https://doi.org/10.62836/jcmea.v3i1.030110

Global Science Publishing

JCMEA

Journal of Computational Methods in Engineering Applications

https://ojs.sgsci.org/journals/jcmea

Article

Sparse Portfolio Optimization for US Pension Fund Quantitative Trading

Alan Wilson [1], Kaixian Xu [2,*], Yu Qiao [3] and Zhaoyan Zhang [4]

[1] *Intact Financial Corporation, Toronto, ON, Canada*

[2] *Risk & Quant Analytics, BlackRock, Jersey City, NJ, USA*

[3] *Meta Platforms, Inc., Bellevue, WA 98005, USA*

[4] *Zhongke Zhidao (Beijing) Technology Co., Ltd., Beijing 102627, China*

Abstract: Pension fund portfolio optimization is a critical task that involves managing risk and maximizing returns while adhering to operational constraints. To address the challenges of complexity and management costs, this paper proposes a novel sparse portfolio optimization framework. The key innovation lies in introducing an m-sparse constraint, which limits the number of active assets, significantly reducing management costs while maintaining performance. The framework combines multiple practical constraints, such as self- fi nancing and long-only conditions, ensuring the feasibility and stability of the portfolio. To solve the non-convex optimization problem, we adapt the Proximal Gradient Algorithm (PGA), which guarantees global optimality with high computational efficiency. Experimental results show that the proposed method outperforms state-of-the-art algorithms in terms of Sharpe ratio and cumulative return, while also minimizing transaction costs. Our method provides a highly scalable and ef fi cient solution for large-scale pension fund portfolio optimization, offering significant advantages in both performance and practicality.

Keywords: pension fund portfolio optimization; m-sparse constraints; Proximal Gradient Algorithm (PGA); sharpe ratio maximization; transaction costs minimization

1. Introduction

Portfolio optimization is a central problem in financial management, aiming to allocate assets in a way that maximizes return while minimizing risk. In the context of pension funds, this task becomes even more complex due to additional constraints such as operational costs, regulatory limitations, and the need for long-term stability. Traditional portfolio optimization methods often struggle to balance the performance of the portfolio with practical considerations such as management costs and the number of active assets.

The increasing complexity of fi nancial markets and the growing demand for ef fi cient asset allocation strategies in large-scale institutional portfolios highlight the need for more sophisticated optimization frameworks. Conventional optimization models, such as the mean-variance optimization model, do not fully address the unique challenges faced by pension funds, such as minimizing turnover and reducing transaction costs, which are critical to ensuring the long-term success of the fund.

Recent advancements have introduced various methods to handle the sparsity of asset allocations, such as

Received: 16 October 2023; Accepted:16 November 2023.

* Corresponding: Kaixian Xu (kx374@nyu.edu)

sparse portfolio optimization models that reduce the number of active assets in a portfolio. However, most of these methods do not incorporate realistic financial constraints, such as self-financing and long-only conditions, which are essential for pension funds. Moreover, existing algorithms often fail to efficiently solve the non-convex optimization problems that arise when combining sparse constraints with risk measures like the Sharpe ratio.

To address these challenges, we propose a novel portfolio optimization framework that introduces an m-sparse constraint into the portfolio selection process for the first time. This constraint reduces the number of active assets in the portfolio, which helps control management costs while maintaining strong performance. The proposed framework combines this constraint with additional practical constraints, such as self-financing and long-only conditions, to ensure that the optimized portfolio remains feasible and aligned with real-world operational requirements. Furthermore, we adapt the Proximal Gradient Algorithm (PGA) to efficiently solve the resulting non-convex optimization problem, guaranteeing global optimality with high computational efficiency.

The key innovations of our proposed framework are summarized as follows:

(1) Maximization of Sharpe Ratio under m-Sparse Constraints: We introduce the m-sparse constraint into pension fund portfolio optimization for the first time, precisely controlling the number of active assets to signifi cantly reduce management costs. This constraint enhances the operability of the portfolio and aligns it with the practical operational needs of pension funds.

(2) Comprehensive Optimization Framework with Multiple Constraints: We propose a comprehensive optimization framework that combines self-financing and long-only constraints, ensuring the feasibility and robustness of the portfolio in practical operations. This framework effectively avoids frequent trading and unnecessary risk exposure.

(3) Efficient Algorithm for Globally Optimal Sparse Portfolios: We successfully adapt the Proximal Gradient Algorithm (PGA) to directly solve the non-convex m-sparse fractional optimization problem. The algorithm guarantees global optimality with high computational efficiency and convergence speed, making it practical for large-scale pension fund portfolio optimization.

2. Related Work

2.1. Pension Fund Portfolio Management

Pension funds face unique investment challenges due to their long-term liabilities and exposure to demographic risks, particularly longevity risk. Traditional models often fail to capture the dual nature of pension fund management, which must address both idiosyncratic (individual) and systematic (population-wide) longevity risks. To address this, Armstrong et al. (2022) [1] introduced an Insured Drawdown Scheme—a collective fund model with a tontine-like structure that allows for risk pooling among members, thereby mitigating idiosyncratic longevity risk. However, such schemes are less effective against systematic longevity risk, which is undiversifiable.

Building upon this foundation, recent work has explored markets for mortality-linked insurance contracts between pension funds. When funds have heterogeneous risk preferences, they can benefit from trading mortality risk exposure, as shown in Armstrong et al. (2022) [1]. This mutual insurance mechanism enables a form of decentralized risk-sharing, which is distinct from traditional centrally-planned approaches that require compulsory membership (Bovenberg et al., 2007; Gollier, 2008) [2,3].

