NExt ApplicationS of Quantum Computing

D5.7: Update of review of state-of-the-art for Pricing and Computation of VaR

Document Properties

Contract Number	951821
Contractual Deadline	30-09-2024
Dissemination Level	Public
Nature	Report
Editors	Alberto Manzano, UDC Gonzalo Ferro, CESGA
Authors	Gonzalo Ferro, CESGA Alberto Manzano, UDC Andrés Gómez, CESGA Carlos Vázquez, UDC María R. Nogueiras, HSBC
Reviewers	Vedran Dunjko, ULEI Mohamed Hibti, EDF
Date	24-07-2023
Keywords	Quantum Finance, Amplitude Amplification and Estimation, Quantum Accelerated Monte Carlo, Pricing, VaR
Status	Submitted
Release	1.0

This project has received funding from the European Union's Horizon 2020 research and innovation program under Grant Agreement No. 951821

History of Changes

Release	Date	Author, Organisation	Description of Changes
0.1	15/05/2023	Gonzalo Ferro (CESGA), Alberto Manzano (UDC), Andrés Gómez (CESGA), Carlos Vázquez (UDC)	First draft version
0.2	15/06/2023	Gonzalo Ferro (CESGA), Alberto Manzano (UDC), Andrés Gómez (CESGA), Carlos Vázquez (UDC)	Second draft version
0.3	06/07/2023	Gonzalo Ferro (CESGA), Alberto Manzano (UDC), Andrés Gómez (CESGA), Carlos Vázquez (UDC), María Nogueiras (HSBC)	Final full version for review
1.0	24/07/2023	Gonzalo Ferro (CESGA), Alberto Manzano (UDC), Andrés Gómez (CESGA), Carlos Vázquez (UDC), María Nogueiras (HSBC)	Final reviewed version

Table of Contents

1.Executive Summary

Derivatives contracts are one of the fundamental pillars of modern financial markets and are routinely traded by both financial institutions and traders with a variety of objectives, such as financial risk hedging. For this reason, the fair valuation of financial derivatives, known as pricing, and the computation of various risk measures, such as the Value at Risk (VaR), have become two of the tasks that consume a great amount of computational resources in financial institutions.

Classically, the problems of derivatives pricing and the computation of VaR are mainly solved by means of Monte Carlo-simulation (MC) techniques or numerical algorithms for solving partial differential equations (PDE). The key advantage of MC technique is that it is easy to implement, very general and scales well with the dimension of the problem. Therefore, it has become the *de facto* standard by financial institutions to tackle both problems.

In this context, NExt ApplicationS of Quantum Computing (NEASQC) Use Case 5 (UC5) works on the development and evaluation of quantum algorithms for derivatives pricing and VaR problems. Inspired by the aforementioned classical algorithms, the quantum computing community has developed different strategies to speed up the classical techniques.

This report summarises the advances on derivatives pricing and the computation of VaR since the publication of the deliverable D5.3 review. The report is divided in three main sections. The first section corresponds to the main line of research on quantum computing for pricing and VaR, namely, Quantum Accelerated Monte Carlo (QAMC). This technique promises to be as flexible and general as its classical counterpart while requiring quadratically fewer computational steps. The second section corresponds to quantum algorithms for solving Partial Differential Equations (PDEs). In this case it is possible to find some methods that promise to obtain even exponential speedups. However, such speedups can only be attained for very specific cases which makes it less attractive for practitioners. The last section is devoted to different techniques which do not fall in the previous two categories. Thus, we include recent advances on quantum machine learning techniques applied to finance, financial modeling via quantum mechanics and quantum assets.

It is important to note that some of the proposed techniques require far more quantum computational resources than are currently available. This includes limitations in the number of qubits, in the depth of the circuits and in the levels of “noise”. Hence, the current deliverable pays special focus to quantum algorithms in Noisy Intermediate-Scale Quantum (NISQ) computers.

2.Introduction

This document reviews several techniques for derivatives pricing and the computation of VaR that have been introduced in the literature since the elaboration of report D5.3, published in (Gómez et al., 2022). In quantitative finance, various strategies exist to solve the problem of pricing and VaR, each with its unique advantages and limitations. Without aiming to be exhaustive, Figure 1 schematically depicts some of the most popular ones.

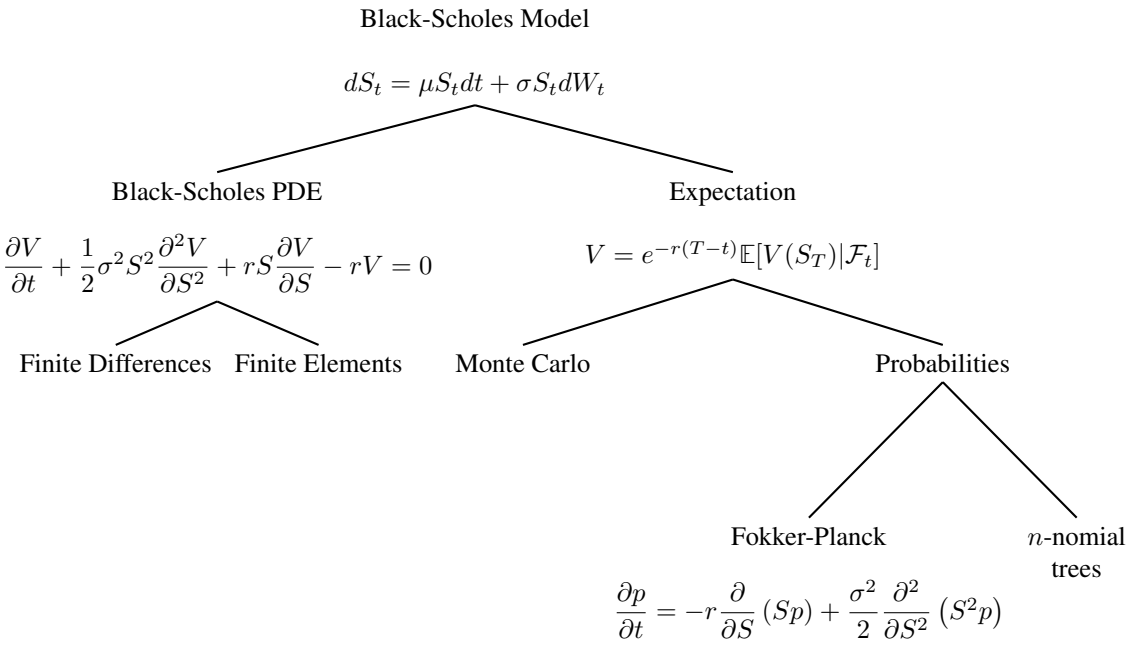

***Figure 1**: Some mathematical strategies for the pricing of financial derivatives.*

Monte Carlo techniques are the most commonly used methods in the industry for pricing and VaR. This is due to their ability to handle a variety of setups, good convergence properties in multi-dimensional problems, and simple implementation. Therefore, this report primarily focuses on quantum algorithms designed to speed up classical MC, which are generally known as QAMC techniques.

Another contender in the industry for solving pricing problems is through the direct solution of the PDE that arises from the derivatives pricing problem. Typically, Finite Differences methods are used to solve the PDE because the boundaries in pricing problems tend not to be very complex. These methods offer a faster convergence rate than MC methods for low-dimensional problems.

Lastly, a distinct research line aims to create quantum-inspired models of the financial industry, rather than attempting to accelerate classical techniques with quantum tools. This approach offers the advantage that, if the dynamics of the market dynamics followed a quantum-mechanics evolution, we could use the full potential of quantum computers to address them, rather than being constrained to speed up classical techniques. Nonetheless, there is no compelling reason to assume that the dynamics of the market follow the rules of quantum mechanics.

This manuscript is structured as follows. Section 3 provides a detailed discussion on the latest developments in QAMC, with Sections 3.2 and 3.1 delving into the advancements made in the building blocks of QAMC. Section 4 then examines the latest progress in solving pricing PDEs while Section 5 highlights quantum modelling techniques that have been applied to finance. Finally, the conclusions in Section 6 wrap up the report.

3.Quantum Accelerated Monte Carlo

In this section we will focus on the recent advances made in the field of QAMC and its applications in derivatives pricing and VaR computations with particular emphasis on their implementation on NISQ architectures. More specifically, our review will only cover the publications that have been released between the previous report and the current one.

QAMC is a promising approach that offers a quadratic speedup compared to Classical Monte Carlo (CMC) methods, making it an attractive option for solving complex financial problems. This technique was first introduced in (Montanaro, 2015). Later, in (Rebentrost et al., 2018) it was proposed for the first time to the task of pricing financial contracts. However, it was not until the publication of (Egger et al., 2020; Stamatopoulos et al., 2020), that QAMC was implemented for an extremely simplified version of pricing and VaR on a real quantum computer.

All of the techniques grouped into the category of QAMC share a common structure or pipeline:

- The construction of an oracle U which is able to load the path probability distribution.
- The application of a second unitary V which encodes the payoff of the target contract.
- The measurement of a state which encodes the solution of the problem by means of Amplitude Estimation (AE) techniques.

