qubit-lab.ch

Quantum computing for financial services: why, and why now.
Which of your problems could quantum methods matter for, measured on your own data, in weeks.

qubit-lab.ch

In short. In physics, quantum advantage has been demonstrated and posted for anyone to refute. In finance it is an open question, use case by use case. The institution that measures its own cases now owns the yardstick for every claim that follows.

Why

1. **A different method, not a faster computer.** Quantum methods are an alternative to classical ones with scaling of their own. For Monte Carlo pricing and risk, a provable quadratic speedup in sampling. For optimization, possibly new routes into problems with hard limits on how many positions you may hold. For classification, a feature space whose dimension doubles with every added qubit (2^n for n qubits), built differently from any classical kernel; whether your labels separate in it is the measurable question, and the classical twin stays in the room. Valuable now as a second opinion, powerful once the crossover is reached.
2. **The problem shapes are yours.** Constrained allocation with pairwise effects, estimation by sampling, classification with hidden structure: that is the computational core of portfolio construction, pricing, risk and fraud. Quantum methods are built for exactly these shapes and for little else; everything linear stays classical.
3. **The upside is estimated, unmeasured for you, and whoever benchmarks decides.** McKinsey puts the value at stake in finance at USD 400 to 600 billion by 2035; at the sizes a referee can certify today the result is parity, and which of your problems crosses the line first is measurable. Every result you will be shown, ours included, had its benchmark chosen by whoever published it; the only one that counts is your problem against the strongest classical method, with a referee allowed to say "no".

Why now

1. **Your peers are already on the map.** 67 published quantum initiatives at 34 financial institutions, 31 of them from 2025 and 2026, 18 in 2026 alone: Intesa Sanpaolo, JPMorgan, HSBC, Wells Fargo, Erste, BBVA and UBS (2024) among them. Real projects at prototype scale, each graded by its own publisher, and since July 2026 joined by hardware demonstrations in physics posted for public refutation. The question your board will ask is not whether, but where you are on the map (position paper p. 4, full matrix on qubit-lab.ch).
2. **Your board is likely to see a quantum-risk roadmap in 2027.** Of 60 Swiss institutions FINMA surveyed, 8% had a roadmap and about 70% no PQC measures, and recommends a roadmap by mid-2027. The board will then ask what the same technology could do for the firm; whoever answers needs a measured number, not a vendor slide: a first verdict before the mid-2027 roadmap.
3. **The entry is lean, the yardstick is strict.** A half-day briefing, five hands-on sessions and a few weeks per use case on your own data: sized at CHF 40,000 to 80,000 all in, under 100,000 in any configuration. The natural owner is the AI or data team, and the size fits a team-level innovation budget. Whether it pays off for your firm is known in weeks, not years; what the weeks leave behind is a certified benchmark on your own constraints, rerun against every new device and every vendor pitch.

Our record so far, published against ourselves. Allocation modelled on a private-bank mandate (demo case): 100% of the certified optimum. S&P 500 with 20 names: certified tracking error matched. Pricing: error falling as 1/N instead of 1/sqrt(N), the quadratic speedup, at simulation scale. Classification: advantage shown on a dataset built to have the structure (1.45x); the one-minute pre-screen has said no on six market datasets so far; bring yours to the first call. At certifiable sizes parity is the expected result; the benchmark measures how the gap moves as your instance grows. qubit-lab.ch/evidence

The entrance, in three steps

1 · Quantum Awareness Briefing	90 minutes to half a day for leadership, risk and technology.
2 · Quantum Capability Kickstart	Five half-day sessions, typically within one week; the team leaves with its own three to four candidate use cases.
3 · RQP usage	Monthly team access to the Rapid Quantum Prototypes: workbook in, certified report out, circuits exported as Qiskit, PennyLane, Cirq, CUDA-Q, Classiq or OpenQASM.

"It was the first time I fully grasped how quantum computing can be applied to real-world problems." Simone, Risk Quant, Kickstart pilot October 2026

Your data. Workbooks masked, processed in the Zurich region, run artifacts expire after 24 hours; a synthetic book and your own in-house rerun where positions cannot leave the bank. NDA on request.

