Portfolio Tail-Risk Estimation (VaR/CVaR) from Log-Price Time Series:
Classical Monte Carlo Baselines and Quantum Amplitude Estimation on a Simulator

Superpositions Studio — Final Report

February 15, 2026

Abstract

We study the estimation of portfolio tail-risk metrics, Value at Risk (VaR) and Conditional Value at Risk (CVaR), from two equity log-price time series (AAPL and MSFT). The data exhibit an irregular trading-day-like calendar, a contiguous missing block for MSFT in the raw row order (terminal in the table layout), moderate positive cross-asset dependence on the overlapping window, and substantial non-Gaussianity in empirical log returns. We implement a classical baseline that fits a bivariate Gaussian model to aligned per-day log returns and estimates VaR/CVaR at the 95% and 99% confidence levels via Monte Carlo simulation. In parallel, we demonstrate a quantum-inspired workflow in which the fitted Gaussian is discretized on a small grid, the resulting probability mass function is loaded into a quantum state, and tail probabilities and tail expectations are estimated with a QFT-free, iterative amplitude-estimation routine (maximum-likelihood estimation over Grover powers) executed on the PennyLane state-vector simulator. On the chosen discretization (three qubits per asset, 8×8 grid) and shot budget (2000 shots per circuit), the quantum estimates are close to the Monte Carlo baselines at 95% confidence and show larger deviations at 99%, consistent with discretization and finite-shot effects.

1 Introduction

Measuring and managing the probability and severity of extreme portfolio losses is a central task in quantitative risk management. Two commonly used tail-focused summaries are Value at Risk (VaR) and Conditional Value at Risk (CVaR, also known as Expected Shortfall), which quantify, respectively, a high quantile of the loss distribution and the average loss beyond that quantile [20, 8, 9]. Classical estimation pipelines frequently rely on Monte Carlo simulation under parametric or nonparametric return models [10, 29]. Quantum amplitude estimation (QAE) is a canonical algorithmic primitive that, under coherent oracle access assumptions, can provide a quadratic improvement in the precision scaling of Monte Carlo-style expectation estimation [2, 3]. QAE and its iterative or maximum-likelihood variants have been proposed for financial applications, including risk analysis [6, 4, 5].

This report combines exploratory data analysis, algorithmic considerations, and an end-to-end implementation on a quantum simulator for the problem of estimating VaR and CVaR of an equal-weight AAPL/MSFT portfolio. Because the empirical return distributions display heavy tails and volatility clustering [19, 32, 30], we emphasize that Gaussian modeling is a simplifying assumption used here to provide a well-defined baseline and to match the discretized quantum demonstration.

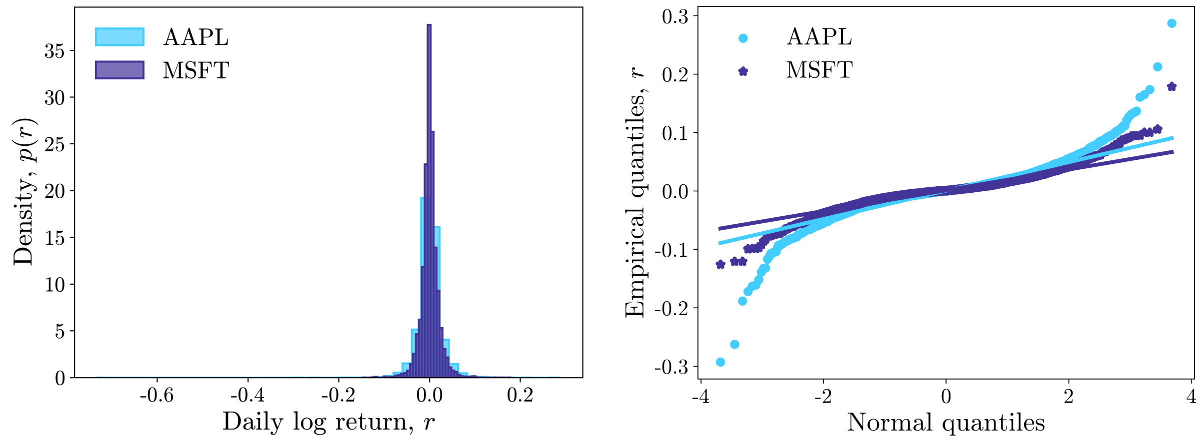

(a) Return histograms. (b) Gaussian QQ-plots.

