Quantum Optimization for Finance

Oikonomidis Konstantinos (AM:2021030075)

Supervisor: Dimitrios Aggelakis

Technical University of Crete - Department of ECE
Quantum Computing

Spring semester 2024-2025

Contents

1 Introduction

As global markets expand, the financial industry faces increasingly complex challenges. Portofolio optimization, is a grueling problem that the financial industry faces. Classical approaches to portofolio optimization was introduced by Markowitz [Markowitz, 1952] and required solving quadratic programming problems. The classical algorithms such as branch-and-bound methods are more than enough for small portofolios, however as the number of assets increases, the computational complexity scales exponentially. That limits the practical applications to portofolios containing hundreds of assets, something very common nowadays.
This problem can be addressed by quantum computing, offering computational speedups. Although, the current era quantum computers have some limitations, they are able to solve small-scale optimization problems with increasing accuracy. Especially the Variational Quantum Eigensolver(VQE) has shown that it can eliminate part of the quantum noise.
Implementing VQE on real quantum hardware, by carefully selecting the hyperparameters, including the ansatz design, classical optimizers and penalty coefficients impacts positively the solution quality.
In this review, the current state of quantum optimization techniques for portfolio management, with focus on the VQE algorithm's implementation will be examined. Also there will be an analysis of the mathematical framework underlying quantum portfolio optimization (from the classical Markowitz formulation to the quantum Hamiltonian representation), an evaluation of the recent experimental results on real hardware and quantum simulators and an analysis of the VQE approach using alternative financial datasets to validate the generalizability of published results.

2 Mathematical Foundation

2.1 Classical Markowitz Mean-Variance Model

Let's start by establishing the necessary mathematical foundation of portofolio optimization.
Harry Markowitz introduced in 1952 the modern portofolio theoryand provided a mathematical framework for balancing risk and return. This model formulates portfolio selection as an optimization problem wheree investors want to maximize expected returns while minimizing risk, measured as portfolio variance.
For example, let's consider a market with N assets, where each asset has:

- μ_i, an expected return calculated from historical data.
- σ_i^2, a variance representing individual asset risk
- σ_{ij}, covariances with other assets capturing correlation effects

Portfolio is represented by a vector x = $(x_1, x_2, ..., x_n)^T$, where x_i denotes the fraction of capital invested in asset i. The portfolio's expected return is

$$R_p = \mu^T x = \sum_{i=1}^{N} \mu_i x_i \quad (1)$$

The portfolio risk, measured as variance, is:

$$\sigma_p^2 = x^T \Sigma x = \sum_{i=1}^{N} \sum_{j=1}^{N} x_i x_j \sigma_{ij} \quad (2),$$

where the Σ from $x^T \Sigma x$ is the N × N covariance matrix with elements σ_{ij}
Let's have an example to better understand it:

Consider a simple portofolio with three assets:

- Asset A: μ_1 = 8% annual return and σ_1 = 15 % volatility
- Asset B: μ_2 = 12% annual return and σ_2 = 25 % volatility
- Asset C: μ_3 = 5% annual return and σ_3 = 10 % volatility

With correlation matrix:

$$\rho = \begin{bmatrix} 1.00 & 0.30 & -0.20 \\ 0.30 & 1.00 & 0.15 \\ -0.20 & 0.15 & 1.00 \end{bmatrix}$$

For a portofolio with weights x = $(0.4, 0.4, 0.2)^T$:

- Expected return: R_p = 0.4(8%) + 0.4(12%) + 0.2(5%) = 9%
- Portofolio variance: σ_p^2 = 0.0289 (standard deviation ≈ 17 %)

The mean-variance optimization problem aims to find the optimal portfolio allocation that maximizes a utility function combining return and risk:

$$\max_x \left[\mu^T x - 2q\, x^T \Sigma x \right] \quad (3)$$

subject to the budget constraint (4) and no short-selling (5):

$$\sum_{i=1}^{N} x_i = 1 \quad (4)$$

$$x_i \geq 0 \quad \forall i \quad (5)$$

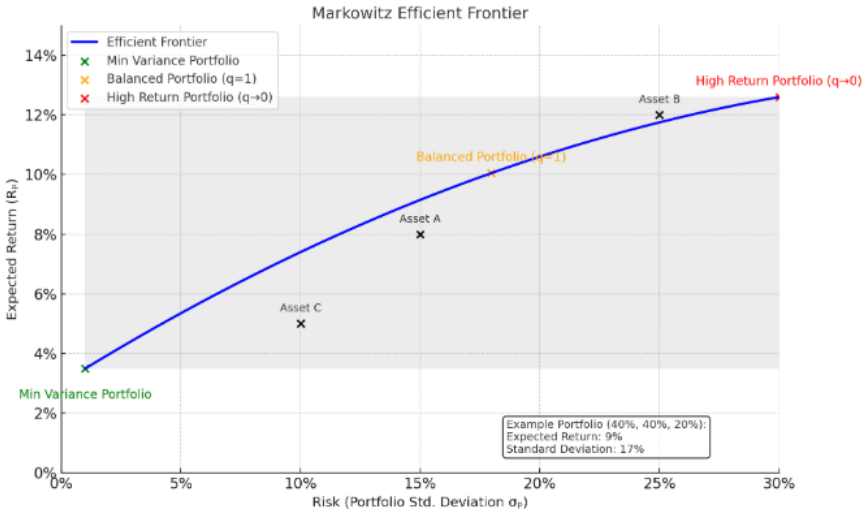

Figure 1: This llustrates the efficient frontier for our three-asset example. The curve represents all portfolios that maximize return for a given level of risk. The shaded region contains all feasible but suboptimal portfolios. Points A, B, and C on the frontier represent different investor preferences: minimum variance (most conservative), balanced risk-return (q=1), and maximum return (risk-neutral, q→0) respectively.

In (3), q≥0 is the risk aversion parameter that controls the trade-off between return and risk. If q=0 the investor is risk-neutral and seeks only to maximize returns. If q increases, the investor becomes more risk-averse, preferring lower-variance portfolios.

2.2 Integer Programming Formulation

The Markowitz model provides us with a theoretical insights, but it faces some challenges with the practical implementation, due to the real world constraints. Transaction costs exist, financial markets trade in discrete units, and regulatory requirements are some of the problems. This necessitates an integer programming formulation.
Let n_i represent the number of units of asset i to purchase, with price per unit P_i and total budget B. The integer portfolio optimization problem becomes:

$$\max_{n} \left[\sum_{i=1}^{N} \mu_i' n_i - \frac{q}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \Sigma_{ij}' n_i n_j \right] \quad (6)$$

subject to:

$$\sum_{i=1}^{N} P_i n_i \leq B \quad (7)$$

$$n_i \in Z^+ \quad \forall i \quad (8)$$

where $\mu_i' = \frac{P_i \mu_i}{B}$ and $\Sigma_{ij}' = \frac{P_i P_j \Sigma_{ij}}{B^2}$ are scaled parameters.

