Computers, Materials & Continua

Tech Science Press

https://doi.org/10.32604/cmc.2026.083120

ARTICLE

A Novel Multi-Objective Quantum-Inspired Algorithm for Portfolio Optimization with Short-Selling in Real-World Market

Yun-Ting Lai, Ming-Ho Chang and Yao-Hsin Chou*

Department of Computer Science and Information Engineering, National Chi Nan University, Puli, Taiwan

*Corresponding Author: Yao-Hsin Chou. Email: yhchou@ncnu.edu.tw

Received: 29 March 2026; Accepted: 07 July 2026; Published: 13 August 2026

ABSTRACT: Portfolio optimization is inherently a multi-objective problem that aims to maximize expected return while minimizing investment risk, while also facing exponential growth in the search space and increasing market complexity. Existing multi-objective optimization approaches often struggle to balance convergence and diversity, particularly under realistic trading conditions such as short-selling. To address these challenges, this paper proposes a novel Multi-objective Quantum-inspired Tabu Search (MoQTS) framework for portfolio optimization with short-selling strategies. The proposed method incorporates a quantum-inspired superposition mechanism to enhance global exploration and introduces an entanglement-driven neighborhood search strategy that systematically generates structured local perturbations by modifying one or two asset-selection states. This mechanism enables effective exploration of the neighborhood of non-dominated solutions, thereby improving both convergence accuracy and solution diversity. In addition, a trend ratio (TR)-based evaluation model is adopted to jointly capture return and risk dynamics under real-world market fluctuations. Experiments are conducted on the U.S. stock market using Dow Jones Industrial Average (DJIA) data from 2013 to 2025. The proposed MoQTS is compared with several state-of-the-art multi-objective algorithms, including NSGA-II, MOEA/D, SMS-EMOA, and MOPSO. Experimental results demonstrate that MoQTS can obtain high-quality Pareto-optimal solutions with strong convergence and diversity performance. Results over multiple independent periods further support the effectiveness and robustness of MoQTS. In addition, computational cost analysis shows that MoQTS substantially reduces the number of evaluations and execution time compared with the comparison algorithms while maintaining high-quality Pareto-optimal solutions.

KEYWORDS: Multi-objective optimization; multi-objective quantum-inspired tabu search; real-world application; portfolio optimization; short-selling strategy; Pareto front

1 Introduction

Optimization techniques have become essential for addressing complex decision-making problems in real-world applications, including circuit synthesis, industrial optimization, and financial portfolio optimization. Among these, portfolio optimization is one of the most important problems in practice. According to Modern Portfolio Theory (MPT) [1], introduced by Harry Markowitz, the objective is to achieve an optimal balance between expected return and risk. In portfolio optimization, expected return and investment risk are the two most critical criteria. When formulated as a single-objective problem, traditional approaches typically yield a single solution representing a trade-off between return and risk. However, investors often exhibit diverse risk preferences, and a single solution may not fully capture such variability. Therefore, portfolio optimization is more appropriately formulated as a multi-objective optimization (MOO) problem, aiming to simultaneously maximize return and minimize risk.

In contrast to single objective optimization, MOO aims to identify a set of trade-off solutions, known as the Pareto-optimal set. These solutions form the Pareto front (PF). Constructing a well-distributed and convergent PF is crucial, as it provides investors with diverse investment choices under different risk-return preferences. However, identifying the PF in portfolio optimization is highly challenging due to the exponential growth of the search space and the complex interactions among assets. In the portfolio optimization problem, each stock in the market can exist in two possible states: selected or unselected. As a result, the number of possible portfolio combinations grows exponentially with the number of available assets, leading to high computational complexity. For example, when 30 stocks are considered, the search space contains 2^{30} possible portfolio combinations. Such a large search space makes it difficult for traditional optimization methods to efficiently identify high-quality solutions. To address this problem, metaheuristic algorithms have been widely adopted as effective approaches for solving complex optimization problems.

Metaheuristic algorithms, such as genetic algorithms (GA) [2], particle swarm optimization (PSO) [3], and differential evolution (DE) [4], are widely used for solving large-scale and complex optimization problems. To further improve search capability, various quantum-inspired metaheuristic algorithms have been proposed, including the quantum-inspired evolutionary algorithm (QEA) [5], quantum-inspired particle swarm optimization (QPSO) [6], quantum-inspired gravitational search algorithm (QGSA) [7], and quantum-inspired tabu search (QTS) [8].

However, MOO requires handling multiple conflicting objectives, which significantly enlarges the search space and increases problem complexity. As a result, achieving a good balance between convergence and diversity remains a challenging task in many real-world applications. Several algorithms, such as on-dominated sorting genetic algorithm II (NSGA-II) [9], multi-objective evolutionary algorithm based on decomposition (MOEA/D) [10], and S-metric selection evolutionary multi-objective optimization algorithm (SMS-EMOA) [11], have been developed for MOO problems. To better handle the complexity of MOO and improve the balance between convergence and diversity, this study proposes a Multi-objective Quantum-inspired Tabu Search (MoQTS) algorithm that incorporates quantum-inspired mechanisms to obtain a high-quality Pareto front.

MoQTS incorporates the concepts of superposition and entanglement to effectively explore the search space and obtain a high-quality PF. Initially, MoQTS leverages the quantum superposition mechanism to extensively explore the search space, followed by a targeted search to identify a diverse set of non-dominated solutions (NDS). Subsequently, MoQTS performs a local search based on the constructed PF. By leveraging quantum entanglement states, the algorithm establishes logical correlations between decision variables, allowing for a more refined exploitation of the neighborhood around NDS. This mechanism effectively guides the search toward higher-quality regions of the objective space, thereby enhancing the convergence and precision of the identified Pareto set.

Beyond its theoretical advancements, the practical significance of MoQTS lies in its adaptability to complex, real-world optimization challenges. To empirically demonstrate its utility, this paper evaluates the performance of MoQTS in dynamic financial market environments. Considering that stock markets are inherently volatile, investment strategies must remain robust across diverse market cycles instead of being confined to profiting solely from upward trends. Therefore, this study extends the application of MoQTS to short-selling transactions and presents the extension application of MoQTS to this problem. To evaluate the performance of the proposed method, the trend ratio (TR) [8] is adopted as the primary metric for the objective functions. Recent research [8] has demonstrated that TR can effectively assess portfolio performance across various markets. Therefore, the trend ratio is employed in this study to evaluate both return and risk.

The main contributions of this paper are summarized as follows:

- This study adopts a novel MoQTS algorithm for portfolio optimization, which integrates quantum-inspired superposition and an entanglement-move local search mechanism to improve both convergence performance and solution diversity.
- To the best of our knowledge, this paper is the first to apply MoQTS to portfolio optimization in a short-selling setting, thereby enabling more realistic portfolio optimization in dynamic financial environments.
- This paper introduces an entanglement-move local search mechanism combined with the trend ratio evaluation metric to improve the exploration of non-dominated solutions and achieve a more accurate and well-distributed Pareto front.

The remainder of this paper is organized as follows. Section 2 reviews the related work. Section 3 presents the proposed method. Section 4 reports the experimental results and analysis. Finally, Section 5 concludes the paper.

2 Related Work

The main objective of portfolio optimization is to maximize the return and minimize risk. MPT [1] is widely used to identify portfolio performance. Under the MPT model, investment risk is quantified using the variance and covariance, and the return is defined as the expected return. The Sharpe ratio [12] is one of the most widely used indicators for evaluating risk-adjusted portfolio performance, and its risk measurement is closely related to the MPT framework. However, this approach often identifies overly conservative portfolios with stagnant growth, and is thus limited in its ability to capitalize on significant market opportunities during periods of high market volatility. Thus, this paper adopts the trend ratio to assess the portfolio return and risk more accurately. In [8], the TR effectively quantifies portfolio performance by identifying sustained upward momentum, thereby demonstrating a superior capacity for evaluating investment quality.

In practical applications, optimization problems often involve multiple conflicting conditions that must be addressed simultaneously. Recent studies have demonstrated the practical applicability of multi-objective and intelligent optimization frameworks in real-world scenarios, including portfolio optimization [13]. Portfolio optimization is a representative example, since investors generally seek to maximize expected return while minimizing investment risk. These two objectives are inherently conflicting, as portfolios with higher expected returns are often associated with higher levels of risk. Consequently, portfolio optimization is naturally formulated as a multi-objective optimization problem, where the goal is to identify a set of non-dominated solutions (NDS) that represent optimal trade-offs between return and risk. These solutions collectively form the efficient frontier, providing investors with multiple portfolio alternatives according to their risk preferences and investment goals. Recent studies have increasingly applied multi-objective optimization methods to portfolio selection problems. For example, Ref. [14] proposed a multi-objective minimax-based model for ESG portfolio construction, while Ref. [15] developed an integrated MCDM and multi-objective optimization framework for socially responsible portfolio selection. In [16], the paper proposed a multi-objective genetic algorithm integrated with the fuzzy method to solve the cardinality-constrained portfolio optimization problem.

