MatchPoint Research Working Paper No. 2026/T02

Quantum Computing for Asset Management: Portfolio Optimisation Beyond Classical Limits

An evidence-led framework for hybrid experimentation, validation and quantum-safe readiness

Chennakeshav (CK) Adya
Independent Researcher
July 2026

ABSTRACT

Background. Portfolio construction becomes computationally difficult when institutional mandates add cardinality, lot-size, turnover, liquidity and multi-period constraints. Quantum methods offer new approaches to combinatorial optimisation and high-precision risk estimation, while current evidence spans theory, classical simulation and physical-hardware experiments. **Objective.** This paper assesses the present evidence for quantum computing in asset management and develops a controlled operating model for institutional experimentation and quantum-safe readiness. **Approach.** The paper reviews 16 primary and authoritative sources, including foundational portfolio theory, quantum-algorithm studies, hardware experiments, BIS analysis and NIST standards. Results are classified as classical production, quantum hardware demonstration, quantum simulator evidence, theoretical result or Matchpoint scenario analysis. **Findings.** Current hardware has produced credible simplified portfolio demonstrations, including a 109-qubit IBM-Vanguard experiment and a trapped-ion hybrid pipeline using real market data. These studies rely materially on classical reduction or post-processing and do not establish broad production quantum advantage. A benchmark-first hybrid programme should use strong classical comparators, holdout instances, complete resource accounting, deterministic risk checks and accountable approval. **Implications.** Asset managers can establish reproducible experimentation capabilities and advance cryptographic inventory and post-quantum migration now. Production adoption requires decision-relevant evidence across solution quality, reliability, resources, security and governance.

Highlights

- Current evidence does not establish broad production quantum advantage in asset management.
- Physical-hardware, simulator and theoretical results require separate evidence labels.
- Every experiment needs a strong, reproducible classical baseline and holdout instances.
- Hybrid workflows require complete attribution of quantum and classical contributions.
- Cryptographic inventory and post-quantum migration are present-day resilience actions.

JEL Classification: C61, C63, G11, G17, O33

Keywords: quantum computing, portfolio optimisation, asset management, QAOA, variational quantum algorithm, quantum risk analysis, hybrid computing, post-quantum cryptography

This paper is educational research. It is not investment, legal, tax, cybersecurity or regulatory advice; it is not an offer or solicitation. It contains no client information. Hardware, simulator and theoretical results are identified separately. Public sources and standards were reviewed to 31 July 2026; each institution must verify current evidence, contracts, configurations and obligations.

1. INTRODUCTION

Asset management converts mandates, beliefs and market data into portfolios subject to return objectives, risk budgets and operating constraints. The familiar mean-variance framework begins with expected returns, a covariance matrix and a trade-off between return and variance [1]. Institutional implementation adds cardinality limits, minimum and maximum holdings, transaction costs, turnover, liquidity, tax, regulatory, concentration and multi-period constraints. These additions can produce mixed-integer, nonlinear or otherwise combinatorial problems whose difficulty grows rapidly with the size and detail of the investable universe.

Quantum computing has therefore attracted attention as a possible route to new optimisation and risk-analysis methods. The relevant evidence spans different technical categories. Some papers propose algorithms and asymptotic results. Others simulate quantum circuits on classical computers. A smaller group executes simplified portfolio problems on physical quantum processors. These categories answer different questions and should not be combined into a single claim of readiness.

This paper examines the current evidence for quantum computing in asset management and develops a controlled operating model for institutional experimentation. Its title, *Portfolio Optimisation Beyond Classical Limits*, describes the research ambition: identifying portfolio problems whose scale, structure or precision may ultimately exceed practical classical resources. The current evidence set does not establish broad production quantum advantage over leading classical portfolio solvers. The immediate institutional opportunity is to build a benchmark-first hybrid capability, preserve reproducible classical baselines, prepare quantum-compatible formulations and advance post-quantum cryptographic migration.

The paper asks four questions. First, which portfolio and risk problems have structures that quantum methods could plausibly address? Second, what has been demonstrated on physical hardware, what has been simulated, and what remains theoretical? Third, how should an asset manager test a quantum workflow against strong classical alternatives? Fourth, which security and governance actions are justified before useful large-scale quantum optimisation is available?

The analysis uses five evidence labels. **Classical production** denotes established portfolio mathematics or current operational practice. **Quantum hardware demonstration** denotes execution on a physical quantum processor. **Quantum simulator evidence** denotes results produced through classical simulation of quantum circuits. **Theoretical result** denotes a mathematical or asymptotic result under stated assumptions. **Matchpoint scenario analysis** denotes an illustrative operating case developed for this paper; it is not observed investment performance, a forecast or an allocation recommendation.

The evidence supports a measured conclusion. Physical-hardware experiments have become larger and more relevant to finance. A 2025 IBM-Vanguard preprint reports a simplified bond exchange-traded fund construction experiment using 109 qubits and circuits of up to 4,200 gates [9]. A 2026 trapped-ion preprint reports an end-to-end hybrid portfolio-selection pipeline using real market data, with quantum kernels reaching 78 qubits and 1,016 two-qubit gates [10]. Both studies depend materially on classical computation, problem reduction or post-processing. The latter reports that standalone QAOA failed to find the optimum for two larger indices in the tested results [10]. These findings are valuable demonstrations; they do not prove a general production advantage.

Current management action can proceed on two tracks. The first is decision technology: select a small number of computation-heavy use cases, define rigorous classical baselines, run reproducible hybrid experiments and stop programmes that fail their pre-agreed gates. The second is resilience: inventory cryptographic dependencies and plan migration towards the post-quantum standards finalised by the US National Institute of Standards and Technology in 2024 [15]. The resilience track concerns classical cryptography designed to resist future quantum attacks. It can advance independently of quantum portfolio optimisation.

1.1 Contribution and scope

The paper contributes an evidence ladder, a use-case screen, a benchmark protocol and an operating model designed for investment organisations. It covers gate-model methods relevant to constrained optimisation and risk estimation, including the Quantum Approximate Optimization Algorithm, variational quantum algorithms,

quantum-walk approaches and amplitude-estimation methods [2-5, 12-13]. It also reviews two recent portfolio experiments on physical hardware [9-10]. Quantum annealing is acknowledged as a related optimisation route, while the detailed evidence review focuses on the gate-model and hybrid literature in the source register.

The paper does not prescribe securities, portfolio weights or a quantum-vendor selection. It does not forecast alpha, drawdowns or financial returns. No budget, staffing or timing estimate is represented as market evidence. The readiness sequence in Section 11 is Matchpoint scenario analysis and uses qualitative gates so that an institution can apply its own approved resources and risk tolerances.

1.2 Executive conclusion

Four conclusions follow from the evidence.

First, formulation quality governs relevance. A quantum experiment should represent the actual institutional decision, including its constraints and validation rules. A small unconstrained proxy may test circuit behaviour while providing limited information about an investable workflow.

