

Adopting the R Software Engine: `binom.test()`

Coding Corner: `binom.test`

Head to Section 9.5.1 in your textbook to learn the syntax of the function, `binom.test`.

Barry the Statypus-in-Training says:

Wait, I just calculated standard deviations under the null hypothesis, sketched Normal curves, and found areas to get a single p -value. Are you telling me there's an easier way?

Bill the Statypus says: You had to experience the “Friction of the Pencil” to learn how sample evidence compresses our decision structures. But in real research, data scientists don't draw smooth bell curves by hand; they hand the counts over to the computer.

Sally the Statypus says: Exactly! `binom.test()` is specifically engineered to perform hypothesis tests and calculate confidence intervals for a population proportion all at once.

1. Evaluating Melody's Footwear Predicament

Recall the little girl Melody, who struggled to put her shoes on the correct feet. Her father observed a sample of 40 attempts where Melody was successful 13 times. He claimed that her success rate was strictly less than 0.5 (meaning $H_0 : p = 0.5$ vs. $H_a : p < 0.5$). Write down the line of Base R code required to execute this entire hypothesis test against the null value of $p_0 = 0.5$. Make sure to specify the direction of Melody's alternative hypothesis! Record the displayed p -value as well.

Reflection: Textbook Investigation Checkpoint

Melody's manual calculation using a the Code Template gave her a slightly different probability of $p\text{-value} = 0.02888$. Go to Section 9.4 of your textbook and read the commentary regarding exact methods, and summarize why the `binom.test` is superior to our previous estimates.

Bill the Statypus says: Time to tell the full story, go read Section 9.4 of your book.

2. Reinvestigating the 2024 Gender Ratio

In the last worksheet, we saw that in 2024 there were 1,774,965 female babies born out of a total of 3,628,934 births. Place these 2 values into `binom.test` as `x` and `n`.

1. Copy the line of code you used below and run the code.
2. Copy the two lines of output that begin with “number of successes” and “alternative hypothesis.”
3. Circle the p -value found and translate the syntax given for H_a .

Bill the Statypus says: That p -value result should look familiar.

4. Copy the next two lines of the output that you haven’t copied down yet.
5. Interpret the meaning in context and rounding values to the nearest hundredth of a percent.

Barry the Statypus-in-Training says:

Wait, it just did both a confidence interval and hypothesis test with that one easy line of code?

Sally the Statypus says: Yep! You know the theory, so use the “Exact binomial test” to do inference on a population proportion. Make sure to use `alternative` and `conf.level` when necessary!

Barry the Statypus-in-Training says:

That means I just have to focus on interpreting what the output means!

Bill the Statypus says: (coyly smiling) That’s right... computers compute and we think.