Unit 5: Inference for categorical variables
Lecture 2: Inference for proportions - simulation

Statistics 104

Mine Çetinkaya-Rundel

November 5, 2013

Announcements

-

Back of the hand

There is a saying "know something like the back of your hand". Describe an experiment to test if people really do know the backs of their hands.

In the MythBusters episode, 11 out of 12 people guesses the backs of their hands correctly.

Hypotheses

What are the hypotheses for evaluating if people are capable of recognizing the back of their hand at a rate that is better than random guessing. Remember, in the MythBusters experiment, there were 10 pictures to choose from, and only 1 was correct.

H_0 : $p = 0.10$ (random guessing)

H_A : $p > 0.10$ (better than random guessing)

Conditions

1. *Independence:* We can assume that each person guessing is independent of another.
2. *Sample size:* The number of expected successes is *smaller than 10*.

$$12 \times 0.1 = 1.2$$

So what do we do?

Since the sample size isn't large enough to use CLT based methods, we use a simulation method instead.

Clicker question

Pick the simulation scheme for testing if results of this experiment suggest that people are capable of recognizing the back of their hand at a rate that is better than random guessing. ($H_0 : p = 0.10$; $H_A : p > 0.10$; $\hat{p} = 11/12 = 0.9167$; $n = 12$)

(a) Roll a 10-sided die 12 times, record the proportion of 1s, repeat many times: p-val = proportion of simulations where the proportion of 1s ≥ 0.9167.

(b) Roll a 10-sided die 12 times, record the proportion of 1s, repeat many times: p-val = proportion of simulations where the proportion of 1s ≥ 0.10.

(c) Roll a 12-sided die 10 times, record the proportion of 1s, repeat many times: p-val = proportion of simulations where the proportion of 1s ≥ 0.9167.

(d) Flip a coin 12 times, record the proportion of heads, repeat many times: p-val = proportion of simulations where the proportion of 1s ≥ 0.9167.

(e) From a bag containing 10 chips (1 blue, 9 black), draw with replacement 12 chips, record the proportion of blue chips, repeat many times: p-val = proportion of simulations where the proportion of 1s ≥ 0.9167.

Simulation results

- In the next slide you can see the results of a hypothesis test (using only 100 simulations to keep things simple).
- Each dot represents a simulation proportion of success. There were 25-30 simulations where the success rate ($\hat{p}$) was 10%, 40-45 simulations where the success rate was slightly less than 10%, about 20 simulations where the success rate was slightly less than 20% and 1 simulation where the success rate was more than 30%.
- There are no simulations where the success rate is as high as the observed success rate of 91.67%.
- Therefore we conclude that the observed result is near impossible to have happened by chance (p-value = 0).
- And hence that these data suggest that people are capable of recognizing the back of their hand at a rate that is better than random guessing.


```
# create "back of the hand" dataset
back = c(rep("correct", 11), "wrong")

# inference
inference(back, est = "proportion", type = "ht", method = "simulation", success = "correct",
        null = 0.1, alternative = "greater", seed = 654, nsim = 100)
```

```
Single proportion -- success: correct
Summary statistics: p_hat = 0.9167 ;  n = 12
Randomizing, please wait...
H0: p = 0.1
HA: p > 0.1
p-value =  0
```

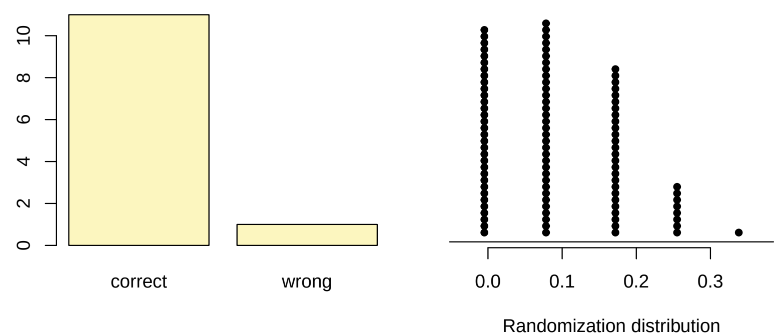

Simulation scheme for bootstrap CI for a proportion

What about constructing a confidence interval for the proportion of people who can guess the back of their hand correctly? How are the centers of the sampling distributions for a HT and CI different?

Simulation scheme for bootstrap CI for a proportion

Clicker question

Pick the simulation scheme for constructing the bootstrap distribution for the proportion of people who are capable of recognizing the back of their hand at a rate that is better than random guessing. ($\hat{p} = 11/12 = 0.9167; \quad n = 12$)

(a) Roll a 12-sided die 12 times, record the proportion of 1-11s, repeat many times.

