Quantum computing for multidimensional option pricing: End-to-end pipeline

Julien Hok (Investec Bank)   Álvaro Leitao (CITIC, Univ. A Coruña)

BofA Quant Speaker Series Seminar, London, July 2026

Disclaimer

“The opinions expressed in this presentation and on the following slides are solely those of the presenter and not necessary those of INVESTEC. INVESTEC does not guarantee the accuracy or reliability of the information provided herein.”

Outline

Why quantum computing for multi-asset options?

- **Classical challenges:**
  - High-dimensional integration $\rightarrow$ Monte Carlo (CMC) converges as $O(1/\varepsilon^2)$ (slow).
  - Realistic marginals require arbitrage-free densities (skew, fat tails).
  - Copula construction adds complexity.
- **Quantum opportunity:**
  - Quantum Amplitude Estimation (QAE) achieves $O(1/\varepsilon)$ convergence (quadratic speedup).
  - End-to-end pipeline: market data $\rightarrow$ NIG marginals $\rightarrow$ copula $\rightarrow$ QAE pricing.
- **This work:**
  - Calibrate NIG marginals to real equity options (Crédit Agricole, AXA, Michelin).
  - Couple with Gaussian copula.
  - Apply QAE to both density reconstruction and multi-asset option pricing.

Normal Inverse Gaussian (NIG) model

- Under $\mathbb{Q}$, asset dynamics:

$$\frac{\mathrm{d}S(t)}{S(t)} = (r-q)\mathrm{d}t + \mathrm{d}X(t) \tag{1}$$

where $X(t) \sim \mathrm{NIG}(\alpha, \beta, \delta t, \mu t)$ has a density function (location μ fixed to 0 for simplicity) given by

$$f_{\mathrm{NIG}}(x; \alpha, \beta, \delta t, 0) = \frac{\alpha(\delta t)}{\pi} e^{\delta t\sqrt{\alpha^2-\beta^2}+\beta x} \frac{K_1\left(\alpha\sqrt{(\delta t)^2 + x^2}\right)}{\sqrt{(\delta t)^2 + x^2}} \tag{2}$$

The solution is

$$S(T) = S(t)\exp\left((r - q + \omega)\tau + X(\tau)\right), \quad \tau := T - t, \tag{3}$$

with the martingale compensator

$$\omega = \delta\left(\sqrt{\alpha^2 - (\beta+1)^2} - \sqrt{\alpha^2 - \beta^2}\right). \tag{4}$$

- **Properties:**
  - δ controls the volatility level, β the skewness and α the tail heaviness (larger $\alpha \to$ lighter tails).
  - Arbitrage-free per time slice.

Calibration of NIG marginals

Regularized least-squares

$$\mathcal{I}(\theta) = \sum_{m=1}^{M} w_m \left(V_m^{\mathrm{NIG}}(\theta) - \bar{V}_m\right)^2 + \lambda\|\theta - \theta_0\|^2, \quad \theta = (\alpha, \beta, \delta) \tag{5}$$

Theoretical guarantees (Propositions 2.1, 2.3)

- Option prices independent of μ (set $\mu = 0$).
- Continuity of $V_m^{\mathrm{NIG}}(\theta)$ in θ.
- Existence of minimizer on compact admissible set.
- Tikhonov regularization ensures stability.

Other approaches for volatility surface fitting exist (SVI, SABR, local volatility), but we adopt the time-slice NIG method for its analytic tractability and arbitrage-free marginals with intuitive parameters interpretation.

Calibration results: Crédit Agricole (1Y expiry)

