

<https://doi.org/10.1038/s44335-025-00050-4>

# Quantum stochastic walks for portfolio optimization: theory and implementation on financial networks

Check for updates

Yen Jui Chang<sup>1,2</sup> , Wei-Ting Wang<sup>3</sup>, Yun-Yuan Wang<sup>4</sup>, Chen-Yu Liu<sup>5</sup>, Kuan-Cheng Chen<sup>6,7</sup> & Ching-Ray Chang<sup>2,3</sup>

Classical mean-variance optimization is powerful in theory but fragile in practice, often producing highly concentrated, high-turnover portfolios. Naive equal-weight (1/N) portfolios are more robust but largely ignore cross-sectional information. We propose a quantum stochastic walk (QSW) framework that embeds assets in a weighted graph and derives portfolio weights from the stationary distribution of a hybrid quantum-classical walk. The resulting allocations behave as a “smart 1/N” portfolio: structurally close to equal-weight, but with small, data-driven tilts and a controllable level of trading. On recent S&P 500 universes, QSW portfolios match the diversification and stability of 1/N while delivering higher risk-adjusted returns than both mean-variance and naive benchmarks. A comprehensive hyper-parameter grid search shows that this behavior is structural rather than the result of fine-tuning and yields simple design rules for practitioners. A 34-year, multi-universe robustness study with rolling re-optimization further demonstrates that the QSW optimizer preserves these advantages across market regimes. Overall, the QSW framework improves risk-adjusted performance while maintaining strong diversification and moderate turnover.

Modern investment management operates in an arena of high-dimensional, tightly coupled, and rapidly evolving data. Conventional quantitative techniques—most notably the seminal mean-variance framework of Markowitz<sup>1</sup>—encapsulate the trade-off between expected return and risk through a quadratic objective that minimizes portfolio variance for a target mean return. While mathematically elegant, this quadratic form becomes fragile when confronted with today’s realities: estimation error in large covariance matrices, non-Gaussian return distributions, time-varying correlations, and a proliferation of real-world constraints such as turnover limits, liquidity frictions, and regulatory capital charges<sup>2–6</sup>.

The classical foundations of portfolio optimization are well established. The work of Markowitz<sup>1</sup> introduced the mean-variance framework, which remains the cornerstone of modern portfolio theory. In this setting, expected returns  $\mu$  and the covariance matrix  $\Sigma$  jointly determine the efficient frontier through problems such as minimum-variance optimization and maximum-Sharpe optimization<sup>7–9</sup>. Over the decades, numerous refinements have been proposed: the Black–Litterman model blends equilibrium returns with subjective views in a Bayesian framework<sup>3</sup>; robust

optimization formulations introduce uncertainty sets for  $\mu$  and  $\Sigma$ <sup>2,10,11</sup>; and alternative risk measures such as coherent risk and CVaR have been explored to better capture tail risks<sup>12,13</sup>.

Despite seven decades of refinement, however, mean-variance optimization faces severe challenges in practice. Estimation error and dimensionality are the first obstacles: the covariance matrix for  $n$  assets contains  $\frac{1}{2}n(n+1)$  parameters, and for  $n = 500$  this already exceeds  $10^5$  parameters, with condition numbers often above  $10^6$ . With finite samples, small errors in  $\mu$  and  $\Sigma$  can produce extreme and unstable portfolios<sup>4,14</sup>. Markets are also non-stationary: correlations spike toward unity during crises<sup>15</sup>, precisely when diversification is most needed, and correlation structures exhibit asymmetric behavior between up and down markets<sup>16</sup>. Empirical returns display heavy tails, skewness, and volatility clustering<sup>17</sup>, making variance alone an incomplete proxy for risk. Real-world implementation adds further complexity: cardinality constraints, transaction costs, market impact, and liquidity constraints turn the convex quadratic program into a mixed-integer nonlinear program that is NP-hard<sup>6,18</sup>.

<sup>1</sup>Master Program in Intelligent Computing and Big Data, Chung Yuan Christian University, Taoyuan City, Taiwan. <sup>2</sup>Quantum Information Center, Chung Yuan Christian University, Taoyuan City, Taiwan. <sup>3</sup>Department of Physics, National Taiwan University, Taipei, Taiwan. <sup>4</sup>NVIDIA AI Technology Center, NVIDIA Corp., Taipei, Taiwan. <sup>5</sup>Graduate Institute of Applied Physics, National Taiwan University, Taipei, Taiwan. <sup>6</sup>Department of Electrical and Electronic Engineering, Imperial College London, London, UK. <sup>7</sup>Centre for Quantum Engineering, Science and Technology (QuEST), Imperial College London, London, UK.

e-mail: [aceest@cycu.edu.tw](mailto:aceest@cycu.edu.tw)

Perhaps most critically, small perturbations in expected returns can cause dramatic rebalancing; errors in means dominate those in covariances by roughly an order of magnitude<sup>19,20</sup>, leading to concentrated, counter-intuitive portfolios that often underperform simple equal-weight (1/N) strategies out of sample<sup>14</sup>.

To mitigate these issues, a growing body of work has turned to neural networks (NNs). One line of research follows a predict-then-optimize paradigm: deep learning models are trained to forecast inputs such as expected returns or risk premia more accurately than classical econometric models<sup>21</sup>, with the forecasts then fed into a conventional optimizer. A second, more recent line uses end-to-end deep reinforcement learning (DRL) to learn portfolio policies directly: agents map market states to portfolio weights to maximize cumulative reward functions such as the Sharpe ratio or risk-budgeting objectives<sup>22–24</sup>. These approaches can, in principle, bypass some of the sensitivity of mean-variance optimization to input estimates. However, NN-based approaches introduce their own challenges. The resulting models are typically highly non-linear and opaque, making it difficult to understand or audit how they respond to changes in market conditions<sup>25</sup>. This lack of interpretability is a serious concern for risk management, regulatory review, and practitioner trust. Moreover, deep models are data-hungry and prone to overfitting noisy, non-stationary financial time series<sup>21</sup>, and DRL agents can suffer from training instability and sensitivity to reward specification. In practice, these methods can behave as powerful but opaque “black boxes”, which limits their adoption in settings where transparency and stability are as important as raw performance.

Parallel to these developments, the econophysics literature has highlighted a complementary perspective based on financial networks. Early work showed that a covariance matrix can be filtered into sparse structures such as minimum-spanning trees or planar maximally filtered graphs, whose topology reveals hidden economic sectors and contagion channels<sup>26–28</sup>. Subsequent studies demonstrated that risk propagation and systemic fragility travel along these edges much like a diffusion process<sup>29,30</sup>. From this viewpoint, the quadratic risk term  $w^\top \Sigma w$  can be interpreted as a diffusion of capital on a weighted network, where assets are vertices, covariances set edge weights, and self-loops reward each asset’s own risk-adjusted return. Classical solvers glimpse this structure only indirectly through dense linear algebra; they do not embed diversification in the dynamics of the flow.

Recent progress in quantum information science enables this network view to be dynamic and adaptive. Quantum stochastic walks (QSWs)<sup>31,32</sup> and quantum network-ranking algorithms<sup>33–36</sup> merge coherent quantum evolution with classical random walks on the same graph. In this framework, a coherent channel governed by a Hermitian Hamiltonian explores the network through superposition and interference, while a stochastic channel, implemented via a Google-type generator, ensures ergodicity and convergence. By tuning a quantum-classical mixing parameter, a QSW interpolates between fully quantum exploration and purely classical diffusion, providing an adaptive mechanism that can respond to both structured dependencies and stochastic shocks. The coherent channel can survey multiple highly covariant clusters in parallel, while the stochastic channel damps cyclic probability within tight clusters and nudges capital towards under-represented regions of the network.

Building on these insights, we develop a quantum graph-theoretic framework for portfolio optimization with four main contributions. First, we map expected returns and covariances to a weighted, directed graph in which self-loops encode asset-specific Sharpe information and inter-node edges penalize correlation, providing a network-based representation of the mean-risk trade-off. Second, we introduce a dual-channel QSW whose density matrix evolves on this graph and whose diagonal converges to a steady-state weight vector; the classical and quantum channels jointly realize a “smart 1/N” allocation, structurally close to equal-weight but with small, data-driven tilts and a controllable level of trading. Third, we provide a practical implementation of the QSW framework (including GPU acceleration for matrix operations), making it feasible to optimize and rebalance

portfolios of order 100 assets in a realistic time. Fourth, we conduct a comprehensive empirical evaluation, benchmarking QSW-based portfolios against maximum-Sharpe mean-variance solutions, naive 1/N portfolios, and the S&P 500 index. The tests span fixed parameter presets, a 625-point hyper-parameter grid search, and a 34-year multi-universe robustness study across 30 randomly sampled 100-stock S&P 500 universes. Across these experiments, QSW-based portfolios consistently retain 1/N-like diversification while improving risk-adjusted returns and maintaining moderate turnover. The grid search shows that these gains are structural, not the result of fine-tuning, and yield simple design rules that link the QSW parameters to diversification, turnover, and efficiency.

The remainder of the paper proceeds as follows. We first present the empirical results in the “Results” section, followed by a Discussion of implications. Finally, we describe the QSW formulation and experimental design in the “Methods” section.

## Results

This section presents a comprehensive empirical validation of the QSW framework. The experimental design has two phases: (1) a recent-period parameter exploration (2018–2024) to understand and tune the model, and (2) a long-horizon, multi-universe robustness study (1990–2024) to test generalizability across regimes and universes. The first phase uses a fixed “top-100-by-market-cap” S&P 500 universe and includes both a preset-based analysis and a systematic grid search. The second phase runs 30 independent trials on dynamically maintained, point-in-time S&P 500 universes to guard against selection and survivorship bias. Across all experiments, QSW-based portfolios are compared against three benchmarks: the maximum-Sharpe MPT portfolio, the naive 1/N portfolio, and the S&P 500 index.

### Experimental overview

Our experimental framework is designed to evaluate the QSW methodology under both controlled and realistic conditions.

In Phase 1, parameter exploration (2018–2024), we use daily data from 2018 to 2024 on the top 100 S&P 500 companies by market capitalization as a fixed universe. This phase consists of two experiments. Experiment 1 (preset analysis) defines six QSW strategy presets that represent different investment philosophies, from ultra-diversified to high-activity trading styles. Each preset is evaluated across five quantum-mixing values  $\omega \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$  to study the impact of the quantum-classical spectrum. Experiment 2 (comprehensive grid search) systematically explores a 625-point grid over the four QSW parameters ( $\alpha, \beta, \lambda, \omega$ ) on the same top-100 universe. This grid search identifies robust high-performing regions of the parameter space and clarifies how the four control knobs interact.

In Phase 2, robustness validation (1990–2024), we test the generalizability of the approach over a 34-year backtest with rolling 2-year training windows. Experiment 3 (multi-universe testing) runs 30 independent trials on distinct 100-stock universes. As described in the Methods section (“Data and experimental setup”), each universe is constructed point-in-time by sampling from contemporaneous S&P 500 constituents and is dynamically maintained to avoid survivorship bias. This design validates that the QSW results are not dependent on a single, favorable asset selection.

In all experiments, QSW-based portfolios are evaluated against three benchmarks: the maximum-Sharpe MPT portfolio (classical mean-variance optimizer), the naive 1/N portfolio on the same universe, and the S&P 500 index as a market-cap-weighted passive benchmark.

### Experiment 1: preset-based analysis

This experiment evaluates 30 distinct configurations, which are derived by crossing the six parameter presets in Table 1 with five quantum-classical mixes ( $\omega \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ ). To test the model’s sensitivity to the length of historical data, we evaluate every configuration twice: (i) using a 1-year (“short-memory”) training window, and (ii) using a 2-year (“long-memory”) training window. Both sets of models are then out-of-sample back-tested from 2018-01-02 to 2024-12-31 with quarterly rebalancing.

**Table 1 | QSW strategy presets used in phase 1**

Preset	$\alpha$	$\beta$	$\lambda$	Investment philosophy
Ultra-diversified	1.0	100.0	10.0	Prioritizes maximum diversification
Moderate-balanced	10.0	10.0	10.0	Balanced emphasis on risk, return, and turnover
Stability-focused	1.0	10.0	100.0	Emphasizes position stability
Balanced-active	10.0	1.0	100.0	Moderate return focus with stable positions
Sharpe-maximizer	100.0	1.0	10.0	Aggressive pursuit of high-Sharpe assets
High-activity	100.0	10.0	1.0	Maximum rebalancing frequency

A combined analysis of the cumulative wealth trajectories (Fig. 1) and the detailed performance metrics (Table 2) reveals several critical insights into the QSW model's behavior relative to the benchmarks.

