How to Interpret a Proportion Confidence Interval

A confidence interval for a proportion behaves differently from one for a mean, especially near the boundaries of 0 and 1. This page explains why and how to read it.

What this answers

This page answers "why does my proportion confidence interval use the Wilson method, and how do I read it, especially when I observed zero or all events?" A proportion is bounded between 0 and 1, which creates edge behavior that a simple normal-approximation interval handles poorly.

Why the Wilson interval is the default

The simplest proportion interval formula, based directly on the normal approximation, can produce a lower bound below 0 or an upper bound above 1, values that are impossible for a proportion, and it performs especially poorly with small samples or proportions near the boundaries. The Wilson interval adjusts the calculation to stay within the valid 0 to 1 range and maintains closer to the stated confidence level across a much wider range of sample sizes and observed proportions, which is why it is the default method here rather than the simpler normal approximation.

Reading events, trials, and the interval

A proportion confidence interval is built from a count of events and a count of trials, for example 12 successes out of 50 attempts. The resulting interval expresses a range of proportions reasonably compatible with that observed rate at the chosen confidence level. A narrower interval reflects either a larger number of trials or a proportion further from 0.5, since variability is highest for proportions near 0.5 and lowest near the boundaries.

The zero-events case

Observing zero events out of a number of trials does not mean the true rate is exactly zero; it means the interval's lower bound is 0 and the upper bound reflects the largest rate still plausible given that no events occurred in this many attempts. A Wilson interval, unlike the simple normal approximation, produces a sensible, non-degenerate upper bound even in this edge case, which is one of the clearest reasons to prefer it.

Worked example

Observing 3 successes in 40 trials gives a Wilson 95 percent interval of roughly 0.010 to 0.198 for the true proportion, comfortably within the valid range. Observing 0 successes in 40 trials gives an interval of roughly 0.000 to 0.088, correctly showing that a true rate up to about 9 percent remains plausible even though no events were observed, rather than implausibly asserting the rate is exactly zero.

Assumption audit

> **Calculated from your data:** the exact Wilson interval bounds, once you enter your event and trial counts into the linked calculator.

> **Evidence to review:** whether your trials are independent and identically distributed, an assumption behind the binomial model this interval relies on.

> **You must verify:** that your sample of trials was collected in a way that fairly represents the population you want to describe, since no interval method can correct for a biased sampling process.

Source

This guidance follows the Wilson score interval treatment in the NIST/SEMATECH e-Handbook of Statistical Methods and the shared statistical reasoning contract every StatReason engine is built against.

Limitations

This page covers a single proportion; comparing two proportions or estimating a difference between them uses a related but distinct interval, covered on its own engine page.

Next action: use the **Proportion Confidence Interval Calculator**, or run an exact test with the **Exact Binomial Test**.