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

Julien Hok (Investec) Álvaro Leitao (CITIC, UDC)

Applied Quantum Methods in Computational Science and Engineering (AQMCSE) 2026
Delft (Netherlands)

Motivation & Proposal

- **Classical challenges:**
  - High-dimensional integration $\rightarrow$ Monte Carlo $\rightarrow$ slow convergence $O(1/\varepsilon^2)$.
  - Realistic marginals require arbitrage-free densities (skew, fat tails).
  - Copula construction adds complexity.
- **Quantum opportunity:**
  - Quantum Amplitude Estimation (QAE) achieves $O(1/\varepsilon)$, so quadratic speedup.
  - QAE applicability to both marginal reconstruction and option pricing.
- **This work:**
  - Quantum marginal density estimation.
  - End-to-end pipeline: Market data $\rightarrow$ Marginals $\rightarrow$ Copula $\rightarrow$ QAE pricing.
  - Theoretical results on the quantum complexity for multi-dimension option pricing.

Outline

1. Basics on option pricing
2. Market-consistent marginals: NIG model
3. Copula-based joint distribution
4. Quantum accelerated Monte Carlo
5. Experimental results
6. Conclusions & outlook

Basics on option pricing (I)

- An option is a financial derivative: financial contract whose value depends on the future performance of some underlying assets

- Option pricing consists of obtaining the fair price of the derivative at any previous date to maturity date.

- For this purpose, the uncertain future dynamics of the price of the underlying asset must be taken into account, which is usually modelled in terms of stochastic differential equations (SDEs).

- Lets denote by S_t and V_t the value of the underlying asset and the price of the derivative at time $t \in [0, T]$, where T is the maturity date.

Basics on option pricing (II)

- Assuming that the derivative can be exercised only at maturity and by using standard mathematical finance tools (martingale properties, Itô's lemma and the Feynmann-Kàc theorem), the price of the derivative at time t, is

$$V_t = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[V_T | \mathcal{F}_t] = e^{-r(T-t)} \int_{\Omega_S} f_S(x) V_T(x) dx, \tag{1}$$

  where $\mathbb{E}^{Q}$ denotes the expectation under the risk neutral measure $\mathbb{Q}$, r is the risk-free interest rate, and $\mathcal{F}_t$ denotes the σ-algebra containing the information until time t.

- The value V_T is defined via the so-called **payoff function**, i.e., $V_T = h(S_T)$.

- Then, the derivative price is the discounted value of the expected payoff, conditioned to the current information of the market.

Normal Inverse Gaussian (NIG) model

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

$$\frac{\mathrm{d}S(t)}{S(t)} = (r-q)\mathrm{d}t + \mathrm{d}X(t) \tag{2}$$

where $X(t) \sim \mathrm{NIG}(\alpha, \beta, \delta t, \mu t)$ has a density function (with $\mu = 0$) given by

$$f_{\mathrm{NIG}}(x; \alpha, \beta, \delta t, 0) = \frac{\alpha(\delta t)}{\pi} e^{\delta t\sqrt{\alpha^2-\beta^2}+\beta x} \frac{K_1\left(\alpha\sqrt{(\delta t)^2 + x^2}\right)}{\sqrt{(\delta t)^2 + x^2}} \tag{3}$$

The solution is

$$S(T) = S(t)\exp\left((r - q + \omega)\tau + X(\tau)\right), \quad \tau := T - t, \tag{4}$$

with the martingale compensator $\omega = \delta\left(\sqrt{\alpha^2 - (\beta+1)^2} - \sqrt{\alpha^2 - \beta^2}\right)$.

Parameter interpretation: δ controls the volatility level, β the skewness and α the tail heaviness (larger $\alpha \to$ lighter tails).

From marginals to joint density

- Sklar's theorem: $F(x_1, \ldots, x_N) = \mathcal{C}\big(F_1(x_1), \ldots, F_N(x_N)\big)$.
- Joint density:

$$f(\mathbf{x}) = c\big(F_1(x_1), \ldots, F_N(x_N)\big) \prod_{i=1}^{N} f_i(x_i), \tag{5}$$

where $c(u_1, \ldots, u_N) = \frac{\partial^N}{\partial u_1 \cdots \partial u_N} \mathcal{C}(u_1, \ldots, u_N)$.