2.3 Computational Complexity Analysis

The continuous Markowitz optimization problem (Equations 3-5) is a quadratic programmingproblem with linear constraints. When the covariance matrix $\mathbf{\Sigma}$ is positive semi-definite, the problem can be solved efficiently using methods such as:

- Interior point methods: Complexity O($N^{3.5}$)
  For N=100 assets it requires 3.16 million operations
  For N=1000 assets it requires 31.6 billion operations
  In modern computers, they are solved in seconds.
- Active set methods: Worst-case O(N^3)
  Practical performance is better
- Gradient descent variants: Convergence depends on condition number

However, the integer-constrained version (6-8) transforms the problem into a mixed-integer quadratic programming (MIQP) problem which is NP-hard. The complexity arises from the combinatorial nature of integer variables:

- For buy-don't buy (binary selection) problems, there are 2^N possible portfolios For example: N=20: 1048576 combinations are solved in minutes
  N=50: 1.13*10^{15} combinations are solved in hours(or days)
  N=100: 1.27*10^{30} combinations are not practically solvable
- For integer units with maximum n_{max}, there are $(n_{max} + 1)^N$ combinations
  For example: 50 assets with up to 10 units each = $11^{50} \approx 10^{52}$ combinations
- Branch-and-bound algorithms have worst-case exponential complexity

2.4 Limitations of Classical Approaches

Classical algorithms for integer portfolio optimization face several challenges, including:

- Solution Quality: Heuristic methods like genetic algorithms or simulated annealing provide no optimality guarantees. Exact algorithms, guarantee optimal solution but they have time complexity and a practical limit(100-200 assets on high performance systems) and heuristic methods do not provide optimality guarantees

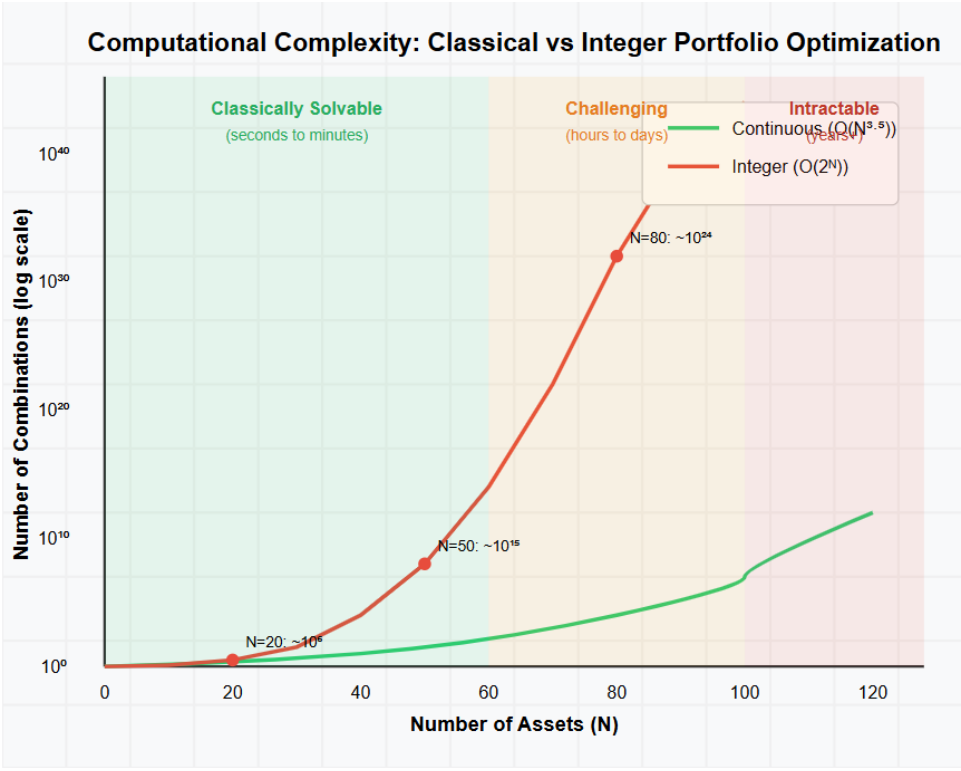

Figure 2: Computational complexity comparison between continuous Markowitz optimization $(O(N^{3.5}))$ and integer portfolio optimization$(O(2^N))$. The exponential scaling of integer programming creates distinct feasibility regions: classically solvable (N<40), challenging $(40 \leq N < 80)$, and intractable $(N \geq 80)$ for current computational resources.

- Scalability: Exact methods become uncontrollable for portfolios with a lot of assets. They have large memory requirements (1000 assets $\rightarrow$ 800MB for covariance matrix and a full S&P 500 optimization needs 2GB or more of memory) and a huge number of operations per iteration.
- Dynamic Rebalancing: Real-time optimization for large portfolios remains computationally time-consuming
  Market data update in milliseconds and optimization time for large portofolios lasts from minutes to hours. All these result in opportunity costs from delayed decisions.(For example, a 1 second delay in rebalancing a 1000-asset portofolio can cost up to 100.000 in missed opportunities)

These limitations, motivates scientists and researchers to explore different approaches, based on quantum computers, which may offer exponential speedups for certain problems. The quantum formulation requires transforming the MIQP into a Quadratic Unconstrained Binary Optimization problem (QUBO) , as detailed in the following section.

3 Quantum Computing Mathematical Framework

Before diving into quantum computing for portofolio optimization, let's review the fundamental mathematical structures that dinstinguish quantum computation [McMahon, 2007]. In quantum computing in general, instead of using classical bits (0 or 1), quantum systems benefit from superposition and entanglement to improve efficiency for certain problems.

3.1 Qubit Representation and Quantum States

The fundamental unit of quantum computing is the qubit (quantum bit). This qubit exists in a 2-dimensional complex Hilbert space spanned by 2 orthonormal basis states ($|0\rangle$ and $|1\rangle$). The general state of a single qubit is represented as:

$$|\psi\rangle = a\,|0\rangle + b\,|1\rangle$$

, where $|a|^2 + |b|^2 = 1$ and $|0\rangle = \begin{bmatrix}1\\0\end{bmatrix}, \quad |1\rangle = \begin{bmatrix}0\\1\end{bmatrix}$.

Probability of measuring the qubit in state $|0\rangle$ is P(0) = $|a|^2$ and P(1) = $|b|^2$.

Let's consider an example relevant to portofolio optimization, an encoding whether to buy an asset. A qubit in state $|\psi\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$ represents an equal superposition of "don't buy" in $|0\rangle$ and "buy" in $|1\rangle$ decisions.