Since it is not possible to cover all the different combinations of such sub-techniques, in this introduction we will briefly discuss the applications of QAMC to finance while the advances in the generation of an oracle U and in the AE techniques will be discussed in Sections 3.1 and 3.2, respectively. The works that have applied QAMC to the problem of pricing and VaR can be categorised in three blocks.

The first block encompasses the advances made in pricing problems which require more complex models. In this regard we have (Tang et al., 2022) where they compute the prices of market caps using the LIBOR market model. In a technical level, it does not offer many differences from the foundational QAMC. However, it is worth mentioning that they explore the possibility of loading the probability distribution of the underlying factors in a per underlying basis and in a multidimensional way.

The second block focuses on proposals to extend the current methodology to different financial products. In this setting, there are four different proposals. In the first one (see (Tang et al., 2021)), they price collateralized debt obligations using copula models. The main difference with other works is that they improve in the techniques to generate distributions of correlated assets. The second one (see (Pracht et al., 2022)), encompasses the pricing of Asian options. Here they need to solve the optimal stopping problem and they use the quantum computer to do the heavy lifting. Nevertheless, the overall pipeline is classical. Moreover, they perform an analysis of the different sources of error and they observe that the biggest source of error does not come from the discretization of the distribution or the payoff but from the QAE, which is aligned with what is expected theoretically. The third proposal (see (Miyamoto, 2022)), does a similar job to the previous one, this time with Bermudan options. Again, they need to solve the optimal stopping problem in order to price the options and they give a theoretical discussion on how can it be done. Last, we have (Pracht, 2023) where they give a strategy to price American options on a quantum computer that differs from previous proposals in that they parametrise the optimal stopping rule as a quantum circuit. In summary, we see that there are no major breakthroughs in the field but incremental proposals. Furthermore, there is an increasing interest in the pricing of path-dependent options.

Lastly, the third block pertains to new convergence results. Starting with (Han & Rebentrost, 2022), they give a bound on the convergence of QAMC when it is applied to Credit Valuation Adjustment (CVA) problems which is tighter than the general one. Next, in (Li & Neufeld, 2023) they perform a thorough study of the different errors that affect the QAMC algorithm. This study carefully analyses how the discretization and other errors might affect the overall convergence of the QAMC. The conclusion is that taking into account all the different errors arising in QAMC, it is not possible to achieve a quadratic improvement with respect to the ideal CMC. Nonetheless, this analysis does not put under the same conditions the QAMC and the CMC. While in the former one all the different sources of error are taken into account, in the second one it is assumed that there are no other sources of error. For this reason, it is not possible to conclude that the speedup won't be quadratic in practice. Finally, in (Udvarnoki et al., 2023) they perform an empirical assessment to understand which level of noise is admissible to start getting an actual speedup in a quantum computer.

3.1. Initialization

The problem of vector initialization is usually defined as follows. Let $\vec{v} = [v_0, \ldots, v_{2^n-1}]^T$ be a vector with $v_i \in \mathrm{C}$. The task is to find the decomposition into a set of quantum gates of a unitary U_v such that:

$$|v\rangle = U_v|0\rangle = \frac{1}{\|\vec{v}\|_{\ell^2}} \sum_{i=0}^{2^n-1} v_i \, |i\rangle \,, \tag{3.1}$$

with n being the number of qubits. How $|i\rangle$ is computationally understood (an integer, float fixed point, or floating point representation) depends on the context.

A special case of the vector initialization problem is the function initialization problem, which tries to find a decomposition in a set of quantum gates of the unitary U_f which performs the following mapping:

$$|f\rangle = U_f|0\rangle = \frac{1}{\|\vec{f}\|_{\ell^2}} \sum_{i=0}^{2^n-1} f(i) \, |i\rangle \,, \tag{3.2}$$

where $f(i)$ is the evaluation of the target function f on the points i. The connection between the vector and the function initialization problems is straightforward by making v_i correspond to $f(i)$. This formulation is particularly useful for financial applications since two of the most common tasks in QAMC are the ones of loading a Probability Distribution Function (PDF) and a payoff function. In this case, the PDF or the payoff are usually defined in the real interval $[a, b]$ and $|i\rangle$ typically corresponds to the index of the grid element $x_i = \frac{(b-a)*(i+\frac{1}{2})}{2^n} + a$.

So far, the initialization problem was discussed in the context of mapping from the initial state $|0\rangle$ to any vector in the Hilbert space. However, there is a more general problem of finding the unitary U_V that maps from a family of quantum states to another family of quantum states (see (Szasz et al., 2023)). More specifically, let $V = \{\vec{x}_1, ..., \vec{x}_m\}$ and $W = \{\vec{y}_1, ..., \vec{y}_m\}$ be two sets of m orthonormal vectors each. The unitary synthesis problem is defined as finding a decomposition into a set of quantum gates of the operator U_V such that:

$$U_V|x_j\rangle = |y_j\rangle, \, \forall \, j \in \{1, ..., m\}. \tag{3.3}$$

Note that Equation (3.1) is a particular case of Equation (3.3) when $m = 1$, $|x_0\rangle = |0\rangle$ and $|y_0\rangle = U_v \, |0\rangle$. The unitary synthesis methods are interesting for the pricing and VaR problems because they could help in defining an evolution operator $U_{\Delta t}$ that, when repeatedly applied to the same register, could give the target PDF p at any given point in time $t_k = k\Delta t$:

$$\left(U_{\Delta t}\right)^k |0\rangle = \sum_{i=0}^{2^n-1} \sqrt{p(x_i; t_k)} \, |i\rangle \,. \tag{3.4}$$

However, the unitary synthesis methods are restricted by the fact that U_V is a unitary operator, which limits their applicability to price estimation to a few cases where the time step evolution preserves unitarity (Chefles et al., 2004).

Following (Iaconis et al., 2023), some interesting properties for an initialization method are:

1. being efficient, i.e., scaling polynomially in the number of qubits,
2. being robust to noise,
3. having controllable precision,
4. being deterministic, i.e, not having any post-selection, and
5. having efficient circuit construction, i.e. that the compilation of the circuit is efficient.

Moreover, since the initialization step is typically used as part of a more general routine, there is an additional requirement that is often overlooked by some of the proposals in the literature: the algorithm should recover the target relative phases in the amplitudes. For example, changing $f(i)$ to $-f(i)$ for some of the elements in (3.2) should not be allowed. This change of the sign could change significantly the result of the overall algorithm, even when the measured probabilities do not change. Also, because NISQ Quantum Processing Units (QPU) have usually a limited number of qubits, these characteristics may include also that it should not consume many auxiliary qubits.

Since the publication of the first report (Gómez et al., 2022), there are many new techniques that tackle the initialization problem in its various flavours. In the following text they are organised in eight different categories.

The first category covers specific methods for finance. In (Pracht, 2022), the authors propose a Quantum Binomial Tree method that is the quantum analogous to the classical binomial tree simulation method. Although the probability loading seems efficient, the number of required gates makes it unlikely to be implemented in computers of the NISQ era. Furthermore, since the Gaussian steps are approximated by a discrete distribution with only two values, the method is only capable of loading a coarse approximation of the target distribution.

The second category includes general methods for unitary synthesis. In (Sun et al., 2023) they explore new algorithms and their scalability when many qubits are available. Instead of using the amplitudes of binary encoded states (as it is typically done to exploit the size of the Hilbert space), the vector is encoded in the amplitudes of a unitary encoding:

$$U|0\rangle = \frac{1}{\|\vec{v}\|} \sum_{i=0}^{2^n-1} v_i \, |e_i\rangle \, , \tag{3.5}$$

where e_i is a unitary vector of dimension 2^n, where the i-th component is equal to 1 and the rest are 0. The paper by (Zhang et al., 2022) works in a similar environment, they propose an algorithm that can initialize n qubits, with a circuit depth scaling as $\mathcal{O}(n)$. In order to do so, they need to classically compute an angle associated to each leaf of the binary tree. Since this strategy only requires that each qubit is connected to other 3 qubits, it is better suited than the previous algorithm for NISQ architectures. However, the total number of gates still scales as $\mathcal{O}(2^n)$.

The third category deals with specific methods based on matrix decomposition. In (McArdle et al., 2022), they propose an algorithm based on the Quantum Singular Value Transformation (QSVT) (Gilyén et al., 2019). This method allows to load the values of any function of the eigenvalues of a Hermitian matrix into a quantum circuit when the function can be approximated by a polynomial (for example, a Taylor decomposition) or a Fourier series. In (Araujo, Blank, et al., 2021), the authors introduce two new algorithms: the *Low-rank state preparation* (LRSP) algorithm and the *Bounded Approximation Algorithm* (BAA). The first one uses a Singular Value Decomposition (SVD) of the unitary that exactly prepares the initial vector $\vec{v}$, based on the previous work of (Plesch & Brukner, 2010). The Low-rank name comes from the fact that the algorithm allows a small loss in the fidelity of the generated state. The loss in controlled by selecting the rank (r) of the SVD matrix of rank k associated with the state $|v\rangle$. This idea can be executed recursively, looking at each level the best Low-rank approximation such that the error is bounded. As the number of partitions for each number of qubits is not unique, the classical search could be exponential. An interesting result of this paper is that, on NISQ architectures, loading an approximate initial state yields a lower error than loading the exact one. The reason for this is that the proposed scheme requires a smaller number of gates, hence reducing the amount of noise, which turns out to be the main source of error in loading techniques.

