"99 percent of all statistics only tell 49 percent of the story"
Ron DeLegge II (Economist)

Statistics in its best -
revealing unexpected effects

THE
WOW Factor!

General linear models (not Generalized linear model)

- Traditionally, authors tend to separate t-tests, ANOVA, Regression and Analysis of Covariance (ANCOVAs). However, because they share the same calculations (and theories and assumptions), we often classify these methods under the category of ***General Linear Models***.

- *General Linear Models* (unlike Generalized linear models) assume that response variables are normally distributed.

General linear models

Linear Model	Common name
$Y = \mu + X$	Simple linear regression
$Y = \mu + A_1$	One-factorial (one-way) ANOVA
$Y = \mu + A_1 + A_2 + A_1 \times A_2$	Two-factorial (two-way) ANOVA
$Y = \mu + A_1 + X + (A_1 \times X)$	Analysis of Covariance (ANCOVA)
$Y = \mu + X_1 + X_2 + X_3$	Multiple regression
$Y = \mu + A_1 + g + A_1 \times g$	Mixed model ANOVA
$Y_1 + Y_2 = \mu + A_1 + A_2 + A_1 \times A_2$	Multivariate ANOVA (MANOVA)

Y (response) is a continuous variable
X (predictor) is a continuous variable
A represents categorical predictors (factors)
g represents groups of data (more on this later)

$Y = \mu + A_1 + X_1 + (A_1 \times X_1)$	Analysis of Covariance (ANCOVA)

- Test for differences in slopes among groups (treatments).

- Adjust for the effects of a covariate X_1 (continuous) in an ANOVA design (response variable Y and a categorical variable A_1).

- μ the grand mean (i.e., mean of the response across all observations independent of their groups).

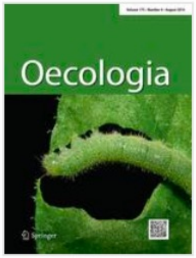

Oecologia
June 1992, Volume 90, Issue 3, pp 435–444 | Cite as

The effects of grazers on the performance of individuals and populations of scarlet gilia, *Ipomopsis aggregata*

Authors | Authors and affiliations

Joy Bergelson, Michael J. Crawley

I. aggregata exhibits considerable powers of regrowth following removal of its primary shoot by herbivores, but we found no evidence of overcompensation (i.e. of significantly higher plant performance where plants were exposed to ungulate herbivory) in a comparison between individuals on grazed and ungrazed sides of exclosure fences

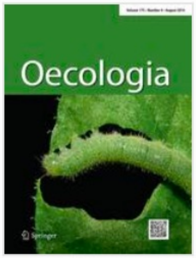

Oecologia
June 1992, Volume 90, Issue 3, pp 435–444 | Cite as

The effects of grazers on the performance of individuals and populations of scarlet gilia, *Ipomopsis aggregata*

Authors | Authors and affiliations

Joy Bergelson, Michael J. Crawley

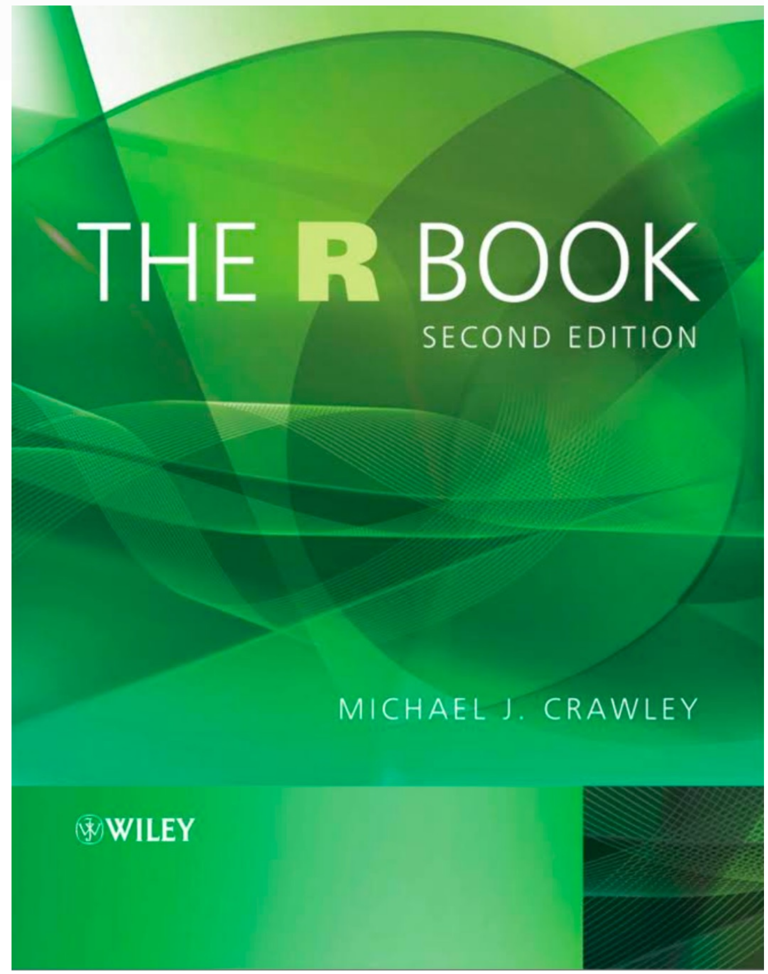

Study design and data

- 40 plants (*Ipomopsis*) were allocated to a grazing factor (two levels: grazed or ungrazed):

1) Grazed plants were exposed to grazers for the first two weeks of stem elongation (initial plant size measured as diameter of the rootstock top).

2) After two weeks, fence was built to prevent grazing.

3) At the end of the growing season, fruit production (dry weight in *mg* was recorded for each plant.

4) Initial plant size (diameter of the *rootstock* top, i.e., root size) was thought to influence fruit production and it was also measured.



Data structure

Fruit production (dry weight in mg).

Grazing (Yes / No).

Initial plant size (diameter of the rootstock top).

Data Structure

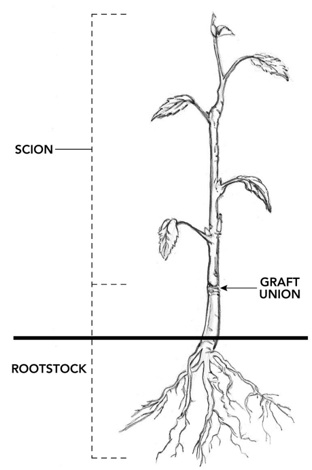

B	C	D
Fruit	**Grazing**	**Root**
59.77	Ungrazed	6.225
60.98	Ungrazed	6.487
14.73	Ungrazed	4.919
19.28	Ungrazed	5.13
34.25	Ungrazed	5.417
35.53	Ungrazed	5.359
87.73	Ungrazed	7.614
63.21	Ungrazed	6.352
24.25	Ungrazed	4.975
64.34	Ungrazed	6.93
52.92	Ungrazed	6.248
32.35	Ungrazed	5.451
53.61	Ungrazed	6.013
54.86	Ungrazed	5.928
64.81	Ungrazed	6.264
73.24	Ungrazed	7.181
80.64	Ungrazed	7.001
18.89	Ungrazed	4.426
75.49	Ungrazed	7.302
46.73	Ungrazed	5.836
116.05	Grazed	10.253
38.94	Grazed	6.958
60.77	Grazed	8.001
84.37	Grazed	9.039
70.11	Grazed	8.91
14.95	Grazed	6.106
70.7	Grazed	7.691
80.31	Grazed	8.988
82.35	Grazed	8.975
105.07	Grazed	9.844
73.79	Grazed	8.508
50.08	Grazed	7.354
78.28	Grazed	8.643
41.48	Grazed	7.916
98.47	Grazed	9.351
40.15	Grazed	7.066
52.26	Grazed	8.158
46.64	Grazed	7.382
71.01	Grazed	8.515
83.03	Grazed	8.53

The data: initial thoughts?

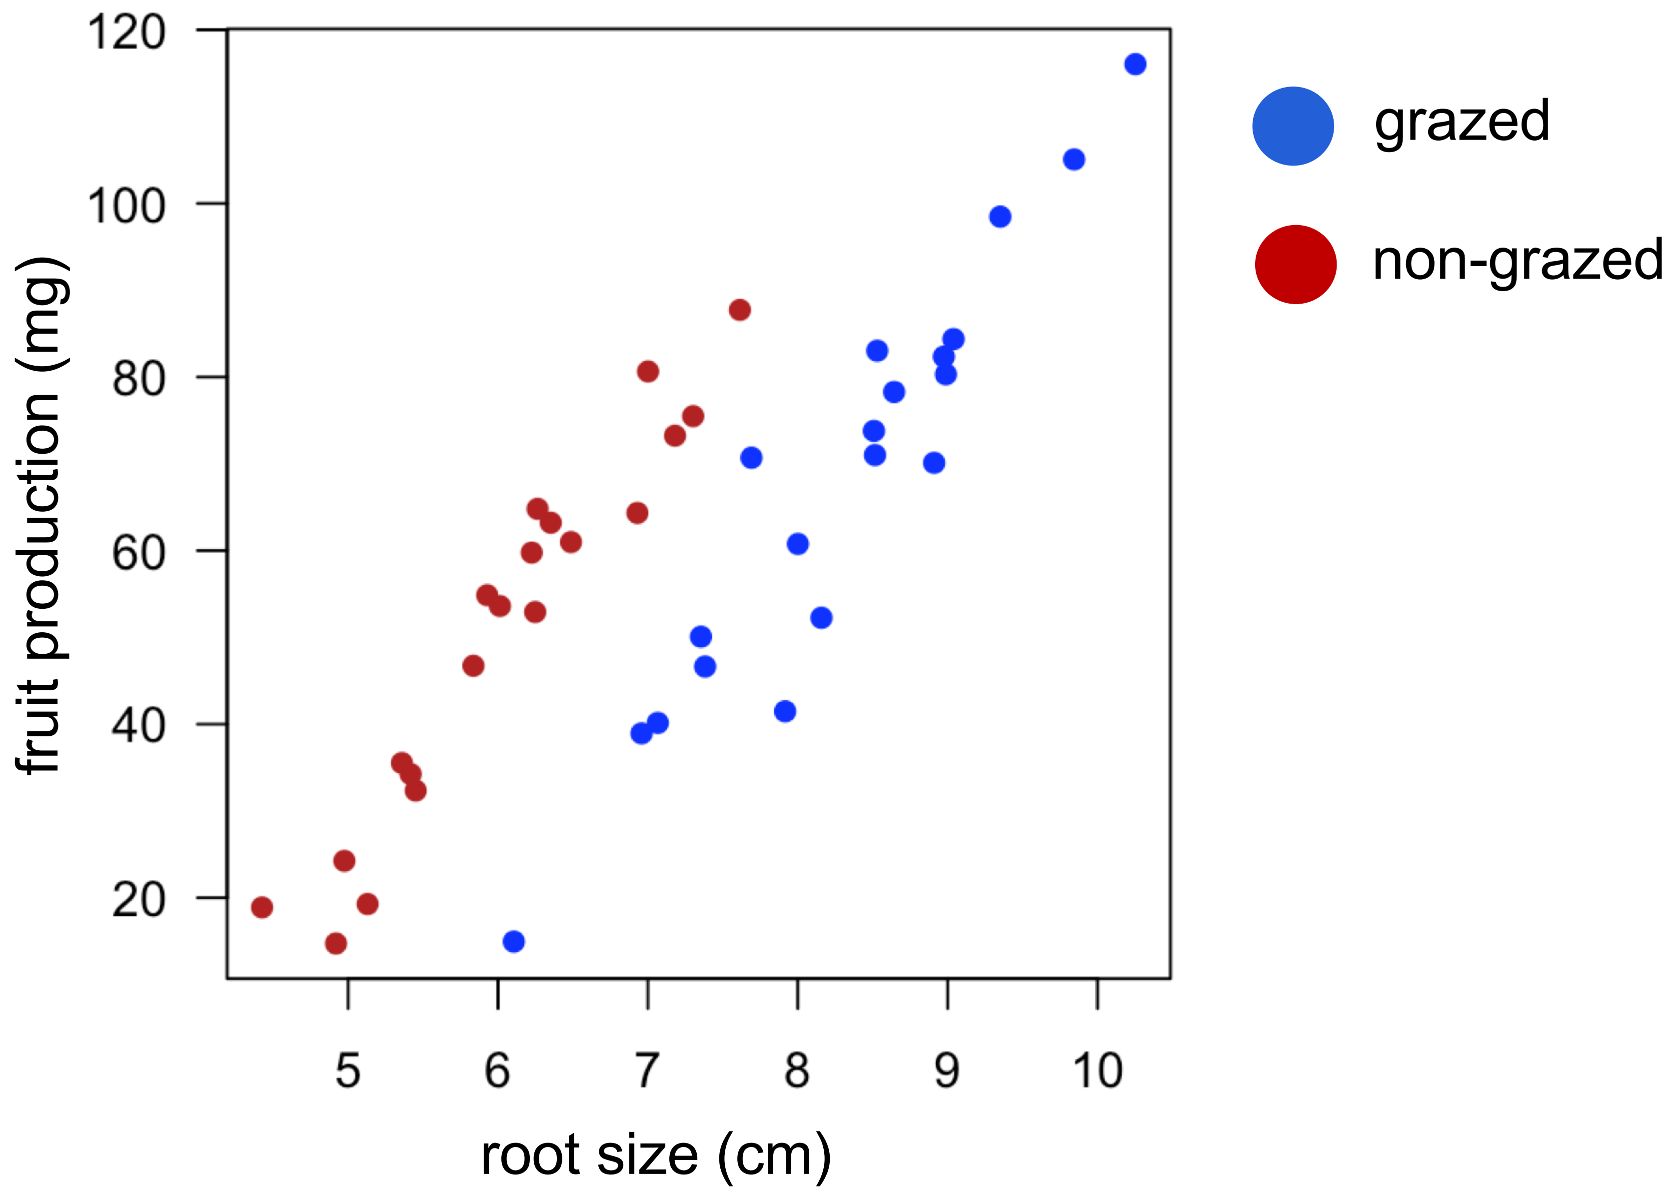

Key initial observations about the grazing experiment data

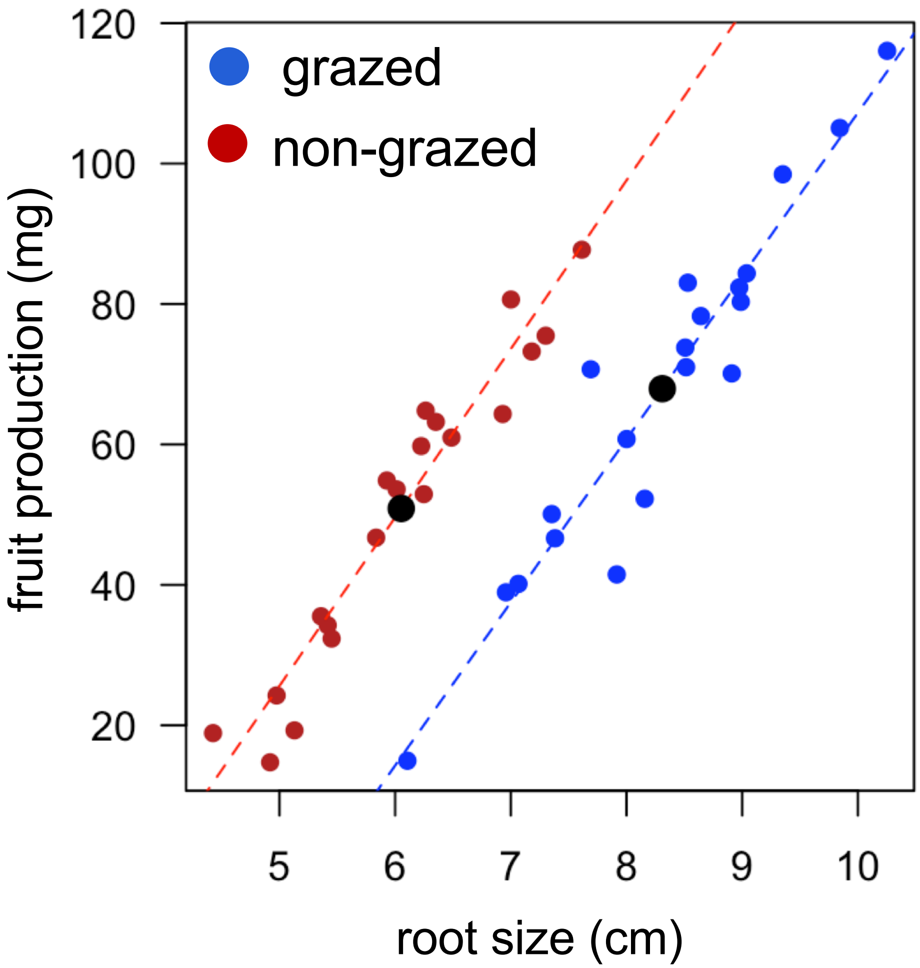

1) Plants with different initial sizes were allocated to the two treatments, i.e., The regression line for the non-grazed plants is above the line for the grazed plants.

