

Lab 7 - Math 58B: RR & OR

due Tuesday March 10, 2026

your name here

```
library(tidyverse)
library(openintro) # dog data is in openintro
library(infer)     # computational inference
library(mosaic)    # xpnorm and xqnorm
library(praise)
```

Lab Goals

Just like any statistic which is measured on a dataset, the variability of the sample relative risk (\widehat{RR}) and the sample odds ratio (\widehat{OR}) is described by a sampling distribution. The goals for today include being able to:

- describe the null sampling distribution of \widehat{RR} and \widehat{OR} under the null hypothesis using a randomization test
- describe the sampling distribution of \widehat{RR} and \widehat{OR} for general setting using bootstrapping
- describe the sampling distribution of \widehat{RR} and \widehat{OR} (under the null or the alternative) using the Central Limit Theorem (normal distribution!)
- use the randomization distribution to run a complete hypothesis test
- use the bootstrap distribution to create a confidence interval
- use the normal distribution to run a complete hypothesis test

Advice for turning in the assignment

- render early and often. In fact, go ahead and render your .qmd file right now. Maybe set a timer so that you render every 5 minutes. Do **not** wait until you are done with the assignment to render.
- The **assignment** part of the lab is **ONLY** the last six questions at the very bottom. However, the commands in the first half of the assignment are key to doing the second half.

- Save the .qmd file somewhere you can find it. Don't keep everything in your downloads folder. Maybe make a folder called **StatsHW** or something. That folder could live on your Desktop. Or maybe in your Dropbox.

Getting started

- the infer package will be used for both the randomization test and the bootstrapping
- the following two applets may help with forming intuition
 - Randomization test applet: <http://www.rossmanchance.com/applets/2021/chisqshuffle/ChiSqShuffle>
 - Bootstrapping applet: <http://www.rossmanchance.com/applets/2021/twopopprop/twopopprop.html>

Load packages / data

In this lab we will continue to use the **infer** package, along with the applets. The infer syntax is meant to focus understanding on the randomization and bootstrapping processes. For each line, pay attention to what the code is doing.

A study in 1994 examined 491 dogs that had developed cancer and 945 dogs as a control group to determine whether there is an increased risk of cancer in dogs that are exposed to the herbicide 2,4-Dichlorophenoxyacetic acid (2,4-D).

Hayes HM, Tarone RE, Cantor KP, Jessen CR, McCurnin DM, and Richardson RC. 1991. Case- Control Study of Canine Malignant Lymphoma: Positive Association With Dog Owner's Use of 2, 4- Dichlorophenoxyacetic Acid Herbicides. Journal of the National Cancer Institute 83(17):1226-1231.

The dog data are in the **openintro** package.

```
data("cancer_in_dogs")

cancer_in_dogs |>
  table()
```

order	response	
	cancer	no cancer
2,4-D	191	304
no 2,4-D	300	641

Variability of \widehat{OR} when H_0 is true

The randomization process (which leads to the randomization test) creates different possible values of \widehat{RR} and \widehat{OR} assuming the null hypothesis is true.

See: Randomization test applet: <http://www.rossmanchance.com/applets/2021/chisqshuffle/ChiSqShuffle.htm>

Note that the variability of the statistics (the sampling distribution) can be displayed by applying the same ideas that were used on the difference in sample proportions.

What are the steps needed to create a sampling distribution?

Remember to change to #| eval: true !

```
cancer_in_dogs |>
  ___(response ~ order, success = "cancer") |>
  ___(null = "independence") |>
  ___(reps = 1000, type = "permute") |>
  ___(stat = "odds ratio", order = c("2,4-D", "no 2,4-D")) |>
  ___() +
  xlab("statistic = odds ratio")
```

QUESTIONS

- What is the computer doing to generate the null sampling distribution of \widehat{OR} ?
- Can you describe a tactile method that could mimic what the computer code does?

Hypothesis test of OR when H_0 is true

Because the data come from a case-control study, we'll choose to investigate the OR (instead of the RR). Note that the test is one-sided because the researchers suspected that the herbicide was associated with higher cancer rates.

$$H_0 : OR = 1$$

$$H_A : OR > 1$$

I got a p-value of 0.003 (you might have gotten a slightly different answer if you set a different seed). We can reject the null hypothesis and claim that the true odds (DO NOT SAY THE WORD "risk" or "probability") of cancer is higher for those dogs with exposure to the 2, 4-D herbicide.

Remember to change to #| eval: true !

