




Research Proposal

How Much Can a Portfolio Optimization QUBO Be Simplified While Preserving Solution Quality?

Sparsification and Variable Reduction of Portfolio Optimization QUBOs for QAOA

Varun Ananthakrishnan
Ethan Cui
Ryan Guo

October 1, 2026

Contents

1	Introduction	2
1.1	Intellectual Merit	2
1.2	Broader Impact	2
2	Background and Prior Work	3
3	Research Goals	4
3.1	Goal 1: Build a Portfolio Optimization Benchmark	4
3.2	Goal 2: Compare Sparsification and Variable Reduction	4
3.3	Goal 3: Measure Portfolio Quality and QAOA Circuit Cost	4
4	Proposed Methods and Research Design	5
5	Expected Results and Data Analysis	6
6	Feasibility and Resources	7
7	Safety and Risk Considerations	7
8	Potential Pitfalls and Alternative Strategies	8
9	Research Timeline	8
10	Group Responsibilities	9
	Annotated Bibliography	9

1 Introduction

Our project began with a simple question: if some covariances between stocks are very small, do we need to keep all of them when optimizing a portfolio? In a quantum circuit, retaining a pairwise interaction can add gates. Removing enough interactions might make the calculation easier, but it could also change which stocks are selected. We want to measure that tradeoff and investigate better ways to decide what can be removed.

We will study a portfolio that selects a fixed number of stocks and gives them equal weights. Each stock has a binary variable indicating whether it is included. Expected returns, the covariance matrix, and a penalty for selecting the wrong number of stocks can be combined into a quadratic unconstrained binary optimization problem, or QUBO. QAOA, the Quantum Approximate Optimization Algorithm, will be used to search for solutions on small instances.

Mr. Hannum suggested starting with a manageable mutual fund and considering whether reduction could go far enough to shrink the matrix itself. That adds a second question to the project. Removing covariance terms makes the matrix sparser, but it does not reduce the number of stock variables. To use fewer qubits in this encoding, we would need to fix or eliminate some of those variables. We will investigate both kinds of reduction and judge the resulting portfolios using the original objective.

1.1 Intellectual Merit

We will study how sparsification and asset screening interact in fixed-size portfolio selection with QAOA. In particular, we will test whether their order changes the result: should we remove weak interactions before deciding which assets to exclude, or screen the assets using the full model first? Pruning can alter the portfolios used to rank assets, so the two orders may retain different stocks even when they produce problems of the same size. We will compare both orders with each reduction applied alone and measure the resulting portfolio loss and circuit cost.

The proposed contribution is an empirical study of this interaction and its effect on QAOA. Both reductions have precedents, so we will compare the order experiment with the closest existing methods. Exact solutions on small cases and compiled circuit measurements will help distinguish information lost during preprocessing from limitations of QAOA.

1.2 Broader Impact

Smaller circuits could make portfolio optimization experiments more practical with limited quantum resources. Our code and documented experiments could give later QLab groups a starting point. Identifying which assets or interactions cannot be removed without substantial loss would also help clarify when sparsification is useful.

2 Background and Prior Work

The covariance matrix is the part of our model that describes how stock returns vary together. It supplies the risk terms, while the QUBO also contains expected-return terms and any constraint penalties. QAOA is the optimization algorithm used after this model has been constructed. He et al. study how the initial state and mixer affect QAOA for constrained problems, including portfolio selection and hardware experiments [1]. Their work is relevant because our portfolios must contain a fixed number of assets.

Choi et al. compare covariance estimators for portfolio optimization [2]. We will use that work to understand the construction of the original model. During the main sparsification comparison, the original estimator will stay fixed so that we can isolate the effect of removing its terms.

Buhler and Benson directly study portfolio dimension reduction and covariance sparsification [3]. Their model uses continuous portfolio weights and classical optimization. Our proposed comparison concerns fixed-cardinality binary selection and the resulting QAOA circuits. We will use their work as a reference for the reduction methods.

Other relevant work includes Liu et al.’s sparsified QAOA phase operator for MaxCut [4] and Lewis and Glover’s variable-fixing rules for QUBOs [5]. Lozano’s recent preprint studies sparsification and constraint penalties in portfolio optimization on a quantum annealer [6]. Moondra et al. also connect graph sparsification with compilation cost and noise in QAOA for MaxCut [8].