Second, the correct comparator is the strongest practical classical method available for the same problem. Classical solvers, heuristics, decompositions and warm starts form part of the production baseline. A hybrid method should be assessed end to end, including data preparation, encoding, queue time, circuit execution, measurement, classical optimisation, repair and risk validation.

Third, current quantum evidence belongs in an experimental technology portfolio. Hardware scale alone is insufficient. Connectivity, gate fidelity, circuit depth, shot count, optimiser behaviour, reproducibility and post-processing materially affect the result [6-10]. An investment committee needs decision-grade evidence on solution quality, feasibility, stability, total wall-clock time, operating cost and control burden.

Fourth, cryptographic readiness is an immediate operational programme. BIS analysis identifies both financial applications and quantum-related risks, while NIST has issued final standards for post-quantum key establishment and digital signatures [14-15]. Asset managers hold long-lived confidential data, use external administrators and custodians, and depend on extensive software and communications chains. A cryptographic inventory and dependency map therefore belong in present technology-risk planning.

2. THE PORTFOLIO-CONSTRUCTION PROBLEM

2.1 Classical foundation

Markowitz portfolio selection treats a portfolio as a vector of weights and relates those weights to expected return and variance [1]. In its simplest continuous form, a manager selects weights subject to a budget constraint and possibly long-only bounds. Under standard assumptions, this form can be solved efficiently with established convex-optimisation methods. The resulting efficient frontier shows portfolios that deliver the highest expected return for a given variance, or the lowest variance for a given expected return.

Production portfolios rarely end at this formulation. A mandate can restrict the number of holdings, set minimum lot sizes, impose sector and issuer caps, limit factor exposures, require minimum liquidity, account for trading costs and taxes, and constrain turnover from the current portfolio. Liability-aware or multi-period mandates add cash-flow matching, path-dependent restrictions and future rebalancing choices. Some constraints preserve convexity. Others introduce binary variables, non-convex objectives or discontinuities.

The computational question therefore depends on the precise formulation. Portfolio size alone does not determine difficulty. A large convex problem may be tractable using a classical optimiser, while a smaller discrete problem with tight interacting constraints may require extensive branch-and-bound, decomposition or heuristic search. A useful quantum programme begins with a problem inventory that records variables, objective terms, constraint types, data dependencies, required precision and the accepted classical solution method.

Problem class	Typical institutional features	Computational character	Required evidence before quantum testing
Continuous allocation	Full investment, bounds, linear exposure constraints	Often convex under a quadratic risk model	Reproducible classical optimum and sensitivity analysis
Cardinality-constrained selection	Maximum holdings, minimum position sizes, inclusion decisions	Mixed-integer or combinatorial	Feasible classical incumbent, bound or approximation gap
Turnover-aware rebalancing	Current holdings, trading costs, tax or market-impact terms	Can become nonlinear, non-convex or mixed-integer	Complete cost model and trade-feasibility checks
Multi-period allocation	Scenario paths, liabilities, cash flows and future decisions	Rapid state-space growth	Scenario governance and out-of-sample validation
Risk estimation	VaR, CVaR, sensitivities and tail events	Simulation-intensive at high precision	Classical Monte Carlo error, runtime and state-preparation cost

2.2 The discrete formulation

A common route into quantum optimisation maps a constrained selection problem into binary variables. Each variable can represent whether an asset is selected, whether a lot is held, or whether a constraint-related auxiliary state is active. The objective and penalties are then written as a quadratic unconstrained binary optimisation problem or as an Ising Hamiltonian. The mapping makes the problem compatible with several quantum optimisation approaches.

This translation creates practical choices. Penalty coefficients must be large enough to discourage infeasible portfolios and measured enough to preserve useful energy differences among feasible candidates. Binary encodings can require multiple qubits per asset or additional variables for constraints. Correlation matrices may create dense interactions that are expensive to embed on hardware with limited connectivity. Constraint repair after measurement can improve feasibility while adding classical work and changing the end-to-end comparison.

The objective used by the quantum circuit also needs to match the portfolio decision. An expected-return and variance score may omit turnover, liquidity or mandate restrictions that dominate the real implementation. A demonstration on a simplified objective is evidence about the simplified problem. Its relevance to production depends on a documented bridge between the demonstration and the actual mandate.

2.3 What a classical baseline contains

The classical baseline is a controlled benchmark suite. It should include a commercially or academically credible exact solver where the problem permits one, a strong heuristic for larger instances, and the current production method used by the investment team. The suite should record hardware, software versions, tolerances, random seeds, warm starts and stopping criteria. For approximate solutions, it should retain the best-known objective, feasibility result and any certified optimality gap.

Classical preprocessing is legitimate within a hybrid workflow. Warm-starting research explicitly uses a classical relaxation to initialise quantum optimisation at low circuit depth [5]. The evaluation must therefore account for every stage on both sides. If a quantum-labelled pipeline gains most of its performance from a classical reduction or repair routine, that contribution should be reported. The same routine should be tested with classical candidate generators where possible.

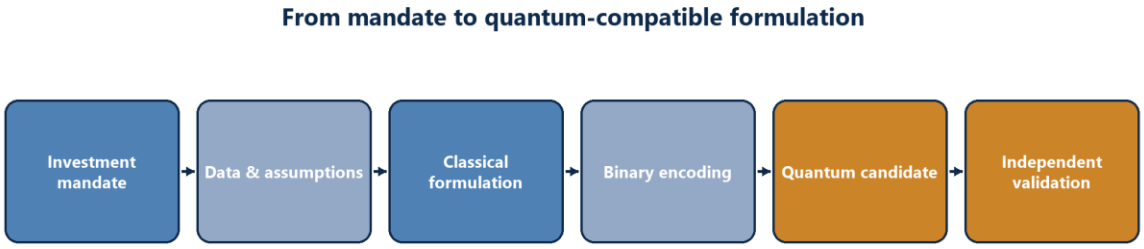

Figure 1. From investment mandate to quantum-compatible formulation
Source: Matchpoint framework.

2.4 Defining a decision-grade result

An asset manager needs more than a low objective value from one run. A decision-grade result is feasible under all hard constraints, stable enough to survive reasonable input perturbations, explainable to the investment and risk functions, reproducible under a defined protocol and operationally deliverable within the rebalancing window. It must pass independent calculations of exposures, risk, turnover and implementation costs.

Solution quality can be expressed through the objective gap relative to a best-known classical solution, a certified bound or an agreed high-quality incumbent. Feasibility should be reported before and after any classical repair. Distributional information matters because variational algorithms produce samples. The report should show success probability, dispersion across seeds, tail outcomes and sensitivity to hyperparameters. Total wall-clock time should include queueing and communications where these affect practical use.