First, regarding model stability, the “two starkly different regimes” visible in Fig. 1 are driven almost entirely by the instability of the MPT benchmark, not the QSW model. This highlights MPT's extreme sensitivity to estimation error, a key limitation we identified in the Introduction. With 1-Year Training, MPT produces a poor portfolio (Sharpe 0.86) with extremely high volatility (31.61%). The short, noisy training window likely results in a highly unstable covariance matrix, leading to a flawed optimization. In contrast, all QSW strategy presets (Sharpe  $\approx$  0.96–0.98) are stable and easily outperform MPT. With 2-Year Training, with a more stable window, MPT's volatility drops to 20.94% and its Sharpe ratio jumps by 58% to 1.36. The QSW and 1/N models, however, are almost completely unaffected by the change in window length, with their Sharpe (0.98) and Volatility ( $\approx$  17.2%) remaining rock-solid. This demonstrates the robustness of the QSW method; its performance is not dependent on a “lucky” or perfectly stable estimation window.

Second, the QSW model converges to a “Smart 1/N”. The most striking finding is that the QSW model, across all QSW strategy presets, produces a risk and diversification profile that is nearly identical to the naive 1/N benchmark. As seen in Table 2, the QSW's Sharpe (0.96–0.98), Volatility ( $\approx$  17%), Max Drawdown ( $\approx$  19%), and HHI ( $\approx$  0.01) are all in lockstep with the 1/N portfolio. This suggests the QSW's quantum-graph dynamics naturally find the highly diversified, stable 1/N state. The key difference is that QSW acts as a “smart 1/N”, initiating a very low (2–35%) but non-zero turnover, in contrast to MPT's hyper-active 320–480% turnover.

Third, while MPT achieves the highest “paper” Sharpe ratio (1.36) in the 2-year case, Table 2 shows this is operationally unfeasible. This performance is achieved by making an extreme, concentrated bet ( $HHI \geq 0.25$ , or  $\approx$  4–5 effective stocks) and by turning over the entire portfolio 3–5 times per year (320–480% turnover). The transaction costs from this churn would completely erase its “paper” alpha. QSW, by contrast, generates its alpha with minimal turnover and maximum diversification ( $HHI \approx 0.01$ ,  $\approx$  90–100 effective stocks), making it a practical and cost-effective strategy.

Finally, Fig. 1 shows that in both 1-year and 2-year setups, all QSW configurations (final wealth  $\approx$  3.0 $\times$ ) and the 1/N benchmark ( $\approx$  3.0 $\times$ ) consistently and significantly outperform the passive S&P 500 index (final wealth  $\approx$  2.2 $\times$ ).

Comparing the 1-year (Table 2a) and 2-year (Table 2b) results reveals the core weakness of MPT and the core strength of QSW. The MPT benchmark is fundamentally unstable. When the training window changes from 1-year to 2-years, MPT's Sharpe ratio leaps 58% (from 0.86 to 1.36) and its volatility drops from 31.6% to 20.9%. This demonstrates that MPT is extremely sensitive to the estimation error in its inputs, a key classical limitation. In stark contrast, the QSW and 1/N models are rock-solid. Their key metrics (Sharpe  $\approx$  0.98, Vol  $\approx$  17.2%) are almost completely unaffected by the change in training data. This provides strong empirical evidence that our QSW framework, like the 1/N benchmark, is structurally robust and

successfully insulated from the parameter instability that plagues classical MPT.

To understand the behavior within each preset, we analyze the effect of the quantum-classical mixing parameter,  $\omega$ , and the portfolio's concentration over the 2018–2024 backtest period. We present the results for the 1-year training window here; the corresponding analysis for the 2-year window, which confirms the same structural conclusions, is provided in the Supplementary Information. Figure 2 shows the performance metrics for all six QSW strategy presets as a function of the quantum-classical balance,  $\omega$ . The results reveal a clear and powerful trend: the  $\omega$  parameter acts primarily as a “turnover dial.” In all six QSW strategy presets, increasing  $\omega$  (making the model more “classical”) leads to a near-linear increase in portfolio turnover. For example, in the “Ultra-Diversified” preset, turnover rises from 4 to 30% as  $\omega$  goes from 0.2 to 1.0. Crucially, this increase in turnover has almost no impact on the Sharpe ratio or volatility. These metrics remain remarkably flat across the entire  $\omega$  spectrum, and always superior to MPT's 1-year performance. Most importantly,  $\omega$  does not break the model's diversification. The HHI remains at its 1/N floor ( $\approx$  0.01) in all but the most active presets, and even then, it stays an order of magnitude more diversified than MPT. This demonstrates that the QSW model's core benefits (high Sharpe, low volatility, and extreme diversification) are structurally inherent and robust, while the  $\omega$  parameter provides a simple, interpretable knob to control the “cost” (i.e., turnover) of the strategy.


A core structural difference between QSW and MPT is revealed by analyzing their portfolio concentration. We compare the models over the 2018–2024 backtest using four standard metrics, as shown in Fig. 3. The QSW model, across all presets, behaves as a “smart 1/N”. Its concentration is tunable and directly linked to the quantum-classical mix,  $\omega$ . As seen in Fig. 3, when  $\omega$  is low (more quantum), the model's HHI is locked at around the theoretical floor of 0.01. As  $\omega$  increases, the model is permitted to become slightly more concentrated, and the HHI, effective number of stocks, max single stock weight, and top 5 holdings all show a small, controlled deviation from the 1/N state. However, even at its most “classical” ( $\omega = 1.0$ ), the QSW model remains exceptionally well-diversified. The MPT benchmark (red line) is defined by its extreme and erratic concentration. Its HHI is dangerously unstable. While its average HHI is  $\approx$  0.28, the Top 5 Holdings chart shows that MPT's typical state is to hold  $\approx$  87–89% of the entire portfolio in just five names. This confirms MPT is a structurally unstable and highly concentrated strategy, while QSW is robustly diversified by design.

Portfolio annual turnover directly translates into implementation costs that can significantly erode theoretical performance gains. These costs compound with trading frequency, making turnover control a critical factor in real-world portfolio management. Institutional studies such as<sup>37</sup> place the all-in round-trip cost of trading cash equities at: large-cap stocks (5–15 bp), mid-/small-cap (15–30 bp), and market impact (a further 10–50 bp). In this work, we apply a conservative 20 bp all-in round-trip cost for large-cap equities (10 bp commission & spread + 10 bp market impact). Table 3 converts realized turnover into ex-ante implementation drag using this rate. Even under the more forgiving 2-year training calibration, MPT forfeits  $\sim$  64 bp of alpha to dealing costs versus  $\sim$  7.4 bp for the QSW preset mean.

*Is a 64 bp vs. 7 bp gap economically meaningful?* Yes. Compare each strategy's “paper” alpha (CAGR in excess of the S&P 500) with, and without, implementation drag (Table 4).

$$\text{Retention}_{\text{MPT}} = \frac{1155}{1220} \approx 95\%, \text{Retention}_{\text{QSW}} = \frac{592}{600} \approx 99\%.$$

Even after deducting realistic dealing costs, QSW preserves virtually all of its back-tested edge, whereas MPT forfeits about 5% of its “paper” alpha. Because the classical maximum-Sharpe solution trades eight to thirteen times more than QSW, its cost drag—while “only” 64 bp—is still an order of magnitude larger in proportional terms. QSW's subdued trading footprint additionally mitigates price pressure, timing risk, and information leakage, making its low-turnover profile a structural implementation advantage. In



**Fig. 1 | Cumulative value of \$1 invested in 2018.** Top: 1-year training. Bottom: 2-year training. Red dashed line = MPT benchmark; dotted grey = S&P 500. Solid QSW lines denote configurations that *beat* MPT on Sharpe ratio, dotted lines underperform.

Experiment 3, we show that, even when the QSW parameters are re-optimized every quarter over 1990–2024, the dynamic QSW optimizer still trades substantially less than the long-memory MPT benchmark while preserving its risk-adjusted edge.

In the previous analysis, we demonstrated that MPT is structurally unstable at the macro level: changing the training window from 1 to 2 years led to dramatic changes in performance and risk profile (a 58% jump in Sharpe). This section's Monte Carlo test confirms this fragility on a micro level. A fundamental flaw of classical MPT is its extreme

sensitivity to input parameters. A small, statistically insignificant change in the expected return vector  $\mu$  can cause the optimizer to produce a radically different portfolio, leading to massive, unforced turnover. To test the robustness of our QSW model to this, we conduct a Monte Carlo sensitivity analysis. We take a single period's data and apply a 20% random “shock” to the returns of a single, randomly chosen asset (i.e.,  $r_i \times \text{random}(0.8, 1.2)$ ). We then re-run the optimization for both QSW and MPT and measure the total absolute change in portfolio weights (i.e.,  $\sum_i |w_{i,\text{new}} - w_{i,\text{old}}|$ ). This process is repeated 1000 times for

**Table 2 | Headline metrics for the six QSW strategy presets (best mix  $\omega = 0.2$ )**

(a) 1-year training							
Preset	Sharpe	Vol. [%]	MDD [%]	Turn. [%]	Eff.	HHI	Eff.# stk
Ultra-diversity	0.9846	17.17	-19.27	4.00	19.64	0.0100	99.96
Moderate-balanced	0.9800	17.10	-19.15	31.92	2.95	0.0102	98.81
Stability-focused	0.9837	17.05	-19.19	35.28	2.65	0.0102	97.37
Balanced-active	0.9803	17.03	-19.12	50.49	1.87	0.0105	94.84
Sharpe-maximizer	0.9615	16.80	-19.56	79.71	1.19	0.0107	92.98
High-activity	0.9615	16.80	-19.56	79.71	1.19	0.0107	92.98
1/N benchmark	0.98	17.17	-19.27	0	$\infty$	0.01	100
MPT benchmark	0.86	31.61	-16.03	480.85 (-)	0.18	0.268	3.7
(b) 2-year training							
Preset	Sharpe	Vol. [%]	MDD [%]	Turn. [%]	Eff.	HHI	Eff.# stk
Ultra-diversity	0.98	17.18	-19.20	2.11	31.68	0.0100	99.98
Moderate-balanced	0.98	17.15	-18.59	17.66	5.25	0.0101	98.80
Stability-focused	0.98	17.10	-18.13	26.06	3.64	0.0102	97.59
Balanced-active	0.98	17.10	-17.75	34.39	2.77	0.0104	95.95
Sharpe-maximizer	0.96	16.99	-18.25	71.24	1.33	0.0109	91.99
High-activity	0.96	16.99	-18.25	71.24	1.33	0.0109	92.00
1/N benchmark	0.98	17.17	-19.27	0	$\infty$	0.01	100
MPT benchmark	1.36	20.94	-17.5	320.08	0.42	0.254	4.65

Vol. annualized volatility, MDD max drawdown, Turn. annual turnover, Eff. sharpe/turnover.

each of the six QSW strategy presets. Figure 4 plots the 1000-run distribution of this weight change, while Table 5 quantifies the results.

The results are clear and highly suggestive. The MPT benchmark (red distributions in Fig. 4) exhibits extreme fragility. As quantified in Table 5, while MPT's *mean* weight change (1.72%) is moderate, its standard deviation (7.82%) is massive—over 2 $\times$  to 100 $\times$  larger than any QSW preset. This highlights how MPT amplifies minor input errors into significant portfolio changes. The QSW model, by contrast, is both structurally robust and tunable. Its stability is directly governed by its parameters, which function as an accelerator and a brake: presets with high  $\alpha$  ("Sharpe-Maximizer", "High-Activity") show the highest mean weight change (3.61%), as they are designed to chase returns, while presets with a high diversification penalty ( $\beta$ ) or holding coefficient ( $\lambda$ ) (such as "Ultra-Diversified") show exceptionally low mean weight changes (0.07–1.90%). This test provides clear empirical evidence that the QSW framework not only overcomes the fragility of MPT but also provides an interpretable, tunable mechanism to control the model's sensitivity to input noise. The next experiment broadens the evidence: a full 625-point grid search pinpoints globally robust parameter regions, and a 30-universe robustness study confirms that these findings generalize beyond this fixed top-100 universe.

## Experiment 2: comprehensive grid search


Having established in Experiment 1 that all six QSW strategy presets behave robustly and outperform MPT's 1-year model, we now search the full four-dimensional parameter space  $\theta = (\alpha, \beta, \lambda, \omega)$  to answer three key questions: (i) Does an even better configuration exist than the presets? (ii) How sensitive is performance to each hyper-parameter? (iii) Can we extract simple design rules for practitioners?

This experiment performs a systematic grid search over the 625 parameter combinations defined by:  $\alpha, \beta, \lambda \in \{0.1, 5, 50, 100, 500\}$  and  $\omega \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ , yielding 625 unique combinations. All 625 combinations are evaluated using the Parameter Exploration Setup described in the "Methods" (2018–2024 backtest, top-100 universe, and quarterly rebalancing). The entire 625-point grid search is performed twice: once using the 1-year ("short-memory") rolling training window, and a second time using the 2-year ("long-memory") window. This results in a total of  $625 \times 2 = 1250$

full, independent back-tests. For each of the 1250 back-tests, we record the full set of performance metrics (Sharpe ratio, volatility, MDD, turnover, efficiency  $\mathcal{E}$ , and HHI).

