Probability, amplitudes and normalisation

Probabilities add to one. Amplitudes can do something more interesting.

16 min read Dr. Charith Jayasekara

By the end, you'll be able to

- Convert complex amplitudes into measurement probabilities.
- Validate and normalise a nonzero state vector.

The rule that connects state to outcome

For $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, the Born rule gives $P(0) = |\alpha|^2$ and $P(1) = |\beta|^2$. The state must be normalised:

$$|\alpha|^2 + |\beta|^2 = 1.$$

An amplitude can be negative or complex. A probability is real and lies between zero and one.

A complete numerical example

Take amplitudes $\alpha = 1/2$ and $\beta = i\sqrt{3}/2$. Their squared magnitudes are $1/4$ and $3/4$, so the vector is normalised. Across 1,000 independent preparations, the **expected** counts are 250 and 750. A sampled result will usually differ.

If instead someone gives you the vector $(1, i)^T$, its squared norm is 2. Divide both amplitudes by $\sqrt{2}$ to obtain a valid unit vector. You cannot normalise the zero vector because its norm is zero.

Add amplitudes before squaring

When indistinguishable paths contribute to the same outcome, their amplitudes add. Squaring afterwards introduces interference terms:

$$|a + b|^2 = |a|^2 + |b|^2 + 2\,\mathrm{Re}(a^* b).$$

For $a = 1/2$ and $b = -1/2$, the total amplitude is zero. Adding the two separate squared magnitudes would incorrectly give $1/2$.

Practise

In the qubit explorer, find a state with $P(1) = 3/4$. Change only its phase. Then switch to the interference lab and explain why relative phase can affect a later measurement even when the current Z probabilities do not change.

Use a probability audit before applying gates

WORKED DIAGRAM · F06

Different signs, the same squared magnitudes



The first two panels show real amplitudes; the last shows probabilities. Squaring removes the sign. Later interference can reveal the relative sign that a Z measurement misses.

The two states in the diagram have the same Z probabilities. One has a positive second amplitude and the other a negative one. The probability audit passes for both, but the states are still different. Normalisation answers whether a vector is a valid candidate state; it does not tell you that two valid vectors describe the same physical preparation.

For any statevector, calculate the squared magnitude of every entry and add the results. The sum must be one within a justified numerical tolerance. A tiny deviation caused by floating-point rounding is different from a badly normalised input. If the total is zero, normalisation is impossible: dividing the zero vector by a norm cannot create a direction. If the total is positive but not one, divide all entries by the same norm, not by the total squared norm.

Worked example: normalise a complex vector

Take the unnormalised column with entries one plus i and two. The first squared magnitude is two and the second is four, giving a total of six. The norm is therefore the square root of six. Dividing both amplitudes by that norm gives a valid state:

$$|\psi\rangle = \frac{1+i}{\sqrt{6}}|0\rangle + \frac{2}{\sqrt{6}}|1\rangle, \qquad P(0) = \frac{1}{3}, \quad P(1) = \frac{2}{3}.$$

Notice that dividing each entry by six would leave the squared norm equal to one sixth. Also notice that dividing each component by its own magnitude would destroy the relative weights. Normalisation is a single common rescaling of the whole vector. It preserves the ratios and relative phases that determine the state's direction.

This operation is a mathematical way to describe a state, not a free physical gate that can arbitrarily amplify selected amplitudes. A deterministic quantum operation on an unknown input must satisfy the relevant physical constraints. Conditional filtering can change relative weights, but the success probability and failure outcomes then belong in the model. A simulator should not silently renormalise every mistake and hide a nonunitary gate implementation.

Expected counts do not specify a sample

If the probability of one is three quarters and you perform one thousand fresh preparations, the expected number of ones is 750. The count is a random variable. Its standard deviation in the independent binomial model is the square root of the shot count times the probability times its complement, approximately 13.7 here. A count such as 742 is compatible with the model; a demand for exactly 750 misunderstands expectation.

