

Unit 4: Inference for numerical data

2. Bootstrapping

Sta 101 – Spring 2019

Duke University, Department of Statistical Science



Dr. Ellison

Slides posted at
<https://www2.stat.duke.edu/courses/Spring19/sta101.001/>

Outline

1. Housekeeping

2. Main ideas

Problem: We can't make CLT confidence intervals or hypothesis tests for MEDIANS.

Solution: Use Bootstrapping (a simulation method).

1. 📄 🖱️ Bootstrapping = sampling with replacement from the observed sample
2. 📄 Bootstrap percentile intervals: middle XX% of the bootstrap distribution
3. 📄 Bootstrap SE intervals: point estimate \pm ME
4. 📄 Bootstrap testing for a single numerical variable requires shifting the bootstrap distribution to be centered at the null value

Announcements

Coming up...

- ▶ Lab Assignment 6 is due **Thursday just before your lab section time.**
- ▶ Peer Evaluations is due **Thursday 2/28 11:55pm** (part of your participation grade)
- ▶ Read over project statement before **Thursday 2/28** (see Sakai Resources)
- ▶ Project Stage 1 is due **Thursday 3/7**

1

Outline

1. Housekeeping


2. Main ideas

Problem: We can't make CLT confidence intervals or hypothesis tests for MEDIANS.

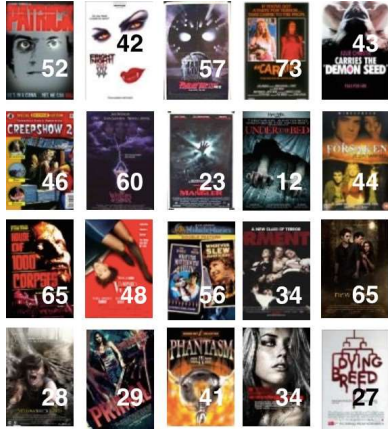
Solution: Use Bootstrapping (a simulation method).

1. 📄 🖱️ Bootstrapping = sampling with replacement from the observed sample
2. 📄 Bootstrap percentile intervals: middle XX% of the bootstrap distribution
3. 📄 Bootstrap SE intervals: point estimate \pm ME
4. 📄 Bootstrap testing for a single numerical variable requires shifting the bootstrap distribution to be centered at the null value

Rotten horrors



is a movie aggregator, where the audience is also able to review and score the movies. We want to *estimate* the **median** audience score of horror movies on RottenTomatoes.com. We start with a **random sample of 20** horror movies.



2

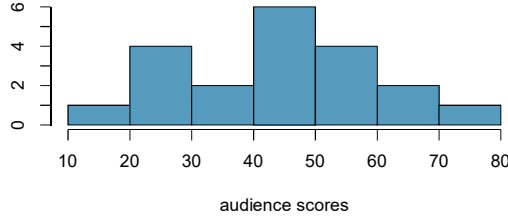
Data

	title	audience_score
1	Patrick	52
2	Demon Seed	43
3	Tormented	34
4	Under the Bed	12
5	Phantasm IV: Oblivion	41
6	Fright Night Part 2	42
7	House of 1000 Corpses	65
8	Creepshow 2	46
9	The Forsaken	44
10	All the Boys Love Mandy Lane	34
11	Jason Lives: Friday the 13th Part VI	57
12	Vampire's Kiss	48
13	The Witches of Eastwick	60
14	Yellowbrickroad	28
15	Dying Breed	27
16	Carrie	73
17	Whoever Slew Auntie Roo?	56
18	The Mangler	23
19	Primal	29
20	The Twilight Saga: New Moon	65

3

First look

The histogram below shows the distribution of the audience scores of these movies (ranging from 0 to 100). The median score in the sample is 43.5. Can we apply CLT based methods we have learned so far to construct a confidence interval for the median RottenTomatoes score of horror movies. Why or why not?



4

Outline

Recap... what we can do so far...

	Do we know how to create a confidence interval/hypothesis test yet? If so, how?	
Population Parameter	Confidence Interval	Hypothesis Testing
$\mu_1 - \mu_2$	CLT Confidence Interval (Unit 4)	Randomization Testing (Unit 1) CLT Hypothesis Testing (Unit 4)
$p_1 - p_2$		Randomization Testing (Unit 1)
Median ₁ -Median ₂		Randomization Testing (Unit 1)
μ	CLT Confidence Interval (Unit 3 and 4)	CLT Hypothesis Testing (Units 3 and 4)
μ_{diff}	CLT Confidence Interval (Unit 4)	CLT Hypothesis Testing (Units 3 and 4)
Median		
p		

Outline

What we can do after today...