The grid search results provide a powerful and conclusive answer to our three research questions by revealing the QSW's "smart" hybrid nature. The model is not a sequential "two-step" process, but a simultaneous dual-channel optimizer that intelligently balances its classical and quantum components to achieve a robust, efficient portfolio. As noted, Tables 6 and 7 provide a strong and conclusive answer to our first research question. First, QSW is structurally robust. The descriptive statistics in Table 6 show that the QSW framework is inherently stable. In the 1-year window, the worst-performing QSW configuration (Sharpe 0.85) is still comparable to the MPT benchmark (0.86). Furthermore, the median QSW Sharpe ( $\approx 0.96$ ) is vastly superior, proving that the model delivers strong performance across the majority of its parameter space, in stark contrast to the brittle MPT solution. Second, QSW converges to a smart 1/N-like portfolio. The most critical finding comes from Table 7. In both the 1-year and 2-year windows, the Top-10 best-performing QSW configurations are effectively identical: they all have minimal return-chasing ( $\alpha = 0.1$ ) and maximal diversification penalty ( $\beta = 500$ ). The QSW model, when allowed to search the entire parameter space, consistently identifies an almost infinitesimally traded, 1/N-like portfolio as the most robust solution in this universe. Its performance (Sharpe 0.989, HHI 0.010, Turnover  $< 3\%$ ) is a near-perfect match for the 1/N benchmark. Third, QSW correctly identifies MPT's "paper" victory. The difference between the 1-year and 2-year results perfectly frames the QSW's value. In the 1-Year case, MPT (0.86) fails due to estimation error. QSW (0.989) discovers that 1/N (0.980) is the superior strategy and outperforms both benchmarks. In the 2-Year case, MPT (1.36) stabilizes and finds a high-turnover, high-concentration "paper" victory. QSW, faced with the same data, still identifies the 1/N-like state as its optimal solution. It correctly avoids MPT's operationally unfeasible (320% turnover) solution and selects the most practical and efficient strategy.

Our analysis of the tables has answered our first question: the optimal strategy is the "smart 1/N" state. To visualize the global impact of each hyper-parameter, we plot the full 625-run parameter space. The 2D heatmaps (Fig. 5) and the full correlation matrix (Fig. 6) answer our second and




**Fig. 2 | Impact of the quantum-classical mixing parameter ( $\omega$ ) on key metrics for all six QSW strategy presets (1-year training, 2018–2024 backtest).**  $\omega$  primarily functions as a “turnover dial” without degrading the portfolio’s Sharpe ratio or diversification.

third questions, showing how the model finds this state by using its dual-channel design. The 2-year data, which are nearly identical, are provided in the Supplementary Information and confirm that these conclusions are robust to the training window choice.

Figure 5 makes the QSW’s structural behavior immediately visible. Most strikingly, the concentration panel in Fig. 5b shows that, across the entire 625-run grid, the HHI remains confined to an extremely narrow band: from a floor of 0.010 (dark blue) to a maximum of 0.014 (bright green). In other words, even the “worst-case” combination of  $(\alpha, \beta, \lambda, \omega)$  never produces anything close to a concentrated portfolio. By comparison, the MPT benchmark sits at  $\text{HHI} = 0.268$ , more than  $20 \times$  higher. This gap is not a tuning artifact but a structural property: the QSW framework is incapable of generating the kind of brittle, single-bet allocations that MPT routinely produces. The Sharpe and Efficiency panels (Fig. 5a, c) complete this picture by revealing how the hyper-parameters act as behavioral knobs

rather than sources of instability. Sharpe ratios remain tightly clustered in the 0.94–0.99 range over most of the grid, indicating that performance is remarkably insensitive to moderate perturbations of  $(\alpha, \beta, \lambda, \omega)$ . Efficiency, by contrast, varies much more strongly along the  $\lambda$  and  $\omega$  axes: higher values of  $\lambda$  (the “holding” parameter) and  $\omega$  systematically reduce  $\mathcal{E}$  by lowering Sharpe and/or increasing turnover, while leaving the underlying diversification almost untouched.

These visual impressions are quantified by the full correlation matrix in Fig. 6. Within our tested range  $\omega \in [0.2, 1.0]$ ,  $\omega$  exhibits a strong negative correlation with Sharpe and a positive correlation with HHI, confirming that it primarily acts as a noise parameter: smaller  $\omega$  values lead, on average, to more diversified, higher-quality portfolios, whereas larger  $\omega$  values inject additional stochasticity that slightly raises concentration and erodes risk-adjusted returns. In contrast,  $\lambda$  shows a pronounced negative correlation with Efficiency while having only a mild effect on Sharpe, confirming its role



**Fig. 3 | Evolution of portfolio concentration (Herfindahl-Hirschman Index, HHI) for QSW-based portfolios (blue and purple) compared with the classical MPT benchmark (red) during the 2018–2024 backtest period.** The MPT portfolio exhibits structural instability, alternating between highly concentrated and diffuse

allocations, whereas all QSW configurations maintain stable, near-uniform diversification ( $HHI \approx 0.01$ ) even through periods of market stress such as the 2020 COVID-19 crash.

**Table 3 | Illustrative annual implementation drag assuming a 20 bp round-trip cost**

Strategy/preset	1-year training		2-year training	
	Turn. [%]	Cost [bp]	Turn. [%]	Cost [bp]
<i>QSW presets</i>				
Ultra-diversified	4.0	0.8	2.1	0.4
Moderate-balanced	31.9	6.4	17.7	3.5
Stability-focused	35.3	7.1	26.1	5.2
Balanced-active	50.5	10.1	34.4	6.9
Sharpe-maximizer	79.7	15.9	71.2	14.2
High-activity	79.7	15.9	71.2	14.2
QSW preset mean	46.9	9.4	37.1	7.4
1/N (naive)	0.0	0.0	0.0	0.0
MPT (max-Sharpe)	480.9	96.2	320.1	64.0

Costs are computed as Turnover  $\times$  0.20%. 1 basis point (bp) = 0.01 percentage points.

as a trade-off knob between turnover and implementability rather than a driver of raw performance. Together, the heatmaps and the correlation matrix show that QSW's smart 1/N solution is not a single fine-tuned point, but a broad, structurally robust region of the parameter space.


**Table 4 | After-cost alpha retention**

	Paper $\alpha$ [bp]	Cost drag [bp]	Net $\alpha$ [bp]
MPT (2-yr calibration)	1 220	65	1 155
QSW (6-preset mean)	600	8	592

1 basis point (bp) = 0.01 percentage points.

To visually confirm these findings and answer our second research question, we analyze the global impact of each hyper-parameter. Figure 7 plots the Pearson correlations between the four control knobs ( $\alpha, \beta, \lambda, \omega$ ) and our key metrics. We present the plots for both the 1-year “noisy” regime and the 2-year “stable” regime, as their comparison provides insight into the QSW model’s adaptive, structural behavior.

The correlation plots provide four practical design rules for practitioners. First,  $\omega$  is the main “noise” dial, with a mild impact on concentration. Among the four knobs,  $\omega$  shows the strongest and most stable pattern across both horizons. Within our tested range  $\omega \in [0.2, 1.0]$ , it is strongly negatively correlated with Sharpe (1Y:  $-0.495$ , 2Y:  $-0.556$ ), and positively correlated with HHI (1Y:  $0.225$ , 2Y:  $0.393$ ). In other words, higher  $\omega$  injects additional stochastic noise into the walk, which slightly increases concentration and systematically erodes risk-adjusted returns. Conversely, lower  $\omega$  values tend, on average, to produce more diversified, higher-quality



**Fig. 4 | Portfolio stability under a 20% random shock to a single asset's return, repeated 1000 times.** Each plot corresponds to a QSW preset (blue) vs. the MPT benchmark (red).

**Table 5 | Portfolio stability under a Monte Carlo shock test (1000 runs)**

Strategy/ preset	$\alpha$	$\beta$	$\lambda$	Mean % Wt change	Std. Dev. % Wt change
QSW presets					
Ultra-diversified	1.0	100.0	10.0	0.07%	0.08%
Moderate-balanced	10.0	10.0	10.0	0.80%	0.69%
Stability-focused	1.0	10.0	100.0	1.40%	2.22%
Balanced-active	10.0	1.0	100.0	1.90%	2.45%
Sharpe-maximizer	100.0	1.0	10.0	3.61%	3.76%
High-activity	100.0	10.0	1.0	3.61%	3.76%
MPT (max-Sharpe)	—	—	—	1.72%	7.82%

**Table 6 | Descriptive statistics across the full 625-point grid**

Metric	1-year window			2-year window		
	Min	Median	Max	Min	Median	Max
Sharpe	0.85	0.96	0.99	0.86	0.95	0.99
Turnover (%)	0.03	56.42	248.42	0.01	48.56	212.27
Efficiency (Sharpe/Turn.)	0.35	1.71	95.72	0.41	1.94	97.50

portfolios, although the magnitude of the HHI effect is small relative to the structural diversification enforced by the QSW dynamics. Second,  $\lambda$  is the “turnover” knob, trading off implementability against efficiency. As expected from its construction,  $\lambda$  behaves as a holding/inertia parameter. Across both windows, it exhibits a moderate positive correlation with HHI (1Y: 0.225, 2Y: 0.242) and a moderate negative correlation with Efficiency (1Y: -0.199, 2Y: -0.219), while its correlation with Sharpe is weak. Forcing the model to “stay where it is” (large  $\lambda$ ) suppresses trading and reduces turnover, but at the cost of higher concentration and lower efficiency. In other words,  $\lambda$  does not change the core performance level so much as it controls the trade-off between practical implementability and how aggressively the strategy rebalances back toward its preferred low-HHI state. Third,  $\alpha$  is a regime-dependent return bias, not a standalone source of alpha.

The return-preference parameter  $\alpha$  shows only weak and inconsistent correlations with Sharpe (positive in the 1-year window, slightly negative in the 2-year window), while maintaining a tendency to reduce HHI and Efficiency when increased. This indicates that, over the 2018–2024 sample, classical return signals behave more like a noisy, regime-dependent bias than a robust source of outperformance: raising  $\alpha$  tilts the walk toward past winners and can slightly improve Sharpe in the short-memory regime, but it does not produce a reliable, horizon-independent Sharpe uplift. Fourth,  $\beta$  looks neutral globally, but is decisive in the winning region. At first glance, the global bar charts in Fig. 7 suggest that the diversification penalty  $\beta$  has almost no correlation with HHI or Sharpe (correlations near zero). Taken in isolation, this would seem to imply that  $\beta$  does not matter. However, the Top-10 list in Table 7 tells a very different story: every optimal configuration sets  $\beta = 500$ , its maximum value. This is not a contradiction but a selection effect: once  $\alpha$  and  $\omega$  are in a favorable regime (low return-chasing, low noise), the QSW is already close to fully diversified ( $HHI \approx 0.01$ ), so further changes in  $\beta$  have little marginal impact and its global correlation is washed out. Conditioning on the high-Sharpe region, however, reveals the true design rule: strong diversification pressure (large  $\beta$ ) is a prerequisite for ending up in the structurally robust, “smart 1/N” regime.


