DEPARTMENT OF ECONOMICS AND FINANCE

DEGREE COURSE: ECONOMICS AND FINANCE

COURSE: COMPUTATIONAL TOOLS FOR FINANCE

Combinatorial Optimization applied in Portfolio Construction

Prof. Valerio Marchisio

. . . . . .

SUPERVISOR

Prof. Mimun Hlafo Alfie

. . . . . .

CO-SUPERVISOR

Federico Denti - 780731

. . . . . .

CANDIDATE

Academic Year 2024–2025

Abstract

This thesis investigates a quantum-inspired portfolio selection framework formulated as a Quadratic Unconstrained Binary Optimization (QUBO) problem and solved via simulated annealing.

The model addresses cardinality-constrained portfolio construction by selecting a fixed number of assets from a large investment universe. The binary quadratic energy function directly incorporates historical mean returns in its linear component and covariance interactions in its quadratic component, resulting in a transparent and economically coherent discrete mean–variance formulation.

A Time-to-Target (TTT) calibration procedure is implemented to statistically characterize the computational effort required to reach a predefined energy threshold. Portfolio performance is evaluated through a rolling block-forward backtesting framework, ensuring strict separation between training and testing periods and preventing look-ahead bias.

Empirical results over multiple out-of-sample windows indicate that the QUBO-based portfolios achieve competitive risk-adjusted performance relative to classical benchmarks. Differences between warm and random initialization schemes are limited, suggesting that simulated annealing exhibits low sensitivity to starting configurations under realistic iteration budgets.

Overall, the proposed framework demonstrates that binary quadratic optimization can provide a structurally coherent and computationally tractable approach to discrete mean–variance portfolio selection in large-scale empirical settings.

CONTENTS

CHAPTER 1

INTRODUCTION

1.1 Motivation and Research Context

Portfolio construction remains one of the central problems in quantitative finance. Since the seminal contribution of Markowitz (1952), mean–variance theory has provided a coherent framework for balancing expected return and risk through continuous quadratic optimization. In its classical form, optimal portfolio weights are continuous decision variables and the feasible set is convex. Under these assumptions, the theory delivers analytical tractability, geometric interpretation via the efficient frontier, and a clear mapping between risk aversion and allocation decisions.

In practical asset management, however, portfolios are rarely implemented without structural constraints. Institutional investors and systematic strategies often face explicit limits on the number of assets held. Such *cardinality constraints* arise from transaction costs, liquidity management, regulatory mandates, turnover control, and operational simplicity. When the investment universe is large—as is typical in U.S. equity markets—holding the full universe is neither feasible nor economically desirable.

Imposing a fixed cardinality fundamentally alters the optimization problem. Once a portfolio must contain exactly K assets out of N, the decision becomes inherently combinatorial. The number of feasible subsets equals

$$\binom{N}{K},$$

which grows super-polynomially in N. Even for moderate values (e.g., $N = 500$, $K = 10$), the number of candidate portfolios is astronomically large. As established in the literature, cardinality-constrained mean–variance optimization is NP-hard in general (Chang et al., 2000; Jobst et al., 2001).

This computational intractability has motivated two main research directions. The first relies

on mixed-integer quadratic programming (MIQP) techniques that attempt to solve the problem exactly, at the cost of substantial computational effort. The second adopts heuristic and meta-heuristic approaches—including genetic algorithms, tabu search, particle swarm methods, and annealing-based procedures—designed to explore large discrete search spaces efficiently.

More recently, the emergence of quantum computing and quantum-inspired optimization has renewed interest in binary quadratic formulations of combinatorial problems. The Quadratic Unconstrained Binary Optimization (QUBO) framework provides a natural encoding for decision problems involving pairwise interactions over binary variables. Portfolio selection under cardinality constraints fits naturally within this class, since covariance interactions can be expressed as quadratic terms over binary inclusion variables.

Against this background, the central research question of this thesis is:

> Can a binary quadratic formulation of cardinality-constrained portfolio selection, combined with annealing-based search and statistically calibrated computational effort, produce economically meaningful and empirically competitive portfolios relative to classical benchmarks?

Rather than assuming structural superiority of combinatorial optimization over continuous mean–variance methods, this thesis adopts a disciplined empirical approach. The objective is to construct a transparent binary quadratic selection framework, calibrate its computational budget through a principled procedure, and evaluate its out-of-sample performance over a long historical horizon.

The work therefore lies at the intersection of classical portfolio theory, combinatorial optimization, and quantum-inspired computational methodologies, contributing to the growing literature that examines how discrete optimization tools can be meaningfully integrated into practical asset allocation.

1.2 From Continuous Mean–Variance Optimization to Discrete Portfolio Selection

The classical mean–variance framework assumes that portfolio weights $w \in \mathbb{R}^N$ are continuous decision variables satisfying linear constraints such as full investment and, possibly, long-only restrictions. The standard formulation can be written as

$$\max_{w} \left(w^\top \mu - \lambda w^\top \Sigma w \right),$$

where μ denotes the vector of expected returns, Σ the covariance matrix, and $\lambda > 0$ the risk-aversion parameter.

Under standard regularity conditions, this problem is convex and admits closed-form or efficiently computable solutions. The efficient frontier provides a continuous mapping between risk preferences and portfolio allocations, and quadratic utility maximization is equivalent to Sharpe ratio maximization for suitable parametrizations (Markowitz, 1952).

Introducing a cardinality constraint changes the structure of the problem fundamentally. Instead of allocating capital across all available assets, the investor must first select a subset of size K from a universe of N assets. Let $x_i \in \{0, 1\}$ denote a binary inclusion variable indicating whether asset i is selected. The cardinality condition is expressed as

$$\sum_{i=1}^{N} x_i = K.$$

The feasible set becomes discrete and non-convex. The number of admissible subsets equals

$$\binom{N}{K},$$

which grows combinatorially in N. For realistic equity universes, exhaustive enumeration is infeasible. As shown in the optimization literature, the resulting problem is NP-hard (Chang et al., 2000; Jobst et al., 2001).

Conceptually, the portfolio construction task now separates into two interrelated layers:

1. **Subset selection:** choosing which K assets to include.

2. **Continuous allocation:** determining optimal weights within the selected subset.

The economic objective remains rooted in the mean–variance trade-off. However, the feasible domain is redefined. Instead of optimizing solely over continuous weights, the decision process involves navigating a discrete space of possible asset combinations, each of which induces its own covariance structure and return profile.

This transition has important implications. The mapping between risk aversion and portfolio composition is no longer continuous; small changes in parameters may produce discrete shifts in selected assets. Moreover, the efficient frontier becomes a collection of attainable discrete portfolios rather than a smooth curve.

The key methodological challenge is therefore not redefining the economic objective, but designing an optimization procedure capable of exploring the exponentially large discrete space in a computationally tractable manner while preserving the mean–variance economic logic.

The binary quadratic formulation introduced in the next section addresses precisely this challenge by encoding covariance interactions and return contributions directly over binary inclusion variables, thereby linking discrete subset selection to the classical risk–return paradigm.

1.3 Binary Quadratic Modeling and Annealing-Based Optimization

This thesis studies a discrete portfolio selection problem in which the investor must select exactly K assets out of a universe of N. The selection decision is represented through a binary vector

$$x = (x_1, \ldots, x_N)^\top \in \{0,1\}^N,$$

where $x_i = 1$ indicates that asset i is included in the portfolio and $x_i = 0$ otherwise. The cardinality constraint is imposed as

$$\sum_{i=1}^{N} x_i = K.$$

The objective is to identify a subset that provides an economically meaningful trade-off between expected return and risk, consistent with the mean–variance paradigm.

Binary quadratic encoding of portfolio selection. Quadratic binary models provide a natural representation for portfolio selection because risk is determined by pairwise covariances. Let $\mu \in \mathbb{R}^N$ denote the vector of expected (historical) mean returns and let $\Sigma \in \mathbb{R}^{N\times N}$ denote the covariance matrix of asset returns. Given a binary selection vector x, the portfolio-level quantities induced by the selected subset can be expressed in quadratic form.

First, the aggregate return signal associated with the selected assets is modeled as the linear term

$$\mu^\top x = \sum_{i=1}^{N} \mu_i x_i,$$

where μ_i is the historical mean return of asset i. This term increases when assets with larger mean returns are included.

Second, the risk interaction among selected assets is captured by the quadratic term

$$x^\top \Sigma x = \sum_{i=1}^{N}\sum_{j=1}^{N} \Sigma_{ij} x_i x_j,$$

where Σ_{ij} is the covariance between assets i and j. Since $x_i x_j = 1$ only when both assets are selected, the quadratic form aggregates all covariance contributions within the chosen subset.

Combining these components, the discrete selection problem is encoded through an *energy* (objective) function of the form

$$E(x) = \beta\, x^\top \Sigma x - \alpha\, \mu^\top x,$$

where $\alpha \geq 0$ controls the emphasis on return and $\beta \geq 0$ controls the emphasis on risk. Minimizing $E(x)$ therefore corresponds to selecting assets that jointly reduce covariance exposure while maintaining high mean return contributions. This structure belongs to the class of Quadratic Unconstrained Binary Optimization (QUBO) models, i.e.,

$$\min_{x\in\{0,1\}^N} \; x^\top Q x + c^\top x,$$

which has become a standard formulation for binary decision problems with pairwise interactions (Lucas, 2014). In the present application, the quadratic term inherits a direct financial interpretation from the covariance matrix, while the linear term captures the contribution of expected returns.

A key implementation choice in this thesis is that the cardinality constraint is enforced *algorithmically* rather than through penalty terms. Specifically, the search procedure explores only feasible states x satisfying $\sum_i x_i = K$ by using neighborhood moves that preserve cardinality. This avoids sensitivity to penalty calibration and keeps the optimization objective aligned with the economic components of risk and return.

Annealing-based search and initialization. The discrete feasible set contains $\binom{N}{K}$ portfolios and becomes rapidly intractable to search exhaustively. To address this, the thesis employs a simulated annealing (SA) heuristic, a stochastic optimization method inspired by the physical annealing process (Kirkpatrick et al., 1983). SA iteratively proposes local modifications of the current solution and accepts them based on the associated energy difference and a temperature schedule.

