Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Quantum Algorithms for Computational Economics: A Critical Review of Simulation, Optimization, and Near-Term Evidence

Diego A. Diaz

diegodiaz@uchicago.edu

Universidad del Desarrollo

Research Article

Keywords: quantum computing, computational economics, quantum Monte Carlo, QUBO, dynamic programming, quantum optimization

Posted Date: October 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9961695/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Quantum Algorithms for Computational Economics: A Critical Review of Simulation, Optimization, and Near-Term Evidence

Diego A. Diaz[1*]

[1*]School of Business and Economics, Universidad del Desarrollo, Santiago, Chile.

Corresponding author(s). E-mail(s): diegodiaz@uchicago.edu;

Abstract

Quantum computing is increasingly discussed as a possible source of new computational methods for economics, but its relevance depends on whether economic bottlenecks match quantum algorithmic primitives. This critical review evaluates that match for computational economics. It synthesizes work on quantum amplitude estimation and quantum Monte Carlo for simulation and risk measurement; quadratic unconstrained binary optimization (QUBO), Ising formulations, quantum annealing, and the quantum approximate optimization algorithm for discrete optimization; and quantum linear-algebra and machine-learning primitives for econometrics and model solution. The review argues that there is no general-purpose quantum algorithm for economics analogous to Shor's algorithm for factoring. Credible claims are problem-class specific and depend on state preparation, output requirements, conditioning, hardware noise, embedding overhead, and comparison with strong classical and quantum-inspired baselines. The most mature economic applications are derivative pricing and financial risk, payment-system liquidity optimization, quantum Monte Carlo for stress testing, and early quantum-annealing formulations of dynamic programming. Transportation and infrastructure applications are relevant mainly as networked economic optimization problems. The article proposes an evidence ladder for quantum computational economics and a research agenda centered on reproducible benchmarks, economically meaningful objective functions, end-to-end resource accounting, and evaluation by economic value rather than hardware novelty.

Keywords: quantum computing; computational economics; quantum Monte Carlo; QUBO; dynamic programming; quantum optimization

JEL Classification: C02; C15; C61; C63; C88; E37; G11; G17; R41

1 Introduction

Computational economics studies economic questions whose analysis depends on numerical approximation, simulation, optimization, algorithmic modeling, and increasingly large-scale computation. Macroeconomic models use perturbation, projection, simulation, deep learning, and dynamic programming. Finance relies on Monte Carlo methods, portfolio optimization, scenario generation, and risk measurement. Industrial organization, market design, public finance, supply chains, payments systems, and urban transportation involve integer choices, networks, strategic interaction, congestion, and equilibrium constraints. These are not merely implementation details. Computational feasibility shapes the economic questions that can be asked, the models that can be estimated, and the policies that can be evaluated.

Quantum computing has entered this methodological landscape because it offers a different model of computation, not simply a faster version of classical hardware. Shor's algorithm gives an exponential quantum speedup for integer factorization and discrete logarithms (Shor, 1994). Grover's algorithm gives a quadratic improvement for unstructured search (Grover, 1996). Quantum algorithms for amplitude estimation, linear systems, Hamiltonian simulation, and combinatorial optimization have created a broader algorithmic toolbox (Brassard et al., 2002; Harrow et al., 2009; Montanaro, 2015, 2016; Hadfield et al., 2019). Yet these primitives do not automatically imply useful methods for economics. A model can be economically important and still lack the input-output structure required for quantum advantage. Conversely, a small subroutine inside a larger economic workflow may be a plausible quantum target even if the surrounding model remains classical.

The central claim of this review is that quantum computing for economics should be understood as a set of mappings between economic computational bottlenecks and quantum algorithmic primitives. It is not a unified "quantum economics" paradigm and it is not a general solution to non-convexity, high dimensionality, or equilibrium computation. This distinction matters for computational economics because the relevant unit of evaluation is not a quantum circuit in isolation. The relevant unit is the end-to-end economic workflow: data preparation, model encoding, numerical approximation, solution extraction, sensitivity analysis, and interpretation of an output in economic terms.

Several streams of research now make this assessment possible. Quantum finance studies derivative pricing, risk measurement, portfolio optimization, arbitrage, stochastic modeling, and machine learning (Orus et al., 2019; Egger et al., 2020; Herman et al., 2023; Bunescu and Vartei, 2024; Zhou, 2026). Quantum Monte Carlo methods have been proposed for economic stress testing and macroeconomic deep-learning workflows (Skavysh et al., 2023). Dynamic programming has been reformulated for quantum annealing in macroeconomic applications (Fernandez-Villaverde and Hull, 2023, 2026). Payments-system studies use hybrid quantum optimization for liquidity-saving payment reordering (McMahon et al., 2024; Priazhkina et al., 2024). Operations and transportation papers formulate traffic control, vehicle routing, and transport-network design problems as QUBO or Ising models (Dixit and Niu, 2023; Singh et al., 2021; Mohammed et al., 2025; Salloum et al., 2025). Expository work has also begun to translate quantum technology for economists (Hull et al., 2024).

This literature has moved beyond pure speculation. It includes peer-reviewed articles in economics and management outlets, including *Journal of Economic Dynamics and Control*, *Management Science*, *Quantitative Economics*, and *Computational Economics*. It also includes implementation papers and reviews in quantum information, finance, and engineering venues. The emergence of these papers justifies a critical review for computational economists. At the same time, the field remains prone to overstatement. Terms such as "speedup" and "advantage" are often used for asymptotic query complexity under idealized assumptions. Economic computation, however, is end-to-end. Quantum amplitude estimation may reduce the number of samples required to estimate an expectation, but the analyst must still encode distributions, implement payoff functions reversibly, manage circuit depth, and extract a useful statistic. Quantum annealing may sample low-energy configurations of an Ising Hamiltonian, but the economist must formulate the original economic problem, choose penalty weights, embed the problem on hardware, compare with integer programming and heuristics, and interpret feasibility in the original model.

This review contributes to computational economics in four ways. First, it identifies the quantum primitives that appear most often in economic applications and states the limitations attached to each. Second, it maps economic problem classes to those primitives, distinguishing credible problem-specific pathways from broad claims about quantum advantage. Third, it evaluates the maturity of evidence across finance, macroeconomics, dynamic programming, payments, and networked public-sector applications. Fourth, it proposes an evidence ladder and benchmark agenda for future work. The aim is neither to promote nor dismiss quantum computing, but to provide computational economists with criteria for judging when a quantum or hybrid quantum-classical method is economically meaningful.

The article proceeds as follows. Section 2 defines the review scope. Section 3 introduces the relevant quantum primitives. Section 4 maps economic problem classes to those primitives. Section 5 examines QUBO and Ising formulations. Section 6 reviews finance and risk applications. Section 7 evaluates macroeconomics, stress testing, payments, and dynamic programming. Section 8 discusses networked public-sector and infrastructure applications. Section 9 explains why there is no Shor-like algorithm for economics. Section 10 summarizes maturity and limitations. Section 11 proposes a research agenda.

