Deloitte.

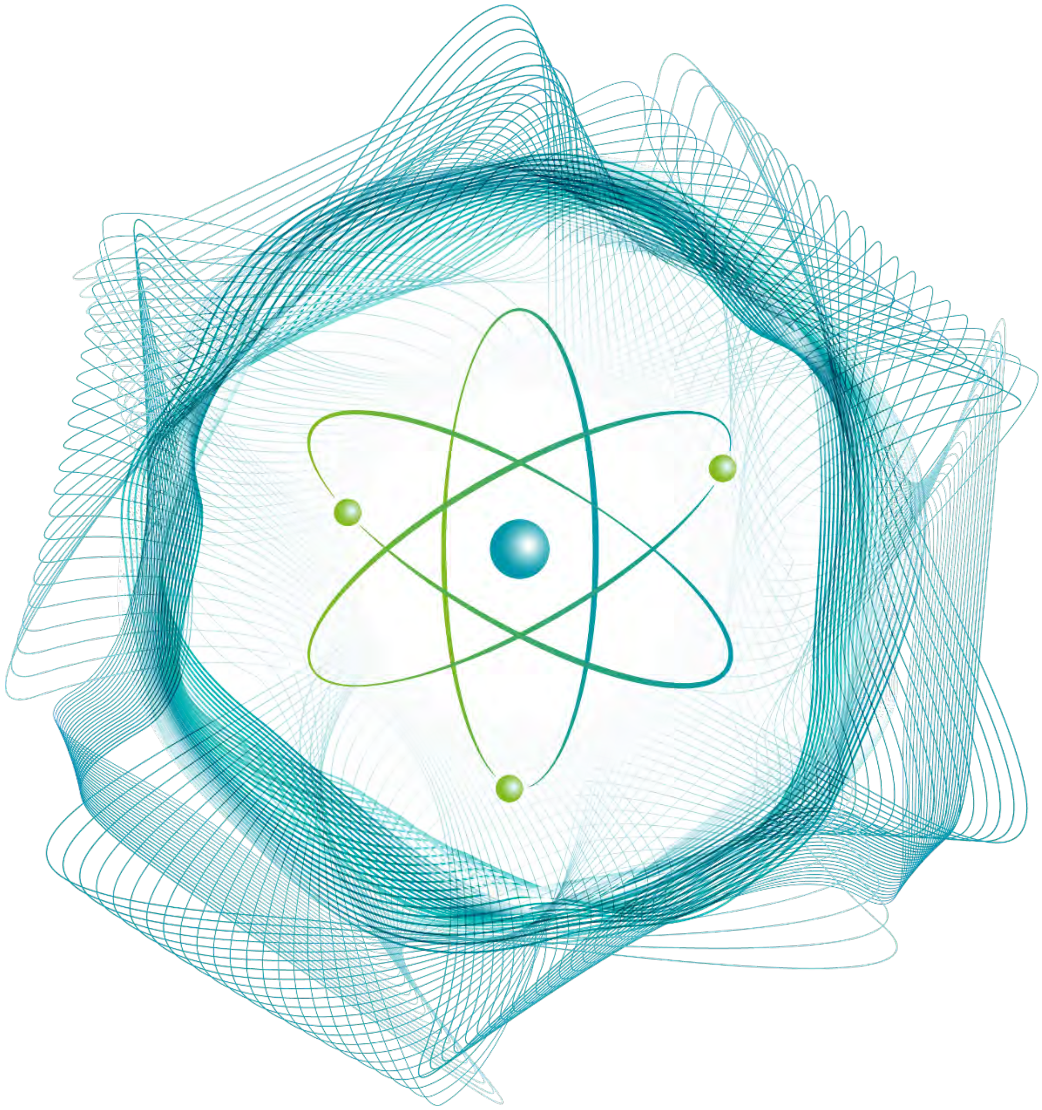

Financial Services Industry Market Research

Market Analysis of Quantum Computing Use Cases

September 2024

Prepared by: Deloitte Consulting LLP

Table of Contents

Executive Summary

Quantum computing, a transformative computational approach, is on the brink of ushering in a new era of advanced computing. By performing calculations differently than today's computers, quantum computers may enable companies across the financial services industry to improve the identification of fraud, accelerate complex calculations for valuation, better optimize portfolios and liquidity, and much more. This report offers an accessible perspective on how quantum computing can be applied in financial services. The report prioritizes use cases by impact and feasibility which is useful to both financial services companies and quantum technology companies as they seek to understand and invest in quantum technology.

The financial services industry is particularly well-positioned to reap the potential benefits of these quantum computing advancements. There is a wealth of both existing and untapped use cases in this sector, which can be addressed by quantum solutions. These include, but are not limited to, applications in optimization, machine learning, and simulation.

To better understand the nature of these opportunities, we evaluated over 50 potential use cases, consolidating them into a dozen classes of opportunities. For each of these opportunity classes, we evaluated a combination of business factors (i.e., potential profitability improvement, size of industry, degree of regulation of the problem) and technical factors (i.e., whether current alternatives are scalable enough, the complexity of running the problem on a quantum computer, known likelihood of speedup). Opinions and feedback from subject matter experts and industry leaders were incorporated to validate the findings, and we compared our results to surveys conducted with industry executives. This report summarizes and evaluates the most attractive financial services opportunities for quantum computing as determined by weighted analysis combining the highest business impact and technical feasibility.

The six classes of use cases which have the highest business impact and technical feasibility include: derivative pricing, liquidity optimization, portfolio optimization, risk analysis, supervised machine learning, and unsupervised machine learning. Each of these is further detailed in Section V, including the anticipated business benefits across the different aspects of the financial services industry.

