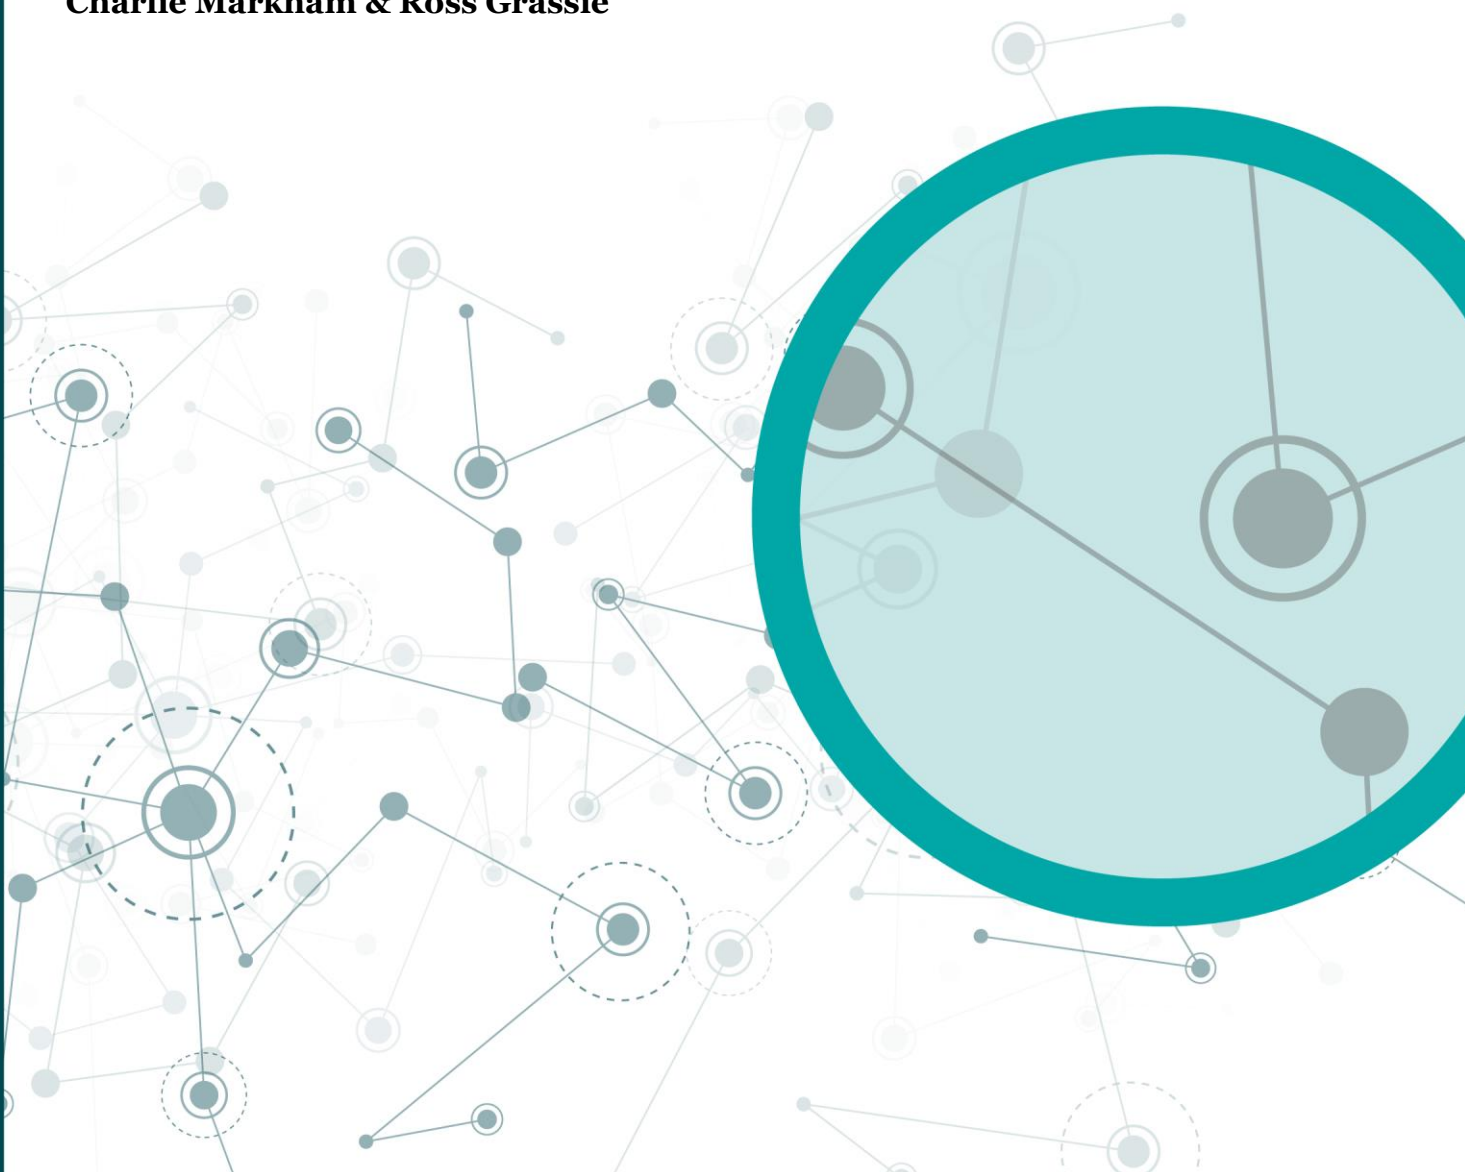


Research Note

03/10/2025

Quantum Computing Applications in Financial Services

Charlie Markham & Ross Grassie



FCA research notes in financial regulation

The FCA research notes

The FCA is committed to encouraging debate on all aspects of financial regulation and to creating rigorous evidence to support its decision-making. To facilitate this, we publish a series of Research Notes, extending across economics and other disciplines.

The main factor in accepting papers is that they should make substantial contributions to knowledge and understanding of financial regulation. If you want to contribute to this series or comment on these papers, please contact David Stallibrass (David.Stallibrass@fca.org.uk).

Disclaimer

Research notes contribute to the work of the FCA by providing rigorous research results and stimulating debate. While they may not necessarily represent the position of the FCA, they are one source of evidence that the FCA may use while discharging its functions and to inform its views. The FCA endeavours to ensure that research outputs are correct, through checks including independent referee reports, but the nature of such research and choice of research methods is a matter for the authors using their expert judgement. To the extent that research notes contain any errors or omissions, they should be attributed to the individual authors, rather than to the FCA.

Authors

Charlie Markham (Financial Conduct Authority) & Ross Grassie (formerly the Quantum Software Lab)

Acknowledgements

We would like to thank colleagues from the Quantum Software Lab and the National Quantum Computing Centre for their contributions in delivering this final research paper. We would also like to acknowledge the UK academics, quantum computing companies and financial services firms who took part in our interviews, workshops, and provided the insights required to produce this research note.

All our publications are available to download from www.fca.org.uk. If you would like to receive this paper in an alternative format, please call 020 7066 9644 or email publications_graphics@fca.org.uk or write to Editorial and Digital Department, Financial Conduct Authority, 12 Endeavour Square, London E20 1JN.

Equality and diversity considerations

We have considered the equality and diversity issues that may arise from the proposals in this Research Note.

Overall, we do not consider that the findings in this Research Note adversely impact any of the groups with protected characteristics i.e. age, disability, sex, marriage or civil

partnership, pregnancy and maternity, race, religion and belief, sexual orientation and gender reassignment.

Contents

1	Executive Summary	6
	Purpose	6
	Key findings	6
2	Quantum computing and financial services: A UK growth opportunity	8
3	Understanding the quantum computing landscape	10
	Foundations: from information to quantum information	10
	From quantum information to quantum computing	11
	The quantum computing stack today	12
	Layer 1: Hardware	13
	Layer 2: Error correction & decoders	13
	Layer 3: Circuits compilers & orchestration	14
	Layer 4: High-level abstractions (programming languages)	14
	Layer 5: Quantum algorithms and applications	14
4	The quantum computing landscape in financial services	16
	Building quantum readiness across financial services	16
5	Problem domain 1: Optimisation	18
	Defining optimisation	18
	The challenges for classical computing	18
	The potential for quantum advantage	19
	Many options all at once	19
	Capturing relationships between variables	19
	Emerging use case: Portfolio optimisation	20
	Why portfolio optimisation matters	20
	Classical approaches and their Limits	20
	Emerging quantum approaches	20
	Reframing the problem for Quantum approaches	20
	Leveraging quantum resources: superposition and entanglement	21
	The potential for quantum advantage	21
	Sentiment landscape	22
	Barriers to quantum advantage	23
6	Problem domain 2: Machine learning	25
	Defining machine learning	25
	The challenges for classical computing	25
	Data collection & pre-processing: Data quality and structure	25
	Training a model: Complex data, complex model	26
	Testing a model & evaluating the model: Generalisation	26

	The potential for quantum advantage	27
	Pre-processing: Quantum feature engineering	27
	Quantum models	28
	Emerging use case: Fraud detection	28
	Why fraud detection matters	28
	The computational challenge	29
	Classical approaches and their limits	29
	Emerging quantum advantage	29
	Reframing the problem for quantum approaches	30
	The potential for quantum advantage	30
	Sentiment landscape	31
	Barriers to quantum advantage	32
7	Problem domain 3: Stochastic modelling	34
	Defining stochastic modelling	34
	Focus on Monte Carlo	34
	An introduction to Monte Carlo	34
	The challenge for classical computing	35
	Computational Expense	35
	The potential for quantum advantage	35
	Reducing samples to accelerate convergence	35
	Emerging use case: Pricing models	36
	Why asset pricing matters	36
	Classical approaches	36
	Reframing Monte Carlo for quantum advantage	36
	Obtaining the final value	37
	The reality of the ideal	37
	Sentiment landscape	38
	Barriers and Limitations	39
8	Regulatory considerations	40
	Regulatory Readiness for Quantum Computing	40
	An early approach to surfacing regulatory insights	40
	Cross-cutting considerations	41
	Quantum model explainability	41
	Validation, benchmarking, and replicability	42
	Outsourcing, concentration, and operational resilience	42
	Market integrity, fairness, and competition	42
	Data governance, security, and confidentiality	42
	Regulatory Coordination	43
	Timing, proportionality, and triggers	43
9	The quantum opportunity: Going for growth	44
	Considerations on building quantum readiness	44
	Financial services firms	45
	Quantum computing vendors	46
	UK regulators	46

1 Executive Summary

Purpose

This research was undertaken to explore the potential applications of quantum computing in UK financial services, and to assess how firms and regulators can prepare for their emergence. It was triggered by the continued advancement of quantum technologies, and by the need to understand what their deployment could mean for financial markets, consumers, and the UK's position as a global financial centre. Now is a critical moment: while commercial applications are still emerging, progress is accelerating, and proactive regulatory engagement can help the UK capture opportunities while managing risks.

Key findings

- **Quantum computing in financial services presents a national growth opportunity:** The UK is uniquely positioned to reap the benefits of quantum computing applications in financial services, but realising this opportunity will require coordinated effort across industry, government, academia, and regulators.
- **Near-term commercial applications may be viable, but the full quantum stack needs attention:** Progress on qubit numbers, stability, and error correction is enabling potential early applications, but more focus must shift to software, algorithms, and integration to develop commercial applications.
- **Leading firms are already building quantum readiness strategies:** Some adopt a more use-case first approach (identifying business needs), while others take a technology-first approach (testing areas of potential quantum advantage). However, a blend of both is typically adopted, with firms exploring both quantum and quantum-inspired approaches.
- **Quantum applications for financial services broadly fall into three problem domains:** Optimisation, machine learning, and stochastic modelling are the most relevant areas, with all current approaches seeking advantage via hybrid quantum-classical methods.
- **The sentiment landscape varies across these domains:** Optimisation is seen as potentially promising but requiring more mature hardware; quantum machine learning is at an exploratory stage; and stochastic modelling offers a clear theoretical speed-up, which may not suffice for commercial advantage.
- **New regulations are unlikely to be required in the near term:** Potential, future quantum applications might intersect with established themes of explainability, fairness, and operational resilience that are already central in regulatory frameworks.
- **Quantum readiness relies on interdependent actions across the ecosystem:** The actions taken by financial services firms, quantum computing firms, and

regulators will not only shape their individual success but also influence the ability of other stakeholders to progress.

- **For financial services firms, readiness strategies can balance direction with flexibility:** Firms have different perspectives on the routes to value (first-mover advantage vs. second-mover advantage). However, pursuing “no” or “low” regret actions, such as skills development, can deliver near-term benefits while helping to manage long-term uncertainty.
- **Regulatory awareness is low among quantum computing vendors:** There is an opportunity for quantum computing vendors to build early relationships with regulators. These relationships can help establish technical and regulatory clarity, reassure end-users, and create a more predictable environment to support innovation.
- **Regulators could adapt or create innovation tools to support quantum development:** Existing mechanisms, such as sandboxes, could be expanded to give firms safe spaces for experimentation. Regulators should also consider whether specialist quantum computing teams are required.
- **A new Applications Regulatory Readiness Framework could support effective regulatory engagement:** This framework would assess the maturity of individual applications and match them with proportionate regulatory tools at each stage. It would help regulators support innovation, while signalling the UK’s leadership in adaptive, proportionate regulation of emerging technologies.

2 Quantum computing and financial services: A UK growth opportunity

The UK has set its sights on delivering strong, sustainable, and internationally competitive growth in the decade ahead. The government's Industrial Strategy commits to creating the conditions for growth by unlocking private investment, accelerating innovation, strengthening skills, and building the infrastructure for the UK's highest-potential sectors to thrive.ⁱ

The financial services sector is central to this vision. It is one of the UK's most productive sectors, contributing around 9% to the UK's total economic output.ⁱⁱ But its impact extends far beyond the sector itself, it provides both the capital and infrastructure other industries need to grow. Viewed from this perspective, the financial services sector is not only a driver of economic activity; it is a critical enabler of innovation across the wider economy.

