Cardinality-Constrained Portfolio Selection as a QUBO: VQE (QAOA-style) on a Simulator with Simulated Annealing Baseline

Superpositions Studio

February 17, 2026

Abstract

We study cardinality-constrained portfolio selection as a quadratic unconstrained binary optimization (QUBO) problem derived from a penalized mean–variance objective. Using a built-in dataset of daily asset returns spanning 2018–2024 ($T = 1759$ dates) for $N = 10$ large-cap equities, we estimate expected returns and covariances on the first 80% of the time series and select exactly $K = 4$ assets with risk–return trade-off parameter $\alpha = 0.5$. The resulting QUBO is mapped to an Ising Hamiltonian and minimized using a variational quantum eigensolver (VQE) instantiated with a depth-$p = 2$ QAOA-style ansatz optimized for 400 Adam steps on a statevector simulator; the classical baseline is simulated annealing applied to the identical QUBO, and exact enumeration over all $\binom{10}{4} = 210$ feasible subsets is used as an oracle check. All three optimizers recover the same selection (AAPL, GOOGL, MSFT, NVDA), with QUBO energy -10.8209 and no constraint violation under the chosen penalty. Sampling the trained VQE circuit yields a feasible-sample rate of 0.736 for the equality constraint. Evaluated as an equal-weight portfolio on the held-out 20% test window, the selected subset achieves annualized return 0.493, annualized volatility 0.223, Sharpe ratio 2.207 (risk-free rate assumed zero), and maximum drawdown -0.169. Compared to equal-weighting all assets, the selected portfolio increases return but also increases risk, resulting in a slightly lower Sharpe on this particular test period.

1 Introduction

Modern portfolio construction is often framed as a trade-off between expected return and risk, most classically formalized by Markowitz mean–variance theory [1, 2]. In practical settings, however, investors frequently impose discrete constraints—such as limiting the number of holdings—which converts an otherwise convex optimization into a combinatorial subset-selection problem. Cardinality-constrained variants are known to be computationally challenging and motivate heuristic solvers as well as reformulations into quadratic unconstrained binary optimization (QUBO) models [12, 32].

QUBO objectives admit a direct mapping to Ising Hamiltonians over Pauli-Z operators [15], enabling the use of both classical heuristics such as simulated annealing [13, 14] and quantum variational optimization methods. Variational quantum eigensolver (VQE) algorithms [16, 17, 18] minimize the expected energy of a parameterized quantum state with respect to a cost Hamiltonian; when the ansatz alternates exponentials of a diagonal cost Hamiltonian and a non-commuting mixer, the resulting circuit family corresponds to the quantum approximate optimization algorithm (QAOA) [19] and its generalizations (quantum alternating operator ansatz) [20]. These methods are natural candidates for QUBO-encoded finance problems and have been explored in the broader context of quantum computing for finance [30, 31].

This paper presents an end-to-end study of cardinality-constrained portfolio selection using a QAOA-style instantiation of VQE on a classical simulator, benchmarked against classical simulated annealing on the identical QUBO. In addition to summarizing dataset diagnostics, we report optimization convergence, constraint satisfaction under a penalty encoding, and out-of-sample portfolio performance on a chronological hold-out period. Because the instance considered here is small ($N = 10$), we also compute the exact optimum by enumerating all feasible subsets, providing a correctness oracle that isolates algorithmic effects from modeling and estimation choices.

2 Dataset and exploratory analysis

The dataset consists of a wide multivariate time series of daily arithmetic returns for $N = 10$ equities (AAPL, AMZN, GOOGL, JPM, META, MSFT, NVDA, PG, V, XOM) observed over $T = 1759$ dates from 2018-01-03 to 2024-12-30. Standard integrity checks confirm that the panel is complete (zero missing values across all assets) and contains no duplicate dates. The distribution of calendar gaps between adjacent rows is dominated by 1-day and 3-day intervals, consistent with trading-day sampling with weekend breaks.

For descriptive statistics, we treat each column as a daily simple return series $r_{t,i}$ and summarize its sample mean and standard deviation with a 252-trading-day annualization convention [6]. As is typical in equity returns [7], the distributions exhibit heavy tails and heterogeneous volatilities across assets. In this sample, NVDA has the largest annualized volatility (approximately 0.515), while PG has the smallest (approximately 0.201). Such heterogeneity is important for subset selection because a mean–variance objective combines both expected returns and covariances.

The dependence structure was examined via correlation and covariance matrices, both on the full sample and on the first 80% chronological window used later as an estimation period. On the 80% window, the strongest correlation is observed between GOOGL and MSFT (about 0.78) and the weakest between AMZN and XOM (about 0.20), indicating that the risk term is meaningfully non-diagonal and that diversification opportunities exist across sectors. The covariance eigen-spectrum on the same window yields an estimated condition-number proxy $\kappa(\Sigma) = \lambda_{\max}/\lambda_{\min} \approx 42$, suggesting moderate ill-conditioning; for larger universes, shrinkage estimators can improve robustness [8].

3 Methods

3.1 Problem formulation

Let $r_{t,i}$ denote the daily arithmetic return of asset $i \in \{1, \ldots, N\}$ on date $t \in \{1, \ldots, T\}$. We adopt a chronological 80/20 split to avoid look-ahead bias: the first 1407 dates define an estimation window used to compute the empirical mean return vector $\mu \in \mathbb{R}^N$ and covariance matrix $\Sigma \in \mathbb{R}^{N\times N}$, while the remaining 352 dates are reserved for out-of-sample evaluation.

