Research Square

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

Expressivity limits in variational quantum portfolio construction

José Requena Martínez

jrequena99@gmail.com

Universidad de Granada

José Manuel Zurita López

Universidad de Granada

Research Article

Keywords: Variational quantum algorithms, portfolio optimisation, dequantisation, benchmarking, quantum machine learning.

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11156291/v1

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

Additional Declarations: No competing interests reported.

Expressivity limits in variational quantum portfolio construction

José Requena Martínez[1]
José Manuel Zurita López[2]

Department of Computer Science and Artificial Intelligence, University of Granada, Granada, Spain.

Abstract

A recurring construction in quantum portfolio optimisation obtains weights by normalising the Pauli-Z expectation values $\langle Z_i \rangle$ of a variational circuit. We show that its standard one-layer ansatz admits the closed form $\langle Z_k \rangle = \prod_{i \le k} \cos \theta_i$, verified to 4×10^{-8} against simulators and complete backtests. The result generalises: in the static case, any per-qubit readout normalised into weights reparametrises the $(n-1)$-simplex and therefore admits no expressivity advantage.

It admits an expressivity deficit. Because the readout is a cumulative product, $|\langle Z_k \rangle|$ cannot increase along the chain, and reachability reduces to whether $n-1$ half-lines on a single scalar intersect: at nine assets the construction attains fewer than one portfolio in twenty thousand, and from twelve none we could sample in two million draws. Over forty training windows the circuit attains the matched classical optimum in every window where that optimum is reachable and in none where it is not ($p = 2.6 \times 10^{-7}$), so what separates the circuit from its control is reach rather than optimisation dynamics. Depolarising noise contracts $\langle Z \rangle$ by $(1 - 4p/3)^k$ and therefore degrades the portfolio toward the equal-weight portfolio, not toward randomness. The failure is invisible to the usual metrics: with the circuit 95 % of the way to equal weighting, its Sharpe ratio still exceeds the benchmark. Across five universes and 1,313 accounted configurations, none of the twenty-five quantum variants survives Holm–Bonferroni, and the largest statistic in the sweep, 3.12, falls short of the deflated threshold of 3.83. In a pre-registered sweep over eight universe sizes, the construction beats the equal-weight portfolio at all eight and the best classical competitor at none. We give the evaluation protocol that follows.

[1] Corresponding author: jrequena99@correo.ugr.es

[2] zurita@ugr.es

Keywords

Variational quantum algorithms, portfolio optimisation, dequantisation, benchmarking, quantum machine learning.

1 Introduction

Reports of quantum advantage in portfolio construction are accumulating faster than the tools needed to evaluate them. The consequence is that it has become difficult to tell a result that measures something real from one that measures an undeclared configuration choice, and that difficulty is not resolved by accumulating further comparisons but by examining what the proposed circuit actually computes. That is what this paper does: it takes one concrete construction, derives in closed form what it computes, determines which portfolios it can and cannot represent, and tests empirically whether the two routes that remain open — a better inductive bias, or easier training — ever materialise. From the whole there follows an evaluation protocol that generalises well beyond the case studied.

The construction we analyse works as follows. First, a parameterised quantum state is prepared by a variational circuit, that is, a circuit whose gates depend on a set of angles θ that a classical optimiser adjusts. Next, a single Pauli-Z expectation value is measured per qubit on that state, so that a register of n qubits yields n real numbers. Finally, that vector of n numbers is transformed, by a fixed classical function with no free parameters of its own, into the weights of a portfolio. The general pattern — parameterised circuit plus classical optimiser — is the one surveyed by Cerezo *et al.* [17]. Of its instances in asset allocation we have verified one in the published literature [1], and we rest nothing on there being more: the reason to examine the construction is not how often it has been repeated, but that the closed form derived below constrains every map of this shape, whoever writes it down. A growing line of work on quantum reinforcement learning and variational policies makes the circuit conditional on an observed market state, which changes the function class without changing the readout; Section 2.4 fixes how far the argument reaches in each case. Of the construction we ask three things: what that map computes, what determines its output, and what would have to be measured for a claim of advantage over it to mean anything.

It is worth fixing from the outset what lies outside. There is a second and currently dominant formulation that encodes asset selection as a combinatorial QUBO problem and solves it with VQE, QAOA or quantum annealing. That problem is of a different nature and the structural arguments of this paper do not apply to it; Section 2.4 draws the boundary precisely.

2 Background

2.1 The construction

Let $U(\theta)$ be a variational circuit on n qubits, one per asset, and let $\hat{B} = (Z_1, \dots, Z_n)$ be the list of the n Pauli-Z observables, one per qubit. The construction measures

$$z(\theta) \;=\; \langle 0|U^{\dagger}(\theta)\,\hat{B}\,U(\theta)|0\rangle \;\in\; [-1,1]^n. \qquad (1)$$

It is worth reading equation (1) piece by piece, because the whole paper rests on it. The vector θ collects the rotation angles of the circuit, which are the parameters the classical optimiser adjusts; $|0\rangle$ is the initial state of the n qubits, all set to zero; $U(\theta)$ is the unitary operation the circuit applies to that state, and $U^{\dagger}$ its adjoint. The full expression $\langle 0|U^{\dagger}\hat{B}\,U|0\rangle$ is the *expectation value* of each observable: the average one would obtain by measuring Z_k an infinite number of times on the prepared state. Each component z_k therefore lies in $[-1,1]$, with $+1$ and -1 the two deterministic outcomes and 0 the case in which both are equally likely. The result is a vector of n bounded real numbers, one per asset.

That vector is carried to weights by a fixed classical function with no parameters of its own. Two variants appear in the literature: the long-only map $w \propto (1 - z)/2$ followed by ℓ_1 normalisation — dividing by the sum so that the weights add to one — and direct ℓ_1 normalisation of z, which by preserving signs admits short positions. Both are, in the end, maps from n one-body observables into a set of dimension $n - 1$, and that drop in dimension is where the argument of Section 2.3 will come from.

It remains to make $U(\theta)$ concrete, and that is where the term *ansatz* enters. A variational circuit does not fix which gates it applies or in what order: that is decided by whoever designs it, and that choice — the concrete functional form proposed as a candidate before knowing whether it is the right one — is called the ansatz, from the German usage of the word as an *approach* or trial solution. Choosing the ansatz is choosing which family of states the circuit can reach, and it is therefore a modelling decision as consequential as the architecture of a neural network.

Throughout this paper, **the standard ansatz** designates a specific one: the one used by the audited construction and, with minor variations, by the line of work in which it appears. It has two parts. First, a layer of rotations $R_Y(\theta_i)$, one per qubit, applied to the initial state $|0\rangle$; these introduce the n free parameters, one per asset. Second, a *chain* of CNOT gates connecting each qubit to the next, $\mathrm{CNOT}(0,1)$, $\mathrm{CNOT}(1,2)$ and so on to the last; this is the part that generates entanglement and the only thing that distinguishes the circuit from n independent rotations. Section 3.1 derives exactly what this construction observes, and the rest of the paper is concerned with the consequences.

2.2 The classical benchmark that is usually missing

The whole mean-variance family, from Markowitz [7] onward, needs two ingredients estimated from historical data. The first is the vector of *expected returns* μ, whose component μ_i is the average return attributed to asset i. The second is the *covariance matrix* Σ, whose entry Σ_{ij} measures how far assets i and j move together: the diagonal holds the variance of each asset, that is its risk in isolation, and the off-diagonal terms record how much they cancel or reinforce one another, which is what makes diversification possible. With μ and Σ estimated, the method chooses the weights w that maximise return per unit of risk.

The problem is that neither is known: both are estimated with error, and the error in μ is notoriously large, because estimating an average return with any precision requires far more data than exist. DeMiguel, Garlappi and Uppal [6] established the consequence: the equal-weight portfolio $1/N$, which splits capital evenly and estimates nothing, is not a straw man. Out of sample it beats mean-variance optimisation across a wide range of datasets, precisely because estimation error outweighs the benefit of optimising against poorly known quantities. Any claim that a new method beats mean-variance must therefore also clear $1/N$, or it has cleared nothing.

Although the result of DeMiguel and coauthors has been known for over a decade, that benchmark is frequently absent from the literature at issue here, and its absence is not innocuous, because it is precisely the bar that separates a real improvement from an improvement over a method already known to perform poorly. In the instance of the construction we verified [1], for example, the comparators used are static mean-variance, Black–Litterman, DDPG and DQN, all of them methods that estimate parameters, and none of them the equal-weight portfolio, which does not appear at any point.

2.3 Why expressivity advantage is unavailable

A long-only, fully invested portfolio is a weight vector satisfying two conditions: no weight is negative, because short selling is not allowed, and all weights sum to one, because the whole capital is invested and no more. The set of all vectors meeting both is called the *simplex*:

$$\Delta^{n-1} = \left\{ w \in \mathbb{R}^n : w_i \geq 0,\ \sum_i w_i = 1 \right\}. \qquad (2)$$

Its geometry is what matters here. With $n = 2$ assets the simplex is a segment joining "everything in the first" to "everything in the second"; with $n = 3$ it is a triangle whose vertices are the three portfolios concentrated in a single asset and whose centre is the equal-weight one; in general it is an object with n vertices in which every interior point is a mixture. What is decisive is its **dimension**: although it is written with n coordinates, the constraint $\sum_i w_i = 1$ removes one degree of freedom, so the simplex has dimension $n - 1$.

Fix $n-1$ weights and the remaining one is determined. Put another way, $n-1$ real numbers parameterise the whole set: there is no long-only portfolio that needs more.

The argument follows immediately. The circuit produces n numbers, which a fixed function carries into that set of dimension $n-1$. Therefore, for any θ there exists a classical parameterisation with the same number of free parameters producing exactly the same portfolio, and no advantage can come from the circuit *reaching* portfolios that the classical one does not.

This is a statement about representability in one direction only: it bounds what the circuit can do from above, not from below. The converse fails, and Section 3.3 shows how badly: the map reaches the vertices of the simplex but almost none of its interior, so the containment is strict. What it also does not do is reach the portfolios it can represent by gradient descent from random initialisation within the budget employed.

2.4 The boundary of the argument

The simplex argument covers per-qubit readout to a continuous weight vector in the *static* case. It does not cover:

- **QUBO with binary asset selection.** The target is a discrete feasible set and the classical analogue is NP-hard. This is where a quantum advantage is at least conceivable, and where the extensive benchmark of Stopfer and Wagner [2] finds "only very limited room" for it: all 250 instances of up to 1,000 assets solve to proven optimality by mixed-integer programming in seconds, and a problem-tailored heuristic consistently outperforms the quantum approaches at fixed runtime. That benchmark is over solvers; the branch has since produced innovations at the level of the encoding itself, and Section 2.5 takes up two of them. The exclusion covers those as well, and for the same reason: their target is a discrete feasible set, not the simplex.