2) Grazed plants produced more fruits.

3) The two regression lines are almost parallel.

Two major statistical potential biases tackled by ANCOVA:

1) Were plants assigned to treatments randomly according to initial size?

2) How could the assignment influence the interpretation of results?

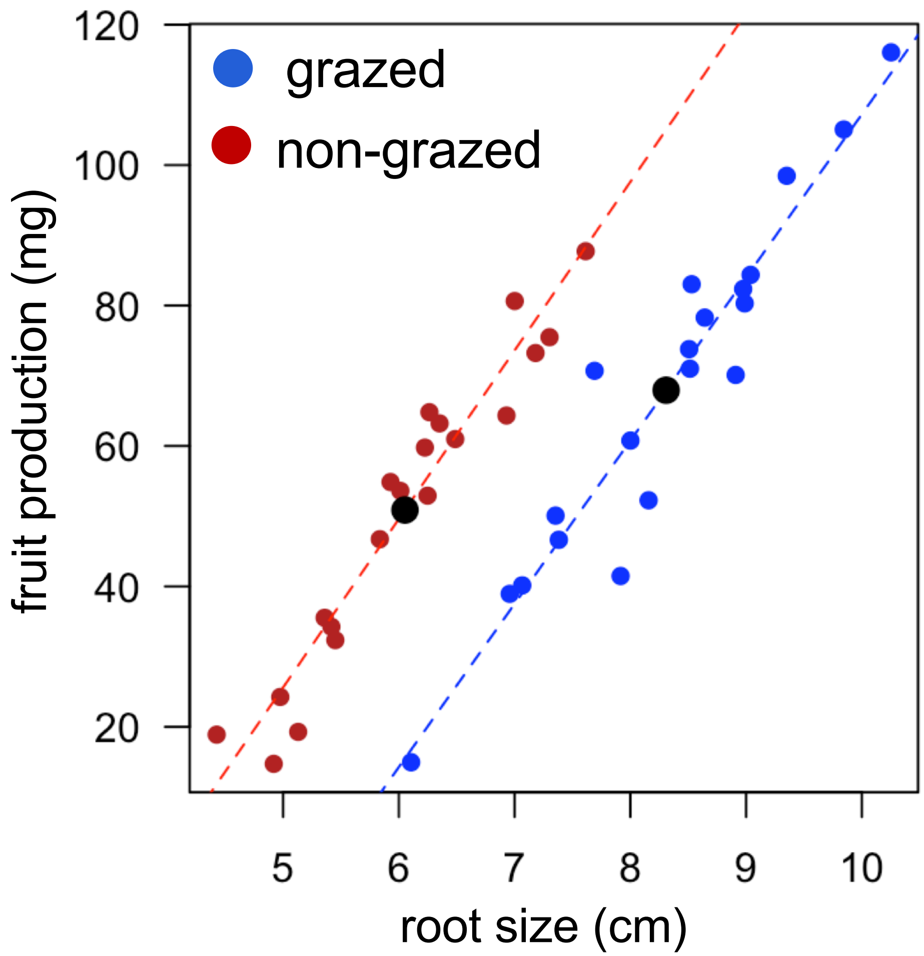

1) Plants with different initial sizes were allocated to the two treatments, i.e., The regression line for the non-grazed plants is above the line for the grazed plants.

2) Grazed plants produced more fruits.

3) The two regression lines are almost parallel.

Let's start with a simple ANOVA (a two-sample *t* test could had been used as well; $t^2 = F$) comparing the fruit production as a function of grazing (i.e., grazed, non-grazed).

```
> anova(lm(Fruit~Grazing))
Analysis of Variance Table

Response: Fruit
          Df  Sum Sq Mean Sq F value  Pr(>F)
Grazing    1  2910.4 2910.44  5.3086 0.02678 *
Residuals 38 20833.4  548.25
```

What's the conclusion?

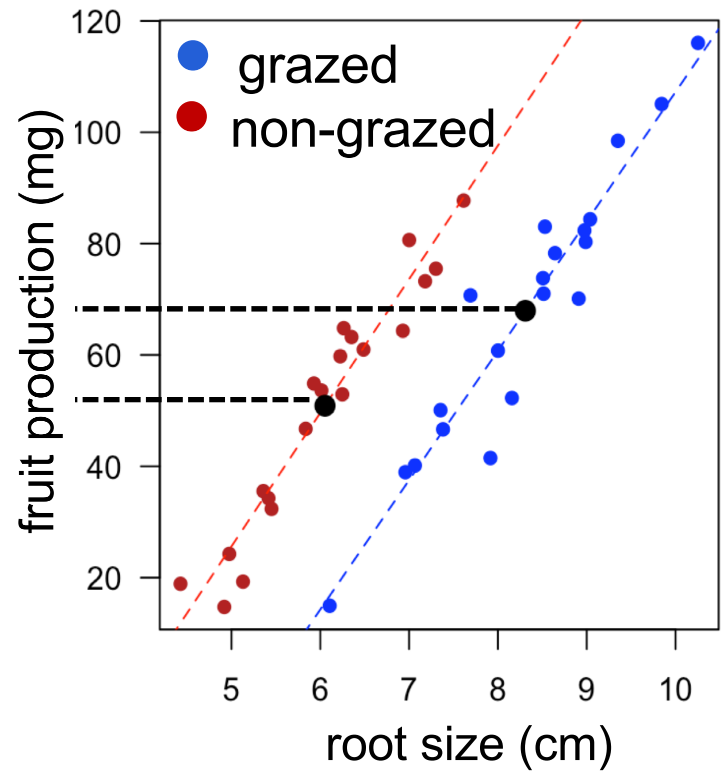

Let's start with a simple ANOVA (a two-sample *t* test could had been used as well; t^2 = F) comparing the fruit production as a function of grazing (i.e., grazed, non-grazed).

```
> anova(lm(Fruit~Grazing))
Analysis of Variance Table

Response: Fruit
          Df  Sum Sq Mean Sq F value  Pr(>F)
Grazing    1  2910.4 2910.44  5.3086 0.02678 *
Residuals 38 20833.4  548.25
```

What's the conclusion?

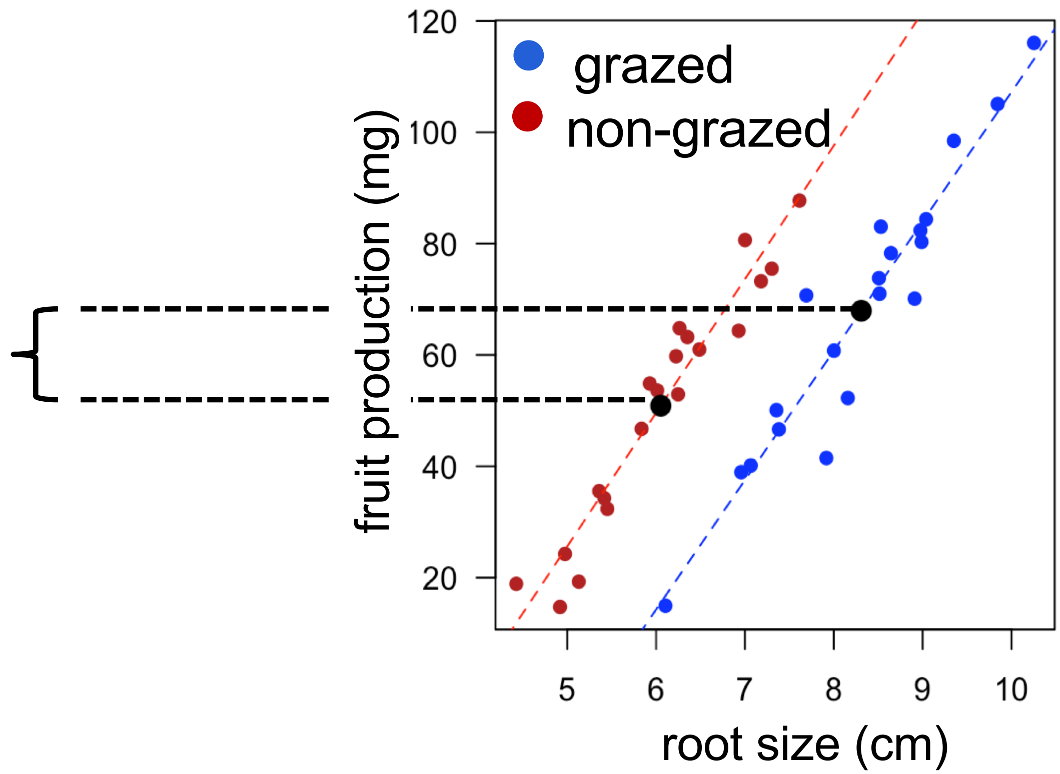

Greater fruit production under grazing!

Greater fruit production under grazing! **Can this conclusion be justified given that the initial root sizes in grazed plants are larger than non-grazed plants?**

```
> anova(lm(Root~Grazing))
Analysis of Variance Table

Response: Root
          Df Sum Sq Mean Sq F value    Pr(>F)
Grazing    1 50.918  50.918  56.087 5.411e-09 ***
Residuals 38 34.498   0.908
```

```
> anova(lm(Fruit~Grazing))
Analysis of Variance Table

Response: Fruit
          Df  Sum Sq Mean Sq F value  Pr(>F)
Grazing    1  2910.4 2910.44  5.3086 0.02678 *
Residuals 38 20833.4  548.25
```

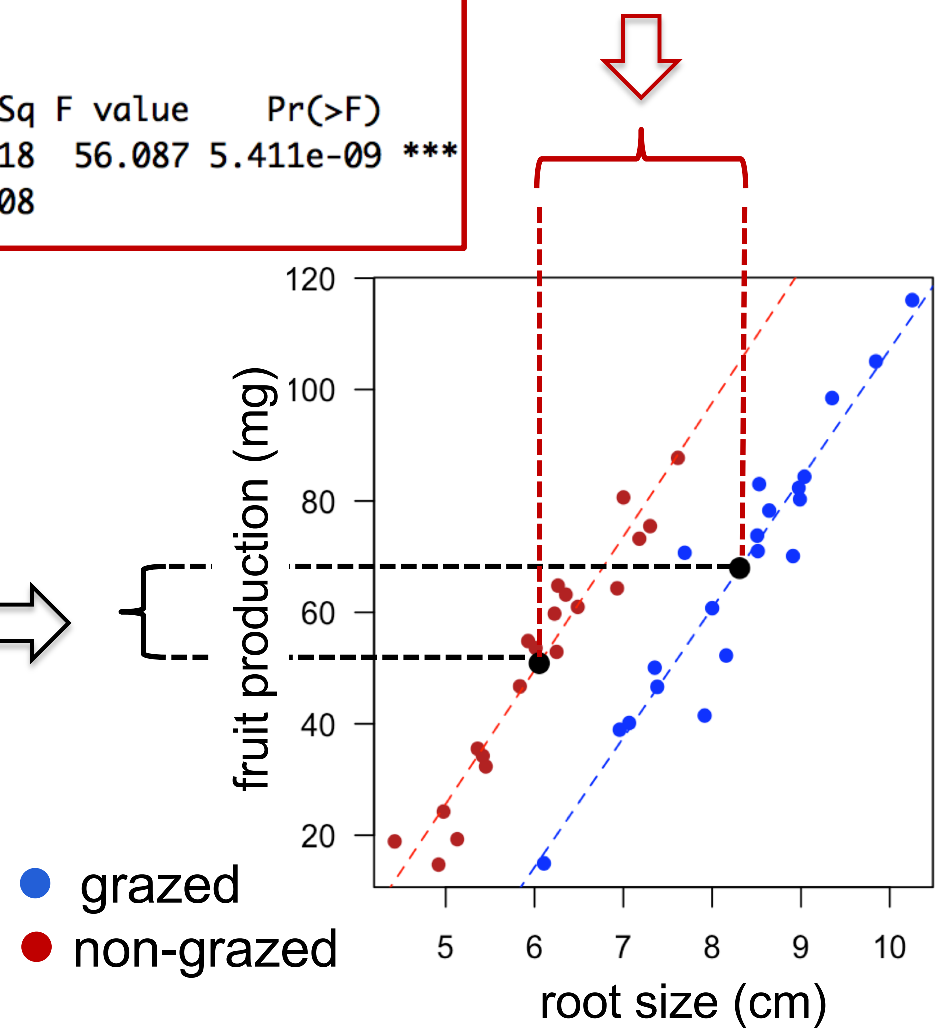

Greater fruit production under grazing! **Can this conclusion be justified given that the initial root sizes in grazed plants are larger than non-grazed plants?**