While there is significant optimism about the future use of quantum computers in financial services, at the time of this writing, production workloads are largely outside the bounds of the capabilities of today's hardware. The quantum computing industry is continuing to combine hardware and software to overcome the current challenges; this process is sometimes referred to as trying to construct "logical qubits" or fault-tolerant quantum computers (FTQC.) Moreover, future research may uncover software algorithms or approaches that substantively change the relative attractiveness of different use cases.

In the short term, we expect that quantum computing will positively impact our operations through combinatorial optimization. In the longer-term, we expect the strongest impact to be on our approach to modeling with differential equations and matrix inversions.

– FSI Industry Executive

Quantum Computing in Financial Services

By operating in a different way than traditional computers, quantum computers (QC) may enable financial services institutions to solve computationally hard problems. As the capability of quantum computers continues to mature, there is growing interest in understanding how these new devices can be applied to industrial applications.

Today, despite continued advancements, quantum computers still face many technical challenges to running production workloads. Although today's quantum computers have left the lab and are now commercially available, the current state of the art is often described as "Noisy Intermediate Scale Quantum" (NISQ) computing. That is, the nature of problems that can be solved is still limited and the volume of data that can be utilized is relatively low. The industry is currently focused on using Quantum Error Correction (QEC) to improve the repeatability and reliability of calculations, in addition to working on hardware and software improvements. The ultimate goal is to transcend the current limitations and create, though various means, a fully fault-tolerant quantum computer with enough capacity (measured in "qubits") to support complex calculations.

> *Quantum computing may enhance fraud detection and prevention in financial services by rapidly analyzing large datasets and identifying patterns with greater precision, helping us stay ahead of emerging threats.*
>
> – FSI Industry Executive

The Financial Services Sector

The financial services industry serves as a pivotal cornerstone in the economic landscape, offering a breadth of financial solutions to individual consumers and diverse businesses. Encompassing a range of entities such as banking institutions, investment establishments, insurance providers, and finance corporations, the industry extends its reach across various specialized sub-sectors, each catering to distinct financial needs and demands:

Banking

This sector is the backbone of the financial industry, featuring commercial banks offering services like deposit accounts, loans, and credit cards. Credit unions provide similar services but are member-owned. Internet banks, fully digital entities, offer online services including bill payment, loan applications, and brokerage services. Savings and loan associations primarily focus on mortgages, thus diversifying the sector's services.

Insurance

Insurance companies serve as financial guardians, safeguarding individuals and businesses from potential fiscal adversities. Life insurance pays out upon the policyholder's death, while health insurance covers medical costs. Property and casualty insurance protect against property loss or damage. Reinsurance companies further manage risk by providing insurance to other insurers.

Investment Services and Asset Management

Asset management professionals facilitate the management of clients' portfolio through developing strategies to meet client's individual needs, such as increasing value and mitigating risk wherever possible. Other investment services include private equity firms who invest in companies to increase their value for profitable sales, and venture capital firms fund startups with high growth potential.

Financial Planning and Advisory Services

Financial planners serve as personal finance architects, assisting individuals in devising strategies to achieve their long-term financial aspirations. On the digital front, robo-advisors deliver automated, algorithm-powered financial planning services, operating with minimal to zero human intervention. Complementing these are wealth management firms, which provide investment counsel and comprehensive financial planning services, specifically catering to the needs of high-net-worth individuals.

Capital Markets

This sector facilitates the buying and selling of securities, including stocks, bonds, and derivatives. Stock exchanges are marketplaces where securities are bought and sold, and brokerages execute trades on behalf of clients.

Real Estate

Real estate companies deal with the buying, selling, renting, and management of properties. This includes residential, commercial, and industrial properties. Some companies also invest in real estate on behalf of clients.

> *Quantum computing is expected to have a major impact in addressing complex mathematical problems in the banking industry. Despite the uncertainty as to when a quantum computer will be able to tackle production-scale problems, some near-term advantages can be obtained from developing internal talent in researching new algorithmic approaches, such as quantum-inspired solutions. This approach can help identify and mitigate suboptimal approaches to problems like those in the field of combinatorial optimization.*
>
> – FSI Industry Executive

Complex Problems in Financial Services

As quantum computing technology matures, the potential applications within the financial services industry are starting to take shape. These applications are largely centered around three key areas: combinatorial optimization, machine learning, and simulations.

Combinatorial Optimization

Combinatorial Optimization is the process of analyzing and processing many potential outcomes, resulting in making the "best" choice when there are many options to choose from, utilizing computationally intense models to process many variables very quickly. Combinatorial optimization is at the heart of many financial services operations. In financial services, combinatorial optimization helps solve complex problems such as portfolio optimization, fraud detection, and resource allocation, to name a few. From asset pricing to portfolio management and logistical planning, finding the most efficient solutions to complex problems is a constant challenge.

Quantum computers enable accurate and expedited optimization calculations. Traditional computers can struggle with these tasks, but quantum computers, with their ability to process complex calculations rapidly, could revolutionize these processes. For instance, in portfolio management, a quantum computer could simultaneously assess all asset combinations and their prospective returns, thereby optimizing the portfolio for peak returns with the least risk.

Today, organizations utilize a variety of commercial and open-source solvers, as well as custom optimization solutions for certain problem classes such as Branch-and-Price that are tailored for Linear and Mixed Integer Programming (MIP) problems but can also solve quadratic problems (either quadratic terms in the objective or constraints).