- **The conditional case.** A policy mapping state to weights is a different object: the argument covers its readout but not its function class. What does transfer is the empirical analysis of the readout — order dependence and the bias toward $1/N$ affect the output in every state, whatever policy produces it.

2.5 Related work

Our closest relative is Lozano [3], an operational decomposition audit of D-Wave's hybrid portfolio solver. It reports that the hybrid solver matches Gurobi's proven optimum on all 54 head-to-head instances at $N \leq 120$ while consuming 0.034 s of QPU access, 0.68% of its nominal five-second budget; that a matched-budget classical Tabu sampler reaches the same objectives; and — independently of us, in the other branch — that the resulting portfolios yield mean Sharpe 1.94 against 2.22 for $1/N$.

The two audits are complementary and their agreement is informative. Lozano asks *where the quantum computation is* in an annealing pipeline and answers operationally, by accounting for time. We ask *what the quantum circuit computes* in a variational pipeline and answer structurally, by deriving its closed form. Both branches of quantum portfolio construction have now been audited by independent groups using unrelated methods, and neither beats equal weighting.

Carrascal *et al*. [14] established the methodological precedent for backtesting classical and quantum portfolio algorithms under equivalent conditions, over 10,000 experiments on simulators and hardware. Their object is VQE on a binary formulation, so it sits in the branch this paper excludes; what we borrow is the insistence on equivalent conditions, which Section 5.3 extends from equal runtime to equal optimiser, initialisation and function-evaluation budget.

Two 2026 entries in the excluded branch are worth naming precisely, because both innovate at the level of the *encoding* rather than the solver, so the solver benchmark of Stopfer and Wagner [2] does not cover them. Scursulim *et al*. [4] replace penalty terms with a tensor product of parametrised Dicke states, one per asset class, so that the ansatz spans only the fixed-Hamming-weight subspace and the multiclass diversification constraints hold by construction with $\lambda = 0$. Their instances are 10 and 25 assets, simulated without noise, and the reference mixed-integer solver returns the proven optimum in 0.01–0.03 s; the paper compares ansaetze and optimisers, and claims nothing else. Soloviev and Krompiec [5] take the opposite route and compress: under Pauli correlation encoding each variable is read as $x_i = \mathrm{sgn}\langle \Pi_i \rangle$ for a k-local Pauli string, giving $m \leq 3\binom{n}{k}$ variables from n qubits, so that 250 assets fit in nine qubits at $k = 3$. Their quantum objective is a weighted cut on the market graph rather than a mean–variance Hamiltonian: the variational circuit recursively bipartitions the graph into correlated clusters, and the portfolio is then formed classically by taking the highest trailing-return asset from each cluster. The gain they report is measured against buy-and-hold over all m assets, a portfolio of different cardinality; $1/N$ at matched cardinality is not among their comparators, so the result is not evidence in either direction on the question this paper asks.

This paper belongs to the genre of *dequantisation*. The term names the work that takes a quantum algorithm credited with an advantage and shows that a classical algorithm achieves the same, so that the advantage disappears: not because the quantum algorithm is wrong, but because the problem never required a quantum computer in the first place. The founding case is Tang's [12], which dequantised recommendation systems by constructing a classical algorithm with the same asymptotic cost as the quantum one. Where this paper differs from the genre is in timing: dequantisation is usually *retrospective*, applied to published algorithms with a claimed advantage, whereas here it is *prospective*, applied to a construction that is spreading before anyone has established what it does.

The nearest methodological relatives outside portfolios are the line that asks when a classical model matches a quantum one on the same data [20] and the line that asks what

a parameterised circuit can express as a function of how it encodes its input [19]; Section 3.3 is that second question asked of the output map instead of the input map, where it turns out to have a sharper answer. We note in support of that reading that we were unable to identify a canonical source for the construction. The instance we verified presents the readout as its own contribution and cites no precedent, and the two works usually invoked in this connection address a different problem [13] or a binary formulation [14]. A construction with no citable origin and several independent appearances is being reinvented, not inherited.

2.6 Contributions

With the preceding elements now defined, the contributions of the paper can be stated without ambiguity. They are these.

1. **The standard ansatz has an elementary closed form.** A layer of $R_Y(\theta_i)$ on $|0\rangle$ followed by a CNOT chain yields $\langle Z_k \rangle = \prod_{i \le k} \cos \theta_i$. The fifteen-qubit circuit is `np.cumprod(np.cos(theta))`, agreeing with the simulator to machine precision and running two orders of magnitude faster.

2. **No per-qubit readout can have an expressivity advantage here.** The output is a point in an $(n-1)$-dimensional simplex, which $n-1$ real numbers parameterise completely. Advantage, if any, must be inductive bias or optimisation dynamics.

3. **The construction has an expressivity *deficit*, and it is what the measurements were detecting.** Because $|\langle Z_k \rangle|$ is a product of cosines it is non-increasing in k, so the readout is confined to an ordered cone. Reachability reduces to whether $n-1$ half-lines on a single scalar intersect: at $n = 9$ the construction attains fewer than one portfolio in twenty thousand, and from $n = 12$ none we could sample (Section 3.3). Over forty training windows the circuit attains the classical optimum in every window where that optimum is reachable and in none where it is not, with complete separation ($p = 2.6 \times 10^{-7}$).

4. **Neither remaining route survives measurement.** Under 126 optimiser configurations applied equally to every method compared, the circuit never attains the classical control's training optimum, and its poorer in-sample fit does not convert into better out-of-sample performance.

5. **Feasible-subspace ansaetze and compressed readouts do not compose.** In any fixed-Hamming-weight subspace every odd-locality X and Y Pauli correlator vanishes identically, so a constraint-preserving ansatz destroys two thirds of the encoding budget of a $k = 3$ Pauli correlation encoding (Section 7). The result is elementary, and it concerns the branch this paper otherwise excludes.

6. **Hardware noise biases the portfolio toward $1/N$, not toward randomness.** Derived, verified against a density-matrix simulator, and confirmed on complete

backtests in two universes. Its consequence is a failure mode invisible to the metrics currently reported.

7. **What determines the published result is a set of undeclared choices.** The alphabetical order of ticker symbols, the optimiser budget, the random seed and the optimiser each move the outcome by an amount comparable to or larger than the difference between the methods being compared — by an order of magnitude more in the survivorship-biased universe, and still more in one of the two clean ones.

8. **An evaluation protocol**, given in Section 6.5.

3 Structural results

This section establishes, without recourse to any experiment, what the construction computes and which portfolios it can produce. It proceeds in four linked steps, and it helps to keep them in view because each rests on the one before.

First the closed form is *derived*: propagating the observable backward through the circuit shows that the expectation value of qubit k is nothing but the product of the cosines of the first k angles. Second that formula is *verified* against simulators and against a complete backtest, to rule out its being an approximation. Third it is exploited: since a cumulative product of cosines cannot grow in magnitude, the readout is confined to a very small subset of the simplex, and that *reach* restriction turns out to be the central finding. Fourth the *two consequences* the formula has for the shape of the portfolios are read off: that they select contiguous blocks of the asset list, and that they flatten toward $1/N$ as the universe grows.

The reader in a hurry may keep the third step, which is the one the following sections depend on.

3.1 Derivation

To derive what the standard ansatz measures it is best not to compute the final state, which has 2^n amplitudes, but to run the circuit backward. In the Heisenberg picture the state is left alone and the observable is propagated backward through the gates, which is especially convenient here because CNOTs map products of Z to products of Z. Propagating Z_k backward through the chain $\mathrm{CNOT}(0,1), \mathrm{CNOT}(1,2)$, …, each gate adds the previous qubit to the observable's support, so that what began as Z_k ends as the product $Z_0 Z_1 \cdots Z_k$.

The rest is immediate. Before the chain only independent rotations have been applied, one per qubit, so the state at that point is a product state with no correlations between qubits. The expectation value of a product of observables on a product state factorises into the

product of the expectation values, and that of a single Z after a rotation $R_Y(\theta_i)$ equals $\cos\theta_i$. Multiplying the first k:

$$\langle Z_k \rangle = \prod_{i \le k} \cos\theta_i. \qquad (3)$$

3.2 Verification

Table 1 Verification of the closed form. The state-vector check runs over $n = 2 \ldots 15$ and the density-matrix check at $p = 0{,}0.02{,}0.1$; both backtests are complete runs over the sector universe

Against	Discrepancy
State-vector simulator	10^{-12}
Density-matrix simulator	10^{-10}
Backtest, weights	4.3×10^{-8}
Backtest, monthly returns	2.8×10^{-9}
Cosine control vs. circuit	11 sig. figures

A derivation can be correct on paper and still fail to describe what the code actually executes, so we checked it by five independent routes, which Table 1 collects. The first two compare the formula with simulators, without noise and with it; the next two compare it with the output of a complete backtest, both in the weights it produces and in the monthly returns derived from them. In all of them the discrepancy stays within floating-point rounding, between 10^{-12} and 10^{-8}.

The fifth check arrived unbidden, and is the most telling for exactly that reason. In an architecture sweep over a clean universe, a purely classical readout and a simulated circuit returned Sharpe ratios of 0.8034428993133284 and 0.8034428993071027, agreeing to eleven significant figures without anyone having arranged the comparison. The conclusion of all this is that equation (3) is not a convenient approximation but an identity: the circuit and the formula are the same object.

That identity has, moreover, a practical consequence worth recording, because it is the one contribution of this paper that directly serves anyone wishing to keep using the construction. Evaluating the formula costs 4.2 seconds on the same backtest in which simulating the circuit costs $1{,}951.9$, a factor of 459 computed on unrounded timings. Whoever uses this ansatz can therefore replace the simulator with three `numpy` operations and obtain exactly the same result two orders of magnitude faster.

3.3 The reachable set is a vanishing slice of the simplex

The argument of Section 2.3 bounded the circuit from above, by showing that any portfolio it produces is also produced by a classical map with the same parameters. What follows

bounds it from below, which is a different restriction and one the closed form reveals while a dimension count cannot see. We start from the fact that, taking absolute values in (3),

$$|\langle Z_k \rangle| \;=\; \prod_{i \le k} |\cos\theta_i|, \qquad (4)$$

and that each factor $|\cos\theta_i|$ necessarily lies in $[0,1]$. A product to which factors no larger than one keep being appended can only hold or shrink, never grow, so the sequence $|\langle Z_k \rangle|$ is *non-increasing in* k. That monotonicity, which looks like a minor observation, turns out to be the entire restriction: it forces the readout to be confined to the set

$$\mathcal{R} \;=\; \{\, z \in [-1,1]^n \,:\, |z_1| \ge |z_2| \ge \cdots \ge |z_n| \,\},$$

and every $z \in \mathcal{R}$ is attained, by $\cos\theta_1 = z_1$ and $\cos\theta_k = z_k / z_{k-1}$ for $k > 1$. The containment is exact, not a bound.

