R08 - Experimental design

STAT 5870 (Engineering)
Iowa State University

December 4, 2024

Random samples and random treatment assignment

Recall that the objective of data analysis is often to make an inference about a population based on a sample. For the inference to be statistically valid, we need a random sample from the population.

In order to make a causal statment, the levels of the explanatory variables need to be randomly assigned to the experimental units.

- random assignment → randomized experiment
- non-random assignment → observational study

Data collection

	Treatment randomly assigned?	
	No	Yes
Sample	Observational study	Randomized experiment
Not random	No inference to population No cause-and-effect	No inference to population Yes cause-and-effect
Random	Yes inference to population No cause-and-effect	Yes inference to population Yes cause-and-effect

Strength of wood glue

You are interested in testing two different wood glues:

- Gorilla Wood Glue
- Titebond 1413 Wood Glue

On a scarf joint:

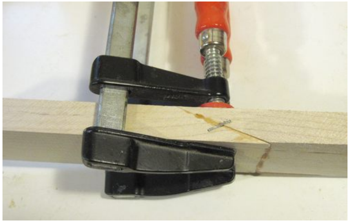

So you collect up some wood, glue the pieces together, and determine the weight required to break the joint. (Lots of details are missing.)

Inspiration: `https://woodgears.ca/joint_strength/glue.html`

Completely Randomized Design (CRD)

Suppose I have 8 pieces of wood laying around. I cut each piece and randomly use either Gorilla or Titebond glue to recombine the pieces. I do the randomization in such a way that I have exactly 4 Gorilla and 4 Titebond results, e.g.

```
# A tibble: 8 x 2
  woodID glue
  <chr>  <chr>
1 wood1  Gorilla
2 wood2  Titebond
3 wood3  Gorilla
4 wood4  Titebond
5 wood5  Titebond
6 wood6  Gorilla
7 wood7  Titebond
8 wood8  Gorilla
```

This is called a completely randomized design (CRD). Because all treatment (combinations) have the same number of replicates, the design is balanced. Because all treatment (combinations) are repeated, the design is replicated.

Visualize the data

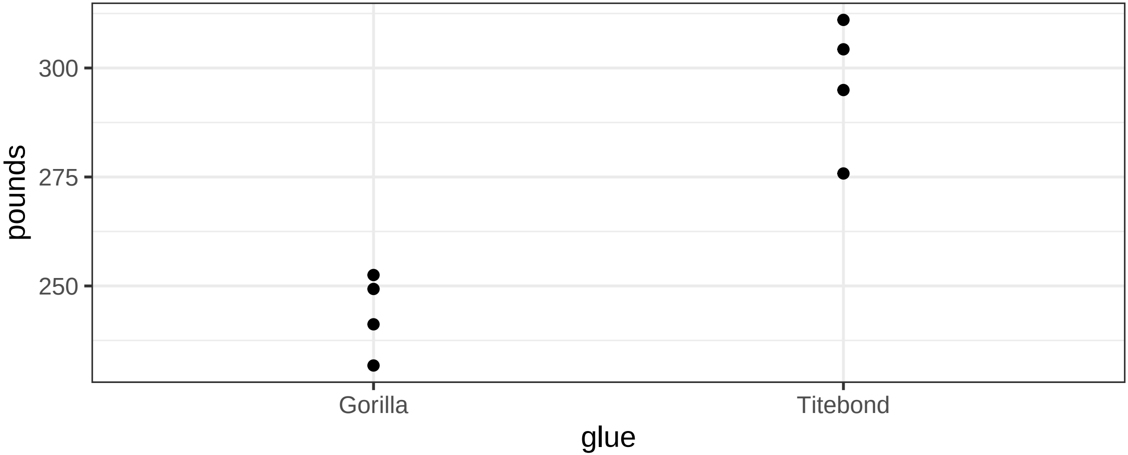

Model

Let

- P_w be the weight (pounds) needed to break wood w,
- T_w be an indicator that the Titebond glue was used on wood w, i.e.