The UK government is taking action to unleash this potential. Through the Financial Services Growth and Competitiveness Strategy, and the Leeds Reforms, it has set a clear course to rewire the financial system to support investment and drive growth across the UK.^{iiiiv}

At the same time, the FCA has set out an ambitious strategy for how we will work over the next five years to deepen trust, rebalance risk, support growth and improve lives.^v We will focus on 4 priorities: helping consumers fighting crime, supporting growth, and being a smart regulator. Success in one priority spurs success on others. By being a smarter regulator, ensuring we are predictable, purposeful, and proportionate we can also support growth by providing the sector with the clarity it needs to invest and innovate.

But enabling growth is not just about today's markets. It is about ensuring the UK is ready for the technologies that will shape tomorrow. Among all these potential opportunities, quantum computing stands out. It offers capabilities that could redefine how we solve complex problems beyond the reach of classical computing. And while competitors are striving to establish global leadership, the UK is entering this race from a position of real strength.

In 2024, the UK's National Quantum Strategy committed £2.5 billion over the next decade to position the UK as a quantum superpower.^{vi} It set out bold missions to scale quantum hardware, build a world-class supply chain, and develop globally leading talent. This commitment was further reinforced by an additional £670 million through the Industrial Strategy to accelerate not only the development – but crucially, the adoption – of quantum technologies.^{vii}

This investment is energising an already thriving UK quantum ecosystem. The UK is home to both world-leading universities and a specialised Quantum Software Lab that seeks to transform existing computational challenges into research problems. In addition, the UK has invested in a new generation of quantum hubs to drive industry-academic collaboration on real world use cases. At the core of this ecosystem, is the National

Quantum Computing Centre which hosts cutting-edge testbeds driving the evolution of different quantum computing hardware.

While the UK is building on a world-leading research, infrastructure, and public investment, the true measure of success will be whether the UK translates these foundations into real-world impact. That impact will only materialise if quantum computing is applied in areas of the economy where it can deliver advantages, and where there is both the scale and appetite to adopt it.

The UK's financial services sector is well placed to support this transition. Its expertise in complex modelling, optimisation, and high-performance computing make it an ideal proving ground for emerging quantum applications. At the same time, its appetite for innovation and competitive advantages can accelerate the development, testing, and commercialisation of this technology.

The application of quantum computing within financial services is not only a potential driver of economic growth, but an opportunity for the UK to consolidate its position as a global financial hub and demonstrate international leadership in financial innovation.

Realising this opportunity will require a coordinated effort from a diverse range of stakeholders. The UK quantum ecosystem – from researchers to vendors – must continue bridging the gap between scientific breakthrough and commercial application. To maintain global competitiveness, the UK financial services sector must invest in, experiment with, and test early applications. Meanwhile, UK regulators can act as catalysts for innovation. Early regulatory engagement can provide the clarity and confidence these stakeholders need to sustain investment, pursue experimentation, and ultimately adopt quantum computing applications when advantage materialises.

Success will require a shared understanding of the opportunities and challenges that lie ahead. This research paper supports that effort by providing an overview of perspectives from UK academia, quantum computing vendors, leading financial services firms, and UK regulators. Each brings a unique viewpoint: scientific foundations, technological feasibility, commercial opportunity, and regulatory consideration. By integrating these perspectives, this paper offers a collective view of the growth opportunity that quantum computing represents for UK financial services.

This paper begins with an introduction to quantum computing, outlining the core concepts and the current maturity of the quantum computing stack. It then explores how financial services firms are preparing for quantum computing, providing an overview of the strategies and approaches firms are taking to build quantum readiness. From there, the paper examines three problem domains where the financial services sector is exploring the potential for quantum advantage. These sections include emerging use cases, firms' sentiment on their viability, and the barriers preventing further adoption. It then looks ahead to how these future use cases might interact with the FCA's existing regulatory framework. Finally, the paper concludes with a set of immediate considerations for financial services firms, quantum computing vendors, and UK regulators in building quantum readiness. Taken together, it provides insights to help stakeholders across the ecosystem prepare for – and ultimately realise – the growth opportunity that quantum computing represents for financial services.

3 Understanding the quantum computing landscape

Quantum computing is complex. As a field of study, it draws upon aspects of advanced mathematics, quantum physics, and computer science. The complexity can easily give rise to uncertainty. Consequently, this uncertainty can quickly result in misunderstandings; specifically, misunderstandings about how quantum computers work, what they can do, and the impact they may have.^{viii}

To prevent this, audiences are often introduced to quantum computing through simplified analogies. In other cases, quantum computing is explained through domain-specific applications, enabling individuals to interpret quantum computing through its potential impacts.

This paper seeks to take an alternative approach. Whilst most of this paper focuses on quantum computing applications in financial services, it begins by outlining a set of foundations. The aim is not to replace the substantial education and training required to work in the field of quantum computing. Instead, these opening subsections aim to provide readers with an entry point: a collection of key concepts required to engage with this paper and in broader discussions on the topic of quantum computing in financial services.

The first in these foundations is an important distinction. The computers in use every day – from smartphones to High Performance Computing – are increasingly referred to as classical computers. This label serves to differentiate familiar digital technology from the new class of devices referred to as quantum computers.

The distinction reflects how each type of computer uses physics. Classical computers follow the familiar rules of classical physics, while quantum computers exploit quantum mechanics – the physics of the very small. The following subsections build on this distinction to provide readers with a foundational understanding of quantum computing.

Foundations: from information to quantum information

All computation relies on two primary factors: the ability to store information, and the capacity to process that information to produce an outcome.

In classical computing, information is stored in units, and the smallest possible unit of information is called a binary bit, or a bit for short. Bits are binary because they can only ever store one of two values: 0 or 1. This store of information is referred to as a state: if the bit holds the value of 1 it is referred to as being in the 1 state. Typically, these binary states are physically represented by an electrical signal: being either on for the 1 state or off from the 0 state.

When these units of information are grouped together, we refer to this information storage as memory. Readers may be more familiar with the units of megabytes (MB) or gigabytes (GB). These terms simply describe how many bits are grouped to store information: millions in the case of a megabyte, and billions for a gigabyte.

At first glance, quantum computing appears similar. In quantum computers, the store of information is a quantum bit, more commonly referred to as a qubit. But, as a unit of information, qubits fundamentally differ from binary bits. While binary bits can only ever exist in one of two states, a qubit can exist in a superposition of these states. This means a qubit can represent a combination of 0 and 1 states rather than being fixed to one value. Crucially, this does not mean the qubit is *uncertain* or *somewhere in between* these states. Instead, it holds both probabilities until it is measured.

When a qubit is measured, it resolves into one definite outcome – either a 0 or 1 state – but until that point, it contains the information of both states. By holding multiple potential values until measurement, qubits provide a flexibility that classical bits do not.

But superposition is not the only property that sets qubits apart, another is the principle of entanglement. In classical computing, the state of one bit can be described independently. There is no direct relationship between bits at the hardware level, so interactions are programmed at the software level.

In contrast, when qubits become entangled, these units of information form a direct relationship. The result is a single system, where the state of one qubit cannot be described without reference to the others. Where bits can be described independently – in isolation – entangled qubits can only be described in relation to the combined system they form. As a result, entangled qubits can become correlated in ways that classical bits cannot. Quantum computers use entanglement to generate complex correlations between their units of information that are prohibitively difficult to simulate classically at scale.

Together, superposition and entanglement define how qubits store and relate information. In turn, quantum computers and classical computers solve problems in fundamentally different ways. For specific classes of problems, classical computers can be inefficient or even unable to find solutions. In such cases, quantum computers may have an advantage by providing an alternative computational approach.

From quantum information to quantum computing

As introduced in the previous subsection, computation involves both storing and processing information. Having examined how qubits differ from classical bits as stores of information this subsection turns to how information is processed.

In classical computing, processing takes place through logic gates. These are simple operations, such as AND, OR, or NOT, that take one or more bits as inputs and produce a new bit as an output. Combining these gates builds a series of circuits that enable classical computers to perform any computation. Bits represent information as a 0 or 1, logic gates process that information, and the result is computation.

Quantum computers work in a similar way, but instead of logic gates, they use quantum gates. These act directly on qubits to change their quantum state. For example, a Hadamard gate can place a qubit into state of superposition, while a controlled-NOT (CNOT) gate, used in combination with a Hadamard, can entangle two qubits. By combining gates in sequence, quantum computers build circuits that form the instructions for running algorithms – just as classical computers do with logic gates.

Together, qubits and quantum gates provide the building blocks for quantum computation. But unlike their classical counterparts, qubits are extremely fragile. Classical bits are stored in relatively stable physical systems which can be isolated and controlled with high reliability. Qubits, in contrast, rely on delicate quantum states, which are susceptible to disturbances from their environment. Miniscule amounts of heat, vibration, or stray electromagnetic noise can disrupt a qubit's state; even the gates used to manipulate and measure qubits can introduce errors.

To address this fragility, researchers have developed methods of quantum error correction. The idea is to not rely on a single qubit, which is too unstable, but to use several qubits together to protect one unit of information. By arranging qubits into carefully constructed patterns, the system can detect which qubits have been affected and reconstruct the original information through error-correction schemes.

This process is not about eliminating errors entirely but about managing them so that the underlying quantum information remains intact for computation. In practice, this means that a single "useful" qubit of information requires many physical qubits to support it.

This introduces an important distinction between the physical qubits built into quantum computing hardware and logical qubits which are constructed through error correction. Logical qubits are fundamental to achieving useful quantum computing. Their development depends both on scaling the number of physical qubits and advancing error-correction techniques to reduce the physical qubit overheads required.

The ultimate goal is to develop enough logical qubits to construct a fault tolerant quantum computer – a system that runs algorithms reliably at scale to solve useful problems. Today's devices fall short of this vision, but they are rapidly evolving along multiple hardware pathways. To understand this progress, the next subsection explores some key parts of the quantum computing stack today: the layers of technology being developed to move from quantum computing theory to practice.

The quantum computing stack today

Turning quantum computing from a laboratory experiment into technology with real-world applications relies on developing an entire computing stack, including software, algorithms, and integration layers.

The quantum computing stack can be understood in five broad layers, each critical to making these systems commercially viable:^{ixx}

- 1. Hardware:** The physical qubits and controls systems that run quantum operations
- 2. Error correction and decoders:** The classical algorithms needed to make qubit operations reliable at scale.
- 3. Circuit compilers and orchestration:** The processes that translate high-level user-defined programmes into instructions that can run on hardware.
- 4. High-level abstractions (programming languages):** Tools that enable users to describe how a problem should be solved rather than how the solution will run on the device.
- 5. Quantum algorithms and applications:** The user-defined sequence of steps used to solve a problem on a quantum computer.

The following paragraphs examine each of these layers in more detail, drawing on research and industry engagement to assess their current state of development. Each layer is considered in terms of its role in the stack, its importance for commercial adoption, the progress achieved so far, and the challenges that remain. Altogether, these insights help illustrate the maturity of the quantum computing stack and highlight where continued effort will be required to move from experimental demonstrations to commercially viable solutions.

Layer 1: Hardware

To date, the vast majority of quantum computing development has concentrated on improving hardware: demonstrating stable qubits and scalable architecture. So far, there appears to be no singular route to achieving universal quantum computing. Instead, there is a large amount of investment into a few different physical systems – referred to as modalities – to generate qubits. These modalities include superconducting systems, trapped ion systems, photonic systems, and even more recently neutral atoms and silicon dots.

Whether one system becomes dominant will depend on factors, such as, their ability to produce a large number of stable qubits for calculation; whether they integrate effectively with existing computational infrastructure; and, the extent to which they can solve commercially relevant problems faster or more accurately than classical alternatives.

In recent years, the quantum computing hardware landscape has been described as noisy and intermediate scale (NISQ).^{xi} In other words, these machines produce a relative low number of qubits which are prone to error and instability, while still demonstrating some utility. However, more recent advancements in both quantum hardware and approaches to error correction suggest that current devices are moving towards a level of scale and stability to, in principle, tackle meaningful problems.^{xii,xiii,xiv}

These recent advancements are in part why there is a growing sense of optimism surrounding the nearer-term application of quantum computers in industrial settings.

Layer 2: Error correction & decoders

As discussed in the previous subsection, quantum systems are inherently fragile: environmental noise, imperfect control systems, and undesired interactions between qubits introduce errors into calculations. Quantum Error Correction (QEC) is the process of detecting and correcting errors without destroying the quantum state required for useful computation.

To function, these QEC schemes require decoders: classical algorithms that identify errors and determine which corrective transformations are required. These decoding algorithms run on classical devices that are closely integrated with quantum computers.^{xv}

While there has been rapid progress in QEC, even in the most promising demonstrations, the overheads remain immense. Achieving full fault-tolerance could require orders of magnitude more physical qubits to create a single error-corrected logical qubit.

Therefore, continued advancements in QEC requires more efficient decoders and the potential reduction of the physical-to-logical qubit ratio.^{xvi,xvii} Until this is achieved, quantum advantage will remain limited to specialised ‘noisy’ algorithms that can tolerate higher error rates.

Layer 3: Circuits compilers & orchestration

Moving up the stack, quantum programs are typically expressed as abstract circuits – a sequence of quantum gates and measurements – that assume arbitrary gates and connectivity between qubits. In practice, real devices have limited connectivity between qubits and constrained gate libraries. Circuit compilation is the process of mapping an abstract quantum circuit onto specific hardware architecture, which can require additional gates and operations, increasing execution time and error rates. This is a non-trivial problem requiring significant overheads and is currently performed on classical processors connected to the quantum device.^{xviii}

Orchestration is the next level of sophistication that logically flows from circuit compilation. In these hybrid quantum-classical workflows, orchestration determines which parts of the application (subroutines) should run on quantum hardware and which on classical resources. While orchestration frameworks are starting to emerge, they remain immature.^{xix} Achieving low-latency, tightly integrated execution will continue to be a critical component for achieving practical application and further adoption of quantum computing.