The portfolio-selection decision is encoded by a binary inclusion vector $x \in \{0, 1\}^N$, where $x_i = 1$ indicates that asset i is included. The user-specified cardinality constraint enforces exactly $K = 4$ selected assets:

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

Following mean–variance theory [1], we use a quadratic risk term and a linear return term, combined via a trade-off parameter $\alpha \in [0, 1]$. To obtain a QUBO, we incorporate the equality constraint

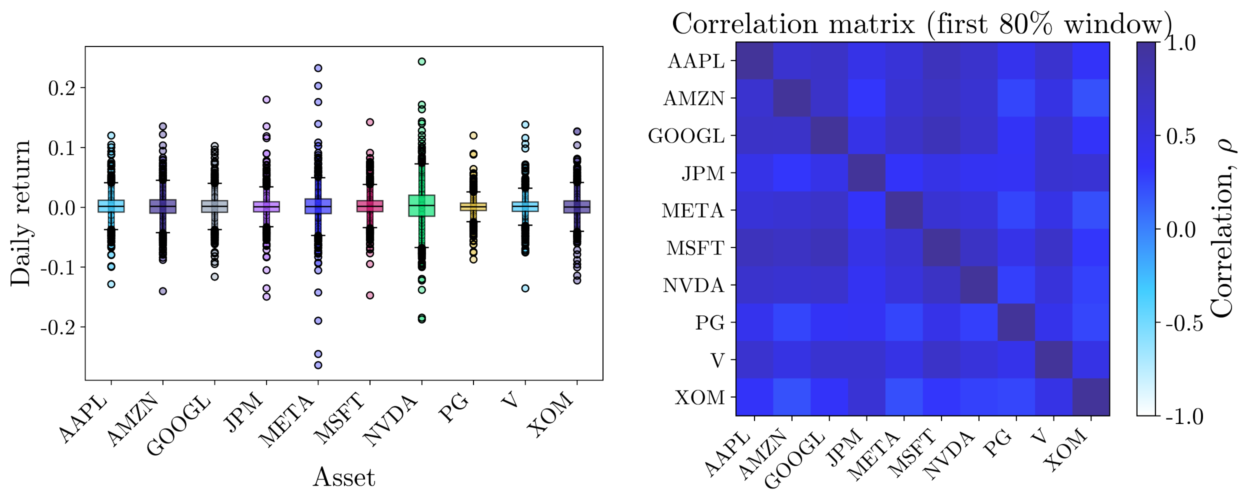

(a) Per-asset daily return distributions (boxplots).

(b) Correlation matrix on the first 80% (estimation) window.

Figure 1: Key exploratory summaries supporting mean–variance/QUBO modeling: heterogeneous marginal volatilities and a non-trivial correlation structure.

through a squared penalty, resulting in the objective minimized in this work:

$$F(x) = \frac{1-\alpha}{K^2} x^\top \Sigma x - \frac{\alpha}{K} \mu^\top x + A \left(\sum_{i=1}^{N} x_i - K \right)^2 . \tag{2}$$

We set $\alpha = 0.5$ and annualize μ and Σ by multiplying them by 252, consistent with standard empirical-finance conventions [6]. The penalty strength $A > 0$ must be chosen large enough that violating the constraint is suboptimal; in the executed run, A was selected by an automatic scaling rule (relative to the magnitudes of the risk and return terms), yielding $A \approx 0.66785$.

3.2 QUBO and Ising mapping

Because $x_i \in \{0, 1\}$ implies $x_i^2 = x_i$, expanding the penalty term in Eq. (2) yields a quadratic polynomial in the binary variables and can be written in standard QUBO form [32]:

$$F(x) = x^\top Q x + \text{const.} \tag{3}$$

A convenient coefficient choice (absorbing symmetric terms into the upper triangle) is

$$Q_{ii} = \frac{1-\alpha}{K^2} \Sigma_{ii} - \frac{\alpha}{K} \mu_i + A(1 - 2K), \tag{4}$$

$$Q_{ij} = \frac{1-\alpha}{K^2} \Sigma_{ij} + 2A, \quad i < j. \tag{5}$$

In this convention, $x^\top Q x$ is evaluated as $\sum_i Q_{ii} x_i + \sum_{i<j} Q_{ij} x_i x_j$.

To obtain a cost Hamiltonian suitable for gate-model variational optimization, we map x to Ising spins $z_i \in \{+1, -1\}$ via $x_i = (1 - z_i)/2$ [15]. Substituting and promoting z_i to Pauli-Z operators produces a diagonal Ising Hamiltonian

$$H_C = c_0 I + \sum_{i=1}^{N} h_i Z_i + \sum_{i<j} J_{ij} Z_i Z_j. \tag{6}$$

The constant shift c_0 does not affect the minimizing bitstring and is therefore ignored when comparing candidate solutions; for feasible solutions, different conventions for absorbing constants may shift reported energies by a fixed offset without changing the argmin.

3.3 Variational quantum optimization (VQE / QAOA-style)

We follow the VQE paradigm [16, 17, 18], minimizing the expected cost energy

$$E(\theta) = \langle \psi(\theta) | H_C | \psi(\theta) \rangle \tag{7}$$

with respect to variational parameters θ that specify a parameterized quantum state $|\psi(\theta)\rangle$. For QUBO/Ising objectives, a natural choice is the QAOA ansatz [19] of depth p,

$$|\psi(\gamma, \beta)\rangle = \prod_{\ell=1}^{p} \exp\big(-i\beta_\ell H_M\big) \exp\big(-i\gamma_\ell H_C\big) |+\rangle^{\otimes N}, \tag{8}$$

with mixer Hamiltonian $H_M = \sum_i X_i$. In this work we use $p = 2$, resulting in four continuous parameters $(\gamma_1, \gamma_2, \beta_1, \beta_2)$.

