

We told a quantum computer to cherry-pick a CHF 60m loan book.

Then we made it prove itself.

45

facilities

committed book, bundles,
exclusions

8.4m

possible baskets

23 acquisition decisions

1

certified referee

an exact solver that
proves the optimum

What honest quantum benchmarking looks like – in nine slides.

THE USE CASE

A bank exits. You pick which loans to take.

- **37 facilities offered** as one package, on top of your committed book of 8.
- **Connected clients transfer whole or not at all:** 8 groups, decided as single yes/no calls.
- **The board's capital envelope** covers ~62% of what you would like to buy. A third stays with the seller.
- **Sector caps** apply to the combined book: max 25% commercial real estate.

Cherry-picking a loan book is a correlation problem, not a ranking problem.

THE REFEREE

Every answer faces a proof.

Before comparing anything, an exact solver certifies **the** optimal basket, with a mathematical proof of optimality.

23

binary decisions

15 single facilities + 8
client groups

8,388,608

baskets in scope

every combination counts

146 s

to certified proof

SCIP branch & bound

No verdict without a certificate.

ROUND 1

Everyone misses.

- The **classical heuristic** lands **+0.19%** above the proven optimum. In a spreadsheet it would look perfect.
- **QAOA with default settings** lands **+1.07%** above, with a nearly flat sampling distribution.
- Only the certificate reveals either gap.

Nobody's answer looks wrong until you can prove that it is.

ROUND 2

Tune the physics, not the story.

- The budget penalty dominated the energy landscape. **Softening it:** gap 1.07% → **0.20%**.
- More depth, sharper CVaR tails, new seeds: **plateau at 0.18%**.
- Two independently trained circuits converged on the **same runner-up basket**.

The plateau was a measurement problem in disguise.

THE DIAGNOSIS

So we computed the truth.

Rebuilding the trained circuit exactly and asking: what amplitude does the **certified optimum** actually carry?

7×

above uniform

the circuit had lifted the optimum, not suppressed it

#21,126

probability rank

of 8,388,608 states

8.3e-7

exact amplitude

verified against the engine to 13 digits

The needle was lifted. The readout just wasn't looking long enough.

ROUND 3 — THE HONEST FIX

8 million shots found it. Then we did the arithmetic.

- A finite readout of **8,000,000 shots** returned the certified optimum — 0.00%.
- The self-audit: at that budget, a **uniform random sampler** finds it **61% of the time**. Eight million shots on 8.4 million baskets is enumeration, not evidence.
- So we fixed the **circuit**, not the readout: **warm-started CVaR training** — classically seeded, disclosed.

4,900×

above uniform

the optimum's probability,
retrained

~8,000

shots for a 99% hit

under a second of QPU
time

+0.19%

the heuristic's miss

found what classical
search missed

We printed our own overclaim — then beat it.

THE FULL LADDER · ALL CERTIFIED

From 1.07% to zero, step by measured step.

QAOA, default settings



+1.07%

Classical heuristic, same run



+0.19%

QAOA, tuned landscape



+0.20%

QAOA, warm-started CVaR (seeded, disclosed)



0.00%

The zero row is sampled at **thousands of shots, not millions** — concentration did the work. And clearly labeled in the report: the loss-tail panel (VaR, expected shortfall) is **classical** Monte Carlo.

That is what evidence-based quantum exploration looks like.

TEST IT ON YOUR OWN BOOK

Credit Allocation RQP Pro

- Facilities, connected-client bundles, sector caps, capital budget: **one Excel workbook.**
- Certified exact referee on **every run**. Credit metrics a committee reads: EL, UL, RAROC, loss tail.
- Full acquisition report, workbook, and tuning data in the client resource room.

The referee model for quantum in finance:

no wins without certificates, no claims without numbers.

qubit-lab.ch/rqp