The analysis of covariance (ANACOVA)

The analysis of covariance is an extension of the analysis of variance technique. Like ANOVA, several factors may be compared but in addition, all the observations are "adjusted" by (regressed on) a number of continuous variables.

The simplest analysis of covariance model involves one nominal (categorical) and one continuous variable (covariate, or concomitant variable) and equal number of observations per cell as follows:

That is, there is a part of the model that is exactly like a (one-way) analysis of variance, and then there is an adjustment for the relationship of *y* on *x*.

For example, consider the following figure from an industrial experiment. There are three machines that test the strength of a fiber. Strength is also associated however, with the thickness (or diameter) of the fiber. Notice how the deviations from the regression lines are dramatically reduced compared to those from the mean (horizontal) lines that correspond to a one-way analysis of variance among the three machines (main effects).

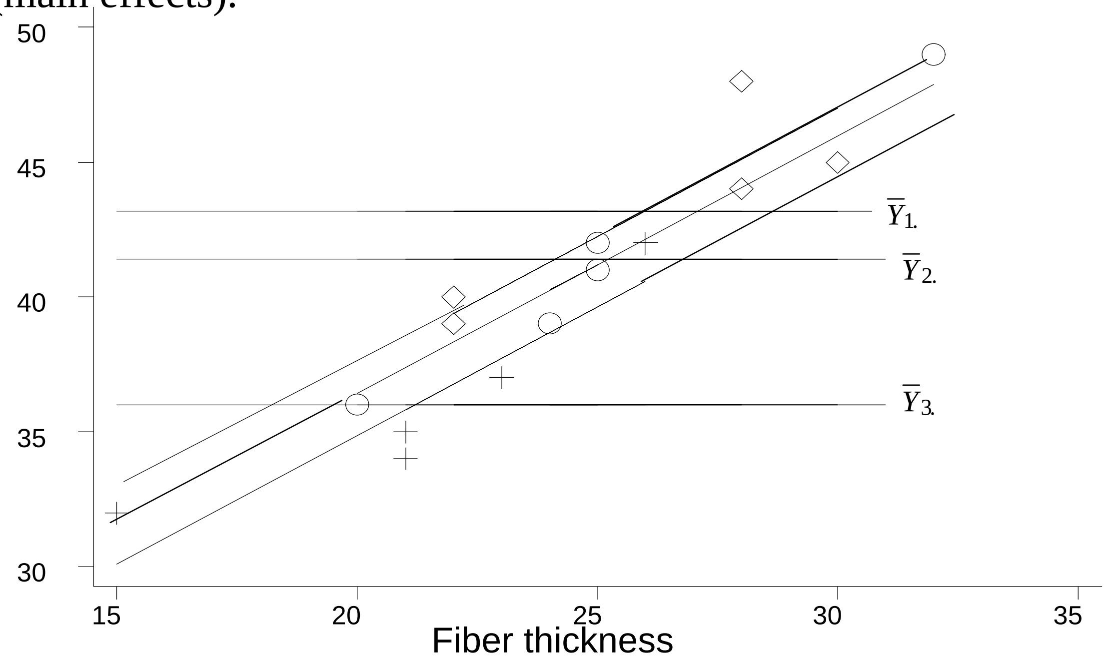

I

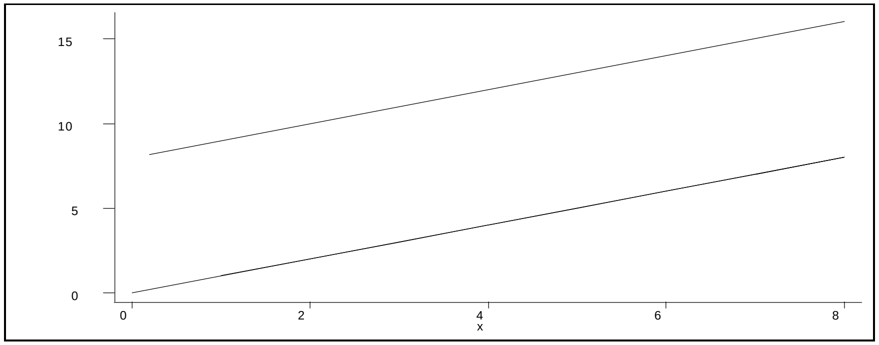

II

III

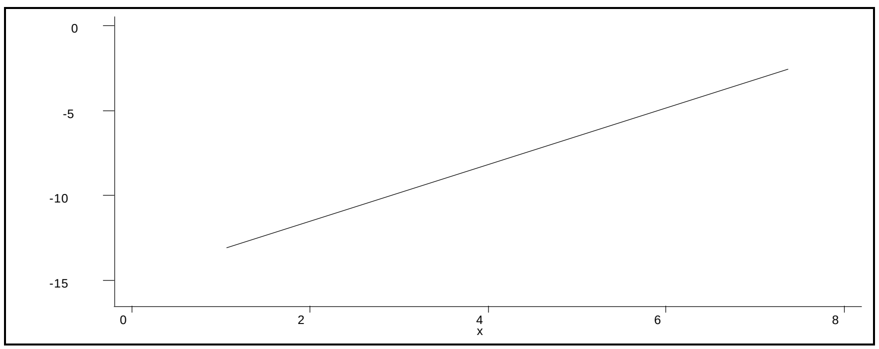

IV

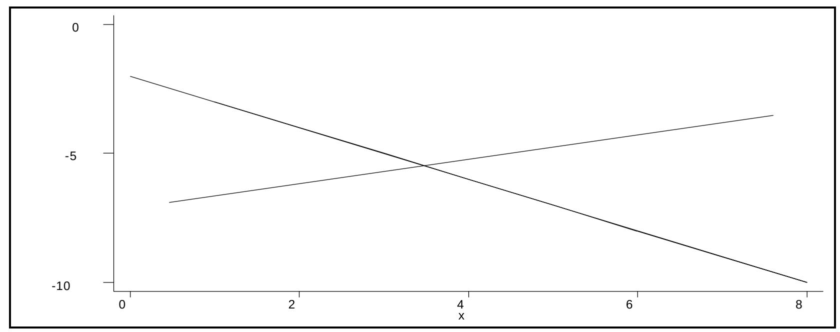

I. *Parallel lines*. In this case, intercepts $\beta_{1o} \neq \beta_{2o}$, but the slopes are equal; i.e., $\beta_{11} = \beta_{21}$

II. *Non-parallel lines.* In II, $\beta_{1o} = \beta_{2o}$ but $\beta_{11} \neq \beta_{21}$

III. *Coincident lines*. In this case, both $\beta_{1o} = \beta_{2o}$ and $\beta_{11} = \beta_{21}$.

IV. *Non-parallel lines*. Both the intercept and slope are different

Relevant scientific questions in the analysis of covariance

Frequently asked questions in the analysis of covariance are usually expressed as follows:

1. *Parallelism*. This question has to do with whether the relationship between the dependent variable and the covariate is the same in all groups considered (cases I and III above). If this is not the case, then we say that there is *interaction* between the groups and the covariate (see cases II and IV above).
2. *Differences among groups*. Usually the covariate is of secondary importance. Its primary use is as a necessary adjustment, in order to increase the precision of inference. The primary question is whether the mean values of the dependent variable in all groups are equal *after adjustment for the covariate.*

Comments:

1. In cases **I** and **III** the slopes are equal, which means that the relationship between the dependent variable and the covariate is identical in both groups. Parallelism implies that the difference between groups is constant and is equal to the differences in the intercepts. Thus, the statistical test that investigates differences among the groups is based on the differences of the group intercepts
2. In cases **II** and **IV**, the slopes are not equal, that is interaction is present, then the relationship between the dependent variable and the covariate is different in each group. Thus, we are no longer able to talk about adjustment for the covariate and separate analyses should be carried out. The conclusions from the statistical analysis must now focus on the different relationships between the dependent variable and covariate in the two groups rather than the difference in the adjusted mean values of the dependent variable in the groups. Thus, in that case, we will have several regression analyses rather than an analysis of variance of adjusted mean group values.