The executed run used an analytic statevector simulator to evaluate $E(\theta)$ and gradients via automatic differentiation, implemented with PennyLane [39] and PyTorch [42]. Parameters were optimized for 400 steps using Adam [41]. After optimization, the circuit was measured in the computational basis for 6000 shots to generate candidate portfolios; among the sampled bitstrings, we selected the best feasible solution (Hamming weight exactly K) according to the original objective $F(x)$. While penalty encodings do not enforce feasibility at the circuit level, feasibility can be improved by quantum alternating-operator constructions with constraint-preserving mixers [20]; we revisit this point in the outlook.

3.4 Classical baselines and oracle

As a classical baseline, we applied simulated annealing (SA) [13, 14] directly to the same QUBO objective, ensuring an apples-to-apples comparison in which only the optimizer differs. The SA run used 60,000 update steps and a geometric temperature schedule from $T_0 = 5.0$ to $T_f = 0.01$.

Because the instance is small ($N = 10$ and $K = 4$), we additionally enumerated all feasible subsets of size K (210 portfolios) and selected the best by direct evaluation of $F(x)$. This exact enumeration is not intended as a scalable baseline, but serves as a correctness oracle for validating the QUBO/Ising encoding and the quality of heuristic solutions.

3.5 Out-of-sample evaluation

To connect discrete subset selection to realized performance, each selected subset was evaluated as an equal-weight portfolio over the train and test windows (i.e., weights $1/K$ on selected assets, 0 otherwise). From daily portfolio returns r_t, we report annualized return and volatility as

$$\text{AnnRet} = \Big(\prod_t (1 + r_t)\Big)^{252/T_{\text{win}}} - 1, \qquad \text{AnnVol} = \text{Std}(r_t)\sqrt{252}, \tag{9}$$

where T_{win} is the number of days in the evaluated window. The Sharpe ratio is reported as Sharpe = AnnRet/AnnVol with risk-free rate assumed zero [4]. We also compute maximum drawdown as the most severe peak-to-trough decline of the cumulative growth curve [10]. For reference, we compare against a simple diversified benchmark that equal-weights all $N = 10$ assets.

Table 1: Dataset and solver configuration for the executed instance.

Assets (N)	10
Dates (T)	1759 (2018-01-03 to 2024-12-30)
Split	Chronological 80/20 (train 1407, test 352)
Cardinality constraint	$K = 4$
Risk–return parameter	$\alpha = 0.5$
Penalty strength	$A = 0.6678503772$
VQE/QAOA depth	$p = 2$
VQE optimizer	Adam, 400 steps, learning rate 0.18
VQE sampling	6000 shots; feasible-sample rate 0.736
Classical baseline	Simulated annealing, 60,000 steps ($T_0 = 5.0$, $T_f = 0.01$)
Oracle	Exact enumeration over $\binom{10}{4} = 210$ feasible subsets

Table 2: Solution quality and runtime. The "risk" and "return" columns correspond to the first two terms of Eq. (2); for all reported solutions the penalty term is zero (feasible portfolios).

Method	QUBO energy	Risk	Return	$F(x)$ total	Runtime (s)
VQE (QAOA, $p = 2$)	−10.820878	0.034973	−0.170245	−0.135272	77.67
Simulated annealing	−10.820878	0.034973	−0.170245	−0.135272	3.50
Exact enumeration	−10.820878	0.034973	−0.170245	−0.135272	0.01

4 Results

Table 1 summarizes the main configuration choices for the executed run. We report solver agreement, convergence behavior, and out-of-sample performance in turn.

4.1 Optimization outcomes, convergence, and runtimes

All solvers converged to the same portfolio, selecting {AAPL, GOOGL, MSFT, NVDA}. The corresponding bitstring (ordered as AAPL, AMZN, GOOGL, JPM, META, MSFT, NVDA, PG, V, XOM) is

$$(1, 0, 1, 0, 0, 1, 1, 0, 0, 0). \tag{10}$$

Table 2 reports the objective decomposition and wall-clock runtimes.

For convergence diagnostics, Figure 2 compares the VQE optimization trace with the best-so-far energy trajectory of simulated annealing. Figure 3 shows the distribution of sampled Hamming weights from the trained VQE circuit. Under the penalty encoding and the selected A, 4417 out of 6000 samples satisfy the equality constraint (feasible rate 0.736).

4.2 Out-of-sample portfolio performance

Because VQE, SA, and enumeration return the same subset, their realized performance is identical. Table 3 reports performance on the estimation (train) and hold-out (test) windows for the selected subset and for the equal-weight-all benchmark.

Figure 4 visualizes cumulative growth on the test window. On this period, the selected subset achieves higher cumulative return but experiences larger drawdowns than equal-weighting the full universe.