Figure 1: Diagnostics for per-day log returns. Large deviations from normality motivate caution when using Gaussian baselines for tail-risk estimation.

2 Dataset and exploratory analysis

2.1 Raw data structure and missingness

The dataset contains 9909 observations and four floating-point columns: time in days and log price for each asset. AAPL is observed over the full horizon, while MSFT contains 1325 missing rows (13.37%), concentrated in one contiguous block at the end of the raw table. In calendar time, MSFT starts later than AAPL; therefore, portfolio-level modeling should focus on the overlapping window, and missingness handling is primarily an alignment decision rather than sporadic imputation.

2.2 Time grid and return definition

Both time axes are strictly increasing and show typical step sizes of one day, with frequent multi-day gaps consistent with weekends and holidays. Because Δt is not always exactly one, we define per-day log returns via

$$r_t = \frac{\Delta \log P_t}{\Delta t} = \frac{\log P_t - \log P_{t-\Delta t}}{\Delta t}, \tag{1}$$

which normalizes return increments by the time step and reduces sensitivity to irregular sampling.

2.3 Alignment, dependence, and non-Gaussianity

Per-day returns for AAPL and MSFT are computed on their native time grids and then inner-joined on rounded time stamps (three decimals), yielding 8583 aligned paired returns. The resulting empirical correlation is approximately $\rho \approx 0.389$, indicating moderate positive dependence. Both series exhibit significant deviations from Gaussianity, with AAPL in particular showing strong negative skewness and extremely high excess kurtosis, consistent with stylized facts of equity returns [19, 32]. Figure 1 illustrates the heavy-tailed behavior using histograms and Gaussian QQ-plots.

3 Problem setup and risk measures

Let $r \in \mathbb{R}^2$ denote the vector of per-day asset returns and let $w = (0.5, 0.5)$ be the equal-weight portfolio. The portfolio return is $R_p = w^\top r$ and we define the loss as $L = -R_p$. For a confidence level $\alpha \in (0, 1)$, VaR and CVaR (Expected Shortfall) are defined as

$$\mathrm{VaR}_\alpha(L) = \inf\{\ell \in \mathbb{R} : \mathbb{P}(L \leq \ell) \geq \alpha\}, \tag{2}$$

$$\mathrm{CVaR}_\alpha(L) = \mathbb{E}[L \mid L \geq \mathrm{VaR}_\alpha(L)], \tag{3}$$

with CVaR providing a coherent alternative to VaR in many settings [9, 8].

4 Methods

4.1 Classical baseline: Gaussian model and Monte Carlo

On the aligned return pairs, we fit a bivariate Gaussian model with mean vector μ and covariance matrix Σ. Losses are then sampled by drawing $r \sim \mathcal{N}(\mu, \Sigma)$ and computing $L = -w^\top r$. VaR is estimated by empirical quantiles and CVaR by averaging losses exceeding VaR. This approach is a standard Monte Carlo baseline [10] and provides a reference for later comparisons.

4.2 Quantum path: discretization, state preparation, and amplitude estimation

Discretized distribution. To construct a finite-dimensional quantum state, the fitted Gaussian distribution is discretized on an 8×8 grid spanning $\mu_i \pm 5\sigma_i$ for each asset dimension, where σ_i are marginal standard deviations. The joint probability mass function is approximated by evaluating the Gaussian density on the grid and multiplying by the cell area, followed by normalization.

State preparation. Let p_x denote the discretized probabilities over $N = 64$ grid points. We prepare

$$|\psi\rangle = \sum_{x=0}^{N-1} \sqrt{p_x}\, |x\rangle \tag{4}$$

using amplitude embedding on $n = 6$ data qubits (three per asset). In general, distribution loading is a key cost driver in quantum Monte Carlo workflows [11, 43].