In recent years, quantum-inspired optimization algorithms have gained attention for improving search efficiency in complex optimization problems. Several studies have introduced quantum-inspired methods for solving complex optimization problems. MoQTS is developed based on QTS in [8], the TR effectively quantifies portfolio performance by identifying sustained upward momentum, thereby demonstrating a superior capacity for evaluating investment quality. In [8], the TR effectively quantifies portfolio performance by identifying sustained upward momentum, thereby demonstrating a superior capacity for evaluating investment quality. QTS demonstrated superior performance in constructing high performance portfolios across multiple stock markets. Therefore, MoQTS inherits the advantages of QTS for solving multi-objective optimization problems. MoQTS has proven to be a highly effective method for constructing high-quality Pareto fronts under long-selling scenarios. However, in financial markets, both upward and downward trends coexist, and the ability to generate returns during declining market conditions is equally important. Accordingly, this study extends MoQTS to portfolio optimization with short-selling strategies to provide a more realistic portfolio optimization framework for dynamic financial environments.

3 Proposed Method

This section presents the workflow of MoQTS approach for portfolio optimization with short-selling. Fig. 1 illustrates its overall framework. The objective is to identify a diverse set of portfolios that achieve maximum returns across different risk levels, thereby facilitating the construction of the Pareto front.

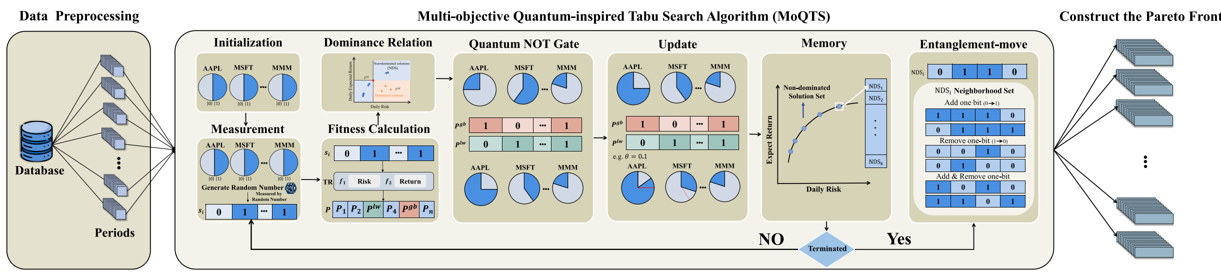

Figure 1: The structure of the proposed MoQTS method.

3.1 Portfolio Representation with Short-Selling Strategy

In this paper, the proposed framework is developed under a short-selling strategy, where all selected assets correspond to short positions in the portfolio. Each asset is represented by a binary decision variable x_i, where $x_i = 1$ indicates that asset i is selected for short-selling and $x_i = 0$ otherwise. The variable n denotes the number of assets in the stock pool and k denotes the number of selected assets. The number of selected assets is constrained by $1 \leq k \leq n$. Since an equal-weight allocation scheme is adopted, each selected asset receives the same short weight, denoted by $w_i = 1/k$. Therefore, the total weight of the short-selling portfolio is normalized to one, i.e., $\sum_{i=1}^{n} |w_i| = 1$, which implies that the portfolio is constructed under a no-leverage setting. In this paper, short-selling is assumed to be feasible for all assets in the stock pool, and no explicit upper bound is imposed on individual positions due to the equal-weight structure. Furthermore, the proposed framework considers a trading environment based on online brokerage platforms, where transaction fees and borrowing costs in non-leveraged account are minimal or not explicitly charged. Therefore, these costs are not incorporated into this study. In summary, the short-selling setting in this paper is based on equal-weight short positions, assumes no leverage, and does not incorporate transaction fees, borrowing costs, or margin-related costs.

Additionally, to facilitate a more intuitive analysis of price trends and returns, the negative return pattern associated with short-selling is transformed into a positive return representation. The transformed values are then used to compute the trend ratio, which serves as the basis for evaluating portfolio performance. In the proposed multi-objective optimization framework, the return and risk derived from the trend ratio are used as two optimization objectives. Specifically, MoQTS searches the solution space to identify promising Pareto-front solutions by maximizing the trend-based return while minimizing the corresponding trend-based risk.

After the main search process, an entanglement-based local search is further applied to the obtained non-dominated solutions to explore neighboring portfolios and identify solutions with better performance.

The corresponding decision variables and constraints are formulated as follows:

$$x_i = \begin{cases} 1, & \text{if asset } i \text{ is selected for short-selling,} \\ 0, & \text{otherwise,} \end{cases} \qquad i = 1, 2, \ldots, n \tag{1}$$

$$k = \sum_{i=1}^{n} x_i \tag{2}$$

$$1 \leq k \leq n \tag{3}$$

$$w_i = \begin{cases} \frac{1}{k}, & \text{if } x_i = 1 \\ 0, & \text{if } x_i = 0 \end{cases} \tag{4}$$

$$\sum_{i=1}^{n} |w_i| = 1 \tag{5}$$

3.2 Trend Ratio

This paper utilizes the trend ratio to evaluate the two main objectives, return and risk. The TR is calculated based on the funds standardization (FS) [17]. The FS converts stock price movements directly into funds fluctuations. By representing price changes in terms of capital variation, investors can more accurately understand the actual changes in investment funds during the trading period. The TR is calculated according to the volatility of FS to assess the investment's expected return (ER) and daily risk. Eq. (6) defines the TR, which evaluates portfolio performance by jointly considering return and risk. The trend ratio is used to derive the return and risk measures associated with each portfolio for the multi-objective optimization framework. The TR can be interpreted as the return generated per unit of risk, providing an intuitive measure of investment efficiency. When incorporated into a multi-objective optimization framework, the TR enables the identification of portfolios that achieve higher returns under different risk levels, thereby facilitating the construction of efficient portfolios along the Pareto frontier. The derived return and risk values are then treated as two separate objectives in the multi-objective optimization framework. Accordingly, the portfolio optimization problem is formulated with two objective functions: the maximization objective f_1, which represents the expected return, and the minimization objective f_2 represents the investment risk. Eq. (7) defines the expected return of the solution $\mathcal{P}$, while Eq. (8) defines the corresponding investment risk. The variables used in these equations are defined as follows, where x_j is the jth day, d denotes the total transaction day, C represents the initial funds. The two objective functions are defined as $f_1, f_2 : \mathbb{R}^n \to \mathbb{R}$. The $Y_{\mathcal{P}}$ is the regression line of the portfolio fluctuation, whose intercept is fixed at the initial fund, is formulated in Eq. (9). In the equation, $m_{\mathcal{P}}$ represents the slope of the trend line, which corresponds to the expected daily return. Eq. (10) is the actual daily portfolio FS of jth day ($y_{\mathcal{P},j}$), where c_i is allocated fund, $\rho_{i,j}$ is the jth day stock price for stock i, l_i indicates the number of stock i holding share $(l_i > 0)$, $\delta_{i,j}$ is the transaction cost for stock i at jth day, and n is the number of stocks considered in the portfolio.

$$\text{Trend Ratio (TR)} = \frac{Expected\ Return\ (ER)}{Daily\ Risk} \tag{6}$$

$$\max\ f_1(\mathcal{P}) : \text{Expected Return} = \frac{\sum_{j=1}^{d} (x_j y_{\mathcal{P},j} - x_j C)}{\sum_{j=1}^{d} (x_j^2)} \tag{7}$$

$$\min\ f_2(\mathcal{P}) : \text{Daily Risk} = \sqrt{\frac{\sum_{j=1}^{d}(Y_{\mathcal{P},j} - y_{\mathcal{P},j})^2}{d}} \tag{8}$$

$$Y_{\mathcal{P},j} = m_{\mathcal{P}} x_j + C \tag{9}$$

$$y_{\mathcal{P},j} = \sum_{i=1}^{n}(\rho_{i,j}\ell_i - \delta_{i,j}\ell_i) + C \tag{10}$$

3.3 *Multi-Objective Quantum-Inspired Tabu Search Algorithm (MoQTS)*

This paper extends the MoQTS framework to portfolio optimization under short-selling strategy in the real-world market. MoQTS incorporates the superposition mechanism with the Q-gate to accelerate convergence, while the entanglement move expands the non-dominated solution set, thereby facilitating the discovery of the Pareto front. In the proposed MoQTS framework, the tabu mechanism is realized through a quantum-inspired probabilistic update process. Specifically, promising solutions are progressively reinforced during the search process, while inferior solutions are gradually suppressed through probability adjustment, thereby guiding the search toward high-quality regions of the solution space. Algorithm 1 presents the procedure of MoQTS.