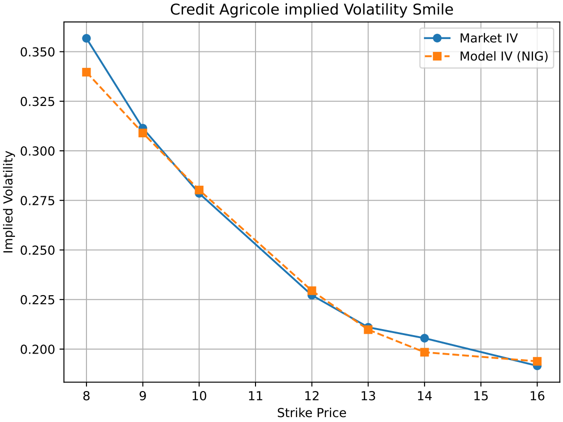

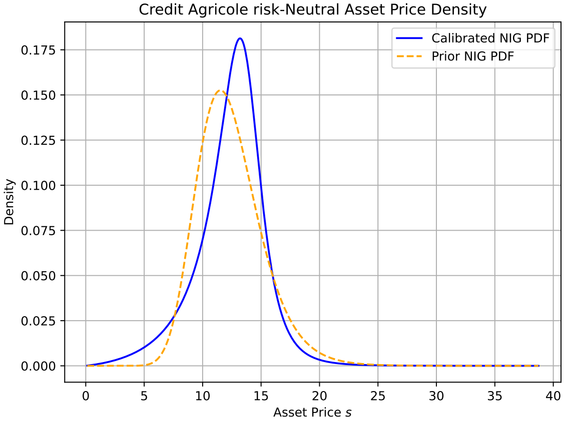

Spot $S_0 = 12.91$€ (as of 24/12/2024), $\quad \bar{\alpha} = 4.69,\ \bar{\beta} = -3.06,\ \bar{\delta} = 0.18, \quad \lambda = 5 \times 10^{-7}$

Max pricing error ≈ 21 bps (tail).

Calibration results: AXA (1Y expiry)

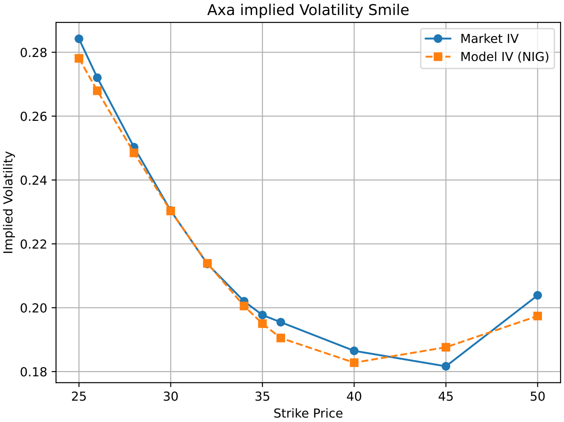

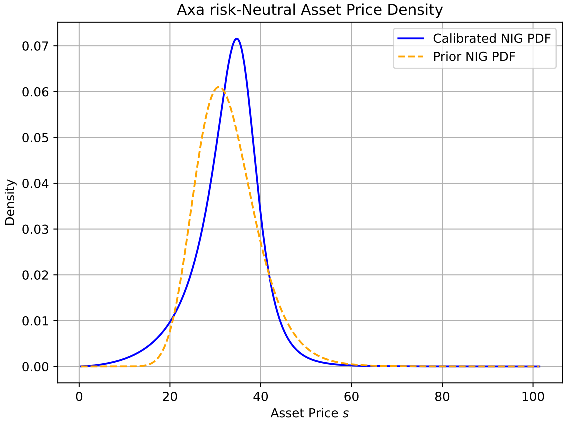

Spot $S_0 = 33.80$€ (as of 24/12/2024), $\bar{\alpha} = 5.24,\ \bar{\beta} = -3.26,\ \bar{\delta} = 0.18, \quad \lambda = 5 \times 10^{-7}$
Max pricing error ≈ 10 bps.

Calibration results: Michelin (1Y expiry)