Payoff encoding for tail probability and tail expectation. For a loss threshold ℓ, define the tail indicator $g_\ell(x) = \mathbf{1}[L(x) \geq \ell]$. Tail probabilities take the form $a(\ell) = \mathbb{E}[g_\ell(X)] = \mathbb{P}(L \geq \ell)$. To estimate CVaR we also require a tail-loss numerator, which can be written as an expectation $\mathbb{E}[L(X)\, \mathbf{1}[L(X) \geq \ell]]$. In the implementation, such expectations are encoded by applying, conditional on x, an ancilla rotation that maps $|0\rangle \mapsto \sqrt{1 - g(x)}|0\rangle + \sqrt{g(x)}|1\rangle$; measuring the ancilla yields a Bernoulli random variable with success probability $\mathbb{E}[g(X)]$.

Grover operator and QFT-free amplitude estimation. Given a state-preparation and payoff operator A such that

$$A|0\rangle^{\otimes(n+1)} = \sqrt{1-a}|\Phi_0\rangle|0\rangle + \sqrt{a}|\Phi_1\rangle|1\rangle, \tag{5}$$

QAE estimates a using Grover-style reflections and repeated applications of a Grover iterate Q [1, 2]. Instead of phase estimation, we use a QFT-free strategy in the spirit of maximum-likelihood amplitude estimation [4], in which circuits with $m \in \{0, 1, 2, 4\}$ Grover powers are

Table 1: Key configuration parameters for the baseline and simulated QAE experiments.

Quantity	Value
Aligned paired returns	8583
Portfolio weights (w_1, w_2)	(0.5, 0.5)
Confidence levels α	0.95, 0.99
Monte Carlo samples	200,000
Discretization grid	8×8
Qubits per asset	3
Shots per circuit	2000
Grover powers m	0, 1, 2, 4
Bounds for discretization	$\mu \pm 5\sigma$

sampled and the amplitude parameterization $a = \sin^2\theta$ is fit by maximizing the likelihood of observing ancilla outcomes with probabilities $\sin^2((2m+1)\theta)$. This procedure is related to iterative amplitude estimation approaches that avoid deep QFT circuits [5, 13].

Recovering VaR and CVaR. VaR at confidence α is obtained by searching for a threshold ℓ whose estimated tail probability matches $1-\alpha$. Once a VaR estimate is obtained, CVaR is computed as

$$\mathrm{CVaR}_\alpha \approx \frac{\mathbb{E}[L\,\mathbf{1}(L \geq \mathrm{VaR}_\alpha)]}{\mathbb{P}(L \geq \mathrm{VaR}_\alpha)}, \tag{6}$$

where both numerator and denominator are estimated by amplitude estimation on appropriately encoded payoffs.

5 Experimental configuration

The classical and quantum routines were executed on a local simulator. The Monte Carlo baseline used 200,000 samples. The quantum simulation used three qubits per asset (six data qubits total), a single payoff ancilla, and 2000 shots per circuit; Grover powers were chosen as $m \in \{0, 1, 2, 4\}$. All quantum circuits were executed on the PennyLane `default.qubit` simulator [17].

6 Results

6.1 Baseline Monte Carlo estimates

Under the fitted bivariate Gaussian model, the Monte Carlo baseline yields the tail-risk estimates reported in Table 2. These values should be interpreted as risk under a Gaussian return model; the EDA indicates heavier empirical tails, implying potential model risk [19, 20].

6.2 Simulated QAE estimates and comparison

Table 2 also reports the simulated QAE estimates. At 95% confidence, the relative errors versus the Monte Carlo baseline are about 2.3% for VaR and 0.7% for CVaR. At 99% confidence the deviations are larger (about 10.1% for VaR and 5.8% for CVaR), reflecting the combined effects of coarse discretization and finite-shot amplitude estimation in the deeper tail.

Table 2: Tail-risk estimates for the equal-weight AAPL/MSFT portfolio loss. "Quantum" refers to the discretized Gaussian model evaluated with a QAE-like amplitude-estimation routine on a simulator. Relative error is computed with respect to the Monte Carlo baseline.