3. WHAT "BEYOND CLASSICAL LIMITS" MEANS

3.1 Four distinct limits

The phrase can refer to at least four different limits. **Complexity limits** arise where exact solution time grows rapidly with problem size. **Memory limits** arise when scenarios, states or matrices exceed available resources. **Precision limits** arise where a simulation requires many samples to estimate a tail probability or sensitivity. **Decision-window limits** arise when a high-quality solution cannot be produced within the time allowed for rebalancing or risk action.

A quantum method can be relevant only when it addresses a defined limit and when data loading, error management and output interpretation preserve the claimed benefit. An asymptotic improvement in one subroutine may be outweighed by state preparation, circuit depth, repeated measurements or classical post-processing. Theoretical speed-up and practical advantage are separate claims.

3.2 Advantage claims

Three advantage concepts should be distinguished. **Computational advantage** means a quantum system performs a specified computation beyond the practical reach of a classical comparator under a defined resource measure. **Economic advantage** means the complete workflow produces sufficient value after cost, latency and control burdens. **Investment advantage** would mean improved realised investment outcomes. The cited portfolio literature establishes neither broad economic advantage nor investment advantage.

The evidence supports narrower statements. QAOA and related methods provide quantum-compatible approaches to combinatorial optimisation [2-7]. Physical devices have executed increasingly large simplified problems [8-10]. Quantum amplitude-estimation and gradient algorithms have theoretical error-scaling properties that may be relevant to risk computation under demanding assumptions [12-13]. Each result advances a component of the research programme.

3.3 Evidence discipline

Every result in this paper is classified by where it was produced. A physical-hardware run provides information about device execution, noise, circuit compilation and sampling on that platform. A simulator study can explore algorithm design and larger logical structures while omitting some device constraints. A theoretical paper can establish a scaling relationship under assumptions. A production claim would require an end-to-end comparison on a decision-relevant problem under operational conditions.

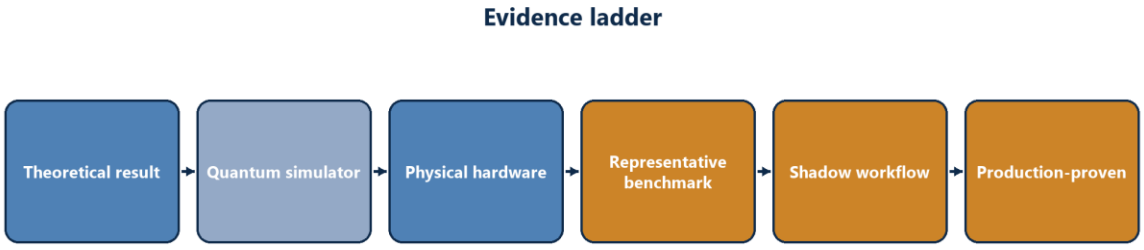

Figure 2. Evidence ladder from theory to production

Source: Matchpoint evidence framework. Current portfolio evidence reviewed in this paper does not reach the production-proven rung.

The final rung, production-proven, is intentionally demanding. It requires a representative portfolio problem, a strong classical baseline, repeatable out-of-sample performance, a complete resource account and investment-governance acceptance. The evidence reviewed here does not reach that rung. This classification preserves the value of current research while defining the work still required.

4. QUANTUM METHODS RELEVANT TO ASSET MANAGEMENT

4.1 Variational optimisation and QAOA

QAOA is a hybrid quantum-classical method proposed for approximate combinatorial optimisation [2]. A quantum circuit alternates operations associated with the problem objective and a mixing operator. A classical optimiser adjusts circuit parameters using measured objective values. The circuit produces a distribution of candidate bit strings rather than a single deterministic portfolio.

For portfolio selection, the bit strings can represent selected assets or discrete positions. Constraint handling can enter through penalty terms, a constraint-preserving mixer, encoding choices or classical repair. Each route affects the circuit and benchmark. Penalties can make the energy landscape difficult to navigate. Constraint-preserving circuits may be deeper. Repair can improve the measured candidates while moving computation back into the classical stage.

The Conditional Value-at-Risk variational objective focuses parameter training on a selected tail of favourable measured outcomes. Barkoutsos et al. report improved results for tested combinatorial instances, including portfolio optimisation examples [3]. The finding supports CVaR as a candidate training objective for specific variational settings. It does not establish a universal improvement across portfolio structures, devices or optimisers.

Warm-start QAOA incorporates a classical relaxation into the initial quantum state and mixer design [5]. The method can improve low-depth performance on tested problems, including portfolio examples. This is operationally attractive because current noisy hardware limits circuit depth. It also reinforces the hybrid nature of the workflow: classical optimisation supplies information that the quantum circuit attempts to refine or sample.

4.2 Variational quantum eigensolvers

Variational quantum eigensolver methods also search for low-energy states of a Hamiltonian through parameterised circuits and classical optimisation. Buonaiuto et al. study portfolio optimisation using a four-asset formulation encoded with 12 qubits across simulators and real devices [8]. Their experiments demonstrate material sensitivity to the ansatz, optimiser, hardware topology and device quality. The small problem is useful for evaluating design choices; it is not representative of an institutional universe.

The variational loop has several failure modes. The optimiser may encounter a flat or noisy landscape. Measurement noise introduces objective uncertainty. Hardware noise can shift the output distribution. A deeper, more expressive circuit can become less executable. A shorter circuit can lack the flexibility required to represent high-quality solutions. A benchmark therefore needs multiple seeds, repeated runs and transparent reporting of failed or infeasible outcomes.

4.3 Quantum walks

Quantum-walk optimisation uses structured movement over a graph of candidate solutions. Slate et al. formulate a portfolio problem and compare a quantum-walk approach with variational alternatives [4]. The method can encode movement among feasible or near-feasible solutions in a way that reflects problem structure. Its institutional relevance depends on scalable state preparation, circuit construction, hardware execution and comparison with classical graph-search methods.

4.4 Quantum amplitude estimation and risk

Risk calculation provides a different potential route. Classical Monte Carlo estimation typically reduces sampling error at a rate proportional to the inverse square root of the number of samples. Quantum amplitude-estimation methods can offer improved asymptotic error scaling under stated assumptions [12]. Woerner and Egger develop methods for expected value, value at risk and conditional value at risk, with a toy Treasury-bill model on hardware and a simulated two-asset portfolio [12].

The resource boundary is central. A useful speed-up requires the probability distribution and payoff to be encoded efficiently, circuits to execute at the required depth and errors to remain controlled. Current demonstrations use small models. Stamatopoulos et al. extend the theoretical discussion to market-risk gradients and provide resource estimates under fault-tolerant assumptions [13]. Their numerical studies are simulator-based. These papers motivate long-term research into high-precision risk engines and sensitivities; they do not establish present production readiness.