Algorithm 1: Multi-objective quantum-inspired tabu search algorithm

Input: Historical stock price data
Initialize superposition state Q matrix
for $t = 1$ to Gen **do**
    Generate a population $\mathcal{P}$ by measuring Q
    Evaluate objective values of $\mathcal{P}$
    Applying a quantum NOT gate (Eq. (13))
    Update by global best ($\mathcal{P}^{gb}$) and local worst solutions ($\mathcal{P}^{lw}$) (Eq. (15))
    Memory mechanism
end for
while S^* keeps improving **do**
    for all $S \in S^*$ **do**
        Generate entangled neighbor $\mathcal{N}_{\text{entangled}}$ (Eq. (18))
        if $\mathcal{P}^{nb}$ improves S^* **then**
            Update S^*
        end if
    end for
end while
return S^*

3.3.1 Superposition Encoding

MoQTS incorporates quantum superposition as its encoding mechanism, as formulated in Eq. (11), where the basis states $|0\rangle$ and $|1\rangle$ represent the exclusion and inclusion of a stock, respectively. α and β are the amplitudes of the corresponding states $|0\rangle$ and $|1\rangle$. The squared magnitudes of the amplitudes represent the probabilities of observing the corresponding states. In the initialization, $\alpha = \beta = \frac{1}{\sqrt{2}}$. The selection probability of stock i is given by β_i^2, which determines the value of q_i used to construct the Q-matrix. A solution $\mathcal{P}$ is defined as the vector $[s_1, s_2, \ldots, s_n]$, where each element $s_i \in \{0, 1\}$ denotes the selection status of the ith stock, and n represents the total number of stocks in the search space.

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad \text{where } ||\alpha||^2 + ||\beta||^2 = 1 \tag{11}$$

3.3.2 Measurement

In this process, MoQTS generates a random number $r_i \in [0, 1]$ for each stock $i \in \{1, 2, \ldots, n\}$. Subsequently, the measurement process compares r_i with the quantum probability q_i to determine the binary state s_i, as formulated in Eq. (12). Specifically, s_i takes the value of 1 provided that q_i exceeds r_i, and 0 otherwise. The binary variables $s_i = 1$ and $s_i = 0$ represent the selection and non-selection state of the ith stock for the portfolio, respectively. Afterwards, the MoQTS employs the trend ratio (Eq. (6)) to evaluate the performance of each portfolio combination. Based on the resulting fitness values, the algorithm identifies the global best solution ($\mathcal{P}^{gb}$) and the local worst solution ($\mathcal{P}^{lw}$).

$$s_i = \begin{cases} 1, & \text{if } q_i \geq r_i \\ 0, & \text{if } q_i < r_i \end{cases}, \quad s_i \in \{0, 1\} \tag{12}$$

3.3.3 Quantum Not Gate

MoQTS incorporates the quantum NOT gate (Pauli-X gate) to enhance its ability to escape local optima, as formulated in Eq. (13). The quantum NOT gate swaps the probability amplitudes of the basis states, allowing the search particle to escape local optima and move toward better solutions. In this process, the gate operation is determined according to the stock selection states of the global best (s_i^{gb}) and the local worst (s_i^{lw}), together with the probability q_i. Under the conditions defined in Eq. (14), q_i is updated to $1 - q_i$, thereby reversing the selection tendency of the corresponding stock.

$$\alpha|0\rangle + \beta|1\rangle \xrightarrow{\text{Pauli-}X\text{ gate}} \beta|0\rangle + \alpha|1\rangle \tag{13}$$

$$\text{if } \begin{cases} s_i^{gb} = 1 \,\&\, s_i^{lw} = 0 \,\&\, q_i < 0.5 \\ s_i^{gb} = 0 \,\&\, s_i^{lw} = 1 \,\&\, q_i > 0.5 \end{cases}, \quad q_i = 1 - q_i \tag{14}$$

3.3.4 Update within Qmatrix

The update process utilizes the parameter θ to adjust selection probabilities, aiming to guide the population toward the optimal solution while avoiding local worst outcomes, as demonstrated in Eq. (15). The probability q_i within the Q-matrix is updated according to the selection states of the highest TR portfolio ($\mathcal{P}^{gb}$) and the lowest TR portfolio in the current iteration ($\mathcal{P}^{lw}$). The update mechanism naturally drives the value of q_i toward boundary values. When $q_i + \theta > 1$ or $q_i - \theta < 0$, the probability reaches a saturated state, indicating that the corresponding decision variable converges to a deterministic selection, where the asset is consistently selected or consistently not selected.

$$q_i = \begin{cases} q_i + \theta, & \text{if } s_i^{gb} = 1 \,\&\, s_i^{lw} = 0, \\ q_i - \theta, & \text{if } s_i^{gb} = 0 \,\&\, s_i^{lw} = 1, \\ q_i, & \text{if } s_i^{gb} = s_i^{lw}. \end{cases} \tag{15}$$

3.3.5 Memory Mechanism

Before introducing the memory mechanism, the dominance relationship and the concept of Pareto optimality are defined as follows.

Definition 1: A solution S_1 is said to dominate another solution S_2, denoted by $S_1 \succ S_2$, if and only if the following conditions are satisfied: (a) $\forall i \in \{1, 2\},\ f_i(S_1) \le f_i(S_2)$, and (b) $\exists j \in \{1, 2\},\ f_j(S_1) < f_j(S_2)$.

Definition 2: Two solutions S_1 and S_2 are said to be non-dominated with respect to each other if neither $S_1 \succ S_2$ nor $S_2 \succ S_1$ holds.

Definition 3: The set of non-dominated solutions is referred to as the Pareto optimal set, denoted by $S^* = \{S_1, S_2, \ldots, S_z\}$, where $|S^*|$ denotes the number of solutions in the set. All solutions in S^* are not dominated by any other feasible solution in the search space.

Following Eq. (16), the non-dominated solution set S^* is iteratively refined to facilitate the convergence of MoQTS toward the true PF, where S denotes an existing solution in the archive S^*. The archive functions as a memory mechanism by retaining historically discovered non-dominated solutions, allowing the algorithm to utilize previous search information to guide subsequent exploration toward the true Pareto front. As the algorithm continuously updates the solution set during the search process, dominated solutions are removed, and only non-dominated solutions are retained in the archive, ensuring the quality of the approximation. To fully explore all potential Pareto-optimal solutions, no explicit size limit is imposed on the archive. All non-dominated solutions encountered during the search process are preserved to maintain solution diversity and avoid the premature loss of potentially valuable candidates. In addition, duplicate solutions are avoided during the update process to prevent unnecessary growth of the archive.

$$S^* \leftarrow \begin{cases} S^* \cup \{\mathcal{P}\} \setminus \{S \in S^* \mid \mathcal{P} \text{ dominates } S\}, & \text{if } \nexists S \in S^* \text{ such that } S \text{ dominates } \mathcal{P}, \\ S^*, & \text{otherwise.} \end{cases} \tag{16}$$

3.3.6 Entanglement-Move Local Search

Local search is incorporated to further enhance the performance of the proposed algorithm. After the Pareto front is constructed, an entanglement-inspired local search mechanism is applied to further refine the obtained non-dominated solutions. This mechanism is inspired by the concept of entanglement in quantum mechanics. In the proposed method, this concept is translated into a heuristic search strategy for portfolio optimization: the relationships among stocks are represented through the entanglement state, and neighboring portfolio solutions are explored by jointly adjusting the selection states of one or two related stocks. Through this mechanism, the local search can investigate potentially improved portfolio structures around the current non-dominated solutions. In this study, the concept of entanglement is adopted in a quantum-inspired heuristic sense to facilitate solution exploration, without involving physical quantum entanglement or quantum hardware implementation.