$$T_w = \mathrm{I}(\text{glue}_w = \text{Titebond}).$$

Then a regression model for these data is

$$P_w \overset{ind}{\sim} N(\beta_0 + \beta_1 T_w, \sigma^2).$$

Check model assumptions

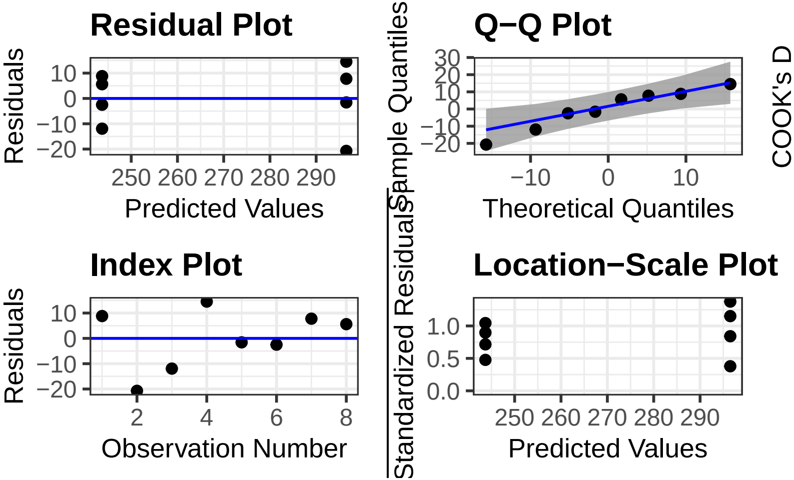

Obtain statistics

```
coefficients(m)

 (Intercept) glueTitebond
    243.6971      52.8206

summary(m)$r.squared

[1] 0.8531122

confint(m)

                 2.5 %    97.5 %
(Intercept)  228.21529 259.17885
glueTitebond  30.92606  74.71514

emmeans(m, ~glue)

 glue     emmean   SE df lower.CL upper.CL
 Gorilla     244 6.33  6      228      259
 Titebond    297 6.33  6      281      312

Confidence level used: 0.95
```

Interpret results

A randomized experiment was designed to evaluate the effectiveness of Gorilla and Titebond in preventing failures in scarf joints cut at a 20 degree angle through 1" $\times$ 2" spruce with 4 replicates for each glue type. The mean break weight (lbs) was 244 with a 95% CI of (228,259) for Gorilla and 297 (281,312) for Titebond. Titebond glue caused an increase in break weight of 53 (31,75) lbs compared to Gorilla Glue. This difference accounted for 85 % of the variability in break weight.

Randomized complete block design (RCBD)

Suppose the wood actually came from two different types: Maple and Spruce. And perhaps you have reason to believe the glue will work differently depending on the type of wood. In this case, you would want to block by wood type and perform the randomization within each block, i.e.

```
# A tibble: 8 x 3
  woodID woodtype glue
  <chr>  <fct>    <chr>
1 wood1  Spruce   Gorilla
2 wood3  Spruce   Gorilla
3 wood2  Spruce   Titebond
4 wood4  Spruce   Titebond
5 wood6  Maple    Gorilla
6 wood8  Maple    Gorilla
7 wood5  Maple    Titebond
8 wood7  Maple    Titebond
```

This is called a randomized complete block design (RCBD). If all treatment combinations exist, then the design is complete. If a treatment combination is missing, then the design is incomplete. This is experiment is replicated and balanced because each combination of woodtype and glue has more than 1 observation and the number of observations for each combination is the same, respectively.