Machine Learning

Machine learning, a subset of artificial intelligence, involves the creation of algorithms that learn and improve over time. This technology, which is already reshaping the financial services sector, employs algorithms to discern patterns and forecast outcomes. Often, these algorithms are executing a process known as 'classification,' where they categorize inputs into two or more groups, such as potential fraud versus non-fraud cases.

Quantum machine learning, which combines machine learning and quantum physics, could take machine learning to the next level. Quantum machine learning algorithms could potentially handle larger datasets with far fewer trained parameters and create complex models more efficiently than their classical counterparts. This could enhance the predictive capabilities of financial models, leading to more accurate risk assessment and decision-making. In fraud detection, for instance, quantum machine learning could help identify anomalies that might be missed by traditional algorithms.

Currently, organizations employ an extensive array of machine learning models that vary in their explainability and predictive prowess. These models span a range of techniques, from gradient-boosted ensemble models and support vector machines to deep learning neural networks. For pattern detection or unsupervised machine learning, methods such as Isolation Forest, Local Outlier Factor, and neural networks like autoencoders and generative adversarial networks (GANs) are commonly utilized.

Simulations

In financial services, simulations are utilized to attempt to model the real world, to calculate the most likely outcomes, and more. Some of the most advanced simulations use techniques like Monte Carlo calculations which repeatedly perform random sampling to obtain numerical results, such as statistics of unknown evolving distributions. Monte Carlo simulations can be utilized to estimate the likelihood of someone's retirement savings being sufficient, to value complex derivative instruments, or to simulate complex financial systems.

Quantum computers can model complex systems that traditional computers struggle with, allowing for a more nuanced understanding of these systems and their dynamics. Quantum algorithms allow for simulating stochastic and evolving processes with many interacting scenarios to estimate statistics of interest more quickly. For example, quantum simulations could improve asset pricing, where the quantum algorithms could factor in more variables and scenarios than traditional models, potentially leading to more accurate pricing in a much shorter time.

Today, organizations use a variety of high performance custom code to calculate the results of simulations. Numerical methods are often implemented in compiled codebases using languages such as C or C++.

Use Case Ranking

Ranking Approach

To rank applicable financial services use cases, we developed an analytical approach that included both quantitative and qualitative aspects. These criteria are categorized as either business or technical: business criteria describe the industry of the use case as well as potential benefits and costs associated with implementing the quantum solution for the use case, and technical criteria describe the level of effort and technical expertise required to implement the quantum solution for the use case.

Each ranking criterion was assigned a relative weight or importance, and every identified use case was evaluated on a quantitative scale. Weighted totals were used to determine the final use case ranking, with the top 6 use cases being selected for further in-depth analysis within this report.

We believe reliable and fast fraud detection will get massive improvements from quantum computers.

– FSI Industry Executive

Ranking Criteria

Number of Entities in the Industry: We have assumed that the more entities that are impacted, the higher the chance of a solution being adopted.	**Industry Risk Profile:** We have assumed the more tolerant an industry is to taking risks, the higher the chance of a solution being adopted.
Impacted Dollar Size of Industry: We have assumed that the larger the overall potential financial impact, the higher the chance of a solution being adopted.	**Complexity of Quantum Processor Unit (QPU) Processing[1]:** We have estimated the complexity of the quantum computing algorithm (a combination of qubits and gates), with lower complexity solutions expected to be more valuable in the short term than higher complexity solutions.
Profitability Impact: We have assumed that the larger the absolute profitability impact, the higher the chance of a solution being adopted.	**Estimated Improvement Impact[1]:** We have estimated the degree of improvement and impact that quantum-based solutions may generate. Higher degrees of improvement were assumed to result in higher chances of a solution being adopted.
Cost of Ownership: We have estimated the potential switching, implementation, and ongoing costs of a quantum computing-based solution, with high ownership costs lowering the chance of a solution being adopted.	**Certainty of Improvement from Quantum Solution[1]:** We have estimated the current likelihood of a quantum-based solution providing an improvement, based on the best scientific understanding of current algorithms and approaches. Higher likelihood results in higher chances of a solution being adopted.
Degree of Existing Regulation: We have assumed that problems with a high degree of existing regulation will require additional time and effort to make a case for change, lowering the chance of a solution being adopted.	**Scalability of Classical Methodologies:** We have estimated the degree of scalability of existing classical techniques, with lower potential scalability likely increasing the chance of a quantum-based solution being adopted.

[1] Given the rate of change in algorithm development, new approaches may arise which may change the required algorithmic complexity, likelihood, and degree of improvements.

Additional Assumptions

Current Quantum Calculation Limitations

Quantum computers use quantum bits or "qubits" as the fundamental element of computation. The qubits that exist on today's machines are often referred to as "physical qubits" and the theoretical construct of error-free, fault tolerant qubits as "logical qubits." Current approaches to building logical qubits often assume that multiple physical qubits will be combined to improve the quality of output from today's quantum computers.

The number of physical qubits required to produce a logical qubit is currently a subject of debate. Estimates can range from a couple of dozen qubits to 1,000 or 10,000 qubits. The current scientific consensus is that it may require approximately 1,000 physical qubits. Recent research, however, suggests that that number may be as low as 24. While more research and validation is required, we assume that the field of error correction will evolve rapidly. For the purposes of this report, we assume a conservative estimate of several hundred physical qubits.

In addition, qubits must interact with one another in a process commonly referred to as "gates." The number of gates that need to operate sequentially is often referred to as "gate depth." Having a large gate depth can be problematic for today's hardware for reasons similar to why it can be problematic to need large numbers of qubits. Calculations can require relatively low or high gate depth and relatively low or high qubit counts. Some problems require relatively fewer qubits, but relatively higher gate depth, leading to different types of complexity.

