t-distributions

November 2025

Motivation

The last two slide decks of been almost completely unreasonable.

Why?

Motivation

The last two slide decks of been almost completely unreasonable.

Why?

We know the population variance but not the population mean

Motivation

That's not really reasonable.

- Generally if we know one parameter we know the other
- Alternatively, if we don't even know the average value how would we know the spread?

We did it initially as a simplification step because it let's us use the normal distribution

- Only unknown in our sampling distribution was the mean, μ

Where to go?

Ideas on what we can do?

Where to go?

Ideas on what we can do?

Plug in our sample variance, s^2?

Don't forget to ask the question Vinny!

Where to go?

Ideas on what we can do?

Plug in our sample variance, s^2?

Yes! Buuut....

- Our sample variance itself is an estimate (of the population variance)
- Estimates (test statistic, confidence intervals) are built on top of estimates (estimated variance)
- Used the same data for both (eg double dipped)
  - AND we even have to use the sample mean to calculate the sample variance to then test the sample mean.....

Is that a problem?

Is that a problem?

Yes, yes it is a problem.

After a million simulations my 95% confidence interval calculated using the sample variance my coverage rate (intervals that overlap the mean) is down to 94.01%.

It's a small difference but it's a consistent problem

In steps...William "Student" Gosset

Solution

We can't use the normal distribution because it's optimistic

- Too optimistic (overly small intervals)
- Assumes we know a parameter that we actually have to estimate

Gosset's big breakthrough...

- We can't use the normal distribution
- t-distribution is better (mathematically correct)
  - t-distribution is normal's shorter, fatter sibling
  - Requires "degrees of freedom" (discussed in a minutes)

Student's t-distribution

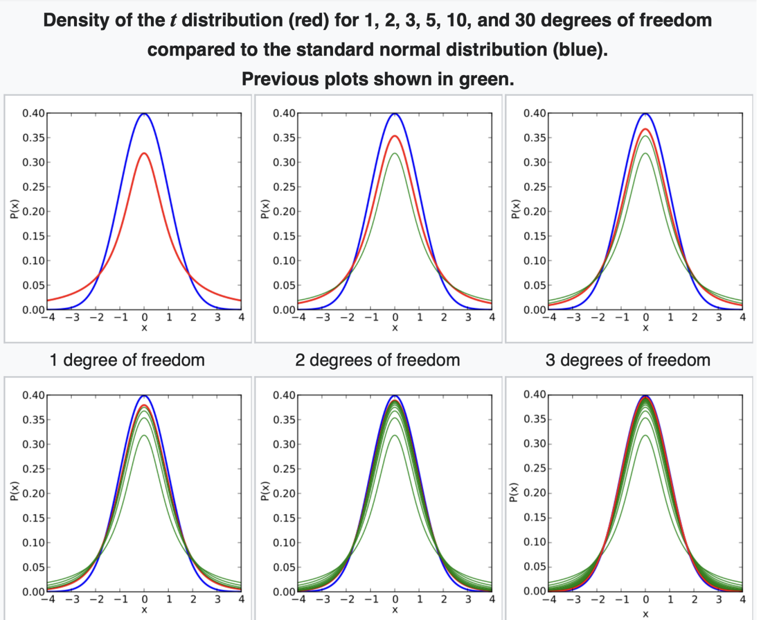

Degrees of Freedom?

So what *is* degrees of freedom?

- It describes the shape of the t-distribution
  - Similar how the normal distribution needs a mean and variance
- It's an esoteric concept but there is a classic example
  - Sample = (-5, 2, 5, ??)
  - $\bar{x} = 0$
  - You can calculate ??
- For the t-distribution for a single mean (what we have) it's

$$d.f. \; = \; n \; - \; 1$$

- Literally just 1 less than your sample size
  - Does get more complicated later on

So what changes for us?

Not a whole lot....

Hypothesis Testing

- Our test static now has a t dist and not a normal dist.
- We plug in s^2 for σ^2
- Have to calculate degrees of freedom

$$\frac{mean - hypothesis}{\sqrt{s^2/n}} \tag{1}$$

Confidence intervals

$$estimate \pm (distributional\ value)(st.error)$$

$$\bar{x} \pm t^*_{df}(\sqrt{s^2/n})$$

t-test Example: Penguins! Again

Over the course of a few years in the mid-2000's a collection of body measurements were taken on 344 penguins, randomly selected, on some islands off of Antarctica.

t-test Example: Set Up

My uncle: *The average penguin flipper length? I don't know Vinny I'd guess a foot long maybe*

Let's try to disprove my uncle's (unwilling) claim that a penguins flipper is about a foot long (305 mm).

Hypothesis: Same as for a t-test

You must list exactly two hypothesis statements, the null and the alternative.

$$\begin{aligned} H_0 : \ \mu_{flipper} &= 305 \\ H_A : \ \mu_{flipper} &\neq 305 \end{aligned}$$

Here we are trying to show there is a difference; that RJ's claim they the flippers are a foot long is just wrong.

Sampling Distribution's Mean for a Hypothesis Test

Critical: During a hypothesis test we assume we *know* the population mean (via H_0)

- Z-test or t-test
- We assume the sampling distribution is centered at the population's mean
- We make some claim about the pop's mean in H_0
- So plug in the hypothesized value from H_0 into the sampling distriubtion

When we judge whether the results we got are unlikely or not; we are judging in effect how believable H_0 is.

Ex: Penguin Sampling Distribution

The real sampling distribution then is N(μ, σ^2/n) which, when we plug in s^2 for σ^2 is

$$\bar{X} \sim N(\mu, 195/344)$$

where

- μ is the mean of our population
  - We can estimate it with out sample mean (200)
  - We can hypothesize it's value
- 195 is our SAMPLE variance
  - Assumed known before we ever started talking about penguins
  - Very unrealistic short of an ornithologist savant
- 344 is our sample size

Ex: Penguin's Test Statistic

For a t-test, the test statistic is the standardized observation so....

- The sample mean was 200
- Hypothesized mean is 305
- And the standard deviation for our sampling distribution is $(\frac{s^2}{n})^{1/2} = (195/334)^{1/2} = .7341$
  - This is called the estimated **standard error**; it's the standard deviation of the sampling distribution

$$\frac{\bar{X} - \mu}{\frac{s}{\sqrt{n}}} = \frac{200 - 305}{.753} = -139.46$$

p-value and decision

Need our degrees of freedom...

$$n - 1 = 344 - 1 = 343$$

pt(-139.26, df = 343) = 3.172628e-304 ≈ 0

Decision: We have very strong evidence the mean flipper length is not 305mm.