

POSITION PAPER

Quantum Computing for Financial Services

Why, why now, and how to find out what it means for your institution. Two clocks are running at once: regulators ask for post-quantum roadmaps while the same technology moves from physics demonstrations toward finance. This paper sets out the structural case, the 2026 timeline on both sides, the evaluation problem every bank faces, and a low-barrier way to answer it on your own data, with evidence instead of promises.

November 2026

qubit-lab.ch

qubit-lab.ch
contact@qubit-lab.ch

© 2026 qubit-lab.ch · All rights reserved · qubit-lab.ch/legal

Contents

Quantum cuts both ways: the one-page version . . . 2

1. Why now, the risk side . . . 2

2. Why now, the opportunity side . . . 3

3. Why: where quantum can plausibly matter in finance, and where not . . . 4

4. The evaluation problem: whose benchmark? . . . 5

5. A low-barrier entrance: Rapid Quantum Prototypes . . . 7

6. What this means for your role . . . 8

7. The next 90 days . . . 8

References . . . 9

Scope of this paper . . . 9

Quantum cuts both ways: the one-page version

Quantum computing reaches a bank from two directions at once, and the two arrive on different clocks.

Why quantum at all. A different method, not a faster computer: quantum finance methods are an alternative to classical methods 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 may be held. 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, and useless for most other problem shapes.

The risk clock is already set by regulators. The Swiss Financial Market Supervisory Authority (FINMA) surveyed 60 Swiss financial institutions and published its guidance on quantum computing in July 2026. It recommends that a post-quantum cryptography (PQC) roadmap be drawn up by mid-2027 at the latest. Germany's federal guideline of September 2026 sets an initial crypto inventory for the end of 2026 and quantum-safe high-risk systems for 2030. The G7 Cyber Expert Group's roadmap puts discovery in 2027 to 2028 and migration of critical systems in 2030 to 2032. None of these is a binding deadline. All of them are the same institutions, with the same threat model, saying the same thing: the technology is taken seriously enough to plan around.

The opportunity clock runs in parallel, on the same technology. In July 2026 IBM and partners published three quantum-advantage demonstrations with an open challenge to refute them. They are physics simulations, not finance, and Swiss experts rightly called them solid progress rather than a breakthrough. Investment rose 6.3 times to USD 12.6 billion in 2025, according to McKinsey's 2026 monitor. Your peers are already on the map: 67 published quantum initiatives at 34 financial institutions, 31 of them from 2025 and 2026 and 18 from 2026 alone, Intesa Sanpaolo, JPMorgan, HSBC, Wells Fargo, Erste, BBVA and UBS (2024) among them (map in section 2). McKinsey's monitor puts the value at stake in finance at USD 400 to 600 billion by 2035, and names unclear use cases as a leading barrier to adoption.

The state in one sentence: quantum advantage has been demonstrated, and is open to refutation, in physics; in finance it is still an open question, use case by use case, and that is exactly why an institution benchmarks its own use cases instead of waiting.

The decision this paper argues for is not a quantum programme. It is a measured answer, on your own data, to three questions: where could quantum methods matter for us, where do they not yet, and what would have to change for that to move. The Rapid Quantum Prototypes (RQPs) described in section 5 are built to give that answer in weeks: an Excel workbook in, a report with a certified classical referee and a verdict out, every run exportable as code your own engineers can rerun.

Why this year. FINMA recommends a post-quantum roadmap by mid-2027, and most Swiss boards are likely to see one. The board will then ask what the same technology could do for the firm, and whoever answers needs a measured number, not a vendor slide. The work to produce it, a briefing, a Kickstart for a small team and a first use case on your own data, fits a current innovation budget and takes weeks, not years.

1. Why now, the risk side

This paper does not sell post-quantum cryptography. The risk side is here for one reason: it is the clearest external signal that the technology has moved from speculation to planning, and it sets the calendar.