Method	Candidate asset-management use	Current evidence in this paper	Principal constraint
QAOA	Discrete selection and constrained allocation	Algorithm, simulation and hardware studies [2-3, 6-7, 9-10]	Depth, noise, optimiser behaviour and constraint encoding
Warm-start QAOA	Refining or sampling from a classical relaxation	Algorithm and tested portfolio examples [5]	Quality and cost of the classical start; instance dependence
Variational eigensolver	Low-energy portfolio formulations	Small real-device and simulator studies [8]	Ansatz, topology, sampling and scale
Quantum walk	Search over structured solution spaces	Portfolio algorithm study [4]	Scalable circuit construction and hardware evidence
Amplitude estimation	Expected loss, VaR, CVaR and sensitivities	Theory, toy hardware and numerical simulation [12-13]	State preparation, fault tolerance and circuit resources

5. EVIDENCE FROM CURRENT EXPERIMENTS

5.1 Hardware and algorithm constraints

Preskill's NISQ framing describes processors with noisy operations, limited circuit depth and no full fault tolerance [16]. Current experiments therefore combine quantum circuits with substantial classical computation. Device topology affects how logical interactions are mapped to physical qubits. Added routing gates increase depth and error. Gate fidelity, execution speed, sampling count and calibration drift all affect the observed distribution.

Weidenfeller et al. examine QAOA scaling on superconducting hardware with problems containing 7 and 27 decision variables [7]. The study identifies gate fidelity, gate speed, connectivity and shot count as important constraints. Brandhofer et al. benchmark portfolio optimisation with QAOA and report sensitivity to the variational form, classical optimiser, finite sampling, gate errors and readout errors [6]. Their results also distinguish easier and harder problem instances. These studies support a design principle: an institutional benchmark should report the full experimental configuration and the distribution of outcomes.

Buonaiuto et al. compare design choices on real quantum devices for a small portfolio problem [8]. Their findings show that qubit count alone cannot describe useful capability. An asset manager evaluating hardware should examine native gate performance, connectivity, compilation, calibration, queueing, circuit depth and the quality of the final feasible portfolio.

5.2 IBM-Vanguard 109-qubit experiment

Agliardi et al. report a sampling-based variational quantum scheme for a simplified bond ETF portfolio-construction problem [9]. The preprint, revised in November 2025, describes execution on 109 qubits of an IBM Heron processor with circuits reaching 4,200 gates. It reports a relative solution error of 0.49 percent for the tested experiment and uses classical local-search post-processing.

The experiment is significant for three reasons. It executes a finance-labelled optimisation problem at a larger physical-qubit scale than earlier small examples. It uses a shallow sampling-oriented design intended to work within current hardware constraints. It evaluates the quantum samples together with classical post-processing and compares results with classical methods, including CPLEX and local search [9].

Its boundary is equally important. The portfolio problem is simplified. The reported metric applies to the tested construction and experimental configuration. Classical local search contributes to the final result. The study remains a preprint at the source-review date. The evidence therefore supports a large hardware demonstration for a specific simplified problem. It does not establish general superiority over leading classical portfolio solvers.

5.3 Trapped-ion hybrid selection with real market data

Yalovetzky et al. report a 2026 end-to-end pipeline tested with real market data from four equity indices containing up to 225 assets [10]. The study uses a hybrid method, qReduMIS, in which classical graph reductions decompose or simplify the portfolio-selection problem and quantum routines address remaining kernels. Reported trapped-ion executions reach 78 qubits and 1,016 two-qubit gates [10].

The use of real market data and the separation of hybrid and standalone results improve decision relevance. The paper reports that standalone QAOA failed to find the optimum for the S&P 100 and Nikkei 225 cases in the reported experiments [10]. The hybrid pipeline achieves stronger results through material classical reductions. This evidence supports classical problem reduction as part of the current operating model. It also shows why labels should identify the contribution of each component.

The result remains a preprint and does not establish production advantage. A full institutional assessment would need independent replication, comparison with best-in-class classical solvers on matched resources, stability across dates and universes, complete wall-clock accounting and implementation-level constraints.

5.4 Simulator scale and encoding efficiency

Soloviev and Krompiec study a Pauli-correlation encoding that represents multiple portfolio variables per qubit [11]. Their simulator experiments address a problem with more than 250 variables. The work is relevant because encoding overhead can otherwise make one-variable-per-qubit mappings impractical. The reported implementation is a classical simulation of quantum circuits. Real-device execution and noise analysis are identified as future work [11]. The result belongs on the simulator rung of the evidence ladder.

5.5 Evidence summary

Study	Problem and scale	Evidence label	Reported result	Decision boundary
Barkoutsos et al. [3]	Variational optimisation including portfolio instances	Simulator and experimental algorithm evidence	CVaR objective improved tested optimisation outcomes	Instance and implementation dependent
Weidenfeller et al. [7]	QAOA with 7 and 27 variables	Physical hardware study	Scaling constraints tied to gates, connectivity and shots	Portfolio advantage not tested
Buonaiuto et al. [8]	Four assets encoded with 12 qubits	Hardware and simulator study	Device and design choices materially affected results	Example is far below institutional scale
Agliardi et al. [9]	Simplified bond ETF problem; 109 qubits; up to 4,200 gates	Physical hardware preprint	Reported 0.49 percent relative solution error with classical local search	No general quantum advantage established
Yalovetzky et al. [10]	Real data; up to 225 assets; quantum kernels up to 78 qubits	Physical hardware preprint and hybrid pipeline	Hybrid qReduMIS reported separately; standalone QAOA missed two larger optima	Material classical reduction; independent production validation required
Soloviev and Krompiec [11]	More than 250 variables through Pauli-correlation encoding	Quantum simulator evidence	Multi-variable-per-qubit formulation tested in simulation	Hardware and noise testing remain future work

The Matchpoint evidence synthesis is that no broad production quantum advantage in asset management is established by sources [3-14]. This conclusion is limited to the reviewed evidence and source-review date. It does not state that future devices or algorithms cannot produce an advantage.

6. THE HYBRID OPERATING MODEL

6.1 End-to-end workflow

The current practical architecture is hybrid. Classical systems validate data, formulate the problem, create baselines and often reduce the instance. A quantum processor generates or evaluates candidates for a defined subproblem. Classical routines decode, repair and improve those candidates. Independent risk systems recalculate the portfolio. An authorised investment process makes the final decision.

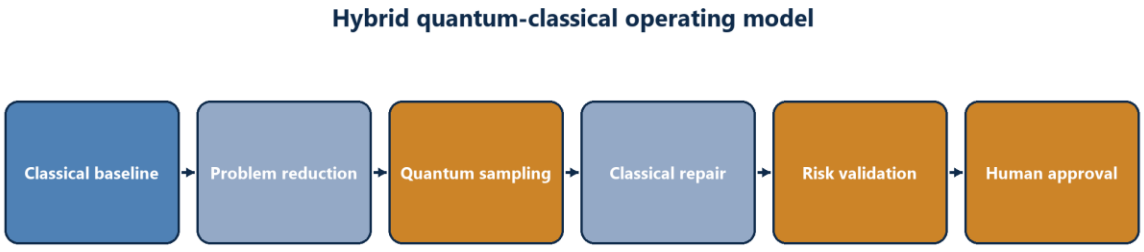

Figure 3. Hybrid quantum-classical operating model

Source: Matchpoint framework, informed by the hybrid methods in [5, 9, 10].

