Measuring and simulating quantum circuits

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

16 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.

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

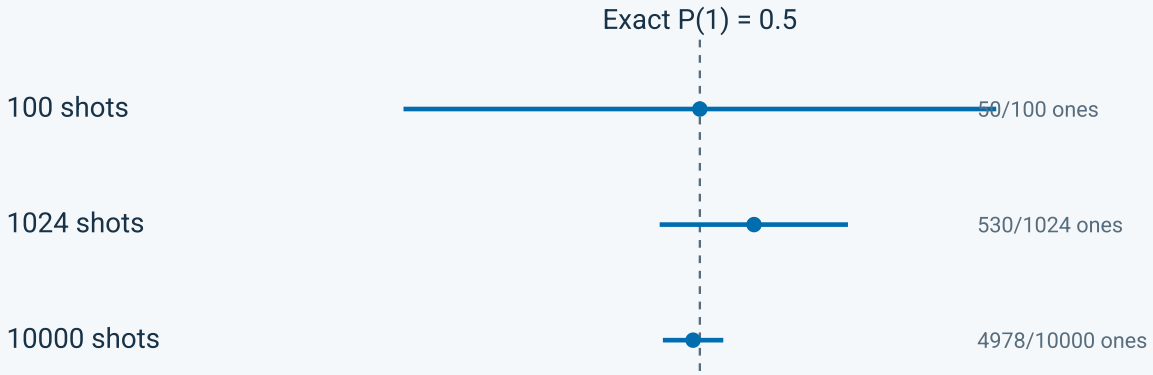
Generic pure-state storage grows as 2^n. The lab caps its input sizes and performs only small calculations so it remains responsive on a phone. Scaling a classical model is different from estimating a quantum device's useful performance.

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

Match each claim to its evidence

WORKED DIAGRAM · F21

More shots narrow uncertainty; they do not change the state



Seeded samples from |+⟩, with 95% Wilson intervals for P(1). The vertical reference is the exact value 0.5. These intervals model independent shots, not systematic hardware error.

The diagram uses a fixed ideal probability and three shot budgets. Each point is a seeded observed frequency, and each horizontal segment is a Wilson interval. More shots generally narrow uncertainty, but the observed point need not approach the exact value monotonically. A larger experiment can fluctuate to the other side of the prediction while still providing a more precise estimate.

A statevector table is evidence about a classical calculation under stated conventions. A shot sample tests the sampler and illustrates statistical variation. Hardware observations add preparation, control, environment and readout effects. Moving from one level to another changes what has been validated. An ideal simulation agreeing with algebra is valuable, but it does not establish the fidelity of a physical implementation.

Worked example: estimate an uncertainty scale

For independent binary outcomes with probability one half, the standard deviation of the observed frequency is one divided by twice the square root of the shot count. At one hundred shots it is five percentage points. At ten thousand shots it is half a percentage point. Multiplying the shot budget by one hundred improves this typical error scale by a factor of ten, not one hundred.

That square-root relationship is why precision can be expensive. Reducing ordinary sampling error by another decimal place generally needs one hundred times as many independent shots. Quantum amplitude-estimation algorithms address a related estimation task using coherent subroutines, but their resource model is different from merely pressing Run with a larger budget. They require controlled operations or other repeated coherent access assumptions.

An interval is useful because a point estimate alone hides the sample size. Seeing 52 percent ones in one hundred shots and in one million shots calls for very different interpretations. The first is unsurprising under a half-probability model; the second would require examining whether the model, independence assumptions or implementation are wrong. Report both the estimate and the evidence supporting its precision.

Reproducibility has several layers

The seed fixes the pseudorandom stream used by this implementation. Keeping the state, shot budget, seed and code version fixed reproduces its counts. Another sampler may use a different random-number generator and produce different valid counts from the same numeric seed. Exact probability agreement and identical sampled records are different reproducibility claims.

Floating-point calculations also need tolerances. A theoretically zero amplitude may be represented by a tiny residual after trigonometric operations. Tests should compare with a justified tolerance and check the norm, rather than demand bit-for-bit equality for every algebraically equivalent sequence. Conversely, tolerances should not be so loose that a wrong gate sign or swapped bit order passes unnoticed.

Understand the classical simulation limit

A general pure statevector stores two complex numbers per qubit-basis index, and the number of indices grows exponentially with qubit count. A double-precision representation commonly uses sixteen bytes per complex entry before overhead. Thirty qubits would already require roughly sixteen gibibytes just for those amplitudes. This is a storage illustration, not a benchmark for this browser tool, which deliberately limits register sizes.

Special structure can change the cost. Product states, stabiliser circuits and some tensor-network-friendly circuits admit more efficient descriptions. Therefore “contains entanglement” and “has many qubits” are not sufficient simulation-hardness arguments. State what family of circuits and outputs a method supports before comparing it with a full statevector simulator.

Run a reproducible sampling study

Start from the default plus preparation. Write an expected count range and run one thousand or roughly one thousand shots. Record the exact probability, observed count, seed and interval. Repeat with another seed while keeping the preparation fixed. Then increase the shot budget and compare interval widths across several seeds rather than selecting only the closest-looking result.

L15 · INTERACTIVE LAB

Runs on your device

Shots & statistical uncertainty

This lab starts with |+⟩. Each shot is a fresh preparation. Exact probabilities, observed frequencies and a Wilson interval describe different aspects of the experiment.

MODEL · Ideal pure-state vector, complex double precision

ORDER · q0 is the leftmost bit

Initial setup · reference state

Exact probability



Statevector before readout · ideal simulation

Basis	Amplitude a + bi	Magnitude	Phase (rad)	Probability
|0⟩	`0.707107 + 0i`	0.707107	0.000000	50.0%
|1⟩	`0.707107 + 0i`	0.707107	0.000000	50.0%

Signs and complex phases are part of the state. Phase is undefined for a zero amplitude. A shared global phase changes no probabilities; relative phases can affect a later operation. Values are rounded to six decimal places.

Prediction to try: Predict approximate counts and plausible fluctuations at this shot budget. Do not demand exact equality.

Challenge: Keep the state fixed, change the seed, then increase the shots. Compare interval widths rather than demanding identical counts.

Open the live experiment studio

Next prepare polar angle 120 degrees. Predict three quarters ones and repeat the study. A relative-phase change alone should not alter these Z probabilities. If you want to test phase, add a suitable gate before readout and predict its effect. An experiment must be sensitive to the feature it is supposed to validate.

Practice: diagnose a false noise claim

Two ideal simulator runs with different seeds give different histograms. A student concludes that the simulator is modelling decoherence. Explain why that inference is unsupported and propose a comparison that separates finite-shot variation from a changed probability model.

Check your reasoning

Different finite samples are expected from a fixed distribution. Compare the exact statevector probabilities before sampling, and repeat many runs at increasing shot budgets. Sampling variation shrinks in typical frequency scale as shots increase. A physical-noise model would need to specify how it changes preparation, evolution or readout; a changing seed alone does not introduce that model.

Reference: IBM Quantum — Quantum mechanics basics.

Check your understanding

PAUSE & THINK

What happens to typical frequency uncertainty when shots increase by a factor of 100?

- ☐ It falls by a factor of 100
- ☐ It falls by about a factor of 10
- ☐ Every sample becomes exact

Answer: It falls by about a factor of 10. Independent sampling uncertainty scales approximately as 1/√shots. A particular run need not move monotonically closer.