

BIS Papers

No 149

Quantum computing and the financial system: opportunities and risks

by Raphael Auer, Angela Dupont, Leonardo Gambacorta, Joon Suk Park, Koji Takahashi, and Andras Valko

Monetary and Economic Department

October 2024

JEL classification: C19, C63, C8, M15, G1, G17

Keywords: quantum computing, quantum algorithm, quantum cryptography, quantum-resilient cryptography, artificial intelligence, computational finance, Project Leap.

The views expressed are those of the authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2024. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1682-7651 (online)
ISBN 978-92-9259-794-8 (online)

Quantum computing and the financial system: opportunities and risks*

Raphael Auer, Angela Dupont, Leonardo Gambacorta, Joon Suk Park, Koji Takahashi, and Andras Valko*

Abstract

Quantum computers are still in an experimental phase, but in the future, they may have a profound impact on the financial system. By providing faster and potentially more efficient solutions, quantum computers have the potential to solve certain complex problems that are of paramount interest in the field of economics and finance. For example, quantum simulation algorithms can be leveraged in stress testing and macroeconomic analysis, and quantum optimisation can be used in asset pricing. Meanwhile, the advent of quantum computers also introduces a potential threat to financial stability, especially through their ability to breach some of the most widely used cryptographic algorithms. Despite the nascent state of quantum computing development, the potential for sensitive data to be stored now with the intention to be decrypted later necessitates immediate preparation. This paper explores the transformative potential of quantum mechanics and its applications to the financial system, including the potential benefits as well as the main risks. It also highlights current actions within the central bank community to address these potential risks, including Project Leap, started by the Bank for International Settlements Innovation Hub, the Banque de France and Deutsche Bundesbank.

Keywords: quantum computing, quantum algorithm, quantum cryptography, quantum-resilient cryptography, artificial intelligence, computational finance, Project Leap.

JEL classification: C19, C63, C8, M15, G1, G17.


* Raphael Auer and Leonardo Gambacorta are with the Bank for International Settlements (BIS) and affiliated with the Center for Economic and Policy Research. Koji Takahashi, Joon Suk Park and Andras Valko are with the BIS. Angela Dupont is with the BIS and the Banque de France. We thank Giulio Cornelli and Ilaria Mattei for assistance with data and figures, and Christophe Laforge and David Koepfer for valuable comments on earlier drafts. The authors acknowledge the use of GPT-4 for proofreading and style edits. The views expressed are those of the authors and not necessarily those of the BIS and the Banque de France.

1. Introduction

In recent years, the pace of research and development (R&D) in quantum computing has accelerated. Following the initial idea in the 1980s (Feynman (1982)), the concept was brought closer to reality in 1998, when researchers at Oxford University created the first two-qubit quantum computer (QC), demonstrating the potential of leveraging quantum physics to perform calculations (Jones et al (1998)). Between 2023 and 2024, experimental QCs reached 1,000 qubits in size (Atom Computing (2023), IBM (2024)), marking a significant increase in computational capabilities. As of 2024, though QCs are still in experimental phase, technology developers have further extended the horizon with roadmaps to systems with 2,000 logical qubits (IBM (2024)), aiming to provide the full power of quantum computing beyond 2033 (Graph 1).

Evolution of quantum computing over time

Graph 1



Sources: Vandersypen et al (2001); Gartner; IBM Quantum Roadmap; NIST; Stanford Encyclopedia of Philosophy Archive; authors' elaboration.

A qubit, or quantum bit, is the basic unit of information in QCs, similar to bits in classical computing. While a classical bit must always be in one of two possible states, 0 or 1, a qubit can exist in a superposition of those two states, with a certain probability of being a 0 and a certain probability of being a 1. This property is counterintuitive to those used to classical physics, but it is fundamental in quantum mechanics. The ability of qubits to be in superposition, combined with another property called entanglement, meaning that multiple qubits can share a common quantum state, opens the possibility to build computers working fundamentally differently from classical computers.

While a single qubit can represent a combination of two states, due to superposition, a 1,000-qubit QC has theoretically the capability to represent 2^{1000} different states simultaneously.¹ QCs utilise the entanglement between qubits and

¹ This doesn't mean a QC stores all 2^{1000} states simultaneously in the classical sense, but rather that it can represent a quantum superposition involving all those states.

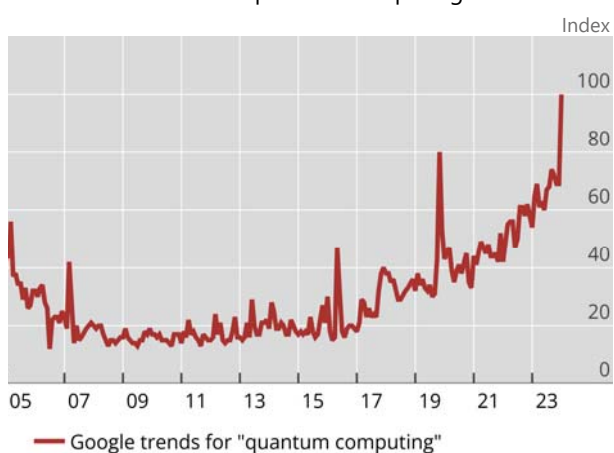
probabilities relating to superpositions to perform a series of operations. This allows them to perform some calculations much more efficiently than classical computers, creating the possibility of faster and more accurate computations in some cases than with any computer available today.

This potential and the rapid development in the field have attracted the attention of all relevant actors: researchers, investors, financial market operators and policymakers (Graph 2.A). The increasing interest in the new opportunities provided by QCs gives rise to a reciprocal interaction between surging investments by venture capitalists in quantum computing-related projects and new discoveries in this field of research (Graph 2.B).

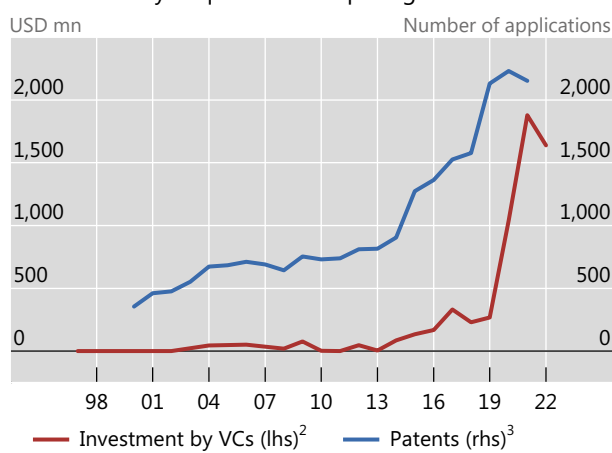
Increasing attention and R&D activities in quantum computing

Graph 2

A. Public attention to "quantum computing"¹



B. R&D activity in quantum computing



¹ Google trends index for "quantum computing" retrieved on 15 December 2023. ² Investment by VCs indicates the total value invested by venture capitalists in firms involved in quantum computing projects. ³ The number of the patents is based on the patents that include "quantum" in their abstract and are applied to US and European patent offices.

Source: Google Trends; PATENTSCOPE; PitchBook.

Quantum machines operate on fundamentally different principles and bring new solutions for intractable challenges.² Some of the most promising fields include simulations, optimisation and search problems. Indeed, one of the first application areas will likely be the simulation of quantum mechanical phenomena. In addition, quantum computing has potential applications in many other sectors, such as chemical and biological engineering and complex manufacturing (Bova et al (2021)).

This paper examines the impact of QCs on the financial system, with a focus on potential applications, including opportunities and risks. Indeed, QCs hold the

² Quantum computers will not merely be faster versions of today's general purpose computers. In practice, classical and quantum computers are expected to complement each other. Classical computers will continue to handle routine tasks efficiently, while quantum computers will tackle specific, complex problems that are infeasible for classical computing.

promise of benefiting the financial system through their ability to implement novel algorithms, which can significantly improve complex pricing and risk evaluation processes. Conversely, as highlighted above, QCs pose a threat to the security of the financial system by potentially breaking current cryptographic protocols that are securing financial data. The paper highlights both of these aspects and discusses strategies that central banks and financial intermediaries can employ to mitigate the threats arising from the advancement of QCs.

In the financial sector, QCs could offer new possibilities by enhancing risk assessment, optimising portfolio management, improving macro analysis and many other areas. Combined with artificial intelligence, it may also provide significant value, for example, in fraud detection or trading strategies. Further, new applications will continue to emerge, as evidenced for example by the strong growth of literature on quantum algorithms for financial applications.

At the same time, QCs pose a serious threat to the financial system due to their expected ability to break some commonly used encryption techniques. In digital communications and in financial systems, cryptographic schemes are based on mathematical problems which are prohibitively difficult for today's computers to solve. However, some of these problems will be solvable by future QCs, opening the possibility for malicious actors to hack encryption. Without proper action, authentication and encryption methods reliant on current cryptographic techniques will become ineffective, compromising confidentiality and integrity and potentially leading to substantial financial and reputational damages (BIS (2023)).

