Analysis of Matched Pairs Experiments

An Experiment

Methylphenidate is FDA-approved for treating attention deficit hyperactivity disorder (ADHD) in children and adults and as a second-line treatment for narcolepsy in adults (more info here). In one experiment investigating the impact of methylphenidate on cognitive function, subjects were asked to complete a delay of gratification (DOG) task both on a dose of the drug and not on the drug (the order being randomized). Children were told that a star would appear on the computer screen if they waited "long enough" to press a response key. If a child responded in less than four seconds after their previous response, they did not earn a star, and the 4-second counter restarted. The DOG differentiates children with and without ADHD. On RShiny, the dataset is called "delay of gratification."

What type of experiment is this?

- A matched pairs experiment
- If interested, you can read more about the study here.
- You can download the original data from the experiment here.

The Data from the Experiment

Patient	D0	D60	Diff
1	57	62	5
2	27	49	22
3	32	30	-2
4	31	34	3
5	34	38	4

Variables:

- Patient = Patient Number
- D0 = Number correct on 0 Dose
- D60 = Number correct on 60mg Dose
- Diff = Difference of D60-D0
- $n = 24$

Analyzing the Data

Steps of statistical analysis:

1. Identify the population and parameter you are interested in.
    - Population:
    - Parameter:

- With matched pairs experiments, we analyze the difference column because we focus on the difference the treatment made in each patient.

Analyzing the Data

Steps of statistical analysis:

1. Identify the population and parameter you are interested in.
    - Population: Children with ADHD
    - Parameter: The mean difference in number correct on the drug minus not on the drug for all children with ADHD. We will use μ_d to denote this difference.
2. Collect data - did the experiment use all 3 principles of good experimental design (randomization, replication and control)?

Analyzing the Data

Steps of statistical analysis:

1. Identify the population and parameter you are interested in.
    - Population: Children with ADHD
    - Parameter: The mean difference in number correct on the drug minus not on the drug for all children with ADHD. We will use μ_d to denote this difference.
2. Collect data - did the experiment use all 3 principles of good experimental design (randomization, replication and control)?
    - Yes - replication (more than 1 subject), randomization (randomized the order) and control (zero dose).
3. Posit a Statistical Model - how do we this?

- Exploratory Data Analysis

EDA

You want to explore the differences:

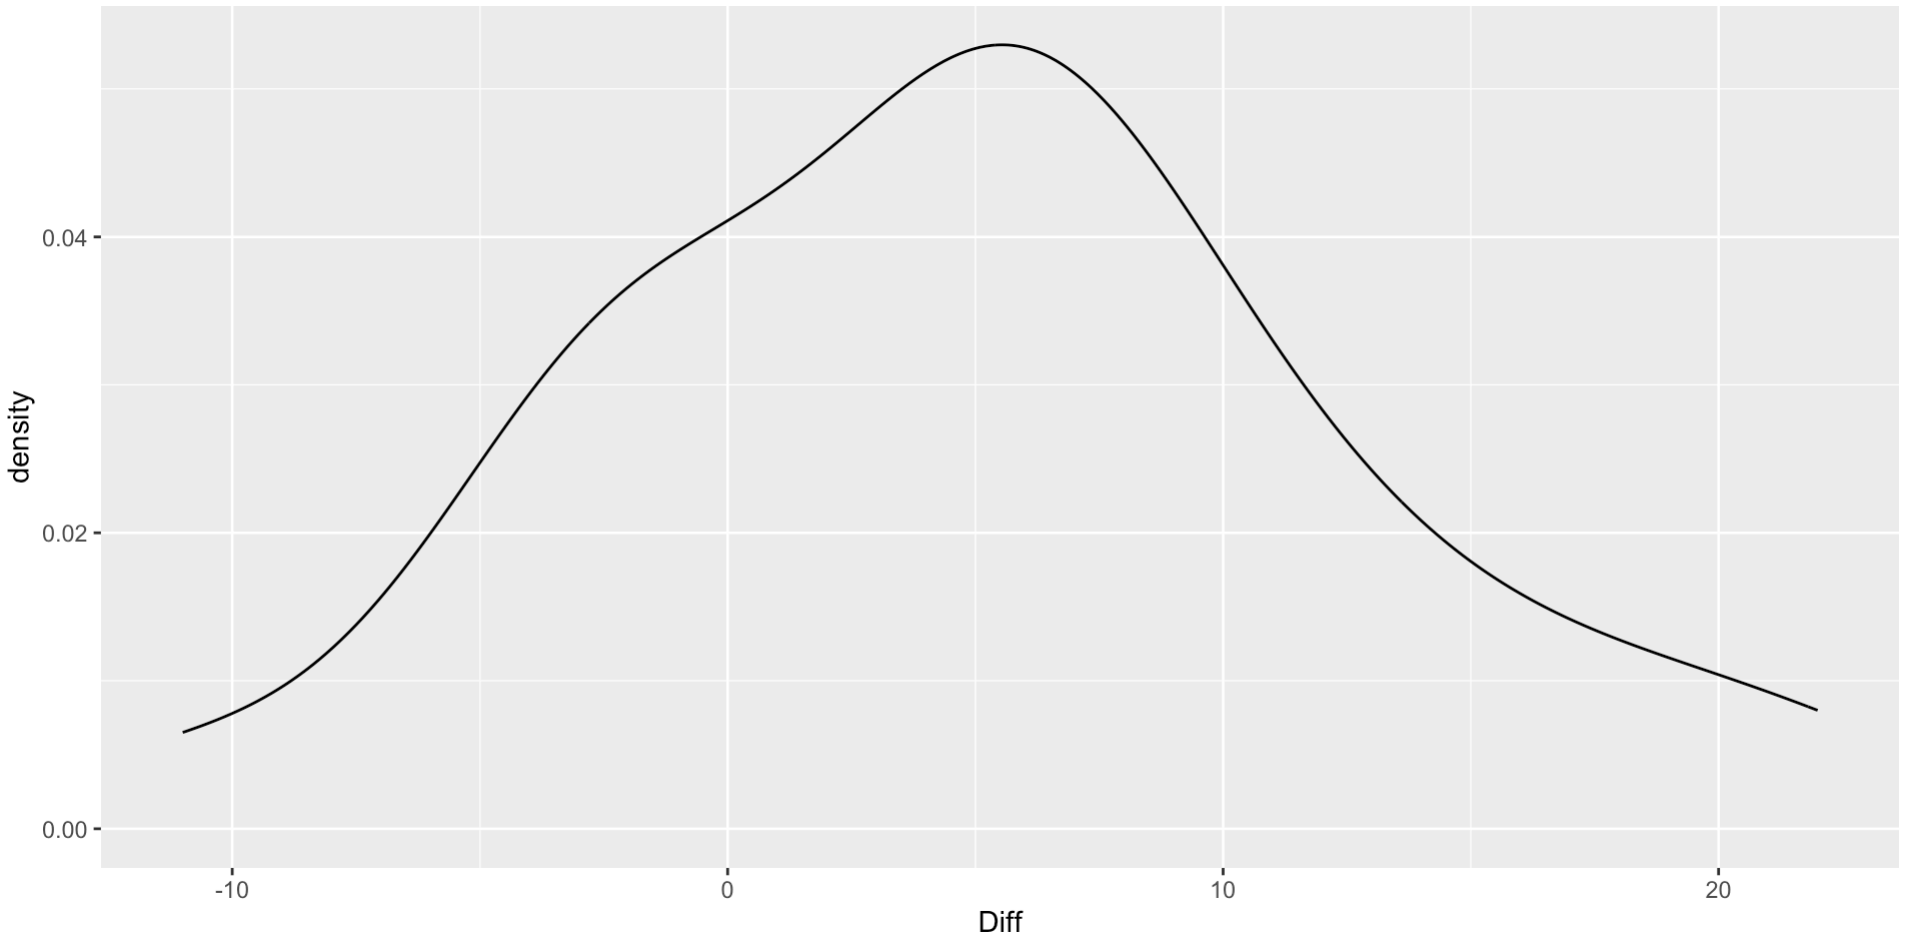

Mean: 4.958 SD: 7.538 Skew: 0.226 Median: 5

Analyzing the Data

Steps of statistical analysis:

3. Posit a Statistical Model

$$Y \sim N(\mu_d, \sigma_d)$$

How do we interpret μ_d and σ_d?

- μ_d: the average difference between the number correct on the drug minus the number correct not on the drug for all ADHD children
- σ_d the standard deviation of the differences

Analyzing the Data

4. Draw inference about the population using your model.

What are our point estimates for μ_d and σ_d?

- $\bar{y} = 4.958$
- $s = 7.538$
- Does $\bar{y} = \mu_d$ and $s = \sigma_d$?
- No! Our analysis would be better if we had a conclusion about μ_d rather than simply $\bar{y}$.

- Lets carry out a hypothesis test!

t-test

- Step 1 - Write out the Hypothesis (we did this earlier)

$$H_0 : \mu_d = 0$$
$$H_a : \mu_d > 0$$

t-test

- Step 2 - See if our data matches (or doesn't match) the null hypothesis
  - First, is the t-distribution appropriate to use here?

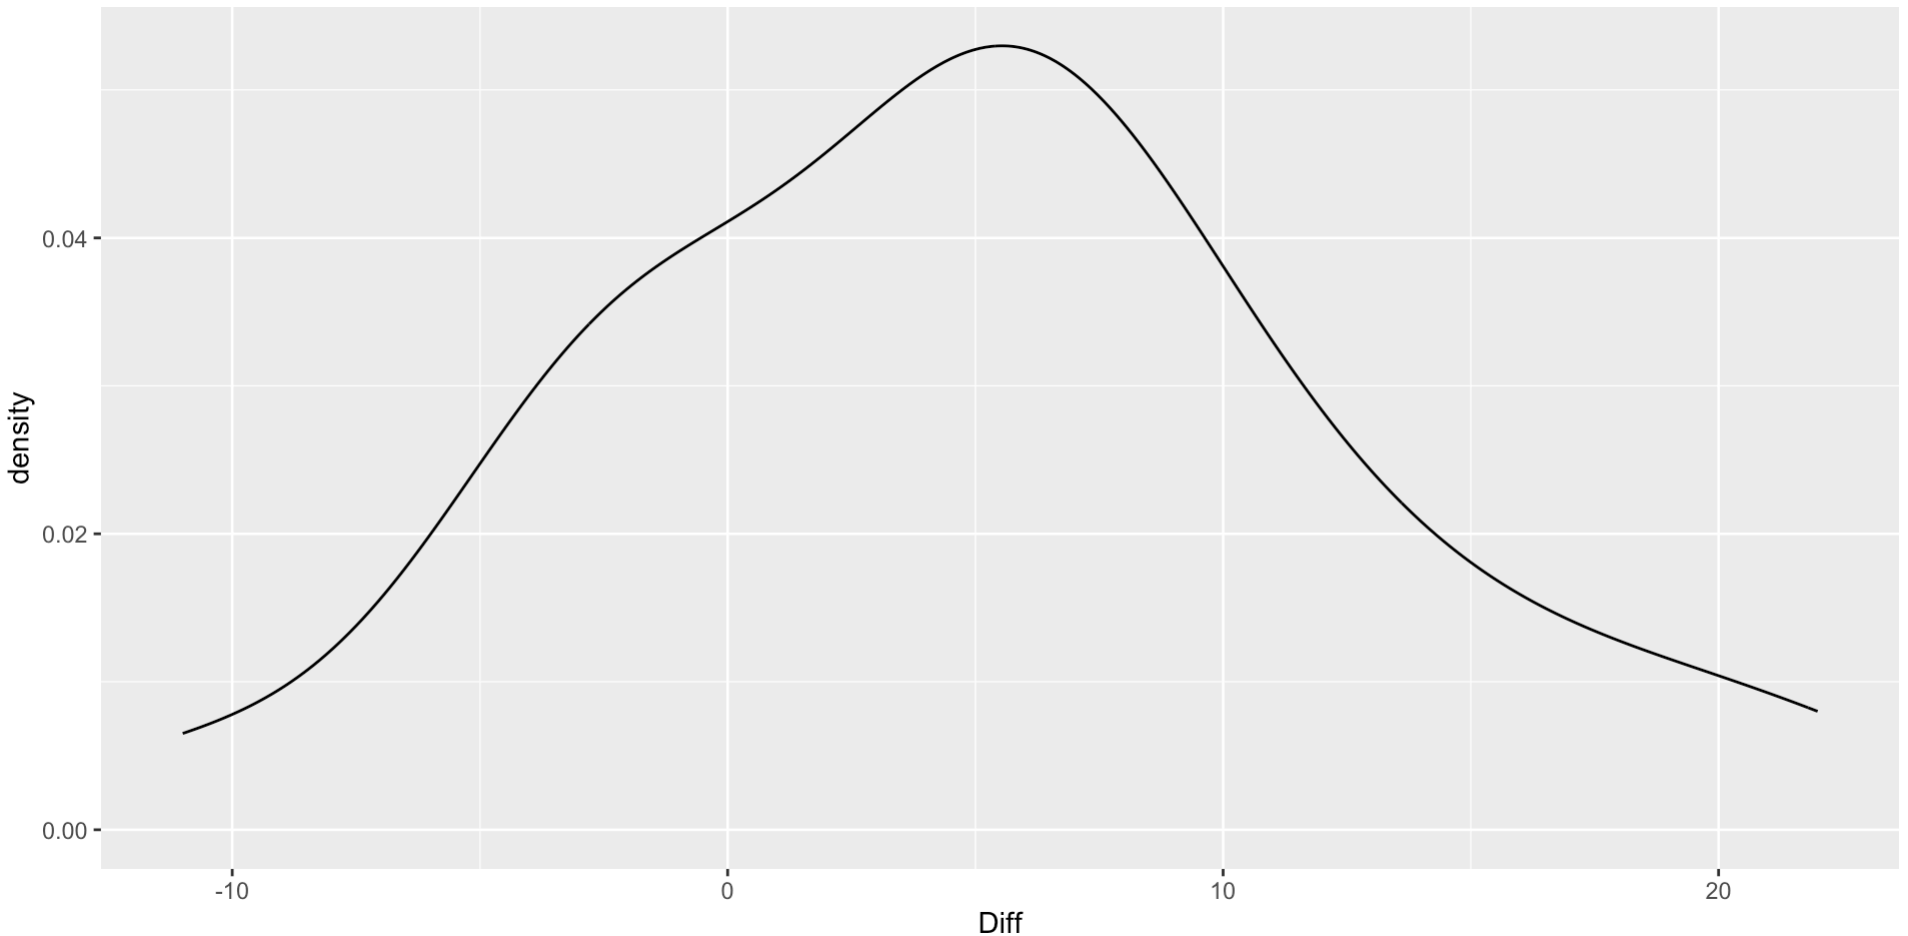

Confidence Interval

The drug helps but how much of a difference does the drug make (on average)? A 95% confidence interval for μ_d is

$$\bar{y} \pm t^\star \frac{s}{\sqrt{n}}$$

which is (1.775, 8.141). How do we interpret this interval?

- We are 95% confident that ADHD kids, on average, on a 60mg dose of Methylphenidate get between 1.775 and 8.141 more correct on the DOG task than when not on the drug.
- On the basis of this interval alone, can we say that $\mu_d = 0$?
- No because 0 is not in the interval.