Under the weight map, write $u_i = (1 - z_i)/2 \in [0,1]$ and $c = \sum_i u_i$, so that $u = c\,w$. A portfolio w is reachable if and only if some $c \in (0,\, 1/\max_i w_i]$ makes $|c\,w_i - 1/2\,|$ non-increasing in i. Squaring consecutive pairs and cancelling $c > 0$ turns each pair into a half-line,

$$(w_i - w_{i+1})[\, c\,(w_i + w_{i+1}) - 1\,] \;\ge\; 0, \qquad (5)$$

that is, $c \ge (w_i + w_{i+1})^{-1}$ where $w_i > w_{i+1}$ and $c \le (w_i + w_{i+1})^{-1}$ where $w_i < w_{i+1}$. **The portfolio is reachable exactly when the $n - 1$ half-lines intersect.** One free scalar must satisfy $n - 1$ constraints, so the system is overdetermined and the feasible interval closes as n grows (Table 2).

Table 2 Fraction of the simplex the readout attains

n	Reachable fraction	95 % interval
3	5.41×10^{-1}	$[5.41,5.42] \times 10^{-1}$
4	1.95×10^{-1}	$[1.94,1.95] \times 10^{-1}$
5	5.18×10^{-2}	$[5.15,5.21] \times 10^{-2}$
6	1.08×10^{-2}	$[1.07,1.10] \times 10^{-2}$
7	1.86×10^{-3}	$[1.80,1.92] \times 10^{-3}$
8	2.94×10^{-4}	$[2.71,3.18] \times 10^{-4}$
9	$\mathbf{3.50 \times 10^{-5}}$	$[2.77,4.42] \times 10^{-5}$
10	3.5×10^{-6}	$[1.7,7.2] \times 10^{-6}$
12	0	$[0,\ 1.9 \times 10^{-6}]$
15	0	$[0,\ 1.9 \times 10^{-6}]$
32	0	$[0,\ 1.9 \times 10^{-6}]$

Two million uniform draws from Δ^{n-1} per size, evaluated with the exact criterion (5); intervals are Wilson, which remains meaningful at zero hits. An independent grid search

over c agrees draw by draw, not merely in aggregate, on every case tested. At the nine assets of the sector universe the construction attains under one portfolio in twenty thousand; each further asset costs another factor of five to ten; and from twelve upward we observed no reachable draw at all, which bounds the fraction below 2×10^{-6} rather than establishing that it is zero.

This does not weaken the simplex argument, it sharpens it in the other direction. Section 2.3 states that the circuit has no expressivity *advantage*: its output is a point in Δ^{n-1} and $n-1$ classical parameters reproduce it. Equation (4) adds that it carries an expressivity *deficit*: the matched classical control $z = \cos\theta$ places no ordering constraint on $|z|$ and attains the whole simplex, while the chain attains a vanishing slice of it. The two compose into a strict containment, and the CNOT chain — the only thing distinguishing the circuit from n independent rotations — is what removes the freedom.

3.4 Two structural consequences

Two consequences for the shape of the resulting portfolios follow from the closed form. Neither requires any experiment, and both are observed afterwards in the data.

The first is that portfolios organise into contiguous blocks. The reasoning starts from $\langle Z_k \rangle$ being a cumulative product, which imposes two behaviours at once. On one hand, the sign changes only when a negative cosine appears, so the chain is partitioned into runs within which all assets share a sign. On the other, and by virtue of (4), the magnitudes decay monotonically along each run, so the tail of a long run ends up as a set of near-identical values. Combining the two effects, the resulting portfolio is a union of contiguous blocks of the asset list, with near-uniform weight inside each block, and Figure 1 shows it as horizontal bands. In terms of what the construction actually does: it does not select assets by their risk or their return, but selects intervals of a list.

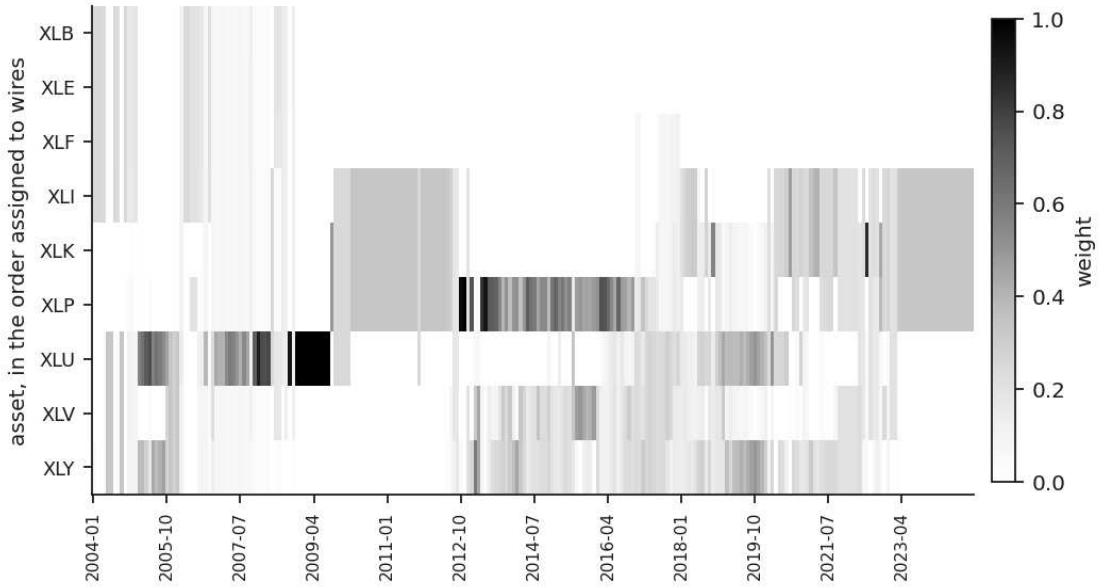


Fig. 1 Weights by rebalancing date, assets ordered as the data vendor returns them. The visible structure is horizontal bands: the map selects contiguous runs of the list rather than individual assets on their properties.

The second is a scaling law that fixes how the problem worsens as the universe grows. If the angles are initialised in the usual way, that is $\theta \sim \mathcal{N}(0, \sigma^2)$, the expectation of the cumulative product admits a closed expression:

$$\mathbb{E}[\langle Z_k \rangle] \;=\; e^{-k\sigma^2/2}. \qquad (6)$$

What that exponential says is that the tail of the chain flattens with position, and does so quickly. Consequently, the larger the universe, the more assets fall into that tail and the more indistinguishable their weights become, until the final stretch of the list effectively receives the same weight for all. The reading that matters is this: the bias toward the equal-weight portfolio is not an artefact of having worked with few assets, of the kind one might expect to dissipate at scale. It grows with n, so the construction drifts away from its promise precisely at the sizes where an advantage would be interesting.

4 Experimental protocol

4.1 Registration

The analysis protocol was closed before the main battery was run, and three further contrasts — the strong-VQA contrast, the universe-size sweep and the effort sweep — were specified before their experiments were executed. Each specification fixes three things in advance: the metric on which the result will be judged, the threshold it must clear, and the concrete prediction expected. Deviations from the plan are recorded in an addendum rather than by rewriting it.

One analysis falls outside that ordering and is labelled as such where it appears: the investigation of the order dependence in Section 5.2 was designed after the threshold had been declared, in response to earlier results, and is therefore post hoc. We flag it because the distinction between what was pre-specified and what was exploratory is exactly the one this paper demands of the literature it audits.

4.2 Design

Before listing the design it is worth fixing a term that recurs throughout what follows. We call each of the methods entering the comparison an **arm**, borrowing the word from clinical trials, where each arm is a group of patients receiving a different treatment under otherwise identical conditions. Here the arms are portfolio-construction strategies: the variational circuit is one arm, each parameter-matched classical control is another, and equal weight, minimum variance or risk parity are further arms. All see exactly the same assets, the same months and the same backtest engine, so any difference between their results is attributable to the method and not to the conditions. When we speak below of "126 configurations per arm" or of "twenty-five quantum arms", we are counting in that unit: one strategy, evaluated on equal terms with the rest.

- **Universes.** Five. Three are inherited from the instance of the construction we verified in the literature [1] — which we henceforth call *the audited configuration*, and from which we take universe, period and optimiser budget so as to reproduce its starting point before departing from it —: they have 5, 10 and 15 assets and carry survivorship bias, because their lists were chosen among securities still traded today. The other two are clean and constructed by us: nine sector SPDRs (2004–2024, 252 months) and 32 country and sector ETFs (2006–2024, 228 months).

- **Engine.** Monthly rebalancing, five-year rolling estimation window, geometric compounding within the month, transaction costs at 0 and 10 bp.

- **Primary metric.** The *aggregate Sharpe ratio*, declared in advance because it is the objective the method optimises explicitly. The Sharpe ratio measures return obtained per unit of risk taken: one takes the series of monthly portfolio returns, computes its mean $\bar{r}$ and its standard deviation σ — the dispersion of those returns around the mean, which is the usual measure of risk — and the quotient $\bar{r}/\sigma$ says how much return is obtained per unit of variability borne. Higher is better. We call it *aggregate* because it is computed once over the whole sample rather than by averaging yearly ratios, and it is annualised by multiplying by $\sqrt{12}$, the convention when passing from monthly data to yearly figures. One warning on units: the test statistics are computed on the un-annualised monthly quantity, so a Sharpe difference quoted in that context is $\sqrt{12} \approx 3.46$ times smaller than the same difference in aggregate units.

- **Tests.** Comparing two Sharpe ratios is not comparing two means: the statistic depends on four moments of two mutually correlated series, and the naive test

rejects far more often than it should. We therefore use the test of Ledoit and Wolf [8], designed specifically for Sharpe differences. Two corrections appear inside it. The first is a *HAC* variance, for *heteroskedasticity and autocorrelation consistent*: monthly returns neither have constant variance nor are independent from one month to the next, and a HAC variance estimates the dispersion of the statistic without assuming either. The second is a block *bootstrap*, which instead of resampling individual months — which would destroy temporal dependence — resamples contiguous stretches, thereby preserving the correlation structure over time; it is used in its circular, studentised variant. For confidence intervals we use the stationary bootstrap of Politis and Romano [9], which randomises the length of those stretches rather than fixing it. Finally, when several arms are tested at once the probability that *some* of them comes out significant by chance grows with their number, so we correct with Holm's procedure [10] within each family of comparisons declared in advance.