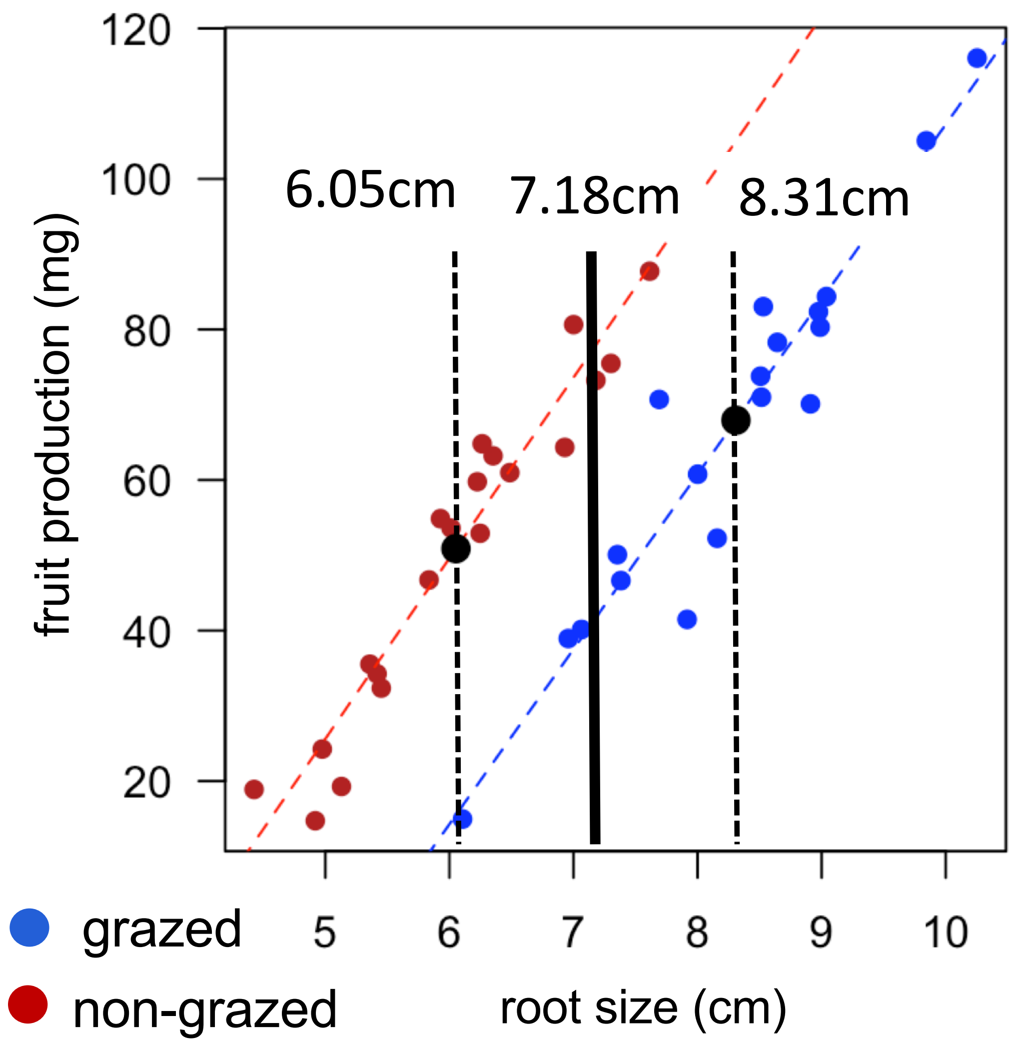

What would the effect of grazing be if all plants had started with the same initial (root) size?
Say the average value for the two treatments = $(6.05 + 8.31)/2 = 7.18$cm

Assuming a common initial mean size of root size (7.18cm), we can adjust (predict) the fruit production for each treatment (grazed / non-grazed), i.e., as if root size would had been the same in the beginning of the experiment.

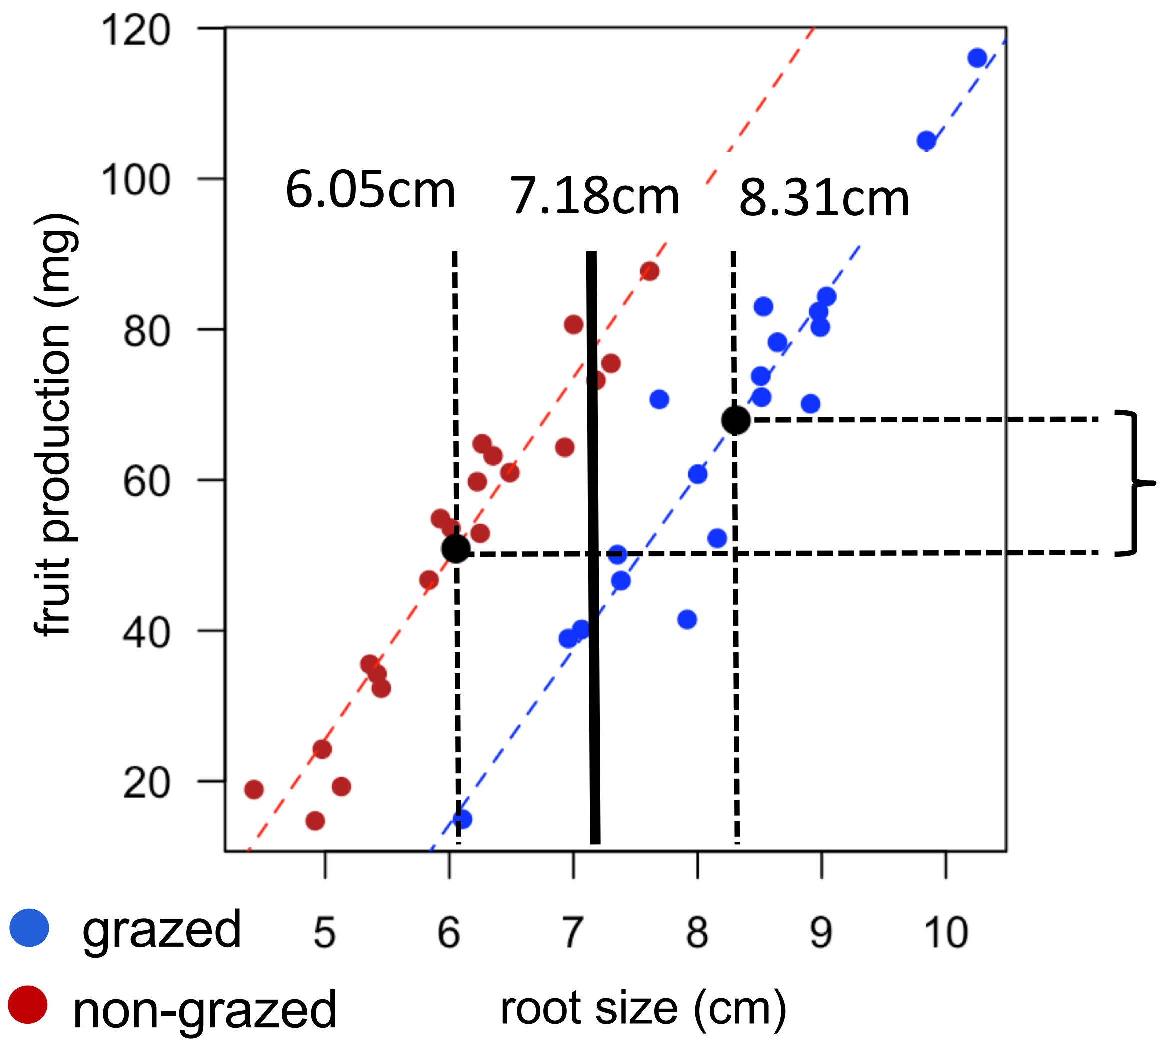

Initial conclusion:
grazed > non-grazed

Assuming a common initial mean size of root size (7.18cm), we can adjust (predict) the fruit production for each treatment (grazed / non-grazed), i.e., as if root size would had been the same in the beginning of the experiment.

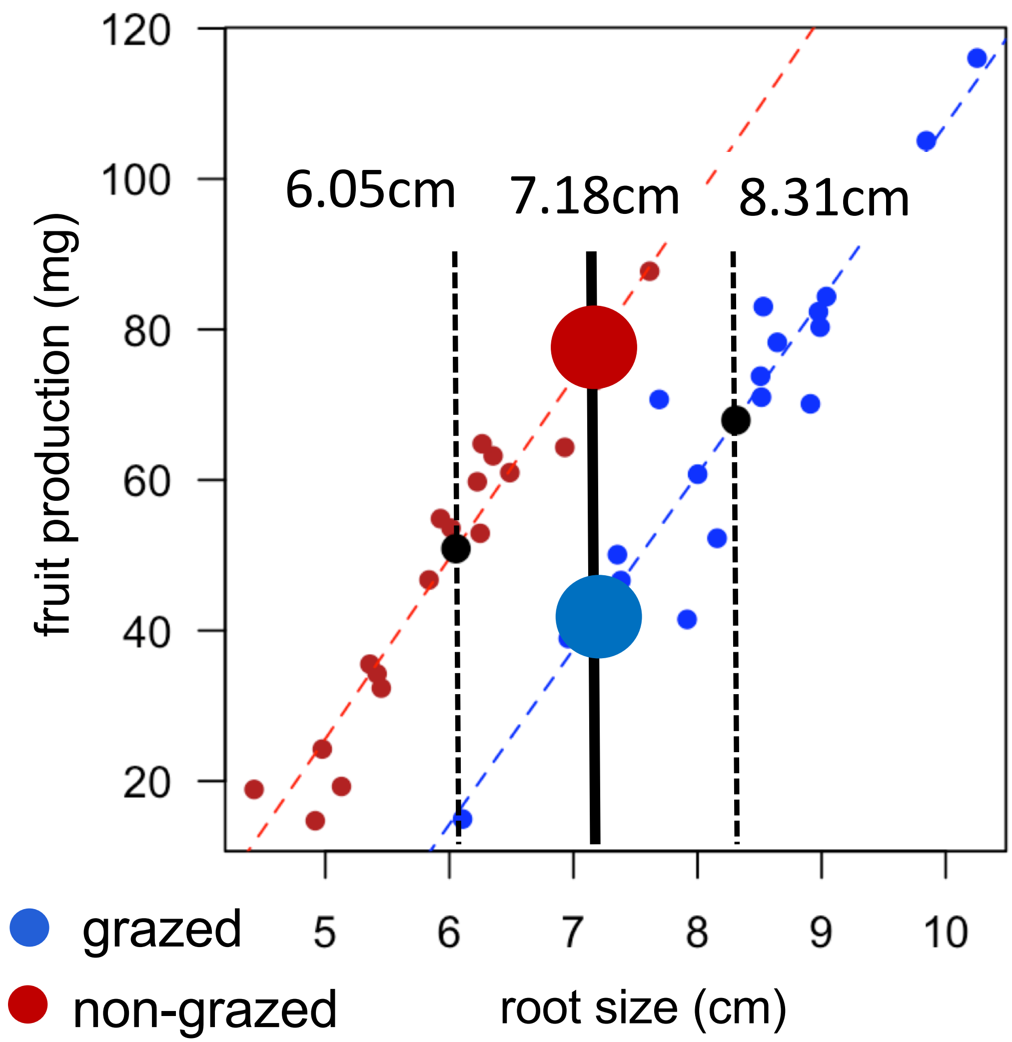

What would the effect of grazing be if all plants had started with the same initial (root) size?
Say the average value for the two treatments = $(6.05 + 8.31)/2 = 7.18$cm

Let's adjust each group's average by predicting their expected values.

Assuming a common initial mean size of root size (7.18cm), we can adjust (predict) the fruit production for each treatment (grazed / non-grazed), i.e., as if root size would had been the same in the beginning of the experiment.

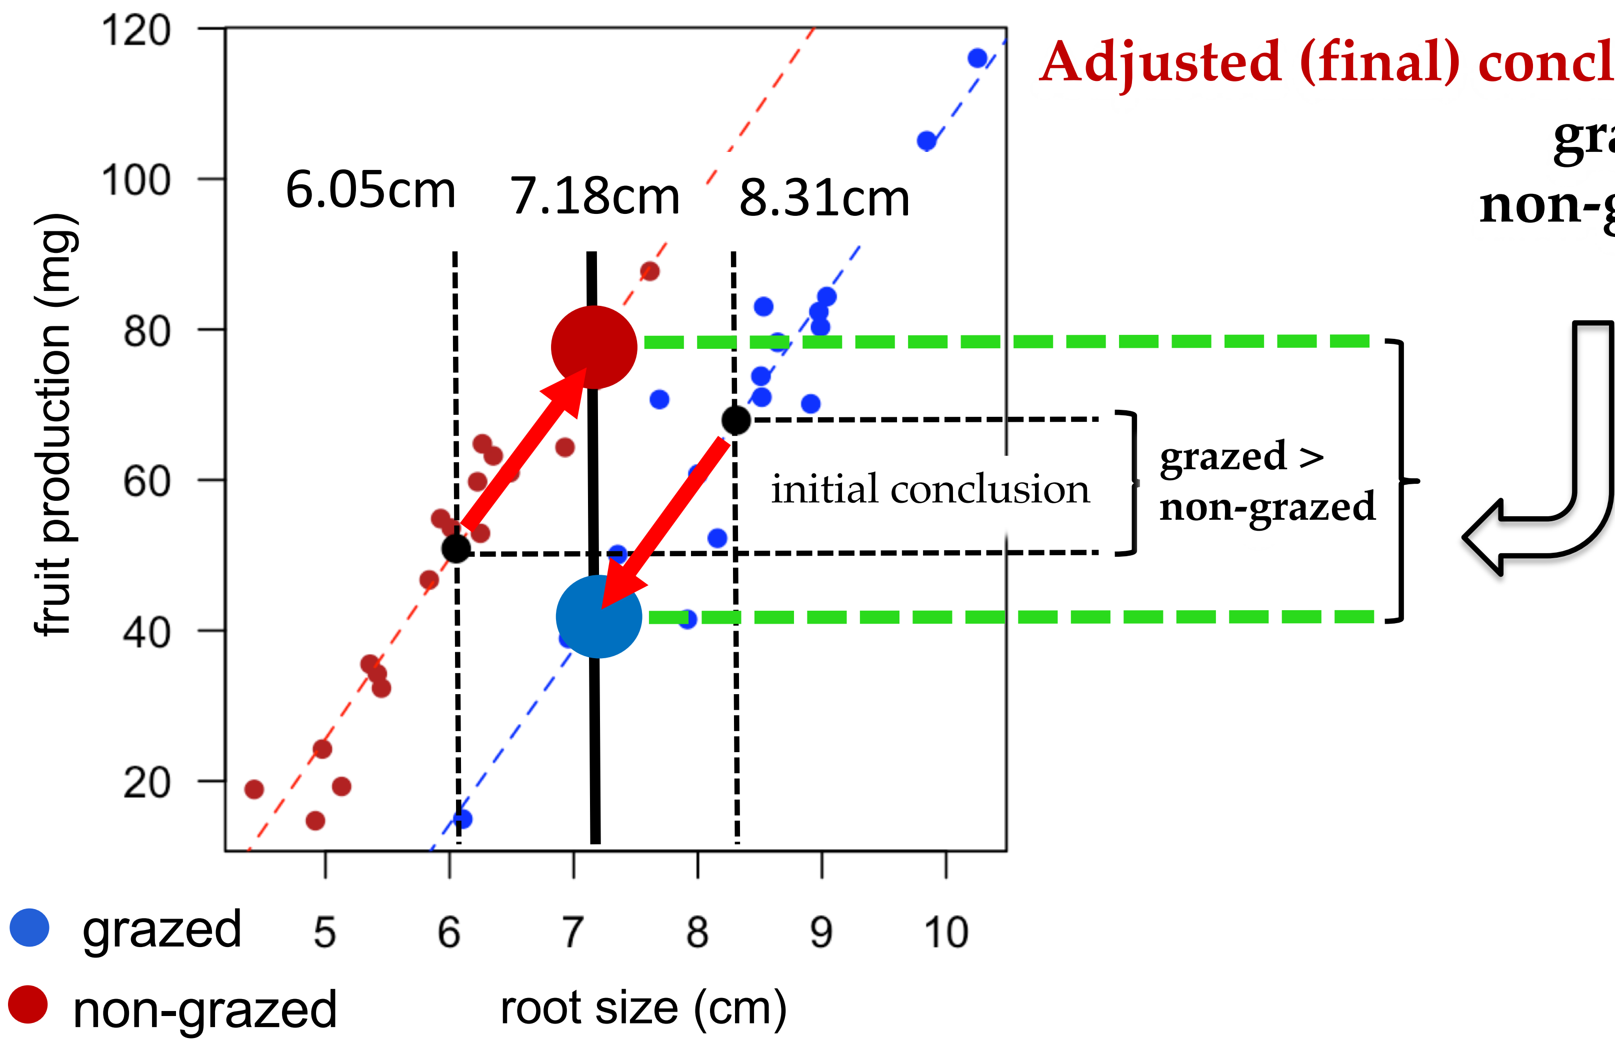

Adjusted (final) conclusion:

grazed < non-grazed

Statistics in its best - revealing “unexpected” effects

Initial conclusion: grazed > non-grazed.