The grid search results provide empirical support for the dual-channel framework introduced earlier. Taken together, the Top-10 list, the full-grid statistics, and the visual diagnostics reveal a genuinely hybrid design: a classical channel that pins the portfolio to a structurally robust baseline, and a quantum channel that fine-tunes how this baseline is implemented in practice. The Top-10 list in Table 7 is the most direct piece of evidence for the classical channel acting as a stabilizer. Across both the 1-year and 2-year windows, every optimal configuration chooses the same classical settings: a minimal return-chasing coefficient ( $\alpha = 0.1$ ) and the strongest possible diversification penalty ( $\beta = 500$ ). In other words, the model “learns” that the 2018–2024 market is effectively noisy from a classical perspective and responds by turning the accelerator  $\alpha$  almost off, while slamming the diversification brake  $\beta$  to its maximum. The 2D heatmaps in Fig. 5 and the correlation plots in Fig. 7 confirm this picture: once  $\alpha$  is kept small and  $\beta$  is large, the portfolio is structurally forced into an almost perfectly diversified “smart 1/N” state ( $HHI \approx 0.01$ ) across a broad region of the grid. The classical channel, therefore, acts as a stabilizer that locks in a robust risk-return profile by overpowering unstable return signals. Given this stabilized baseline, the quantum channel, controlled by  $\omega$ , is then used to fine-tune how the same robust allocation is implemented. This can be seen most clearly in Table 7(a) when we fix  $(\alpha, \beta, \lambda) = (0.1, 500, 0.1)$  and only vary  $\omega$ : as  $\omega$  decreases from 1.0 to 0.2, the Sharpe ratio remains essentially unchanged (0.989 → 0.988); turnover decreases (from 2.9% to 1.5%); efficiency

**Table 7 | Top 10 QSW configurations from the 625-point grid search, compared against benchmarks**

(a) 1-year training (sorted by Sharpe ratio)									
Rank	$\alpha$	$\beta$	$\lambda$	$\omega$	Sharpe	Final val.	HHI	Turn. (%)	Efficiency
1	0.1	500	0.1	1.0	0.989++	3.02+	0.010	2.9	25.6
2	0.1	500	5.0	1.0	0.989++	3.02+	0.010	3.3	23.0
3	0.1	500	0.1	0.8	0.988++	3.01+	0.010	1.9	33.8
4	0.1	500	0.1	0.6	0.988++	3.01+	0.010	1.7	36.3
5	0.1	500	5.0	0.6	0.988++	3.01+	0.010	1.9	34.5
6	0.1	500	0.1	0.4	0.988++	3.00+	0.010	1.6	37.6
7	0.1	500	5.0	0.8	0.988++	3.01+	0.010	2.1	31.5
8	0.1	500	5.0	0.4	0.988++	3.00+	0.010	1.7	36.3
9	0.1	500	0.1	0.2	0.988++	3.00+	0.010	1.5	39.6
10	0.1	500	5.0	0.2	0.988++	3.00+	0.010	1.6	38.6
1/N	—	—	—	—	0.980	2.98	0.010	0.0	$\infty$
MPT	—	—	—	—	0.860	4.56	0.268	480.9	0.18

(b) 2-year training (sorted by Sharpe ratio)									
Rank	$\alpha$	$\beta$	$\lambda$	$\omega$	Sharpe	Final val.	HHI	Turn. (%)	Efficiency
1	0.1	500	0.1	1.0	0.989+	3.01+	0.010	1.4	41.6
2	0.1	500	50.0	0.2	0.989+	2.99+	0.010	10.4	8.7
3	0.1	500	5.0	1.0	0.989+	3.01+	0.010	1.6	38.4
4	0.1	500	0.1	0.8	0.988+	3.00+	0.010	0.9	50.7
5	0.1	500	5.0	0.8	0.988+	3.00+	0.010	1.1	47.9
6	0.1	500	0.1	0.6	0.988+	3.00+	0.010	0.9	53.4
7	0.1	500	5.0	0.6	0.988+	3.00+	0.010	0.9	51.6
8	0.1	500	0.1	0.4	0.988+	2.99+	0.010	0.8	54.2
9	0.1	500	100.0	0.2	0.988+	2.98+	0.010	25.2	3.8
10	0.1	500	5.0	0.4	0.988+	2.99+	0.010	0.9	53.1
1/N	—	—	—	—	0.980	2.98	0.010	0.0	$\infty$
MPT	—	—	—	—	1.360	5.84	0.254	320.1	0.42



**Fig. 5 | Hyper-parameter heatmaps (1-year training).** These plots reveal the “design rules” of the QSW model, showing how  $\alpha$ ,  $\beta$ , and  $\lambda$  control the portfolio’s behavior. a Sharpe ratio. b HHI (concentration). c Efficiency ratio (Sharpe/Turnover).

increases (from 25.6 to 39.6); while HHI stays fixed at 0.010 in all cases. In other words, once the classical channel has pinned the portfolio to a highly diversified, smart 1/N-like state, the quantum channel does not search for a different target allocation; instead, it searches for cheaper ways to reach and maintain that allocation. Consistent with the global correlations, lower  $\omega$  reduces noisy fluctuations and turnover, improving cost-efficiency without sacrificing Sharpe or diversification. This is the practical “quantum exploration” advantage: the QSW simultaneously explores multiple paths to the same robust state and selects the one with the most attractive implementation profile.

Experiment 2 shows that, for a fixed period (2018–2024), the QSW can reliably discover a structurally robust, smart 1/N-like regime and fine-tune

its implementation cost through its hybrid quantum-classical channels. The natural next question is whether this behavior is truly adaptive: can the same mechanism re-learn an appropriate balance between the classical channel ( $\alpha$ ,  $\beta$ ,  $\lambda$ ) and the quantum channel ( $\omega$ ) as market conditions change? Experiment 3 addresses this question by moving from a static to a fully dynamic setting. Instead of identifying a single “best” parameter set over 2018–2024, we run the entire 625-point grid search at *every* quarterly rebalancing date from 1990 to 2024. At each quarter, the QSW optimizer re-estimates its optimal mix of classical and quantum parameters for that specific market regime. This 34-year, rolling experiment serves as the ultimate test of the QSW’s robustness and “smartness” over multiple cycles, crises, and structural shifts.

**Fig. 6 | Full hyper-parameter and metric correlation matrix (1-year training).** Note the near-zero correlation for Sharpe, confirming its robustness, and the strong, opposing correlations for  $\alpha$  and  $\beta$  on HHI and Efficiency.



**Fig. 7 | Hyper-parameter correlation bar charts for 1-year (top row) and 2-year (bottom row) training.** These plots reveal the “design rules” of the QSW model and how its behavior adapts to the input data. Parameter correlations with Sharpe ratio (a, d), HHI (b, e), and efficiency ratio (c, f) for 1-year (top row) and 2-year (bottom row) training.

### Experiment 3: robustness validation (1990–2024)

This Phase 2 experiment stress-tests the QSW as a dynamic hybrid optimizer over a 34-year period (1990–2024) that spans multiple major crises, including the Dot-com bubble, the 2008 global financial crisis, and the 2020 COVID crash. Building directly on the insights from Phase 1, we now ask a harder question: can the same dual-channel mechanism continuously re-

learn an appropriate balance between its classical channel ( $\alpha, \beta, \lambda$ ) and quantum channel ( $\omega$ ) as market conditions change, and can it do so robustly across different equity universes? To answer this, Experiment 3 combines a dynamic dual-channel strategy with multi-universe testing.

The experimental setup involves a dynamic adaptive model that re-optimizes channels every quarter. At each quarterly rebalancing date from

**Table 8 | Summary statistics for Experiment 3 (1990–2024), averaged over 30 random S&P 500 universes**

Strategy	CAGR	Volatility	Sharpe	Calmar	Max DD	Turnover
QSW	$23.3\% \pm 3.4\%$	$20.6\% \pm 1.9\%$	$0.979 \pm 0.082$	$0.63 \pm 0.20$	$38.0\% \pm 5.0\%$	$201.1\% \pm 11.2\%$
MPT	$14.6\% \pm 2.4\%$	$25.8\% \pm 9.7\%$	$0.480 \pm 0.114$	$0.28 \pm 0.07$	$55.1\% \pm 9.7\%$	$333.1\% \pm 13.8\%$
1/N	$17.2\% \pm 3.5\%$	$19.0\% \pm 2.2\%$	$0.742 \pm 0.083$	$0.41 \pm 0.19$	$43.8\% \pm 5.5\%$	$0.1\% \pm 0.1\%$
S&P 500	$8.3\% \pm 0.0\%$	$18.0\% \pm 0.0\%$	$0.368 \pm 0.000$	$0.17 \pm 0.00$	$47.7\% \pm 0.0\%$	N/A

1990 to 2024, the model runs a full 625-point grid search on the most recent 2-year training window. It then selects the parameter set that best balances the classical and quantum channels ( $\alpha, \beta, \lambda, \omega$ ) under that specific market regime. In contrast to Experiment 2, which identified a single “best-fit” configuration for 2018–2024, this experiment allows the QSW optimizer to re-learn its settings every quarter over 34 years. We perform multi-universe testing using 30 independent trials. To ensure that the results are not an artifact of a single, “lucky” 100-stock universe, we repeat the entire 34-year dynamic backtest 30 times. As described in the Methods (“Data and experimental setup”), each trial starts from a different, randomly sampled 100-stock subset of the point-in-time S&P 500 constituents in 1990Q1 and is dynamically maintained: at each quarterly date, delisted stocks are removed and immediately replaced by new names drawn from the contemporaneous S&P 500 membership. This multi-universe design ensures that our robustness conclusions do not rely on a particular asset selection or survivorship bias.

Table 8 reports the cross-universe means and standard deviations. Three patterns stand out. First, QSW delivers the highest average CAGR and Sharpe, with markedly lower dispersion than MPT. Second, this is not achieved by taking more risk: QSW’s volatility lies between 1/N and MPT, its average maximum drawdown is smaller than both MPT and the S&P 500, and its Calmar ratio is materially higher than all benchmarks. Third, QSW attains these outcomes with less turnover than MPT, indicating that the QSW optimizer is not simply “over-trading” its way to higher returns but improves the overall risk-return-cost profile.

Figure 8 complements Table 8 by showing the full cross-universe distributions. The Sharpe violins in Fig. 8(a) show that QSW dominates not only on average but in most individual universes: the bulk of its mass lies above both 1/N and MPT, with little overlap. Figure 8(b) shows the same ordering for Calmar ratios, confirming that the advantage persists after normalizing by maximum drawdown. Finally, the final-value distribution in Fig. 8(c) illustrates how these differences compound over 34 years: starting from \$1, QSW typically reaches several times the terminal wealth of 1/N and an order of magnitude more than MPT or the index. The medians and interquartile ranges closely track the means, indicating that the out-performance is systematic rather than driven by a few outlier universes.

Figure 9 examines how QSW earns its outperformance. On volatility, QSW sits between 1/N and MPT, indicating that its higher Sharpe is driven by better risk pricing rather than simply taking more overall risk. On drawdowns, QSW again occupies the favorable middle ground: its maximum drawdowns are consistently shallower than those of MPT and the index, and are either slightly better or comparable to 1/N. Finally, QSW strikes a pragmatic balance on trading activity: it trades substantially less than MPT while, of course, trading more than the nearly buy-and-hold 1/N portfolio. Taken together, these distributions show that the QSW optimizer improves the overall risk-return-cost trade-off rather than relying on extreme leverage, concentration, or turnover.

It is worth noting that the average QSW turnover in Experiment 3 (around 200% per year) is higher than in the static Phase 1 experiments (typically 20–80% for the best presets in Experiment 1 and a preset mean of about 40% in Experiment 2). This is an expected consequence of moving from a fixed parameter configuration to a fully dynamic hybrid optimizer. At each quarter, the model not only updates the portfolio weights in response to new returns, but may also switch to a different point in the  $(\alpha, \beta, \lambda, \omega)$  grid. Even if the selected configurations are individually low- or

moderate-turnover, the act of re-optimizing and jumping between configurations introduces an additional layer of trading. Crucially, however, QSW still trades substantially less than the MPT benchmark (about 200% vs. 330% annual turnover on average) while delivering markedly higher Sharpe and Calmar ratios. Hence, the higher turnover in Experiment 3 reflects the cost of adaptivity rather than a loss of control.

The win-rate bars in Fig. 10a summarize the robustness of these results. QSW outperforms MPT, 1/N, and the S&P 500 in the vast majority of universes across all key metrics (Sharpe, Calmar, CAGR, maximum drawdown, and final value), with win rates close to 100% in most comparisons. Figure 10b shows the same message in risk-return space: QSW outcomes form a tight cluster along a high-Sharpe ridge, while 1/N sits on a lower ridge and MPT outcomes are both more dispersed and systematically below the QSW cloud. The index lies at the bottom-left corner, with the lowest CAGR and Sharpe. This confirms that QSW’s advantage is not the result of a few extreme outliers, but a stable shift of the entire risk-return distribution.


Figure 11 shows that the shape of the QSW curves closely tracks that of the 1/N benchmark across all 30 universes. Whenever the equal-weight portfolio delivers a higher final wealth, Sharpe, or Calmar ratio in a given universe, the QSW strategy tends to sit slightly above it, with a very similar risk and drawdown profile and a moderate increase in turnover. In contrast, the MPT series is both noisier and structurally different, with large swings in volatility, drawdown, and trading activity. This behavior is exactly what one would expect from the “smart 1/N” structure identified in Experiment 2: the dynamic QSW optimizer repeatedly re-discovers an almost equal-weight allocation—driven by small  $\alpha$ , large  $\beta$ , and low-to-moderate  $\omega$ —and then adds a thin, data-driven tilt on top of that baseline. Across universes, QSW therefore behaves not as an unstable maximizer, but as a robust, smart version of 1/N: it inherits the diversification and resilience of the naive equal-weight portfolio, while systematically lifting its risk-adjusted performance.

While the cross-universe results establish statistical robustness, it is also instructive to inspect a single, representative universe in detail. Figure 12 illustrates the time evolution of the dynamic QSW strategy on a randomly sampled 100-stock universe, together with the corresponding path of the optimal hyper-parameters.