- Enabling equivalent pricing reformulations:

$$\begin{aligned} V_0 &= e^{-rT}\mathbb{E}[h(\mathbf{S})] = e^{-rT} \int_{\Omega} f(\mathbf{S}) h(\mathbf{S})\, d^N\mathbf{S} \\ &= e^{-rT} \int_{\Omega} h(\mathbf{S}) c\left(F_1(S_1), \ldots, F_N(S_N)\right) \cdot \prod_{i=1}^{N} f_i(S_i) d^N\mathbf{S} \\ &= e^{-rT}\mathbb{E}^{\text{ind}}[h(\mathbf{S}) c\left(F_1(S_1), \ldots, F_N(S_N)\right)], \end{aligned} \tag{6}$$

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

Cosine Series for Marginals (I)

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

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

$$\gamma_k(x) = \begin{cases} \frac{1}{\sqrt{b-a}}, & k=0, \\ \sqrt{\frac{2}{b-a}}\cos\left(\frac{k\pi(x-a)}{b-a}\right), & k\geq 1. \end{cases} \tag{7}$$

Orthonormality:

$$\int_a^b \gamma_k(x)\gamma_l(x)\,dx = \delta_{kl}. \tag{8}$$

Projection (Fourier) coefficients:

$$a_k = \int_a^b f(x)\gamma_k(x)\,dx = \mathbb{E}[\gamma_k(X)], \quad X\sim f. \tag{9}$$

Cosine Series for Marginals (II)

Let X_i, f_i and F_i, denote the marginal random variable and its density and distribution functions, respectively.

Truncated Series. Given estimated coefficients $\hat{a}_k^{X_i} \approx a_k^{X_i} := \mathbb{E}[\gamma_k(X_i)]$,

$$f_i \approx \hat{f}_i(x) := \sum_{k=0}^{\mathcal{K}_i - 1} \hat{a}_k^{X_i} \, \gamma_k(x). \tag{10}$$

By integration, we get an approximation $\hat{F}_i$ for F_i,

$$\hat{F}_i(x) := \begin{cases} 0, & x < a_i, \\ \displaystyle\sum_{k=0}^{\mathcal{K}_i} \hat{a}_k^{X_i} \, \Gamma_{k,[a_i,b_i]}(x), & a_i \leq x < b_i, \\ 1, & x \geq b_i, \end{cases} \tag{11}$$

where $\Gamma_{k,[a_i,b_i]}(x) := \int_{a_i}^{x} \gamma_k(t) \, dt$ can be readilly solved analytically.

Cosine Series for Marginals (III)

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:

$$\sup_{x\in[a,b]} |f(x) - f^{\mathcal{K}}(x)| \leq \zeta\, e^{-\nu\mathcal{K}}, \tag{12}$$

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.
Convenient representation for quantum loading: Cosine basis functions can be natively loaded in a quantum system (Fourier-state preparation).

Quantum accelerated Monte Carlo

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

Classical Monte Carlo (CMC): L samples $\Rightarrow \varepsilon \sim 1/\sqrt{L} \Rightarrow O(1/\varepsilon^2)$.

By Brassard et al. (2002) and Montanaro (2015): Given the superposition state

$$|\psi\rangle = \mathcal{U}|0\rangle = \sqrt{\mu}\,|\psi_1\rangle|1\rangle + \sqrt{1-\mu}\,|\psi_0\rangle|0\rangle \tag{13}$$

Quantum Amplitude Estimation (QAE) returns $\tilde{\mu}$ with $|\tilde{\mu} - \mu| \leq \varepsilon$ using $O(1/\varepsilon)$.

- **Key assumption:** Querying a quantum oracle one time is equivalent to draw a single sample $\rightarrow$ Quantum-Accelerated Monte Carlo (QAMC).
- **Requirement:** Construct $\mathcal{U}$ such that it encodes the integrand and the density into a quantum state (state preparation).