(b) Roll a 12-sided die many times, record the proportion of 1s.

(c) From a bag containing 12 chips (1 blue, 11 black), draw 12 chips, with replacement, record the proportion of black chips, repeat many times.


```
inference(back, est = "proportion", type = "ci", method = "simulation",
     success = "correct", seed = 547, nsim = 100)
```

```
Single proportion -- success: correct
Summary statistics: p_hat = 0.9167 ;  n = 12
Bootstrapping, please wait...
95 % Bootstrap interval = ( 0.7896 , 1 )
```

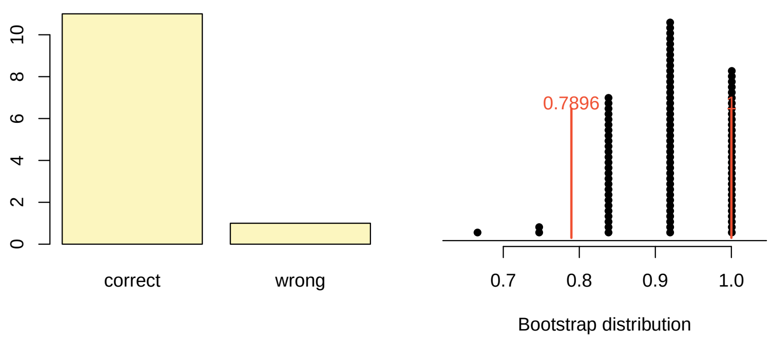

Recap

	HT	CI
simulation technique	Randomization	Bootstrapping
probability of success	null value (0.10)	observed proportion of success (0.9167)
center of resulting distribution	null value (0.10)	observed proportion of success (0.9167)
goal	p-value proportion is at least as high as the observed, or higher), and to use this p-value to make a decision about the hypotheses	estimate true population parameter middle XX% of the bootstrap distribution, or calculated using the SE method

Comparing back of the hand to palm of the hand

MythBusters also asked these people to guess the palms of their hands. This time 7 out of the 12 people guesses correctly. The data are summarized below.

	Palm	Back	Total
Correct	11	7	18
Wrong	1	5	6
Total	12	12	24

Proportion of correct guesses

	Palm	Back	Total
Correct	11	7	18
Wrong	1	5	6
Total	12	12	24

- Proportion of correct in the back group: $\frac{11}{12} = 0.916$
- Proportion of correct in the palm group: $\frac{7}{12} = 0.583$
- Difference: 33.3% more correct in the back of the hand group.

Based on the proportions we calculated, do you think the chance of guessing the back of the hand correctly is higher than palm of the hand?

Hypotheses

What are the hypotheses for comparing if the proportion of people who can guess the backs of their hands correctly is greater than the proportion of people who can guess the palm of their hands correctly?

H_0: $p_{back} = p_{palm}$

H_A: $p_{back} > p_{palm}$

Simulation scheme

1. Use 24 index cards, where each card represents a subject.
2. Mark 18 of the cards as “correct” and the remaining 6 as “wrong”.
3. Shuffle the cards and split into two groups of size 12, for back and palm.
4. Calculate the difference between the proportions of “correct” in the back and palm decks, and record this number.
5. Repeat steps (3) and (4) many times to build a randomization distribution of differences in simulated proportions.

Simulation results

- In the next slide you can see the result of a hypothesis test (using only 100 simulations to keep the results simple).
- Each dot represents a difference in simulated proportion of successes. We can see that the distribution is centered at 0 (the null value).
- We can also see that 9 out of the 100 simulations yielded simulated differences at least as large as the observed difference (p-value = 0.09).


```
palm = c(rep("correct", 7), rep("wrong", 5))
hand = c(back, palm)
group = c(rep("back", 12), rep("palm", 12))

inference(hand, group, est = "proportion", type = "ht", null = 0, alternative = "greater",
          success = "correct", method = "simulation", seed = 879, nsim = 100)
```

```
Response variable: categorical, Explanatory variable: categorical
Two categorical variables
Difference between two proportions -- success: correct
Summary statistics:
         group
data      back palm Sum
  correct   11    7  18
  wrong      1    5   6
  Sum       12   12  24
Observed difference between proportions (back-palm) = 0.3333

H0: p_back - p_palm = 0
HA: p_back - p_palm > 0
Randomizing, please wait...
p-value =  0.07
```