Several features are worth noting. First, the dynamic QSW strategy consistently stays ahead of both MPT and 1/N, with the performance gaps widening after each major crisis episode (Dot-com, 2008, and COVID-19). This confirms that the gains observed in the cross-universe statistics are not confined to a particular bull market regime, but persist across markedly different environments. Second, the bottom panel shows that the optimizer does not lock into a single parameter setting. Instead, it repeatedly re-learns a familiar pattern:  $\alpha$  is kept low most of the time,  $\beta$  is frequently driven to its upper bound, and  $\omega$  toggles between more classical and more quantum regimes depending on market conditions. During volatile or stressed periods, the model tends to reinforce the diversification channel (large  $\beta$ , low-to-moderate  $\omega$ ), while in calmer regimes it allows slightly more classical exposure. This behavior is fully consistent with the design rules inferred from the 2018–2024 grid search in Phase 1.


Figure 13 further shows that QSW experiences shallower and faster-recovering drawdowns than MPT and the S&P 500, while maintaining volatility in a similar range to the 1/N benchmark. The QSW optimizer, therefore, does not achieve its higher CAGR by leaning into extreme tail risk; instead, it improves the balance between return and drawdown over the full

**Fig. 8 | Distribution of risk-adjusted performance and final wealth across 30 random S&P 500 universes (1990–2024).** Each violin summarizes the cross-universe distribution for one strategy; dots show individual universes, and horizontal bars mark the mean. **a** Sharpe ratio distribution. **b** Calmar ratio (CAGR/Max DD) distribution. **c** Final portfolio value (log scale).



34-year horizon. For brevity, we only report a representative subset of the 30-universe results and one detailed case study in this section. Additional diagnostic figures—including per-experiment performance panels for all 30 universes (final portfolio values, Sharpe ratios, CAGRs, Calmar ratios, maximum drawdowns, and annual turnover), as well as extended single-universe comparisons—are provided in an online supplementary image

archive at <https://github.com/aceest/quantum-stochastic-walk-for-Portfolio-Optimization-Theory-and-Implementation-on-Financial-Networks>, in the `30_experiments` subdirectory. These supplementary figures are fully consistent with the results reported here and allow interested readers to visually inspect the behavior of the QSW strategy across all individual universes and metrics.



**Fig. 9 | Risk and trading-cost profile across 30 universes.** QSW achieves higher risk-adjusted returns than 1/N and MPT while keeping volatility and drawdowns at moderate levels and trading substantially less than the MPT optimizer. **a** Annual turnover distribution. **b** Annual volatility distribution. **c** Maximum drawdown distribution.

## Discussion

This section synthesizes the empirical findings from the Results section, discusses their implications for portfolio construction, and outlines limitations and directions for future work. Our experiments were designed to answer three core questions: (i) whether QSW-based portfolios can improve on classical MPT and naive 1/N benchmarks; (ii) how sensitive the QSW performance is to its hyperparameters; and (iii) whether these properties persist over long horizons and across different equity universes.

Across all Phase 1 experiments on the top-100 S&P 500 universe (2018–2024), QSW-based portfolios consistently exhibited risk and diversification characteristics closely aligned with the naive 1/N benchmark. In Experiment 1, all six QSW strategy presets achieved Sharpe ratios around 0.96–0.98, volatilities near 17%, and HHI values  $\approx 0.01$ , nearly identical to 1/N, while systematically outperforming the MPT maximum-Sharpe


portfolio in the 1-year setting. Beyond raw risk-adjusted returns, the Monte Carlo shock analysis in Experiment 1 reveals a critical structural advantage of the QSW framework: intrinsic resistance to the “error maximization” that plagues classical optimization. While the MPT benchmark exhibited extreme fragility at the micro level—amplifying minor input perturbations into massive weight fluctuations with a standard deviation of nearly 8%—the QSW allocations remained remarkably stable. This stability is architectural rather than accidental; the QSW parameters act as tunable dampers, where the diversification penalty ( $\beta$ ) and holding coefficient ( $\lambda$ ) effectively insulate the portfolio from input noise. Unlike classical solvers that require artificial constraints to prevent extreme rebalancing in response to estimation error, the QSW dynamics naturally suppress unforced turnover, providing a transparent mechanism to control the model’s sensitivity to market noise. Experiment 2 generalized this behavior: the full 625-point grid search showed that the entire QSW parameter space yields HHI in the tight range [0.010, 0.014], more than 20x more diversified than the MPT benchmark ( $HHI \approx 0.26$ ), and Sharpe ratios that are both high and remarkably stable. Taken together, these results indicate that the QSW dynamics are structurally anchored to a highly diversified, equal-weight-like regime, while still allowing for mild, data-driven tilts.


The grid search in Experiment 2 revealed that the Top-10 configurations (by Sharpe) all converge to the same pattern: minimal return-chasing ( $\alpha = 0.1$ ) and maximal diversification pressure ( $\beta = 500$ ), with low-to-moderate  $\lambda$  and  $\omega$ . Heatmaps and correlation analyzes showed that  $\beta$  acts as a structural “brake” that pins HHI at the 1/N floor across large regions of the grid, while  $\omega$  and  $\lambda$  act as behavioral knobs that trade off turnover and cost-efficiency without destabilizing performance. Importantly, the Sharpe surface is flat and robust: even the worst QSW configuration in the grid has a Sharpe comparable to, or better than, MPT under a 1-year window, and the median QSW Sharpe is substantially higher. This suggests that the QSW framework does not rely on a single fine-tuned parameter vector, but instead defines a broad, robust region of “good” configurations.

Phase 2 (Experiment 3) stress-tested the framework as a dynamic QSW optimizer from 1990 to 2024, with quarterly re-optimization over rolling 2-year windows and 30 independent 100-stock universes randomly sampled from point-in-time S&P 500 constituents. Over 34 years and 30 universes, QSW remains a high-Sharpe, low-fragility strategy. On average, QSW achieved a CAGR of about 23%, a Sharpe ratio near 1.0, and a Calmar ratio materially above those of MPT, 1/N, and the S&P 500, while maintaining volatility between 1/N and MPT, and maximum drawdowns smaller than or comparable to 1/N and well below MPT and the index. The dynamic optimizer accepts a higher average turnover (around 200% per year) than in the static Phase 1 experiments (typically 20–80% for the best presets and about 40% on average in Experiment 2) as the cost of adaptivity, but still trades substantially less than MPT’s 300–400% levels. Distributional plots, win-rate charts in Fig. 10a, and risk-return scatter plots in Fig. 10b all support the same conclusion: QSW shifts the entire risk-return cloud upward relative to the benchmarks, rather than merely producing a few lucky outliers.

The single-universe case study in Experiment 3 shows how the rolling grid search behaves through the Dot-com bubble, the 2008 crisis, and the COVID shock. The parameter paths exhibit a recurring pattern:  $\alpha$  stays low,  $\beta$  is often at or near its upper bound, and  $\omega$  toggles between more classical and more quantum regimes as volatility and market stress change. The dynamic hybrid optimizer repeatedly re-discovers the smart 1/N regime. The resulting portfolios retain a near-1/N diversification profile but adjust turnover and minor tilts to match the current regime. This behavior is consistent with the design rules inferred in Experiment 2, confirming that the same “smart 1/N” logic persists when the optimizer is allowed to adjust itself over three decades.

The empirical results have several implications for practical portfolio construction and for the role of quantum-inspired methods in finance. First, equal-weight remains a powerful baseline, but can be improved in a structured way. Our findings reinforce existing evidence that naive 1/N is a





**Fig. 11 | Performance metrics across all 30 universes (1990–2024).** Each panel plots one metric (final value, Sharpe, CAGR, Calmar ratio, maximum drawdown, and annual turnover) for QSW, MPT, 1/N, and the S&P 500 as a function of the universe index.


concentrated portfolios—a failure mode that classical MPT is particularly prone to, especially under short or noisy training windows.

Third, turnover and implementation constraints must be treated as first-class citizens. A recurring theme in our experiments is that nominal Sharpe or CAGR numbers can be deeply misleading if turnover and implementation costs are ignored. Classical MPT often wins on “paper Sharpe” but loses in any realistic cost-adjusted comparison due to extreme turnover and concentration. By contrast, the QSW design rules naturally push the framework into regions of the parameter space where efficiency (Sharpe per unit of turnover) is high and diversification is maintained. The comparison between Phase 1 and Phase 2 makes this trade-off explicit. In the static 2018–2024 setting of Experiments 1 and 2, the best QSW presets typically operate with annual turnover in the 20–80% range (around 40% on average), reflecting a low-footprint, “smart 1/N” implementation. In the dynamic robustness experiment (Experiment 3), the QSW optimizer accepts a higher average turnover of roughly 200% per year in exchange for quarterly re-optimization of the classical and quantum channels over a rolling 2-year window. Even so, QSW still trades substantially less than the MPT maximum-Sharpe benchmark (about 330% on average) while delivering markedly higher


Sharpe and Calmar ratios. In this sense, turnover is not eliminated but budgeted: the extra trading capacity is allocated to regime adaptation rather than to fragile, high-frequency re-optimization of a highly concentrated portfolio.

Fourth, robustness to estimation error is more valuable than local optimality. From the Monte Carlo shock tests in Experiment 1 to the 625-point grid search and the 30-universe robustness study, a consistent pattern emerges: QSW prioritizes stability and robustness over extreme local optimality. In many institutional contexts, especially where mandates are sizeable and capital is sticky, such robustness is more valuable than squeezing out a few extra basis points of back-tested Sharpe from a highly fragile solution. The QSW framework, therefore, appears particularly well-suited to applications where estimation error, regime shifts, and market frictions are unavoidable.

Despite the encouraging results, several limitations of the present study should be acknowledged. First, our experiments focus on U.S. large-cap equities (S&P 500 constituents) and a specific 34-year historical period. While this horizon includes multiple crises and structural shifts, the conclusions may not directly transfer to other asset classes (e.g., credit, commodities, FX), to strongly illiquid markets, or to regimes with very different



**Fig. 12 | Single-universe case study (1990–2024).** Top: Cumulative portfolio values (log scale) for QSW, MPT, 1/N, and the S&P 500, with major crisis periods highlighted. Bottom: Evolution of the QSW parameters ( $\alpha$ ,  $\beta$ ,  $\lambda$ ,  $\omega$ ) selected by the rolling 2-year grid search at each quarter.



**Fig. 13 | Risk path for the single-universe case study.** Left: cumulative drawdowns over time; Right: rolling 8-quarter volatility.

microstructure and regulation. Extending the analysis to multi-asset portfolios and to non-U.S. markets is a natural next step. Second, implementation costs are treated via stylized turnover-based approximations (e.g., a constant round-trip rate). In practice, transaction costs depend on trade size, volatility, depth, and contemporaneous order flow, among other factors. A more detailed cost model, calibrated to actual execution data, could refine the comparison between QSW and MPT and may alter the optimal choice of  $\alpha$ ,  $\beta$ ,  $\lambda$ , and  $\omega$  in the optimizer. Third, the 625-point grid search provides a transparent way to explore the QSW parameter space, but it is not necessarily optimal from a computational or statistical standpoint. More sophisticated search methods (Bayesian optimization, bandit-style exploration, or meta-learning across universes) could lead to better parameter selection with fewer evaluations. Moreover, we have not explicitly optimized the re-optimization frequency; quarterly updates are a pragmatic choice, not a theoretically justified optimum. Fourth, the particular QSW formulation used here is only one of many possible quantum-inspired stochastic processes. Other variants—including different graph constructions, alternative decoherence models, or multi-particle walks—may yield

different behavior. Similarly, we have not yet incorporated factor structures, sector constraints, or more complex risk models into the QSW dynamics. These design choices could be important in practical settings and warrant further exploration in future work.

The framework introduced in this paper opens several directions for further research. A natural extension is to apply QSW-based optimization to multi-asset portfolios that combine equities, fixed income, commodities, and alternative investments, potentially with a factor-based representation (value, momentum, quality, macro factors, etc.). In such settings, the graph structure could encode both security-level and factor-level relationships, allowing the quantum-walk dynamics to propagate information across multiple layers. Future work could also integrate more realistic transaction-cost models, liquidity constraints, and regulatory or mandate-specific constraints (e.g., tracking error, sector caps, ESG limits) directly into the QSW optimization. This would allow a joint design of the quantum and classical channels that explicitly targets cost-adjusted utility or drawdown-sensitive objectives. In particular, replacing the pure Sharpe objective used in the quarterly grid search of Experiment 3 with a cost-aware criterion—for

example, maximizing the efficiency score  $\mathcal{E}$  or Sharpe subject to an explicit turnover cap—could reduce the average turnover from its current ~200% level while preserving most of the adaptive benefit. Another promising direction is to embed QSW within an online-learning or change-point detection framework, allowing it to respond more intelligently to sudden regime shifts—for example, by adjusting the re-optimization frequency or temporarily increasing diversification pressure when uncertainty is high. Finally, although our implementation is entirely classical and quantum-inspired, the QSW framework is designed to be compatible with future quantum hardware. Exploring hardware-efficient encodings of the QSW dynamics and co-designing algorithms that leverage early quantum devices for the most computationally intensive subroutines (e.g., sampling, large-scale graph propagation) is an interesting avenue as quantum technologies mature.