Layer 4: High-level abstractions (programming languages)

Most current quantum software packages available today require developers to write programmes at the circuit level, manually specifying individual gates. This is roughly equivalent to asking classical programmers to write in assembly code, which may be feasible for small examples but unsustainable for large-scale applications. The evolution of quantum hardware is already at a scale where this approach to programming is a bottleneck to further practical development.

High-level abstractions, such as domain-specific languages, libraries, and software development kits (SDKs), are required to enable developers to write software and algorithms at a conceptual level, leaving compilers to handle hardware-specific operations. This advancement will be essential to widen participation beyond quantum specialists and to enable applications that use thousands of logical qubits in the future. While progress is being made, these languages must evolve alongside hardware. On the one side, this co-evolution will ensure compatibility, performance, and ease of use as hardware matures. On the other, more accessible abstractions of quantum computing can improve accessibility, expanding the population of potential users.

Layer 5: Quantum algorithms and applications

The final layer of the stack is arguably the most important for the practical application of quantum computers: algorithm development. Quantum computers will not outperform classical computers because they are inherently faster or more accurate. Any advantage depends on marrying problems to quantum computational models. Only a small number of algorithms, such as Shor's algorithm for factoring numbers or Grover's search algorithm, have so far provided a theoretical advantage.^{xx,xxi}

Since the 1990's, relatively few new algorithms with clear quantum advantage over their classical counterparts have been discovered, and even fewer have subsequently been connected with commercially relevant problems to produce practical solutions. As quantum computing enters an era of hardware-driven commercial advantage, a more concerted effort is required to drive algorithmic innovation.

While each of these layers has its own development trajectory, they are deeply interdependent. Advances in hardware without matching progress in QEC, compilers, and algorithms will not translate into real-world impact. Despite the challenges evident in the current state of development across the quantum computing stack, there is cautious optimism surrounding near-term application of quantum computing and the prospect of commercial advantage in the coming years. Understanding where and how that advantage could first emerge is therefore essential.

Financial services is one of the sectors already exploring where this advantage may arise and how it could shape commercial operations. The next section examines the state of play for quantum readiness among UK financial services firms. Building on this overview, it turns to the problem domains where quantum advantage may develop, the emerging use cases being pursued, the broader sentiment landscape surrounding their deployment, and the barriers preventing wider adoption.

4 The quantum computing landscape in financial services

Building quantum readiness across financial services

For financial services firms, the central question with any emerging technology is not simply *can* it be used, but *when* should it be used. In the case of quantum computing, the consensus is that the technology is not yet ready for deployment. However, the potential commercial advantage is significant enough that leading firms are investing now to ensure they are prepared.

This preparation phase can be understood as building *quantum readiness*, and it involves developing specialist capabilities, running experiments, and forming an educated perspective on when and where quantum computing can provide value.

Most firms engaged in this research have established relatively small teams to lead this preparatory work. These teams focus on both building technical understanding – and as a byproduct – exploring the potential for *quantum-inspired approaches*. These methods utilise quantum principles to develop algorithms and methods that can be adapted to run on classical hardware. Firms pursuing these approaches are seeking to capture practical near-term performance improvements, and in the process build the skills and knowledge required to identify and capture future advantage as real quantum hardware matures.

In some cases, financial institutions are collaborating with quantum technology providers and academic partners to improve their understanding and accelerate quantum readiness. In these cases, some firms are conducting fundamental research into specific algorithmic components that may have present and future relevance in finance.^{xxii,xxiii} In others, firms focus on delivering targeted proofs-of-concept (PoCs) tackling use cases such as portfolio optimisation, fraud detection, or derivative pricing.^{xxiv,xxv,xxvi} These PoCs do not seek to demonstrate quantum computing's immediate commercial advantage. Instead, they build an organisation's understanding of how to capture the potential short- and long-term advantages of quantum computing for specific use cases.

This research indicates that approaches to PoC vary by institution, but two broad strategies have emerged. The slightly more favoured approach is technology-first. This approach starts with known or hypothesised quantum advantages from academic literature and works backwards to identify potential financial applications. Alternatively, a use-case-driven approach begins with a defined business problem and explores whether quantum computing could provide a better solution than classical methods. While the technology-first approach carries more risk, advocates argue that beginning with quantum's inherent strengths is more likely to yield considerable results when the technology matures.

Some institutions, particularly those interested in Quantum Machine Learning (QML) and the convergence of quantum with AI, are embedding quantum teams within their broader AI groups. This integration fosters knowledge exchange between the two fast-evolving fields and enables parallel testing of quantum and AI solutions on the same business

problems. These parallel runs of quantum, quantum-inspired, and classical solutions help to outline what the potential quantum advantage is, how it can be captured, and when quantum maturity may make this advantage commercially viable.

Other institutions are taking a more cautious approach to building quantum readiness. For these firms, the focus is on exploring a range of modalities to understand which may best suit their needs, while avoiding potential vendor lock-in through experimentation. In most cases, however, quantum-focused teams are also involving their legal departments early to address contractual and intellectual property issues that could otherwise create bottlenecks or friction in potential collaboration opportunities.

Relative to most sectors across the economy, leading firms in financial services have developed an advanced understanding of quantum computing's potential and limitations. As quantum computing hardware continues to mature, there are pockets of optimism surrounding the near-term application of quantum computing in financial services use cases. However, the scarcity of quantum algorithms that can effectively and efficiently solve commercially relevant financial problems remains as an additional obstacle. Overcoming this will require continued academic collaboration, targeted PoCs, and critically drawing on advances from other sectors where quantum research is also progressing rapidly. Firms that can work horizontally across industries to identify transferable quantum methods may be best placed to capture early quantum-advantages as the technology reaches maturity.

While firms are building quantum capabilities and experimenting with early-stage solutions, commercial adoption will depend on quantum computing addressing business-critical challenges more effectively than existing tools. Identifying these *problem domains* – areas where quantum computing has the potential to deliver tangible advantage – is the key to bridging the gap between readiness and real-world deployment.

In financial services, three domains stand out: optimisation, machine learning, and stochastic modelling. The next sections explore each in turn, setting out their relevance to the sector, how quantum methods might provide an advantage, emerging use cases and the sentiment landscape surrounding their deployment, and finally the barriers that continue to hold back their progress.

5 Problem domain 1: Optimisation

Defining optimisation

Optimisation problems involve finding the best possible configuration from a set of decisions, under a given set of constraints. In financial services, these problems tend to recur across various aspects of the business, from selecting the ideal mix of assets in a portfolio, to optimising trading schedules, to managing liquidity across the global financial network.

These optimisation problems are often expressed as combinatorial optimisation problems. In other words, the number of potential solutions grows exponentially with the size of the problem. This growth in complexity means that finding the optimal solution is often computationally impractical.

The challenges for classical computing

For optimisation problems in financial services, the challenges for classical computing approaches tend to occur across scale, complexity, and solution quality:

- **Scale:** For most combinatorial problems, the number of possible solutions grows exponentially as the variables increase, making exhaustive search for an optimal solution impractical. In many cases, poor scaling is not simply a matter of computing power, it is a symptom of how the problem is encoded. Some encodings require a vast number of steps or comparisons to explore the search space, which forces practitioners to simplify the model.
- **Complexity:** Financial problems often involve intricate, non-linear relationships between variables, including correlations, constraints, and market behaviours that are difficult to capture in simplified models.
- **Non-optimal solutions:** Due to complexity and scale, classical heuristic algorithms produce *good enough* solutions that are not truly optimal. These simplifications are often necessary to provide solutions within idealised timeframes but may miss opportunities. In the case of financial services, these opportunities may amount to risk-adjusted returns or cost savings.
- **Time constraints:** Even if better solutions can be found, they may take too long to find. In scenarios where speed is a factor for commercial advantage, the potential benefits of a more optimal answer are undermined.

Over decades of development, classical computing approaches have become remarkably effective at managing these trade-offs. In most cases, these models produce *good enough* solutions that meet practical and commercial needs by balancing speed of return, confidence in the result, and low-cost at scale. In many financial contexts, these solutions are not just adequate but highly refined, enabling firms to operate with confidence in complex markets.

The potential for quantum advantage

Regardless of their efficacy, classical approaches to optimisation are still bounded by computational constraints. These limitations force trade-offs and stem from the way optimisation problems are formulated and solved classically. Therefore, the core opportunity for quantum advantage in optimisation lies in tackling the scaling challenge for classical approaches.

Poor scaling is not simply a matter of insufficient computing power; it reflects the limits of how classical approaches encode optimisation problems and search for solutions. Quantum computing offers a new way to frame and solve optimisation problems, with the potential to overcome these limitations.^{xxvii}

Many options all at once

Central to this approach is superposition. In classical approaches one candidate solution is evaluated at a time, a process that becomes increasingly time-consuming and computationally expensive as the number of possibilities increases. In contrast, quantum approaches use qubits in superposition to encode and represent many of the possible solutions simultaneously.

Each potential solution is assigned a small probability of being the outcome of the computation. Quantum circuits are designed to diminish the probability of unpromising outcomes while amplifying the probability of promising solutions. This mechanism, which is inaccessible to classical systems, allows for a more efficient exploration of complex search spaces. For example, Grover's search algorithm demonstrates how a quantum system can reduce the number of search queries needed to identify a marked item within an unstructured dataset – a scaling solution that could be applied to broader optimisation tasks.

Capturing relationships between variables

A second and equally important resource is entanglement. Classical optimisation methods must often trade-off model accuracy and computational feasibility as correlations and constraints between variables grow more complex. In contrast, quantum computers can exploit entanglement to capture how variables are related to one another.

When qubits are entangled, the state of one qubit depends on the state of another. This makes it possible to represent correlations and constraints between variables within the quantum state itself. Because these relationships are encoded into the quantum state rather than approximated through additional computational steps, quantum approaches can account for greater complexity without incurring the computational costs associated with classical approaches. In principle, this means that more of the real-world complexity can be encoded without rendering the problem computationally infeasible.

Together, these features present a pathway for potential advantage. By addressing scale and complexity in ways that classical computing cannot, quantum approaches have the potential to produce more optimal solutions and reduce the time required to find them. The challenge, then, is understanding how these theoretical advantages might translate into practice within financial services.

Emerging use case: Portfolio optimisation

Why portfolio optimisation matters

Portfolio optimisation is a core function of investment and asset management, shaping how firms allocate capital, manage risk, and ultimately achieve returns. It is the process of selecting the right mix of assets and deciding how much to invest in each. The objective is to strike a balance between maximising expected returns while controlling for risks, transaction costs, and other constraints.

For financial institutions, achieving optimal results can translate directly into competitive advantage. More effective portfolio optimisation can boost revenue through improved risk-adjusted performance and the efficient allocation of capital. It also reduces costs by limiting unnecessary trades while minimising operational expenses.

Portfolio optimisation is a highly complex computational problem. It must account for the intricate relationships between assets, which can shift over time and are influenced by a mix of market dynamics, macroeconomic forces, and behavioural factors. This combination of scale, complexity, and business impact is why portfolio optimisation has emerged as one of the most actively explored use cases for quantum computing in financial services.

Classical approaches and their Limits

Classical approaches typically frame portfolio optimisation as a continuous optimisation problem. This means the task is not to simply decide whether to include an asset, but to determine what fraction of the portfolio to allocate to each. To do this well, models must also account for interactions between assets. This is because the risk of a portfolio is not merely the sum of each individual asset's risks, it also depends on how those assets move in relation to one another. If two assets are highly correlated, holding both increases risk; if they are negatively correlated, the risk of holding one may offset the other and reduce overall risk.

Finding the optimal portfolio can require searching across all possible combinations of allocations while respecting constraint such as cost or risk. As the number of assets grows, the number of possible combinations can increase dramatically, and the correlation between these assets adds a further layer of complexity.^{xxviii}

Emerging quantum approaches

Early experiments with quantum approaches to portfolio optimisation focus on the industry pain points where existing tools hit their classical limits. Specifically, quantum approaches target portfolios with complex asset interactions that classical solvers cannot handle efficiently.

Reframing the problem for Quantum approaches

To do this, quantum approaches reframe portfolio optimisation into a quadratic unconstrained binary optimisation (QUBO).^{xxix} When formulated in this way, each asset is represented as a binary variable where 1 represents its inclusion in the portfolio and 0 its exclusion.^{xxx}

A QUBO model then assigns scores to each combination of these 1s and 0s. Part of this score reflects the value of each individual asset, such as expected return. The score is

also adjusted based on how pairs of assets move together: if two assets are highly correlated, including both lowers the score, while negatively correlated assets are awarded a higher score for diversification. The optimisation task then becomes finding the combination of 0s and 1s (asset inclusions and exclusions) that deliver the best trade-off between maximising returns and limiting risks. Classical computers can solve limited QUBO problems, but as the number of assets grows the calculations become increasingly resource-intensive, as the number of potential combinations increases exponentially.

Leveraging quantum resources: superposition and entanglement

In contrast, QUBO problems naturally align with how quantum computers represent information. When the decision to include an asset in a portfolio is formulated as a binary yes or no (1 or 0 respectively) a qubit can be used to represent that choice.

Crucially, unlike a classical bit, a qubit can exist in a state of superposition. This allows quantum computers to hold and manipulate many possible combinations of 1s and 0s, or, in the case of portfolio optimisation different configurations for a portfolio, at the same time. This creates the potential to explore multiple portfolio configurations in parallel.

An equally important feature in this use case is entanglement. As discussed in the *Potential for Quantum Advantage* subsection, entanglement allows quantum computers to encode correlations into the quantum state. In the case of portfolio optimisation, this enables the quantum system to represent how different assets amplify or offset each other's risk. As a result, diversification effects are built into the representation of portfolios, rather than approximated through additional computation.

In short, by reformulating portfolio optimisation as a QUBO, quantum computers can use superposition to represent many possible portfolio configurations at once, and entanglement to capture interdependencies between assets.

