Analysis of Covariance (ANCOVA)

- ANOVA relates a *continuous response* of interest to a set of *categorical* explanatory variables.
- Analysis of Covariance (ANCOVA) extends the ANOVA framework to allow to account for *continuous* explanatory variables as well.
- This is *NOT* the same thing as regression. In particular, ANCOVA does *not* allow you to estimate the *effect* of a continuous explanatory variable on a continuous response; it only *removes* the variation explained by the continuous explanatory variable, thus:
  - reducing residual error.
  - allowing better estimates of the categorical marginal and interaction effects of interest.
- In an ANCOVA, the continuous explanatory variable is *never* of interest. It is merely a *nuisance* variable to be eliminated.

Analysis of Covariance (ANCOVA) rationale

- Let Y_i be the response of interest for sample unit i. Let X_i be the covariate (continuous explanatory variable) for sample unit i
- First, find the "best fitting" line through the points (X_i, Y_i):

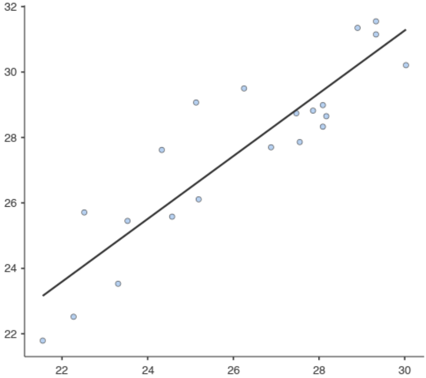

Analysis of Covariance (ANCOVA) rationale

- There are many ways to define "best fitting," but here we take the classical definition; i.e. the *ordinary least squares* (OLS) fitted line obtained by *minimizing the sum of the squared errors*.
- That is, if we write

$$Y_i = \beta_0 + \beta_1 X_i + \varepsilon_i,$$

  for some random error $\varepsilon_i \sim N(0, \sigma^2)$, we can find numbers $\widehat{\beta}_0$ for β_0 and $\widehat{\beta}_1$ for β_1 that minimize

$$\sum_{i=1}^{n} \varepsilon_i^2 = \sum_{i=1}^{n} (Y_i - \beta_0 - \beta_1 X_i)^2$$

- This is a simple calculus exercise and yields the OLS estimators:

$$\widehat{\beta}_0 = \overline{Y} - \widehat{\beta}_1 \overline{X}, \qquad \widehat{\beta}_1 = \frac{S_{XY}}{S_X^2}$$

Analysis of Covariance (ANCOVA) rationale

- Now, with the "best fitting" (OLS regression) line estimated, we can plug in the OLS estimators and rearrange the equation:

$$\begin{aligned} Y_i &= \widehat{\beta}_0 + \widehat{\beta}_1 X_i + \varepsilon_i \\ &= \overline{Y} - \widehat{\beta}_1 \overline{X} + \widehat{\beta}_1 X_i + \varepsilon_i \\ &= \overline{Y} + \widehat{\beta}_1 (X_i - \overline{X}) + \varepsilon_i \end{aligned}$$

  Thus,

$$Y_i - \widehat{\beta}_1 (X_i - \overline{X}) = \overline{Y} + \varepsilon_i$$

- Denote the lefthand side of this equation by

$$Y_i^{adj} := Y_i - \widehat{\beta}_1 (X_i - \overline{X})$$

  This is our response of interest, Y, adjusted for the effect of the covariate X.

Analysis of Covariance (ANCOVA) rationale

- So, we now have a *transformed* version of Y that we can fit ANOVA models to. For example, if W is some categorical explanatory factor of interest for Y, we can now estimate the ANOVA model:

$$Y^{adj} = \mu + \tau_W + \delta$$

- This will give us information about the *effect of W on Y adjusted for the effect of X*.
- The classic (and most common) application: estimating the effect of some intervention Y *adjusting for baseline X* over groups of W.
- Note: we can adjust for *multiple covariates* by using the same "best fit" adjustment procedure for each covariate.

RM-ANOVA vs. ANCOVA

- Suppose we have a pre-test and post-test measurement on 21 people subjected to one of three experimental treatments (a *nested* design).
- Performing a RM-ANOVA, we could address the question of whether or not the average change in pre and post-test measurement differs among the three experimental groups.
- Or, treating the pre-test measurement as a nuisance variable, we can perform an ANCOVA to address the question of whether or not the average post-test measurement, adjusted for baseline differences in pre-test measurements, differs among the three experimental groups.
- RM-ANOVA quantifies *change from pre-test to post-test* between groups.
- ANCOVA quantifies *differences of post-test means* between groups, adjusted for baseline; this is functionally analogous to a post hoc t-test on Group$\times$Time interaction in an RM-ANOVA.