Ultimately, the number of qubits and gate depth required to solve problems are important because they limit how soon existing quantum computers might be available to solve useful business challenges.

Current State of Quantum Algorithm Development

Technical considerations and technical criteria used to evaluate these use cases were developed within the context of the most recent understanding of classical and quantum approaches to the relevant application, as of April 2024. The fast-evolving nature of quantum computing may substantially change qubit and gate depth requirements, algorithmic approach, and outlook. Therefore, future analysis may differ significantly from this point-in-time report.

Where possible, we assume that use cases that have multiple avenues forward, either through different algorithmic approaches or low qubit/gate depth thresholds, have in general a higher potential to be realized in the coming years.

Summary Of Use Cases

TOP FSI USE CASES			
Class of Use Case	**Solution Type**	**Primary Sector**	**Description**
Derivative Pricing	Simulation	Capital Markets	Simulation to enhance pricing of derivatives, options and CDOs through Value at Risk (VaR) calculations.
Liquidity Optimization	Optimization	Banking	Optimization solution that seeks to find the optimal order of transaction processing to maximize the liquidity efficiency for a payment system.
Portfolio Optimization	Optimization	Investment Management	Optimal solution seeking to maximize return while minimizing risk and considering additional factors.
Risk Analysis	Simulation	All Sectors	Simulation designed to enhance risk analysis of instruments and counterparty credit risk through Value at Risk (VaR.)
Supervised Anomaly Detection	Machine Learning	Banking	Machine learning classification task that is trained on labeled cases of fraud/financial anomalies.
Unsupervised Anomaly Detection	Machine Learning	Banking	Machine learning algorithm that does not have labeled training data but learns to identify anomalies.
Algorithmic Trading	Machine Learning & Optimization	Capital Markets	Optimization solution seeking to enhance speed of automated trading platforms to place trades with optimal consideration of timing and/or price.
CAT Modeling	Simulation	Insurance	Simulation solution seeking to increase speed and depth of catastrophic event modeling to increase profitability and customer base.
Collateral Optimization	Optimization	Capital Markets	Optimization solution seeking to find the best mix of bonds to respond to margin calls. This reduces complexity and cost.
Creditworthiness	Optimization	Banking	Identification of independent features that are influential in determining borrower credit worthiness.
Episodic Insurance	Simulation	Insurance	Simulation seeking to enhance speed of insurance calculations to allow for individual trip or event-based insurance analysis.
Financial Crash Prediction	Optimization	Banking	Optimization solution seeking to enhance modeling of financial network for optimized financial crash prediction.

Top Classes of Use Cases

Derivative Pricing

Description

Derivatives are synthetic financial instruments whose value is derived from one or more underlying assets. Derivatives are primarily used for hedging risks, speculating on future price movements, and gaining access to otherwise hard-to-trade assets or markets. There are many types of derivatives, ranging from call or put options on stocks to exotic derivatives such as Target Accrual Redemption Forwards (TARF) or Collateralized Debt Obligations (CDOs). For more complex derivatives, calculating the current valuation can be mathematically complex and time-consuming. For instance, many complex derivatives are valued based on Monte Carlo methods which rely on repeated random sampling and are highly computationally intensive.

Applications in Industry

Complex derivative products are found in capital markets, primarily driven by large sell-side institutions and large buy-side funds or sophisticated investors. Derivatives are pivotally important in providing enhanced risk mitigation strategies for hedging other positions. The global market for derivatives stands at $204 trillion dollars[8]. While a smaller portion of this market includes complex derivatives that would benefit from speed up, the fact remains that more quickly and accurately valuing derivatives could significantly improve profitability and reduce risk.

Industrial Importance

It is generally believed that quantum computing-based solutions can offer a quadratic speedup, meaning that if the classical simulation method requires 500,000 samples, the equivalent quantum algorithm may require only 750 samples. This anticipated speedup has garnered industry-wide attention and interest. Thus, quantum computers present a future opportunity to reduce the time necessary to perform calculations while also potentially offering higher accuracy by allowing more calculations to be run. Valuating derivatives quicker and more accurately equate to better risk management and an edge in trading, both of critical importance for players in the capital markets.

Technical Considerations

A key strategy under consideration is the use of Quantum Monte Carlo (QMC) methods to expedite existing calculations. The primary advantage of QMC methods lies in the theoretical quadratic acceleration brought about by Quantum Amplitude Estimation (QAE), a subroutine based on Grover's algorithm. At present, it's thought that many thousands of logical qubits will be necessary to surpass current classical techniques, and these extensively deep and broad circuits must be executed within seconds. Further efforts are required to prevent inefficient methods from diminishing the algorithm's speed, such as enhancing the loading of time series data distributions and performing arithmetic efficiently.

One of the key challenges in our industry is the complexity of pricing financial derivatives. Quantum computing has the potential to solve this challenge by enabling more accurate and efficient pricing models.

– FSI Industry Executive

Liquidity Optimization

Description

Liquidity optimization is the process of settling payments during the real-time gross settlement process between financial institutions. A prime example is in interbank payment systems, such as the Federal Reserve's Fedwire Funds Transfer System, where a financial institution needs to have the necessary level of liquidity before payments are processed. This means that they may have to wait for an incoming payment before making their own outgoing payment, decreasing the efficiency of the entire financial system. By reordering sequences of payments, a payment system can reduce liquidity burdens on its constituents and thereby lower opportunity costs. This reordering of payments is known to be a computationally difficult optimization problem that grows exponentially in size with the number of payments considered. This use case may also be used to optimize any cash flow intensive business or when sequencing debits and credits by an institution is pivotal to minimize opportunity costs.

