QUANTUM COMPUTING FOR FINANCE

Oswaldo Zapata, PhD
Cofounder, The Quantum Finance Boardroom

Introduction

This chapter aims to illustrate how quantum computing is expected to transform the future of finance. It provides a concise overview of the fundamental concepts underlying quantum computing, along with several of its most prominent applications in finance, such as portfolio optimization and quantum-enhanced machine learning.

Quantum computing encompasses three related but distinct quantum technologies: quantum computation, quantum communication, and quantum sensing. In this chapter, I primarily focus on quantum computation—the idea that quantum computers can solve certain problems more efficiently and accurately than classical computers. Quantum communication, although currently of interest to financial organizations, will be mentioned briefly. Quantum sensing, which has no application in finance, will not be discussed.

The chapter is organized as follows: First, I provide an overview of quantum computing, emphasizing current capabilities in light of the present state of hardware development. Next, I review applications of machine learning in finance and present the fundamental ideas behind applying quantum computing to financial problems. Then, I discuss quantum communication, and I conclude with a summary and a brief discussion of other topics of interest.

Quantum Computation

Quantum computation is a quantum-mechanical approach to solving computational problems. Instead of relying on traditional Boolean algebra, where information is represented in binary form—either 0 or 1—quantum computation uses the principles and mathematical formalism of quantum mechanics, such as state vectors, unitary operators, and measurements, to arrive at logical solutions to computational problems. A *quantum computer* is the physical device that implements the quantum computational processes of interest.

Quantum mechanics describes nature in a fundamentally different way from classical physics, which relies on such concepts as force, mass, velocity, and electrical currents. With this caveat in mind, however, drawing an analogy between classical and quantum circuits can still be useful.

A standard electronic circuit consists of electrical currents and a series of electronic components known as logic gates. A circuit can be designed specifically to solve a given problem. This arrangement of currents and gates—that is, the circuit—constitutes an algorithm. These circuits are deterministic in the sense that, given a set of input currents and a sequence of gates, the output is uniquely determined. The efficiency of the algorithm is evaluated on the basis of the number of gates and layers it uses—a field known as circuit complexity.

The basic elements of a quantum circuit are the quantum analogues of electric currents and electronic components. The "quantum electric currents" are called *qubits*, and the "quantum electronic components" are called *quantum gates*. An arrangement of quantum gates forms a *quantum circuit*. When a quantum circuit is designed to solve a particular computational problem, it is referred to as a *quantum algorithm*. In contrast to classical circuits, quantum circuits are generally nondeterministic because the *measurement* process—an integral part of a quantum circuit—typically produces a probabilistic output. For example, in a classical circuit, one always measures either the presence or absence of a current in a predictable way. In contrast, in a quantum circuit, measuring the same qubit can yield different outcomes—sometimes indicating the presence of a current and other times indicating its absence—even if the circuit is unchanged.

In mathematical terms, qubits are represented by vectors in a normed complex vector space, and quantum gates are realized by unitary transformations. The most basic qubit is a *1-qubit*. This quantum state is represented by a complex linear combination of two nonparallel vectors (usually taken as orthogonal). Quantum gates acting on single qubits change the configuration of the qubit by modifying the complex coefficients. For example, a quantum gate can exchange one of the basis vectors for the other or simply make it vanish. The measurement process at the end provides a probabilistic result based on the new combination created by the quantum gate, or gates, depending on the complexity of the problem. For more general qubits, we use the term *n-qubit*. The vector describing this quantum state is more complicated: It involves 2^n nonparallel vectors, usually orthogonal, and 2^n complex coefficients. Quantum gates are *unitary transformations* that change the coefficients of the input qubit. Quantum gates and circuits more generally leverage the *entanglement property* of quantum mechanics, which involves creating an output qubit that contains more information than the individual input qubits used to create it.

A final word about quantum circuits: It has been known since the early days of modern computing that any classical circuit can be constructed using a small set of electronic components. Such a set is known as a universal gate set. Similarly, any quantum operation can, in theory, be approximated using a finite set of basic gates—known as a *universal quantum gate set*.

The motivation for developing quantum computers stems from two central beliefs:

- They may significantly outperform classical devices on certain tasks, providing exponential or polynomial increases in speed.
- Quantum machines might be able to solve problems deemed intractable for classical computers—offering not only faster computation but also fundamentally new capabilities.

Although this potential remains largely theoretical at present, growing evidence and experimental progress suggest that quantum computation could revolutionize the current approach to complex problems across various scientific and technological fields, including finance. Significant challenges remain, however, and here I highlight two: (1) the difficulty of preserving the quantum properties of qubits and quantum gates and (2) the challenge of constructing larger, more complex quantum circuits.

