Statistical Methods and Data Analysis I
Lecture 12: Confidence Intervals.

Oleg Goldshmidt
oleg.goldshmidt@post.idc.ac.il

Arison School of Business
Interdisciplinary Center (IDC)
Herzliya, Israel

April 28, 2019

Confidence Interval: Analysis (I)

Confidence Interval: Analysis (I)

- What does *“95% confidence”* mean?

Confidence Interval: Analysis (I)

- What does *"95% confidence"* mean?
  - There is a *5%* probability that the actual result will lie *outside* of the confidence interval — *at random*! (Think of p-value.) Is it good enough?

Confidence Interval: Analysis (I)

- What does *"95% confidence"* mean?
  - There is a *5%* probability that the actual result will lie *outside* of the confidence interval — *at random*! (Think of p-value.) Is it good enough?
- Elections are rare. Consider a business problem: you are launching a new product. You know that the gender ratio of your customers is 1:1 with a high precision. At the end of the first day 4390 customers — 2,262 women and 2,128 men — have bought the new product. Can you conclude that the new product is more appealing to women than to men?

Confidence Interval: Analysis (I)

- What does *"95% confidence"* mean?
  - There is a *5%* probability that the actual result will lie *outside* of the confidence interval — *at random*! (Think of *p*-value.) Is it good enough?
- Elections are rare. Consider a business problem: you are launching a new product. You know that the gender ratio of your customers is 1:1 with a high precision. At the end of the first day 4390 customers — 2,262 women and 2,128 men — have bought the new product. Can you conclude that the new product is more appealing to women than to men?

$$P_s = \frac{2262}{4390} = 0.515$$

$$P_p = 0.515 \pm 1.96\sqrt{\frac{0.515 \times 0.485}{4390}} = 0.515 \pm 0.0148 = (50\%, 53\%)$$

Confidence Interval: Analysis (I)

- What does *"95% confidence"* mean?
  - There is a *5%* probability that the actual result will lie *outside* of the confidence interval — *at random*! (Think of *p*-value.) Is it good enough?
- Elections are rare. Consider a business problem: you are launching a new product. You know that the gender ratio of your customers is 1:1 with a high precision. At the end of the first day 4390 customers — 2,262 women and 2,128 men — have bought the new product. Can you conclude that the new product is more appealing to women than to men?

$$
\begin{aligned}
P_s &= \frac{2262}{4390} = 0.515 \\
P_p &= 0.515 \pm 1.96\sqrt{\frac{0.515 \times 0.485}{4390}} = 0.515 \pm 0.0148 = (50\%, 53\%)
\end{aligned}
$$

- Rule of thumb: daily measurements will fall outside of 95% CI every 4 weeks if we are closed on weekends, every 3 weeks if we work weekends. Is it good enough?

Confidence Interval: Analysis (I)

- What does *"95% confidence"* mean?
  - There is a *5%* probability that the actual result will lie *outside* of the confidence interval — *at random*! (Think of *p*-value.) Is it good enough?
- Elections are rare. Consider a business problem: you are launching a new product. You know that the gender ratio of your customers is 1:1 with a high precision. At the end of the first day 4390 customers — 2,262 women and 2,128 men — have bought the new product. Can you conclude that the new product is more appealing to women than to men?

$$P_s = \frac{2262}{4390} = 0.515$$
$$P_p = 0.515 \pm 1.96\sqrt{\frac{0.515 \times 0.485}{4390}} = 0.515 \pm 0.0148 = (50\%, 53\%)$$

- Rule of thumb: daily measurements will fall outside of 95% CI every 4 weeks if we are closed on weekends, every 3 weeks if we work weekends. Is it good enough?
- $95\% \rightarrow 2\sigma$... $3\sigma(99.7\%)$ — about once a year... $4\sigma(99.994\%)$ — twice in a lifetime...

Computing CI in R (file `ci_simple.R`)

```
ci.simple <- function(n,p,conf=0.95) {
    boot::norm.ci(t0=p,var.t0=p*(1-p)/n,conf=conf) # What's normal about it?
}
ci.simple.width <- function(n,p,conf=0.95) {
    interval <<- ci.simple(n,p,conf)
    apply(interval,1,function(x) x[3]-x[2])
}
```

Toolkit: Computing Confidence Intervals

Computing CI in R (file `ci_simple.R`)

```
ci.simple <- function(n,p,conf=0.95) {
    boot::norm.ci(t0=p,var.t0=p*(1-p)/n,conf=conf) # What's normal about it?
}
ci.simple.width <- function(n,p,conf=0.95) {
    interval <<- ci.simple(n,p,conf)
    apply(interval,1,function(x) x[3]-x[2])
}
```

```
> source("ci_simple.R")
> ci.simple(4390,0.515,0.95)
     conf
[1,] 0.95 0.5002161 0.5297839
> ci.simple.width(4390,0.515,0.95)
[1] 0.02956789
> ci.simple(4390,0.515,c(0.95,0.99,0.999,0.9999,0.99999))
        conf
[1,] 0.95000 0.5002161 0.5297839
[2,] 0.99000 0.4955706 0.5344294
[3,] 0.99900 0.4901797 0.5398203
[4,] 0.99990 0.4856534 0.5443466
[5,] 0.99999 0.4816814 0.5483186
> ci.simple.width(4390,0.515,c(0.95,0.99,0.999,0.9999,0.99999))
[1] 0.02956789 0.03885879 0.04964067 0.05869321 0.06663719
```

