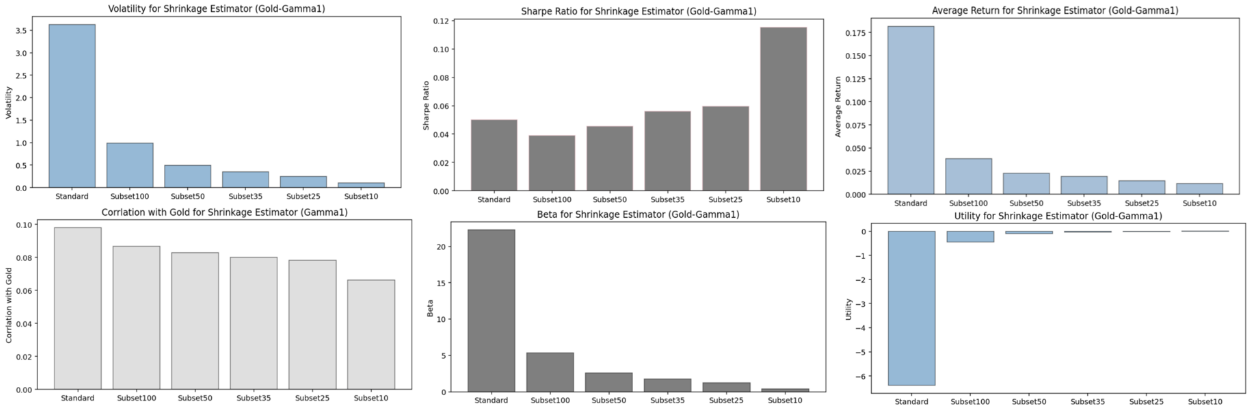

***Figure 2S:* Performance Metrics for Shrinkage inputs, Gold Constraint, Gammas 1**

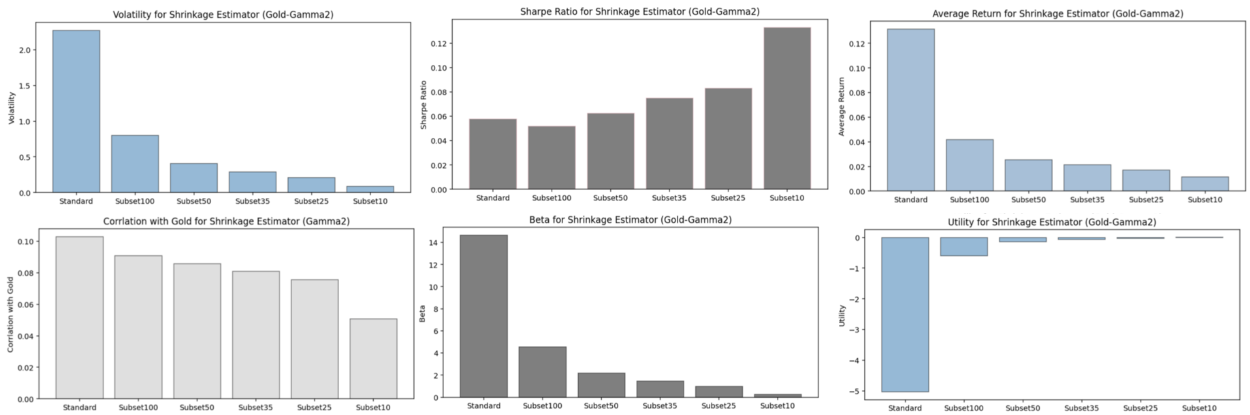

***Figure 2T:* Performance Metrics for Shrinkage inputs, Gold Constraint, Gammas 2**