The fourth category continues the seminal work of Grover (L. K. Grover, 2000), this time applied to initialisation. In (Bausch, 2022), the authors propose to load the data using a method similar to the Grover proposal that uses Amplitude Amplification (L. K. Grover, 2000; Sanders et al., 2019), but in this case the auxiliary qubits are really qudits, so they will need to be simulated by qubits in the NISQ current computers. (Ashhab, 2022) is a Grover like algorithm that uses $n(L+1)+1$ qubits to encode an approximation of the vector $\vec{x}$ of N elements into a register of $n = \log_2 N$ qubits. Here, L is the number of bits used to represent the elements x_i in binary format. Due to the need of additional qubits, this method has limited applicability in the current NISQ era. (Gleinig & Hoefler, 2022) increases the probability of measuring the correct qubit, multiplying the amplitudes by a constant. It needs to measure an auxiliary qubit, so the algorithm can not be used efficiently as a subroutine to be applied several times. (Gonzalez-Conde et al., 2023) presents two different methods for loading a function f on the amplitudes of a register. The first method is based on the Grover-Rudolph method for loading a probability distribution (L. Grover & Rudolph, 2002), although reducing the needs of controlled operations. If the fidelity can be relaxed, they propose the so called multicontrolled Grover-Rudolph method which replaces the multicontrolled rotations for a single rotation in order to reduce the number of gates. For a normal distribution, this method requires only less than 13% of the total number of gates needed by Grover-Rudolph's algorithm for achieving a fidelity of 0.99961 for 8 qubits.

The fifth category uses algorithms based on trees. In (Mozafari et al., 2022), they propose an efficient algorithm that uses $n+1$ qubits to prepare the state when $\vec{x}$ is sparse, i.e., the number of non-zero amplitudes (m) is much less than the total number of states ($m \ll N$). In this case, by using Decision Diagrams plus an additional qubit, it is possible to load the vector into the amplitudes in $O(kn)$, k being the number of paths to traverse in the Decision Diagram (DD). In (Araujo, Park, Petruccione, et al., 2021), the authors present an algorithm based on the classical divide-and-conquer paradigm that can be executed in parallel, reducing the depth of the circuit but increasing the number of needed qubits in order to store partial results. However, the final state is not the same as (3.1), being entangled with additional (possibly orthonormalized) states. Hence, the algorithm is only valid when the produced state does not affect later calculations. (Araujo, Park, Ludermir, et al., 2021) proposes to use a mixed methodology, traversing the tree in two

directions. If n is the number of levels of the tree, the idea is to select a mid point s. From this point, the algorithm is divided to upload firstly the 2^{n-s} sub-trees in parallel using as many auxiliary qubits as needed. In a second step, a botton-up method is done, to go from s to n. When N is large, by selecting $s = \log_2(N)/2$, the algorithm generates a circuit with sublinear depth and width of $O(2^{\frac{n}{2}})$.

The sixth category covers variational algorithms. In the aforementioned work of (Gonzalez-Conde et al., 2023) when the function satisfies certain requirements (for example, not having singular values), they propose to use Parametric Quantum Circuits (PQC). With their strategy they can approximate functions that take negative values (for example, $\sin(x)$) or they can approximate functions of interest on the financial sector, as the probability distributions generated by the Black-Scholes model. In (Gabor et al., 2022), the authors make a first experiment of using reinforcement learning methods to train a circuit to make a specific initial state or a general case. By using a simulator, they made the experiment in a system of 2 qubits, reaching a modest fidelity and a minimum number of gates of 15 although remaining unclear how to address the problem for a higher number of qubits. (Creevey et al., 2023) uses a similar approach to (Rindell et al., 2022) using genetic algorithms. Results are similar to the previous case: the fidelity of the generated state in a noisy QPU is better than the Qiskit synthetic algorithm for a low number of qubits. But still, the fidelity is far from the needs, decreasing very fast when the number of qubits increases. In (Dasgupta & Paine, 2022), they propose to use a variational algorithm to load probability distributions in the amplitudes of a quantum state, where the structure of the circuit comes from an analysis of the structure of the distributions and the asymmetry they have. They tested the method up to 6 qubits for different probability distributions with results strongly dependent on the connectivity of the quantum processors. In (Agliardi & Prati, 2022a, 2022b), the authors study the difficulty of training quantum Generative Adversarial Network (qGAN) for probability distribution loading. They found that the generated system blocks frequently in local minima, needing several trainings for achieving a good result. So, they conclude that it is better to train several times starting from different points and select the best result than making a single detailed training. In fact, they assess that this process is more important than the selection of the optimizer or their hyper-parameters. (Rindell et al., 2022) uses the algorithms of (Araujo, Blank, et al., 2021) to explore the possibility of finding optimal quantum circuits to load approximated states with genetic algorithms. For 5 qubits, they found that there are circuits that perform better than the LRSP-like ones, and also that there exists an optimal circuit that gives the best fidelity. Although still far from the needs, they show maximum fidelities lower than 0.8) with a lower number of CNOTs. (Chang et al., 2023) proposes the Split Step Quantum Walk (SSQW) to load probability distributions (or real+ valuated amplitudes) into the state using a modification of the quantum walk, using a parameterized coin operator. In order to load a specific distribution, a variational algorithm is used to choose the right parameters. So, the algorithm is a mix between the quantum walk and the pure variational algorithm. They apply this method to the European call option, obtaining a reasonable result with only 5 qubits and in less than 20 seconds.

The seventh category leverages the power of Matrix Product States (MPS). The works (Iaconis et al., 2023) and (Johri et al., 2023) are special experimental cases of (Ran, 2020) commented in the previous report D5.3, which presented an iterative method to approximate a MPS circuit using low-depth circuits. As the entanglement of states $|f\rangle$ increases slowly with the number of qubits when the function is smooth, some smooth probability distributions can be approximated using MPS methods. They show that they achieve only valid results up to 20 qubits with circuits with one MPS layer, when more than one is actually needed to achieve a good approximation.

The last category includes different methods to load general functions. In (Moosa et al., 2023), the authors use the efficiency of the Quantum Fourier Transform (QFT) to load the values of periodic and non periodic functions in the interval $[0, 1]$ into the amplitudes. The method is based on four steps, with the first one requiring a classical computational cost of $O(n2^n)$. However, as only dominant Fourier terms are needed, sparse techniques can be applied to decrease this cost. Another point to decrease the quantum circuit complexity is to use the approximated QFT. Experiments using Quantinuum H1-1 and H1-2 quantum computers give qualitative good results (they do not provide a metric to compare the theoretical and experimental distributions, only a picture).

Among the papers analysed in this report plus the ones that were analysed in the previous one (Gómez et al., 2022), there are no algorithms that fulfill all the requirements at once. Some of the algorithms require an excessively large number of auxiliary qubits or an excessive number of steps (i.e., depth of the circuit) which makes them unpractical with current processors. Under this constraints of the NISQ era, variational algorithms have become a very appealing option. Nonetheless, some of them have the risk of adding arbitrary phases into the amplitudes, which are difficult to detect when this strategy is scaled up in the number of qubits, as it would require to do a tomography of the state which becomes computationally prohibitive. Furthermore, variational algorithms require a training phase which can become extremely difficult because it is very easy to get stuck in local minima or barren plateaus. Despite all the difficulties, one of the most interesting features is the possibility of adjusting the desired precision in the final PDF. This adjustment can be very advantageous in the NISQ era where the depth of the circuit is limited and there are very big levels of noise. As a consequence, the relaxation of the precision usually achieves better final results than exact

methods, because the overall effect of noise is partially mitigated.