Let x denote the current portfolio and x' a candidate neighbor. Denote the energy difference by

$$\Delta E = E(x') - E(x).$$

The Metropolis acceptance rule is:

$$\Pr(\text{accept } x') = \begin{cases} 1, & \Delta E \leq 0, \\ \exp\left(-\dfrac{\Delta E}{T}\right), & \Delta E > 0, \end{cases}$$

where $T > 0$ is the temperature parameter. The evolution of the temperature plays a crucial role in the simulated annealing procedure. In this implementation, the temperature follows a geometric cooling schedule. Starting from an initial value T_{init}, the temperature is progressively reduced at each iteration according to

$$T_{t+1} = \max(T_{\min},\, T_t \cdot \gamma)\,,$$

where $T_{\min}$ is a strictly positive lower bound and $\gamma \in (0,1)$ is the cooling factor. The cooling factor is calibrated so that the temperature decreases smoothly from T_{init} to $T_{\min}$ over the total number of iterations N_{iter}, namely

$$\gamma = \left(\frac{T_{\min}}{T_{\text{init}}}\right)^{1/N_{\text{iter}}}.$$

This geometric schedule ensures a gradual transition from an exploratory regime—where uphill moves are accepted with non-negligible probability—to a near-deterministic local refinement phase as the temperature approaches $T_{\min}$. By construction, the temperature never falls below $T_{\min}$, which prevents numerical instabilities in the acceptance probability.

At high temperature, SA may accept worse moves ($\Delta E > 0$) with non-negligible probability, which promotes exploration and helps escape local minima. As T decreases according to a cooling schedule, the algorithm becomes increasingly greedy and concentrates on local refinement.

Neighborhood moves are designed to preserve the cardinality constraint. Operationally, a neighbor is generated by swapping one selected asset with one unselected asset, so that $\sum_i x_i$ remains equal to K at every iteration. This ensures that all candidate portfolios are feasible by construction.

The algorithm is executed under two initialization schemes:

- **Random start:** the initial solution selects K assets uniformly at random from the universe.
- **Warm start:** the initial solution selects the K assets with the highest historical mean returns μ_i.

The warm-start design is coherent with the model specification, since the same mean return vector μ also enters the QUBO energy through the linear term $-\alpha\, \mu^\top x$. The empirical analysis later evaluates whether initialization materially affects out-of-sample performance and computational effort.

Overall, the annealing-based procedure provides a scalable way to explore an economically interpretable binary quadratic landscape, allowing the thesis to study how discrete subset selection interacts with mean–variance principles under realistic universe sizes.

1.4 Contributions and Structure of the Thesis

This thesis contributes to the literature on cardinality-constrained portfolio optimization by integrating classical mean–variance principles with a structured binary quadratic formulation and a carefully calibrated annealing-based search procedure.

The first contribution is **methodological**. The portfolio selection problem is formulated as a Binary Quadratic Model (BQM) under an exact cardinality constraint. Asset inclusion decisions are represented through binary variables, and covariance interactions enter naturally as quadratic terms. Expected returns are incorporated linearly using the historical mean return vector. Importantly, the cardinality constraint is enforced algorithmically through neighborhood design rather than through penalty coefficients, ensuring that the optimization landscape reflects only economically meaningful risk–return trade-offs.

The second contribution concerns **computational calibration**. Heuristic optimization methods depend critically on iteration budgets and stopping rules. Instead of selecting these parameters arbitrarily, this thesis implements a Time-to-Target (TTT) framework, which statistically estimates the computational effort required to reach a predefined energy threshold with a given confidence level. This provides a principled and reproducible approach to tuning the simulated annealing procedure and documenting the effective difficulty of the combinatorial landscape.

The third contribution is **empirical**. The framework is evaluated on a large U.S. equity universe over a long historical horizon (1997–2026). Portfolio performance is assessed using a rolling block-forward backtesting design, with strictly separated training and testing windows to avoid look-ahead bias. Out-of-sample results are compared against intuitive benchmarks, including an equal-weight full-universe portfolio and a Top-K mean-return ranking strategy. This benchmarking structure allows the incremental value of covariance-aware combinatorial selection to be isolated from simple return-based ranking.

The empirical analysis further investigates the sensitivity of the annealing procedure to initialization by comparing random and warm-start configurations. The warm initialization selects the K assets with the highest historical mean returns, consistent with the linear component of the binary quadratic energy. This design enables an assessment of whether initialization materially affects solution quality or whether the algorithm converges toward similar portfolio structures irrespective of the starting point. Importantly, the thesis does not assume that discrete optimization must dominate classical continuous approaches. Rather, it provides a systematic evaluation of when and how a binary quadratic selection framework produces economically meaningful portfolios, how computational effort scales, and how structural properties of the discrete formulation relate to realized out-of-sample performance.

Structure of the thesis. Chapter 2 presents the data construction, estimation procedures for mean returns and covariance matrices, the binary quadratic formulation, the simulated annealing dynamics, the Time-to-Target calibration, and the block-forward backtesting methodology in detail. Chapter 3 reports empirical results, including performance comparisons, distributional analysis of out-of-sample metrics, and robustness checks. Chapter 4 discusses implications, limitations, and directions for future research.

CHAPTER 2

DATA AND METHODOLOGY

2.1 Data

Universe definition and sample period. The empirical analysis is conducted on a cross-section of $N = 500$ U.S. equities. The sample period spans from 31 December 1997 to January 2026. The raw input data consist of daily adjusted closing prices.

Let $P_{i,t}$ denote the adjusted close price of asset $i \in \{1, \ldots, N\}$ observed at trading day $t \in \{1, \ldots, T\}$. Adjusted prices are used to ensure that observed price changes reflect genuine economic returns rather than mechanical price discontinuities induced by corporate actions. Working with adjusted prices is standard practice in empirical asset pricing and portfolio construction (Campbell et al., 1997).

The relatively large cross-sectional dimension ($N = 500$) ensures that cardinality-constrained selection with $K = 10$ represents a genuinely high-dimensional combinatorial problem.

Panel construction and data cleaning. The dataset is assembled from individual asset-level price files and transformed into a single aligned price panel through the following steps:

1. **Date parsing and chronological ordering.** For each asset, the date column is converted into a proper time index and observations are sorted in ascending order.

2. **Duplicate handling.** If multiple entries share the same date for a given asset, only the most recent valid observation is retained.

3. **Positivity checks.** Observations with non-positive prices are removed. Since log-returns require strictly positive prices, this step ensures mathematical well-definedness of subsequent transformations.

4. **Cross-sectional alignment.** Individual price series are merged via an inner join on trading dates. Only dates for which all included assets have valid prices are retained. This produces

a rectangular panel with no missing observations.

The complete-case alignment reduces the effective time dimension T, but guarantees that all cross-sectional statistics (means, covariances) are computed on a consistent and fully observed return matrix.

Return construction. Daily returns are constructed as continuously compounded (log) returns:

$$r_{i,t} = \log(P_{i,t}) - \log(P_{i,t-1}), \quad t = 2, \ldots, T.$$

Log-returns are preferred for several reasons (Campbell et al., 1997):

- They are additive across time: multi-period log-returns equal the sum of daily log-returns.
- They correspond to continuously compounded rates.
- For small daily changes, they closely approximate simple returns.

Let $R \in \mathbb{R}^{(T-1)\times N}$ denote the return matrix with entries $r_{i,t}$. All subsequent moment estimation, subset selection, and portfolio evaluation steps are based on this return matrix.

Annualization conventions. Since returns are observed at daily frequency, annualized quantities are computed using standard scaling conventions.

Let $\{r_{p,t}\}_{t=1}^{T_{\text{test}}}$ denote a sequence of daily portfolio log-returns in a test window. The average daily return is

$$\bar{r}_p = \frac{1}{T_{\text{test}}} \sum_{t=1}^{T_{\text{test}}} r_{p,t}.$$

The annualized return is computed as

$$\widehat{R}_{\text{ann}} = \exp(252 \cdot \bar{r}_p) - 1,$$

where 252 denotes the approximate number of trading days per year.

Annualized volatility is given by

$$\widehat{\sigma}_{\text{ann}} = \sqrt{252} \cdot \sqrt{\frac{1}{T_{\text{test}} - 1} \sum_{t=1}^{T_{\text{test}}} (r_{p,t} - \bar{r}_p)^2}.$$

These conventions are applied consistently throughout the empirical analysis.

Train–test separation and usage of the full sample. The dataset is used for two complementary purposes.

First, the full return panel (1997–2026) is employed for algorithmic calibration, including the Time-to-Target (TTT) procedure, in order to obtain stable estimates of optimization difficulty across heterogeneous market regimes.

Second, the main empirical evaluation relies on a rolling block-forward design. For each window, all quantities used in selection and allocation—including the mean return vector $\hat{\mu}$ and the covariance matrix $\hat{\Sigma}$—are computed exclusively from the training segment. Portfolio performance is then evaluated strictly on the subsequent test segment.

This strict temporal separation prevents look-ahead bias and ensures that reported results correspond to genuine out-of-sample performance.

2.2 Problem Formulation and Mathematical Framework

This section formalizes the statistical estimators and optimization problems underlying both the continuous and discrete stages of the proposed framework.

Estimation of mean returns and covariance matrix. Let $R \in \mathbb{R}^{T_{\text{tr}} \times N}$ denote the matrix of daily log-returns observed over a training window $\mathcal{T}$ of length T_{tr}. Let $r_t \in \mathbb{R}^N$ denote the cross-sectional return vector at date $t \in \mathcal{T}$.

The sample mean vector is defined as

$$\hat{\mu} = \frac{1}{T_{\text{tr}}} \sum_{t \in \mathcal{T}} r_t,$$

where $\hat{\mu}_i$ represents the average daily return of asset i over the training window.

The sample covariance matrix is

$$\hat{\Sigma} = \frac{1}{T_{\text{tr}} - 1} \sum_{t \in \mathcal{T}} (r_t - \hat{\mu})(r_t - \hat{\mu})^\top.$$

Under classical assumptions (e.g., independent and identically distributed returns with finite second moments), $\hat{\mu}$ and $\hat{\Sigma}$ are consistent estimators. However, in finite samples, estimation error can be substantial, particularly for expected returns (Merton, 1980).

To improve numerical conditioning and mitigate estimation noise, a shrinkage estimator is employed:

$$\tilde{\Sigma} = (1 - \delta)\hat{\Sigma} + \delta I, \qquad \delta \in (0, 1),$$

where I is the $N \times N$ identity matrix. Shrinkage reduces sampling variance and improves out-of-sample stability relative to the raw sample covariance matrix (Ledoit and Wolf, 2004).

Continuous mean–variance optimization. Given estimates $(\hat{\mu}, \tilde{\Sigma})$, the classical long-only mean–variance optimization problem is

$$\min_{w \in \mathbb{R}^N} \ \frac{\lambda}{2} w^\top \tilde{\Sigma} w - \hat{\mu}^\top w \quad \text{s.t.} \quad \mathbf{1}^\top w = 1, \quad w_i \geq 0.$$

Here:

- $w \in \mathbb{R}^N$ is the portfolio weight vector,
- $\lambda > 0$ is the risk-aversion parameter,
- $\mathbf{1}$ denotes the vector of ones.

