Null distribution = sampling dist. for your sample statistic in world where $H_0$ holds true.

Stat 145, Mon 4-Oct-2021 -- Mon 4-Oct-2021

---

Monday, October 4th 2021

---

Wk 6, Mo
Topic:: Hypothesis test intro
Read:: Lock5 4.1
HW:: Moodle Quiz Ch. 3 ends Wed.

R syntax obervations:
- only a few of commands we've seen operate directly on data frames
  nrow()
  head()
  names()
  dim()
  with(), filter()  (both of these need extra stuff)

Hypothesis test
- another inference procedure
- idea cultivated using single proportion
  multiple choice test, 5 answers available
  let p = proportion (in population) that have correct answer B

  you'd expect p = 1/5 (null hypothesis)

  data from a sample of questions used to assess strength of evidence against

  Say, in 40 questions, 10 have B for correct answer
    this provides sample data
    identify possible sample statistics

- construction of approximate null distribution
  null distribution is a special case of sampling distribution
    assumes null hypothesis is true
  for univariate data (p-hat or x-bar), it is similar to bootstrapping
    not actually null dist; Lock's call it a "randomization distribution"
  slips in bag are given proportion equal to value in null hypothesis

- example of doing this using R
  nullDistPhat <- do(5000) * rflip(40, prob=.2)

```
    gf_density(~prop, data=nullDistPhat)
    nrow( filter(nullDistPhat, prop >= .25) )
       counts the number of dots in right tail out at/beyond phat=0.25
       divide this by 5000 and double it to get P-value
```

Statistical significance

- corresponds to having a result that is unlikely under null hypothesis

Another example: mean body temperature

# Hypothesis testing

General Framework:

1. State hypotheses
2. Collect data, compute test statistic
3. Construct approximate null distribution

   Locate your test stat. on that distribution

   Compute how frequently something as extreme as your test statistic arises (P-value)
4. Draw a conclusion (Is it statistically significant or not?)

Examples from yesterday: Univariate

1. Categorical data - answers A-E to multiple choice questions

   Focused on B as correct answer (yes/no)

   (i) Hypotheses: $H_0: p = 1/5 = 0.2$

   $H_a: p \neq 1/5$

   (ii) Gather data (sample)

   $n = 40$, get that B is correct $\hat{p} = 10/40 = 0.25$

   (iii) Construct approx null dist.

   RStudio: rflip(40, p = 0.2)

   so as to be sampling from world where $H_0$ holds

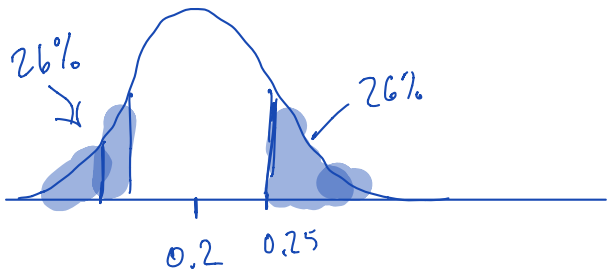


P-value ~ 0.52

Not statistically significant

Looks consistent w/ $H_0$.

Predetermined level $\alpha$

typical values are $\alpha = 0.1,\ 0.05,\ 0.01$

Select it prior to gathering data, agreeing to reject $H_0$ if your P-value is below $\alpha$

Otherwise, say "we fail to reject $H_0$."

2. Body temps

(i) $H_0: \mu = 98.6$

$H_a: \mu \neq 98.6$

(ii) data $n = 50,\ \bar{x} = 98.26$

(iii) Construct approx. null dist.

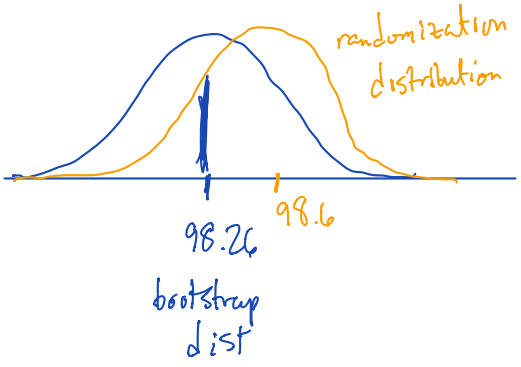


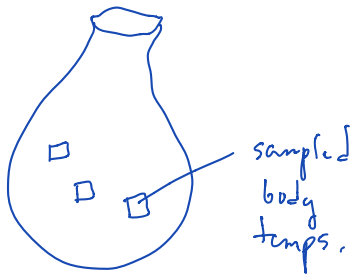


One side amounts to 0.00005

$\rightarrow$ P-value is $2(0.00005) = 0.0001$

Bivariate scenarios

Is there a difference in mean number of finger taps

(i) $H_0: \mu_1 - \mu_2 = 0$

$H_a: \mu_1 - \mu_2 > 0$