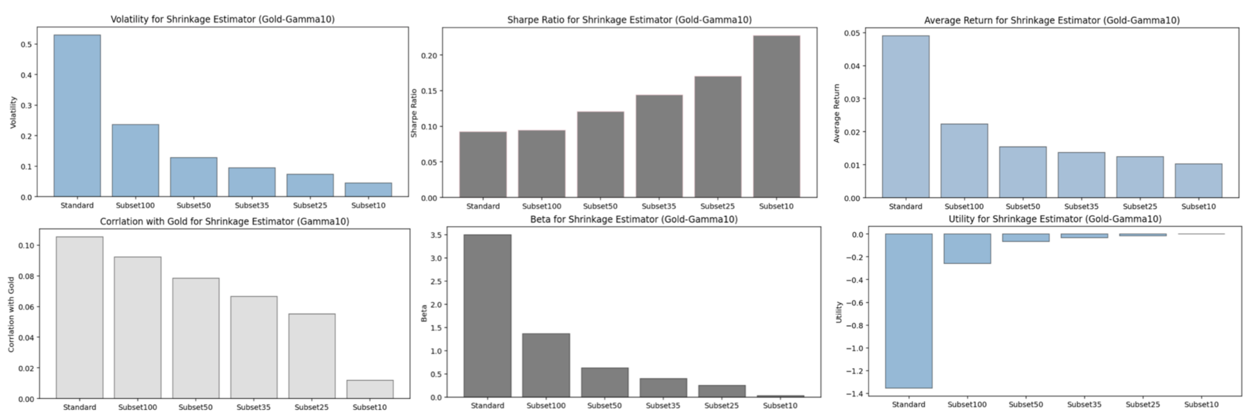

***Figure 2U:* Performance Metrics for Shrinkage inputs, Gold Constraint, Gammas 10**

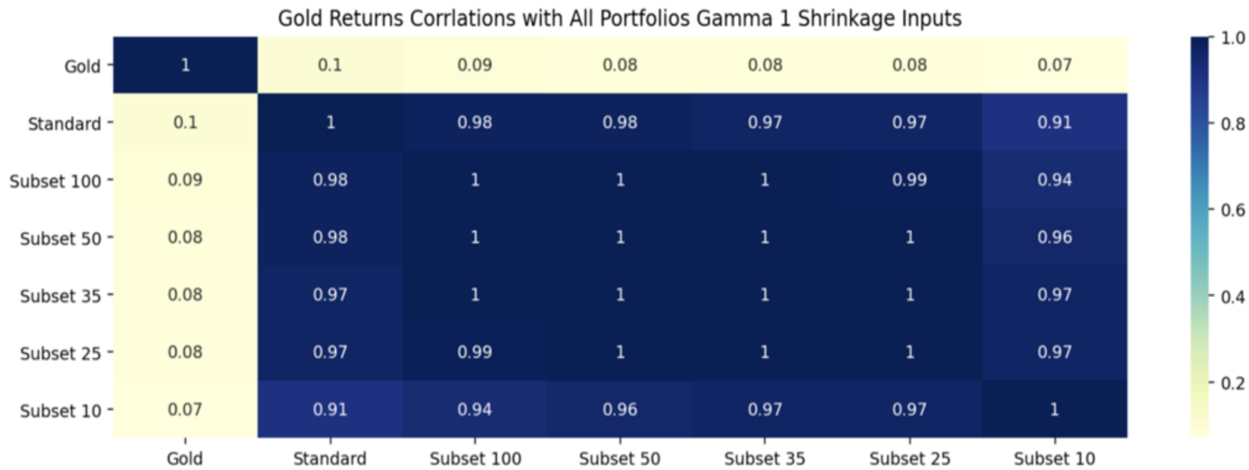

Figure 2V: ***Gold Return Correlations with All Portfolios Gamma 1 Shrinkage Inputs***