Portfolio decomposition is another close line of work. Acharya et al. divide large portfolio problems into smaller constrained subproblems [9], and Gomez Cadavid et al. use hardware-aware decomposition with a counterdiabatic quantum optimizer [10]. Our initial study will produce one reduced asset-selection problem. It will examine the order of interaction pruning and asset screening, with both stages assessed against the same original objective.

3 Research Goals

3.1 Goal 1: Build a Portfolio Optimization Benchmark

First, we need a problem whose solution we can check. We will construct small portfolio QUBOs from historical returns and from synthetic data with known correlation patterns. Exhaustive search will give the optimal portfolios for the smallest cases and allow us to verify the QUBO coefficients. These full, unreduced problems will serve as the reference throughout the study.

3.2 Goal 2: Compare Sparsification and Variable Reduction

We will start by removing the smallest covariance terms and tracking the change in the portfolio's original objective value. Then we will compare this baseline with a more selective rule and investigate asset-variable removal. The reason for excluding an asset is still an open part of the project: weak correlations alone are not sufficient, since they may make a stock useful for diversification.

One candidate for the initial experiments is to screen assets according to how often they occur in good portfolios found by a limited classical search. We will test this as a heuristic, compare it with random removal, and examine whether any established variable-fixing rules apply. We will compare screening on the original model with screening after covariance pruning. This will let us test how much reduction is possible and whether changing the order causes useful assets to be excluded.

3.3 Goal 3: Measure Portfolio Quality and QAOA Circuit Cost

Finally, we will run QAOA on the original and reduced problems. The main comparison will be the portfolio quality obtained for the required qubits, gates, and circuit depth. We will test interaction removal and variable reduction separately as well as together, using repeated runs and comparable optimization budgets. This will show whether the reductions help the complete computation once the cost of preprocessing and the rest of the circuit are included.

4 Proposed Methods and Research Design

Starting dataset. We plan to begin with 8–12 candidate stocks. As a first pilot, we propose selecting four stocks from a set of ten; all 210 feasible selections can then be checked. This pilot will help us check the code and estimate runtime before moving to a wider set of instances.

For a larger candidate pool, we will look for a mutual fund with accessible, dated holdings and enough historical price data for its stocks. Vanguard Dividend Growth Fund is one option [7]. We have not finalized the fund or observation period. We will document missing-data decisions and adjustments for splits and dividends. If a present-day holdings list is used with older prices, the results will be described as a retrospective computational benchmark.

Portfolio model. Let $x_i = 1$ mean that stock i is selected and $x_i = 0$ mean that it is excluded. We require $\sum_i x_i = K$, where K is the number of selected stocks. Equal weighting gives $w_i = x_i/K$. With expected returns μ , covariance matrix Σ , and risk-aversion parameter $\lambda > 0$, we minimize

$$F(x) = \frac{\lambda}{K^2} x^\top \Sigma x - \frac{1}{K} \mu^\top x, \quad \mathbf{1}^\top x = K. \quad (1)$$

The first term measures portfolio variance and the second rewards expected return. All comparisons will use consistent return periods, the same K , and the same value of λ . The QUBO adds a penalty for violating the selection constraint:

$$E(x) = F(x) + A(\mathbf{1}^\top x - K)^2. \quad (2)$$

Here $A > 0$ controls the penalty. Since $x_i^2 = x_i$, we can write $E(x) = x^\top Q x + AK^2$, with

$$Q = \frac{\lambda}{K^2} \Sigma - \frac{1}{K} \text{diag}(\mu) + A\mathbf{1}\mathbf{1}^\top - 2AKI. \quad (3)$$

This uses a symmetric QUBO matrix. We will verify the coefficients against direct calculations of risk and return on the pilot instances.

The portfolio-size constraint. Expanding the penalty adds an interaction between every pair of variables. As a result, pruning covariance terms alone may leave a dense QUBO. Our proposed QAOA implementation will address this with an initial state containing exactly K selected assets and an XY mixer, which preserves that number [1]. The circuit can then use F as its cost function without the penalty. Strictly, this implementation solves the constrained binary quadratic problem corresponding to the QUBO above. The initial-state preparation and mixer will both count toward the circuit cost.