Definitions

To understand what happens by the adjustment with the continuous covariate we define the following quantities:

1. $S_{yy} = \sum_{i=1}^{a}\sum_{j=1}^{n}\left(y_{ij} - \bar{y}_{..}\right)^2$, $S_{xx} = \sum_{i=1}^{a}\sum_{j=1}^{n}\left(x_{ij} - \bar{x}_{..}\right)^2$, $S_{xy} = \sum_{i=1}^{a}\sum_{j=1}^{n}\left(x_{ij} - \bar{x}_{..}\right)\left(y_{ij} - \bar{y}_{..}\right)$
2. $T_{yy} = \sum_{i=1}^{a}\sum_{j=1}^{n}\left(\bar{y}_{i.} - \bar{y}_{..}\right)^2$, $T_{xx} = \sum_{i=1}^{a}\sum_{j=1}^{n}\left(\bar{x}_{i.} - \bar{x}_{..}\right)^2$, $T_{xy} = \sum_{i=1}^{a}\sum_{j=1}^{n}\left(\bar{x}_{i.} - \bar{x}_{..}\right)\left(\bar{y}_{i.} - \bar{y}_{..}\right)$
3. $E_{yy} = \sum_{i=1}^{a}\sum_{j=1}^{n}\left(y_{ij} - \bar{y}_{i.}\right)^2 = S_{yy} - T_{yy}$ $E_{xx} = \sum_{i=1}^{a}\sum_{j=1}^{n}\left(x_{ij} - \bar{x}_{i.}\right)^2 = S_{xx} - T_{xx}$

   $E_{xy} = \sum_{i=1}^{a}\sum_{j=1}^{n}\left(x_{ij} - \bar{x}_{i.}\right)\left(y_{ij} - \bar{y}_{i.}\right) = S_{xy} - T_{xy}$

In general, the prefixes *S*, *T* and *E* refer to sums of squares and cross-product terms for total, treatments and error respectively, and usually, $S=T+E$.

ANACOVA as an "adjusted" ANOVA

Recall that in the absence of treatment groups or, equivalently, in the absence of a statistically significant treatment effect, the least squares estimate of β is $\hat{\beta} = \dfrac{\sum_{i=1}^{a}\sum_{j=1}^{n}\left(x_{ij} - \bar{x}_{\cdot\cdot}\right)\left(y_{ij} - \bar{y}_{\cdot\cdot}\right)}{\sum_{i=1}^{a}\sum_{j=1}^{n}\left(x_{ij} - \bar{x}_{\cdot\cdot}\right)^2} = \dfrac{S_{xy}}{S_{xx}}$.

Also, the unadjusted (for the treatments) regression sum of squares is $SSR' = \left(S_{xy}\right)^2 / S_{xx}$, while $SSE' = S_{yy} - \left(S_{xy}\right)^2 / S_{xx}$ with 1 and $a(n\text{-}1)\text{-}1$ degrees of freedom respectively.

ANACOVA as an "adjusted" ANOVA (continued)

On the other hand, if treatment groups do exist, or, equivalently, there is a significant treatment effect, then the least squares estimate of β is a weighted average of the least squares estimates of b in each treatment group

$$\hat{\beta} = \frac{\sum_{j=1}^{n}\left(x_{1j}-\bar{x}_{1.}\right)^2\hat{\beta}_1+\sum_{j=1}^{n}\left(x_{2j}-\bar{x}_{2.}\right)^2\hat{\beta}_2+\cdots+\sum_{j=1}^{n}\left(x_{aj}-\bar{x}_{a.}\right)^2\hat{\beta}_a}{\sum_{j=1}^{n}\left(x_{1j}-\bar{x}_{1.}\right)^2+\sum_{j=1}^{n}\left(x_{2j}-\bar{x}_{2.}\right)^2+\cdots+\sum_{j=1}^{n}\left(x_{aj}-\bar{x}_{a.}\right)^2} = \frac{E_{xy}}{E_{xx}} = \frac{S_{xy}-T_{xy}}{S_{xx}-T_{xx}}$$

By similar arguments as in the case of simple linear regression, it can be shown, that the sum of squares for the regression adjusted for the treatment effect is $SSR = \left(E_{xy}\right)^2 / E_{xx}$, while the residual sum of squares is $SSE = E_{yy} - \left(E_{xy}\right)^2 / E_{xx} = \left(S_{yy}-T_{yy}\right) - \left(E_{xy}\right)^2 / E_{xx} = SSY - SST - SSR$, where SST is the usual treatment sum of squares.