At this stage, however, all the possible configurations are still weighted the same regardless of their quality. To transform this breadth of possibilities into an optimisation output, the system must be guided so that portfolios with better risk-return trade-offs are favoured.

One approach for this is the Quantum Approximate Optimisation Algorithm (QAOA), which provides a method of steering the system towards stronger portfolio configurations.^{xxxi} Over repeated iterations, QAOA increases the probability of sampling stronger portfolios, while weaker ones become less likely to appear. When the system is finally measured, the portfolio that emerges most often is the closest to the optimal balance of risk and return.

The potential for quantum advantage

The potential for quantum advantage comes down to how we reformulate optimisation problems to match the strengths of quantum computers. Classical approaches search sequentially through a solution space, a process that becomes increasingly demanding as complexity grows. By contrast, quantum approaches rethink the problem and reformulate it to leverage the fundamentals of superposition and entanglement.

The idea is that this reformulation could deliver two forms of advantage. First, as problems scale, quantum approaches may reduce time to solution by enabling a more

efficient exploration of the search space. Second, it offers the potential to improve solution quality, finding answers close to the true optimum where classical solutions may settle for *good enough*.

As with many of the emerging use cases explored in this paper, at this stage, quantum approaches do not seek to replace classical methods, but rather aim to complement them. For example, classical optimisation solvers remain best for continuous allocation decisions, such as fine-tuning the weights of chosen assets in a given a portfolio. Where quantum approaches add value is in the combinatorial core of the problem: deciding which subset of assets to include in the portfolio. The result is a hybrid workflow, where quantum approaches help to identify the most promising asset combinations for a portfolio, and classical methods refine their exact allocation.^{xxxii}

Sentiment landscape

Given the promise of potential advantages presented by quantum approaches to optimisation, it is unsurprising that this area has attracted sustained attention from the financial services sector. Yet, the current reality is more measured: with firm interest tempered by pragmatism. Whilst some firms see value in continued exploration, most recognise that meaningful deployment is likely to be a medium- to long-term prospect.

Although current quantum optimisation algorithms running on present quantum hardware are believed to outperform brute force search, they fall short of the best classical solvers in financial contexts. This is to be expected: these classical tools have been honed over decades, combining advanced classical algorithms with deep domain expertise to deliver solutions that are fast, cost-effective, and highly tailored to specific problems.

The bar for quantum is therefore high, and it is unlikely that a generalised, quantum optimisation solution will outperform these mature classical approaches in the short-term. Instead, progress in quantum optimisation solutions may depend on continued research into quantum algorithms to discover new approaches that are specifically designed for financial use cases.

In the absence of new algorithms, the pathway for quantum advantage will rely on identifying target brute-force subroutines embedded within classical optimisation methods. In theory, applying quantum approaches directly to these subroutines could improve performance. In practice, however, identifying the right target subroutines and integrating quantum solutions without degrading the overall efficiency of the solver is a significant challenge. Again, this is especially true in financial services, where classical optimisation models are highly tuned to specific problems and delicately balance speed, accuracy, and cost.

The limitations of this approach highlight that integration remains a significant barrier to the real-world application of hybrid quantum-classical approaches. To make these approaches commercially viable, data must move between quantum and classical systems with minimal latency, supported by robust middleware and orchestration tools to automate the workflow. Without this level of integration, any theoretical quantum speed-up is quickly outweighed by computational delays, coordination costs, and upfront investment.

Across all these potential avenues for quantum advantage, cost remains a decisive factor. A theoretical quantum improvement in accuracy or speed is only valuable if it translates into tangible gains, such as higher returns, better risk-adjusted outcomes, or operational savings that outweigh initial costs. In the current market, classical solvers already deliver these benefits at a relatively low cost.

Ultimately, from the perspective of financial services firms, the case for adoption hinges less on whether quantum is better in principle. What matters is whether it can deliver competitive results in practice. The challenge is that classical solvers already deliver these advantages at a relatively low cost. Bridging the gap from practical advantage to commercial advantage is a critical step in the pathway from development to adoption. With classical approaches holding a significant head start, and continuing to evolve, the bar for quantum adoption remains high.

Given that quantum advantage has yet to be demonstrated in practice, some financial institutions have adopted a more cautious stance. These firms are actively exploring perceived, nearer-term quantum applications while monitoring progress in quantum optimisation. Even so, proof-of-concept projects continue to play an important role. These partnerships between financial services firms and quantum technology help to build quantum readiness by clarifying where quantum approaches might eventually offer a competitive advantage.

Barriers to quantum advantage

Despite early interest in quantum portfolio optimisation, several persistent barriers continue to prevent meaningful progress.

- **Hardware scale and error rates:** Current quantum computers are constrained by the number of logical qubits they can reliably maintain for computation. In principle, the number of qubits required for portfolio optimisation scales linearly with the number of assets. In practice, however, each logical qubit requires many physical qubits to manage error and noise. Current hardware falls short by several orders of magnitude, meaning that even modestly sized portfolios are beyond reach until both error rates improve and qubit counts significantly increase.
- **Encoding asset weights:** Portfolios rarely involve simple yes/no decisions about whether to hold an asset. Instead, they require deciding how much to invest in each asset. Converting these continuous weights into the discrete form required for quantum algorithms introduces significant complexity. Many of the early experiments tend to under-analyse this conversion step and focus on where quantum approaches to optimisation may provide an advantage. However, this conversion is not trivial. This process adds a notable overhead that reduces the potential performance gains expected from quantum approaches.
- **Connectivity constraints:** Effective portfolio optimisation accounts for the correlations between all assets in the portfolio. In quantum terms, this means that each qubit may need to interact with every other qubit, this is referred to as “all-to-all connectivity”. While certain hardware modalities such as ion traps and neutral atoms can provide this type of interaction between qubits, others such as superconducting and photonic cannot. This narrows the range of suitable platforms and could constrain the future development of this use case.

Together, these barriers mean that portfolio optimisation remains a use case that is theoretically attractive, but with significant gaps between what is possible with current hardware and what is required for financial institutions to adopt and deploy these practices.

6 Problem domain 2: Machine learning

Defining machine learning

As a problem domain, machine learning (ML) is about building models that use data to make predictions about new, unseen data. These models are trained on historical data, refined to capture relevant information, such as patterns or anomalies, and then applied to new data to make predictions.

ML is a deeply embedded process across financial services, powering applications from fraud detection and anti-money laundering to portfolio analytics and risk forecasting.

ML can take many different forms such as, supervised, unsupervised, generative, or reinforcement learning. This variety means it can be difficult to pinpoint where specific challenges arise for classical approaches, and where quantum computing might provide an advantage. To make this discussion clear and structured, this report uses a simplified supervised classification workflow as an illustrative example.

This workflow can be broken down into five high-level stages:^{xxxiii}

- 1. Data collection:** Gathering data that will be used to train and test the model.
- 2. Pre-processing:** Cleaning and preparing the data and attaching labels (e.g. A/B outcomes).
- 3. Training a model:** Using the labelled data to develop a model (in this case a classifier).
- 4. Testing a model:** Applying the trained model to unseen data to assess predictive accuracy.
- 5. Evaluating the model:** Comparing predictions against actual outcomes to assess the model's performance and utility.

This section returns to this simplified workflow throughout to help anchor discussions on the challenges for classical computing and where quantum approaches could provide an advantage.

The challenges for classical computing

Despite widespread success, ML in finance services still faces many persistent challenges. When viewed from the simplified, five-stage workflow introduced above, three challenges stand out.

Data collection & pre-processing: Data quality and structure

At the earliest stages, the quality of the dataset defines the limits of what a model can achieve. Financial datasets are often noisy, sparse, and imbalanced, with many financial services applications seeking to identify rare but high-impact events, such as fraudulent transactions or sudden market crashes.

When the data is noisy, sparse or imbalanced, models risk classifying noise as if it were a real pattern (overfitting) or fail to capture rare-but-critical events as they appear too infrequently in the training data (underfitting). These imbalances in the dataset make it harder for models to learn meaningful patterns and produce desired outcomes.

Training a model: Complex data, complex model

If data quality and structure define what a model can learn, model training defines how effectively a model can perform. The challenges at this stage originate from the complexity of financial datasets which are often high-dimensional, with features that are interdependent and non-linear in their relationships. Complex data leads to complex models which require more computational power and time to train.

Turning back to the sources of complexity, financial datasets are often high-dimensional. For example, in a financial dataset, each row may represent a transaction and each column a characteristic of that transaction, such as, the amount, the sending and receiving account numbers, and the time of transaction.

These characteristics are referred to as features, and financial datasets may contain thousands of features. To produce accurate and useful outputs, a model has to identify and learn the underlying patterns and relationships that exist between features. However, in financial data, these relationships are interdependent and non-linear. This means the effect of one feature may depend on the others, and these effects can be disproportionate or unexpected.

Capturing these interactions requires large, sophisticated models capable of representing complex dependencies across many dimensions. Training such models creates a computational burden, involving a great number of calculations and repeated training iterations. As models grow in size and complexity, bottlenecks in training time, cost, and performance begin to emerge. In practice, these computational constraints can force trade-offs between model accuracy and development speed.

Testing a model & evaluating the model: Generalisation

Even when a model performs well on training data, the critical test is whether it can produce accurate results on new, unseen data. This ability – referred to as generalisation – determines whether a model can provide useful insights beyond the training set.

This is particularly challenging in financial services, where data evolves in real time, and patterns can shift rapidly as new behaviours emerge. As models are trained on historical data, they may not account for events or behaviours that are either absent or poorly represented in the historical training data. This limitation creates a tension for models: on one side rigidity produces accurate but time limited outputs, on the other flexibility risks overfitting for noise instead of identifying accurate patterns. As a result, models in financial services must often undergo regular testing and retraining to ensure that they continue to provide useful outputs as data distribution change.

The emergence of quantum computing presents the potential to address some of these classical challenges. The field that explores whether quantum approaches can overcome these limitations, and where they may provide an advantage, is known as Quantum Machine Learning (QML). Typically, QML seeks to identify the potential advantages of using quantum algorithms to process classical datasets.

The potential for quantum advantage

Advantage in the context of quantum approaches to machine learning could mean several things: higher accuracy in predictions, shorter training times, faster inference, the ability to model more complex relationships, or even improved generalisation. Compared to portfolio optimisation, performance improvements are both more nuanced and spread across a larger range of potential metrics.

For this reason, the following subsections return to the simplified ML workflow as a structure for exploring where quantum computing approaches may offer advantages.

Pre-processing: Quantum feature engineering

Introducing feature engineering

As explored in the *Challenges for classical computing* subsection, building effective models depends both on the quality of training data, and how that data is represented to the model. While increasing the amount of features can improve a model's predictive power, this comes at a computational cost.

This trade-off is managed through feature engineering, the process of creating, transforming, or selecting features to input into a model. ^{xxxiv} Broadly, the aim of feature engineering is to create a set of features that reduce the computational overheads of model training, while preserving the model's predictive power.

There are many different approaches to feature engineering. In its simplest form, it may amount choosing a smaller subset of features from the available population- referred to as feature selection. More sophisticated approaches may combine features together to capture the underlying patterns more effectively; this is called as feature extraction. In either case, the end goal is the same, reduce the number of features while maintaining the model's performance.

The open question is whether the problem of feature engineering can be reframed to exploit the strengths of quantum computing, and if so, whether quantum approaches can provide an advantage.

Recasting feature engineering for quantum

One potential approach for quantum advantage may lie in the transformation problem associated with feature extraction. In this approach, the aim is to construct features so that the patterns in the data become easier to detect. Classically, this is achieved by applying mathematical transformations to features, casting them into new spaces where output classes are more separable.

Quantum approaches take this further by encoding features into complex quantum states. These transformations are both non-linear and high-dimensional, enabling features to be represented in spaces that are unachievable for classical methods. The intuition is that by embedding features into these richer spaces, subtle patterns that are hard to separate classically become more distinct.

Ultimately, quantum approaches to feature engineering seek to provide a new set of features that benefit a classical model. These features may improve models by enhancing generalisation or lowering retraining costs. In practice, early experiments demonstrate that quantum approaches can generate feature sets that differ from classical approaches.

Whether this difference translates into advantage, in terms of model speed, accuracy, or training efficiency, is still an open area of research.

Quantum models

Early industry and academic exploration have also focused on the later stages of the ML workflow, most notably on developing quantum models. These are models that run natively on quantum hardware rather than using quantum computers to augment classical methods. These approaches encode classical data into quantum states where a quantum model can exploit the properties of the quantum system to both represent and process information differently from classical models. The key question is whether these models provide a performance advantage compared to their classical counterparts.

Quantum models depart from approaches to quantum feature engineering in how they use a quantum computer. In quantum feature engineering a quantum computer is used to generate a new set of features that feed back into a classical ML workflow. In contrast, quantum models use the quantum computer for inference. A set of features are not just encoded into a quantum state but also processed by a quantum model to produce outputs. When operating in this state space, quantum models may have access to richer representations of the data and thus the potential to capture more complex relationships at lower resource costs.

Several categories of quantum models are under active study, including, but not limited to: Quantum Support Vector Machines (QSVMs), Quantum Neural Networks (QNNs), and Quantum Autoencoders.^{xxxv,xxxvi,xxxvii} The aim of these quantum models is to embed complexity into quantum states to produce models that are comparatively more compact, expressive and generalise better than their classical counterparts.

The open question of quantum advantage

However, it is not yet conclusive whether quantum models can- or will- outperform classical approaches. Part of this uncertainty is due restricted experimentation. Current quantum computers lack the stability and circuit depth required to run effective quantum models. Beyond hardware, the deployment of quantum models faces various barriers across the quantum computing stack. For example, loading large-scale classical datasets into quantum states remains a persistent challenge for a several quantum computing use cases including quantum modelling.

These technical limitations mean that most of the early experimentation amounts to classical simulation of quantum models. While this is an important step in developing an understanding of how quantum models may perform, it does not provide a useful context for performance comparison. Taking these technical limitations into account, quantum models are generally seen as a long-term prospect, requiring substantial scientific and engineering experimentation to materialise.