3.2. Amplitude Estimation

Typically, after applying one of the initialization routines from Section 3.1 plus some additional steps, the solution of the financial problem is encoded in a quantum state of the form:

$$|\Psi\rangle = \mathcal{A}\,|0\rangle^{\otimes n} = \sqrt{a}\,|\psi_1\rangle + \sqrt{1-a}\,|\psi_0\rangle\,, \tag{3.6}$$

where $\mathcal{A}$ is the unitary operator that encodes the solution of the financial problem in a quantum state, $|0\rangle^{\otimes n}$ is n-qubit state, $|\psi_1\rangle$ and $|\psi_0\rangle$ are two orthonormal states and a is proportional to the solution of the financial problem at hand. Now, the task becomes estimating the value of a from the quantum circuit. More precisely, given a precision level $\epsilon_a > 0$ and a confidence level $1-\delta$, with $\delta \in (0,1)$, the goal is to find an estimator $\tilde{a}$ of the probability $a \in [0,1]$, such that the probability of a being within an interval of length $2\epsilon_a$ centred around $\tilde{a}$ is greater than $1-\delta$, or equivalently:

$$\mathbb{P}\left[|a-\tilde{a}| \geq \epsilon_a\right] \leq \delta. \tag{3.7}$$

Classically, this problem can be tackled with $N_{\mathcal{A}}$ calls to the oracle $\mathcal{A}$:

$$N_{\mathcal{A}}^{\mathrm{MC}} \geq \frac{1}{2\epsilon_a^2}\log\left(\frac{2}{\delta}\right). \tag{3.8}$$

However, the same task can be achieved using a lower number of calls to the oracle $\mathcal{A}$ by means of a wide range quantum computing techniques grouped under the name of Amplitude Estimation (AE). Although there are many variants, they all share the same core: they leverage the power of a Grover-like operator $\mathcal{G}_{\mathcal{A}}$:

$$\mathcal{G}_{\mathcal{A}}^{k}\,|\Psi\rangle = \sin\left((2k+1)\theta\right)|\psi_1\rangle + \cos\left((2k+1)\theta\right)|\psi_1\rangle\,, \tag{3.9}$$

where $\sin(\theta) = \sqrt{a}$. This section summarises the main contributions in the literature since the previous report (Gómez et al., 2022). Here, three main lines of research are distinguished:

- Parallel algorithms: in parallel algorithms the estimations from different independent executions of the Grover oracle $\mathcal{G}_{\mathcal{A}}^{m_k}\,|\Psi\rangle$ are combined in a clever fashion in order to obtain a better estimate of the probability a. More specifically, these types of algorithms are characterised by two main ingredients. On the one hand, a "schedule" i.e. a set of ordered pairs $\{(m_1,n_1),...,(m_K,n_K)\}$, where m_k specifies the number of applications of the Grover operator in the k-th experiment and n_k the number of shots in the same experiments. On the other hand, an aggregation rule that consistently combines the results of the different executions. The beginning of this line of research can be traced back to the Maximum Likelihood Amplitude Estimation (MLAE) algorithm, which was introduced in (Suzuki et al., 2020). In this paper, the authors proposed two different schedules: the Linear Increasing Schedule (LIS), where m_k is proportional to k, and the Exponential Increasing Schedule (EIS), where m_k is proportional to 2^k. In both cases, the number of shots n_k is kept constant. Later in (Giurgica-Tiron et al., 2022) the authors developed the QoPrime AE algorithm which is the first variant of MLAE to prove rigorous bounds for the maximum number of calls to the oracle $N_{\mathcal{A}}$ in order to get a certain precision ϵ_a with confidence $1-\gamma$. Section 3.2.1 gives an overview of some of the most relevant algorithms that work in this way since the publication of (Gómez et al., 2022).

- Iterative algorithms: in iterative algorithms the information of each step determines how the next step is going to be performed. On each iteration, the circuit $\mathcal{G}_{\mathcal{A}}^{m_k}\,|\Psi\rangle$ is executed n_k times. From the collected measurements, upper and lower bounds for the target amplitude a are computed. In the next iteration, the number of amplifications m_{k+1} and the number of shots n_{k+1} are chosen in order to reduce the difference between the upper and lower bounds. The most popular algorithm of this style is the Iterative Quantum Amplitude Estimation (IQAE) (Grinko et al., 2021). One of the most relevant features of this algorithm is that it has a proof of convergence. Section 3.2.2 gives an overview of some of the most relevant algorithms that work in this way since the publication of (Gómez et al., 2022).

- Generalized Grover: since the cornerstone in the AE algorithms is the Grover operator, there are works which try to generalize this operator. To provide the reader with some context, in (Wang et al., 2021) they replace the original reflections of the Grover operator $\mathcal{G}_{\mathcal{A}}$ with parametrized operators $Q_{\mathcal{A}}(\vec{x})$ referred to as "generalized reflections", where $\vec{x}$ is a vector of tunable parameters. These tunable operators induce engineered likelihood

functions (ELFs) in opposition with the Chebyshev likelihood function (CLF) ($\vec{x} = \left(\frac{\pi}{2}, \frac{\pi}{2}, \cdots, \frac{\pi}{2}\right)$) used in other AE approaches like MLAE or IQAE. For a fixed circuit depth L, their results demonstrate that the ELF likelihood functions combined with a Bayesian inference strategy produces better results compared to the conventional likelihood functions (CLF) (resulting from conventional reflections) even in the presence of noise. Section 3.2.3 gives an overview of some of the most relevant algorithms in this direction.

The main advantage of the sequential algorithms over the parallel ones is that in general they offer better theoretical guarantees of performance. The reader can find a comparison of the theoretical performance measured in terms of number of calls to the oracle $N_{\mathcal{A}}$ of the different algorithms presented in this document in Table 1[1]. Nevertheless, parallel algorithms have, in general, more flexibility to control the depth of the circuit, thus making them more suitable for NISQ devices. For a direct comparison of the performance of MLAE and IQAE in simulations, the reader is referd to (Yu et al., 2020) where they show that both approaches offer the same performance in accuracy.

Algorithm	Performance	Reference
Monte Carlo	$N_{\mathcal{A}}^{MC} = \mathcal{O}\left(\frac{1}{\epsilon_a^2}\right)$	
QPE	$N_{\mathcal{A}}^{QPE} = \mathcal{O}\left(\frac{1}{\epsilon_a}\right)$	(Brassard et al., 2002)
MLAE-LIS	$N_{\mathcal{A}}^{LIS} = \mathcal{O}\left(\frac{1}{\epsilon_a^{\frac{4}{3}}}\right)$	(Suzuki et al., 2020)
MLAE-EIS	$N_{\mathcal{A}}^{EIS} = \mathcal{O}\left(\frac{1}{\epsilon_a}\right)$	(Suzuki et al., 2020)
PLAE	$N_{\mathcal{A}}^{PLAE} = \mathcal{O}\left(\frac{1}{\epsilon_a^{1+\beta}}\right), d = \mathcal{O}\left(\frac{1}{\epsilon_a^{1-\beta}}\right)$	(Giurgica-Tiron et al., 2022)
Improved MLAE	$N_{\mathcal{A}}^{imp\ EIS} = \mathcal{O}\left(\frac{1}{\epsilon_a}\frac{1}{d}\log(\frac{1}{\delta})\right), d = 2^{q-2}$	(Callison & Browne, 2022)
IQAE	$N_{\mathcal{A}}^{IQAE} < \frac{50}{\epsilon_a}\log\left(\frac{2}{\delta}\log_2\frac{\pi}{4\epsilon_a}\right)$	(Grinko et al., 2021)
Modified IQAE	$N_{\mathcal{A}}^{mIQAE} < \frac{123}{\epsilon_a}\log\frac{6}{\delta}$	(Fukuzawa et al., 2023)
QCoin	$N_{\mathcal{A}}^{QCoin} = \mathcal{O}\left(\frac{1}{a}\frac{1}{\epsilon_a}\log\frac{1}{\delta}\right), k \geq 2, 1 \geq q \geq (k-1)$	(Abrams & Williams, 1999)
QoPrime	$N_{\mathcal{A}}^{QoPrime} < C\lceil\frac{k}{q}\rceil\frac{1}{\epsilon_a^{1+q/k}}\log\left(\frac{4}{\delta}\lceil\frac{k}{q}\rceil\right), d = \mathcal{O}\left(\frac{1}{\epsilon_a^{1-q/k}}\right)$	(Giurgica-Tiron et al., 2022)
FasterAE	$N_{\mathcal{A}}^{fasterAE} < \frac{4.1\cdot 10^3}{\epsilon_a}\log\left(\frac{4}{\delta}\log_2\left(\frac{2\pi}{3\epsilon_a}\right)\right)$	(Nakaji, 2020)
AdaptiveAE	$N_{\mathcal{A}}^{adaptiveAE} < \frac{1}{\epsilon_a}\log\left(\frac{\pi^2(T+1)}{3\delta}\right), T = \lceil\frac{\log\frac{\pi}{K\epsilon_a}}{\log K}\rceil$	(Zhao et al., 2022)
RQAE	$N_{\mathcal{A}}^{RQAE} < \frac{70\pi}{6\arcsin(\epsilon_a)+1}\log\left[\frac{3.3}{\delta}\log_2\left(\frac{\pi}{2\arcsin(2\epsilon_a)}\right)\right]$	(Manzano et al., 2022)

***Table 1**: Performance of the different AE algorithms presented. $N_{\mathcal{A}}$ denotes the number of calls (or queries) to the operator $\mathcal{A}$, ϵ_a is the target precision and $1-\delta$ is the confidence level. Other parameters appearing in the table are related to each specific algorithm. For a full description of their meaning the reader is refered to the associated references. The $\sim$ symbol indicates that the algorithm has an asymptotic behaviour, while the $<$ indicates that the performance is proved rigorously.*

3.2.1. Parallel algorithms

Since the report (Gómez et al., 2022), the seminal work of (Suzuki et al., 2020) has produced four main lines of research. The first one studies the performance of MLAE and its variants when noise is taken into account. The second one addresses the impact of limited depth on the amplitude estimation algorithms. The third one deals with particular setups where MLAE fails to work. The last line is an extension of the original algorithm to estimate multiple parameters.

