Quantum unary approach to option pricing

Sergi Ramos-Calderer,^{1,2,*} Adrián Pérez-Salinas,^{1,3} Diego García-Martín,^{1,3,4} Carlos Bravo-Prieto,^{1,3} Jorge Cortada,⁵ Jordi Planagumà,⁵ and José I. Latorre^{1,2,6}¹*Departament de Física Quàntica i Astrofísica and Institut de Ciències del Cosmos (ICCUB), Universitat de Barcelona, Martí i Franquès 1, 08028 Barcelona, Spain.*²*Quantum Research Centre, Technology Innovation Institute, Abu Dhabi, UAE.*³*Barcelona Supercomputing Center (BSC), Spain.*⁴*Instituto de Física Teórica, UAM-CSIC, Madrid, Spain.*⁵*Caixabank, Barcelona, Spain.*⁶*Center for Quantum Technologies, National University of Singapore, Singapore.*

We present a quantum algorithm for European option pricing in finance, where the key idea is to work in the unary representation of the asset value. The algorithm needs novel circuitry and is divided in three parts: first, the amplitude distribution corresponding to the asset value at maturity is generated using a low depth circuit; second, the computation of the expected return is computed with simple controlled gates; and third, standard Amplitude Estimation is used to gain quantum advantage. On the positive side, unary representation remarkably simplifies the structure and depth of the quantum circuit. Amplitude distributions uses quantum superposition to bypass the role of classical Monte Carlo simulation. The unary representation also provides a post-selection consistency check that allows for a substantial mitigation in the error of the computation. On the negative side, unary representation requires linearly many qubits to represent a target probability distribution, as compared to the logarithmic scaling of binary algorithms. We compare the performance of both unary vs. binary option pricing algorithms using error maps, and find that unary representation may bring a relevant advantage in practice for near-term devices.

I. INTRODUCTION

Quantum computing provides new strategies to address problems that nowadays are considered difficult to solve by classical means. The first quantum algorithms showing a theoretical advantage over their classical counterparts are known since the 1990s, such as integer factorization to prime numbers [1] or a more efficient unstructured database search [2]. Nevertheless, current quantum devices are not powerful enough to run quantum algorithms that are able to compete against state-of-the-art classical algorithms. Indeed, available quantum computers are in their Noisy Intermediate-Scale Quantum (NISQ) stage [3], as errors due to decoherence, noisy gate application or error read-out limit the performance of these new machines. These NISQ devices may nonetheless be useful tools for a variety of applications due to the introduction of hybrid variational methods. Some of the proposed applications include quantum chemistry [4–6], simulation of physical systems [7–9], combinatorial optimization [10], solving large systems of linear equations [11–13], state diagonalization [14, 15] or quantum machine learning [16–18]. Some exact, non-variational, quantum algorithms are also well suited for NISQ devices [19–22].

A field that is expected to be transformed by the improvement of quantum devices is quantitative finance [23–27]. In recent years, there has been a surge of new methods and algorithms dealing with financial problems

using quantum resources, such as optimization problems [28–31] which are in general hard.

Notably, pricing of financial derivatives is a prominent problem, where many of its computational obstacles are suited to be overcome via quantum computation. In this paper we will deal with options, which are a particular type of financial derivatives. Options are contracts that allow the holder to buy (*call*) or sell (*put*) some asset at a pre-established price (*strike*), or at a future point in time (*maturity date*). The payoff of an option depends on the evolution of the asset's price, which follows a stochastic process. A simple, yet successful model for pricing options is the Black-Scholes model [32]. This is an analytically-solvable model that predicts the asset's price evolution to follow a log-normal probability distribution, at a future time t . Then, a specified payoff function, which depends on the particular option considered, has to be integrated over this distribution to obtain the expected return of the option. Current classical algorithms rely on computationally-costly Monte Carlo simulations to estimate the expected return of options.

A few quantum algorithms have been proposed to improve on classical option pricing [33–35]. It has been shown that quantum computers can provide a quadratic speedup in the number of quantum circuit runs as compared to the number of classical Monte Carlo runs needed to reach a certain precision in the estimation. The basic idea is to exploit quantum Amplitude Estimation [36–38]. Nonetheless, this can only be achieved when an efficient way of loading the probability distribution of the asset price is available. The idea of using quantum Generative Adversarial Networks (qGANs) [39, 40] to address this issue has been analyzed [41].

* sergi.ramos@tii.ae

In the following, we propose a quantum algorithm for option pricing. The key new idea is to construct a quantum circuit that works in the unary basis of the asset's value, *i.e.*, in a subspace of the full Hilbert space of n qubits. Then, the evolution of the asset's price is computed using an amplitude distributor module. Furthermore, the computation of the payoff greatly simplifies. A third part of the algorithm is common to previous approaches, namely it uses Amplitude Estimation. The unary scheme brings further advantage since it allows for a post-selection strategy that results in error mitigation. Let us recall that error mitigation techniques are likely to be crucial for the success of quantum algorithms in the NISQ era. On the negative side, the number of qubits in the unary algorithm scales linearly with the number of bins, while in the binary algorithm it is logarithmic with the target precision. This results in a worse asymptotic scaling for the unary algorithm. Yet, our estimates for the number of gates indicate that the crossing point between these two is located at a number of qubits that renders a good precision ($< 1\%$) for real-world applications. Moreover, the performance of the unary algorithm is more robust to noise, as we show in simulations. Hence, our proposal seems to be better suited to be run on NISQ devices. Unary representations have also been considered in previous works [42–44].

We will illustrate our new algorithm focusing on a simple European option, whose payoff is a function of only the asset's price at maturity date, the only date the contract can be executed at. This straightforward example has been chosen as a proof of concept for this new approach. We will compare the performance of our unary quantum circuit with the previous binary quantum circuit proposal, for a fixed precision or binning of the probability distribution.

The paper is organized as follows. We first introduce the basic ideas on option pricing, both classical and quantum, in Sec. II. The unary quantum algorithm is presented and analyzed in Sec. III. We devote Sec. IV to outline the circuit specifications and compare them for the unary and binary quantum algorithms. Sec. V is dedicated to describe the results obtained by means of classical simulations for both algorithms. Lastly, conclusions are drawn in Sec. VI. Further details on several topics are described in the Appendices.

II. BACKGROUND

There are three main pieces that lay the groundwork needed for our algorithm. They are a) the economical model employed in European-option pricing, known as the Black-Scholes model; b) the Amplitude Estimation technique that provides a quadratic quantum advantage over classical Monte Carlo methods; and c) a quantum algorithm for option pricing in the binary basis, as proposed in [33].

A. Black-Scholes model

The evolution of asset prices in financial markets is usually computed using the model established by F. Black and M. Scholes in Ref. [32]. This evolution is governed by two properties of the market, the interest rate and the volatility, which are incorporated into a stochastic differential equation.

The Black-Scholes model for the evolution of an asset's price at time T , S_T , is based on the following stochastic differential equation,

$$dS_T = S_T r dT + S_T \sigma dW_T, \quad (1)$$

where r is the interest rate, σ is the volatility and W_T describes a Brownian process. Let us recall that a Brownian process W_T is a continuous stochastic evolution starting at $W_0 = 0$ and consisting of independent gaussian increments. To be specific, let $\mathcal{N}(\mu, \sigma_s)$ be a normal distribution with mean μ and standard deviation σ_s . Then, the increment related to two steps of the Brownian processes is $W_T - W_S \sim \mathcal{N}(0, T - S)$, for $T > S$.

The stochastic differential equation (1) can be approximately solved analytically to first order, yielding the solution

$$S_T = S_0 e^{(r - \frac{\sigma^2}{2})T} e^{\sigma W_T} \sim e^{\mathcal{N}((r - \frac{\sigma^2}{2})T, \sigma\sqrt{T})}, \quad (2)$$

which corresponds to a log-normal distribution. The details of this procedure are outlined in App. A.

To obtain the expected return of an option, a payoff function has to be integrated over the resulting probability distribution. This is usually solved using classical Monte Carlo simulation.

In the case of European options, the payoff function is

$$f(S_T, K) = \max(0, S_T - K), \quad (3)$$

yielding an expected payoff given by

$$C(S_T, K) = \int_K^\infty (S_T - K) dS_T, \quad (4)$$

where K is the strike. European options can only be executed at a fixed pre-specified time, called *maturity date*. This is the reason why the payoff is computed using only the probability distribution of S_T at time T .

Our algorithm employs a quantum circuit that generates a probability distribution following Eq. (2), and then encodes the expected payoff of a European option, Eq. (4), into the amplitudes of an ancilla qubit.

B. Amplitude Estimation

Amplitude Estimation (AE) is a quantum technique that allows to estimate the probability of obtaining a certain outcome from a quantum state (with a given precision), with up to a quadratic speedup in the number of function calls as compared to direct sampling [36, 45].

AE with Quantum Phase Estimation

Let us take an algorithm $\mathcal{A}$ such that

$$\mathcal{A}|0\rangle_n|0\rangle = \sqrt{1-a}|\psi_0\rangle_n|0\rangle + \sqrt{a}|\psi_1\rangle_n|1\rangle, \quad (5)$$

where the last qubit serves as an ancilla qubit and the states $|\psi_{0,1}\rangle_n$ can be non-orthogonal. The ancilla qubit is a flag which enables to identify the states as *good* ($|1\rangle$) or *bad* ($|0\rangle$). The state $\mathcal{A}|0\rangle_n|0\rangle$ can be directly sampled N times, and the estimate for probability of finding a good outcome will be $\bar{a}$, with

$$|a - \bar{a}| \sim \mathcal{O}(N^{-1/2}), \quad (6)$$

as dictated by the sampling error of a multinomial distribution.

However, AE can improve this result. Let us first define the central operator for AE [36]

$$\mathcal{Q} = -\mathcal{A}\mathcal{S}_0\mathcal{A}^\dagger\mathcal{S}_{\psi_0}, \quad (7)$$

where the operators $\mathcal{S}_0$ and $\mathcal{S}_{\psi_0}$ are inherited from Grover's search algorithm [2], being

$$\mathcal{S}_0 = \mathbf{I} - 2|0\rangle_n\langle 0|_n \otimes |0\rangle\langle 0|, \quad (8)$$

$$\mathcal{S}_{\psi_0} = \mathbf{I} - 2|\psi_0\rangle_n\langle \psi_0|_n \otimes |0\rangle\langle 0|. \quad (9)$$

The $\mathcal{S}_0$ operator changes the sign of the $|0\rangle_n|0\rangle$ state, while $\mathcal{S}_{\psi_0}$ takes the role of an oracle and changes the sign of all bad outcomes. The operator $\mathcal{Q}$ has eigenvalues $e^{\pm i2\theta_a}$, with $a = \sin^2(\theta_a)$. The procedure of Quantum Phase Estimation (QPE) is then applied to extract an integer number $y \in \{0, 1, \dots, 2^m - 1\}$ such that $\bar{\theta}_a = \pi y / 2^m$ is an estimate of θ_a , with m the number of ancilla qubits. Recall that a Quantum Fourier Transform (QFT) is required to perform QPE.

