Maximizing Returns with Quantum-Entangled Portfolio Optimization: A Case Study on the Top 4 U.S. Banks

Risk-Adjusted Portfolio Strategies Using Quantum Simulation and Historical Closing Prices (Jan 2023 – Jan 2025)

Author: Dean Jay Mathew (www.lordllandaff.com)
Publication Date: 10/09/2025

Abstract

Purpose of the Study

The purpose of this study is to explore the use of quantum-inspired simulation techniques to assist in investment decision-making for the top four U.S. banks. By leveraging real historical closing prices and simulating all possible portfolio combinations using a 4-qubit quantum circuit, this research provides a systematic, data-driven approach for evaluating risk-adjusted returns. The study demonstrates how quantum simulation can offer insights and confidence in selecting the most promising bank portfolios. Put simply, following the simulation's recommendations could provide investors with greater confidence in achieving more reliable profits while making better-informed financial decisions.

Methodology

Data Collection
Historical stock data for the top four U.S. banks—JPMorgan Chase (JPM), Bank of America (BAC), Citigroup (C), and Wells Fargo (WFC)—was collected from Yahoo Finance using adjusted closing prices from January 1, 2023, to January 1, 2025. The data was processed to calculate daily returns and the covariance matrix, providing a foundation for evaluating portfolio risk and expected return.

Quantum Simulation
A 4-qubit quantum circuit was constructed using AWS Braket's LocalSimulator, where each qubit represents one of the banks. Hadamard gates were applied to create superposition, allowing the simulation to explore all possible portfolio combinations simultaneously. CNOT gates introduced

entanglement to capture correlations between assets. Each circuit was run for 1,000 shots to obtain a statistical distribution of portfolio bitstrings.

Portfolio Evaluation
For each portfolio configuration represented by a qubit bitstring, expected return and risk (variance) were calculated. A cost function was defined as:

Cost = −Expected Return + λ × Risk

where λ is a tunable risk-aversion parameter. Portfolios were ranked based on this cost function, allowing identification of optimal combinations for different risk preferences.

Key Findings

Key Insights from Lambda (λ) Sensitivity Analysis

1. **λ = 0.0 (No Risk Aversion – Focus on Maximum Return)**
Best Portfolio: All four banks – JPM, BAC, C, WFC
Interpretation: With no penalty for risk, the simulation favors including all assets to maximize expected return. Ideal for highly aggressive investors willing to accept full market variability.

2. **λ = 0.5 (Low Risk Aversion – Slight Emphasis on Risk)**
Best Portfolio: All four banks remain optimal
Interpretation: Introducing a small risk penalty slightly reduces expected return but encourages broad diversification. Suitable for investors seeking growth with moderate caution.

3. **λ = 1.0 (Moderate Risk Aversion – Balanced Return/Risk Trade-off)**
Best Portfolio: JPM, C, WFC
Interpretation: At moderate risk aversion, portfolios begin to drop one bank (BAC) to reduce overall risk while maintaining strong expected returns. Suited for balanced investors.

4. **λ = 1.5 (High Risk Aversion – Risk Reduction Prioritized)**
Best Portfolio: JPM, WFC
Interpretation: Higher risk aversion favors smaller portfolios with fewer assets. Focus is on minimizing risk over maximizing returns, highlighting strongest individual performers.

5. **λ = 2.0 (Very High Risk Aversion – Minimal Risk Focus)**
Best Portfolio: JPM only
Interpretation: At very high risk aversion, only the single least risky bank is selected. Expected returns are lower, but portfolio volatility is minimized. Suitable for highly conservative investors prioritizing capital preservation.

Business Takeaway

- **Lower λ:** Aggressive strategy, favoring higher returns with diversified assets.
- **Medium λ:** Balanced strategy, maintaining strong returns while reducing risk slightly.
- **Higher λ:** Conservative strategy, focusing on fewer, lower-risk assets to protect capital.