The first line of research studies the performance of MLAE under different noise models. For example, in (Herbert et al., 2021), they study the impact of a noise of the form $\mathcal{G}_{\mathcal{A}}^k(\theta_e)|\Psi\rangle = \sin((2k+1)\theta+\theta_e)|\psi_1\rangle + \cos((2k+1)\theta+\theta_e)|\psi_1\rangle$, where θ_e follows a Gaussian distribution. The authors conduct experiments on four IBM quantum computers and Honeywell's H1 trapped ion device and demonstrate that their model of noise fits the experimental data. In addition, once the noise is properly modelled, they suggest increasing the number of shots n_k on each experiment to achieve a variance similar to that of the noiseless case. Moreover, they obtain results that better resemble the expected result in the noiseless case. In a smiliar spirit, in (Brown et al., 2020), MLAE is analyzed under 3 different noisy channels: depolarising, dephasing and amplitude damping, and using several different schedules:

[1] Here it is considered equivalent a call to the oracle or a call to its adjoint.

linear, quadratic, cubic and exponential ones. After simulating several noise strengths in the three noise channels for the enumerated schedules, they show that even in low-strength noise, the quadratic speedup for the exponential schedule is not achievable. Furthermore, they illustrate that, in general, the linear schedule shows the best performance. Moreover, they claim that a linear schedule can obtain a speed up over the CMC one if the noise level per $CNOT$ is approximately $\sim 10^{-5}$ (three orders of magnitude improvement over the best currently available hardware). Also dealing with noise, we have the work (Tanaka et al., 2021), where the one-parameter estimation problem from MLAE is converted into a two-parameter estimation problem by adding an additional parameter related to the depolarising noise channel. Under this model, they compute the Fisher information matrix and derive a lower limit for the error of the estimator that depends on the two parameters. Under their noise model, they show that there is a maximum number of oracle calls $N_{\mathcal{A}}$, beyond which the precision ϵ_a does not improve (even in low-strength noisy scenarios). In addition, they conduct a series of experiments on a real quantum computing device (*ibm_valencia* device from IBM Quantum Systems) obtaining good agreement with their noise model. Finally, in (Uno et al., 2021), the authors compare the amplification operators $\mathcal{G}$ and $\mathcal{Q}$ in both noiseless and depolarizing noise channel scenarios. They analyze these operators using the framework of Fisher information and the Cramér-Rao inequality. Based on the results they present, the authors claim that the $\mathcal{Q}$ amplification operators are superior to the $\mathcal{G}$ operators for actual quantum computers as they would exhibit better performance in the presence of noise, providing lower error rates.

The second line of research focuses on reducing the depths of the circuits. This is important since deeper circuits are more sensible to error. An example of the effect of the depth of the circuit on AE is given in (Certo et al., 2022), where the authors extensively benchmark the MLAE algorithm on the IBM quantum computer and show that, for a two-qubit circuit and beyond seven applications of the Grover operator, the error obtained with the MLAE algorithm does not decrease further. Moreover, by increasing the number of qubits to four and applying three times the Grover operator, the MLAE algorithm exhibits significantly higher error compared to the CMC approach. In order to decrease the depths of the MLAE in (Plekhanov et al., 2022), they propose a variational strategy that approximates the circuit $\mathcal{G}^k\mathcal{A}\,|0\rangle \approx |\phi_{VAR}(\tilde{\lambda})\rangle$ using a PQC. Nevertheless, this approach is very computationally demanding due to the optimization of the PQC. Thus, they propose an adaptive version of Variational Quantum Amplitude Estimation (VQAE) that requires fewer parameters for the variational circuit, and simulations show an improvement in performance compared to the Monte Carlo approach. In the same direction, we have the first part of (Callison & Browne, 2022), where the authors introduce a modification of the EIS that limits the maximum depth in exchange of performance.

The third line of research deals with anomalies that arise in MLAE-like algorithms from their very foundations. In the second part of (Callison & Browne, 2022), they show that there exist values of a where the desired target precision cannot be achieved because the Bernstein-von Mises theorem cannot be satisfied. To mitigate this effect, the authors propose a heuristic method called "depth jittering" scheduler which empirically solves the problem. In (Lu & Lin, 2023), the authors investigate the problem of critical points or exceptional values of a in MLAE for different m applications of the Grover operator. They demonstrate that these critical points of order m lead to biased estimators in MLAE, and the error cannot be reduced beyond a certain degree, even in the absence of noise. In order to mitigate the impact of these critical points, the authors propose a new algorithm called Random Quantum Amplitude Estimation (RandomQAE), which combines the results from the Grover operator with the results from a second amplification operator $\mathcal{Q}(\mathcal{A})$, thus solving the problem of the critical points.

The last line that we want to address is developed in (Tanaka et al., 2022), where the authors present a generalization of MLAE for multi-parameter estimation. They introduce a model with $M+1$ parameters, denoted as $\theta, \beta_1, \beta_2, \cdots, \beta_M$. Here, β_k represents the noise parameter associated with the k-th application of the Grover operator, specifically $G^{m_k}A\,|0\rangle$. By using the theory of nuisance parameters and the parameter orthogonalization method, the authors establish a mapping between the original parameters and a new set of parameters, where the Fisher information matrix elements are zero for all parameters except for θ. This mapping is derived through an equation provided in the paper. In addition to the standard amplification circuits used in MLAE, the authors introduce auxiliary circuits to avoid degeneracy in the Fisher information matrix. Their model enables the reproduction of a noise-depolarizing channel. They validate their approach by testing it on an actual IBM quantum computer, where they observe a reduction in the parameter estimation error with an increasing number of oracle calls. However, the error behavior follows that of the Monte Carlo case and remains far from achieving a quadratic speedup.

3.2.2. Sequential algorithms

Since the publication of (Gómez et al., 2022) there have been three main lines of research in sequential algorithms. The first one tries to produce tighter bounds on the performance measured in the number of calls to the oracle $N_{\mathcal{A}}$ of the IQAE algorithm. The second one includes different algorithms that claim to obtain better results than IQAE. The third line leverages amplitude estimation to compute not only the amplitude a but also retrieve some information

about the phase.

In the first line of research we include the work (Fukuzawa et al., 2023) where they reduce the confidence $1-\gamma_k$ at each iteration to provide tighter bounds on the performance $N_{\mathcal{A}}$ of the IQAE algorithm. The improvement is non-trivial since they are able to remove logarithmic factors with respect to the desired precision ϵ_a.

In the second line of research we have different contributions. For example, in (Nakaji, 2020) they introduce a new AE algorithm called Faster Amplitude Estimation (FAE) where they use two different amplification circuits on each iteration to avoid ambiguities in the determination of the amplitude a. The authors claim that with their strategy they are capable of reducing the number of calls to the oracle $N_{\mathcal{A}}$ by a factor of 100 with respect to the original IQAE. The other work that is included in this review is the one by (Zhao et al., 2022), where they introduce the Adaptive Quantum Amplitude Estimation (AdaptQAE) algorithm. This algorithm works by introducing a new operator $\mathcal{A}'$ and its corresponding Grover operator $\mathcal{G}'$ on each step. With this strategy they show that AdaptQAE is approximately 40 times faster on average than IQAE and MLAE, although the number of calls to the oracle $N_{\mathcal{A}}$ is slightly higher than that in the other two.

The last line of research is presented in (Manzano et al., 2022) where the authors aim to estimate $\sqrt{a}$ with a sign, i.e. $\sqrt{a}e^{i\phi}$ with $\phi \in \{0, \pi\}$. For this purpose, they developed the Real Quantum Amplitude Estimation, RQAE algorithm which uses a shifted version $\mathcal{A}_b$ of the original operator $\mathcal{A}$ in order to estimate both the amplitude and the sign. Moreover, the paper provides a rigorous analysis of the algorithm and its performance in terms of number of calls to the oracle $N_{\mathcal{A}}$ showing that is on a par with the rest of algorithms in the literature.

3.2.3. Generalised Grover

Since the seminal work of (Wang et al., 2021) that first proposed the idea of using generalized Grover reflections is very recent and there are no differentiated lines of research among all the papers which exploit this idea, they will all be discussed altogether.

In (Koh et al., 2022) an exhaustive characterization of the ELFs is provided. Additionally, explicit expansions of the ELFs, and their associated Fisher information, in Taylor and Fourier series are computed. With these expansions, they can study the ELFs and obtain posterior distributions analytically. Additionally, they claim that other prior distributions than Gaussian can be used in the Bayesian inference procedure. In (Dalal & Katabarwa, 2022) they derive an expression for CLF with different sources of noise from the ELF provided in (Wang et al., 2021). This model successfully handles incoherent noises such as amplitude damping or dephasing, although it fails when dealing with coherent noises. In order to address these types of errors, the authors propose the use of a randomized compiling (RC) technique and show in different simulations that for moderate coherent noises, their proposal yields estimators with lower bias and root mean square error (RMSE) compared to the Monte Carlo approach. In (Rall & Fuller, 2022), the concept of *generalized reflections* is utilized to develop a new framework called quantum signal processing (QSP). In this framework, they use alternative applications of very specific reflections to produce estimators $\tilde{a}$ of a by sampling over special types of polynomials called *Pellian* and *semi-Pellian*. With their framework, they recover an improved version of IQAE and devise two AE estimation algorithms; one which produces an unbiased estimator $\tilde{a}$ and one which controls the maximum depth in exchange with an increasing number of queries. Moreover, given an AE algorithm which produces an estimate $\tilde{a}$, using their framework they can construct another algorithm that yields the same estimator $\tilde{a}$ as the original one.