Metric	Baseline MC	Quantum (sim.)	Abs. diff.	Rel. error
$\mathrm{VaR}_{0.95}$	0.030012	0.029328	-0.000683	2.28%
$\mathrm{CVaR}_{0.95}$	0.037797	0.037549	-0.000248	0.66%
$\mathrm{VaR}_{0.99}$	0.042770	0.047072	0.004302	10.06%
$\mathrm{CVaR}_{0.99}$	0.049156	0.046311	-0.002845	5.79%

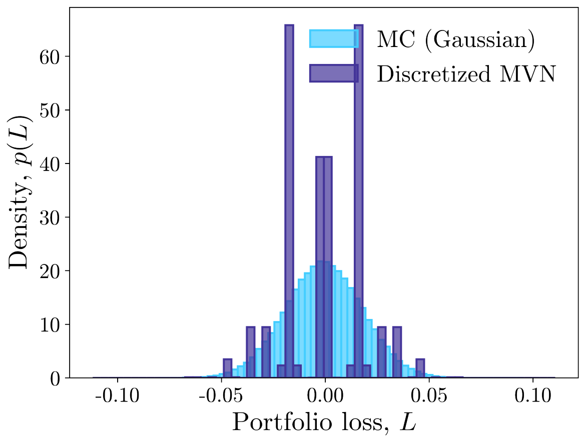

(a) Loss distribution.

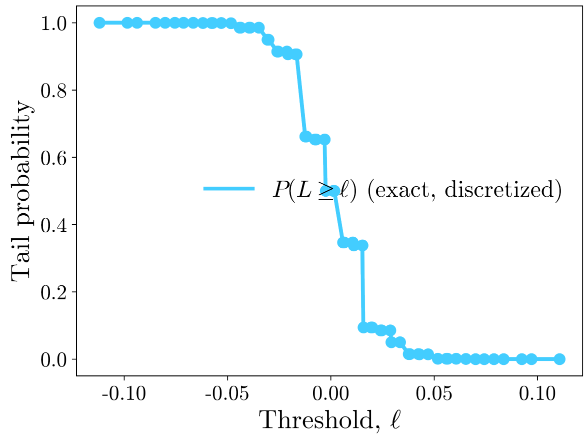

(b) Exact tail probability curve on the discretized grid.

Figure 2: Diagnostics for the fitted Gaussian model and its discretized approximation used by the amplitude-estimation circuits.

6.3 Discretization effects

The discretized model induces approximation error relative to the continuous fitted Gaussian. In particular, quantile definitions on a finite grid can be sensitive to how thresholds are chosen: the discrete VaR obtained from the cumulative distribution function may differ from the threshold that achieves a target tail probability when the threshold is allowed to lie between grid loss values. This phenomenon is most pronounced at 95% confidence for the coarse 8×8 grid, and it highlights the importance of resolution selection and careful quantile conventions when mapping classical risk measures to finite-dimensional quantum oracles.

7 Complexity and scaling discussion

Classical Monte Carlo estimation of tail probabilities and expectations converges with root-mean-square error $\mathcal{O}(1/\sqrt{N})$, implying that achieving accuracy ε typically requires $N = \Theta(1/\varepsilon^2)$ samples [10, 34]. In contrast, QAE achieves an oracle-query complexity scaling of $\mathcal{O}(1/\varepsilon)$ in idealized settings [2, 3], and iterative/QFT-free variants reduce circuit depth requirements [5, 4]. Figure 3 summarizes the asymptotic scaling behavior.

In the present experiment, the quantum workflow is slower in wall-clock time than the classical Monte Carlo baseline because state-vector simulation and explicit payoff-unitary construction dominate runtime. These costs do not negate the theoretical scaling of QAE, but they emphasize that practical advantage depends on efficient state preparation, compact oracles, and sufficiently low noise, which are challenging in the NISQ regime [16, 15].