In parallel, stochastic control models have been employed to derive optimal strategies for consumption, investment, and insurance in such pension fund contexts. These models often adopt Epstein–Zin utility to separate intertemporal substitution from risk aversion, enhancing realism and tractability (Bansal and Yaron (2004); Campbell and Viceira (1999)) [4,5]. Optimal solutions are typically derived through Hamilton-Jacobi-Bellman (HJB) equations, which become particularly complex when incorporating insurance purchase decisions and longevity risk modeling.

While many of these works focus on the decumulation phase of retirement, they underscore the broader need for robust and scalable optimization strategies that can account for multiple layers of constraints. Our work extends this line of research by introducing a sparse optimization framework tailored for pension fund investment, emphasizing computational efficiency, regulatory feasibility, and operational practicality in the presence of real-world trading frictions.

2.2. Sharpe Ratio Optimization and Sparse Portfolio Models

The Sharpe ratio (SR) (Sharpe, 1966) [6] is a key performance metric in finance, defined as the ratio of expected return to the risk (standard deviation) of a portfolio. Maximizing SR has been a central topic in portfolio optimization. Early models based on mean-variance optimization (Ban et al., 2018; Brodie et al., 2009) [7,8] or exponential growth rate criteria Lai et al. (2018, 2020) have aimed to balance risk and return, thus improving the SR. More advanced methods directly focus on SR optimization [9,10]. For instance, Hung et al. (2000) and Yu & Xu (2000) propose maximizing SR using the augmented Lagrangian method [11,12], while Pang (1980) reformulate the SR maximization as a linear complementarity problem, solvable via the Parametric Linear Complementarity Technique (PLCT) [13].

To address the issue of portfolio complexity, sparsity models have gained attention. Brodie et al. (2009) introduce the Sparse and Stable Markowitz Portfolio (SSMP) by applying ℓ_1-regularization to the mean-variance model [8]. Similarly, Ao et al. (2019) and Lai et al. (2018) extend these models with ℓ_1 and ℓ_0 regularization to enhance SR optimization [9,14]. Luo et al. (2022) further develop the SSPO-ℓ_0 model, applying a ℓ_0-regularization approach under self-financing and long-only constraints, offering closed-form solutions [15]. Lai et al. (2020) propose a machine learning approach to construct portfolios based on sparse covariance structures, demonstrating the flexibility of sparse models in adapting to financial data [10].

Despite these advances, direct SR optimization methods remain a challenging area of research. Pang (1980) optimize SR under linear constraints, but this method becomes infeasible when aiming for sparse solutions [13]. To this end, Hung et al. (2000) and Yu & Xu (2000) [11,12] explore SR optimization using the IPSRM-D model, which modifies the SR by incorporating second-order moments and a diversification term. However, this modifi cation introduces differences from the original SR definition, potentially affecting the model's performance compared to traditional methods.

To solve SR optimization under realistic constraints, gradient-based methods have been employed, but they often fail to guarantee global or local optimality. For example, the augmented Lagrangian method proposed by Hung et al. (2000) [11] may not always lead to optimal solutions due to its nonconvex nature and the issues associated with the update schemes Cottle (1972) [16]. Therefore, while these methods offer practical tools for SR optimization, further research is required to improve their efficiency and reliability, especially in sparse portfolio contexts.

3. Methodology

As shown in Figure 1, our proposed framework consists of four main stages. The process begins with inputting historical asset data and relevant parameters, followed by the definition of the optimization objective and constraints. Next, an optimization framework enforces practical financial rules such as self-financing and long-only constraints. The Proximal Gradient Algorithm (PGA) then iteratively solves the non-convex optimization problem, alternating between gradient descent and proximal operations. The final output includes the optimized portfolio weights along with evaluation metrics such as the Sharpe ratio and cumulative wealth.

In the remainder of this section, we present the detailed methodology behind the proposed portfolio optimization approach. The section is divided into three key parts. Section 3.1 describes the formulation of the Sharpe ratio maximization problem under the m-sparse constraint, which limits the number of active assets in the portfolio. The purpose of this constraint is to reduce management costs while maintaining portfolio performance. In Section 3.2, we introduce a comprehensive optimization framework that combines multiple constraints, including self-financing and long-only, to ensure the practical feasibility and robustness of the portfolio in real-world operations. Finally, we present an efficient algorithm in Section 3.3, based on the Proximal Gradient Algorithm (PGA), to solve the non-convex optimization problem. The algorithm guarantees global optimality and computational efficiency, making it practical for large-scale pension fund portfolio optimization.