4.Quantum algorithms for solving PDEs models

In this section some recent proposals to use quantum algorithms to solve PDE models for pricing financial derivatives are analyzed. This analysis only covers the new papers since the publication of (Gómez et al., 2022).

In the recent article (Radha, 2021), the authors propose a quantum formulation of the classical option pricing setting based on PDEs. For this purpose, they start by arguing that extensions of classical Monte Carlo to quantum settings are mainly based on quantum amplitude estimation algorithms, which is not so practical for current NISQ processors, although some recent advances have been recently done. More precisely, after some transformations in the Black-Scholes PDE, the option price can be obtained from the evolution of an initial quantum state according to an appropriate Hermitian Hamiltonian after a Wick's imaginary time rotation. Therefore, the proposed quantum algorithm contains two steps: 1) encoding the initial condition into a quantum state and 2) evolve this state according to the Hamiltonian that corresponds to the pricing method in imaginary time. Moreover, a suitable encoding of the initial state for NISQ circuits is based on a parameterized quantum circuit, while the imaginary time evolution associated to the Hamiltonian is based on a variational Hilbert space evolution algorithm. Numerical results are described for a European vanilla call option and the different sources of error are briefly described, also pointing some possibilities to reduce each one of them.

More recently, in (Miyamoto & Kubo, 2022), the authors address the pricing of multiasset derivatives by means of a finite differences numerical solution of the high dimensional PDE. Basically, the time discretization leads to a system of ODEs, which is solved by a quantum algorithm proposed in (Berry et al., 2017). However, the extraction of the derivative price associated to the spot prices of the asset from the quantum state vector would require to obtain the amplitude of a specific computational basis state in that quantum state. As an exponential number of grid points in the high dimensions leads to exponentially small amplitudes, it requires large time complexity to read it, thus loosing the potential quantum advantage. In order to overcome this drawback, the authors propose to compute the quantum state in a time instant after the pricing time and recover the particular value of the price at initial time in terms of an expectation. This strategy requires the appropriate determination of this future time, the generation of all prices and associated probabilities at that future time, as well as estimating the price at initial time in terms of expectation (which is obtained as an inner product of prices and probabilities at future times). So, the complexity of the method involves a factor $\text{poly}(\epsilon^{-1}, d)$ instead of $((\epsilon^{-1})^{\mathcal{O}(d)})$, with d being the dimension and ϵ the precision. The article analyzes the case of the European barrier multi-asset option. For a particular example, the time complexity of the proposed algorithm is computed in terms of the number of Toffoli gates.

Next, in (Kubo et al., 2022) the authors aim to make it possible to run the algorithm proposed in (Miyamoto & Kubo, 2022) on a small quantum computer. For this purpose, mainly a variational quantum algorithm is proposed. More precisely, there are three parts in the procedure: 1) embedding the probability distribution of the asset prices into the quantum state, 2) solving the PDE and 3) compute the inner product to get the expectation. For the first part, a Fokker Plank equation is solved by a quantum variational simulation, for solving the discretized PDE with finite differences a quantum variational formulation is used and the square of the inner product is obtained by a SWAP test. Finally, the algorithm requires quantum circuits with $\mathcal{O}(\text{poly}(d \log(\epsilon^{-1})))$ few-qubit gates and can be implemented in a small quantum computer. An example on a double knock-out barrier option on a single asset is shown, for which an analytical formula is available.

5.Other approaches

In this section, alternative approaches to those addressed in previous sections are discussed. One of them connects pricing PDEs with quantum mechanics via Schrödinger's equation. Another one addresses the use of quantum reinforcement learning in finance for portfolio hedging. Also an approach based on martingale asset pricing for incomplete markets is reported. The last line of research corresponds to the original exploration of hypothetical quantum financial products and the corresponding markets. As for the previous sections, the list of discussed papers does not attempt to be exhaustive and probably does not include all the recent and not so recent interesting lines of research arising in the literature that could be classified in this section.

The article (Baaquie, 2019) is a continuation of the ideas present in the book (Baaquie, 2020) on the use of the quantum oscillator for pricing risky corporate coupon bonds. This paper is not about quantum computing for solving classical pricing models but a proposal of quantum models for option pricing. More precisely, a Merton-Shrodinger equation with a Merton Hamiltonian is proposed as a more general framework than the Black-Scholes PDE. Two choices are proposed for the Merton Hamiltonian. In order to satisfy the put-call parity (which is a consequence of no arbitrage assumption), the associated kernel of the Hamiltonian needs to satisfy two conditions: martingale condition and correct discounting. So a first choice does guarantees the martingale condition but not discounting, while in the second option the reverse holds. In both case, the call option price is computed. A pending analysis is the calibration with market data to check if the proposed models fit better the market prices than the classical Black-Scholes model, although they violate the put-call parity, which can be related to the presence of arbitrage opportunities. Note that the classical model can be obtained as a limit case of both models.

Another approach is related to the use of deep and reinforcement learning techniques for portfolio hedging, which referred to as deep hedging. More precisely, deep hedging techniques are initially developed in (Buehler et al., 2019), which aims to hedge a portfolio of derivatives in the presence of transaction costs, liquidity constraints or risk thresholds by means of reinforcement learning techniques. Moreover, a numerical example based on Heston model for the hedging under transaction costs is presented. In (El Amine Cherrat, 2023), the authors develop quantum reinforcement learning algorithms for the so called deep hedging techniques, the classical version of which is mainly developed in (Buehler et al., 2019). For this purpose, based on the previous concept of deep hedging, in (El Amine Cherrat, 2023) quantum reinforcement learning techniques are applied by using orthogonal and compound layers. Moreover, a novel distributional actor-critic algorithm that takes advantage of the fact that quantum states and operations deal with large distributions. The quantum implementation reduces the number of required trained parameters and provides better performance that other alternative classical and quantum alternatives. A trapped ion-quantum processor is used and circuits with up to 16 qubits are considered. A major limitation comes from the number of qubits, as in the simplest case we need 1 qubit per time step to simulate the evolution of the asset prices. Another one is that the implementation corresponds to a rough approximation of the Black-Scholes model, i.e. a Geometric Brownian Motion (GBM) for the price evolution.

In (Rebentrost et al., 2022) the authors propose the use of three algorithms different from quantum Monte Carlo to price derivatives contracts on assets. For this purpose, they assume that they can have quantum access to the asset prices and to the payoff matrix in a single period market. Under this assumptions, the proposed algorithms can extract the martingale risk neutral measure from market variables and estimate the value of the derivative. After posing the risk neutral pricing as a linear program problem, a quantum zero-sum game algorithm which speeds up the classical one is proposed to obtain the price with a relative error. The approach is applied to a Black-Scholes model for pricing call options by adding a regularization technique. Moreover, in the Appendix of (Rebentrost et al., 2022), a quantum simplex algorithm is used for solving the linear programming problem with a certain relative error, which is compared with the quantum zero-sum game one. On the other hand, in order to efficiently use quantum linear systems algorithms, the concept of least-square market is introduced which is weaker than complete markets. This concept allows the use of pseudoinverse matrix of the payoff that obtains the minimum l_2-norm of the residual of the system, so that a positive solution of the system and every positive payoff is replicable even when the market is not complete. In this setting, quantum pseudoinverse algrithms can be applied.

Finally, the recent article (Bao & Rebentrost, 2022) presents a very original problem setting and approach. There, an hypothetical market and assets exhibiting quantum properties is considered. For this purpose, the article starts describing a possible scenario where an institution acts as the hypothetical quantum market maker with access to a quantum computer. After describing the market scenario, a finite sample space is replaced by Hilbert bases, so that each event is understood as a vector of this basis. Next, the probability measure appearing in classical finance is replaced by a density operator, so that the expected value of quantum assets can be defined. Later, the analogous tools used in classical finance are introduced in quantum markets: arbitrage theory, change of quantum measure and

the first fundamental theorem of the pricing of quantum assets. Finally, the derivatives of quantum assets are defined as positive semidefinite Hermitian matrices with a certain payoff. Valuation of derivatives consist in obtaining an arbitrage free price which is given by the trace associated to the risk free density operator and the derivative and whose calculus requires the use of quantum computers.

6.Conclusions

The aim of this document is to provide an updated review of the advances on quantum computing for derivatives pricing and VaR since the publication of our previous review (Gómez et al., 2022). For a more complete overview on the topic, the reader is highly encouraged to consider together both (Gómez et al., 2022) and the current document.