FINMA Guidance 05/2026, published 9 July 2026. Based on a survey of 60 Swiss banks, insurers, asset managers and market infrastructures, FINMA reports that about 70% had not yet planned or implemented any PQC measures, that only 8% had a specific roadmap, and that about two thirds expect quantum-related cyber risks to become directly relevant within seven years. About two thirds expect a quantum computer able to break RSA-2048 within ten years at the latest. FINMA recommends a PQC roadmap by mid-2027 at the latest, a cryptographic inventory, protection of long-lived critical data against "store now, decrypt later" collection, crypto-agility, and attention to external service providers. It is guidance, not a binding deadline, and it is deliberately limited to cryptography. It says nothing about quantum applications. The point for this paper is that a supervisor has put a date on the subject.

Germany, September 2026. The joint guideline of the federal ministries and the Federal Office for Information Security (BSI) adopts the EU coordinated roadmap for Germany: planning and an initial crypto inventory by the end of 2026, all high-risk use cases quantum-safe by the end of 2030, migration largely complete by 2035. The BSI's own working assumption is a cryptographically relevant quantum computer within at most 15 years, 2040 or earlier. The deadlines are urgent recommendations, to be made binding for critical infrastructure. For a Swiss institution with German or EU entities, or with German service providers, they apply through the supply chain before they apply in law.

The G7, January 2026. The G7 Cyber Expert Group's coordinated roadmap for the financial sector sets discovery and inventory for 2027 to 2028, risk assessment for 2028 to 2029, migration of critical systems for 2030 to 2032, and completion by 2035. It is explicitly non-binding. Switzerland is not a G7 member and is affected in practice, through dollar and euro clearing, through SWIFT and through FINMA's alignment.

What this means for the opportunity side. Three things. First, the roadmap work creates an inventory, a governance structure and a named owner for "quantum" inside the firm; the capability track can use all three. Second, when the roadmap reaches the board in 2027, the board will ask what the same technology could do for the firm. The question to put to your roadmap owner now: who answers that, and with what number? A measured number from the firm's own books, not a vendor slide. Third, the dates give the opportunity side a calendar it would otherwise lack.

2. Why now, the opportunity side

Hardware: advantage demonstrated in physics, open in finance. On 30 July 2026 IBM, with the University of Chicago, Qedma and Algorithmiq, announced three simultaneous quantum-advantage demonstrations: verifiable sampling on about 70 error-corrected logical qubits, an error-mitigated materials simulation that a national supercomputer could not match, and a disordered-matter simulation that stood unmatched for eight months on a public tracker with an open challenge to refute it. The results are posted openly for anyone to attempt to beat classically. The Swiss press ran a sceptical take two days later. Joseph Tindall of the Flatiron Institute, who advised one of the studies early on, said the material was "90% artificially constructed": chosen to be hard for classical methods, easy for quantum ones, with practical relevance secondary. Johannes Knörzer of ETH Zurich called it solid progress, not yet the great leap. Both readings are correct, and the second one matters most for a bank: all three demonstrations are physics simulation and sampling. No finance advantage was demonstrated. What was demonstrated is that the machines, the error handling and, notably, the open-challenge culture of verification have arrived.

Money and attention. Investment in quantum companies rose about 6.3 times year on year to USD 12.6 billion in 2025, and vendor revenue passed USD 1 billion, according to McKinsey's 2026 monitor. McKinsey's 2026 monitor estimates the value at stake in financial services at USD 400 to 600 billion by 2035, led by portfolio allocation in asset and wealth management, and names unclear use cases and limited executive buy-in among the leading barriers. The barrier is not computers. It is the absence of an answer to "which of our problems".

Banks are publishing. Intesa Sanpaolo has published peer-reviewed work on quantum amplitude loading for rainbow and autocallable option pricing. HSBC and IBM reported in September 2025 that quantum-processed inputs improved the accuracy of a bond-trade fill-probability model by up to 34% (relative gain in area under the curve, on 1.7 million historical requests for quote, processed offline, no live trading), a figure the partners reported themselves and that no independent referee has certified. The pattern is the same everywhere: the work is real, the scale is small, and the benchmark is chosen by whoever publishes. The map below counts 67 published initiatives at 34 institutions on the opportunity side; 31 of them date from 2025 and 2026, 18 from 2026 alone, against 15 from the four years 2018 to 2021 together. The opportunity clock has a tempo of its own.