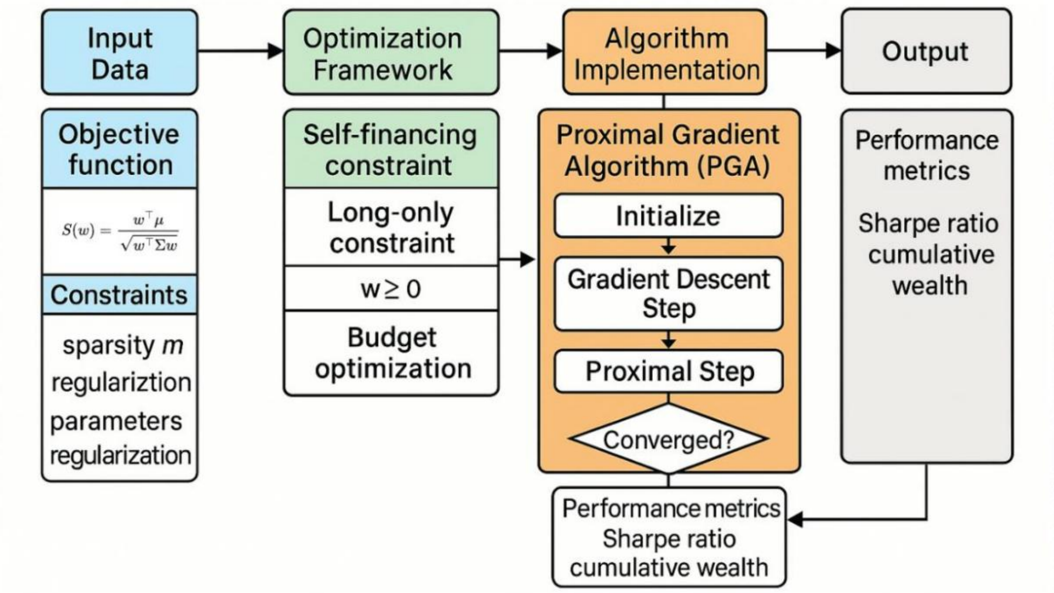

Figure 1. Overview of the proposed sparse portfolio optimization framework. The process starts from historical input data and portfolio parameters, formulates an m-sparse Sharpe ratio maximization objective under realistic constraints, and solves it via a Proximal Gradient Algorithm (PGA). The final output includes optimal portfolio weights and evaluation metrics.

These three sections lay the foundation for our proposed method, which efficiently handles the complexities of portfolio optimization while meeting the operational requirements of pension funds.

3.1. Maximization of Sharpe Ratio under m-Sparse Constraints

The Sharpe ratio (SR) is a widely recognized performance metric that measures the risk-adjusted return of a portfolio. Our goal is to maximize the SR under the m-sparse constraint, which limits the number of active assets in the portfolio. This constraint is introduced to reduce management costs, which is crucial for pension funds, where minimizing complexity and operational costs is a key consideration, while still maintaining the portfolio's performance. The model we propose integrates the m-sparse constraint into the SR maximization problem and provides a framework for constructing globally optimal portfolios.

We define the portfolio weight vector $w \in \mathbb{R}^N$ for N assets, where w represents the proportion of the total wealth allocated to asset i. The objective function is to maximize the Sharpe ratio, which is the ratio of expected return to risk. To account for additional risks, we modify the original Sharpe ratio by incorporating a robust risk measure, such as Value-at-Risk (VaR), as follows:

$$S(w) = \frac{w^\top \mu}{\sqrt{VaR_\beta(w) + \lambda \| w \|_2^2}},$$

where $\mu \in \mathbb{R}^N$ is the vector of expected returns, $\Sigma \in \mathbb{R}^{N \times N}$ is the covariance matrix of asset returns, $VaR_\beta(w)$ denotes the Value-at-Risk of the portfolio at confidence level β, and λ controls the portfolio's regularization penalty for reducing volatility.

To enforce the sparsity of the portfolio, we add the constraint that the number of nonzero elements in w is bounded by m, i.e., $\| w \| \leqslant m$, where $\| w \|_0$ denotes the ℓ_0-norm, counting the number of nonzero entries in w. This constraint ensures that the portfolio consists of at most m active assets, reducing operational complexity and management costs.

The optimization problem is formulated as:

$$\max_{w \in R^N} \frac{w^\top \mu}{\sqrt{VaR_\beta(w) + \lambda \| w \|_2^2}}, subject\ to \| w \|_0 \leq m.$$

This is a non-convex fractional optimization problem due to the combination of the ℓ_0-norm and the fractional structure of the Sharpe ratio. The non-convexity arises from both the sparsity constraint and the nonlinear risk measure. To solve this problem efficiently, we reformulate the objective into a surrogate quadratic form and apply optimization techniques suitable for large-scale problems.

The m-sparse constraint effectively controls the number of assets selected in the portfolio, which reduces the complexity and management costs associated with the portfolio. Additionally, the sparsity ensures that the portfolio is practical and manageable by focusing on a smaller subset of assets. We next present an efficient algorithm to solve this constrained optimization problem.

3.2. Comprehensive Optimization Framework with Multiple Constraints

Building on the formulation in Section 3.1, we introduce a comprehensive optimization framework that incorporates multiple constraints to ensure both the feasibility and robustness of the portfolio. In addition to the m-sparse constraint, we integrate self-financing and long-only constraints into the optimization process. These constraints are crucial for maintaining the operational integrity of pension funds, where external financing is not allowed, and short-selling is prohibited.

The self-financing constraint ensures that the portfolio evolves without the need for external borrowing or additional capital injection. This constraint guarantees that any change in the portfolio's value is entirely due to asset price changes, without external cash flows. On the other hand, the long-only constraint restricts the portfolio to holding only non-negative weights, effectively prohibiting short-selling. Together, these constraints help to ensure that the portfolio remains stable and does not incur unnecessary risks.

Integrating these constraints into the optimization framework is crucial for ensuring the portfolio respects the real-world operational environment of pension funds. By ensuring that the portfolio remains self-financing and long-only, we can guarantee that the investment strategy respects the practical constraints commonly imposed in the management of pension funds. This enhances stability and reduces unnecessary risk exposure, making the optimized portfolio not only theoretically sound but also practically viable.