Analysis of Covariance table

(expressed as an adjusted analysis of variance)

Source of variability	Sums of squares (SS)	Df	Mean squares (MS)	F	Prob > F
Regression	$SSR = (S_{xy})^2 / S_{xx}$	1	$MSR = SSR$	$F = \frac{MSR}{MSE}$	
Treatments	$SSE' - SSE = S_{yy} - (S_{xy})^2 / S_{xx} - \left[E_{yy} - (E_{xy})^2 / E_{xx}\right]$	a-1	$MST = \frac{SSE' - SSE}{a-1}$	$F = \frac{MST}{MSE}$	$p = P(F > F_{a-1,\ a(n-1)-1;p})$
Error	$SSE = E_{yy} - (E_{xy})^2 / E_{xx}$	$a(n$-1)-1	$MSE = \frac{SSE}{a(n-1)-1}$		
Total	S_{yy}	an-1			

Comments.

1. The error sum of squares $SSE^{'} > SSE$ since the latter results from the addition of the treatment effect. The F test is based on $F = \frac{(SSE' - SSE)/(a-1)}{SSE/[a(n-1)-1]} \sim F_{a-1, a(n-1)-1}$ is a partial F test that investigates the statistical significance of adding the treatment effect after the regression model.

2. The F test involving MSR, $F = \frac{(MSR)}{SSE'/[a(n-1)-1]} \sim F_{1, a(n-1)-1}$ merely tests whether there is a statistically significant relationship between the covariate and the dependent variable. This is not an appropriate test for testing parallelism between the lines corresponding to the treatment groups.

3.

ANACOVA for a single-factor experiment with one covariate

In this simplest case, both the significance of the regression (adjusted for the effect of treatment) and the treatment effect (adjusted for the covariate) are important. The ANACOVA table in this case is as follows:

Source of Variability	Sums of squares (SS)	Df	Mean squares (MS)	F	Prob > F
Regression	$SSR = (E_{xy})^2 / E_{xx}$	1	$MSR = SSR$	$F = \dfrac{MSR}{MSE}$	$p = P(F > F_{1,\ a(n-1)-1;p})$
Treatments	$SSE' - SSE = S_{yy} - (S_{xy})^2 / S_{xx} - \left[E_{yy} - (E_{xy})^2 / E_{xx}\right]$	a-1	$MST = \dfrac{SSE' - SSE}{a-1}$	$F = \dfrac{MST}{MSE}$	$p = P(F > F_{a-1,\ a(n-1)-1;p})$
Error	$SSE = E_{yy} - (E_{xy})^2 / E_{xx}$	$a(n$-1)-1	$MSE = \dfrac{SSE}{a(n-1)-1}$		
Total	S_{yy}	an-1			

1. <u>Test for the significance of regression</u> To assess the significance of regression, the hypotheses tests that are constructed are as follows:

   a. H_o: $\beta=0$
   The alternative hypotheses are constructed as follows:

   i. H_a: $\beta\neq 0$ (two-sided tests) ii. $\left.\begin{matrix} H_a:\beta > 0 \\ H_a:\beta < 0 \end{matrix}\right\}$ (one-sided tests)

   b. The test statistic is $F = \frac{MSR}{MSE} = \frac{(E_{xy})^2 / E_{xx}}{\left[E_{yy} - (E_{xy})^2 / E_{xx}\right] / [a(n-1)-1]} \sim F_{1,a(n-1)-1}$. This a Type III

partial F test of the regression in the presence of treatments and is equivalent to the t test based on the statistic $T = \frac{\hat{\beta}}{\text{s.e.}(\hat{\beta})} \sim t_{a(n-1)-1}$.

   c. The null hypothesis of non-significant regression is rejected if $F > F_{1,\ a(n-1)-1;1-\alpha}$ which is equivalent to rejecting the null hypothesis in favor of the two-sided alternative with $|T| > t_{a(n-1)-1;1-\alpha/2}$. On the other hand, the null hypothesis is rejected in favor of the one-sided alternative hypotheses, if $T > t_{a(n-1)-1;1-\alpha}$ and $T < t_{-a(n-1)-1;\alpha}$ respectively.