The value of $\bar{\theta}_a$ leads to an estimate of $\bar{a}$, such that

$$|a - \bar{a}| < \frac{2\pi\sqrt{a(1-a)}}{2^m} + \frac{\pi^2}{2^{2m}} \sim \mathcal{O}\left(\frac{\pi}{2^m}\right) \quad (10)$$

with probability at least $8/\pi^2 \approx 81\%$.

The original Amplitude Estimation procedure requires the implementation of QPE, which is highly resource demanding. Hence, the complexity of the circuit precludes its feasibility in the NISQ era.

AE without Quantum Phase Estimation

Recently, there has been a new proposal for Amplitude Estimation that does not require QPE, and therefore, is less resource-demanding. This approach is based on iterative procedures [45]. The key fact allowing to circumvent the use of QPE is that

$$\mathcal{Q}^m\mathcal{A}|0\rangle = \cos((2m+1)\theta_a)|\psi_0\rangle_n|0\rangle + \sin((2m+1)\theta_a)|\psi_1\rangle_n|1\rangle. \quad (11)$$

An integer m is chosen to prepare the state in Eq. (11) and its outcome is measured N_{shots} times, so that the value of $\sin^2((2m+1)\theta_a)$ is estimated with a precision of $\sim N_{\text{shots}}^{-1/2}$. This process is repeated several times with different values of m extracted from a set of $\{m_j\}$. At the end of the procedure, the precision achieved is bounded by $\sim N_{\text{shots}}^{-1/2}M^{-1}$, with $M = \sum_{j=0}^J m_j$, where J is the last index. The exact scaling of the precision depends on the choice of m_j 's. In App. D the full method is explained in further detail.

C. Binary algorithm

We now present a binary algorithm for option pricing, as introduced in Ref. [33]. This algorithm is divided in three parts.

- (a) *Amplitude distributor*: it encodes the underlying probability distribution of an asset's price at maturity date into a quantum register. The operator representing this piece will be denoted by $\mathcal{D}$. This algorithm uses a quantum Generative Adversarial Network (qGAN) [39–41] in order to fulfill this part. Classical knowledge of the probability distribution is required at this stage.
- (b) *Payoff calculation*: it computes the expected payoff of the option, which is encoded into the amplitude of an ancillary qubit. The operators that perform this step will be a comparator $\mathcal{C}$, that separates the state as above or below the strike, and a set of controlled rotations $\mathcal{R}$, that encode the expected payoff into the probability of measuring an ancilla.
- (c) *Amplitude Estimation*: it extracts the expected payoff calculation encoded in the amplitude of the ancilla, reducing the number of circuit calls needed to reach a desired precision. It is based on the operator $\mathcal{Q}$, which may be applied several times.

A sketch of a quantum circuit implementing the full algorithm is shown in Fig. 1. For a detailed description of each part, refer to App. B.

III. UNARY ALGORITHM

We present now a quantum algorithm that prices European options according to the Black-Scholes model, as outlined in Sec. II A. The key new idea is to construct a quantum circuit that works in the unary basis of the asset's value. The structure of the algorithm is inherited from the one explained in Sec. II C, namely: amplitude distributor module, computation of the payoff and Amplitude Estimation. Furthermore, the implementation of all different pieces is greatly simplified with respect to the binary case. The unary scheme brings further advantage in practice, since it allows for a post-selection

Figure 1. Full circuit for the binary algorithm for option pricing that include all steps, namely, the amplitude distributor $\mathcal{D}$, payoff estimator comprised of the comparator and payoff estimator $\mathcal{C}$ and $\mathcal{R}$ respectively, followed by components of Amplitude Estimation, $\mathcal{Q}$. The operator $\mathcal{Q}$ is repeated m times, where m depends on the AE algorithm. The payoff is indirectly measured in the last qubit.

strategy that results in error mitigation. Although our unary algorithm requires more qubits than a binary one, its performance is more robust to noise and probably better suited to be run on NISQ devices.

A. Unary representation

The main feature of the algorithm is that it works in the *unary representation* of the asset value encoded on the quantum register. That means that for every element of the basis only one qubit will be in the $|1\rangle$ state, whereas all others will remain in $|0\rangle$. A quantum register $|\psi\rangle_n$ made of n qubits in the unary representation can be written as

$$\begin{aligned} |\psi\rangle &= \sum_{i=0}^{n-1} \psi_i |i\rangle_n = \sum_{i=0}^{n-1} \psi_i \left(\bigotimes_{j=0}^{n-1} |\delta_{ij}\rangle \right) = \\ &= \psi_0 |00\dots 01\rangle_n + \psi_1 |00\dots 10\rangle_n + \dots \\ &\quad + \psi_{n-2} |01\dots 00\rangle_n + \psi_{n-1} |10\dots 00\rangle_n, \end{aligned} \quad (12)$$

where $|i\rangle_n$ corresponds to the i -th element of the unary basis, δ_{ij} is the Kronecker delta and $\sum_{i=1}^n |\psi_i|^2 = 1$. A well-known example of a state in the unary representation is the W state, which defines a class of three-qubit multipartite entanglement [46]. Depicted in Fig. 2 is a visual representation of how the unary algorithm would map the outcomes of a Monte Carlo simulation of the asset's price to a quantum register. The ratio of Monte Carlo paths leading to each of the bins will translate into the amplitudes of the corresponding unary basis states.

Given a fixed number of qubits, the unary scheme allows for a lower precision than the binary one. Indeed,

only n out of the 2^n basis elements of the Hilbert space are used. However, due to the natural mapping between the unary representation and the asset's price evolution, we will find that the probability distribution loading and the expected payoff calculation can be carried out with much simpler quantum circuits. On real devices, the potential gain of the unary representation translates into a shallower circuit depth and simpler connectivity requirements. Furthermore, the unary scheme provides means to post-select results so as to increase the faithfulness of the computation. This is due to the fact that the unary representation resides within a restricted part of the Hilbert space, and that extra space can be used as an indicator of the appearance of errors. As a matter of fact, given a realistic precision goal ($<1\%$), it may well turn out to be advantageous to move to the unary representation on NISQ devices, as it simplifies the complexity of the circuit and mitigates errors.

In most cases, a quantum computation does not start in a quantum state that belongs to the unary representation, but in the $|0\rangle_n$ state instead. To solve this issue, we act with a single Pauli X-gate on any qubit. In our algorithm, this qubit is chosen to be the central one to improve overall circuit depth. At this point, the register displays a single qubit in $|1\rangle$ while all others are in $|0\rangle$. This register is an element of the unary basis.

B. Implementation of the algorithm

The basic structure of the unary algorithm is directly inherited from the structure of the binary one. All three independent parts are *Amplitude distributor*, *Payoff calculator* and *Amplitude Estimation*. We discuss them now in further detail.

Amplitude distributor

The probability distribution predicted by the Black-Scholes model is based on the one in Eq. (2). For a given number of qubits, that is of precision, the asset price at any time can be mapped to a fixed depth circuit that distributes probabilities according to the final desired result. The unary representation is akin to the value of the asset. In other words, for every element in the superposition describing the quantum register, the qubit which is flipped into $|1\rangle$ determines the value of the asset. The classical Monte Carlo spread of asset values will be mapped into the probability of measuring each unary basis element.

The quantum circuit generating the final register operates as a distributor of probability amplitudes. The initial state of the algorithm is given by $|0\dots 010\dots 0\rangle_n$, i.e., the element of the unary basis with $|1\rangle$ in the middle qubit. Then, the coefficients of the register in the next step of the circuit are generated using partial-SWAP gates (also called parametrized-SWAP or SWAP power gate) between the middle qubit and its neighbors. The

Figure 2. Scheme for the quantum representation of a given asset price at maturity date. For a given number of Monte Carlo paths, a binning scheme must be applied in such a way that the prices of the asset are separated according to its value. Different Monte Carlo paths that end up in the same bin are color coded accordingly. Each bin is mapped then to an element of the unary basis, whose coefficient is the number of Monte Carlo paths in this bin. The quantum representation of the asset price at maturity contains all possible Monte Carlo paths simultaneously. The precision is then bounded by the numbers of bins we can store on a quantum state, *i.e.*, how many qubits are available.

partial-SWAP gate is defined as

$$\text{partial-SWAP}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta/2 & \sin \theta/2 & 0 \\ 0 & -\sin \theta/2 & \cos \theta/2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (13)$$

Moreover, the partial-SWAP gate could be substituted with a partial-iSWAP gate which performs the same purpose of amplitude sharing. This partial-iSWAP gate,

$$\text{partial-iSWAP}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta/2 & -i \sin \theta/2 & 0 \\ 0 & -i \sin \theta/2 & \cos \theta/2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (14)$$

is a universal entangling gate that comes naturally from the capacitive coupling of superconducting qubits [47, 48]. As a matter of fact, Google's Sycamore chip in which the supremacy experiment was performed [49] allows for this type of gates as they are of great importance as well for quantum chemistry applications [50, 51] or combinatorial optimization [52–54].

This provides the first step to distribute the probability amplitude from the middle qubit to the rest. The procedure is repeated until the edge of the system is reached, as illustrated in Fig. 3. Specific angles can be fed into each partial-SWAP gate to obtain the target probability distribution in the unary representation. The detailed procurement of these angles is described in App. C.

Let us note that any final probability distribution at time t can be obtained with this circuit whose depth is independent of time, since all the necessary information is carried in the angles of the partial-SWAP gates. To be precise, given n qubits, the circuit will always be of depth $\lfloor n/2 \rfloor + 1$. The time dependency of the solution is encoded in the angles determining the partial-SWAP

gates. This idea is reminiscent of the quantum circuits that describe the exact solution of the Ising model [19, 55, 56].

Figure 3. Quantum circuit for loading any probability distribution in the unary representation $\mathcal{D}$. The circuit works as a distributor of amplitude probabilities from its middle qubit to the ones in the edges, using partial-SWAP gates that act only on nearest neighbors. Time dependence is encoded in the angles determining the gates.

The mapping of a known probability distribution function to the unary system is dependant on $(n - 1)$ angles that need to be introduced in the partial-SWAP gates. There are two distinct situations depending on whether the final distribution probability is known exactly or not. The first case can be addressed solving an exact set of n equations with $n - 1$ parameters after computing the probability distribution classically. In case only the differential equation is known, but not its solution, other methods should be employed [57].

Figure 4. Quantum circuit that encodes the expected payoff in an ancillary qubit in the unary representation $\mathcal{C} + \mathcal{R}$. Each qubit with a mapped option price higher than the designated strike controls a cR_y gate on the ancilla, where the rotation angle is a function of its contribution to the expected payoff. The comparator $\mathcal{C}$ is constructed through the control wires, while the $\mathcal{R}$ piece is performed by rotations in the last qubit.

Payoff calculator

The expected payoff calculation circuit builds upon the action of the amplitude distributor to encode the expected return on an ancillary qubit. The unary representation allows for a simple algorithm to accomplish this task. The procedure will prepare an entangled state in the form

$$|\Psi\rangle = \sqrt{1-a} |\psi_0\rangle_n |0\rangle + \sqrt{a} |\psi_1\rangle_n |1\rangle, \quad (15)$$

where $|\psi_{0,1}\rangle$ are states in a superposition of the basis elements below and above the strike, respectively. The payoff is encoded within the amplitude $\sqrt{a}$, with $|a| \leq 1$, ready for Amplitude Estimation [38].