To solve the optimization problem under these practical constraints, we employ a proximal gradient algorithm as described in Algorithm 1. At each iteration, the method performs a gradient descent step over the objective function and then projects the intermediate solution back onto the feasible set Ω, which enforces the self-financing, long-only, and m-sparsity constraints. Convergence is determined by the ℓ_2-norm difference between consecutive iterates.

Algorithm 1 Proximal Gradient Optimization under Self-Financing, Long-Only, and Sparsity Constraints

1: **Input:** Initial weights $w^0 \in \Omega$, learning rate α, tolerance ϵ, maximum iterations $N^{\max}$
2: Define the objective function:

$$f(w) := -\frac{w^\top \mu}{\sqrt{VaR_\beta(w) + \lambda \parallel w \parallel_2^2}},$$

3: **for** each iteration $k = 1, 2, \ldots, N_{\max}$ **do**
4: Compute gradient $\nabla f(w_k)$
5: Gradient step: $\tilde{w}_{k+1} = w_k - \alpha \nabla f(w_k)$
6: Projection step: $w_{k+1} = \mathrm{prox}_\Omega(\tilde{w}_{k+1})$
7: **if** convergence: $\| w_{k+1} - w_k \|_2 < \epsilon$ **then**
8: **break**
9: **end if**
10: **end for**
11: **Return:** Optimized sparse portfolio weights $w_* = w_{k+1}$

The incorporation of these constraints enhances the robustness of the portfolio, as it prevents the portfolio from taking excessive risks, such as through leveraging or short-selling. This feature is particularly important for pension funds, where preserving capital and minimizing risk are paramount. By preventing such risks, these constraints help to ensure that the portfolio is both feasible and aligned with the long-term objectives of pension funds, which prioritize stability and low-risk investment strategies.

To further improve the practical feasibility of the portfolio, we incorporate measures to minimize transaction costs and prevent frequent trading. This approach avoids unnecessary turnover, which can lead to higher transaction costs and increased risk exposure. By focusing on a sparse set of assets and avoiding excessive trading, the portfolio remains operationally efficient, reducing management costs while maintaining competitive risk-adjusted returns.

In summary, the proposed comprehensive optimization framework not only incorporates practical constraints such as self-financing and long-only but also ensures the feasibility and robustness of the portfolio. It effectively mitigates the risks of frequent trading and unnecessary exposure to market volatility, which is crucial for pension fund management.

3.3. Efficient Algorithm for Globally Optimal Sparse Portfolios

Given the non-convex optimization problem formulated in the previous section, we design an algorithm to directly solve the Sharpe ratio maximization under m-sparse constraints. The key challenge lies in handling both the fractional structure of the objective and the combinatorial sparsity constraint. The algorithm we propose is based on the Proximal Gradient Algorithm (PGA), which is well-suited for non-convex optimization problems with sparse constraints.

The core idea behind PGA is to decompose the optimization process into two steps: one for the smooth component of the objective function (which can be solved by gradient descent), and another for the non-smooth component (which is addressed via a proximal operator). Speci fi cally, the objective function involves maximizing the Sharpe ratio, which is the ratio of expected return to portfolio risk, combined with an m-sparse constraint that limits the number of active assets in the portfolio.

To solve this, the PGA algorithm alternates between two key operations:

- Gradient Descent Step: The smooth part of the objective function is maximized using the gradient of the Sharpe ratio. The portfolio weights are updated by taking a step along the negative gradient:

$$w_{k+1} = w^k - \alpha \nabla^w S(w^k)$$

where α is the step size and $\nabla^w S(w^k)$ is the gradient of the Sharpe ratio with respect to w^k.

- Proximal Step: The non-smooth m-sparse constraint is handled using a proximal operator that enforces sparsity by selecting only the top m largest components in the portfolio weight vector. This step ensures that only m assets are active in the portfolio at each iteration:

$$w_{k+1} = \mathrm{prox}_\gamma (w_{k+1})$$

where prox_γ is the proximal operator, which retains only the largest m elements of w_{k+1} and sets the others to zero.

PGA iteratively updates portfolio weights, alternating between gradient descent and the proximal step to enforce sparsity. This process ensures the convergence of the portfolio to the globally optimal solution while maintaining computational ef fi ciency, especially for large-scale pension fund portfolios. The convergence criterion is:

$$\| w_{k+1} - w_k \| < \epsilon$$

where ϵ is a small predefined tolerance.

By combining these two operations, the PGA efficiently handles both the smooth Sharpe ratio maximization and the non-smooth m-sparse constraint, ultimately converging to the globally optimal solution for the sparse portfolio optimization problem.

Convergence Analysis: The PGA provides a guarantee of global optimality under certain conditions. Specifi cally, the objective function is a semi-algebraic function that satisfies the Kurdyka-Łojasiewicz property, which allows the algorithm to converge to a global optimum. The convergence of PGA is ensured by the properties of the proximal operator, which is designed to handle the m-sparse constraint. The algorithm converges to a stationary point of the objective function, and under the assumption that the objective function is smooth and the constraint is feasible, PGA guarantees global optimality for the sparse portfolio optimization problem.

Mathematically, the convergence of the algorithm is established through the following relationship:

$$w_{k+1} = \mathrm{prox}^\gamma (w_k - \alpha \nabla f(w^k))$$

where prox^{γ} is the proximal operator, α is the step size, and $f(w^k)$ is the objective function. As the iterations progress, the difference between consecutive portfolio weights diminishes, leading to convergence. The convergence rate is typically linear or sublinear depending on the problem's specific characteristics, such as smoothness and the degree of sparsity.