2. Test of parallelism In order to test for parallelism, the following model is assumed:

$$y_{ij} = \mu + \tau_i + \beta\left(x_{ij} - \bar{x}_{..}\right) + \gamma_i\left(x_{ij} - \bar{x}_{..}\right) + \varepsilon_{ij} \begin{cases} i = 1,2,\ldots,a \\ j = 1,2,\ldots,n \end{cases}$$

The slope for the i^{th} treatment is given by $\beta^* = \beta + \gamma_i$. The term γ_i is called the *interaction* term. Just like the main (categorical) effect, a usual assumption is that $\sum_{i=1}^{a} \gamma_i = 0$. The test for parallelism is constructed as follows:

a. H_o: $\gamma_1 = \gamma_2 = \ldots = \gamma_a = 0$ versus H_a: at least some of the γ's are not zero.

b. The test statistic $F = \dfrac{[\text{SSR}(X,Z,XZ) - \text{SSR}(X,Z)]/(a-1)}{\text{MSE}(X,Z,XZ)} \sim F_{a-1,a(n-1)-(a-1)-1}$ defines a Type III partial F test, where SSR(X,Z,XZ) is the regression sum of squares of the full model that includes the interaction term, while SSR(X,Z) is the model with only treatment and regression effects and no interaction, and MSE(X,Z,XZ) is the mean square error from the full model.

c. The null hypothesis (of no interaction, or of parallel lines across treatment groups) is rejected if $F > F_{a-1,a(n-1)-(a-1)-1;\alpha}$

3. Test of no difference in the slopes

If the test for no-interaction (parallelism) is rejected, we can consider the regression lines as parallel. Geometrically, parallel lines have a constant distance between them that corresponds to the difference between their slopes. In the model above, the difference of slopes *i* and *j* amounts to comparison of treatment τ_i versus τ_j, i.e., $\beta_{oi} - \beta_{oj} = \tau_i - \tau_j$. In the simplest case of two lines this amounts to an independent sample *t* test, while in the case of *k*>2 lines, the usual (Type III) *F* test (in the presence of the covariate) is used (with possibly a number of *post-hoc* comparisons). Note that this test is given by $F = \frac{SS(X,Z) - SS(Z)}{MSE(X,Z)} \sim F_{a-1,a(n-1)-1}$.

Example: Age-systolic blood pressure data. The systolic blood pressure is compared among males and females (Table 11-1 in text). The expectation is that males have different blood pressure from females, but the effect of age on blood pressure is also an important factor. The data set follows:

```
. by gender: list SBP age
-> gender=     Male                                -> gender=    Female
      SBP   age         SBP   age                        SBP   age         SBP   age
  1.  158    41   21.   150    38    1.   175    69   21.   120    39
  2.  185    60   22.   156    52    2.   158    67   22.   136    36
  3.  152    41   23.   134    41    3.   170    67   23.   110    34
  4.  159    47   24.   134    18    4.   162    65   24.   130    29
  5.  176    66   25.   174    51    5.   162    64   25.   125    25
  6.  156    47   26.   174    55    6.   144    63   26.   120    21
  7.  184    68   27.   158    65    7.   140    59   27.   116    20
  8.  138    43   28.   144    33    8.   150    56   28.   124    19
  9.  172    68   29.   139    23    9.   154    56   29.   114    17
 10.  168    57   30.   180    70   10.   158    53
 11.  176    65   31.   165    56   11.   142    50
 12.  164    57   32.   172    62   12.   130    48
 13.  154    61   33.   160    51   13.   145    47
 14.  124    36   34.   157    48   14.   142    46
 15.  142    44   35.   170    59   15.   135    45
 16.  144    50   36.   153    40   16.   138    45
 17.  149    47   37.   148    35   17.   160    44
 18.  128    19   38.   140    33   18.   124    42
 19.  130    22   39.   132    26   19.   128    42
 20.  138    21   40    169    61   20.   144    39
```