In conclusion, this work introduces and empirically validates a QSW-based framework for portfolio construction that combines the robustness of naive 1/N with the flexibility of a parameterized, hybrid quantum-classical optimizer. Across multiple experiments, time horizons, and equity universes, the QSW model consistently delivers higher risk-adjusted returns than classical MPT and naive 1/N, while maintaining strong diversification and moderate turnover. The central message is not that quantum-inspired models magically generate alpha, but that they can enforce desirable structural properties—in particular, diversification and robustness to estimation error—in a principled way. As financial markets continue to grapple with noisy data, regime uncertainty, and implementation frictions, such structurally robust, “smart 1/N” approaches may offer a valuable alternative to both naive diversification and fragile, high-dimensional optimization. We hope that the QSW framework presented here can serve as a foundation for further work at the intersection of quantum-inspired algorithms and practical portfolio management.

## Methods

### QSW formulation on financial graphs

Quantum stochastic walks (QSWs) extend classical random walks by blending coherent quantum evolution with controlled decoherence on the same graph. The coherent channel, driven by a Hamiltonian, allows superposition and interference across neighboring nodes, while the stochastic channel, implemented via a Lindblad dissipator, guarantees ergodicity and convergence to a unique stationary state. This dual-channel design overcomes the shortcomings of purely classical diffusion (no interference, fragile reliance on local correlations) and purely quantum walks (no stationary distribution), making QSWs a natural engine for portfolio allocation on financial networks.

Regarding classical random walks and the PageRank component, consider a graph  $G = (V, E)$  with  $n$  nodes representing financial assets. A simple continuous-time random walk can be described by the master equation

$$\frac{d\vec{p}}{dt} = (P - I)\vec{p}, \quad (1)$$

for a fixed unit of excess return where  $\vec{p} = (p_1, \dots, p_n)^\top$  is the probability distribution over nodes,  $P$  is a row-stochastic transition matrix, and  $(P - I)$  plays the role of an infinitesimal generator. The PageRank algorithm<sup>38</sup> enhances this framework by introducing a “teleportation” step:

$$\mathcal{G} = \alpha_{\text{damp}}P + (1 - \alpha_{\text{damp}})\frac{\mathbf{1}\mathbf{1}^\top}{n}, \quad (2)$$

where  $\alpha_{\text{damp}} \in (0, 1)$  is the damping parameter (typically  $\approx 0.85$ ) and  $\mathbf{1}\mathbf{1}^\top$  is the  $n \times n$  all-ones matrix. This construction mixes graph-based transitions with a uniform teleportation step, ensuring that  $\mathcal{G}$  is irreducible and aperiodic and therefore admits a unique stationary distribution even for directed graphs with dangling nodes.

A continuous-time quantum walk (CTQW) replaces the classical distribution  $\vec{p}$  with a quantum state  $|\psi\rangle$  evolving under the Schrödinger equation<sup>39,40</sup>

$$i\hbar \frac{d|\psi\rangle}{dt} = H|\psi\rangle, \quad (3)$$

where  $H$  is a Hermitian Hamiltonian. The state evolves as  $|\psi(t)\rangle = e^{-iHt/\hbar}|\psi(0)\rangle$ , and the probability of finding the walker at node  $j$  is  $P_j(t) = |\langle j|\psi(t)\rangle|^2$ . While CTQWs can achieve speedups on certain graph problems<sup>41</sup>, they lack a stationary state and are highly sensitive to decoherence.

QSWs interpolate between the classical master equation (1) and the Schrödinger-type CTQW (3) via the Gorini–Kossakowski–Lindblad–Sudarshan (GKLS) master equation<sup>36</sup>:

$$\frac{d\rho}{dt} = -i(1 - \omega)[H, \rho] + \omega \sum_{i,j} c_{ij}(L_{ij}\rho L_{ij}^\dagger - \frac{1}{2}\{L_{ij}^\dagger L_{ij}, \rho\}), \quad (4)$$

where  $\rho$  is the density matrix,  $\omega \in [0, 1]$  is the quantum-classical mixing parameter,  $H$  is the Hamiltonian,  $L_{ij} = |i\rangle\langle j|$  are Lindblad jump operators representing transitions from node  $j$  to node  $i$ , and  $c_{ij} \geq 0$  are transition rates. The generator consists of two competing terms: a coherent term  $-i(1 - \omega)[H, \rho]$  that drives unitary evolution, and a dissipative term  $\omega \sum_{i,j} c_{ij}(\cdot)$  that introduces decoherence and classical jumps. The mixing parameter  $\omega$  controls the balance:  $\omega = 0$  recovers a pure CTQW,  $\omega = 1$  a classical stochastic process, and  $0 < \omega < 1$  a hybrid dynamics. For optimization applications we require convergence to a unique stationary state. Following<sup>42</sup>, if we set  $c_{ij} = \mathcal{G}_{ij}$  where  $\mathcal{G}$  is the Google matrix (2), then for any  $\omega > 0$  the quantum dynamical semigroup generated by Eq. (4) is primitive: a unique, full-rank stationary state  $\rho_\infty$  exists and  $\rho(t) \rightarrow \rho_\infty$  exponentially fast. We define the final portfolio weights as  $w_i = \rho_{\infty,ii}$  i.e., the diagonal of the stationary density matrix.

We encode financial data on the graph by taking  $n$  financial assets with daily returns  $\{r_{i,t}\}$  over a training window of length  $T$ . Each node  $i \in \{1, \dots, n\}$  is characterized by its mean return  $\mu_i = \frac{1}{T} \sum_{t=1}^T r_{i,t}$ , its mean excess return  $\mu_i^{(\text{ex})} = \mu_i - r_f^{(d)}$ , its volatility  $\sigma_i$ , and its daily Sharpe ratio  $SR_i = \mu_i^{(\text{ex})}/\sigma_i$ , where  $r_f^{(d)}$  is the daily risk-free rate. The graph weight matrix  $W$  encodes both individual asset quality and pairwise relationships:

$$W_{ij} = \begin{cases} \exp(\alpha SR_j - \beta \Sigma_{ij}) & \text{if } i \neq j, \\ \exp(\lambda SR_i) & \text{if } i = j, \end{cases} \quad (5)$$


where  $\alpha > 0$  controls the preference for high-Sharpe destinations,  $\beta > 0$  penalizes transitions between highly correlated assets,  $\Sigma_{ij}$  is the sample covariance between assets  $i$  and  $j$  over the same window, and  $\lambda$  is the holding coefficient for self-loops. Row-normalizing  $W$  yields a stochastic transition matrix

$$P_{ij} = \frac{W_{ij}}{\sum_k W_{ik}}, \quad (6)$$

which is then embedded into the Google matrix  $\mathcal{G}$  in Eq. (2). The same financial information is encoded in the Hamiltonian via


$$H_{ij} = \begin{cases} -\gamma_1 SR_i & \text{if } i = j, \\ \gamma_2 \widehat{\Sigma}_{ij} & \text{if } i \neq j, \end{cases} \quad (7)$$

where  $\widehat{\Sigma}_{ij}$  is a normalized covariance  $\widehat{\Sigma}_{ij} = (\Sigma_{ij} - \Sigma_{\min})/(\Sigma_{\max} - \Sigma_{\min})$  and  $\gamma_1, \gamma_2$  are scaling factors chosen for numerical stability. This construction yields a dual-channel encoding of financial data: the coherent channel explores highly correlated clusters through  $H$ , while the stochastic channel enforces diversification through  $\mathcal{G}$  and the rates  $c_{ij} = \mathcal{G}_{ij}$ . The steady-state diagonal of



**Fig. 14** | Conceptual illustration of the quantum stochastic walk (QSW) as a dual-channel engine on a financial graph: a coherent channel (purple) explores correlated clusters through Hamiltonian evolution, while a stochastic channel (blue)

implements Google-type jumps that enforce diversification and guarantee convergence to a unique stationary state.



**Fig. 15** | Four-stage quantum-stochastic-walk (QSW) portfolio framework employed in all subsequent experiments.

$\rho_\infty$  then defines a well-behaved, graph-based mapping from return/covariance inputs to portfolio weights. Figure 14 summarizes how the QSW operates on the financial graph. At each step, the portfolio state (represented by node occupancies on the left-hand graph) is updated by two complementary channels acting on the same network. The coherent channel (purple arrow) applies Hamiltonian evolution, allowing the walk to tunnel within highly correlated clusters and to explore multiple paths in superposition. This captures nonlinear dependency structure that is invisible to purely classical diffusion. The stochastic channel (blue arrow) then applies a Google-type update based on the weighted transition matrix, which damps probability mass inside tight clusters and redistributes it toward under-represented nodes. Iterating these two channels in alternation drives the system toward a unique stationary density matrix, whose diagonal elements define the final portfolio weights on the right-hand graph. In this way, the QSW framework implements portfolio allocation as a dual-channel process that jointly encodes return, risk and correlation directly in the dynamics of the walk.

### QSW algorithm implementation

We implement the QSW-based portfolio optimizer in four stages, summarized in Fig. 15.

Stage 1: Data processing - Ingest daily adjusted close prices. For each rebalancing date, compute inputs from a rolling window of  $T_w$  trading days of simple returns  $r_{i,t} = P_{i,t}/P_{i,t-1} - 1$ :

- Daily statistics: sample mean  $\mu_i$ , sample standard deviation  $\sigma_i$ , and the  $N \times N$  sample covariance matrix  $\Sigma$  over the window.
- Sharpe ratios: daily Sharpe for each asset

$$SR_i = \frac{\mu_i - r_f^{(d)}}{\sigma_i},$$

where we assume an annual risk-free rate of 3% and set  $r_f^{(d)} = 0.03/252$ .

Stage 2: Dual-channel graph—Embed these statistics in a dual network:

- Stochastic channel: construct the weight matrix  $W_{ij} \propto \exp(\alpha SR_j - \beta \Sigma_{ij})$ , row-normalize it to obtain a transition matrix  $P$ , and form the Google matrix  $\mathcal{G}$  as in Eq. (2); this favors moves toward high-Sharpe, low-covariance destinations.
- Coherent channel: build a Hamiltonian with  $H_{ii} \propto SR_i$  and  $H_{ij} \propto \hat{\Sigma}_{ij}$  for  $i \neq j$ , where  $\hat{\Sigma}_{ij}$  is the normalized covariance, capturing quantum interference across correlated assets.

Stage 3: QSW solver—Evolve the density operator under the quantum-classical mix  $\omega$  (GKLS generator) until convergence to the stationary state  $\rho_\infty$ .

Stage 4: Weights—Extract the diagonal of  $\rho_\infty$  and normalize it to obtain fully invested portfolio weights  $\mathbf{w}$  with  $\sum_i w_i = 1$ , which are then fed into the back-test loop.

We use efficient Kraus-operator evolution to integrate the GKLS master equation (4) because state-vector unraveling increases the state dimension from  $N$  to  $N^2$  and is computationally expensive. To accelerate the simulation, we map the Hamiltonian and Lindblad operators  $\{H, L_{ij}\}$  to a set of Kraus operators  $\{K_0, K_{ij}\}$  and implement the quantum channel in operator-sum form:

$$p(t + \Delta t) = K_0(\Delta t) p(t) K_0^\dagger(\Delta t) + \sum_{i,j} K_{ij}(\Delta t) p(t) K_{ij}^\dagger(\Delta t), \quad (8)$$

with the trace-preserving condition

$$K_0^\dagger K_0 + \sum_{i,j} K_{ij}^\dagger K_{ij} = I. \quad (9)$$

Discretizing Eq. (4) yields

$$\rho(t + \Delta t) = \rho(t) - i(1 - \omega)[H, \rho(t)]\Delta t + \omega \sum_{i,j} c_{ij}(L_{ij}\rho(t)L_{ij}^\dagger - \frac{1}{2}\{L_{ij}^\dagger L_{ij}, \rho(t)\})\Delta t + \mathcal{O}(\Delta t^2). \quad (10)$$

Comparing Eqs. (10) and (8) leads to the approximate Kraus operators

$$K_{ij} = \sqrt{1 - e^{-\omega c_{ij}\Delta t}} L_{ij}, \quad K_0 = \sqrt{I - \sum_{i,j} K_{ij}^\dagger K_{ij}} e^{-i(1-\omega)H\Delta t}. \quad (11)$$

In our setting each  $K_{ij}$  has the form  $\gamma_{ij} |i\rangle\langle j|$  and is therefore extremely sparse. The double-sided product simplifies to

$$K_{ij}\rho K_{ij}^\dagger = \gamma_{ij}^2 \rho_{jj} |i\rangle\langle i|, \quad (12)$$

so the dissipative term only involves the diagonal of  $\rho$ . Defining  $P = (\rho_{11}, \rho_{22}, \dots, \rho_{NN})^\top$  and a matrix  $\Gamma$  with entries  $\Gamma_{ij} = \gamma_{ij}^2$ , we obtain

$$\sum_{i,j} K_{ij}\rho K_{ij}^\dagger = \Gamma P, \quad (13)$$

and the operator-sum update can be written as

$$\rho' = K_0\rho K_0^\dagger + \Gamma P. \quad (14)$$

Thus, the expensive sum of many matrix-matrix products reduces to a single matrix-vector multiplication.