Significance / Implications for Portfolio Management

- **Dynamic Risk-Return Trade-off:** Portfolio composition should adapt to investor risk appetite. Low λ values favor diversified, high-return portfolios; high λ values favor smaller, lower-risk portfolios.
- **Prioritization of Asset Selection:** As λ increases, the model highlights the most stable assets, helping managers identify "core holdings."
- **Strategic Decision-Making Tool:** Quantum simulation explores all portfolio combinations quickly, providing confidence in portfolio selection.
- **Transparency and Confidence:** Quantifying expected return and portfolio risk gives clear insights for defensible investment decisions.
- **Innovation and Competitive Advantage:** Leveraging quantum simulation signals advanced analytical capability, differentiating firms in a data-driven investment landscape.

Abstract Summary

The study confirms that quantum-assisted simulations can be a practical, scientifically grounded tool for optimizing portfolios according to risk preference, providing clarity and confidence for investment decisions.

Introduction

Portfolio optimization is a core activity in investment management, guiding decisions to balance risk and return. Quantum computing offers a novel approach, enabling simultaneous exploration of all portfolio combinations and more nuanced correlation capture. This report explores 4-bank portfolios using AWS Braket LocalSimulator to find risk-adjusted optimal combinations.

Comparison with Classical Portfolio Modeling

- **Exploration of All Portfolios Simultaneously:** Quantum-inspired simulation evaluates all portfolios at once; classical approaches must evaluate sequentially or rely on heuristics.

- **Capturing Correlations via Entanglement:** Entanglement encodes asset correlations directly; classical methods require covariance matrices.

- **Probabilistic Output vs Deterministic Output:** Quantum simulation outputs probability distributions; classical methods provide a single solution or Monte Carlo approximation.

- **Risk-Return Trade-Off via Lambda:** Adjusting λ in a quantum circuit produces simultaneous distributions for all portfolios; classical approaches require recomputation for each λ.

- **Scalability:** Quantum circuits scale exponentially with qubit count; classical combinatorial optimization becomes infeasible as assets increase.

Business Interpretation: Even with four banks, quantum-inspired simulation quickly explores all risk-return scenarios, revealing a spectrum of viable portfolios. For larger asset pools, classical brute-force methods would be too slow, while quantum simulation can handle them efficiently.

Methodology

- **Data:** JPM, BAC, C, WFC; period 01/01/2023 – 01/01/2025; source Yahoo Finance; adjusted close prices.

- **Quantum Simulation:** AWS Braket LocalSimulator; 4 qubits; Hadamard gates for superposition; CNOT gates for entanglement; 1,000 shots.

- **Portfolio Evaluation:** Expected return = mean of historical returns; risk = portfolio variance; cost function = −Expected Return + λ × Risk; λ adjustable from 0 (ignore risk) to 2 (risk-weighted heavily).

Results

Portfolio Bitstring Frequencies

The simulation generated all 16 possible portfolio combinations for the four banks. The frequency distribution over 1,000 shots for each bitstring shows which portfolios were more likely to be optimal given the cost function and λ value. This provides a probabilistic ranking rather than a single deterministic result.

The following tables show the *frequency for λ = 0.0, 0.5, 1.0, 1.5, and 2.0. That is to say, the following tables show the optimal investment strategy across all Top 4 U.S. banks based on an investor's risk appetite from λ = 0.0 (large risk appetite) to 2.0 (low risk appetite).*

```
lambda_risk          0.00

Top Portfolios (lambda_risk=0.0):
1111 -> ['JPM', 'BAC', 'C', 'WFC'], Cost: -0.004584
1011 -> ['JPM', 'C', 'WFC'], Cost: -0.003810
1101 -> ['JPM', 'BAC', 'WFC'], Cost: -0.003430
1110 -> ['JPM', 'BAC', 'C'], Cost: -0.003277
0111 -> ['BAC', 'C', 'WFC'], Cost: -0.003237

Best Portfolio Overall: ['JPM', 'BAC', 'C', 'WFC'] (Cost: -0.004584)
```