2 Review scope and approach

This article is a critical review rather than a systematic review or formal meta-analysis. The literature does not yet contain enough comparable empirical performance studies to support pooled estimates of effect sizes. The purpose is interpretive: to identify the conceptual terrain, categorize problem types, evaluate maturity, and distinguish credible algorithmic pathways from speculative claims. The review therefore focuses on problem-algorithm fit and end-to-end economic usefulness rather than on bibliometric enumeration. It does not report database search counts or PRISMA-style screening statistics, because those would imply a systematic-review protocol that was not used.

The corpus was assembled from three bodies of work. The first consists of foundational quantum-algorithm papers and reviews on amplitude estimation, quantum Monte Carlo, quantum approximate optimization, quantum annealing, QUBO or Ising formulations, and quantum linear systems. The second consists of papers that explicitly apply quantum computing to economic models, finance, payments, macroeconomics, dynamic programming, or economic theory. The third consists of papers on optimization problems that are central to economic or policy computation even when the venue is engineering, physics, operations research, or quantum information. Transportation optimization is included not because this review is primarily an urban-policy article, but because routing, traffic control, and network design are economically important constrained optimization problems with congestion externalities and public-sector decision value.

Several areas are deliberately excluded or treated only briefly. Quantum decision theory, quantum cognition, and quantum probability models of preferences use quantum mathematical formalisms to represent ambiguity, bounded rationality, or contextual probability. Those literatures are relevant to social science but are not primarily about using quantum computers to solve economic optimization, simulation, or estimation problems. Post-quantum cryptography, quantum money, and quantum communication are also important for financial infrastructure, but they fall outside the computational focus except where payment systems are optimized using quantum methods. Claims about quantum hardware markets and investment forecasts are included only when they help clarify the economic concept of quantum advantage.

A key organizing distinction is between fault-tolerant and near-term quantum computation. Fault-tolerant devices would use error correction to execute deep circuits reliably. Many theoretically attractive algorithms, including coherent amplitude estimation and linear-system algorithms, are most naturally interpreted in that regime. Near-term noisy intermediate-scale quantum (NISQ) devices have limited qubit counts, limited connectivity, noise, and shallow circuit depth (Preskill, 2018). Quantum annealers are specialized devices for optimization rather than universal gate-model computation (Johnson et al., 2011). Many economic demonstrations use simulators, annealers, or hybrid algorithms. This distinction matters because a speedup theorem for an ideal algorithm does not imply practical superiority on current hardware.

The review also distinguishes quantum, quantum-inspired, and hybrid methods. Quantum methods run circuits or annealing processes on quantum hardware. Quantum-inspired methods use mathematical ideas from quantum information but run classically. Hybrid methods combine classical decomposition, preprocessing, feasibility repair, or parameter optimization with a quantum subroutine. Most plausible economic workflows are hybrid: a classical modeler formulates a constrained problem, maps part of it to QUBO or a quantum circuit, uses a quantum device or simulator to generate candidate solutions or estimates, and then applies classical validation and economic interpretation. This hybridity complicates evaluation because performance gains may come from classical preprocessing rather than from the quantum component.

The evaluative standard used here is economic relevance. A quantum method is more relevant when it improves a bottleneck that prevents economists from answering a substantive question, when the improvement survives comparison with high-quality

classical methods, and when the output is usable in an economic decision. A lower QUBO energy is not necessarily a better policy if the corresponding bitstring violates feasibility, distributional, institutional, or regulatory constraints. A faster expectation estimate is not necessarily valuable if model misspecification dominates sampling error. A toy hardware demonstration can be scientifically useful, but it should not be interpreted as operational readiness.

3 Quantum primitives relevant to computational economics

The primitives most relevant to computational economics are amplitude amplification and estimation, quantum Monte Carlo, quantum linear algebra, quantum approximate optimization, and quantum annealing. Quantum simulation in the physical-science sense is less central to the economic literature reviewed here, although it may matter indirectly for energy, materials, and climate-economy modeling.

Amplitude amplification generalizes Grover search. If a desirable event has probability a, classical repeated sampling generally requires order $1/a$ trials to observe it with high probability. Amplitude amplification can reduce this to order $1/\sqrt{a}$ oracle uses under appropriate assumptions (Grover, 1996; Brassard et al., 2002). Amplitude estimation goes further by estimating a probability or expectation encoded in a quantum amplitude. In idealized settings, quantum amplitude estimation can achieve error scaling of order $O(M^{-1})$ in the number of quantum queries, compared with $O(M^{-1/2})$ for classical Monte Carlo (Brassard et al., 2002; Montanaro, 2015). This is why quantum Monte Carlo appears repeatedly in finance and economics. Derivative pricing, expected losses, Value-at-Risk (VaR), Conditional Value-at-Risk (CVaR), simulated moments, and welfare under stochastic shocks are all expectation-estimation tasks.

For economists, the first catch is that amplitude estimation is not a drop-in replacement for Monte Carlo. The random variable must be encoded into a quantum state or circuit, the payoff or loss function must be evaluated coherently, and the relevant amplitude must be estimated without losing the asymptotic gain through preparation or readout costs. Iterative and maximum-likelihood variants reduce some circuit-depth requirements, but practical value remains hardware dependent (Suzuki et al., 2020; Grinko et al., 2021). In economic terms, amplitude estimation reduces one margin of computational cost while potentially increasing fixed costs of model encoding and validation.

Quantum linear-algebra algorithms, especially the Harrow-Hassidim-Lloyd (HHL) algorithm, are relevant because econometrics and macroeconomics use matrix inversion, least squares, projection methods, filtering, principal components, and high-dimensional approximation (Harrow et al., 2009; Hull et al., 2024). The theoretical promise is that certain sparse, well-conditioned systems can be solved in time polylogarithmic in dimension under assumptions about input access and output requirements. For empirical work, the output requirement is the hard part. Researchers often need explicit coefficient vectors, standard errors, diagnostics, and interpretability, not merely a quantum state proportional to a solution vector. If reading out all entries eliminates the speedup, the application must be reformulated around global functionals or downstream

quantum subroutines. Quantum linear algebra may therefore be more useful inside larger quantum workflows than as a direct substitute for ordinary least squares.

Quantum approximate optimization algorithms address combinatorial optimization. The quantum approximate optimization algorithm (QAOA) and related quantum alternating-operator ansatz methods apply alternating problem and mixing unitaries, with classical optimization of variational parameters, to produce approximate solutions for discrete problems such as MaxCut (Hadfield et al., 2019). Quantum annealing uses a different computational model: a system is initialized in an easy ground state and evolved toward a problem Hamiltonian whose low-energy states encode good solutions. Commercial quantum annealers implement programmable Ising models with limited connectivity and precision (Johnson et al., 2011). Both QAOA and annealing are relevant to economics because many decisions are discrete: select assets, allocate projects, schedule payments, assign workers, route vehicles, or choose policy actions on a grid.