	Do we know how to create a confidence interval/hypothesis test yet? If so, how?	
Population Parameter	Confidence Interval	Hypothesis Testing
$\mu_1 - \mu_2$	CLT Confidence Interval (Unit 4)	Randomization Testing (Unit 1) CLT Hypothesis Testing (Unit 4)
$p_1 - p_2$		Randomization Testing (Unit 1)
Median ₁ -Median ₂		Randomization Testing (Unit 1)
μ	CLT Confidence Interval (Unit 3 and 4) Bootstrap Confidence Interval (Unit 4)	CLT Hypothesis Testing (Units 3 and 4) Bootstrap Hypothesis Testing (Unit 4)
μ_{diff}	CLT Confidence Interval (Unit 4) Bootstrap Confidence Interval (Unit 4)	CLT Hypothesis Testing (Units 3 and 4) Bootstrap Hypothesis Testing (Unit 4)
Median	Bootstrap Confidence Interval (Unit 4)	Bootstrap Hypothesis Testing (Unit 4)
p	Bootstrap Confidence Interval (Unit 4)	Bootstrap Hypothesis Testing (Unit 4)

Outline

What are some reasons to use bootstrapping methods to create a confidence interval or conduct a hypothesis test instead of CLT-based methods?

a) The population parameter of interest is _____.

b) The conditions for CLT-based methods are _____ strict. In bootstrapping methods you don't need:

- i. _____
- ii. _____

Therefore you can create a confidence interval or conduct a hypothesis test on population parameters like _____, _____, _____ under less strict conditions.

Outline

What are some reasons to use bootstrapping methods to create a confidence interval or conduct a hypothesis test instead of CLT-based methods?

a) The population parameter of interest is the median.

b) The conditions for CLT-based methods are more strict. In bootstrapping methods you don't need:

- i. $n > 30$ OR
- ii. The population distribution to be approx. normal

Therefore you can create a confidence interval or conduct a hypothesis test on population parameters like μ, μ_{diff}, p (others exist) under less strict conditions.

Outline

From videos...

bootstrapping limitations

- ▶ Not as rigid conditions as CLT based methods
- ▶ If the bootstrap distribution is extremely skewed or sparse, the bootstrap interval might be unreliable
- ▶ A representative sample is still required — if the sample is biased, the estimates resulting from this sample will also be biased.

Some CLT conditions still *do* apply however: **Sample should be random.**

Bootstrapping

- ▶ An alternative approach to constructing confidence intervals is *bootstrapping*.

5


Bootstrapping

- ▶ An alternative approach to constructing confidence intervals is *bootstrapping*.
- ▶ This term comes from the phrase “pulling oneself up by one’s bootstraps”, which is a metaphor for accomplishing an impossible task without any outside help.

5

Bootstrapping

- ▶ An alternative approach to constructing confidence intervals is *bootstrapping*.
- ▶ This term comes from the phrase “pulling oneself up by one’s bootstraps”, which is a metaphor for accomplishing an impossible task without any outside help.
- ▶ In this case the impossible task is *estimating a population parameter*, and we’ll accomplish it *using data from only the given sample*.

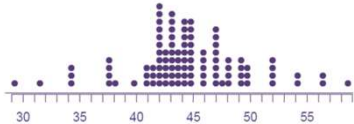


5

Outline

How do we make a confidence interval using bootstrapping?

Step 1: Create a bootstrap distribution.



Bootstrapping

Step 1: Create the bootstrap distribution

	Creating a sampling distribution (approximation)	Creating a bootstrap distribution
<i>Do the following many many times..</i>		
Step 1:		
Step 2:		
Step 3:		

For comparison

6

Bootstrapping

Step 1: Create the bootstrap distribution

	Creating a sampling distribution (approximation)	Creating a bootstrap distribution
<i>Do the following many many times..</i>		
Step 1:	replacement, take a random sample of size n from the...	
Step 2:		
Step 3:		

For comparison

6

Bootstrapping

Step 1: Create the bootstrap distribution

	Creating a sampling distribution (approximation)	Creating a bootstrap distribution
<i>Do the following many many times..</i>		
Step 1:	with replacement, take a random sample of size n from the... population.	sample (of size n).
Step 2:		
Step 3:		