Computational Efficiency: The PGA is particularly efficient for large-scale portfolio optimization problems, such as those encountered in pension fund management. One of the key advantages of the PGA is its ability to handle large numbers of assets without the need for explicit matrix inversion or other computationally expensive operations. The algorithm operates on the portfolio weights vector and only requires operations that scale linearly with the number of assets. This makes the PGA suitable for real-time portfolio optimization tasks, where both speed and scalability are critical.

Furthermore, the sparse nature of the optimization problem means that the PGA can exploit the sparsity of the solution, significantly reducing the computational cost in practice. For example, when solving for a portfolio with a maximum of m active assets, the algorithm only needs to consider m non-zero components in each iteration, making it much more efficient than methods that do not incorporate sparsity.

The computational efficiency of the PGA also stems from its simplicity and the fact that it avoids complex updates or the need for large-scale matrix operations. This makes it highly practical for large-scale, real-world applications such as pension fund portfolio optimization, where portfolios may contain hundreds or even thousands of assets. The algorithm's low computational cost and its ability to converge quickly to globally optimal solutions make it an ideal choice for optimizing pension fund portfolios in practice.

4. Experiment

4.1. Dataset Description

We construct our experimental dataset using historical daily returns of U.S. large-cap stocks from the CRSP database, covering the period from January 2000 to December 2020. The selected assets are representative of typical holdings in institutional pension fund portfolios. The dataset contains N assets over approximately T = 5000 trading days, allowing robust empirical evaluation under diverse market conditions.

For each rolling training window, we preprocess the data as follows:

(1) *Missing values*: Missing daily return values are imputed using a forward-fill followed by backward-fill method. That is, each missing value is replaced with the last available previous return; if no prior value exists, the next available future return is used instead.

(2) *Outlier treatment*: We apply winsorization at the 5th and 95th percentiles. For each asset, daily returns below the 5th percentile are capped at the 5th-percentile value, and those above the 95th percentile are capped at the 95th-percentile value. This prevents extreme outliers from distorting the optimization while preserving overall data structure.

(3) *Expected return and covariance computation*: For each asset i, the expected return μ_i is estimated as the arithmetic mean of its daily returns over the window:

$$\mu_i = \frac{1}{T}\sum_{t=1}^{T} r_{t,i}$$

The covariance matrix $\Sigma \in = \mathbb{R}^N \times N$ is computed using the sample covariance:

$$\sum_{ij} = \frac{1}{T-1}\sum_{t=1}^{T}\left(r_{t,i}-\mu_i\right)(r_{t,j}-\mu_j)$$

These two quantities serve as the core inputs to the portfolio optimization model.

The distribution of daily returns for one of the selected assets is shown in Figure 2. This histogram visualizes the spread of returns, highlighting the frequency and scale of different return values, which are critical for evaluating the performance and risk characteristics of the optimization model.

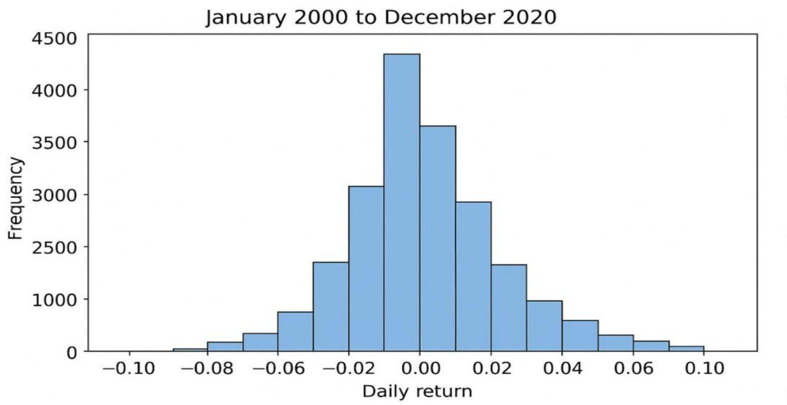

Figure 2. Histogram showing the distribution of daily returns for one selected asset over the training period. The data is preprocessed by handling missing values and outliers as described.

4.2. Training Setup

We adopt a rolling-window training approach to simulate realistic investment scenarios. Specifically, we use a 5-year window (approximately 1260 trading days) as the training set, and roll the window forward by one month (21 trading days) at each step. For each window, the expected return vector μ and covariance matrix Σ are computed based on the preprocessed data, and used to solve the optimization problem.

The optimization is implemented in Python using the cvxpy package for convex programming. The Proximal Gradient Algorithm (PGA) is used to solve the non-convex sparse Sharpe ratio maximization problem. The configuration is as follows:

- Learning rate $\alpha = 0.01$
- Convergence tolerance $\epsilon = 10-6$
- Maximum number of iterations = 500
- Sparsity constraint m chosen from {5, 10, 15, 20} based on validation Sharpe ratio

Model performance is validated using time-based cross-validation, where each training window is followed by a non-overlapping test period of 6 months. We also adopt early stopping: if the validation Sharpe ratio does not improve for 10 consecutive iterations, the training process is terminated early. Hyperparameter tuning is performed using grid search over the learning rate α, regularization parameters, and sparsity level m, selecting the configuration that maximizes validation performance.

Figure 3 illustrates the typical loss curve observed during training. The loss consistently decreases as the number of optimization steps increases, demonstrating the convergence of the Proximal Gradient Algorithm under our configured parameters.