The remainder of this paper is structured as follows. Section 2 provides an overview of quantum computing and quantum algorithms. Section 3 delves into the potential applications of QCs within the financial sector, including advancements at the confluence of quantum computing and artificial intelligence. In Section 4, we discuss the economic effect of these potential applications. In Section 5, we describe the effects of quantum mechanics on cryptography, including risks and opportunities. Based on this, section 6 discusses potential impacts on the financial system and possible actions to mitigate risks. Section 7 summarises the main conclusions.

2. Quantum computing and quantum algorithms

Quantum computing is a groundbreaking technology that utilises the principles of quantum mechanics to perform computations. Unlike classical computers, which process information in binary digits (bits) that can be either 0 or 1, a QC uses quantum bits, or qubits. A qubit can exist in a state of *superposition*, with a certain probability of being a 0 and a certain probability of being a 1. In addition, two qubits can be in *entanglement*, meaning that they share a common quantum state. These phenomena of superposition and entanglement defy explanation through classical physics and often seem counterintuitive without a deep understanding of quantum mechanics. In a nutshell, these two properties enable QCs to perform calculations in radically new ways compared with classical computers.

While the concept of QCs has been known since the 1980s and the first experimental small-scale QCs were built in the late 1990s, practical QC implementations are not yet available, and currently it is impossible to know if or when they will be. A main challenge is related to keeping qubits in a quantum state for the time of computations. If the qubits are not sufficiently isolated from their environment, their quantum state can collapse, introducing noise in the calculations.³ To mitigate this issue, QCs require an even larger number of qubits than would otherwise be required. This allows for redundancy in data representation inside the QC and ensures that even if some qubits are compromised by noise, others can continue to function correctly.⁴

A key research direction relates to increasing the number of qubits in a QC. A significant milestone in the evolution of this technology was Google's announcement in 2019 that its 53-qubit QC had surpassed classical computers in a specific task (Arute et al (2019)).⁵ IBM's introduction of a 433-qubit QC in 2022 marks another considerable advance in QC capabilities. At the time of writing, the largest experimental QCs have just over 1,000 physical qubits, while according to some estimates, 1 million physical qubits are needed for QCs to start addressing practical problems (NEC (2024)).

Yet another line of research is looking at computational techniques and algorithms that can take advantage of a future QC. Classical computers have operating systems to hide the details of the physical implementation of bits, so that a programmer can implement an algorithm without thinking about the physical details of a computer. Similarly, the field of quantum computing is developing its own set of algorithms and computational techniques that can be applied to a future QC, irrespective of exactly how qubits will be implemented.

QCs will not be faster versions of today's computers, and they will not be applicable to all tasks performed by classical computers today. In fact, as understood today, quantum computing is suited for a specific, narrow set of computational tasks only. For these tasks, a QC could outperform classical supercomputers by several orders of magnitude, making it possible to solve some computational problems that are practically unsolvable today.

Quantum supremacy is the milestone of using a QC to solve a problem that no classical computer can solve within a realistic time frame. The term does not have an exact definition; for example, there are different views on whether the problem must be a useful real-world application or whether it can be an artificial test case.

³ Building QCs presents several complex challenges, such as the need for quantum systems to be cooled to near absolute zero temperature to maintain coherence. Developing a large number of precisely controlled qubits is a challenge that is being actively addressed by researchers. But significant advancements are required before reaching large-scale, practical QCs.

⁴ The term "logical qubit" is sometimes used to represent a hypothetical error-free qubit implemented using a collection of physical qubits.

⁵ Other researchers have pointed out that the selected use case involved a highly specialised problem with limited applicability beyond the demonstration. Additionally, it is feasible to improve classical algorithms to achieve comparable results within a reasonable time frame (Kalai et al (2022)).

Correspondingly, there is no consensus on whether quantum supremacy has already been achieved (Kalai et al (2022)). While some research teams claim they have already achieved quantum supremacy, many experts believe that a practical QC is many years away, and some argue that it might never be realised (Dyakonov (2018)). Based on a survey of 37 leading global experts in quantum computing, Mosca and Piani (2023) estimate that a cryptographically relevant QC could emerge within five to 30 years. An “optimistic” view of the responses suggests an 11% likelihood of a cryptographically relevant QC within five years (up from 6% in 2022) and a 31% chance within 10 years (up from 27% in 2022). Even under a “pessimistic” scenario, there is an estimated 33% likelihood of a disruptive quantum threat within the next 15 years.

In theory, without time and computing limitations, any computational problem that is solvable by a QC is also solvable by a classical computer, and any problem solvable by classical computers can also be solved by future QCs (Nielsen and Chuang (2010)). However, certain problems are significantly quicker to solve using a QC, because of its specific computational approach. A subset of problems, though theoretically solvable by classical computers, would require a prohibitively large computational time, even using today’s largest supercomputers, but would be solvable by a QC in reasonable time. Quantum algorithms is a field of research looking into these problems and into how they would be addressed by a future QC of sufficient size.


A number of algorithm families have emerged in the literature, such as quantum Fourier transform, amplitude amplification and quantum walks (Montanaro (2016)). A specific sub-class of algorithms uses a combination of quantum and classical computing and is supposed to be used in mixed computing environments. Since QCs are still in an experimental phase, most of this work is currently theoretical. These algorithms are either only described mathematically and not yet implemented, or they are tested in an emulated QC which does not provide the quantum speedup that future real implementations promise.

3. Quantum computing applications in the financial system

The financial system is characterised by complex calculations for analysis, forecasting and optimisation. Given the high dimensionality of the underlying issues, today’s high-performance computers can struggle to handle the computational complexity of these tasks. Quantum computing, with its potential ability to process vast numbers of solution states and perform complex calculations efficiently, holds great promise for addressing some of these challenges. Indeed, finance is projected to be one of the first industries to reap the benefits from quantum computing (Ménard et al (2020), Pistoia et al (2021)). Interest in applications in finance is thus mounting, as also evidenced by the increasing number of publications in the field (see Graph 3).

Quantum computers and finance: Google scholar results

Graph 3



Note: The y-axis of the graph represents the number of search results retrieved from Google Scholar, quantified in absolute terms. Each data point corresponds to the total count of entries (eg articles, papers, theses) returned by Google Scholar for a given search query. The x-axis delineates the specific time periods under investigation, with each corresponding point reflecting the cumulative number of indexed results. For 2024, data up to 19 June 2024 were used.

Source: Google Scholar.

Graph 4 illustrates the broad range of quantum computing applications in finance. One area where the technology could be useful is evaluating financial risk or performing more accurate stress tests in order to improve **risk management**. Another field where the impact of quantum computing may be significant is **investments and portfolio management**, including financial market risk evaluation. More accurate simulations and analysis could improve investment strategies and asset pricing. It is expected that **artificial intelligence (AI) and machine learning** will gain substantial improvement with better fraud detection mechanisms or even better trading strategies. In addition, other areas such as **payments and settlement** and **macro modelling** would be positively impacted. In what follows, we discuss each of these areas in more detail.

Risk management

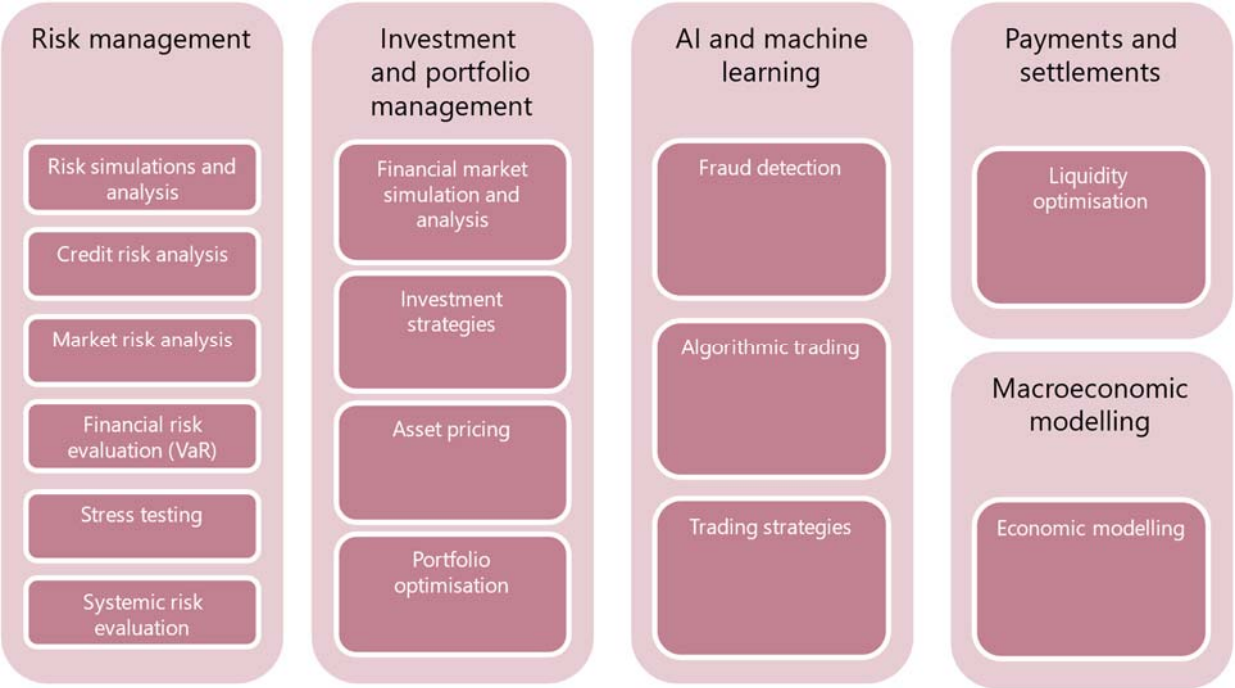
Risk management quantum algorithms for certain financial use cases are expected to offer a quadratic speedup⁶ compared with classical algorithms. This covers applications from risk simulations to systemic risk evaluation. This approach promises to transform how financial risks are modelled and mitigated. In some cases, QCs are capable of performing tasks that are beyond the reach of traditional computers. They