Visualize the data

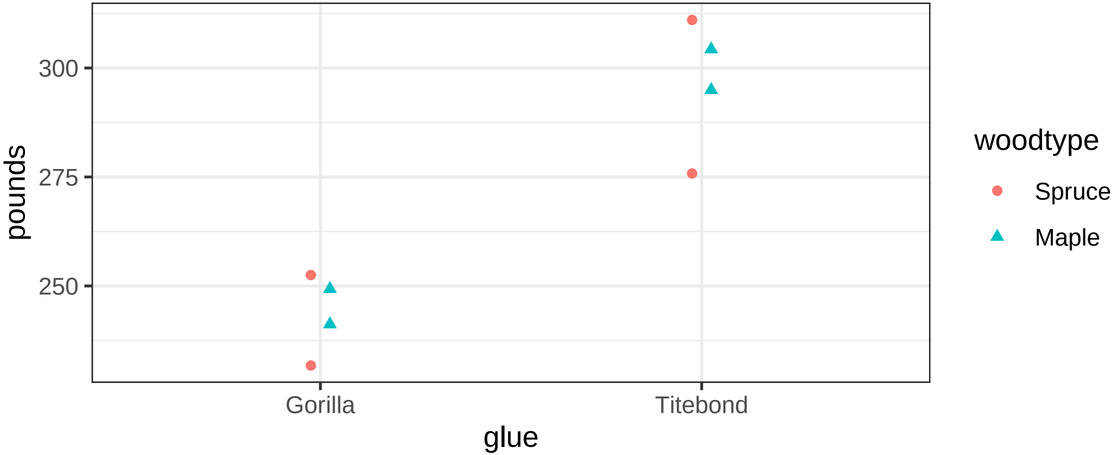

Visualize the data - a more direct comparison for glue

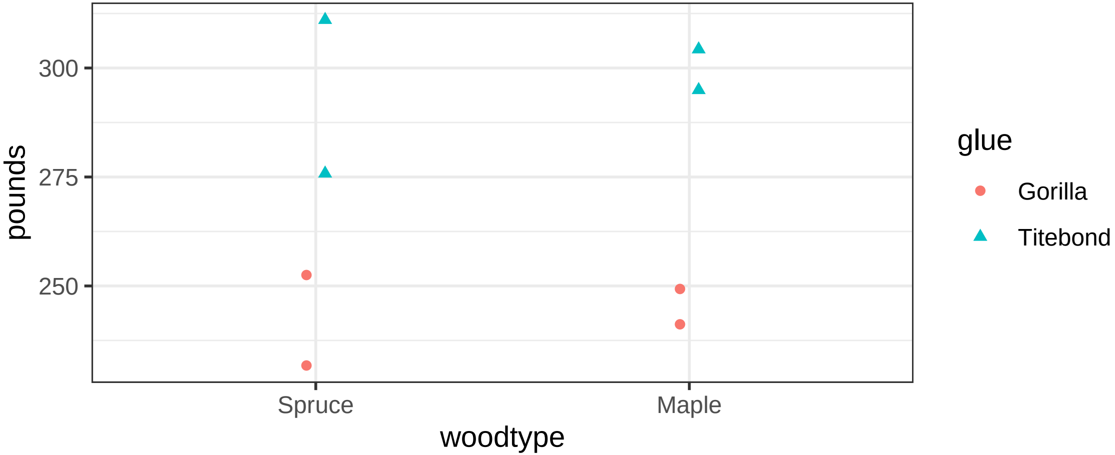

Main effects model

Let

- P_w be the weight (pounds) needed to break wood w
- T_w be an indicator that Titebond glue was used on wood w, and
- M_w be an indicator that wood w was Maple.

Then a main effects model for these data is

$$P_w \stackrel{ind}{\sim} N(\beta_0 + \beta_1 T_w + \beta_2 M_w, \sigma^2)$$

Perform analysis

```
m <- lm(pounds ~ glue + woodtype, data = d)
summary(m)


Call:
lm(formula = pounds ~ glue + woodtype, data = d)

Residuals:
      1       2       3       4       5       6       7       8 
 11.146 -18.384  -9.611  16.849  -3.902  -4.822   5.437   3.286 

Coefficients:
              Estimate Std. Error t value Pr(>|t|)    
(Intercept)    241.366      8.294  29.100 8.98e-07 ***
glueTitebond    52.821      9.578   5.515  0.00268 ** 
woodtypeMaple    4.662      9.578   0.487  0.64702    
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 13.54 on 5 degrees of freedom
Multiple R-squared:  0.8598,Adjusted R-squared:  0.8037 
F-statistic: 15.33 on 2 and 5 DF,  p-value: 0.007365

confint(m)


                  2.5 %    97.5 %
(Intercept)   220.04467 262.68760
glueTitebond   28.20070  77.44051
woodtypeMaple -19.95804  29.28177
```