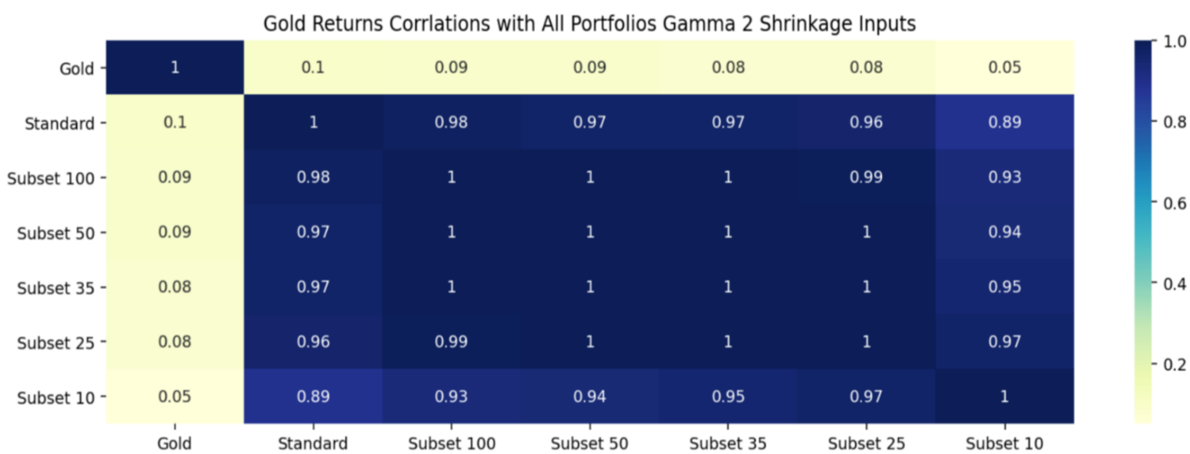

Figure 2W: ***Gold Return Correlations with All Portfolios Gamma 2 Shrinkage Inputs***

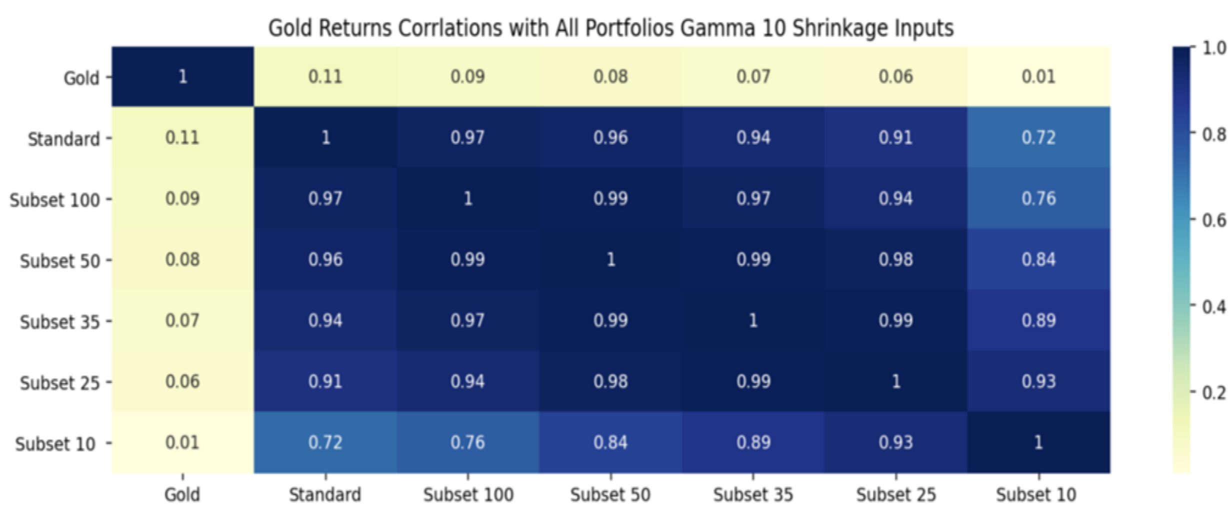

Figure 2X: ***Gold Return Correlations with All Portfolios Gamma 10 Shrinkage Inputs***

Table B7: Performance Statistics for Factors Inputs Portfolios Gamma 1, Oil Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0899	0.0309	0.0210	0.0182	0.0151	0.0116
Volatility	2.8127	0.7595	0.4147	0.3018	0.2231	0.0967
Sharpe Ratio	0.0320	0.0407	0.0506	0.0602	0.0677	0.1205
Utility	-3.8658	-0.2575	-0.0650	-0.0274	-0.0098	0.0070
Correlation	0.0410	0.0059	0.0186	0.0265	0.0388	0.0863
Beta	2.5893	0.1010	0.1732	0.1799	0.1946	0.1874
Alpha	0.0933	0.0311	0.0207	0.0179	0.0149	0.0114
Alpha t-Stat	0.5737	0.7067	0.8646	1.0267	1.1515	2.0457
Beta-t-stat	-0.7080	-0.1021	0.3210	0.4582	0.6707	1.4945