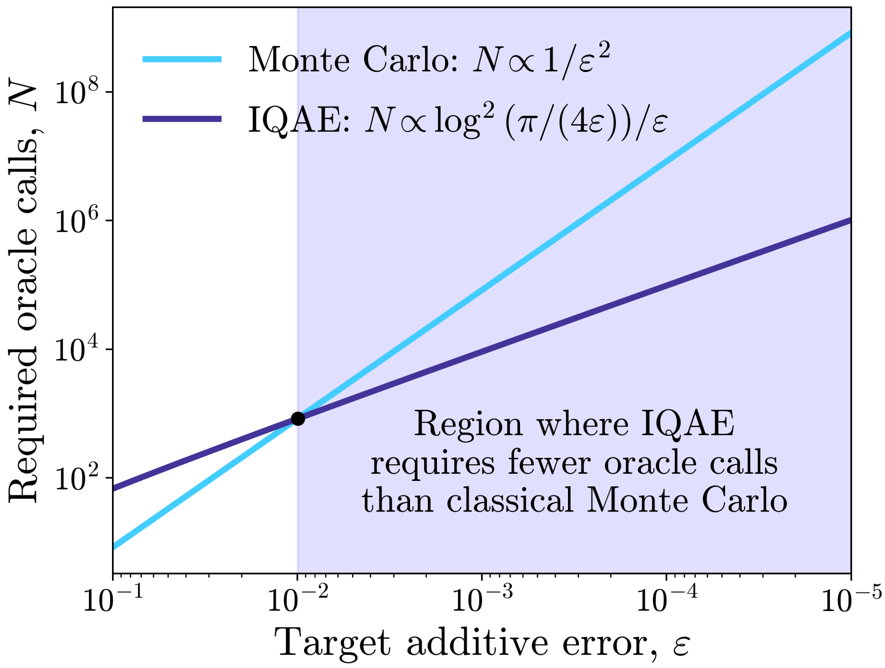

Figure 3: Reference scaling comparison between classical Monte Carlo and QAE-style estimators. The quantum advantage refers to the number of coherent oracle calls as a function of target precision, not end-to-end wall-clock time in a simulator.

8 Limitations and future outlook

The reported values are conditional on a Gaussian model for joint returns and on an aggressive discretization to keep the quantum simulation tractable. Empirically, the data show heavy tails and volatility clustering, suggesting that more realistic risk estimation may require heavier-tailed marginals, time-varying volatility models, or copula-based dependence structures [19, 30, 20]. From a quantum-algorithmic perspective, improving accuracy would require increasing discretization resolution (more qubits) and allocating more shots or adopting tighter amplitude-estimation schedules, which in turn increases circuit depth and sensitivity to noise. Near-term hardware demonstrations may therefore rely on error-mitigation techniques [46, 47, 48], while larger-scale advantage is widely expected to require fault-tolerant quantum computing and careful resource analysis [37, 49, 50].

9 Conclusion

We presented an end-to-end pipeline for estimating VaR and CVaR of a two-asset portfolio from log-price time series, combining classical Monte Carlo baselines with a simulated QAE workflow for tail probability and tail expectation estimation. While the quantum routine does not provide practical runtime benefits on a simulator and is sensitive to discretization choices, it illustrates how risk measures can be expressed as expectation values amenable to amplitude-estimation subroutines and provides a reproducible starting point for future work on higher-resolution discretizations, improved return models, and eventual hardware execution.

References

[1] Lov K. Grover. *A fast quantum mechanical algorithm for database search*. 1996. doi: 10.1145/237814.237866.

[2] G. Brassard, P. Høyer, M. Mosca, others. *Quantum Amplitude Amplification and Estimation*. 2000. doi: 10.1090/conm/305/05215.

[3] A. Montanaro. *Quantum speedup of Monte Carlo methods*. 2015. doi: 10.1098/rspa.2015.0301.

[4] Yohichi Suzuki, Shumpei Uno, Rudy Raymond, others. *Amplitude estimation without phase estimation*. 2019. doi: 10.1007/s11128-019-2565-2.

[5] Dmitry Grinko, Julien Gacon, Christa Zoufal, others. *Iterative quantum amplitude estimation*. 2019. doi: 10.1038/s41534-021-00379-1.

[6] Stefan Woerner, D. Egger. *Quantum risk analysis*. 2018. doi: 10.1038/s41534-019-0130-6.