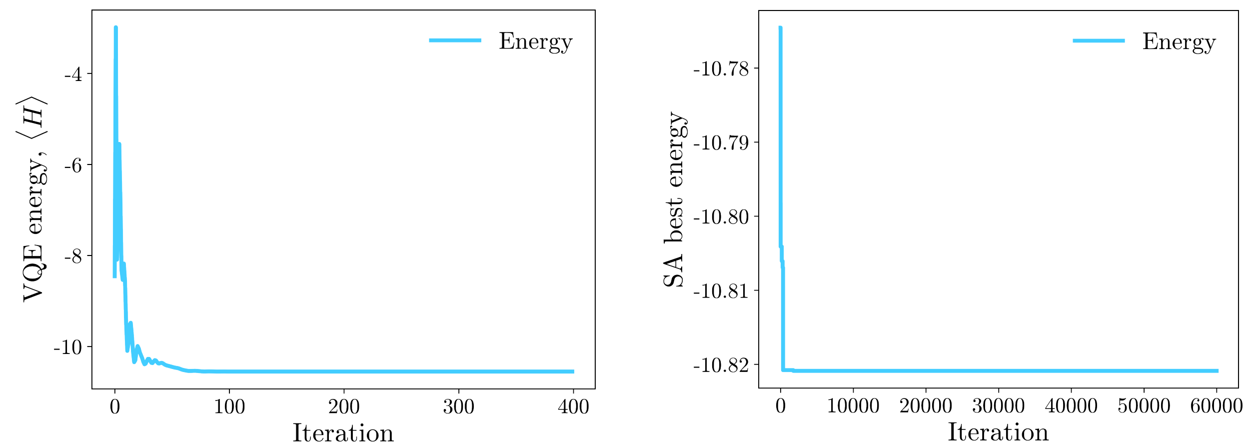

(a) VQE/QAOA energy during training.

(b) Best-so-far QUBO energy during SA.

Figure 2: Optimization traces for the variational quantum solver and simulated annealing baseline.

Table 3: Equal-weight realized performance for the selected subset and the equal-weight-all benchmark.

Window	Strategy	Ann. return	Ann. vol.	Sharpe	Max drawdown
Train	Selected (VQE/SA/Exact)	0.333963	0.322680	1.034967	−0.401651
Train	EqualWeight(ALL)	0.230889	0.253457	0.910958	−0.316858
Test	Selected (VQE/SA/Exact)	0.492752	0.223244	2.207237	−0.169171
Test	EqualWeight(ALL)	0.369430	0.145988	2.530550	−0.103736

5 Discussion

For the studied instance ($N = 10$, $K = 4$), the variational quantum solver reproduces the exact optimum, demonstrating that the end-to-end pipeline—from empirical estimation of (μ, Σ) to QUBO/Ising construction and variational optimization—is internally consistent. At the same time, the runtime comparison underscores an important practical point: on classical hardware, statevector-based simulation of VQE/QAOA is substantially more expensive than classical heuristics such as simulated annealing, while exact enumeration is trivial at this scale. Consequently, the present experiment should be interpreted primarily as a correctness and workflow demonstration rather than evidence of computational advantage.

From a constraint-handling perspective, the penalty formulation yields a substantial mass of feasible samples (0.736). Nevertheless, a non-negligible fraction of measurement outcomes violate the equality constraint and must be discarded during decoding. For larger problems, this post-selection overhead can be reduced by using constraint-preserving mixers within the alternating-operator framework [20], or by considering hybrid strategies such as biased initial states, penalty schedules, or risk measures that encourage feasibility.

The relationship between the optimization objective and realized performance is also non-trivial. The selected subset attains higher annualized return on the test period than the equal-weight-all benchmark, but with higher volatility and deeper drawdowns; accordingly, the benchmark achieves a slightly higher Sharpe ratio on this specific out-of-sample window. This highlights that the QUBO objective in Eq. (2) captures only one particular risk–return trade-off and depends on noisy estimates

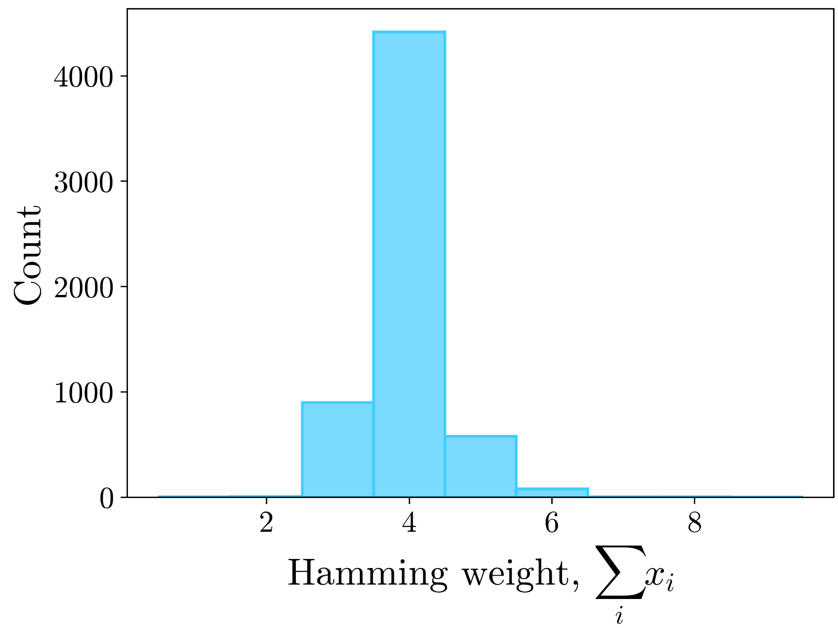

Figure 3: Histogram of sampled Hamming weights $\sum_i x_i$ after VQE/QAOA training. Feasible solutions correspond to weight $K = 4$.

of μ and Σ from a fixed window. In applied settings, one would typically incorporate additional modeling elements such as transaction costs, turnover limits, and more robust return/covariance estimators (e.g., shrinkage or Bayesian priors) [8, 5].

Finally, like other variational methods, QAOA-style circuits can exhibit optimization pathologies (e.g., barren plateaus) for certain ansatz choices or scalings [21]. While the present problem is small enough to be well-behaved, scaling studies should monitor trainability, shot noise, and the effect of noise and error mitigation [24, 25].