- **Search accounting.** If many configurations are tried on the same data, the best of them will look good even when none is, simply because the maximum of many random quantities is large. To *deflate* a statistic is to correct it for that effect: instead of asking whether the best observed result clears the threshold of a single isolated test, one asks whether it clears what would be expected of the *maximum* of as many tests as were actually run, under the hypothesis that none works. The required threshold rises, and rises the more one has searched. Applying it requires keeping count: every evaluation producing a return series comparable with $1/N$ is entered into a configuration ledger, now 1,313. The best statistic is then deflated in the manner of the deflated Sharpe ratio of Bailey and López de Prado [15], against the distribution of the maximum under the null, using the mean pairwise correlation estimated from the configurations' own monthly series — not assumed, because similar configurations give similar results and treating them as independent would over-deflate. A block-structured estimator that drops the equicorrelation assumption entirely is implemented alongside it. The need to deflate is not specific to quantum work: Harvey, Liu and Zhu [16] show that the accumulated search behind the published cross-section of returns raises the credible threshold to $t > 3$, and the ledger here plays the same role for a much smaller literature.

- **Reproducibility.** A suite of 200 automated tests over the analysis instruments, including an empirical size study of the Sharpe-difference test under the null hypothesis — that is, how often it rejects when there is no effect, which must be the nominal 5 %; all dependency versions pinned; and automatic regeneration of every figure and table from the analysis code, with no manual steps.

4.3 Sensitivity to our own free choices

The protocol above contains two numbers that no datum obliges us to choose, and which we therefore chose ourselves. The first of them is the *block length* of the bootstrap, that is, how many consecutive months are resampled together to preserve the temporal

dependence of the series, and its choice carries a trade-off at each end: if blocks are too short the dependence one meant to preserve is destroyed, whereas if they are too long so few distinct blocks remain that resampling stops generating variety. The second is the *bandwidth* of the HAC variance, that is, how many lags are taken into account when estimating autocorrelation, for which there exists a customary rule of $T^{1/3}$ that, while reasonable, remains a convention rather than a deduction.

The risk of such a parameter is evident: if the conclusion depended on having fixed it at a particular value, the result would be ours and not the data's. That is why we swept them rather than defending them. Running the block length from 1 to 24 months and the bandwidth from 1 to 12 lags, the test statistic moves between 1.086 and 1.208. It is a narrow range, well below the significance threshold at either end, so no choice within that margin changes a single conclusion of the paper.

5 Results

5.1 What determines the outcome

The question organising this section is simple to state: when a paper in this literature reports that its method beats mean-variance, how large is that advantage compared with the size of the decisions its author made without declaring them? If the latter move the result as much as the former, the reported advantage is not evidence about the method, because any other equally defensible configuration would have given a different figure.

Answering it requires a common yardstick, and we choose it so that it is the audited work's own: the difference in annual return between the equal-weight portfolio $1/N$ and mean-variance optimisation, which is exactly the comparison on which that literature builds its claims. Against that yardstick we measure the undeclared decisions, also expressed in annual return as the range they open when swept.

Now, only two of the four admit that direct measurement, and it is worth saying why before showing it. Asset ordering and optimiser budget can be swept in isolation: everything else is held fixed, one of them is varied, and the range of annual return that opens is attributable to it and nothing else. The other two do not isolate as cleanly. The random seed produces not a range in annual return but a dispersion across runs that reads better in Sharpe, and the optimiser appears in our sweep entangled with budget and initialisation, so its marginal effect is not separately identified. We return to those two further down, with the metric appropriate to each.

Table 3 therefore collects the first two, and does so in all three universes rather than in one, because **the comparison does not survive intact outside the universe that carries survivorship bias**.

Table 3 Sources of variation in annual return (percentage points), by universe. Ordering is the range over 200 *random relabellings; budget is the range over the five budgets swept.*

Source	15-asset (biased)	Sector (clean)	32-asset (clean)
Asset ordering	**27.3**	4.55	3.31
Optimiser budget	**14.8**	1.32	3.29
$\lvert 1/N -$ mean-variance$\rvert$	3.45	1.94	**6.98**

The reading of the table changes with the universe, and that variation is itself the result. In the survivorship-biased universe, the one inherited from the audited configuration, the two undeclared decisions dominate the measured effect by close to an order of magnitude: 27.3 and 14.8 points against 3.45. In the sector universe, already clean, the decisions' lead shrinks considerably but remains larger than the effect. In the 32-asset clean universe, by contrast, the relation inverts, because the method difference reaches 6.98 points and comfortably doubles either source of variation.

We report that inversion rather than keeping the two thirds that favour us, and it is worth being explicit about what stands after it. The strong formulation — that undeclared decisions always weigh more than the effect being measured — does not survive all three universes, and we therefore do not maintain it. The one that does survive is weaker and still sufficient for the argument: in all three cases these decisions move the outcome by an amount *comparable to or larger than* the measured effect, and in none of the three is the decision declared. That is enough for a figure reported without its configuration to be uninterpretable, even if it is not always directly misleading.

The other two decisions remain, and as anticipated they are measured otherwise. The random seed is best seen as dispersion: across five seeds of a single architecture the aggregate Sharpe opens a range of 0.1201, and in two of the three universes that range contains the value of $1/N$, so that not merely the magnitude but the very *sign* of the reported advantage is decided by which number was used to initialise. The optimiser, for its part, is reported in Section 5.3 and in training cost rather than annual return, for the reason given above: the effort sweep varies it jointly with budget and initialisation, so isolating its contribution would require a factorial design we did not run.

It is worth adding, finally, that order dependence shows not only in performance but directly in the portfolios constructed. In the sector universe, relisting exactly the same assets in a different order changes the resulting portfolio by a mean ℓ_1 distance of 0.7032, against a maximum possible of $2(1 - 1/9) = 1.78$ that would correspond to concentrating all capital in a single different asset. That is, 40 % of the maximum traversal, just from reordering the list (Figs. 2 and 3).

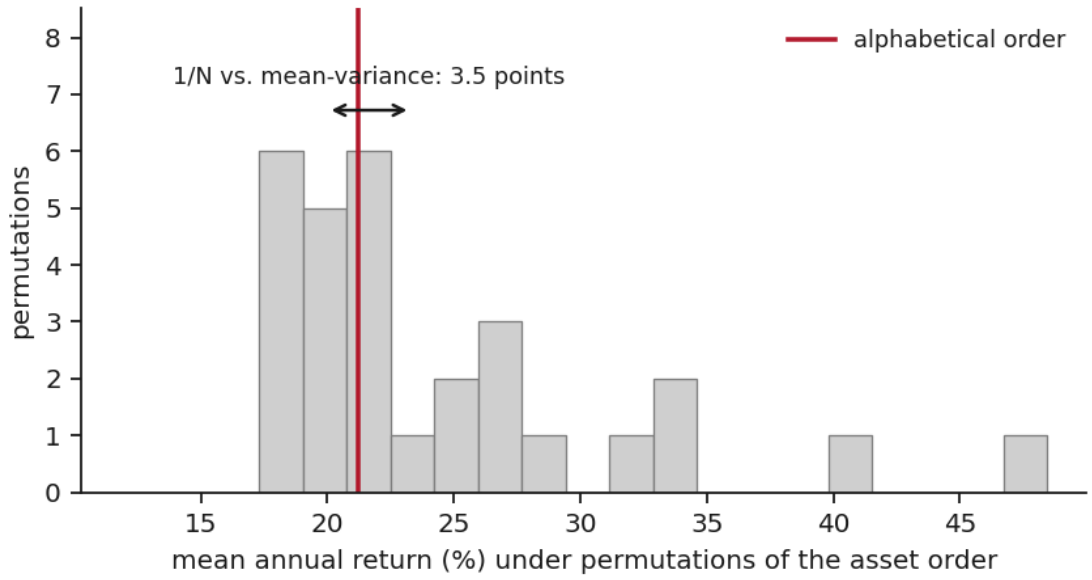


Fig. 2 Distribution of annual return under random permutations of the asset order, with the alphabetical ordering marked and, for scale, the span that separates the methods being compared.

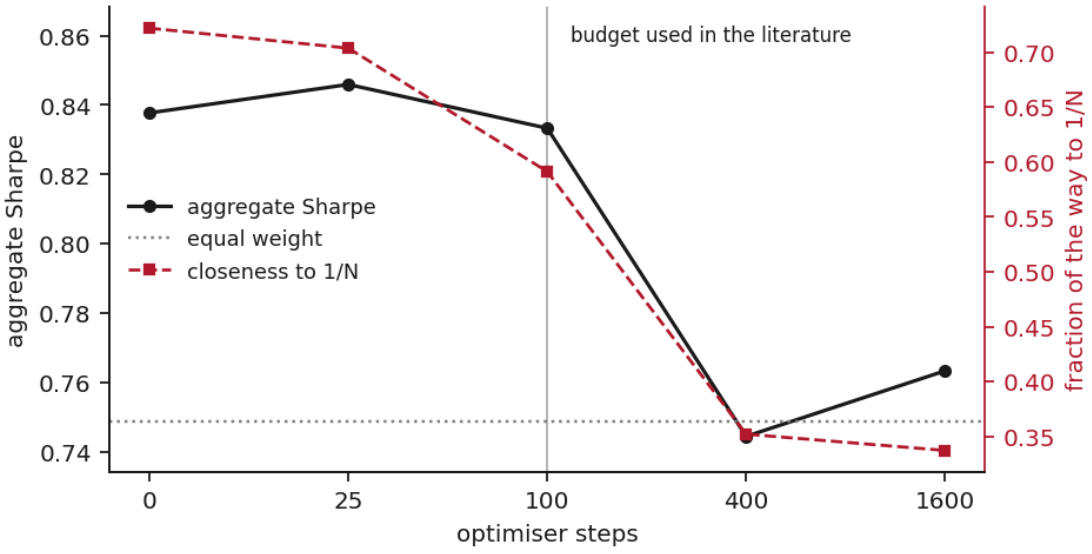


Fig. 3 The optimiser budget moves the portfolio continuously from near-$1/N$ to concentrated. The budget used in the literature is one arbitrary point on that continuum, not an optimum.

5.2 The mechanism behind the order dependence, and what we could not find

The alphabetical ordering is not a typical ordering. Table 4 places it against 200 random relabellings of the same assets.

Table 4 Aggregate Sharpe of the alphabetical ordering against 200 *random relabellings of the same assets, by universe*

Universe	Alphab.	Mean	s.d.	Pctile.
sector (9)	0.8333	0.7791	0.0494	0.86
32-asset	0.6382	0.5135	0.0257	$\mathbf{1.00}$
survivorship (15)	1.2526	1.2504	0.0962	0.52

In the 32-asset universe the alphabetical ordering beats *all two hundred* random orderings, at 4.85 standard deviations above their mean. That is not arbitrariness. It is an effect of a size the rest of this paper would call decisive, and we report it because we could not explain it.

