Which quantum method? It depends on what you know going in.

We measured how much probability QAOA puts on the **proven optimum** — not whether it eventually finds it. The answer is not a method. **It's a decision tree.**

~12,400
certified runs
every one scored against a proven optimum

24
instances
20 synthetic + 4 real portfolio books, 7–16 qubits

2
null results
printed, on purpose

THE FRAME

Everyone reaches 0%. So measure concentration.

With a certified referee in the loop, **every reasonable QAOA method eventually finds the optimum** on prototype-size problems. A 0% gap alone doesn’t rank methods.

What differs is **concentration**: how much probability the trained circuit puts on the answer, whether it’s the most likely state — and what that means in shots. Exact statevector, matched budgets, pre-specified design with amendments recorded.

Not “did it find it” — “how hard does it point at it”.

FINDINGS 1–2 · THE PRIOR

Warm start: perfect when the seed is right. Locked when it’s wrong.

Right seed: 100%

The proven optimum is the most likely portfolio every time — every method, synthetic and real.

Wrong seed: 0 of 96 escapes

On the four real books the circuit never escaped a wrong classical seed — even at **7× the training budget**. The lock sits in the starting state itself, not the mixer (we swapped it — nothing changed; null result, printed).

And the seed was wrong — on every real book

The classical heuristic’s answer: **3 / 3 / 2 assets off** the proven optimum, a certified gap of **0.65–1.60%** that looks perfect in a spreadsheet. 8,000 extra classical starts didn’t fix it.

A prior that fails quietly.

FINDING 3 · THE SECOND OPINION

The adaptive circuit escapes what the warm start can’t.

Standard warm start, wrong real seeds: 0 of 96

Any budget up to 20,000 evaluations. Probability on the optimum: 0.4%.

Adaptive, deep & well-fed: 3 of 3

The real 16-qubit wrong seed, cap 12, 10,000 evaluations — escaped every run, **13.7%** of probability on the optimum. Starve it of depth or budget and it fails; it never wins on cost.

Independence, not economy.

FINDING 4 · THE READOUT TRAP

A tallest grain is not a peak.

Unseeded at 16 qubits (65,536 portfolios), the trained circuit put **0.008%** of its probability on the optimum — about **5× a blind coin flip**. You'd need ~50,000 shots to see it once with 99% confidence; by then you have nearly read every portfolio anyway.

The trap: on **5 of 6 runs** that optimum was still the *single most likely portfolio*. A report that prints only "most likely state = optimum ✓" calls this a win. The ten most likely portfolios together held 0.08% — a flat readout, and the optimum happens to be the tallest grain of dust.

Our reports print probability mass and shots-to-99% — not the flag.

THE ANSWER

What you know going in decides the method.

TRUST THE SEED

Warm start

Optimum on top **100%** of the time; tighten the blur for more concentration.

DOUBT THE SEED

Adaptive, deep & well-fed

Cap 12, hundreds of iterations per layer — **escapes what the prior can’t**.

KNOW NOTHING

Measure first

At scale it won’t concentrate on its own — **the referee tells you**.

And before any of it: **soften the penalty.** Across 20 instances the single biggest concentration lever is the penalty-to-risk ratio — correlation **−0.91** — not the ansatz, not the objective.

The referee doesn’t care which method wins. That’s the point.

QUBIT-LAB.CH · STUDY

The 18-page study is public.

Pre-specified design, amendments and **both null results printed**. Every switch it tested — objective, ansatz, classical seed with blur, penalty weights — is a switch in our QAOA-based Rapid Quantum Prototypes, and every report prints concentration next to the certified gap.

Read the study · qubit-lab.ch/evidence **qubit-lab.ch/evidence**

The full case story behind this study — three rounds against a proof on a real loan book — is told in the study itself. Sample reports and demo books in the Resource Room, access on request. qubit-lab.ch — Evidence, not promises.