6 Future outlook and recommendations

6.1 Algorithmic improvements for constrained portfolio QUBOs

A direct penalty formulation with the standard X-mixer is attractive for its simplicity, but it does not restrict the quantum state to the feasible Hamming-weight-K subspace. For larger N and tighter constraints, incorporating constraint-preserving mixers within the alternating-operator framework [20] is a principled way to increase feasible-sample rates and to reduce post-selection overhead. Beyond feasibility, practical QAOA deployments often benefit from refined objective surrogates; for example, risk measures based on conditional value-at-risk (CVaR) and related functionals can be used either at the finance-model level [9, 36] or as a variational objective that emphasizes low-energy tails of the measurement distribution [26].

6.2 Modeling robustness and backtesting

The objective in Eq. (2) depends on empirical estimates of μ and Σ, which are noisy and non-stationary in real markets. Two immediate extensions are (i) rolling-window or expanding-window backtests to quantify temporal stability, and (ii) robustness improvements to covariance estimation

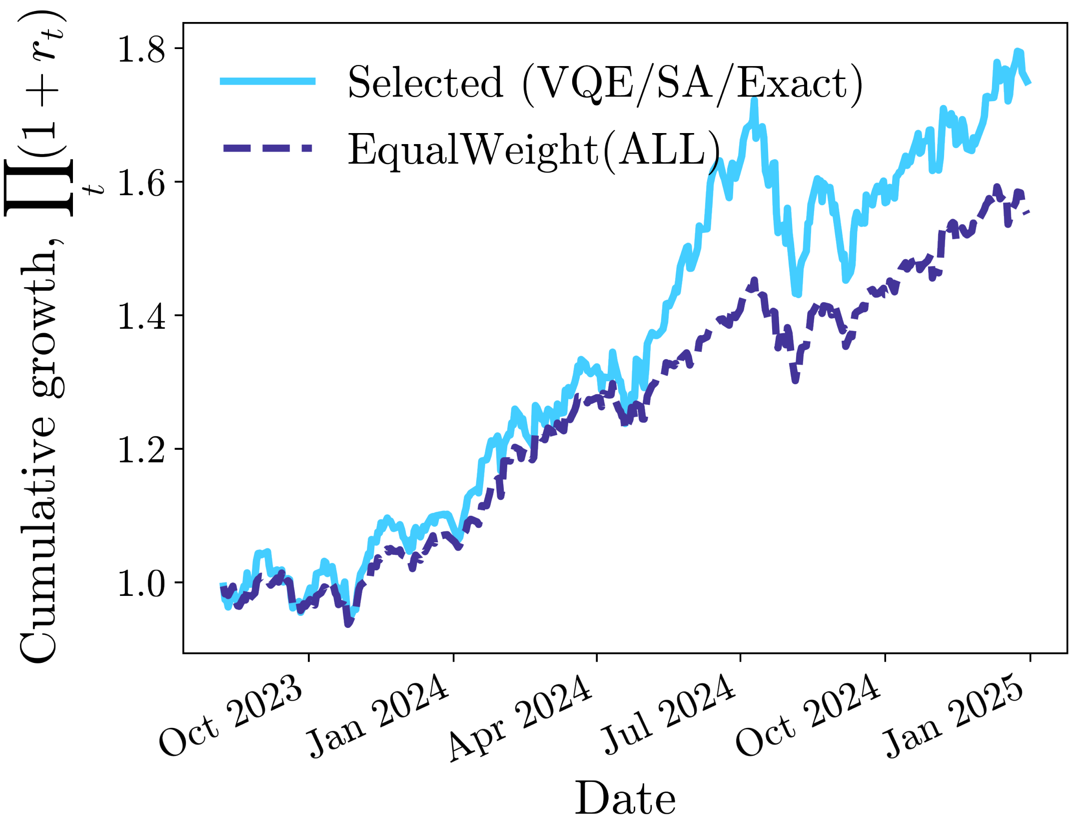

Figure 4: Cumulative growth on the test window for the selected subset (VQE/SA/Exact, identical) versus equal-weighting all assets.

(e.g., shrinkage) [8]. For expected returns, incorporating equilibrium priors and analyst "views" via Black–Litterman [5] can reduce sensitivity to sampling error. Sensitivity sweeps over (α, K) can also reveal whether the selected subset is stable to modest changes in the risk–return trade-off and diversification level.

6.3 Scaling considerations and hardware outlook

Even for classical simulation, the dense mean–variance risk term induces $O(N^2)$ pairwise couplings, and a direct gate decomposition typically requires many entangling operations. On quantum hardware, both the number of entangling gates and the number of measurement shots (to estimate energies) become the dominant costs. Figure 5 reproduces a representative error-threshold scaling law for variational algorithms, illustrating that the tolerable per-gate error decreases approximately inversely with the number of entangling gates [49]. This trade-off aligns with the broader view that near-term devices can support only relatively shallow circuits before noise dominates [24], motivating shallow ansätze, problem-structure exploitation, and error-mitigation techniques [25].