Replication

Since there are more than one observation for each woodtype-glue combination, the design is replicated:

```
d |> group_by(woodtype, glue) |> summarize(n = n())

# A tibble: 4 x 3
# Groups:   woodtype [2]
  woodtype glue         n
  <fct>    <chr>    <int>
1 Spruce   Gorilla      2
2 Spruce   Titebond     2
3 Maple    Gorilla      2
4 Maple    Titebond     2
```

When the design is replicated, we can consider assessing an interaction.

Interaction model

Let

- P_w be the weight (pounds) needed to break wood w
- T_w be an indicator that Titebond glue was used on wood w, and
- M_w be an indicator that wood w was Maple.

Then a model with the interaction for these data is

$$P_w \stackrel{ind}{\sim} N(\beta_0 + \beta_1 T_w + \beta_2 M_w + \beta_3 T_w M_w, \sigma^2)$$

Assessing an interaction using a t-test

```
m <- lm(pounds ~ glue * woodtype, data = d)
summary(m)


Call:
lm(formula = pounds ~ glue * woodtype, data = d)

Residuals:
      1       2       3       4       5       6       7       8 
 10.379 -17.616 -10.379  17.616  -4.670  -4.054   4.670   4.054 

Coefficients:
                           Estimate Std. Error t value Pr(>|t|)    
(Intercept)                 242.134     10.680  22.671 2.24e-05 ***
glueTitebond                 51.285     15.104   3.395   0.0274 *  
woodtypeMaple                 3.127     15.104   0.207   0.8461    
glueTitebond:woodtypeMaple    3.070     21.361   0.144   0.8927    
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 15.1 on 4 degrees of freedom
Multiple R-squared:  0.8605,Adjusted R-squared:  0.7558 
F-statistic: 8.223 on 3 and 4 DF,  p-value: 0.03475
```

Assessing an interaction using an F-test

```
anova(m)

Analysis of Variance Table

Response: pounds
              Df Sum Sq Mean Sq F value   Pr(>F)
glue           1 5580.0  5580.0 24.4582 0.007786 **
woodtype       1   43.5    43.5  0.1905 0.685012
glue:woodtype  1    4.7     4.7  0.0207 0.892654
Residuals      4  912.6   228.1
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

drop1(m, test='F')

Single term deletions

Model:
pounds ~ glue * woodtype
              Df Sum of Sq    RSS    AIC F value Pr(>F)
<none>                     912.58 45.895
glue:woodtype  1     4.714 917.30 43.936  0.0207 0.8927
```

What if this had been your data?