State preparation

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

$$\mathcal{U}_{f,\phi}|0\rangle^{\otimes Nn}|0\rangle = \sum_{j=0}^{2^{Nn}-1} \sqrt{f(\mathbf{x}_j)}\,|\mathbf{x}_j\rangle \left(\sqrt{\phi(\mathbf{x}_j)}|1\rangle + \sqrt{1-\phi(\mathbf{x}_j)}|0\rangle\right) \quad (14)$$

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

- The oracle $\mathcal{U}$ is commonly composed of two operators, one loading the density, $\mathcal{A}_X$, and one loading the function of interest, W_ϕ.

- The probability of measuring $|1\rangle$ estimates the expected value $\mathbb{E}[\phi(\mathbf{x})]$.

Quantum computation of series coefficients

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 $1 - \varrho_i$, achieving error ε_i in the CDF F_i requires

$$O\left(\frac{\sqrt{b_i - a_i}\, \mathcal{K}_i^2}{\varepsilon_i} \log \frac{\mathcal{K}_i}{\varrho_i} \right) \tag{15}$$

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

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

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

Quantum Pricing of Multi-Asset Options (I)

Joint formulation – direct density loading

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

$$\mathcal{A}_{\mathbf{X}}|0\rangle^{\otimes Nn} = \sum_{j=0}^{2^{Nn}-1} \sqrt{\hat{f}(\mathbf{x}_j)}\,|\mathbf{x}_j\rangle \tag{16}$$

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

$$W_h : |\mathbf{x}_j\rangle|0\rangle \;\mapsto\; |\mathbf{x}_j\rangle \left(\sqrt{h(\mathbf{x}_j)}|1\rangle + \sqrt{1-h(\mathbf{x}_j)}|0\rangle\right) \tag{17}$$

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

$$\mathcal{U}_{f,h}|0\rangle^{\otimes Nn}|0\rangle = \sum_j \sqrt{\hat{f}(\mathbf{x}_j)}\,|\mathbf{x}_j\rangle \left(\sqrt{h(\mathbf{x}_j)}|1\rangle + \sqrt{1-h(\mathbf{x}_j)}|0\rangle\right) \tag{18}$$

The probability of measuring $|1\rangle$ estimates $\sum_j \hat{f}(\mathbf{x}_j)h(\mathbf{x}_j) \approx \mathbb{E}[h(\mathbf{X})]$ with $O\left(\frac{1}{\varepsilon}\log\frac{1}{\varrho}\right)$, for a prescribed error ε and confidence $1-\varrho$ (QAE theorem).

Quantum Pricing of Multi-Asset Options (II)

Independent formulation – copula absorbed into payoff

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

$$\mathcal{A}_{\mathbf{X}}^{\text{ind}}|0\rangle^{\otimes Nn} = \sum_j \sqrt{\prod_i f_i(x_{i,j})}\,|\mathbf{x}_j\rangle \tag{19}$$

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

$$\hat{H}(\mathbf{x}) = \frac{1}{c_{\max}}\, h(\mathbf{x})\, c\big(\hat{F}_1(x_1), \ldots, \hat{F}_N(x_N)\big) \tag{20}$$

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

Why this works:

$$\mathbb{E}^{\text{ind}}[\hat{H}(\mathbf{X})] = \int \Big(\prod_i f_i(x_i)\Big) \frac{h(\mathbf{x})\, c(F_1, \ldots, F_N)}{c_{\max}}\, d\mathbf{x} = \frac{1}{c_{\max}}\, \mathbb{E}[h(\mathbf{X})] \tag{21}$$

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

Complexity Breakdown (I)

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

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

$$O\left(\frac{N^2\, c'_{\max} \sqrt{I_{\max}}\, \mathcal{K}_{\max}}{\varepsilon_c} \log \frac{N\mathcal{K}_{\max}}{\varrho_c}\right) \tag{22}$$

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

$$O\left(\underbrace{\frac{N^2 c'_{\max} I_{\max} \mathcal{K}^2_{\max}}{\varepsilon_c} \log \frac{N\mathcal{K}_{\max}}{\varrho_c}}_{\text{marginal estimation}} + \underbrace{\frac{N c_{\max}}{\varepsilon_c} \log \frac{1}{\varrho_c}}_{\text{pricing phase}}\right) \tag{23}$$