⁶ A quadratic speedup refers to a significant improvement in the computational efficiency of an algorithm or problem-solving process. The time required to solve a problem using a specific algorithm is reduced by a factor proportional to the square of the input size. For instance, assume solving a problem necessitating a brute-force search and the prevailing algorithm exhibits a quadratic time complexity. If a novel algorithm can solve the same problem with linear time complexity, it would be considered a quadratic speedup.

enable more precise risk metric evaluations, such as value at risk (VaR), which is widely used specifically for the capital requirements of financial institutions.

Potential advantages of quantum mechanics in the financial sector

Graph 4



Source: Authors' elaboration.

Incorporating quantum computing into risk management enables a re-evaluation of how risk management is operated. It offers the potential to perform complex calculations at an unprecedented speed. Quantum algorithms developed for risk estimation in financial asset portfolios have the potential to accelerate calculations. In other words, as the input size of the problem increases, the best classical algorithms require an exponentially growing number of steps, whereas a quantum algorithm needs, at most, a linear number of steps. Egger et al (2020) estimate that a QC could perform the VaR calculations for a 1 million-asset portfolio in a run-time of 30 minutes. This is, however, dependent on several factors, such as the number of scenarios, the complexity of the portfolio, the selected VaR methodology and the efficiency of the implemented algorithms. In contrast, these specific calculations could take several hours using traditional computing.

Financial risk is frequently assessed through the utilisation of models such as Monte Carlo simulations. Zhang et al (2023) highlight the benefits of quantum computing in these simulations. These are crucial for modelling processes with inherent randomness. With quantum computing, financial risk assessment models can be enhanced, and hence prediction accuracy can be improved.

Quantum algorithms have been tested for accelerating pricing and risk analysis of financial derivatives. Stamatopoulos et al (2022) compare quantum, classical and semi-quantum circuits that allow a quadratic speedup of derivative pricing. The authors also investigate a method requiring fewer resources than existing methods. This advancement will significantly improve the process of evaluating potential risks in terms of required precision and run-time.

Quantum computing brings advantages to the financial system that could result in more accurate estimations and simulations. It also allows some complex problems to be performed that are intractable with the classical computing available today. Financial networks are so interconnected that it is exceedingly difficult to assess the impact of potential financial crashes such as asset price perturbations. Orus et al (2019) highlight this challenge, stating that the mathematical problem of forecasting a damaging event is unsolvable with classical algorithms, even for a small network. They argue that a problem involving a network composed of 20 to 30 organisations would take more than the age of the universe to be solved with classical computing.

Aboussalah et al (2023) envisage a quantum approach to systemic risk analysis in financial networks. They demonstrate how quantum computing techniques can provide insights into the resilience of financial systems. While quantum risk management is evolving at a fast pace and concrete applications are under active research, the potential improvements depend on the development of QCs' capacities. It is expected that further research will allow new quantum algorithms to be developed, making risk management one of the key application areas.

Investment and portfolio management

Portfolio optimisation is a delicate balancing act. It involves minimising risk and maximising return simultaneously. The complexity of these optimisation problems can result in highly complex calculations. In some cases, these can take classical computers days or months to compute, or even exceed their capabilities altogether. The emergence of QCs may offer new possibilities to solve these challenging problems, presenting a competitive advantage for investment institutions with access to quantum computing.

Bova et al (2021) introduce the concept of quantum advantage. They highlight the potential for QCs to generate more profitable outcomes for the most intractable problems. For instance, option pricing involves calculating the fair value of options based on various factors, including the underlying asset price, volatility, time to expiration and interest rates. Traditional methods provide approximate solutions but may not capture all market complexities accurately. With the enhanced computational speed of quantum systems, investment institutions could achieve superior results in portfolio optimisation.

Several researchers have proposed quantum algorithms for various portfolio management tasks. Doriguello et al (2021) focus on stochastic optimal stopping problems. Ghysels et al (2023) look at derivatives pricing. Both research teams demonstrate the accelerated computations enabled by quantum computing.

Rebentrost and Lloyd (2018) also showcase the applicability of quantum algorithms in portfolio optimisation.

The time it takes to complete a complex calculation is a variable to consider when assessing potential benefits of quantum computing. However, it is equally relevant to evaluate its potential economic advantage. While QCs offer enhanced computational speed, classical supercomputers may still provide a more cost-effective solution, as Bova et al (2022) point out. The demanding computational requirements of certain optimisation problems must be carefully evaluated. Building a QC is costly compared with classical computing, and this cost must be taken into account.

Payments and settlements

The payments sector is another area where quantum computing can bring significant advantages

McMahon et al (2022) explore an innovative use of quantum computing to improve liquidity when settling payment transactions. Their proposal could lead to faster, more efficient settlements. The experiment operated payment transactions on a 30-day sample from a Canadian payment system, providing a practical example of these benefits. They use a hybrid quantum annealing solver to test potential improvements in daily liquidity. The results are promising, with significant liquidity savings estimated between CAD 239.93 million and CAD 275.70 million. This demonstrates the tangible benefits that quantum computing can bring to the payments and settlements sector.

Macroeconomic modelling

Dynamic economic models, which provide a representation of the evolution of economic factors over time, pose a significant computational challenge. To provide an analogy, consider the challenge of solving a large and intricate puzzle. The larger and more complex the puzzle, the longer it takes to fit in each piece, given the need to test each piece individually against each of the others. This complexity is a point emphasised by Hull et al (2020), who note that computational costs rise exponentially with the number of variables. This challenge makes it difficult to accurately represent and solve these models with powerful traditional computing resources. In contrast, the ideal solution would involve testing all pieces simultaneously to understand how they fit together. This is akin to what a QC can perform.

Fernández-Villaverde and Hull (2023) demonstrate that today, quantum annealers can already solve certain dynamic programming problems an order of magnitude faster than classical computers. These authors implement a well-known economic model, namely the real business cycle (RBC), on a quantum annealer. This quantum machine – already available today, albeit with limited capacity – solved the RBC model in 3% of the computation time that a traditional computer would take. The paper introduces a new method to solve performing dynamic programming, displaying the potential of quantum hardware for economic modelling.

Quantum computing and AI: a new frontier in machine learning?

Researchers are actively exploring synergies between quantum computing and AI. A burgeoning field of research relates to the use of quantum computing algorithms to enhance AI and machine learning, which rely on classical computing. Potential domains within AI that could benefit from the integration of quantum computing include reinforcement learning, deep learning and support vector machines, among other areas.

Reinforcement learning is a type of machine learning that allows an agent (or an AI-driven system) to learn to make decisions by interacting with its environment (Li (2017)). The agent determines its next action based on its current state and receives feedback in the form of rewards or penalties after each action. Since the aim of the agent is to maximise rewards, its learning process includes finding actions that lead to the highest rewards through trial and error and developing a strategy that maximises rewards. Quantum reinforcement learning algorithms can leverage the principles of quantum mechanics to accelerate the learning and decision-making processes. These algorithms could introduce new approaches to exploring reinforcement learning problems, potentially resulting in more effective strategies for finding optimal policies. Lockwood and Si (2020) examine specifically the use of quantum circuits in reinforcement learning, highlighting how quantum computing can offer innovative methods for policy optimisation. Saggio et al (2021) provide empirical evidence of a quantum speedup in a reinforcement learning task, underscoring the potential of quantum algorithms to enhance reinforcement learning.