Emerging use case: Fraud detection

Why fraud detection matters

Fraud poses a persistent and evolving threat to the integrity of global financial systems, and a major opportunity for firms to reduce losses. Strategies to address fraud typically fall into three categories: prevention, detection, and mitigation. Fraud detection is the

most readily automatable, which explains the sustained investment in algorithmic solutions seeking to improve efficiency and accuracy in identifying suspicious or fraudulent activity.

The computational challenge

At its core, fraud detection is a binary classification problem, it is about categorising datapoints into two classes: fraudulent and legitimate transactions. Fraudulent transactions, however, are an anomaly, typically amounting to less than 1% of transactions in a dataset. The challenge, then, is that there are limited samples of fraud to train the classification model. This difficulty can manifest in the model missing threats – referred to as underfitting – or miscategorising normal transactions as fraud, known as overfitting.

Fraud is also an adaptive behaviour. This means that new types of fraud can emerge rapidly in an attempt to outmanoeuvre detection systems. This is a constant game of cat and mouse where fraud detection models are continuously retrained to capture new forms of fraud and maintain the model's performance. As such, effective fraud detection strategies almost always rely on the ability to quickly retrain models to account for the ever-shifting patterns of fraudulent behaviour.

Therefore, most fraud detection strategies require two central components: models that can distinguish rare events within a large, noisy dataset, and the ability to quickly retrain models to capture the ever-shifting patterns of fraudulent activity.

Classical approaches and their limits

Classical approaches to fraud detection typically use ML models such as LightGBM and XGBoost.^{xxxviii,xxxix} These models are preferred because they offer a favourable balance of predictive accuracy and speed. However, even marginal improvements in detection can translate into significant financial value, and there is evidence that quantum approaches could unlock such gains. By leveraging quantum state space, QML has the potential to capture subtle, non-linear relationships between features more effectively than classical methods.

Retraining is also a serious constraint. Because fraud typologies evolve quickly, models must be updated frequently. Yet retraining at scale is computationally expensive and time intensive, creating cost and performance trade-offs. Research suggests that quantum-engineered features could reduce retraining times, allowing fraud models to adapt faster and sustain effectiveness.

Emerging quantum advantage

Quantum approaches to fraud detection seek to provide advantages either in reducing the retraining overheads for classical models or developing quantum models which, under certain conditions, might yield performance improvements.

In the case of reducing retraining times and resources, quantum approaches focus on the pre-processing stage of the ML workflow, specifically feature engineering. As discussed in the *Potential for Quantum Advantage* subsection, the key idea is to identify which approaches to feature engineering can be translated into problems well suited for quantum computers. In one such example, feature engineering may seek to generate a smaller subset of features that maintain the levels of descriptive accuracy while reducing

the complexity of the model. In this case, feature engineering becomes an optimisation problem, which is well-suited to leverage quantum approaches for a potential advantage.

Beyond pre-processing, academic literature has explored the potential to reformulate some fraud detection models into quantum models. The approaches vary widely, both in the type of advantage under investigation and in the quantum resources required to develop the model. The common thread, however, is the exploration of whether unique feature spaces available to quantum models could amount an advantage in the model's performance.

Reframing the problem for quantum approaches

The recurring foundation for potential quantum advantage arises from the use of quantum feature maps. Quantum feature maps embed data into a higher-dimensional space, making patterns easier to detect. Classically, this might involve plotting a transaction with two features on a two-dimensional chart. In the quantum case, the same variables are used as parameters of a quantum circuit. The quantum state generated by the circuit is then used to represent the transaction. This representation allows the application of uniquely quantum approaches to determine whether a transaction is fraudulent.

When transactions are represented in a quantum state, quantum models such Quantum Support Vector Machine (QSVM) can be used to classify between fraudulent and non-fraudulent transactions.^{xi} Support Vector Machines (SVMs) classify data by drawing decision boundaries, guided by a kernel function that measures the distance between points in a feature space.

In the classical example, two transaction data points are mapped into the feature space represented by the two-dimensional chart. The kernel function measures the distance between the two points in this space, and this information is used to draw a decision boundary in the two-dimensional chart. If a new data point is on one side of the boundary, they are labelled fraudulent. On the other side, they are labelled non-fraudulent.

Higher dimensional feature spaces can simplify the decision boundaries. For example, if all fraudulent transactions are in a circle around the non-fraudulent transactions in two dimensions, a circular decision boundary is needed. However, if a new feature is added, which allows for a three-dimensional model of the data, and the fraudulent and non-fraudulent transactions separate along this dimension, it allows for a simple linear decision boundary to be drawn.

By using quantum states to describe transactions, QSVMs embed transactions into a quantum feature space that is exponentially large in the number of qubits. In theory, this richer representation could enable a more straightforward and more effective separation of fraudulent and non-fraudulent transactions.

The potential for quantum advantage

Any potential advantage of quantum feature maps is significantly constrained by a phenomenon known as exponential concentration. As the number of qubits increases, the values used to compare data points tend to converge, making it increasingly difficult to extract meaningful distinctions. This issue can arise from overly complex feature maps or

hardware noise, resulting in a model that produces trivial outputs unrelated to the training data.

In practice, this means that any quantum advantage depends on highly specific conditions. Feature maps must be carefully designed: hard for classical systems to simulate yet structured enough to avoid concentration and tailored to solving the problem at hand. Without this balance, the benefits of a quantum approach are unlikely to materialise.

Sentiment landscape

QML provides a wide range of possible entry points for quantum methods to provide a potential advantage within existing computational workflows. Whilst this remains an area of active academic and scientific research, there is a growing interest from financial services, thanks to the many potential avenues for quantum advantage. However, unlike optimisation or stochastic modelling, QML does not yet offer a collection of known algorithms with targeted speed-ups for quantum approaches.

Instead, the landscape can currently be described as exploring “the art of the possible”: applying quantum approaches to different stages in the ML workflow, observing outcomes, and assessing whether these could constitute meaningful advantage.

As discussed in the *Potential for quantum advantage* subsection, initial optimism had centred on the ideal of building quantum models that could potentially offer richer expressivity and better generalisation than classical models.

In practice, this vision has been slowed by hardware limitations. Training ML models requires large datasets, but current quantum computers lack the capacity to process data at this scale. Even before training begins, encoding classical data into quantum information is a major bottleneck. Existing strategies either demand qubit numbers far beyond what is available or circuit depths that current hardware cannot sustain.

Beyond data loading and scalability, inference time is another critical limitation. Many financial services applications, such as fraud detection, require predictions in microseconds. Meeting this standard would demand continuous access to a quantum computer with both the stability and throughput to rival classical systems. Current quantum computers offer neither the uptime nor the seamless integration with broader financial services infrastructure to make this viable.

These real-world constraints have pushed firms to instead focus on earlier stages of the ML workflow. In particular, the pre-processing and model training stages where the requirements are less demanding. The working hypothesis is that quantum methods may surface different or richer features, and in some cases reduce training or retraining time. This is the current focus for PoC approaches, with partnerships between financial services firms, academia, and quantum computing providers testing whether quantum feature engineering or quantum feature mapping can produce small but consistent improvements in model accuracy or reductions in retraining costs.

To some extent, these PoCs are constrained by limits in data availability. Many quantum teams sit within innovation units with limited access to sensitive – but incredibly useful – financial data, forcing reliance on open or synthetic datasets. While this is largely a self-

imposed compliance precaution rather than a regulatory barrier, it creates the risk that PoCs fail to reflect real-world patterns or produce misleading results.

The prevailing stance among firms interested in QML is one defined by exploration. Experimentation focuses on hybrid approaches: introducing quantum methods at targeted stages while maintaining efficient classical stacks across the rest of the pipeline. QML is therefore seen less as a pathway to near-term deployment and more as a testbed for experimentation. The value of this exploration is not merely scientific; it is also viewed as return on investment in terms of quantum readiness. Through this exploration, firms hope to improve their organisational capabilities, identify potential partners in the quantum ecosystem, and develop a practical understanding of where quantum advantage may emerge as the technology stack matures.

Barriers to quantum advantage

While QML has emerged as one of the leading problem domains for exploration by financial services firms, both theoretical and practical barriers still stand in the way of adoption. These limitations explain why activity remains in the proof-of-concept stage and why the path to production will require advancement across hardware, software, and algorithm design.

- **Theoretical constraints:** The academic community is increasingly examining where QML might realistically deliver an advantage.^{xlix} While quantum models can, in theory, capture patterns that classical models cannot, their greater complexity also makes them harder to train effectively. At the same time, researchers are finding ways to replicate some of the potential benefits of QML using classical methods- a process known as “dequantization”.^{xliii} These developments do not eliminate the possibility of quantum advantage, but they suggest the window where it can be achieved is narrower than once assumed.
- **Data Loading:** QML also has practical constraints that limit its development to deployment. Effective machine learning typically depends on large amounts of data, yet today’s quantum hardware can only handle limited capacity.^{xliv,xlv} Moving classical data into quantum systems is currently constrained by hardware that has not yet reached the maturity for scale. Current approaches involve trade-offs: either requiring a greater number of qubits with simplified processing, or fewer qubits with more complex processing. In both cases, the balance is dictated by hardware limitations rather than business needs, making scalability an unresolved challenge.
- **Slow Quantum-Classical Interconnects:** Most QML solutions today also rely on hybrid approaches, where classical optimisers guide quantum circuits.^{xlvi} This structure introduces long training cycles, as the system must repeatedly switch between classical and quantum resources. The data exchange between these environments is not yet efficient, creating additional bottlenecks.^{xlvii} While similar overheads exist in conventional high-performance computing, the issue is more acute in quantum systems because the physical connections themselves are still under development.

QML remains a leading area of exploration in financial services; however, the theoretical and practical barriers, as well as their interaction, make its trajectory difficult to track. On the theoretical side, identifying where genuine quantum advantage can be achieved is

an actively researched area, with the picture changing rapidly. On the practical side, data loading limits, hardware-driven trade-offs, and inefficient quantum–classical interconnects keep activity in the proof-of-concept stage for now. Yet progress is accelerating on many fronts, making QML a dynamic and closely watched field.

7 Problem domain 3: Stochastic modelling

Defining stochastic modelling

Stochastic modelling is a mathematical framework used to capture uncertainty in complex systems.^{xlviii} It achieves this by inserting probabilistic elements to represent uncertainty in inputs and dynamics, enabling the exploration of potential outcomes and the quantification of risks, rather than relying on single-point forecasts. In financial services, this toolkit is utilised for various applications, including pricing derivatives, modelling interest rates, and assessing credit risk.

Quantum computing presents a natural conceptual alignment with stochastic modelling. Quantum systems are inherently probabilistic, with measurement outcomes determined by the probability amplitudes: values that describe the likelihood of observing one quantum state or another.

This shared probabilistic foundation raises the question of whether quantum methods could provide new and potentially more powerful ways to represent, simulate, and reason about uncertainty.

Focus on Monte Carlo

Stochastic modelling covers a broad range of approaches used widely in financial services. Whilst this research has identified early exploration into how quantum methods might extend to approaches such as Markov chains and random walks, the remainder of this section concentrates on Monte Carlo simulations.

This focus reflects the prominence of Monte Carlo in quantum research on stochastic modelling in financial services. It also enables a more focused discussion on how quantum approaches may offer tangible advantages.

An introduction to Monte Carlo

Monte Carlo methods are widely used in financial services, including pricing,^{xlix} Value at Risk (VaR) calculations,ⁱ and stress testing.ⁱⁱ These applications are significant as they directly influence firms' risk management practices and capital requirements.

Monte Carlo works by introducing randomness to generate possible future scenarios.ⁱⁱⁱ Starting from a model of how a value may evolve, the method repeatedly formulates random steps to build simulated pathways over time. Each individual path provides limited insight, but when aggregated, they form an approximate probability distribution of outcomes. From this distribution, measures such as expected values and associated risk can be estimated.

The underlying statistical assumption of Monte Carlo is simple; it is the law of large numbers. This mathematical law states that the average of results obtained from a large number of independent samples converges to a true value. In the context of Monte Carlo

approaches, this means that each individual simulation in isolation is uncertain, but their combined averages become increasingly reliable as more are added.

The challenge for classical computing

Computational Expense

The classical challenge in Monte Carlo is that convergence toward accurate estimations is computationally expensive. Monte Carlo simulation requires running vast numbers of randomised trials to approximate outcomes. For example, estimating the value (or price) of a complex derivative may involve millions of simulations across multiple time steps and risk factors. Mathematically, the error of a Monte Carlo estimate reduces in proportion to the square root of the number of simulations. In simpler terms, halving the margin of error requires running four times as many simulations, and reducing it by a factor of ten requires one hundred times as many. This slow rate of improvement means that Monte Carlo is one of the most computationally expensive models in the financial services industry.

Although expensive, Monte Carlo methods remain widely used as they are flexible and can generate the full probability distribution rather than just point estimates. This capability is particularly valuable in risk management scenarios. In addition, Monte Carlo simulations are also relatively easy to parallelise, meaning that scaling up compute power is often a more practical route to improved results rather than developing entirely new algorithms.

The potential for quantum advantage

Classical Monte Carlo methods are powerful but inherently inefficient, and firms face diminishing returns from additional classical computing power. These limitations present an opportunity for quantum advantage. Quantum approaches to Monte Carlo simulation are often referred to as Quantum Monte Carlo Integration (QMCI).^{liii} These methods aim to accelerate the rate at which Monte Carlo estimates converge. Unlike several other potential applications of quantum techniques, QMCI benefits from a proven theoretical speed-up.

Reducing samples to accelerate convergence

At the core of QMCI lies Quantum Amplitude Estimation (QAE), the algorithmic approach that enables this improvement. QAE reframes the estimation process in a way that can deliver a quadratic speed-up under idealised conditions. Classically, halving the margin of error in a Monte Carlo estimate requires four times as many simulations. QAE reduces this requirement to only twice as many, highlighting the potential efficiency advantage of quantum methods.