[7] N. Stamatopoulos, D. Egger, Yue Sun, others. *Option Pricing using Quantum Computers*. 2019. doi: 10.22331/q-2020-07-06-291.

[8] R. Rockafellar, S. Uryasev. *Optimization of conditional value-at risk*. 2000. doi: 10.21314/jor.2000.038.

[9] Philippe Artzner, F. Delbaen, J. Eber, others. *Coherent Measures of Risk*. 1999. doi: 10.1111/1467-9965.00068.

[10] P. Glasserman. *Monte Carlo Methods in Financial Engineering*. 2003. doi: 10.1007/978-0-387-21617-1.

[11] Lov K. Grover, T. Rudolph. *Creating superpositions that correspond to efficiently integrable probability distributions*. 2002. url: https://www.semanticscholar.org/paper/07f4621702954efa955f8877f338c3dfe58c7bae.

[12] S. Aaronson, Patrick Rall. *Quantum Approximate Counting, Simplified*. 2019. doi: 10.1137/1.9781611976014.5.

[13] Shion Fukuzawa, C. Ho, S. Irani, others. *Modified Iterative Quantum Amplitude Estimation is Asymptotically Optimal*. 2022. doi: 10.1137/1.9781611977561.ch12.

[14] Tomoki Tanaka, Yohichi Suzuki, Shumpei Uno, others. *Amplitude estimation via maximum likelihood on noisy quantum computer*. 2020. doi: 10.1007/s11128-021-03215-9.

[15] P. Rao, Kwangmin Yu, Hyun-Kyung Lim, others. *Quantum amplitude estimation algorithms on IBM quantum devices*. 2020. doi: 10.1117/12.2568748.

[16] J. Preskill. *Quantum Computing in the NISQ era and beyond*. 2018. doi: 10.22331/q-2018-08-06-79.

[17] V. Bergholm, J. Izaac, M. Schuld, others. *PennyLane: Automatic differentiation of hybrid quantum-classical computations*. 2018. url: https://www.semanticscholar.org/paper/38d8230a7aeae6554497b253848ad5bf677e4fb3.

[18] D. Egger, Ricardo García Gutiérrez, Jordi Cahu'e Mestre, others. *Credit Risk Analysis Using Quantum Computers*. 2019. doi: 10.1109/tc.2020.3038063.

[19] R. Cont. *Empirical properties of asset returns: stylized facts and statistical issues*. 2001. doi: 10.1080/713665670.

[20] Alexander J. McNeil, Rdiger Frey, P. Embrechts. *Quantitative Risk Management: Concepts, Techniques, and Tools*. 2005. url: https://www.semanticscholar.org/paper/b3877f996625ef0467a0a2843970b119bfecc5b3.

[21] F. Black, Myron S. Scholes. *The Pricing of Options and Corporate Liabilities*. 1973. doi: 10.1086/260062.

[22] Michael A. Nielsen, Isaac L. Chuang. *Quantum Computation and Quantum Information (10th Anniversary edition)*. 2011. url: https://www.semanticscholar.org/paper/3bfb6c2d7a301ab094db12ceaa8812782aaac23b.

[23] R. Rockafellar, S. Uryasev. *Conditional Value-at-Risk for General Loss Distributions*. 2001. doi: 10.2139/ssrn.267256.

[24] C. Acerbi, Dirk Tasche. *Expected Shortfall: a natural coherent alternative to Value at Risk*. 2001. doi: 10.1111/1468-0300.00091.

[25] Peter F. Christoffersen. *Evaluating Interval Forecasts*. 1998. doi: 10.2307/2527341.

[26] Jose A. Lopez. *Methods for Evaluating Value-at-Risk Estimates*. 1998. doi: 10.2139/ssrn.1029673.

[27] Chris Marshall, Michael Siegel. *Value at Risk*. 1996. doi: 10.3905/jod.1997.407975.

[28] M. Phelan. *Probability and Statistics Applied to the Practice of Financial Risk Management: The Case of J.P. Morgan's RiskMetrics™*. 1997. doi: 10.1023/a:1007926803627.