An iterative solver initializes the density matrix at the maximally mixed state  $\rho^{(0)} = I/n$ , chooses a time step  $\Delta t$  and tolerance  $\epsilon$ , and iterates Eq. (14) until convergence. Algorithm 1 summarizes the implementation used in our experiments.

#### Algorithm 1. GPU-accelerated QSW portfolio optimizer

**Require:** Historical returns, QSW parameters  $\{\omega, \alpha, \beta, \lambda\}$

**Ensure** Portfolio weights  $\mathbf{w}$

- 1: Compute  $\mu_i, \sigma_i, \text{SR}_i, \Sigma_{ij}$  from the training window
- 2: Construct weight matrix  $W$  and Hamiltonian  $H$ ; form Google matrix  $\mathcal{G}$
- 3: Initialize  $\rho^{(0)} = I/n$ , choose  $\Delta t$ , tolerance  $\epsilon$
- 4: **repeat**
- 5: Build Kraus operators  $K_0, K_{ij}$  from Eq.(11)
- 6:  $\tilde{\rho} \leftarrow K_0\rho^{(t)}K_0^\dagger$
- 7:  $P \leftarrow \text{diag}(\tilde{\rho})$
- 8:  $\rho^{(t+1)} \leftarrow \tilde{\rho} - \text{diag}(P) + \text{diag}(\Gamma P)$
- 9: Normalize  $\rho^{(t+1)} \leftarrow \rho^{(t+1)} / \text{Tr}(\rho^{(t+1)})$
- 10: **until**  $\|\rho^{(t+1)} - \rho^{(t)}\|_1 < \epsilon$  or  $t$  reaches a maximum
- 11: Set  $w_i = \rho_{ii}^{(\infty)} / \sum_j \rho_{jj}^{(\infty)}$  for  $i = 1, \dots, n$
- 12: **return**  $\mathbf{w}$

The computational efficiency is ensured by GPU acceleration. The implementation leverages cuPyNumeric<sup>43</sup> for linear-algebra kernels, an NVIDIA-developed library that serves as a drop-in replacement for CPU-based NumPy. It accelerates all core primitives in the QSW solver, including matrix exponentials for unitary evolution, element-wise operations for transition rates, and diagonal extraction and reconstruction. In practice, a single NVIDIA A100 GPU can evaluate the full 625-parameter grid of Experiment 2 in roughly three hours of wall-clock time (about 17 s per configuration), including data loading, graph construction, QSW convergence, and metric computation. This enables real-time or near-real-time portfolio optimization for universes of order  $10^2$ – $10^3$  assets on modern GPU hardware.

#### Data and experimental setup

For the asset universe and data collection, we consider two types of equity universes. In Phase 1 (parameter exploration), we use the 100 largest U.S.

stocks by market capitalization (“top-100” universe), including technology names (e.g., AAPL, MSFT, GOOGL, AMZN, NVDA), financials (e.g., JPM, BAC, WFC), healthcare (e.g., JNJ, UNH, PFE), and other sectors. This universe provides sufficient cross-sectional diversity while keeping the QSW simulations computationally tractable. To assess robustness to stock selection, Phase 2 employs a multi-universe design based on 30 random subsets of the S&P 500. Each trial defines a distinct 100-stock universe that is maintained in a point-in-time fashion to avoid survivorship bias: at the initial date (1990Q1) we draw 100 constituents from the S&P 500 members active at that time; at each quarterly rebalance we remove any delisted names (e.g., due to mergers or bankruptcies) and replace them by randomly sampling from the S&P 500 constituents active on that date. Historical daily adjusted close prices from 1988-01-01 to 2024-12-31 are obtained from Yahoo Finance.

All strategies are evaluated in a common backtesting framework with quarterly rebalancing. The specific backtest horizon and training window depend on the experiment:

- Phase 1: parameter exploration (2018–2024).
  - Backtest period: 2018-01-02 to 2024-12-31 on the top-100 universe.
  - Training window: rolling 1–2 years of daily data (252, 378, or 504 trading days) to study the impact of lookback length.
- Phase 2: multi-universe robustness (1990–2024).
  - Backtest period: 1990-01-02 to 2024-12-31.
  - Training window: fixed 2-year (504 trading days) rolling window.

In both phases the rebalancing frequency is quarterly (every three months), and transaction costs are incorporated indirectly via portfolio turnover statistics. This frequency reflects common institutional practice, balancing responsiveness to changing market conditions against implementation costs.

Regarding computational performance, all QSW simulations and grid searches are implemented in Python using GPU-accelerated linear algebra (see above). On a single NVIDIA A100 GPU, the full 625-parameter grid of Phase 1, Experiment 2 can be evaluated in roughly three hours of wall-clock time (about 17 s per configuration), including data loading, graph construction, QSW convergence, and performance metric computation.

In Phase 1: parameter exploration, we use two complementary approaches on the top-100 universe over 2018–2024. First, we define six QSW strategy presets that capture different investment styles (Table 1) by fixing  $(\alpha, \beta, \lambda)$ :

Each preset is combined with  $\omega \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$  to explore the quantum-classical spectrum. Second, we perform a comprehensive grid search over the four QSW parameters:  $\alpha, \beta, \lambda \in \{0.1, 5, 50, 100, 500\}$  and  $\omega \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ , yielding 625 parameter combinations per training-window choice. This grid is used to identify robust high-performing regions of the parameter space.

In Phase 2: multi-universe robustness, to ensure that our findings are not artifacts of a particular stock universe, we run the full 625-point grid search on each of the 30 point-in-time 100-stock universes over 1990–2024. For each trial we (a) construct the initial universe at 1990Q1 from contemporaneous S&P 500 constituents, (b) maintain the universe through quarterly replacement of delisted names, and (c) backtest all QSW configurations and benchmarks. We then analyze the cross-universe distribution of performance and risk metrics. This multi-universe design mitigates survivorship and selection bias and demonstrates that the observed “smart 1/N” behavior and robustness of the QSW optimizer are not driven by a single favorable sample.

#### Performance metrics and evaluation

We evaluate portfolio performance using a set of complementary metrics that capture returns, risk, drawdown behavior, concentration, and trading efficiency.

For return and risk metrics, we track both level and risk-adjusted returns.

## 1. Cumulative return.

$$V_T = V_0 \prod_{t=1}^T (1 + r_{p,t}), \quad (15)$$

where  $r_{p,t} = \sum_{i=1}^n w_{i,t} r_{i,t}$  is the portfolio return at time  $t$ .

2. Compound annual growth rate (CAGR). Let  $Y$  denote the investment horizon in years. The CAGR is defined as

$$\text{CAGR} = \left( \frac{V_T}{V_0} \right)^{1/Y} - 1, \quad (16)$$

which measures the long-run average growth rate of portfolio wealth.

3. Annualized Sharpe ratio. We compute the Sharpe ratio using daily excess returns and annualize it with 252 trading days:

$$SR = \frac{\mu_p^{(\text{excess})}}{\sigma_p} \sqrt{252}, \quad (17)$$

where  $\mu_p^{(\text{excess})}$  and  $\sigma_p$  are the mean and standard deviation of daily portfolio excess returns (portfolio return minus the daily risk-free rate). In all experiments we assume a constant annual risk-free rate of 3%.

## 4. Maximum drawdown (MDD).

$$\text{MDD} = \max_{t \in [0, T]} \frac{\max_{s \in [0, t]} V_s - V_t}{\max_{s \in [0, t]} V_s}, \quad (18)$$

which captures the worst peak-to-trough loss in percentage terms over the evaluation period.

5. Calmar ratio. To relate long-run growth to downside risk, we use the Calmar ratio,

$$\text{Calmar} = \frac{\text{CAGR}}{|\text{MDD}|}, \quad (19)$$

which measures annualized return per unit of worst drawdown. Unless otherwise noted, MDD is computed over the full evaluation horizon of the experiment (e.g., 1990–2024 for Experiment 3); rolling Calmar ratios in the figures are defined analogously using window-specific MDD values. This metric is used extensively in Experiment 3 to compare strategies on a drawdown-adjusted basis.

For portfolio concentration metrics, we quantify how capital is distributed across names using three standard measures.

## 1. Herfindahl-Hirschman Index (HHI).

$$\text{HHI} = \sum_{i=1}^n w_i^2, \quad (20)$$

where  $w_i$  is the weight of asset  $i$ . A higher HHI indicates greater concentration.

## 2. Effective number of stocks.

$$N_{\text{eff}} = \frac{1}{\text{HHI}}, \quad (21)$$

which can be interpreted as the number of equally weighted assets that would yield the same HHI.

## 3. Top-5 concentration.

$$C_5 = \sum_{i=1}^5 w_{[i]}, \quad (22)$$

where  $w_{[i]}$  denotes the  $i$ -th largest portfolio weight.  $C_5$  measures how much of the portfolio is concentrated in the five largest positions.

We also track trading and efficiency metrics. Portfolio turnover represents the frequency of rebalancing activity and directly impacts implementation costs, market impact, and strategy scalability. High-turnover strategies face substantial transaction costs, price impact during execution, and capacity constraints that can erode theoretical performance gains in practice<sup>44,45</sup>.

## 1. Average annualized turnover.

$$\text{TO}_{\text{ann}} = 4 \times \frac{1}{N_{\text{quarters}}} \sum_{q=1}^{N_{\text{quarters}}} \sum_{i=1}^n |w_{i,q} - w_{i,q-1}|, \quad (23)$$

which expresses the proportion of portfolio value traded *per year*. In buy-side practice, annual turnover above roughly 200% is often considered prohibitively expensive because of spread and market-impact costs, whereas turnover below about 50% is usually viewed as readily implementable<sup>46</sup>. Classical mean-variance optimization frequently produces 300–500% annual turnover, posing significant implementation challenges<sup>14,19</sup>.

## 2. Sharpe-turnover efficiency ratio.

$$\mathcal{E} = \frac{SR}{\text{TO}_{\text{ann}} + 0.01}, \quad (24)$$

which captures the trade-off between risk-adjusted returns and trading intensity by measuring how much Sharpe ratio is achieved per unit of annual turnover<sup>45</sup>. The small constant 0.01 prevents division by zero for extremely low-turnover strategies. Higher efficiency ratios indicate strategies that deliver superior risk-adjusted performance without excessive trading activity—a critical requirement for real-world implementation, where transaction costs, market impact, and capacity constraints dominate practical considerations<sup>44,47</sup>. For example, a strategy achieving a Sharpe ratio of 1.0 with 20% annual turnover ( $\mathcal{E} \approx 5.0$ ) is significantly more valuable than one achieving 1.2 with 400% turnover ( $\mathcal{E} \approx 0.3$ ) once implementation frictions are taken into account<sup>46</sup>.

## Benchmark strategies

We compare QSW-based portfolios against three primary benchmarks.

The classical benchmark is the maximum-Sharpe (MSR) portfolio based on Modern Portfolio Theory (MPT). While the MSR is defined as a non-convex ratio, it can be computed via an equivalent convex quadratic program. Since our inputs  $\mu$  are excess returns (over the risk-free rate), we solve<sup>48</sup>

$$\mathbf{y}^* = \arg \min_{\mathbf{y}} \frac{1}{2} \mathbf{y}^T \Sigma \mathbf{y} \quad (25)$$

subject to  $\mathbf{y}^T \mu = 1$  and  $\mathbf{y} \geq \mathbf{0}$ . This finds the portfolio that minimizes variance for a fixed unit of excess return. The fully invested MPT weights are then  $\mathbf{w}_{MPT} = \mathbf{y}^* / \sum_i y_i^*$ . We implement this using the PyPortfolioOpt library, feeding it the same sample mean  $\mu$  and sample covariance  $\Sigma$  estimated from the rolling window as for the QSW optimizer.