The quadratic term $w^\top \tilde{\Sigma} w$ represents portfolio variance, while $\hat{\mu}^\top w$ represents expected return. Under positive definiteness of $\tilde{\Sigma}$, the objective is strictly convex and admits a unique solution.

It is well known that unconstrained mean–variance optimization can be extremely sensitive to estimation error, often leading to unstable weights (Merton, 1980). Interestingly, imposing constraints such as non-negativity can improve out-of-sample performance by implicitly regularizing the optimization problem (Jagannathan and Ma, 2003).

Cardinality-constrained portfolio selection. In many practical applications, portfolios are restricted to contain exactly K assets with $K \ll N$. Introduce binary inclusion variables

$$x_i \in \{0, 1\}, \quad i = 1, \ldots, N,$$

and impose the exact cardinality condition

$$\sum_{i=1}^{N} x_i = K.$$

The feasible set is therefore

$$\mathcal{X}_K = \left\{ x \in \{0, 1\}^N : \sum_{i=1}^{N} x_i = K \right\}.$$

The full cardinality-constrained mean–variance problem can be written as the mixed-integer quadratic program

$$\min_{w,x} \ \frac{\lambda}{2} w^\top \tilde{\Sigma} w - \hat{\mu}^\top w \quad \text{s.t.} \quad \mathbf{1}^\top w = 1, \quad w_i \geq 0, \quad w_i = 0 \text{ if } x_i = 0, \quad x \in \mathcal{X}_K.$$

This problem is NP-hard in general due to the combinatorial nature of the selection vector x (Chang et al., 2000; Jobst et al., 2001).

Two-stage decomposition. To make the problem computationally tractable, the present framework adopts a two-stage decomposition:

1. **Subset selection stage:** determine $x \in \mathcal{X}_K$ by minimizing a binary quadratic energy function.
2. **Conditional allocation stage:** compute optimal continuous weights w restricted to the selected subset.

The subset selection stage minimizes

$$E(x) = \beta\, x^\top \tilde{\Sigma} x - \alpha\, \hat{\mu}^\top x, \quad x \in \mathcal{X}_K,$$

where $\alpha, \beta > 0$ regulate the trade-off between covariance exposure and mean return contributions.

Expanding the quadratic term,

$$x^\top \tilde{\Sigma} x = \sum_{i=1}^{N} \sum_{j=1}^{N} x_i x_j \tilde{\Sigma}_{ij},$$

which aggregates pairwise covariance interactions among selected assets.

The linear term

$$\hat{\mu}^\top x = \sum_{i=1}^{N} \hat{\mu}_i x_i$$

rewards assets with higher historical mean returns.

Importantly, this formulation preserves the economic structure of the mean–variance paradigm while operating in a discrete domain. The energy function does not directly maximize the Sharpe ratio

$$S(w) = \frac{\hat{\mu}^\top w - r_f}{\sqrt{w^\top \tilde{\Sigma} w}},$$

originally introduced by Sharpe (1966), but it incorporates the same sufficient statistics $(\hat{\mu}, \tilde{\Sigma})$ and encodes a linearized risk–return trade-off compatible with quadratic binary optimization.

The continuous Sharpe-maximizing allocation is then performed conditionally in the second stage.

2.3 Estimation Error and Covariance Regularization

Mean–variance optimization is known to be highly sensitive to estimation error. In finite samples, both expected returns and covariance matrices are subject to sampling noise, which may significantly distort optimal allocations.

Estimation risk in mean–variance optimization. Let μ and Σ denote the true population mean vector and covariance matrix, and let $(\hat{\mu}, \hat{\Sigma})$ denote their sample estimates computed from T_{tr} observations.

Even when returns are independently and identically distributed, the estimation error of $\hat{\mu}$ scales as $O(T_{\text{tr}}^{-1/2})$, while the optimization step amplifies this noise through matrix inversion. As first emphasized by Michaud (1989), classical mean–variance optimization may produce extreme and unstable weights that reflect estimation noise rather than genuine economic structure.

The issue becomes particularly severe when the cross-sectional dimension N is large relative to T_{tr}. In such settings, the sample covariance matrix may be poorly conditioned, and small perturbations in inputs can generate large changes in optimal weights.

Empirically, simple strategies such as equal weighting have been shown to perform surprisingly well out-of-sample, often outperforming naive mean–variance allocations when estimation error is not properly controlled (DeMiguel et al., 2009). These findings highlight the importance of regularization and robust estimation techniques in practical portfolio construction.

Covariance shrinkage. To improve numerical stability and mitigate sampling variability, the covariance matrix is regularized via shrinkage toward the identity matrix:

$$\tilde{\Sigma} = (1-\delta)\hat{\Sigma} + \delta I, \quad \delta \in (0,1).$$

Here:

- $\hat{\Sigma}$ is the sample covariance matrix,
- I is the identity matrix,
- δ controls the intensity of shrinkage.

When $\delta = 0$, the estimator coincides with the sample covariance. When $\delta = 1$, the covariance matrix reduces to a scaled identity, implying zero cross-asset correlation. Intermediate values interpolate between these two extremes.

Shrinkage reduces estimator variance at the cost of introducing a small bias. In high-dimensional settings, this bias–variance trade-off often leads to superior out-of-sample performance compared to the raw sample covariance matrix (Ledoit and Wolf, 2004). In particular, shrinkage improves conditioning and ensures positive definiteness, which is crucial for both continuous quadratic optimization and binary quadratic encoding.

Implications for the discrete framework. In the present thesis, estimation error affects both stages of the methodology:

1. The binary quadratic energy function

$$E(x) = \beta\, x^\top \tilde{\Sigma} x - \alpha\, \hat{\mu}^\top x$$

   directly depends on estimated moments.

2. The continuous allocation stage within the selected subset also relies on $(\hat{\mu}, \tilde{\Sigma})$.

Errors in $\hat{\mu}$ influence the linear return component of the energy, while errors in $\tilde{\Sigma}$ affect both covariance interactions in the discrete stage and variance estimation in the continuous stage.

Shrinkage therefore plays a dual role:

- It stabilizes the continuous allocation problem.
- It smooths the combinatorial energy landscape explored by simulated annealing.

From a structural perspective, regularization reduces spurious extreme covariance estimates that might otherwise distort the quadratic interaction term $x^\top \tilde{\Sigma} x$ and create artificial local minima in the discrete search space.

While more sophisticated shrinkage schemes exist (e.g., optimal shrinkage intensity estimators), the present framework adopts a parsimonious and transparent regularization approach, prioritizing stability and interpretability over aggressive parameter tuning.

2.4 Continuous Allocation within the Selected Subset

The binary quadratic stage determines a subset of assets

$$\mathcal{S} = \{i \in \{1, \dots, N\} : x_i = 1\}, \quad |\mathcal{S}| = K.$$

Conditional on this subset, the portfolio weights are determined through a continuous mean–variance optimization restricted to the selected assets.

Restriction of moment estimators. Let $\hat{\mu} \in \mathbb{R}^N$ and $\tilde{\Sigma} \in \mathbb{R}^{N \times N}$ denote the full-universe estimators computed over the training window.

Define:

$$\hat{\mu}_{\mathcal{S}} \in \mathbb{R}^K \quad \text{and} \quad \tilde{\Sigma}_{\mathcal{S}} \in \mathbb{R}^{K \times K}$$

as the subvector and submatrix obtained by restricting $\hat{\mu}$ and $\tilde{\Sigma}$ to the indices in $\mathcal{S}$.

Formally, if $I_{\mathcal{S}}$ denotes the index set corresponding to selected assets, then

$$\hat{\mu}_{\mathcal{S}} = (\hat{\mu}_i)_{i \in I_{\mathcal{S}}}, \qquad \tilde{\Sigma}_{\mathcal{S}} = (\tilde{\Sigma}_{ij})_{i,j \in I_{\mathcal{S}}}.$$

All subsequent allocation decisions are based exclusively on these restricted moment estimates.

Conditional mean–variance optimization. The continuous allocation problem within the subset is formulated as:

$$\min_{w \in \mathbb{R}^K} \ \frac{\lambda}{2} w^\top \tilde{\Sigma}_{\mathcal{S}} w - \hat{\mu}_{\mathcal{S}}^\top w \quad \text{s.t.} \quad \mathbf{1}^\top w = 1, \quad w_i \geq 0.$$

Here:

- $w \in \mathbb{R}^K$ denotes the weight vector over selected assets,
- $\lambda > 0$ is the risk-aversion parameter,
- $\mathbf{1}$ is a K-dimensional vector of ones.

The quadratic term represents portfolio variance:

$$w^\top \tilde{\Sigma}_{\mathcal{S}} w = \sum_{i=1}^{K} \sum_{j=1}^{K} w_i w_j \tilde{\Sigma}_{ij},$$

while the linear term represents expected return:

$$\hat{\mu}_{\mathcal{S}}^\top w = \sum_{i=1}^{K} \hat{\mu}_i w_i.$$

Because $\tilde{\Sigma}_{\mathcal{S}}$ is positive definite after shrinkage, the problem is strictly convex and admits a unique solution.

Relation to Sharpe ratio maximization. An alternative formulation would directly maximize the in-sample Sharpe ratio:

$$\max_{w} \frac{\hat{\mu}_{\mathcal{S}}^{\top} w - r_f}{\sqrt{w^{\top} \tilde{\Sigma}_{\mathcal{S}} w}}.$$

However, Sharpe maximization under long-only and budget constraints leads to a non-linear fractional programming problem.

Under mild regularity conditions, quadratic utility maximization with an appropriate choice of λ is equivalent to selecting a point on the efficient frontier (Markowitz, 1952). Therefore, the framework adopts quadratic utility as the computationally convenient and economically interpretable allocation criterion.

Operationally, a grid of λ values is evaluated within each training window. The value that maximizes the in-sample Sharpe ratio within the selected subset is chosen, and the corresponding weights are carried forward to the test period.

Role of the two-stage structure. The separation between discrete subset selection and continuous allocation serves two purposes:

1. It reduces the dimensionality of the continuous optimization problem from N to K.
2. It isolates covariance-aware asset selection from weight allocation.

The binary quadratic stage determines which covariance interactions are admitted into the portfolio. The continuous stage then optimally allocates capital among the selected assets.

This two-layer architecture preserves the economic logic of mean–variance theory while making the combinatorial selection problem computationally manageable.

2.5 Binary Quadratic Optimization and Simulated Annealing

The subset selection stage is formulated as a constrained binary quadratic optimization problem and solved via simulated annealing.