The relevant point to encode the payoff of a European option in an ancillary qubit is to distinguish in the quantum register whether the option price S_i is above or below the strike K . This task turns out to be very simple when working in the unary representation, as opposed to the binary one where a comparator $\mathcal{C}$ needs to be introduced. To be explicit, the computation of the payoff can be achieved by applying controlled Y rotations (cR_y gates), whose control qubits are those encoding a price higher than the accorded strike K , namely the operator $\mathcal{R}$. These cR_y gates will only span over those qubits that represent asset values larger than the strike. Note that the depth of the circuit will be $n-k$, where k is the unary label of the strike K , see Fig. 4.

The rotation angle for each cR_y depends on the contribution of the qubit to the expected payoff. This can be achieved using

$$\phi_i = 2 \arcsin \sqrt{\frac{S_i - K}{S_{max} - K}}, \quad (16)$$

Figure 5. Quantum circuit representation of $\mathcal{S}_{\phi_0}$ (Left) and $\mathcal{S}_0$ (Right) required to perform Amplitude Estimation in the unary basis. Notice that operator $\mathcal{S}_0$ is much simpler in the unary representation as it does not require multi-controlled CNOT gates.

where the denominator inside the arcsin argument is introduced for normalization.

Applying the payoff calculator to a quantum state representing the probability distribution, as depicted in Fig. 2, results in

$$\begin{aligned} |\Psi\rangle = & \sum_{S_i \leq K}^{n-1} \sqrt{p_i} |i\rangle_n |0\rangle + \sum_{S_i > K}^{n-1} \sqrt{p_i} \cos(\phi_i/2) |i\rangle_n |0\rangle + \\ & + \sum_{S_i > K}^{n-1} \sqrt{p_i} \sqrt{\frac{S_i - K}{S_{max} - K}} |i\rangle_n |1\rangle. \end{aligned} \quad (17)$$

The state is now in the form of Eq. (15). It is straightforward to see that the probability of measuring $|1\rangle$ in the ancillary qubit is

$$P(|1\rangle) = \sum_{S_i > K} p_i \frac{S_i - K}{S_{max} - K}. \quad (18)$$

In order to recover the encoded expected payoff, we need to measure the probability of obtaining $|1\rangle$ for the ancilla and then multiply it by the normalization factor $S_{max} - K$. Note that the form of the state is such that further Amplitude Estimation can be performed.

Amplitude Estimation

Let us now move to the application of Amplitude Estimation to our unary option pricing algorithm. As described in Sec. IIB, Amplitude Estimation is performed by concatenating the operators $\mathcal{A}$ and its inverse $\mathcal{A}^\dagger$ with operators $\mathcal{S}_0$ and $\mathcal{S}_{\psi_0}$. In the following, we will describe how to implement these operators in the unary algorithm. Detailed implementation can be seen as well in Fig. 5.

The oracle operator $\mathcal{S}_{\psi_0}$ acts by identifying those coefficients corresponding to accepted outcomes and inverting their sign. In this problem, the task of identifying the element of the basis has been already done by the algorithm $\mathcal{A}$. Accepted outcomes are labelled with $|1\rangle$ in

Figure 6. Full circuit for the option pricing algorithm in the unary representation. The gate $\mathcal{D}$ is the probability distributor, and $\mathcal{C} + \mathcal{R}$ represent the computation of the payoff. After applying the algorithm, the oracle $\mathcal{S}_{\psi_0}$, the reverse algorithm and $\mathcal{S}_0$ follow. The last step is applying the algorithm again. This block $\mathcal{Q}$ is to be repeated for Amplitude Estimation. Measurements in all qubits is a requirement for post-selection. The qubit labelled as q_3 is the one starting the unary representation.

the ancilla qubit. Therefore, the function of this oracle can be achieved by local operations in the ancilla qubit. Explicitly, such operation is

$$\mathcal{S}_{\psi_0} = (I^{\otimes n} \otimes (XZX)). \quad (19)$$

Notice that the X -gates can be deleted since they only provide a global sign.

For the case of the operator $\mathcal{S}_0$ we must remark a detail that greatly simplifies this computation. The operator $\mathcal{S}_0$ is normally defined using $|0\rangle$ since most quantum algorithms start on that state, as depicted in Eq. (15). However, a more apt definition should instead include $|initial\rangle$ as a basis for operator $\mathcal{S}_0$, the state onto which the algorithm $\mathcal{A}$ is first applied. For the unary case, if we isolate the first extra X gate, we can consider the algorithm as starting in that state of the unary basis, heavily simplifying the overall construction. That being the case, $\mathcal{S}_0$ can be constructed out of 2 single-qubit gates and one entangling gate.

With the operator $\mathcal{Q}$ constructed, Amplitude Estimation schemes can be performed. Since the unary algorithm is aimed towards NISQ devices, we use an Amplitude Estimation scheme without Quantum Phase Estimation, explained in detail in App. D. The main idea consists in applying operator $\mathcal{Q}$ a different amount of m times, and process the data in order to get an advantage over ordinary sampling.

C. Error mitigation

NISQ era algorithms need to be resilient against gate errors and decoherence, since fault-tolerant logical qubits are still far from being a reality. Error mitigation techniques have been studied in past literature, see Refs.

[58, 59], and some of them might find valid applications in the unary algorithms as well. However, the unary representation we are proposing here turns out to offer an additional, native, post-selection strategy that manages to mitigate different types of errors. This feature is not present in its binary counterpart.

The key idea behind the possibility of accomplishing error mitigation is that unary algorithms should ideally work within the unary subspace of the Hilbert space. As a consequence, the read-out of any measurement should reflect this fact. It is then possible to reject any outcome that does not fulfil this requirement. As a matter of fact, a number of failed repetitions of the experiment could be discarded, what results in a trade-off between reduction of errors and loss of accepted samples.

A scheme for the full circuit is depicted in Fig. 6. In summary, the circuit is composed by one first X gate that initializes the unary basis, one set of amplitude distributor ($\mathcal{D}$) and payoff calculator ($\mathcal{C} + \mathcal{R}$), and m rounds of Amplitude Estimation $\mathcal{Q} = \mathcal{A}\mathcal{S}_0\mathcal{A}^\dagger\mathcal{S}_{\psi_0}$. Read-out in all qubits is a requirement for post-selection to reduce errors.

We will investigate in detail the performance of unary vs. binary circuits for option pricing in Sec. V. There we will find out that the unary representation is advantageous to the binary one, when targeting the same realistic precision and errors are taken into account.

IV. UNARY AND BINARY COMPARISON

We compare here the unary algorithm for option pricing described in Sec. III to the binary one stated in Sec. II C, in terms of the necessary circuit design as well as the number of gates required to apply the algorithm and successfully perform Amplitude Estimation.

A. Ideal chip architecture

The structure of the unary algorithm allows for a simple chip design. In order to upload the desired probability distribution to the quantum register, only local interactions between first-neighbor qubits are required. Therefore, qubits can be arranged on a single 1D line with two-local interactions. Such a connectivity is perfectly suited to carry out the algorithm. In order to compute the expected payoff, the ancillary qubit needs to interact with the rest of the quantum register. This structure is outlined in Fig. 7 for an arbitrary number of qubits.

The simplicity of the architecture needed to implement the unary algorithm might yield an advantage over alternative algorithms in NISQ computers. Note also that superconducting qubits allow for a natural implementation of the partial-iSWAP gate [47]. This realization of the quantum circuit would result in a decrease in the number of needed gates by factor of 6 in the amplitude distributor module.

On the other hand, the binary algorithm for payoff calculation needs non-local chip connectivity. For the sake of comparison with the simplest chip architecture presented for the unary algorithm, the most basic connectivity needed to perform the steps described for the binary scheme is displayed in Fig. 8. The number of necessary qubits for the binary algorithm scales better asymptotically than in the unary approach, despite the increasing number of ancillary qubits required. Nevertheless, the need for Toffoli gates and almost full connectivity may eliminate this advantage in practical problems for NISQ devices.

B. Gate count

The unary algorithm needs $\mathcal{O}(n)$ partial-SWAP gates in order to distribute the amplitude and $\mathcal{O}(\kappa n)$ controlled- R_y gates to encode the payoff in an ancillary qubit, where $0 \leq \kappa \leq 1$ depends on the strike price K . However, actual quantum devices operate using a native set of gates that are used to construct any other unitary. We present in Table I, left, the gate count of the full circuit as a function of the number of qubits, using either CNOT or partial-iSWAP as the native entangling gate. The gate count assumes the simple ideal chip structure, see Fig. 7, that requires first-neighbor interactions and an ancilla connected to the rest of the qubits.

The partial-iSWAP gate introduces a substantial gain for the amplitude distributor but requires more gates in order to implement the payoff calculation. If both partial-iSWAP interaction between nearest neighbors and CNOT-based connection with the single ancilla are implemented, the best scaling of the full algorithm would be achieved. To be precise, the total number of gates would be $(4\kappa + 1)n + 1$, and the depth of the circuit would become $(4\kappa + \frac{1}{2})n$.

The scaling of the gate count for the binary algorithm

Figure 7. Ideal chip architecture to implement the unary algorithm for option pricing. Only a single ancilla qubit, labelled as a in the image, has to be non-locally controlled by the rest of the qubits. All other interactions are first-nearest-neighbor gates.

Figure 8. Ideal chip architecture to implement the binary algorithm for option pricing with 4 qubits of precision, q_0, q_1, q_2, q_3 , where a and c stand for ancillary and carrier qubit, respectively, and b is another ancilla. The algorithm requires a number of ancillary and carrier qubits equal to the number of precision qubits plus two, $4+2$ in this example. Full connectivity is needed between the precision qubits and two ancillas.

is displayed in Table I, right. We compare the scaling when using CNOT or partial-iSWAP as native gates. The CNOT gate turns out to be more convenient for the binary algorithm. These results include the part of the algorithm that produces the uploading of the probability distribution into the quantum register (hence the dependence on the number of layers of the variational circuit), but it does not take into account the training required by the qGAN. In both unary and binary cases, the counting for single-qubit gates was made compiling several successive single-qubit gates into a single one.

This gate count is performed assuming full connectivity, or at least the connectivity presented in Fig. 8. Existing quantum devices need to implement extra SWAP gates to account for insufficient connections, which are not taken into account in this calculations. Therefore, the gate counting on a computer with less than this ideal connectivity will result in a worse scaling.