Table B8: Performance Statistics for Factors Inputs Portfolio Gamma 2, Oil Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0674	0.0319	0.0223	0.0185	0.0166	0.0122
Volatility	2.0029	0.6033	0.3264	0.2387	0.1797	0.0834
Sharpe Ratio	0.0336	0.0528	0.0684	0.0775	0.0924	0.1468
Utility	-3.9442	-0.3321	-0.0842	-0.0385	-0.0157	0.0053
Correlation	0.0444	0.0202	0.0035	0.0189	0.0265	0.0755
Beta	1.9999	0.2733	0.0259	0.1011	0.1069	0.1415
Alpha	0.0700	0.0322	0.0223	0.0184	0.0165	0.0121
Alpha t-Stat	0.6045	0.9230	1.1812	1.3299	1.5839	2.5061
Beta-t-stat	-0.7680	-0.3482	0.0610	0.3255	0.4575	1.3073

Table B9: Performance Statistics for Factors Inputs Portfolio Gamma 10, Oil Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0299	0.0170	0.0134	0.0127	0.0115	0.0101
Volatility	0.4832	0.1914	0.1028	0.0779	0.0618	0.0420
Sharpe Ratio	0.0619	0.0886	0.1307	0.1628	0.1865	0.2412
Utility	-1.1375	-0.1662	-0.0394	-0.0177	-0.0076	0.0013
Correlation	0.0316	0.0129	0.0121	0.0419	0.0471	0.1140
Beta	0.3431	0.0554	0.0279	0.0733	0.0655	0.1076
Alpha	0.0304	0.0170	0.0134	0.0126	0.0114	0.0100
Alpha t-Stat	1.0867	1.5374	2.2529	2.7948	3.2032	4.1382
Beta-t-stat	-0.5458	-0.2224	0.2088	0.7236	0.8146	1.9803

Table B10: Performance Statistics for Shrinkage Inputs Portfolios Gamma 1, Oil Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.1815	0.0383	0.0246	0.0185	0.0157	0.0120
Volatility	3.6266	0.9897	0.4977	0.3462	0.2483	0.0990
Sharpe Ratio	0.0500	0.0387	0.0494	0.0533	0.0632	0.1212
Utility	-6.3946	-0.4514	-0.0993	-0.0415	-0.0151	0.0071
Correlation	0.0095	0.0357	0.0402	0.0433	0.0500	0.0852
Beta	0.7723	0.7932	0.4492	0.3366	0.2791	0.1896
Alpha	0.1804	0.0373	0.0240	0.0180	0.0153	0.0118
Alpha t-Stat	0.8600	0.6517	0.8338	0.9005	1.0675	2.0596
Beta-t-stat	0.1636	0.6162	0.6941	0.7476	0.8649	1.4768

Table B11: Performance Statistics for Shrinkage Inputs Portfolios Gamma 2, Oil Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.1315	0.0417	0.0254	0.0215	0.0171	0.0116
Volatility	2.2721	0.8024	0.4053	0.2869	0.2067	0.0865
Sharpe Ratio	0.0579	0.0520	0.0627	0.0750	0.0829	0.1336
Utility	-5.0309	-0.6021	-0.1389	-0.0608	-0.0256	0.0041
Correlation	0.0076	0.0283	0.0333	0.0437	0.0495	0.0880
Beta	0.3879	0.5107	0.3029	0.2817	0.2297	0.1710
Alpha	0.1310	0.0411	0.0250	0.0212	0.0168	0.0113
Alpha t-Stat	0.9968	0.8851	1.0680	1.2759	1.4101	2.2726
Beta-t-stat	0.1312	0.4892	0.5745	0.7554	0.8551	1.5251

Table B12: Performance Statistics for Shrinkage Inputs Portfolios Gamma 10, Oil Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0490	0.0223	0.0155	0.0137	0.0125	0.0102
Volatility	0.5299	0.2366	0.1279	0.0950	0.0733	0.0450
Sharpe Ratio	0.0925	0.0944	0.1209	0.1438	0.1706	0.2274
Utility	-1.3549	-0.2575	-0.0664	-0.0315	-0.0144	0.0001
Correlation	0.0617	0.0620	0.0841	0.0955	0.1064	0.1326
Beta	0.7342	0.3293	0.2418	0.2039	0.1752	0.1341
Alpha	0.0481	0.0219	0.0151	0.0134	0.0123	0.0101
Alpha t-Stat	1.5706	1.6034	2.0536	2.4478	2.9115	3.8972
Beta-t-stat	1.0666	1.0718	1.4573	1.6560	1.8467	2.3096