The common bridge from economic optimization to quantum optimization is the Ising or QUBO formulation. A QUBO problem has the form

$$\min_{x \in \{0,1\}^n} x^\top Q x, \tag{1}$$

where x is a vector of binary variables and Q encodes linear and pairwise quadratic effects. Through the transformation $x_i \in \{0,1\}$ to $z_i \in \{-1,1\}$, the same problem can be written as an Ising Hamiltonian,

$$H(z) = \sum_i h_i z_i + \sum_{i<j} J_{ij} z_i z_j. \tag{2}$$

Ising formulations exist for many NP-hard problems (Lucas, 2014). QUBO is attractive because it is expressive and maps naturally to annealing and some gate-model optimization algorithms (Kochenberger et al., 2014; Glover et al., 2019). The danger is that almost any binary problem can be forced into QUBO by adding penalties, but not every formulation is computationally useful. Penalty weights can distort the energy landscape, embedding can require many physical qubits for each logical variable, and constraints that are natural in mixed-integer programming may become awkward quadratic penalties.

The final primitive is conceptual: quantum economic advantage. In computer science, quantum advantage usually means completing a computational task that classical computers cannot perform within feasible resources. Bova et al. (2023) shift the discussion toward cost structures, market competition, and use cases. That shift is useful for economics. A bank, central bank, exchange, firm, or public agency may care less about a proof of asymptotic dominance than about lower cost, lower latency, better solution quality under a time limit, or better risk estimates for an operational decision. Computational economics should adopt a similar standard: value should be measured in welfare accuracy, risk reduction, capital savings, liquidity savings, improved allocation, or better policy robustness, not only in wall-clock time or qubit count.

4 Economic problem classes and quantum mappings

Economic problems are not good quantum targets merely because they are complex. They become plausible targets when their mathematical structure matches a quantum primitive and when the output can be used economically. Table 1 summarizes the main mappings.

Table 1: Economic problem classes and plausible quantum primitives

Problem class	Economic examples	Quantum primitive	Current maturity
Expectation estimation	Derivative pricing, stress testing, simulated moments, risk measures, welfare under stochastic policy	Amplitude estimation; quantum Monte Carlo	Strong theoretical primitive; finance and early economics demonstrations
Discrete constrained optimization	Portfolio selection, combinatorial auctions, matching, project selection, payment scheduling, facility location	QUBO/Ising, quantum annealing, QAOA	Active demonstrations; practical advantage not generally established
Dynamic programming	Bellman equations, asset accumulation, dynamic discrete choice, heterogeneous-agent approximations, dynamic games	Annealing for discrete policy or value iteration; amplitude estimation inside expectations	Emerging computational-economics publications; small-scale examples
Network optimization	Payment systems, supply chains, transportation, routing, congestion, infrastructure choices	QUBO/Ising, annealing, hybrid decomposition	Active in operations and infrastructure; economic integration uneven
Linear algebra and regression	Least squares, filtering, projection, principal components, large sparse systems	Quantum linear systems; quantum PCA	Theoretically relevant; input-output barriers significant
Machine learning and sampling	Macro deep learning, credit risk, classification, generative scenario modeling	Quantum kernels, variational circuits, quantum sampling	Experimental; economic interpretation and benchmarking unresolved

The first class is expectation estimation under uncertainty. Classical Monte Carlo estimates an expected value by sampling and averaging. Many economic tasks have this form: expected utility, asset pricing, insurance loss, option pricing, default probabilities,

simulated method of moments, expected welfare under policy experiments, and stress-test losses. Quantum amplitude estimation is relevant because it offers a quadratic improvement in sampling complexity under ideal assumptions (Brassard et al., 2002; Montanaro, 2015). This does not solve model misspecification or data limitations, but it can matter when repeated evaluation of a well-specified stochastic model is the binding computational constraint.

The second class is discrete constrained optimization. Economics contains many binary or integer decisions: choose assets, accept bids, schedule payments, assign workers, clear markets, allocate seats, open facilities, and choose infrastructure projects. Some objectives are naturally quadratic because they involve pairwise interactions, such as covariance risk, complementarities, conflicts, and congestion. Others can be transformed into QUBO by adding auxiliary variables and penalties. Quantum annealing and QAOA are relevant here, as are classical and quantum-inspired QUBO solvers. The strongest economic examples are portfolio optimization, payment reordering, and discretized dynamic programming (Phillipson and Bhatia, 2021; McMahon et al., 2024; Fernandez-Villaverde and Hull, 2026).

The third class is non-convex optimization and global search. Economic models often contain non-convexities caused by indivisibilities, fixed costs, increasing returns, network externalities, discrete choices, complementarity, or multiple equilibria. Quantum optimization is sometimes described as useful for non-convex landscapes, but this claim must be narrowed. Quantum methods do not make arbitrary non-convex optimization easy. They may help for specific encodings, especially binary quadratic forms or structured oracle problems. Practical performance depends on instance families and benchmarks.

The fourth class is dynamic programming. Many economic models solve Bellman equations of the form

$$V(s) = \max_{a \in A(s)} \{u(s,a) + \beta E[V(s') \mid s, a]\}. \tag{3}$$

The curse of dimensionality arises because states, actions, shocks, and heterogeneity grow rapidly with economic realism. Quantum methods can enter dynamic programming through expectation calculations inside the Bellman operator, through discrete action-selection subproblems, or through QUBO formulations of policy or value iteration. The work of Fernandez-Villaverde and Hull (2023, 2026) is important because it turns dynamic programming itself into a quantum-annealing target and demonstrates the idea on economic examples.

The fifth class is network and graph optimization. Payment systems, supply chains, transportation systems, and production networks involve flows, cuts, paths, matching, congestion, and cascading effects. Some graph problems have direct Ising formulations; others require decomposition or mixed-integer reformulation. Payments-system liquidity, for example, can be improved by reordering indivisible payments, a combinatorial problem over sequences (McMahon et al., 2024). Transportation-network design can be formulated as a bilevel problem in which planners choose infrastructure or controls and users respond through route choice; quantum approaches typically target simplified or upper-level discrete subproblems (Dixit and Niu, 2023; Mohammed et al., 2025).

The sixth class is statistical estimation and machine learning. Quantum machine learning is broad and contested. For computational economics, the question is not whether a quantum classifier performs well on generic benchmark data, but whether quantum methods improve high-dimensional empirical tasks where classical econometrics or machine learning is a binding bottleneck. Potentially relevant primitives include quantum kernels, quantum linear algebra, quantum principal components, and quantum sampling (Hull et al., 2024). Economic adoption will require causal interpretation, uncertainty quantification, robustness to non-stationarity, and transparent model comparison. At present, this stream is less mature for economics than Monte Carlo and QUBO optimization.

A practical implication follows. The best early opportunities are not necessarily the largest models in economics. They are high-value problem fragments with repeated solution, clear objective functions, structure that maps cleanly to quantum primitives, and outputs that can be summarized compactly. Risk metrics, stress-test expectations, payment batches, portfolio-selection subproblems, and discrete policy choices may be better targets than full general-equilibrium models.

5 QUBO, Ising models, and economic discrete choice

QUBO is the most concrete near-term bridge between economic optimization and quantum hardware. Many economic decisions are binary, and many economic interactions are pairwise. In a QUBO model, diagonal elements of Q represent linear costs or benefits associated with choosing option i, while off-diagonal elements represent interactions between choices i and j. Positive interactions can encode penalties for incompatible choices; negative interactions can encode complementarities. Constraints are incorporated by penalty terms. For example, a budget constraint can be represented by adding $\lambda(\sum_i c_i x_i - B)^2$ to the objective, which expands into linear and quadratic terms. Cardinality, assignment, and capacity constraints can be treated similarly.

