

Testing Worksheet 4

Question 1 – Conceptual

Part A: What does a test-statistic quantify?

Part B: What does the ‘significance level’ refer to in hypothesis testing?

Part C: What is another name for a Type I error?

Part D: What is another name for a Type II error?

Part E: Explain why using $\alpha = 0.05$ as a strict cut-off for significance is problematic.

Question 2 – Decision Error

Part A: What type of error occurs if H_0 is False but we fail to reject it?

Part B: What type of error occurs if H_0 is True but we do reject it?

Part C: How does the sample size affect the test-statistic, and as a consequence the p-value, with everything else held constant?

Part D: How does the standard error affect the test-statistic, and as a consequence the p-value, with everything else held constant?

Conclusions for the ‘Decision Making’ Approach

Making conclusions with this approach will differ only slightly from what we did using ‘Strength of evidence.’ Use the significance level α you are given in a problem and compare the p-value to this. (If a significance level is not given use best judgement to pick one)

You will answer the following question:

Do we reject H_0 at significance level α ?

If p-value $\leq \alpha \rightarrow$ **Yes**

If P-value $> \alpha \rightarrow$ **No**

Then we apply this answer to the context (what does H_0 and H_A actually say?) to answer the research question. Usually I like to see the following in the conclusion:

- a statement of the test-statistic value and the p-value
- takeaway from the p-value comparison to α
- what that actually means in context of the study

Example: Idea: With a test-statistic value of $Z = 1.62$ and a p-value of .104, we fail reject the null hypothesis that the coin is biased.

Write up: The data yields a test-statistic of $Z = 1.62$ and a p-value of .104. The data is not sufficient to indicate that the coin is biased at a 0.05 significance level.

Alternative: The data yields a test-statistic of $Z = 1.62$ and a p-value of .104. There is no statistically discernible bias in the coin at the 0.05 significance level.

Question 3 – Piano Revisited

Georgianna claims that in a small city renowned for its music school, the average child takes less than 5 years of piano lessons. We have a random sample of 35 children from the city, with a sample mean of 4.6 years of piano lessons and a sample standard deviation of 2.2 years.

Evaluate Georgianna’s claim using a hypothesis test at the 0.05 significance level.

The test-statistic and p-value from the previous worksheet are

$$T = \frac{4.6-5}{2.2/\sqrt{35}} = -1.07 \text{ and } 0.145 \text{ respectively.}$$

Question: Do we reject the null hypothesis? Write a conclusion to summarize our results. What are type-I and type-II errors in the context of this study?

Practical vs Statistical Significance

We have so far been talking about ‘statistical discernibility’ when it comes to test-statistics and p-values. What we are doing is basically going ‘our results are so many standard errors away from the hypothesized value.’

In essence, we are saying there is (or there is not) a difference based on our data. But even if we reject a null hypothesis, it does not mean that an effect has any practical difference in the real world.

Examine the following scenarios related to coin flipping we’ve seen before. For both we have the same null and alternate hypotheses:

H_0 : $p = .5$ (the coin is fair), H_A : $p > .5$ (coin is biased towards heads)

Scenario 1

We flip the coin 20 times and get 16 heads. (ignore conditions for now)

$$Z = \frac{.8 - .5}{\sqrt{\frac{.5 \times .5}{20}}} = 2.683, \text{ p-value} = 0.0036$$

Scenario 2

We flip the coin 1 million times and get 501342 heads.

$$Z = \frac{.501342 - .5}{\sqrt{\frac{.5 \times .5}{1,000,000}}} = 2.684, \text{ p-value} = 0.0036$$

In both scenarios we get nearly the same test-statistic and the same p-value rounded to 2 decimal places. In the real world, are you going to notice the affect of the biased coin in scenario 1? Probably. What about scenario 2? Probably not.

If we want an idea of practical significance – “HOW biased is the coin,” we can pair hypothesis tests with confidence intervals.

Scenario 1

```
CI_lower = .8 - 1.96*sqrt(.8 * .2 / 20)
CI_upper = .8 + 1.96*sqrt(.8 * .2 / 20)
paste(round(CI_lower, 4), round(CI_upper, 4))
```

```
## [1] "0.6247 0.9753"
```

Scenario 2

```
CI_lower = .501342 - 1.96*sqrt(.501342 * (1-.501342) / 1000000)
CI_upper = .501342 + 1.96*sqrt(.501342 * (1-.501342) / 1000000)
paste(round(CI_lower, 4), round(CI_upper, 4))
```

```
## [1] "0.5004 0.5023"
```

Both scenarios have ‘statistical discernibility’ since there are low p-values. In scenario 1 we can see an estimated proportion of heads much larger than the .5 proportion of heads for a fair coin.

But in scenario 2 the actual difference between the CI estimates and a fair coin proportion of 0.5 is basically nothing. You will not notice the effects of the coin bias. A bias this small won’t affect the # of heads flipped compared to a fair coin until you flip the coin, on average, 435 times.

```
1 / (.5023 - .5)
```

```
## [1] 434.7826
```

Effect Size

The actual difference between the statistic and the hypothesized value (after removing the effect of sample size) is called the **effect size**. In lieu of having confidence intervals that give you a range of estimates of the actual parameter, seeing reports of the effect size can help us judge practical significance too. There are a few different ways to calculate effect sizes, but all seek to remove the effect of sample size values from saying results are meaningful. We are comparing differences in raw variability now. The formula for proportions is... not easy. For means just leave out n in the standard error.

effect size (means): $\frac{\bar{x} - \mu_0}{s}$ or $\frac{\bar{x}_1 - \bar{x}_2}{s_{pool}}$

effect size (proportions): $2[\sin^{-1}(\sqrt{\hat{p}}) - \sin^{-1}(\sqrt{p_0})]$ or $2[\sin^{-1}(\sqrt{\hat{p}_1}) - \sin^{-1}(\sqrt{\hat{p}_2})]$

The following are some rough estimates for interpreting effect size magnitudes (ignoring signs).

- 0: negligible
- 0.2: small
- 0.5: medium
- 1+: large

Scenario 1: The effect size is $2[\sin^{-1}(\sqrt{0.8}) - \sin^{-1}(\sqrt{0.5})] = 0.64$, which would be a medium effect size, noticeable in the real world

Scenario 2: the effect size is $2[\sin^{-1}(\sqrt{.501342}) - \sin^{-1}(\sqrt{0.5})] = 0.003$, which is negligible