# Quantum Computing for Portfolio Optimization

**Where Quantum meets Finance:**
Possibilities, Progress, and Practical Limits

qBraid

SOUTH CAROLINA QUANTUM

qbraid.com | scquantum.org

# Overview

Portfolio optimization plays a central role in finance. It involves identifying an investment mix that maximizes returns while managing risk, all within a set of predefined constraints. This work includes both static allocation decisions and dynamic strategies that adjust over time. As portfolios grow in complexity, with more assets, tighter constraints, and longer time horizons, the difficulty of solving these problems increases. Classical computing methods can address many of these challenges, but they become less efficient when real-world factors such as transaction costs, liquidity, and volatility are introduced.

Quantum computing offers a possible alternative. Algorithms developed for combinatorial optimization, including the Variational Quantum Eigensolver (VQE) and the Quantum Approximate Optimization Algorithm (QAOA), are being tested for their ability to manage larger and more complex problems. These approaches could support portfolio managers, quantitative researchers, and institutional investors who need to make faster and more informed decisions. Although the technology is still in the early stages of adoption, research and early applications suggest that quantum methods may eventually provide valuable tools for tackling high-dimensional portfolio optimization in financial services.



This white paper is part of a commissioned series from SC Quantum examining real-world case studies of quantum technologies deployed in industry. Developed in collaboration with qBraid, the series offers insight into how forward-looking organizations are integrating quantum solutions into their operations today. South Carolina Quantum (SC Quantum), a 501(c)(3) established in 2022, brings together academia, entrepreneurs, industry, and government to develop collaborative solutions that strengthen the state's position in a rapidly evolving field. https://qbraid.com | https://scquantum.org

# 01 - Introduction to Portfolio Optimization

Portfolio Optimization (PO) is a foundational practice in modern finance. It is widely used by investment funds, pension plans, and both institutional and individual investors. The goal is to allocate a fixed budget across a group of assets in a way that seeks to maximize returns while managing risk. This approach applies to a range of use cases, from long-term pension management to more active strategies that require regular rebalancing in response to changing market conditions.

At its core, portfolio optimization involves balancing two competing priorities: increasing expected returns and minimizing risk. Risk is often measured as the variance or volatility of returns, which reflects how much actual performance may differ from projections. To navigate this tradeoff, investors rely on statistical inputs such as each asset's expected return, its return variance, and how it interacts with other assets in the portfolio. These values are typically estimated using historical market data.

Mathematically, portfolio optimization is often framed as a constrained quadratic optimization problem. The objective is to find a mix of assets that provides strong expected returns while limiting overall portfolio risk. These two elements, return and risk, are weighted using a risk-aversion parameter that reflects how much volatility an investor is willing to accept. This formulation is the foundation of the mean variance optimization model introduced by Harry Markowitz, which remains one of the most widely used approaches in the field. The solutions define what is known as the efficient frontier, a collection of portfolios that deliver the highest possible return for a given level of risk. An investor's position on this frontier depends on their goals and tolerance for risk.

Although the basic model assumes a single time period and focuses on two objectives, return and risk, real-world applications are often more complex. Constraints such as budget limits, short selling restrictions, sector exposure caps, and turnover limits are commonly applied. In many cases, investors also need to manage portfolios over multiple time horizons. As the number of assets and constraints increases, the computational demands grow as well. These factors make portfolio optimization a strong candidate for exploring the potential advantages of quantum computing.

# 02 - Using Classical Computing for Portfolio Optimization

Classical portfolio optimization is often framed as a constrained quadratic optimization problem, as outlined in the previous section. In the continuous case, where asset weights are expressed as fractions of the total budget, the problem can be solved exactly using standard mathematical methods. These include Lagrange multipliers for equality constraints and the Karush–Kuhn–Tucker (KKT) conditions for handling inequality constraints. Since the covariance matrix of asset returns is positive semi-definite, the optimization reduces to solving a linear system, typically using Cholesky decomposition. This approach involves computing the inverse of an N-by-N matrix, where N is the number of assets. The computational complexity scales on the order of $O(N^3)$, which is manageable for small portfolios but becomes increasingly costly as the number of assets grows.

When the optimization problem includes integer or binary decision variables, such as discrete units of assets or on-off investment decisions, it becomes a combinatorial problem. The number of possible portfolio configurations increases exponentially with the number of assets, which makes brute-force search infeasible. Classical exact methods like branch and bound can theoretically solve these problems, but their runtimes become prohibitive as market size grows.