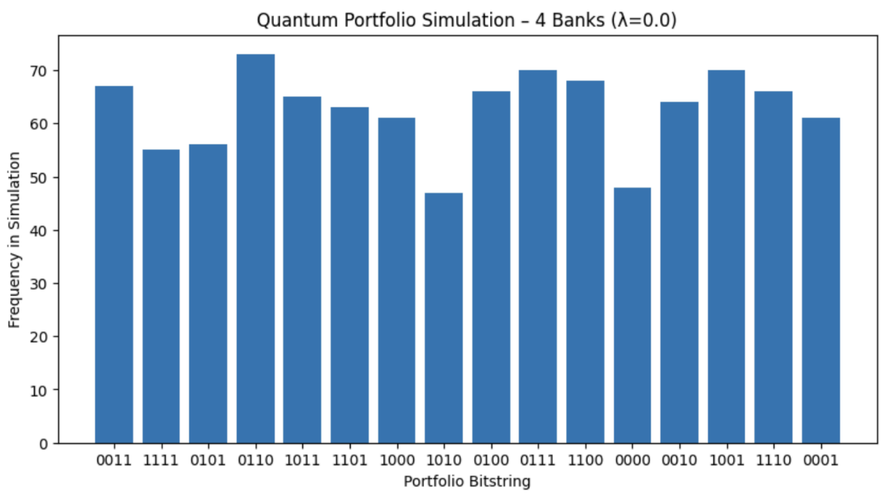

Figure 1: Portfolio frequency distribution for λ = 0.0 (No Risk Aversion), showing all four banks included to maximize expected return.

lambda_risk 0.50

```
Top Portfolios (lambda_risk=0.5):
1111 -> ['JPM', 'BAC', 'C', 'WFC'], Cost: -0.002941
1011 -> ['JPM', 'C', 'WFC'], Cost: -0.002872
1101 -> ['JPM', 'BAC', 'WFC'], Cost: -0.002487
1110 -> ['JPM', 'BAC', 'C'], Cost: -0.002405
1001 -> ['JPM', 'WFC'], Cost: -0.002218

Best Portfolio Overall: ['JPM', 'BAC', 'C', 'WFC'] (Cost: -0.002941)
```

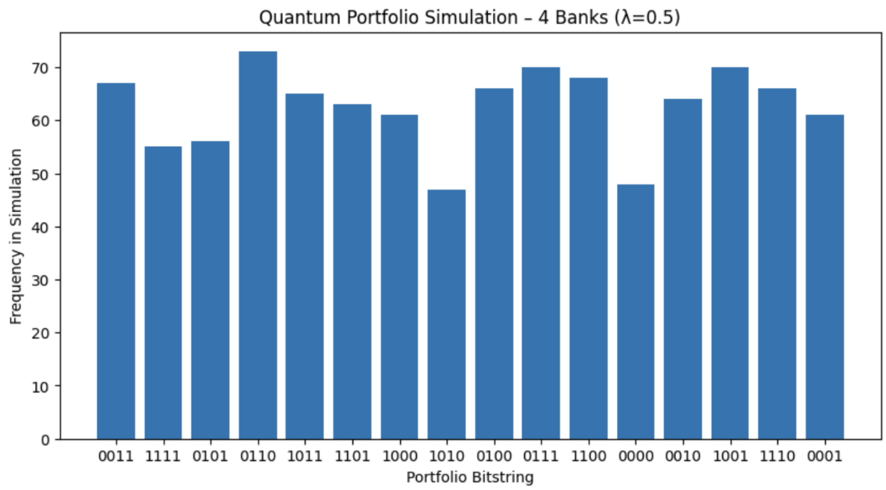

Figure 2: Portfolio frequency distribution for λ = 0.5 (Low Risk Aversion), illustrating slight risk weighting while still favoring all banks.