For comparison

6

Bootstrapping

Step 1: Create the bootstrap distribution

	Creating a sampling distribution (approximation)	Creating a bootstrap distribution
<i>Do the following many many times..</i>		
Step 1:	with replacement, take a random sample of size n from the... population.	sample (of size n).
Step 2:	Calculate the sample statistic (ie: sample mean/median/proportion) of this random sample from (step 1).	
Step 3:		

For comparison

6

Bootstrapping

Step 1: Create the bootstrap distribution

	Creating a sampling distribution (approximation)	Creating a bootstrap distribution
<i>Do the following many many times..</i>		
Step 1:	with replacement, take a random sample of size n from the... population.	sample (of size n).
Step 2:	Calculate the sample statistic (ie: sample mean/median/proportion) of this random sample from (step 1).	
Step 3:	Plot this sample statistic in a histogram/dotplot. This histogram/dotplot is your... sampling distribution.	bootstrap distribution.

For comparison

6

Bootstrapping

Step 1: Create the bootstrap distribution

	Creating a sampling distribution (approximation)	Creating a bootstrap distribution
<i>Do the following many many times..</i>		
Step 1:	with replacement, take a random sample of size n from the... population.	sample (of size n). aka: bootstrap sample
Step 2:	Calculate the sample statistic (ie: sample mean/median/proportion) of this random sample from (step 1). aka: bootstrap statistic	
Step 3:	Plot this sample statistic in a histogram/dotplot. This histogram/dotplot is your... sampling distribution.	bootstrap distribution.

For comparison

6

Outline

How do we make a confidence interval using bootstrapping?

Step 1: Create a bootstrap distribution.
Step 2: Use the bootstrap distribution to create your confidence interval.

Bootstrapping

Step 2: Use the bootstrap distribution to create a confidence interval.

- **Percentile Method**
 - XX% bootstrap confidence interval = the cutoff values for the middle XX% of the bootstrap distribution

6

Bootstrapping

*Step 2: Use the **bootstrap distribution** to create a confidence interval.*

- Percentile Method**
 - XX% bootstrap confidence interval = the cutoff values for the middle XX% of the bootstrap distribution
- Standard Error Method**
 - XX% bootstrap confidence interval = $point\ estimate \pm t^*SE_{boot}$

6

Bootstrapping

*Step 2: Use the **bootstrap distribution** to create a confidence interval.*

- Percentile Method**
 - XX% bootstrap confidence interval = the cutoff values for the middle XX% of the bootstrap distribution
- Standard Error Method**
 - XX% bootstrap confidence interval = $point\ estimate \pm t^*SE_{boot}$

Standard deviation of all the bootstrap statistics you calculated

When the point estimate is: median, μ , μ_{diff} , p

- degrees of freedom = $n-1$ (n = size of the original sample = each bootstrap sample)
- area under t-distribution curve = $(1-XX/100)\%$ TWO TAILS.

6

Outline

Let's make a bootstrap confidence interval for the population median audience score!

Step 1: Create a bootstrap distribution.



7

Bootstrap sample 1

(1) Take a bootstrap sample:

7

Bootstrap sample 1

(1) Take a bootstrap sample:

	title	audience_score
1	Vampire's Kiss	48
2	Phantasm IV: Oblivion	41
3	House of 1000 Corpses	65
4	Dying Breed	27
5	Whoever Slew Auntie Roo?	56
6	The Forsaken	44
7	The Twilight Saga: New Moon	65
8	The Twilight Saga: New Moon	65
9	Whoever Slew Auntie Roo?	56
10	The Twilight Saga: New Moon	65
11	The Mangler	23
12	Dying Breed	27
13	Creepshow 2	46
14	House of 1000 Corpses	65
15	Whoever Slew Auntie Roo?	56
16	Tormented	34
17	Jason Lives: Friday the 13th Part VI	57
18	Vampire's Kiss	48
19	Primal	29
20	The Witches of Eastwick	60

(2) Calculate the median of the bootstrap sample:

7

Bootstrap sample 1

(1) Take a bootstrap sample:

	title	audience_score
1	Vampire's Kiss	48
2	Phantasm IV: Oblivion	41
3	House of 1000 Corpses	65
4	Dying Breed	27
5	Whoever Slew Auntie Roo?	56
6	The Forsaken	44
7	The Twilight Saga: New Moon	65
8	The Twilight Saga: New Moon	65
9	Whoever Slew Auntie Roo?	56
10	The Twilight Saga: New Moon	65
11	The Mangler	23
12	Dying Breed	27
13	Creepshow 2	46
14	House of 1000 Corpses	65
15	Whoever Slew Auntie Roo?	56
16	Tormented	34
17	Jason Lives: Friday the 13th Part VI	57
18	Vampire's Kiss	48
19	Primal	29
20	The Witches of Eastwick	60

(2) Calculate the median of the bootstrap sample:
 23, 27, 27, 29, 34, 41, 44, 46, 48, 48, 56, 56, 56, 57, 60, 65, 65, 65, 65, 65
 median = (48 + 56) / 2 = 52

7

Bootstrap sample 1

(1) Take a bootstrap sample:

	title	audience_score
1	Vampire's Kiss	48
2	Phantasm IV: Oblivion	41
3	House of 1000 Corpses	65
4	Dying Breed	27
5	Whoever Slew Auntie Roo?	56
6	The Forsaken	44
7	The Twilight Saga: New Moon	65
8	The Twilight Saga: New Moon	65
9	Whoever Slew Auntie Roo?	56
10	The Twilight Saga: New Moon	65
11	The Mangler	23
12	Dying Breed	27
13	Creepshow 2	46
14	House of 1000 Corpses	65
15	Whoever Slew Auntie Roo?	56
16	Tormented	34
17	Jason Lives: Friday the 13th Part VI	57
18	Vampire's Kiss	48
19	Primal	29
20	The Witches of Eastwick	60

(2) Calculate the median of the bootstrap sample:
 23, 27, 27, 29, 34, 41, 44, 46, 48, 48, 56, 56, 56, 57, 60, 65, 65, 65, 65, 65
 median = (48 + 56) / 2 = 52

(3) Plot this value ●

7

Bootstrap sample 1

(1) Take a bootstrap sample:

	title	audience_score
1	Vampire's Kiss	48
2	Phantasm IV: Oblivion	41
3	House of 1000 Corpses	65
4	Dying Breed	27
5	Whoever Slew Auntie Roo?	56
6	The Forsaken	44
7	The Twilight Saga: New Moon	65
8	The Twilight Saga: New Moon	65
9	Whoever Slew Auntie Roo?	56
10	The Twilight Saga: New Moon	65
11	The Mangler	23
12	Dying Breed	27
13	Creepshow 2	46
14	House of 1000 Corpses	65
15	Whoever Slew Auntie Roo?	56
16	Tormented	34
17	Jason Lives: Friday the 13th Part VI	57
18	Vampire's Kiss	48
19	Primal	29
20	The Witches of Eastwick	60

(2) Calculate the median of the bootstrap sample:
 23, 27, 27, 29, 34, 41, 44, 46, 48, 48, 56, 56, 56, 57, 60, 65, 65, 65, 65, 65
 median = (48 + 56) / 2 = 52

(3) Plot this value ● → A bootstrap statistic (median)

7

Bootstrap sample 2

(1) Take another bootstrap sample:

8

Bootstrap sample 2

(1) Take another bootstrap sample:

	title	audience_score
1	Fright Night Part 2	42
2	Carrie	73
3	The Forsaken	44
4	The Mangler	23
5	Primal	29
6	Patrick	52
7	Jason Lives: Friday the 13th Part VI	57
8	The Mangler	23
9	Vampire's Kiss	48
10	All the Boys Love Mandy Lane	34
11	The Twilight Saga: New Moon	65
12	All the Boys Love Mandy Lane	34
13	Yellowbrickroad	28
14	Vampire's Kiss	48
15	Tormented	34
16	The Mangler	23
17	Phantasm IV: Oblivion	41
18	Patrick	52
19	House of 1000 Corpses	65
20	The Twilight Saga: New Moon	65

8

Bootstrap sample 2

(1) Take another bootstrap sample:

	title	audience_score
1	Fright Night Part 2	42
2	Carrie	73
3	The Forsaken	44
4	The Mangler	23
5	Primal	29
6	Patrick	52
7	Jason Lives: Friday the 13th Part VI	57
8	The Mangler	23
9	Vampire's Kiss	48
10	All the Boys Love Mandy Lane	34
11	The Twilight Saga: New Moon	65
12	All the Boys Love Mandy Lane	34
13	Yellowbrickroad	28
14	Vampire's Kiss	48
15	Tormented	34
16	The Mangler	23
17	Phantasm IV: Oblivion	41
18	Patrick	52
19	House of 1000 Corpses	65
20	The Twilight Saga: New Moon	65