By exploring the neighborhood of candidate solutions, the local search mechanism refines current solutions and improves their quality, enabling the algorithm to exploit promising regions of the search space more effectively. In multi-objective optimization problems, this mechanism also plays an important role in improving both the convergence and diversity of the solution set. By refining NDS and exploring their neighboring regions, the algorithm can identify additional trade-off solutions and expand the coverage of the S^*. Consequently, the obtained solution set can better approximate the true PF while maintaining a well-distributed set of solutions. MoQTS simulates the characteristics of quantum entanglement states to capture the relationships among stocks within portfolios in S^*. Through this mechanism, the algorithm can efficiently expand the S^* and identify additional promising NDS, as described in Eq. (17). The equation characterizes the quantum correlations between two portfolios and is used to model the relationships among stocks in Pareto-optimal portfolios. By utilizing the properties of quantum entanglement, the proposed method considers the interactions among stocks when constructing candidate portfolios. In particular, the entangled states simulate the portfolio configurations in which one stock is either added to or removed from the portfolio. This mechanism enables the algorithm to explore neighboring portfolio structures while preserving the underlying relationships among assets, thereby facilitating the identification of portfolios that either increase returns or reduce risk.

The entangled neighborhood is constructed as defined in Eq. (18). The entangled neighborhood $\mathcal{N}_{\text{entangled}}(S^*)$ is defined as the union of neighboring solutions generated from each non-dominated solution in S^*, where each neighboring solution differs from its corresponding solution by one or two asset selection states. The number of neighboring solutions generated during the entanglement move is adaptively determined according to the number of non-dominated solutions maintained in S^* and the asset selection states of each solution. For each non-dominated solution, neighboring portfolios are generated by modifying one or two asset selection states according to the entanglement operation. Since the entanglement move local search continuously updates S^*, newly generated non-dominated solutions may further produce additional neighboring solutions. This iterative process continues until no additional non-dominated solutions can be obtained. Therefore, the total number of generated neighbors is adaptively determined by the current Pareto-front approximation and the portfolio configurations encountered during the search process.

$$|\psi^+\rangle = \frac{1}{\sqrt{2}}\left(|01\rangle + |10\rangle\right), \quad |\phi^+\rangle = \frac{1}{\sqrt{2}}\left(|00\rangle + |11\rangle\right) \tag{17}$$

$$\mathcal{N}_{\text{entangled}}(S^*) = \bigcup_{S\in S^*}\left\{S \,\middle|\, \sum_{i=1}^{n}\left(s_i^S \oplus s_i^{\mathcal{P}^{nb}}\right) \in \{1,2\}\right\} \tag{18}$$

To establish a clear mapping between the entanglement states and portfolio modifications, the neighborhood solutions are constructed by explicitly modifying the binary decision vector. Given a portfolio solution $\mathcal{P} = [s_1, s_2, \dots, s_n]$, the entanglement move generates neighboring solutions by applying bit-wise operations to the binary variables. The construction of the entangled neighborhood follows two deterministic rules:

1. **Single-bit modification:** One asset selection state is flipped:

   $$s_i \leftarrow 1 - s_i, \tag{19}$$

   which corresponds to adding or removing one asset from the portfolio.
2. **Two-bit modification:** One unselected asset and one selected asset are simultaneously modified:

   $$s_i \leftarrow 1, \quad s_j \leftarrow 0, \tag{20}$$

   where $s_i = 0$ and $s_j = 1$. This operation maintains the number of selected assets while adjusting the portfolio composition.

The entanglement operator determines which bits are selected for modification, and the resulting solutions constitute the entangled neighborhood $\mathcal{N}_{\text{entangled}}$. Through this mapping, the quantum-inspired entanglement mechanism is translated into a deterministic neighborhood generation process for binary portfolio optimization.

4 Experimental Result

This section presents a comprehensive empirical evaluation of portfolio optimization with short selling under different market conditions in the U.S. equity market. The analysis focuses on convergence quality, computational efficiency, and practical investment effectiveness from both optimization and financial perspectives. The following subsections present the experimental environment, performance indicators, sensitivity analysis, comparative results with classical multi-objective algorithms, computational cost analysis, ablation study, and representative portfolio performance.

4.1 Experimental Environment

The investment target consists of the 30 constituents of the Dow Jones Industrial Average (DJIA). The experimental dataset is collected from Yahoo Finance, and the investment period spans monthly intervals from 2013 to 2025, yielding a total of 157 periods. In all experiments, the initial capital of every portfolio is set to 10 million. The proposed method is compared with four classical multi-objective algorithms: NSGA-II, MOEA/D, SMS-EMOA, and MOPSO. For MoQTS, the population size is set to 10, and 30 independent runs are conducted with 3,600 iterations per run. The key parameter θ in MoQTS is set to 0.0007, while the parameter settings of the other algorithms are summarized in Table 1. All algorithms are implemented in C++ and executed on a system with an Intel Core i9-14900K processor and 32 GB of RAM.

Table 1: Parameter settings of the comparison algorithms.

Algorithm	Main Parameters	Additional Parameters
NSGA-II	$P = 100, I = 1000, R = 30, N = 30$	$p_c = 1.0, p_m = 1/N$
SMS-EMOA	$P = 100, I = 1000, R = 30, N = 30$	$p_c = 0.9, p_m = 1/N$
MOEA/D	$P = 100, I = 1000, R = 30, N = 30$	$p_c = 1.0, p_m = 1/N, T = 10, H = 10$
MOPSO	$P = 100, I = 1000, R = 30, N = 30$	$W = 0.4$, share_fitness = 10, archive = 100, grid = 30

Note: P denotes population size, I denotes the maximum number of iterations, R denotes the number of independent runs, and N denotes the number of candidate stocks. In this study, $N = 30$, so $p_m = 1/N$ corresponds to 1/30.

4.2 Performance Indicators

In multi-objective optimization, two important criteria are usually considered: convergence and diversity. Convergence evaluates how close the obtained non-dominated solutions are to the true Pareto front, whereas diversity measures the distribution of the obtained solutions. This study uses four indicators to evaluate algorithm performance: Generational Distance (GD), Inverted Generational Distance (IGD), Hypervolume (HV), and Hypervolume rate (HV rate).

In real-world optimization problems, the solution space is often extremely large, making it difficult to obtain the true Pareto front. In the considered portfolio optimization problem, even under the equal capital allocation setting with 30 candidate stocks, the stock selection space contains at least 2^{30}, i.e., 1,073,741,824, possible portfolio combinations for each investment period. If the problem is further extended to include capital allocation decisions, the solution space becomes substantially larger. In addition, considering multiple investment periods further increases the computational complexity. Under such circumstances, obtaining the true Pareto front through exhaustive enumeration or constructing a stronger independent approximation becomes computationally impractical. Therefore, for each investment period, this paper constructs an empirical reference Pareto front, denoted as PF^*. The PF^* is not a true PF, which is based on the best-known non-dominated solutions obtained by different independent approximation methods. Due to the difficulty of obtaining the exact true Pareto front in such a large solution space, PF^* is used as a substitute for the true Pareto front. This empirical reference front provides a fair and consistent benchmark for comparing the performance of all algorithms. Specifically, this paper first constructs a combined non-dominated set by integrating the non-dominated solutions produced by all compared algorithms in that period. The PF^* is then generated by extracting the final non-dominated set from this combined solution set. This construction is not intended to represent the exact true Pareto front. Instead, it serves as a practical best-known empirical reference front for large-scale real-world portfolio optimization problems where the true Pareto front is unavailable. By combining the best non-dominated solutions found by different independent approximation methods, the constructed PF^* provides a feasible and consistent benchmark for performance evaluation. More importantly, using the same PF^* for all algorithms ensures that the convergence- and diversity-related indicators are computed under a unified and fair comparison basis. Therefore, the constructed PF^* provides a common empirical benchmark for evaluating the relative performance of all compared algorithms.

4.2.1 Generational Distance (GD)

GD is an indicator used to evaluate convergence performance. It measures the distance between the obtained non-dominated solution set and the reference Pareto front. A smaller GD value indicates that the obtained solutions are closer to the reference Pareto front and therefore exhibit better convergence quality.