Deep learning, a subset of machine learning techniques, represents another area where quantum computing may enhance computational efficiency and thereby increase the capabilities of AI. Deep learning involves training artificial neural networks to analyse data and make decisions. These networks consist of layers, including an input layer, one or more hidden layers and an output layer. Each node acts as an artificial neuron, forming connections with others through assigned weights and thresholds. The term “deep” in deep learning refers to the network’s multiple layers. Despite its widespread success, deep learning faces challenges related to memory and time efficiency, which quantum computing could address. Quantum deep learning holds promise for enhancing the speed and efficiency of the training process, with potential advancements in quantum neural networks through the use of quantum gradient descent algorithms. These algorithms aim to expedite training by estimating gradients more efficiently through quantum principles (Wiebe et al (2014)).

Although integrating AI with quantum computing presents significant challenges, researchers are actively investigating this avenue to potentially augment AI capabilities. These two domains could bolster each other’s development, as highlighted by the growing body of research on the application of AI techniques in quantum computing development (Dunjko and Briegel (2018)). While realising practical applications involves surmounting considerable obstacles, such as costs and

technical limitations, the potential impact of the synergy between these two fields warrants attention.


4. Economic effect of QCs: from short-term benefits to a new general purpose technology?

Despite growing interest and new quantum algorithm proposals, the practical benefits of quantum computing in real-life problems are difficult to assess. Choi et al (2024) present a framework to understand where quantum computing will provide an advantage over classical computing and conclude that real-life use cases with a practical quantum advantage will be few. They argue that classical computers solve problems of small to moderate complexity faster than QCs. The authors point out that only problems where the algorithmic gain is large will benefit from quantum computing. They assert that, since very large algorithmic gains are "rare" in practice and theorised to be rare even in principle, their analysis suggests that the benefits from quantum computing will flow to a small set of use cases, and mostly to those processing very large amounts of data. They illustrate this with the example of an unstructured search where quantum advantages appear if the problem size exceeds 10^{12} (Graph 5.A).

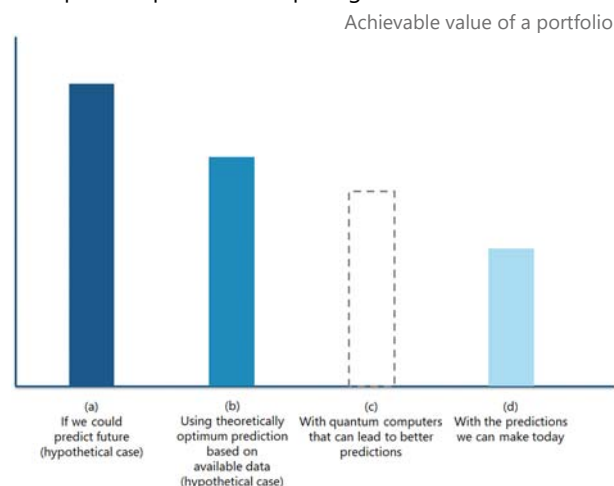
Impact of quantum computing

Graph 5

A. Quantum algorithm superiority for large problem sizes



B. Impact of quantum computing



Sources: Adapted from Choi et al (2024); authors' graphical analysis.

In addition to the algorithmic challenge, Biamonte et al (2017) identify two additional obstacles: data initialisation and output measurement. To illustrate these challenges, following Castelvechi (2024), the quantum computing process could be decomposed into three steps: (1) initialising the QC with data-encoded quantum states, (2) performing operations on these configured qubits and (3) extracting information from the final quantum state. While most quantum algorithms focus on

optimising step (2), steps (1) and (3) can be time-consuming. Initialising large data sets by converting them into quantum states can be inefficient,⁷ and extracting meaningful information from the final qubit string can also be computationally expensive, potentially overshadowing the speed gains achieved with quantum algorithms (Aaronson (2015)). Many of the use cases studied in the current literature address problems that are being solved today by classical computers. For such problems, the short-term benefits of quantum computing often come in the form of faster computation with similar results. For instance, Sanz-Fernández et al (2021) utilise a quantum computing algorithm to solve dynamic portfolio optimisation problems and find that quantum annealing completes the task in approximately three minutes, while the computational time using classical computers exceeds one day.

The short-term economic value of such a speedup varies between use cases. As a simple upper bound on the potential benefits of QC for any specific use case, one can consider how much value we would be able to realise with classical computers of infinite computing capacity. For example, in the case of an optimisation problem, this corresponds to finding the global optimum in zero computational time. For some use cases, this would have a large economic value. But in many practical use cases, the immediate value is limited, for example because of intrinsic uncertainties and noisy input data, which reduce the value of more accurate calculations, or because the computations need to be performed rarely, and a speedup does not directly convert to immediate economic value. These considerations also apply to the economic benefits achievable by solving the same problems with a QC (Graph 5.B).

One rough intuitive gauge of the short-term impact of QCs is the market that is most directly affected by them – the one for high-performance computing. This market is expected to grow to around USD 90 billion in 10 years, although there is high uncertainty in this projection (see Graph 6).⁸ By assuming an average increase in performance derived from quantum computing compared with classical computing, we can estimate potential savings. For example, if a 10% reduction in costs is assumed, this would imply a “saving” of USD 0.9 billion per year.⁹

⁷ Alternative approaches explore the utilisation of inherently quantum data obtained through methods such as quantum sensing. See Cerezo et al (2022) for details.


⁸ Analysts also estimate that the banking, financial services and insurance sector has the third-highest share as an end user in the market.

⁹ Indeed, this rough back-of-the-envelope calculation is well grounded in economic theory. Hulten (1978) shows that in production networks, incremental (ie up to a first order) technological improvements in different sectors contribute to overall economic growth in proportion to the size of the affected sectors and to sectoral productivity growth. The key underlying idea is that even if technological advancements occur in just a few sectors, they can still have significant impacts on the aggregate economy, depending on the economic weight of those sectors.

Projection for the global market size of a high-performance computer

In billions of US dollars

Graph 6



Note: "Average" indicates the average projections from the reports published by market survey companies.

Sources: Fortune Business Insights; Future Market Insights Inc; Grand View Research; MarketsandMarkets; Precedence Research; SNS Insider; The International Market Analysis Research and Consulting Group; Verified Market Research; author's calculations.

However, one needs to differentiate the short-term benefits of the technology and the long-run implications. Historical technology transitions offer an insightful parallel to the long-run economic potential of quantum computing and suggest that much larger economic effects may be expected in the long run. One case in point is the shift from steam engines to electric power in industrial settings. Initially, the adoption of electricity may bring modest improvements if a factory replaces a centralised steam engine without reconfiguring the broader infrastructure, ie it still distributes power in the factory via belts and shafts. However, the long-term effects can be much more profound, as wiring allows the factory to dramatically reshape the distribution of power, with electric motors automating a large set of smaller steps in the production process (Devine (1983), David (1990)). Another example is the internet, which provided short-term gains such as improved communication and information access. However, its long-term impacts have been monumental, enabling teams to work in distant settings, and hence also facilitating the evolution of global value chains (Cairncross (2002)).

The long-term implications of quantum computing could also be transformative. While many existing use cases see "only" a QC speedup, computational tasks that are not yet performed today because they are not feasible using classical computers may become possible with QCs. Through these use cases, quantum computing may revolutionise data processing and algorithmic calculations, leading to entirely new applications across multiple domains, including materials science, cryptography, AI and finance.

Thus, while the immediate benefits of quantum computing are modest, its potential to drive fundamental transformations in computational capabilities could mirror the widespread impacts seen with the electrification of factories, marking another significant leap in technological evolution.

5. Quantum implications on cryptography

Quantum computing and other emerging applications of quantum physics simultaneously present both challenges and opportunities for cryptography, and thereby for the security and stability of financial systems. In what follows, we first provide an overview of the cryptographic methods used in today's systems, followed by a description of the most important opportunities and challenges represented by quantum physics.

Cryptography in today's financial systems

Conventional cryptography uses algorithms based on complex mathematical problems to protect sensitive data. Two communicating entities usually rely on a shared secret key to protect their messages from eavesdropping. This is called symmetric key encryption. The sender uses the shared key to encode the information, and the receiver uses the same key to decode the received information and restore the original message. The algorithms used for encryption and decryption are standardised, and only the key is kept secret.

One challenge in symmetric cryptography is the need to create the shared secret between the communicating entities. Traditionally, they would have had to agree on a shared key in advance by physically transmitting the key from one endpoint to the other, for example on paper or on some electronic device. In today's digital systems, this initial step is usually replaced by creating a shared secret using public key cryptography, also known as asymmetric key cryptography.

Graph 7 illustrates the difference between symmetric and asymmetric key cryptography. In a symmetric key scenario, shown in the upper part of the figure, Alice and Bob agree to use a shared key for encrypting and decrypting their secret messages. Alice encrypts the secret messages using the shared key and sends them to Bob. Upon receiving the encrypted messages, Bob decrypts them using the same shared key.