To manage this complexity, practitioners often use heuristics or metaheuristics to steer the search toward high-quality solutions without exhaustively evaluating every possibility. Common approaches include Genetic Algorithms (GA), Simulated Annealing (SA), and Particle Swarm Optimization (PSO). These techniques are especially helpful for large-scale or non-convex problems where exact solvers struggle. Although they can produce useful results, they do not guarantee global optimality and may require extensive tuning. Scalability also remains a challenge, since most classical methods begin to break down when applied to portfolios with more than a few hundred assets.

In real-world settings, the choice of method depends on portfolio size, the nature of the constraints, and the level of precision required. Exact solvers such as quadratic programming are well suited for small to medium-sized continuous problems. Heuristic methods are more common in large, discrete, or heavily constrained cases. As portfolios grow in complexity and scale, even the most advanced classical techniques face practical limits. These constraints suggest that quantum approaches may offer meaningful gains in speed or scalability.

# 03 - Potential of Quantum Solutions

Quantum computing introduces a different model of computation. It uses principles such as superposition, entanglement, and interference to explore large solution spaces in ways that may outperform classical methods for certain types of problems. In theory, this capability allows quantum algorithms to offer speed advantages for tasks that are considered intractable on classical systems, including NP-hard problems like large-scale portfolio optimization. Although quantum computers may not always produce provably optimal solutions, they have the potential to generate high-quality approximations in far less time. As the underlying hardware continues to evolve, the practical relevance of these methods is expected to grow.

Several quantum methods have been explored for portfolio optimization. The Quantum Approximate Optimization Algorithm (QAOA) is seen as a strong candidate because it can encode complex constraints and search efficiently for near-optimal solutions. Variational Quantum Eigensolver (VQE) techniques have also shown potential, particularly when the problem is reformulated as a Quadratic Unconstrained Binary Optimization (QUBO) and mapped to an Ising Hamiltonian. Research has shown that, with careful parameter tuning, VQE can produce results close to classical optima even on today's small-scale quantum devices. As hardware improves, performance is expected to scale as well.

Quantum annealers, such as those developed by D-Wave, have also been applied to portfolio problems and have demonstrated competitive performance in benchmark studies. Other approaches, like the Harrow–Hassidim–Lloyd (HHL) algorithm for solving linear systems, may be relevant when the problem is expressed in linear algebraic form, although current noisy intermediate-scale quantum (NISQ) devices are not yet capable of implementing them effectively. In the meantime, hybrid quantum–classical strategies are gaining traction. These methods assign the most computationally demanding subproblems to quantum processors while relying on classical systems for the rest of the workflow.