Before calculating GD, risk and expected return are normalized to avoid scale bias between the two objectives. Let r denote risk and e denote expected return. The normalized risk and expected return are defined as

$$\tilde{r} = \frac{r - r_{\min}}{1.01 r_{\max} - r_{\min}}, \qquad \tilde{e} = \frac{e - 0.99 e_{\min}}{e_{\max} - 0.99 e_{\min}},$$

where $r_{\min}$, $r_{\max}$, $e_{\min}$, and $e_{\max}$ are obtained from all compared solution sets and the reference Pareto front within the same interval. The factors 1.01 and 0.99 are used to introduce a small margin in the normalization range. Specifically, 1.01 slightly enlarges the upper bound of risk, whereas 0.99 slightly decreases the lower bound of expected return. This treatment prevents boundary solutions from coinciding exactly with the normalization limits and makes the distance calculation more numerically stable. The Euclidean distance used in GD is then computed in the normalized objective space $(\tilde{e}, \tilde{r})$.

$$GD = \frac{1}{|S^*|} \left(\sum_{i=1}^{|S^*|} \left[\min_{p \in PF^*} d\left(\tilde{S}_i, \tilde{p}\right) \right]^2 \right)^{1/2} \tag{21}$$

In Eq. (21), S^* denotes the obtained non-dominated solution set, $|S^*|$ is the number of solutions in S^*, PF^* denotes the reference Pareto front, $\tilde{S}_i$ is the normalized objective vector of the ith solution in S^*, and $\tilde{p}$ is the normalized objective vector of a reference point $p \in PF^*$. The term $d(\tilde{S}_i, \tilde{p})$ denotes the Euclidean distance between $\tilde{S}_i$ and $\tilde{p}$. Therefore, $\min_{p \in PF^*} d(\tilde{S}_i, \tilde{p})$ represents the shortest distance from an obtained solution to the reference Pareto front in the normalized objective space.

4.2.2 Inverted Generational Distance (IGD)

IGD is similar to GD and is used to evaluate both convergence and diversity performance. It measures the distance from the reference Pareto front to the obtained non-dominated solution set. A smaller IGD value indicates that the obtained solutions are closer to the reference front and provide better coverage of the objective space.

The same normalization procedure used for GD is also applied before calculating IGD. Specifically, the same small margins introduced by the factors 1.01 and 0.99 are retained so that the normalized objective space remains consistent across GD and IGD calculations. The Euclidean distance is therefore computed in the normalized objective space $(\tilde{e}, \tilde{r})$, where $\tilde{e}$ and $\tilde{r}$ denote the normalized expected return and normalized risk, respectively.

$$IGD = \frac{1}{|PF^*|}\left(\sum_{j=1}^{|PF^*|}\left[\min_{S\in S^*} d\left(\tilde{p}_j, \tilde{S}\right)\right]^2\right)^{1/2} \tag{22}$$

In Eq. (22), PF^* denotes the reference Pareto front, $|PF^*|$ is the number of reference points in PF^*, $\tilde{p}_j$ is the normalized objective vector of the j^{th} reference point in PF^*, and $\tilde{S}$ denotes the normalized objective vector of a solution $S \in S^*$. The term $d(\tilde{p}_j, \tilde{S})$ denotes the Euclidean distance between $\tilde{p}_j$ and $\tilde{S}$. Therefore, $\min_{S\in S^*} d(\tilde{p}_j, \tilde{S})$ represents the shortest distance from a reference point to the obtained solution set in the normalized objective space.

4.2.3 Hypervolume (HV)

HV is an indicator used to evaluate both convergence and diversity performance by measuring the volume in the objective space dominated by the obtained non-dominated solution set with respect to a reference point. A larger HV value indicates that the solution set is closer to the reference Pareto front and covers a wider region of the objective space.

$$HV(S^*) = \lambda\left(\bigcup_{S\in S^*}[f_1(S), R_1] \times [f_2(S), R_2] \times \cdots \times [f_N(S), R_N]\right) \tag{23}$$

In Eq. (23), S^* denotes the obtained non-dominated solution set, N is the number of objectives, $f_k(S)$ denotes the value of the kth objective for solution S, and R is a predefined reference point corresponding to the worst objective values in the normalized objective space. Before computing HV, the objective values are normalized to remove scale differences between return and risk. In our implementation, the normalization uses $refpoint_{risk} = MaxRisk$ and $refpoint_{return} = MinReturn$, with normalization denominators $normal_{risk} = refpoint_{risk} - MinRisk = MaxRisk - MinRisk$ and $normal_{return} = MaxReturn - refpoint_{return} = MaxReturn - MinReturn$. After normalization, HV is computed by using $(MaxRisk,\ MinReturn)$ as the reference point. Here, $\lambda(\cdot)$ denotes the measure of the covered region. Therefore, $HV(S^*)$ represents the size of the objective-space region dominated by the obtained solution set with respect to the reference point. As shown in Eq. (24), the HV rate measures the relative hypervolume coverage of the obtained solution set. It is defined as the ratio of the hypervolume of the non-dominated solution set generated by an algorithm to the hypervolume of the reference Pareto front. A larger HV rate indicates that the obtained solution set covers a larger portion of the reference Pareto front.

$$HVrate(S^*) = \frac{HV(S^*)}{HV(PF^*)} \tag{24}$$

4.3 Sensitivity Analysis

This subsection analyzes the sensitivity of MoQTS to the parameter θ. The other parameters of MoQTS and the comparison algorithms follow the settings described in the experimental environment. To examine the influence of θ, a two-stage tuning procedure is conducted. In the first stage, a coarse search is performed by testing five values of θ from 0.1 to 0.00001, decreasing by a factor of 10 each time. For each setting, the non-dominated solutions obtained from all tested values are combined to construct a reference co-front, and the performance is evaluated using HV, HV rate, GD, and IGD. The results are shown in Table 2.

Table 2: Sensitivity analysis of the parameter θ in the coarse search.

θ	HV	HV Rate (%)	GD	IGD
0.1	0.624177	98.6851	1.1175×10^{-4}	5.6287×10^{-4}
0.01	0.624183	98.6858	1.0869×10^{-4}	3.4270×10^{-4}
0.001	**0.624223**	**98.6915**	$\mathbf{1.9196 \times 10^{-5}}$	$\mathbf{1.7258 \times 10^{-5}}$
0.0001	0.624033	98.0305	1.2127×10^{-4}	1.5325×10^{-3}
0.00001	0.624032	98.0305	9.8015×10^{-5}	2.4745×10^{-3}

Note: Bold values indicate the best result for each performance metric.

The coarse search results indicate that $\theta = 0.001$ achieves the best overall performance among the tested values, especially in terms of HV, HV rate, GD, and IGD. Based on this observation, a fine search is further conducted within the interval from 0.0001 to 0.001, increasing θ by 0.0001 each time. Similarly, the non-dominated solutions obtained from the tested settings are combined to construct a reference co-front, and the four performance indicators are computed. The results are reported in Table 3.

Table 3: Sensitivity analysis of the parameter θ.

θ	HV	HV Rate	GD	IGD
0.0001	0.624191	98.0303	1.2315×10^{-4}	1.5476×10^{-3}
0.0002	0.624368	98.6895	5.8796×10^{-5}	1.8504×10^{-4}
0.0003	0.624386	98.6921	2.2066×10^{-5}	4.8720×10^{-5}
0.0004	0.624381	98.6912	2.7904×10^{-5}	3.4825×10^{-5}
0.0005	0.624383	98.6915	1.6641×10^{-5}	2.3499×10^{-5}
0.0006	0.624384	98.6916	1.5123×10^{-5}	$\mathbf{2.2755 \times 10^{-5}}$
0.0007	**0.624388**	**98.6924**	$\mathbf{1.0163 \times 10^{-5}}$	2.6910×10^{-5}
0.0008	0.624383	98.6916	1.7773×10^{-5}	2.5930×10^{-5}
0.0009	0.624381	98.6912	1.4911×10^{-5}	3.0855×10^{-5}
0.0010	0.624382	98.6913	2.2111×10^{-5}	3.3015×10^{-5}

Note: Bold values indicate the best result for each performance metric.

The fine search results show that $\theta = 0.0006$ and $\theta = 0.0007$ both achieve strong performance. Although $\theta = 0.0006$ obtains the best IGD value, $\theta = 0.0007$ achieves the best HV, HV rate, and GD values. Therefore, $\theta = 0.0007$ is selected as the θ parameter for the subsequent experiments in this study.

4.4 Comparing MoQTS with Classical Multi-Objective Algorithms

This subsection presents the comparative performance of MoQTS among NSGA-II, MOEA/D, SMS-EMOA, and MOPSO based on the four indicators introduced above. Overall, MoQTS shows stable and strong behavior across all metrics, and the experimental results indicate that MoQTS consistently outperforms the other compared algorithms. These results suggest that the quantum-inspired search and memory mechanisms contribute to robust solution quality over repeated runs.

For convergence-related indicators, the GD and IGD results show that MoQTS remains the closest to the reference Pareto front among the compared algorithms. As shown in Fig. 2, MoQTS performs better than the other classical algorithms overall, with NSGA-II being the closest competitor on both GD and IGD.

Compared with SMS-EMOA, MOEA/D, and especially MOPSO, the advantage of MoQTS becomes more clear. The improvement reaches up to 5,897,298.60% for GD and 3,136,896.46% for IGD.