Portfolio optimization is the most familiar economic example. In a binary mean-variance model, x_i indicates whether asset i is selected, μ_i is expected return, and Σ is the covariance matrix. A stylized objective is

$$\min_{x\in\{0,1\}^n} \lambda x^\top \Sigma x - \mu^\top x + \rho \left(\sum_i x_i - K \right)^2 , \tag{4}$$

where K is a cardinality target and ρ is a penalty weight. The variance term is quadratic, the return term is linear, and the cardinality penalty is quadratic. This maps naturally to QUBO. It is therefore unsurprising that portfolio optimization is a common quantum-finance demonstration (Phillipson and Bhatia, 2021; Herman et al., 2023; Zhou, 2026). However, the standard continuous Markowitz problem is not the difficult part when convex constraints are used (Markowitz, 1952). The harder cases arise with cardinality, transaction costs, lot sizes, taxes, turnover limits, minimum holdings, sector constraints, and nonlinear risk measures. QUBO can encode some of these features, but may also create large penalty-dominated models.

Combinatorial auctions and market design provide another natural QUBO setting. Bids over bundles contain complementarities: the value of obtaining two goods together may exceed the sum of their separate values. Winner determination is a hard integer optimization problem. QUBO can represent conflicts among incompatible bids and benefits from accepting bids. School assignment, spectrum allocation, electricity-market unit commitment, and procurement problems contain related structures. The challenge is that economic mechanisms are not merely optimization problems. They also require incentive compatibility, fairness, legal constraints, and robustness against manipulation. A quantum optimizer that finds a high-value allocation does not by itself design a mechanism.

Payment systems offer a particularly compelling QUBO-style application because the decision is operational and recurring. High-value payments are indivisible, liquidity is costly, and settlement delays are costly. Reordering a batch of payments can reduce liquidity needs while preserving acceptable delays. McMahon et al. (2024) develop a hybrid quantum-annealing algorithm for payment reordering and apply it to Canadian high-value payment data. The importance of this work is not only reported liquidity savings, but also the structure of the application: a clear objective function, real institutional data, repeated operational use, and a compact output.

Dynamic programming can also be discretized into binary decision structures. On a finite grid, policy choices over states can be represented by binary variables indicating which action is selected in each state. Bellman residuals, feasibility conditions, and consistency constraints can be translated into penalties. Fernandez-Villaverde and Hull (2023, 2026) show how quantum annealing can be used to solve dynamic programming problems in economics, including a real-business-cycle model. This is conceptually significant because dynamic programming is usually associated with recursive numerical analysis, not QUBO. The broader lesson is that QUBO can serve as a representation layer for parts of dynamic economic computation, although current examples remain small.

The limits of QUBO are as important as its appeal. First, QUBO is unconstrained only after constraints have been transformed into penalties. If penalties are too low, the solver may return infeasible solutions. If penalties are too high, the energy landscape becomes dominated by feasibility rather than objective quality, making it hard to distinguish good feasible solutions. Second, hardware connectivity matters. Dense QUBO models may require minor embedding on sparse annealer graphs, using chains of physical qubits to represent one logical variable. This reduces effective problem size and may introduce chain breaks. Third, classical solvers are strong. Fair baselines include commercial mixed-integer solvers, branch-and-bound, cutting planes, simulated annealing, tabu search, local branching, decomposition, and problem-specific heuristics. A quantum method must be compared against these baselines, not against exhaustive search or naive local search.

A further difficulty is interpretability. In economics, a solution is rarely accepted merely because it optimizes a scalar objective. Researchers and policymakers ask which constraints bind, who gains and loses, what happens under alternative parameters, and whether the solution is robust. QUBO formulations can obscure these questions

because economically meaningful constraints are embedded as penalty terms. Computational economists should therefore report the original economic model, the QUBO transformation, penalty calibration rules, scaling choices, feasibility rates, solution repair procedures, and sensitivity. The transformation is a modeling decision, not a neutral technical step.

6 Finance, risk, and simulation

Finance is the most developed area of quantum computing applied to economic problems. Financial computation contains several tasks that map cleanly to quantum primitives. Derivative pricing and risk analysis are expectation-estimation problems. Portfolio selection and arbitrage are optimization problems. Credit risk, fraud detection, and scenario generation involve high-dimensional data and stochastic processes. Reviews by Orus et al. (2019), Egger et al. (2020), Herman et al. (2023), Bunescu and Vartei (2024), and Zhou (2026) organize this terrain around simulation, optimization, and machine learning.

Derivative pricing is the canonical Monte Carlo application. A derivative payoff depends on future stochastic variables, and the price is an expected discounted payoff under an appropriate measure. Classical Monte Carlo converges slowly: reducing error by a factor of ten generally requires one hundred times as many samples. Quantum amplitude estimation promises a quadratic reduction in sample complexity. Rebentrost et al. (2018) present a quantum computational-finance algorithm for Monte Carlo pricing of derivatives. Woerner and Egger (2019) extend this logic to risk analysis, including VaR and CVaR, and discuss implementation trade-offs between convergence rate and circuit depth. Matsakos and Nield (2024) move toward end-to-end risk-factor modeling by incorporating scenario generation into quantum circuits for equity, interest-rate, and credit risk factors.

The economic relevance of these papers is substantial because Monte Carlo simulation is a genuine bottleneck in financial risk management. Banks, insurers, clearinghouses, and regulators compute risk under many scenarios and confidence levels. Tail-risk quantities are especially expensive because rare events require many samples. A quadratic sampling improvement could be valuable if it reduces capital-model runtimes, enables more stress scenarios, or improves intraday risk management. The assumptions are demanding, however. The stochastic process must be represented in a quantum circuit, the payoff or loss function must be implemented reversibly, and the statistic must be extracted. Near-term demonstrations remain simplified relative to production risk engines.

Portfolio optimization is the leading optimization application. Quantum-finance papers often formulate portfolio selection as QUBO or Ising optimization, especially when assets are selected discretely or constraints create non-convexity. Phillipson and Bhatia (2021) use a D-Wave annealer and hybrid solvers for portfolio optimization, comparing outcomes to classical commercial solvers. This is representative of the field: realistic integer portfolio problems can be hard, and annealing or QAOA may provide approximate search heuristics. Yet classical portfolio methods are highly developed. A

persuasive claim requires hard instances, strong baselines, realistic constraints, and end-to-end cost accounting.

Risk measurement and capital allocation sit between simulation and optimization. VaR and CVaR estimation can use amplitude estimation, while portfolio allocation under risk constraints can use QUBO or other discrete formulations. A future quantum workflow could combine quantum Monte Carlo for risk evaluation with quantum or hybrid optimization for balance-sheet decisions. The catch is architectural: outputs from one quantum subroutine would have to become inputs to another quantum or classical decision system, multiplying validation and error-accounting problems.

Machine learning in finance is less settled. Quantum kernels and variational circuits have been studied for classification, anomaly detection, and prediction. Finance has many high-dimensional prediction problems, but the bar for usefulness is high. Predictive gains may be unstable, data may be non-stationary, and regulation often requires explainability. Data loading can dominate any theoretical gain. The more credible near-term use may be quantum-inspired classical methods or hybrid feature maps rather than direct deployment on quantum hardware.

