Six customers, one campaign

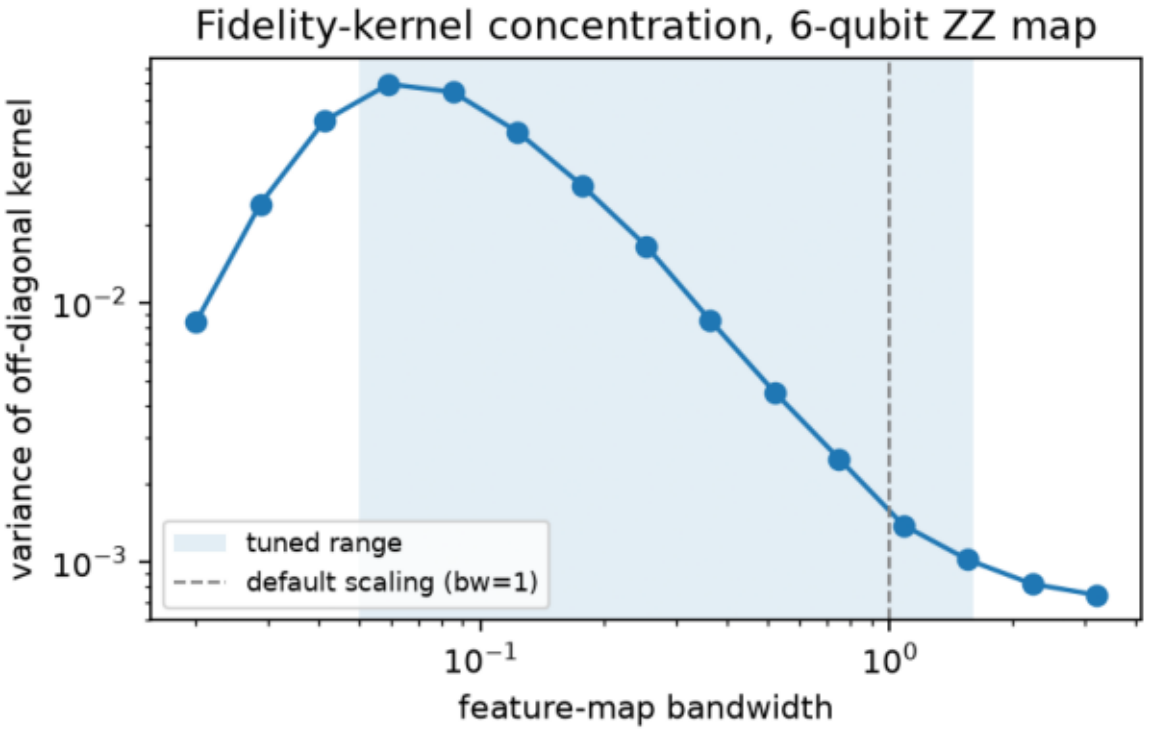
- Answer: contact C0031, C0081, C0079, C0087, C0397, C0374.
- P(at least 3 of 6 subscribe) = 0.1971 from the quantum circuit (exact 0.1969, Stage-1 bootstrap 95% band 0.059 to 0.456).
- Stage 2 optimality gap 0.0%; selection checked against all 210 subsets.
- Every number on these slides is read from the result files in the repository.

Pipeline: what is quantum and what is classical

- Stage 1 (hybrid): ZZ feature-map kernels (fidelity and projected) on 6 qubits, blended with logistic regression. Classical: preprocessing, PCA, SVM training, Platt calibration, the blend weight.
- Stage 2 (quantum optimiser, classically verified): QAOA with a Dicke-state start and an XY ring mixer, so 'exactly 6' holds by construction. Classical: COBYLA, presolve, brute-force check.
- Stage 3 (quantum sampler, classically verified): one RY-rotated qubit per customer; count outcomes with >= 3 ones. Classical: exact Poisson-binomial and Monte Carlo references.
- Stage 1 probabilities flow unchanged into Stages 2 and 3; scripts/check_submission.py verifies this.