Two hypotheses were tested and both failed. The first was that ticker symbols encode asset class — iShares country ETFs begin `EW`, SPDR sector funds begin `XL` — so alphabetical ordering groups by class without anyone deciding it, and the map selects contiguous blocks. Prediction: grouped $\geq$ alphabetical $>$ random $>$ interleaved. Observed: grouped 0.5504, alphabetical 0.6382, interleaved 0.5235, random mean 0.5214. Both halves fail. The contrast is clean because the two constructed orders are *equally* grouped — three contiguous class blocks each — and differ only in where a three-ticker block sits.

The second was therefore positional. Sweeping that block along the chain gives Spearman $+0.564$ ($p = 0.071$) over eleven positions, with the best position mid-chain rather than at the alphabetical one and three positions beating it. Moving three tickers along the chain disperses the Sharpe ratio by s.d. 0.0462 over eleven positions, against 0.0257 for permuting all thirty-two over two hundred. We give standard deviations rather than ranges because the two designs have different numbers of draws and the range is not comparable across them. Relocating three of thirty-two tickers thus moves the result more per draw than reshuffling the whole list, which is consistent with a positional effect but does not identify one: eleven single-seed points with no error bars cannot carry that weight.

What we can state is therefore weaker than a mechanism and stronger than arbitrariness. There is a large, reproducible effect of asset ordering; in one of three universes the alphabetical ordering sits at the extreme of its own null distribution; and neither of the two structural explanations we could formulate survives contact with the data. The consequence for the literature is unchanged and arguably sharper: a result reported at the alphabetical ordering is reported at a draw that, in the largest universe we tested, no random relabelling matched. Whether that is a mechanism nobody has described or a coincidence in a single universe, it is not a property of the circuit, and it is not declared.

5.3 Optimisation dynamics, with the control treated identically

The structural argument leaves two routes open: better inductive bias, or easier training. Both were tested with the rule that every advantage given to the circuit is given identically to its matched classical controls, and with the budget counted in **function evaluations**

rather than steps, because central finite differences cost $2n + 1$ evaluations per step while SPSA costs two [11]. 126 configurations per arm; 378 in total; 4,536 optimisations.

The circuit arm is evaluated through the closed form (3) rather than a simulator. Table 1 licenses the substitution — the two agree to eleven significant figures — and it is what makes a sweep of this size affordable. One consequence must be stated: the closed form is noiseless, so no finite-sampling noise enters any arm. This is the condition most favourable to the circuit, and it is the condition under which the results below are obtained.

Not one of the 126 circuit configurations attains the best control (Fig. 4).

Table 5 Best training cost by budget (lower is better). †: the audited configuration uses approximately this budget

Budget	Cosine	Tanh	Circuit
200	-0.9298	-0.9458	-0.7963
2,000†	$\mathbf{-0.9495}$	-0.9486	$\mathbf{-0.9010}$
20,000	-0.9495	-0.9494	-0.9397

Table 6 Best training cost by optimiser and by initialisation

	Cosine	Tanh	Circuit
Adam	-0.9495	-0.9494	-0.9397
Nelder–Mead	-0.9495	-0.9468	-0.9251
SPSA	-0.9494	-0.8730	$\mathbf{-0.7954}$
`randn`	-0.9495	-0.9494	-0.9397
small ($\sigma = 0.1$)	-0.9495	-0.9494	-0.8815
warm start at $1/N$	-0.9495	-0.9494	-0.9251

The two tables answer different questions, and it helps to read them separately. Table 5 asks whether the problem is that the circuit lacks budget, and the answer is no. The classical controls reach their best cost already at 2,000 evaluations and do not improve thereafter, a sign that they have converged; the circuit, by contrast, keeps improving when its budget is multiplied by ten, up to 20,000, and even so ends below the point where the controls have been sitting from the start. Giving it more compute does not close the gap: it only makes the gap more visible.

Table 6 asks something else, namely whether the result depends on how one searches. Its six rows are not a cross of conditions but six marginals: the first three give, for each optimiser, the best cost that arm achieves with it whatever the rest of the configuration, and the last three do the same for each initialisation. Read that way, the contrast is immediate. The cosine control sits at -0.9495 in five of the six marginals and at -0.9494 in the remaining one, a difference in the fourth decimal: however it gets there, it ends in the

same place, which indicates that the optimum is easy to find from any starting point. The circuit does not behave that way, but spreads its six marginals between -0.9397 and -0.7954. Put another way, choosing one optimiser or another moves its result by more than the whole distance separating the best circuit from the best control.

Within that spread, the circuit degrades especially under SPSA, and a warm start placed at the benchmark itself does not help it either. The first has a tempting reading, since SPSA is the optimiser typically used on physical hardware, but we stop short of drawing a hardware conclusion from that row: our SPSA runs see a noiseless objective, whereas on a real device its stochastic gradient estimate would compound with shot noise, and there is no reason for the two cases to order the same way.

The warm start deserves a caveat of the same kind, and the closed form supplies it. Setting every $\theta_i = \pi/2$ puts all $\langle Z_k \rangle$ at zero, which is exactly $1/N$ after the weight map, so in portfolio space it is the most generous starting point available. In parameter space it is the least: with $\cos\theta_1 \approx 0$ the cumulative product is $O(\varepsilon^k)$ and so is every partial derivative, giving a flat region of the same character as the barren plateaus of McClean *et al.* [18], though arising here from the ansatz's algebraic form rather than from Haar-random concentration. Identical treatment of the arms in portfolio space is therefore *not* identical treatment in optimisation terms: the same start is well-conditioned for the cosine control, whose gradient there is maximal, and degenerate for the circuit. We report the row as measured and do not lean on it; the reach result of Section 3.3 does not depend on it.

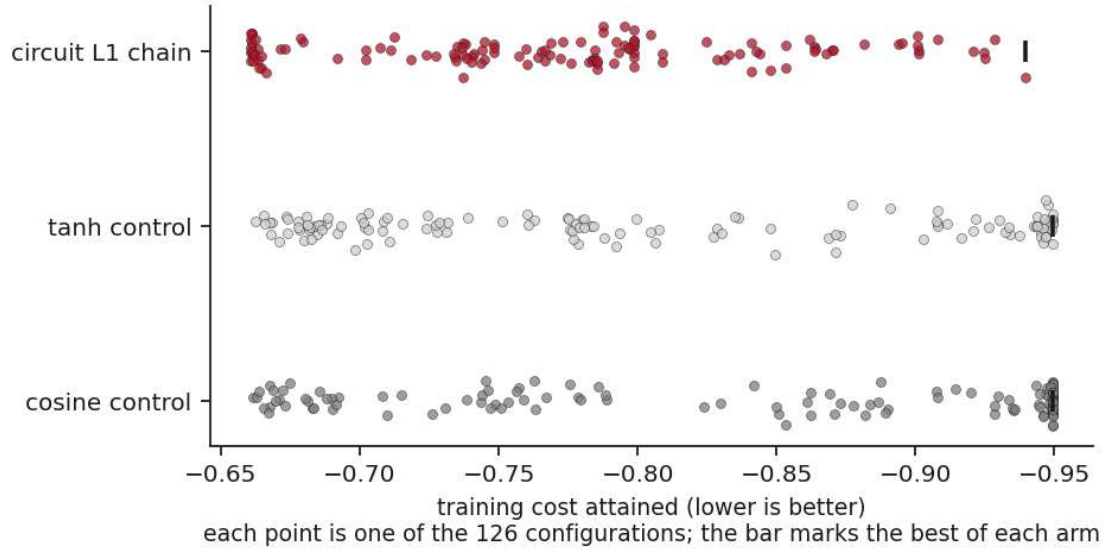


Fig. 4 Training cost attained by each of the 126 configurations per arm; the bar marks the best of each. Both classical controls concentrate against the optimum and reach it from every optimiser and initialisation. The circuit is dispersed and its best falls short of theirs.

Carrying the best configuration of each arm — selected by *training* cost, never by out-of-sample result — to the full backtest over three seeds gives mean Sharpe 0.8473 (s.d. 0.0001) for the tanh control, 0.8472 (s.d. 0.0000) for the cosine control, 0.8264

(s.d. 0.0780) for the circuit, against 0.7489 for equal weight. Three seeds estimate a standard deviation poorly, and we lean on these only for the ordering of magnitudes, which spans three orders. Zero of six paired comparisons favour the circuit under Holm. Well optimised, the classical controls are deterministic; the circuit remains a lottery, with a mean below both.

5.4 Depth buys nothing, and a single-seed table said otherwise

This subsection deserves telling in the order it happened, because the episode illustrates the failure mode this paper describes better than any argument. We swept the circuit's architectures along two axes, number of layers and type of entangler, with one seed per cell, and the outcome was that the best circuit reached a Sharpe of 0.8883 against 0.8491 for the best classical control. Presented that way, with a single run per cell and no standard error attached, that pair of figures reads inevitably as an advantage for the circuit, and so it would have entered a paper that asked no further questions.

Repeating with five seeds inverts the reading completely. It turns out that the 0.8883 was not the circuit's behaviour but its best seed of five, and that the worst of those five falls to 0.7683, below all ten control runs. The paired contrast confirms it from several angles at once: none of the ten comparisons survives Holm's correction, the smallest observed p is 0.084, the circuit wins in sign only six times out of ten, and its mean advantage over the controls — $+0.0033$ in monthly Sharpe, which annualised to the units of the rest of this section is $+0.0115$ — turns out to be ten times smaller than the 0.1201 range its own seed opens. On top of that, its seed dispersion is twice that of either control, so not only is the advantage small but the circuit is also the least stable arm of the three.

It remains to explain the structure the sweep seemed to show, a peak at two layers inviting the conclusion that an optimal depth exists. It does not survive either, and the comparison that shows it is direct: the range opened by the nine architectures measured with one seed is 0.111, while the range opened by a *single* architecture measured across five seeds is 0.120. The whole visible structure of the sweep therefore fits inside the noise of a single cell. There was no depth optimum to find; there was initialisation noise arranged in the shape of a table.

5.5 Hardware noise degrades toward $1/N$

We insert a single-qubit depolarising channel on each wire after its rotation, in the convention $\rho \mapsto (1-p)\rho + p/3\,(X\rho X + Y\rho Y + Z\rho Z)$, which contracts any single-qubit Pauli expectation by $1 - 4p/3$ and is fully mixing at $p = 3/4$. The contraction accumulated along the chain gives

$$\langle Z_k \rangle \;=\; (1 - 4p/3)^k \prod_{i \le k} \cos\theta_i, \qquad (7)$$

verified against the density-matrix simulator to 10^{-10}. Placing two-qubit channels on the CNOTs instead changes the constant — $1 - 16p/15$ for a full-weight Pauli — but not the

geometry: the contraction remains uniform across wires and its exponent remains the chain position, which is all the argument below uses. Because the contraction is uniform across wires, and because the weight map sends a uniformly contracted z to the uniform portfolio, **noise does not degrade the portfolio toward randomness: it degrades it toward** $1/N$.