The comparison between males and females is conducted as follows:

```
. by gender: sum SBP age

-> gender=     Male
Variable |     Obs        Mean   Std. Dev.       Min        Max
---------+-----------------------------------------------------
     SBP |      40      155.15   16.56618        124        185
     age |      40      46.925   14.87105         18         70

-> gender=   Female
Variable |     Obs        Mean   Std. Dev.       Min        Max
---------+-----------------------------------------------------
     SBP |      29    139.8621   17.50454        110        175
     age |      29    45.06897   15.56078         17         69

. sum SBP age

Variable |     Obs        Mean   Std. Dev.       Min        Max
---------+-----------------------------------------------------
     SBP |      69    148.7246   18.47565        110        185
     age |      69    46.14493   15.07947         17         70
```

The analysis points out a significant difference in mean systolic blood pressure between males and females. However, we must adjust for the age effect in both groups. The mean in the two groups is: $\overline{Y}_{\text{M}} = 155.15$ for the males, and $\overline{Y}_{\text{F}} = 139.86$ for the females.

The regression of systolic blood pressure (SBP) on age is conducted as follows:

```
. reg SBP age

  Source |       SS       df       MS                  Number of obs =      69
---------+------------------------------               F(  1,    67) =  121.27
   Model |  14951.2546     1  14951.2546               Prob > F      =  0.0000
Residual |  8260.51351    67  123.291246               R-squared     =  0.6441
---------+------------------------------               Adj R-squared =  0.6388
   Total |  23211.7681    68  341.349531               Root MSE      =  11.104

------------------------------------------------------------------------------
     SBP |      Coef.   Std. Err.       t     P>|t|       [95% Conf. Interval]
---------+--------------------------------------------------------------------
     age |   .9833276   .0892947     11.012   0.000       .8050947    1.161561
   _cons |   103.3491   4.331896     23.858   0.000       94.70256    111.9956
------------------------------------------------------------------------------
```

Analysis of variance adjusted for the effect of age:

```
. anova SBP age gender, continuous(age) seq

                         Number of obs =      69     R-squared     =  0.7759
                         Root MSE      = 8.87795     Adj R-squared =  0.7691

               Source |    Seq. SS     df       MS            F     Prob > F
           -----------+----------------------------------------------------
                Model |  18009.7794     2  9004.88968     114.25     0.0000
                      |
                  age |  14951.2546     1  14951.2546     189.69     0.0000
               gender |  3058.52475     1  3058.52475      38.80     0.0000
                      |
             Residual |  5201.98876    66  78.8180115
           -----------+----------------------------------------------------
                Total |  23211.7681    68  341.349531
```

The plot of the data and the necessary STATA commands are as follows:

```
. predict sbphat
. gen shat0=sbphat if gender==0
. gen shat1=sbphat if gender==1
. gen sbp0=SBP if gender==0
. gen sbp1=SBP if gender==1
. label var sbp0 "Systolic Blood Pressure (females)"
. label var sbp1 "Systolic Blood Pressure (males)"
. graph shat0 shat1 sbp0 sbp1, c(ll..) s(iido) xlab ylab
     l1(Mean Systolic Blood Pressure)
```