6.4 QPU execution outlook (not executed in this work)

All results in this report were produced on classical simulators. This run does not include a hardware-QPU experiment; if such execution is added in future runs, the main changes would include

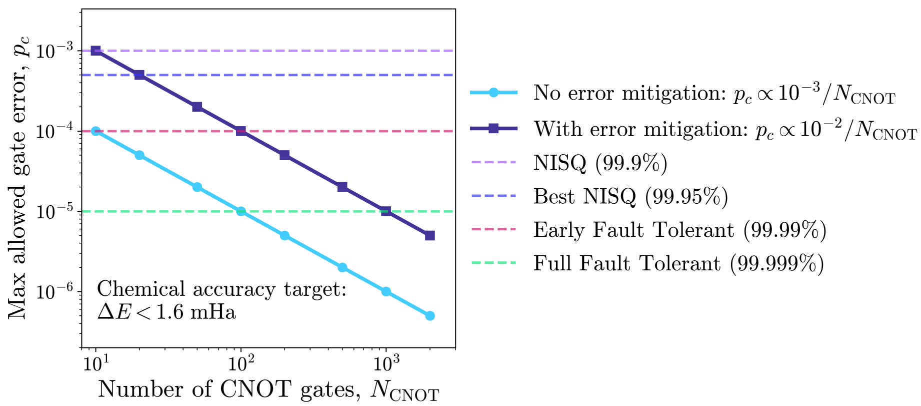

Figure 5: Illustrative scaling relationship between entangling-gate count and tolerable error rates for variational algorithms, adapted from an error-threshold analysis [49]. Although derived in a quantum-chemistry context, the qualitative message (deeper circuits demand proportionally lower error rates or stronger mitigation) applies broadly.

shot-based energy estimation, careful choice of measurement-grouping strategies, and the use of error mitigation to compensate for noise. Near-term roadmaps anticipate substantial improvements as systems transition toward error-corrected computation [50]; in that regime, deeper QAOA circuits and larger asset universes may become feasible.

7 Conclusion

We presented a complete workflow for cardinality-constrained portfolio selection formulated as a penalized mean–variance QUBO and solved using a QAOA-style instantiation of VQE, benchmarked against classical simulated annealing on the same objective. On a 10-asset universe with a $K = 4$ constraint, the variational quantum solver, simulated annealing, and exact enumeration all identify the same optimal subset (AAPL, GOOGL, MSFT, NVDA). Out-of-sample evaluation on a chronological hold-out period shows strong realized returns for the selected subset, accompanied by increased volatility and drawdown relative to an equal-weight-all benchmark.

While the small size of the instance precludes claims of computational advantage, the study validates the modeling and optimization pipeline and highlights concrete directions for scaling: improving constraint handling (e.g., constraint-preserving mixers), strengthening statistical estimation and backtesting, and accounting for the hardware costs of dense quadratic objectives.

References

[1] Markowitz, Harry. "Portfolio selection." *The Journal of Finance* 7(1) pp. 77–91 1952.

[2] Tobin, James. "Liquidity Preference as Behavior towards Risk." *The Review of Economic Studies* 25(2) pp. 65–86 1958. doi: 10.2307/2296205. https://www.jstor.org/stable/2296205

[3] Sharpe, William F.. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *The Journal of Finance* 19(3) pp. 425–442 1964. doi: 10.1111/j.1540-6261.1964.tb02865.x. https://www.semanticscholar.org/paper/a7da090fdb85b79cd8a52aadc9ff4715814656e5

[4] Sharpe, William F.. "Mutual fund performance." *The Journal of Business* 39(1) pp. 119–138 1966.

[5] Black, Fischer, Litterman, Robert B.. "Global Portfolio Optimization." *Financial Analysts Journal* 48(5) pp. 28–43 1992. doi: 10.2469/faj.v48.n5.28. https://www.semanticscholar.org/paper/aeda5ef81f1468cd5005683fa2029abc5988594a

[6] Campbell, John Y., Lo, Andrew W., MacKinlay, A. Craig. "The Econometrics of Financial Markets." *Princeton University Press* 1997.

[7] Cont, Rama. "Empirical properties of asset returns: stylized facts and statistical issues." *Quantitative Finance* 1(2) pp. 223–236 2001.

[8] Ledoit, Olivier, Wolf, Michael. "A well-conditioned estimator for large-dimensional covariance matrices." *Journal of Multivariate Analysis* 88(2) pp. 365–411 2004. doi: 10.1016/S0047-259X(03)00096-4.

[9] Rockafellar, R. Tyrrell, Uryasev, Stanislav. "Optimization of conditional value-at-risk." *Journal of Risk* 2(3) pp. 21–42 2000. doi: 10.21314/jor.2000.038. https://www.semanticscholar.org/paper/58444c142b6ea5c71a435cac7a0b4c66d6c68869

[10] Chekhlov, Alexei, Uryasev, Stanislav P., Zabarankin, Michael. "Drawdown Measure in Portfolio Optimization." 2003. doi: 10.2139/ssrn.544742. https://www.semanticscholar.org/paper/792dfbbe299e9890f7f28d17ac1a17485de300ab

[11] Konno, Hiroshi, Yamazaki, Hiroaki. "Mean-absolute deviation portfolio optimization model and its applications to Tokyo stock market." *Management Science* 37(5) pp. 519–531 1991. doi: 10.1287/mnsc.37.5.519. https://www.semanticscholar.org/paper/16053126f5027144fe23feebde9c7cea23193644

[12] Gao, Jianjun, Li, Duan. "Optimal Cardinality Constrained Portfolio Selection." *Operations Research* 2013. doi: 10.1287/opre.2013.1170. https://www.semanticscholar.org/paper/8dc2597725923d29baf3c1f0ea316b114f8cbc7f

[13] Kirkpatrick, Scott, Gelatt, C. Daniel, Vecchi, Mario P.. "Optimization by simulated annealing." *Science* 220(4598) pp. 671–680 1983.