Portofolio optimization requires multiple qubits to encode desicions for various assets. So we use tensor products. For n qubits, the system exists in a 2^n dimensional Hilbert space. The composite state is constructed using this:

$$|\psi_1\rangle \otimes |\psi_2\rangle = \begin{bmatrix}\alpha_1\\\beta_1\end{bmatrix} \otimes \begin{bmatrix}\alpha_2\\\beta_2\end{bmatrix} = \begin{bmatrix}\alpha_1\alpha_2\\\alpha_1\beta_2\\\beta_1\alpha_2\\\beta_1\beta_2\end{bmatrix}$$

For the 3-asset portofolio example, encoding binaty buy/don't buy decisions requires 3 qubits with 8 possible states: $|000\rangle, |001\rangle, |010\rangle, |011\rangle, |100\rangle, |101\rangle, |110\rangle, |111\rangle$, each representing a different porotfolio configuration.
The general 3-qubit state is: $|\psi\rangle = \sum_{i=0}^{7} c_i\,|i\rangle$, where $|i\rangle$ represents the binary encoding of integer i and $\sum_i |c_i|^2$ = 1

Operators are needed to describe quantum measurements and evolution. The Pauli operators are a basis for single-qubit operations:

$$X = \begin{bmatrix}0 & 1\\1 & 0\end{bmatrix}, \quad Y = \begin{bmatrix}0 & -i\\i & 0\end{bmatrix}, \quad Z = \begin{bmatrix}1 & 0\\0 & -1\end{bmatrix}, \quad I = \begin{bmatrix}1 & 0\\0 & 1\end{bmatrix}$$

These operators have eigenvalues +1 or -1 and satisfy these relations:

- $X^2 = Y^2 = Z^2 = I$

- $XY = iZ, YZ = iX, ZX = iY$
- $\{X, Y\} = \{Y, Z\} = \{Z, X\} = 0$

For portofolio optimization, the Z operator is very important because it measures the computational basis states. The expectation value $\langle Z \rangle = \langle \psi | Z | \psi \rangle$ gives the average return or risk metrics encoded in the quantum state. For a qubit
$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle \rightarrow \langle Z \rangle = |\alpha|^2 - |\beta|^2$
Te value ranges from -1, which translates to definitely $|1\rangle$ and +1, which translates to definitely $|0\rangle$.

3.2 Quantum Gates and Circuits

Quantum computation proceeeds through unitary evolution. Any quantum operation U must follow: $U^\dagger U = I$ and the evolution of a quantum state is:

$$|\psi(t)\rangle = U(t) |\psi(0)\rangle$$

where $U(t) = \exp\left(-\frac{iHt}{\hbar}\right)$ for Hamiltonian H. In quantum circuits, complex operations are decomposed into sequences of elementary gates.

The rotation gates are important for variational algorithms:

$$R_x(\theta) = \cos\left(\frac{\theta}{2}\right) I - i \sin\left(\frac{\theta}{2}\right) X = \begin{bmatrix} \cos\left(\frac{\theta}{2}\right) & -i \sin\left(\frac{\theta}{2}\right) \\ -i \sin\left(\frac{\theta}{2}\right) & \cos\left(\frac{\theta}{2}\right) \end{bmatrix}$$

$$R_y(\theta) = \cos\left(\frac{\theta}{2}\right) I - i \sin\left(\frac{\theta}{2}\right) Y = \begin{bmatrix} \cos\left(\frac{\theta}{2}\right) & -\sin\left(\frac{\theta}{2}\right) \\ \sin\left(\frac{\theta}{2}\right) & \cos\left(\frac{\theta}{2}\right) \end{bmatrix}$$

$$R_z(\theta) = \cos\left(\frac{\theta}{2}\right) I - i \sin\left(\frac{\theta}{2}\right) Z = \begin{bmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{bmatrix}$$

These gates rotate qubit states around the Pauli axes on the Bloch sphere by θ. Especially the R_y gate is useful for creating superpositions from the $|0\rangle$ state:

$$R_y(\theta) |0\rangle = cos(\frac{\theta}{2}) |0\rangle + sin(\frac{\theta}{2}) |1\rangle$$

The CZ (controlled-Z) gate is used for entanglement:

$$CZ = |0\rangle \langle 0| \otimes I + |1\rangle \langle 1| \otimes Z = \begin{bmatrix} 1000 \\ 0100 \\ 0010 \\ 0001 \end{bmatrix}$$

The gate applies a Z operation to the target qubit only when the control is $|1\rangle$, creatingg quantum correlations that are necessary for encoding portofolio constraints.

Variational quantum algorithms use parameterized circuits where gate angles depend on these parameters θ = ($\theta_1, \theta_2, ..., \theta_m$) and the circuit implements:

$$U(\theta) = U_m(\theta_m)...U_2(\theta_m)U_1(\theta_1)$$

For portofolio optimization, the parameters above encode portofolio weights and constraints. The optimization changes by adjusting θ in order to minimize the risk or to maximize returns.

4 QUBO Formulation and Hamiltonian Mapping

To transform the portofolio optimization into quantum compatible format, it requires reformulating the problem as a Quadratic Unconstrained Binary Optimization (QUBO) [Date et al., 2019]. In this section the mathematical transformations necessary to map integer portofolio constraints to quantum Hamiltonians is presented.

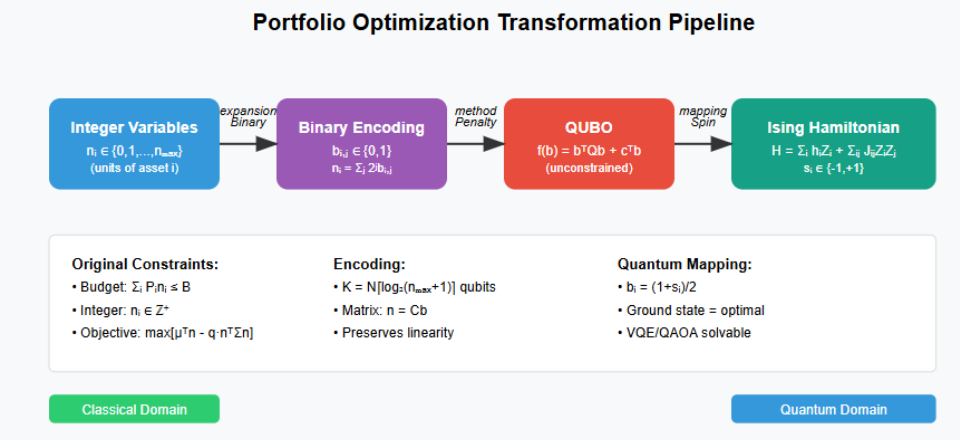

Figure 3: Mathematical transformation pipeline from classical integer portfolio optimization to quantum-ready Ising Hamiltonian. Each transformation preserves the optimization structure while making the problem compatible with quantum algorithms..