Applications in Industry

This use case applies to depository banks and other institutions that maintain accounts with real-time gross settlement transfer systems, as well as those involved in clearing processes where payments are settled in real time. Due to the volume and frequency that every bank processes to central organizations daily, this use case is of fundamental importance in maximizing the efficiency of our economic system. It could potentially also apply more generally to any capital-intensive institution to optimize cash flows and liquid assets.

Industrial Importance

Existing research has demonstrated the potential for liquidity optimization to generate significant cost savings and lowered risk for payment systems. Net settlements, where credits and debits are collected for a period, netted, and settled at the end, have the advantage of requiring minimal liquidity from participants at the expense of increasing settlement and credit risk. Gross settlement removes these risks but drastically increases liquidity requirements. By optimizing the gross settlement process, participants can strike a balance between risk management and decreased capital requirements.

Technical Considerations

Two of the most promising quantum algorithms for liquidity optimization problems are Quantum Approximate Optimization Algorithm (QAOA) as well as others such as Grover's Adaptive Search (GAS). Both QAOA and GAS necessitate logical qubits due to the depth of the circuits, where a minimum of 1 logical qubit represents one variable. Given the current algorithms available for solving combinatorial optimization problems, there is likely a need for a few thousand logical qubits for successful implementation of this use case. Generally, the best classical solvers can solve these combinatorial optimization problems efficiently, ranging from a few hundred to a few thousand variables. Therefore, the scale of the industrial applications needs to be adequately large to justify a quantum approach.

We see quantum computing as a promising avenue for addressing some of the longstanding challenges like optimizing pricing models and improving customer experience. While the road ahead may be uncertain, we're eager to explore the possibilities.

– FSI Industry Executive

Portfolio Optimization

Description

Portfolio Optimization is the process of selecting the best combination of assets to maximize return and minimize risk, subject to constraints (i.e., position type, investment bands, turnover constraints, cardinality constraints, sector constraints, transaction costs). With the inclusion of thousands of exotic assets or constraints such as the impact of trades, transaction costs, and multiperiod rebalancing, portfolio optimization can become computationally intractable. The frequency of recalculation is also crucial, as it may or may not be time dependent (e.g., if portfolios are constructed intraday versus only once every few weeks).

Applications in Industry

Portfolios of assets arise in nearly every sector of finance, not only in typical areas like asset management. For example, real estate can include portfolios of properties or land that each have a measure of return and risk that need to be balanced. Portfolio optimization is equally critical in insurance, either through the direct selection of insurance contracts and markets to hold or through selecting assets whose future income streams match forecasted liability outflows. Banks and capital market institutions also have many portfolios that need to be optimized, including non-financial instruments such as selecting sets of geographic markets in which to do business or selecting which combinations of retail branches or ATMs to operate.

Industrial Importance

Given the large amount of assets under management (AUM) by investment professionals, along with relatively lighter regulation in certain subsectors such as hedge funds, even a small increase in the quality of portfolio optimization solutions can result in substantial increases in profitability. Enhanced allocation of asset weights, better forecasting of asset returns, and deeper assessment of both volatility and risk are all potential benefits to be realized from quantum portfolio optimization techniques. Classical solvers can account for hundreds – not tens of thousands – of exotic assets, and their ability to incorporate complex constraints like market impact of trades and rebalancing costs are currently limited.

Technical Considerations

Variants of the Quantum Approximate Optimization Algorithm (QAOA) and Grover's Adaptive Search (GAS) offers quadratic speedup from a computational perspective, although the wall-time may be slower given the overhead of quantum error correction. GAS has several limitations as well, including the complexity of loading data and constructing an "oracle" to recognize the best solutions. Grover's Algorithm, while providing a quadratic speedup, has significant circuit depth requirements and it is unclear how many repeated applications of GAS would be required to find high-quality solutions.

While quantum computing holds promise for solving complex computational problems, its integration into our existing systems will be challenging. We are focusing on understanding the technology before considering its adoption. One future use case could be for processing volumes of financial data in real time, enabling us to identify market trends and investment opportunities more quickly and accurately.

– FSI Industry Executive

Risk Analysis

Description

Risk analysis is the calculation of various statistics and measurements used to mitigate present and future financial risk, such as Value at Risk (VaR) calculations. Numerous types of related calculations include Conditional Value at Risk (CVaR), Credit Value Adjustment, Economic Capital Requirement, and Counterparty Risk calculations. Existing calculation methods, such as Monte Carlo simulations, are resource intensive as the analysis requires precise estimation of the distribution tails. Importance sampling has been used to attempt to diminish resource requirements, however, the calculations remain resource intensive.

Industrial Importance

Risk analysis has been heavily researched and studied as a potential use case for quantum computing. Having a more holistic view of the risk profile of any financial decision, as well as quantifying the risk faster, represents a significant competitive advantage for institutions. Research has identified a theoretical quantum improvement in speed. There is also the possibility of enhancing the number of factors that can be realistically incorporated in multivariance models such as CVaR, as well as improvements in model speed or quality that would be highly valuable given the large dollar amount and number of industries impacted.

Applications in Industry

Risk mitigation techniques are applicable to every sector in finance, as forecasting and understanding of an organization's risk posture is unequivocally the most important aspect in ensuring success. As the core requirement in finance is to make decisions today based on future uncertainties, financial entities have a vested interest in being able to more accurately, holistically, and efficiently estimate the probability and magnitude of adverse future events. Common use cases include capital requirements, counterparty and default risk, disaster and life expectancy for insurance, and demographic shifts for real estate.