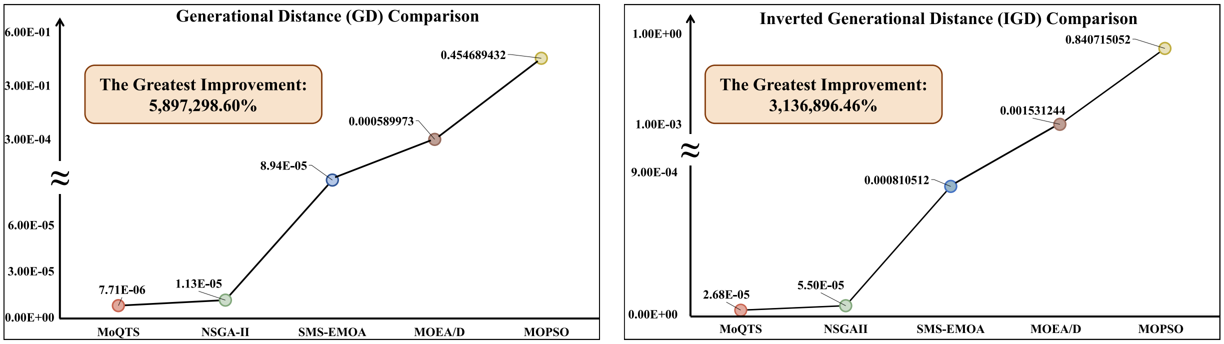

Figure 2: Comparison results of MoQTS and classical algorithms for GD (**left**) and IGD (**right**).

For coverage-related indicators, HV and HV rate further support the effectiveness of MoQTS in obtaining broad and high-quality dominated regions in the objective space, as shown in Fig. 3. The results indicate that MoQTS achieves superior performance compared with other classical algorithms, including NSGA-II, SMS-EMOA, and MOEA/D, on both HV and HV rate. When compared with MOPSO, the improvement becomes even more significant, reaching up to 74.29% for HV and 84.93% for HV rate.

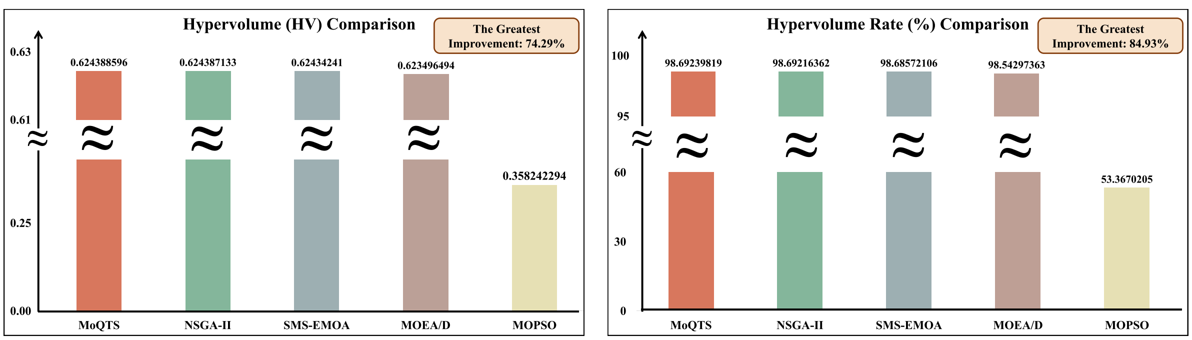

Figure 3: Comparison results of MoQTS and classical algorithms for HV (**left**) and HV rate (**right**).

To further compare the distribution quality of non-dominated solutions, Fig. 4 presents the PF comparisons of the five algorithms in global and zoomed-in views. In the left panel (global view), MoQTS spans a broader portion of the Pareto front and reaches both lower-risk and higher-return regions than the other algorithms, indicating a more favorable overall trade-off distribution. In the right panel (zoomed-in view), MoQTS maintains a denser set of non-dominated solutions across the low- to medium-risk region, whereas MOPSO is generally located below the leading front and NSGA-II and SMS-EMOA cover a smaller range. MOEA/D remains competitive in part of the zoomed-in region, but MoQTS still provides a wider and more consistently favorable spread. These observations indicate that MoQTS achieves stronger search effectiveness and better Pareto-front coverage than the compared algorithms.

4.5 Statistical Significance Analysis

The statistical significance of the performance differences among the compared algorithms is assessed using non-parametric statistical tests on the four performance indicators. Table 4 reports the average ranks obtained by the Friedman test. Since a smaller average rank indicates better performance, the results show that MoQTS achieves the best average rank on all four indicators, namely HV, HV rate, GD, and IGD. Moreover, the Friedman p-values for all four indicators are extremely small, indicating that statistically significant differences exist among the compared algorithms.

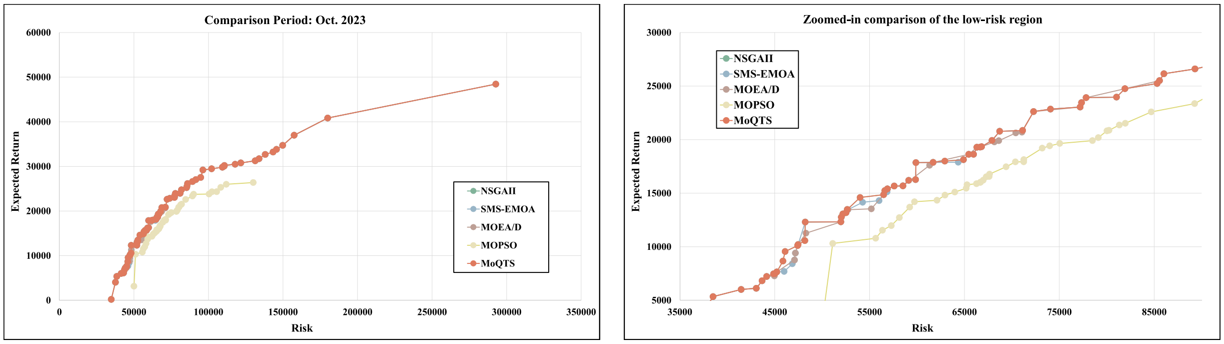

Figure 4: Comparison of Pareto-front distributions for MoQTS and other algorithms in global (**left**) and zoomed-in (**right**) views.

Table 4: Statistical comparison results based on the Friedman test.

Metric	MoQTS	NSGA-II	SMS-EMOA	MOEA/D	MOPSO	p-Value
HV	**1.9869**	2.0523	2.3693	3.6176	4.9739	1.779×10^{-95}
HV rate (%)	**1.9673**	2.0588	2.3758	3.6242	4.9739	6.993×10^{-96}
GD	**2.1503**	2.1601	2.3758	3.3399	4.9739	1.230×10^{-81}
IGD	**1.9510**	2.0229	2.4804	3.5523	4.9935	1.185×10^{-92}

Note: Bold values indicate the best average rank for each performance metric.

Table 5 summarizes the results of the post-hoc Wilcoxon signed-rank test with Holm correction, in which the proposed method, MoQTS, is separately compared with each competing algorithm. The corrected p-values show that MoQTS significantly outperforms SMS-EMOA, MOEA/D, and MOPSO on all four indicators. Meanwhile, the corrected p-values between MoQTS and NSGA-II are all larger than 0.05, suggesting that these two methods exhibit comparable performance across the four indicators at the 5% significance level.

Table 5: Post-hoc Wilcoxon signed-rank test with Holm correction using MoQTS as the control method.

Metric	NSGA-II	SMS-EMOA	MOEA/D	MOPSO
HV	2.652×10^{-1}	1.176×10^{-6}	2.696×10^{-20}	6.301×10^{-26}
HV rate (%)	2.087×10^{-1}	7.421×10^{-7}	1.263×10^{-20}	6.301×10^{-26}
GD	5.940×10^{-1}	7.253×10^{-5}	7.653×10^{-16}	6.301×10^{-26}
IGD	8.753×10^{-2}	8.723×10^{-8}	1.263×10^{-20}	2.957×10^{-26}

4.6 Computational Cost Analysis

To further evaluate the efficiency of the proposed framework, this subsection analyzes the computational cost of MoQTS and the comparison algorithms. Table 6 summarizes the computational cost of the compared algorithms, including the average execution time, average evaluation count, average entanglement move local search count, and average total evaluation count. The average evaluation count represents the number of evaluations performed during the optimization process, whereas the average entanglement move local search count denotes the additional evaluations introduced by the entanglement move local search mechanism in MoQTS. As shown in Table 6, MoQTS requires substantially fewer evaluations than the other comparison algorithms. Even after including the additional entanglement move local search evaluations, MoQTS still requires only about 36% of the total evaluation count used by NSGA-II, SMS-EMOA, and MOEA/D, and about 18% of that used by MOPSO. Furthermore, under the same experimental setup, the actual execution time of MoQTS is also considerably shorter. These results show that MoQTS can identify a better Pareto front than the traditional algorithms while using only a fraction of their computational cost, highlighting its strong advantage in computational efficiency.