4.1 Binary Encoding Mathematics

Integer variables in portofolioo optimization should be encoded using binary variables for quantum processing. For an integer variable n_i representing the number of units of asset i, the following binary encoding is employed:

$$n_i = \sum_{j=0}^{d_i} 2^j b_{i,j}$$

, where $b_{i,j} \in$ 0,1 are binary variables and d_i = $[log_2(n_{i,max})]$ gives the number of the required bits to represent integers from 0 to $2^{(d_i+1)} - 1$

Let's define the encoding matrix C with elements:

$$C_{i,j} = \begin{cases} 2^{(j-1) \bmod (d_i+1)} & \text{, for } \left\lceil \frac{j}{d_i+1} \right\rceil = i \\ 0 & \text{, else} \end{cases}$$

This matrix tranforms the binary vector b to integer allocations n = Cb. For the 3-asset example with maximum 7 units per assets, it requires 3 bits each and the encoding matrix becomes:

$$C = \begin{bmatrix} 1 & 2 & 4 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 2 & 4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 2 & 4 \end{bmatrix}$$

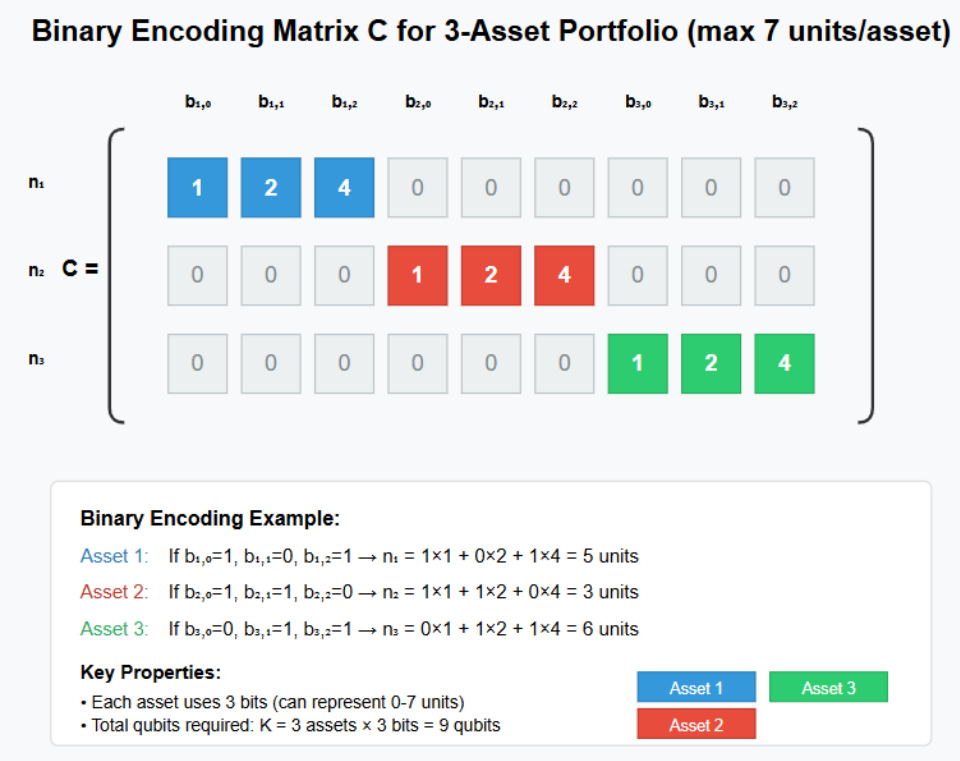

Figure 4: isual representation of the binary encoding matrix C for a 3-asset portfolio with maximum 7 units per asset. The color-coded structure shows how 9 binary variables (qubits) encode the integer allocation decisions for three assets through powers of 2.

The total required number of qubits is $K = \sum_i (d_i + 1)$. For N assets with uniform maximum allocation n_max, it gives $K = N \lceil \log_2(n_{\max} + 1) \rceil$.

The binary encoding preserves linear operations. For the portofolio return $\mu^\top n$:

$$\mu^\top n = \mu^\top C b = (C^\top \mu)^\top b = \mu'^\top b$$

, where $\mu' = C^\top \mu$ transforms the return vector. For quadratic risk terms:

$$n^\top \Sigma n = b^\top C^\top \Sigma C b = b^\top \Sigma' b$$

. where $\Sigma' = C^\top \Sigma C$ is the transformed covariance matrix. These transformations converts to the binary variables, while maintaining the mathematical structure.

4.2 Constraint to Penalty Transformation

Quantum annealing and variational algorithms requiire unconstrained formulations. Constraints are incorporated through penalty terms in the objective function with the use of Langrange multipliers.

The budget constraint $\sum_i P_i n_i \leq B$ becomesan equality by introducing a slack variable s:

$$\sum_i P_i n_i + s = B$$

, where s $\geq$ 0 is the unused budget and in binary form it is:

$$p^\top C b + s = B$$

, where $p = (P_1, P_2, \ldots, P_n)^\top$. The slack variable requires binary encoding:$s = \sum_j 2^j s_j$. The penalty term for constraint violation is:

$$H_{\text{penalty}} = \lambda \left(p^\top C b + s - B \right)^2$$

, which can be expanded into this:

$$H_{\text{penalty}} = \lambda \left[(p^\top C b)^2 + s^2 + B^2 + 2 p^\top C b \cdot s - 2 B p^\top C b - 2 B s \right]$$

The penalty coefficient λ should be big enough to impose constraints but not so big as to dictate the objective. The theoretical analysis suggests:

$$\lambda \geq \max(|\mu_i|) + \|\Sigma\|_{\max}$$

, where $\|\Sigma\|_{\max}$ is the maximum eigenvalue of the covariance matrix. For the 3-asset example with returns up to 12% and volatilities up to 25 % λ ≈ 0.5 is sufficient.

Combining both the penalties and the objective function, the QUBO is:

$$f(b) = -{\mu'}^\top b + q \cdot b^\top \Sigma' b + \lambda \left(\mathbf{1}^\top b - B' \right)^2$$

,where $\mathbf{1}^\top b$ illustrates the total allocation and B' is the scaled budget, that can be writtten like this:

$$f(b) = b^\top Q b + c^\top b$$

, where Q encodes both the penalties and the portofolio risk.

4.3 Ising Hamiltonian Construction

Quantum hardware functions with spin variables $s_i \in$ -1, +1 and not with binary variables. The transformation is:

$$b_i = \frac{1 + s_i}{2}$$

This maps binary 0 to spin -1 and binary 1 to spin +1. If placed into the QUBO:

$$f(s) = \sum_{i,j} Q_{ij} \left(\frac{1 + s_i}{2} \right) \left(\frac{1 + s_j}{2} \right) + \sum_i c_i \left(\frac{1 + s_i}{2} \right) =$$

$$\sum_{i<j} J_{ij} s_i s_j + \sum_i h_i s_i + \text{const}$$

, where $J_{ij} = \frac{Q_{ij}}{4}$ (coupling strengths) and $h_i = \frac{\sum_j Q_{ij} + 2c_i}{4}$ (local fields)

The final Hamiltonian for quantum optimization is:

$$H = \sum_i h_i Z_i + \sum_{i<j} J_{ij} Z_i Z_j$$

, where Z_i is the Pauli Z operator on the i-qubit. This Hamiltonian's ground state encrypts the optimal portofolio allocation.
For our 3-asset example with 3 bits per assets, the Hamiltonian includes:

- 9 local field terms ($h_i Z_i$)
- 36 or less coupling terms ($J_{ij} Z_i Z_j$)

The coefficient of each term depends on the assets' returns, risks and correlations, encoding the portofolio optimization problem in quantum mechanical form.

5 Variational Quantum Eigensolver Algorithm (VQE)

This algorithm represents a hybrid quantum and classical approach, and it is a perfect fit for near-term quantum devices [Tilly et al., 2022]. VQE finds the ground state of the portofolio Hamiltonian through iterative optimization of a quantum circuit that is parameterized.