Technical Considerations

Quantum algorithms that are applicable to risk analysis include Quantum Accelerated Monte Carlo (QMC). QMC had been successfully applied to risk analysis problems and has demonstrated an elevated level of accuracy and a quadradic speedup in calculations. QMC presents a theoretical quadratic speedup over classical Monte Carlo methods. It is dependent on the accurate and fast loading of probability distributions and functions to be effective and can utilize amplification techniques such as quantum amplitude estimation to achieve speedups and lower quantum resource requirements.

> *It (quantum computing) is intriguing, but its practical applications in our industry are uncertain. We are closely monitoring developments but are cautious about overestimating its near-term impact.*
>
> – FSI Industry Executive

Supervised Anomaly Detection

Description

Supervised anomaly detection is a machine learning application that utilizes extensive datasets where data records of interest, such as fraudulent or suspicious transactions or claims, have been pre-labeled. The model is trained on this labeled historical data to understand the patterns and feature types that can forecast future anomalies. Frequently, the speed and volume of data are substantially large, while the number of anomalies is extremely small. This poses a significant challenge in both obtaining sufficient training data and in the accurate and efficient training of a machine learning model.

Applications in Industry

Supervised anomaly detection is often used in fraud detection but has applications in cybersecurity, quality control, predictive maintenance, and more. Anomaly detection has been used in financial services to identify transactional fraud, abnormal trading patterns, and falsified insurance claims.

Industrial Importance

Fraud is a particularly important use case in financial services, where small percentage improvements can represent significant improvements in profitability and customer satisfaction. Financial institutions seek to identify fraud as quickly and accurately as possible, while minimizing false positives (cases where legitimate transactions are mistakenly identified as fraudulent). Yet fraud schemes are growing increasingly sophisticated and classical machine learning models can struggle to handle the complexity of financial transactions and high-dimensional data. Anomaly detection covers a broad space of use cases. While some cases require fast calculation speed, other use cases, such as over time payment fraud or insurance fraud, may prioritize result quality over calculation speed.

Technical Considerations

Quantum methods can utilize quantum-native effects which simplify the process of creating accurate models, providing a potential path to quantum advantage in the coming years. Today, several quantum algorithms appear to provide meaningful improvement compared to classical methods, for instance Quantum Neural Networks (QNN) and Quantum Support Vector Machines (QSVM) via quantum kernels have shown promising near-term results. These models have been shown to require less training data than their classical counterparts and can be implemented with relatively low-depth circuits which make the use case more feasible from a technical perspective. Newer approaches, such as Quantum Reservoir Computing (QRC), may also prove useful. QRC has a low barrier to entry and maps data features to a higher dimensional space so that the resulting model can more easily separate datapoints. These current algorithms require moderate quantum circuit depth and width, providing potential avenues for quantum advantage in the near term.

Quantum computing has the potential to transform our business model in unprecedented ways. Some of the most promising use cases to date include portfolio optimization and Quantum Machine Learning for improved fraud detection. Although the hardware is still evolving, we believe our financial future is not just digital, but quantum.

– FSI Industry Executive

Unsupervised Anomaly Detection

Description

Unsupervised anomaly detection is often used in fraud detection, but has applications in cybersecurity, application management, and more. It is a type of machine learning that identifies unusual patterns or outliers in a dataset without prior knowledge. It works by learning the typical patterns in the data and identifies items that don't fit those patterns. The advantage of these algorithms is that they can identify patterns that may not be readily apparent to humans, enabling the detection of more sophisticated and subtle issues and challenges.

Applications in Industry

Using unsupervised machine learning to detect anomalies has widespread applicability across the financial services industry. It can be used to detect suspicious trading patterns or market movements in capital markets, uncover networks of malicious actors in retail banks, as well as identify potential insurance fraud.

Industrial Importance

Unsupervised anomaly detection is done today, although there are issues with the complexity, scalability, and time required to perform the calculations. Decreasing costs or uncovering previously unseen relationships in data may provide large upsides for profitability. For financial institutions, correctly identifying anomalies can have large impacts in reducing regulatory, compliance, operational and reputational risk. As the velocity of data continues to proliferate, having capabilities to identify anomalies without labeled data will become increasingly important and directly affect the bottom line.

Technical Considerations

As it stands today, there are many classical unsupervised methods, including Autoencoders, Generative Adversarial Networks (GANs), Isolation Forest algorithms, Restricted Boltzmann Machine (RBMs), and more. These approaches are powerful, but often require extremely large datasets, careful tuning, and significant computational resources. Two quantum-based approaches include Quantum Generative Adversarial Networks (QGANs) or Quantum Restricted Boltzmann Machines (QRBMs.) These machine learning models show promise in being able to be trained on orders of magnitude less data than their classical counterparts with equal or better performance. By using properties native to quantum computers, these machine learning techniques could use less data to more accurately pinpoint anomalies.

> *Quantum computing can revolutionize machine learning and artificial intelligence in financial services to enable more accurate predictions and personalized services for our clients.*
>
> – FSI Industry Executive

Additional Classes of Use Cases

The financial services industry provides financial services to individuals and businesses. The sector includes firms such as banks, investment houses, insurance companies, and finance companies and can be subdivided into the following sub-sectors:

Collateral Optimization: Collateral optimization is the process of finding the best combination of assets to minimize opportunity costs while satisfying the increasing regulatory and margin requirements. Quantum computing could increase the speed and quality of identifying optimized collateral pools which would reduce the administrative and holding cost of collateral pools.