If a qubit or quantum gate unpredictably changes its quantum properties, the theoretical predictions will no longer align with the experimental results. This discrepancy can compromise the entire quantum computation. Such unpredictability is referred to as an *error*. Fortunately, some

errors can be corrected. This challenge is not entirely new. In the early days of classical computing, digital components were also imperfect, and *error correction methods* were essential. Over time, classical hardware became so reliable that such corrections became largely unnecessary. In contrast, quantum systems remain highly sensitive and difficult to isolate from their environments. Even minor interactions—collectively known as *noise*—can disrupt quantum behavior. To tackle this problem, engineers work on techniques to shield quantum circuits from environmental noise, and theorists develop procedures to detect and correct various types of errors.

A common example is a bit-flip error, where the state of a qubit flips unexpectedly during transmission or processing. In classical computing, such errors are corrected by duplicating the bit and using majority voting; for example, instead of using 0, we use 000. If one copy is flipped, the system can identify and fix the mistake by comparing the values. This method assumes that the probability of an error occurring is extremely low; if multiple errors occur, the correction fails. Quantum computing applies a similar idea but with important differences. Instead of copying qubits directly (impossible because of quantum rules), the information is spread across multiple qubits in a way that allows for error detection and correction. These groups of qubits are called *logical qubits*, and each individual qubit in the group is referred to as a *physical qubit*.

Detecting errors in quantum systems is more delicate than in classical ones because directly measuring a qubit destroys its state. Instead, quantum computers use indirect methods—such as *parity checks*—to detect errors without collapsing the quantum information. Once an error is identified, correction techniques are applied to restore the original state. Beyond simple bit flips, other types of errors can affect qubits, including those caused by faulty gates. If a qubit enters a faulty gate, the error may propagate throughout the quantum circuit. Worse still, the process of correcting errors also involves quantum components—which means it can introduce new errors. This dynamic creates a paradox: Fixing errors can sometimes cause more of them. As a result, building reliable quantum computers requires scaling up the system, which further increases the chances of something going wrong. Fortunately, researchers have proven that if certain conditions are met, error correction codes can reduce error rates to very low levels. A *fault-tolerant quantum computer* is one that continually detects and corrects errors in its logical qubits throughout the computation, ensuring the final result is reliable.

In recent years, however, scientists have come to accept that fault-tolerant quantum computers will not be available anytime soon. As a result, they began looking for more realistic algorithms that could be implemented on near-term quantum computers, which are characterized by a moderate amount of noise and a relatively small number of qubits and gates. We are currently in this stage, known as the *noisy intermediate-scale quantum* (NISQ) era. According to experts, we will likely remain in this phase for several years (even decades) before achieving fault tolerance.

The algorithms expected to be implemented in the near term are known as *hybrid quantum–classical algorithms*. These combine quantum and classical parts: The quantum subroutines address problems that are hard for classical computers, while the classical computer handles tasks where its efficiency is well established. For example, *variational quantum algorithms* (VQAs) are NISQ algorithms designed to demonstrate *quantum advantage* (practical *quantum supremacy*) in the near future. Because many problems—not only in physics and chemistry but also in finance—share a common underlying structure, the techniques used in VQAs can be applied to a wide variety of situations. VQAs are considered *heuristic*, which means that

although there is currently no rigorous proof that they outperform known classical algorithms, there are theoretical reasons for optimism. The hope is that future results will demonstrate their advantage.

Machine Learning for Finance

As mentioned in several chapters of this book, *artificial intelligence* (AI) is transforming many industries, including finance. Banks and public institutions are using AI to detect fraud, assess credit risk, and identify investment opportunities. This section focuses on *machine learning* (ML)—a subfield of AI concerned with algorithms that seek to uncover patterns in data. In finance, where data such as stock movements and customer behavior are abundant, ML models are used to analyze this information and generate actionable insights. In this section, I will review several ML algorithms currently used in finance—especially those seen as promising candidates for enhancement through quantum computing. But first, it is essential to clarify what is meant by "data."

The term *data* refers to information associated with physical objects or abstract concepts. In ML, such information comes in various forms and is often classified into two major types. *Structured data* are organized and easily represented in tabular formats, such as matrices. In contrast, *unstructured data* lack this inherent organization; examples include raw text, images, and audio recordings. Despite being more prevalent in the real world, unstructured data must first undergo cleaning and formatting before they can be used in ML models. If the data are incomplete, inconsistent, or poorly selected, models built on such foundations may yield inaccurate or misleading predictions. Therefore, the preprocessing stage is not a peripheral step but a foundational component of the machine learning pipeline.