Let us emphasize that the gate overhead for the unary representation is much lower than the one for the binary

Unary	CNOT				partial-iSWAP			
	$\mathcal{D}$	$\mathcal{C} + \mathcal{R}$	$\mathcal{S}_{\psi_0}$	$\mathcal{S}_0$	$\mathcal{D}$	$\mathcal{C} + \mathcal{R}$	$\mathcal{S}_{\psi_0}$	$\mathcal{S}_0$
1-qubit gates	$2n$	$2\kappa n$	1	4	1	$\kappa 10n$	1	9
2-qubit gates	$4n$	$2\kappa n$	0	1	n	$\kappa 5n$	0	2
Circuit depth	$3n$	$4\kappa n$	1	5	$n/2$	$15\kappa n$	1	10

Binary	CNOT				partial-iSWAP			
	$\mathcal{D}$	$\mathcal{C} + \mathcal{R}$	$\mathcal{S}_{\psi_0}$	$\mathcal{S}_0$	$\mathcal{D}$	$\mathcal{C} + \mathcal{R}$	$\mathcal{S}_{\psi_0}$	$\mathcal{S}_0$
1 qubit gates	$3nl$	$(16+5\kappa)n$	1	$20n - 23$	$8nl$	$(86+5\kappa)n$	1	$80n - 113$
2 qubit gates	nl	$14n$	0	$12n - 18$	$2nl$	$28n$	0	$24n - 36$
Circuit depth	$nl+1$	$(27+2\kappa)n$	1	$24n - 30$	$6nl+1$	$(97+2\kappa)n$	1	$90n - 129$

Table I. Scaling of the number of 1- and 2-qubit gates and circuit depth as a function of the number of qubits n representing the asset value in unary and binary representations, for the amplitude distributor $\mathcal{D}$, payoff estimator $\mathcal{C} + \mathcal{R}$ and Amplitude Estimation operators $\mathcal{S}_{\psi_0}$ and $\mathcal{S}_0$. Ideal chips architectures are assumed. We compare this scaling in case CNOT or partial-iSWAP gates are implemented. In case the experimental device can implement both CNOT and partial-iSWAP basic gates, the total amount of gates and total depth would be reduced. For the unary circuit, the parameter $0 \leq \kappa \leq 1$ depends on the position of the strike in the qubit register. In the binary case, note the large overheads due to the use of Toffoli gates. The parameter $0 \leq \kappa \leq 1$ characterizes the number of 1s in the binary representation of the strike price. For the amplitude distributor, l is the number of layers of the qGAN.

Figure 9. Scaling of the number of gates required for the full algorithm, including a step, $m = 1$, of Amplitude Estimation, with the number of bins, for different native gates: CNOT gates (Left), partial-iSWAP gates (Center) and the best possible combination (Right), in which one is allowed both CNOT and iSWAP gates as native to the device. The scaling is calculated assuming ideal connectivity, which would largely hinder the binary implementation were that not the case.

case. This is due to the fact that the unary circuit does not require any three-qubit gate. This simplification is eclipsed by the gain in precision for large n , provided an efficient uploading of probability distributions is found for the binary case. The detailed gate count comparing unary *vs.* binary circuits is shown in Fig. 9, where we have taken $\kappa = \frac{1}{2}$ and $l = \frac{\log_2 n}{2}$, where l is the number of layers of the qGAN. In order to compare like with like, the comparison of scaling is made as follows. For a given number of n bins, which directly relate to precision, we take n qubits in the unary representation and only $\log_2 n$ in the binary one. Note that the overhead in the binary representation makes the unary one more convenient for a number of bins less than ~ 100 . This scaling behaviour confirms that this unary representation would get outperformed by the binary one for a large number of bins, provided the devices performed gates with no error. However, if quantum resources are limited, as in NISQ devices, circumstances are favorable for the unary representation. Moreover, in practise, the connectivity requirements further benefits the unary representation over the binary one.

V. SIMULATIONS

The circuits we present in this paper can be simulated using the tools provided by the Python package *Qiskit* [60]. We first consider the unary and binary algorithms in ideal conditions, that is, we verify the performance of the quantum circuits in the absence of any noise. Then, we test them both under increasing amounts of different sources of noise in order to assess which of the two procedures may be more advantageous for NISQ devices.

The simulations in this work were carried out using a simple yet descriptive model. In the case of single-qubit and two-qubit gate errors, we consider depolarizing noise. That is equivalent to transforming the state after each gate application by $\rho \rightarrow (1 - \epsilon)\rho + \epsilon \text{Tr} \rho_d^I$, with d the dimension of the Hilbert space. Measurement errors are ten times more likely to happen (10ϵ), and they are symmetric, *i.e.*, the probability of measuring an incorrect $|0\rangle$ or $|1\rangle$ is identical. Let us remark here that we have not included thermal relaxation or thermal dephasing. The reason is that, given the shallow depth of the simulated circuits, the execution times are far below current coherence times of qubits (the latter being ~ 1000 times the duration of a single-qubit gate), and thermal errors are therefore negligible. This description was adjusted to be comparable to state-of-the-art quantum devices [49].

The accuracy of the expected payoff estimation is used

to benchmark both algorithms against the aforementioned errors, as a function of the interpolation parameter ϵ . The simulations were performed with 8 and 3 qubits in the unary and binary basis, respectively, using their ideal chip structures, see Sec. IV. Notice that both cases correspond to 8 bins. Recall as well that the unary approach includes post-selection, which results in a clear improvement of the algorithm's performance.

The results presented in this section consider depolarizing and measurement errors together. A separate analysis of these errors can be found in App. E. The code is publicly available in Ref. [61]. It allows to perform simulations with different combinations of several errors, namely *bitflip*, *phaseflip*, *bitphaseflip*, thermal and measurement errors, isolated or as part of custom error models.

A. Amplitude distribution loading

The log-normal probability distribution used for the simulations is generated in accordance with the Black-Scholes model discussed in Sec. II A. We work with a particular example, chosen such that the asset price at $T = 0$ is $S_0 = 2$, the volatility of the asset is $\sigma = 40\%$, the risk-free market rate is $r = 5\%$, the maturity time is $T = 0.1$ years, and the accorded strike price for the asset is $K = 1.9$. The simulation of the asset price ranges up to three standard deviations from the mean value of the distribution.

The capability of quantum computers to approximate a given probability distribution in the presence of noise can be quantified by the Kullback-Leibler (KL) divergence [62]. This quantity measures the distance between two probability distributions, vanishing when they are indistinguishable. Fig. 10 plots the KL divergence for the unary and binary approximations to the log-normal distribution. For the maximum allowed error, the KL divergence of the binary algorithm is one order of magnitude larger than that of the unary one.

B. Expected payoff calculation

In terms of payoff calculation, the algorithms diverge slightly. Classically, with a precision of 10^4 bins, the estimated payoff for this financial option is 0.1595, that we take as the exact value for comparison with the quantum strategies. Recall that in order to compare like with like, on the quantum side we have 8 unary qubits and 3 binary qubits, that both correspond to 8 bins.

In Fig. 11, we show the error of the expected payoff as a function of the number of bins in the probability distribution, for the classical computation. This precision depends on the binning and the position of the strike. Therefore, at a large enough number of bins, the results fall within a reasonable percentage of the actual value. At 100 bins, errors for the option price go well below 1%.

Figure 10. Kullback-Leibler divergence between the target probability distribution and those achieved by the quantum algorithms, for equivalent 8 unary qubits and 3 binary qubits, and different levels of depolarizing error. Crosses stand for average results, and the shaded regions encompass the central 70% of the instances. Each probability distribution is estimated using 100 experiments with 10^4 samples each. For noiseless computers, the KL divergence almost vanishes, but gets larger as noise is added. For the maximum allowed error, the KL divergence of the binary algorithm is one order of magnitude larger than that of the unary one.

Figure 11. Percentage error from the exact value of the expected payoff, for the classical computation, as a function of the number of bins in the probability distribution. With only ~ 50 bins, errors for the option price below 0.5% are already reached.

This shows that the unary algorithm can be implemented in the range where it uses less quantum gates than the binary algorithm, and still have low discretization errors coming from the binning.

Figure 12. For equivalent 8-qubit unary and 3-qubit binary algorithms, percentage error in the payoff calculation for depolarizing and measurement errors, up to 0.5% for single-qubit gates, 1% for two qubit gates and 5% for read-out errors, consistent with state-of-the-art devices. The calculations were averaged over 100 repetitions with 10^4 runs each. The shaded regions encompass the central 70% of the instances in each case. The unary algorithm is more robust against these errors.

Robustness against noise

We show in Fig. 12 the average of the relative error of the expected payoff computation when compared to the classical value. The x axis of Fig. 12 depicts the single-qubit gate error percentage ϵ , but two-qubit and read-out errors are also included following the model explained previously. The shaded region includes 70% of the total instances used for the average. It can be seen that the unary algorithm, in general, has less deviation from the mean value than the binary algorithm.

C. Amplitude Estimation

This section comprises results obtained for the Amplitude Estimation algorithm, that can be divided in three parts. First, results are shown for the case of noiseless devices, converging to the expected value within errors due to binning and Taylor approximation, the latter only in the binary case. All results were obtained for 8 bins, unless stated otherwise. Second, an analysis on the effect that quantum errors induce on the estimated value of a has been performed, both for unary and binary approaches. Third, an analysis on the statistical uncertainty incurred in the estimation has been also included.

Only Amplitude Estimation without phase estimation can be performed on NISQ devices. In these simulations, we have used a procedure based on weighted averages that consider both mean values and uncertainties, for a given series of AE steps, see App. D for further details. In our results, every instance has been repeated 100 times.

The choice of m_j is linear, $m_j = j$, with $j = \{0, 1, 2, \dots\}$, in order to control how the performance evolves. The confidence level was adjusted to $1 - \alpha = 0.95$.

In Fig. 13 it is shown how Amplitude Estimation increases the precision of the measured outcome, converging to the actual value as more iterations of AE are used. The results of this simulation unveil that Amplitude Estimation reduces with every iteration the uncertainty in the value of the expected payoff.

The next step of the analysis is to assess the robustness of both the unary and binary representation against noisy circuits. The results for the deviations in the outcomes of a obtained for noisy circuits are depicted in Fig. 14, taking into account depolarizing and read-out errors together. The number of iterations has been limited to $M = 4$. Two very different behaviors can be observed. In the case of the unary approach, the outcomes endure the noise of the device for $M = 0, 1, 2, 3, 4$ and for low error rates, while entering into an erratic regime for large ones. For instance, at $M = 2$, a result that is very close to the optimal value and with low uncertainty is obtained up to error parameter $\epsilon \sim 0.3\%$. In contradistinction, the binary approach loses its robustness for very small noise levels and $M \geq 1$. This can be attributed to the post-selection scheme and to the lower number of applied gates, that benefit the unary algorithm significantly. The decrease of the uncertainties is detailed in Fig. 21 in the Appendix.

From these results we can infer that the Amplitude Estimation procedure, when performed on NISQ devices, provides a quantum advantage only for the unary representation and for limited noise levels in the device; specifically, similar to those present in available state-of-the-art machines [49].