Adjusted (final) conclusion: grazed < non-grazed.

THE
WOW Factor!

When can we use ANCOVA to adjust for a continuous predictor (here initial plant size)?

$Y = \mu + A_1 + X_1 + (A_1 \times X_1)$	Analysis of Covariance (ANCOVA)

- It "combines" ANOVA and regression into one analysis (we know they are the same models).
- As such, it includes at least one categorical predictor (factor, e.g., Grazing) and one continuous predictor (e.g., initial root size).
- The goal of an ANCOVA (in general) is to test for the effect of a categorical predictor while adjusting (controlling) for the effect of a continuous predictor.
- The continuous predictor is called covariate.

Analysis of Covariance (ANCOVA) - possible model outcomes
Which cases allow to control (adjust) for a covariate (root size)?

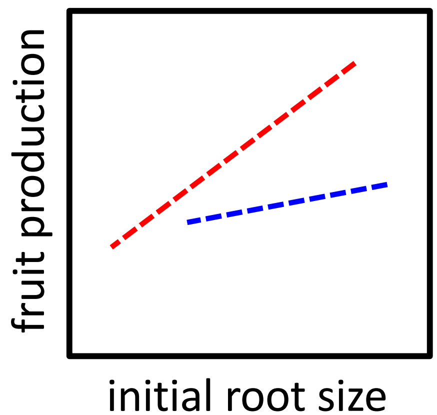

grazed

non-grazed

Analysis of Covariance (ANCOVA) - possible model outcomes
Which cases allow to control (adjust) for a covariate (root size)?

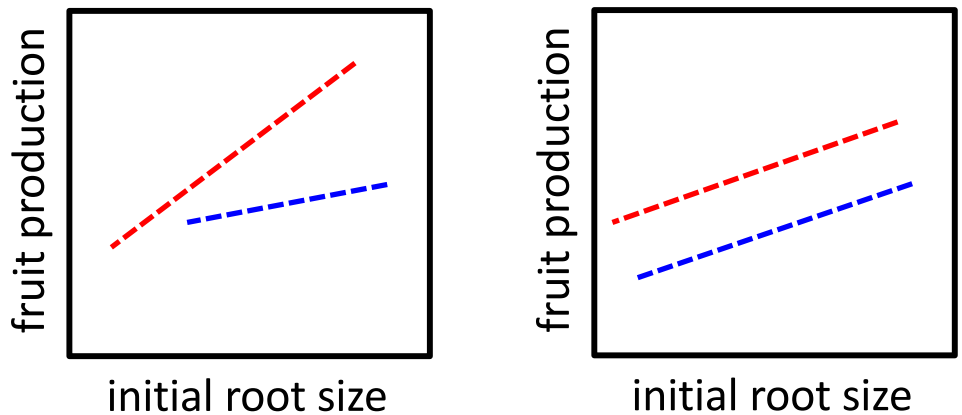

grazed

non-grazed

Analysis of Covariance (ANCOVA) - possible model outcomes
Which cases allow to control (adjust) for a covariate (root size)?

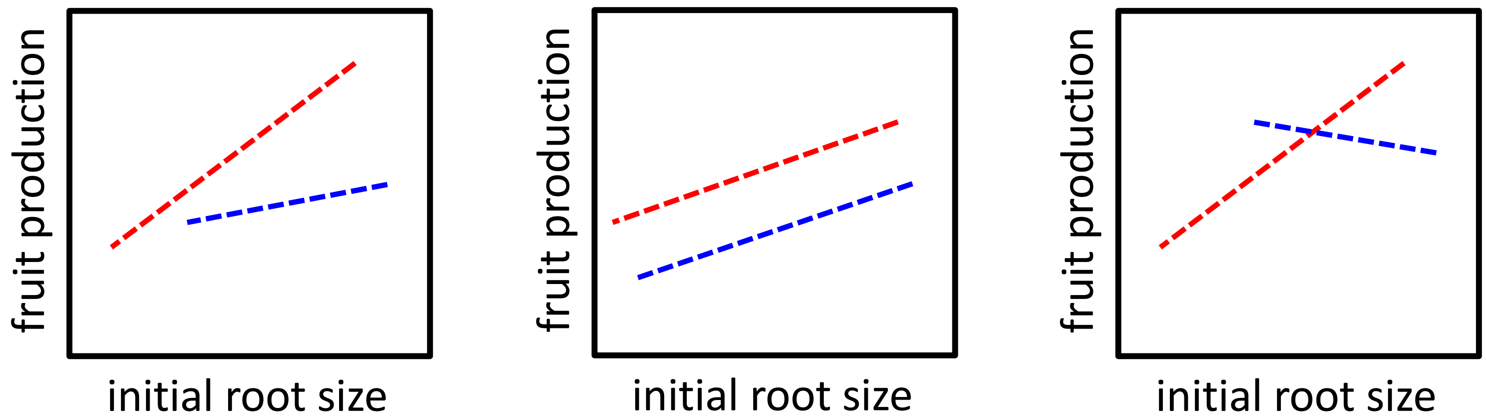

grazed
non-grazed

Analysis of Covariance (ANCOVA) - possible model outcomes
Which cases allow to control (adjust) for a covariate (root size)?

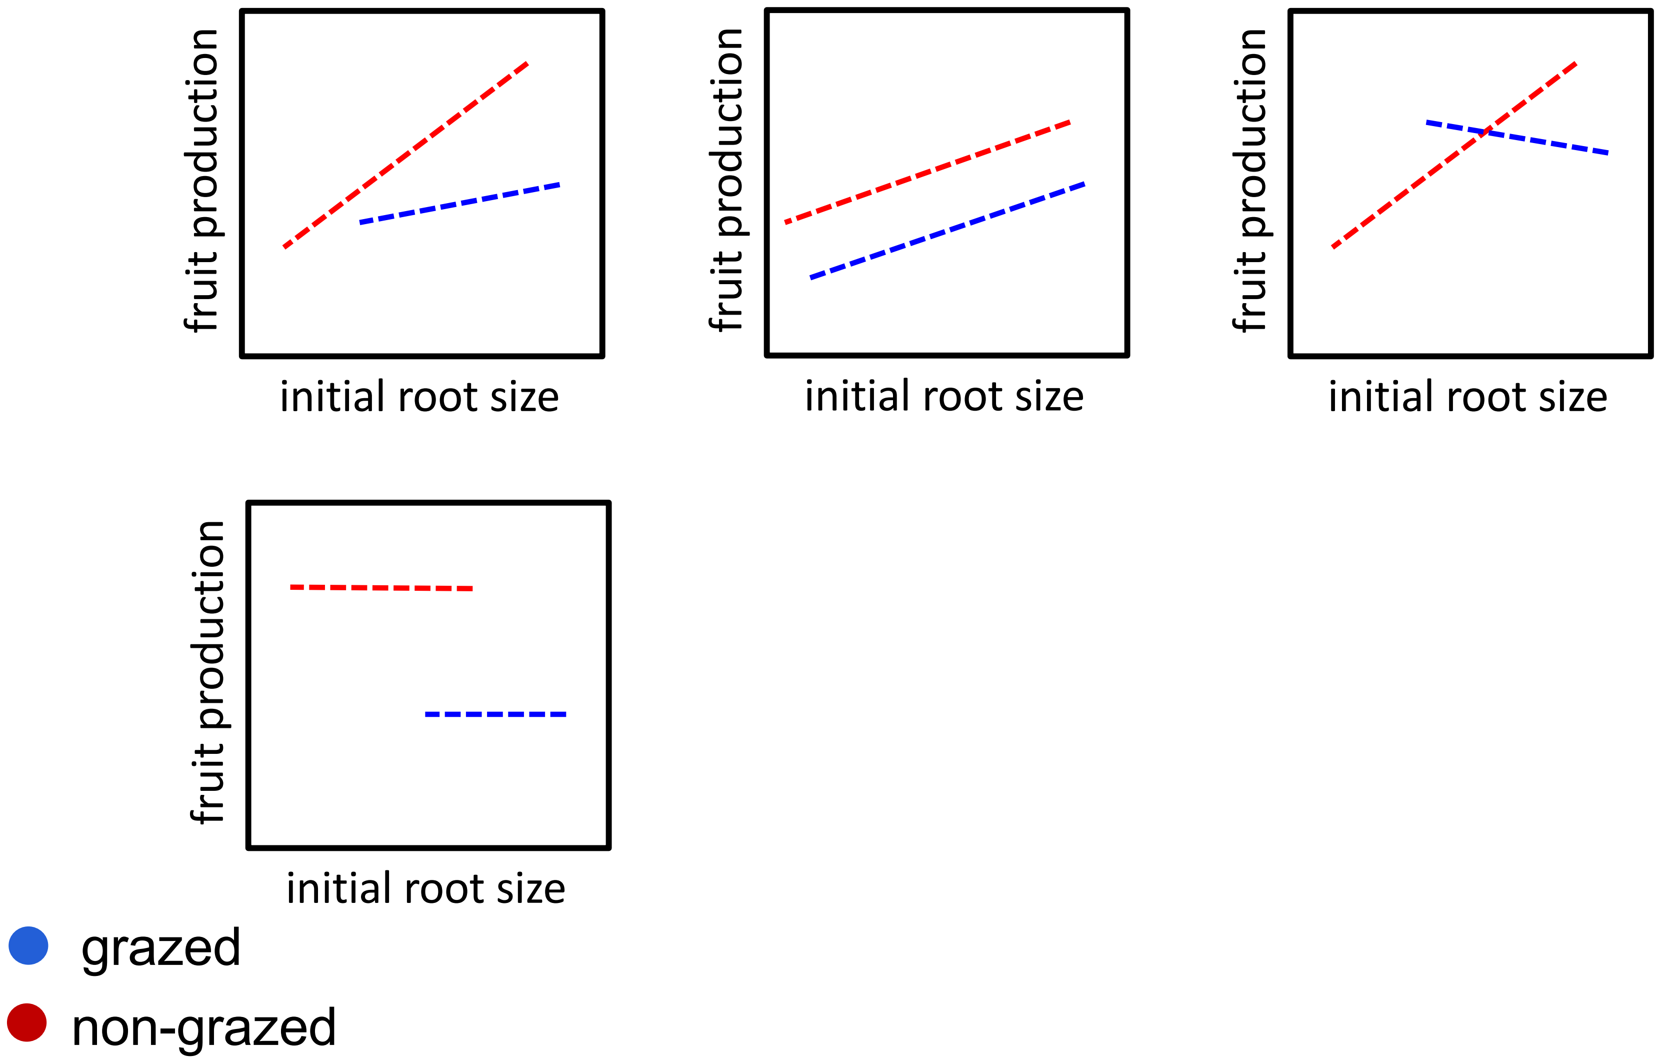

Analysis of Covariance (ANCOVA) - possible model outcomes
Which cases allow to control (adjust) for a covariate (root size)?