Supervised learning is among the most widely used ML paradigms. In this approach, models are trained on *labeled data*, where the inputs and their corresponding outputs are known in advance. The objective is to learn a mapping from inputs to outputs that generalizes well to new, unseen data. During the *training phase*, the model iteratively adjusts itself to minimize the discrepancy between its predictions and the actual labels. This adjustment process is typically achieved by minimizing a *loss function*, and the model's ability to generalize is then evaluated using a separate *test set*. Supervised learning tasks are commonly divided into regression and classification problems. In *regression*, the model predicts continuous outcomes. *Classification* tasks, in contrast, involve assigning discrete labels to data points.

Unsupervised learning diverges from supervised learning in a fundamental way: It operates on datasets that lack labeled outputs. The aim of unsupervised learning algorithms is to uncover hidden structures or patterns in the data without the aid of explicit guidance. In these cases, the input consists solely of features, and the algorithm is tasked with identifying intrinsic relationships among them. One common challenge in such scenarios is managing high-dimensional data, where the number of features is large. Reducing the dimensionality of the feature space becomes an essential step—not only to simplify the data but also to enhance the efficiency of subsequent analysis. One systematic method to perform such feature reduction is *principal component analysis* (PCA), a key technique under the broader category of *dimensionality reduction*.

Neural networks represent a class of ML models that can be supervised or unsupervised. Inspired by the structure of the human brain, these models consist of layers of interconnected

nodes, or *neurons*. Each neuron performs a computation based on its inputs and transmits the result to the next *layer*. The network is trained by adjusting internal parameters, such as weights and biases, to optimize an objective function—often using such algorithms as *gradient descent*. For tasks involving sequential data, *recurrent neural networks* are often used. These networks differ from traditional feed-forward models in that they maintain a hidden state, enabling them to incorporate information from previous inputs. However, recurrent networks are prone to such issues as the *vanishing gradient problem*, which hinders their ability to capture long-term dependencies. *Long short-term memory networks* address this limitation by introducing gating mechanisms that preserve important information over extended sequences.

ML applications in finance are vast and span various domains, such as credit scoring, fraud detection, portfolio management, and market analysis. Some explicit examples are covered next. The improvement of these ML methods with quantum computers is expected to enhance financial performance and services.

Supervised models, both regression and classification, are particularly useful in credit scoring, where they predict the probability of a borrower repaying a loan or defaulting. The dataset typically includes personal information about borrowers—such as their age, gender, and financial data, including loan amount, credit history, and repayment records. Once the model is trained on this enriched dataset, it can predict the likelihood of loan repayment for new applicants. Financial institutions can then use this prediction to decide whether to approve a loan or set credit limits based on the applicant's perceived creditworthiness. Risk assessment in finance is a more comprehensive task than credit risk evaluation alone. In this context, risk assessment involves predicting both the probability and potential cost of adverse events that could impact a company's financial health. The specific dataset and features used depend on the type of risk being assessed, which could include market risk, credit risk, operational risk, or country risk. If the goal is to assess the risk an event poses to a company's market valuation, relevant data are gathered and the supervised model is trained to uncover patterns. By analyzing the relationships between the input features and target variables (such as changes in market valuation or earnings volatility), financial institutions and corporations can make more-informed decisions to mitigate risks.

The *k-nearest neighbors* (kNN) algorithm is a classification method that is particularly well suited for such tasks as credit risk assessment and fraud detection. In credit risk assessment, for example, kNN can predict whether a potential borrower will repay a loan or default. The algorithm works by training on data from previous borrowers, where each data point consists of personal and financial information along with a repayment history. Once trained, the model compares a new applicant's data to the *k*-nearest neighbors in the training dataset and, through majority voting, predicts whether the applicant is likely to repay or default. Similarly, in transaction fraud detection, kNN compares new transactions with past ones, determining whether new transactions appear fraudulent based on their similarity to previous fraudulent or legitimate transactions. The kNN algorithm can also be used to identify money laundering patterns, although money laundering detection often requires more complex feature engineering and domain-specific knowledge.

PCA is another powerful tool used in finance, particularly for such tasks as credit risk assessment and portfolio management. In credit risk assessment, many features in the dataset are highly correlated, such as income, debt-to-income ratio, and credit utilization. PCA helps by reducing the dimensionality of the dataset while preserving the most important variance

in the data. This process allows risk analysts to focus on the most significant factors influencing credit risk, such as identifying borrower clusters with similar risk profiles or spotting outliers that may represent unique risks. PCA can also be applied to portfolio management, where it helps in reducing data complexity, focusing on the most relevant risk factors for asset allocation. By simplifying the covariance matrix of assets, PCA identifies and removes highly correlated assets, allowing portfolio managers to maintain diversification without unnecessarily complicating the portfolio. The efficiency improvements brought by PCA make it a valuable tool in market analysis as well, where it can help focus on key market drivers while discarding less relevant factors.