Each stage should produce a controlled artefact. The input stage records data provenance, valuation time and transformations. The formulation stage versions the objective, constraints, penalties and encoding. The baseline stage stores solver configuration and results. The quantum stage stores circuit, compilation, device and sampling evidence. The repair stage records changes to measured candidates. The validation stage recalculates exposures and risk without relying on the experimental code path. The release stage records human approval or rejection.

6.2 Role of the investment team

Portfolio managers define the economic objective and acceptable trade-offs. Quantitative researchers translate these requirements into mathematical formulations and test sensitivity. Technology teams manage environments, integration and reproducibility. Independent risk verifies exposures and constraints. Information-security and operational-risk functions assess data, cryptography, vendor access and concentration. Model-risk or equivalent governance reviews limitations, monitoring and change control.

The quantum component does not own the investment decision. Its output is a candidate set or analytical estimate within the established governance framework. A manager remains accountable for the mandate, market context, implementation and approval.

6.3 Reproducibility package

Every material experiment should be reproducible from a sealed package containing the input-data fingerprint, formulation version, source code, dependencies, random seeds, classical solver settings, circuit definition, compiler settings, hardware identifier, calibration snapshot where available, shot count, optimiser history, raw measurements, repair logic and independent validation output. Vendor-managed services may restrict access to some device details. Those restrictions should be documented as evidence limitations.

6.4 Measurement and failure handling

Variational results are distributions. Reporting only the best measured sample creates selection bias and obscures operational reliability. The scorecard should show feasibility before and after repair, best and median objective gap, success probability at a defined threshold, dispersion across runs, and the share of runs that fail to produce a usable candidate.

Failure is an expected research outcome. A gate should stop or redesign an experiment when the formulation cannot represent the mandate, the best classical baseline remains materially superior, results cannot be reproduced, data or security controls are inadequate, or the total control burden exceeds the potential decision value. Recorded negative results reduce repeated expenditure and improve future use-case selection.

7. USE-CASE READINESS

7.1 Constrained portfolio construction

Discrete selection has the richest portfolio-specific experimental base in the reviewed sources [3-11]. It is therefore a reasonable research use case when the institution already faces a hard combinatorial problem and can define a strong classical comparator. Readiness remains experimental. Suitable first instances are non-production shadow calculations with complete feasibility and risk checks.

7.2 Risk estimation and sensitivities

Amplitude estimation and quantum gradient methods offer theoretically attractive error scaling for high-precision calculations [12-13]. Present readiness is lower because state preparation and fault-tolerant circuit resources are substantial. An asset manager can prepare data models, resource estimates and small verification experiments. Production capital or risk decisions should continue to use validated classical calculations unless and until the complete quantum workflow passes the institution's standards.

7.3 Scenario generation and multi-period decisions

Scenario-rich allocation and dynamic hedging can experience large state spaces. Quantum algorithms may eventually contribute to sampling, optimisation or simulation subroutines. The evidence register for this paper does not establish a production method for these uses. They belong in a longer-horizon research portfolio with explicit theoretical and simulator labels.

7.4 Quantum-safe cryptography

Cryptographic migration has a different readiness profile. NIST published FIPS 203 for ML-KEM key establishment, FIPS 204 for ML-DSA signatures and FIPS 205 for SLH-DSA signatures in 2024 [15]. BIS analysis describes the need to prepare for quantum-related risks before a cryptographically relevant machine exists [14]. The programme concerns present classical systems, protocols and suppliers. It can begin with inventory, classification, vendor engagement and crypto-agility design.

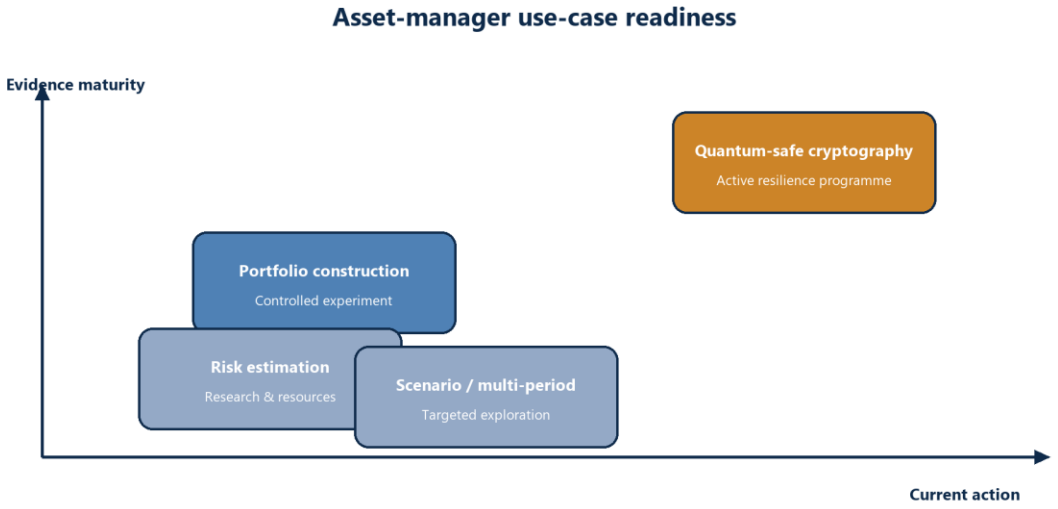

Figure 4. Asset-manager use-case readiness matrix

Source: Matchpoint evidence synthesis based on [3-15]. Readiness labels apply to the evidence reviewed to 31 July 2026.

The readiness matrix separates action types. Portfolio construction warrants controlled experiments. Quantum risk algorithms warrant research and resource estimation. Scenario and multi-period methods warrant targeted exploration. Post-quantum cryptography warrants an active resilience programme. These labels describe the evidence reviewed here and should be revisited as hardware, algorithms and standards develop.

8. BENCHMARK AND VALIDATION PROTOCOL

8.1 Start with a falsifiable hypothesis

A useful experiment begins with a claim that can fail. An example is: for a specified cardinality-constrained portfolio family, a named hybrid quantum-classical pipeline will produce feasible candidates within an agreed objective-gap threshold and decision window at an acceptable total resource cost. The hypothesis defines the problem family, quality threshold, comparator, resource measure and test period. It avoids a general promise of quantum advantage.

The protocol should be registered before results are observed. Pre-registration records instance selection, data splits, baselines, metrics, hyperparameter-search limits and stop/go rules. It reduces the risk of selecting favourable instances or changing the comparator after seeing quantum outputs.

8.2 Representative and holdout instances

Training or development instances support formulation and parameter tuning. Holdout instances provide a more credible test. They should span realistic differences in universe size, correlation density, constraint tightness, turnover pressure and market regime. Synthetic instances can isolate specific computational characteristics. Historical instances can test decision relevance. Both should avoid using future information in portfolio inputs.