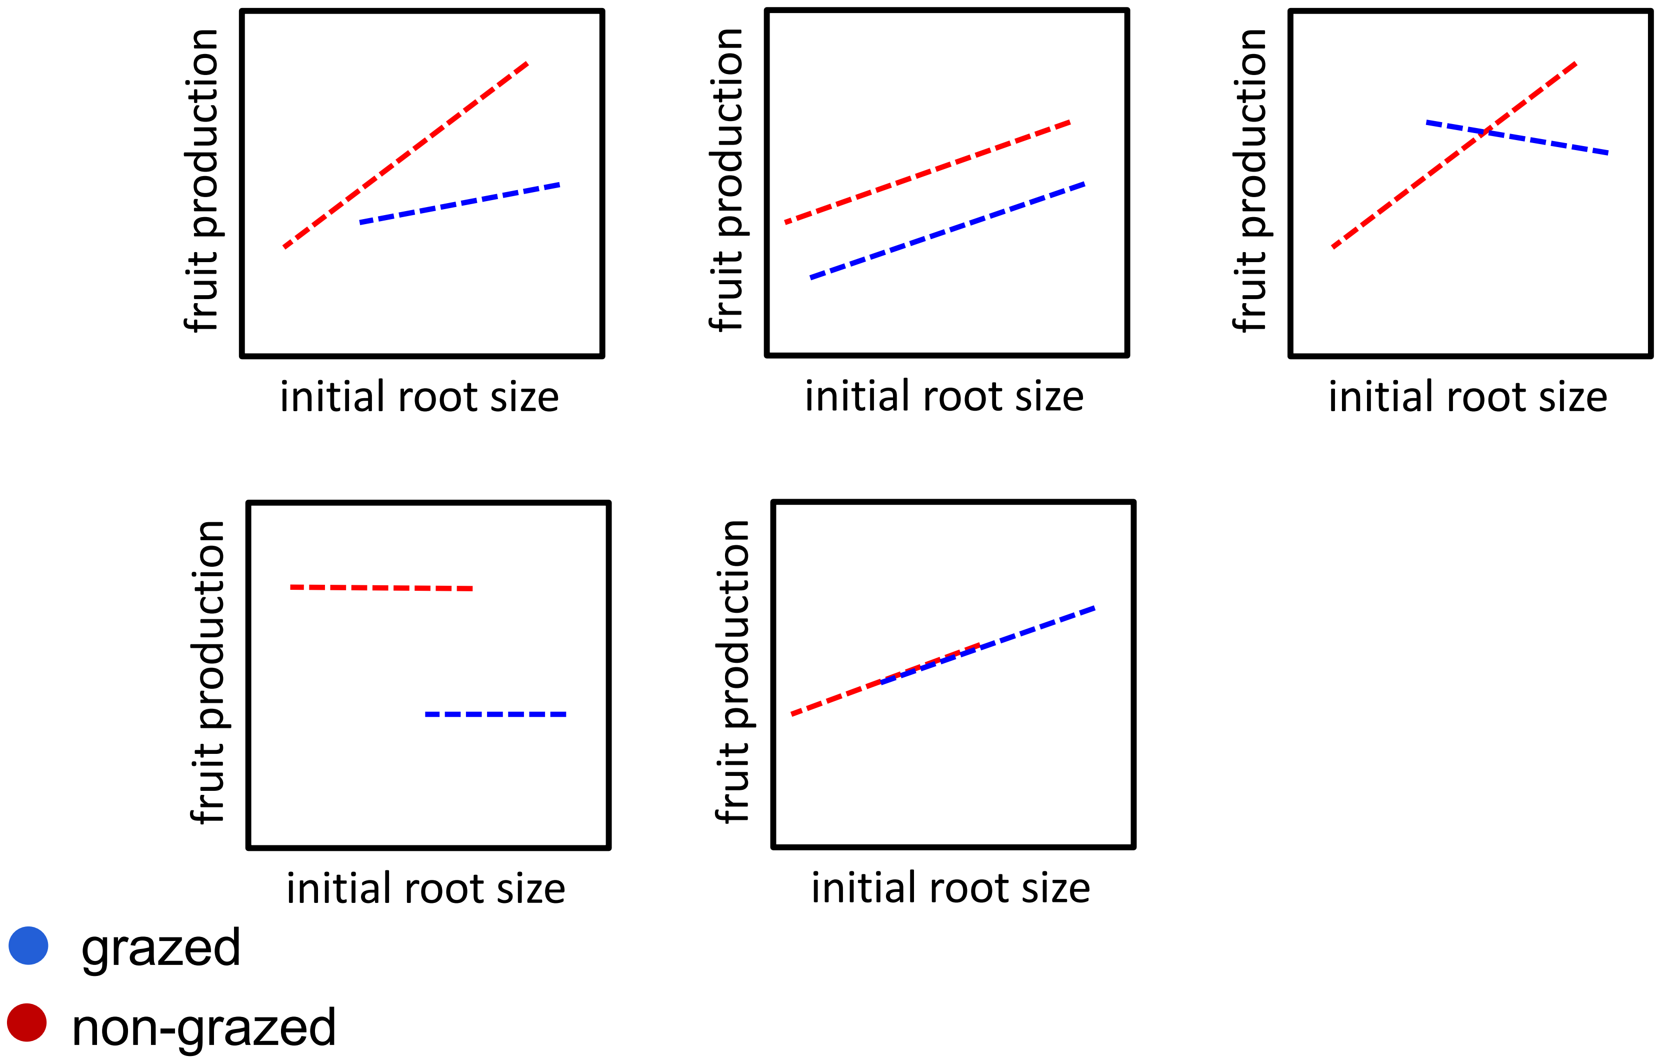

Analysis of Covariance (ANCOVA) - possible model outcomes
Which cases allow to control (adjust) for a covariate (root size)?

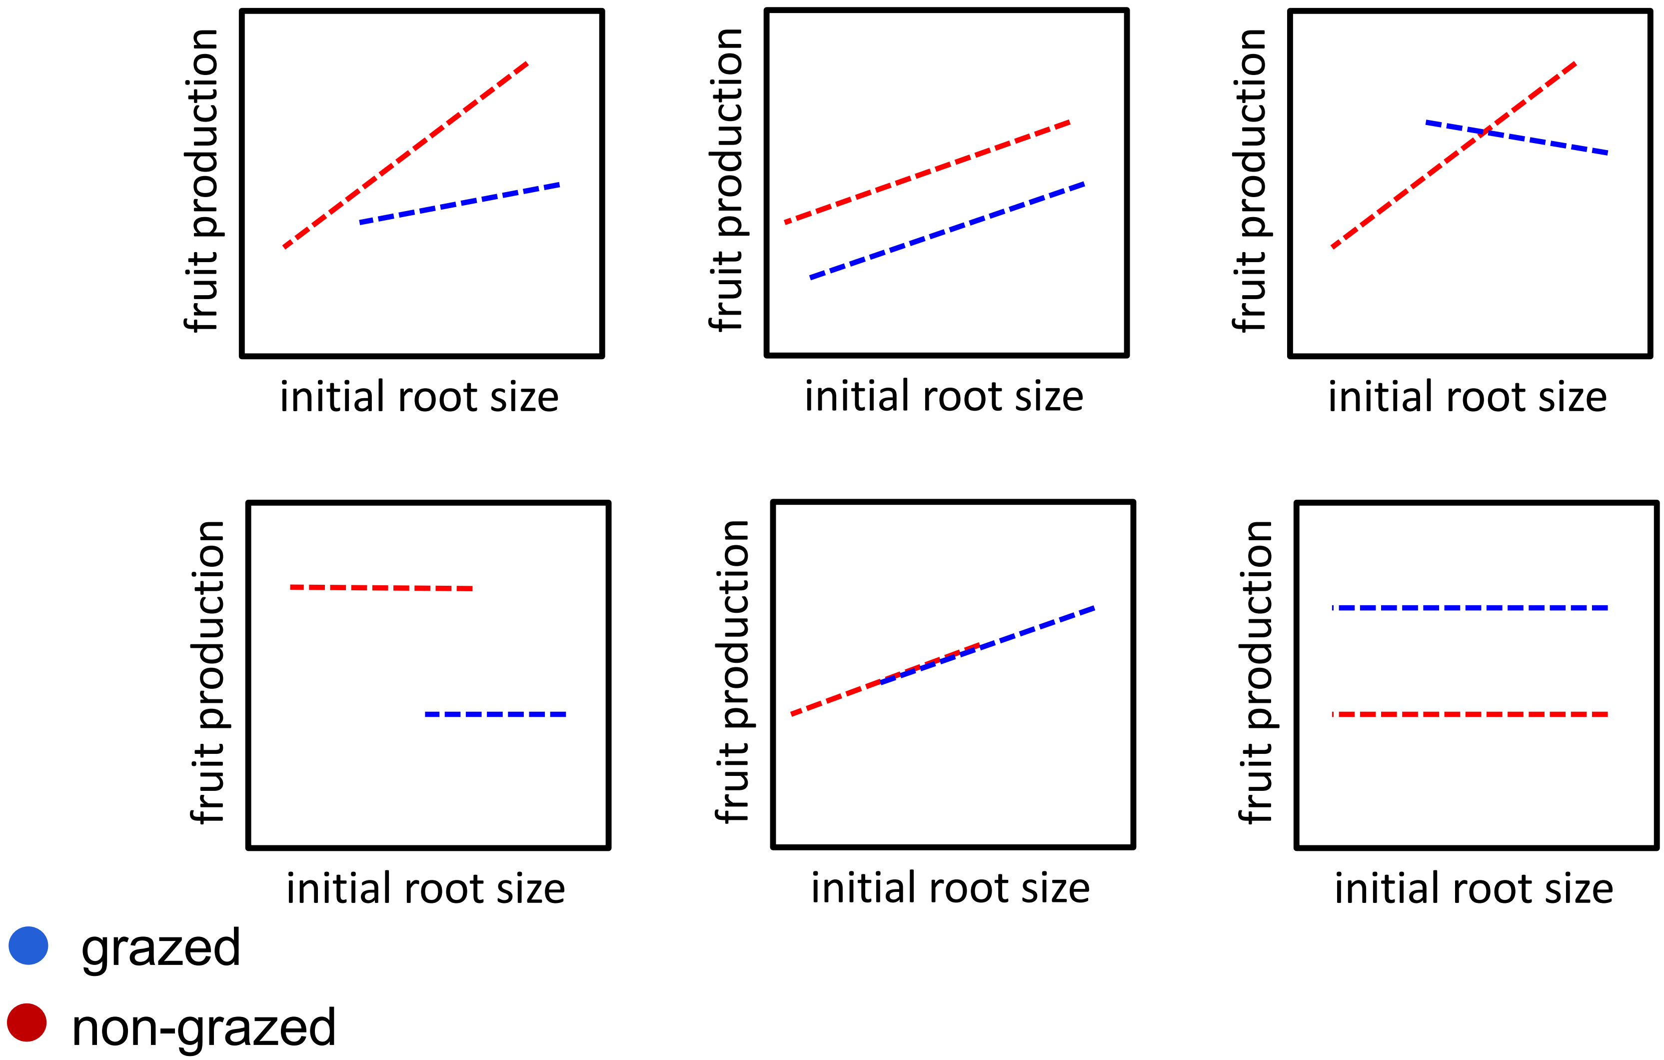

Analysis of Covariance (ANCOVA) - possible model outcomes
Which cases allow to control (adjust) for a covariate (root size)?

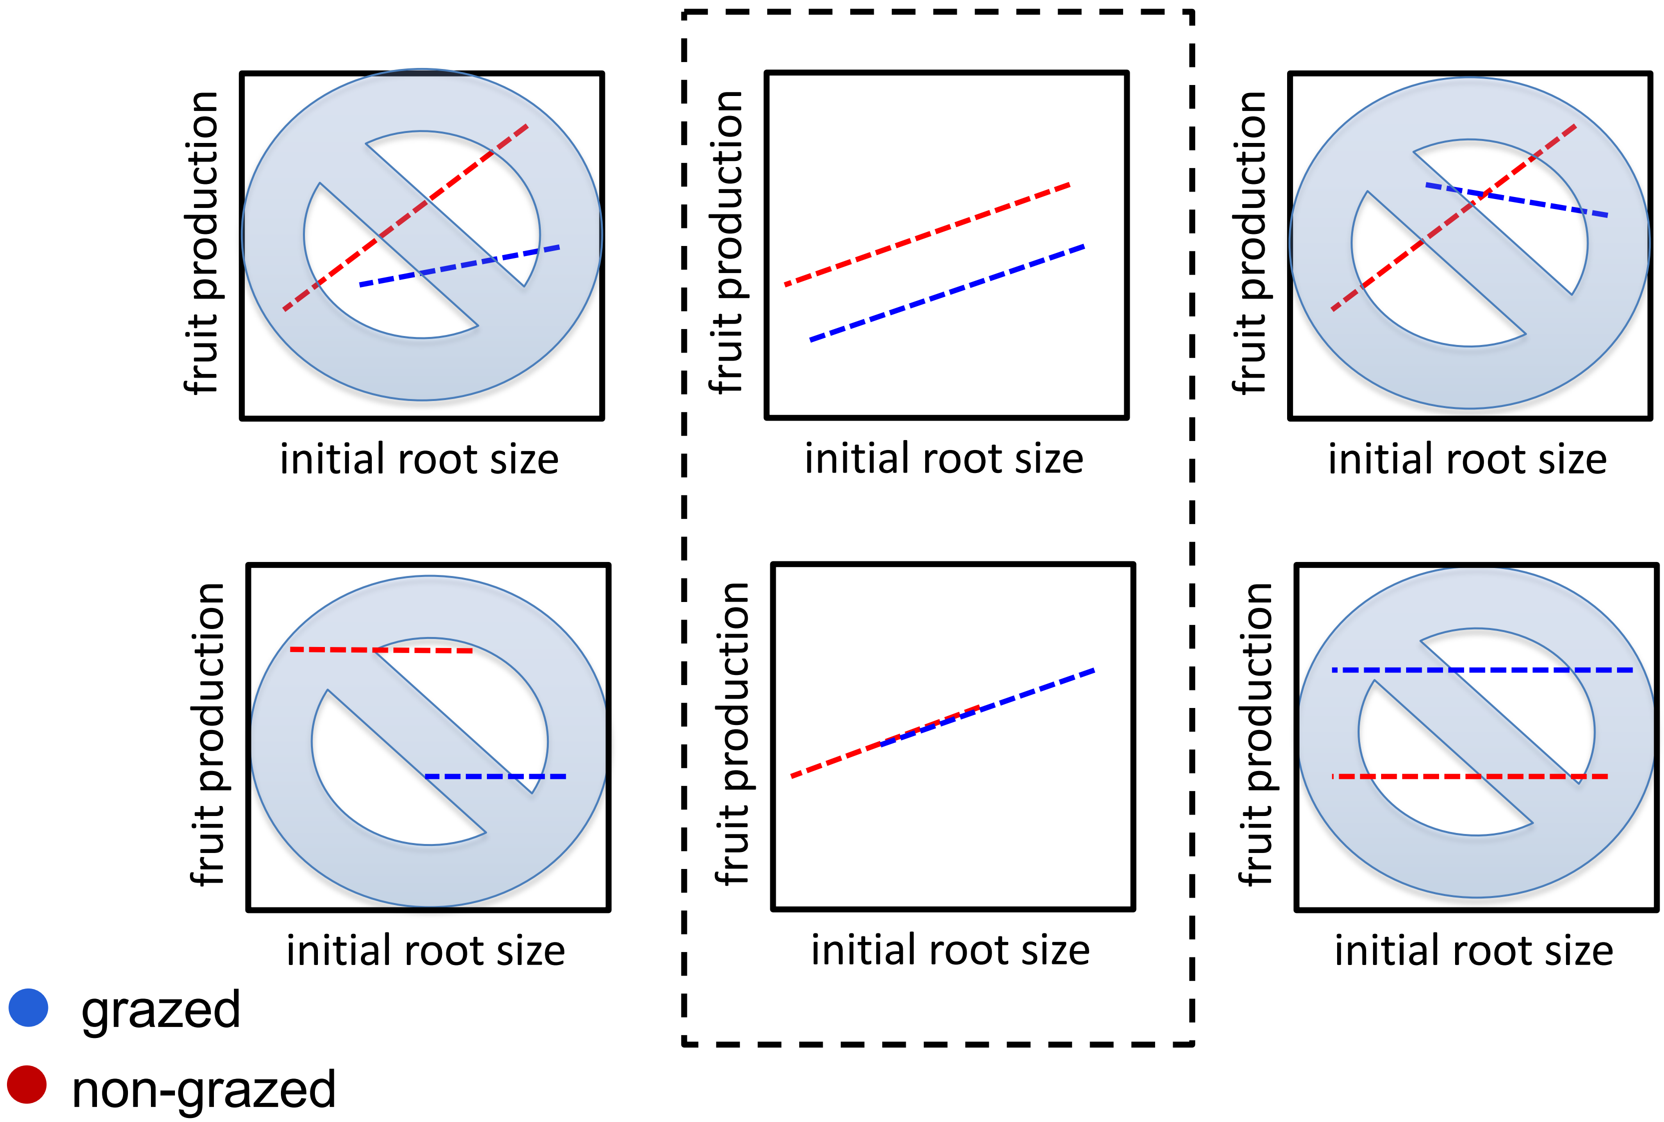

No need for adjustments as initial root size do not differ between grazed and non-grazed treatments