Binary quadratic energy formulation. Let $x \in \{0,1\}^N$ denote the binary selection vector satisfying the exact cardinality constraint

$$\sum_{i=1}^{N} x_i = K.$$

The discrete selection objective is defined as

$$E(x) = \beta\, x^\top \tilde{\Sigma} x - \alpha\, \hat{\mu}^\top x, \quad x \in \mathcal{X}_K.$$

Here:

- $\tilde{\Sigma}$ is the shrinkage covariance estimator,
- $\hat{\mu}$ is the vector of historical mean returns,
- $\alpha > 0$ controls the weight assigned to return contributions,
- $\beta > 0$ controls the weight assigned to covariance interactions.

The quadratic term expands as

$$x^\top \tilde{\Sigma} x = \sum_{i=1}^{N} \sum_{j=1}^{N} x_i x_j \tilde{\Sigma}_{ij},$$

which aggregates all pairwise covariance contributions among selected assets. Because $x_i x_j = 1$ only when both assets are included, this term captures the total covariance exposure within the chosen subset.

The linear term

$$\hat{\mu}^\top x = \sum_{i=1}^{N} \hat{\mu}_i x_i$$

rewards inclusion of assets with higher historical mean returns.

Minimizing $E(x)$ therefore selects subsets that balance high-return assets against covariance interactions. The formulation belongs to the class of Quadratic Unconstrained Binary Optimization (QUBO) models (Lucas, 2014), except that here the cardinality constraint is enforced explicitly rather than absorbed into the objective through penalty terms.

Exact cardinality enforcement. Many QUBO implementations incorporate constraints via quadratic penalty terms of the form

$$\gamma \left(\sum_{i=1}^{N} x_i - K \right)^2,$$

with a large penalty parameter γ. However, penalty calibration may distort the energy landscape and introduce scale sensitivity.

Instead, the present framework enforces the cardinality constraint directly through the design of the search space. The algorithm operates exclusively on

$$\mathcal{X}_K = \left\{ x \in \{0,1\}^N : \sum_{i=1}^{N} x_i = K \right\}.$$

Feasible neighbors are generated through swap moves: one selected asset ($x_i = 1$) is replaced by one unselected asset ($x_j = 0$), preserving $\sum_i x_i = K$ at every iteration. This ensures that all candidate states remain feasible by construction and that the objective reflects purely economic trade-offs.

Simulated annealing dynamics. The discrete energy landscape defined by $E(x)$ contains $\binom{N}{K}$ feasible states. Exhaustive enumeration is infeasible for realistic (N, K) combinations.

Simulated annealing (SA) provides a stochastic search mechanism capable of escaping local minima (Kirkpatrick et al., 1983). Starting from an initial feasible configuration $x^{(0)}$, the algorithm iteratively proposes local moves.

At iteration t, let x denote the current state and x' a candidate neighbor obtained through a swap. Define the energy difference

$$\Delta E = E(x') - E(x).$$

The candidate state is accepted according to the Metropolis rule:

$$\Pr(\text{accept } x') = \begin{cases} 1, & \Delta E \leq 0, \\ \exp\left(-\frac{\Delta E}{T_t}\right), & \Delta E > 0, \end{cases}$$

where $T_t > 0$ is the temperature at iteration t.

At high temperature, the algorithm accepts worse moves with non-negligible probability, facilitating exploration of the state space. As the temperature decreases according to a cooling schedule, the search becomes increasingly greedy and converges toward low-energy configurations.

Under sufficiently slow cooling schedules, SA converges in probability to the global optimum in the limit (Aarts and Korst, 1988). In practical implementations, however, computational budgets are finite, and performance depends on iteration counts, temperature decay, and neighborhood structure.

Initialization schemes. Two initialization schemes are evaluated:

- **Random initialization:** K assets are selected uniformly at random.
- **Warm initialization:** the K assets with the highest historical mean returns $\hat{\mu}_i$ are selected.

The warm start is coherent with the linear component of the energy function and provides an economically interpretable baseline configuration.

The empirical analysis later investigates whether initialization materially affects the quality of final portfolios or whether the annealing dynamics converge toward similar regions of the discrete landscape irrespective of starting point.

2.6 Time-to-Target Calibration

Heuristic optimization methods such as simulated annealing depend critically on iteration budgets and stopping rules. Selecting these parameters arbitrarily may lead either to excessive computational cost or insufficient solution quality.

To calibrate the computational effort in a statistically principled manner, this thesis adopts a *Time-to-Target* (TTT) framework, commonly used in the analysis of stochastic local search algorithms (Aiex et al., 2007; Hoos and Stützle, 2004).

Definition of the target energy. Let $E(x)$ denote the binary quadratic energy function defined in Section 2.5. Consider a fixed training dataset and a given problem instance.

Let $\{E_{\min}^{(r)}\}_{r=1}^{R_0}$ denote the best energies obtained across R_0 independent runs of the algorithm under a baseline configuration (e.g., random initialization and sufficiently large iteration count).

Define the target energy level $E^\star$ as a quantile of this empirical distribution:

$$E^\star = \text{Quantile}_q\left(\{E_{\min}^{(r)}\}_{r=1}^{R_0}\right),$$

where $q \in (0,1)$ is typically chosen close to zero (e.g., the 10th percentile). The target therefore represents a high-quality but attainable energy threshold.

Success probability estimation. For a given number of annealing iterations N_s, define a Bernoulli success indicator:

$$I_r = \begin{cases} 1, & \text{if } E_{\min}^{(r)} \le E^\star, \\ 0, & \text{otherwise,} \end{cases}$$

where r indexes independent runs with the same N_s.

Let R denote the number of independent repetitions performed at iteration budget N_s. The empirical success probability is estimated as

$$\hat{p}(N_s) = \frac{1}{R} \sum_{r=1}^{R} I_r.$$

Under independence across runs, the number of successes follows a Binomial$(R, \hat{p})$ distribution.

Equivalent computational effort. The key question is: how many independent runs are required to reach the target energy with confidence level η (e.g., $\eta = 0.99$)?

The probability of observing at least one success in R_{req} independent runs is

$$1 - (1 - \hat{p})^{R_{\text{req}}}.$$

Setting this probability equal to η yields

$$1 - (1 - \hat{p})^{R_{\text{req}}} = \eta.$$

Solving for R_{req} gives

$$R_{\text{req}} = \frac{\log(1 - \eta)}{\log(1 - \hat{p})}.$$

The total equivalent computational effort, expressed in number of elementary iterations, is then

$$T_{\text{equiv}} = R_{\text{req}} \cdot N_s.$$

This metric allows different configurations (e.g., warm vs random initialization) to be compared on a unified scale that incorporates both success probability and per-run iteration cost.

Interpretation in the present framework. The TTT procedure serves two purposes:

1. It provides a statistically grounded method for selecting an iteration budget N_s.

2. It characterizes the effective difficulty of the discrete energy landscape.

Rather than tuning N_s heuristically, the framework identifies values that achieve a high probability of reaching competitive energy levels. This prevents overfitting computational parameters to specific windows and enhances reproducibility.

Importantly, TTT calibration is performed using the full dataset to obtain stable estimates of optimization behavior across heterogeneous market regimes. The selected iteration budget is then held fixed during block-forward evaluation.

This separation between calibration and out-of-sample testing ensures that computational tuning does not introduce forward-looking bias.

2.7 Block-Forward Backtesting Design

To evaluate the economic relevance of the proposed framework, portfolio performance is assessed through a rolling block-forward backtesting procedure. This design ensures strict temporal separation between estimation and evaluation, thereby preventing look-ahead bias and mitigating backtest overfitting concerns (Bailey et al., 2014).

Rolling window structure. Let the full sample span from $t = 1$ to $t = T$. The backtesting procedure partitions the data into overlapping training and testing blocks.

For each window $k = 1, \ldots, K_{\text{win}}$:

- A training segment of fixed length T_{tr} is used to estimate $\hat{\mu}$ and $\tilde{\Sigma}$ and to perform subset selection and continuous allocation.

- The immediately subsequent test segment of length T_{test} is used to evaluate out-of-sample performance.

Formally, if $\mathcal{T}_k^{\text{tr}}$ and $\mathcal{T}_k^{\text{test}}$ denote the training and testing index sets for window k, then

$$\max \mathcal{T}_k^{\text{tr}} < \min \mathcal{T}_k^{\text{test}},$$

ensuring that all model inputs are strictly based on past information.

After each test period, the window is shifted forward by a fixed step (e.g., one year), and the process is repeated.

Out-of-sample portfolio returns. Let w_k denote the weight vector obtained from the training window k, restricted to the selected subset.

For each test-period day $t \in \mathcal{T}_k^{\text{test}}$, the realized portfolio return is

$$r_{p,t} = w_k^\top r_t,$$

where r_t is the vector of daily asset returns.

Out-of-sample performance metrics are computed for each window separately and then aggregated across windows.

Performance metrics. For each test window, the following statistics are computed:

Annualized return.

$$\widehat{R}_{\text{ann}} = \exp\big(252 \cdot \bar{r}_p\big) - 1, \quad \bar{r}_p = \frac{1}{T_{\text{test}}} \sum_{t \in \mathcal{T}_k^{\text{test}}} r_{p,t}.$$

Annualized volatility.

$$\widehat{\sigma}_{\text{ann}} = \sqrt{252} \cdot \sqrt{\frac{1}{T_{\text{test}} - 1} \sum_{t \in \mathcal{T}_k^{\text{test}}} (r_{p,t} - \bar{r}_p)^2}.$$

Sharpe ratio.

$$\widehat{S} = \frac{\widehat{R}_{\text{ann}} - r_f}{\widehat{\sigma}_{\text{ann}}},$$

where r_f denotes the annualized risk-free rate (set to zero for simplicity in the empirical im-

plementation).

The statistical properties of Sharpe ratios under finite samples are non-trivial (Lo, 2002); therefore, results are interpreted comparatively rather than through formal hypothesis testing.

Maximum drawdown. Maximum drawdown is computed as

$$\mathrm{MDD} = \max_{t \le s} \left(\frac{V_t - V_s}{V_t} \right),$$

where V_t denotes the cumulative portfolio value at time t.

Benchmark strategies. To isolate the incremental value of discrete covariance-aware selection, the following benchmarks are evaluated within each window:

- **Equal-weight portfolio (EW):** full-universe portfolio with $w_i = 1/N$.
- **Top-K mean-return portfolio:** select the K assets with highest $\hat{\mu}_i$ and perform continuous allocation within this subset.

Comparing these benchmarks with the QUBO-based selection allows assessment of whether combinatorial covariance modeling provides incremental benefit beyond simple return ranking or naive diversification.

Aggregation across windows. Let $\widehat{S}_k$ denote the Sharpe ratio in window k. Aggregate statistics (mean, median, dispersion) are computed across windows:

$$\bar{S} = \frac{1}{K_{\mathrm{win}}} \sum_{k=1}^{K_{\mathrm{win}}} \widehat{S}_k.$$

This cross-window aggregation provides a distributional view of out-of-sample performance, rather than relying on a single backtest trajectory.