Table 6: Computational cost of the comparison algorithms.

Algorithm	Avg. Execution Time (s)	Avg. Evaluation Count	Avg. Entanglement Move Local Search Count	Avg. Total Evaluation Count
MoQTS	**6.72**	**1,084,755**	4755	**1,089,510**
NSGA-II	10.34	3,003,000	–	3,003,000
SMS-EMOA	19.94	3,003,000	–	3,003,000
MOPSO	84.84	5,997,000	–	5,997,000
MOEA/D	95.24	3,003,000	–	3,003,000

Note: Bold values indicate the best result for each computational cost metric.

4.7 Ablation Study

This subsection presents the ablation experiments of MoQTS to analyze the contributions of its main components. In particular, the proposed framework is evaluated by selectively removing the quantum NOT gate mechanism and the entanglement move local search, and then comparing the resulting performance in terms of HV, HV rate, GD, and IGD. For clarity, four variants are considered in this experiment: MoQTS w/o NOT & EL denotes the variant without both the quantum NOT gate and the entanglement move local search, MoQTS w/o EL denotes the variant without the entanglement move local search only, MoQTS w/o NOT denotes the variant without the quantum NOT gate only, and MoQTS denotes the complete proposed method with both mechanisms enabled.

Table 7 shows the ablation results for HV and HV rate. When both mechanisms are removed, the worst performance is observed and is therefore used as the baseline. After adding back only the quantum NOT gate, that is, in MoQTS w/o EL, both HV and HV rate improve slightly. A much larger improvement is observed for MoQTS w/o NOT, where the entanglement move local search is retained while the quantum NOT gate is removed. Finally, MoQTS achieves the best HV and HV rate among all variants. These results indicate that both mechanisms improve solution quality, with the entanglement move local search providing the larger contribution.

Table 8 further reports the ablation results for GD and IGD, for which smaller values indicate better performance. A similar trend is observed. MoQTS w/o NOT & EL again yields the worst results. After reintroducing only the quantum NOT gate, MoQTS w/o EL improves both GD and IGD, although the improvement remains limited. When only the quantum NOT gate is removed, namely in MoQTS w/o NOT, the results improve substantially, indicating that retaining the entanglement move local search is much more beneficial for convergence performance. The best GD and IGD values are finally achieved by MoQTS. Overall, the ablation results consistently show a step-by-step improvement from the variant without both mechanisms to the complete method, confirming that both components are beneficial, with the entanglement move local search playing the more dominant role.

Table 7: Ablation study results for HV and HV rate.

Variant	HV	Imp. (%)	HV Rate (%)	Imp. (%)
MoQTS w/o NOT & EL	0.616778	–	97.6499	–
MoQTS w/o EL	0.617048	0.0438	97.6605	0.0108
MoQTS w/o NOT	0.624227	1.2079	98.6899	1.0650
MoQTS	**0.624245**	**1.2107**	**98.6925**	**1.0677**

Note: Bold values indicate the best result for each performance metric and the highest improvement.

Table 8: Ablation study results for GD and IGD.

Variant	GD	Imp. (%)	IGD	Imp. (%)
MoQTS w/o NOT & EL	5.5237×10^{-4}	–	2.3203×10^{-3}	–
MoQTS w/o EL	4.7159×10^{-4}	14.6233	2.0806×10^{-3}	10.3335
MoQTS w/o NOT	3.4783×10^{-5}	93.7028	1.7292×10^{-4}	92.5478
MoQTS	6.2596×10^{-6}	**98.8668**	1.8166×10^{-5}	**99.2171**

Note: Bold values indicate the best result for each performance metric and the highest improvement.

To further examine whether the performance differences among the ablation variants are statistically significant, non-parametric statistical tests are conducted. Table 9 presents the statistical comparison results based on the Friedman test. The mean ranks show that MoQTS obtains the best rank for all four evaluation indicators, including HV, HV rate, GD, and IGD. Specifically, MoQTS achieves the lowest mean rank among all variants, indicating that the complete framework with both the quantum NOT gate and the entanglement move local search generally provides the most competitive performance. In contrast, MoQTS w/o NOT & EL obtains worse ranks in most indicators, confirming that removing both proposed mechanisms degrades the overall performance. Moreover, the p-values reported in Table 9 are all smaller than 0.05, indicating that statistically significant differences exist among the four variants for all evaluation metrics. This result further supports the effectiveness of the proposed components in MoQTS.

Table 9: Statistical comparison results of the ablation study based on the Friedman test.

Variant	MoQTS w/o NOT & EL	MoQTS w/o EL	MoQTS w/o NOT	MoQTS	p-Value
HV	2.8627	2.7908	2.1863	**2.1601**	2.823×10^{-29}
HV rate (%)	2.8627	2.7908	2.1863	**2.1601**	2.823×10^{-29}
GD	2.7386	2.6993	2.2941	**2.2680**	1.660×10^{-29}
IGD	2.8235	2.8366	2.1732	**2.1667**	4.540×10^{-29}

Note: Bold values indicate the best average rank for each performance metric.

After confirming the existence of significant differences among the variants, a post-hoc Wilcoxon signed-rank test with Holm correction is further performed by using MoQTS as the control method. The results are shown in Table 10. For HV and HV rate, MoQTS shows statistically significant differences compared with MoQTS w/o NOT & EL and MoQTS w/o EL, with very small p-values. Similar results can also be observed for GD and IGD. These results indicate that removing the entanglement move local search leads to a significant performance degradation, demonstrating that this mechanism plays an important role in improving both convergence and diversity. The quantum NOT gate is mainly designed to enhance the exploration ability of MoQTS. By changing the search direction, this mechanism increases solution diversity and provides additional opportunities for the search process to escape from local optima. In the ablation results, the performance gap between MoQTS and MoQTS w/o NOT is relatively limited. In the considered experimental setting, the entanglement moves the local search, effectively refining the obtained Pareto front and thereby reducing the likelihood of severe local-optimum stagnation. Consequently, the contribution of the quantum NOT gate is mainly reflected in maintaining exploration capability and preventing premature convergence, rather than causing a large performance difference. Overall, the statistical analysis indicates that the entanglement move local search effectively improves Pareto front refinement, while the quantum NOT gate further supports the robustness and exploration ability of the complete MoQTS framework.

Table 10: Post-hoc Wilcoxon signed-rank test with Holm correction for the ablation study, using MoQTS as the control method.

Variant	MoQTS w/o NOT & EL	MoQTS w/o EL	MoQTS w/o NOT
HV	3.013×10^{-10}	3.520×10^{-10}	0.3343
HV rate (%)	3.013×10^{-10}	3.520×10^{-10}	0.3343
GD	2.482×10^{-7}	3.660×10^{-7}	0.3270
IGD	2.557×10^{-10}	3.519×10^{-10}	0.4724

4.8 Portfolio Performance

This subsection presents representative portfolio trends under continuous investment to illustrate the distinct behaviors of different portfolio types over a longer horizon. Two representative cases, corresponding to 2017 and 2024, are provided. In both cases, the lowest-risk portfolio, the highest-TR portfolio, and the highest-ER portfolio are compared.

Fig. 5 presents the representative portfolio performance in 2017 and 2024. In 2017, the lowest-risk portfolio remains the most stable and exhibits the smallest fluctuation range. The highest-TR portfolio shows a steady upward trend with a favorable balance between growth and stability, whereas the highest-ER portfolio achieves the strongest growth but also exhibits the largest fluctuation range. In 2024, the differences among the three portfolios become even more pronounced. The highest-TR portfolio maintains a strong upward trend while preserving relatively stable fluctuations, whereas the highest-ER portfolio again delivers the strongest growth together with the highest level of volatility.

Table 11 further summarizes the corresponding portfolio statistics. In both periods, the lowest-risk portfolio has the smallest risk value. In 2017, it also records the smallest maximum drawdown, whereas the highest-ER portfolio achieves the largest cumulative return but also suffers the highest drawdown. In 2024, the highest-TR portfolio attains the highest trend ratio and delivers a strong cumulative return with slightly lower drawdown than the lowest-risk portfolio, indicating an especially favorable balance between return and stability. By contrast, the highest-ER portfolio again achieves the largest cumulative return, but at the cost of the highest drawdown. These observations are consistent with the trajectory patterns shown in Fig. 5.