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

Complexity Breakdown (II)

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

$$O\left(\frac{c_{\max}}{\varepsilon_c}\log\frac{1}{\varrho_c}\right) \tag{24}$$

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

$\boxed{\text{Both formulations achieve } O(1/\varepsilon_c) \text{ scaling -- quadratic speedup over CMC}}$

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

Experimental setup

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

- Credit Agricole ($S_0 = 12.91€$), AXA ($S_0 = 33.80€$), Michelin ($S_0 = 31.76€$)
- NIG marginals calibrated to 1-year expiry European options

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
- Computation of 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)

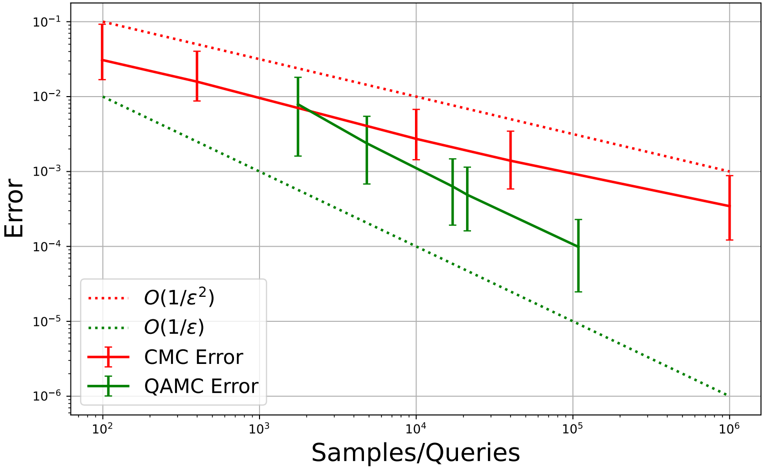

- Average error across all coefficients with 90% confidence intervals shown
- Demonstrates overall QAMC advantage

Density & CDF reconstruction (AXA)

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

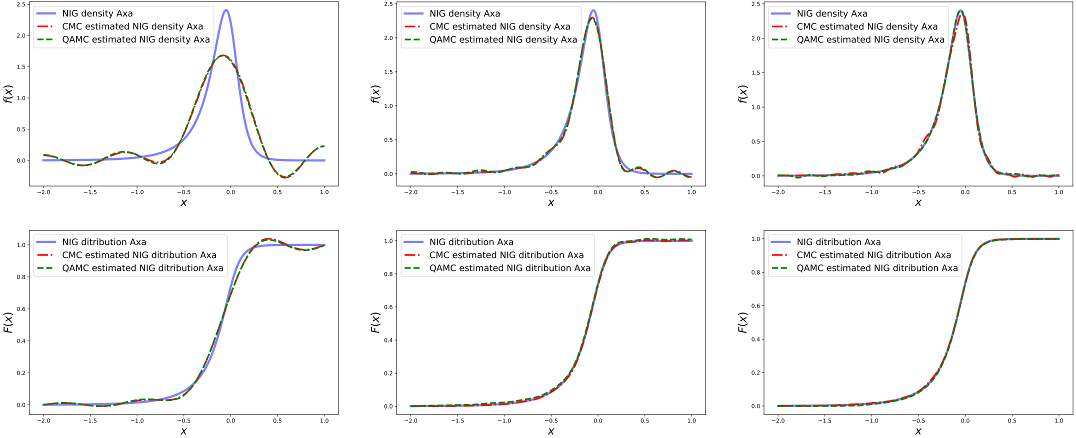

Option pricing: Spread call (AXA & Michelin)

Settings:

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

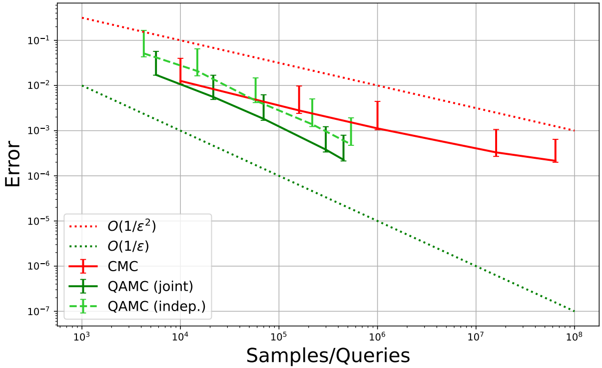

Observations:

- 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.