Why this adds up to "now". Each of these signals alone could be waited out. Together they describe a technology whose risk side has regulatory dates, whose hardware side has crossed from claims to open, refutable demonstrations in physics, and whose finance side consists of small, real, self-benchmarked results. An institution that wants to tell the difference between the three when the next vendor arrives needs a yardstick of its own. That yardstick is built by running the firm's own problems, not by reading about other people's.

Current quantum-finance initiatives by banking domain and algorithm family, November 2026

	Annealing	QAOA	QMC / QAE	VQC (QML)	QSVM (QML)	General QML	Other / unspecified
Derivatives Pricing			UOB 73 2026; Scotiabank 54 2019; J.P.Morgan 17 2020; ERSTE Group 93 2025; BMO 53 2019				UBS 33 2024; Goldman Sachs 32 2022; Crédit Agricole 62 2023; citi 21 2022; BBVA 42 2025
Trading / FX						WELLS FARGO 81, 82 2024, 2026; HSBC 10, 41 2025; Fidelity 58, 59 2021, 2023	Standard Chartered 31 2020; HSBC 78 2026
Customer Retention				itaú 6 2022	ERSTE Group 94 2025		
Fraud & AML	AMERICAN EXPRESS 80 2025			LLOYDS BANKING GROUP 12 2025; Intesa Sanpaolo 5 2024	itaú 20 2022	Mastercard 89 2026; LLOYDS BANKING GROUP 67 2026; Bci 90 2026	Standard Chartered 30 2025; NatWest 74 2026; HSBC 19 2025
Risk Optimization	Standard Chartered 29 2020; Intesa Sanpaolo qi 83 2026; Bank of Canada 84 2026; BBVA 15 2020		Deutsche Börse 57 2021				Goldman Sachs 26 2022; BARCLAYS 25 2023
Credit Risk			MIZUHO 64 2025; Bank of Canada 4 2022	Intesa Sanpaolo 18 2023; BARCLAYS 24 2023		AMERICAN EXPRESS 79 2025	Crédit Agricole 63, 72 2023, 2026; BBVA 13 2020
Portfolio Allocation	SMBC qi 88 2026; Rabobank 37 2023; NatWest qi 52 2018; COMMERZBANK qi 55 2019; CaixaBank 60 2022; ally 61 2022	J.P.Morgan 11, 77 2024, 2026; citi 1 2021		Vanguard 39 2025			UBS 27 2024; J.P.Morgan 75 2026; BBVA 3 2020
Cross-Domain						WELLS FARGO 56 2020; ERSTE Group 95 2025	Standard Chartered 65 2026; RBC 85 2026; J.P.Morgan 71 2026; Crédit Mutuel 35, 36 2022, 2023; TD Bank 86 2026; BMO 68 2026

Opportunity side only: 67 published projects, pilots and papers at 34 financial institutions, 31 of them from 2025 and 2026. Rows are banking domains, columns algorithm families, each card an institution with the year. Published projects, pilots and papers only; no regulator guidance, no consultancy studies; the quantum-risk side (post-quantum cryptography, QKD) is left out here. Card numbers refer to the entries of the online matrix; logos identify the publishing institution only and remain their owners' trademarks; entries marked quantum-inspired ran on classical hardware. Full matrix on qubit-lab.ch.

3. Why: where quantum can plausibly matter in finance, and where not

Quantum computing is not a faster computer. It is a different one, and quantum finance methods are an alternative to classical methods 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 may be held (cardinality constraints). 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, and useless for most other problem shapes. A fair "why" has to say all three.

Where the shape fits.

- **Constrained combinatorial optimization with genuine pairwise structure.** Portfolio construction under budgets, bands and bundles; index replication with a limited number of names; credit and collateral allocation under concentration limits; feature selection for a model. These are quadratic problems: the value of choosing one position depends on which others are chosen. Quantum optimization methods (the QAOA family) are built for exactly this shape. The qualification is that the pairwise terms must be essential, not decoration; a problem whose objective is really linear belongs to a classical solver forever.
- **Estimation by sampling: pricing and risk.** Monte Carlo pricing of derivatives and the tail-risk measures built on it. Quantum amplitude estimation offers a provable improvement in how the estimation error falls with the number of

samples. The improvement is real and measurable at small scale today; what is not yet known is the break-even budget on realistic payoffs once hardware overheads are included. That is a number to measure, not to assert.