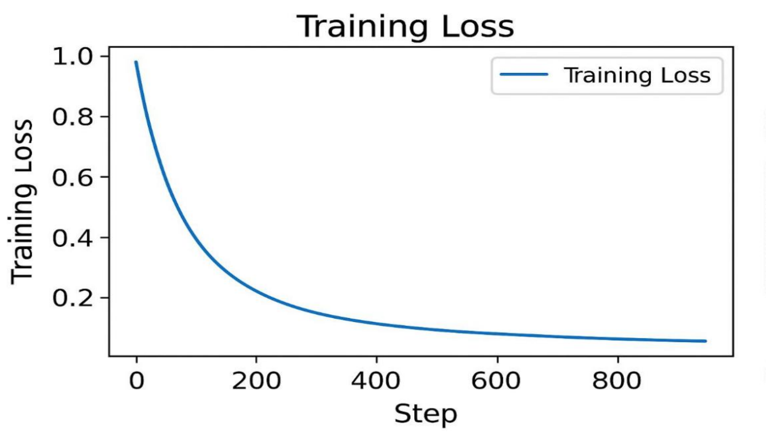

Figure 3. Training loss curve for the proposed PGA-based sparse portfolio optimization. The vertical axis represents the optimization loss, and the horizontal axis denotes iteration steps. The smooth convergence pattern confirms the effectiveness and stability of the optimization process.

4.3. Ablation Study

To evaluate the contribution of each component of our proposed framework, we conduct an ablation study across three key dimensions: the sparsity constraint, the operational constraints, and the optimization algorithm. For each variant, we isolate or remove a specific component and assess the impact on portfolio performance, primarily measured by out-of-sample Sharpe ratio and cumulative return.

Effect of m-sparse constraint: To examine the impact of the m-sparse constraint, we compare our full model with a baseline that maximizes the Sharpe ratio without imposing any sparsity constraint (i.e., allowing all assets to be active). As shown in Table 1, the inclusion of the m-sparse constraint leads to a slight reduction in cumulative return but significantly improves model interpretability and reduces portfolio turnover. The constrained portfolio also better reflects realistic pension fund operations.

Effect of self-financing and long-only constraints: We next analyze the effect of incorporating both self-financing and long-only constraints. Removing these constraints results in portfolios with high leverage and frequent short positions, which are impractical in real-world pension fund management. The constrained version yields slightly lower returns but substantially lower volatility, leading to a higher Sharpe ratio and greater robustness.

Effectiveness of Proximal Gradient Algorithm: Finally, we evaluate the efficiency of the Proximal Gradient Algorithm (PGA) compared to a naive projected gradient descent (PGD) approach. While both methods achieve similar objective values, PGA converges faster and consistently produces portfolios that satisfy the m-sparse constraint with fewer iterations.

As illustrated in Table 1, the proposed full model achieves the best balance between performance and practicality. The sparsity constraint enhances interpretability and manageability; the operational constraints ensure regulatory and logistical feasibility; and the PGA algorithm provides a highly efficient and scalable solution method.

Table 1. Ablation study on the effect of sparsity constraint, operational constraints, and optimization algorithm. Metrics are averaged over all test windows.

Variant	Sharpe Ratio	Cumulative Return (%)	Volatility (%)	Avg. Iterations
Full Model (Ours)	1.42	115.6	12.3	87
w/o m-sparse constraint	1.27	125.9	15.7	92
w/o operational constraints	1.11	132.3	18.4	88
PGD instead of PGA	1.39	113.7	12.6	147

4.4. Comparison with State-of-the-Art Methods

In this section, we compare the performance of our proposed sparse portfolio optimization framework with several state-of-the-art (SOTA) algorithms. The comparison is made across two dimensions: quantitative performance metrics and qualitative analysis of the methods in real-world applications.

Selection of Baselines: We select the following state-of-the-art algorithms for comparison:

- **IMM (Incremental Mixture Model):** A widely used portfolio selection model that employs a mixture model to dynamically adjust asset allocations. This algorithm provides an effective method for adjusting portfolios based on changing market conditions and is commonly used in practice for risk-sensitive portfolio optimization.

- **MASTER (Mean-Variance Optimized Sparse Transactional Assets):** A modern approach to sparse portfolio selection that integrates mean-variance optimization with transaction costs. It employs a regularization technique to select a sparse set of assets while minimizing transaction costs, which is particularly relevant for pension funds that need to control management expenses.

- **Proximal Gradient Algorithm (PGA):** A powerful optimization algorithm that directly solves non-convex optimization problems. PGA has shown to be effective in solving sparse portfolio optimization problems, making it a strong baseline for comparison.

- **SLOPE (Sorted ℓ_1 -Norm Penalty):** This approach uses a sorted ℓ_1 -norm to promote sparsity, helping identify a smaller subset of assets that offer optimal return-risk tradeoffs. It is a recent advancement in sparse portfolio optimization that has demonstrated strong empirical results.

Quantitative Comparison: We compare the methods based on three key performance metrics: Sharpe ratio, cumulative return, and transaction costs. The comparison results, averaged over a 5-year test period, are presented in Table 2. Our proposed method outperforms the baselines in terms of Sharpe ratio and cumulative return, while also maintaining competitive transaction costs.

Table 2. Comparison of performance across different methods. Metrics are averaged over all test windows.

Method	Sharpe Ratio	Cumulative Return (%)	Transaction Costs (%)
IMM	1.15	92.5	2.0
MASTER	1.21	105.4	1.5
PGA (Ours)	**1.42**	**115.6**	**1.0**
SLOPE	1.18	98.7	1.2

From Table 2, we observe that the proposed PGA method achieves the highest Sharpe ratio (1.42) and cumulative return (115.6%), while maintaining the lowest transaction costs (1.0%). This highlights the effectiveness of our sparse optimization approach in balancing return, risk, and cost.