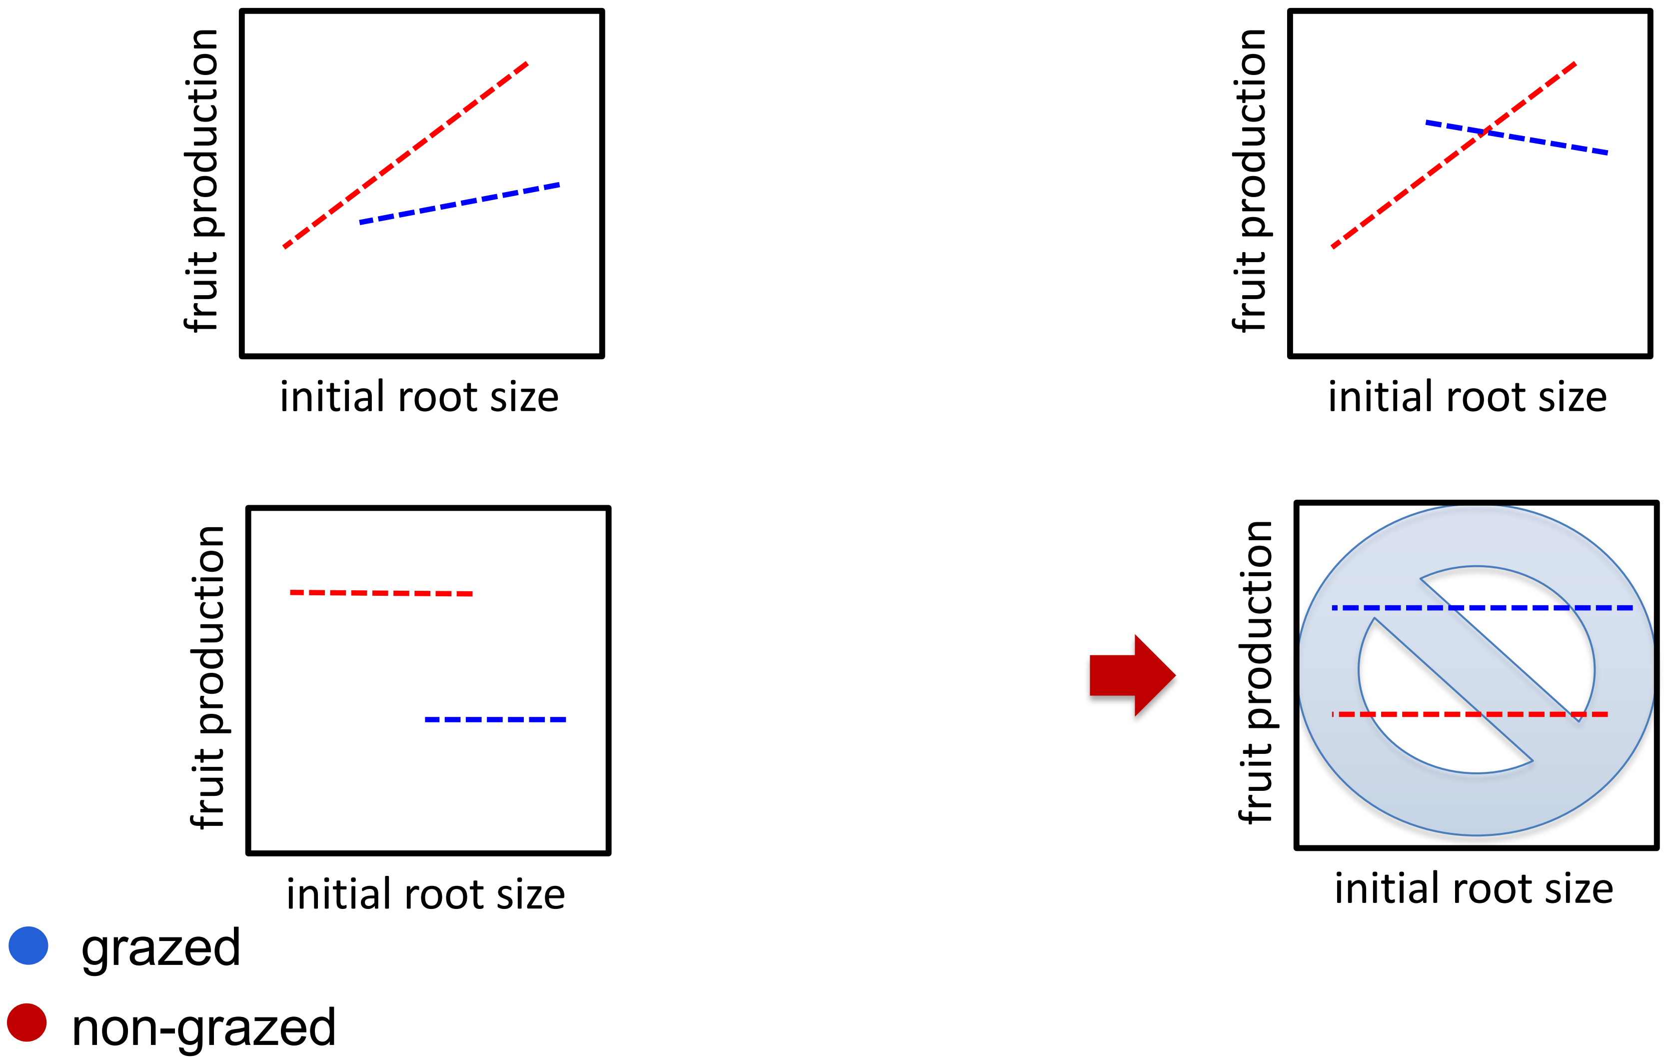

Can’t be adjusted because the fruit production does not change as a function of initial root size, so root size can’t be used to predict fruit production on the basis of a common mean

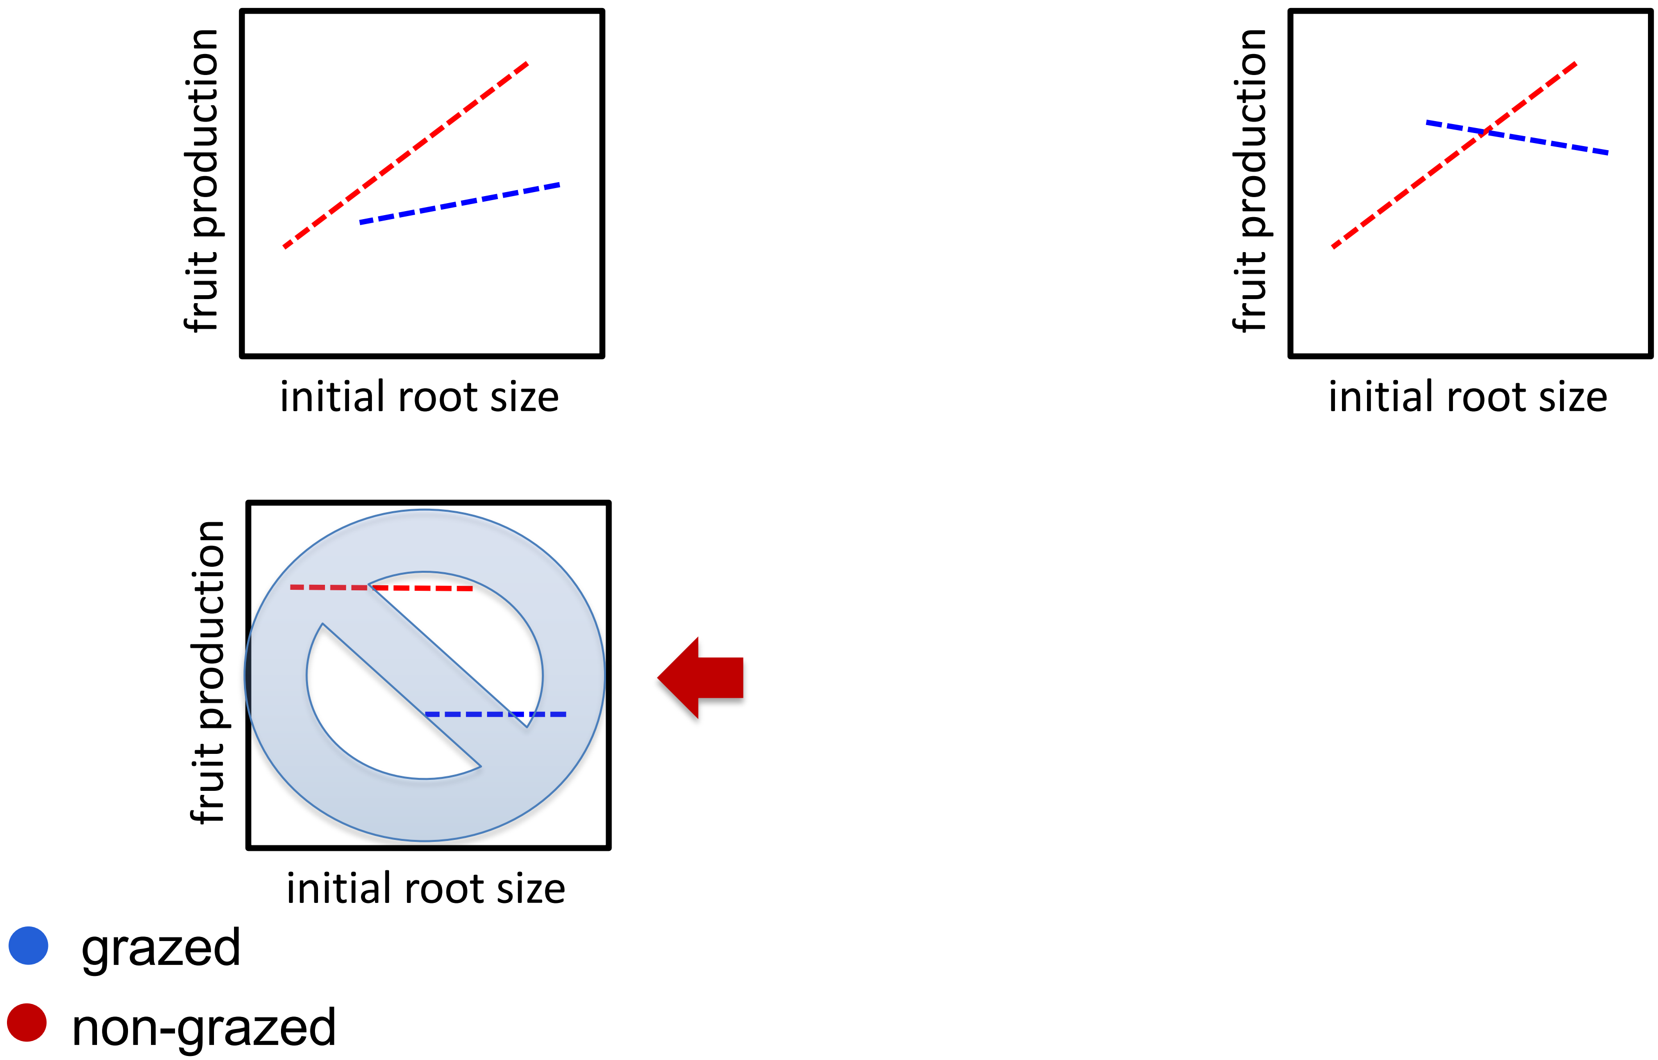

Can't be adjusted because an interaction between initial root size and grazing treatment. When there is an interaction, the differences in mean values between treatments vary as a function of the covariate, so we can't generalize to all initial root sizes.

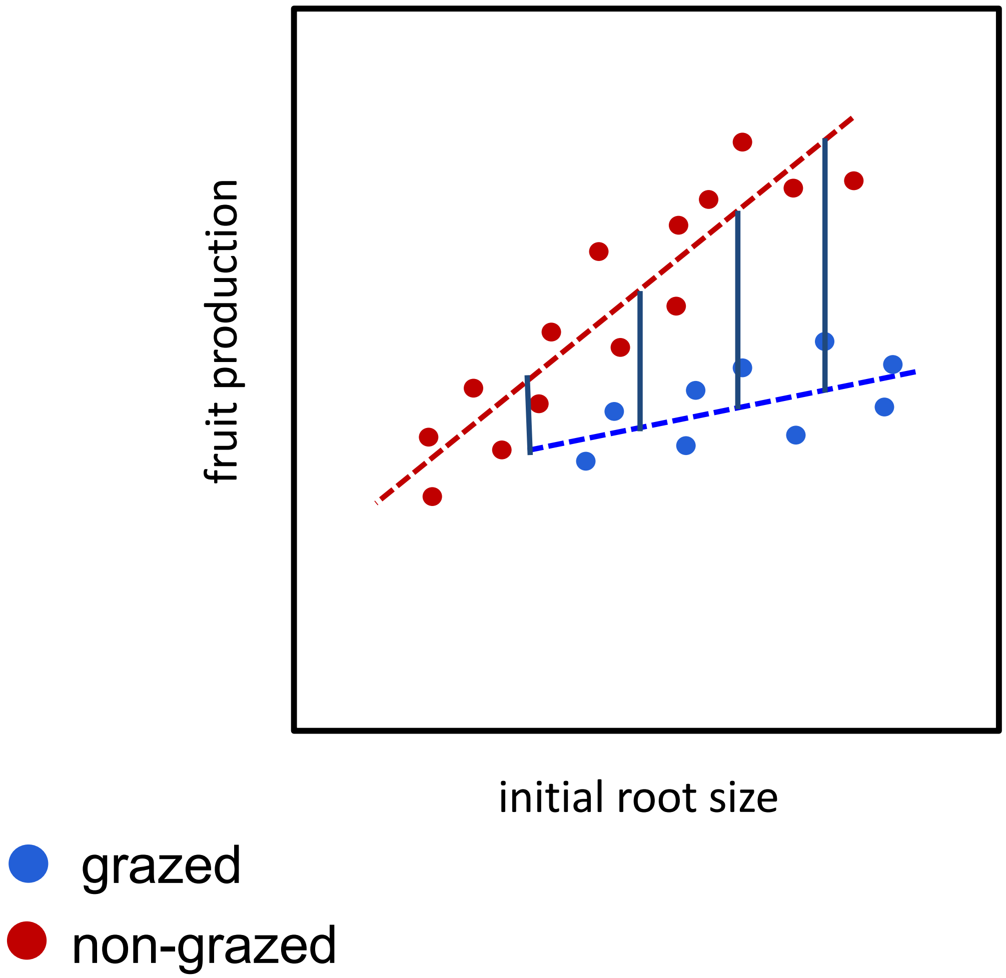

Can't be adjusted because an interaction between initial root size and grazing treatment. When there is an interaction, the differences in mean values between treatments vary as a function of the covariate, so we can't generalize to all initial root sizes.