Table 7 Depolarising noise, sector universe, 252 months

p	Fraction to $1/N$	Sharpe	Turnover
0	0.591	0.8333	0.128
0.005	0.608	0.8352	0.126
0.01	0.624	0.8285	0.118
0.02	0.654	0.8166	0.108
0.05	0.719	0.8123	0.085
0.10	0.793	0.8161	0.047
0.20	0.872	0.8002	0.020
0.40	**0.954**	**0.7688**	**0.003**
Equal weight	1.000	0.7489	0.000

All eight noise levels are shown. Two of the three declared predictions hold strictly (Table 7, Fig. 5): the fraction toward $1/N$ rises monotonically across all eight and turnover falls monotonically across all eight. The third, that the Sharpe gap to equal weight closes, holds between the endpoints — from 0.084 to 0.020 — but not monotonically: the Sharpe ratio rises at $p = 0.005$ and again at $p = 0.10$. The non-monotonicity is of the order of the seed noise floor established in Section 5.3 and we do not interpret it; it is shown rather than omitted because the two quantities that do move monotonically are the ones the argument uses.

The noise levels beyond $p \approx 0.05$ are far above current gate error rates and are included to trace the limit, not to model a device. The mechanism does not require them: at $p = 0.05$ the portfolio has already travelled 72% of the way to equal weight, and because the contraction enters as λ^k the same displacement is reached at proportionally smaller p as the universe grows.

The consequence is the size of the degradation, not its existence. At $p = 0.40$ the portfolio has travelled 95.4% of the way to equal weight and its turnover is effectively zero — the circuit has stopped selecting — and its Sharpe has fallen only 7.7%, remaining *above* the benchmark. An experiment at that noise level would report a working quantum portfolio optimiser with low turnover that beats its benchmark. All three statements would be true. The portfolio would be $1/N$.

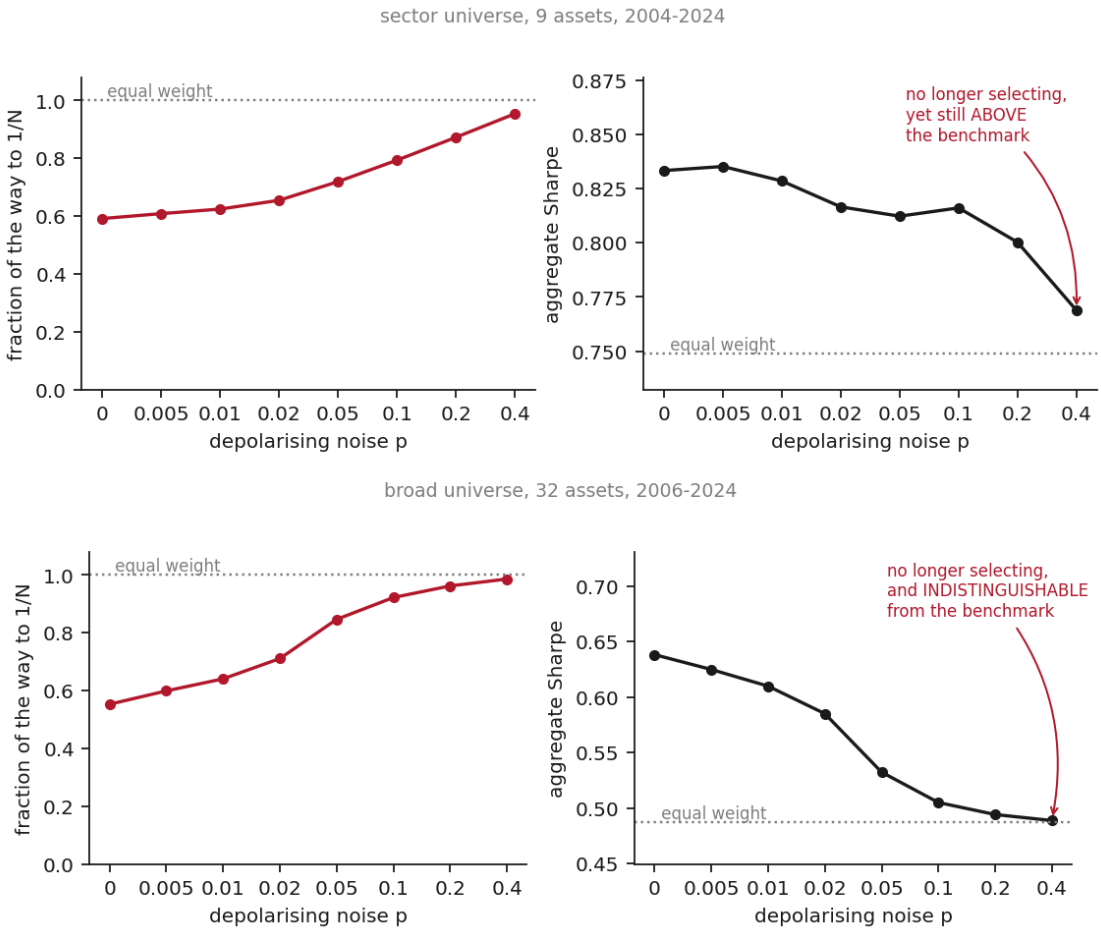


Fig. 5 Depolarising noise, sector universe (top) and 32-asset universe (bottom). Left: the portfolio travels monotonically toward equal weight. Right: the Sharpe ratio barely notices. At the highest noise level the circuit has stopped selecting, turnover 0.003*, and in the sector universe it still sits above the benchmark. The metric that is reported does not see the failure.*

In the 32-asset universe the convergence completes: the Sharpe gap closes from 0.151 to 0.002. The difference between universes is predicted, and it is size, not universe: the contraction accumulated by $\langle Z_k \rangle$ is λ^k, so at $p = 0.05$ the tail is 0.537 with nine assets and 0.110 with thirty-two. This is the scaling law (6) confirmed by an independent route — decay of the noise rather than decay of the cosines — and the two causes compose.

5.6 Universe size, pre-registered

Nested subsets of a single clean universe, eight sizes, five draws each, assets alphabetical within each subset so that size is not confounded with order.

“Best cl.” in Table 8 is the maximum over the four classical arms *within each draw*, averaged over draws. That is a selected statistic and by the standard of Section 4 it should be treated as one: it favours the classical side by construction, and the circuit receives no

corresponding maximisation. The selection is worth at most 0.03 here, because minimum variance wins on its own in 4 of 5 draws below $n = 16$ and 5 of 5 from $n = 16$ upward; against minimum variance alone the eight gaps read -0.088, -0.162, -0.182, -0.249, -0.245, -0.309, -0.294, -0.281, which changes no sign and no conclusion. A pre-registered replication should declare minimum variance as the single comparator and report the maximum as secondary.

Table 8 Universe size sweep (aggregate Sharpe)

n	Frac. $1/N$	Circuit	Best cl.	$1/N$	Circ.−cl.
4	0.475	0.649	0.744	0.534	-0.095
8	0.619	0.587	0.782	0.519	-0.194
12	0.640	0.586	0.786	0.516	-0.200
16	0.615	0.548	0.797	0.500	-0.249
20	0.638	0.550	0.795	0.479	-0.245
24	0.685	0.506	0.815	0.467	-0.309
28	0.676	0.536	0.830	0.479	-0.294
32	0.728	0.548	0.829	0.487	$\mathbf{-0.281}$

The three predictions we had declared before running the sweep hold, and it is worth stating each alongside the datum that confirms it. The first was that the portfolio would approach the equal-weight one as the universe grew, and indeed the fraction of the way to $1/N$ rises monotonically with n (Spearman $+0.857$, $p = 0.0065$), just as the scaling law (6) demands. The second was that the advantage would narrow with scale, and that too is observed, with a correlation of -0.881 ($p = 0.0039$) between size and advantage. The third was that the advantage would in no case exceed the noise floor, and it holds in the most emphatic way possible, since the advantage never becomes positive at any size. The result is summarised in two lines:

> **The construction beats equal weighting at 8 of 8 sizes.**
> **It beats the best classical competitor at 0 of 8.**

The winning classical method is the same at every size: minimum variance, in five of five draws from $n = 16$ upward (Fig. 6).

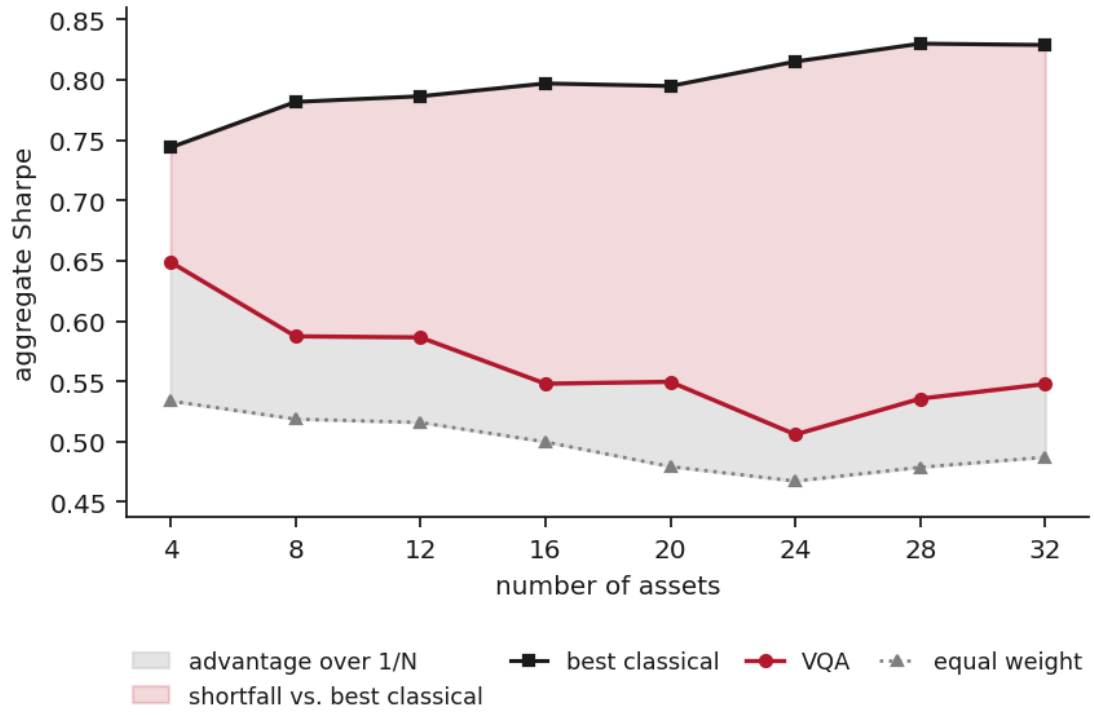


Fig. 6 Nested subsets of a single clean universe. The construction sits between the benchmark it beats and the classical competitor it does not, at every size, and the gap to the competitor widens with n.