In asymmetric key cryptography, shown in the bottom part of Graph 7, Alice and Bob each possess two different keys: a public key and a corresponding private key. As implied by its name, a public key is accessible to anyone, while a private key remains confidential. It is not feasible to infer a private key by analysing its public counterpart. Alice encrypts secret messages using Bob's public key. When Bob receives encrypted messages from Alice, he decrypts them using his private key. Public key cryptography does not require a shared secret between the communicating endpoints; hence, it avoids the need to transmit the key prior to secret communication. Instead, the communicating entities can simply publish their public keys, to be used by anyone willing to send them encrypted messages.

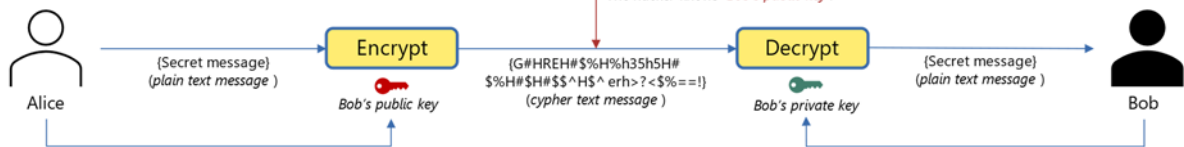
Symmetric key encryption

Encryption and decryption use the same key



Asymmetric key encryption

Encryption and decryption use different public key and private key



Source: Adapted from Deodoro et al (2021).

In public key cryptography, the challenge is: how can Alice be certain that the public key she uses belongs to Bob? What if a hacker, masquerading as Bob, sends their own public key to Alice to intercept the secret messages?¹⁰ To verify the authenticity of the public key, a public key infrastructure (PKI) that manages digital certificates and public key encryption is used. Digital certificates are electronic documents issued by a certificate authority to prove the match between an entity's identity and its public key. If Alice wishes to verify that Bob's public key indeed belongs to Bob, she needs to obtain a digital certificate that authenticates Bob's public key as legitimate. For Bob to receive his digital certificate, he must submit his identity and public key to a certificate authority, a trusted third-party organisation that issues Bob's certificate. This digital certificate includes a digital signature that validates Bob's ownership of his public key. The digital signature is created by encrypting the message using the private key of the certificate authority. Alice and others can confirm that the digital signature is authentic by decrypting the signature with the public key of the certificate authority, which is stored on their IT devices.

In today's financial systems, and indeed in most IT systems, security relies on a combination of symmetric key cryptography, public key cryptography and PKI. Most of the messages are encrypted by symmetric key cryptography because it is more efficient than asymmetric key cryptography. However, asymmetric key cryptography is used to create the shared secret between the communicating entities before a communication session starts. And PKI is used to verify the authenticity of public keys.

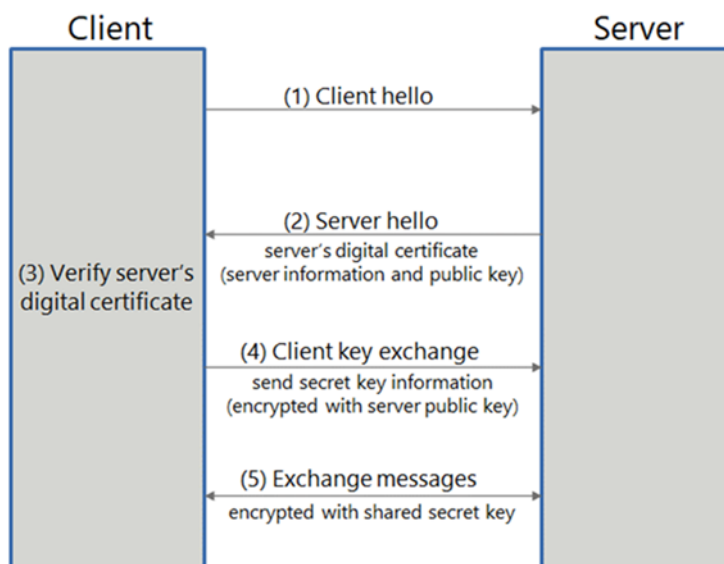
¹⁰ The risk of a hacker stealing secret messages using a fraudulent public key purported to be Bob's is inherent to the protocol illustrated in this example and does not necessarily stem from the development of QCs.

Graph 8 illustrates this combined solution using the example of HTTPS, a protocol used to secure online and mobile banking transactions. The graph shows a simplified sequence of establishing a secure channel between a client and a server within HTTPS, where the client is typically a customer, and the server represents a bank. The process unfolds as follows:

- The client initiates a secure connection by sending a “hello” message to the server.
- The server responds with its own “hello” message to the client, including its digital certificate, which contains the server information and public key.
- The client verifies the server’s digital certificate by decrypting it using the public key of the central authority, ensuring the certificate is indeed signed by the central authority.
- The client then sends information regarding the secret symmetric key, encrypted with the server’s public key.
- The server retrieves the shared symmetric key by decrypting it with its private key. Subsequently, both the client and the server can exchange messages encrypted with the shared symmetric key, completing the establishment of a secure connection.

Overview of an HTTPS connection

Graph 8



Source: Authors' graphical analysis.

Quantum threat to legacy cryptography

Both symmetric and asymmetric key cryptography rely on mathematical algorithms to encode and decode messages using the respective keys. The algorithms are constructed such that by examining the encoded version of the message, it is practically impossible to deduce the original message or the secret key. The emphasis

is on the word “practically”. In theory, there are algorithms that can break cryptography and find out the secret key, but with sufficiently large keys, this would require prohibitively large computation time, even for the largest supercomputers. Hence, conventional cryptography is generally considered sufficiently secure for today’s needs.

The emergence of QCs brings a change in this respect because they will be able to perform some of these computations faster than classical computers, potentially rendering current security obsolete. The problem impacts asymmetric key cryptography more than symmetric key schemes.

For symmetric key cryptography, a quantum algorithm called Grover’s algorithm can be used for attacks (Grover (1996)). However, this algorithm can achieve only limited speedup compared with attacks using a traditional computer. Therefore, security experts believe that legacy symmetric key encryption can continue to be used, provided that the key length is increased (Rao et al (2017)).

QCs represent a more profound impact to asymmetric key cryptography, where some existing algorithms will no longer be secure if scalable QCs exist. Specifically, public key cryptographic methods such as Rivest-Shamir-Adleman (RSA) and Elliptic Curve Cryptography (ECC), which are prevalent in today’s financial systems, are at risk of becoming obsolete with the advent of QCs.

RSA encryption is based on a mathematical problem called factorisation. To create the keys, the algorithm randomly generates two large prime numbers and then multiplies them together. The product of this multiplication forms part of the public key, while the two large prime numbers are used to generate the private key. The security of this algorithm comes from the one-way nature of this operation: multiplying two large numbers is easy, but factoring large numbers is very hard using classical computers, and with large enough numbers, it becomes practically impossible. However, the factoring problem is one of those areas where QCs will potentially significantly outperform today’s computers (Van Meter et al (2008)).

Mathematician Peter Shor demonstrated that a sufficiently large QC could perform factorisation in reasonable time and be used to break currently widely used encryption algorithms (Shor (1994)). In 2001, researchers demonstrated the practical feasibility of Shor’s algorithm, which further emphasised the potential of QCs to render some existing cyber security measures unsecure.

The ability of QCs to hack today’s cryptographic schemes has a profound impact on the security of financial systems. Deodoro et al (2021) identify the most vulnerable areas as: (1) online/mobile banking, (2) payment transactions (including cash withdrawals), (3) business-to-business privacy and (4) virtual private network (VPN) communications. In online/mobile banking, an attacker utilising a QC could compromise and eavesdrop on communications between users and financial institutions during authentication and authorisation processes. Similarly, during payment transactions or ATM withdrawals, such an attacker could intercept these exchanges. Corporations often use secure communication channels based on public key encryption; however, these channels could be fully accessed by attackers if

compromised with quantum computing techniques. VPNs used by employees working remotely could also be vulnerable to these attacks.

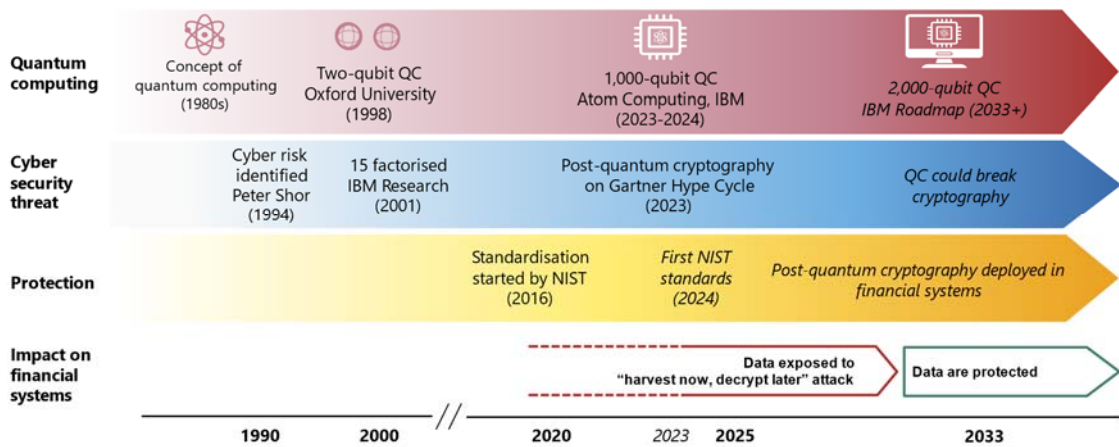
Quantum-resistant cryptography

The ongoing use of certain asymmetric key cryptographic methods in financial communications, specifically RSA and ECC, represents a looming risk to the security of financial systems amid rapid advances in quantum computing. The theoretical capacity of QCs to effectively breach these cryptographic algorithms calls for pre-emptive actions to safeguard the enduring security of financial operations.