- **Classification with structure a classical model does not see.** Fraud detection, credit scoring, downgrade prediction: quantum kernel and variational classifiers map data into a space where some label structures become linearly separable. Whether a given dataset has such structure is an empirical question with a measurable answer.

Where it does not. Anything a linear or mixed-integer solver already handles at scale, which is most of settlement, netting, cash-flow matching and scheduling. Anything whose bottleneck is data movement rather than computation. Anything that must run in production latency this decade. And, to be plain about it, the cryptographic threat is a risk to manage, not an application to buy.

What our own evidence says so far. Since August 2026 qubit-lab.ch has published six method studies and twelve evidence records on qubit-lab.ch/evidence, every run scored against a certified classical optimum or a tuned classical baseline. The short version by problem shape:

- Optimization: on an allocation modelled on a Swiss private bank's 78-asset mandate, whose policy leaves 23 free decisions, 23 qubits (a demo case, report public) the quantum method reached 100% of the certified optimum; on the S&P 500 replicated with 20 names it matched the certified tracking error; on a 23-line collateral inventory it matched the referee that enumerated every allocation. All three are filed as Parity. On today's hardware, circuits deeper than three layers returned noise, not answers.
- Pricing: on a two-asset basket option the quantum estimator's error fell as 1/N instead of $1/\sqrt{N}$, the quadratic speedup, with break-even at several thousand samples at simulation scale. Filed as a Method study.
- Classification: on six datasets no quantum kernel beat the classical bar on market data. On a dataset built to have the right structure, a six-qubit quantum classifier beat the best classical model by a factor of 1.45, with the whole confidence interval above parity, filed as Advantage demonstrated. So quantum territory exists; the open question is whether any of your data lie inside it. The one-minute pre-screen that answers that has said no on six market datasets so far: it checks whether a quantum kernel on your features is measurable at all and separates your labels better than its classical twin. Bring yours to the first call.

The result so far is plain: at the problem sizes today's hardware can handle, quantum approaches are a credible alternative, and the classical benchmarks still win or tie. The real question is scale. Exact certification stops being practical as the instance grows; from there on the classical answer is a heuristic without a proof, and that is where the crossover, if any, would appear, and where the referee's own limit becomes the measurement. At the sizes a referee can certify, parity is the expected result; the benchmark measures how the gap and the classical effort move as the instance grows toward the referee's limit, and leaves the workbook for next year's rerun on larger hardware. Measuring that movement, rather than assuming the crossover, is the point. The full figures, run by run, are on qubit-lab.ch/evidence.

4. The evaluation problem: whose benchmark?

Every quantum result a bank will be shown in the next three years has the same weakness, including ours: somebody chose the problem, the baseline and the metric. The IBM demonstrations chose problems built for the machine. HSBC chose its own metric. Our own published cases chose instance sizes small enough for an exact referee to certify. The question for a decision-maker is therefore never "is the claim true" but "true on what, against what, and would it hold on ours".

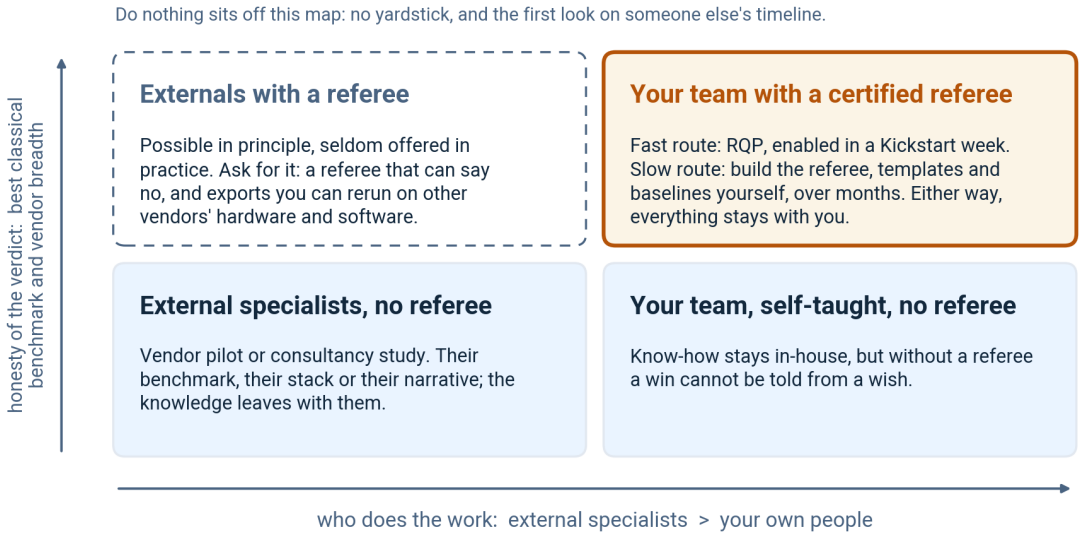
Three rules follow, and they are the whole of our methodology:

1. **No quantum method grades its own homework.** Every quantum result is scored against the strongest classical attempt we can mount on the identical data and objective: a certified exact solver where one exists, a tuned classical baseline where it does not, including a classical twin built to imitate the quantum model. A "0.00% gap" is then a mathematical statement.
2. **The instrument must be able to say no.** An assessment runs before any paid quantum run and is allowed to conclude that quantum is not worth prioritizing on this data. The public evidence page files records under all verdicts, Parity and No advantage demonstrated included, next to the one Advantage demonstrated record. A tool that can never say no is marketing.

3. **Every number must be reproducible by the client.** Reports, notebooks and machine-readable records that rerun offline, in the client's own framework. Trust should never rest on the vendor's word, ours included.

What "certified referee" means. A conventional solver that provably finds the best possible answer to the same spreadsheet, so that the quantum result is graded against the truth, not against our own baseline. Where no exact solver exists, the referee is the strongest tuned classical model we can build on the same data.

The alternatives, side by side. Two axes decide the shape of a first step: who does the work, your own people or external specialists, and whether a certified referee grades the result so that the verdict can say no. RQP is the fast route into the top-right field, not the only one: a quant team can build its own referee over months, and a vendor could bring one. The options combine: external specialists or a self-taught team can use the RQP referee as their instrument, which is how we work with partner consultancies.



External specialists: vendor pilot or consultancy study

Vendor neutrality. The vendor or the consultancy picks the benchmark; ask for a referee.

Can it say no? ✘ Only if the brief asks for one.

Knowledge transfer. Code on one stack, or a market view and a narrative.

Learning curve and risk. Months; their benchmark, and the knowledge leaves with them.

Your team with RQP, enabled in a Kickstart week

Vendor neutrality. A certified classical referee and framework-neutral exports.

Can it say no? ✔ By design.

Knowledge transfer. Reports, benchmarks, exportable circuits, trained people.

Learning curve and risk. Weeks to a verdict; prototype scale, and the likely first verdict is parity, with the yardstick as what you keep.

Do nothing

Vendor neutrality. Nobody picks a benchmark.

Can it say no? ✘ Never asked.

Knowledge transfer. Nothing.

Learning curve and risk. Never a verdict; the first serious look comes on someone else's timeline.

Your team, self-taught

Vendor neutrality. Your team, learning as it goes; a referee of your own is possible.

Can it say no? ✘ Until the referee exists.

Knowledge transfer. Know-how; referee and templates only if you build them.

Learning curve and risk. Months to a referee; shelved after the first disappointment.

Two costs sit on either side of a wrong answer. Building too early means engineering spend on circuits that classical methods beat, and a team that has learned to distrust the subject. Dismissing too early means that the firm's first serious look happens when a competitor's result forces it, on the competitor's timeline, with no internal yardstick. A measured answer at prototype scale costs weeks and is the cheapest insurance against both.

5. A low-barrier entrance: Rapid Quantum Prototypes

The Rapid Quantum Prototyping suite is the measuring instrument this paper has been describing. It exists so that an institution can answer the "which of our problems" question without a research programme, a vendor commitment or a hiring plan.

What it is. Seven prototypes, each for one finance use case: portfolio optimization, index replication, credit allocation, collateral allocation, feature selection, classification for fraud and credit scoring, and derivatives pricing by quantum Monte Carlo. All seven run on the client's own data, uploaded as an Excel workbook: positions, loan books, collateral inventories, feature tables, payoff scenarios. The workbook carries the investment policy too, so fixed holdings, exclusions, bundles, sector budgets and bands are entered in Excel and respected by the quantum step. Workbook in, certified report out.