Creditworthiness: The identification of independent features for creditworthiness of individual or business applications. Quantum computing could provide improved identification of more factors that could impact the creditworthiness of an application.

Episodic Insurance: Episodic insurance or insurance for individual trips or events is a compelling use case. Although these products are not widespread today, quantum computing may provide the calculation speed and depth required to insure micro events.

CAT Modeling: Catastrophic event modelling, as utilized by the insurance industry, could be improved by quantum computing. Quantum improvements of this optimization problem through increasing the number of considered variables could lead to enhanced pricing accuracy, greater risk recognition, and the ability to effectively insure in higher-risk areas.

Financial Crash Prediction: Financial crash prediction for financial markets determines the cascading effects of severe changes in asset prices or black swan events. Detailed understanding of the interconnectivity of financial institutions and its effects on an organization can be enhanced by quantum computing. The application of quantum computing could improve the quality of predictive factors and analysis.

Algorithmic Trading: Algorithmic trading enhancement through machine learning of timing and price of trades presents significant improvement through quantum technologies. Quantum solutions could provide enhanced depth of algorithmic factor consideration and potentially enhanced solution speed.

As global efforts mature and the devices become more reliable, I feel confident that quantum optimization heuristics will be some of the first winners for use cases that do not currently scale well. Such examples include combinatorial optimization problems such as portfolio optimization, collateral loan obligation optimization, etc. Grover-like algorithms will also have an impact, with quantum replacements to Monte Carlo simulations, offering quadratically faster response and decision times for a variety of financial modelling scenarios.

– FSI Industry Executive

Conclusion

Quantum computing is a transformational technology that may provide compelling value propositions to the financial services industry by improving business outcomes for critical problems. As we stand on the brink of what may represent a new era in advanced computing, it is important that financial services leaders actively engage with the possibilities this technology presents.

From the improvement in the identification of fraud, acceleration of complex calculations for valuation, to better optimization of portfolios and liquidity, quantum computing may offer unprecedented opportunities for future growth and innovation. Despite the current limitations of quantum computers, the rapid pace of research and development suggests that these challenges are not insurmountable.

The analysis of over 50 potential use cases, and the identification of the six most impactful and feasible ones, provides a clear roadmap for investment and exploration. The potential of quantum computing to improve outcomes in derivative pricing, liquidity optimization, portfolio optimization, risk analysis, and supervised and unsupervised machine learning for anomaly detection, is a prospect too significant to overlook.

The future of quantum computing is likely to continue to be dynamic. Financial services leaders should plan to remain agile and open-minded, to not just adapt to, but shape the future use of quantum computers. As quantum technology continues to mature, the viability of individual use cases will evolve, presenting new opportunities and challenges. It is not just about understanding and investing in quantum computing as a one-time activity, but about shaping and leading its trajectory.

The development of quantum computers and their capabilities is continuing at a rapid pace. As we navigate this exciting frontier, let us not forget the potential it holds, not just for financial services, but for the world at large. Quantum computers are not an evolution from our traditional computers, they are something new and different. Working together, we can help shape the future of both quantum computing and the financial services industry.

We think that risk analysis based on Monte Carlo will be where quantum computing will have the greatest impact on the industry. According to Basel 3, accurate risk analysis is important. It might be mandatory to apply quantum in the future; therefore, the risk is that if we don't do it, we will also lose our right to do business. Basel 3 is already negatively impacting the Japanese financial industry.

– FSI Industry Executive

Glossary

1. Collateralized Debt Obligation (CDO):

A collateralized debt obligation (CDO) is a complex structured finance product that is backed by a pool of loans and other assets and sold to institutional investors. A CDO is a particular type of derivative because, as its name implies, its value is derived from another underlying asset. These assets become the collateral if the loan defaults. (*Source: Investopedia*)

2. Conditional Value at Risk (CVaR)

While VaR represents a worst-case loss associated with a probability and a time horizon, CVaR is the expected loss if that worst-case threshold is ever crossed. CVaR, in other words, quantifies the expected losses that occur beyond the VaR breakpoint. *(Source: Investopedia)*

3. Generative Adversarial Networks (GANs)

A generative adversarial network (GAN) is a class of machine learning frameworks and a prominent framework for approaching generative AI. In a GAN, two neural networks contest with each other in the form of a zero-sum game, where one agent's gain is another agent's loss. *(Source: Wikipedia)*

4. Grover's Adaptive Search (GAS)

Grover adaptive search (GAS) for binary optimization (BO) problems is a quantum algorithm for iteratively finding optimal solutions using Grover's search algorithm. In GAS iterative oracle constructions are needed, and a high computational demand is required. *(Source: InspireHEP)*

5. Mixed Integer Programming (MIP)

A mixed integer programming problem is a mathematical optimization problem in which some or all of the variables are restricted to be integer. MIPs are more difficult than linear programs due to the integrality constraints, and are commonly solved through Branch and Bound methods. *(Source: MILP, Wikipedia)*

6. Noisy Intermediate-Scale Quantum (NISQ):

The current state of quantum computing is referred to as the noisy intermediate-scale quantum (NISQ) era characterized by quantum processors containing up to 1,000 qubits which are not advanced enough for fault-tolerance or large enough to achieve quantum advantage. *(Source : ScienceDaily, TechSpot, Wikipedia)*

7. Parallel Tempering with Isoenergetic Clusters (PT-ICM)

PT-ICM is a computer simulation method to find the lowest energy state of a many-body system. A type of Monte Carlo method, it uses many copies of the systems, each evolving under different environments to continuously search for the best solution. At various stages, certain sets of variables are swapped to avoid local minima and continue searching for globally optimal solutions (*Source: Deloitte SMEs*)..