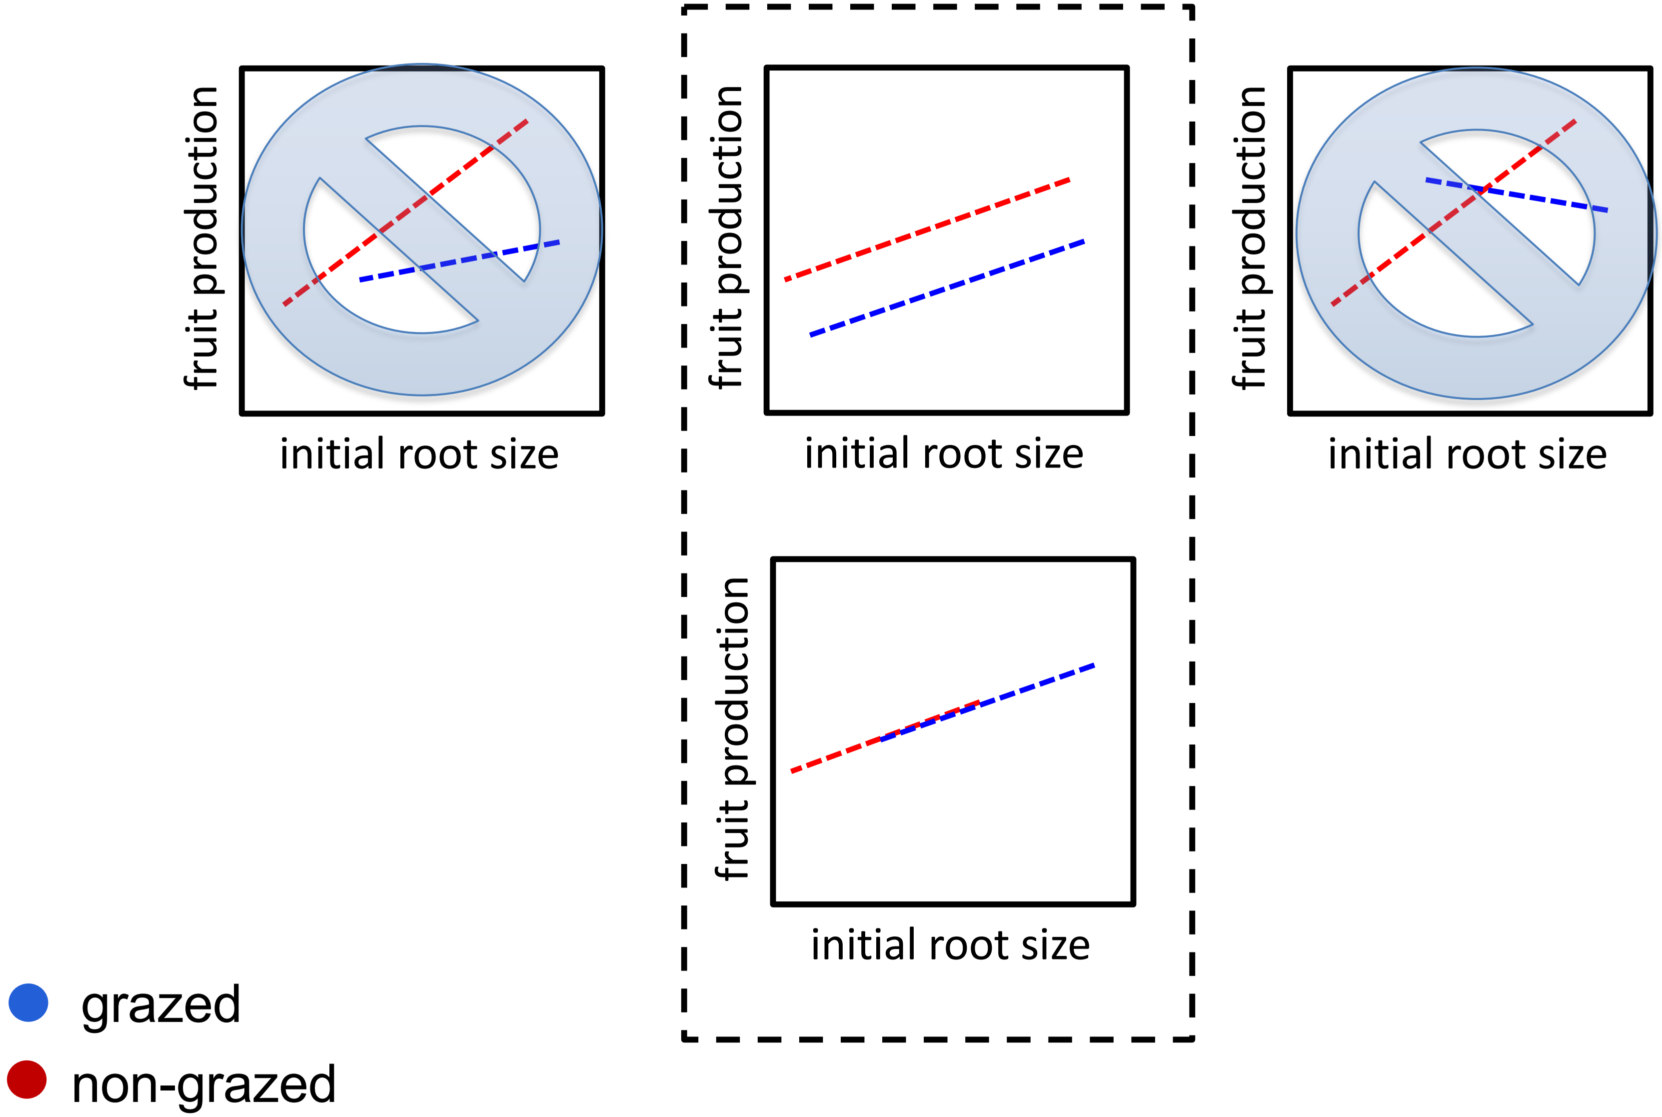

There are approaches for the more complex cases when slopes differ between groups or when the response (Y) does not depend on the covariate (initial root size) – next lecture.

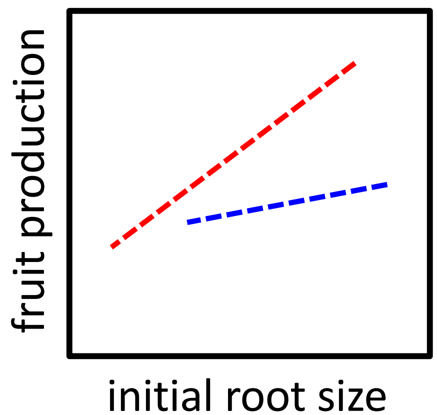

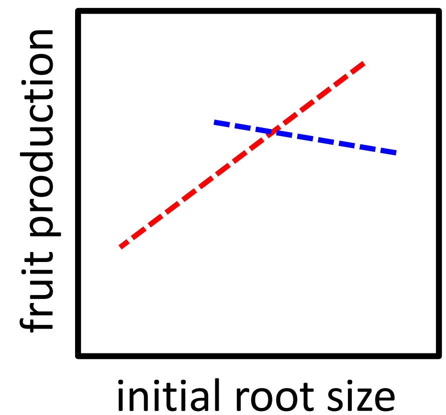

- grazed
- non-grazed

A common (equal) slope (parallel curves) between groups implies that mean differences between groups in their response (fruit production) are the same regardless of the value of the covariate (initial root size).

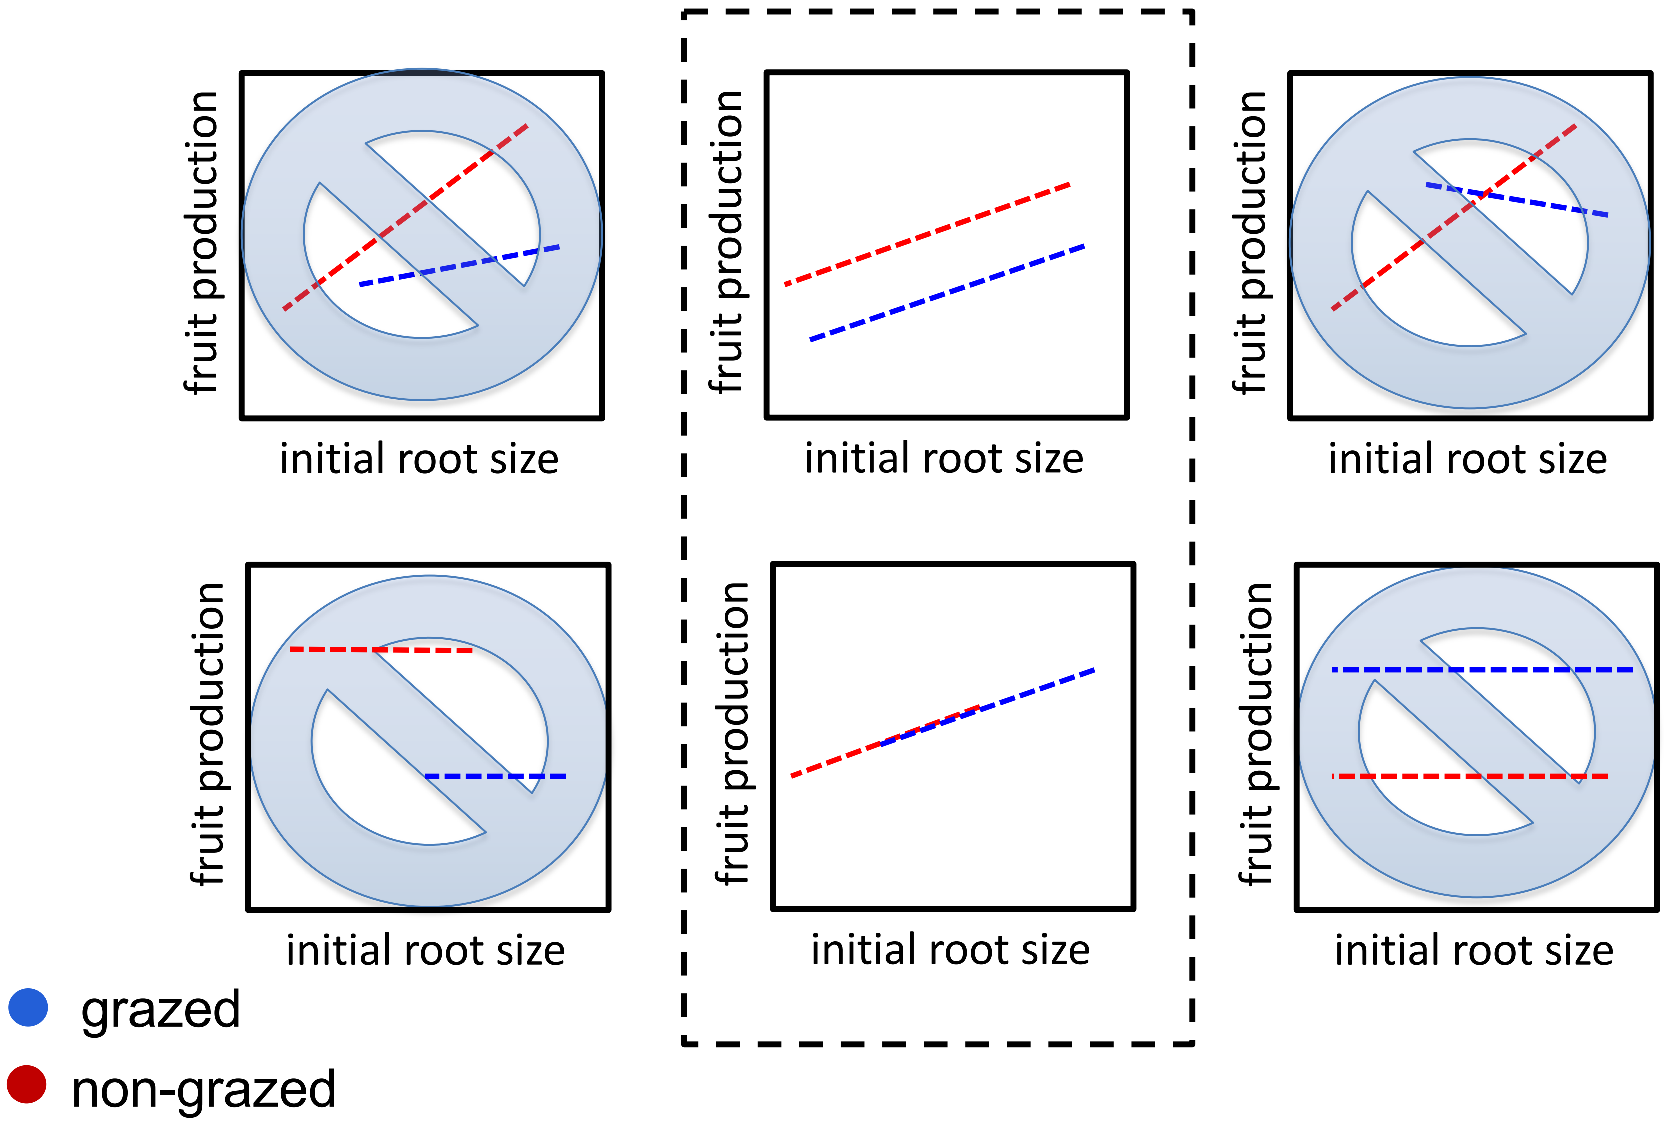

When can ANCOVA adjustments be used? Statistical assessments

First assessment –
Can the covariate predict the response?