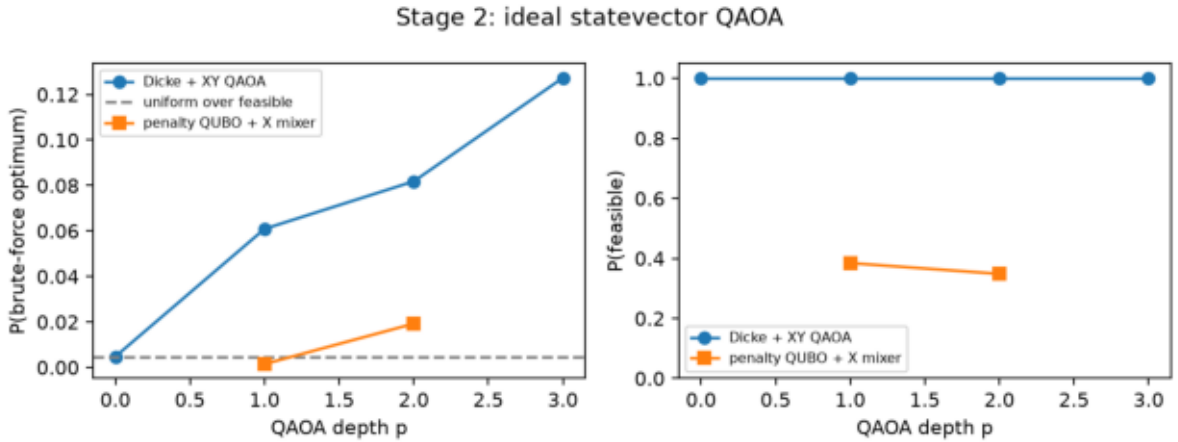
Stage 1: an honest quantum classifier

- Nested CV (10 x 5 folds, train only). ROC-AUC: LR 0.651, quantum fidelity kernel 0.510, projected kernel 0.497, hybrid 0.649.
- Brier: LR 0.0965, hybrid 0.0976, constant prior 0.1031.
- Pre-registered rule shipped: Hybrid (final blend weight on the quantum kernel w = 0.1; most CV folds chose 0 to 0.3). The quantum kernels alone do not beat LR.
- Why plain fidelity kernels fail: kernel values concentrate as bandwidth grows (right). We tune bandwidth inside CV.
- Plain LR shares 6 of our six (rank correlation 1.00); CV cannot say which list is better.