These two cases show clear differences among the three portfolio types. The lowest-risk portfolio is more stable, the highest-ER portfolio has the strongest growth potential, and the highest-TR portfolio provides a balanced compromise between return and stability.

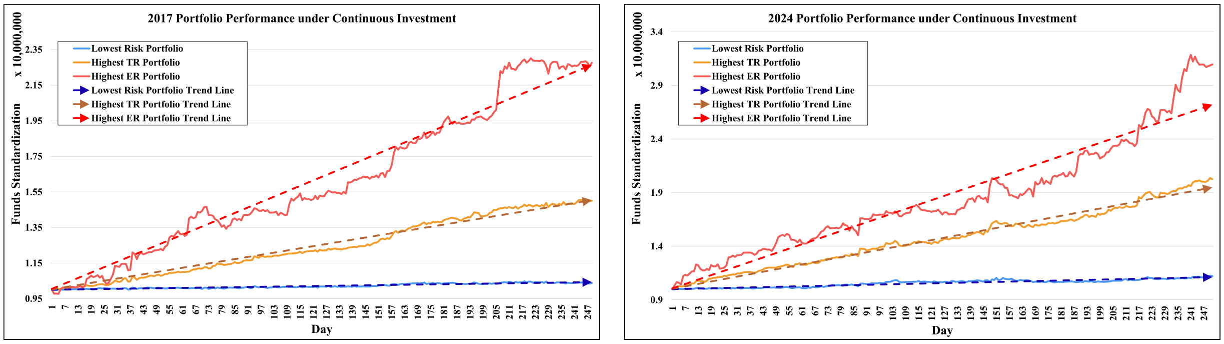

Figure 5: Representative portfolio performance under continuous investment in 2017 (**left**) and 2024 (**right**).

Table 11: Portfolio performance in 2017 and 2024.

Portfolio Type	2017					2024				
	Risk	Return	Trend Ratio	Cum. Return (%)	MDD (%)	Risk	Return	Trend Ratio	Cum. Return (%)	MDD (%)
Lowest Risk	41,819.8635	1,715.3173	0.041017	3.7858	1.2423	124,671.5354	4,471.0366	0.035863	11.9935	3.8605
Highest TR	233,992.6481	20,236.5518	0.086484	50.0510	1.7866	351,084.7977	37,905.5095	0.107682	102.5213	3.7401
Highest ER	796,879.8898	50,795.3290	0.063743	127.6663	8.3745	1,529,073.4810	68,581.0659	0.044851	209.6449	9.2169

5 Conclusion

This study proposes a Multi-objective Quantum-inspired Tabu Search algorithm for portfolio optimization with short-selling strategies. The proposed framework integrates quantum-inspired techniques and entanglement move local search approach to improve both convergence and solution diversity in multi-objective optimization. Under the TR formulation, the optimization objective jointly considers expected return and investment risk, enabling the construction of high-quality non-dominated portfolios for different investor preferences.

Experimental evaluations on DJIA constituent data from 2013 to 2025 show that MoQTS can obtain high-quality Pareto-optimal solutions with strong convergence and diversity performance. PF comparisons also show that MoQTS generates a favorable non-dominated distribution, with a greater concentration of solutions in low-risk, high-return regions. The computational cost analysis demonstrates that MoQTS can efficiently identify high-quality Pareto fronts while requiring substantially fewer computations than the comparison algorithms. After accounting for the additional evaluations introduced by the entanglement move local search, MoQTS still requires only about 36% of the total evaluation count used by NSGA-II, SMS-EMOA, and MOEA/D, and about 18% of that used by MOPSO, while also achieving the shortest execution time. The ablation study and statistical analysis further support the positive contributions of both the quantum NOT gate and the entanglement move local search, particularly in improving Pareto front refinement. These findings suggest that the proposed quantum-inspired search mechanism is effective and highly efficient for practical portfolio optimization in dynamic market environments.

Future work can extend MoQTS to larger assets, incorporate additional realistic markets, and evaluate adaptive parameter-control strategies to further improve robustness and scalability.

Acknowledgement: The authors acknowledge the support from the National Center for Theoretical Sciences, Taiwan.

Funding Statement: This work was supported by the National Science and Technology Council, Taiwan, under Grants 113-2221-E-260-014-MY2 and 115-2119-M-033-001.

Author Contributions: The authors confirm contributions to the paper as follows: Yun-Ting Lai performed the system implementation, conducted the experimental analysis, and wrote the draft manuscript; Ming-Ho Chang jointly verified the experimental results, conducted comparative experiments, analyzed the experimental data, and assisted in manuscript writing; Yao-Hsin Chou proposed the model architecture and supervised the project. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: All data generated or analyzed during this study are included in this published article.

Ethics Approval: Not applicable.

Conflicts of Interest: Given his role as guest editor of this journal, Yao-Hsin Chou had no involvement in the peer review of this article and had no access to information regarding its peer review. Full responsibility for the editorial process for this article was delegated to another journal editor. The authors declare no other conflicts of interest.

References

1. Markowitz H. Portfolio selection. J Financ. 1952;7(1):77–91. doi:10.2307/2975974.
2. Holland JH. Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence. Ann Arbor, MI, USA: University of Michigan Press; 1975.
3. Kennedy J, Eberhart R. Particle swarm optimization. In: Proceedings of the IEEE International Conference on Neural Networks (ICNN); 1995 Nov 27–Dec 1; Perth, WA, Australia. p. 1942–8.
4. Storn R, Price K. Differential evolution: a simple and efficient heuristic for global optimization over continuous spaces. J Glob Optim. 1997;11(4):341–59.
5. Han KH, Kim JH. Quantum-inspired evolutionary algorithm for a class of combinatorial optimization. IEEE Trans Evol Comput. 2002;6(6):580–93. doi:10.1109/tevc.2002.804320.
6. Sun J, Feng B, Xu W. Particle swarm optimization with particles having quantum behavior. In: Proceedings of the IEEE Congress on Evolutionary Computation (CEC); 2004 Jun 19–23; Portland, OR, USA. p. 322–7.
7. Soleimanpour-Moghadam M, Nezamabadi-Pour H, Farsangi MM. A quantum inspired gravitational search algorithm for numerical function optimization. Inf Sci. 2014;267(6):83–100. doi:10.1016/j.ins.2013.09.006.
8. Kuo SY, Lai YT, Jiang YC, Chang MH, Wu KM, Chen PC, et al. Entanglement local search-assisted quantum-inspired optimization for portfolio optimization in G20 markets. In: Proceedings of the Companion Conference on Genetic and Evolutionary Computation (GECCO Companion); 2023 Jul 15–19; Lisbon, Portugal. p. 2232–40.
9. Deb K, Pratap A, Agarwal S, Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans Evol Comput. 2002;6(2):182–97. doi:10.1109/4235.996017.
10. Zhang Q, Li H. MOEA/D: a multiobjective evolutionary algorithm based on decomposition. IEEE Trans Evol Comput. 2007;11(6):712–31. doi:10.1109/tevc.2007.892759.
11. Beume N, Naujoks B, Emmerich M. SMS-EMOA: multiobjective selection based on dominated hypervolume. Eur J Oper Res. 2007;181(3):1653–69.
12. Sharpe WF. The sharpe ratio. J Portf Manag. 1994;21(1):49–58. doi:10.3905/jpm.1994.409501.
13. Hoang SD, Dey SK, Nguyen THH, Nguyen PND. The role of AI recommendations in extending the black-litterman portfolio. Int J Intell Comput Cybern. 2026;19(1):115–38. doi:10.1108/ijicc-03-2025-0137.
14. Xidonas P, Essner E. On ESG portfolio construction: a multi-objective optimization approach. Comput Econ. 2024;63(1):21–45.
15. Wu Q, Liu X, Qin J, Zhou L, Mardani A, Deveci M. An integrated multi-criteria decision-making and multi-objective optimization model for socially responsible portfolio selection. Technol Forecast Soc Change. 2022;184(4):121977. doi:10.1016/j.techfore.2022.121977.
16. Deliktaş D, Ustun O. Multi-objective genetic algorithm based on the fuzzy MULTIMOORA method for solving the cardinality constrained portfolio optimization. Appl Intell. 2023;53(12):14717–43. doi:10.1007/s10489-022-04240-6.
17. Chou YH, Kuo SY, Lo YT. Portfolio optimization based on funds standardization and genetic algorithm. IEEE Access. 2017;5:21885–900. doi:10.1109/access.2017.2756842.