

# Inference and simulation

---

Two-sided hypothesis tests using `infer`



# Download and load the dataset

You can follow along by downloading and loading the dataset by placing the following *setup* code block at the top of a R Markdown file.

```
```{r setup, include = FALSE}
# Load required packages
library(tidyverse)
library(infer)
# Load datasets
college_apps <- read_rds(
  url("http://data.cds101.com/college_applications.rds")
)
applicants_data <- read_rds(
  url("http://data.cds101.com/gender_discrimination.rds")
)
```
```

# Simulations recap

In the previous lecture, we ran the following simulations:

```
# College applications null distribution
college_apps_null <- college_apps %>%
  specify(formula = number_colleges ~ NULL) %>%
  hypothesize(null = "point", mu = 8) %>%
  generate(reps = 10000, type = "bootstrap") %>%
  calculate(stat = "mean")

# Gender discrimination null distribution
simulation_results <- applicants_data %>%
  specify(outcome ~ sex, success = "Promoted") %>%
  hypothesize(null = "independence") %>%
  generate(reps = 10000, type = "permute") %>%
  calculate(stat = "diff in props", order = combine("Male", "Female"))
```

# Two-sided hypothesis testing with p-values

If the research question was "Do the data provide convincing evidence that the average amount of schools that GMU students apply to is **different** than the national average?", the alternative hypothesis would be different.

$$H_0 : \mu = 8$$

$$H_A : \mu \neq 8$$

The above is an example of a two-sided hypothesis test. In principle, the  $p$ -value of a two-sided test can be different from a one-sided test.

Computing the  $p$ -value of a two-sided hypothesis test is simple to do with `infer`, just change the `direction` keyword in `get_p_value()` from `"right"` to `"both"`.

# Two-sided hypothesis testing with p-values

If the research question was "Do the data provide convincing evidence that the average amount of schools that GMU students apply to is **different** than the national average?", the alternative hypothesis would be different.

$$H_0 : \mu = 8$$

$$H_A : \mu \neq 8$$

The above is an example of a two-sided hypothesis test. In principle, the  $p$ -value of a two-sided test can be different from a one-sided test.

Computing the  $p$ -value of a two-sided hypothesis test is simple to do with `infer`, just change the `direction` keyword in `get_p_value()` from `"right"` to `"both"`.

```
college_apps_p_value_two_sided <- college_apps_null %>%  
  get_p_value(obs_stat = 9.7, direction = "both")
```

# Two-sided hypothesis testing with p-values


If the research question was "Do the data provide convincing evidence that the average amount of schools that GMU students apply to is **different** than the national average?", the alternative hypothesis would be different.

$$H_0 : \mu = 8$$

$$H_A : \mu \neq 8$$

The above is an example of a two-sided hypothesis test. In principle, the  $p$ -value of a two-sided test can be different from a one-sided test.

Computing the  $p$ -value of a two-sided hypothesis test is simple to do with `infer`, just change the `direction` keyword in `get_p_value()` from "right" to "both".



$p\text{-value} = 0$

# Two-sided hypothesis testing with p-values

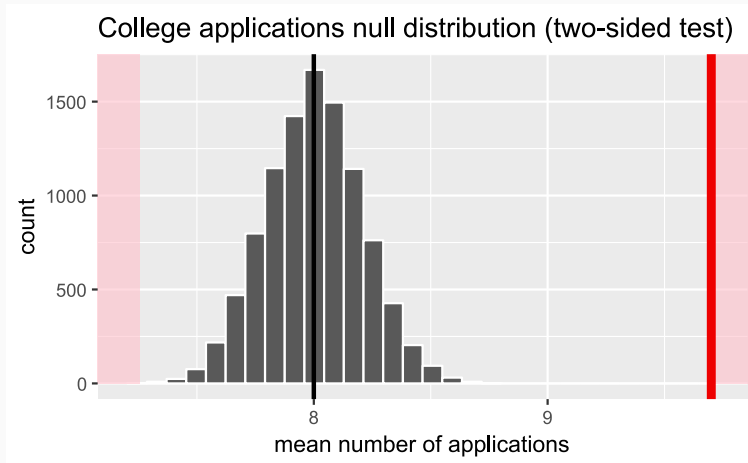
If the research question was "Do the data provide convincing evidence that the average amount of schools that GMU students apply to is **different** than the national average?", the alternative hypothesis would be different.