A hardware result on one hand-selected instance remains useful engineering evidence. It offers limited support for generalisation. Repeated results across a documented instance distribution provide stronger evidence. Production consideration requires shadow testing through live operational cycles and independent validation.

8.3 Baseline hierarchy

The benchmark should compare at least three classical references where applicable: an exact or bound-producing solver on tractable instances, a strong heuristic on larger instances, and the existing production process. The exact solver establishes solution quality for smaller cases. The heuristic tests realistic scale. The production process measures practical relevance.

Warm starts, reductions and local search must be attributed. If a hybrid pipeline uses a classical routine, the benchmark should test that routine with classical candidate generators. This component-level comparison reveals where quality originates. The goal is a fair account of the full decision workflow.

8.4 Metric stack

The primary metric is feasible solution quality. Secondary metrics cover reliability, resources and controls. A minimum stack includes objective gap, hard-constraint violations, success probability, dispersion, total wall-clock time, quantum-processing time, classical-processing time, number of circuit executions, shot count, queue time, failure rate and cost under the applicable vendor contract. Investment-specific validation adds expected return assumptions, ex ante risk, turnover, transaction-cost estimates, liquidity and exposure stability.

Timing should begin when an approved input is available and end when a validated candidate is ready for the authorised decision process. Excluding compilation, queueing, data loading or repair can overstate practical performance. A separate technical metric may isolate quantum-processing time, provided it is labelled clearly.

8.5 Stop/go gates

Gate	Evidence required	Go condition	Stop or redesign condition
Formulation	Mandate-to-model traceability; independent constraint review	Experimental objective represents the decision usefully	Material mandate features cannot be represented or validated
Classical baseline	Versioned exact solver, heuristic and production comparison where applicable	Comparator is reproducible and accepted by quantitative review	Weak or undocumented baseline
Quantum execution	Circuit, device, shots, optimiser, seeds and raw measurements	Runs are reproducible within a stated tolerance	Results depend on undisclosed settings or cannot be repeated
Portfolio quality	Feasibility, objective gap and stability across holdouts	Pre-registered quality threshold is met	Repaired or unrepaired candidates remain materially inferior
Operations	End-to-end time, cost, access and resilience evidence	Fits the decision window and approved control model	Queueing, cost, concentration or integration burden is unacceptable
Governance	Independent validation and named accountable owner	Risks are within approved appetite	No accountable approval or unresolved material issue

The gates create evidence for continuation decisions. They do not predict future hardware progress. A stopped experiment can be retained as a reproducible benchmark and rerun if an algorithm or device change addresses the recorded failure.

9. GOVERNANCE AND VENDOR DILIGENCE

9.1 Model and investment governance

Quantum optimisation introduces familiar model risks in a new implementation. The economic assumptions, objective, covariance model and transaction-cost model may be misspecified. Penalty choices can distort preferences. Measurement and optimisation can be unstable. Classical repair can introduce undocumented biases. The final portfolio can be feasible in code while breaching a separately interpreted mandate rule.

Governance should therefore treat the complete hybrid pipeline as a model or decision-support system according to the institution's policies. The inventory entry should identify purpose, owner, users, data, theory, limitations, validation, monitoring, change thresholds and contingency arrangements. Independent validation should reproduce core results and challenge both the quantum and classical components.

9.2 Data control

Portfolio holdings, orders, strategies and client data can be highly sensitive. A vendor assessment should establish where data is processed, retained and logged; which personnel or subcontractors can access it; how tenancy is separated; how encryption keys are managed; and whether data is used for service improvement. Initial experiments can use synthetic, public or de-identified inputs while security review progresses.

Data minimisation remains useful even after approval. A circuit may require only an encoded optimisation instance rather than account identifiers or raw client records. The mapping between investment data and quantum inputs should be controlled, versioned and reversible only where necessary for validation.

9.3 Operational resilience and concentration

Cloud quantum services create dependencies on a small number of hardware and platform providers. Queue availability, regional service, API changes, calibration, pricing and support can affect an experiment. Production consideration would require recovery arrangements, an approved fallback and a tested route to the classical process.

Portability is limited by differences in native gates, connectivity, compilers and runtime services. Reproducible intermediate representations, open experiment records and modular adapters can reduce avoidable lock-in. They do not remove hardware-specific tuning. Concentration should be evaluated at the service, software-development-kit and skills levels.

9.4 Cybersecurity and access

Access should use named identities, least privilege, strong authentication and logging. Credentials must remain outside notebooks and shared research files. Code and dependencies need scanning and version control. Output channels should prevent experimental results from entering production order or risk systems without explicit approval.

Quantum-safe cryptography is related to security governance and distinct from running quantum algorithms. A vendor's use of quantum hardware does not make its communications quantum-safe. The institution should request protocol and certificate roadmaps separately from optimisation claims.

9.5 Vendor diligence questions

Area	Decision-grade question	Evidence requested
Hardware	Which physical processor, topology, native gates and calibration window produced the result?	Job record, compilation report and device metadata
Algorithm	Which stages are quantum, and which are classical reduction, optimisation, repair or validation?	Architecture, source code or independently reviewable method description
Benchmark	Which exact solvers and heuristics were used, with which settings and resource limits?	Reproducible benchmark package and complete negative results
Data	Where are inputs, logs and outputs stored; who can access them; what retention applies?	Data-flow map, contract terms and control assurance
Security	Which cryptographic protocols and post-quantum transition plans apply?	Cryptographic inventory, standards mapping and migration roadmap
Operations	What are queue, availability, cost, support, portability and exit arrangements?	Service metrics, pricing schedule, incident process and export capability
Claims	Which claims are peer reviewed, preprint, simulator-based or hardware-demonstrated?	Claim-to-source register with evidence labels

10. QUANTUM-SAFE DATA AND CRYPTOGRAPHY

10.1 Present-day exposure

Public-key cryptography supports authentication, secure communications, certificates, software signing and key establishment across the asset-management ecosystem. A sufficiently capable future quantum computer could threaten widely used public-key schemes. The date of such a machine is uncertain. Long-lived confidential information creates a present concern because an adversary could retain encrypted data and attempt decryption later.

BIS Papers No 149 describes opportunities and risks for the financial system and identifies preparation for quantum-related cyber risk as a current policy and operational issue [14]. NIST finalised its first three post-quantum standards in 2024 [15]. This combination supports a structured migration programme without relying on a forecast date for cryptographically relevant quantum computing.

10.2 Standards baseline

FIPS 203 specifies ML-KEM for key establishment. FIPS 204 specifies ML-DSA for digital signatures. FIPS 205 specifies SLH-DSA, a stateless hash-based signature scheme [15]. These are classical algorithms designed for resistance to known quantum attacks. Implementation requires protocol support, tested libraries, key and certificate management, interoperability and vendor coordination.