This improvement is based on how qubits store and process information. When in a state of superposition, a qubit can simultaneously represent both 0 and 1 values. In the context of QMCI, this allows quantum computers explore a vast number of outcomes in parallel, rather than simulating them one by one. When these qubits are entangled, operations on one affect the state of the others. This interdependence enables algorithms such as QAE to increase the probability of sampling more accurate estimates. Taken

together, these factors enable reliable results to be reached with fewer samples, forming the basis for a potential quantum advantage in Monte Carlo methods.

Emerging use case: Pricing models

Why asset pricing matters

Asset pricing is a core function in portfolio and risk management. Accurately valuing securities, such as stocks, bonds, and derivatives, enables more accurate forecasting of risk and return, allowing for better asset allocation in investment portfolios and a deeper understanding of operational risk. Due to the vast sums of money handled by these business functions, minor improvements in timing or accuracy of pricing solutions can result in significant returns or savings.

Classical approaches

The exact pricing model used by a given business function typically depends on how quickly the data is needed and how it will be utilised. When seeking to optimally price derivatives, only cases involving multi-dimensional (multi-asset) and path-dependent (e.g. Asian options) derivatives typically require Monte Carlo simulations.

One common topic is pricing exotic derivatives, which may be multi-asset, volatility-sensitive, or path-dependent.^{liv,lv} It should be noted that typical investors do not generally use exotic derivatives. These instruments are primarily tailored to clients with specific hedging or investment needs, such as:

- Hedge Funds for arbitrage, speculative strategies, and tailored exposure to market dynamics,
- Asset Managers to enhance yields or protect against downside in complex portfolios, or
- Corporations that are hedging against specific risks, such as FX volatility, interest rate changes, and commodity price fluctuations.

For these types of exotic derivatives, Monte Carlo simulation remains the standard approach. It is often the only practical method for capturing the complexity of multi-dimensional and path-dependent payoffs. However, the large computational expense makes pricing both costly and slow. This limitation highlights why derivatives pricing is viewed as a potentially promising application where quantum Monte Carlo methods could eventually provide an advantage.

Reframing Monte Carlo for quantum advantage

Classical approaches to Monte Carlo simulation for pricing generate outcomes one by one and averages them to ascertain an expected value. In contrast, quantum approaches prepare qubits in a state of superposition. In this state, a single qubit encodes the probability amplitudes for being measured as either 0 or 1, allowing a compact representation of the probability distribution over both outcomes. When multiple qubits are combined, the number of possible inputs that can be represented grows exponentially. In the context of Monte Carlo, this means that instead of simulating each random pathway separately, a register of qubits encodes many possible pathways simultaneously.

In the classical approach, each of these simulated pathways produces an outcome. These outcomes are stored one by one and then averaged to estimate the expected value. QMCI compresses this computation. Instead of storing outcomes separately, all simulated pathways are encoded so that the expected value is mapped to an amplitude of an additional qubit, referred to as an ancilla qubit.

The ancilla qubit captures all the information about the simulated pathways in its probability amplitude. Each pathway in the superposition influences the ancilla qubit, with the size of the influence reflecting the outcome of that pathway. When all pathways are considered together, these rotations accumulate and the ancilla qubit's probability amplitude encodes the expected value.

Obtaining the final value

The desired output of a Monte Carlo simulation for pricing is an expected value. But, at this point, that number is encoded into a probability amplitude of the ancilla qubit. If this qubit is measured directly, it will only return a 0 or 1, never the probability itself. Estimating that probability would require repeated measurements and the averaging of outcomes, thus reintroducing the same inefficiency as classical Monte Carlo methods.

To avoid this, QMCI relies on an additional step called Quantum Phase Estimation (QPE).^{lvi} This method avoids repeated sampling by entangling the ancilla qubit to a register of extra qubits that act as a readout. Through a series of quantum operations, the ancilla's probability is recorded in the qubit register. When this register is finally measured, it outputs a binary number that directly estimates the expected value. Increasing the number of qubits in this register improves the precision of the final estimation.

The reality of the ideal

The approach described above shows how QMCI can provide speed-ups under idealised conditions. In reality, implementation challenges associated with current quantum computers have prevented this advantage from materialising.

QMCI requires encoding large numbers of pathways into quantum states. This requirement necessitates many more qubits and deeper circuit depth than current quantum computers can provide. The process also depends on highly stable qubits and precise gate operations as even small errors can distort the delicate probability amplitudes that carry the expected value.

In addition, Quantum Phase Estimation – the process of converting a probability amplitude into an expected value – is a resource-intensive process. To reach desirable levels of precision, the readout register must contain multiple qubits, and each additional qubit increases both the length of quantum circuit needed and the reliance on error correction.

As a result, while the theoretical advantage of quantum Monte Carlo integration is well established, the hardware and algorithmic requirements place it beyond the reach of near-term devices. Additionally, there is growing scepticism over whether a quadratic speed-up will translate into commercially viable benefits. Classical Monte Carlo methods, although computationally expensive, already deliver accurate results for mainstream use cases such as pricing within acceptable timeframes.

For financial services, quantum approaches to Monte Carlo simulation are perceived as a long-term ambition rather than a near-term likelihood. With a clear theoretical foundation and the potential for significant impact on capital-generating activities, this remains one of the most promising applications of quantum computing in finance. However, with commercial viability increasingly uncertain, QMCI is beginning to resemble a “false dawn”: transformative in theory but limited in practice.

Sentiment landscape

Stochastic modelling, particularly Quantum Monte Carlo Integration (QMCI) methods, is a typical early project for financial services firms interested in quantum computing. The primary reason for the continued popularity of this class of solutions is the clear connection between quantum algorithms and the potential for a speedup against computationally expensive classical counterparts.

Insights from interviews with leading financial services firms and quantum computing experts indicate that there is general agreement that quantum Monte Carlo methods are not yet ready to replace classical Monte Carlo methods for asset pricing. This observation is due to several factors; in particular, quantum computers are not yet large enough to enable QMCI on real-world datasets.^{lvii} However, it is still interesting to consider quantum-inspired approaches for relatively more straightforward derivatives with a single underlying asset.^{lviii}

With the improved devices that will become available in the medium term, the current methods being run in simulation will be able to be tested in practice. In particular, qubit control should become sufficient to enable amplitude encoding, allowing small-scale QAE to occur, which will in turn allow for small-scale QMCI.^{lix} Amplitude encoding can require extremely long circuits; however, algorithmic improvements can significantly reduce the depth needed for these circuits, again shortening the timeframe to possible utility.^{lx} Therefore, a combination of enhanced devices and improved algorithms could make QMCI feasible within the next few years.

One method of bridging the gap between the current methods and systems and what could be achievable in the future is to utilise tensor network simulations of the current algorithms. These are classical methods which approximate the transformation of quantum states, run on GPUs, and demonstrate some of the advantages expected from quantum approaches. Quantum software companies have utilised these techniques to explore new derivative pricing models in collaboration with major banks.

While the *Emerging Use Case* section focused solely on how asset pricing models are incorporated into derivative pricing, it is also worth noting that a similar discussion applies to their use in risk management. In risk management, Monte Carlo simulations are essential for incorporating high levels of uncertainty and accounting for many correlated defaults, as well as for performing rigorous stress testing for regulatory purposes. Calculations such as Value at Risk (VaR) and Conditional Value at Risk (CVaR) typically occur overnight, intraday, or daily, and are needed by a broader class of institutions than the derivative pricing solutions. For these reasons, VaR and CVaR were also commonly cited as potential applications in interviews.

Barriers and Limitations

While the theory behind the speedup in QMCI is sound, this speedup has not been borne out in practice due to several factors.

- **Hardware Limitations:** Quantum systems have not yet reached the qubit numbers and qubit stability necessary to observe this quadratic speedup. It is generally believed that to achieve these requirements, quantum computers will need to reach fault tolerance and have a fully developed software stack, at which point we might see this quadratic advantage.
- **Encoding:** In addition, these methods do not currently scale particularly well, typically requiring one qubit per decision variable in the problem. While efforts have been made to alleviate this issue, further work will still be needed to make quantum stochastic algorithms practical.
- **Does quadratic equal commercial?** Although it is generally agreed that the quadratic advantage could be seen by overcoming hardware limitations, it is unclear whether the speedup will result in a commercial advantage.
- **Use case fit:** Ensuring quadratic speedup results in commercial advantage requires a more precise mapping between theory and practice, tying use cases directly to algorithmic quantum advantage.

Over the longer term, quantum computers are likely to reach the maturity needed to execute these algorithms effectively. The critical question, however, is whether QMCI will translate into a genuine commercial advantage. Realising such an advantage is far from guaranteed and will depend on sustained progress in use-case exploration, algorithm design, and the underlying hardware.

8 Regulatory considerations

Regulatory Readiness for Quantum Computing

The emergence of quantum computing presents a potentially substantial growth opportunity for the UK. In financial services, harnessing the innovative potential of this technology could transform critical operations, create competitive advantage, and help position the UK as a world-leading hub for financial innovation. Realising this opportunity, however, will require coordination across multiple stakeholders.

Within this, the FCA has an important role to play. Not in regulating the technology itself, but in ensuring that the UK financial market remains safe, resilient, and supportive of innovation that can drive growth, competitiveness, and better outcomes for consumers.

For the FCA, quantum computing also represents an opportunity to demonstrate what it means to be a smarter regulator: predictable, purposeful, and proportionate. By engaging early, we can send clear signals about what is expected, identify areas where closer collaboration with industry will be required, and highlight questions that require further exploration. In doing so, we help create a regulatory environment that both supports innovation and builds trust.

At the same time, early engagement is essential to give firms the confidence to invest and experiment. If regulatory engagement evolves in parallel with technology, it can provide clarity and reduce uncertainty. But if engagement falls behind, the regulatory environment risks becoming a bottleneck that slows innovation and creates a barrier to realising the opportunity.

The purpose of this section is therefore to provide the earliest possible insights into how quantum applications might interact with the FCA's remit, and to highlight areas where future regulatory thinking may be required as applications mature. These insights are not exhaustive, nor do they represent regulatory guidance. Instead, they surface some initial considerations to help shape a direction of travel that supports safe, responsible innovation in quantum applications for financial services.

An early approach to surfacing regulatory insights

This project combined insights from industry stakeholders and regulatory colleagues to build an early view of how quantum applications could interact with the FCA's remit. Technical experts from academia, financial services, and the quantum sector were interviewed and convened in workshops to explore the current state of quantum technology, its likely development trajectory, and emerging use cases in financial services. These insights then informed a dedicated workshop with colleagues from across the FCA, designed to discuss how emerging applications might interact with the FCA's regulatory remit.

The regulatory workshop was deliberately designed to be technology-neutral. Participants considered future scenarios derived from emerging use cases, but without explicit reference to quantum computing.

This approach achieved two outcomes. First, it mitigated against novelty bias, grounding the discussion in the FCA's existing outcomes-focused framework and ensuring that quantum computing is not treated as an object of regulation in itself. Second, it lowered technical barriers, allowing FCA colleagues to engage directly with the regulatory implications without needing specialist knowledge of quantum technologies. Whilst not an end-state, this approach provides a useful entry point for regulatory colleagues to engage with emerging technologies.

The workshop was structured around four steps. Participants began by mapping scenarios to the FCA's remit, identifying which existing rules, principles, or governance requirements might apply. They then explored potential regulatory implications, including possible tensions or gaps where current expectations may be stretched by the scenarios. Next, they assessed the innovation itself, considering whether the applications might be beneficial or harmful, and under what conditions they could be acceptable in UK markets. Finally, they reflected on how regulators might shape an effective response.

The insights generated through this process are not exhaustive, but they provide a first-step methodology for surfacing areas of potential regulatory relevance and setting a direction for how thinking could mature in line with the development of quantum applications. The following subsection outlines the cross-cutting considerations that emerged from this work: areas of the evolving regulatory environment that will require deeper exploration to ensure that frameworks continue to support safe, responsible innovation as quantum computing applications emerge in UK financial services.

Cross-cutting considerations

The regulatory workshop explored themes that cut across multiple use cases and problem domains. These themes do not represent definitive regulatory positions or gaps, but rather areas where further consideration may be beneficial as quantum applications mature. By highlighting them here, the aim is not to prescribe new rules or frameworks, but to signal early points of interaction between quantum applications and the FCA's regulatory remit.

These considerations matter because quantum computing is unlikely to introduce entirely new categories of regulatory challenge. Instead, it may amplify existing dynamics – such as data use, model transparency, or market conduct – in ways that stretch current expectations or require a new perspective.

Quantum model explainability

Participants repeatedly highlighted that explainability sits at the heart of fair outcomes and regulatory confidence. For example, under Consumer Duty and related obligations, firms must not only make sound decisions but also be able to demonstrate how those decisions were reached and explain them to consumers and the regulator. Quantum models, like complex AI systems, risk compounding the “black box” challenge. Outcomes may appear accurate, but without a clear record of input and transformations firms may struggle to provide meaningful explanations in cases of consumer harm or disputes. This

risk is exacerbated by the gap in specialist knowledge between those firms exploring and leveraging quantum and regulators such as the FCA. Over time, a concerted effort is required to ensure that regulators can understand and engage with the application of quantum computing in UK financial services to reduce the chance of risks materialising.

Validation, benchmarking, and replicability

Trust in financial models relies in part on their validation and reproducibility. Workshop participants pointed out that while supervisors rarely expect perfect replication of models, they do require clear benchmarks and evidence that outputs are robust and stable. With quantum systems, reproducibility will be complicated by the probabilistic nature of outputs as well as hardware variability. Firms may therefore need to consider how to demonstrate result equivalence, showing, for example that a quantum model's output aligns with expected distributions. While this is not a new principle it will require further consideration on adapting its application for quantum systems.