(2) Calculate the median of the bootstrap sample:

8

Bootstrap sample 2

(1) Take another bootstrap sample:

	title	audience_score
1	Fright Night Part 2	42
2	Carrie	73
3	The Forsaken	44
4	The Mangler	23
5	Primal	29
6	Patrick	52
7	Jason Lives: Friday the 13th Part VI	57
8	The Mangler	23
9	Vampire's Kiss	48
10	All the Boys Love Mandy Lane	34
11	The Twilight Saga: New Moon	65
12	All the Boys Love Mandy Lane	34
13	Yellowbrickroad	28
14	Vampire's Kiss	48
15	Tormented	34
16	The Mangler	23
17	Phantasm IV: Oblivion	41
18	Patrick	52
19	House of 1000 Corpses	65
20	The Twilight Saga: New Moon	65

(2) Calculate the median of the bootstrap sample:
 23, 23, 23, 28, 29, 34, 34, 34, 41, 42, 44, 48, 48, 52, 52, 57, 65, 65, 65, 73
 median = (42 + 44) / 2 = 43

8

Bootstrap sample 2

(1) Take another bootstrap sample:

	title	audience_score
1	Fright Night Part 2	42
2	Carrie	73
3	The Forsaken	44
4	The Mangler	23
5	Primal	29
6	Patrick	52
7	Jason Lives: Friday the 13th Part VI	57
8	The Mangler	23
9	Vampire's Kiss	48
10	All the Boys Love Mandy Lane	34
11	The Twilight Saga: New Moon	65
12	All the Boys Love Mandy Lane	34
13	Yellowbrickroad	28
14	Vampire's Kiss	48
15	Tormented	34
16	The Mangler	23
17	Phantasm IV: Oblivion	41
18	Patrick	52
19	House of 1000 Corpses	65
20	The Twilight Saga: New Moon	65

(2) Calculate the median of the bootstrap sample:
 23, 23, 23, 28, 29, 34, 34, 34, 41, 42, 44, 48, 48, 52, 52, 57, 65, 65, 65, 73
 median = (42 + 44) / 2 = 43

(3) Plot this value ● → ...Another bootstrap statistic (median)₈

Bootstrap sample 3

(1) Take another bootstrap sample:

9

Bootstrap sample 3

(1) Take another bootstrap sample:

	title	audience_score
1	Tormented	34
2	The Witches of Eastwick	60
3	The Witches of Eastwick	60
4	The Witches of Eastwick	60
5	The Mangler	23
6	The Witches of Eastwick	60
7	Patrick	52
8	Phantasm IV: Oblivion	41
9	Yellowbrickroad	28
10	Jason Lives: Friday the 13th Part VI	57
11	Yellowbrickroad	28
12	Jason Lives: Friday the 13th Part VI	57
13	Fright Night Part 2	42
14	Primal	29
15	Fright Night Part 2	42
16	Whoever Slew Auntie Roo?	56
17	Fright Night Part 2	42
18	Fright Night Part 2	42
19	Under the Bed	12
20	Phantasm IV: Oblivion	41

9

Bootstrap sample 3

(1) Take another bootstrap sample:

	title	audience_score
1	Tormented	34
2	The Witches of Eastwick	60
3	The Witches of Eastwick	60
4	The Witches of Eastwick	60
5	The Mangler	23
6	The Witches of Eastwick	60
7	Patrick	52
8	Phantasm IV: Oblivion	41
9	Yellowbrickroad	28
10	Jason Lives: Friday the 13th Part VI	57
11	Yellowbrickroad	28
12	Jason Lives: Friday the 13th Part VI	57
13	Fright Night Part 2	42
14	Primal	29
15	Fright Night Part 2	42
16	Whoever Slew Auntie Roo?	56
17	Fright Night Part 2	42
18	Fright Night Part 2	42
19	Under the Bed	12
20	Phantasm IV: Oblivion	41

(2) Calculate the median of the bootstrap sample:

9

Bootstrap sample 3

(1) Take another bootstrap sample:

	title	audience_score
1	Tormented	34
2	The Witches of Eastwick	60
3	The Witches of Eastwick	60
4	The Witches of Eastwick	60
5	The Mangler	23
6	The Witches of Eastwick	60
7	Patrick	52
8	Phantasm IV: Oblivion	41
9	Yellowbrickroad	28
10	Jason Lives: Friday the 13th Part VI	57
11	Yellowbrickroad	28
12	Jason Lives: Friday the 13th Part VI	57
13	Fright Night Part 2	42
14	Primal	29
15	Fright Night Part 2	42
16	Whoever Slew Auntie Roo?	56
17	Fright Night Part 2	42
18	Fright Night Part 2	42
19	Under the Bed	12
20	Phantasm IV: Oblivion	41

(2) Calculate the median of the bootstrap sample:
 12, 23, 28, 28, 29, 34, 41, 41, 42, 42, 42, 42, 52, 56, 57, 57, 60, 60, 60, 60
 median = $(42 + 42) / 2 = 42$

9

Bootstrap sample 3

(1) Take another bootstrap sample:

	title	audience_score
1	Tormented	34
2	The Witches of Eastwick	60
3	The Witches of Eastwick	60
4	The Witches of Eastwick	60
5	The Mangler	23
6	The Witches of Eastwick	60
7	Patrick	52
8	Phantasm IV: Oblivion	41
9	Yellowbrickroad	28
10	Jason Lives: Friday the 13th Part VI	57
11	Yellowbrickroad	28
12	Jason Lives: Friday the 13th Part VI	57
13	Fright Night Part 2	42
14	Primal	29
15	Fright Night Part 2	42
16	Whoever Slew Auntie Roo?	56
17	Fright Night Part 2	42
18	Fright Night Part 2	42
19	Under the Bed	12
20	Phantasm IV: Oblivion	41

(2) Calculate the median of the bootstrap sample:
 12, 23, 28, 28, 29, 34, 41, 41, 42, 42, 42, 42, 52, 56, 57, 57, 60, 60, 60, 60
 median = $(42 + 42) / 2 = 42$

(3) Plot this value ...and another bootstrap statistic (median)

9

Many more bootstrap samples

... repeat

Bootstrap Distribution

bootstrap medians

10

Clicker question

The dot plot below is the bootstrap distribution of medians constructed using 100 simulations. What does each dot on the dot plot represent?

Bootstrap Distribution

bootstrap medians

- (a) Score of a horror movie in the original sample
- (b) Score of a horror movie in the population
- (c) Median from one bootstrap sample from the original sample
- (d) Median from one sample from the population

11

Clicker question

The dot plot below is the bootstrap distribution of medians constructed using 100 simulations. What does each dot on the dot plot represent?

(a) Score of a horror movie in the original sample (*sample dist*)
 (b) Score of a horror movie in the population (*population dist*)
 (c) Median from one bootstrap sample from the original sample (*bootstrap dist*)
 (d) Median from one sample from the population (*sampling dist*)

11

Outline

1. Housekeeping

2. Main ideas

Problem: We can't make CLT confidence intervals or hypothesis tests for MEDIANS.

Solution: Use Bootstrapping (a simulation method).

1. Bootstrapping = sampling with replacement from the observed sample
2. Bootstrap percentile intervals: middle XX% of the bootstrap distribution
3. Bootstrap SE intervals: point estimate \pm ME
4. Bootstrap testing for a single numerical variable requires shifting the bootstrap distribution to be centered at the null value

Outline

Let's make a bootstrap confidence interval for the population median audience score!

Step 1: Create a bootstrap distribution.
*Step 2: Use the bootstrap distribution to create a confidence interval. Let's use the **percentile method**!*

Clicker question

The dot plot below shows the distribution of 100 bootstrap medians. Estimate the 90% bootstrap confidence interval for the median RT score of horror movies using the **percentile method**.

(a) (29, 58.5) (c) (37.5, 52)
 (b) (34, 57) (d) (40, 49.5)

12

Clicker question

The dot plot below shows the distribution of 100 bootstrap medians. Estimate the 90% bootstrap confidence interval for the median RT score of horror movies using the **percentile method**.

90% confidence Interval should "hug" the middle 90% of points.

Bootstrap Distribution

bootstrap medians

(a) (29, 58.5) (c) (37.5, 52)

(b) (34, 57) (d) (40, 49.5)

12

Outline

1. Housekeeping

2. Main ideas

Problem: We can't make CLT confidence intervals or hypothesis tests for **MEDIANS**.

Solution: Use **Bootstrapping** (a simulation method).