What comes back. A report that opens with the verdict and the referee's number, then the quantum result, then the full benchmark table. The vocabulary is closed, Advantage demonstrated, Parity, No advantage demonstrated, Inconclusive, Method study, so that it cannot inflate. Every run exports rerun and retrain notebooks for Qiskit, PennyLane, Cirq, CUDA-Q and Classiq, plus the circuit in OpenQASM, so that the client's engineers can reproduce the result in the framework they already use and carry the logic forward without depending on us or on any one hardware vendor. Runs execute on simulators with a noise ladder and, for the optimization tools, on IBM hardware.

Why the barrier is low. No code to write: the interface is a spreadsheet and a report. No vendor to choose: the exports are framework-neutral. No budget to defend beyond a monthly team licence with a minimum of two months, fixed fee, agreed up front. No claim to take on faith: every number has a referee, and the methodology is published as the RQP Evidence Engine paper. The suite can also be driven by the client's own scripts or an assistant through a standard interface (MCP), switched on only on request and only for the client's own runs. The natural owner is the AI or data team, and the size fits a team-level innovation budget. A first engagement is sized at CHF 40,000 to 80,000 depending on scope (briefing, Kickstart, two to three months of usage), and under CHF 100,000 in any configuration.

Your data. Workbooks are processed in a Google Cloud environment in the Zurich region; run artifacts expire after 24 hours; the workbook needs no client identifiers, so positions and counterparties can be masked; and nothing from a client run enters our published evidence. Where position data cannot leave the bank, the first look runs on a sample or synthetic book, and the exported notebooks rerun inside your own environment. A non-disclosure agreement is available on request.

How an institution enters. The Lean Quantum Exploration Process has three steps, and a firm can enter at any of them:

1. **Quantum Awareness Briefing.** Ninety minutes to half a day for leadership, risk and technology: what quantum can realistically do today and which preparation is proportionate.
2. **Quantum Capability Kickstart.** Five half-day sessions, typically within one week, for a team of up to eight: first circuits in the leading frameworks, the seven RQP use-case templates, and the team's own three to four candidate use cases. An October 2026 pilot cohort, three mornings and so the first two days of the format, brought together practitioners from a wealth-tech firm, a consultancy, a family office and a bank.
3. **RQP usage.** Monthly team access to the suite: the team runs its candidate use cases on its own data, trial and error, every result benchmarked by the referee, with onboarding, weekly support and a closing stop-or-go verdict per candidate. Usage continues month by month as hardware and algorithms improve, so the same workbooks can be rerun against next year's devices.

Advisory at any step: build versus buy, vendor shortlists, and the reading of other people's benchmarks.

Who we are. qubit-lab.ch is founded and run by Daniel Hug: a career in financial-services technology and transformation at UBS, Credit Suisse and SIX, an MSc in physics from ETH Zurich, an MBA. Methodology and evidence have been published since August 2026. The firm is deliberately small: one counterpart, fixed scope, a Swiss sole proprietorship in Zurich. Every result leaves as reports and code in your own frameworks, complete at hand-over, so nothing you build depends on our continued existence.

What participants say. From the Kickstart pilot of October 2026: "It was the first time I fully grasped how quantum computing can be applied to real-world problems." Simone, Risk Quant. "I started with no quantum background and left with a real understanding of the theory and a clear idea of where it can be applied in finance." Marco Zietzling, Manager at a consulting firm.

What the first engagement is designed to produce. Two or three of the firm's own use cases with a verdict each, a benchmark the quant team can rerun each quarter, exported circuits the engineers have opened, and a one-page answer for the board: where quantum matters here, where not yet, what would have to change.

6. What this means for your role

Board, C-level, heads of business lines

Is there money in this, when, and what is the decision?

A measured answer on your own books before the next planning cycle, and a yardstick for every vendor claim after it: the decision stops being a bet on a slide. Not a programme, not this year in production. Finding out costs weeks and a first engagement is sized at CHF 40,000 to 80,000 and can sit in an existing AI budget. First step: fund one measured use case with a verdict you can quote.

Head of quant or trading desk

Do these methods beat what we run?