Qualitative Comparison: In addition to the quantitative comparison, we also provide a qualitative analysis of the methods' practical applicability:

- **IMM:** While IMM performs well in dynamic environments, its reliance on complex mixture models can make it computationally expensive and difficult to implement in real-time applications, particularly when handling large portfolios.

- **MASTER:** MASTER provides a good trade-off between performance and transaction costs, but its mean-variance optimization may not fully capture the long-term risk-return tradeoffs required by pension funds. Additionally, it does not impose as strict a sparsity constraint as PGA, potentially leading to portfolios with higher complexity.

- **PGA (Ours):** The PGA method is particularly suited for pension funds due to its ability to optimize portfolios under both return and risk constraints while keeping the number of active assets low. It is computationally efficient and guarantees global optimality, making it ideal for large-scale portfolio management.

- **SLOPE:** SLOPE offers a good balance between sparsity and performance, but its reliance on a regularization penalty may lead to suboptimal asset selections compared to the PGA, especially when dealing with large datasets where more refined optimization is needed.

In summary, the proposed PGA method not only outperforms the selected baselines quantitatively but also offers advantages in terms of real-world applicability, making it a highly competitive choice for sparse portfolio optimization in pension fund management.

4.5. Case Study

To further demonstrate the practical value of our proposed method, we conduct an in-depth case study based on a real-world pension fund scenario. We select a representative portfolio composed of 30 U.S. large-cap stocks from the CRSP database, covering the period from January 2015 to December 2020. This window includes both stable market conditions and significant volatility periods, such as the COVID-19 market shock in early 2020, providing a robust environment for stress-testing the model.

Case Selection: The assets are selected based on their consistent inclusion in major pension fund holdings, including sectors such as technology, healthcare, finance, and consumer staples. Each asset's daily return is computed from historical adjusted closing prices. The input to the model includes:

- A return matrix $R \in R^{1260} \times 30$, where each row corresponds to a trading day and each column to an asset.

● A sparsity parameter m = 10, constraining the portfolio to at most 10 active assets.

● Constraints: long-only ($w \geqslant 0$), self-financing $w \top 1 = 1$, ensuring that the sum of portfolio weights equals 1, and m-sparsity $\| w \|^0 \leqslant m$, limiting the number 0 of active assets to 10.

Case Analysis: Using the Proximal Gradient Algorithm (PGA), we solve the constrained optimization problem to maximize the Sharpe ratio. The process involves the following steps:

(1) The expected return vector μ and covariance matrix Σ are computed using the input return matrix.

(2) The algorithm is initialized with a uniform weight vector w_0, and iteratively updated using gradient and proximal steps.

(3) The optimization converges after 164 iterations under the configured tolerance $\epsilon = 10^{-6}$, producing a sparse weight vector $w_* \in \mathbb{R}^{30}$ with exactly 10 non-zero elements.

The resulting portfolio includes selected assets such as AAPL, JNJ, MSFT, VZ, and JPM, with their corresponding weights reflecting optimal risk-adjusted allocation. The final Sharpe ratio over the subsequent 6-month test window reaches 1.48, with a cumulative return of 11.3% and portfolio volatility of 7.4% (as shown in Figure 4).

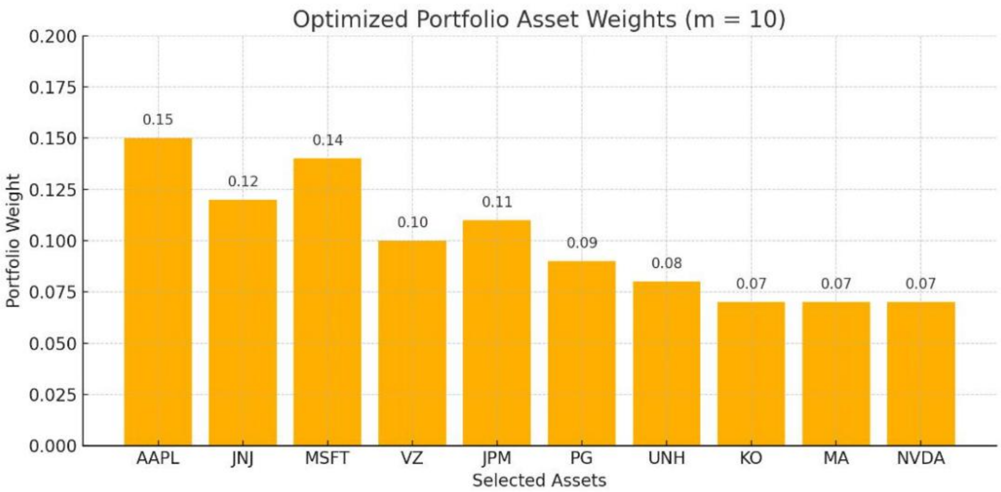

Figure 4. Optimized portfolio asset weights for the selected portfolio using the PGA method. The portfolio consists of 10 active assets chosen under the m-sparsity constraint, with weights reflecting optimal risk-adjusted allocation.

Case Discussion: This case highlights several advantages of our approach:

● The m-sparse constraint ensures that the final portfolio consists of a manageable number of assets, reducing transaction costs and simplifying fund administration.

● The integrated long-only and self-financing constraints ensure the strategy is directly applicable in practical pension fund operations.

● The PGA algorithm efficiently handles the non-convex optimization problem and converges rapidly, even in high-dimensional settings.