lambda_risk 1.00

```
Top Portfolios (lambda_risk=1.0):
1011 -> ['JPM', 'C', 'WFC'], Cost: -0.001935
1001 -> ['JPM', 'WFC'], Cost: -0.001781
1010 -> ['JPM', 'C'], Cost: -0.001721
1101 -> ['JPM', 'BAC', 'WFC'], Cost: -0.001545
1110 -> ['JPM', 'BAC', 'C'], Cost: -0.001532

Best Portfolio Overall: ['JPM', 'C', 'WFC'] (Cost: -0.001935)
```

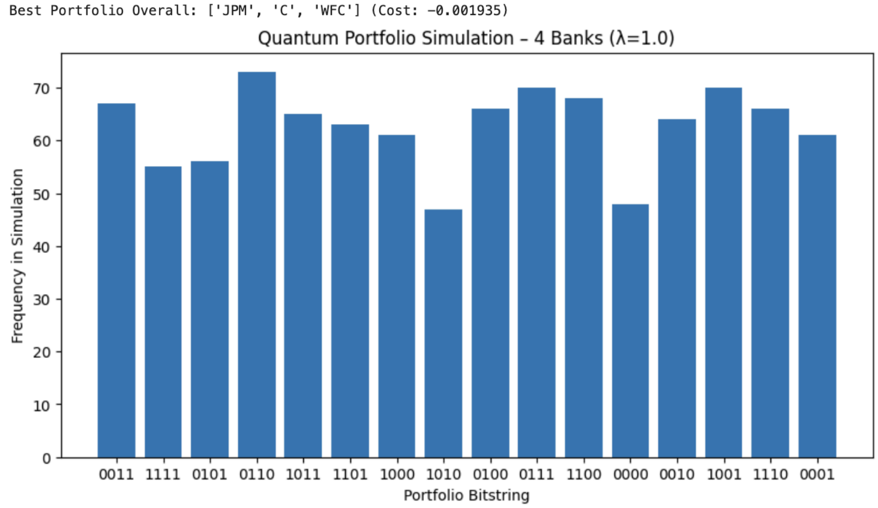

Figure 3: Portfolio frequency distribution for λ = 1.0 (Moderate Risk Aversion), highlighting the exclusion of BAC to reduce portfolio risk while maintaining strong returns.

lambda_risk 1.50

```
Top Portfolios (lambda_risk=1.5):
1001 -> ['JPM', 'WFC'], Cost: -0.001344
1010 -> ['JPM', 'C'], Cost: -0.001330
1000 -> ['JPM'], Cost: -0.001055
1011 -> ['JPM', 'C', 'WFC'], Cost: -0.000998
1100 -> ['JPM', 'BAC'], Cost: -0.000987

Best Portfolio Overall: ['JPM', 'WFC'] (Cost: -0.001344)
```