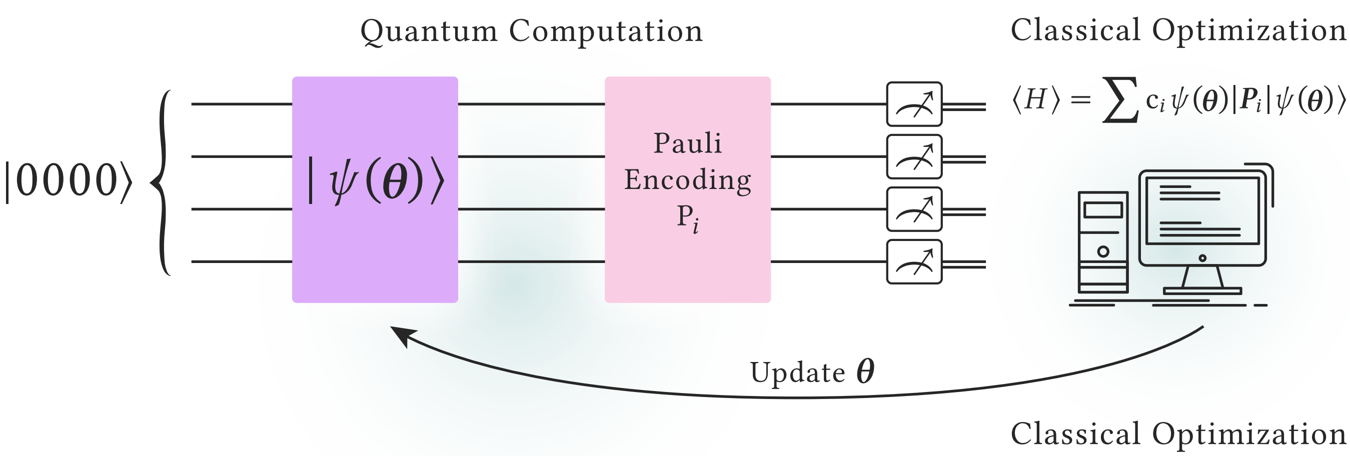


Figure 1: Schematic of VQE algorithm for an abstracted optimization problem.
Original image from Scientific Reports volume 13, Article number: 19434 (2023)

The practical value of these algorithms depends on thorough, end-to-end assessments of their real-world performance. Important evaluation factors include the use of realistic input models, the ability to load data efficiently into quantum systems, and evidence of quantum speedup at problem sizes that matter to industry. Most near-term quantum algorithms trade formal guarantees for experimental feasibility, but ongoing research and steady hardware progress suggest that quantum approaches could eventually outperform classical methods for large, complex portfolio optimization problems, especially in settings where speed and scalability are essential.

## 04 - Current Developments in Quantum Portfolio Optimization

Several major financial institutions and technology companies are actively exploring quantum approaches to portfolio optimization. Many of these efforts involve partnerships that combine domain-specific expertise with advanced quantum hardware and algorithm development.

**IBM** has developed a Quantum Portfolio Optimizer built on the Variational Quantum Eigensolver (VQE). The approach reformulates the optimization problem as a Quadratic Unconstrained Binary Optimization (QUBO), maps it to an Ising Hamiltonian, and executes the process on quantum hardware. The workflow incorporates noise-aware post-processing to improve result quality, producing optimized investment trajectories based on historical asset prices and user-defined constraints.

**AWS** and **Goldman Sachs** jointly conducted a full resource analysis of the Quantum Interior Point Method (QIPM) applied to portfolio optimization. Although QIPM showed initial promise, the study found it unlikely to be viable in the near term without major improvements in quantum random access memory (QRAM). For now, more practical alternatives involve reformulating portfolio problems into binary structures and applying variational algorithms such as QAOA, or using quantum annealing.

**J.P. Morgan** has made progress using Hybrid HHL++, a modified version of the Harrow–Hassidim–Lloyd (HHL) algorithm adapted for current hardware. By compressing circuit structures and improving the scaling of linear systems, the team successfully solved small-scale portfolio optimization problems on trapped-ion quantum computers. This work represents one of the largest experi mental implementations of HHL so far.

Together, these initiatives point to growing momentum in applying quantum computing to real-world financial optimization challenges. Although most experiments remain limited by hardware scale, they are helping refine algorithms, estimate resource needs, and establish workflows that could support larger, more practical applications as quantum systems continue to evolve.

## 05 - South Carolina Landscape

South Carolina's financial landscape includes a substantial base of institutional investment activity, most notably the $40 billion portfolio managed by the South Carolina Retirement System Investment Commission (RSIC). As a public pension fund, RSIC must meet long-term return targets while managing risk, maintaining liquidity, and adapting to shifting market conditions. These challenges make it a strong candidate for advanced portfolio optimization methods. With allocations across public equities, fixed income, private markets, and alternatives, the complexity of rebalancing and scenario planning is significant, especially in periods of heightened volatility.

Beyond the public pension system, South Carolina supports a broad and active financial services sector. Banks, credit unions, and insurers manage portfolios and balance sheets under a range of regulatory constraints related to capital, liquidity, and risk. In these environments, optimization tools can enhance decisions related to asset allocation, duration matching, and credit risk management. Smaller investment firms and wealth advisors, particularly in growing high-net-worth markets like Charleston, Greenville, and Columbia, also stand to benefit from more systematic approaches to portfolio construction that improve risk-adjusted returns and downside protection.

Within this broader context, portfolio optimization brings clear value to South Carolina's financial ecosystem. Public pensions can use these tools for strategic asset allocation, liability-driven investment planning, and scenario analysis. Private-sector institutions can apply them to improve risk management, meet regulatory demands, and allocate capital more effectively. Classical optimization methods remain the standard today. However, rising complexity and scale suggest that more advanced approaches, including quantum-based models, may play a growing role as the technology becomes more accessible.