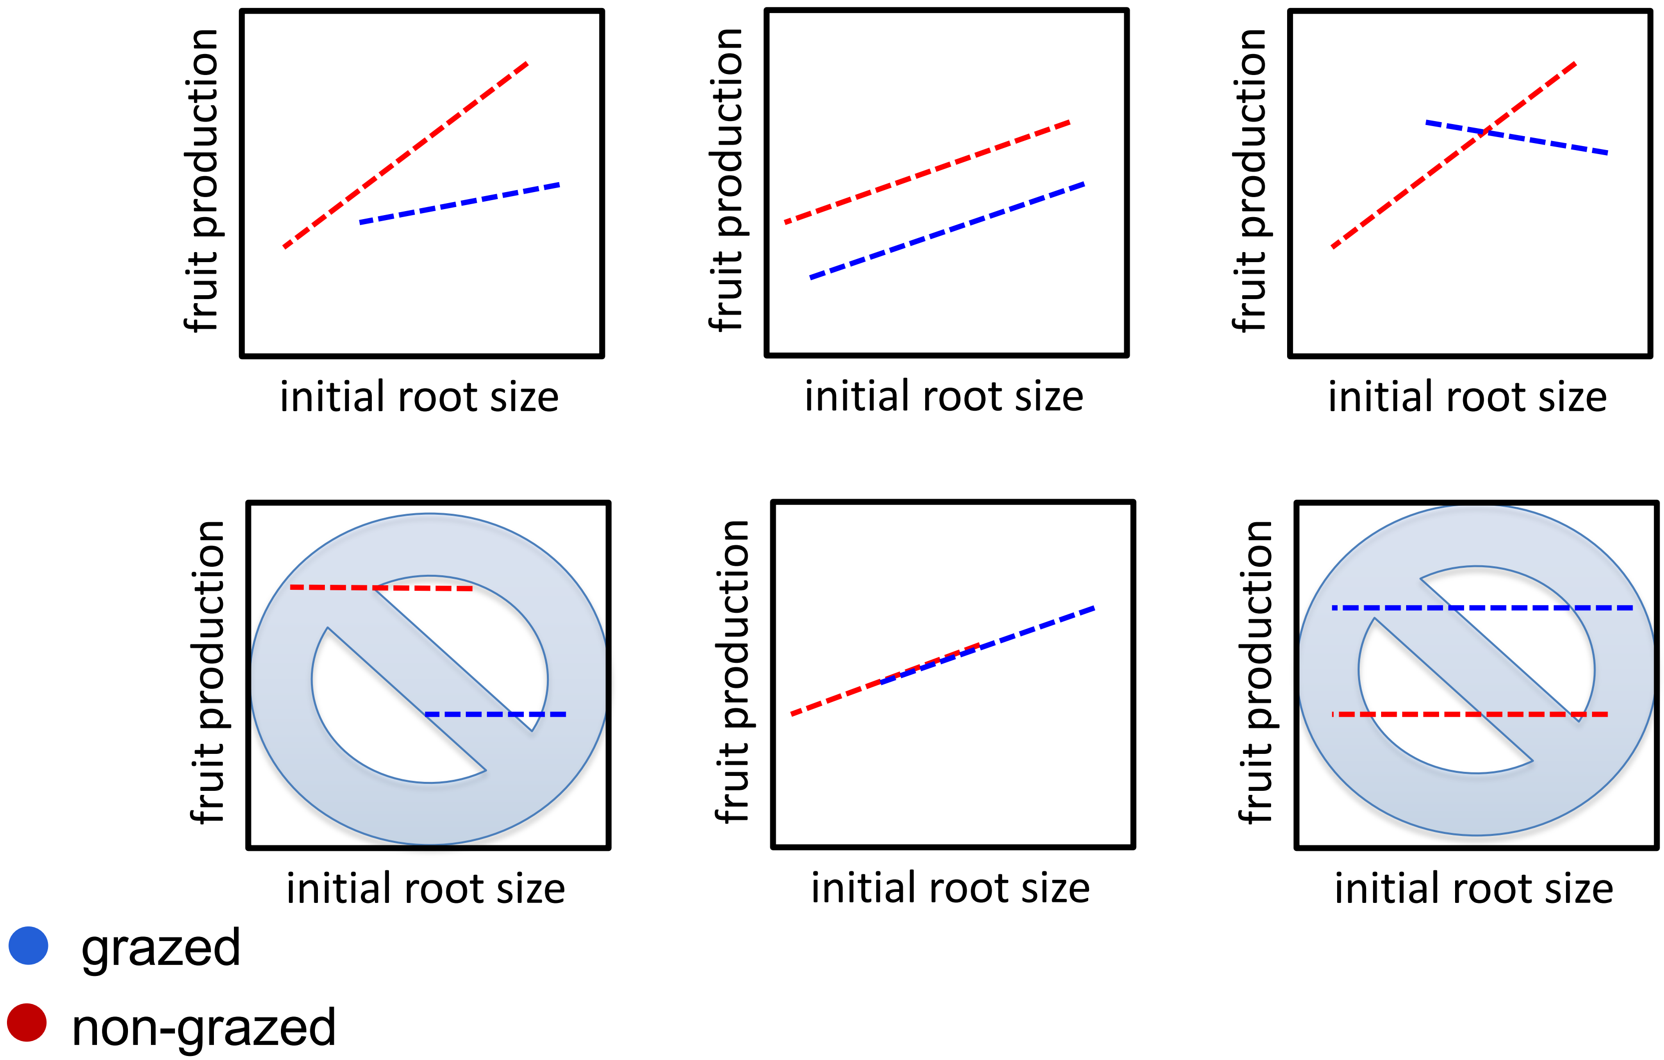

First assessment –
Can the covariate predict the response?

H_0: The slope of the regression of fruit production on initial root size is zero ($\beta = 0$).

H_A: The slope of the regression of fruit production on initial root size is not zero ($\beta \neq 0$).

```
> anova(lm(Fruit ~ Root))
Analysis of Variance Table

Response: Fruit
          Df  Sum Sq Mean Sq F value    Pr(>F)
Root       1 16795.0 16795.0  91.844 1.099e-11 ***
Residuals 38  6948.8   182.9
---
```

Second assessment – Do groups share a common slope?

Test whether the interaction between covariate (Initial root size) and factor (grazing) is significant; if significant, they don't share a common slope and initial root size can't be used for adjusting fruit production.

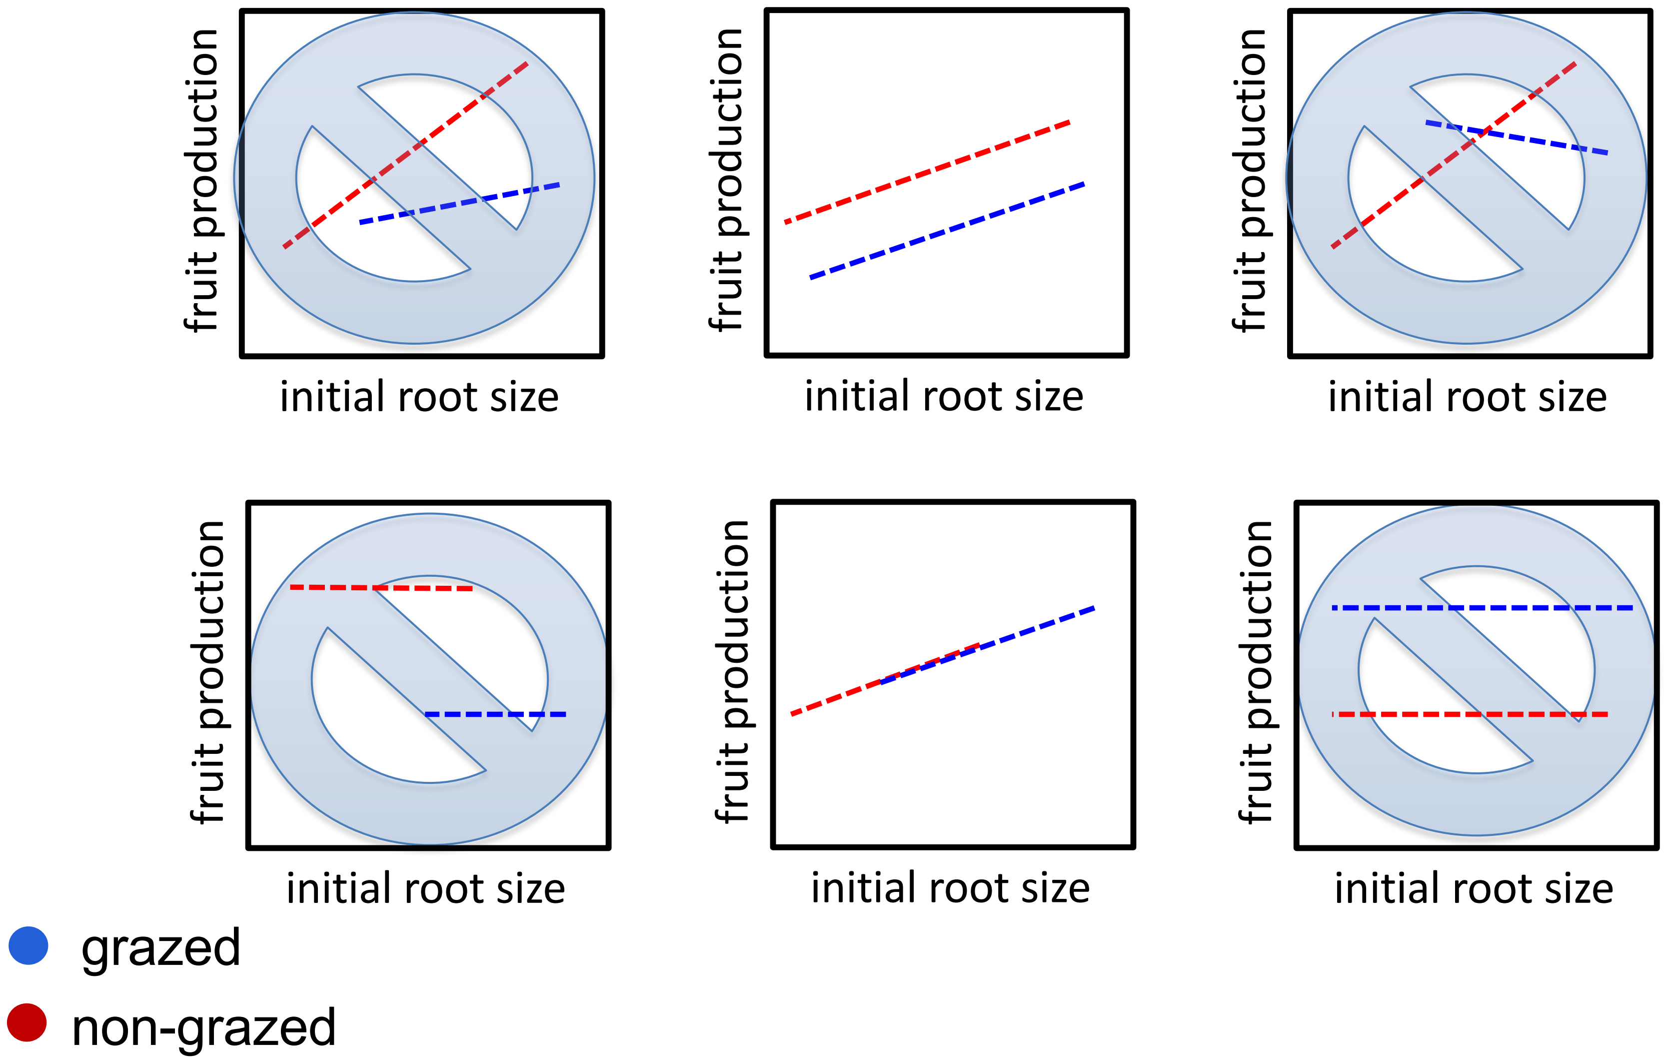

Second assessment – Do groups share a common slope?

Test whether the interaction between covariate (Initial root size) and factor (grazing) is significant.

H_0: There is no interaction between grazing treatment and initial root size (i.e., grazing/no-grazing (groups) do not differ in their slopes).

H_A: There is an interaction between grazing treatment and initial root size (i.e., grazing/no-grazing differ in their slopes).

```
> anova(lm(Fruit ~ Root*Grazing))
Analysis of Variance Table

Response: Fruit
             Df  Sum Sq Mean Sq  F value    Pr(>F)    
Root          1 16795.0 16795.0 359.9681 < 2.2e-16 ***
Grazing       1  5264.4  5264.4 112.8316 1.209e-12 ***
Root:Grazing  1     4.8     4.8   0.1031      0.75    
Residuals    36  1679.6    46.7                       
```

Second assessment – Do groups share a common slope?

Test whether the interaction between covariate (Initial root size) and factor (grazing) is significant.

H_0: There is no interaction between grazing treatment and initial root size (i.e., grazing/no-grazing (groups) do not differ in their slopes).

H_A: There is an interaction between grazing treatment and initial root size (i.e., grazing/no-grazing differ in their slopes).

Note that testing for differences in slopes (Y on X) between groups (e.g., grazed versus non-grazed), i.e., testing the interaction between the categorical (groups) and X, is interesting in itself.

In the problem analysed here we don't want to have them different but in other cases we may (e.g., allometric differences).

Second assessment – Do groups share a common slope?

Test whether the interaction between covariate (Initial root size) and factor (grazing) is significant.

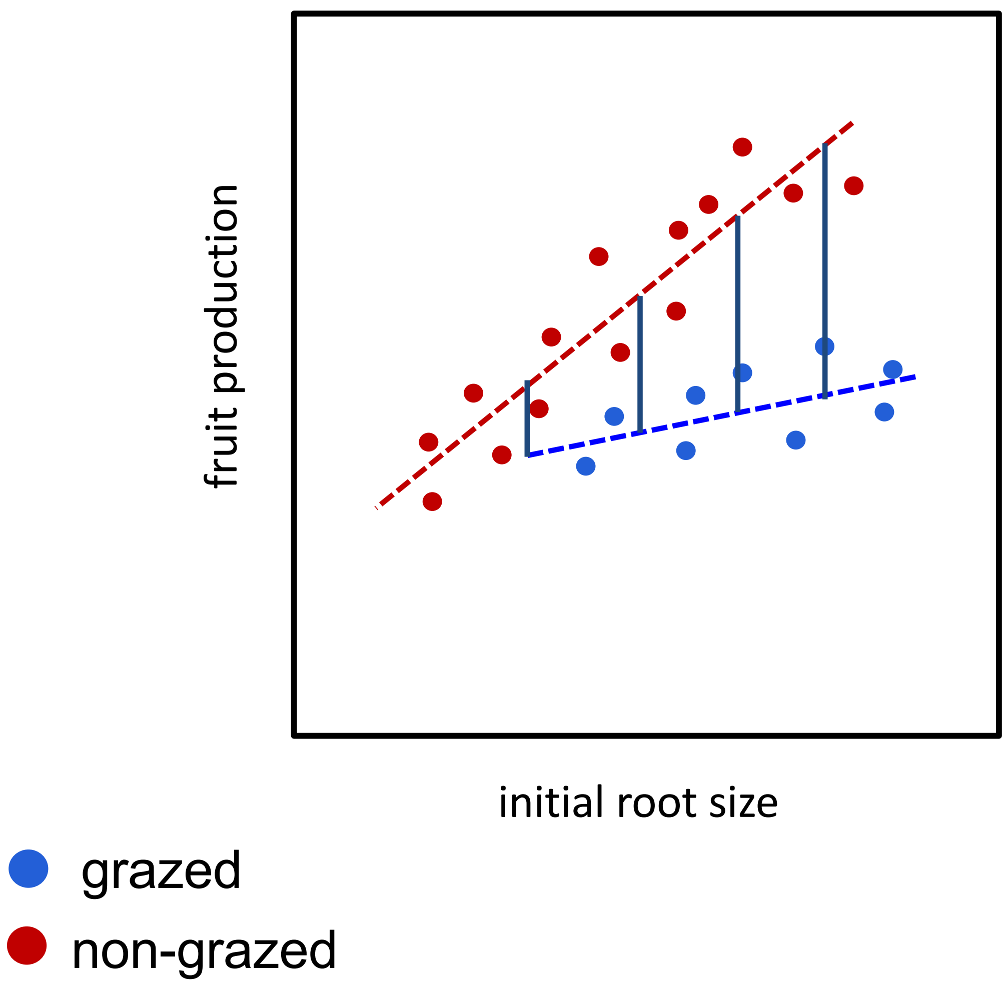

Remember: when there is an interaction, then the differences in mean values between treatments vary as a function of the covariate, so we can't generalize to all initial root sizes (more on this later).

So far, we have:

Covariate can predict the response

```
> anova(lm(Fruit ~ Root))
Analysis of Variance Table

Response: Fruit
          Df  Sum Sq Mean Sq F value    Pr(>F)
Root       1 16795.0 16795.0  91.844 1.099e-11 ***
Residuals 38  6948.8   182.9
---
```

Groups share a common slope

```
> anova(lm(Fruit ~ Root*Grazing))
Analysis of Variance Table

Response: Fruit
             Df  Sum Sq Mean Sq  F value    Pr(>F)
Root          1 16795.0 16795.0 359.9681 < 2.2e-16 ***
Grazing       1  5264.4  5264.4 112.8316 1.209e-12 ***
Root:Grazing  1     4.8     4.8   0.1031      0.75
Residuals    36  1679.6    46.7
```

Now we can test for differences in adjusted means; but before that:

Critical statistical issues underlying General Linear Models (including ANCOVAs)

Lecture 10 - a pedagogical guide (Type I and III sum-of-square)