The naive diversification (1/N) strategy assigns an equal weight  $w_i = 1/N$  to each of the  $N$  available assets at every rebalance. It completely avoids estimating  $\mu$  and  $\Sigma$  and is therefore immune to the estimation errors that plague MVO, especially in high dimensions<sup>14</sup>. Despite its simplicity, 1/N is a strong benchmark and serves as the natural reference for evaluating whether QSW provides a genuinely robust improvement.

Finally, the capitalization-weighted S&P 500 index serves as a passive market benchmark, representing the opportunity cost of active management in large-cap U.S. equities.

## Analysis framework

We analyze QSW performance along several complementary dimensions. First, we study parameter sensitivity by analyzing correlations between

QSW parameters ( $\alpha, \beta, \lambda, \omega$ ) and key performance metrics (Sharpe, Calmar, turnover, concentration), and identify robust regions of the parameter space. Second, we examine concentration-performance trade-offs by analyzing the efficient frontier in Sharpe-HHI space, comparing concentration levels versus MPT and 1/N, and assess how effective diversification varies under different parameter settings.

To ensure robust conclusions, we use simple but informative diagnostics. We perform a win-rate analysis by computing the percentage of QSW configurations that outperform each benchmark on a given metric (e.g., Sharpe, Calmar, final value). We also examine cross-universe consistency. In the multi-universe setting (Experiment 3), we examine the distribution of performance across all 30 random 100-stock universes, focusing on how often and by how much QSW dominates MPT, 1/N, and the index on risk and return metrics.

### Summary of methods

The implementation and evaluation methodology is designed to bridge the gap between theoretical innovation and realistic portfolio practice. In particular, it:

1. adopts a two-phase design: a short-horizon parameter exploration (2018–2024, top-100 S&P 500) and a long-horizon robustness phase (1990–2024, rolling 2-year windows);
2. combines three complementary experiment types: fixed strategy presets, a full 625-point hyper-parameter grid search, and a dynamic, multi-universe study;
3. uses quarterly rebalancing and Sharpe ratios computed on excess returns over a 3% annual risk-free rate, in line with institutional practice;
4. evaluates portfolios with a rich set of metrics, including Sharpe, Calmar, maximum drawdown, turnover, and multiple concentration measures (HHI, effective number of stocks, top-5 weights);
5. benchmarks QSW against maximum-Sharpe MPT portfolios, naive 1/N portfolios on the same universe, and the S&P 500 index;
6. validates robustness through extensive parameter-sensitivity analysis and a 34-year, 30-universe experiment based on point-in-time S&P 500 constituents.

Together, these elements provide a rigorous framework for evaluating the QSW framework as a quantum-inspired portfolio construction tool, ensuring that the reported results are not an artifact of a particular sample, universe, or parameter choice but reflect stable behavior under realistic market and implementation conditions.

### Data availability

All numerical results in this study are generated by simulating the QSW optimizer on equity return series as described in the “Methods” section. Daily adjusted close prices for S&P 500 constituents over 1988–2024 were obtained from Yahoo Finance (<https://finance.yahoo.com>); no proprietary or restricted datasets were used. The raw price data can be re-downloaded from the same public source, and all preprocessing steps (return calculation, training windows, universe construction) are fully specified in the text. Aggregated backtest outputs (CAGR, Sharpe, Calmar, drawdown, turnover, and concentration metrics for all configurations and universes), as well as the per-universe diagnostic figures shown in the Results and Supplementary Information, are available from the corresponding author upon reasonable request.

### Code availability

All algorithms and parameter settings required to reproduce the results are described in the Methods section. The custom code used to implement the QSW optimizer and run the backtests is not publicly archived at this time, as it is part of an ongoing research project. Results from specific simulations, as well as pseudo-code and representative implementation details, are available from the corresponding author upon reasonable request for academic and non-commercial purposes.

Received: 4 July 2025; Accepted: 24 December 2025;

Published online: 02 February 2026

### References

1. Markowitz, H. Portfolio selection. *J. Finance* **7**, 77–91 (1952).
2. Fabozzi, F. J., Kolm, P. N., Pachamanova, D. A. & Focardi, S. M. *Robust Portfolio Optimization and Management* (John Wiley & Sons, 2007).
3. Black, F. & Litterman, R. Global portfolio optimization. *Financ. Anal. J.* **48**, 28–43 (1992).
4. Mantegna, R. N. & Stanley, H. E. *Introduction to Econophysics: Correlations and Complexity in Finance* (Cambridge University Press, 1999).
5. Jegadeesh, N. & Titman, S. Returns to buying winners and selling losers: implications for stock market efficiency. *J. Finance* **48**, 65–91 (1993).
6. Fletcher, R. *Practical Methods of Optimization* 2nd edn. (John Wiley & Sons, 2013).
7. Tobin, J. Liquidity preference as behavior towards risk. *Rev. Econ. Stud.* **25**, 65–86 (1958).
8. Sharpe, W. F. Capital asset prices: a theory of market equilibrium under conditions of risk. *J. Finance* **19**, 425–442 (1964).
9. Merton, R. C. An analytic derivation of the efficient portfolio frontier. *J. Financ. Quant. Anal.* **7**, 1851–1872 (1972).
10. Goldfarb, D. & Iyengar, G. Robust portfolio selection problems. *Math. Oper. Res.* **28**, 1–38 (2003).
11. Garlappi, L., Uppal, R. & Wang, T. Portfolio selection with parameter and model uncertainty: a multi-prior approach. *Rev. Financ. Stud.* **20**, 41–81 (2007).
12. Artzner, P., Delbaen, F., Eber, J.-M. & Heath, D. Coherent measures of risk. *Math. Finance* **9**, 203–228 (1999).
13. Rockafellar, R. T. & Uryasev, S. Optimization of conditional value-at-risk. *J. Risk* **2**, 21–42 (2000).
14. DeMiguel, V., Garlappi, L. & Uppal, R. Optimal versus naive diversification: how inefficient is the 1/N portfolio strategy? *Rev. Financial Stud.* **22**, 1915–1953 (2009).
15. Longin, F. & Solnik, B. Extreme correlation of international equity markets. *J. Finance* **56**, 649–676 (2001).
16. Ang, A. & Chen, J. Asymmetric correlations of equity portfolios. *J. Financ. Econ.* **63**, 443–494 (2002).
17. Cont, R. Empirical properties of asset returns: stylized facts and statistical issues. *Quant. Financ.* **1**, 223–236 (2001).
18. Chang, T.-J., Meade, N., Beasley, J. E. & Sharaiha, Y. M. Heuristics for cardinality constrained portfolio optimisation. *Comput. Oper. Res.* **27**, 1271–1302 (2000).
19. Best, M. J. & Grauer, R. R. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: Some analytical and computational results. *Rev. Financ. Stud.* **4**, 315–342 (1991).
20. Chopra, V. K. & Ziemba, W. T. The effect of errors in means, variances, and covariances on optimal portfolio choice. *J. Portf. Manag.* **19**, 6–11 (1993).
21. Gu, S., Kelly, B. & Xiu, D. Empirical asset pricing via machine learning. *Rev. Financ. Stud.* **33**, 2223–2273 (2020).
22. Jiang, Z., Xu, D. & Liang, J. A deep reinforcement learning framework for the financial portfolio management problem. Preprint at arXiv <https://doi.org/10.48550/arXiv.1706.10059> (2017).
23. Zhang, Z., Zohren, S. & Roberts, S. Deep learning for portfolio optimization. *J. Financ. Data Sci.* **2**, 8–20 (2021).
24. Han, S. Risk budgeting portfolio optimization with deep reinforcement learning. *J. Financ. Data Sci.* **5**, 86–101 (2023).
25. Barredo Arrieta, A. et al. Explainable artificial intelligence (XAI): concepts, taxonomies, opportunities and challenges. *Inf. Fusion* **58**, 82–115 (2020).
26. Mantegna, R. N. Hierarchical structure in financial markets. *Eur. Phys. J. B* **11**, 193–197 (1999).

27. Onnela, J.-P., Chakraborti, A., Kaski, K. & Kertész, J. Dynamic asset trees and black monday. *Phys. A Stat. Mech. Appl.* **324**, 247–252 (2003).

28. Tumminello, M., Lillo, F. & Mantegna, R. N. Correlation, hierarchies, and networks in financial markets. *J. Econ. Behav. Organ.* **75**, 40–58 (2010).

29. Pozzi, F., Di Matteo, T. & Aste, T. Spread of risk across financial markets: better to invest in the peripheries. *Sci. Rep.* **3**, 1665 (2013).

30. Kenett, D. Y., Preis, T., Gur-Gershgoren, G. & Ben-Jacob, E. Dependency network and node influence: application to the study of financial markets. *Int. J. Bifurc. Chaos* **22**, 1250181 (2012).

31. Whitfield, J. D., Rodriguez-Rosario, C. A. & Aspuru-Guzik, A. Quantum stochastic walks: a generalization of classical random walks and quantum walks. *Phys. Rev. A* **81**, 022323 (2010).

32. Attal, S., Petruccione, F., Sabot, C. & Sinayskiy, I. Open quantum random walks. *J. Stat. Phys.* **147**, 832–852 (2012).

33. Sánchez-Burillo, E., Duch, J., Gómez-Gardeñes, J. & Zueco, D. Quantum navigation and ranking in complex networks. *Sci. Rep.* **2**, 605 (2012).

34. Paparo, G. D., Müller, M., Comellas, F. & Martin-Delgado, M. A. Quantum Google in a complex network. *Sci. Rep.* **2**, 444 (2013).

35. Wang, Y., Xue, S., Wu, J. & Xu, P. Continuous-time quantum walk based centrality testing on weighted graphs. *Sci. Rep.* **12**, 6001 (2022).

36. Chruściński, D. & Pascazio, S. A brief history of the GKLS equation. *Open Syst. Inf. Dyn.* **24**, 1740001 (2017).

37. Frazzini, A., Israel, R. & Moskowitz, T. J. Trading costs. *SSRN Electron. J.* <https://doi.org/10.2139/ssrn.3229719> (2018).

38. Brin, S. & Page, L. The anatomy of a large-scale hypertextual web search engine. *Comput. Netw. ISDN Syst.* **30**, 107–117 (1998).

39. Farhi, E. & Gutmann, S. Quantum computation and decision trees. *Phys. Rev. A* **58**, 915 (1998).

40. Kempe, J. Quantum random walks: an introductory overview. *Contemp. Phys.* **44**, 307–327 (2003).

41. Childs, A. M. et al. Exponential algorithmic speedup by a quantum walk. In *Proc. 35th Annual ACM Symposium on Theory of Computing* 59–68 (Association for Computing Machinery, 2003).

42. Spohn, H. Algebraic conditions for the approach to equilibrium of an open  $n$ -level system. *Lett. Math. Phys.* **2**, 33–38 (1977).

43. NVIDIA Corporation. cuPyNumeric: a NumPy-compatible library for NVIDIA GPU calculations. <https://developer.nvidia.com/cupynumeric> (2024).

44. Almgren, R. & Chriss, N. Optimal execution of portfolio transactions. *J. Risk* **3**, 5–40 (2001).

45. Grinold, R. C. & Kahn, R. N. *Active Portfolio Management: a Quantitative Approach for Producing Superior Returns and Controlling Risk*, 2nd edn (McGraw-Hill, 1999).

46. Kissell, R. *The Science of Algorithmic Trading and Portfolio Management* (Academic Press, 2013).

47. Perold, A. F. The implementation shortfall: paper versus reality. *J. Portf. Manag.* **14**, 4–9 (1988).

48. Boyd, S. & Vandenberghe, L. *Convex Optimization* (Cambridge University Press, 2004).

010-MY3 (Applications of Quantum Computing in Financial Market Simulation and Optimization) and NSTC 114-2119-M-033-001 (Applications of Quantum Computing in Optimization and Finance) We also thank the NVAITC and the NVIDIA Quantum team for their continued technical guidance. Our gratitude further extends to the NVIDIA Academic Grant Program for providing the GPU resources essential to this research, and to the NVIDIA Strategic Researcher Engagement team for their steadfast support.

## Author contributions

Yen Jui Chang conceived the study, designed and executed the experiments, and wrote the manuscript. Wei-Ting Wang developed the Quantum Stochastic Walk software. Yun-Yuan Wang implemented GPU acceleration for the QSW framework. Chen-Yu Liu and Kuan-Cheng Chen analyzed and visualized the experimental results, contributing to their interpretation. Ching-Ray Chang discussed the results and reviewed the manuscript. All authors have read and approved the final version of the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s44335-025-00050-4>.

**Correspondence** and requests for materials should be addressed to Yen Jui Chang.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026

## Acknowledgements

We gratefully acknowledge support from the National Science and Technology Council (NSTC), Taiwan, under Grants NSTC 114-2112-M-033-