Table B1: Performance Statistics for Factors Inputs Portfolio Gamma 1, Gold Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0899	0.0308	0.0206	0.0186	0.0160	0.0117
Volatility	2.8127	0.7578	0.4169	0.3019	0.2239	0.0957
Sharpe Ratio	0.0320	0.0407	0.0495	0.0615	0.0716	0.1225
Utility	-3.8658	-0.2563	-0.0663	-0.0270	-0.0090	0.0071
Correlation	0.0931	0.1056	0.0979	0.0878	0.0896	0.0575
Beta	16.3996	5.0127	2.5561	1.6616	1.2573	0.3448
Alpha	0.0527	0.0194	0.0148	0.0148	0.0132	0.0109
Alpha t-Stat	0.3219	0.4416	0.6124	0.8417	1.0122	1.9602
Beta-t-stat	1.6136	1.8329	1.6976	1.5224	1.5536	0.9944

Table B2: Performance Statistics for Factors Inputs Portfolio Gamma 2, Gold Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0674	0.0318	0.0224	0.0185	0.0166	0.0122
Volatility	2.0029	0.6033	0.3264	0.2387	0.1797	0.0832
Sharpe Ratio	0.0336	0.0527	0.0686	0.0777	0.0926	0.1462
Utility	-3.9442	-0.3322	-0.0841	-0.0384	-0.0157	0.0052
Correlation	0.0951	0.0973	0.0988	0.0978	0.0822	0.0557
Beta	11.9276	3.6771	2.0195	1.4618	0.9259	0.2903
Alpha	0.0403	0.0235	0.0178	0.0152	0.0145	0.0115
Alpha t-Stat	0.3459	0.6689	0.9392	1.0975	1.3892	2.3707
Beta-t-stat	1.6484	1.6874	1.7132	1.6956	1.4243	0.9624

Table B3: Performance Statistics for Factors Inputs Portfolio Gamma 10, Gold Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0299	0.0169	0.0134	0.0127	0.0115	0.0101
Volatility	0.4832	0.1915	0.1028	0.0779	0.0618	0.0420
Sharpe Ratio	0.0619	0.0884	0.1306	0.1628	0.1865	0.2410
Utility	-1.1375	-0.1664	-0.0394	-0.0177	-0.0076	0.0013
Correlation	0.1122	0.1068	0.0954	0.0829	0.0717	0.0200
Beta	3.3980	1.2815	0.6143	0.4049	0.2774	0.0526
Alpha	0.0222	0.0140	0.0120	0.0118	0.0109	0.0100
Alpha t-Stat	0.7916	1.2607	2.0132	2.5942	3.0251	4.0775
Beta-t-stat	1.9500	1.8549	1.6541	1.4364	1.2401	0.3447

Table B4: Performance Statistics for Shrinkage Inputs Portfolio Gamma 1, Gold Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.1815	0.0385	0.0227	0.0195	0.0147	0.0116
Volatility	3.6266	0.9900	0.4979	0.3468	0.2466	0.1008
Sharpe Ratio	0.0500	0.0389	0.0455	0.0562	0.0597	0.1154
Utility	-6.3946	-0.4515	-0.1013	-0.0407	-0.0157	0.0066
Correlation	0.0980	0.0868	0.0829	0.0801	0.0783	0.0664
Beta	22.2623	5.3823	2.5854	1.7400	1.2092	0.4190
Alpha	0.1309	0.0263	0.0168	0.0155	0.0120	0.0107
Alpha t-Stat	0.6210	0.4566	0.5795	0.7689	0.8335	1.8185
Beta-t-stat	1.6996	1.5037	1.4358	1.3870	1.3552	1.1484