Evidence, not promises.

Two ways in: a 30-minute call to pick a first use case, or guest access to the suite on a sample workbook for a first look at your own pace (personal code by mail, subject "Guest access"). Daniel Hug, founder: ETH physicist, MBA, formerly UBS, Credit Suisse, SIX. · daniel.hug@qubit-lab.ch

Book a call

Write to us

qubit-lab.ch

Quantum computing for financial services: four questions we always get.
Too early? Buy hardware? Lock-in? Physicists? Short answers, and what each means for the first step.

qubit-lab.ch

Four questions come up in almost every first conversation about quantum computing with a financial institution. The short answers are no, no, yes and no. The longer answers are below, and each ends with what it means for the first step.

No.

Isn't it too early? Are there working quantum computers at all? Should we wait?

Too early to deploy, not too early to test. Quantum is an alternative approach next to the classical ones, not a replacement, and how much it can contribute today and in the near future depends on the problem and the data, not on the calendar. Working machines exist and are reachable: cloud devices with more than a hundred qubits from several vendors, and since July 2026 advantage demonstrations in physics, one of them on error-corrected logical qubits, posted for public refutation. What does not exist is a machine that beats classical methods on a finance workload; on today's hardware, our own noise-ladder study found, circuits deeper than three layers returned noise, not answers. Hardware improves on the vendors' roadmaps whether you wait or not. Your experience, your radar on the field, and your ability to answer a board, a regulator or a client with evidence instead of opinion do not.

In practice. Build experience now on exact simulators, at a scale a certified referee can grade. Keep the radar out and the workbooks ready for next year's rerun. Every stakeholder question then gets a measured answer.

No.

Do we need to buy a quantum computer?

No bank needs to own one. Simulators on ordinary servers reproduce quantum circuits exactly up to a few dozen qubits, which is where method comparisons with a certified classical referee are possible today. Real devices are reached through the cloud, pay per use, from several vendors. Our own published hardware rung ran 100,000 shots on an IBM processor for 44 seconds of processor time, within a standard monthly allowance. The machine is not the cost; the question is.

In practice. Start on simulators, with your own data. Move a result to cloud hardware only once the simulator says it is worth measuring there.

Yes.

Is vendor lock-in a real threat?

Yes, and more so than in classical IT, because the hardware is not settled. Superconducting circuits, trapped ions, neutral atoms, photonics, annealers and spin qubits are different machines with different strengths: gate speed, connectivity, fidelity, qubit count, and how far each is from error correction. Nobody knows which will carry the first finance workload. An institution that writes its logic for one vendor's stack has bet on an answer. One that keeps its problems, its benchmarks and its circuits portable has not.

In practice. Stay open, challenge every capability claim, and test it on your own problems. Insist on exports in the common frameworks (Qiskit, PennyLane, Cirq, CUDA-Q, Classiq, OpenQASM), so the same circuit can be rerun on the next vendor's device.

No.

Do we need a vendor's physicists to start?

No. External specialists first have to understand your problems, your constraints and your data, and that transfer is slow and expensive in both directions, and when they leave, the knowledge leaves with them. It is easier to train some of your own people on quantum than to train outsiders on your business. A data scientist with the linear algebra an AI team uses every day, vectors, matrices, a little complex arithmetic, can learn the basics within a hands-on Kickstart and start testing on the firm's own use cases. No physics degree is required. Becoming an expert takes months or more, and that is the argument for starting with your own people early.

In practice. Pick two or three of your own people, run them through a hands-on Kickstart, and let them test the firm's own candidate use cases with a referee that grades the result. What they do not need to build: the certified referee, the seven workbook templates and the published baselines to compare against.

Evidence, not promises.

Two ways in: a 30-minute call to pick a first use case, or guest access to the suite on a sample workbook for a first look at your own pace. Daniel Hug, founder: ETH physicist, MBA, formerly UBS, Credit Suisse, SIX. · daniel.hug@qubit-lab.ch

Book a call

Write to us