Option pricing: Arithmetic basket call

Settings:

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

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

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

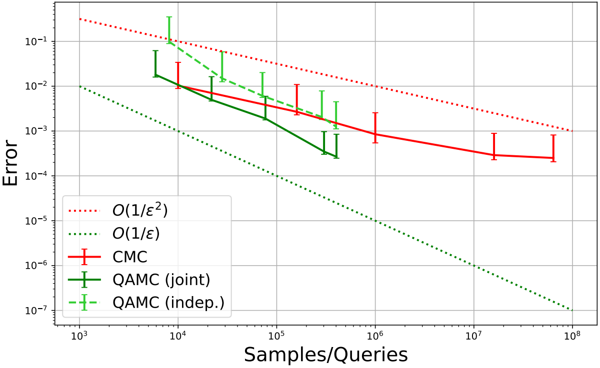

Conclusion: Quantum advantage persists in higher dimensions; joint formulation more efficient.

Summary

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

Future work

- Quantum calibration of the marginals under NIG.
- 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).

Selected references

- Hok, J., Leitao, Á. (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* 305, 53-74.
- Montanaro, A. (2015). Quantum speedup of Monte Carlo methods. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 471*.
- Fang, F., Oosterlee, C.W. (2009). A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM Journal on Scientific Computing* 31, 826-848.
- 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.

Acknowledgements & Questions

Bedankt!

More:

`alvaro.leitao@udc.gal`

`alvaroleitao.github.io`

Thanks to the funding received from the MCIU of Spain through the Ramón y Cajal 2022 and PID2022-141058OB-I00 grants, and from the Xunta de Galicia through the programs ED451C 2022/047 and ED431F 2025/032, as well as the support from CITIC under the agreement for the development of strategic R&D&I actions 2025-2026 and as a centre accredited for excellence within the Galician University System and a member of the CIGUS Network, receiving subsidies from the Xunta de Galicia, additionally, co-financed by the EU through the FEDER Galicia 2021-27 (60%) in the framework of the operational program ED451G 2023/01.

Rede CIGUS
Centros de investigación do Sistema Universitario de Galicia

Cofinanciado por la Unión Europea

Fondos Europeos

Calibration of NIG marginals

Regularized least-squares

$$\mathcal{I}(\theta) = \sum_{m=1}^{M} w_m \left(V_m^{\mathrm{NIG}}(\theta) - \bar{V}_m\right)^2 + \lambda \|\theta - \theta_0\|^2, \quad \theta = (\alpha, \beta, \delta) \tag{25}$$

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)

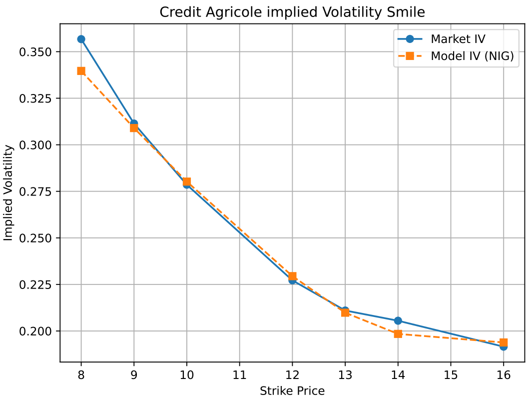

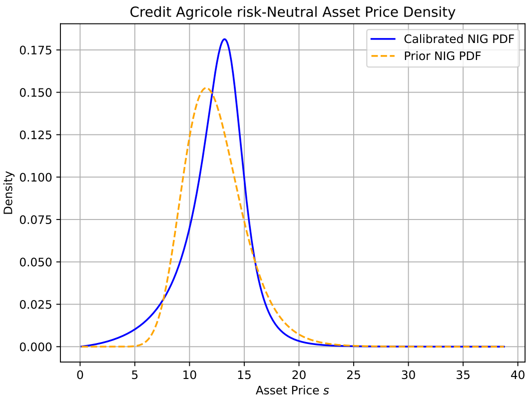

Spot $S_0 = 12.91€$ (as of 24/12/2024), $\bar{\alpha} = 4.69$, $\bar{\beta} = -3.06$, $\bar{\delta} = 0.18$, $\lambda = 5 \times 10^{-7}$

Max pricing error ≈ 21 bps (tail).