[14] Geman, Stuart, Geman, Donald. "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images." *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-6(6) pp. 721–741 1984. doi: 10.1109/TPAMI.1984.4767596. https://www.semanticscholar.org/paper/459b30a9a960080f3b313e41886b1aa0e51e882c

[15] Lucas, Andrew. "Ising formulations of many NP problems." *Frontiers in Physics* 2 pp. 5 2014.

[16] Peruzzo, Alberto, McClean, Jarrod, Shadbolt, Peter, Yung, Man-Hong, Zhou, Xiao-Qi, Love, Peter J., Aspuru-Guzik, Alán, O'Brien, Jeremy L.. "A variational eigenvalue solver on a photonic quantum processor." *Nature Communications* 5 pp. 4213 2014.

[17] McClean, Jarrod R., Romero, Jonathan, Babbush, Ryan, Aspuru-Guzik, Alán. "The theory of variational hybrid quantum-classical algorithms." *New Journal of Physics* 18(2) pp. 023023 2016.

[18] Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., Coles, P. J.. “Variational quantum algorithms.” *Nature Reviews Physics* 3 pp. 625–644 2021.

[19] Farhi, Edward, Goldstone, Jeffrey, Gutmann, Sam. “A quantum approximate optimization algorithm.” *arXiv preprint arXiv:1411.4028* 2014. https://arxiv.org/abs/1411.4028

[20] Hadfield, Stuart, Wang, Zhihui, O’Gorman, Bryan, Rieffel, Eleanor G., Venturelli, Davide, Biswas, Rupak. “From the quantum approximate optimization algorithm to a quantum alternating operator ansatz.” *Algorithms* 12(2) pp. 34 2019.

[21] McClean, Jarrod R., Boixo, Sergio, Smelyanskiy, Vadim N., Babbush, Ryan, Neven, Hartmut. “Barren plateaus in quantum neural network training landscapes.” *Nature Communications* 9 pp. 4812 2018. doi: 10.1038/s41467-018-07090-4. https://www.semanticscholar.org/paper/d699e0958fe1d8a4c1d691765f7e11b823fa606f

[22] Schuld, Maria, Bergholm, Ville, Gogolin, Christian, Izaac, Josh, Killoran, Nathan. “Evaluating analytic gradients on quantum hardware.” *Physical Review A* 99(3) pp. 032331 2019. doi: 10.1103/PhysRevA.99.032331. https://www.semanticscholar.org/paper/efb241f6d4209b1202d2b64072ed89dc9e93f73a

[23] Crooks, Gavin E.. “Gradients of parameterized quantum gates using the parameter-shift rule.” *arXiv preprint arXiv:1905.13311* 2019. https://arxiv.org/abs/1905.13311

[24] Preskill, John. “Quantum Computing in the NISQ era and beyond.” *Quantum* 2 pp. 79 2018. doi: 10.22331/q-2018-08-06-79. https://www.semanticscholar.org/paper/f3d594544126e202dbd81c186ca3ce448af5255c

[25] Endo, Suguru, Benjamin, Simon C., Li, Ying. “Practical Quantum Error Mitigation for Near-Future Applications.” *Physical Review X* 8 pp. 031027 2017. doi: 10.1103/PhysRevX.8.031027. https://www.semanticscholar.org/paper/9bbf679d0fd8d3e802842b0318347946b4fd9d1a

[26] Barkoutsos, Panagiotis, Nannicini, Giacomo, Robert, Anirudh, Tavernelli, Ivano, Woerner, Stefan. “Improving Variational Quantum Optimization using CVaR.” *Quantum* 2019. doi: 10.22331/q-2020-04-20-256. https://www.semanticscholar.org/paper/03ac37bb0df538a121198ad6ace15057a3e106c6

[27] Arute, Frank, Arya, Kunal, Babbush, Ryan, others. “Quantum approximate optimization of non-planar graph problems on a planar superconducting processor.” *Nature Physics* 2020. doi: 10.1038/s41567-020-01105-y. https://www.semanticscholar.org/paper/91c10ab4c59843ecb405018b13f792719ce7009c

[28] Kadowaki, Tadashi, Nishimori, Hidetoshi. “Quantum annealing in the transverse Ising model.” *Physical Review E* 58(5) pp. 5355–5363 1998. doi: 10.1103/PhysRevE.58.5355. https://www.semanticscholar.org/paper/af38085ac160e7f89d6e83751252154681ff4c59

[29] Farhi, Edward, Goldstone, Jeffrey, Gutmann, Sam, Lapan, Joshua, Lundgren, Andrew, Preda, Daniel. “A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem.” *Science* 292(5516) pp. 472–475 2001. doi: 10.1126/science.1057726. https://www.semanticscholar.org/paper/d2d5e71b19740bf91b1704570958bf235967857c

[30] Egger, Daniel J., Gambella, Claudio, Mareček, Jakub, McFaddin, Scott, Mevissen, Max, Raymond, Rudy, Simonetto, Andrea, Woerner, Stefan. "Quantum computing for finance: state of the art and future prospects." *IEEE Transactions on Quantum Engineering* 1 pp. 1–24 2020.

[31] Woerner, Stefan, Egger, Daniel J.. "Quantum risk analysis." *npj Quantum Information* 5(15) 2019.