5.7 The corrected verdict

Gathering every comparison under the multiplicity correction, the tally is as follows. Of the 45 arms evaluated at each of the two transaction-cost levels, exactly one clears the Holm–Bonferroni threshold on the primary metric declared in advance, and that survivor is classical: risk parity on the sector universe, with $p = 0.0027$ against the 0.00556 threshold it faces. **None of the twenty-five quantum arms survives at either cost level**, and the smallest p any of them reaches is 0.0080, nearly twice the threshold it would need to beat.

It is worth stating how the correction was grouped, because its severity depends on it. Each Holm family is the nine arms of one universe at one cost level, and that grouping was declared before looking at results. The ten families so formed are not further corrected against each other, which in principle would be a laxity; in practice it is not, because the deflation of the next paragraph accounts for the full search and therefore absorbs that higher-order multiplicity.

One secondary result must be reported alongside this, because omitting it would be the omission this paper is about. On the difference in *mean monthly return* — a different quantity from the declared Sharpe metric, and not the one the method optimises — the 32-asset quantum arm at seed 42 does survive Holm, at both cost levels ($t = 2.96$, $p = 0.0035$ at 0 bp; $t = 2.86$, $p = 0.0047$ at 10 bp), as does one seed on a 2016–2024 subwindow of

the 15-asset universe. It changes no conclusion, for the reason the next paragraph gives, but it is the strongest result any quantum arm obtains anywhere in this study and the reader is entitled to it.

Accounting for the full search — 1,313 configurations at mean pairwise correlation 0.272 — the deflated threshold is $z = 3.83$, and the largest statistic observed anywhere is 3.1202. The $t = 2.96$ of the previous paragraph is smaller still. Two summaries of the same search should not be confused: the 5% critical value of the correlated maximum is $z = 3.83$, whereas the *median* of that maximum matches the median of 200 independent trials, a number we quote as an effective count and not as a threshold. The threshold reported here uses the equicorrelation approximation at the estimated $\bar{\rho}$. A block-structured estimator that drops the approximation is implemented alongside it; as argued under Limitations, relaxing equicorrelation raises the threshold rather than lowering it, and nothing survives the lower one.

An obvious objection remains to be met: one might think the construction fails because we evaluated it in its naive version, and that a careful implementation would obtain a different result. To answer it we specified in advance a “strong” version correcting the four weaknesses the analysis itself had identified, averaging over 64 permutations of the order to neutralise the dependence on the list, quadrupling the optimiser budget, replacing the sample covariance with the Ledoit–Wolf estimator, and selecting among restarts by training cost rather than by outcome. With those four improvements accumulated, and contrasted against a threshold fixed before running it, it reached $z = +0.837$. Not only does it fall far short of the deflated threshold of 3.310, it does not even clear the nominal 1.645 that would apply to a single isolated test. And it does so while consuming 486 times the computational cost of a classical mean-variance.

6 Analysis of results

6.1 What the construction is doing

The results assemble into a single mechanism. The map computes a cumulative product of cosines, which decays exponentially along the chain, so the weights at the end of the asset list become indistinguishable. The portfolio is therefore a union of contiguous blocks of the list, and which blocks it selects is determined by where assets happen to sit in that list. Since the list is alphabetical because that is what the data vendor returns, the selection is determined by an orthographic accident.

On top of that mechanism act two independent forces pushing the same way, and the fact that they are independent is what makes them hard to escape. The first is the cosine decay, which flattens the tail of the chain as a function of position and operates even if the hardware is perfect. The second is physical noise, which contracts every expectation value by a factor $(1 - 4p/3)^k$ whose exponent is again the position in the chain. Since both carry the portfolio toward $1/N$ and both worsen as the universe grows, their effects do not offset

but accumulate. The size sweep confirms that composite prediction directly, and with a conclusion that inverts the usual intuition: large universes, where an advantage would be most valuable, do not give the circuit room but take it away.

That is why tuning does not rescue it. Section 5.3 gave the circuit three optimisers, ten times the budget, an exact warm start at the benchmark and restarts, and gave the same to a classical control with the same parameter count. The control converged to the same optimum from everywhere; the circuit did not converge to it from anywhere.

Section 3.3 says why, and the answer is not the one we expected. Over forty rolling five-year training windows of the sector universe — the same window construction and moment estimator as the effort sweep — we classified each by whether the control's optimum lies in the circuit's reachable set, and compared the training cost each arm attains (Table 9).

Table 9 Training-cost gap against reachability of the control's optimum, over forty five-year windows of the sector universe

Control optimum	Gap $= 0$	Gap > 0
reachable	6	0
not reachable	0	34

The separation is complete (Fisher exact, $p = 2.6 \times 10^{-7}$). Where the optimum lies inside the reachable set the circuit attains it exactly — to 2×10^{-16}, with the same optimiser, budget and restarts as the control; where it lies outside, the circuit never attains it, by a median of 0.0030 and up to 0.053. **What separates the arms is reach, not conditioning.** No window needed the tenfold restart budget we held in reserve for the case where a reachable optimum went unfound, so we can put no part of the gap down to search.

Six of forty is far above the 3.75×10^{-5} of Table 2, and the reason is instructive: the feasible interval (5) has width zero in all six. A Sharpe optimum is frequently a corner of the simplex, corners are reachable, and the reachable optima here are exactly those degenerate cases. The circuit matches the control only where the answer happens to be a corner.

That reading makes a prediction. Table 2 puts the reachable fraction at 32 assets below 2×10^{-6}, so in the 32-asset universe the contingency table should collapse to a single column. It does: across twenty-nine windows not one control optimum is reachable, every window shows a positive gap, and the median gap is 0.215 against 0.0024 in the sector universe. Ninety times the shortfall on three and a half times the assets, with the reachable set as the only thing that changed.

6.2 Why the reported advantage is real and uninformative

The preceding results admit a reading that is worth making explicit, because otherwise they read as a plain refutation and they are not one. The construction *does* beat the equal-weight portfolio, and does so at every size tested; what it never does, at any of them, is

beat the best classical competitor. Those two facts are compatible, and they explain why the literature reports advantages in good faith.

Indeed, a study that compares itself against mean-variance and its variants but does not include $1/N$ among its comparators will find an advantage, and that advantage will be real in its own terms: its figures are not fabricated nor its comparisons rigged. The problem is of another kind, namely that the advantage does not constitute evidence about the circuit. A classical map with the same degrees of freedom obtains it just as well, and a minimum-variance portfolio proposed in 1952 [7] obtains considerably more. What the comparison measures, therefore, is not the contribution of the quantum component but the well-known difficulty mean-variance has in performing out of sample.

6.3 The failure mode that current metrics cannot see

Of everything above, Section 5.5 carries the most serious practical consequence, because it describes a way for a hardware experiment to be at once well executed, honestly reported and wrong. The mechanism is as follows. Noise does not degrade this construction toward randomness, as one would expect and as would be easy to detect, but toward a portfolio that happens to be *good*. Consequently the Sharpe ratio — which is precisely what the literature reports — stays high while the circuit has stopped contributing anything, and no alarm goes off.

What closes the trap is that the two quantities capable of revealing it, namely the fraction of the distance travelled toward $1/N$ and the turnover actually realised, are reported by none of the works in this line that we examined. An experimenter acting in good faith can thus publish an entirely correct table in which the three claims — that turnover is low, that the Sharpe ratio beats the benchmark, and that the method works — are each true separately, while the portfolio they describe is the equal-weight one.

Within its scope this converts an argument about methodology into a checkable prediction: as hardware noise falls, portfolios produced by this construction should move measurably away from $1/N$. If they do not, the circuit was not contributing at the higher noise level either.

6.4 The instability of the sign is itself a result

One declared hypothesis did not resolve. It held that proximity to $1/N$ arises from stopping the optimiser early, so that a larger budget should push the portfolio away from $1/N$ *and* worsen out-of-sample performance. The first half holds everywhere. The second holds in one universe and fails in two, and survivorship bias does not explain the pattern, because one of the two failures is in a clean universe.

We report this as unresolved rather than resolving it by choice of universe. But the instability carries information of its own: **a study reporting one universe is reporting the sign of a coin flip.** The same applies to the noise control, which reproduces the circuit's

Sharpe in two universes and not in the third, and to the seed, whose range contains the benchmark in two of three universes.

Any single-universe, single-seed, single-budget result in this literature is therefore uninterpretable, independently of whether its authors did anything wrong. That is a stronger claim than "this method does not work", and it follows from measurements we would rather have come out cleanly.

6.5 An evaluation protocol

The following would have detected every effect reported here, and none of it requires quantum hardware.

1. **Benchmark against** $1/N$, not only against mean-variance.
2. **Report the fraction of the distance travelled toward** $1/N$ alongside the Sharpe ratio. A method converging on the benchmark can outperform it.
3. **Report realised turnover**, measured against drifted weights rather than target weights, since a portfolio that has stopped selecting has near-zero turnover.
4. **Report the order-invariance error**: relist the assets and measure the ℓ_1 change in weights. For the construction studied here it is 40% of the theoretical maximum.
5. **Report dispersion across seeds**, not a single run.
6. **Match the classical control** on parameter count, optimiser, initialisation and compute budget, with the budget counted in function evaluations rather than iterations.
7. **Account for the configurations evaluated** and deflate the best statistic accordingly [15, 16].

The seven points neither cost the same nor weigh the same, and it is worth saying so. The first five cost nothing at all: they are figures the researcher already has computed and need only print, so omitting them is a decision and not a limitation. The sixth does demand work, because building a classical control matched in parameters, optimiser, initialisation and budget takes as long as building the quantum arm; but it is exactly the difference between presenting a comparison and presenting a demonstration. The seventh demands keeping count of everything that was tried, which is inconvenient precisely when one has searched the most, and it is the difference between reporting a result and reporting a maximum.

6.6 The protocol is not about portfolios

Nothing in Section 6.5 is specific to asset allocation. The object it audits is a *per-qubit readout normalised into a decision vector*, and that object recurs wherever a variational circuit is asked to produce a distribution, an allocation or a ranking rather than a scalar: portfolio weights here, but equally mixing coefficients, resource splits, or the policy head

of a variational reinforcement learner. Wherever it appears, four questions have answers obtainable before any hardware is involved. Does the readout admit a closed form? Does the map into the decision set restrict which decisions are representable, as Section 3.3 finds here? Does a classical control with the same parameter count and optimiser reach what the circuit reaches? And does the reported metric move when the circuit stops contributing?