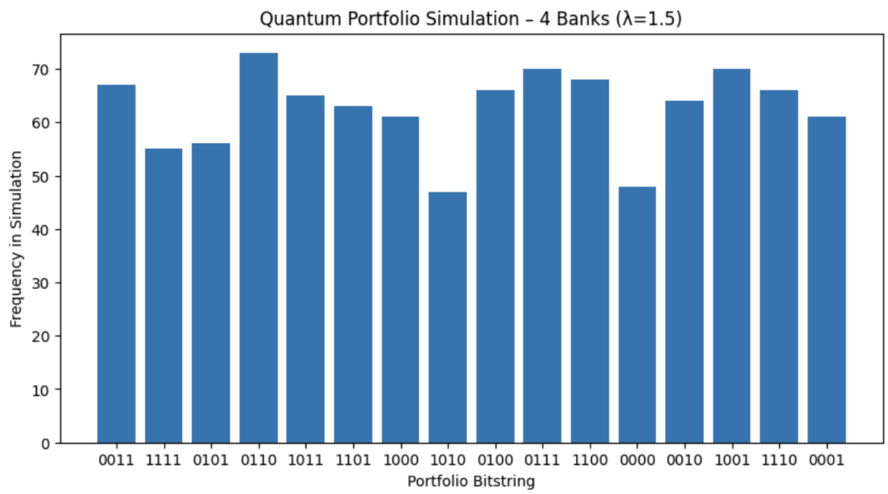

Figure 4: Portfolio frequency distribution for λ = 1.5 (High Risk Aversion), showing selection concentrated in JPM and WFC to prioritize risk reduction.

lambda_risk 2.00

```
Top Portfolios (lambda_risk=2.0):
1000 -> ['JPM'], Cost: -0.000958
1010 -> ['JPM', 'C'], Cost: -0.000940
1001 -> ['JPM', 'WFC'], Cost: -0.000907
0001 -> ['WFC'], Cost: -0.000675
0010 -> ['C'], Cost: -0.000623

Best Portfolio Overall: ['JPM'] (Cost: -0.000958)
```

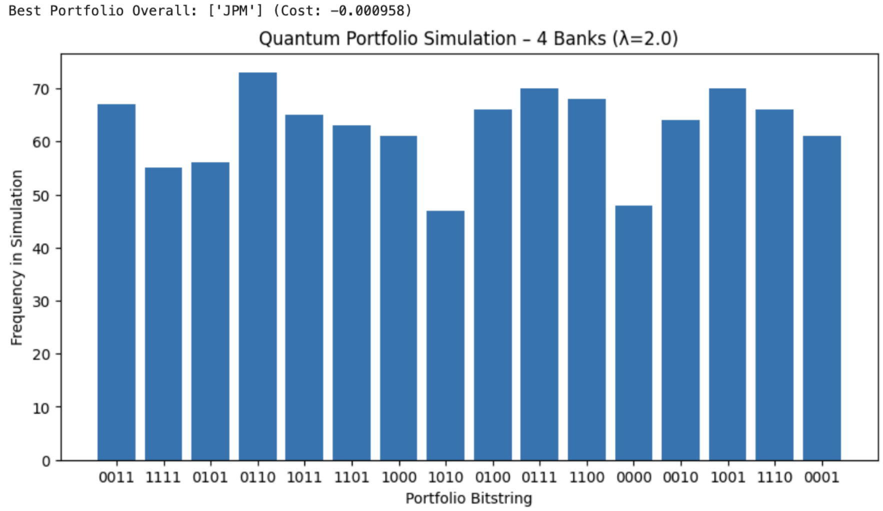

Figure 5: Portfolio frequency distribution for λ = 2.0 (Very High Risk Aversion), with only JPM selected to minimize volatility and preserve capital.

Recommendations

Diversification Patterns and Risk-Return Trade-offs

1. **Low λ (0.0 – 0.5): Highly Aggressive / Growth-Oriented Portfolios**

   Quantum Assessment Results: Include **all four banks (JPM, BAC, C, and WFC).**

   When risk is not penalized, or only slightly considered, the optimal portfolios include **all four banks (JPM, BAC, C, and WFC)**. This approach maximizes expected returns by capturing the upside potential across the entire sector. However, it comes with the highest exposure to market fluctuations. Such a strategy would appeal to highly aggressive investors who prioritize growth and are comfortable absorbing significant volatility in pursuit of higher profits.

2. **Moderate λ (1.0): Balanced Risk-Return Portfolios**

   Quantum Assessment Results: Exclude **Bank of America (BAC).**

   At a mid-level risk aversion, the model begins to filter out higher-risk assets. In this case, **Bank of America (BAC)** is excluded, leaving a portfolio composed of JPM, Citigroup (C), and Wells Fargo (WFC). This adjustment reduces overall portfolio variance while maintaining strong return potential. The outcome represents a **balanced strategy**, suitable for investors seeking to participate in sector growth but with a controlled level of downside risk.