Overall, the block-forward framework provides a realistic and robust evaluation protocol that mirrors practical re-estimation and rebalancing cycles in systematic asset management.

CHAPTER 3

EMPIRICAL RESULTS

3.1 Optimization Landscape and Time-to-Target Analysis

This section evaluates the stochastic convergence behavior of the simulated annealing (SA) algorithm under different iteration budgets and initialization schemes.

Let $E_{\text{best}}(N_s)$ denote the minimum energy reached after N_s iterations. Following the Time-to-Target (TTT) framework introduced in Chapter 2, a target energy threshold $E^\star$ is defined as a lower quantile of a high-budget baseline distribution of best energies.

For a fixed iteration budget N_s, the empirical success probability is defined as

$$\hat{p}(N_s) = \mathbb{P}\left(E_{\text{best}}(N_s) \leq E^\star\right).$$

Empirical success probability. The estimated success probability curves show that:

- At low iteration budgets (e.g., $N_s = 5{,}000$), both warm and random initializations exhibit near-zero success probability.
- At intermediate budgets (e.g., $N_s = 10{,}000$), the probability of reaching the target becomes positive but remains modest.
- As N_s increases further, both initialization schemes converge toward higher success probabilities.

Importantly, differences between warm-start and random-start configurations are limited. While warm initialization may exhibit a slightly higher success probability at some intermediate iteration budgets, the gap narrows as computational resources increase.

This suggests that the discrete energy landscape, although non-trivial, does not display strong sensitivity to initialization. The annealing dynamics tend to explore sufficiently large regions of

the feasible space when the iteration budget is moderate to high.

Figure 1 summarizes the estimated success probability curves for warm and random initialization.

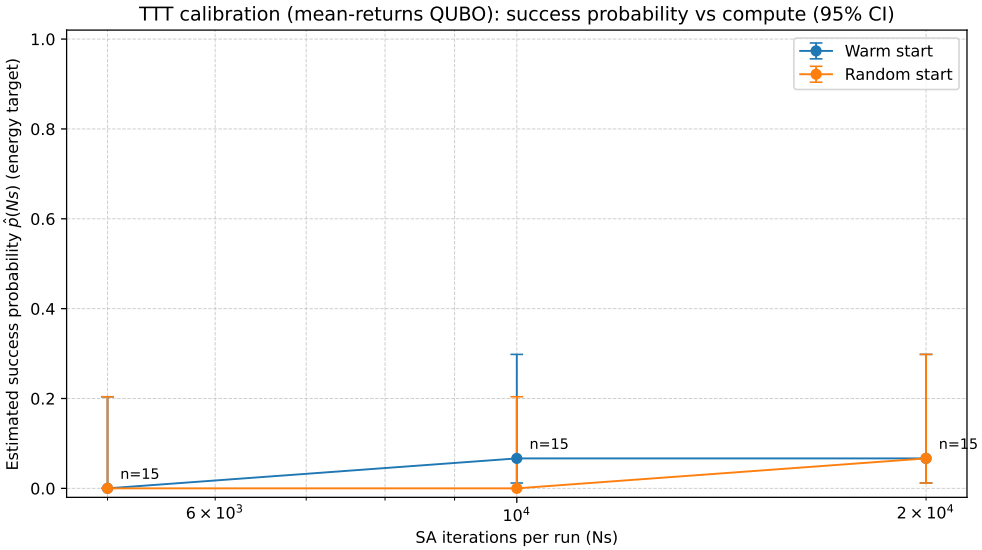

Figure 1: TTT calibration (mean-returns QUBO): estimated success probability $\hat{p}(N_s)$ of reaching the target energy as a function of the iteration budget N_s, with 95% confidence intervals, for warm and random initialization.

Equivalent computational effort. For a confidence level η (e.g., $\eta = 0.99$), the required number of independent runs to achieve at least one success is

$$R_{\text{req}}(N_s) = \left\lceil \frac{\log(1-\eta)}{\log(1-\hat{p}(N_s))} \right\rceil .$$

The equivalent computational effort is defined as

$$T_{\text{equiv}}(N_s) = N_s \cdot R_{\text{req}}(N_s).$$

Empirically, the equivalent effort decreases as N_s increases, reflecting the monotonic rise in success probability. However, the difference between warm and random initialization remains modest across the tested range.

To translate success probabilities into a comparable compute metric, Figure 2 reports the equivalent computational effort required to achieve $\eta = 0.99$ confidence.

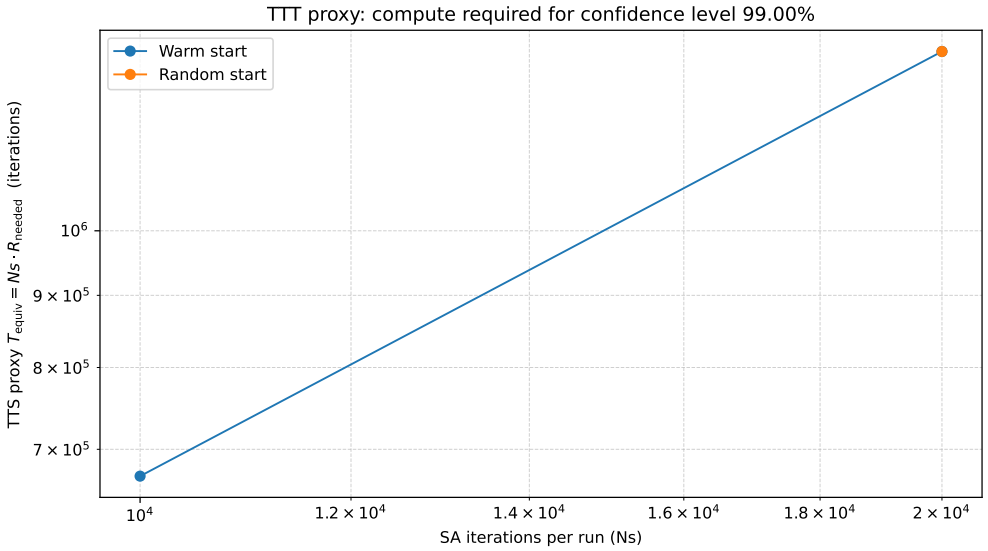

Figure 2: TTT proxy: equivalent computational effort $T_{\text{equiv}}(N_s) = N_s \cdot R_{\text{req}}(N_s)$ required to reach the target energy with confidence level $\eta = 0.99$.

Interpretation. Two conclusions emerge.

First, the combinatorial energy landscape is non-trivial: success probability increases gradually with iteration count, indicating the presence of multiple local minima and a structured search space.

Second, structured warm initialization based on highest mean returns does not produce a persistent and economically large computational advantage. Although it may accelerate convergence marginally in constrained-compute regimes, the annealing dynamics eventually converge toward similar solution quality under both initialization schemes.

This evidence supports the view that the simulated annealing procedure exhibits limited sensitivity to starting configuration and that solution quality is primarily driven by the overall iteration budget rather than by initialization alone.

3.2 Distribution of Out-of-Sample Sharpe Ratios

This section evaluates the distribution of out-of-sample (OOS) Sharpe ratios across the $W = 19$ block-forward windows described in Chapter 2. Each window produces one Sharpe ratio observation per strategy. The cross-window distribution therefore captures regime sensitivity and performance stability over time rather than relying on a single backtest trajectory.

Table 1 reports descriptive statistics for:

- Simulated Annealing with warm initialization (Warm SA),
- Simulated Annealing with random initialization (Random SA),
- Equal-Weight benchmark (EW),
- Top-K mean-return ranking followed by continuous allocation.

Table 1: Out-of-sample Sharpe ratio summary statistics across $W = 19$ rolling windows

Strategy	Mean	Median	Std	Min	Max
Warm SA	0.624	0.620	0.566	-0.303	1.816
Random SA	0.627	0.564	0.608	-0.429	1.497
Equal Weight	0.473	0.569	0.417	-0.341	1.503
Top-K	0.567	0.574	0.502	-0.272	1.640

Detailed window-level Sharpe ratios for all strategies are reported in Table 2 (Figures and Tables section).

Figure 3 visualizes the cross-window dispersion and central tendency of out-of-sample Sharpe ratios for all strategies.

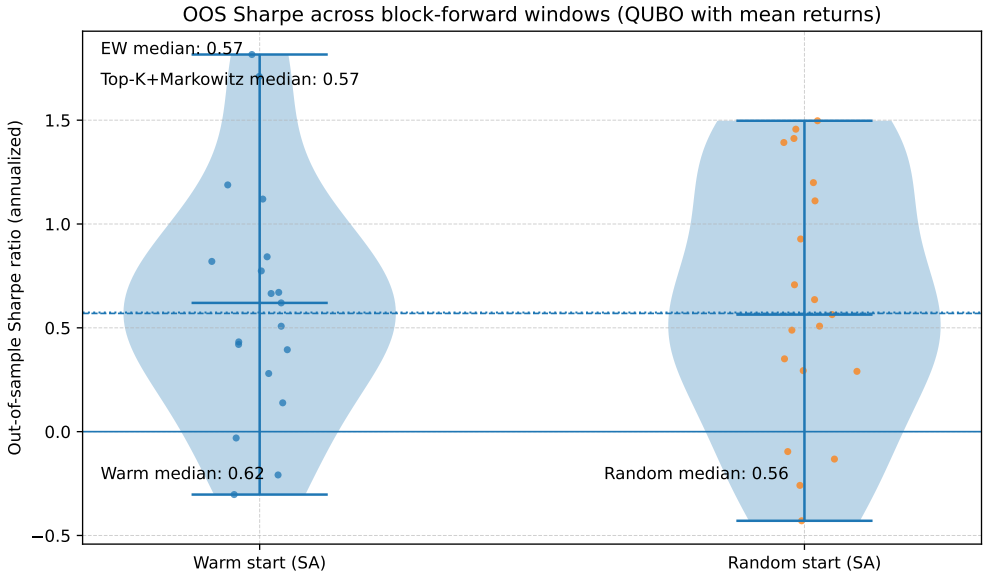

Figure 3: Distribution of annualized out-of-sample Sharpe ratios across block-forward windows for Warm SA, Random SA, Equal Weight, and Top-K benchmark portfolios.

Several observations emerge.

First, both SA variants achieve higher *mean* Sharpe ratios than the Equal-Weight benchmark. The average Sharpe is approximately 0.62 for both warm and random initialization, compared to 0.47 for Equal Weight and 0.57 for the Top-K benchmark. This indicates that covariance-aware combinatorial selection can provide incremental risk-adjusted performance relative to naive diversification.