$$H_0 : \mu = 8$$

$$H_A : \mu \neq 8$$

The above is an example of a two-sided hypothesis test. In principle, the  $p$ -value of a two-sided test can be different from a one-sided test.

Computing the  $p$ -value of a two-sided hypothesis test is simple to do with `infer`, just change the `direction` keyword in `get_p_value()` from "right" to "both".



$p\text{-value} = 0$

In this example, the  $p$ -value does not change.

# Gender discrimination dataset: two-sided hypothesis test

We can use the same null distribution that we generated earlier to compute the two-sided  $p$ -value for the gender discrimination experiment.



# Gender discrimination dataset: two-sided hypothesis test

We can use the same null distribution that we generated earlier to compute the two-sided  $p$ -value for the gender discrimination experiment.

Conducting a two-sided hypothesis test means we are counting when the difference in the men and women hiring fractions is **larger** than 0.292 and also when it is in the opposite extreme, which would be when the bias is towards hiring more women than men.

# Gender discrimination dataset: two-sided hypothesis test

We can use the same null distribution that we generated earlier to compute the two-sided  $p$ -value for the gender discrimination experiment.

Conducting a two-sided hypothesis test means we are counting when the difference in the men and women hiring fractions is **larger** than 0.292 and also when it is in the opposite extreme, which would be when the bias is towards hiring more women than men.

```
sim_p_value_two_sided <- simulation_results %>%  
  get_p_value(obs_stat = experiment_result, direction = "both")
```