3. **High λ (1.5 – 2.0): Defensive / Capital-Preservation Portfolios**

   Quantum Assessment Results: Exclude **Citigroup (C)** and **Bank of America (BAC)**. Additionally exclude **Wells Fargo (WFC)** at λ = 2.0, leaving only **JPMorgan Chase & Co. (JPM).**

   As risk aversion increases further, the portfolios become more concentrated in the most stable performers. At λ = 1.5, the portfolio narrows to **JPM and WFC**, and at λ = 2.0 it reduces further to **JPM alone**. These configurations prioritize **risk reduction and capital preservation** over return maximization. Expected profits are lower, but volatility is minimized, making these portfolios attractive to conservative investors—such as institutions, pension funds, or risk-averse individuals—who prioritize protecting capital even at the expense of growth.

Why Certain Banks Were Favored

JPM consistently appears across all λ levels due to strong historical returns and lower volatility. WFC and C appear selectively; BAC is excluded at higher λ levels due to higher variance.

Limitations

While the findings of this study are promising, several important limitations must be acknowledged:

1. **Restricted Asset Universe**
   This experiment considered only four assets—the top U.S. banks by market capitalization (JPMorgan Chase, Bank of America, Citigroup, and Wells Fargo). Although this makes the problem tractable and illustrative, it does not capture the broader diversification opportunities available in a real investment portfolio that might include additional banks, sectors, or global equities. Results may therefore understate the complexity of real-world optimization.

2. **Simulation Environment**
   The analysis was conducted using AWS Braket's **LocalSimulator**, which provides an idealized environment free from the noise, decoherence, and gate errors that affect actual quantum hardware. While this enables clean testing of portfolio optimization logic, it does not reflect the current practical limitations of running algorithms on real quantum devices. Future work should validate results on actual quantum hardware to assess robustness.

3. **Dependence on Historical Data**
   Portfolio expected returns and risks were calculated using historical adjusted closing prices between January 2023 and January 2025. This assumes that past performance and correlations are reliable indicators of the future, which may not hold in practice—particularly during periods of structural change in banking, regulation, or macroeconomic conditions. Thus, the outcomes should be interpreted as illustrative rather than predictive.

4. **Simplified Cost Function**
   The cost function used, *Cost = –Expected Return + λ × Risk*, while effective for demonstrating the principle of risk-return trade-off, is a simplified abstraction of real-world portfolio optimization. In practice, managers may need to incorporate transaction costs, liquidity constraints, regulatory capital requirements, and other non-linear factors not captured here.

Conclusion

This study demonstrates that quantum-assisted simulations can provide a practical and innovative approach to portfolio optimization, particularly in modeling explicit investor risk preferences. By leveraging the probabilistic nature of quantum circuits, it becomes possible to evaluate all potential portfolio combinations simultaneously, offering investors and managers a broader and more transparent view of trade-offs between risk and return.

The results confirm a clear pattern: as risk aversion (λ) increases, the composition of the optimal portfolio shifts from fully diversified, high-return allocations to more concentrated, lower-risk selections. For aggressive investors, low λ values support broad exposure across all four banks, maximizing expected return. For more conservative investors, higher λ values progressively narrow the portfolio to focus on the most stable banks, ultimately prioritizing capital preservation over growth.

From a business standpoint, this provides **actionable guidance**. Portfolio managers can use the quantum framework not only to identify the most appropriate holdings for each client's risk profile but also to communicate these decisions transparently with data-backed evidence. By quantifying expected returns, risks, and likelihoods of outcomes, the simulation offers **probabilistic confidence**—a defensible foundation for investment decisions that enhances both client trust and internal decision-making.