The last question generalises furthest. A construction that degrades toward a *good* default under noise — $1/N$ here, a uniform distribution elsewhere — keeps reporting a respectable headline number while the quantum component does nothing, and only a distance-to-default statistic and a realised-activity statistic will show it. That failure mode is a property of the readout, not of finance. The expressivity question is likewise general: any readout whose components are constrained to a fixed order, or to a monotone envelope, carries the same kind of deficit that Section 3.3 quantifies, and the deficit is invisible to a dimension count of the kind that Section 2.3 performs.

6.7 Limitations

- **Scope.** The construction analysed is not the dominant formulation. The QUBO branch is untouched and the simplex argument does not apply to it.
- **The conditional case.** The argument covers the readout, not the function class of a state-to-weights policy, which is the form the live literature takes.
- **Noise model.** Depolarising noise is treated analytically and in a density-matrix simulator, but finite-sampling (shot) noise is not modelled, and no experiment was run on physical hardware.
- **Data.** Five universes, all exchange-traded equity instruments from a single vendor snapshot. Three are free of survivorship bias; two are not, and are retained and labelled rather than discarded.
- **Deflation.** The threshold extrapolates the observed correlation structure to the full ledger. Families not represented in the observed block are likely less correlated, which would raise the threshold further; nothing survives the lower one.
- **Engine.** Turnover is measured against target weights in the headline tables; the drift-corrected measurement changes turnover materially but changes the Sharpe ratio by at most 0.0018, twenty-five times less than the seed noise floor.
- **Reachability.** Table 2 is a Monte Carlo estimate under the uniform measure on the simplex. It bounds how much of the simplex the readout attains, not how far the nearest attainable portfolio lies from a given target, which is what determines the loss in any particular problem; the window-by-window contrast in Section 6.1 measures that consequence directly, on two universes of nine and thirty-two assets.

- **Selected comparators.** The classical side of the size sweep is a per-draw maximum over four methods and is therefore an optimistically selected statistic; the gap against a single pre-declared comparator is smaller but unchanged in sign at all eight sizes.

- **Order dependence.** The mechanism is unidentified. The alphabetical ordering is extreme in one universe of three, and we cannot distinguish an undescribed structural effect from a coincidence with the evidence available.

7 A constraint on encoding design, outside the scope above

This result does not follow from the closed form. It concerns the branch excluded in Section 2.4, and we record it because it is elementary and because the two innovations of Section 2.5 otherwise invite being combined.

One line of work enforces feasibility in the state: a parametrised Dicke ansatz spans only the fixed-Hamming-weight subspace W_k, so cardinality constraints hold by construction and no penalty term is needed [4]. Another compresses the readout: Pauli correlation encoding recovers $m \leq 3\binom{n}{k}$ binary variables from n qubits as $x_i = \operatorname{sgn}\langle \Pi_i \rangle$, with Π_i a k-local Pauli string [5]. The two address the same obstacle from opposite sides, and applying them together is the obvious next construction. It does not work.

Let Π_S be a Pauli string of X or Y letters on a support S, and let $|\psi\rangle = \sum_x c_x |x\rangle \in W_k$. Since Π_S acts on the computational basis as the bit flip $|x\rangle \mapsto |x \oplus 1_S\rangle$ up to site-wise phases,

$$\langle \Pi_S \rangle \;=\; \sum_x \omega_x \; c^*_{x \oplus 1_S} \, c_x, \qquad |\omega_x| = 1, \qquad (8)$$

and a term survives only if x and $x \oplus 1_S$ both lie in W_k. Since $|x \oplus 1_S| = |x| + |S| - 2|x \cap S|$, that requires $|S| = 2|x \cap S|$. For $|S|$ odd no such x exists, so $\langle \Pi_S \rangle = 0$ identically — for every state in the subspace and every choice of parameters.

A feasibility-preserving ansatz therefore annihilates the odd-locality X and Y correlator families entire. At the locality $k = 3$ used in [5] for 250 assets those are two of the three families, so two thirds of the encoding budget decodes to $\operatorname{sgn}(0)$ and the construction loses the compression that motivated it. The argument is not specific to Dicke states: it holds for any ansatz preserving Hamming weight, including XY mixers and particle-number-conserving circuits. Enforcing feasibility in the state and compressing the readout are not independent design choices, and only even locality escapes.

8 Conclusions and next steps

8.1 Conclusions

The variational construction studied here computes a cumulative product of cosines. It is a reparametrisation of the simplex, and no per-qubit readout to a normalised weight vector can be otherwise, so an expressivity advantage is unavailable by construction rather than absent by measurement. The containment is moreover strict in the other direction: the cumulative product forces $|\langle Z_k \rangle|$ to decrease along the chain, and the portfolios compatible with that constraint are a vanishing fraction of the simplex from about nine assets upward.

That deficit is what the optimisation experiments were measuring. Under 126 matched optimiser configurations per arm the circuit never reaches the training optimum that a classical control with the same parameters and compute reaches from every starting point; window by window, it reaches that optimum whenever the optimum lies inside its reachable set and never when it does not. The limitation is representational, not dynamical, and its poorer fit does not convert into better generalisation. Under a pre-registered sweep over universe size, it beats equal weighting at every size and the best classical competitor at none. Of twenty-five quantum arms across five universes and two cost levels, none survives multiple-comparison correction on the metric declared in advance; the single survivor in the entire study is a classical risk parity heuristic.

What determines the outcome of this construction is not the circuit. It is the alphabetical order of ticker symbols, the optimiser budget, the random seed and the optimiser, none of which is declared in the literature and each of which moves the result by an amount comparable to or larger than the difference between the methods being compared. The order dependence we could measure but not explain: two candidate mechanisms were tested and both failed, while in the largest universe the alphabetical ordering outperformed all two hundred random relabellings of the same assets, 4.85 standard deviations above their mean. We report that as an open question rather than resolving it, and note that either answer — an undescribed mechanism or a coincidence — leaves the reported result resting on an undeclared choice.

The finding with the widest consequence is that hardware noise degrades this construction toward the equal-weight portfolio rather than toward randomness. A noisy experiment can therefore report low turnover, a Sharpe ratio above benchmark and a working method, and be describing $1/N$.

8.2 Next steps

The conditional case. The live instances of this construction are policies mapping market state to weights. Our impossibility argument covers their readout but not their function class, and closing that gap — or measuring how much of the order dependence and $1/N$

bias survives in a conditional policy — is the single extension that would change what can be claimed.

Shot noise. The depolarising channel we have modelled is only one of the mechanisms by which a real device departs from the ideal, and finite sampling is another quite different one: since each expectation value is estimated by averaging a limited number of measurements, it carries a variance of $(1 - \langle Z \rangle^2)/N$ per estimate. The relevant question is whether that second channel pushes the portfolio in the same direction as the first or in another, and answering it needs no new method, because it suffices to apply that of Section 5.5 to a simulator with finite shots, at negligible computational cost.

Alternative readouts. Correlator readouts $\langle Z_i Z_j \rangle$ produce $n(n-1)/2$ numbers from n qubits and therefore escape the premise of the simplex argument as stated, though not its conclusion, since the target set is unchanged. Premise and conclusion part company only when the target set changes with them, as in Pauli correlation encoding [5], where the correlators are thresholded by sgn into a discrete assignment instead of being normalised into weights. That construction lies outside this paper in both directions; Section 7 records the one structural constraint we can state about it.

Funding

The authors declare that no funds, grants or other support were received during the preparation of this manuscript.

References

1. V. Gurgul, “Variational quantum circuit-based reinforcement learning for dynamic portfolio optimization,” *Preprint*, arXiv:2601.18811, 2026.

2. E. Stopfer and F. Wagner, “Quantum portfolio optimization: an extensive benchmark,” *Preprint*, arXiv:2509.17876, 2025 (rev. 2026).

3. L. Lozano, “Where the quantum lives in D-Wave hybrid portfolio optimization: an operational decomposition audit,” *Preprint*, arXiv:2605.17623, 2026.

4. J. V. S. Scursulim, G. M. Langeloh, V. L. Beltran, and S. Brito, “Multiclass portfolio optimization via variational quantum eigensolver with Dicke state ansatz,” *Sci. Rep.*, vol. 16, art. 6208, 2026.

5. V. P. Soloviev and M. Krompiec, “Large-scale portfolio optimization using Pauli correlation encoding,” *Sci. Rep.*, vol. 16, art. 25158, 2026.

6. V. DeMiguel, L. Garlappi, and R. Uppal, “Optimal versus naive diversification: how inefficient is the $1/N$ portfolio strategy?,” *Rev. Financ. Stud.*, vol. 22, no. 5, pp. 1915–1953, 2009.

7. H. Markowitz, “Portfolio selection,” *J. Finance*, vol. 7, no. 1, pp. 77–91, 1952.

8. O. Ledoit and M. Wolf, “Robust performance hypothesis testing with the Sharpe ratio,” *J. Empirical Finance*, vol. 15, no. 5, pp. 850–859, 2008.

9. D. N. Politis and J. P. Romano, “The stationary bootstrap,” *J. Amer. Statist. Assoc.*, vol. 89, no. 428, pp. 1303–1313, 1994.

10. S. Holm, “A simple sequentially rejective multiple test procedure,” *Scand. J. Statist.*, vol. 6, no. 2, pp. 65–70, 1979.

11. J. C. Spall, “Multivariate stochastic approximation using a simultaneous perturbation gradient approximation,” *IEEE Trans. Autom. Control*, vol. 37, no. 3, pp. 332–341, 1992.

12. E. Tang, “A quantum-inspired classical algorithm for recommendation systems,” in *Proc. 51st ACM Symp. Theory of Computing*, 2019, pp. 217–228.

13. S. Woerner and D. J. Egger, “Quantum risk analysis,” *npj Quantum Information*, vol. 5, art. 15, 2019.

14. G. Carrascal, P. Hernamperez, G. Botella, and A. del Barrio, “Backtesting quantum computing algorithms for portfolio optimization,” *IEEE Trans. Quantum Eng.*, 2023.

15. D. H. Bailey and M. López de Prado, “The deflated Sharpe ratio: correcting for selection bias, backtest overfitting and non-normality,” *J. Portfolio Management*, vol. 40, no. 5, pp. 94–107, 2014.

16. C. R. Harvey, Y. Liu, and H. Zhu, “… and the cross-section of expected returns,” *Rev. Financ. Stud.*, vol. 29, no. 1, pp. 5–68, 2016.

17. M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, “Variational quantum algorithms,” *Nature Reviews Physics*, vol. 3, pp. 625–644, 2021.

18. J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature Communications*, vol. 9, art. 4812, 2018.

19. M. Schuld, R. Sweke, and J. J. Meyer, “Effect of data encoding on the expressive power of variational quantum-machine-learning models,” *Phys. Rev. A*, vol. 103, no. 3, art. 032430, 2021.

20. H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, “Power of data in quantum machine learning,” *Nature Communications*, vol. 12, art. 2631, 2021.