Finance offers a useful precedent rather than a template. It progressed by identifying standard computational bottlenecks, mapping them to quantum primitives, and testing simplified but recognizable versions. Macro, public finance, industrial organization, and urban economics can borrow that discipline, but not its metrics wholesale. In finance, a risk estimate or portfolio objective may be relatively compact. In other fields, researchers may need welfare comparisons, policy functions, equilibrium diagnostics, and distributional consequences.

7 Macroeconomics, dynamic programming, banking, and payments

Outside finance narrowly defined, the clearest economics applications involve macroeconomic simulation, banking stress testing, payments systems, and dynamic programming. These areas matter because they are core computational-economics problems rather than generic optimization benchmarks.

The most direct quantum Monte Carlo application in economics is Skavysh et al. (2023). The paper studies whether quantum Monte Carlo can improve the runtime of economic applications, with examples in bank stress testing with credit shocks and fire sales, a neoclassical investment model solved with deep learning, and a realistic macro model solved with deep neural networks. Its contribution is partly pedagogical, because it introduces quantum Monte Carlo to economists, and partly methodological, because it shows how recognizable economic workflows can be encoded for quantum computation.

Stress testing is a natural application. Regulators and banks evaluate losses under correlated defaults, market shocks, liquidity stress, and fire sales. Classical stress testing can be computationally intensive, especially when institutions run many scenarios or when models include nonlinear feedbacks. Quantum Monte Carlo could reduce sampling costs for expected losses or tail quantities. But stress testing is also a governance process.

It requires model validation, supervisory judgment, transparency, and communication. A faster sampling engine is useful only if it fits within that framework.

Dynamic programming is even more central to computational economics. Many models involve agents or planners choosing actions over time under uncertainty. The computational burden grows with state dimensionality, action dimensionality, heterogeneity, and aggregate shocks. Numerical economics has addressed this burden with discretization, interpolation, perturbation, projection, simulation, and approximation (Judd, 1998). Deep learning has recently become important for high-dimensional dynamic models (Maliar et al., 2021). Quantum computing enters this landscape as another possible approximation technology, not as a replacement for economic structure.

The papers by Fernandez-Villaverde and Hull (2023, 2026) are notable because they explicitly target dynamic programming in economics. They introduce algorithms for solving dynamic programming problems on a quantum annealer and apply them to economic examples such as a real-business-cycle model. The result does not imply that quantum annealers can solve large heterogeneous-agent economies. It shows that Bellman problems can be reformulated in a way that exposes discrete optimization structure to annealing hardware. This opens research questions that are squarely in computational economics: which discretizations produce useful QUBO structures, how Bellman residuals should be penalized, how approximation errors compare with projection or deep learning, and whether annealing can produce economically interpretable policy functions.

Payments systems provide an operational application rather than a model-solution application. High-value payment systems settle large payments among financial institutions. Gross settlement reduces credit risk but requires liquidity. Liquidity-saving mechanisms reorder or net payments subject to delay and risk constraints. This is a classic economic-engineering problem: the objective is clear, the decisions are discrete, and the institutional consequences are measurable. McMahon et al. (2024) use a hybrid quantum annealing solver to reorder payments and seek liquidity savings without excessive delay. The paper is important because it uses real institutional data and reports economically interpretable outcomes.

A related central-bank literature considers digital payments in firm networks. Priazhkina et al. (2024) study digital payment adoption and a quantum algorithm in firm networks. This connects quantum optimization with network economics and payment adoption rather than only with computational acceleration. It also suggests that the economic impact of quantum computing may appear in infrastructures and platforms before it appears in standard academic model solution. Central banks, clearinghouses, and exchanges manage repeated optimization problems under strict operational constraints. If quantum or hybrid methods find better solutions under time limits, the welfare implications could be real even without a clean exponential speedup theorem.

Compared with finance, the macroeconomic and central-bank literature is smaller. Its strength is that the best papers are tied to recognizable economic institutions or model classes. Stress tests, Bellman equations, and payment systems are not arbitrary benchmarks. They are areas where economists already spend computational effort. The next step is scale, reproducibility, and comparison with strong classical numerical

methods. Papers should report economic metrics such as welfare error, Euler-equation error, policy-function accuracy, liquidity savings, or capital-model runtime, not only circuit depth, qubit count, or QUBO energy.

8 Networked public-sector and infrastructure applications

Transportation and infrastructure applications are relevant to *Computational Economics* insofar as they are networked economic optimization problems. Traffic control, vehicle routing, fleet assignment, transit scheduling, evacuation routing, and transport-network design involve discrete choices, congestion externalities, stochastic demand, and public objectives. They also illustrate a broader lesson for economic applications: quantum optimization is most plausible when it is embedded in a decomposed, hybrid workflow rather than used as a one-shot solver for a full system.

Transportation optimization has a growing quantum-annealing literature. Mohammed et al. (2025) review quantum-annealing-related technologies in transportation optimization and categorize studies into traffic-signal control, traffic-flow optimization, and routing problems. Salloum et al. (2025) study a mini-scale traffic-flow optimization approach using iterative QUBO formulations and quantum processing units. Singh et al. (2021) formulate real-time traffic control as QUBO. Dixit and Niu (2023) address transport-network design, a planning problem involving infrastructure or network decisions. The review literature also emphasizes that stable and reproducible quantum advantage has not yet been demonstrated for real transport systems and that hybrid frameworks are more realistic than standalone quantum deployment (Mohammed et al., 2025).

Traffic-signal control is an intuitive QUBO target. At each intersection and time interval, signal phases must be chosen subject to safety constraints. The objective may include queue lengths, delays, stops, emissions, or transit priority. Pairwise penalties can encode incompatible phases or coordination among nearby intersections. A city-scale signal system is too large for current quantum hardware if naively encoded, but local or corridor-level subproblems can be decomposed. The relevant economic question is not whether a quantum annealer can optimize an entire city in one step. It is whether a hybrid solver can improve local control decisions within operational time limits relative to classical adaptive-control systems.

Vehicle routing and fleet management are also natural targets. Urban applications include delivery, waste collection, paratransit, school buses, emergency response, and shared mobility. QUBO formulations encode route choices, visitation constraints, capacity constraints, and travel costs. Potential economic value is high because routing affects congestion, emissions, labor costs, service quality, and equity of access. But classical routing solvers are extremely strong. Quantum claims must therefore compare against state-of-the-art heuristics, commercial solvers, decomposition methods, reinforcement learning, and realistic constraints such as time windows, driver rules, stochastic travel times, and service priorities.

Transport-network design is especially relevant because it concerns strategic planning rather than only operations. Planners may choose which links to expand, which

lanes to dedicate to transit, where to add infrastructure, or how to price roads. Classic network design is often bilevel: the planner chooses interventions and travelers respond through route choice, producing equilibrium congestion. Dixit and Niu (2023) formulate an upper-level problem as QUBO and solve it with a D-Wave annealer. The approach is promising because infrastructure decisions are discrete and interactions are nonlinear. It also reveals a gap between optimization and policy. Real network design involves induced demand, land use, environmental review, distributional effects, budget cycles, and political constraints. A quantum optimizer can at most help search or evaluate a subset of alternatives.