8. Quantum Error Correction (QEC):

Quantum error correction (QEC) is used in quantum computing to protect quantum information from errors due to decoherence and other quantum noise. Quantum error correction is theorized as essential to achieve fault tolerant quantum computing that can reduce the effects of noise on stored quantum information, faulty quantum gates, faulty quantum preparation, and faulty measurements. This would allow algorithms of greater circuit depth. *(Source: Fundamental Research, 2020, Wikipedia)*

9. Quantum Monte Carlo (QMC)

Quantum Monte Carlo methods are algorithms that involve several subroutines to achieve a quadratic speedup in number of samples to estimate statistics of unknown distributions. Used for derivative pricing and risk analysis, QMC relies on Grover's Algorithm through amplitude estimation to achieve the speedup. QMC methods typically involve loading or evolving distributions of interest, calculating and embedding desired quantities of interest into a qubit's amplitude, and retrieving the amplitude through variations of Grover's algorithm. (*Source: Deloitte SMEs*)

10. Quantum Approximate Optimization Algorithm (QAOA)

The Quantum Approximate Optimization Algorithm (QAOA) is a variational quantum algorithm for solving combinatorial optimization problems. QAOA encodes the optimization problem data into parameterized rotational gates. Some variants encode problem constraints into the circuit. It is currently thought many successive layers of the algorithm are required to find good solutions, making the overall circuit very deep *(Source: IBM Quantum Learning)*

11. Quantum Generative Adversarial Networks (QGANs)

The generative adversarial network (GAN), an important machine learning invention that excellently solves generative tasks, has also been extended with quantum versions. Many implementations of QGANs have been suggested. A QGAN may have a fully quantum or a hybrid quantum–classical architecture, which may need additional data processing in the quantum–classical interface. Similarly to classical GANs, QGANs are trained using a loss function in the form of max likelihood, Wasserstein distance, or total variation. *(Source: MDPI)*

12. Quantum Neural Networks (QNN)

Quantum neural networks are computational neural network models which are based on the principles of quantum mechanics. typical research in quantum neural networks involves combining classical artificial neural network models (which are widely used in machine learning for the important task of pattern recognition) with the advantages of quantum information to develop more efficient algorithms. *(Source: Wikipedia)*

13. Quantum Processing Unit (QPU)

Quantum Processing Units are quantum circuit and quantum logic gate-based model of computing, in addition to annealing quantum processors and analog quantum processors. *(Source: Wikipedia)*

14. Quantum Reservoir Computing (QRC)

Quantum Reservoir Computing (QRC) is an extension of the classical Reservoir Computing (RC) paradigm to the quantum domain. It leverages the principles of quantum mechanics to create a reservoir that can process and map information in a high dimensional space, potentially offering unique advantages in computational power and efficiency for downstream tasks. *(Source: QuEra)*

15. Quantum Restricted Boltzmann Machines (QRBMs)

Quantum Restricted Boltzmann Machines are single-layer quantum neural networks that can be used for implementing quantum computation tasks such as compute the wave functions for the notable cases of physical interest (i.e., the ground state) as well as the Gibbs state (i.e., thermal state) of molecules on the superconducting quantum chip. *(Source: Arxiv)*

16. Quantum Support Vector Machines (QSVM)

SVMs are a type of supervised machine learning algorithm that can be used for classification and regression tasks. They identify the hyperplane in a high-dimensional space that maximally separates different classes and are known for their robustness and ability to process large amounts of data. Quantum support vector machines (QSVMs) are a variant of SVMs that use quantum computers to perform the optimization required to find the hyperplane. *(Source: DZone)*

17. Restricted Boltzmann Machine (RBMs)

A restricted Boltzmann machine (RBM) (also called a restricted Sherrington–Kirkpatrick model with external field or restricted stochastic Ising–Lenz–Little model) is a generative stochastic artificial neural network that can learn a probability distribution over its set of inputs. RBMs have found applications in dimensionality reduction] classification, collaborative filtering, feature learning, topic modelling, immunology, and even many-body quantum mechanics. *(Source: Wikipedia)*

18. Value at Risk (VaR):

Value at risk (VaR) estimates how much a set of investments might lose (with a given probability), given normal market conditions, in a set time period such as a day. VaR is typically used by firms and regulators in the financial industry to gauge the amount of assets needed to cover possible losses. *(Source: Wikipedia)*

Citations

1. Enrique Lizaso, et. al, “Quantum Computing for Finance: Overview and Prospects,” Reviews in Physics, 2019
2. Andrés Gómez, et. al, “A Survey on Quantum Computational Finance for Derivatives Pricing and VaR.” Archives of Computational Methods in Engineering, 2022
3. Cody Googin, et. al, “A Survey of Quantum Computing for Finance.” JPMorgan Chase Bank, 2022
4. Alexander M. Dalzell, et. al, “Quantum Algorithms: A Survey Of Applications And End-To-End Complexities.” (2023): AWS Center for Quantum Computing, 2023
5. Daniel J. Egger, et. al, “Quantum Computing for Finance: State-of-the-Art and Future Prospects.” IEEE Transactions on Quantum Engineering, 2020
6. Unidentified Authorship "Quantum Computing And The Implications For The Securities Industry." FINRA, 2023
7. Ajit Desai, et. al, "Improving the Efficiency of Payments Systems Using Quantum Computing,". Staff Working Papers 22-53, Bank of Canada, 2022
8. Office of the Comptroller of the Currency, Washington D.C., “Quarterly Report on Bank Trading and Derivatives Activities,” 2023

Deloitte.