The main conclusion is that there are no new major breakthroughs since (Gómez et al., 2022). More specifically, the vast majority of the reviewed documents are a continuation of previous works where the performance is improved up to a constant factor, with the exception of a very few papers which do not change the game either. Note that this does not prevent the field from having valuable contributions, it just means that there are lines of research that are theoretically well established and that have been consolidated over time. As a result, most of the work that has been recently done is in the direction of closing the gap between the theory, where fault tolerant quantum computers are assumed to be available, and the practice, where only NISQ hardware is available. Those are good news for the financial use cases, since the interests of this project are well aligned with those of the academia.

In the context of developing algorithms which are better suited for NISQ, there is an increasing concern about the role of noise in quantum computation in contrast with number of qubits and depth. More and more papers which perform experiments in real hardware reported that noise was the main limiting factor found in the experiments. Moreover, researchers are not only concerned about noise with the current hardware limitations but also about the potential effects that noise will have with a more advanced technology. Even if quantum computers with quantum correcting codes are available, the implementation of such error correcting codes makes quadratic speedups ineffective since the instances where an speedup is realizable are impractically large.

The issues raised by the presence of noise in quantum computer directly impact this use case. The reason for this is twofold. On the one hand, on the short term the high levels of noise prevents to implement efficient algorithms in NISQ technology. On the other hand, in the long run the quadratic speedups promised in quantum computing for option pricing and the computation of VaR will be obliterated by the error correcting codes. For this reason, the different stakeholders from the industry and academia should moderate the expectations of what can be achieved with quantum computers for realistic financial applications. Nevertheless, this does not mean that quantum computing will not be able to bring a valuable contribution in the financial field. In the near future, it means that the improvements that can be devised may come from a relevant decrease in energy consumption and not from a speed-wise improvement. In the more distant future (mid and long term), it is extremely difficult to forecast what will come up and the promising directions that this field will follow.

List of Acronyms

Term	Definition
AdaptQAE	Adaptive Quantum Amplitude Estimation
AE	Amplitude Estimation
BAA	Bounded Approximation Algorithm
CLF	Conventional Likelihood Function
CMC	Classical Monte Carlo
CQPEAE	Amplitude Estimation based on Classical Quantum Phase Estimation
CVA	Credit Valuation Adjustment
DD	Decision Diagram
EIS	Exponential Increasing Schedule
ELF	Engineered Likelihood Function
FAE	Faster Amplitude Estimation
FT3	FinisTerrae III
GBM	Geometric Brownian Motion
IQPEAE	Amplitude Estimation based on Iterative Quantum Phase Estimation
IQAE	Iterative Quantum Amplitude Estimation
LIBOR	London Interbank Offered Rate
LIS	Linear Increasing Schedule
LRSP	Low-Rank State Preparation
MC	Monte Carlo
MLAE	Maximum Likelihood Amplitude Estimation
MPS	Matrix Product States
NEASQC	NExt ApplicationS of Quantum Computing
NISQ	Noisy Intermediate-Scale Quantum
PDE	Partial Differential Equation
PDF	Probability Distribution Function
PE	Phase Estimation
PQC	Parametric Quantum Circuit
QAMC	Quantum Accelerated Monte Carlo
QFT	Quantum Fourier Transformation
qGAN	Quantum Generative Adversarial Network
QLM	Quantum Learning Machine
QPU	Quantum Processing Unit
QQuantLib	Quantum Quantitative Finance Library
QSP	Quantum Signal Processing
QSVT	Quantum Singular Value Transformation
RandomQA	Random Quantum Amplitude Estimation
RC	Randomised Compiling
RMSE	Root Mean Square Error
RQAE	Real Quantum Amplitude Estimation
SDE	Stochastic Differential Equation
SSQW	Split Step Quantum Walk
SVD	Singular Value Decomposition
UC	Use Case
VQAE	Variational Quantum Amplitude Estimation
VaR	Value at Risk

***Table 2**: Acronyms and Abbreviations*

List of Figures

List of Tables

Bibliography

Abrams, D. S., & Williams, C. P. (1999). Fast quantum algorithms for numerical integrals and stochastic processes.

Agliardi, G., & Prati, E. (2022a). Optimal tuning of quantum generative adversarial networks for multivariate distribution loading. *Quantum Reports*, *4*(1), 75–105.

Agliardi, G., & Prati, E. (2022b). Optimized quantum generative adversarial networks for distribution loading. *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 824–827.

Araujo, I. F., Blank, C., & da Silva, A. J. (2021). Approximated quantum-state preparation with entanglement dependent complexity. *arXiv preprint arXiv:2111.03132*.

Araujo, I. F., Park, D. K., Ludermir, T. B., Oliveira, W. R., Petruccione, F., & da Silva, A. J. (2021). Configurable sublinear circuits for quantum state preparation. https://doi.org/10.1007/s11128-023-03869-7

Araujo, I. F., Park, D. K., Petruccione, F., & da Silva, A. J. (2021). A divide-and-conquer algorithm for quantum state preparation. *Scientific Reports*, *11*. https://doi.org/10.1038/s41598-021-85474-1

Ashhab, S. (2022). Quantum state preparation protocol for encoding classical data into the amplitudes of a quantum information processing register's wave function. *Physical Review Research*, *4*(1), 013091.

Baaquie, B. E. (2019). Merton's equation and the quantum oscillator ii: Option pricing. *Physica A: Statistical Mechanics and its Applications*, *532*, 121792. https://doi.org/https://doi.org/10.1016/j.physa.2019.121792

Baaquie, B. E. (2020). Merton's equation and the quantum oscillator i: Risky coorporate coupon bonds. *Physica A: Statistical Mechanics and its Applications*, *541*, 123367. https://doi.org/https://doi.org/10.1016/j.physa.2019.123367

Bao, J., & Rebentrost, P. (2022). Fundamental theorem for the pricing of quantum assets.

Bausch, J. (2022). Fast black-box quantum state preparation. *Quantum*, *6*, 773.

Berry, D., Childs, A., Ostrander, A., & Wang, G. (2017). Higher order quantum algorithms for solving linear differential equations. *Journal of Physics, A*, *47*(10), 105301.

Brassard, G., Høyer, P., Mosca, M., & Tapp, A. (2002). Quantum amplitude amplification and estimation. https://doi.org/10.1090/conm/305/05215

Brown, E. G., Goktas, O., & Tham, W. K. (2020). Quantum amplitude estimation in the presence of noise. https://doi.org/10.48550/ARXIV.2006.14145

Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2019). Deep hedging. *Quantitative Finance*, *19*(8), 1271–1291.

Callison, A., & Browne, D. E. (2022). Improved maximum-likelihood quantum amplitude estimation. https://doi.org/10.48550/ARXIV.2209.03321

Certo, S., Pham, A. D., & Beaulieu, D. (2022). Benchmarking amplitude estimation on a superconducting quantum computer. https://doi.org/10.48550/ARXIV.2201.06987

Chang, Y.-J., Wang, W.-T., Chen, H.-Y., Liao, S.-W., & Chang, C.-R. (2023). Preparing random state for quantum financing with quantum walks. https://doi.org/10.48550/ARXIV.2302.12500

Chefles, A., Jozsa, R., & Winter, A. (2004). On the existence of physical transformations between sets of quantum states. *International Journal of Quantum Information*, *02*(01), 11–21.

Creevey, F. M., Hill, C. D., & Hollenberg, L. C. L. (2023). Gasp – a genetic algorithm for state preparation. https://doi.org/10.48550/ARXIV.2302.11141

Dalal, A., & Katabarwa, A. (2022). Noise tailoring for robust amplitude estimation. https://doi.org/10.48550/ARXIV.2208.11797

Dasgupta, K., & Paine, B. (2022). Loading probability distributions in a quantum circuit. *arXiv preprint arXiv:2208.13372*.

Egger, D. J., Gutiérrez, R. G., Mestre, J. C., & Woerner, S. (2020). Credit risk analysis using quantum computers. *IEEE Transactions on Computers*.

El Amine Cherrat, e. a. (2023). Quantum deep hedging [arXiv:2023.16585v1].

Fukuzawa, S., Ho, C., Irani, S., & Zion, J. (2023). Modified iterative quantum amplitude estimation is asymptotically optimal. In *2023 proceedings of the symposium on algorithm engineering and experiments (alenex)* (pp. 135–147). https://doi.org/10.1137/1.9781611977561.ch12

Gabor, T., Zorn, M., & Linnhoff-Popien, C. (2022). The applicability of reinforcement learning for the automatic generation of state preparation circuits. *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 2196–2204.

Gilyén, A., Su, Y., Low, G. H., & Wiebe, N. (2019). Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics. *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 193–204. https://doi.org/10.1145/3313276.3316366

Giurgica-Tiron, T., Kerenidis, I., Labib, F., Prakash, A., & Zeng, W. (2022). Low depth algorithms for quantum amplitude estimation. *Quantum*, *6*, 745. https://doi.org/10.22331/q-2022-06-27-745

Gleinig, N., & Hoefler, T. (2022). Circuits for measurement based quantum state preparation. *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 328–333.

Gómez, A., Leitao Rodriguez, A., Manzano, A., Nogueiras, M., Ordóñez, G., & Vázquez, C. (2022). A survey on quantum computational finance for derivatives pricing and VaR. *Archives of Computational Methods in Engineering*, *29*, 4137–4163. https://doi.org/10.1007/s11831-022-09732-9

Gonzalez-Conde, J., Marin-Sanchez, G., & Sanz, M. (2023). Quantum algorithms for approximate function loading. *Bulletin of the American Physical Society*.