In response to this threat, researchers have been developing new cryptographic algorithms designed to be resistant to QCs (Graph 9, yellow arrow). Somewhat counterintuitively, it does not take a QC to protect against attacks by a QC. It is possible to encrypt data using classical computers such that QCs will not be able to break the encryption. QCs will be very efficient at solving certain types of problems (such as factorisation), so the solution lies in designing new cryptographic algorithms based on mathematical problems that do not fall into these categories.

Quantum computer threats to financial systems

Graph 9



Sources: Vandersypen et al (2001); Gartner; IBM Quantum Roadmap; NIST; Stanford Encyclopedia of Philosophy Archive.

In 2016, the US National Institute of Standards and Technology (NIST) initiated a standardisation process for such cryptographic algorithms.¹¹ To denote such algorithms, the terms quantum-resistant or post-quantum cryptography are often used interchangeably. As part of this process, NIST announced in 2022 the first four encryption and digital signature algorithms selected for standardisation, with the first

¹¹ See NIST, "Post-quantum cryptography standardization", Cyber Security Resource Center, 13 August 2024, csrc.nist.gov/Projects/Post-Quantum-Cryptography/Post-Quantum-Cryptography-Standardization.

standards published in August 2024.¹² This ongoing process will proceed with the selection of additional algorithms offering varied approaches to encryption.

Though QCs do not yet exist, there is an urgent need to deploy quantum-resistant algorithms, specifically for highly sensitive data. This is due to the potential threat of so-called “harvest now, decrypt later” (HNDL) attacks (Graph 9, white arrow). Even though QCs capable of breaking current encryption are not yet available, malicious entities might still intercept financial data today, with the intention of decrypting it in the future. Until quantum-secure solutions are fully deployed, the vulnerability of highly sensitive data to HNDL attacks remains a critical concern. This looming threat highlights the urgent need for pre-emptive measures to safeguard the authenticity, integrity and confidentiality of data. Findings from a survey of central banks conducted in 2021 by the Global Cyber Resilience Group at the Bank for International Settlements (BIS) (Doerr et al (2022)) indicate that the financial system could face substantial losses if stakeholders do not prepare for HNDL type attacks.

Quantum cryptography

While QCs represent a threat to legacy cryptography, other applications of quantum physics bring a new opportunity to strengthen the security of IT systems, including the financial system. Quantum cryptography utilises the inherent principles of quantum mechanics such as superposition, uncertainty, coherence and the no-cloning theorem entanglement to create secured communications.

A pivotal element of quantum cryptography is quantum key distribution (QKD), which is based on quantum properties of particles, typically photons. This enables the parties of a communication channel to obtain a shared secret key with the assurance that any interception attempts will be detected (Aditya and Shankar Rao (2005)). A study conducted by Bank of Italy explored the potential integration of QKD into payment systems, shedding light on its potential implementation to enhance the security of current cryptographic systems (Tiberi and Buccioli (2023)). Nevertheless, the practical implementation of this promising technology currently poses some difficulties in terms of scalability and integration into existing infrastructures. The transmission of quantum information over long distances is still a challenge due to the lack of network infrastructures and significant implementation costs (BSI et al (2024)). However, this subfield of quantum information science is rapidly evolving and is considered a promising technology for the security of IT systems. Different government and regional initiatives have been launched with the construction of fibre or even free space quantum networks (Aditya and Shankar Rao (2005)).

Other applications of quantum cryptography include, for example, quantum digital signatures and quantum money. While these and other quantum cryptography

¹² See NIST, “NIST announces first four quantum-resistant cryptographic algorithms”, 5 July 2022, www.nist.gov/news-events/news/2022/07/nist-announces-first-four-quantum-resistant-cryptographic-algorithms and “NIST Releases first 3 finalised post-quantum encryption standards”, 13 August 2024, <https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards>.

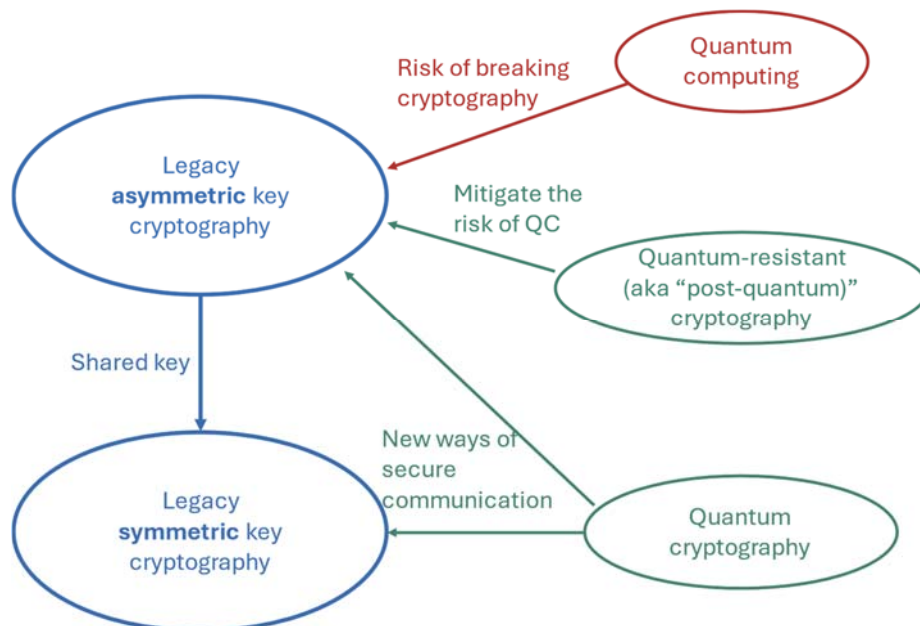
areas may represent new opportunities for the financial system in the long run, they are currently conceptual, and their potential impact is difficult to assess.

The terminology around quantum technologies and their impact on security can be somewhat confusing. For example, the names might suggest that “post-quantum cryptography” is something that comes after “quantum cryptography”, but this is not the case. Post-quantum cryptography refers to algorithms that can be implemented on classical computers, and these should be introduced before QCs become available, to replace vulnerable cryptography used today. In contrast, quantum cryptography relies on quantum technology and is likely to be further away in time.

In Graph 10 we summarise the most widely used terms and their interrelations. Legacy cryptography, shown with blue lines and fonts, consists of a combination of asymmetric key and symmetric key cryptography, where the former creates the shared secret to be used by the latter. Quantum computing represents a threat by its potential ability to break existing versions of asymmetric key cryptography. Quantum-resistant (aka “post-quantum”) cryptography is a new family of asymmetric key cryptographic algorithms which rely on mathematical problems that are hard for QCs to crack. This is a way to mitigate the risk coming from QCs. Quantum cryptography consists of new cryptographic methods, built on the concepts of quantum physics, which might support or replace some of the legacy techniques.

The impact of quantum technologies on cyber security

Graph 10



Source: Authors' graphical analysis.

6. Implications on financial system stability


The potential impact on QC cyber attacks

The stability of the financial system is fundamentally dependent on cryptographic security. The emergence of quantum machines, with their potential for advanced decryption capabilities, introduces significant risks that could lead to considerable economic losses. Although it is challenging to quantify these risks due to the uncertain scope and magnitude of the quantum threat, estimating potential losses is crucial for understanding its impact. For example, a particularly alarming scenario is a quantum-enabled cyber attack on critical financial infrastructures, such as Fedwire Funds Service, whose users include government agencies (Butler and Herman (2023)).

Projected economic losses from a systemically relevant cyber attack enabled by quantum computing¹

As a percentage of GDP

Graph 11



¹ The expected losses are derived by multiplying the magnitude of potential losses with the probability of quantum computing breaching RSA-2048 encryption within 24 hours over various time horizons. The estimation of loss size is based on a survey question directed to central banks, asking, "In your opinion, what is the maximum loss in % of annual GDP that a systemically relevant cyber-attack on traditional financial institutions could cause?" This query was posed to a sample comprising seven advanced economies (AEs) and 11 emerging market economies (EMEs). The probability estimates are obtained from a professional survey of field experts. The x-axis represents the projection horizon in years.