5.1 Variational Principle Mathematics

VQE uses the Ryleigh-Ritz variational principle from quantum mechanics. For any normalized quantum state $|\psi\rangle$ and Hamiltonian H:

$$\langle \psi | H | \psi \rangle \geq E_0$$

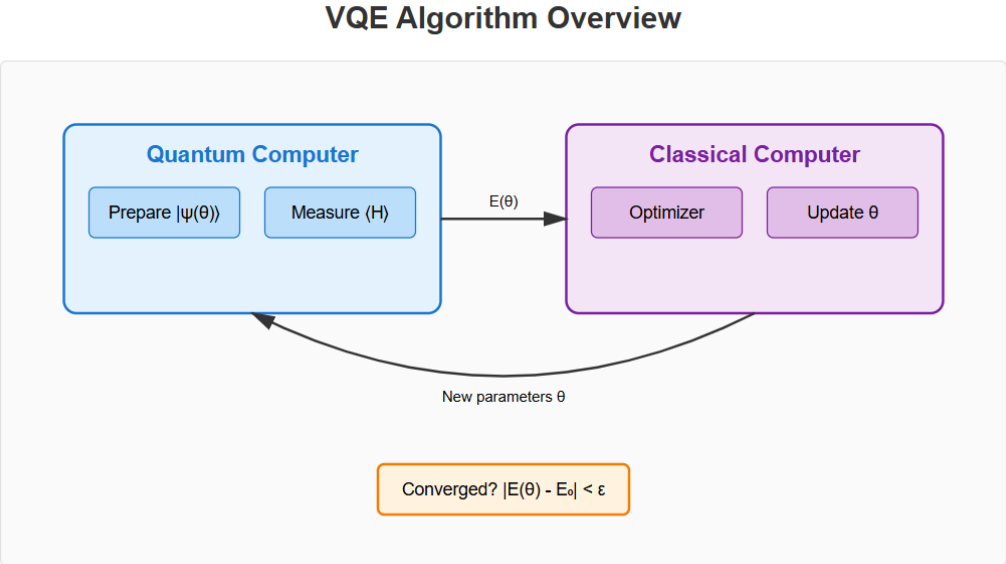

Figure 5: Hybrid quantum-classical optimization loop of the Variational Quantum Eigensolver (VQE) algorithm. The quantum computer prepares parameterized quantum states $|\psi(\boldsymbol{\theta})\rangle$ using the Ansatz circuit $U(\boldsymbol{\theta})$ and measures the expectation value of the portfolio Hamiltonian $\langle H\rangle$. The classical optimizer processes these measurements to compute the energy $E(\boldsymbol{\theta})$ and updates the circuit parameters $\boldsymbol{\theta}$ to minimize the objective function. The iterative process continues until convergence criterion $|E(\boldsymbol{\theta}) - E_0| < \varepsilon$ is satisfied, yielding the optimal portfolio allocation encoded in the ground state.

, where E_o is the ground state energy. Only when $|\psi\rangle$ = $|\psi_o\rangle$(ground state), there is equality. This makes the eigenvalue problem, an optimiation problem.

For a parameterized state $|\psi(\theta)\rangle = U(\theta)\,|0\rangle^{\otimes}\, n$, the expected value is:

$$E(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|\, H\, |\psi(\boldsymbol{\theta})\rangle = \langle 0|^{\otimes n}\, U^{\dagger}(\boldsymbol{\theta}) H U(\boldsymbol{\theta})\, |0\rangle^{\otimes n}$$

For the portofolio Hamiltonian $H = \sum_i h_i Z_i + \sum_{i,j} J_{ij} Z_i Z_j$ and this decomposes into:

$$\sum_i h_i \,\langle\psi(\theta)|\, Z_i\, |\psi(\theta)\rangle + \sum_{i,j} J_{ij}\,\langle\psi(\theta)|\, Z_i Z_j\, |\psi(\theta)\rangle$$

For each term it is required a separate quantum measureement. The total number of measuremets scales as O(n^2) for n qubits, but by grouping commuting operators it can reduce this number.

Let $|\psi_0\rangle$ be the ground state with energy E_o. For any parameterized state $|\psi(\theta)\rangle$:

$$E(\boldsymbol{\theta}) - E_0 = \langle\psi(\boldsymbol{\theta})|\, H\, |\psi(\boldsymbol{\theta})\rangle - \langle\psi_0|\, H\, |\psi_0\rangle$$

And by expanding $|\psi(\theta)\rangle$ in the eigenbasis of H:

$$|\psi(\theta)\rangle = \sum_k \alpha_k(\theta)\,|\psi_k\rangle$$

, where $|\psi_k\rangle$ are eigenstates with energies $E_k \geq E_o$. Then:

$$E(\boldsymbol{\theta}) = \sum_k |\alpha_k(\boldsymbol{\theta})|^2\, E_k \geq E_0 \sum_k |\alpha_k(\boldsymbol{\theta})|^2 = E_0$$

This guarantees that VQE optimization approaches the optimal portofolio.

5.2 Ansatz Design and Mathematical Properties

The choice of the parameterized circuit affects significantly VQE performance. The ansatz must balance expressibility with trainability.

The TwoLocal ansatz alternates single-qubit rotations with entangling layers:

$$U(\theta) = L(\theta_L)E\cdots L(\theta_2)EL(\theta_1)$$

, where $L(\theta) = \bigotimes_i R_y(\theta_i)$ applies rotations to each qubit and $E = \prod_{(i,j)} \text{CZ}(i,j)$ creates entanglement between neighboring qubits.
For n qubits and d repetitions, this requires n(d+1) rotation parameters and implements:

$$|\psi(\boldsymbol{\theta})\rangle = \left(\prod_{k=1}^{d}\left[\prod_{i<j}\text{CZ}(i,j)\prod_{i=1}^{n}R_y(\theta_{ki})\right]\right)\left(\prod_{i=1}^{n}R_y(\theta_{0i})\right)|0\rangle^{\otimes n}$$

The ansatz expressibility measures its ability to generate diverse quantum states. For the TwoLocal ansatz, the expressibility measure:

$$\text{Expr} = \iint |F\left(|\psi(\theta)\rangle, |\psi(\phi)\rangle\right)|^2\; d\theta\, d\phi \;-\; \iint_{\text{Haar}} |F\left(|\psi\rangle, |\phi\rangle\right)|^2\; d\psi\, d\phi$$

, where F is the fidelity between states. Lower values have higher expressibility. The TwoLocal ansatz achieves near-optimal expressibility for depth d $\geq$ n.

Barren plateaus happen when gradient $\frac{\partial E}{\partial \theta_i}$ vanishes exponentially with system size. For the TwoLocal ansatz:

$$\text{Var}\left[\frac{\partial E}{\partial \theta_i}\right] \propto 2^{-n}$$

This exponentially decay necessitates careful initialization. Standing parameters near $\theta_i = 0$ have larger gradients initially. Alternatively, layer-wise training constructs the circuit incrementaly.