Table B5: Performance Statistics for Shrinkage Inputs Portfolio Gamma 2, Gold Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.1315	0.0417	0.0253	0.0215	0.0172	0.0115
Volatility	2.2721	0.8024	0.4053	0.2868	0.2067	0.0866
Sharpe Ratio	0.0579	0.0520	0.0625	0.0751	0.0832	0.1331
Utility	-5.0309	-0.6021	-0.1389	-0.0607	-0.0255	0.0040
Correlation	0.1029	0.0909	0.0856	0.0808	0.0757	0.0507
Beta	14.6405	4.5684	2.1745	1.4519	0.9803	0.2749
Alpha	0.0983	0.0313	0.0204	0.0182	0.0150	0.0109
Alpha t-Stat	0.7446	0.6716	0.8654	1.0923	1.2437	2.1588
Beta-t-stat	1.7850	1.5754	1.4839	1.3993	1.3109	0.8756

Table B6: Performance Statistics for Shrinkage Inputs Portfolio Gamma 10, Gold Constraint

		Subset Size				
	Standard	100	50	35	25	10
Arithmetic Mean	0.0490	0.0224	0.0154	0.0137	0.0125	0.0102
Volatility	0.5299	0.2365	0.1279	0.0950	0.0733	0.0450
Sharpe Ratio	0.0925	0.0945	0.1208	0.1440	0.1704	0.2274
Utility	-1.3549	-0.2574	-0.0663	-0.0315	-0.0144	0.0001
Correlation	0.1053	0.0923	0.0783	0.0665	0.0552	0.0118
Beta	3.4966	1.3682	0.6273	0.3957	0.2533	0.0334
Alpha	0.0411	0.0192	0.0140	0.0128	0.0119	0.0102
Alpha t-Stat	1.3347	1.3988	1.8828	2.3088	2.7877	3.8636
Beta-t-stat	1.8284	1.6007	1.3557	1.1496	0.9536	0.2044

Bibliography

Baur, D. G. (2012). Financial contagion and the real economy. *Journal of Banking & Finance, 36*(10), 2680–2692. https://doi.org/10.1016/j.jbankfin.2011.05.019

Bellman, R. (1957). *Dynamic programming*. Princeton University Press.

Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance, 52*(1), 57–82. https://doi.org/10.1111/j.1540-6261.1997.tb03808.x

Cooper, J. C. B. (2003). Price elasticity of demand for crude oil: Estimates for 23 countries. *OPEC Review, 27*(1), 1–8. https://doi.org/10.1111/1468-0076.00121

Frost, P. A., & Savarino, J. E. (1986). An Empirical Bayes Approach to Efficient Portfolio Selection. *Journal of Financial and Quantitative Analysis, 21*(3), 293–305. https://doi .org/10.2307/2331043

Gillen, B. (2016). *The subset optimization for asset allocation* (Social Science Working Paper No. 1421). California Institute of Technology. https://authors.library.caltech.edu/79336/1/ sswp1421.pdf

https://openknowledge.worldbank.org/bitstream/handle/10986/28040/577210NWP0Box353766 B01PUBLIC10gcwp021web.pdf?sequence=1&isAllowed=y

Kan, R., & Zhou, G. (2007). Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis, 42*(3), 621–656.

Yu, C., & Yang, L. (2017). Implications of sovereign wealth funds on the governance of the Middle East region. Asian Journal of Middle Eastern and Islamic Studies, 11(4), 64–75. https://doi.org/10.1080/25765949.2017.12023318

Kritzman, M. (2006). Are optimizers error maximizers? *Journal of Portfolio Management, 32*(4), 66–69. https://doi.org/10.3905/jpm.2006.644197

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

Public Investment Fund Program. (2017). *The public investment fund program*. https://www.pif.gov.sa/en/PIFContentProgram/PIF%20Program_EN.pdf

Saudi Arabian Government. (2023). *Statistics.* https://www.stats.gov.sa/en/823

Saudi Central Bank. (2021). *Homepage*. Retrieved June 3, 2021, from https://www.sama.gov.sa/en-us/pages/default.aspx

Sovereign Wealth Fund Institute. (2023). *Top 100 largest sovereign wealth fund rankings by total assets*. Retrieved July 12, 2023, from https://www.swfinstitute.org/fund-rankings/sovereign-wealth-fund

Vision 2030. (2022). *Economy and Business*. https://www.vision2030.gov.sa/thekingdom/explore/economy