Standards adoption is a lifecycle programme. Asset managers depend on custodians, administrators, banks, market-data providers, cloud platforms, identity systems and software suppliers. A complete migration cannot be achieved through an isolated internal library change. Contract, procurement and architecture processes should request each critical supplier's inventory, supported algorithms, milestones and fallback arrangements.

10.3 Cryptographic inventory

The inventory should identify cryptography in applications, infrastructure, network protocols, certificates, code-signing, hardware-security modules, backups, archives and third-party integrations. Each item should record algorithm, key size, protocol, library, owner, data classification, confidentiality horizon, replacement path and external dependency.

Priority can be based on data longevity, business criticality, internet exposure and migration lead time. Long-lived client records, investment research, legal agreements and authentication roots deserve early attention. Discovery tools can accelerate inventory, while accountable owners validate findings and remove false positives.

10.4 Crypto-agility

Crypto-agility is the ability to change algorithms and parameters without replacing entire systems. It requires abstraction from hard-coded algorithms, tested configuration changes, certificate and key lifecycle support, vendor interoperability and rollback. New systems should include crypto-agility requirements at design and procurement stages. Legacy systems require prioritised remediation or compensating controls.

Migration testing should cover performance, message size, handshake behaviour, certificate chains, hardware compatibility, monitoring and failure recovery. Hybrid classical and post-quantum modes may be used during transition where approved standards and protocols support them. The institution should follow current official guidance and its regulators when choosing implementation patterns.

11. MATCHPOINT READINESS ROADMAP

11.1 Scenario status

This section is **Matchpoint scenario analysis**. It is an illustrative programme structure. It is not observed industry performance, a budget, a staffing forecast or evidence that a quantum workflow will outperform a classical system. Each institution should approve its own timing, resources and risk appetite.

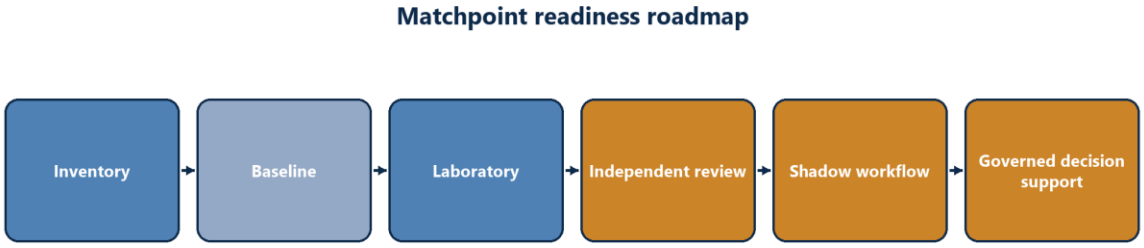

Figure 5. Illustrative readiness roadmap

Source: Matchpoint scenario analysis. The sequence is illustrative; institutional timing and resources require separate approval.

11.2 Phase 1: inventory and baseline

The programme begins by naming an accountable sponsor and separating optimisation research from cryptographic migration. The research workstream creates a portfolio-problem inventory, ranks computation-heavy candidates and records the current classical methods. The security workstream creates a cryptographic inventory and maps critical third parties.

The exit evidence is a selected use case with a falsifiable hypothesis, an accepted classical baseline, an approved data classification, a source and claim register, and a prioritised cryptographic dependency map. No quantum hardware commitment is required to complete this gate.

11.3 Phase 2: reproducible laboratory experiment

The team builds the smallest representative formulation that retains the decision's important constraints. It implements exact or bound-producing classical comparisons on tractable instances and strong heuristics on larger ones. Quantum methods are first tested in simulation for debugging and resource estimation, then on approved hardware if the prior gate is met.

The exit evidence includes a sealed reproducibility package, holdout results, attribution of classical and quantum contributions, complete resource accounting and independent quantitative review. Failed runs and weak results remain in the record.

11.4 Phase 3: shadow workflow

A qualifying method enters a non-production shadow process. It receives time-stamped inputs alongside the existing workflow and produces candidate portfolios or risk estimates. Independent systems test all constraints and exposures. Investment professionals record whether the output is useful, rejected or altered and why.

The gate considers reliability across market conditions, decision-window fit, control effort, security, vendor resilience and explainability. A quantum-labelled method does not receive a lower validation threshold. Material model or operational changes trigger revalidation.

11.5 Phase 4: governed decision support

Production consideration requires approval under the institution's model, technology, operational-risk, security and investment-governance processes. The initial role should remain bounded, reversible and monitored. The classical workflow remains an approved fallback. Monitoring tracks feasibility, solution quality, drift, failure rates, cost, latency and vendor changes.

The roadmap permits a programme to pause at any phase. Research value can come from better problem formulations, stronger classical baselines, improved data lineage and clearer cryptographic dependencies. Those outputs are independently useful and should be retained.

12. LIMITATIONS AND RESEARCH AGENDA

The evidence base changes rapidly and includes preprints. The IBM-Vanguard and trapped-ion case studies had preprint status at the source-review date [9-10]. Their methods and results should be rechecked against subsequent peer review, revisions and independent replications before a material institutional decision.

The reviewed studies use different objectives, instance structures, devices, noise conditions, resource accounting and quality metrics. Direct ranking across them would be misleading. This paper therefore classifies their evidence and boundaries rather than combining results into a meta-analysis.

Classical optimisation also advances. Any future quantum comparison should use contemporary solvers and hardware under a matched protocol. An advantage claim can become obsolete when classical decomposition, approximation, hardware or software improves. Benchmark packages should therefore be rerunnable.

Portfolio models simplify markets. Expected returns and covariances are estimated with uncertainty. Transaction costs, liquidity and market impact change with conditions. A computationally superior solution to a misspecified model can remain economically weak. Research should combine algorithmic measures with robustness, implementation and investment-governance tests.

Future work should prioritise five areas: representative constrained formulations; independent replication of hardware experiments; transparent component attribution in hybrid pipelines; end-to-end resource accounting; and standards-aligned security migration. Fault-tolerant resource estimates for risk and optimisation should state data-loading, precision and error-correction assumptions. Hardware studies should publish unsuccessful runs and calibration context where possible.

13. CONCLUSION

Quantum computing has produced credible experiments relevant to portfolio optimisation. The evidence includes small real-device studies, larger simulator formulations, a 109-qubit IBM-Vanguard hardware experiment and a trapped-ion hybrid pipeline using real market data [8-11]. These results show technical progress and clarify the role of circuit design, device properties, classical reduction and post-processing.

The reviewed evidence does not establish broad production quantum advantage in asset management. The appropriate present programme is benchmark-first, hybrid and governed. It begins with a decision-relevant formulation and a strong classical comparator. It labels theory, simulation and hardware evidence separately. It measures feasible solution quality, reliability, total resources and control burden. It retains deterministic risk validation and accountable human approval.

