

DSC 152: Applied Statistical Data Analysis and Inference

Lecture #11
Introduction to Interactions

Tuesday, May 5
Spring Quarter 2026
Peter Chi

Where have we been?

Thus far, all of our regression examples have been with exactly ONE primary covariate of interest

Some of the Examples

- March Madness
 - bracket score vs. knowledge
- Lung Function
 - FEV vs. smoking status, adjusted for age
- Mental Health Dogs
 - Loneliness vs. treatment group, adjusted for age

What if we had TWO covariates of primary interest (and we care about both of them equally)?

Example: Mental Health Self-Experiment

A person is interested in the impact of different activities on her mental health. Specifically, she believes that going to the gym and meditation are both things that improve her mental health. She wants to formally investigate the impact of each of these things.

The setup

She will perform an experiment that will last for 28 days

- Each day will be randomly assigned for her to:
 - Either go to the gym or not
 - Either meditate or not
- At the end of each day, she will rate her happiness on a scale of 0 to 10

This is another example of A/B testing! As there are two factors, each of which has two levels, this is known as a 2^2 Factorial Experiment.

Example: Mental Health Self-Experiment

The data:

	no meditate	meditate
no gym	7, 4, 2, 8, 4, 5, 8	7, 5, 4, 8, 5, 6, 3
gym	5, 1, 8, 8, 9, 7, 6	9, 6, 8, 9, 9, 8, 10

Into R they go:

```
neither <- c(7, 4, 2, 8, 4, 5, 8)
gym <- c(5, 1, 8, 8, 9, 7, 6)
meditate <- c(7, 5, 5, 8, 5, 6, 5)
both <- c(9, 6, 8, 9, 9, 8, 10)

mh_df <- data.frame(
  score = c(neither, gym, meditate, both),
  gym = c(rep("no", 7), rep("yes", 7),
          rep("no", 7), rep("yes", 7)),
  meditate = c(rep("no", 14), rep("yes", 14))
)
```


Example: Mental Health Self-Experiment

score	gym	meditate
7	no	no
4	no	no
2	no	no
8	no	no
4	no	no
5	no	no
8	no	no
5	yes	no
1	yes	no
8	yes	no
8	yes	no
9	yes	no
7	yes	no
6	yes	no
7	no	yes
5	no	yes
5	no	yes
8	no	yes

This continues on but I cut it off due to space.

Example: Mental Health Self-Experiment


Let's visualize it



Example: Mental Health Self-Experiment

Let's visualize it: this is better


```
ggplot(data=mh_df, mapping=aes(x=gym, y=score, fill=meditate)) +  
  geom_boxplot() + ggtitle("Mental Health Comparisons")
```



Example: Mental Health Self-Experiment

Let's visualize it: another option

Mental Health Comparisons



Example: Mental Health Self-Experiment

Let's visualize it: another option

```
library(dplyr)
tmp <- mh_df %>%
  group_by(_____) %>%
  summarize(mean=mean(score))

# make line interaction plot
ggplot(data=tmp, aes(y=mean, x=gym,
                     group=meditate, color=meditate)) +
  geom_point() + geom_line() +
  ggtitle("Mental Health Comparisons") +
  ylab("mean score")
```

What should go in the `group_by` function here? [pollev.com](https://www.pollevo.com)

Your Turn #1

Now for some real data...

Consider the data from Winter's data collection of my probability theory research project...

score	handwritten	coding
48.2	No	No
NA	No	Yes
46.2	Yes	Yes
NA	Yes	Yes
NA	Yes	No
NA	Yes	Yes
31.3	No	Yes
NA	No	Yes
NA	Yes	Yes

Recall, as mentioned in a previous class, that the study design then was different than it is now.

Your Turn #1

Just like in the Mental Health example, in the Winter study design there are two factors:

- Coding Exercises (Yes/No)
- Handwritten Exercises (Yes/No)

So this is also a 2^2 Factorial Experiment. We can therefore visualize the data in exactly the same manner as the Mental Health example.

In an R Markdown file:

Load the `learn_prob.csv` dataset into R Markdown, and then:

- Make the better version of the boxplots following the code on Slide 7
- Make a line graph following the code on Slide 9 (note that it will get unhappy about the NA values...)

Note: in this dataframe, the outcome variable is `score`; the relevant covariates are `handwritten` and `coding`.

Back to the Mental Health Self-Experiment

Let's analyze it!

What are our null and alternative hypotheses??