Finally, the broader significance of this work lies in its role as a bridge between theoretical quantum finance and real-world portfolio management. While the study is limited to a small asset universe and simulated environment, it illustrates how quantum-inspired techniques can enhance existing methods and open the door to more sophisticated applications. **As quantum hardware matures and asset universes expand, such approaches could evolve from experimental demonstrations into competitive advantages for institutions seeking to maximize returns while managing risk in increasingly complex markets.**

Next Steps

Building on the findings of this study, several clear pathways exist for advancing the research and moving closer to real-world application.

First, the asset universe should be expanded beyond the four largest U.S. banks. By including additional banks, sectors, or exchange-traded funds (ETFs), the simulation could provide a richer and more representative view of diversification benefits and cross-sector dynamics. This would also allow the quantum optimization approach to demonstrate its scalability and value in more complex investment environments.

Second, future work should incorporate **real quantum hardware** rather than relying solely on local simulators. Running these portfolio optimization circuits on actual devices would expose the models to hardware noise and practical limitations, offering insight into how well quantum methods can perform under real-world constraints. This step is critical for moving from proof-of-concept simulations toward production-grade quantum financial applications.

Third, exploring **alternative quantum optimization algorithms** such as the Quantum Approximate Optimization Algorithm (QAOA) or the Variational Quantum Eigensolver (VQE) could yield new efficiencies or improve accuracy in capturing the risk-return trade-off. These methods may also provide different perspectives on how best to encode financial problems into quantum circuits.

Fourth, integrating **live market data streams** would enable the framework to evolve into a dynamic decision-support system. Instead of relying on historical price data, real-time feeds could allow portfolio managers to adjust strategies on the fly, creating a genuinely adaptive, quantum-assisted investment tool.

Finally, the most impactful next step is to partner with a financial institution. Collaborating with a bank or asset manager would allow this research to be validated against real investment practices, tested at scale, and refined in line with industry needs. Such a partnership would not only accelerate the practical deployment of quantum portfolio optimization but also position the participating institution at the forefront of financial innovation. Together, these next steps form a roadmap for advancing from an experimental case study toward a fully deployable quantum-enabled portfolio management solution.

Closing Statement

This study demonstrates the practical potential of quantum-assisted portfolio optimization for modeling investor risk preferences and improving decision transparency. By extending this research with larger asset universes, live market data, and real quantum hardware, financial institutions can unlock a new class of tools for portfolio construction. Real progress will depend on collaboration—bringing together banks, researchers, and technology providers to translate these early results into profitable, scalable strategies. The author acknowledges the contributions of open-source developers, AWS Braket, and the Python research community, whose tools and resources made this work possible.

References

1. Yahoo. (n.d.). *Yahoo Finance: Historical market data*. Retrieved from https://finance.yahoo.com (via *yfinance* Python library: https://pypi.org/project/yfinance/).

2. Amazon Web Services. (n.d.). *Amazon Braket Documentation*. Retrieved from https://docs.aws.amazon.com/braket/.

3. McKinney, W. (2010). *Data structures for statistical computing in Python*. Proceedings of the 9th Python in Science Conference, 51–56. (*pandas* library). Retrieved from https://pandas.pydata.org.

4. Harris, C. R., Millman, K. J., van der Walt, S. J., et al. (2020). *Array programming with NumPy*. Nature, 585, 357–362. Retrieved from https://numpy.org.

5. Hunter, J. D. (2007). *Matplotlib: A 2D graphics environment*. Computing in Science & Engineering, 9(3), 90–95. Retrieved from https://matplotlib.org.

6. Jupyter Development Team. (n.d.). *ipywidgets: Interactive widgets for Jupyter notebooks*. Retrieved from https://ipywidgets.readthedocs.io.

7. Orús, R., Mugel, S., & Lizaso, E. (2019). *Quantum computing for finance: Overview and prospects*. Reviews in Physics, 4, 100028. https://doi.org/10.1016/j.revip.2018.12.002