- 🔍 🖐️ Bootstrapping = sampling with replacement from the observed sample
- 🔍 Bootstrap percentile intervals: middle XX% of the bootstrap distribution
- 🔍 Bootstrap SE intervals: point estimate \pm ME
- 🖐️ Bootstrap testing for a single numerical variable requires shifting the bootstrap distribution to be centered at the null value

🔍 Outline

Let's make a bootstrap confidence interval for the population median audience score!

Step 1: Create a bootstrap distribution.

*Step 2: Use the bootstrap distribution to create a confidence interval. Let's use the **standard error method**!*

🔍 🖐️ Bootstrap interval, standard error

The dot plot below shows the distribution of 100 bootstrap medians. The median of the original sample is 43.5 and the bootstrap standard error is 4.88. Estimate the 90% bootstrap confidence interval for the median RT score of horror movies using the **standard error method** (remember $n = 20$).

Bootstrap Distribution

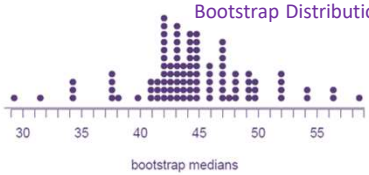
bootstrap medians

$(\text{sample median}) \pm t_{n-1}^* SE_{\text{bootstrap}}$

13

Bootstrap interval, standard error

The dot plot below shows the distribution of 100 bootstrap medians. The median of the original sample is 43.5 and the bootstrap standard error is 4.88. Estimate the 90% bootstrap confidence interval for the median RT score of horror movies using the **standard error method** (remember n=20).



Bootstrap Distribution

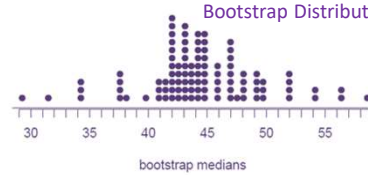
$$(sample\ median) \pm t_{n-1}^* SE_{bootstrap}$$

$$43.5 \pm (1.73 \times 4.88) = (35.1, 51.9)$$

13

Bootstrap interval, standard error

The dot plot below shows the distribution of 100 bootstrap medians. The median of the original sample is 43.5 and the bootstrap standard error is 4.88. Estimate the 90% bootstrap confidence interval for the median RT score of horror movies using the **standard error method** (remember n=20).



Bootstrap Distribution

two tails	0.200	0.100	0.050	
df	16	1.34	1.75	2.12
	17	1.33	1.74	2.11
	18	1.33	1.73	2.10
	19	1.33	1.73	2.09

$$(sample\ median) \pm t_{n-1}^* SE_{bootstrap}$$

$$43.5 \pm (1.73 \times 4.88) = (35.1, 51.9)$$

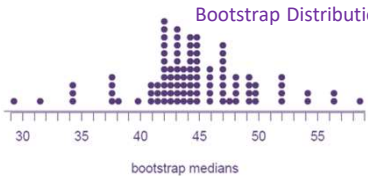
- degrees of freedom = 20-1 (n is the size of the original sample)
- area under t-distribution curve = (1-.90)=0.10 TWO TAILS.

13

Clicker question

How *would* we calculate the bootstrap standard error (SE_{boot}) in the bootstrap distribution below?

a.) $SE_{boot} = s$
 b.) $SE_{boot} = \frac{\sigma}{\sqrt{n}}$
 c.) $SE_{boot} = \frac{s}{\sqrt{n}}$
 d.) $SE_{boot} = \sqrt{\frac{\sum_{i=1}^{100} (boot.median_i - mean\ of\ all\ boot\ medians)^2}{100}}$



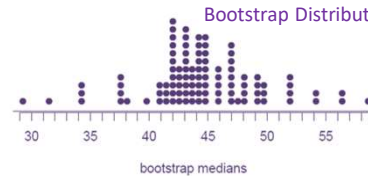
Bootstrap Distribution

6

Clicker question

How *would* we calculate the bootstrap standard error (SE_{boot}) in the bootstrap distribution below?

a.) $SE_{boot} = s$
 b.) $SE_{boot} = \frac{\sigma}{\sqrt{n}}$
 c.) $SE_{boot} = \frac{s}{\sqrt{n}}$
 d.) $SE_{boot} = \sqrt{\frac{\sum_{i=1}^{100} (boot.median_i - mean\ of\ all\ boot\ medians)^2}{100}}$



Bootstrap Distribution

6

Bootstrap vs. sampling distributions

Application exercise: 4.2 Bootstrap intervals

See the course webpage for details.