Post-quantum cryptography warrants parallel action. NIST's final standards and BIS analysis support cryptographic inventory, supplier engagement and crypto-agility work now [14-15]. This resilience programme can advance independently of the timetable for useful quantum portfolio optimisation.

For asset managers, readiness is an evidence capability. A disciplined organisation can identify where quantum methods deserve testing, stop weak experiments, preserve valuable benchmarks and respond coherently as the technology changes. That capability offers a sounder basis for future adoption than hardware scale or vendor claims alone.

DECLARATIONS

Evidence disclosure. Quantum-computing results cited in this paper span physical-hardware experiments, classical simulations of quantum circuits and theoretical algorithms. These categories are identified separately. The cited evidence does not establish broad production quantum advantage in asset management.

Scenario disclosure. Matchpoint scenario analysis is illustrative and is not observed investment performance, a forecast, an allocation recommendation or evidence that a quantum workflow will outperform a classical system.

Advice and conflicts. This paper is educational research. It is not investment, legal, tax, cybersecurity or regulatory advice; it is not an offer or solicitation. No securities are recommended. The author reports no client information in this paper.

Source date. Public sources and standards were reviewed to 31 July 2026. Preprint status and technical specifications should be revalidated before use in an institutional decision.

REFERENCES

[1] Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7(1), 77-91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

[2] Farhi, E., Goldstone, J. and Gutmann, S. (2014). A Quantum Approximate Optimization Algorithm. arXiv:1411.4028. https://arxiv.org/abs/1411.4028

[3] Barkoutsos, P. K. et al. (2020). Improving Variational Quantum Optimization using CVaR. *Quantum*, 4, 256. https://doi.org/10.22331/q-2020-04-20-256

[4] Slate, N. et al. (2021). Quantum walk-based portfolio optimisation. *Quantum*, 5, 513. https://doi.org/10.22331/q-2021-07-28-513

[5] Egger, D. J., Marecek, J. and Woerner, S. (2021). Warm-starting quantum optimization. *Quantum*, 5, 479. https://doi.org/10.22331/q-2021-06-17-479

[6] Brandhofer, S. et al. (2023). Benchmarking the performance of portfolio optimization with QAOA. *Quantum Information Processing*, 22, 25. https://doi.org/10.1007/s11128-022-03766-5

[7] Weidenfeller, J. et al. (2022). Scaling of the quantum approximate optimization algorithm on superconducting qubit based hardware. *Quantum*, 6, 870. https://doi.org/10.22331/q-2022-12-07-870

[8] Buonaiuto, G. et al. (2023). Best practices for portfolio optimization by quantum computing, experimented on real quantum devices. *Scientific Reports*, 13, 19434. https://doi.org/10.1038/s41598-023-45392-w

[9] Agliardi, G. et al. (2025). Portfolio construction using a sampling-based variational quantum scheme. arXiv:2508.13557v2. https://doi.org/10.48550/arXiv.2508.13557

[10] Yalovetzky, R. et al. (2026). Quantum-Informed Portfolio Selection: An End-to-End Pipeline Validated on Trapped-Ion Hardware with Real Market Data. arXiv:2607.01037. https://arxiv.org/abs/2607.01037

[11] Soloviev, V. P. and Krompiec, M. (2026). Large-scale portfolio optimization using Pauli correlation encoding. *Scientific Reports*. https://doi.org/10.1038/s41598-026-54244-2

[12] Woerner, S. and Egger, D. J. (2019). Quantum risk analysis. *npj Quantum Information*, 5, 15. https://doi.org/10.1038/s41534-019-0130-6

[13] Stamatopoulos, N. et al. (2022). Towards Quantum Advantage in Financial Market Risk using Quantum Gradient Algorithms. *Quantum*, 6, 770. https://doi.org/10.22331/q-2022-07-20-770

[14] Auer, R. et al. (2024). Quantum computing and the financial system: opportunities and risks. BIS Papers No 149. Bank for International Settlements. https://www.bis.org/publ/bppdf/bispap149.htm

[15] National Institute of Standards and Technology (2024). NIST Releases First 3 Finalized Post-Quantum Encryption Standards. FIPS 203, FIPS 204 and FIPS 205. https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards

[16] Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum*, 2, 79. https://doi.org/10.22331/q-2018-08-06-79

APPENDIX A. FORMULATION AND EVIDENCE LABELS

A.1 Illustrative binary portfolio formulation

Let binary variable x_i indicate whether asset i is selected. A simplified objective can combine expected return, covariance risk and constraint penalties:

Minimise: risk aversion multiplied by x-transpose Sigma x, less expected return multiplied by x, plus penalty terms.

The formulation can include a cardinality condition requiring a defined number of selected assets. Additional binary variables or penalty terms can represent sector, lot-size or turnover restrictions. This text is a structural illustration. Any production formulation requires approved parameter definitions, scaling, data and independent validation.

A.2 Evidence labels for committee papers

- **Classical production:** established mathematics or a validated operational method.
- **Quantum hardware demonstration:** execution on a named physical processor under a reported configuration.
- **Quantum simulator evidence:** circuit behaviour produced through classical simulation.
- **Theoretical result:** mathematical or asymptotic evidence under stated assumptions.
- **Matchpoint scenario analysis:** an illustrative governance or operating design created for this paper.

APPENDIX B. EXPERIMENT RECORD

An experiment record should contain the decision hypothesis; owner; review date; data fingerprint; objective; constraints; encoding; penalty coefficients; classical comparators; solver settings; random seeds; quantum circuit; compiler; device; calibration context; shots; optimiser; raw measurements; repair method; independent risk calculations; wall-clock components; cost; failures; reviewer comments; exceptions; and stop/go disposition.

The retained record should support rerunning the benchmark when hardware, software or classical methods change. It should also allow an independent reviewer to reproduce the principal claim without relying on a vendor presentation.

APPENDIX C. GLOSSARY

Amplitude estimation. A family of quantum algorithms for estimating an amplitude or probability, with theoretical error-scaling properties relevant to simulation and risk under stated assumptions.

Ansatz. The parameterised circuit form used by a variational quantum algorithm.

Classical repair. A classical routine that transforms a measured candidate into a feasible or improved solution.

Crypto-agility. The capability to replace cryptographic algorithms and parameters through controlled changes.

Hamiltonian. A mathematical operator whose energy structure can encode an optimisation objective.

NISQ. Noisy intermediate-scale quantum; a framing for processors with noise, limited depth and no full fault tolerance [16].

Post-quantum cryptography. Classical cryptographic algorithms designed to resist attacks from future quantum computers.

QAOA. Quantum Approximate Optimization Algorithm; a hybrid variational method for approximate combinatorial optimisation [2].

Qubit. A quantum information unit. Physical-qubit count alone does not measure usable application performance.

Shot. One execution and measurement of a quantum circuit. Repeated shots estimate an output distribution.

State preparation. The process of encoding an input distribution or data structure into a quantum state.

Variational algorithm. A hybrid method in which a classical optimiser updates parameters of a measured quantum circuit.