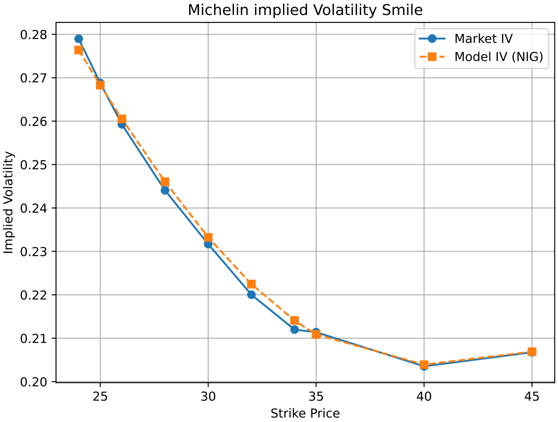

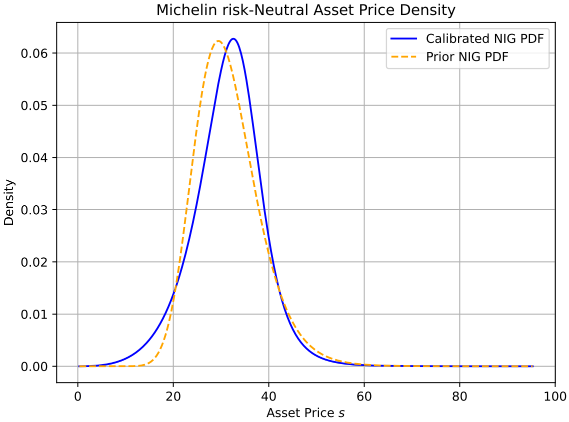

Spot $S_0 = 31.76$€ (as of 24/12/2024), $\bar{\alpha} = 6.20,\ \bar{\beta} = -3.31,\ \bar{\delta} = 0.26, \quad \lambda = 5 \times 10^{-7}$
Max pricing error ≈ 6 bps.

From marginals to joint density (Gaussian copula)

- Sklar's theorem:

$$F(x_1,\ldots,x_N) = \mathcal{C}_\Sigma\big(F_1(x_1),\ldots,F_N(x_N)\big) \tag{6}$$

- Gaussian copula:

$$\mathcal{C}_\Sigma(u) = \Phi_N\big(\Phi^{-1}(u_1),\ldots,\Phi^{-1}(u_N);\Sigma\big) \tag{7}$$

- Joint density:

$$f(\mathbf{x}) = c_\Sigma\big(F_1(x_1),\ldots,F_N(x_N)\big)\prod_{i=1}^{N} f_i(x_i) \tag{8}$$

- Enabling equivalent pricing reformulations:

$$\begin{aligned} V(T,K) &= e^{-rT}\mathbb{E}[h(\mathbf{S}] \\ &= e^{-rT}\mathbb{E}^{\mathrm{ind}}[h(\mathbf{S})c\,(F_1(S_1),\ldots,F_N(S_N))], \end{aligned} \tag{9}$$

where $\mathbb{E}^{\mathrm{ind}}$ uses product density $\prod f_i$ (independent formulation).

Quantum Amplitude Estimation (QAE)

Problem: estimate $\mu = \mathbb{E}[\phi(\mathbf{X})] = \int f(\mathbf{x})\phi(\mathbf{x})\mathrm{d}\mathbf{x}$

- Classical Monte Carlo: L samples $\Rightarrow \varepsilon \sim 1/\sqrt{L} \Rightarrow O(1/\varepsilon^2)$.
- QAE (Brassard et al. 2002): queries oracle $\mathcal{U}$ that prepares the superposition state

$$|\psi\rangle = \mathcal{U}|0\rangle = \sqrt{\mu}\,|\psi_1\rangle|1\rangle + \sqrt{1-\mu}\,|\psi_0\rangle|0\rangle \tag{10}$$

- Amplitude estimation returns $\tilde{\mu}$ with $|\tilde{\mu} - \mu| \leq \varepsilon$ using $O(1/\varepsilon)$ queries (quadratic improvement).

Requirement: encode the integrand and the density into a quantum state (state preparation).

State preparation: Need to define an oracle $\mathcal{U}_{f,\phi}$ to enables to construct the following quantum state:

$$\mathcal{U}_{f,\phi}|0\rangle^{\otimes Nn}|0\rangle = \sum_{j=0}^{2^{Nn}-1} \sqrt{f(\mathbf{x}_j)}\,|\mathbf{x}_j\rangle \left(\sqrt{\phi(\mathbf{x}_j)}|1\rangle + \sqrt{1-\phi(\mathbf{x}_j)}|0\rangle\right) \tag{11}$$

where f is the density, ϕ the integrand function (scaled to $[0,1]$), and $|\mathbf{x}_j\rangle$ encodes discretized integration points. The probability of measuring $|1\rangle$ equals the expected value μ.

