RAPID QUANTUM PROTOTYPING (RQP)

The RQP Suite Overview

One honest methodology, five instruments - portfolio optimization, index replication, feature selection, quantum machine learning, and quantum Monte Carlo. What each tool does, what the runs actually delivered, and how to try it on your own data.

July 2026

qubit-lab.ch

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

Contents

1. The problem this suite solves

Financial institutions receive quantum pitches with claims that are hard to verify. Boards ask "what should we do about quantum" and deserve an answer that survives due diligence - not a vendor slide. The Rapid Quantum Prototyping (RQP) suite exists to produce that answer on your own data: every quantum result is benchmarked honestly against strong classical baselines and, wherever the problem size allows, against a certified exact optimum - **including the cases where the classical method wins**. Real numbers instead of quantum hype.

Four commitments hold across all five products:

- **Like-for-like scoring.** Quantum and classical methods are always judged on the identical objective, on the identical data. No cherry-picked baselines.
- **A referee wherever possible.** Certified exact solvers (SCIP) prove optima for the optimization tools; tuned classical baselines and pre-registered protocols referee the ML and Monte Carlo tools. A "0.00% gap" in an RQP report is a mathematical statement.
- **The gate that says no.** Assessments run before paid quantum runs and are allowed to conclude "Do Not Prioritize Quantum". A tool that can never say no is marketing, not measurement.
- **Verifiable artifacts.** Every claim lands in a downloadable report, notebook, or machine-readable file that reruns offline. No claims you cannot check.

2. The five instruments

Everything runs on your own data, uploaded as a simple Excel workbook: portfolios, index compositions, labelled datasets, payoff scenarios. Current focus is finance; energy and other domains are covered by the shipped demonstrators and can be prepared jointly.

Tool	Tier	The business question
Portfolio Optimization RQP Pro	Pro	Which investment blocks should the portfolio hold, under budget, risk, and type constraints?
Index Replication RQP	Pro	How well can a large index be tracked with a small basket - and what does each mandate rule cost?
Feature Selection RQP	Pro	Which K features may the model use, under fairness, bundle, and budget constraints?
QML RQP Pro	Pro	Can a quantum kernel or variational classifier beat tuned classical models on your labelled data?
QMC RQP Lab	Lab	Where does quantum amplitude estimation stand against classical Monte Carlo on real payoffs?

Pro versus Lab is an honesty distinction, not a quality one: Pro tools create suspense on your own data - the answer is genuinely unknown before the run. The Lab tool delivers transparency on estimator convergence - the scientific state of play, measured cleanly.

3. Portfolio Optimization RQP Pro

Cardinality-style portfolio construction as a QUBO: costs, expected returns, covariance, budget, and up to five subtype budget constraints (sector, region, asset class, sustainability bucket). Four independent routes attack the same objective: a classical heuristic, the QAOA sampler, the certified SCIP referee, and - on request - a real IBM quantum computer as a **hardware second opinion**, whose noisy readout is reported as honestly as the simulator's clean one. Exports include the branded PDF and runnable Qiskit, PennyLane, and Cirq notebooks.

4. Index Replication RQP

Replicate a large benchmark with a small basket - fixed core holdings, basket-traded building blocks, single lines - under exclusions and type budgets. The referee prices every mandate rule: the report states the certified tracking-error cost of each ESG exclusion, fixed holding, and band. The cardinality-conserving XY mixer keeps quantum readout feasible by construction; enhanced-indexing mode adds a return tilt on the same certified ladder.

5. Feature Selection RQP

Upstream of any model: which K of N candidate features may the model use, under must-include and must-exclude policy, all-or-nothing bundles (one-hot families, sin/cos phase pairs), exclusive variant groups, and type budgets - exact or band, conserved by the quantum mixer with zero slack qubits. The solver ladder ends in a certified optimum plus a ranked **solution pool**, because the best generalizer is not always the objective winner - choosing rank 5 is a governance decision, and the manifest records it. One click exports a colored QML handoff pack for the next tool in the chain.

6. QML RQP Pro - quantum machine learning

Quantum kernel methods (QSVM) and variational classifiers (VQC) against seriously tuned classical baselines, on a temporal-honest split. The staged pipeline - prepare, assess, plan, execute - puts two assessments before any paid run, including the **Fourier Wall analysis** (order-matched periodic twin with joint terms to 3rd order, off-grid indicator, entanglement lift, joint-spectrum density, and the five-condition quantum-territory checklist; independent implementation after arXiv:2607.15815) that predicts whether quantum feature maps have anything to grip on your data. The verdict metric is the Quantum Advantage Factor with a full confidence interval; the whole-CI rule is pre-registered, and "Do Not Prioritize Quantum" is a common, correct answer.

7. QMC RQP Lab - quantum Monte Carlo

Amplitude-estimation pricing of option payoffs (calls, puts, digitals, barriers, baskets) benchmarked against classical Monte Carlo on shared budget schedules: absolute error versus oracle calls, both estimators, same axes. The quadratic-speedup claim, measured instead of cited.

8. Worked examples, in numbers

All numbers below come from actual RQP runs; the full reports, workbooks, and notebooks behind them are available on the product pages and on request.

Allocating new capital for a private bank (Portfolio Optimization, demo case). A global equity model portfolio with a Europe/Switzerland anchor: 78 positions evaluated, 23 investment decisions mapped to 23 qubits, a USD 1.013m budget, a technology band of USD 225k-305k, and six pre-approved thematic bundles. QAOA at depth p = 4 with layerwise warm start reached, in CVaR objective mode, **100% of the certified optimum**. The problem size is deliberate: large enough to carry the real mandate structure, small enough that the exact referee can still certify the optimum - which is what makes the benchmark meaningful.