The word “independent” is doing work. Measuring a single collapsed qubit repeatedly is not the same experiment as preparing the original state again for each shot. Correlated noise, changing calibration or a drifting source can also invalidate the simple independent-sampling model. In this book, seeded sampling draws from a fixed ideal distribution. Its uncertainty describes finite sampling rather than a simulation of every physical imperfection.

Phase can survive a probability check

For a real state with amplitudes one half and square root of three over two, applying H gives a zero probability of one half plus square root of three over four, about 93.3 percent. Negating the second input amplitude changes that to one half minus square root of three over four, about 6.7 percent. Before H, both states had probabilities one quarter and three quarters. The large difference afterward comes from retaining signs during amplitude addition.

This is why a useful debugging record contains more than a histogram. Check the input norm, complex entries, gate ordering, final norm and measurement basis. Then compare the resulting probabilities and sampled counts. Each layer catches a different class of error, and a correct total probability alone cannot establish the entire calculation.

Connect the numerical state to the lab

This explorer starts at polar angle 120 degrees and relative phase 90 degrees. Predict the amplitudes one half and i times square root of three over two. Reveal the table and check that the probability column reads one quarter and three quarters. Then change the relative phase to zero while leaving the polar angle fixed. Predict which columns change immediately and which remain fixed.

L01 · INTERACTIVE LAB

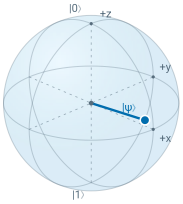
Runs on your device

Qubit explorer

Prepare a pure qubit, apply gates, then compare amplitudes, phases and Z-basis probabilities. Reset restores the initial preparation shown by the lesson’s controls.

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

Initial setup · reference state



Exact probability

|0⟩ 25.0%

|1⟩ 75.0%

Statevector before readout · ideal simulation

Basis	Amplitude a + bi	Magnitude	Phase (rad)	Probability
|0⟩	`0.5 + 0i`	0.500000	0.000000	25.0%
|1⟩	`0 + 0.866025i`	0.866025	1.570796	75.0%

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 the amplitudes and Z-basis probabilities for the preparation and gates you have chosen. If starting from |0⟩, what will H do?

Challenge: Prepare θ = 120°, φ = 90°. Predict both complex amplitudes. Then change φ without changing θ.

Open the live experiment studio

Apply H after the zero-phase preparation and compare the output with the calculation above. Reset before testing another phase, so you are comparing the same original magnitudes. The angle sliders describe a new preparation; a gate button transforms the currently prepared state. Mixing those actions without recording them makes an experiment difficult to reproduce.

Practice: repair a state description

A student writes amplitudes three and negative four, then labels the probabilities three sevenths and four sevenths. Find the correctly normalised amplitudes and probabilities. Explain what information the negative sign retains.

> **Check your reasoning**
>
> The norm is five, so the amplitudes are three fifths and negative four fifths. Their probabilities are nine twenty-fifths and sixteen twenty-fifths. Dividing the absolute amplitudes by their sum would use the wrong rule. The negative sign is a relative phase of pi and can affect a later interference measurement, although it does not make the probability negative.

Reference: IBM Quantum — Quantum information.

Check your understanding

PAUSE & THINK

How do you normalise the nonzero column $(3, -4)^T$?

- [ ] Divide both entries by 7
- [ ] Divide both entries by 5
- [ ] Square the entries and keep their signs

Answer: Divide both entries by 5. The norm is $\sqrt{(9 + 16)} = 5$. The normalised amplitudes are 3/5 and −4/5; probabilities are 9/25 and 16/25.

PAUSE & THINK

A qubit has amplitudes 1/2 and i√3/2. What is P(1)?

- [ ] √3/2
- [ ] 3/4
- [ ] 1/2

Answer: 3/4. Take the squared magnitude: $|i\sqrt{3}/2|^2 = 3/4$. The other outcome has probability 1/4, so the state is normalised.