[32] Glover, Fred, Kochenberger, Gary, Du, Yu. "Quantum Bridge Analytics I: a tutorial on formulating and using QUBO models." *4OR* 2019. doi: 10.1007/s10288-019-00424-y. https://www.semanticscholar.org/paper/2691557c5d9d1e8ace7fd63bb8f056e2407f6d4b

[33] Brandhofer, Sebastian, Braun, Daniel, Dehn, Vanessa, van der Linde, Friedrich, Papp, Carsten, Stehle, Dominik, Woerner, Stefan. "Benchmarking the performance of portfolio optimization with QAOA." *Quantum Information Processing* 2022. doi: 10.1007/s11128-022-03766-5. https://www.semanticscholar.org/paper/27585eba1bf58fcb75b0c170f4e9ec08868f2d40

[34] He, Zichang, Shaydulin, Ruslan, Chakrabarti, Shouvanik, others. "Alignment between Initial State and Mixer Improves QAOA Performance for Constrained Portfolio Optimization." 2023. doi: 10.48550/arXiv.2305.03857. https://www.semanticscholar.org/paper/f565eb76681c4a87f5372d227a909e9311681b12

[35] Lang, Jonas, Zieliński, Sebastian, Feld, Sebastian. "Strategic Portfolio Optimization Using Simulated, Digital, and Quantum Annealing." *Applied Sciences* 2022. doi: 10.3390/app122312288. https://www.semanticscholar.org/paper/98c7dfd93b588dee6f6b78d83fc2e5eeb7552731

[36] Xu, Hanjing, Dasgupta, S., Pothen, A., others. "Dynamic Asset Allocation with Expected Shortfall via Quantum Annealing." *Entropy* 2021. doi: 10.3390/e25030541. https://www.semanticscholar.org/paper/f2debb4adda0522fbb9de16c11f5b1d8ea92cea6

[37] Parizy, Matthieu, Sadowski, Peter, Togawa, Naotaka. "Cardinality Constrained Portfolio Optimization on an Ising Machine." *IEEE International Conference on Smart Computing (SMARTCOMP)/SOCC* 2022. doi: 10.1109/SOCC56010.2022.9908082. https://www.semanticscholar.org/paper/d9cfd11ab77c968584171e896713a86b10ea883e

[38] Rosenberg, Gili, Haghnegahdar, Parisa, Goddard, Phil, Carr, Paul, Wu, K., de Prado, Marcos Lopez. "Solving the Optimal Trading Trajectory Problem Using a Quantum Annealer." *Proceedings of the Workshop on Post-Moore's Era Supercomputing (PMES) at SC* 2015. doi: 10.1145/2830556.2830563. https://www.semanticscholar.org/paper/109f769c706542d8e0f14c94245ae27bb495f057

[39] Bergholm, Ville, Izaac, Josh, Schuld, Maria, Gogolin, Christian, Ahmed, Shahnawaz, Ajith, Mohan, Alam, M. S., Alonso-Linaje, Guillermo, Akash, A., Asadi, A., others. "PennyLane: Automatic differentiation of hybrid quantum-classical computations." *arXiv preprint arXiv:1811.04968* 2018. https://arxiv.org/abs/1811.04968

[40] Asadi, Ali, Dusko, Amintor, Park, Chae-Yeun, others. "Hybrid quantum programming with PennyLane Lightning on HPC platforms." 2024. doi: 10.48550/arXiv.2403.02512. https://www.semanticscholar.org/paper/8d7ee48af14f2ddb985feb7d56f7ccd2b0a7aaeb

[41] Kingma, Diederik P., Ba, Jimmy. "Adam: A Method for Stochastic Optimization." 2014. https://www.semanticscholar.org/paper/a6cb366736791bcccc5c8639de5a8f9636bf87e8

[42] Paszke, Adam, Gross, Sam, Massa, Francisco, Lerer, Adam, Bradbury, James, Chanan, Gregory, Killeen, Trevor, Lin, Zeming, Gimelshein, Natalia, Antiga, Luca, others. "PyTorch: An imperative style, high-performance deep learning library." *Advances in Neural Information Processing Systems* 32 2019.

[43] Harris, Charles R., Millman, K. Jarrod, van der Walt, Stéfan J., Gommers, Ralf, Virtanen, Pauli, Cournapeau, David, Wieser, Eric, Taylor, Julian, Berg, Sebastian, Smith, Nathaniel J., others. "Array programming with NumPy." *Nature* 585(7825) pp. 357–362 2020.

[44] McKinney, Wes. "Data structures for statistical computing in Python." *Proceedings of the 9th Python in Science Conference* pp. 56–61 2010.

[45] Virtanen, Pauli, Gommers, Ralf, Oliphant, Travis E., Haberland, Matt, Reddy, Tyler, Cournapeau, David, Burovski, Evgeni, Peterson, Pearu, Weckesser, Warren, Bright, Jonathan, others. "SciPy 1.0: Fundamental algorithms for scientific computing in Python." *Nature Methods* 17(3) pp. 261–272 2020.

[46] Hunter, John D.. "Matplotlib: A 2D graphics environment." *Computing in Science & Engineering* 9(3) pp. 90–95 2007.

[47] Mitarai, Kosuke, Negoro, Makoto, Kitagawa, Masahiro, Fujii, Keisuke. "Quantum circuit learning." *Physical Review A* 98(3) pp. 032309 2018. doi: 10.1103/PhysRevA.98.032309.

[48] Schuld, Maria, Killoran, Nathan. "Quantum machine learning in feature Hilbert spaces." *Physical Review Letters* 122(4) pp. 040504 2019. doi: 10.1103/PhysRevLett.122.040504.

[49] Dalzell, Alexander M., others. "How much entanglement do quantum chemistry algorithms require?." *npj Quantum Information* 10 pp. 8 2024. doi: 10.1038/s41534-024-00808-x. https://www.nature.com/articles/s41534-024-00808-x

[50] IBM Quantum. "IBM lays out clear path to fault-tolerant quantum computing." 2023. https://www.ibm.com/quantum/blog/large-scale-ftqc