The S&P 500 with 20 stocks (Index Replication, demo case). The full 500-name benchmark; an instrument pool built the way mandates actually trade: five fixed mega-caps, CIO sector blocks, AMC-style theme baskets (climate transition, dividend quality, AI infrastructure), a completion basket of the smallest index weights, and ten single lines - 17 qubits in total. Certified optimum: **4.42% tracking error with exactly 20 ISINs**. The two mandate rules - a technology cap and a climate-transition floor - cost a certified 0.40% per year, and the report shows what the ESG rule displaced (the AI basket). QAOA reproduces the certified optimum exactly - and the referee proves it.

Governed feature selection on one million rows (Feature Selection, real public data). The NeurIPS Bank Account Fraud dataset: 32 raw columns reduced to 15 governed features under a fairness exclusion, a mandated bureau score, one-hot bundles, and a behavioral budget. QAOA matched the certified optimum at **0.00% gap** with every constraint conserved by construction - feasible readout mass 1.0000, zero repairs.

Quantum kernel classification (QML). On an engineered 8-qubit benchmark designed to favour quantum kernels, QSVM reached **0.87 test accuracy and 0.95 ROC-AUC versus 0.73 and 0.85** for the best classical baseline on identical data - so you can see what winning conditions look like. On plain tabular data, classical baselines often still win, and the reports say so openly (see the chain below).

Pricing readout (QMC Lab). Quantum amplitude-estimation readout reproduces the analytic and grid-based references across European, digital, capped, basket, and barrier payoffs, plus rare-event probability estimation, with confidence intervals and convergence curves against classical Monte Carlo on shared oracle budgets.

9. The chain - a complete, honest arc on real data

The suite's tools compose, and the flagship demonstrator runs the full chain twice. **Act one**: on the public Bank Account Fraud dataset, Feature Selection solved the constrained 15-of-32 selection to a certified 0.00% gap and handed off to QML - whose own gate said Do Not Prioritize Quantum. We ran the kernel anyway, for the record: classical won, exactly as predicted. **Act two**: on two years of real German power-market data, phase-encoded cycle features chosen by the assessment's own mechanism reasoning reached **certified statistical parity** - Quantum Advantage Factor 0.99, confidence interval 0.975 to 1.011 on purely future test hours - under a pre-registered six-configuration protocol with a hard stop.

The dataset where the gate said no. The dataset that reached certified parity. The engineered dataset where the quantum kernel wins. All three exhibits are downloadable with reports, workbooks, and rerunnable notebooks - the loss published next to the win, because the loss is what makes the win credible.

10. What you get with every run

- **A branded PDF report** written for a committee, not only for quants: plain-language lead-ins, certified gaps, solver comparisons, circuit resources, and a methodology primer.
- **Runnable notebooks** (Qiskit / PennyLane / Cirq, per tool) with the full result embedded - they rerun offline, survive cloud record expiry, and are execution-verified with pinned package versions.
- **Machine-readable results** (JSON/CSV) for your own tooling, plus the exact workbook and settings for reproduction.
- **An honest verdict** - advantage, parity, or classical win - stated in pre-registered terms you can defend internally.

For the technically curious: how the quantum engine works

Three of the five instruments - portfolio optimization, index replication, and feature selection - run on the same quantum engine, the Quantum Approximate Optimization Algorithm (QAOA). In brief:

Unlike a classical optimizer, QAOA does not examine one candidate portfolio after another. Instead, it creates a single quantum wave that spans the entire search space, giving every possible solution the same initial probability.

The algorithm then repeatedly applies two global transformations to this wave. The cost step causes the wave to evolve according to the optimization objective. The mixer step allows the wave to evolve further, enabling different parts of the wave to interfere. Through constructive and destructive interference, the probability distribution is gradually reshaped. Regions of the search space containing good solutions become increasingly likely, while regions containing poorer solutions become less likely.

After a small number of these alternating steps, the quantum state is measured. Each measurement returns a single portfolio; in practice the circuit is read out many times, yielding a ranked list of candidates. There is no guarantee that the best candidate is the optimum, but it is significantly more likely to be a high-quality solution than one selected at random.

The only quantities that must be adjusted are the strengths of the cost and mixer steps. These are tuned by a conventional classical optimizer, which repeatedly executes the quantum circuit, evaluates the measured results, and updates the parameters until the best setting is found.

A useful intuition is to imagine a hall filled with tuning forks. Nobody walks through the hall striking or tuning individual forks. Instead, the entire hall is exposed to a carefully designed sequence of sound fields. The forks respond collectively. As the sound field changes, resonance and interference gradually reshape the overall pattern of vibration until certain resonances dominate while others fade away. Listening to the hall is like measuring the quantum state: the strongest resonance is now much more likely to correspond to a good solution than to a poor one.

The analogy is not exact, but it captures the essential idea. QAOA does not search through candidate portfolios one by one. It repeatedly reshapes a single collective wave, using interference to make good solutions more likely to be observed.

The full treatment - the picture, the mathematics, and each tool's constraint-conserving variant - is in every product manual's "The algorithm in two passes" section.

11. Trying it and working with qubit-lab.ch

Each product page on qubit-lab.ch carries sample workbooks, sample reports from the runs above, and demo access; an LLM prompt pack lets any capable AI assistant build a valid workbook from your own data against the live validator. For teams, the **Lean Quantum Exploration** process on the services page takes you from a first briefing through a hands-on kickstart to a six-week guided exploration of your own use cases on the suite - fixed fees, agreed up front. Swiss-based, hands-on, one counterpart.

Intro call, demo keys, walkthroughs: contact@qubit-lab.ch