Important reflection

Note that the sampling distribution is not symmetric!

```
set.seed(4774)
obs_or <- cancer_in_dogs |>
  specify(__ ~ ____, ___ = "___") |>
  calculate(stat = "___", order = c("___", "___"))

obs_or

null_or <- cancer_in_dogs |>
  specify(__ ~ ____, ___ = "___") |>
  hypothesize(___ = "___") |>
  generate(___ = ____, ___ = "___") |>
  calculate(stat = "___", order = c("___", "___"))

___ |>
  ___() +
  ___(aes(x = ___)) +
  ___(___ = ____, color = "red") +
  xlab("statistic = odds ratio")

___ |>
  summarize(p_value = ___(___ > ___))
```

Variability of \widehat{OR} with no hypothesis

The bootstrap process (which leads to a bootstrap confidence interval) creates different possible values of \widehat{OR} **without** assuming the null hypothesis is true.

See: Bootstrapping applet: <http://www.rossmanchance.com/applets/2021/twopopprop/twopopprop.html>

Note that the variability of the statistics (the sampling distribution) can be displayed by applying the same ideas that were used on the difference in sample proportions. It is important to observe that the distribution (for \widehat{OR}) is not centered at 1.

Remember to change to #| eval: true !

```
set.seed(5)
boot_or <- cancer_in_dogs |>
  ___(response ~ order, success = "cancer") |>
  ___(reps = 1000, type = "bootstrap") |> # no hypothesize step!!!
```

```

___(stat = "odds ratio", order = c("2,4-D", "no 2,4-D"))

boot_or |>
  ___() +
  ___(aes(x = ___)) +
  xlab("statistic = odds ratio")

```

QUESTIONS

- What is the difference in the **code** from the null curve vs the no hypothesis curve?
- What is the difference in the **look of the histograms** from the null curve vs the no hypothesis curve?

Important reflection

Note that the sampling distribution is not symmetric!

Bootstrap Confidence Interval for OR

Using the histogram above (no hypothesis), the CI for the OR can be taken directly from the bootstrap sampling distributions.

We are 93% confident that the true odds of lymphoma is between 1.11 and 1.66 times higher for those dogs who have been exposed to 2, 4-D than those dogs who have not.

Remember to change to `#| eval: true!`

```

boot_or |>
  ___(lower = ___(___, ___), ___ = ___(___, ___))

boot_or |>
  ___() +
  ___(x = ___) +
  ___(xintercept = quantile(___, ___), color = "red") +
  ___(xintercept = quantile(___, ___), color = "red") +
  xlab("statistic = odds ratio")

```

Variability of \widehat{OR} assuming H_0 is true.

As with other statistics we've seen, the central limit theorem tells us about the distribution of $\ln(\widehat{OR})$, if the sample size is big enough. Indeed, the sampling distribution of $\ln(\widehat{OR})$ can be written as:

$$\ln(\widehat{OR}) \sim N\left(\ln(OR), \sqrt{\frac{1}{A} + \frac{1}{B} + \frac{1}{C} + \frac{1}{D}}\right)$$

It is hard to compare the Normal curve to the (computational) curves above because the computational curves describe the (slightly skewed) distribution of \widehat{OR} and the normal curve describes the distribution of $\ln(\widehat{OR})$. However, you should be able to sketch (with a pencil) the theoretical sampling distribution of $\ln(\widehat{OR})$, given the information above.

Write the null and alternative hypotheses in terms of the true parameter $\ln(OR)$.

Solution:

$$H_0 : \ln(OR) = 0$$

$$H_A : \ln(OR) > 0$$

Find a Z score

Using the data and the null hypothesized value of the parameter $\ln(OR)$, calculate a Z score.

Solution:

$$Z_{score} = \frac{\ln(\widehat{OR}) - 0}{\sqrt{\frac{1}{A} + \frac{1}{B} + \frac{1}{C} + \frac{1}{D}}} = \frac{\ln((191/304)/(300/641))}{\sqrt{\frac{1}{191} + \frac{1}{300} + \frac{1}{304} + \frac{1}{641}}} = 2.54$$

```
(log((191/304)/(300/641)) - 0) / sqrt((1/191) + (1/300) + (1/304) + (1/641))
```