Key assumption: Querying a quantum oracle one time is equivalent to draw a single sample.

Goal: Approximate a density f (supported on $[a, b]$) by an orthonormal series expansion.

Cosine basis $\{\gamma_k\}_{k\geq 0}$ on $[a, b]$:

$$\gamma_k(x) = \begin{cases} \frac{1}{\sqrt{b-a}}, & k = 0, \\ \sqrt{\frac{2}{b-a}} \cos\left(\frac{k\pi(x-a)}{b-a}\right), & k \geq 1. \end{cases} \tag{12}$$

Orthonormality:

$$\int_a^b \gamma_k(x)\gamma_l(x)\,dx = \delta_{kl}. \tag{13}$$

Projection (Fourier) coefficients:

$$a_k = \int_a^b f(x)\gamma_k(x)\,dx = \mathbb{E}[\gamma_k(X)], \quad X \sim f. \tag{14}$$

Truncated series:

$$f(x) \approx f^{\mathcal{K}}(x) = \sum_{k=0}^{\mathcal{K}-1} a_k\,\gamma_k(x). \tag{15}$$

Why cosine? – Exponential convergence (Theorem 3.2)

If f is analytic on $[a, b]$ and its derivatives vanish (or are negligible) at the endpoints, then the cosine series converges exponentially fast:

$$\sup_{x \in [a,b]} |f(x) - f^{\mathcal{K}}(x)| \leq \zeta\, e^{-\nu \mathcal{K}}, \tag{16}$$

where $\zeta, \nu > 0$ are constants independent of $\mathcal{K}$.

Key consequence:

$$\mathcal{K} = O\left(\log \frac{1}{\varepsilon}\right) \quad \text{suffices for accuracy } \varepsilon.$$

This is ideal for quantum computing: Only a logarithmic number of coefficients are needed – even before quantum speedup.

Covenient representation for quantum loading: Cosine basis functions can be natively loaded in a quantum system.

Cosine Series for Marginal Densities (3/3)

Quantum estimation of coefficients $a_k = \mathbb{E}[\gamma_k(X)]$: For each k, use QAE to estimate $\mathbb{E}[\gamma_k(X)]$ with $O(1/\varepsilon)$ queries.

Complexity bound (Theorem 3.4): With probability $1-\rho$, achieving error ε in the CDF F_i requires

$$O\left(\frac{\sqrt{b_i - a_i}\,\mathcal{K}_i^2}{\varepsilon}\,\log\frac{\mathcal{K}_i}{\rho}\right) \tag{17}$$

queries to the oracles $\mathcal{U}_{a_k}$.

Why this is a quantum advantage: - Classical Monte Carlo for density estimation would need $O(1/\varepsilon^2)$ samples. - Here, $\mathcal{K}_i = O(\log(1/\varepsilon))$ (exponential convergence), so total queries are nearly $O(1/\varepsilon)$ – quadratic speedup.

Takeaway: High-accuracy marginal CDFs $\hat{F}_i$ are obtained efficiently with QAE.

Joint formulation – direct density loading

- Prepare a quantum state encoding the joint risk-neutral density $\hat{f}(\mathbf{x})$ (copula):

$$\mathcal{A}_{\mathbf{X}}|0\rangle^{\otimes Nn} = \sum_{j=0}^{2^{Nn}-1} \sqrt{\hat{f}(\mathbf{x}_j)}\,|\mathbf{x}_j\rangle \tag{18}$$

- Apply a controlled rotation based on the payoff h (scaled to $[0,1]$):

$$W_h : |\mathbf{x}_j\rangle|0\rangle \;\mapsto\; |\mathbf{x}_j\rangle\left(\sqrt{h(\mathbf{x}_j)}|1\rangle + \sqrt{1-h(\mathbf{x}_j)}|0\rangle\right) \tag{19}$$

- Full oracle: $\mathcal{U}_{f,h} = W_h \cdot (\mathcal{A}_{\mathbf{X}} \otimes \mathcal{I})$

Resulting state:

$$\mathcal{U}_{f,h}|0\rangle^{\otimes Nn}|0\rangle = \sum_j \sqrt{\hat{f}(\mathbf{x}_j)}\,|\mathbf{x}_j\rangle\left(\sqrt{h(\mathbf{x}_j)}|1\rangle + \sqrt{1-h(\mathbf{x}_j)}|0\rangle\right) \tag{20}$$

The probability of measuring $|1\rangle$ equals $\sum_j \hat{f}(\mathbf{x}_j)h(\mathbf{x}_j) \approx \mathbb{E}[h(\mathbf{X})]$.