---

**p\_value**

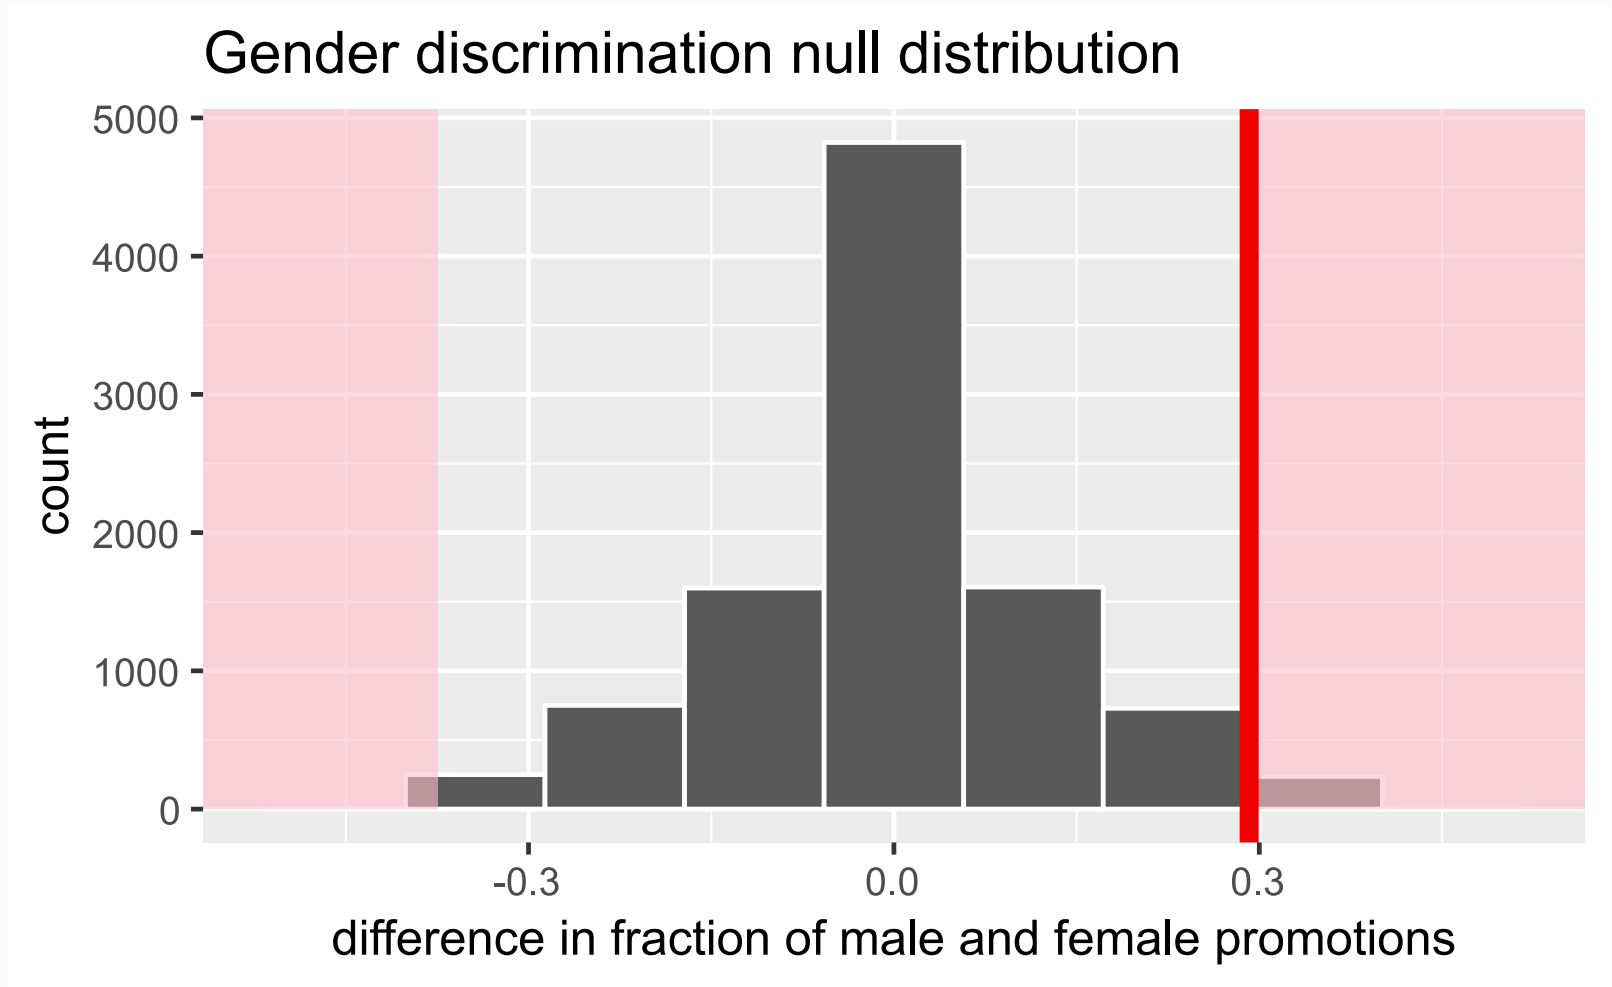
0.0486

---

# Visualization of null distribution (two-sided)

```
simulation_results %>%  
  visualize(bins = 9) +  
  shade_p_value(obs_stat = experiment_result, direction = "both") +  
  labs(  
    x = "difference in fraction of male and female promotions",  
    title = "Gender discrimination null distribution"  
  )
```

# Visualization of null distribution (two-sided)



# Credits

License

Creative Commons Attribution-NonCommerical-ShareAlike 4.0 International

Acknowledgments

Content adapted from the Chapter 3 [OpenIntro Statistics slides](#) developed by Mine Çetinkaya-Rundel and made available under the [CC BY-SA 3.0 license](#).