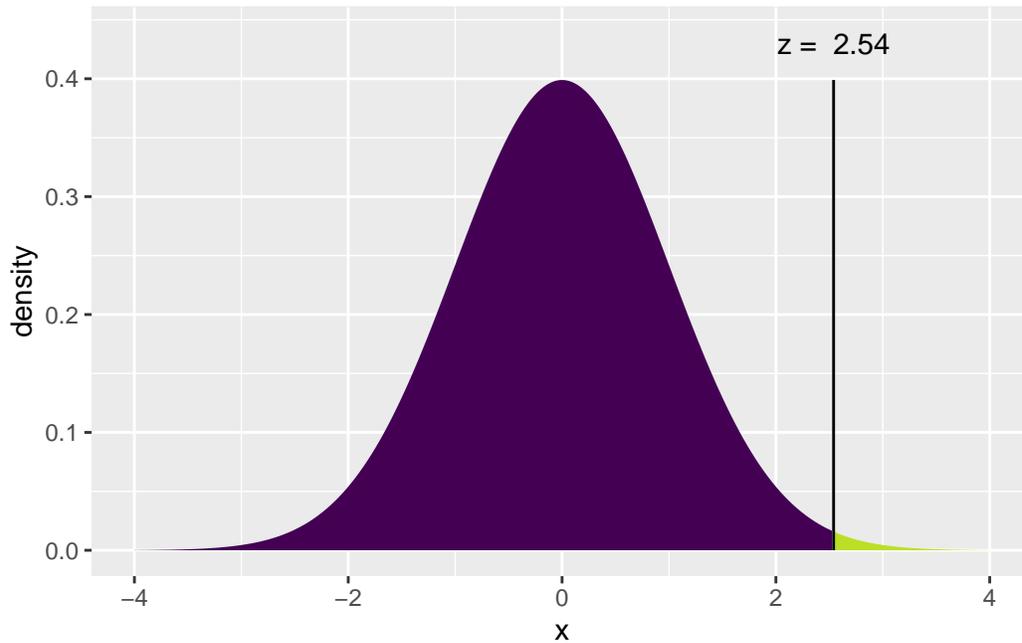
```
[1] 2.542278
```

Using the Z score, find the p-value

Solution:

The normal theory p-value is 0.006. Nowhere near the 0.05 significance level cutoff!

```
1 - xpnorm(2.54)
```



```
[1] 0.005542623
```

Conclude

Using the p-value, conclude the test using words like **odds**, herbicide, dogs, cancer.

QUESTIONS

- What is the difference in the computational sampling distribution vs. the normal theory sampling distribution?
- What is the difference in the hypothesis test conclusions for the computational sampling distribution vs. the normal theory sampling distribution?
- What is the difference between the null sampling distribution and the no hypothesis sampling distribution?

To Turn In

The data

Consider the article (and data therein) on sleepy driving: (Connor et al. “Driver sleepiness and risk of serious injury to car occupants: population based case control study”, British Medical Journal, 2002, <https://www.bmj.com/content/bmj/324/7346/1125.1.full.pdf>)

Although the researchers look at many variables, consider the following two variables: (1) driver sleepiness score of 1-3 vs 4-7, (2) driver involved in an “injury crash” or not. Note that there seems to be some missing information with respect to the sleepiness score. [See Table 1 in the paper.]

You will need to enter the data as you did in Lab 6 (after looking at Table 1 in the paper).

Q1. Variables

In the study on sleepy driving, which is the explanatory and which is the response variable? What is an observational unit?

Q2. Study

(You won't be able to answer Q2 until after class on Monday.)

- How were the data selected, as a case-control study or a cohort study?
- What aspect of the population is it impossible to know given how the data were sampled:
 - is it impossible to estimate the proportion of people who crash (of those who are sleepy)?
 - Or is it impossible to estimate the proportion of people who are sleepy (of those who crash)?

Note: in a case-control study, the participants are selected based on the value of the response variable. In a cohort study, the participants are selected based on the value of the explanatory variable. Ask yourself, how were the people in this study selected?

Q3. OR computational hypothesis test

Using a randomization test with the odds ratio statistic, complete a full hypothesis test: state null and alternative hypotheses with parameter values, the p-value from the randomization test, and the conclusion in words of the problem.

Q4. OR computational confidence interval

Find and interpret a 90% bootstrap percentile CI for the true OR of the variables above. [Note: to complete this question you will need to write out the words which describe the true OR. Do not use words like “probability”, “rate”, or “risk”. Instead, use the word “odds”.]

Q5. OR normal theory hypothesis test

Using central limit theorem, complete a full hypothesis test: state null and alternative hypotheses with parameter values, the Z score, the p-value, and the conclusion in words of the problem.

Q6. “strong” and “evidence”

- Which part(s) of the analysis reveals evidence of a strong association? Explain.
- Which part(s) of the analysis reveals strong evidence of an association? Explain.

```
praise()
```

```
[1] "You are fine!"
```