Independent formulation – copula absorbed into payoff

- Use the product density $\prod_{i=1}^{N} f_i(x_i)$ (independent marginals) – easier to prepare:

$$\mathcal{A}^{\mathsf{ind}}_{\mathbf{X}} |0\rangle^{\otimes Nn} = \sum_j \sqrt{\prod_i f_i(x_{i,j})}\, |\mathbf{x}_j\rangle \tag{21}$$

- Define an **adjusted payoff** that includes the copula density:

$$\hat{H}(\mathbf{x}) = \frac{1}{c_{\max}}\, h(\mathbf{x})\, c_{\Sigma}\big(\hat{F}_1(x_1), \ldots, \hat{F}_N(x_N)\big) \tag{22}$$

where $c_{\max} = \max_{u \in [0,1]^N} c(u)$ guarantees $\hat{H} \in [0,1]$.

Why this works:

$$\mathbb{E}^{\mathsf{ind}}[\hat{H}(\mathbf{X})] = \int \Big(\prod_i f_i(x_i)\Big) \frac{h(\mathbf{x})\, c_{\Sigma}(F_1, \ldots, F_N)}{c_{\max}}\, d\mathbf{x} = \frac{1}{c_{\max}}\, \mathbb{E}[h(\mathbf{X})] \tag{23}$$

Thus the true price is recovered up to the known factor $c_{\max}$.

Quantum complexity (Theorem 3.5) – for accuracy ε_c, confidence $1 - \rho_c$:

1. Cosine basis oracles W_{γ_k} (marginal estimation)

$$O\left(\frac{N^2\, c'_{\max}\, \sqrt{I_{\max}}\, \mathcal{K}_{\max}}{\varepsilon_c} \log \frac{N\mathcal{K}_{\max}}{\rho_c}\right) \tag{24}$$

2. Marginal density oracles $\mathcal{A}_{X_i}$ (used twice)

$$O\left(\underbrace{\frac{N^2 c'_{\max} I_{\max} \mathcal{K}^2_{\max}}{\varepsilon_c} \log \frac{N\mathcal{K}_{\max}}{\rho_c}}_{\text{marginal estimation}} + \underbrace{\frac{N c_{\max}}{\varepsilon_c} \log \frac{1}{\rho_c}}_{\text{pricing phase}}\right) \tag{25}$$

where $c'_{\max} = \|\nabla c\|_\infty$, $I_{\max} = \max_i (b_i - a_i)$, $\mathcal{K}_{\max} = O(\log(1/\varepsilon_c))$.

3. Payoff oracle $W_{\hat{H}}$ (final pricing)

$$O\left(\frac{c_{\max}}{\varepsilon_c}\log\frac{1}{\rho_c}\right) \tag{26}$$

where $c_{\max} = \max_{\mathbf{u}\in[0,1]^N} c(\mathbf{u})$.

Both formulations achieve $O(1/\varepsilon_c)$ scaling – quadratic speedup over CMC

- Joint formulation has, in principle, lower constant overhead
- Matches empirical results (joint outperforms independent)

Market data (Euronext, 24/12/2024):

- Credit Agricole ($S_0 = 12.91$€), AXA ($S_0 = 33.80$€), Michelin ($S_0 = 31.76$€)
- 1-year expiry European call/put options
- NIG marginals calibrated (max pricing errors: 21, 10, 6 bps)

Quantum setup:

- Simulator: myQLM + NEASQC QQuantLib (ideal quantum simulator)
- QAE variants: RealQAE (coefficients), IterativeQAE (option prices)
- Each experiment repeated $2^5 = 32$ times for statistics

Two sets of experiments:

1. Marginal density estimation

- Recover NIG density via cosine series
- Compare CMC vs QAMC for coefficients a_k
- $\mathcal{K} = 16$ coefficients, $n = 5$ qubits

2. Multi-asset option pricing

- Spread call (2 assets, $K = 0$)
- Basket call (3 assets, $K = 25$)
- Gaussian copula, $n = 2$–3 qubits/dim

Convergence for cosine coefficients (AXA NIG)

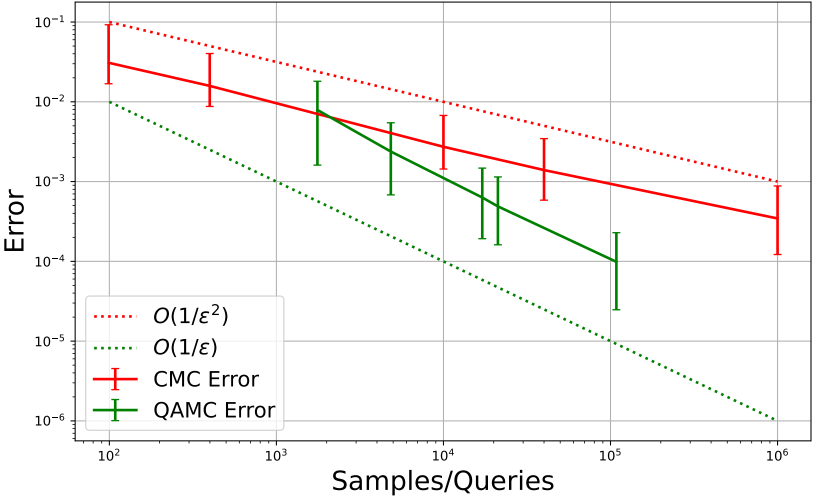

- Average error across all coefficients with 90% confidence intervals shown
- Demonstrates overall QAMC advantage

Density & CDF reconstruction (AXA NIG)

Recovered density (top) and CDF (bottom) for $\mathcal{K} = 2^3, 2^4, 2^5$ (~ 5000 *samples/queries).*

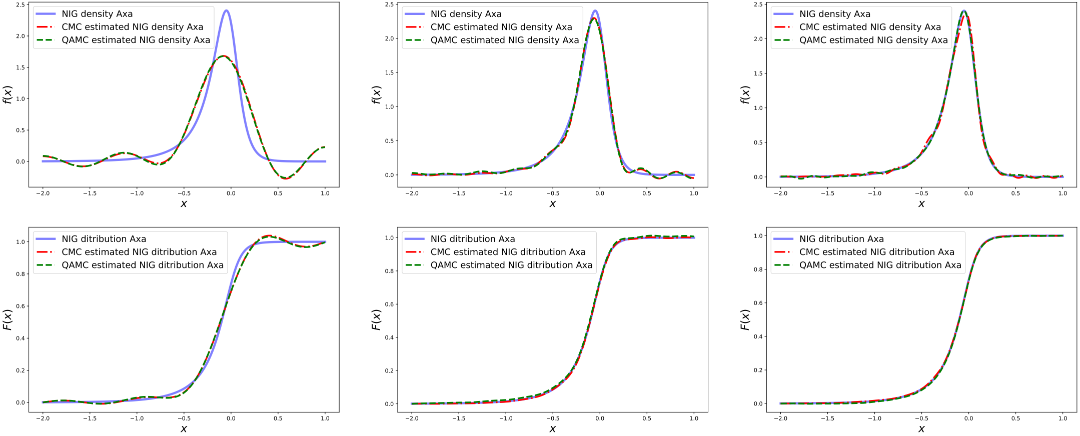

Settings:

- Payoff: $\max(S_1 - S_2 - K, 0)$, $K = 0$.
- Gaussian copula, $\rho = -0.25$.
- $n = 3$ qubits/dim (8 points).