Removing interactions. We will first rank off-diagonal covariance entries by absolute magnitude and remove them in batches. Each entry and its symmetric counterpart will be removed together, and the diagonal will stay unchanged. A proposed pilot comparison will retain 100%, 75%, 50%, and 25% of the off-diagonal pairs. Finer steps can be added around any point where solution quality changes sharply.

An informed alternative will be tested using the effect of a term on sampled feasible portfolios, with particular attention to portfolios that score well. The pilot will determine the sampling and scoring rule. We will then fix it before the main evaluation.

Removing asset variables. The initial screening candidate will use a fixed-budget classical search to produce feasible portfolios and rank assets by their inclusion in the better solutions. We will test whether excluding low-ranked assets by setting $x_i = 0$ retains good solutions. The search budget, definition of a good portfolio, and number of retained assets will be recorded. This rule will not have access to the exact optimum used for evaluation.

Similarity between assets' returns is another possible criterion, but we will not assume that two highly correlated stocks are interchangeable. Their expected returns and relationships with the other assets can differ. We will also check whether the certified fixing rules in Lewis and Glover apply to our formulation [5]. If a variable is fixed to 1, both the remaining objective and the required number of additional selections must be updated. QAOA tests will retain more than K candidate assets so that a selection problem remains.

Experimental comparisons. We will compare the full problem, interaction-only sparsification, variable-only reduction, and two combined versions. In the first combined version, the classical search ranks assets using the full objective, after which we sparsify the retained problem. In the second, the search ranks assets using a sparsified objective before the low-ranked variables are removed. We will use the same search procedure and budget in both cases. Comparisons will match the final asset count and retained interaction count where possible, and report the compiled circuit cost separately. Random removal at matched sizes will provide a basic control. Where an existing preprocessing method can be adapted to the binary model, it will provide a stronger comparison. A classical greedy or local-search solver will also be tested on the original problem, with preprocessing time included when comparing complete workflows.

Reduction settings will be selected using training and validation data. Later data windows will be reserved for evaluation. Every reduced solution will be mapped back to the original stock list and scored using the original covariance matrix and expected returns. This prevents a reduced model from appearing better simply because it leaves some risk terms out of its own score.

QAOA runs. We propose beginning with depths $p = 1, 2, 3$, several parameter initializations, and matched measurement and objective-evaluation budgets. Noiseless simulation will be used first. Circuits will be compiled to a common target connectivity, and we will count the gates needed for the full circuit, including routing. Noise simulations and a small hardware test can follow if the baseline experiments leave enough time. The device or noise model, compiler settings, and optimization budget will be stated in the results.

5 Expected Results and Data Analysis

We expect small amounts of covariance pruning to have less effect than aggressive pruning, but that will depend on the instance. Variable reduction may save more resources because it reduces the number of qubits. It also carries a different risk: excluding one important stock can remove the original optimal portfolio from the feasible set altogether.

For a feasible returned portfolio $\hat{x}$ and a known optimum x^* , our primary measure will be the objective loss

$$\Delta F = F(\hat{x}) - F(x^*). \quad (4)$$

Lower loss is better. When optimality cannot be certified, we will label the reference as the best-known solution and report any available solver bounds. We will not divide by $F(x^*)$ to form a percentage error, since its value may be zero or negative.

On small instances, we will first solve the reduced problem exactly and evaluate that solution under the original objective. We will then run QAOA and measure its error relative to the optimum of the reduced objective it receives. These checks separate the effect of changing the problem from the ability of QAOA to solve it.

The main plots will show objective loss against retained interactions, retained variables, and compiled circuit resources. We will also report portfolio overlap, modeled return and variance, feasibility rate, and variation across repeated runs. Different portfolios can have nearly equal objective values, so a change in selected stocks will not automatically count as a poor result.

The acceptable loss tolerance and repetition count will be chosen after the pilot, before the final tests. We will present the full tradeoff curves as well as results at that tolerance. Performance on later return data will be a secondary check; preserving the estimated objective does not establish future investment performance.

6 Feasibility and Resources

The core study needs Python, historical price data, a classical optimization solver, and tools for simulating and compiling QAOA. It can be completed without a quantum processor. We will start small enough to enumerate the feasible portfolios and increase the size only after measuring runtime and memory use.