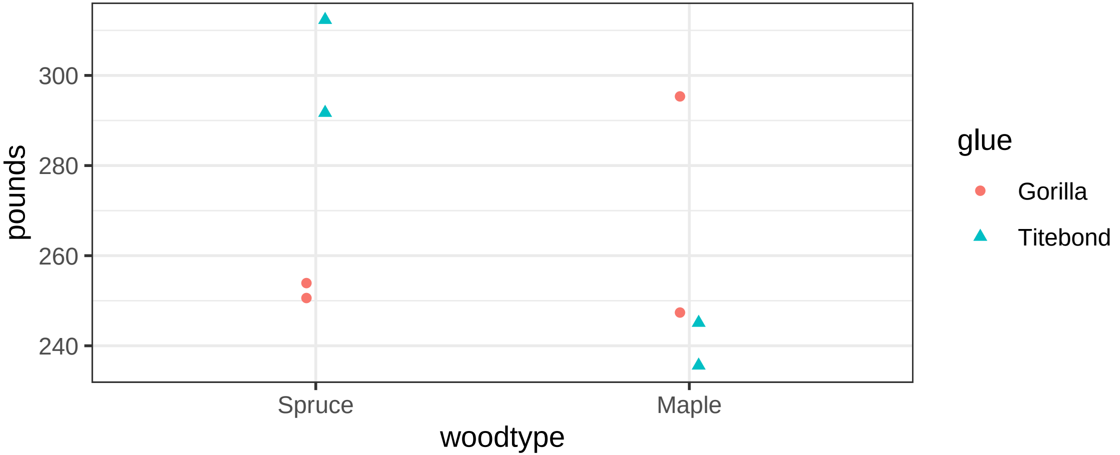

Assessing an interaction using a t-test

```
m <- lm(pounds ~ glue * woodtype, data = d)
summary(m)


Call:
lm(formula = pounds ~ glue * woodtype, data = d)

Residuals:
      1       2       3       4       5       6       7       8
  1.657  -1.657 -10.312  10.312  -4.741  23.986   4.741 -23.986

Coefficients:
                          Estimate Std. Error t value Pr(>|t|)
(Intercept)                 252.26      13.29  18.976 4.54e-05 ***
glueTitebond                 49.76      18.80   2.647   0.0572 .
woodtypeMaple                19.10      18.80   1.016   0.3670
glueTitebond:woodtypeMaple  -80.76      26.59  -3.038   0.0385 *
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 18.8 on 4 degrees of freedom
Multiple R-squared:  0.7544,Adjusted R-squared:  0.5702
F-statistic: 4.095 on 3 and 4 DF,  p-value: 0.1034
```

Unreplicated study

Suppose you now have

- 5 glue choices
- 4 different types of wood with
- 5 samples of each type of wood.

Thus you can only run each glue choice once on each type of wood.

Then you can run an unreplicated RCBD.

Visualize

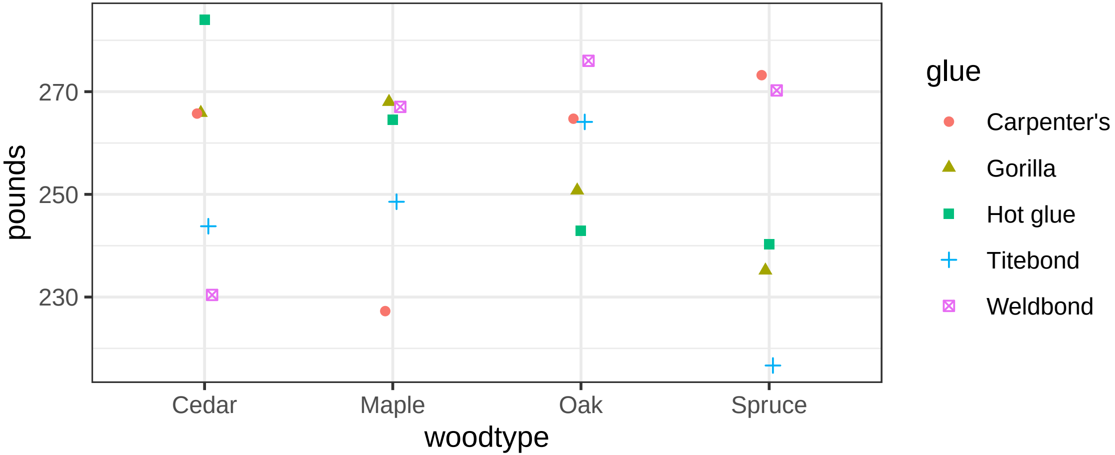
270
250
230
pounds
Cedar
Maple
Oak
Spruce
woodtype
glue
Carpenter's
Gorilla
Hot glue
Titebond
Weldbond