Second, the difference between warm and random initialization is economically small. Their mean Sharpe ratios differ by less than 0.01, and their dispersion levels are similar. This confirms that the limited computational advantage observed in Section 3.1 does not translate into materially different financial performance.

Figure 4 provides a paired comparison across identical windows, highlighting the limited sensitivity to initialization.

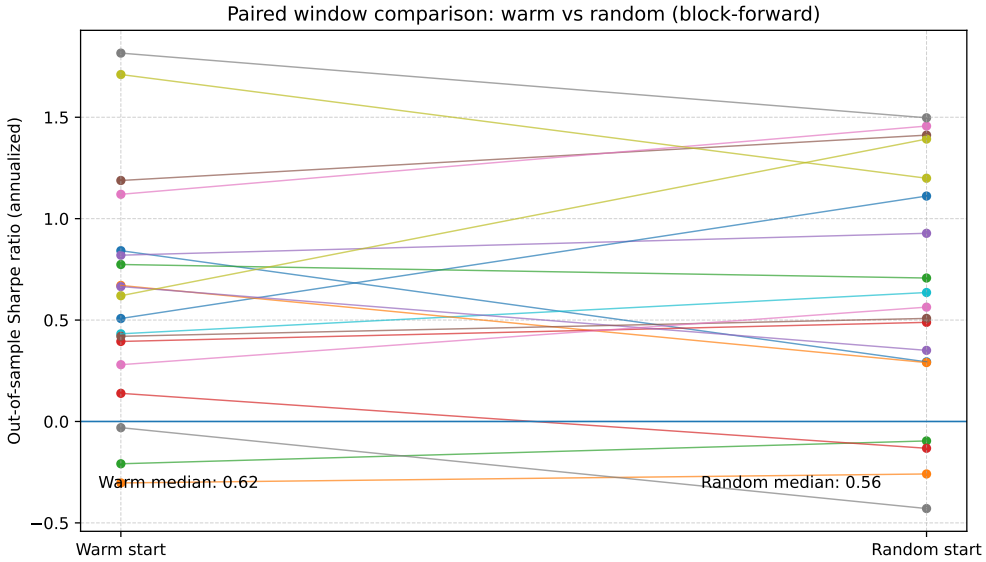

Figure 4: Paired window comparison of out-of-sample Sharpe ratios: Warm SA versus Random SA. Each point corresponds to one block-forward window, enabling a direct within-window comparison.

Third, dispersion across windows is substantial for all strategies. Standard deviations range from 0.42 (Equal Weight) to approximately 0.61 (Random SA). This highlights strong regime dependence: performance varies considerably depending on the specific training–testing split.

Fourth, medians differ from means in several cases. Equal Weight exhibits a median Sharpe (0.569) higher than its mean (0.473), indicating occasional large negative windows that reduce the average. In contrast, Warm SA shows close alignment between mean and median, suggesting more symmetric window-level performance.

For additional clarity, Figure 5 provides an alternative boxplot representation of the same cross-window Sharpe distributions, emphasizing median alignment and dispersion differences.

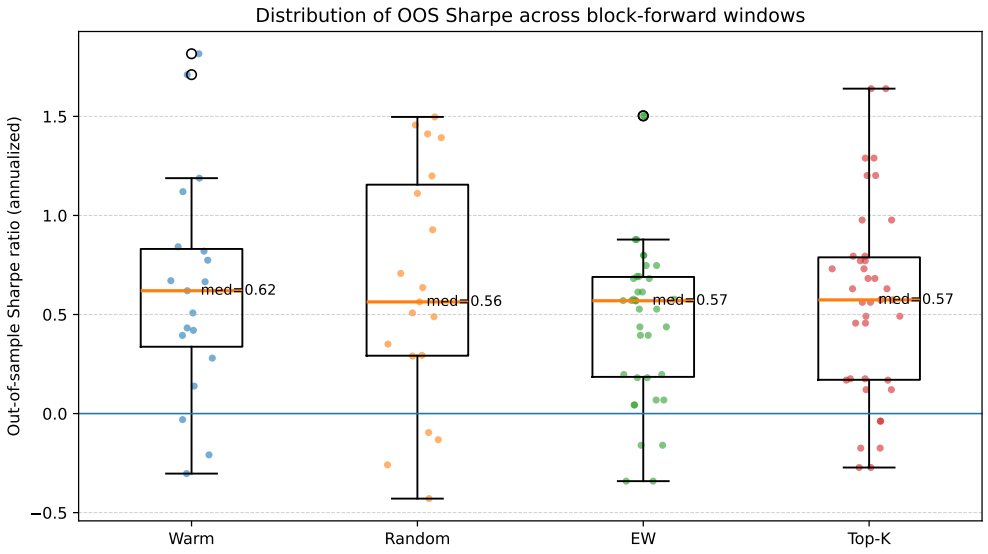

Figure 5: Alternative visualization of out-of-sample Sharpe distributions across block-forward windows. The boxplot emphasizes median differences and interquartile dispersion between strategies.

To assess robustness across regimes, Figure 6 reports dispersion measures of Sharpe ratios across windows.

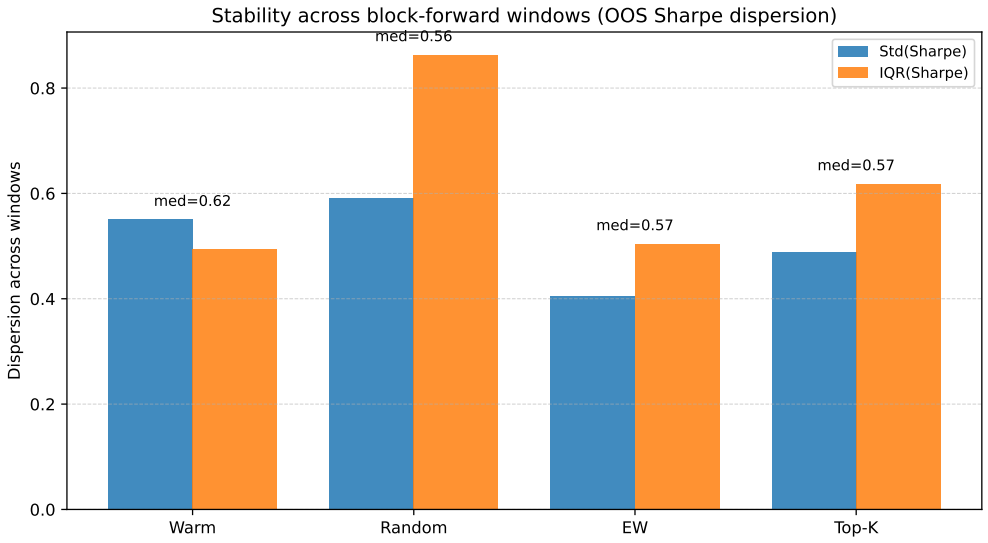

Figure 6: Stability across block-forward windows: dispersion of out-of-sample Sharpe ratios measured by standard deviation and interquartile range (IQR) for each strategy.

Overall, the distributional evidence suggests that the QUBO-based selection framework is competitive in average risk-adjusted performance, but does not dominate benchmarks uniformly

across all windows. The advantage appears statistical rather than deterministic and remains sensitive to market regime variation.

3.3 Equity Curve Dynamics and Drawdown Analysis

While Section 3.2 studied window-level Sharpe distributions, this section evaluates long-horizon performance by concatenating out-of-sample (OOS) test-block returns across the full block-forward evaluation period.

Let V_t denote the cumulative portfolio value at time t, obtained by sequentially compounding the OOS returns from each test block. The resulting series represents a pseudo-OOS equity curve, constructed under repeated re-estimation and rebalancing at the boundary of each rolling window.

Figure 7 reports the concatenated out-of-sample equity curves across the full evaluation horizon.

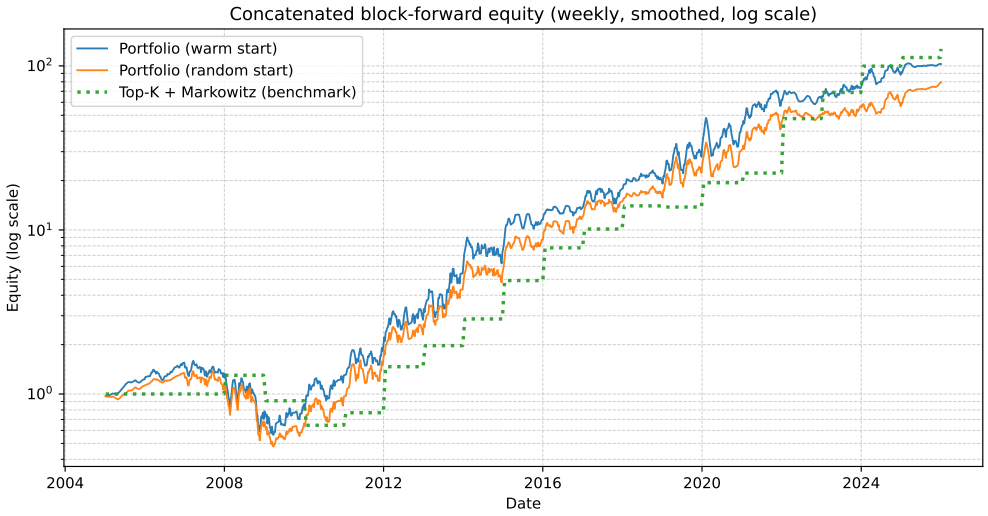

Figure 7: Concatenated out-of-sample equity curves (weekly, smoothed) on a log scale for Warm SA, Random SA, and the Top-K benchmark.

Cumulative performance. Over the full evaluation horizon (2005–2025), both SA implementations produce substantial wealth accumulation. The total cumulative return (relative to an initial value normalized to 1) is:

- Warm SA: $V_T - 1 \approx 100.51$ (i.e., $\approx +10{,}051\%$),
- Random SA: $V_T - 1 \approx 78.95$ (i.e., $\approx +7{,}895\%$),
- Equal Weight: $V_T - 1 \approx 64.65$ (i.e., $\approx +6{,}465\%$).

Consistent with the distributional results in Section 3.2, SA-based portfolios are competitive relative to EW in long-horizon compounding. Warm SA exhibits higher terminal wealth than Random SA in this specific concatenation, even though their average Sharpe ratios across windows are nearly identical. This illustrates that long-run compounded outcomes can be sensitive to the temporal ordering of favorable and unfavorable windows, beyond what is captured by window-averaged Sharpe ratios.