Transportation matters here because it exposes the full set of evaluation issues that also arise in other economic networks: problem formulation, objective choice, decomposition, feasibility repair, comparison with strong classical baselines, and institutional interpretation. A technically elegant QUBO that minimizes aggregate travel time may worsen distributional outcomes if it ignores transit-dependent populations or residential street burdens. Similarly, a payment-reordering algorithm may save liquidity at the cost of unacceptable delays for certain participants. Quantum optimization does not eliminate normative choices; it makes the objective function and constraint transformation more consequential.

9 Why there is no Shor-like algorithm for economics

A recurring question is whether economics has an equivalent of Shor's algorithm. The answer is no, and the reason is instructive. Shor's algorithm exploits the algebraic structure of integer factorization and period finding. The problem is mathematically precise, widely used in cryptography, and reducible to a structure that the quantum Fourier transform can exploit. Economics does not have a single canonical computational problem with comparable universality and structure. Economic computation includes equilibrium computation, dynamic programming, estimation, simulation, prediction, assignment, pricing, and policy evaluation. These tasks differ in data, objectives, constraints, and output requirements.

Many hard economic problems are hard because they combine computation with modeling. A macroeconomic model may be difficult to solve because of high-dimensional states, but it may also be difficult to validate because expectations, frictions, and distributions are misspecified. A transport policy may involve a hard network subproblem, but it also requires demand forecasts, route-choice behavior, values of time, emissions, equity weights, and legal constraints. Quantum computing can address computational bottlenecks; it cannot remove economic uncertainty or institutional disagreement.

The absence of a Shor-like algorithm does not mean the field is empty. Classical operations research improved economics and management through a portfolio of methods: simplex, interior-point methods, branch-and-bound, dynamic programming, decomposition, simulated annealing, tabu search, and stochastic gradient methods. Quantum computing may follow a similar path. A better question is which recurring economic subroutines can be accelerated, approximated, or improved by quantum primitives.

This reframing avoids two errors. The first error is dismissing quantum economics because no exponential speedup has been shown for general equilibrium. That standard is too demanding and misdirected. A quadratic speedup in Monte Carlo could be valuable for risk analytics if the task is repeated frequently and the stakes are high. A modest improvement in payment reordering could save liquidity in a high-value payment system. A hybrid method that finds better feasible solutions under time limits could matter for infrastructure operations. The second error is accepting any quantum mapping as economically important. A QUBO formulation of a toy allocation problem is not enough. It must solve a bottleneck, scale plausibly, and improve decisions relative to strong baselines.

The distinction between asymptotic and practical advantage is central. Shor's algorithm has an asymptotic exponential speedup, but practical cryptographic impact still requires large fault-tolerant machines. Many economic quantum algorithms have weaker theoretical speedups or heuristic claims, so the path to practical advantage is more contingent. Economic value can nevertheless arise before formal asymptotic dominance if a hybrid method improves a constrained workflow. Bova et al. (2023) are useful here because they shift attention from abstract advantage to economic advantage, where costs, competition, and deployment context matter.

10 Evidence of maturity and remaining limitations

Four developments justify taking this literature seriously. Skavysh et al. (2023) connect quantum Monte Carlo to stress testing and macroeconomic deep-learning workflows. McMahon et al. (2024) use real payments-system data and report liquidity-related outcomes. Fernandez-Villaverde and Hull (2026) bring quantum annealing into dynamic programming, a core computational-economics problem. In parallel, finance and transportation reviews organize broader application domains (Herman et al., 2023; Bunescu and Vartei, 2024; Mohammed et al., 2025; Zhou, 2026). This is enough for a critical review, although not enough for broad deployment claims.

The first limitation is hardware. Near-term gate-model devices are noisy and have limited depth. Quantum annealers have more qubits but implement specialized optimization models with connectivity and precision constraints. Fault-tolerant quantum computers would make some algorithms more credible, especially amplitude estimation and linear systems, but they are not yet available at economically relevant scale. Current demonstrations are therefore valuable for formulation, resource estimation, and learning, not for broad replacement of classical workflows.

The second limitation is data loading and state preparation. Many speedups assume efficient access to data or oracle representations. Economic data are classical, noisy, high-dimensional, and often proprietary. Loading a covariance matrix, scenario distribution, transition kernel, or panel dataset may be expensive. If state preparation scales poorly, it can erase the speedup. Quantum papers should therefore report the cost of preparing inputs and not only the cost of the quantum subroutine.

The third limitation is output. Quantum states are not directly readable in full. Measurement produces samples or specific statistics. That is appropriate for expectation estimation and candidate-solution sampling, but less appropriate when economists

require complete policy functions, coefficient vectors, standard errors, or rich diagnostics. A quantum linear-system algorithm may be asymptotically fast but economically unhelpful if the user needs every coefficient explicitly.

The fourth limitation is benchmarking. Classical algorithms are highly optimized. In QUBO and integer optimization, fair baselines include mixed-integer solvers, simulated annealing, tabu search, local search, decomposition, and problem-specific heuristics. In Monte Carlo, fair baselines include variance reduction, quasi-Monte Carlo, multilevel Monte Carlo, and parallel computation. Quantum papers that compare against brute force or plain Monte Carlo alone are not persuasive for deployment. Benchmarking should measure solution quality, runtime, cloud or energy cost, robustness, scaling, and economic objective value.

The fifth limitation is formulation dependence. The same economic problem can produce many QUBO formulations. Different penalty weights, auxiliary variables, scaling choices, and embeddings can change performance dramatically. This creates a reproducibility challenge. A paper should publish the original model, the QUBO transformation, penalty calibration rules, solver parameters, code, and data when possible.

The sixth limitation is economic validity. Optimizing a simplified objective can produce misleading economic conclusions. A portfolio QUBO that ignores transaction costs and regulatory constraints is not a realistic asset-management tool. A stress-test algorithm that speeds up simulation but uses a weak loss model may produce false precision. A transport QUBO that minimizes delay but ignores safety or distributional burdens is not a policy solution. Computational economists should judge quantum methods by the full model and decision context, not only by the computational kernel.

These limitations suggest the maturity ladder in Table 2. Most work in computational economics is between levels 2 and 4. Payment-system optimization is comparatively mature because it uses real data and reports economic outcomes. Dynamic programming and stress testing are methodologically important but earlier in operational readiness. Transportation has many demonstrations, but stable real-system advantage remains unproven.

Table 2: Evidence ladder for quantum computational economics

Level	Evidence standard	Interpretation
1. Conceptual mapping	Economic problem is mapped to a plausible quantum primitive	Useful for agenda setting, not evidence of performance
2. Simulator demonstration	Small instances are solved with quantum circuits or annealing simulators	Shows implementability and reveals encoding costs
3. Hardware proof of concept	Small instances run on quantum hardware or annealers	Shows hardware interaction; does not imply advantage
4. Hybrid benchmark	Realistic instances are tested with strong classical and quantum-inspired baselines	Beginning of persuasive computational-economics evidence

Level	Evidence standard	Interpretation
5. Operational economic value	Repeated real-data workflow improves economic outcome under audited costs and constraints	Evidence of practical economic advantage

11 Research agenda for computational economics

A productive research agenda should be built around economic bottlenecks rather than hardware publicity. Computational economists can make distinctive contributions because they understand the model structures, approximation errors, and economic outputs that matter.