We have two factors

- Going to the gym
- Meditating

So a reasonable null hypothesis is that they BOTH have no impact.
However...

There is also their *interaction*. What does that mean?

In simple terms, it means that if there is a positive interaction, then their impact together is different from the sum of their impacts individually.

Back to the Mental Health Self-Experiment

Let's analyze it!

```
means <- mh_df %>%  
  group_by(gym, meditate) %>%  
  summarize(mean=round(mean(score), 3))
```

gym	meditate	mean
no	no	5.429
no	yes	5.857
yes	no	6.286
yes	yes	8.429

Back to the Mental Health Self-Experiment

Let's analyze it!

gym	meditate	mean
no	no	5.429
no	yes	5.857
yes	no	6.286
yes	yes	8.429

- Among days without meditation, going to the gym results in an increase of approximately $6.286 - 5.429 = 0.857$
- Among days with meditation, going to the gym results in an increase of approximately $8.429 - 5.857 = 2.572$

Roughly speaking, an *interaction* is present when these values are different!

Back to the Mental Health Self-Experiment

Let's analyze it!

```
model1 <- lm(score ~ gym + meditate + gym*meditate, data=mh_df)
summary(model1)
```

```
##
## Call:
## lm(formula = score ~ gym + meditate + gym * meditate, data = mh_df)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -5.2857 -0.8571 -0.0714  1.5714  2.7143
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)      5.4286     0.7469   7.268 1.65e-07 ***
## gymyes           0.8571     1.0562   0.812  0.425
## meditateyes      0.4286     1.0562   0.406  0.689
## gymyes:meditateyes 1.7143     1.4938   1.148  0.262
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 1.976 on 24 degrees of freedom
## Multiple R-squared:  0.2846, Adjusted R-squared:  0.1952
## F-statistic: 3.183 on 3 and 24 DF,  p-value: 0.04208
```

Back to the Mental Health Self-Experiment

Let's analyze it!

On the previous slide, the p-value for the F-statistic is 0.0420791.

This is the overall p-value for the entire model.

That is, the null hypothesis for this p-value, in words, is:

H_0 : Each factor has no impact on the mental health score
and there is no interaction between the two.

What is the null hypothesis in mathematical notation if we let μ_i be the true mean of the i^{th} treatment group?

Interpretation of Regression Model Coefficients

Now, our linear model looks like this:

$$\hat{y} = \hat{\beta}_0 + \hat{\beta}_1 x_1 + \hat{\beta}_2 x_2 + \hat{\beta}_3 x_1 x_2$$

where

- x_1 is `gym` (treated as a 0/1 variable)
- x_2 is `meditate` (treated as a 0/1 variable)
- $x_1 x_2$ is the interaction between `gym` and `meditate`

What are the $\hat{\beta}$ s??

Interpretation of Regression Model Coefficients

Each $\hat{\beta}$ shows up under Estimate in the output table:

##	Estimate	Std. Error	t value	Pr(> t)
## (Intercept)	5.429	0.747	7.268	0.000
## gymyes	0.857	1.056	0.812	0.425
## meditateyes	0.429	1.056	0.406	0.689
## gymyes:meditateyes	1.714	1.494	1.148	0.262

For example:

- $\hat{\beta}_1 = 0.857$ represents the expected increase in mental health score associated with going to the gym, among days that she does NOT meditate.
 - Note that this matches the 0.857 that we calculated on Slide 14
- $\hat{\beta}_2 = 0.429$ follows analogously.

Interpretation of Regression Model Coefficients

Each $\hat{\beta}$ shows up under Estimate in the output table:

##	Estimate	Std. Error	t value	Pr(> t)
## (Intercept)	5.429	0.747	7.268	0.000
## gymyes	0.857	1.056	0.812	0.425
## meditateyes	0.429	1.056	0.406	0.689
## gymyes:meditateyes	1.714	1.494	1.148	0.262

For example:

- $\hat{\beta}_3 = 1.714$ is the *interaction* term. This represents the expected *difference in the effect* of going to the gym, between when she meditates and doesn't meditate.
 - That is, recall that on Slide 14, we calculated $8.429 - 5.857 = 2.572$ as being the effect of going to the gym when she DOES meditate. This is also equal to $0.857 + 1.714$

Interpretation of Regression Model Coefficients

Each $\hat{\beta}$ shows up under Estimate in the output table:

##	Estimate	Std. Error	t value	Pr(> t)
## (Intercept)	5.429	0.747	7.268	0.000
## gymyes	0.857	1.056	0.812	0.425
## meditateyes	0.429	1.056	0.406	0.689
## gymyes:meditateyes	1.714	1.494	1.148	0.262

What about the p-values?