[29] N. Metropolis, S. Ulam. *The Monte Carlo method.*. 1949. doi: 10.1080/01621459.1949.10483310.

[30] Tim Bollerslev. *Generalized autoregressive conditional heteroskedasticity*. 1986. doi: 10.1016/0304-4076(86)90063-1.

[31] Anil K. Bera, M. Higgins. *ARCH Models: Properties, Estimation and Testing*. 1993. doi: 10.1111/j.1467-6419.1993.tb00170.x.

[32] B. Mandelbrot. *The Variation of Certain Speculative Prices*. 1963. doi: 10.1007/978-1-4757-2763-0_14.

[33] S. Shapiro, M. Wilk. *An Analysis of Variance Test for Normality (Complete Samples)*. 1965. doi: 10.1093/biomet/52.3-4.591.

[34] W. Hoeffding. *Probability inequalities for sum of bounded random variables*. 1963. doi: 10.1007/978-1-4612-0865-5_26.

[35] R. Orús, Samuel Mugel, Enrique Lizaso. *Quantum computing for finance: Overview and prospects*. 2018. doi: 10.1016/j.revip.2019.100028.

[36] Samuel Mugel, Carlos Kuchkovsky, E. Sanchez, others. *Dynamic Portfolio Optimization with Real Datasets Using Quantum Processors and Quantum-Inspired Tensor Networks*. 2020. doi: 10.1103/physrevresearch.4.013006.

[37] Shouvanik Chakrabarti, R. Krishnakumar, G. Mazzola, others. *A Threshold for Quantum Advantage in Derivative Pricing*. 2020. doi: 10.22331/q-2021-06-01-463.

[38] G. Low, I. Chuang. *Optimal Hamiltonian Simulation by Quantum Signal Processing.*. 2016. doi: 10.1103/physrevlett.118.010501.

[39] András Gilyén, Yuan Su, G. Low, others. *Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics*. 2018. doi: 10.1145/3313276.3316366.

[40] N. Stamatopoulos, W. Zeng. *Derivative Pricing using Quantum Signal Processing*. 2023. doi: 10.22331/q-2024-04-30-1322.

[41] N. Stamatopoulos, B. D. Clader, Stefan Woerner, others. *Quantum Risk Analysis of Financial Derivatives*. 2024. url: https://www.semanticscholar.org/paper/ed75faca9c0ff4d6a5d48ff91546920235d6c6fc.

[42] N. Stamatopoulos, G. Mazzola, Stefan Woerner, others. *Towards Quantum Advantage in Financial Market Risk using Quantum Gradient Algorithms*. 2021. doi: 10.22331/q-2022-07-20-770.

[43] V. Giovannetti, S. Lloyd, L. Maccone. *Quantum random access memory.*. 2007. doi: 10.1103/physrevlett.100.160501.

[44] V. Bergholm, J. Vartiainen, M. Mottonen, others. *Quantum circuits with uniformly controlled one-qubit gates (7 pages)*. 2004. doi: 10.1103/physreva.71.052330.

[45] V. Shende, S. S. Bullock, Igor L. Markov. *Synthesis of quantum logic circuits*. 2005. doi: 10.1145/1120725.1120847.

[46] K. Temme, S. Bravyi, J. Gambetta. *Error Mitigation for Short-Depth Quantum Circuits.*. 2016. doi: 10.1103/physrevlett.119.180509.

[47] Suguru Endo, S. Benjamin, Ying Li. *Practical Quantum Error Mitigation for Near-Future Applications*. 2017. doi: 10.1103/physrevx.8.031027.

[48] A. Kandala, K. Temme, Antonio Mezzacapo, others. *Experimental implementation of error mitigation for short-depth quantum circuits*. 2018. url: https://www.semanticscholar.org/paper/416e6ed024eecb66ce29650def82555094a498a5.

[49] IBM Quantum. *IBM Quantum development roadmap*. 2025. url: https://www.ibm.com/quantum/roadmap.

[50] IEEE Spectrum. *PsiQuantum's plans for a fault-tolerant quantum computer*. 2025. url: https://spectrum.ieee.org/psiquantum-supercomputer.