Outsourcing, concentration, and operational resilience

In the workshop, participants considered a scenario where a financial services firm outsources computer capabilities to a specialist provider for better results. While today's quantum ecosystem is diverse, the market is likely to consolidate over time, raising familiar risks of concentrations and dependency seen in both cloud and AI technologies. Participants considered the FCA's Operational Resilience framework as highly relevant in the context of these potential dependencies. If quantum tools are embedded in pricing, trading, or risk management, outages or degraded performance could have consumer and market impacts. In such future scenarios, firms will need to ensure they have robust alternatives, whether that comes from switching providers or reverting to classical systems. Participants noted that this switching of services may not be straightforward and that outputs could differ significantly, especially if the vendor market is emerging or solution services are immature. Participants concluded that proactive observation is required to better understand the emerging quantum vendor market and its potential to create systemic vulnerabilities for the financial services sector.

Market integrity, fairness, and competition

The use of quantum for market-facing applications raised questions about fairness and competition. The promise of precision analytics or superior optimisation in financial markets could give the firms with access a short-lived but significant advantage. Workshop participants emphasised that the regulatory concern is not with the technology but with its effects on market outcomes, potentially echoing precedents from high-frequency trading and cloud adoption. How this potential advantage alters the market structure is yet to be seen but participants argued that market surveillance will need to keep pace with the evolution and application of quantum computing to ensure that markets remain fair, orderly, and resilient.

Data governance, security, and confidentiality

Both workshop participants and industry contributors noted the risks with moving sensitive financial data into new environments, often managed by small, specialist teams outside of firms' normally recognised infrastructure. Additionally, the Information Commissioner's Office (ICO) has acknowledged that the novelty and complexity of quantum technology could be an additional barrier to engaging internal teams on basic

data protection questions.^{lxi} However, data governance and security are not uniquely quantum challenges. Therefore, these potential barriers are more likely a result of perception than regulatory restrictions. Firms could benefit from further reference to the ICO's Data Sharing guidance to ensure that they are sharing data safely and to reduce the internal compliance barriers to experimentation.^{lxii}

Regulatory Coordination

The workshop highlighted the importance of aligning with emerging international standards to avoid fragmentation or opportunities for regulatory arbitrage. Engagement with international peers in fora such as International Organization of Securities Commissions (IOSCO) and the Financial Stability Board (FSB) will continue to be essential in shaping a globally consistent regulatory approach to quantum computing in financial services.

Equally, domestic coordination can ensure that the UK's regulatory environment builds confidence and supports the development of quantum computing applications in UK financial services. From this perspective, continued engagement across the FCA, Bank of England, and Prudential Regulatory Authority (PRA) can help to reduce regulatory uncertainty regarding the application of quantum computing in UK financial services.

Beyond this, broader regulatory engagement and coordination for example, at the newly established UK Quantum Regulators Forum, can help to ensure the UK's regulatory environment proactively supports the safe deployment of quantum technologies across the UK economy. As quantum computing becomes increasingly market ready, firms and vendors will require further regulatory clarity to reduce friction and build confidence in market deployment.

Timing, proportionality, and triggers

Early regulatory engagement with future quantum applications demonstrates that these applications do not demand a new regulatory playbook. The principles of transparency, resilience, and fairness are already familiar for financial services firms. What may change is the degree of emphasis, the points of application, and the urgency of preparing for them before quantum computing moves from proof-of-concept to real world application. Insights from the workshop reiterate that the FCA's role is not to regulate the technology itself, but to ensure that as applications mature, the regulatory environment supports safe and beneficial innovation, rather than obstructing it.

9 The quantum opportunity: Going for growth

This report set out to demonstrate how the UK's growth ambitions, a globally competitive financial services sector, and a world-leading quantum ecosystem converge to create a unique opportunity at the intersection of quantum computing and financial services. It began by setting out the foundations of quantum computing, providing an accessible entry point for readers to engage with the subject both in this paper and beyond. It then explored the state of the quantum computing stack today, and how UK financial services firms are beginning to shape their own journeys towards readiness. From there, the analysis turned to three problem domains – optimisation, machine learning, and stochastic modelling – highlighting areas of potential advantage, the emerging use cases, the broader sentiment landscape, and the barriers preventing further adoption. Finally, the report outlined some early regulatory considerations, drawing on the expertise of FCA colleagues and using future scenarios to explore where quantum applications might intersect with the FCA's remit as the technology matures.

Collectively, the insights provided here offer the first shared understanding of the quantum opportunity in financial services; they reflect the unique perspectives of academia, industry, vendors, and regulators. This synthesis underscores the importance of aligned action to realise the growth potential of quantum computing in financial service

Over the course of this research, two overarching insights have emerged. First, while quantum computing remains a scientific pursuit, it is also entering the realm of practical consideration. This transition marks a critical juncture for the UK: impacting the country's growth ambitions, its standing as a global centre of financial innovation, and its role in the international race to commercialise quantum technologies. The decisions taken now by industry, policymakers, and regulators will determine whether the UK converts its early scientific leadership into tangible real-world impact. Second, collaboration and coordination between these stakeholders will be essential to achieving success. This research has shown that bringing together diverse perspectives provides clarity potential quantum advantage, where the barriers remain, and what each stakeholder requires from the others to move forward in realising the opportunity ahead.

While the UK financial services sector is well positioned to support the translation of quantum computing towards commercialisation, it cannot do so alone. Realising this opportunity will require the continued alignment across the financial services sector, the quantum ecosystem, and regulators. The aim of this research has been to help nurture those early connections: to build a shared understanding of the technology, to align on the opportunities ahead, and to leverage collaboration as the primary means of overcoming the challenges to come.

Considerations on building quantum readiness

The first steps in realising this opportunity are tied to the notion of building quantum readiness. This report concludes with immediate considerations for three stakeholder

groups whose actions will be critical to developing quantum computing applications in financial services. Whilst each group faces its own challenges, the actions they take will not only shape their individual success but also influence the ability of other stakeholders to progress. These considerations are therefore presented to demonstrate how individual actions can support the broader ecosystem in achieving shared goals.

Financial services firms

The primary consideration for financial services firms is the design of their quantum readiness strategy. Whilst most firms engaged in this research view readiness as a long-term project, their approaches differ depending on market position, cost-benefit assessments, and perceptions of how best to capture commercial value.

Nearly all firms engaged in this research have already begun participating in small-scale experiments and proofs-of-concept, usually with partners across academia and the quantum computing vendor ecosystem. In some cases, firms have paused this activity, determining that the technology is too immature for further investment. Others have continued, with activities ranging from exploring near-term applications to contributing to the underlying scientific progress, often focusing on developing new quantum algorithms for financial services use cases.

These different strategies reflect the divergent views on how and when value will emerge. Firms that have paused experimentation tend to keep a watching brief and adopt a second-mover approach, assuming that commercial applications will primarily be driven by quantum computing vendors. In contrast, firms pursuing early investment tend to have more optimism regarding near-term deployment and believe that financial services firms will play a leading role in shaping applications that provide commercial advantage. Firms that fall into this category usually see the development of internal knowledge, skills, and intellectual property as core components in securing first-mover advantage and deriving future value.

Whichever approach firms take, it will be important to consider how their quantum readiness strategies can remain effective if the technology develops faster, slower, or in a different direction than expected. Broad strategies provide direction, but the individual components within them may need to be flexible to cope with the uncertainty associated with quantum computing's evolution. Firms may therefore wish to consider which of these actions enable them to benefit from nearer-term commercial applications and help hedge against longer-term risks.

Within the development of quantum readiness strategies, there are likely to be a collection of no regret or low regret actions that firms can take. One such action is the development of quantum skills, where investment in building or accessing expertise is likely to remain valuable across a wide range of development scenarios. Beyond benefiting each individual firm, these efforts can help to raise the level of expertise across the sector and strengthen the ecosystem's ability to progress quantum computing applications. Similarly, continued engagement in collaborative proofs-of-concept with vendors and academia provides firms with essential insights while also directing capital and financial expertise into the ecosystem, accelerating broader scientific and commercial progress. In either case, the approaches firms take to building quantum readiness will have a lasting impact on their own preparedness and the progress of the wider ecosystem.

Quantum computing vendors

For quantum computing vendors, immaturity in both technology and markets remain the most immediate barriers to scale. Many are still refining their understanding of which applications hold commercial potential, while working with financial services end users to test and validate these use cases as they emerge.

At the same time, there is limited awareness of regulatory frameworks and expectations within this community. This is understandable at such an early stage, but it risks creating perceived regulatory barriers that deter experimentation and delay deployment. If left unaddressed, these uncertainties can ripple across the broader ecosystem, discouraging investment and slowing adoption even in areas where commercial applications might be viable.

Therefore, a key consideration for quantum computing vendors is how to approach regulatory engagement as a strategic enabler of market adoption. Proactive engagement with regulators on proofs-of-concept, potential applications, and potential deployments can help to build a mutual understanding, prevent misconceptions and provide clarity on how regulatory frameworks may apply. Building these relationships early may feel secondary to technical development, but over time they become a catalyst for confidence: reducing friction for vendors on their path to market and lowering uncertainty for stockholders awaiting signals of regulatory clarity.

As applications mature and move towards commercial deployment in financial services, confidence in their regulatory alignment will be critical to unlocking investment from end-users. Beyond reducing friction for individual vendors, early engagement can help establish clarity, reassure end-users, and create a more predictable environment for innovation and adoption across the ecosystem.

UK regulators

For regulators, the situation is more complex, and there are a number of important considerations to take forward. Like other stakeholders, regulators will need to develop the awareness and skills required to explore potential applications for their own operations. Unlike others, however, regulators will face the unique challenge of understanding how quantum computing applications intersect with their regulatory remits.

Building knowledge and open dialogues

For regulators, it will be critical to build knowledge of quantum computing and establish open channels for dialogue as the foundation of quantum readiness. Regulators need both the technical understanding to anticipate where applications might emerge and trusted relationships across the broader ecosystem to engage effectively as those applications take shape.

The formation of the UK's Quantum Regulatory Forum is an important first step in this direction.^{lxiii} Its initial work to anticipate future developments, practical applications, and regulatory implications will help authorities understand where quantum applications may interact with their mandates, and on what timescales. But UK regulators should consider how to take this further: using the forum to actively strengthen links with academia, industry, and the broader quantum ecosystem

Knowledge-building cannot be a passive exercise. Regulators should consider creating opportunities for dialogue where technical and regulatory experts can develop a shared language. Through this process, regulators can accelerate their own learning also reduce misconceptions across industry about regulatory expectations.

Beyond improving their own quantum readiness, investment in these relationships and forums has a wider ecosystem value. By helping innovators navigate uncertainty, regulators can build the trust, predictability, and confidence required for industry and vendors to sustain investment, pursue experimentation, and eventually adopt quantum applications.

Supporting development through innovation services

Supporting innovation in emerging technologies does not always require formal regulatory guidance or policy. UK regulators have a longstanding history of developing a wider suite of engagement tools that allow for dialogue, learning, and safe experimentation. Whilst regulatory activity is sometimes perceived as a constraint, these mechanisms are designed to support the journey from idea to application in ways that mitigate risks without stifling innovation.

For quantum computing, regulators should consider how these approaches can be adapted, or whether dedicated, quantum-specific services are needed. Possible options for consideration include:

- 1. Regulatory sandboxes:** Existing sandboxes could be expanded to provide access to quantum computing resources alongside classical compute, integrated development environments, and synthetic data for testing and development. Alternatively, regulators could consider piloting quantum-specific sandboxes, convening interested stakeholders from across the ecosystem. In considering this approach, regulators should seek to engage with the NQCC to gather invaluable insights from their experience on building the UK quantum testbed programme. Such initiatives would not only support regulators to understand the practical challenges, but also lower barriers for firms and researchers by providing a shared infrastructure and a safe space to experiment.
- 2. Observation of proofs-of-concept (PoCs):** Regulators could consider engaging directly with early-stage pilots, observing how applications move from concept to prototype. This would give early visibility of the challenges firms face, highlight potential regulatory barriers, and provide regulators with a low-cost way to deepen their understanding of quantum applications ahead of commercial deployment. At the same time, firms would benefit from early signals of regulatory expectations, reducing uncertainty and avoiding later barriers to deployment.
- 3. Open engagement through events and forums:** Regulators may also consider hosting or attending regular workshops, roundtables, and conferences to engage with a broad set of stakeholders. This would signal regulatory interest, create safe spaces to discuss unanswered questions, and help identify blind spots, gaps, or points of friction before they materialise into barriers. By participating actively, regulators also demonstrate that they are part of the innovation process, not separate from it.
- 4. Establish quantum-focused teams:** Many regulators have created small teams dedicated to other emerging technologies, such as AI or Digital Assets. These teams are essential for understanding how these technologies interact with existing

regulatory frameworks and where new regulation is required. As quantum matures, regulators could consider establishing dedicated teams to build hubs of expertise, raise awareness across regulatory colleagues, and provide innovators with a clear point of contact. These teams could act as conduits for knowledge-sharing and collaboration across the quantum ecosystem.

Together, these considerations demonstrate how regulators can further reduce perceived regulatory barriers, foster trust, and signal that they are actively engaged in supporting innovation while mitigating for the associated risks.

Applications Regulatory Readiness Framework

Ultimately, regulators will need a structured way to decide *when* and *how* to act as quantum applications mature. To support this action, the Regulatory Horizon Council's report on Regulating Quantum Technology Applications was an important step forward.^{lxiv} Its core recommendation, a "regulation by application" approach, has since been adopted by the UK government and is consistent with wider UK regulatory practice.^{lxv}

The report also introduced the Proportionate and Adaptive Governance of Innovative Technologies (PAGIT) framework, which provides a useful starting point for considering how regulators might engage with emerging technologies.

This research, however, highlights two critical limitations when applying the PAGIT framework to quantum computing. First, while many quantum applications share the same underlying technology, their maturity levels differ significantly. Some may be close to commercial deployment in financial services, while others remain firmly in the research stage. This is a crucial nuance for regulators who seek to take a technology-neutral, outcomes-focused, regulation-by-application approach. Second, the PAGIT framework does not reflect the full range of regulatory tools available.