Calibration results: AXA (1Y)

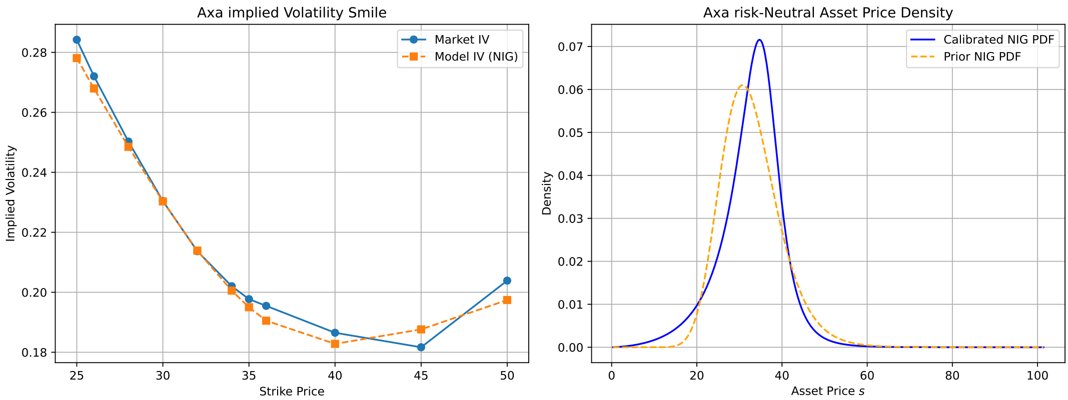

Spot $S_0 = 33.80$€ (as of 24/12/2024), $\bar{\alpha} = 5.24,\ \bar{\beta} = -3.26,\ \bar{\delta} = 0.18, \quad \lambda = 5 \times 10^{-7}$

Max pricing error ≈ 10 bps.

Calibration results: Michelin (1Y)

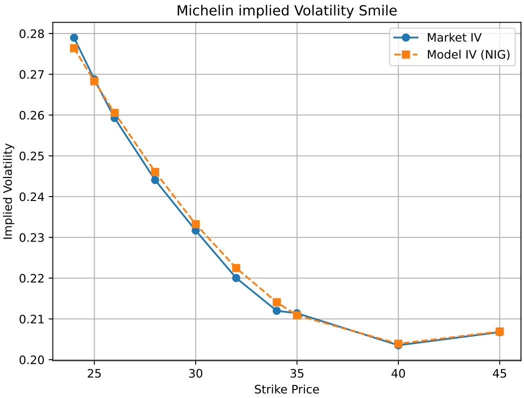

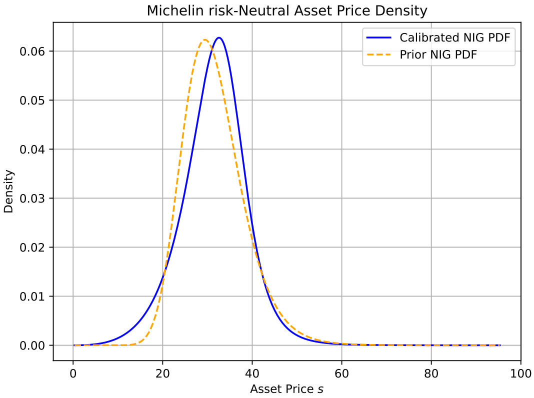

Spot $S_0 = 31.76$€ (as of 24/12/2024), $\bar{\alpha} = 6.20,\ \bar{\beta} = -3.31,\ \bar{\delta} = 0.26, \quad \lambda = 5 \times 10^{-7}$

Max pricing error ≈ 6 bps.