The larger mutual-fund universe will mainly support the classical reduction experiments. We will not assume that its full state vector can be simulated. If preprocessing produces a smaller problem, we can test that problem with QAOA while checking its portfolio against the original model classically.

The minimum complete project will include a checked QUBO baseline, magnitude pruning, one informed reduction candidate, and comparisons of solution quality and circuit cost. A broader search for variable-fixing rules or additional hardware experiments will depend on progress with that core study.

7 Safety and Risk Considerations

This is a computational project using public or licensed market data. It will not involve real trades or private financial information. We will follow the data sources' terms, protect computing credentials, and obtain approval before using paid resources. The results will be presented as an optimization experiment using a simplified portfolio model.

8 Potential Pitfalls and Alternative Strategies

The main technical risk is that sparsification removes information that matters collectively even when individual coefficients are small. We will examine where the loss begins to rise and reduce the pruning level when necessary. Variable screening will be assessed separately, since its errors restrict which portfolios are available.

Thresholding can also make a covariance matrix indefinite. Such a matrix can still define a binary optimization objective, but we will treat it as an approximation and use the original covariance for risk evaluation. Any claim that it is a valid covariance estimator would need a separate check.

Circuit savings may be smaller than the reduction in coefficients suggests. The mixer, state preparation, and routing can account for much of the cost, which is why we will measure the compiled circuit. If larger QAOA runs are impractical, the project can still combine larger classical reduction tests with quantum simulations on smaller verifiable cases.

A final possibility is that the classical preprocessing does most of the useful optimization. Comparing it with direct classical search will help us identify this. It would limit claims about QAOA’s contribution, but would still tell us something useful about the reduction method.

9 Research Timeline

October 2026: Establish the baseline. Review the closest prior work, choose the initial data, and construct the QUBO. Check the pilot instances by enumeration and estimate the time needed for repeated experiments.

November–December 2026: Test reduction methods. Implement magnitude pruning and controls, then test candidate interaction and asset-screening rules. Choose the main rules and evaluation settings after the pilot.

January–March 2027: Run QAOA comparisons. Compare original and reduced problems, measure compiled circuits, and repeat runs under matched budgets. Add noise experiments or hardware tests if resources allow.

April–May 2027: Analyze and report. Complete the held-out evaluations, investigate failure cases, and prepare the paper, figures, and presentation. These dates remain subject to the laboratory’s deadlines.

10 Group Responsibilities

The following assignments are proposed for discussion within the group. All three members will contribute to the literature review, interpretation, writing, and presentation, and will review one another’s implementation.

Varun Ananthakrishnan: Sparsification and variable-reduction methods; analysis of portfolio quality and computational cost.

Ethan Cui: Data preparation, QUBO construction, and classical reference solutions.

Ryan Guo: QAOA implementation, circuit compilation, and resource measurements.

Annotated Bibliography

- [1] Z. He, R. Shaydulin, S. Chakrabarti, D. Herman, C. Li, Y. Sun, and M. Pistoia. “Alignment between initial state and mixer improves QAOA performance for constrained optimization.” *npj Quantum Information*, 9, 121, 2023. <https://doi.org/10.1038/s41534-023-00787-5>.

The authors investigate how the initial state and mixer affect QAOA for constrained optimization, using portfolio selection as a test problem and including quantum-hardware experiments. Our problem also fixes the number of selected assets, so the mixer and initial state are part of the experimental design. We will use this work to guide a feasible-state QAOA implementation and to justify controlling these choices when comparing reduced and unreduced instances.

- [2] Y.-G. Choi, J. Lim, and S. Choi. “High-dimensional Markowitz portfolio optimization problem: empirical comparison of covariance matrix estimators.” *Journal of Statistical Computation and Simulation*, 89(7), 1278–1300, 2019. <https://doi.org/10.1080/00949655.2019.1577855>.

This study compares covariance matrix estimation methods in high-dimensional portfolio optimization, including approaches that impose structural assumptions. It connects the representation of asset dependence to the quality of an optimized portfolio. Sparse covariance estimation and sparse inverse-covariance estimation must be distinguished when interpreting its methods. We will use it as background for choosing and documenting a baseline estimator, while keeping estimation changes separate from the effect of pruning a fixed objective.