First, build benchmark libraries of economically meaningful QUBO and simulation problems. Existing QUBO benchmarks often use random Ising instances or generic graph problems. Economics needs benchmarks derived from portfolio constraints, payment queues, school assignment, procurement, project selection, dynamic programming grids, supply chains, and transport-network design. Each benchmark should include the original economic formulation, data, constraints, welfare or cost objective, QUBO transformation, best-known classical baselines, and economically interpretable metrics.

Second, study the economics of penalty design. In QUBO, constraints become penalties. Penalty weights encode trade-offs between feasibility and objective value. In policy and institutional settings, these trade-offs may be normative: how much payment delay is acceptable for liquidity savings, how much allocation efficiency justifies distributional imbalance, or how much travel-time reduction justifies shifting burdens across neighborhoods. Computational economists should treat penalty design as part of welfare analysis and institutional design rather than as a mere tuning problem.

Third, integrate quantum Monte Carlo with established economic simulation methods. In macroeconomics, finance, and public economics, simulation is often combined with variance reduction, quasi-Monte Carlo, importance sampling, surrogate modeling, and parallel computing. Quantum Monte Carlo should be benchmarked against these methods, not only against plain Monte Carlo. Researchers should identify models where sampling error is the dominant bottleneck and state preparation is feasible.

Fourth, develop error metrics that economists recognize. Quantum papers often report circuit depth, qubit count, approximation ratio, or QUBO energy. Economists also need welfare error, Euler-equation error, policy-function distance, moment error, capital error, liquidity savings, regret, distributional effects, and sensitivity to calibration. Reporting both technical and economic metrics would improve communication across disciplines.

Fifth, study dynamic programming more deeply. The early annealing work on Bellman problems raises open questions. Can monotonicity, concavity, shape restrictions, and envelope conditions be encoded efficiently? Can continuous states be handled through adaptive discretization? Can annealing help with discrete choices inside otherwise continuous models? How does performance compare with endogenous-grid methods, sparse grids, projection, reinforcement learning, and deep learning? Dynamic programming is a natural area for economists to lead because domain structure matters.

Sixth, develop end-to-end resource accounting. Quantum speedup claims should state the input-output model, data-loading assumptions, oracle construction, number of circuit evaluations or annealing calls, embedding overhead, classical preprocessing, postprocessing, and readout cost. For computational economics, the same accounting should include economic model-construction and validation costs.

Seventh, analyze market structure and unequal access to quantum advantage. If quantum computing becomes valuable for finance, logistics, or platforms, access may be concentrated among large firms and governments. Bova et al. (2023) begin this conversation by analyzing quantum economic advantage. Computational economists can extend it to financial stability, industrial concentration, labor markets, platform competition, and public-sector procurement.

Finally, adopt staged claims. Papers should distinguish among formulation, simulator demonstration, hardware proof of concept, hybrid benchmark, and operational economic advantage. A formulation paper is valuable if it maps an economic problem cleanly to a quantum primitive. A hardware demonstration is valuable if it reveals encoding and scaling behavior. An operational advantage claim requires real data, strong baselines, audited costs, and economically meaningful improvement. This vocabulary would reduce hype while recognizing useful early work.

12 Conclusion

Quantum computing in computational economics is doing something, but not what popular analogies suggest. There is no Shor-like algorithm that solves a central problem of economics. The field is instead developing through targeted mappings between economic computational bottlenecks and quantum primitives. Quantum amplitude estimation is relevant to Monte Carlo simulation, risk, and stress testing. QUBO, Ising formulations, quantum annealing, and QAOA are relevant to discrete optimization in portfolios, payments, dynamic programming, routing, and network design. Quantum linear algebra and machine learning are potentially relevant to econometrics and high-dimensional modeling, but their economic use cases are less mature.

The strongest case for the field comes from specific papers rather than from generic claims about quantum speedup. Skavysh et al. (2023) show how quantum Monte Carlo enters stress testing and macroeconomic computation. McMahon et al. (2024) test payment reordering on real high-value payments data. Fernandez-Villaverde and Hull (2026) translate dynamic programming into an annealing framework. Together with finance reviews and transport-network demonstrations, these papers justify treating quantum computational economics as a reviewable field.

The same papers also show why claims should stay narrow. Hardware limitations, state preparation, output constraints, embedding overhead, classical baselines, and economic validity remain major barriers. QUBO formulations are powerful but fragile. Monte Carlo speedups are theoretically attractive but depend on efficient model encoding. Networked public-sector applications need stronger integration of economic objectives, institutional constraints, and classical comparisons. Economists should be wary of demonstrations that optimize simplified objectives without showing that the resulting decision is economically meaningful.

A reasonable standard follows from that evidence. Economists should identify bottlenecks where computation genuinely limits analysis. Quantum researchers should provide transparent formulations, honest resource estimates, and strong comparisons. Institutions should demand auditability before deployment. If these standards are met, quantum algorithms may become a selective part of the computational-economics toolkit. Selective improvements can still matter when they affect high-value, repeated, and computationally difficult decisions.

References

Bova, F., Goldfarb, A., & Melko, R. G. (2023). Quantum economic advantage. *Management Science, 69*(2), 1116–1126. https://doi.org/10.1287/mnsc.2022.4578

Brassard, G., Hoyer, P., Mosca, M., & Tapp, A. (2002). Quantum amplitude amplification and estimation. In *Quantum Computation and Information* (Contemporary Mathematics, Vol. 305, pp. 53–74). American Mathematical Society. https://doi.org/10.1090/conm/305/05215

Bunescu, L., & Vartei, A. M. (2024). Modern finance through quantum computing: A systematic literature review. *PLOS ONE, 19*(7), e0304317. https://doi.org/10.1371/journal.pone.0304317

Dixit, V. V., & Niu, C. (2023). Quantum computing for transport network design problems. *Scientific Reports, 13*, 12267. https://doi.org/10.1038/s41598-023-38787-2

Egger, D. J., Gambella, C., Marecek, J., McFaddin, S., Mevissen, M., Raymond, R., Simonetto, A., Woerner, S., & Yndurain, E. (2020). Quantum computing for finance: State-of-the-art and future prospects. *IEEE Transactions on Quantum Engineering, 1*, 1–24. https://doi.org/10.1109/TQE.2020.3030314

Fernandez-Villaverde, J., & Hull, I. J. (2023). Dynamic programming in economics on a quantum annealer. NBER Working Paper No. 31326. https://doi.org/10.3386/w31326

Fernandez-Villaverde, J., & Hull, I. J. (2026). Dynamic programming in economics on a quantum annealer. *Quantitative Economics, 17*(1), 1–37. https://doi.org/10.3982/QE2555

Glover, F., Kochenberger, G., & Du, Y. (2019). Quantum bridge analytics I: A tutorial on formulating and using QUBO models. *4OR, 17*(4), 335–371. https://doi.org/10.1007/s10288-019-00424-y