The *k-means clustering* algorithm is an unsupervised learning technique useful in detecting previously unseen patterns or suspicious behaviors in data. Unlike supervised learning algorithms, *k*-means does not rely on labeled data, making it more flexible in uncovering unexpected trends. This ability is particularly valuable in evolving scenarios, such as fraud detection and anti-money-laundering, where fraudsters continuously adapt their tactics. By grouping similar data points together and identifying outliers that do not fit well into any cluster, *k*-means can reveal unusual or fraudulent behaviors that might otherwise go unnoticed. Although not explicitly designed as an anomaly detection algorithm, *k*-means' ability to highlight atypical data points makes it an important tool for detecting fraud and preventing money laundering.

Quantum Algorithms for Finance

In the dynamic environment of the modern financial industry—characterized by intense competition and evolving regulations—quantum computing holds great promise because it is expected to surpass classical systems in both efficiency and security. Some experts predict that finance may be one of the first industries to undergo a transformation driven by quantum computing. The timeline for the availability of fully functional quantum computers remains uncertain, however.

Remember that fully reliable, fault-tolerant quantum computers are still many years from being realized. We are in the so-called NISQ era, characterized by quantum devices that are relatively noisy and support only a limited number of quantum gates. As a result, researchers have focused on hybrid quantum–classical algorithms, which combine the strengths of both classical and quantum computing. In these hybrid models, quantum computers tackle the most computationally demanding parts of a problem, while classical computers handle the remaining tasks. This approach offers practical advantages: The quantum subroutines require only a limited number of coherent qubits and shallow circuits, making them feasible with today's quantum technology.

Quantum Portfolio Optimization

One of the key areas where quantum computing can improve on classical methods is portfolio optimization. Traditional portfolio optimization techniques can struggle with large datasets, particularly when the portfolios are vast and require the processing of complex data. Variational quantum algorithms, such as the *variational quantum eigensolver* (VQE) and the *quantum approximate optimization algorithm* (QAOA), are believed to offer improvements by processing large datasets more efficiently than classical algorithms.

The VQE is an algorithm that leverages the variational principle of quantum mechanics to approximate solutions to the time-independent Schrödinger equation. This equation describes the behavior of complex quantum systems, such as molecules or the electronic configurations of materials. The VQE was originally developed for quantum chemistry applications, but researchers have explored adapting it to represent the quantized version of certain classical problems, including portfolio optimization, by reformulating them as Hamiltonian minimization tasks. The QAOA is another variational algorithm, specifically designed to solve classical combinatorial optimization problems, such as portfolio optimization. What makes VQAs particularly compelling is that they are designed for hybrid classical–quantum computers. The quantum subroutine prepares quantum states and computes the Hamiltonian expectation values, while the classical computer performs the optimization process.

This hybrid approach makes VQAs suitable for solving complex problems that are challenging for purely classical systems, especially in such domains as portfolio optimization. Although it is still uncertain whether VQAs will outperform classical algorithms, the potential for significant improvements in processing massive datasets and performing optimization tasks at a faster rate could open up new avenues in financial modeling and analysis.

Quantum Machine Learning

In this section, I briefly address the main challenges quantum computing faces in enhancing and potentially revolutionizing classical ML techniques.

The primary challenge in *quantum machine learning* today lies in effectively encoding classical data into qubits so that the quantum computer can process the data efficiently. Several methods have been developed to facilitate this encoding process, enabling quantum computers to perform computations on classical data.

I will illustrate this challenge with the simplest example. Suppose you have two classical data points—for example, two positive numbers—and you want to insert this information into a quantum computer to process them using the quantum algorithm you have designed. These classical data points must be transformed into quantum information that the quantum computer can understand. Perhaps the simplest way to encode these classical data points into a quantum state is by using the angles of a single qubit, a process known as *angle encoding*. As previously mentioned, a 1-qubit generally requires two complex numbers, which corresponds to four real numbers, to fully specify it. By the principles of quantum mechanics, however, these four real numbers can be reduced to only two. The single qubit can thus be represented as a vector on the surface of a unit sphere. Because the position of any point on the sphere is completely determined by two angles, the azimuthal and polar angles, the 1-qubit is determined by these two angles. By properly rescaling if necessary, the original two classical data points can be encoded in the 1-qubit by rotating it accordingly. That is, the two data points can be encoded in the rotation angles of the 1-qubit. For more complex situations with many more classical data points, general n-qubits are necessary, but the principle remains the same. The real challenge lies in implementing these ideas in real-world scenarios.