Further Considerations in the Analysis

Given a 95% confidence interval of (1.775, 8.141), are these results statistically significant?

- Yes because the 95% interval does NOT contain 0.

Given a 95% confidence interval of (1.775, 8.141), are these results practically significant?

- I don't know. The scientists who published the paper seem to think so.

What would be a Type 1 error for this analysis?

- Saying the drug has an effect when, in fact, it doesn't do anything.

What would be a Type 2 error for this analysis?

- Saying the drug doesn't do anything when, in fact, it does.

Further Considerations in the Analysis

Given the two error types, do you think a Type 1 or a Type 2 error is worse?

- Debatable but I'd go with a Type 1.

Assuming a Type 1 error is worse, what do you think we should set α at (larger or smaller)?

- Smaller because we set α smaller if Type 1 error is worse.

Further Considerations in the Analysis

Using our point estimates as our best guess for the population parameters and assuming a normal population model, what is the probability that a student gets at least 10 more correct on the drug vs. not on the drug?

- 0.252, use the course analysis app

Using our point estimates as our best guess for the population parameters and assuming a normal population model, how many more correct do the top 1% of subjects get on the drug vs. not on the drug?

- 22.495 or more, use the course analysis app

Practice 2.5 Question 5

A study was conducted to investigate the placebo effect on patients. Patients who suffer from arthritis were recruited for a study and were asked to rate their average daily arthritis pain on a scale of 0 to 10 (with 0 being no pain) before being given placebo and after being given a placebo and the difference was recorded (the difference was calculated as "before" minus "after" so that large positive differences mean that the placebo "helped"). Use the Placebo dataset on the course analysis app.

Is there a statistically significant placebo effect? In other words, does thinking you have a drug improve health? Use $\alpha = 0.05$.

a. Yes because we reject the null hypothesis of zero placebo effect.

b. Yes because we fail to reject the null hypothesis of zero placebo effect.

c. No because we reject the null hypothesis of zero placebo effect.

d. No because we fail to reject the null hypothesis of zero placebo effect.

e. We don't know because the sampling distribution of t is inappropriate to use for this example.

Practice 2.5 Question 5 Answer

A study was conducted to investigate the placebo effect on patients. Patients who suffer from arthritis were recruited for a study and were asked to rate their average daily arthritis pain on a scale of 0 to 10 (with 0 being no pain) before being given placebo and after being given a placebo and the difference was recorded (the difference was calculated as "before" minus "after" so that large positive differences mean that the placebo "helped"). Use the Placebo dataset on the course analysis app.

Is there a statistically significant placebo effect? In other words, does thinking you have a drug improve health? Use $\alpha = 0.05$.

a. Yes because we reject the null hypothesis of zero placebo effect.

b. Yes because we fail to reject the null hypothesis of zero placebo effect.

c. No because we reject the null hypothesis of zero placebo effect.

d. No because we fail to reject the null hypothesis of zero placebo effect.

e. We don't know because the sampling distribution of t is inappropriate to use for this example.

Practice

A study was conducted to investigate the placebo effect on patients. Patients who suffer from arthritis were recruited for a study and were asked to rate their average daily arthritis pain on a scale of 0 to 10 (with 0 being no pain) before being given placebo and after being given a placebo and the difference was recorded (the difference was calculated as “before” minus “after” so that large positive differences mean that the placebo “helped”). Use the Placebo dataset on the course analysis app to answer the following questions.

What is a 95% confidence interval for the difference in pain before and after the placebo? In other words, how big is the placebo effect?

- We are 95% confident that patients had between (2.98, 4.32) higher pain score NOT on the placebo than on the placebo. OR, we are 95% confident that the placebo reduces pain levels by between 2.98 and 4.32.

Key Terminology

- Matched pairs experiment
- Mean of the differences (μ_d)
- Standard deviation of the differences (σ_d)