Sources: Survey on Doerr et al (2022); Mosca and Piani (2023)

Assessing the precise economic impact of such events is complex; however, a survey of central banks offers some insight into the scale of a systemically relevant cyber attack on financial institutions (Doerr et al (2022)). Based on a professional assessment of the likelihood of QCs breaking current standard encryption, the estimation indicates expected losses of about 0.1% of GDP on average within five years of the attack (Graph 11). The survey results suggest that, with an increasing likelihood of decryption over time, expected losses could surpass 1% of GDP within 15 to 20 years, with emerging market economies potentially facing even larger losses.

This estimation is expected to rise as financial system digitalisation accelerates, highlighting a significant concern for financial stability. Nevertheless, transitioning to quantum-resistant cryptography could greatly reduce these risks.

Central banks' defences against the quantum threat

The crucial role of central banks in ensuring the proper functioning of financial market infrastructures and in maintaining financial stability inevitably leads to the security and integrity of their operations – a topic of utmost importance (BIS 2023). Securing data, for example related to monetary policy formulation, is highly critical. Given the potential impact of quantum computing on the financial system, central banks are primarily affected and need to proactively prepare against this cyber risk in order to maintain trust and confidence in the financial system. Among other initiatives, BIS's Project Leap helps central banks accelerate this process (see Box A).

Responding to the cyber security threat posed by QCs requires a multifaceted approach. Investing in post-quantum cryptography is one of them, as these new algorithms that are under the NIST standardisation process aim to secure information against the computational capabilities of QCs. Migrating from existing cryptographic protocols to quantum-safe cryptography is a complex, time-consuming and resource-intensive process. A preparation phase needs to be planned in the shortest possible time frame, as a migration plan will take substantial time to be completed. When we consider these significant challenges posed by the implementation of new cryptographic algorithms, this transition to quantum-safe schemes becomes a priority and needs to be integrated into cyber security roadmaps. A quantum-readiness roadmap needs to be defined, including the important milestones of a migration plan and a high level of flexibility. Agility and change management principles will be key for a successful transition. With the perspective of the quantum era, IT systems will need to adapt continuously. Today, the most immediate approach for transitioning to new protocols involves the implementation of post-quantum cryptography, which is a branch of quantum-safe solutions. While this strategy enhances security, it is important to acknowledge its potential vulnerability. These new algorithms need to be extensively examined in concrete applications, possibly requiring ongoing adjustments. Furthermore, as QCs' capabilities continue to increase, the cyber security threat that quantum computing poses may evolve, requiring more agile systems. The ability of systems to rapidly switch from one algorithm or protocol to another is an important parameter to consider in the selection of new security frameworks. Relying on a unique protocol based, for instance, on asymmetric cryptography such as RSA protocols will no longer be possible.

One of the first steps when preparing for this cyber security threat is to create a dedicated team with the right competencies to formulate a quantum-readiness plan that considers the specifics of the organisation. Before concrete implementation, central banks must identify their needs regarding upskill training for experts and cryptography asset and critical data inventories.

Project Leap – quantum-proofing the financial system


Project Leap, a collaborative initiative by the BIS Innovation Hub Eurosystem Centre, Bank of France and the Deutsche Bundesbank is designed to fortify the financial system’s defences against potential cyber threats posed by quantum computers (QCs). Recognising the critical role of central banks in maintaining financial stability, the project aims at proactively preparing against novel cyber risks.

Transitioning to quantum-safe cryptographic protocols is a complex process, requiring both agility and robust change management. The project emphasised the need for a comprehensive quantum-readiness roadmap, including upskilling experts, thorough testing of new algorithms and flexible migration strategies that consider the intricacies of legacy systems.

Project Leap has already made significant strides in its first phase, implementing and testing post-quantum cryptography when setting up a virtual private network (VPN) between Germany and France. Initial testing demonstrated the secure transmission of financial data through a quantum-safe communication channel. However, it also revealed that the time needed to set up the VPN was affected. An example of the experimental results is illustrated in Graph A1 where the leftmost column called “legacy” represents that VPN setup time using classical cryptography, and the other bars correspond to various versions of quantum-resistant cryptography. The time performance differs between different algorithms, but they all represent an additional latency compared with legacy solutions. This impact is not problematic for VPN setup, which happens a few times during a workday, but it will have significant consequences for time-demanding use cases such as instant payment systems, indicating a necessary trade-off between security and performance.

In its forthcoming second phase, Project Leap plans to test quantum-safe cryptography in a more complex environment, closer to real-world applications and including additional partners.

Project Leap: testing the performance impact of post-quantum security when setting-up a VPN¹



¹ The red bar labelled IKE represents the time needed to exchange the asymmetric keys and for the authentication. The blue bar labelled CHILD represents the time needed to exchange keys.
Source: BIS (2023).

A transition plan needs to consider the complexity of legacy systems that often embed non-compatible cryptographic algorithms and protocols. The implementation of new algorithms can be technically complex and may require significant adaptation of existing systems, including fundamental changes to systems. The migration from existing encryption algorithms to quantum-safe cryptographic protocols conceived to resist QC attacks need to be deployed at a large scale. This transition involves updating a broad range of infrastructure components, including software, hardware and cryptographic libraries across various systems, which requires significant coordination within and across organisations. In addition, every impacted cyber security process and risk management framework needs to be adapted to the new protocols.

By proactively addressing these challenges well in advance, central banks can enhance the security of their digital systems and protect sensitive data in a quantum-enabled world. Once the preparation and planning phases are completed, the execution of the transition will follow with the concrete implementation of quantum-resistant cryptographic schemes to protect systems.

7. Conclusion

This paper discusses the new opportunities and challenges that QCs represent for the financial system. This new technology will not replace existing computers, and it will not be applicable in all areas of computation. But in some selected areas, it has the potential to significantly accelerate calculations. We survey possible application areas in the financial system, including risk management, investment and portfolio optimisation, artificial intelligence, payments and settlement, and macroeconomic modelling. Naturally, most use cases described in the current literature address known computational problems that are being solved today using classical computers.

We also discuss the potential economic implications of applying QC in the financial system. We conclude that initial gains are modest, partly because QCs are still at an early stage of development, and experiments are performed with small prototypes, partly because the advantages of QC over classical computers become apparent only in a specific set of computational problems. While early experimental use cases focus on computational problems successfully addressed today with classical computers, the biggest gains can be expected in future use cases that are unsolvable using classical computers. These will be identified gradually, and mostly when QCs are already available. It is possible that we will see only marginal economic gains in the early phases, but a transformational effect in the longer run.

We go on to discuss the specific impact that quantum technology has on cryptography and thereby on the stability of the financial system. Current cryptography is based on mathematical problems that are unsolvable using classical computers, but that can potentially be solved by future QCs. This means that a future QC can break the encryption schemes protecting our financial system, thereby representing a threat to financial stability. While the benefits of QC will take time to develop, the cyber security threat is imminent due to the possibility for malicious

actors to intercept and store data today with the intention to decrypt it in the future when QCs become available.

We discuss central banks' responsibility in addressing this challenge, and the steps that need to be taken to mitigate the risk. Finally, we look at how quantum technology can have a positive impact on security, for example by using quantum key distribution, a promising technology to distribute secret keys to communicating endpoints, solving a major challenge in IT systems today.

References

- Aaronson, S (2015): "Read the fine print", *Nature Physics*, vol 11, no 4, pp 291–3.
- Aboussalah, A M, C Chi and C-G Lee (2023): "Quantum computing reduces systemic risk in financial networks", *Scientific Reports*, vol 13, no 3990.
- Aditya, J and P Shankar Rao (2005): "Quantum cryptography", mimeo, Stanford University.
- Arute, F et al (2019): "Quantum supremacy using a programmable superconducting processor", *Nature*, vol 574, pp 505–10.
- Atom Computing (2023): "Quantum startup Atom Computing first to exceed 1,000 qubits", press release, 24 October.
- Bank for International Settlements (BIS) (2023): *Project Leap: quantum-proofing the financial system*, June.
- Biamonte, J, P Wittek, N Pancotti, P Robentrost, N Wiebe and S Lloyd (2017): "Quantum machine learning", *Nature*, vol 549, pp 195–202.
- Bova, F, A Goldfarb and R Melko (2021): "Quantum computing is coming. What can it do?", *Harvard Business Review*, July.
- (2022): "Quantum economic advantage". *NBER Working Papers*, no 29724.
- BSI, ANSSI, NLNCSA and Swedish Armed Forces (2024): *Position paper on quantum key distribution*.
- Butler, A and A Herman (2023): *Prosperity at risk: the quantum computer threat to the US financial system*, Hudson Institute, April.
- Cairncross, F (2002): "The death of distance", *RSA Journal*, vol 149, no 5502, pp 40–2.
- Castelvecchi, D (2024): "The AI-quantum computing mash-up: will it revolutionize science?", *Nature*, January.
- Cerezo, M, G Verdon, H-Y Huang, L Cincio and P Coles (2022): "Challenges and opportunities in quantum machine learning", *Nature Computational Science*, vol 2, no 9, pp 567–76.
- Choi S, W Moses and N Thompson (2024): "The quantum tortoise and the classical hare: a simple framework for understanding which problems quantum computing will accelerate (and which it will not)", arXiv:2310.15505.
- Chuang, I, N Gershenfeld and M Kubinec (1998): "Experimental implementation of fast quantum searching", *Physical Review Letters*, vol 80, no 15, p 3408.
- David, P (1990): "The dynamo and the computer: an historical perspective on the modern productivity paradox", *American Economic Review*, vol 80, no 2, pp 355–361.
- Deodoro, J, M Gorbanyov, M Malaika and T S Sedik (2021): "Quantum computing and the financial system: spooky action at a distance?", *IMF Working Papers*, no 2021/071, March.

