Noise erases depth on today's hardware.

— and a calibrated simulation predicts the device only as an **optimistic ceiling (about ×3)**. We took five trained circuits down a noise ladder that ends on a real IBM processor: same portfolio, certified optimum, **no error mitigation**.

5
trained circuits
replayed unchanged, from clean simulation to the device

100,000
device shots
ibm_kingston, raw readout

44 s
of processor time
for the entire hardware rung

THE SETUP

A ladder from clean simulation to the real device.

Five circuits — four depths plus the adaptive one that escapes wrong seeds — trained once, then **replayed unchanged** at three simulated noise levels, through a twin built from the device's own calibration data, and finally on **ibm_kingston** itself.

Every probability is measured against a **certified optimum**, proven classically before any circuit ran. No error mitigation, no cherry-picking — the raw readout, scored.

Not a forecast — a measurement.

The first number is the routing overhead.

Every circuit grows ×2.5 on the way in

Mapping onto the processor's lattice multiplied the two-qubit gate count by about 2.5 (182 → 440 … 2,190 → 5,454). **The device behaves like a circuit 2.5× deeper than any simulation assumed.**

Only the shallowest circuit survives

One layer keeps a recognisable, flattened readout. **From three layers upward the device reads statistically uniform** — the trained structure is gone.

The escape artist: 0 hits in 20,000

The adaptive circuit puts the optimum on top in simulation (10.1%). On the device, its advantage is **erased completely** — not weakened, erased.

FINDING 3 · THE PREDICTION

The device's own spec sheet is too optimistic — by about ×3.

We simulated the routed circuits with ibm_kingston's **own reported error rates** from the same morning. The prediction came out **2.6–3.5× rosier** than what the device delivered. Published error rates don't contain everything a real processor does to a deep circuit.

And the economics are lopsided: the calibrated prediction cost **17–26 minutes of simulation per circuit**. The device answered in **44 seconds — total** — with ten times the shots, and its answer is the ground truth.

When measuring is this cheap: measure. Treat every calibrated prediction as a best case.

Depth is free in simulation — and worthless under today’s noise.

A number for the hardware question

The adaptive circuit’s wrong-seed escape is worth paying for only at about **99.9% two-qubit fidelity after routing** — not available for this problem class on today’s processors. Now you know what to watch for.

The products print it

The noise ladder now appears **next to the certified gap** in every report, and the calibration-based noise source is labeled what it is: **an optimistic bound, not a forecast**.

Readiness is measurable on demand

44 seconds of processor time bought the whole placement — a monthly free allowance covers a dozen such checks. **The “is hardware ready for us” question has become a cheap, repeatable measurement.**

PRINTED, ON PURPOSE

This is a study about limits — ours included.

One calibration, one day, one backend: the optimism factor is a **single measurement**, and the study says so. Since Revision 1.3 every headline number carries its **finite-shot interval**, recomputed from the committed raw counts — and one imprecise range in the first edition is corrected and disclosed in the revision notes.

Nothing here says quantum hardware is useless — it says **what today's generation delivers on this problem class, measured**, so the next hardware claim that crosses your desk has a baseline to be checked against.

Verdicts you can check beat verdicts you must believe.

THE STUDY

The full evidence is public.

The complete case: the simulated rungs, the calibrated twin, the device tables with **IBM job ids printed**, the finite-shot intervals, and the routing arithmetic — all in the 16-page study.

The study — 16 pages, unedited

qubit-lab.ch/evidence

Study PDF, published benchmark records and the Rapid Quantum Prototypes are on qubit-lab.ch/evidence. Evidence, not promises.