Grinko, D., Gacon, J., Zoufal, C., & Woerner, S. (2021). Iterative quantum amplitude estimation. *npj Quantum Information, 7*, 52. https://doi.org/10.1038/s41534-021-00379-1

Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (pp. 212–219). https://doi.org/10.1145/237814.237866

Hadfield, S., Wang, Z., O'Gorman, B., Rieffel, E. G., Venturelli, D., & Biswas, R. (2019). From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms, 12*(2), 34. https://doi.org/10.3390/a12020034

Harrow, A. W., Hassidim, A., & Lloyd, S. (2009). Quantum algorithm for linear systems of equations. *Physical Review Letters, 103*, 150502. https://doi.org/10.1103/PhysRevLett.103.150502

Herman, D., Googin, C., Liu, X., Sun, Y., Galda, A., Safro, I., Pistoia, M., & Alexeev, Y. (2023). Quantum computing for finance. *Nature Reviews Physics, 5*(8), 450–465. https://doi.org/10.1038/s42254-023-00603-1

Hull, I., Sattath, O., Diamanti, E., & Wendin, G. (2024). *Quantum Technology for Economists*. Springer. https://doi.org/10.1007/978-3-031-50780-9

Johnson, M. W., Amin, M. H. S., Gildert, S., Lanting, T., Hamze, F., Dickson, N., Harris, R., Berkley, A. J., Johansson, J., Bunyk, P., Chapple, E. M., Enderud, C., Hilton, J. P., Karimi, K., Ladizinsky, E., Ladizinsky, N., Oh, T., Perminov, I., Rich, C., Thom, M. C., Tolkacheva, E., Truncik, C. J. S., Uchaikin, S., Wang, J., Wilson, B., & Rose, G. (2011). Quantum annealing with manufactured spins. *Nature, 473*(7346), 194–198. https://doi.org/10.1038/nature10012

Judd, K. L. (1998). *Numerical Methods in Economics*. MIT Press.

Kochenberger, G., Hao, J.-K., Glover, F., Lewis, M., Lu, Z., Wang, H., & Wang, Y. (2014). The unconstrained binary quadratic programming problem: A survey. *Journal of Combinatorial Optimization, 28*, 58–81. https://doi.org/10.1007/s10878-014-9734-0

Lucas, A. (2014). Ising formulations of many NP problems. *Frontiers in Physics, 2*, 5. https://doi.org/10.3389/fphy.2014.00005

Maliar, L., Maliar, S., & Winant, P. (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics, 122*, 76–101. https://doi.org/10.1016/j.jmoneco.2021.07.004

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

Matsakos, T., & Nield, S. (2024). Quantum Monte Carlo simulations for financial risk analytics: Scenario generation for equity, rate, and credit risk factors. *Quantum, 8*, 1306. https://doi.org/10.22331/q-2024-04-04-1306

McMahon, C., McGillivray, D., Desai, A., Rivadeneyra, F., Lam, J.-P., Lo, T., Marsden, D., & Skavysh, V. (2024). Improving the efficiency of payments systems using quantum computing. *Management Science, 70*(10), 7325–7341. https://doi.org/10.1287/mnsc.2023.00314

Mohammed, M. Q., Meess, H., & Otte, M. (2025). Review of the application of quantum annealing-related technologies in transportation optimization. *Quantum Information Processing, 24*, 296. https://doi.org/10.1007/s11128-025-04870-y

Montanaro, A. (2015). Quantum speedup of Monte Carlo methods. *Proceedings of the Royal Society A, 471*(2181), 20150301. https://doi.org/10.1098/rspa.2015.0301

Montanaro, A. (2016). Quantum algorithms: An overview. *npj Quantum Information, 2*, 15023. https://doi.org/10.1038/npjqi.2015.23

Orus, R., Mugel, S., & Lizaso, E. (2019). Quantum computing for finance: Overview and prospects. *Reviews in Physics, 4*, 100028. https://doi.org/10.1016/j.revip.2019.100028

Phillipson, F., & Bhatia, H. S. (2021). Portfolio optimisation using the D-Wave quantum annealer. In *Computational Science – ICCS 2021* (pp. 45–59). Springer. https://doi.org/10.1007/978-3-030-77980-1_4

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

Priazhkina, S., Palmer, S., Martin-Ramiro, P., Orus, R., Mugel, S., & Skavysh, V. (2024). Digital payments in firm networks: Theory of adoption and quantum algorithm. Bank of Canada Staff Working Paper 2024-17. https://doi.org/10.34989/swp-2024-17

Rebentrost, P., Gupt, B., & Bromley, T. R. (2018). Quantum computational finance: Monte Carlo pricing of financial derivatives. *Physical Review A, 98*, 022321. https://doi.org/10.1103/PhysRevA.98.022321

Salloum, H., Zhanalin, S., Al Badr, A., & Kholodov, Y. (2025). Mini-scale traffic flow optimization: An iterative QUBOs approach converting from hybrid solver to pure quantum processing unit. *Scientific Reports, 15*, 22904. https://doi.org/10.1038/s41598-025-04568-2

Shor, P. W. (1994). Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science* (pp. 124–134). https://doi.org/10.1109/SFCS.1994.365700

Singh, A., Lin, C.-Y., Huang, C.-I., & Lin, F.-P. (2021). Quantum annealing approach for the optimal real-time traffic control using QUBO. In *2021 IEEE/ACIS 22nd International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing*. https://doi.org/10.1109/SNPD51163.2021.9704985

Skavysh, V., Priazhkina, S., Guala, D., & Bromley, T. R. (2023). Quantum Monte Carlo for economics: Stress testing and macroeconomic deep learning. *Journal of Economic Dynamics and Control, 153*, 104680. https://doi.org/10.1016/j.jedc.2023.104680

Suzuki, Y., Uno, S., Raymond, R., Tanaka, T., Onodera, T., & Yamamoto, N. (2020). Amplitude estimation without phase estimation. *Quantum Information Processing, 19*, 75. https://doi.org/10.1007/s11128-019-2565-2

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

Zhou, J. (2026). Quantum finance: Exploring the implications of quantum computing on financial models. *Computational Economics, 67*(2), 1043–1072. https://doi.org/10.1007/s10614-025-10894-4

Statements and Declarations. Funding. The author declares that no funds, grants, or other support were received during the preparation of this manuscript.
Competing interests. The author has no relevant financial or non-financial interests to disclose.
Ethics approval. Not applicable. This article is a review and did not involve human participants, animals, or clinical data.
Consent to participate. Not applicable.
Consent for publication. Not applicable.

Data availability. No new data were generated or analyzed for this review. The article synthesizes cited published articles, conference papers, books, working papers, and publicly available bibliographic records.

Code availability. Not applicable. No original code was developed for this review.

Author contributions. Diego A. Diaz conceptualized the review, selected and interpreted the literature, drafted and revised the manuscript, and approved the final version.

Use of artificial intelligence tools. AI-assisted tools were used for copy editing, formatting, and consistency checks. The author selected the literature, verified the references, revised the analysis, and accepts full responsibility for the manuscript.