Devine, W (1983): "From shafts to wires: historical perspective on electrification", *Journal of Economic History*, vol 43, no 2, pp 347–72.

Doerr, S, L Gambacorta, T Leach, B Legros and D Whyte (2022): "Cyber risk in central banking", *BIS Working Papers*, no 1039, September.

Doriguello, J, A Luongo, J Bao, P Rebenrost and M Santha (2021): "Quantum algorithm for stochastic optimal stopping problems with applications in finance", arXiv:2111.15332v1.

Dunjko, V and H Briegel (2018): "Machine learning & artificial intelligence in the quantum domain: a review of recent progress", *Reports on Progress in Physics*, vol 81, no 7, 074001.

Dyakonov, M (2018): "The case against quantum computing", *IEEE Spectrum*, accessed July 2024.

Egger, D, C Gambella, J Marecek, S Mcfaddin, M Mevissen, R Raymon, A Simonetto, S Woerner and E Yndurain (2020): "Quantum computing for finance: state-of-the-art and future prospects", *IEEE Transactions on Quantum Engineering*, vol 1, no 3101724.

Fernández-Villaverde, J and I Hull (2023): "Dynamic programming on a quantum annealer: solving the RBC Model", *CESifo Working Paper*, no 10500, June.

Feynman, R (1982): "Simulating physics with computers", *International Journal of Theoretical Physics*, vol 21, no 6/7, pp 467–88.

Ghysels, E, J Morgan and H Mohammadbagerpoor (2023): "Quantum computational algorithms for derivative pricing and credit risk in a regime switching economy", arXiv.2311.00825.

Grover, L (1996): "A fast mechanical algorithm for database search", in *STOC'96: proceedings of the twenty-eighth annual ACM symposium on the theory of computing*, pp 212–19.

Hull, I, O Sattath, E Diamanti and G Wendin (2020): "Quantum technology for economists", arXiv:2012.04473.

Hulten, C (1978): "Growth accounting with intermediate inputs", *Review of Economic Studies*, vol 45, no 3, pp 511–18.

IBM (2024): IBM Quantum: Development & Innovation Roadmap.

Jones, J, M Mosca, R Hansen (1998): "Implementation of a quantum search algorithm on a nuclear magnetic resonance quantum computer", *Nature*, vol 393, pp 344–46.

Kalai, G, Y Rinott and T Shoham (2022): "Google's 2019 'quantum supremacy' claims: data, documentation, and discussion", arXiv:2210.12753.

Li, Y (2017): "Deep reinforcement learning: an overview", arXiv:1701.07274.

Lockwood, O and M Si (2020): "Reinforcement learning with quantum variational circuit", in *Sixteenth AAAI conference on artificial intelligence and interactive digital entertainment*, vol 16, no 1, pp 245–51.

Ménard, A, I Ostojic, M Patel and D Volz (2020): "A game plan for quantum computing", *McKinsey Quarterly*, February.

McMahon, C, D McGillivray, A Desai, F Rivadeneyra, J-P Lam, T Lo, D Marsden and V Skavysh (2022): "Improving the efficiency of payments systems using quantum computing", *Bank of Canada Staff Working Papers*, no 2022-53, December.

Montanaro, A (2016): "Quantum algorithms: an overview", *npj Quantum Information*, vol 2, no 15023.

Mosca, M and M Piani (2023): *Quantum threat timeline report 2023*, Global Risk Institute and evolutionQ Inc.

NEC (2024): "Successful demonstration of a superconducting circuit for qubit control within large-scale quantum computer systems", press release, 3 June.

Nielsen, M and I Chuang (2010): *Quantum computation and quantum information*, Cambridge University Press.

Orus, R, S Mugel and E Lizaro (2019): "Forecasting financial crashes with quantum computing", arXiv:1810.07690.

Pistoia, M et al (2021): "Quantum machine learning for finance ICCAD special session paper", in *2021 conference proceedings: international conference on computer-aided design (ICCAD)*, November.

Rao, K R, D Mahto, D K Yadav and D A Khan (2017): "The AES-256 cryptosystem resists quantum attacks", *International Journal of Advanced Research in Computer Science*, vol 8, no 3, pp 404–8.

Rebentrost, P and S Lloyd (2018): "Quantum computational finance: quantum algorithm for portfolio optimization", arXiv:1811.03975.

Saggio, V, B E Asenbeck, A Hamann, T Strömberg, P Schiansky, V Dunjko, N Friis, N C Harris, M Hochberg, D Englund, S Wölk, H J Briegel and P Walther (2021): "Experimental quantum speed-up in reinforcement learning agents", *Nature*, vol 591, pp 229–33.

Sanz-Fernández, C, R Hernández, C Marciniak, I Pogorelov, T Monz, F Benfenati and R Orús (2021): "Quantum portfolio value forecasting", arXiv:2111.14970.

Shor, P (1994): "Algorithms for quantum computation: discrete logarithms and factoring", in *Proceedings of the 35th annual symposium on foundations of computer science*, pp 124–34.

Stamatopoulos, N, G Mazzola, S Woerner and W Zeng (2022): "Towards quantum advantage in financial market risk using quantum gradient algorithms", *Quantum*, vol 6, p 770.

Tiberi, P and E Buccioli (2023): "Quantum safe payment systems", Bank of Italy, *Markets, Infrastructures, Payment Systems Working Paper*, no 35, June.

Van Meter, R, K Itoh and T Ladd (2008): "Architecture-dependent execution time of Shor's algorithm," *Controllable quantum states: mesoscopic superconductivity and spintronics (MS+ S2006)*, pp 183–8.

Vandersypen, L, M Steffen, G Breyta, C Yannoni, M Sherwood and I Chuang (2001): "Experimental realisation of Shor's quantum factoring algorithm using nuclear magnetic resonance", *Nature*, vol 414, pp 883–7.

Wiebe, N, A Kapoor and K Svore (2014): "Quantum deep learning", arXiv:1412.3489.

Yogendran, B., Charlton, D. Beddig, M., Kolotouros, I., Wallden, P., (2024), Big data applications on small quantum computers. arXiv:2402.01529

Zhang, Y, Y Huang, J Sun, D Lv and X Yuan (2023): "Quantum computing quantum monte carlo", arXiv:2206.10431.

Previous volumes in this series

No	Title	Issue date
BIS Papers No 148	Keeping the momentum: how finance can continue to support growth in EMEs	September 2024
BIS Papers No 147	Embracing diversity, advancing together - results of the 2023 BIS survey on central bank digital currencies and crypto	June 2024
BIS Papers No 146	Central bank capital and trust in money: lessons from history for the digital age	June 2024
BIS Papers No 145	Generative artificial intelligence and cyber security in central banking	May 2024
BIS Papers No 144	The economic implications of services in the metaverse	February 2024
BIS Papers No 143	Central banking in the Americas: Lessons from two decades	November 2023
BIS Papers No 142	Inflation and labour markets	November 2023
BIS Papers No 141	Will the real stablecoin please stand up?	October 2023
BIS Papers No 140	Central banks, macro-financial stability and the future of the financial system	October 2023
BIS Papers No 139	Digital safety nets: a roadmap	September 2023
BIS Papers No 138	Financial stability risks from cryptoassets in emerging market economies	August 2023
BIS Papers No 137	Building an integrated surveillance framework for highly leveraged NBFIs – lessons from the HKMA	July 2023
BIS Papers No 136	Making headway – Results of the 2022 BIS survey on central bank digital currencies and crypto	July 2023
BIS Papers No 135	The energy transition and its macroeconomic effects	May 2023
BIS Papers No 134	Global tightening, banking stress and market resilience in EMEs	April 2023
BIS Papers No 133	The two-regime view of inflation	March 2023
BIS Papers No 132	Information governance in sustainable finance	December 2022
BIS Papers No 131	Central banking after the pandemic: challenges ahead	December 2022
BIS Papers No 130	Pricing of climate risks in financial markets: a summary of the literature	December 2022

All volumes are available on the BIS website (www.bis.org).