Simulations have been extended to several different numbers of bins in the unary representation. We show in Fig. 15 how the deviation in the payoff from the exact value evolves when larger quantum systems are taken into account. In this example, the error parameter was adjusted to $\epsilon = 0.3\%$. We have considered depolarizing and measurement errors. Each experiment was repeated only 10 times to reduce computational costs. In Fig. 15 it is possible to see that the deviation in the payoff increases as more bins are added. This corresponds to the expected trend since systems with more qubits require a larger number of gates, and thus errors are more likely to happen. Larger errors are observed for numbers of bins between 13 and 18. This behavior is expected to reach a saturation regime for large enough error rates. In the binary case, since the circuit is prone to large errors, the output becomes indistinguishable from one of a random circuit. However, this regime has not been reached yet in the unary representation. Concerning the sampling uncertainty, it is larger as the number of bins increase. This reflects that the more bins, the more errors may occur, and thus more instances are to be discarded via the post-selection mechanism, which translates into a slower convergence rate. The decrease of the uncertainties is

Figure 13. Left: Mean and uncertainty of the outcomes of the expected payoff, obtained with Eq. (D3) and proper transformations. The dashed lines indicate the exact values. Unary and binary approaches are depicted, and convergence to the optimal values are obtained for both. Notice that these values are not the same since the outcomes of both algorithms are not equally related to the payoff. The shaded regions correspond to the statistical uncertainty. Right: Statistical uncertainties in the expected payoff. The dotted lines indicate the uncertainty given by classical sampling, while the dot-dashed lines represent the optimal uncertainty provided by Amplitude Estimation. Results of the simulations lie in between. In this figure we compare procedures with the same number of applications of the $\mathcal{A}$ or $\mathcal{A}^\dagger$ operators, for noiseless circuits.

Figure 14. Results of the errors in the expected payoff respect to the optimal value, for the unary (Left) and binary (Right) representation, with M iterations of Amplitude Estimation for both the unary and binary approaches. Depolarizing and read-out errors have been considered. Scattering points stand for average values, while the shaded region corresponds to the statistical uncertainties of the results. The behavior of both approaches is very different. In the unary case, the expected payoff is resilient to errors. On the other hand, the binary approach returns acceptable results for $M = 0$, while $M \geq 1$ rapidly saturates to the outcome of a random circuit.

detailed in Fig. 22 in the Appendix.

VI. CONCLUSIONS

Finance stands as one of the fields where quantum computation may be of relevance. We have here presented a quantum algorithm that allows for the pricing of European options whose defining trait is that it works in the unary representation of the asset value.

We have illustrated our algorithm in the particular case of a single European option, whose maturity price for the underlying asset is obtained as the solution of a Black-Scholes equation and its expected return depends on a prefixed strike value. The global structure of our algorithm is divided in three steps: a) generation of the amplitude distribution of the asset value at maturity, b) evaluation of the expected return given the strike value, and c) Amplitude Estimation. Our algorithm relies on several new ideas to make this strategy concrete.

Figure 15. Results for the error of the expected payoff for increasing number of bins for up to M iterations of Amplitude Estimation for the unary approach, considering depolarizing and read-out errors together. Scattering points represent the mean values obtained for the experiment, while shadowed areas include 70% of the instances.

The very first step is to define the level of precision the algorithms should aim at. This precision is related to n , the number of qubits in the circuit. The more qubits, the more resolution we can get.

The next step corresponds to building the amplitude distribution of the asset at maturity. We have proposed to handle this problem using a circuit of depth $n/2$ that operates as a distributor of probability amplitude. Given a classical description of the probability distribution, this step of the algorithm substitutes the classical Monte Carlo generation of probabilities.

The computation of the expected return is particularly simple in the unary representation. It only needs a series of n conditional two-body gates from the original qubit register to an ancilla. It then follows iterative Amplitude Estimation.

The use of unary representation seems at odds with performing precise computations. This is not so, as the precision of the expected return is an average over a sampled probability distribution which does not need to have too high a resolution. We verify this statement in detail to find that less than 100 qubits are enough to have com-

petitive computations.

The unary algorithm aims to be used during a middle stage between current quantum computers and fault-tolerant devices. The algorithm is designed to find applicability with a relatively low, while still useful, number of bins. Thus, we have designed a circuit which is simple in terms of logic operations and requires a much less sophisticated connectivity than its binary counterpart.

Unary representation definitely offers relevant advantages over the binary one. First, it allows for a simple distribution of probability amplitudes. Second, it provides a trivial computation of expected returns. Third, unary representation should only trigger one output qubit, while reading the expected return in the ancilla. This offers a consistency check. If no output, or more than one are triggered, the run is rejected. The ability to post-select faithful runs mitigates errors and increases the performance of the quantum algorithm. In addition, Amplitude Estimation may be performed successfully only in the unary basis, considering error levels in NISQ devices, since the procedure is more resilient to errors than the binary one.

There are a number of further improvements that may be included in the algorithm. It is possible, for instance, to increase precision by taking the qubits to represent non equispaced elements in the probability distribution. It is enough to populate more densely the subtle regions of the sample distribution to gain some precision. Ideas to include multi-asset computations are also available [63].

Finally, let us mention that our unary option pricing algorithm could be tested experimentally on quantum computers recently presented.

ACKNOWLEDGEMENTS

SRC, APS, DGM, CBP and JIL are supported by Projects PGC2018-095862-B-C22 and Quantum CAT (001-P-001644). APS, DGM, CBP and JIL acknowledge CaixaBank for its support of this work through Barcelona Supercomputing Center's project CaixaBank Computación Cuántica.

CODE AVAILABILITY

The code is available in Github [61].

-
- [1] P. W. Shor, SIAM Journal on Computing **26**, 1484–1509 (1997).
 - [2] L. K. Grover, Physical Review Letters **79**, 325–328 (1997).
 - [3] J. Preskill, Quantum **2**, 79 (2018).
 - [4] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, Nature Communications **5**, (2014).
 - [5] O. Higgott, D. Wang, and S. Brierley, Quantum **3**, 156 (2019).
 - [6] T. Jones, S. Endo, S. McArdle, X. Yuan, and S. C. Benjamin, Physical Review A **99**, (2019).
 - [7] Y. Li and S. C. Benjamin, Physical Review X **7**, (2017).
 - [8] C. Kokail *et al.*, Nature **569**, 355–360 (2019).
 - [9] C. Cirstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, "Variational fast forwarding for quan-

- tum simulation beyond the coherence time,” (2019), arXiv:1910.04292 [quant-ph].
- [10] E. Farhi, J. Goldstone, and S. Gutmann, “A quantum approximate optimization algorithm,” (2014), arXiv:1411.4028 [quant-ph].
 - [11] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio, and P. J. Coles, “Variational quantum linear solver,” (2019), arXiv:1909.05820 [quant-ph].
 - [12] X. Xu, J. Sun, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, “Variational algorithms for linear algebra,” (2019), arXiv:1909.03898 [quant-ph].
 - [13] H.-Y. Huang, K. Bharti, and P. Rebentrost, “Near-term quantum algorithms for linear systems of equations,” (2019), arXiv:1909.07344 [quant-ph].
 - [14] R. LaRose, A. Tikku, E. O’Neel-Judy, L. Cincio, and P. J. Coles, npj Quantum Information **5**, (2019).
 - [15] C. Bravo-Prieto, D. García-Martín, and J. I. Latorre, Physical Review A **101**, (2020).
 - [16] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Physical Review A **98**, 032309 (2018).
 - [17] D. Zhu *et al.*, Science advances **5**, (2019).
 - [18] A. Pérez-Salinas, A. Cervera-Liarta, E. Gil-Fuster, and J. I. Latorre, Quantum **4**, 226 (2020).
 - [19] A. Cervera-Liarta, Quantum **2**, 114 (2018).
 - [20] Y. Subasi, L. Cincio, and P. J. Coles, Journal of Physics A: Mathematical and Theoretical **52**, 044001 (2019).
 - [21] S. Bravyi, D. Gosset, and R. König, Science **362**, 308–311 (2018).
 - [22] S. Bravyi, D. Gosset, R. Koenig, and M. Tomamichel, “Quantum advantage with noisy shallow circuits in 3d,” (2019), arXiv:1904.01502 [quant-ph].
 - [23] R. Orús, S. Mugel, and E. Lizaso, Reviews in Physics **4**, 100028 (2019).
 - [24] I. Kerenidis, A. Prakash, and D. Szilágyi, “Quantum algorithms for portfolio optimization,” (2019), arXiv:1908.08040 [math.OC].
 - [25] R. Orús, S. Mugel, and E. Lizaso, Physical Review A **99**, (2019).
 - [26] A. Martín, B. Candelas, Á. Rodríguez-Rozas, J. D. Martín-Guerrero, X. Chen, L. Lamata, R. Orús, E. Solano, and M. Sanz, Physical Review Research **3**, 013167 (2021).
 - [27] D. J. Egger, R. G. Gutiérrez, J. C. Mestre, and S. Woerner, “Credit risk analysis using quantum computers,” (2019), arXiv:1907.03044 [quant-ph].
 - [28] G. Rosenberg, P. Haghnegahdar, P. Goddard, P. Carr, K. Wu, and M. L. de Prado, IEEE Journal of Selected Topics in Signal Processing **10**, 1053–1060 (2016).
 - [29] P. Rebentrost and S. Lloyd, “Quantum computational finance: quantum algorithm for portfolio optimization,” (2018), arXiv:1811.03975 [quant-ph].
 - [30] N. Moll *et al.*, Quantum Science and Technology **3**, 030503 (2018).
 - [31] M. L. de Prado, SSRN Electronic Journal, (2015).
 - [32] F. Black and M. Scholes, Journal of Political Economy **81**, 637 (1973).
 - [33] N. Stamatopoulos, D. J. Egger, Y. Sun, C. Zoufal, R. Iten, N. Shen, and S. Woerner, Quantum **4**, 291 (2020).
 - [34] P. Rebentrost, B. Gupt, and T. R. Bromley, Physical Review A **98**, (2018).
 - [35] S. Woerner and D. J. Egger, npj Quantum Information **5**, (2019).
 - [36] G. Brassard, P. Høyer, and M. Mosca, in *Quantum Computation and Information, Contemporary Mathematics Series. AMS* (2002) p. 0005055.
 - [37] S. Aaronson and P. Rall, Symposium on Simplicity in Algorithms, 24–32 (2020).
 - [38] A. Montanaro, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **471**, 20150301 (2015).
 - [39] S. Lloyd and C. Weedbrook, Physical Review Letters **121**, (2018).
 - [40] P.-L. Dallaire-Demers and N. Killoran, Physical Review A **98**, (2018).
 - [41] C. Zoufal, A. Lucchi, and S. Woerner, npj Quantum Information **5**, (2019).
 - [42] D. Poulin, A. Kitaev, D. S. Steiger, M. B. Hastings, and M. Troyer, Physical Review Letters **121**, (2018).
 - [43] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Physical Review X **8**, (2018).
 - [44] M. Steudtner and S. Wehner, Physical Review A **101**, (2020).
 - [45] Y. Suzuki, S. Uno, R. Raymond, T. Tanaka, T. Onodera, and N. Yamamoto, Quantum Information Processing **19**, (2020).
 - [46] W. Dür, G. Vidal, and J. I. Cirac, Physical Review A **62**, (2000).
 - [47] R. C. Bialczak *et al.*, Nature Physics **6**, 409–413 (2010).
 - [48] N. Schuch and J. Siewert, Physical Review A **67**, (2003).
 - [49] F. Arute *et al.*, Nature **574**, 505–510 (2019).
 - [50] P. K. Barkoutsos *et al.*, Physical Review A **98**, (2018).
 - [51] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, npj Quantum Information **6**, (2020).
 - [52] S. Hadfield, Z. Wang, B. O’Gorman, E. Rieffel, D. Venturelli, and R. Biswas, Algorithms **12**, 34 (2019).
 - [53] Z. Wang, N. C. Rubin, J. M. Dominy, and E. G. Rieffel, Physical Review A **101**, (2020).
 - [54] J. Cook, S. Eidenbenz, and A. Bärttschi, “The quantum alternating operator ansatz on max-k vertex cover,” (2019), arXiv:1910.13483 [quant-ph].
 - [55] F. Verstraete, J. I. Cirac, and J. I. Latorre, Physical Review A **79**, (2009).
 - [56] M. Hebenstreit, D. Alsina, J. I. Latorre, and B. Kraus, Physical Review A **95**, (2017).
 - [57] S. Iblisdir, R. Orús, and J. I. Latorre, Physical Review B **75**, (2007).
 - [58] K. Temme, S. Bravyi, and J. M. Gambetta, Physical Review Letters **119**, (2017).
 - [59] S. Endo, S. C. Benjamin, and Y. Li, Physical Review X **8**, (2018).
 - [60] H. Abraham *et al.*, “Qiskit: An open-source framework for quantum computing,” (2019).
 - [61] S. Ramos-Calderer and A. Pérez-Salinas, “Quantum finance,” <https://github.com/UB-Quantic/quantum-unary-option-pricing> (2020).
 - [62] S. Kullback and R. A. Leibler, The Annals of Mathematical Statistics **22**, 79 (1951).
 - [63] C. Bravo-Prieto, D. García-Martín, J. I. Latorre, A. Pérez-Salinas, and S. Ramos-Calderer, in preparation.
 - [64] K. Itô, Proceedings of the Imperial Academy **20**, 519 (1944).
 - [65] L. Grover and T. Rudolph, “Creating superpositions that correspond to efficiently integrable probability distribu-

