Measuring and simulating quantum circuits

Know which parts of a result come from a model, a sample or a device.

12 min read Dr. Charith Jayasekara

By the end, you'll be able to

- Distinguish statevector predictions, finite-shot samples and hardware evidence.
- Record enough parameters to reproduce a small simulation.

Three layers of evidence

A **statevector simulation** explicitly stores complex amplitudes and applies matrices. It can expose the whole ideal state because it is a classical calculation. A **shot sampler** draws outcomes from its squared amplitudes. **Quantum hardware** adds physical preparation, control, readout and environmental effects.

The first two are available here. The site does not execute hardware jobs, and adding shots to an ideal simulation does not create a hardware noise model.

A reproducible run

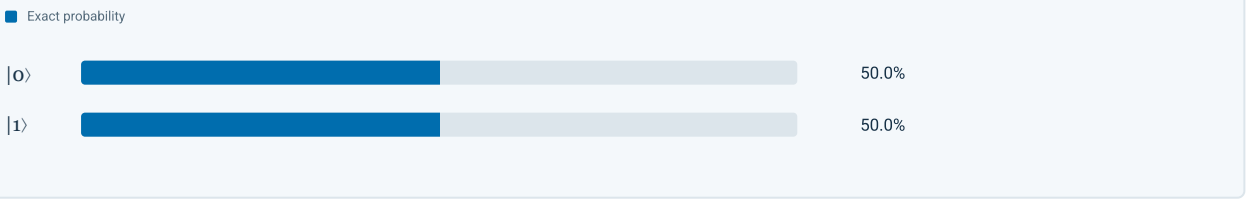
Record the prepared state or circuit, gate convention, qubit order, measured basis, model, number of shots and random seed. Exact amplitudes should agree within numerical tolerance. Seeded counts should reproduce in this implementation.

L15 · INTERACTIVE LAB

Runs on your device

Shots & statistical uncertainty

Exact probabilities are predictions. A finite sample fluctuates. The Wilson interval describes sampling uncertainty for the probability of 1 under independent repeated preparations.



MODEL · Ideal pure-state vector, complex double precision ORDER · q0 is the leftmost bit

Prediction: Will exactly half of 100 shots be 1? What should happen to the uncertainty interval at 10,000 shots? Challenge: Keep the state fixed and change only the seed. Explain why different counts need not indicate a different quantum state. Open the live experiment studio.

A useful statistical check

For a fair binary measurement repeated N times, the expected count is $N/2$ and the count's standard deviation is $\sqrt{N}/2$. At 1,024 shots, that is 16. A difference of a few counts is unsurprising; requiring exactly 512 would be a brittle test.

The displayed Wilson interval estimates uncertainty in P(1). It does not account for correlated errors, drifting preparation or a misspecified physical model.

Know the simulation limit

Generic pure-state storage grows as 2^n. The lab caps its input sizes and performs only small calculations so it remains responsive on a phone. Scaling a classical model is different from estimating a quantum device's useful performance.

Challenge: give one circuit error invisible in a single Z-basis distribution. Explain which additional preparation or readout basis would reveal it.

Reference: IBM Quantum — Quantum mechanics basics.