A certified benchmark on your own constraints that you rerun each quarter and hold up against every vendor pitch: the desk owns the yardstick instead of reacting to claims. Your people learn the methods on your problems, in weeks, without a physics degree. First step: put your constrained allocation or replication problem in the workbook and read the gap.

Head of innovation

What can I show, and what will it prove?

A pilot that is showable and defensible because a referee graded it, with exported circuits the engineers have opened and a one-page verdict for the board. First step: a Kickstart for your team plus one use case with a business-side sponsor.

Head of AI or data science

Is quantum machine learning anything?

A one-minute answer per dataset to whether any of your data lie in quantum territory, with the classical bar kept in the room; if one does, you hold the first market dataset inside it. First step: run one labelled dataset through the assessment gate.

CTO, CIO

What do we build on, what locks us in, and what belongs in the 2027 plan?

A capability track beside the PQC roadmap that costs weeks, not a platform decision: circuits leave as Qiskit, PennyLane, Cirq, CUDA-Q, Classiq or OpenQASM, your engineers rerun them without us, and nothing locks you to a vendor. First step: have one engineer rerun an exported notebook, and name an owner for quantum opportunity as you have for quantum risk.

7. The next 90 days

1. **Decide to measure, not to believe or dismiss.** One sponsor on the business side, one on the technology side, and a mandate to produce the firm's own evidence before the 2027 planning round.
2. **Pick the use cases where the shape fits.** Two or three problems with real pairwise structure, a real estimation task, or a labelled dataset. Size the first instance so that an exact referee can still certify it.
3. **Run them, read the verdicts, and write the one page.** Where quantum matters for us, where it does not yet, what would have to change, and when we look again. Keep the workbooks for next year's rerun.

The next step, two ways in. A 30-minute call to pick the first use case, or guest access to the RQP suite on a sample workbook for a first look at your own pace: write to daniel.hug@qubit-lab.ch with the subject "Guest access" and your institution, and you receive a personal code within two working days. Daniel Hug, founder.

Book a 30-minute call
calendly.com/daniel-hug-qubit-lab

Write to the founder
daniel.hug@qubit-lab.ch

References

1. FINMA, Guidance 05/2026: Quantum Computing (9 July 2026), finma.ch.
2. BMDS, BMI and BSI, Nationale Umsetzung der EU-Roadmap zur Umstellung auf Post-Quanten-Kryptografie, Gemeinsame Leitlinie (11 September 2026), bsi.bund.de.
3. G7 Cyber Expert Group, Coordinated roadmap for the transition to post-quantum cryptography in the financial sector (January 2026), gov.uk.
4. McKinsey, Quantum Technology Monitor 2026: A commercial tipping point (2026), mckinsey.com.
5. IBM Quantum, quantum-advantage announcement and Quantum Advantage Tracker (30 July 2026); Sampling hard circuits with verifiably high fidelity, arXiv:2607.25941; advantage criteria, arXiv:2506.20658.
6. L. Leiva Ariosa, "IBM will die Überlegenheit seiner Quantencomputer nachgewiesen haben", Neue Zürcher Zeitung (1 August 2026).
7. F. Cibrario et al., "Quantum amplitude loading for rainbow options pricing", IEEE QCE 2024; "Autocallable options pricing with integration-based exponential amplitude loading", arXiv:2507.19039.
8. HSBC and IBM, announcement of the quantum-assisted bond RFQ fill-probability trial (25 September 2025), hsbc.com and ibm.com.
9. qubit-lab.ch, The RQP Evidence Engine (August 2026), The RQP Suite Overview (October 2026), evidence page qubit-lab.ch/evidence.

Scope of this paper

A position statement for structured discussion with financial institutions, published by qubit-lab.ch Daniel Hug, a sole proprietorship registered in Zurich in October 2026. Regulatory references are guidance, not binding deadlines; our own figures come from the published studies and records, which remain canonical. Not financial advice; no performance claim beyond what the cited records describe.

Legal notice. © 2026 qubit-lab.ch. This document may be circulated within your institution and shared with your advisers; please keep it complete and attributed. It is a position statement for structured discussion, not financial advice, and it makes no performance claim beyond the cited records. Full legal terms: qubit-lab.ch/legal