In quantum ML, several algorithms have been developed to speed up learning tasks by processing vast amounts of data more efficiently than classical systems. These algorithms can be applied to a variety of ML problems, ranging from classification and regression to clustering

and optimization. In the context of finance, quantum ML holds the potential to improve financial applications, as discussed earlier.

Note that ML is a technology that was only recently incorporated into the financial sector, and quantum ML is still in the early stages of research. As quantum technology progresses, quantum algorithms are expected to become increasingly capable of handling more complex data and providing substantial advantages over classical approaches.

Quantum Cryptography

As mentioned previously, quantum computing is much more than the acceleration of computationally expensive problems. It also encompasses *secure communication*.

The issue is that sufficiently powerful quantum computers could break the current encryption methods used by most public institutions and private organizations, potentially gaining access to sensitive data, such as military and financial information. The threat posed by quantum computers is a reality that governments and financial institutions are taking very seriously. A discussion of these concepts follows.

An *encryption standard* is a method used to transform information into a form that is not easily related to the original. The original form is referred to as *plaintext*, and the transformed version is known as *ciphertext*. Most contemporary digital encryption standards are based on mathematical problems that are difficult for classical computers to solve. Quantum algorithms, however, have the potential to solve some of these problems efficiently. For example, *RSA*, widely used to secure digital data over the internet, relies on the difficulty of integer factorization, and *ECC* (*elliptic curve cryptography*) is based on the discrete logarithm problem. The concern is that *Shor's algorithm*, a quantum algorithm, can efficiently solve both of these problems. Once sufficiently large quantum computers become available—potentially in the next 5–10 years—these encryption systems could be broken in a relatively short period of time. These two examples represent some of the most vulnerable standards in the quantum era. In response, governments worldwide are enacting laws to secure sensitive data.

Post-quantum cryptography (*PQC*) offers one potential solution. It involves the development and, increasingly, the implementation of a set of encryption methods based on mathematical problems designed to remain secure against both classical computers and, more importantly, quantum algorithms running on quantum computers. These algorithms are based on mathematical problems that are not known to be efficiently solvable by quantum algorithms, such as Shor's and Grover's. Examples of such problems include lattice problems, multivariate polynomials, code-based schemes, and hash-based signatures. In the United States, the National Institute of Standards and Technology (NIST) is currently in the process of standardizing several PQC algorithms to either supplement or completely replace existing cryptographic systems. NIST has determined that by 2030, US federal agencies should treat current standard encryption methods as vulnerable, and by 2035, these methods are expected to be phased out. US financial institutions, however, have not yet established a timeline for transitioning to post-quantum cryptography.

Quantum key distribution (QKD) is regarded as one of the most secure encryption methods because it relies not on complex mathematical problems but on the fundamental laws of physics. Information exchange is protected by principles of quantum mechanics—most notably,

the *no-cloning theorem* and the fact that measurement inherently disturbs the system being observed. Although QKD provides strong theoretical security, its practical implementation is currently constrained by high costs, the need for specialized hardware, and distance limitations. As a result, QKD is best suited for niche applications—such as government or high-security financial sectors—that can support dedicated infrastructure.

Transitioning to a fully *quantum-safe* infrastructure—particularly for such institutions as banks—is costly and may take many years to complete. To initiate this transition, experts recommend adopting hybrid cryptosystems that combine traditional (classical) encryption algorithms, which are less expensive and faster to implement, with quantum-resistant algorithms that may require new types of technology and infrastructure. This combination offers protection against both classical and potential future quantum attacks and is especially valuable during the transition phase.

Conclusion

Quantum computing is a technology still in the research phase, awaiting widespread adoption. Although its practical advantages remain limited at this time, it holds the promise of improving the most computationally demanding calculations in such industries as pharmaceuticals, logistics, and finance. In the case of finance, it could help in the process of portfolio optimization, as well as in many services currently tackled by ML techniques. Another active area of research in the finance community is *quantum-enhanced Monte Carlo simulation*. Unlike with variational quantum algorithms and quantum ML, there is mathematical evidence that the quantum version of classical Monte Carlo can enhance parts of the process. Challenges persist, however: Noise still affects quantum systems, and errors remain significant.

Quantum computation and quantum communication are active areas of research, not only in hardware and software but also in real-world applications. The coming years will bring more powerful quantum hardware and a growing interest in applications such as those in finance. The future is certainly exciting for quantum computing and its role in the financial sector.