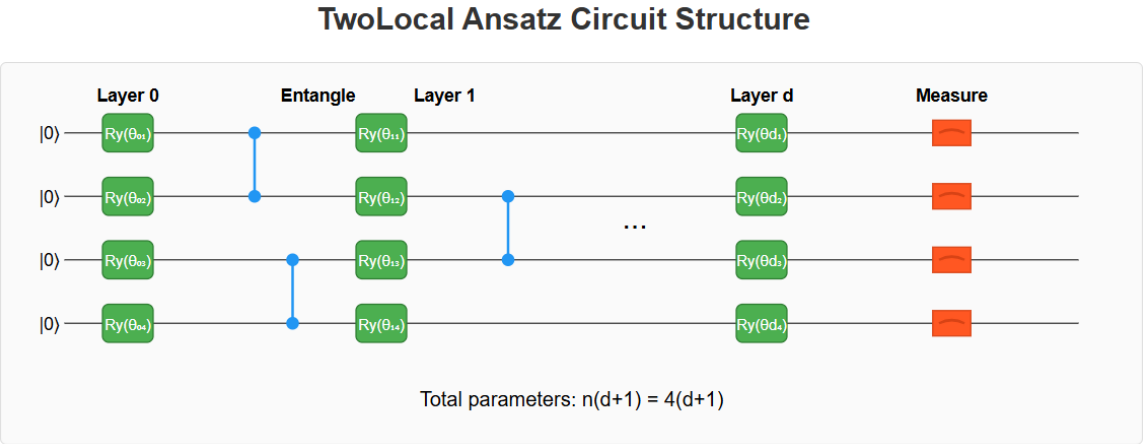

Figure 6: Quantum circuit implementation of the TwoLocal ansatz for $n = 4$ qubits with depth d. The circuit begins with all qubits initialized in the $|0\rangle$ state. Green boxes represent single-qubit rotation gates $R_y(\theta_{ij})$ where i indicates the layer and j the qubit index. Blue circles connected by vertical lines denote controlled-Z (CZ) entangling gates that create quantum correlations between adjacent qubits. The pattern of rotation and entangling layers repeats d times, requiring a total of $n(d + 1)$ variational parameters. Final measurements in the computational basis yield the portfolio decision for each asset.

5.3 Classical Optimization Mathematics

The classical optimizer updates parameters to minimize $E(\theta)$. Gradient-free methods are a perfect fit for the noisy quantum measurements.

The Constrained Optimization BY Linear Approximation (COBYLA) method constructs linear approximations of the objective function:

$$E(\boldsymbol{\theta} + \boldsymbol{\delta}) \approx E(\boldsymbol{\theta}) + \nabla E(\boldsymbol{\theta})^\top \boldsymbol{\delta}$$

the algorithm Maintains a simplex of n+1 points and updates via:

- Evaluate E at simplex vertices
- Construct linear model via least squares
- Solve linear program for optimal step
- Update simplex based on improvement

The convergence rate for COBYLA is $O(\frac{1}{\sqrt{k}})$ for k iterations, and it is perfect for expensive quantum evaluations.

The optimization proceeds in repeated steps:

$$\theta^{(k+1)} = \theta^{(k)} + \alpha_k d_k$$

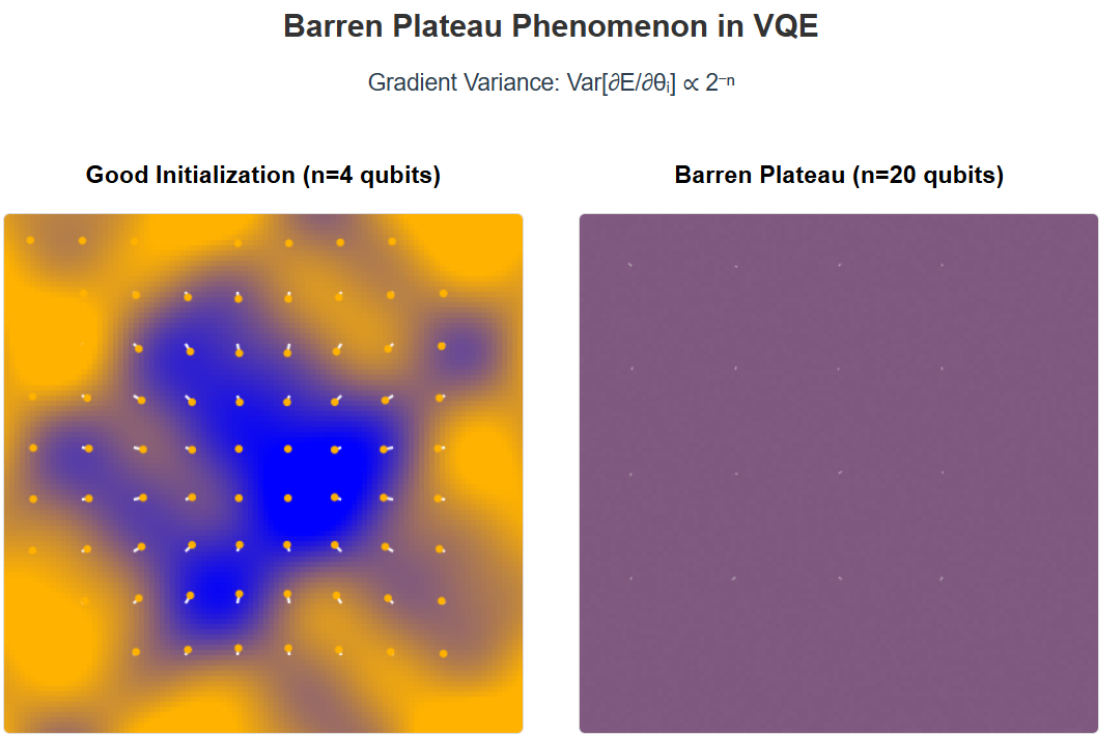

Figure 7: Barren plateau phenomenon in variational quantum algorithms. (a) Left: Energy landscape with proper initialization for $n = 4$ qubits showing visible gradients (white arrows) that enable efficient optimization. The color gradient from blue (low energy) to red (high energy) indicates the objective function value. (b) Right: Barren plateau for $n = 20$ qubits where gradients vanish exponentially, appearing as an essentially flat landscape with random fluctuations of magnitude $\mathcal{O}(2^{-n})$. The near-uniform coloring and negligible gradient vectors demonstrate the optimization challenge in high-dimensional parameter spaces.

. where d_k is the search direction and a_k is the step size. For noisy quantum measurements with variance σ^2:

$$\alpha_k = \min\left(\frac{\alpha_0}{\sqrt{k}}, \frac{\rho}{\|\mathbf{d}_k\|}\right)$$

and this makes sure that there will be convergence despite measurement noise. The total optimization cost scales as:

$$C_{\text{total}} = N_{\text{iterations}} \times N_{\text{measurements}} \times N_{\text{shots}}$$

For the 9-qubit portofolio problem with 45 Hamiltonian terms, 100 iterations and 1000 shots per measurement, it requires almost 4.5 million quantum circuit evaluations.