A critical next step for UK regulators to consider is building on the PAGIT framework to develop an Applications Regulatory Readiness Framework. Such a framework could be used to evaluate the maturity of individual applications and set out what forms of engagement are appropriate at each stage: from early dialogue and observation, through to proof-of-concept, and market deployment

Developing this framework sends a powerful signal that UK regulators are not only preparing for quantum computing but also initiating new approaches to engaging with potentially transformative, emerging technologies. By doing so, the UK could position itself as a global leader in adaptive, proportionate regulatory approaches that that supports growth, builds trust, and translate innovation into tangible economic impact.

Annex 1: References

- ⁱ UK Government. (2025, June 23). *The UK's modern industrial strategy: Policy paper*. GOV.UK. https://assets.publishing.service.gov.uk/media/68595e56db8e139f95652dc6/industrial_strategy_policy_paper.pdf
- ⁱⁱ Ibid.
- ⁱⁱⁱ M Treasury. (2025, July 21). *Financial Services Growth and Competitiveness Strategy: Overview*. GOV.UK. <https://www.gov.uk/government/calls-for-evidence/financial-services-growth-and-competitiveness-strategy/outcome/financial-services-growth-and-competitiveness-strategy-overview>
- ^{iv} HM Treasury. (2025, July 15). *Leeds reforms to rewire financial system, boost investment and create skilled jobs across UK* [News release]. GOV.UK. <https://www.gov.uk/government/news/leeds-reforms-to-rewire-financial-system-boost-investment-and-create-skilled-jobs-across-uk>
- ^v Financial Conduct Authority. (2025). *Our strategy 2025 to 2030* (PDF). <https://www.fca.org.uk/publication/corporate/our-strategy-2025-30.pdf>
- ^{vi} Department for Science, Innovation and Technology. (2023, March 15). *National Quantum Strategy* (PDF). GOV.UK. https://assets.publishing.service.gov.uk/media/6411a602e90e0776996a4ade/national_quantum_strategy.pdf
- ^{vii} UK Government. (2025, June 23). *The UK's modern industrial strategy: Policy paper*. GOV.UK. https://assets.publishing.service.gov.uk/media/68595e56db8e139f95652dc6/industrial_strategy_policy_paper.pdf
- ^{viii} If readers wish to learn more about quantum computing, please read: Decaroli, C. (2025, January 11). *A guide to online resources for learning quantum computing*. NQCC. <https://www.nqcc.ac.uk/a-guide-to-online-resources-for-learning-quantum-computing/>
- ^{ix} Bertels, K. O. E. N., Sarkar, A., Hubregtsen, T., Serrao, M., Mouedenne, A. A., Yadav, A., ... & Almudever, C. G. (2020). Quantum computer architecture toward full-stack quantum accelerators. *IEEE Transactions on Quantum Engineering*, 1, 1-17.
- ^x Resch, S., & Karpuzcu, U. R. (2019). Quantum computing: an overview across the system stack. *arXiv preprint arXiv:1905.07240*.
- ^{xi} Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *arXiv preprint arXiv:1801.00862*.
- ^{xii} Acharya, R., Abanin, D. A., Aghababaie-Beni, L., Aleiner, I., Andersen, T. I., Ansmann, M., ... & Boixo, S. (2024). Quantum error correction below the surface code threshold. *Nature*.
- ^{xiii} Bluvstein, D., Evered, S. J., Geim, A. A., Li, S. H., Zhou, H., Manovitz, T., ... & Lukin, M. D. (2024). Logical quantum processor based on reconfigurable atom arrays. *Nature*, 626(7997), 58-65.
- ^{xiv} AI, G. Q. (2023). Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7949), 676

- ^{xv} Roffe, J., White, D. R., Burton, S., & Campbell, E. (2020). Decoding across the quantum low-density parity-check code landscape. *Physical Review Research*, 2(4), 043423.
- ^{xvi} Hillmann, T., Berent, L., Quintavalle, A. O., Eisert, J., Wille, R., & Roffe, J. (2025). Localized statistics decoding for quantum low-density parity-check codes. *Nature Communications*, 16(1), 8214.
- ^{xvii} Malcolm, A. J., Glaudell, A. N., Fuentes, P., Chandra, D., Schotte, A., DeLisle, C., ... & Simmons, S. (2025). Computing efficiently in QLDPC codes. *arXiv preprint arXiv:2502.07150*.
- ^{xviii} Wille, R., Hillmich, S., & Burgholzer, L. (2020, October). Efficient and correct compilation of quantum circuits. In *2020 IEEE International Symposium on Circuits and Systems (ISCAS)* (pp. 1-5). IEEE.
- ^{xix} Sivarajah, S., Heidemann, L., Lawrence, A., & Duncan, R. (2022, November). Tierkreis: A dataflow framework for hybrid quantum-classical computing. In *2022 IEEE/ACM Third International Workshop on Quantum Computing Software (QCS)* (pp. 12-21). IEEE.
- ^{xx} Shor, P. W. (1999). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2), 303-332.
- ^{xxi} Grover, L. K. (1996, July). A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing* (pp. 212-219).
- ^{xxii} Sun, Z., Boyd, G., Cai, Z., Jnane, H., Koczor, B., Meister, R., ... & Stamatopoulos, N. (2025). Low-depth phase oracle using a parallel piecewise circuit. *Physical Review A*, 111(6), 062420.
- ^{xxiii} Liu, M., Shaydulin, R., Niroula, P., DeCross, M., Hung, S. H., Kon, W. Y., ... & Pistoia, M. (2025). Certified randomness using a trapped-ion quantum processor. *Nature*, 1-6.
- ^{xxiv} Mugel, S., Kuchkovsky, C., Sánchez, E., Fernández-Lorenzo, S., Luis-Hita, J., Lizaso, E., & Orús, R. (2022). Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks. *Physical Review Research*, 4(1), 013006.
- ^{xxv} Kyriienko, O., & Magnusson, E. B. (2022). Unsupervised quantum machine learning for fraud detection. *arXiv preprint arXiv:2208.01203*.
- ^{xxvi} Ghysels, E., Morgan, J., & Mohammadbagherpoor, H. (2023). Quantum computational algorithms for derivative pricing and credit risk in a regime switching economy. *arXiv preprint arXiv:2311.00825*.
- ^{xxvii} Mugel, S., Kuchkovsky, C., Sánchez, E., Fernández-Lorenzo, S., Luis-Hita, J., Lizaso, E., & Orús, R. (2022). Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks. *Physical Review Research*, 4(1), 013006.
- ^{xxviii} Palomar, D. P. (2025). *Portfolio optimization: Theory and application*. Cambridge University Press.
- ^{xxix} Abbas, A., Ambainis, A., Augustino, B., Bärtschi, A., Buhrman, H., Coffrin, C., ... & Zoufal, C. (2024). Challenges and opportunities in quantum optimization. *Nature Reviews Physics*, 1-18.
- ^{xxx} Brandhofer, S., Braun, D., Dehn, V., Hellstern, G., Hüls, M., Ji, Y., ... & Wellens, T. (2022). Benchmarking the performance of portfolio optimization with QAOA: S. Brandhofer et al. *Quantum Information Processing*, 22(1), 25.
- ^{xxxi} Farhi, E., Goldstone, J., & Gutmann, S. (2014). A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*.
- ^{xxxii} Fernández-Lorenzo, S., Porras, D., & García-Ripoll, J. J. (2020). Hybrid quantum-classical optimization for financial index tracking. *arXiv preprint arXiv:2008.12050*.
- ^{xxxiii} Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.
- ^{xxxiv} Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- ^{xxxv} Rebentrost, P., Mohseni, M., & Lloyd, S. (2014). Quantum support vector machine for big data classification. *Physical review letters*, 113(13), 130503.
- ^{xxxvi} Schuld, M. (2021). Supervised quantum machine learning models are kernel methods. *arXiv preprint arXiv:2101.11020*.

- ^{xxxvii} Khoshaman, A., Vinci, W., Denis, B., Andriyash, E., Sadeghi, H., & Amin, M. H. (2018). Quantum variational autoencoder. *Quantum Science and Technology*, 4(1), 014001.
- ^{xxxviii} Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., ... & Liu, T. Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.
- ^{xxxix} Chen, T., & Guestrin, C. (2016, August). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining* (pp. 785-794).
- ^{xl} Kyriienko, O., & Magnusson, E. B. (2022). Unsupervised quantum machine learning for fraud detection. *arXiv preprint arXiv:2208.01203*.
- ^{xli} Gil-Fuster, E., Gyurik, C., Pérez-Salinas, A., & Dunjko, V. (2024). On the relation between trainability and dequantization of variational quantum learning models. *arXiv preprint arXiv:2406.07072*.
- ^{xlii} Larocca, M., Thanasilp, S., Wang, S., Sharma, K., Biamonte, J., Coles, P. J., ... & Cerezo, M. (2025). Barren plateaus in variational quantum computing. *Nature Reviews Physics*, 1-16.
- ^{xliii} Landman, J., Thabet, S., Dalyac, C., Mhiri, H., & Kashefi, E. (2022). Classically approximating variational quantum machine learning with random fourier features. *arXiv preprint arXiv:2210.13200*.
- ^{xliv} Zoufal, C., Lucchi, A., & Woerner, S. (2019). Quantum generative adversarial networks for learning and loading random distributions. *npj Quantum Information*, 5(1), 103.
- ^{xliv} Aaronson, S. (2014). Quantum machine learning algorithms: Read the fine print. *Nature Physics*, 5.
- ^{xlvi} McClean, J. R., Romero, J., Babbush, R., & Aspuru-Guzik, A. (2016). The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2), 023023.
- ^{xlvi} Karalekas, P. J., Tezak, N. A., Peterson, E. C., Ryan, C. A., Da Silva, M. P., & Smith, R. S. (2020). A quantum-classical cloud platform optimized for variational hybrid algorithms. *Quantum Science and Technology*, 5(2), 024003.
- ^{xlvi} Shreve, S. E. (2004). *Stochastic calculus for finance II: Continuous-time models* (Vol. 11). New York: Springer.
- ^{xlvi} Boyle, P., Broadie, M., & Glasserman, P. (1997). Monte Carlo methods for security pricing. *Journal of economic dynamics and control*, 21(8-9), 1267-1321.
- ⁱ Hong, L. J., Hu, Z., & Liu, G. (2014). Monte Carlo methods for value-at-risk and conditional value-at-risk: a review. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 24(4), 1-37.
- ⁱⁱ Sorge, M. (2004). Stress-testing financial systems: an overview of current methodologies.
- ⁱⁱⁱ Jäckel, P. (2002). *Monte Carlo methods in finance*. John Wiley & Sons.
- ⁱⁱⁱ Herbert, S. (2022). Quantum Monte Carlo integration: The full advantage in minimal circuit depth. *Quantum*, 6(arXiv: 2105.09100), 823.
- ^{liv} Cibrario, F., Golan, O. S., Ranieri, G., Dri, E., Ippoliti, M., Cohen, R., ... & Corbellotto, D. (2024, September). Quantum amplitude loading for rainbow options pricing. In *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)* (Vol. 1, pp. 211-220). IEEE.
- ^{iv} Cibrario, F., Cohen, R., Dri, E., Mattia, C., Golan, O. S., Danzig, T., ... & Montrucchio, B. (2025). Autocallable Options Pricing with Integration-Based Exponential Amplitude Loading. *arXiv preprint arXiv:2507.19039*.
- ^{vi} Rebentrost, P., Gupt, B., & Bromley, T. R. (2018). Quantum computational finance: Monte Carlo pricing of financial derivatives. *Physical Review A*, 98(2), 022321.
- ^{vii} Chakrabarti, S., Krishnakumar, R., Mazzola, G., Stamatopoulos, N., Woerner, S., & Zeng, W. J. (2021). A threshold for quantum advantage in derivative pricing. *Quantum*, 5, 463.
- ^{viii} Patel, R. G., Dominguez, T., Dib, M., Palmer, S., Cadarso, A., Contreras, F. D. L., ... & Orus, R. (2023). Application of Tensor Neural Networks to Pricing Bermudan Swaptions. *arXiv preprint arXiv:2304.09750*.

^{lix} Pagni, V., Huber, S., Epping, M., & Felderer, M. (2025). Fast Quantum Amplitude Encoding of Typical Classical Data. *arXiv preprint arXiv:2503.17113*.

^{lx} Sun, Z., Boyd, G., Cai, Z., Jnane, H., Koczor, B., Meister, R., ... & Stamatopoulos, N. (2025). Low-depth phase oracle using a parallel piecewise circuit. *Physical Review A*, 111(6), 062420.

^{lxi} ICO. (2024, October 9). *ICO tech futures: Quantum technologies*. Information Commissioner's Office. <https://ico.org.uk/about-the-ico/research-reports-impact-and-evaluation/research-and-reports/technology-and-innovation/tech-horizons-and-ico-tech-futures/ico-tech-futures-quantum-technologies/>

^{lxii} ICO. (2021, May). *Data sharing: A code of practice*. Information Commissioner's Office. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/data-sharing-a-code-of-practice/>

^{lxiii} Department for Science, Innovation and Technology. (2025, June 26). *Quantum Regulators' Forum: Terms of Reference*. GOV.UK. <https://www.gov.uk/government/publications/quantum-regulators-forum-terms-of-reference/quantum-regulators-forum-terms-of-reference>

^{lxiv} Regulatory Horizons Council. (2024, February). *Regulating quantum technology applications* (PDF). https://assets.publishing.service.gov.uk/media/65ddc83bcf7eb10015f57f9f/RHC_regulating_quantum_technology_applications.pdf

^{lxv} Department for Science, Innovation & Technology; Regulatory Horizons Council. (2024, October 8). *Regulating quantum technology applications: Government response to the RHC*. GOV.UK. <https://www.gov.uk/government/publications/regulating-quantum-technology-applications-government-response-to-recommendations-made-by-the-regulatory-horizons-council/regulating-quantum-technology-applications-government-response-to-the-rhc>