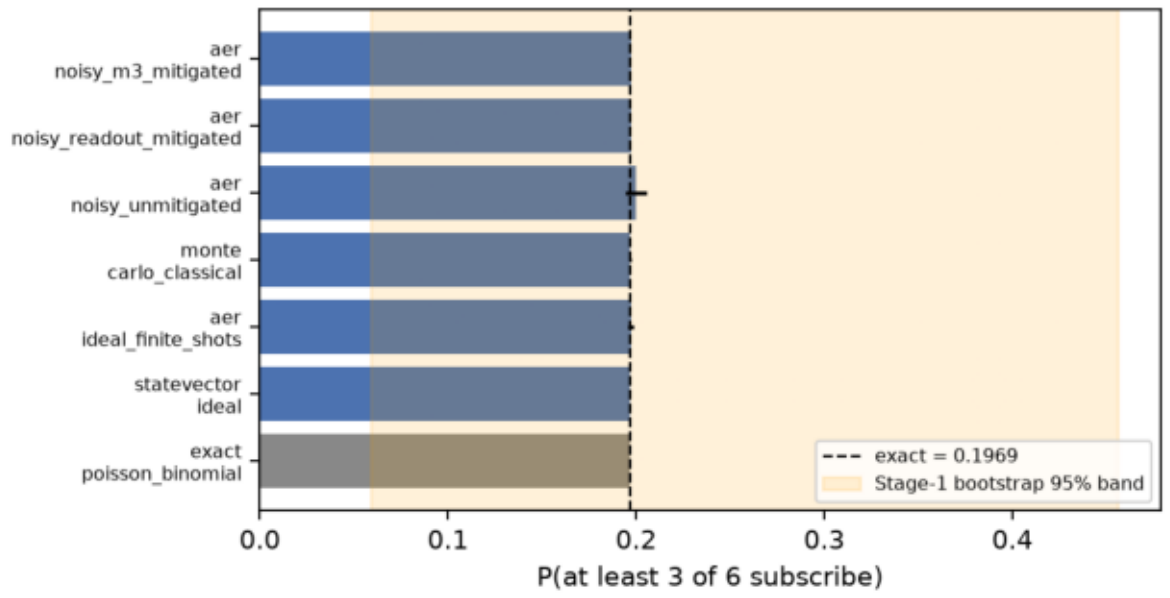
Stage 2: constraint-native QAOA

- max sum_i (p_i D_i - c_i) x_i subject to sum x_i = 6, cost <= 25, capacity <= 22.
- Cardinality: Dicke(10,6) start + XY ring mixer, no penalty. Budget: redundant: the 6 most expensive candidates cost <= 25 (presolve proof). Capacity: redundant (presolve).
- Feasible subsets: 210 of 210. Uniform guess finds the optimum with p = 0.0048.
- QAOA p=3: P(optimum) = 0.127, P(feasible) = 1.000.
- Stress test with binding limits (slack qubits, real ZZ terms): 5 to 8 times more feasible shots than the penalty QUBO, but P(optimum) beats a uniform feasible guess in only 1 of 3 scenarios.



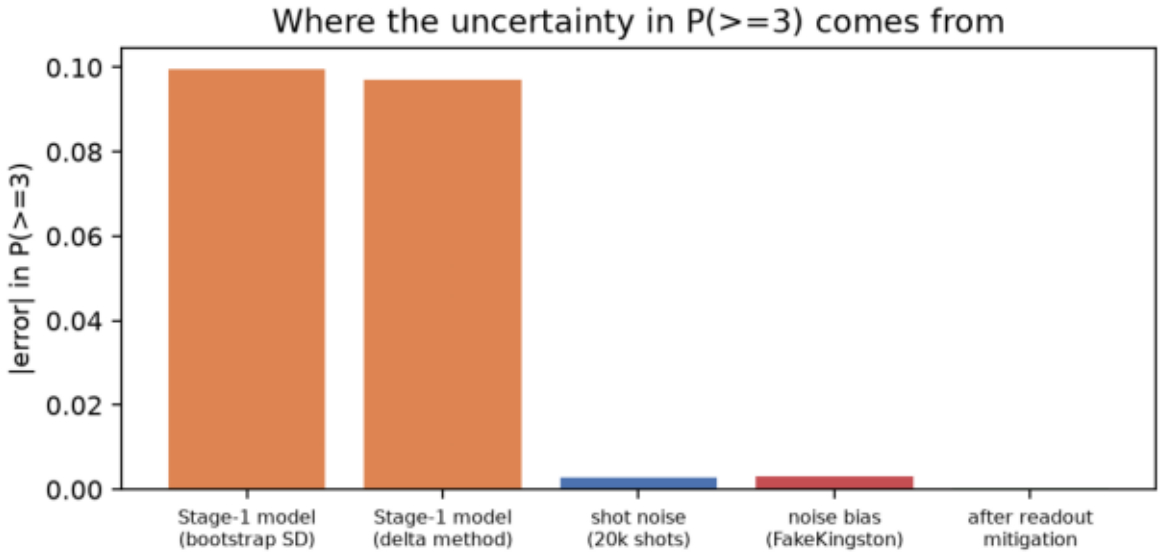
Noise, transpilation and hardware

- Stage 2 p=3 on fake_kingston: depth 1346, 663 two-qubit gates, 119 SWAPs. Under FakeKingston noise, 39% of shots keep weight 6; P(optimum) = 0.031 raw and 0.079 after weight-6 post-selection (uniform guess 0.0048).
- Stage 3 transpiles to depth 5 with no two-qubit gates, so readout error dominates; per-qubit readout inversion removes most of the bias.
- Real-QPU run: scripts/run_hardware.py (one job, about 16k shots) is ready but was not run for this submission.



Error propagation across stages

- Stage-1 model uncertainty dominates: bootstrap SD of P(>=3) = 0.0996 (delta method 0.0970), against shot noise 0.0009 at 200k shots.
- Selection stability: the same 6 customers are chosen in 0% of Stage-1 bootstrap refits; on average 3.5 of 6 overlap.
- So the limit on the answer is 35 positive training examples, not the quantum hardware.



Business reading

- Expected net value of the chosen six: 18.30 (deposit units), using 13/25 budget and 10/22 capacity.
- Hitting 3 conversions from 6 calls is unlikely (about 20%). Plan for 1 to 2 conversions, or widen the list: going down our ranking, 3+ subscribers becomes an even bet at about 17 calls.
- The model is weakly informative with 35 positives; collect outcome data from this campaign before trusting finer rankings.
- Quantum methods here are correct and verifiable, and they match classical answers; we claim no speed-up.