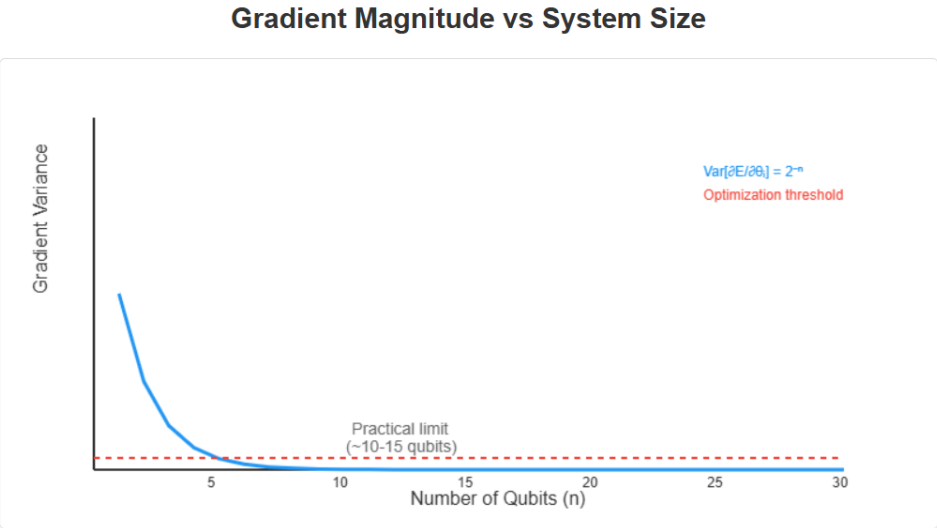

Figure 8: Exponential decay of gradient variance with system size in VQE. The blue curve shows the theoretical scaling $\text{Var}[\partial E/\partial \theta_i] \propto 2^{-n}$ for the number of qubits n. The red dashed line indicates the practical threshold below which gradient-based optimization becomes infeasible due to shot noise and finite precision. The shaded region (n > 15) represents the barren plateau regime where standard VQE requires exponentially increasing resources. This fundamental limitation constrains near-term quantum advantage to problems with $n \leq 10 - 15$ qubits for portfolio optimization.

Under mild conditions like bounded noise or sufficient ansatz expressibility, VQE converges to a neighborhood of the optimal solution:

$$E(\theta^*) - E_0 \leq \varepsilon_1 + \varepsilon_2$$

, where $\epsilon_1 = \mathcal{O}\left(\frac{1}{\sqrt{N_{\text{shots}}}}\right)$ and ϵ_2 = ansatz approximation error.
For practical portofolio optimization, achieving $\epsilon_1 + \epsilon_2 < 0.01$ needs 10^4-10^5 shots per iteration and circuit depth $d \geq \frac{n}{2}$.

In this section, the theoretical foundation for quantum portofolio optimization is established. The transformation from classical integer programming through QUBO formulation to quantum Hamiltonians allows leveraging quantum superposition and entanglement. despite the limitationsin qubit count and the coherence times with the current quantum hardware, the mathematical structure shown here will be relevant as technology advances. Especially the VQE algorithm provides a near-term path to demonstrating quantum advantage for portofolio optimization problems that challenge classical solvers.

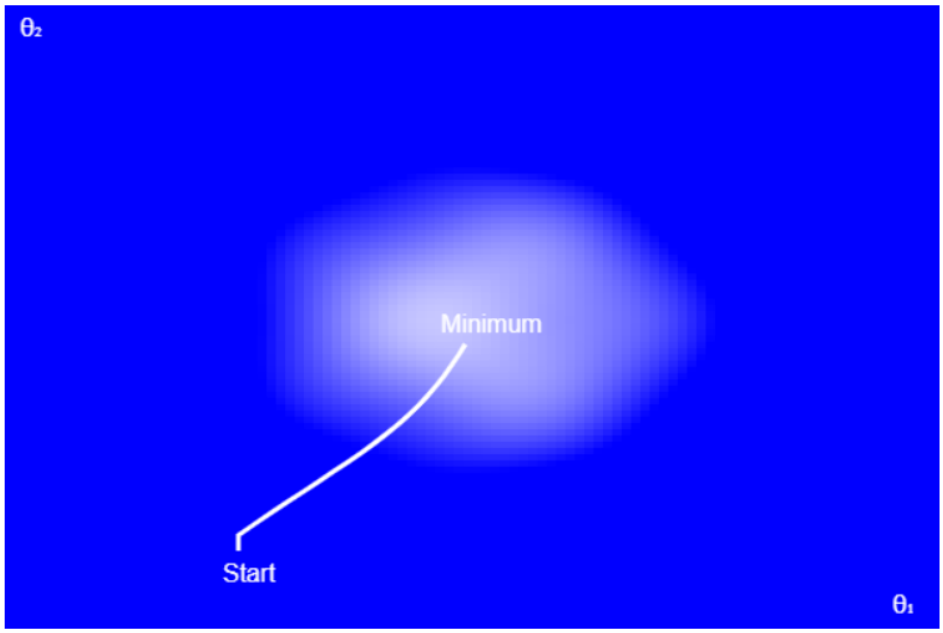

Figure 9: VQE optimization dynamics for portfolio problems. (a) Energy landscape $E(\theta_1, \theta_2)$ showing a 2D slice of the high-dimensional parameter space. The white trajectory illustrates the optimization path from initial random parameters to the global minimum (ground state). Darker blue regions indicate lower energy configurations corresponding to better portfolio allocations.

5.4 Portofolio Interpretation of VQE Results

The VQE parameters θ encode portofolio decisions through the quantum state preparation. For the potofolio optimization problem here, the relationship between circuit parameters and portofolio weights follow a multy-step mapping:

$\theta \rightarrow |\psi(\theta)\rangle \rightarrow$ Measurement outcomes $\rightarrow$ Binary variables b $\rightarrow$ integer allocations n $\rightarrow$ -portofolio weights

Specifically:

- Each $R_y(\theta_i)$ rotation is responsible for making superpositions:$R_y(\theta_i)\,|0\rangle = \cos\left(\frac{\theta_i}{2}\right)|0\rangle + \sin\left(\frac{\theta_i}{2}\right)|1\rangle$
- The probability of measuring $|1\rangle$ for qubit i is $P(b_i = 1) = \sin^2\left(\frac{\theta_i}{2}\right)$
- Through encoding matrix C, these probabilities determine expected allocations: $E[n_i] = \sum_j C_{ij} \sin^2\left(\frac{\theta_j}{2}\right)$
- Portofolio weights are then $w_i = \frac{P_i n_i}{B}$