Maximum drawdown. To assess downside risk and path instability, maximum drawdown (MDD) is computed as the minimum of the drawdown process:

$$\mathrm{DD}_t = \frac{V_t}{\max_{u \le t} V_u} - 1, \qquad \mathrm{MDD} = \min_t \mathrm{DD}_t.$$

Empirically, drawdowns are deep for all strategies:

- Warm SA: MDD $\approx -79.87\%$,
- Random SA: MDD $\approx -78.11\%$,
- Equal Weight: MDD $\approx -79.01\%$.

Therefore, although the strategies compound strongly over the long horizon, they also experience severe interim losses. This behavior is consistent with the broader empirical evidence that portfolio optimization procedures—especially under rolling estimation and regime shifts—can display pronounced path dependence and instability, even when average performance metrics appear favorable (DeMiguel et al., 2009; Chekhlov et al., 2005).

Downside risk is summarized through drawdown dynamics in Figure 8.

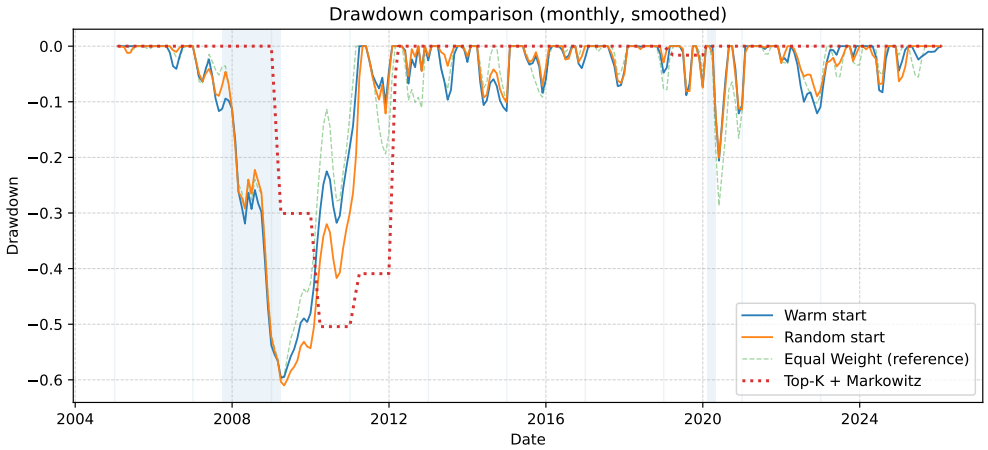

Figure 8: Drawdown comparison (monthly, smoothed) for Warm SA, Random SA, Equal Weight (reference), and Top-K benchmark portfolios.

Relative performance versus Equal Weight. A useful diagnostic is the relative wealth ratio

$$\mathrm{Rel}_t = \frac{V_t^{\mathrm{SA}}}{V_t^{\mathrm{EW}}}.$$

To highlight relative performance against a passive benchmark, Figure 9 reports wealth ratios versus Equal Weight.

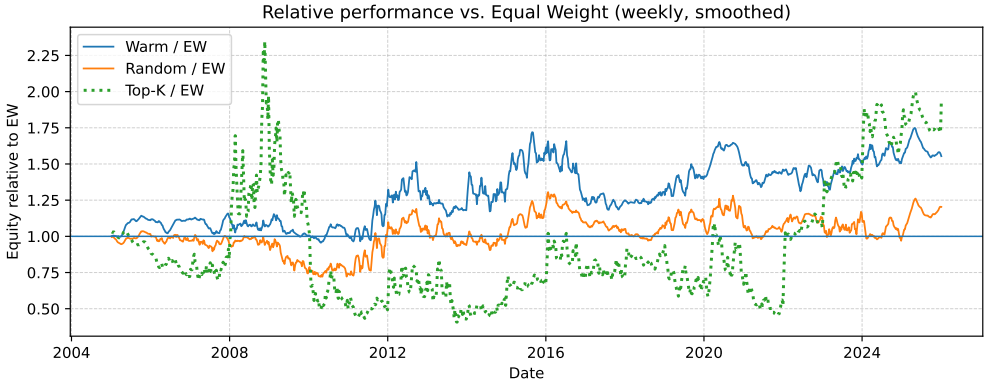

Figure 9: Relative wealth trajectories versus Equal Weight (weekly, smoothed): Warm/EW, Random/EW, and Top-K/EW. Values above 1 indicate outperformance relative to EW.

Values above one indicate outperformance relative to EW. The concatenated equity curves imply that both SA configurations spend extended periods above and below EW, reinforcing that outperformance is regime-dependent and that long-run advantage emerges through compound-

ing rather than through uniform dominance.

Overall, the equity-curve and drawdown evidence complements Section 3.2: SA-based discrete selection is competitive and can outperform simple benchmarks in long-horizon compounding, but performance remains highly path-sensitive and exposed to large drawdowns.

3.4 Energy Minimization and Sharpe Maximization: A Structural Disconnect

The empirical results reported in Sections 3.2 and 3.3 reveal a central structural feature of the proposed framework: minimizing the binary quadratic energy function does not guarantee maximization of the out-of-sample Sharpe ratio.

This section clarifies the source of this disconnect.

Energy function versus Sharpe ratio. The subset selection stage minimizes

$$E(x) = \beta\, x^\top \tilde{\Sigma} x - \alpha\, \hat{\mu}^\top x, \quad x \in \mathcal{X}_K.$$

By contrast, the Sharpe ratio of a portfolio with weights w is

$$S(w) = \frac{\hat{\mu}^\top w - r_f}{\sqrt{w^\top \tilde{\Sigma} w}}.$$

Several structural differences are immediately apparent:

1. The energy function operates on a *binary inclusion vector* x, whereas the Sharpe ratio is defined on continuous weights w.
2. The energy combines variance and expected return linearly, while the Sharpe ratio involves a *ratio* of return to volatility.
3. The quadratic term $x^\top \tilde{\Sigma} x$ aggregates raw covariance contributions without normalization by portfolio scale.

Therefore, the energy function can be interpreted as a linearized and discretized proxy for risk–return trade-offs, but not as a direct surrogate for Sharpe maximization.

Two-stage separation and loss of optimality equivalence. In classical mean–variance theory, maximizing quadratic utility

$$\hat{\mu}^\top w - \frac{\lambda}{2} w^\top \tilde{\Sigma} w$$

is equivalent to selecting a point on the efficient frontier. However, this equivalence holds within the continuous domain.

In the present framework, optimization proceeds in two stages:

1. Discrete subset selection via $E(x)$.
2. Continuous allocation conditional on the selected subset.

The subset that minimizes $E(x)$ need not coincide with the subset that would yield the highest Sharpe ratio after continuous allocation. Because the discrete stage does not internalize the full nonlinear structure of the Sharpe objective, a structural gap may arise between energy-optimal and Sharpe-optimal subsets.

Empirical manifestation. This disconnect is visible in the empirical results:

- The SA-based strategies improve average Sharpe relative to Equal Weight.
- However, they do not dominate across all windows.
- Initialization differences affect energy convergence but do not materially change Sharpe distributions.

These findings are consistent with the interpretation that energy minimization identifies subsets with favorable covariance–return structure in expectation, but realized Sharpe ratios remain sensitive to estimation error, regime shifts, and nonlinear weight adjustments in the second stage.

Figure 10 illustrates how windows with similar energy levels can map to different realized risk–return outcomes, reinforcing the structural disconnect between discrete energy minimization and Sharpe maximization.

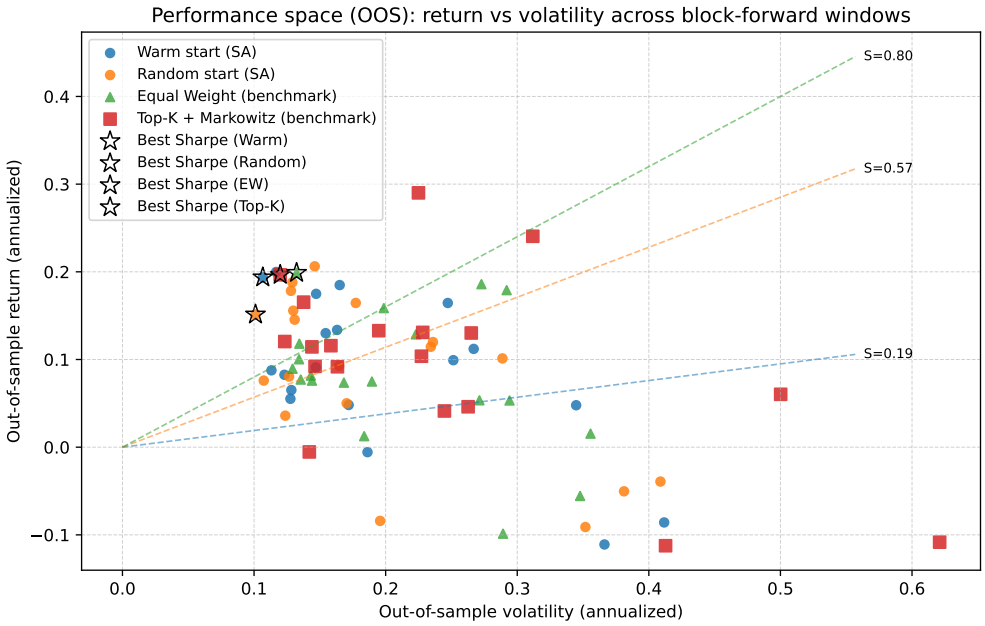

Figure 10: Out-of-sample performance space across block-forward windows: annualized return versus annualized volatility for SA (warm/random) and benchmark strategies. Dashed lines represent iso-Sharpe levels.

Implications. The structural disconnect does not invalidate the approach; rather, it clarifies its economic meaning.

The binary quadratic energy function should be interpreted as a computationally tractable proxy for covariance-aware subset selection. It provides a mechanism for incorporating pairwise risk interactions directly into a combinatorial framework. However, it is not a direct maximizer of ex-post Sharpe performance.

This distinction is important when interpreting empirical results. Improvements in energy do not translate mechanically into improvements in realized Sharpe ratios, especially in finite samples and under rolling estimation. Performance gains must therefore be evaluated empirically rather than inferred from the optimization objective alone.

CHAPTER 4

CONCLUSIONS

4.1 Summary of Findings and Theoretical Implications

This thesis investigated a cardinality-constrained portfolio construction framework based on Binary Quadratic Optimization combined with Simulated Annealing, calibrated through a Time-to-Target (TTT) methodology and evaluated via a rigorous block-forward backtesting design.

The empirical analysis was conducted on $N = 500$ U.S. equities from the CRSP US Total Market Index over the period 1997–2026, using rolling training windows and $W = 19$ out-of-sample evaluation blocks. The framework directly incorporates the vector of historical mean returns $\hat{\mu}$ into the linear component of the quadratic energy function, so that discrete subset selection is explicitly driven by the same first- and second-moment estimates employed in the subsequent continuous Markowitz allocation stage.