Fit the main effects (or additive) model

```
m <- lm(pounds ~ glue + woodtype, data = d)
anova(m)

Analysis of Variance Table

Response: pounds
          Df Sum Sq Mean Sq F value Pr(>F)
glue       4  754.3  188.58  0.4332 0.7822
woodtype   3  465.1  155.04  0.3562 0.7857
Residuals 12 5223.7  435.31
```

Fit the main effects (or additive) model

```
Call:
lm(formula = pounds ~ glue + woodtype, data = d)

Residuals:
    Min      1Q  Median      3Q     Max 
-33.498 -10.327   5.084  10.989  23.325 

Coefficients:
               Estimate Std. Error t value Pr(>|t|)    
(Intercept)    260.7220    13.1956  19.758 1.61e-10 ***
glueGorilla     -2.7764    14.7531  -0.188    0.854    
glueHot glue     0.2159    14.7531   0.015    0.989    
glueTitebond   -14.4517    14.7531  -0.980    0.347    
glueWeldbond     3.1903    14.7531   0.216    0.832    
woodtypeMaple   -2.8726    13.1956  -0.218    0.831    
woodtypeOak      1.7564    13.1956   0.133    0.896    
woodtypeSpruce -10.8349    13.1956  -0.821    0.428    
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 20.86 on 12 degrees of freedom
Multiple R-squared:  0.1893,Adjusted R-squared:  -0.2837 
F-statistic: 0.4002 on 7 and 12 DF,  p-value: 0.8845
```

Fit the full (with interaction) model

```
Warning in anova.lm(m):  ANOVA F-tests on an essentially perfect fit are unreliable


Analysis of Variance Table

Response: pounds
              Df Sum Sq Mean Sq F value Pr(>F)
glue           4  754.3  188.58     NaN    NaN
woodtype       3  465.1  155.04     NaN    NaN
glue:woodtype 12 5223.7  435.31     NaN    NaN
Residuals      0    0.0     NaN
```

Fit the full (with interaction) model

```
Call:
lm(formula = pounds ~ glue * woodtype, data = d)

Residuals:
ALL 20 residuals are 0: no residual degrees of freedom!

Coefficients:
                             Estimate Std. Error t value Pr(>|t|)
(Intercept)                  265.7301        NaN     NaN      NaN
glueGorilla                    0.1451        NaN     NaN      NaN
glueHot glue                  18.2476        NaN     NaN      NaN
glueTitebond                 -21.9394        NaN     NaN      NaN
glueWeldbond                 -35.3158        NaN     NaN      NaN
woodtypeMaple                -38.4658        NaN     NaN      NaN
woodtypeOak                   -1.0001        NaN     NaN      NaN
woodtypeSpruce                 7.4822        NaN     NaN      NaN
glueGorilla:woodtypeMaple     40.6031        NaN     NaN      NaN
glueHot glue:woodtypeMaple    19.0424        NaN     NaN      NaN
glueTitebond:woodtypeMaple    43.2335        NaN     NaN      NaN
glueWeldbond:woodtypeMaple    75.0869        NaN     NaN      NaN
glueGorilla:woodtypeOak      -14.1101        NaN     NaN      NaN
glueHot glue:woodtypeOak     -40.0202        NaN     NaN      NaN
glueTitebond:woodtypeOak      21.3197        NaN     NaN      NaN
glueWeldbond:woodtypeOak      46.5929        NaN     NaN      NaN
glueGorilla:woodtypeSpruce   -38.1789        NaN     NaN      NaN
glueHot glue:woodtypeSpruce  -51.1490        NaN     NaN      NaN
glueTitebond:woodtypeSpruce  -34.6024        NaN     NaN      NaN
glueWeldbond:woodtypeSpruce   32.3448        NaN     NaN      NaN
```

Summary

- Designs:
  - Completely randomized design (CRD)
  - Randomized complete block design (RCBD)
- Deviations
  - Unreplicated
  - Unbalanced
  - Incomplete