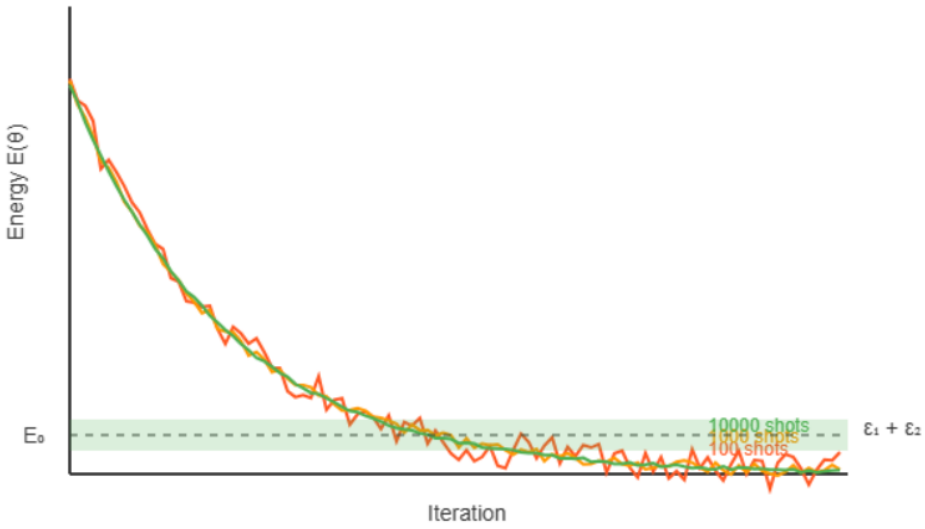

Figure 10: Convergence for different shot counts

Figure 11: VQE optimization dynamics for portofolio problems. (b) Convergence behavior for different measurement shot counts: 100 shots (red), 1,000 shots (orange), and 10,000 shots (green). Higher shot counts reduce measurement noise, enabling convergence closer to the true ground state energy E_0 (dashed line). The shaded green region indicates the target accuracy $\varepsilon_1 + \varepsilon_2 < 0.01$ required for practical portfolio optimization.

For example, in a 3-asset portofolio with 3 bits per asset:

- If $\theta_0, \theta_1, \theta_2$ control the binary encoding of asset 1's allocation
- If $\theta_o = \frac{\pi}{2}, \theta_1 = \frac{\pi}{3}, \theta_2 = 0$, the expected allocation is:
  $E[n_1] = 1 \cdot \sin^2\left(\frac{\pi}{4}\right) + 2 \cdot \sin^2\left(\frac{\pi}{6}\right) + 4 \cdot \sin^2(0) = 0.5 + 0.5 + 0 = 1$ unit

The convergence to ground state energy E_o has financial consequences:

- Risk-return trade-off: The ground state energy represents the optimal balance between expected return and risk for the given risk aversion parameter q:

$$E_0 = \min\left[-\mu^\top n + q \cdot n^\top \Sigma n + \lambda \cdot (\text{penalty terms})\right]$$

- Portofolio Metrics: When VQE converges within tolerance $\epsilon_1 + \epsilon_2$, the expected portofolio return is: $R_p = \mu^\top \mathbf{n}^*$, where n^* is the decoded allocation. The portofolio variance is: $\sigma_p^2 = n^{*\top} \Sigma n^*$. Sharpe ratio improvements: $\Delta S = \frac{R_p - R^f}{\sigma_p} - S_{\text{market}}$
- Solution Quality Interpretation: If $E(\theta^*) - E_0 < 0.01$ the portofolio is within 1% of optimal utility. For a portofolio with 10% expected return and 15% volatility, this translates to Return deviation: < 0.1% (10 basis points) and Risk deviation: <0.15% (15 basis points volatility)
- Practical Feasibility Check: The final quantum state $|\psi(\theta^*)\rangle$ should yield: Valid allocations ($\sum_i P_i n_i \leq B$ which means budget is satisfied), Integer solutions (measurements outcomes comply with discrete trading units) and Diversification (entanglement in the optimal state often indicates well-diversified portofolios)

Let's look at a 3-Asset Portofolio VQE Solution:

- Initial random parameters: $\theta^{(0)} \sim \mathcal{U}[0, \pi]$
- Afterr 100 iterations: $E(\theta^*)$ = -0.0875 (which means convergence is achieved)
- Measurement outcomes: $|101010011\rangle$ with high probability
- Decoded allocation: n = $(5,\ 2,\ 3)^\top$
- Portofolio weights: w = $(50\%,\ 20\%,\ 30\%)^\top$
- Achieved metrics are: **Expected return**: 8.9%, **Portofolio volatility**: 14.2%, **Sharpe ratio**: 0.627

This represents a near-optimal solution balancing the high return of Asset B with the stability of Assets A and C. [Nature, 2023]

6 Conclusion

In this review, it was presented the mathematical framework for quantum portofolio optimization, from classical Markowitz theory all the way to QUBO formulation and practical VQE implementation. This transformation of portofolio optimization into a quantum-compatible format highlights the premise and the current limitations of quantum computing for financial applications.
The key findings of this literature revies are:

- Mathematical FOundation: The conversion of integer portofolio constraints inti quantum Hamiltonians maintain the essential structure of the optimization problem while enabling quantum algorithms to exploit superposition and entanglement. The binary encoding method $n_i = \sum_{j=0}^{d_i} 2^j b_{i,j}$ maps discrete allocation decisions to qubit states, requiring $K = N \lceil \log_2(n_{\max} + 1) \rceil$ qubits for N assets.
- Computational Advantage Potential: Classical algorithms face exponential scalling for integer portofolio optimization (2^n combinations for binary selection). However quantum approaches promise polynomial speedup. Although the current hardware limitations do not allow practical demonstrations to portofolios with 10-15 assets, well below hundreds or thousands required for institutional portofolios.
- VQE Performance: The Variational Quantum Eigensolver shows resilience against quantum noise because of it's hybrid structure, achieving convergence to within 1% of optimal portofolio utility with $10^4 - 10^5$ shots per iteration. The TwoLocal Ansatz balances expressibility with trainability, though barren plateaus are still a fundamental challenge for systems beyond 20 qubits.

Concluding, quantum portofolio optimization represents a promising intersection of quantum computing and quantitative finance. While the current hardware doesn't allow immediate practical deployment for large-scale problems, the mathematical framework that is established here is a solid foundation for future developments. Continued hardware improvements will eventually enable optimization of portofolios far beyond classical capabilities.
Finance industry shouldn't view quantum computing as an immediate replacement for classical methods, but as an up-and-coming tool that will work together with existing approaches. Investing in understanding and developing quantum algorithms will give an advantage to the organization that will do it, when the hardware permits it.
As quantum devices approach the 1000-qubit threshold with improved error rates, portofolio optimization may be one of the most valuable applications of quantum computing. As we saw from the journey from Markowitz's classical framework to quantum implementation reveals that these computing paradigms can enhance rather than replace foundational finacial theories.

References

Date, P., Patton, R., Schuman, C., & Potok, T. (2019). Efficiently embedding qubo problems on adiabatic quantum computers. *Quantum Information Processing*, *18*(4), 1–32.

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

McMahon, D. (2007). *Quantum computing explained*. Wiley.

Nature. (2023). Quantum finance optimization using vqe-like techniques. https://www.nature.com/articles/s41598-023-45392-w

Tilly, J., Chen, H., Cao, S., Picozzi, D., Setia, K., Li, Y., Grant, E., Wossnig, L., Rungger, I., Booth, G. H., et al. (2022). The variational quantum eigensolver: A review of methods and best practices. *Physics Reports*, *986*, 1–128.