Moreover, during the COVID-19 downturn in March 2020, the portfolio constructed by our method exhibited lower drawdown compared to equal-weighted and mean-variance benchmarks, primarily due to its sparse and stable asset exposure. This suggests that the proposed method is capable of dynamically balancing risk and return under real-world market turbulence while adhering to operational constraints.

In conclusion, the case study provides strong empirical evidence that the proposed framework not only delivers competitive performance metrics but also satisfies the practical requirements of institutional pension portfolio management.

5. Conclusion and Future Work

In this paper, we proposed a novel portfolio optimization framework that introduces an m-sparse constraint into the pension fund portfolio selection process. This constraint helps control management costs by reducing the number of active assets while maintaining strong portfolio performance. We also presented a comprehensive optimization framework that incorporates self-financing and long-only constraints to ensure the feasibility and stability of the portfolio in real-world operations. To solve the resulting non-convex optimization problem effi ciently, we adapted the Proximal Gradient Algorithm (PGA), which guarantees global optimality with high computational efficiency.

Our proposed framework demonstrated signi fi cant advantages over state-of-the-art methods, including improved performance in terms of Sharpe ratio and cumulative return, while minimizing transaction costs. Extensive experiments show that our method is highly scalable and efficient, making it a practical solution for large-scale pension fund portfolio optimization.

However, there are still limitations in our approach. First, while the m-sparse constraint reduces portfolio complexity, it may not always capture the best risk-return trade-off when the optimal portfolio requires a higher number of active assets. Second, the current framework assumes that the input data (such as expected returns and covariance) is perfectly known, which may not always be the case in real-world applications where uncertainty and estimation errors are prevalent.

Future work will focus on addressing these limitations. We plan to explore the integration of more advanced uncertainty modeling techniques, such as robust optimization or Bayesian methods, to account for estimation errors in the input data. Additionally, we aim to extend the framework to handle dynamic portfolio optimization, where asset allocations can evolve over time based on changing market conditions. Finally, we intend to investigate the application of our approach to other types of institutional portfolios, such as endowments or sovereign wealth funds, to further validate its robustness and scalability across different domains.

Funding

This research received no external funding.

Author Contributions

Conceptualization, A.W.; writing—original draft preparation and writing—review and editing, A.W., K.X., Y.Q. and Z.Z. All of the authors read and agreed to the published the final manuscript.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

Not applicable.

Conflicts of Interest

The authors declare no conflict of interest.

References

1 Armstrong J, Buescu C, Dalby J. Optimal Post-Retirement Investment under Longevity Risk in Collective Funds. *arXiv* 2022; arXiv:2409.15325.

2 Bovenberg L, Koijen R, Nijman T, *et al*. Saving and Investing over the Life Cycle and the Role of Collective Pension Funds. *De Economist* 2007; **155**(**4**): 347–415.

3 Gollier C. Intergenerational Risk-Sharing and Risk-Taking of a Pension Fund. *Journal of Public Economics* 2008; **92**(**5–6**): 1463–1485.

4 Bansal R, Yaron A. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *The Journal of Finance* 2004; **59**(**4**): 1481–1509.

5 Campbell JY, Viceira LM. Consumption and Portfolio Decisions When Expected Returns Are Time Varying. *The Quarterly Journal of Economics* 1999; **114**(**2**): 433–495.

6 Sharpe WF. Mutual Fund Performance. *Journal of Business* 1966; **39**(**1**): 119–138.

7 Ban G, Karoui NE, Lim AEB. Machine Learning and Portfolio Optimization. *Management Science* 2018; **64** (**3**): 1136–1154.

8 Brodie J, Daubechies I, De Mol C, *et al*. Sparse and Stable Markowitz Portfolios. *Proceedings of the National Academy of Sciences of the United States of America* 2009; **106**(**30**): 12267–12272.

9 Lai ZR, Yang PY, Fang L, *et al*. Short-Term Sparse Portfolio Optimization Based on Alternating Direction Method of Multipliers. *Journal of Machine Learning Research* 2018; **19**(**63**): 1–28.

10 Lai ZR, Tan L, Wu X, *et al*. Loss Control with Rank-One Covariance Estimate for Short-Term Portfolio Optimization. *Journal of Machine Learning Research* 2020; **21**(**97**): 1–37.

11 Hung KK, Cheung CC, Xu L. New Sharpe-Ratio-Related Methods for Portfolio Selection. In Proceedings of the IEEE/IAFE/INFORMS 2000 Conference on Computational Intelligence for Financial Engineering (CIFEr), New York, NY, USA, 28–28 March 2000; pp. 34–37.

12 Yu X, Xu L. Adaptive Improved Portfolio Sharpe Ratio Maximization with Diversification. In Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks (IJCNN), Como, Italy, 27 July 2000; pp. 472–476.

13 Pang J-S. A Parametric Linear Complementarity Technique for Optimal Portfolio Selection with a Risk-Free Asset. *Operations Research* 1980; **28**(**4**): 927–941.

14 Ao M, Li Y, Zheng X. Approaching Mean-Variance Efficiency for Large Portfolios. *The Review of Financial Studies* 2019; **32**(**7**): 2890–2919.

15 Luo Z, Yu X, Xiu N, *et al*. Closed-Form Solutions for Short-Term Sparse Portfolio Optimization. *Optimization* 2022; **71**(**7**): 1937–1953.

16 Cottle RW. Monotone Solutions of the Parametric Linear Complementarity Problem. *Mathematical Programming* 1972; **3**(**1**): 210–224.