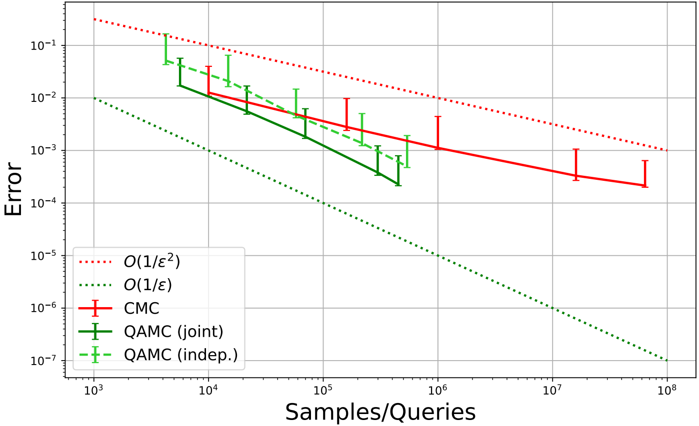

Observations:

- QAMC joint formulation converges as $O(1/\varepsilon)$; CMC as $O(1/\varepsilon^2)$.
- For $\varepsilon = 10^{-3}$, QAMC needs $\sim$10$\times$ fewer queries.
- Independent formulation also quadratic but higher constant.

Settings:

- Payoff: $\max\left(\frac{1}{3}\sum S_i - K, 0\right)$, $K = 25$.
- Gaussian copula correlation matrix:

$$\Sigma = \begin{pmatrix} 1 & -0.2 & -0.25 \\ -0.2 & 1 & -0.15 \\ -0.25 & -0.15 & 1 \end{pmatrix}.$$

- $n = 2$ qubits/dim (4 points).

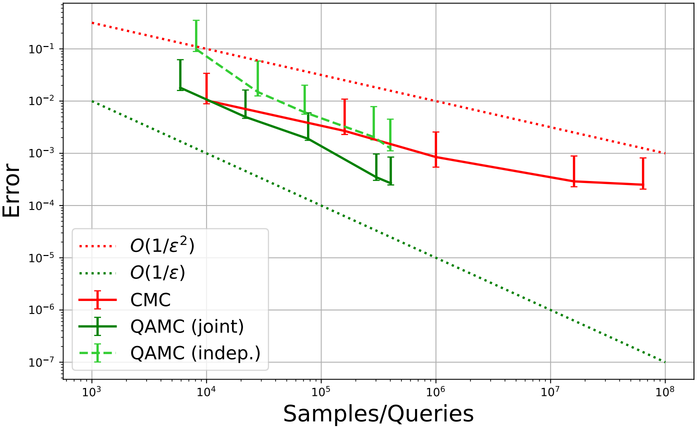

Conclusion: Quantum advantage persists in higher dimensions; joint formulation more efficient.

Summary

- **End-to-end pipeline** from market data to quantum-accelerated pricing:
  - NIG marginals calibrated to real equity options (arbitrage-free, fat tails, skew).
  - Gaussian copula for dependence.
  - QAE applied to both density reconstruction and option valuation.
- **Theoretical contributions:**
  - Existence and continuity of NIG calibration.
  - Query complexity bounds for quantum marginal estimation and multi-asset pricing.
- **Empirical findings:**
  - QAMC achieves quadratic speedup over CMC in practice.
  - 10–100× fewer queries for high accuracy ($\varepsilon \sim 10^{-3}$).
  - Joint copula formulation outperforms independent one.

Future work

- Richer dependence structures (t-copula, vine copulas).
- Path-dependent options (barriers, Asian) – requires more complex state preparation.
- Hardware implementation on current NISQ devices (error mitigation, short depth).
- Integration with classical volatility surface construction (SVI, local volatility).

Thank you! Questions?

Selected references

- Hok, J., Leitao, A. (2026). Quantum computing for multidimensional option pricing: End-to-end pipeline. *Submitted.*
- Brassard, G., Høyer, P., Mosca, M., Tapp, A. (2002). Quantum amplitude amplification and estimation. *Contemporary Mathematics.*
- Fang, F., Oosterlee, C.W. (2009). A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM J. Sci. Comput.*
- Cont, R., Tankov, P. (2004). *Financial Modelling with Jump Processes.* Chapman & Hall/CRC.
- Nielsen, M.A., Chuang, I.L. (2010). *Quantum Computation and Quantum Information.* Cambridge Univ. Press.