Assumptions of ANCOVA

- The usual ANOVA assumptions (independence, homoskedasticity, normality of residuals)
- Relationship between response and covariate is *linear*. **(check plausibility with scatterplots)**
- All regression slopes between the covariate and the response are *equal* across each level of the explanatory factor(s). **(check plausibility with group-wise scatterplots and/or improper ANCOVA)**
- In an RM-ANCOVA framework, the regression slopes are also *equal* over each repeated measurement (virtually *never* satisfied in practice: **avoid this analysis at all costs**).
- Independence of the covariate and the other explanatory factors (often suspect).
- In practice, usually better to perform a *regression analysis* instead of an ANCOVA. ✓✓

ANCOVA Example (covariate adjusting for baseline)

- Suppose we have a pre-test and post-test measurement on 21 people subjected to one of three experimental treatments (a *nested* design).
- We check if the pre-test baseline is linearly related to the post-test measurement:

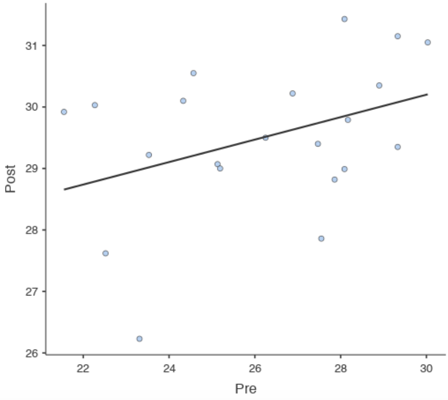

ANCOVA Example

- There's somewhat of a linear relationship between our response of interest (post-test measurement) and nuisance covariate (pre-test measurement), so an ANCOVA approach may be reasonable.
- We estimate the full (improper) ANCOVA model:

$$Y_{post} = \mu + \tau_{groups} + \beta \cdot Y_{pre} + \alpha \cdot \tau_{groups} \cdot Y_{Pre} + \delta$$

- Note: one of the assumptions of the ANCOVA rationale is that $\alpha = 0$. That is, all "best fitting" slopes between the covariate and the response are *equal* across all levels of the explanatory factor(s).
- By specifying the above full model, we can explicitly *test* this assumption.

ANCOVA Example

	Sum of Squares	df	Mean Square	F	p
Pre	3.168	1	3.168	2.017	0.176
Groups	0.846	2	0.423	0.269	0.768
Groups ∗ Pre	0.935	2	0.467	0.298	0.747
Residuals	23.558	15	1.571		

- Notice: no significant effect of 'Groups $\times$ Pre-test'; so no evidence against ANCOVA assumption of equal regression slopes ($\alpha = 0$).
- Not much variation explained by baseline differences ('Pre' sum of squares).
- No evidence of a group effect on the post-test measurements (this is our main effect of interest in the ANCOVA).

ANCOVA Example

- Estimates of the average post-treatment measurement between experimental groups, adjusted for baseline: Group I average = 29.493 Group II average = 29.803, Group III average = 29.165.

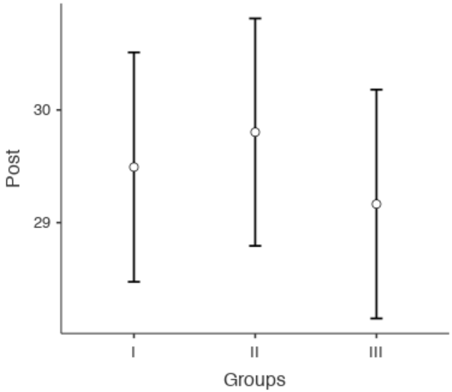

RM-ANOVA vs. ANCOVA

Within Subjects Effects

	Sum of Squares	df	Mean Square	F	p
Pre.vs.Post	114.345	1	114.345	36.074	<.001
Pre.vs.Post ∗ Groups	0.493	2	0.247	0.078	0.925
Residual	57.055	18	3.170		

Note. Type 3 Sums of Squares

Between Subjects Effects

	Sum of Squares	df	Mean Square	F	p
Groups	2.870	2	1.435	0.250	0.782
Residual	103.503	18	5.750		

- Definite evidence for a change in time.
- No significant group effect, marginally or in time.
- Post-hoc tests on interaction (G1T2-G2T2, G1T2-G3T2, G2T2-G3T2) provide analogous (though not equivalent) info as the ANCOVA: here, all post hoc *t*-tests are not close to significant.

RM-ANOVA vs. ANCOVA

So should we perform ANCOVAs or RM-ANOVAs? My advice:

- Never perform ANCOVAs. If you want to account for continuous covariates, use regression modelling instead.
- Regression modelling is a much more statistically defensible and flexible approach than ANCOVA.
- Use RM-ANOVAs for pre-post (i.e., 2 time point) designs.
- Talk to a statistician for designs with more than 2 time points: general mixed effects models.

✓✓