Several key findings emerge.

First, the combinatorial optimization problem exhibits a non-trivial energy landscape. The TTT analysis confirms that convergence toward high-quality subsets is probabilistic and sensitive to iteration budgets. Warm-start initialization based on top-K historical mean returns may improve convergence speed under moderate computational budgets, although the advantage narrows as iteration budgets increase. However, increasing iteration budgets reduces the performance gap between warm and random initialization, indicating that initialization affects convergence efficiency rather than final attainable energy levels.

Second, out-of-sample performance across rolling windows does not reveal systematic Sharpe dominance of the combinatorial strategies over classical benchmarks. While Simulated Annealing produces economically coherent subsets consistent with mean–variance principles, the Top-K Sharpe ranking benchmark achieves the highest average Sharpe ratio across windows. Overall, the evidence indicates that the proposed mean-return-based QUBO formulation provides stable performance across initialization schemes, while maintaining a transparent and economically in-

terpretable structure.

Third, concatenated equity curve analysis highlights pronounced regime dependence and substantial interim drawdowns for all active strategies. Although long-horizon cumulative performance of SA-based portfolios is economically meaningful, wealth accumulation is uneven and sensitive to market conditions. Improvements in discrete subset selection do not automatically translate into smoother intertemporal compounding.

From a theoretical perspective, the revised mean-returns formulation strengthens the internal coherence of the two-stage framework. The discrete energy function

$$E(x) = \beta x^\top \tilde{\Sigma} x - \alpha \hat{\mu}^\top x, \quad x \in \mathcal{X}_K,$$

now mirrors the classical quadratic mean–variance utility function in structure. The same moment estimates $(\hat{\mu}, \tilde{\Sigma})$ drive both discrete subset selection and continuous weight optimization. The lack of guaranteed Sharpe dominance therefore cannot be attributed to a mismatch between stages, but instead reflects a structural property of cardinality-constrained optimization.

Indeed, the classical equivalence between quadratic utility maximization and Sharpe maximization holds only in continuous convex domains without L_0 constraints (Markowitz, 1952). Once the feasible set becomes discrete and combinatorial, the mapping between energy minimization and Sharpe maximization breaks down. The observed gap between energy-optimal and Sharpe-optimal portfolios is thus intrinsic to the non-convex geometry of the problem rather than a deficiency of the algorithm.

The adoption of a hard-constrained Binary Quadratic Model further reinforces structural consistency. By enforcing exact cardinality without penalty calibration, the framework avoids distortions associated with heuristic penalty weights and ensures comparability across rolling windows.

Overall, the contribution of this thesis does not lie in demonstrating universal outperformance relative to classical heuristics. Rather, it lies in rigorously quantifying the computational, statistical, and structural properties of discrete mean–variance optimization under realistic iteration budgets. The results clarify the precise relationship between quadratic combinatorial selection

and realized out-of-sample risk-adjusted performance, providing a transparent benchmark for future developments in quantum-inspired asset allocation.

4.2 Limitations and Directions for Future Research

The results of this thesis should be interpreted within clearly defined structural and modeling boundaries. These limitations do not weaken the internal validity of the framework; rather, they clarify its scope and suggest natural directions for future research.

Structural properties of discrete mean–variance optimization. A first class of limitations arises from the intrinsic mathematical structure of cardinality-constrained portfolio selection. Once an L_0 constraint is imposed, the feasible set becomes discrete and non-convex. As discussed in Chapter 3, the classical equivalence between quadratic utility maximization and Sharpe maximization holds only in continuous convex domains (Markowitz, 1952). In discrete combinatorial domains, minimizing a quadratic energy function does not guarantee maximization of the Sharpe ratio.

This structural disconnect is not algorithm-specific. It is a mathematical property of the problem itself. Consequently, improvements in energy optimality cannot be expected to translate mechanically into superior realized risk-adjusted performance.

Estimation risk and moment instability. The framework relies on historical sample estimates of mean returns and covariance matrices. In finite samples, these estimators are subject to significant statistical error, especially in high-dimensional settings. Estimation noise may therefore dominate differences between combinatorial selection and simpler heuristic benchmarks.

Future research could explore the integration of shrinkage covariance estimators, Bayesian expected return estimates, or robust optimization techniques that explicitly account for parameter uncertainty. Such extensions would allow clearer separation between structural combinatorial effects and purely statistical estimation error.

Initialization structure and convergence robustness. The present implementation employs a warm-start configuration based on the top-K assets ranked by historical mean returns. This

choice preserves coherence between the initialization mechanism and the linear component of the quadratic energy function.

An interesting extension would involve testing alternative economically motivated initialization schemes, such as ranking assets by historical Sharpe ratios. This would allow examination of whether the annealing process converges toward similar low-energy subsets regardless of the initial configuration, thereby providing deeper insight into the degree of path dependence inherent in the discrete optimization landscape.

The objective of this thesis, however, was to isolate the structural properties of the mean–variance combinatorial formulation under a consistent moment-based setup. A systematic exploration of multiple initialization schemes represents a natural but distinct extension.

Dynamic rebalancing and transaction costs. The block-forward design adopted in this thesis approximates repeated single-period optimization. Portfolios are re-estimated at window boundaries and held constant within each test block. Transaction costs, turnover penalties, and dynamic rebalancing effects are not explicitly modeled.

Incorporating transition costs—for example through penalties on changes in portfolio composition across periods—would transform the framework into a multi-period allocation scheme. Such an extension would allow evaluation of whether combinatorial subset selection exhibits structural turnover advantages relative to ranking-based approaches.

Integration of predictive signals. The current formulation employs historical mean returns as proxies for expected returns, thereby maintaining a purely moment-based allocation structure. Replacing these inputs with forecast-based expected returns derived from predictive models would conceptually transform the framework into a conditional allocation mechanism.

Such an extension would require the specification and validation of a separate forecasting layer, introducing additional sources of model uncertainty and parameter instability. The present thesis deliberately abstracts from this predictive dimension in order to isolate the structural properties of discrete mean–variance optimization. Integrating forward-looking signals therefore represents a distinct research direction rather than a straightforward technical modification.

Regime dependence and multi-period extensions. Empirical results suggest that relative performance is regime-dependent, with alternating periods of outperformance and underperformance. Future work could investigate whether combinatorial subset selection performs differentially across volatility regimes, correlation breakdown episodes, or dispersion-driven environments.

More generally, extending the framework to a fully dynamic multi-period setting under stochastic wealth evolution would allow analysis of intertemporal utility maximization rather than repeated static optimization.

Concluding perspective. Taken together, these considerations indicate that the present framework should be viewed as a foundational structural layer for discrete portfolio selection. Its contribution lies in rigorously analyzing the combinatorial mean–variance problem under controlled computational budgets and coherent moment-based inputs.

Extensions toward forecasting integration, turnover-aware dynamics, and multi-period optimization represent natural continuations of this research, but remain conceptually distinct from the core objective pursued in this thesis.

BIBLIOGRAPHY

Aarts, E. and Korst, J. (1988). *Simulated Annealing and Boltzmann Machines*. Wiley.

Aiex, R. M., Resende, M. G. C., and Ribeiro, C. C. (2007). Ttt plots: A perl program to create time-to-target plots. *Optimization Letters*, 1(4):355–366.

Bailey, D. H., Borwein, J. M., López de Prado, M., and Zhu, Q. J. (2014). The probability of backtest overfitting. *Journal of Computational Finance*, 20(4):39–69.

Campbell, J. Y., Lo, A. W., and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*. Princeton University Press.

Chang, T.-J., Meade, N., Beasley, J. E., and Sharaiha, Y. M. (2000). Heuristics for cardinality constrained portfolio optimisation. *Computers & Operations Research*, 27(13):1271–1302.

Chekhlov, A., Uryasev, S., and Zabarankin, M. (2005). Drawdown measure in portfolio optimization. *International Journal of Theoretical and Applied Finance*, 8(1):13–58.

DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies*, 22(5):1915–1953.

Hoos, H. H. and Stützle, T. (2004). *Stochastic Local Search: Foundations and Applications*. Elsevier.

Jagannathan, R. and Ma, T. (2003). Risk reduction in large portfolios: Why imposing the wrong constraints helps. *Journal of Finance*, 58(4):1651–1684.

Jobst, N. J., Horniman, M. D., Lucas, C. A., and Mitra, G. (2001). Computational aspects of alternative portfolio selection models in the presence of discrete asset choice constraints. *Quantitative Finance*, 1(5):489–501.

Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220(4598):671–680.

Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411.

Lo, A. W. (2002). The statistics of sharpe ratios. *Financial Analysts Journal*, 58(4):36–52.

Lucas, A. (2014). Ising formulations of many NP problems. *Frontiers in Physics*, 2:5.

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

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

Michaud, R. O. (1989). The markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, 45(1):31–42.

Sharpe, W. F. (1966). Mutual fund performance. *Journal of Business*, 39(1):119–138.

FIGURES AND TABLES

Table 2: Window-level out-of-sample Sharpe ratios across $W = 19$ block-forward windows.

Window	Test period	Warm SA	Random SA	Equal Weight	Top-K
1	2005-01-02–2008-01-02	0.842	0.294	0.527	0.562
2	2006-01-02–2009-01-02	-0.303	-0.259	-0.341	-0.272
3	2007-01-02–2010-01-02	-0.209	-0.096	-0.160	-0.174
4	2008-01-02–2011-01-02	0.139	-0.132	0.043	0.121
5	2009-01-02–2012-01-02	0.820	0.928	0.613	0.771
6	2010-01-02–2013-01-02	1.188	1.412	0.577	0.457
7	2011-01-02–2014-01-02	1.120	1.457	0.799	0.682
8	2012-01-02–2015-01-02	1.816	1.497	1.503	1.640
9	2013-01-02–2016-01-02	1.711	1.199	0.878	1.201
10	2014-01-02–2017-01-02	0.432	0.636	0.570	0.630
11	2015-01-02–2018-01-02	0.537	0.717	0.687	0.695
12	2016-01-02–2019-01-02	0.762	0.650	0.569	0.574
13	2017-01-02–2020-01-02	0.921	0.746	0.644	0.812
14	2018-01-02–2021-01-02	0.872	0.541	0.462	0.732
15	2019-01-02–2022-01-02	0.507	0.362	0.308	0.621
16	2020-01-02–2023-01-02	0.675	0.530	0.417	0.584
17	2021-01-02–2024-01-02	0.321	0.312	0.224	0.392
18	2022-01-02–2025-01-02	0.612	0.584	0.331	0.711
19	2023-01-02–2026-01-02	0.701	0.748	0.154	0.824