14

Outline

1. Housekeeping

2. Main ideas

Problem: We can't make CLT confidence intervals or hypothesis tests for MEDIANS.

Solution: Use Bootstrapping (a simulation method).

1. 🔍 🖱️ Bootstrapping = sampling with replacement from the observed sample
2. 🔍 Bootstrap percentile intervals: middle XX% of the bootstrap distribution
3. 🔍 Bootstrap SE intervals: point estimate \pm ME
4. 🆕 Bootstrap testing for a single numerical variable requires shifting the bootstrap distribution to be centered at the null value

NEW
Outline

Let's conduct a bootstrap hypothesis test for the population median audience score!

NEW

Do these data provide convincing evidence that the median audience score of horror movies is greater than 40? Remember that the median of the original sample was 43.5.

Original Bootstrap Distribution

Hypotheses:

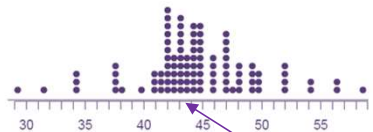
$H_0 : \text{pop. median} = 40$

$H_A : \text{pop. median} > 40$

16

NEW

Do these data provide convincing evidence that the median audience score of horror movies is greater than 40?
Remember that the median of the original sample was 43.5.



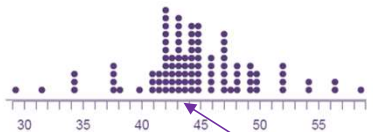
Original Bootstrap Distribution
Centered at: 43.5

Hypotheses:
 $H_0 : \text{pop. median} = 40$
 $H_A : \text{pop. median} > 40$

16

NEW

Do these data provide convincing evidence that the median audience score of horror movies is greater than 40?
Remember that the median of the original sample was 43.5.



Original Bootstrap Distribution
Centered at: 43.5

Hypotheses:
 $H_0 : \text{pop. median} = 40$
 $H_A : \text{pop. median} > 40$

To assess the p-value we want our sampling distribution to:

- Assume H_0 : median = 40

↕

- Be centered at: 40

16

NEW Bootstrap testing for a mean

- ▶ This is very similar to bootstrapping, i.e. we randomly sample with replacement from the sample, but this time we shift the bootstrap distribution to be centered at the null value.

15

NEW Bootstrap testing for a mean

- ▶ This is very similar to bootstrapping, i.e. we randomly sample with replacement from the sample, but this time we shift the bootstrap distribution to be centered at the null value.
- ▶ The **p-value** is then defined as the proportion of simulations that yield a sample statistic at least as extreme as the one observed in the bootstrap distribution that has been shifted to center at the null (aka: the bootstrap distribution that assumes the null value is true).

15

NEW

Do these data provide convincing evidence that the median audience score of horror movies is greater than 40?
Remember that the median of the original sample was 43.5.

Original Bootstrap Distribution

Hypotheses:
 $H_0 : \text{pop. median} = 40$
 $H_A : \text{pop. median} > 40$

Shift all points to the left by 3.5

Original Bootstrap Distribution Shifted to be Centered at Null Value (40)

16

NEW

Do these data provide convincing evidence that the median audience score of horror movies is greater than 40?
Remember that the median of the original sample was 43.5.

Hypotheses:
 $H_0 : \text{pop. median} = 40$
 $H_A : \text{pop. median} > 40$

Original Bootstrap Distribution Shifted to be Centered at Null Value (40)

16

NEW

Do these data provide convincing evidence that the median audience score of horror movies is greater than 40?
Remember that the median of the original sample was 43.5.

Hypotheses:
 $H_0 : \text{pop. median} = 40$
 $H_A : \text{pop. median} > 40$

Original Bootstrap Distribution Shifted to be Centered at Null Value (40)

p-value: proportion of simulations where the simulated bootstrap sample median is at least as extreme as the one observed (43.5). $\rightarrow 20/100 = 0.20$

100=Total number of bootstrap statistics in the bootstrap dist

16

Summary of main ideas

1. Bootstrapping = sampling with replacement from the observed sample
2. Bootstrap percentile intervals: middle XX% of the bootstrap distribution
3. Bootstrap SE intervals: point estimate \pm ME
4. Bootstrap testing for a single numerical variable requires shifting the bootstrap distribution to be centered at the null value

17