The Lean Quantum Exploration Process

Two steps for your own teams: from the first circuit to evidence on your data.

qubit-lab.ch

No quantum expertise needed on your side, no transformation program. Your data science and quant teams learn the craft in one week, then test their own problems on ready-to-use quantum prototypes with honest classical benchmarks. Evidence, not promises.

1 Kickstart week

Enable your team to understand quantum computing and the quantum algorithms for finance.

Five hands-on half-day sessions for quant and data science teams. From qubits and gates, via the first circuits of their own, to the algorithms that matter in finance: optimization, classification, Monte Carlo. Not a coding course, but the understanding to be productive with the RQP tools right after. Python and NumPy are all it takes.

- Your team knows what quantum algorithms can do, what they cannot, and how to tell.
- Real finance examples: portfolio selection, classification, Monte Carlo.
- Ends with your own shortlist of three to four use-case candidates.

2 Rapid Quantum Prototyping (RQP)

Test your own problems and data on the best current quantum algorithms available: from Excel to results, with the best benchmarks.

Your team loads its own data as an Excel workbook. Seven integrated, ready-to-use prototypes cover portfolio optimization, index replication, credit and collateral allocation, classification (fraud, credit scoring) and derivatives pricing. Every run is measured against an exact classical solver that proves the optimum where it can, or against tuned classical baselines on identical data. If classical wins, the report says so.

- One report per run, with proven optimality gaps; stop or go per use case.
- Synthetic or anonymised data is enough for a first pass. Runs on Swiss infrastructure in Zurich; IBM quantum hardware on request.
- Runs in hours, a full use case in a few weeks. Circuits export to Qiskit, PennyLane, Cirq and CUDA-Q.

How it starts

30 minutes	A call to pick the first use case and the team.
Fixed scope	One team, one week, one use case at a time. Fixed fee, agreed up front.
One counterpart	Swiss-based, hands-on, from the first session to the final report.

Scan for studies, run records and sample reports

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

Practical finance use cases ready to test with your own data

Seven prototypes, every result measured against a classical referee.

qubit-lab.ch

The finance questions that today's quantum algorithms address fall into two families. Selection problems (which positions, loans, collateral lines or model features, under budgets and rules) become QUBOs solved with QAOA. Pattern problems (fraud, credit scoring, pricing) become quantum classifiers or amplitude-estimation readouts. In every case an exact solver or a tuned classical baseline sets the bar on identical data, and the report prints one of three verdicts: advantage demonstrated, parity, or no advantage demonstrated.

Finance question	What the prototype does	Classical referee, demo verdict
Portfolio optimization Which positions and thematic bundles make the best portfolio under sector budgets?	QAOA on the cardinality-constrained selection from your Excel workbook of positions, bundles and budget bands.	Exact MIQP (SCIP), certified optimum. Demo: parity, gap 0.0%.
Index replication How closely can 20 names track the S&P 500 under mandate rules?	QAOA over your instrument pool: fixed names, sector blocks, theme baskets; every mandate rule priced in tracking error.	SCIP dual bound, MIP gap 0. Demo: parity, certified 4.473% tracking error.
Credit allocation Which loans from the offered book fill the budget with the best risk-adjusted return?	CVaR QAOA on the loan selection with committed lines and concentration limits.	Exact solver, certified gap on every result.
Collateral allocation Where to post collateral across venues under margin calls, at the lowest cost with stable coverage?	Robust CVaR QAOA over 50 stress scenarios; locked lines and pledge pools respected.	Exhaustive enumeration of 16.8 million allocations. Demo: parity, gap 0.00%.
Feature selection Which K features should the credit or fraud model use, under governance rules?	Constrained mRMR selection as a QUBO: must-include and exclude rules, bundles, type budgets; validated out of sample; hands off to the classifier.	SCIP-certified selection with solution pool.
Classification (fraud, credit scoring) Does a quantum classifier beat our tuned models on this data?	VQC and QSVM with an assessment gate that says whether a quantum run is worth the money before you spend it.	Six tuned baselines incl. an order-matched twin, paired bootstrap CI. Demo: no advantage on market data (0.92x), advantage on quantum-native data (1.45x).
Derivatives pricing How does quantum amplitude estimation converge against Monte Carlo on our payoffs?	European, digital, basket and barrier payoffs from one workbook; convergence readout side by side.	Analytic and grid references, classical Monte Carlo. Lab tier: method study, not an own-data verdict.

What the evidence says

Eleven records on qubit-lab.ch/evidence, unedited solver output on demo data: five product results (three at parity with the certified optimum, one advantage demonstrated on quantum-native demo data, one no advantage demonstrated on market data) and six method studies with pre-registered designs and printed null results. Full reports, workbooks and study PDFs are public.

Scan for the prototypes and demo videos