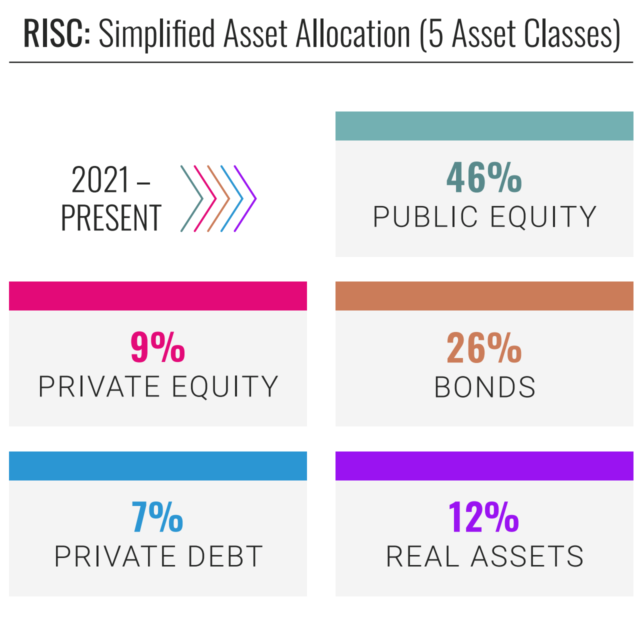


Figure 2: RSIC began implementing a simplified Policy Portfolio consisting of five asset classes in 2020 with underlying complexity. Source: https://rsic.sc.gov/what-we-do/portfolio-assets-chart.html

## 06 - Recommendations

With such a dynamic and rapidly growing landscape, it is important to invest in quantum technologies early. It can be daunting to start, but there are already steps that can be taken.

**Here are a few recommendations, if you're on the business side:**

- Start early: Don't let caution be the enemy of growth. Foster broad awareness and engagement with quantum technologies among leadership teams and encourage all workers to explore quantum technologies.
- Build quantum expertise: Explore and implement training options to ensure your teams are quantum-literate and quantum-ready to create and maintain a competitive advantage through proper use cases.
- Identify concrete applications: Become proficient as an organization in evaluating how quantum techniques can help your company meet your goals. Focus on separating hype from reality, and identify clear use cases.

**If you're on the technical side, here are a few recommendations to get started:**

- Refresh and develop new skills: Refresh your quantum skills or upskill by taking courses in quantum mechanics, programming, and applications.
- Attend conferences: Quantum computing is a fast-evolving landscape, so meet fellow scientists, researchers, and engineers at conferences. Connect with them to explore collaborations and keep up with research.
- Connect with domain experts: Be aware of domain experts and existing work being done in your field. Keep abreast of technological advancements, papers, and patents in your domain.

## 07 - What's Next?

Whether you're building quantum expertise, identifying use cases, refreshing and developing skills, or connecting with domain experts, qBraid offers tools to help you take the next step through interactive courses, a cloud-based sandbox, and ready-to-use tutorials. They're designed to get you moving quickly. To learn more or schedule a conversation, visit qbraid.com. If you want to plug into quantum momentum in South Carolina and across the Southeast, visit scquantum.org. Getting started in quantum can feel complex, but with the right partners, the path forward becomes clear, more collaborative, and more rewarding than it might seem at first.

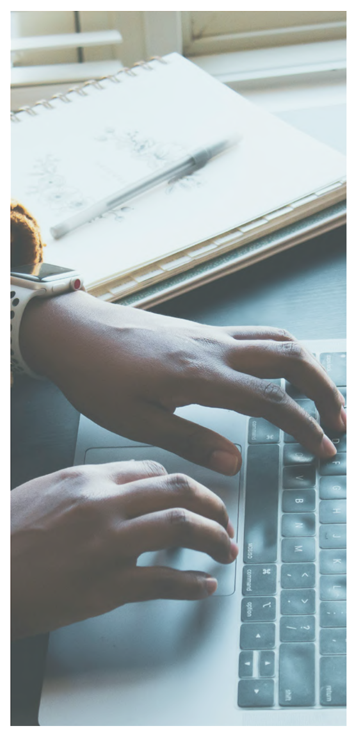

qBraid

**qBraid: Your One-Stop Quantum Platform**

Click here to get started on your quantum journey today. **Questions:** contact@qbraid.com