- [3] C. K. Buhler and H. Y. Benson. “Efficient Solution of Portfolio Optimization Problems via Dimension Reduction and Sparsification.” arXiv preprint, 2023. <https://arxiv.org/abs/2306.12639>.

The authors use asset screening and covariance sparsification to accelerate classical portfolio optimization. Their continuous-weight model is a close precedent for our reduction idea. We will compare the relevant methods and explain what changes in fixed-cardinality binary selection and QAOA.

- [4] X. Liu, R. Shaydulin, and I. Safro. “Quantum Approximate Optimization Algorithm with Sparsified Phase Operator.” *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 133–141, 2022. <https://arxiv.org/abs/2205.00118>.

The authors study a sparsified QAOA phase operator for MaxCut, aiming to simplify the quantum operation while retaining useful optimization behavior. It provides a direct precedent for removing interactions from the QAOA cost operator. We will compare its assumptions with our constrained portfolio setting and distinguish operator simplification from asset-variable removal.

- [5] M. Lewis and F. Glover. “Quadratic Unconstrained Binary Optimization Problem Preprocessing: Theory and Empirical Analysis.” arXiv preprint, 2017. <https://arxiv.org/abs/1705.09844>.

This work develops preprocessing rules for QUBOs that can determine variable values and reduce problem size, supported by computational experiments. Its variable-fixing rules give us a starting point for reducing the number of asset variables. We will examine which rules apply to our encoding and distinguish certified reductions from heuristic asset screening.

- [6] L. Lozano. “A Penalty-Free Pipeline for Direct Quantum-Annealer Portfolio Optimization.” arXiv preprint, version 2, May 31, 2026. <https://arxiv.org/abs/2605.17628v2>.

The preprint examines portfolio optimization on a quantum annealer, including how cardinality penalties, sparsification, and classical feasibility processing affect the computation. It is closely related portfolio-specific work and highlights that a dense penalty can undermine the intended benefit of covariance sparsification. We will use it to motivate explicit constraint handling and separate preprocessing from quantum-solver effects. Its annealing workflow will not be treated as evidence of QAOA performance.

- [7] Vanguard. *Vanguard Dividend Growth Fund: Investment Profile*. Fund profile dated June 30, 2026. https://workplace.vanguard.com/assets/corp/fund_communications/pdf_publish/us-products/investment-profiles/0057.pdf.

This official profile describes one fund we may use to choose a manageable set of stocks. A documented fund provides a more manageable starting universe than a broad index, subject to verifying complete holdings and price-data availability. We will assess the fund as one candidate, obtain and archive a dated full holdings list if selected, and document the final universe. The profile’s largest-holdings list alone will not be treated as complete holdings.

- [8] J. Moondra, P. C. Lotshaw, G. Mohler, and S. Gupta. “Promise of Graph Sparsification and Decomposition for Noise Reduction in QAOA: Analysis for Trapped-Ion Compilations.” *Quantum*, 10, 2185, 2026. <https://doi.org/10.22331/q-2026-08-07-2185>.

This paper studies sparsification and decomposition for compiling QAOA on MaxCut problems. It connects approximation loss with circuit complexity and noise. We will use it to guide our resource measurements while checking which assumptions apply to portfolio objectives.

- [9] A. Acharya et al. “Decomposition pipeline for large-scale portfolio optimization with applications to near-term quantum computing.” *Physical Review Research*, 7, 023142, 2025. <https://doi.org/10.1103/PhysRevResearch.7.023142>.

The authors preprocess correlations and divide portfolio problems into smaller constrained subproblems. Their work provides a precedent for making portfolio problems fit limited quantum resources. We will distinguish that decomposition approach from our proposed asset-screening and interaction-pruning comparison.

- [10] A. Gomez Cadavid et al. “Large-scale portfolio optimization on a trapped-ion quantum computer.” arXiv preprint, 2026. <https://arxiv.org/abs/2602.23976>.

This study combines portfolio decomposition with a counterdiabatic quantum optimizer on trapped-ion hardware. It examines the relationship between executable problem size and portfolio quality. We will use it to place our smaller QAOA study in context; the solver and reduction procedure differ from ours.