Toolkit: Computing Confidence Intervals

Computing CI in R (file ci_simple.R)

```
ci.simple <- function(n,p,conf=0.95) {
    boot::norm.ci(t0=p,var.t0=p*(1-p)/n,conf=conf) # What's normal about it?
}
ci.simple.width <- function(n,p,conf=0.95) {
    interval <<- ci.simple(n,p,conf)
    apply(interval,1,function(x) x[3]-x[2])
}
```

```
> source("ci_simple.R")
> ci.simple(4390,0.515,0.95)
     conf
[1,] 0.95 0.5002161 0.5297839
> ci.simple.width(4390,0.515,0.95)
[1] 0.02956789
> ci.simple(4390,0.515,c(0.95,0.99,0.999,0.9999,0.99999))
        conf
[1,] 0.95000 0.5002161 0.5297839
[2,] 0.99000 0.4955706 0.5344294
[3,] 0.99900 0.4901797 0.5398203
[4,] 0.99990 0.4856534 0.5443466
[5,] 0.99999 0.4816814 0.5483186
> ci.simple.width(4390,0.515,c(0.95,0.99,0.999,0.9999,0.99999))
[1] 0.02956789 0.03885879 0.04964067 0.05869321 0.06663719
```

We need wider intervals to increase confidence!

Confidence ≠ Information

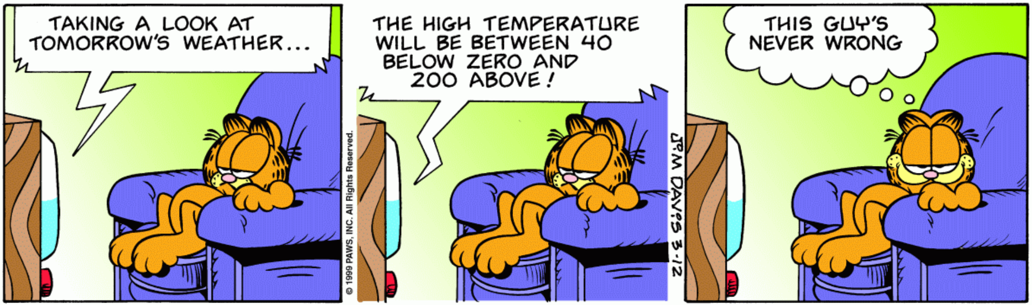

Confidence Interval: Analysis (III)

- CI depends on the observed/measured P_s:
  - widest at $P_s = 0.5$
  - narrowest at the edges ($P_s = 0$ and $P_s = 1$)
  - we are least confident when the result may go either way

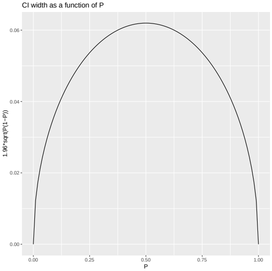

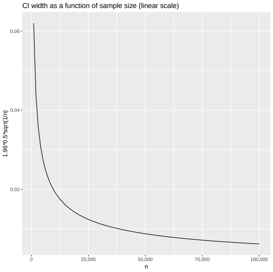

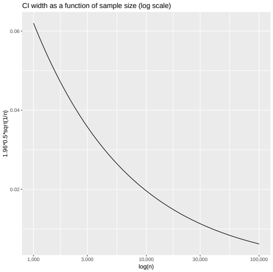

- CI depends on sample size
- but only weakly, as $\sqrt{1/n}$
- rule of thumb: to be 3 times more confident we need a $\sim$10 times larger sample

Accuracy and Precision

Accuracy definition

A measure of *systematic errors* or or *statistical bias*.

Precision definition

A measure of *random errors* or of *statistical variability*.

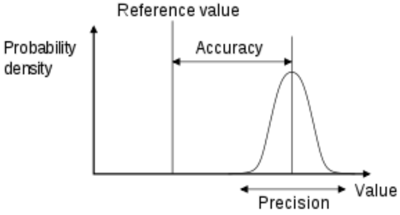

Accuracy and Precision

Accuracy definition

A measure of *systematic errors* or or *statistical bias*.

Precision definition

A measure of *random errors* or of *statistical variability*.

- ISO definition:
  - systematic errors → *trueness*
  - *accuracy* → combination of trueness and precision

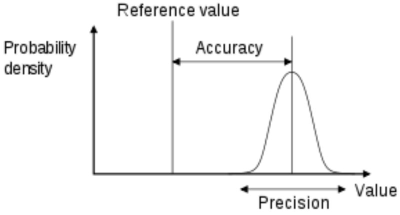

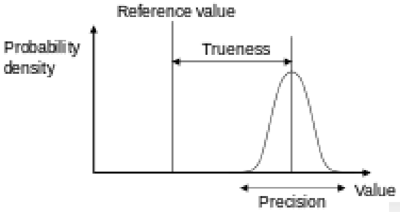

Pop Quiz

Which of the following statements is correct?

- It is always better to take a census than to draw a sample.
- Stopping students on their way out of the campus cafeteria is a good way to sample if we want to collect opinions about the quality of the food there.
- Rather than asking students exiting the cafeteria it is better to set up a "Tell Us What You Think" website.
- We drew a sample of 100 from the 7000 students at IDC. To get the same level of precision for the adult population of Herzliya ($\sim$70,000) we need a sample of 1000.
- A poll on a statistics website garnered 12,357 responses. 8,342 of the respondents said they enjoyed doing statistics homework. With a sample this large we can be quite sure that this is true for most students who take Statistical Methods classes.