- tions,” (2002), arXiv:quant-ph/0208112 [quant-ph].
- [66] A. Montanaro and S. Pallister, Physical Review A **93**, (2016).
- [67] J. J. García-Ripoll, “Quantum-inspired algorithms for multivariate analysis: from interpolation to partial differential equations,” (2019), arXiv:1909.06619 [quant-ph].
- [68] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Physical Review A **52**, 3457–3467 (1995).
- [69] S. Wallis, Journal of Quantitative Linguistics **20**, 178 (2013).

Appendix A: Classical Option Pricing

The evolution of asset prices in financial markets is usually computed using a model established by F. Black and M. Scholes in Ref. [32]. This evolution is governed by two properties of the market, the interest rate and the volatility, which are incorporated into a stochastic differential equation. The equations controlling a set of assets are usually solved using Monte Carlo methods.

1. The Black-Scholes model

The Black-Scholes model for the evolution of an asset is based on the following stochastic differential equation [32]

$$dS_T = S_T r dT + S_T \sigma dW_T, \quad (\text{A1})$$

where r is the interest rate, σ is the volatility and W_T describes a Brownian process. Let us recall that a Brownian process W_T is a continuous stochastic evolution starting at $W_0 = 0$ and made of independent gaussian increments. To be specific, let $\mathcal{N}(\mu, \sigma_s)$ be a normal distribution with mean μ and standard deviation σ_s . Then, the increment related to two steps of the Brownian processes is $W_T - W_S \sim \mathcal{N}(0, T - S)$, for $T > S$.

The above differential equation can be solved analytically up to first order using Ito’s lemma [64], whereby W_T is treated as an independent variable with the property that $(dW_T)^2$ is of the order of dT . Thus, the approximated derivative dS_T can be written as

$$dS_T = \left(\frac{\partial S_T}{\partial T} + \frac{1}{2} \frac{\partial^2 S_T}{\partial W_T^2} \right) dT + \frac{\partial S_T}{\partial W_T} dW_T. \quad (\text{A2})$$

By direct comparison to Eq. (A1), it is straightforward to see that

$$\frac{\partial S_T}{\partial W_T} = S_T \sigma, \quad (\text{A3})$$

$$\frac{\partial S_T}{\partial T} + \frac{1}{2} \frac{\partial^2 S_T}{\partial W_T^2} = S_T r. \quad (\text{A4})$$

Using the initial condition S_0 at $T = 0$, and the Ansatz

$$S_T = S_0 \exp\{(f(T) + g(W_T))\}, \quad (\text{A5})$$

the solution for the asset price turns out to be

$$S_T = S_0 e^{(r - \frac{\sigma^2}{2})T} e^{\sigma W_T} \sim S_0 e^{\mathcal{N}\left((r - \frac{\sigma^2}{2})T, \sigma\sqrt{T}\right)}. \quad (\text{A6})$$

This final result corresponds to a log-normal distribution.

2. European Option

An option is a contract where in its call/put form, the option holder can buy/sell an asset before a specific date or decline such a right. As a particular case, European options can be exercised only on the specified future date, and only depend on the price of the asset at that time. The price that will be paid for the asset is called *exercise* price or *strike*. The day on which the option can be exercised is called *maturity date*.

A European option payoff is defined as

$$f(S_T, K) = \max(0, S_T - K), \quad (\text{A7})$$

where K is the strike price and T is the maturity date. An analytical solution exists for the payoff of this kind of options.

The expected payoff is given by

$$C(S_T, K) = \text{average}_{S_T \geq K} (S_T - K) = \int_{d_1}^{\infty} (S_T - K) \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx, \quad (\text{A8})$$

yielding the analytical solution

$$C(S_T, K) = S_0 \text{CDF}_{\mathcal{N}}(d_1) - K e^{-rT} \text{CDF}_{\mathcal{N}}(d_2), \quad (\text{A9})$$

with

$$d_1 = \frac{1}{\sigma\sqrt{t}} \left(\log \frac{S_0}{K} + \left(r + \frac{\sigma^2}{2} \right) T \right) \quad (\text{A10})$$

$$d_2 = d_1 - \sigma\sqrt{T} \quad (\text{A11})$$

$$\text{CDF}_{\mathcal{N}}(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{u^2}{2}} du. \quad (\text{A12})$$

Appendix B: Details for the binary algorithm

For the sake of completeness, we now present a binary algorithm for option pricing, as introduced in Ref. [33]. The binary algorithm is also divided in three parts, namely (a) amplitude distribution loading, (b) expected payoff computation and (c) Amplitude Estimation. The main difference is that all computational-basis states are used to codify the discretized probability distribution of an asset price at maturity time. This implies steps (a) and (b) will require completely different quantum circuitry. We now proceed to describe these steps for the binary case.

1. Amplitude distribution loading

Uploading probability distributions onto quantum states is a very general problem that was considered in [65]. In this work, it was claimed that any probability distribution that is efficiently integrable on a classical computer, *e.g.* log-concave distributions, can be loaded efficiently onto a quantum state. However, several authors [66, 67] have pointed out that the method proposed there requires pre-calculating a number of integrals that grows exponentially with the number of qubits. In the case of option pricing, a reasonable precision requires a moderate number of qubits in the unary representation and many less in the binary representation. But, as a matter of fact, the reduction to a logarithmic number of qubits in the binary representation is, at least partially, compensated by the effort needed to prepare the probability distribution. That is, in practise, both unary and binary representations require similar effort to pre-process the probability distribution to later encode it in the quantum register.

An alternative method to encode a probability distribution in a quantum state is the use of so-called quantum Generative Adversarial Networks (qGANs) [39–41]. In this scheme, two agents, a generator and a discriminator compete against each other. The generator learns to produce data that mimics the underlying probability distribution, trying to deceive the discriminator into believing that the new data is faithful. On the other hand, the discriminator has to learn how to tell apart the real data from the data produced by the generator. This quantum adversarial game has a unique endpoint: Nash equilibrium is reached when the generator learns to produce states that deliver probability outcomes that are indistinguishable from the desired probability distribution, and the discriminator cannot tell them apart. In order to upload probability distributions onto quantum states using qGANs, a parametrized quantum circuit may play the role of a generator, whereas the discriminator may be a classical neural network.

At present, there is still a lack of precise understanding on how to efficiently upload probability distributions on a quantum computer in binary representation, which makes rigorous complexity analysis in terms of the number of gates difficult.

2. Payoff computation

A useful feature of the unary algorithm is that, given a strike K , one can directly know which qubits will not contribute to the expected return of the option, and therefore adjust the quantum circuit. This is only possible

since the unary representation maps directly to the asset price. In a binary encoded setting one needs to compute explicitly which basis elements will make a non-zero contribution to the expected payoff. Hence the need of a quantum comparator, $\mathcal{C}$, that singles out the values of S_T that are smaller than the strike price K . This comparator requires the use of $n + 1$ ancillary qubits, one of which is retained after the computation. Its action is given by

$$|\psi\rangle|0\rangle \xrightarrow{\mathcal{C}} \sum_{S_i < K} \sqrt{p_i} |e_i\rangle|0\rangle + \sum_{S_i \geq K} \sqrt{p_i} |e_i\rangle|1\rangle, \quad (\text{B1})$$

where $\{|e_i\rangle\}$ is the computational basis and $\{S_i\}$ are the asset values at maturity associated to computational-basis vectors. The quantum circuit implementing Eq. (B1) can be constructed using cNOTs, Toffoli gates and OR gates, see Fig. 16. In order to understand the way this circuit works, let us consider the case where the discretization of the interval $[S_{max} - S_{min}]$ is uniform. In this case, the relation between $\{S_i\}$ and $\{|e_i\rangle\}$ is

$$S_i = S_{min} + \frac{e_i (S_{max} - S_{min})}{2^n}. \quad (\text{B2})$$

This implies that

$$S_i > K \Leftrightarrow e_i > \frac{2^n (K - S_{min})}{S_{max} - S_{min}} \equiv K'. \quad (\text{B3})$$

The idea goes as follows. First, we classically compute the two's complement of K' , *i.e.*, $2^n - K'$, and store it in binary format in a classical array of n bits, $t[j]$ with $j \in [0, 1, \dots, n-1]$. Then, using n ancillas, $|a_0 \dots a_{n-1}\rangle$, initialized to $|0 \dots 0\rangle$, we compute the carry bits of the bitwise addition between t and $\{e_i\}$, and store them in superposition into $|a_0 \dots a_{n-1}\rangle$. If $e_i > K'$, then necessarily $a_{n-1} = 1$.

Figure 16. Quantum comparator, $\mathcal{C}$. The OR gates appearing are three-qubit gates acting on non-adjacent qubits.

The exact circuit needed for a given strike will depend upon the values of the bits in t . If $t[j] = 0$, then there will be a carry bit at position j if and only if there is a carry bit at position $j-1$ and the j -th bit of e_i is 1. This is computed with a Toffoli gate. On the other hand, if $t[j] = 1$, there will be a carry bit at position j if and only if there is a carry bit at position $j-1$ or the j -th bit of e_i is 1. This is computed with an OR gate, shown in Fig.17. Finally, there will be a carry bit at a_0 if and only if $t[0] = 1$ and the first bit of e_i is 1. This is achieved with a simple CNOT gate. As explained above, if $e_i \geq K'$, then a_{n-1} must be equal to 1. Hence, applying a CNOT gate controlled by the qubit $|a_{n-1}\rangle$ and targeted at the ancilla, the desired state in Eq. (B1) is obtained.