Briefly, each of them is for the corresponding coefficient. But we need to exercise some caution when utilizing them...

Global $\mathcal{F}$ -test

```
##  
## Call:  
## lm(formula = score ~ gym + meditate + gym * meditate, data = mh_df)  
##  
## Residuals:  
##      Min       1Q   Median       3Q      Max   
## -5.2857 -0.8571 -0.0714  1.5714  2.7143   
##  
## Coefficients:  
##              Estimate Std. Error t value Pr(>|t|)      
## (Intercept)      5.4286      0.7469   7.268 1.65e-07 ***  
## gymyes           0.8571      1.0562   0.812  0.425      
## meditateyes      0.4286      1.0562   0.406  0.689      
## gymyes:meditateyes 1.7143      1.4938   1.148  0.262      
## ---  
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1  
##  
## Residual standard error: 1.976 on 24 degrees of freedom  
## Multiple R-squared:  0.2846, Adjusted R-squared:  0.1952   
## F-statistic: 3.183 on 3 and 24 DF,  p-value: 0.04208
```

The p-value of 0.0420791...

Global $\mathcal{F}$ -test

The p-value of 0.0420791 is for the overall null hypothesis:

Global $\mathcal{F}$ -test hypotheses

$$H_0 : \beta_0 = \beta_1 = \beta_2 = \beta_3 = 0$$

H_A : at least one is not equal to 0

That is, it is an overall test for whether the model describes the data in any way whatsoever.

Here, with a p-value of 0.0420791, we find statistically significant evidence that some combination of going to the gym and meditating has an impact on her mental health score.

Note: there is a mistake in H_0 here! What is it??

Back to the Individual p-values

```
##  
## Call:  
## lm(formula = score ~ gym + meditate + gym * meditate, data = mh_df)  
##  
## Residuals:  
##      Min       1Q   Median       3Q      Max   
## -5.2857 -0.8571 -0.0714  1.5714  2.7143   
##  
## Coefficients:  
##              Estimate Std. Error t value Pr(>|t|)      
## (Intercept)      5.4286     0.7469   7.268 1.65e-07 ***  
## gymyes           0.8571     1.0562   0.812  0.425      
## meditateyes      0.4286     1.0562   0.406  0.689      
## gymyes:meditateyes 1.7143     1.4938   1.148  0.262      
## ---  
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1  
##  
## Residual standard error: 1.976 on 24 degrees of freedom  
## Multiple R-squared:  0.2846, Adjusted R-squared:  0.1952   
## F-statistic: 3.183 on 3 and 24 DF,  p-value: 0.04208
```

Back to the Individual p-values

Recalling that our linear model looks like this,

$$\hat{y} = \hat{\beta}_0 + \hat{\beta}_1 x_1 + \hat{\beta}_2 x_2 + \hat{\beta}_3 x_1 x_2,$$

the interaction term has a p-value of 0.262. What H_0 is this for?

How about the p-value of 0.425?

So how do we test for main effects?

Since the provided p-value for the “main effect” coefficients are not typically of interest to us, what should we do instead?

More partial F-tests

We generally want to test main effects *in combination* with any interaction terms that involve it. This calls for a partial $\mathcal{F}$ -test!

For example, if we were primarily interested in the impact of meditation, then we would compare the full model against what null model?

pollev.com

Your Turn #2

Consider again the `learn_prob.csv` data that should already be loaded into your R Markdown file.

- First, run a linear model on these data as we saw on the Mental Health data, with an interaction term
- Report the overall p-value and your conclusion from that
- Report the p-value for the interaction term and your conclusion from that
- Suppose we are primarily interested in the impact of coding exercises. Report the p-value from the appropriate partial $\mathcal{F}$ -test for that
- Based on the regression output, what is the estimated impact of coding exercises among those who did not do handwritten exercises? How about among those who DID do handwritten exercises?

Recap and Looking Ahead

Recap

- Interaction terms are used when the impact of one covariate may DIFFER across values of another covariate
- Main effect p-values are typically not of interest; main effects should generally be assessed in combination with any interaction term involving it, using a partial $\mathcal{F}$ -test

Looking ahead

- How do we test other comparisons we might be interested in?
- What if one or both of the covariates is/are not binary?

Today's Daily Check

- Both Your Turns
- A screenshot of the confirmation of completion of the Midway Survey