Grinko, D., Gacon, J., Zoufal, C., & Woerner, S. (2021). Iterative quantum amplitude estimation. *npj Quantum Information*, *7*(1). https://doi.org/10.1038/s41534-021-00379-1

Grover, L., & Rudolph, T. (2002). Creating superpositions that correspond to efficiently integrable probability distributions. https://doi.org/10.48550/ARXIV.QUANT-PH/0208112

Grover, L. K. (2000). Synthesis of quantum superpositions by quantum computation. *Phys. Rev. Lett.*, *85*, 1334–1337. https://doi.org/10.1103/PhysRevLett.85.1334

Han, J. Y., & Rebentrost, P. (2022). Quantum advantage for multi-option portfolio pricing and valuation adjustments. https://doi.org/10.48550/ARXIV.2203.04924

Herbert, S., Guichard, R., & Ng, D. (2021). Noise-aware quantum amplitude estimation. https://doi.org/10.48550/ARXIV.2109.04840

Iaconis, J., Johri, S., & Zhu, E. Y. (2023). Quantum state preparation of normal distributions using matrix product states. https://arxiv.org/abs/2303.01562

Johri, S., Iaconis, J., & Zhu, E. (2023). Experimental loading of normal probability distributions to quantum states. *Bulletin of the American Physical Society*.

Koh, D. E., Wang, G., Johnson, P. D., & Cao, Y. (2022). Foundations for bayesian inference with engineered likelihood functions for robust amplitude estimation. *Journal of Mathematical Physics*, *63*(5), 052202. https://doi.org/10.1063/5.0042433

Kubo, K., Miyamoto, K., Mitarai, K., & Fujii, K. (2022). Pricing multi-asset derivatives by variational quantum algorithms. https://doi.org/10.48550/ARXIV.2207.01277

Li, Y., & Neufeld, A. (2023). Quantum monte carlo algorithm for solving black-scholes pdes for high-dimensional option pricing in finance and its proof of overcoming the curse of dimensionality.

Lu, X., & Lin, H. (2023). Random-depth quantum amplitude estimation. https://doi.org/10.48550/ARXIV.2301.00528

Manzano, A., Musso, D., & Leitao, Á. (2022). Real quantum amplitude estimation. https://doi.org/10.48550/ARXIV.2204.13641

McArdle, S., Gilyén, A., & Berta, M. (2022). Quantum state preparation without coherent arithmetic. *arXiv preprint arXiv:2210.14892*.

Miyamoto, K. (2022). Bermudan option pricing by quantum amplitude estimation and chebyshev interpolation. https://doi.org/10.1140/epjqt/s40507-022-00124-3

Miyamoto, K., & Kubo, K. (2022). Pricing multi-asset derivatives by finite-difference method on a quantum computer. *IEEE Transactions on Quantum Engineering*, *3*, 1–25. https://doi.org/10.1109/TQE.2021.3128643

Montanaro, A. (2015). Quantum speedup of Monte Carlo methods. https://doi.org/http://doi.org/10.1098/rspa.2015.0301

Moosa, M., Watts, T. W., Chen, Y., Sarma, A., & McMahon, P. L. (2023). Linear-depth quantum circuits for loading fourier approximations of arbitrary functions. *arXiv preprint arXiv:2302.03888*.

Mozafari, F., Micheli, G. D., & Yang, Y. (2022). Efficient deterministic preparation of quantum states using decision diagrams. *Physical Review A*, *106*(2). https://doi.org/10.1103/physreva.106.022617

Nakaji, K. (2020). Faster amplitude estimation. *Quantum Information and Computation*, *20*(13&14), 1109–1123. https://doi.org/10.26421/qic20.13-14-2

Plekhanov, K., Rosenkranz, M., Fiorentini, M., & Lubasch, M. (2022). Variational quantum amplitude estimation. *Quantum*, *6*, 670. https://doi.org/10.22331/q-2022-03-17-670

Plesch, M., & Brukner, Č. (2010). Quantum-state preparation with universal gate decompositions. https://doi.org/10.1103/PhysRevA.83.032302

Pracht, R. (2023). The pricing of American options on the quantum computer. https://doi.org/10.2139/ssrn.4350641

Pracht, R. (2022). Quantum binomial tree, an effective method for probability distribution loading for derivative pricing.

Pracht, R., Ryterski, A., Plewa, J., & Stefaniak, M. (2022). Fx asian option pricing using quantum computers. https://doi.org/10.2139/ssrn.4137397

Radha, S. K. (2021). Quantum option pricing using wick rotated imaginary time evolution. https://doi.org/10.48550/ARXIV.2101.04280

Rall, P., & Fuller, B. (2022). Amplitude estimation from quantum signal processing. https://doi.org/10.48550/ARXIV.2207.08628

Ran, S.-J. (2020). Encoding of matrix product states into quantum circuits of one- and two-qubit gates. *Physical Review A*, *101*, 032310.

Rebentrost, P., Gupt, B., & Bromley, T. R. (2018). Quantum computational finance: Monte Carlo pricing of financial derivatives. *Physical Review A*, *98*(2).

Rebentrost, P., Luongo, A., Bosch, S., & Lloyd, S. (2022). Quantum computational finance: Martingale asset pricing for incomplete markets. https://doi.org/10.48550/ARXIV.2209.08867

Rindell, T., Yenilen, B., Halonen, N., Pönni, A., Tittonen, I., & Raasakka, M. (2022). Generating approximate state preparation circuits for nisq computers with a genetic algorithm. *arXiv preprint arXiv:2210.06411*.

Sanders, Y. R., Low, G. H., Scherer, A., & Berry, D. W. (2019). Black-box quantum state preparation without arithmetic. *Phys. Rev. Lett.*, *122*, 020502. https://doi.org/10.1103/PhysRevLett.122.020502

Stamatopoulos, N., Egger, D. J., Sun, Y., Zoufal, C., Iten, R., Shen, N., & Woerner, S. (2020). Option pricing using quantum computers. *Quantum*, *4*, 291.

Sun, X., Tian, G., Yang, S., Yuan, P., & Zhang, S. (2023). Asymptotically optimal circuit depth for quantum state preparation and general unitary synthesis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*. https://doi.org/10.1109/TCAD.2023.3244885

Suzuki, Y., Uno, S., Raymond, R., Tanaka, T., Onodera, T., & Yamamoto, N. (2020). Amplitude estimation without phase estimation. *Quantum Information Processing*, *19*(2). https://doi.org/10.1007/s11128-019-2565-2

Szasz, A., Younis, E., & de Jong, W. (2023). Numerical circuit synthesis and compilation for multi-state preparation. http://arxiv.org/abs/2305.01816

Tanaka, T., Suzuki, Y., Uno, S., Raymond, R., Onodera, T., & Yamamoto, N. (2021). Amplitude estimation via maximum likelihood on noisy quantum computer. *Quantum Information Processing*, *20*, 1–29.

Tanaka, T., Uno, S., Onodera, T., Yamamoto, N., & Suzuki, Y. (2022). Noisy quantum amplitude estimation without noise estimation. *Phys. Rev. A*, *105*, 012411. https://doi.org/10.1103/PhysRevA.105.012411

Tang, H., Pal, A., Wang, T.-Y., Qiao, L.-F., Gao, J., & Jin, X.-M. (2021). Quantum computation for pricing the collateralized debt obligations. *Quantum Engineering*, *3*(4), e84. https://doi.org/https://doi.org/10.1002/que2.84

Tang, H., Wu, W., & Jin, X.-M. (2022). Quantum computation for pricing caps using the libor market model.

Udvarnoki, Z., Fáth, G., & Fogarasi, N. (2023). Quantum advantage of monte carlo option pricing. *Journal of Physics Communications*, *7*(5), 055001. https://doi.org/10.1088/2399-6528/acd2a4

Uno, S., Suzuki, Y., Hisanaga, K., Raymond, R., Tanaka, T., Onodera, T., & Yamamoto, N. (2021). Modified grover operator for quantum amplitude estimation. *New Journal of Physics*, *23*(8), 083031. https://doi.org/10.1088/1367-2630/ac19da

Wang, G., Koh, D. E., Johnson, P. D., & Cao, Y. (2021). Minimizing estimation runtime on noisy quantum computers. *PRX Quantum*, *2*, 010346. https://doi.org/10.1103/PRXQuantum.2.010346

Yu, K., Lim, H., Rao, P., & Jin, D. (2020). Comparison of amplitude estimation algorithms by implementation. https://doi.org/10.48550/ARXIV.2005.05300

Zhang, X.-M., Li, T., & Yuan, X. (2022). Quantum state preparation with optimal circuit depth: Implementations and applications. *Physical Review Letters*, *129*(23), 230504.

Zhao, Y., Wang, H., Xu, K., Wang, Y., Zhu, J., & Wang, F. (2022). Adaptive algorithm for quantum amplitude estimation. https://doi.org/10.48550/ARXIV.2206.08449