Once $\mathcal{C}$ has been applied, the next step is to encode the expected payoff of the option into the amplitudes of a new ancilla. The final state to be created should be

$$\sum_{S_i < K} \sqrt{p_i} |e_i\rangle|0\rangle [\cos(g_0)|0\rangle + \sin(g_0)|1\rangle] + \sum_{S_i \geq K} \sqrt{p_i} |e_i\rangle|1\rangle [\cos(g_0 + g(i))|0\rangle + \sin(g_0 + g(i))|1\rangle], \quad (\text{B4})$$

Figure 17. Decomposition of the OR gate in terms of single-qubit and Toffoli gates.

where

$$g_0 = \frac{\pi}{4} - c \quad , \quad g(i) = \frac{2c(e_i - K')}{e_{max} - K'} \quad , \quad (B5)$$

with c a constant such that $c \in [0, 1]$. Thus, the probability of measuring the second ancilla in the $|1\rangle$ state in (B4) is

$$\text{Prob}(1) = \sum_{S_i < K} p_i \sin^2(g_0) + \sum_{S_i \geq K} p_i \sin^2(g_0 + g(i)) \quad . \quad (B6)$$

Using the approximation

$$\sin^2\left(cf(i) + \frac{\pi}{4}\right) = \frac{1}{2} + cf(i) + O(c^3 f^3(i)) \quad , \quad (B7)$$

to first order, which follows from Taylor-expanding $\sin^2(f(x) + \frac{\pi}{4})$ around $f(x) = 0$, the probability becomes

$$\text{Prob}(1) \simeq \sum_{S_i < K} p_i \left(\frac{1}{2} - c\right) + \sum_{S_i \geq K} p_i \left(\frac{1}{2} + c \left[\frac{2(e_i - K')}{e_{max} - K'} - 1\right]\right) = \frac{1}{2} - c + \frac{2c}{e_{max} - K'} \sum_{S_i \geq K} p_i (e_i - K') \quad . \quad (B8)$$

It is important to note that the approximation made in Eq. (B8) is valid since $cf(i) = c \left[\frac{2(e_i - K')}{e_{max} - K'} - 1\right] \in [-c, c]$. Reversing the change from Eq. (B3), namely

$$\sum_{S_i \geq K} p_i (S_i - K) = \frac{S_{max} - S_{min}}{2^n} \sum_{S_i \geq K} p_i (e_i - K') \quad , \quad (B9)$$

the expected payoff function, *i.e.*, $\sum_{S_i \geq K} p_i (S_i - K)$, can be recovered from the probability of measuring 1 in the ancilla, Eq. (B8), since c, K, S_{max} are all known.

Figure 18. Encoding of the expected return of a European option into the amplitudes of an ancilla qubit in binary representation.

The quantum circuit that produces the final state (B4) from Eq. (B1) is composed of n ccR_y gates, plus one R_y and one cR_y gate, shown in Fig. 18. The decomposition of a ccR_y gate in terms of CNOTs and cR_y gates is shown in Fig. 19.

Figure 19. Decomposition of a ccR_y gate in terms of cNOTs and cR_y gates.

Figure 20. Full circuit for the option pricing algorithm in the binary representation. The gate $\mathcal{D}$ is the probability distributor, $\mathcal{C}$ is the comparator, and $\mathcal{R}$ the rotation step. $\mathcal{C}$ and $\mathcal{R}$ together represent the computation of the payoff. After applying the algorithm, the oracle $\mathcal{S}_{\psi_0}$, the reverse algorithm and $\mathcal{S}_0$ follow. The last step is applying the algorithm again. This block $\mathcal{Q}$ is to be repeated for applying Amplitude Estimation.

3. Amplitude Estimation

The oracle operator $\mathcal{S}_{\psi_0}$ acts on the binary algorithm in the same manner as its unary analogue, as defined in Eq. (19).

The case of the operator $\mathcal{S}_0$ is slightly different. The target state to flip is that with only $|0\rangle$. Thus, a multi-controlled Toffoli gate is required. This multi-controlled gate constitutes the computationally most-costly piece of the circuit. This gate can be decomposed in simpler gates [68].

In order to reduce the complexity of the circuit as much as possible, it is necessary to choose the optimal representation of this gate. The most efficient decomposition consists in a chain of standard Toffoli gates that store their outcomes in ancilla qubits. The number of ancillas required is $c - 2$, where c is the number of control qubits. Depending on whether these ancillas are initialized in 0 or not, the number of Toffoli gates is different. In the case of the binary algorithm, a sufficient number of ancillas is available from the comparator. In addition, the short version of the Toffoli gate can be used. The reason is as follows. The operator $\mathcal{S}_0$ is applied only after the sequence $\mathcal{A}^\dagger \mathcal{S}_{\psi_0} \mathcal{A}$, that leaves the ancillas unchanged as the oracle includes only a phase and does not add any entanglement, making all operations involving the ancillas classical. In other words, there is no need to add any circuit piece that undoes the auxiliary computations stored in the ancillary qubits because the structure of the circuit itself accomplishes this goal.

The full circuit for the binary algorithm, including Amplitude Estimation is depicted in Fig. 20

Appendix C: Details for the Amplitude Distributor in the unary basis

Let us consider Fig. 3. In the unary basis, every qubit represents the basis element in which the qubit is $|1\rangle$. Thus, the coefficient of every element depends on as many angles as partial-SWAP gates are needed to reaching its corresponding qubit. Thus, the central qubits of the circuit will depend only on 2 angles, and the number of dependencies increases one by one as we move to the outer part of the circuit. The very last two qubits depend on the same angles. As we move away from the middle, each qubit inherits the same angle dependency than the previous

ones plus an additional rotation. Starting from the two edges, their coefficients verify the following ratios

$$\left| \frac{\psi_0}{\psi_1} \right|^2 = \tan^2(\theta_1/2) \quad (C1)$$

$$\left| \frac{\psi_{n-1}}{\psi_{n-2}} \right|^2 = \tan^2(\theta_{n-1}/2). \quad (C2)$$

Then we identify $|\psi_i|^2 = p_i$, where $\{p_i\}$ is the target probability distribution of the asset prices at maturity. The next step corresponds to considering the qubits 1 and 2, as well as $n-3$, $n-2$. The relations for their coefficients must obey

$$\left| \frac{\psi_i}{\psi_{i+1}} \right|^2 = \cos^2(\theta_i/2) \tan^2(\theta_{i+1}/2) \quad (C3)$$

$$\left| \frac{\psi_{n-1-i}}{\psi_{n-2-i}} \right|^2 = \cos^2(\theta_{n-i}/2) \tan^2(\theta_{n-1-i}/2). \quad (C4)$$

Then, it is straightforward to back-substitute parameters step by step until we arrive to the central qubits. This procedure fixes all the angles for the partial-SWAP gates used in the amplitude distributor.

The exact algorithm to be followed can be also found in the provided code [61].

Once the exact solution for the angles is inserted into the circuit depicted in Fig. 3, the amplitude distributor algorithm is completed. The quantum register then reads

$$|\Psi\rangle = \sum_{i=0}^{n-1} \sqrt{p_i} |i\rangle. \quad (C5)$$

Note that describing a probability distribution with squared amplitudes of a quantum state allows for a free phase in every coefficient of the quantum circuit. For simplicity, we will set to zero all these relative phases by only operating with real valued partial-SWAP gates.

Let us turn our attention to the gates which are needed in the above circuit. Sharing probability between neighbor qubits can be achieved by introducing a two-qubit gate based on the SWAP and R_y operations. This variant on the SWAP gate performs a partial SWAP operation, where only a piece of the amplitude is transferred from one qubit to another. This operation preserves unarity, that is the state remains as a superposition of elements of the unary basis. This partial-SWAP, can be decomposed using CNOT as the basic entangling gate as

$$\text{partial-SWAP}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta/2 & \sin \theta/2 & 0 \\ 0 & -\sin \theta/2 & \cos \theta/2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = \begin{array}{c} \oplus \text{---} \bullet \text{---} \oplus \\ | \\ \bullet \text{---} \boxed{R_y(\theta)} \text{---} \bullet \end{array}, \quad (C6)$$

where the usual CNOT gate in the center of the conventional SWAP gate has been substituted by a controlled y -rotation, henceforth referred to as cR_y gate. In turn, the cR_y operation can be reworked as a combination of single-qubit gates and CNOT gates [68]:

$$\begin{array}{c} \bullet \\ | \\ \boxed{R_y(\theta)} \end{array} = \begin{array}{c} \bullet \text{---} \bullet \\ | \quad | \\ \boxed{R_y(\theta/2)} \text{---} \oplus \text{---} \boxed{R_y(-\theta/2)} \text{---} \oplus \end{array}. \quad (C7)$$

This decomposition will come into play for the expected payoff calculation algorithm as well, albeit with angle ϕ in the payoff circuit.

For the purposes of this algorithm, both the CNOT and partial-iSWAP basis gates are analogous, but the direct modeling to partial-iSWAPs can economize the total number of required gates for the amplitude distributor. Partial-iSWAP gates can be used to decompose CNOT gates. More explicitly, a CNOT gate can be reproduced with two iSWAP gates, and 5 single qubit gates.

Algorithm 1: Algorithm for Amplitude Estimation based on gaussian distribution of the measurements.

```

1 GaussianAmplitudeEstimation( $N_{\text{shots}}, J, m_j, \alpha$ )
2  $z \leftarrow \text{CDF}_{\mathcal{N}}^{-1}(1 - \alpha/2)$ 
3 Ensure  $m_0 = 0$ 
4  $a \leftarrow |\langle 1 | \mathcal{A} | 0 \rangle|^2$  with  $N_{\text{shots}}$  samples
5  $\theta_a^{(0)} \leftarrow \arcsin \sqrt{a} \quad \Delta\theta_a^{(0)} = \frac{z}{2\sqrt{N_{\text{shots}}}}$ 
6 for  $j \leftarrow 1$  to  $J$  do
7    $a \leftarrow |\langle 1 | Q^{m_j} A | 0 \rangle|^2$  with  $N_{\text{shots}}$  samples
8    $\theta_{\text{array}} \leftarrow \text{MultipleValuesArcsin}(a, m_{j-1})$ 
9    $\theta_a \leftarrow \min(|\theta_{\text{array}} - \theta_a^{(j-1)}|)$ 
10   $\Delta\theta_a \leftarrow \frac{z}{2(2m_j+1)\sqrt{N_{\text{shots}}}}$ 
11   $\theta_a^{(m_j)} \leftarrow \frac{\frac{\theta_a}{\Delta\theta_a^2} + \frac{\theta_a^{(j-1)}}{(\Delta\theta_a^{(j-1)})^2}}{\frac{1}{\Delta\theta_a^2} + \frac{1}{(\Delta\theta_a^{(j-1)})^2}}$ 
12   $\Delta\theta_a^{(m_j)} \leftarrow \left( \frac{1}{\Delta\theta_a^2} + \frac{1}{(\Delta\theta_a^{(j-1)})^2} \right)^{-1/2}$ 
13   $[a_j, \Delta a_j] \leftarrow [\sin^2 \theta_a^j, \sin(2\theta_a^j) \Delta\theta_a^{(j)}]$ 
14 return  $[a_j, \Delta a_j]$ 

```

Algorithm 2: Extracting multiple values for the arcsin, auxiliary function needed in Alg. 1.

```

1 MultipleValuesArcsin( $a, m$ )
2  $\theta_0 \leftarrow \arcsin \sqrt{a}$  // The value of  $\theta_0$  is bounded between 0 and  $\pi/2$ 
3 The arcsin function has several solutions  $\theta_{\text{array}} \leftarrow [0] * (2m+1)$   $\theta_{\text{array}}[0] \leftarrow \theta_0$ 
4 for  $k \leftarrow 1$  to  $m$  do
5    $\theta_{\text{array}}[2k-1] \leftarrow k\pi - \theta_0$ 
6    $\theta_{\text{array}}[2k] \leftarrow k\pi + \theta_0$ 
7  $\theta_{\text{array}} \leftarrow \theta_{\text{array}} / (2m+1)$ 
8 return  $\theta_{\text{array}}$ 

```

Appendix D: Selection of results for the Iterative Amplitude Estimation

We present here a method for obtaining the most probable value of a in an iterative fashion following similar methods as other Amplitude Estimation without QPE algorithms. We base this procedure in the theory of confidence intervals for a binomial distribution assuming normal distributions [69].

Let us consider a binomial distribution with probability a , *i.e.*, for every sample the chance of obtaining 1 is a , while the chance of obtaining 0 is $1 - a$. Then, if an estimate $\hat{a}$ of a was obtained using N samples, the true value of a lies in the interval

$$a = \hat{a} \pm \frac{\text{CDF}_{\mathcal{N}}^{-1}(1 - \alpha/2) \sqrt{\hat{a}(1 - \hat{a})}}{2\sqrt{N}}, \quad (\text{D1})$$

with confidence $(1 - \alpha)$.

From this result we can construct an iterative algorithm returning the optimal value of a using quantum Amplitude Estimation. Let us take a set of m_j for $j = 0, 1, 2, 3, \dots$. For every m_j the probability of obtaining $|1\rangle$ is $\sin^2((2m_j + 1)\theta_a)$, where $a = \sin^2(\theta_a)$. Let us move to the θ space. For a given m the values and error of θ obtained are

$$\theta_a = \frac{1}{2m+1} \arcsin(\sqrt{a}) \quad \Delta\theta_a = \frac{1}{2m+1} \frac{\text{CDF}_{\mathcal{N}}^{-1}(1 - \alpha/2)}{2\sqrt{N}}. \quad (\text{D2})$$

It is important to understand two main properties of Eq. (D2). First, there are $2m+1$ possible values for θ_a within the interval $\theta_a \in [0, \pi/2]$ as the $\sin^2(\cdot)$ function is π -periodical. For every new iteration it will be necessary to choose one of them. It is very important to set $m_j = 0$ at first because this case is the only one for which θ_a corresponds to the expected value for a . Otherwise, several possible values of a arise and it is not possible to tell which one is correct. Combining results for several values of m_j , it is possible to bound the uncertainty to be as small as desired.

The algorithm is based on the following statements. For a given collection of measurements and uncertainties $\{\theta_i, \Delta\theta_i\}$, the weighted average and uncertainty from the first j terms is

$$\tilde{\theta}_j = \frac{\sum_{i=0}^j \theta_i / \Delta\theta_i^2}{\sum_{i=0}^j 1 / \Delta\theta_i^2} \quad \Delta\tilde{\theta}_j = \left(\sum_{i=0}^j 1 / \Delta\theta_i^2 \right)^{-1/2}. \quad (\text{D3})$$

Notice also that this relation is recursive, as $\tilde{\theta}_{j+1}$ can be obtained by combining $\tilde{\theta}_j$ and θ_{j+1} . The same holds for uncertainties. Thus, the interpretation of this algorithm is that for every new step j a new term is added to the series $\{\theta, \Delta\theta\}$. The individual uncertainties decrease as $\sim ((2m+1)^{-1})$, and the final global uncertainty is obtained as

$$\Delta\theta = \frac{\text{CDF}_{\mathcal{N}}^{-1}(1 - \alpha/2)}{\sqrt{N}} \left(\sum_{j=0}^J (2m_j + 1)^2 \right)^{-1/2}, \quad (\text{D4})$$

where J denotes the last iteration performed. The full recipe for the algorithm is described in Algs. 1 and 2.

In the case of a linear selection of m_j , *i.e.*, $m_j = j; j = (0, 1, 2, \dots, J)$, the asymptotic behavior of this uncertainty is $\Delta\theta = \mathcal{O}(N^{-1/2} M^{-3/4})$, with M the sum of all m . For discovering it we just have to compute

$$\sum_{j=0}^J (j+1)^2 = 4 \sum_{j=0}^J j^2 + 4 \sum_{j=0}^J j + \sum_{j=0}^J 1. \quad (\text{D5})$$

We now take the identities $\sum_{j=0}^J j = J(J+1)/2 = M$ and $\sum_{j=0}^J j^2 = J(2J+1)(J+2)/6$. Then, it is direct to check that

$$\Delta\theta = \mathcal{O}(N^{-1/2} J^{-3/2}) = \mathcal{O}(N^{-1/2} M^{-3/4}). \quad (\text{D6})$$

This behavior already surpasses the tendency of the classical sampling, but does not reach the optimal Amplitude Estimation.

In the case of an exponential selection of m_j , *i.e.*, $m_j = \{0\} \cup \{2^j\}; j = (0, 1, 2, \dots, J)$ we can take the identities $\sum_{j=0}^J 2^j = 2^J - 1 = M$ and $\sum_{j=0}^J 2^{2j} = (2^{2J} - 1)/3$. Then it is direct to check that

$$\Delta\theta = \mathcal{O}(N^{-1/2} 2^{-J}) = \mathcal{O}(N^{-1/2} M^{-1}). \quad (\text{D7})$$

1. Extension to error-mitigation techniques

The error-mitigation procedure proposed for the unary algorithm discards some of the algorithm instances to retain outcomes within the unary basis. This reduces the precision achieved in the algorithm with respect to the ones predicted in Eqs. (D6) and (D7) in order to maintain accuracy. This section provides some lower bounds on how many AE iterations can be done while still reaching quantum advantage.

We will work now in the scheme where $m_j = j$. Let us assume that, in every iteration of Amplitude Estimation, only a fraction $\tilde{p}_j$ of the shots are retained. The equivalent version of Eq. (D4) is now

$$\Delta\theta = \text{CDF}_{\mathcal{N}}^{-1}(1 - \alpha/2) \left(\sum_{j=0}^J (2m_j + 1)^2 N \tilde{p}_j \right)^{-1/2}. \quad (\text{D8})$$

As more errors are bound to occur, $\tilde{p}_j$ decreases as m_j increases, we can state a bound for the accuracy as

$$\Delta\theta \leq \frac{\text{CDF}_{\mathcal{N}}^{-1}(1 - \alpha/2)}{\sqrt{N \tilde{p}_J}} \left(\sum_{j=0}^J (2m_j + 1)^2 \right)^{-1/2}, \quad (\text{D9})$$

since the precision is at least as good as the one obtained for the worst-case scenario. Comparing the trends, both in the linear and the exponential case, with the classical scaling, it is possible to see that quantum advantage is still achieved provided

$$\tilde{p}_J \geq M^{1-2\alpha}, \quad (\text{D10})$$

Figure 21. Results of the sampling uncertainties of the expected payoff, for the unary (Left) and binary (Right) representation, with M iterations of Amplitude Estimation for both the unary and binary approaches, considering depolarizing and read-out errors together. Scattering points represent the uncertainties obtained for the experiment while dash-point lines represent theoretical bounds, where each line is accompanied with the corresponding marker. For every color and symbol, the lower bound is for optimal quantum advantage, and the upper bound is for sampling. For $M = 0$, both bounds are equivalent. In every case, a new iteration of the Amplitude Estimation reduces the uncertainty. For the unary case, the scattering points present a tendency to return larger uncertainties as the errors increase, while for the binary case the uncertainties do not explicitly depend on the single-qubit gate error. This difference is a direct consequence of post-selection.

with $\alpha = 3/4$ in the linear case and $\alpha = 1$ in the exponential case.

The probability of retaining a shot is at least the probability of having no errors in the circuit, considering that some double errors may lead to erroneous instances but in the unary basis. This zero-error probability in the worst case scenario, that is, at the last iteration of AE, is written as

$$p_0 = ((1 - p_e)^{an+b})^{m_J}, \quad (\text{D11})$$

where p_e is the error of an individual gate, and a and b are related to the gate scaling, see Tab. I for the details. In principle, one can expand the calculation of p_0 by considering different kinds of errors for different gates, but for the sake of simplicity we will focus on this analysis. Rearranging together the results for Eqs. (D6), (D7) and (D10) it is possible to see that quantum advantage is obtained if the individual gate errors is bounded by

$$p_e < 1 - m_J^{\frac{2-4\alpha}{(an+b)m_J}}. \quad (\text{D12})$$

Appendix E: Independent analysis of errors for Amplitude Estimation

In this section we present results that extend the ones depicted in Sec. V.

It is also interesting to see how the uncertainties of the measurements decrease as more iterations of Amplitude Estimations are allowed. We must remark that those errors are exclusively due to the uncertainty in the sampling. This can be observed in Fig. 21, where the obtained uncertainties are bounded between the classical sampling and the optimal Amplitude Estimation.

Furthermore, a very remarkable behavior of the uncertainties in the unary approach must be noticed. The obtained uncertainties present a tendency to increase as errors get larger. In contradistinction, the binary algorithm does not present this feature. The reason lies in the post-selection that is only applicable in the unary representation. As errors become more likely to happen, the post-selection filter rejects more instances. The direct consequence is that the number of accepted shots drops for large errors, causing less certain outcomes. The joint action of this processes is that the uncertainty decreases more slowly for the unary algorithm than for the binary one. This behaviour contrasts with the error obtained in Fig. 14, where the binary results reflect a very poor performance.

The apparently contradictory result shown in Fig. 22 is related to the distinction between *accuracy* and *precision*. *Accuracy* stands for how close is a measurement to the exact value of a quantity, and *precision* encodes the dispersion of different measurements. Amplitude Estimation is an algorithm to increase the precision of a measurement with respect to the number of samples, but it does not provide any further information regarding the accuracy. Amplitude

Figure 22. Results for the sampling uncertainties of the expected payoff for increasing number of bins for up to M iterations of Amplitude Estimation for the unary approach, considering depolarizing and read-out errors together. Scattering points represent uncertainties obtained and dash-point lines represent theoretical bounds, where each line is accompanied with the corresponding marker. For every color and symbol, the lower bound is for optimal quantum advantage, and the upper bound is for sampling.

Estimation for the binary algorithm reflects the expected tendency for the increase in precision, but comes with very poor results in accuracy. The unary algorithm grows slower in terms of precision, but maintains more accurate results. This decrease in precision might lead to losing the quantum advantage provided by AE in the presence of significant error. We study further in App. D the limit of AE iterations that can be performed given the error rates of the quantum device while still maintaining quantum advantage for the unary representation.