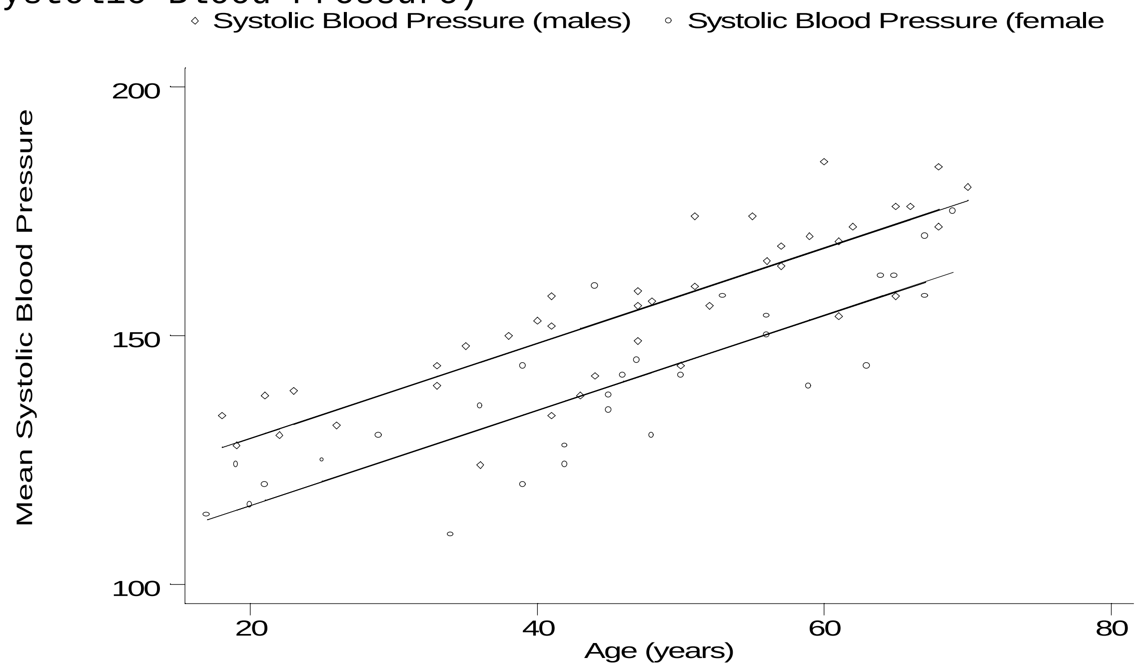

Comments:

1. From the simple linear regression we have $SS(\hat{\beta}_1) = S_{xy}^2 / S_{xx}$ =14951.25, and the sums squares for the residual are $SSE' = S_{yy} - S_{xy}^2 / S_{xx}$ = 3058.52.
2. The overall (from the whole data) estimate of the slope of systolic blood pressure with age is $\hat{\beta}_1 = 0.9833$, meaning that for every year of life, blood pressure increases 0.98 units on average.
3. From the ANOVA table, we see that the $SS(\hat{\beta}_1) = S_{xy}^2 / S_{xx} = 14951.25$. This is the same as the sum squares from a simple linear regression. The Type I sum of squares for the gender comparison is $SS(\text{GENDER} | \hat{\beta}_1) = 3058.52$ leading to a statistically significant partial F=38.80.
4. The $S_y = \sqrt{MSE} = \sqrt{78.8180} = 8.878$. This is smaller than the variance estimate that was produced by the ANOVA model without the age effect that was $S_y' = \sqrt{123.2912} = 11.104$, indicating that the adjustment of the analysis by the age effect was warranted.

Analysis of covariance with one effect and a single covariate

The analysis of covariance where both the gender and age effect are considered as adjustments is as follows:

```
. anova SBP age gender, continuous(age)

                       Number of obs =      69     R-squared     =  0.7759
                       Root MSE      = 8.87795     Adj R-squared =  0.7691

              Source |  Partial SS    df       MS           F     Prob > F
          -----------+----------------------------------------------------
               Model |  18009.7794     2  9004.88968     114.25     0.0000
                     |
                 age |  14080.5595     1  14080.5595     178.65     0.0000
              gender |  3058.52475     1  3058.52475      38.80     0.0000
                     |
            Residual |  5201.98876    66  78.8180115
          -----------+----------------------------------------------------
               Total |  23211.7681    68  341.349531
```

In order to obtain the regression estimates for `age` we add the option `regress` in the previous `anova` command as follows:

```
. reg

  Source |       SS       df       MS                  Number of obs =      69
---------+------------------------------               F(  2,    66) =  114.25
   Model |  18009.7794     2  9004.88968               Prob > F      =  0.0000
Residual |  5201.98876    66  78.8180115               R-squared     =  0.7759
---------+------------------------------               Adj R-squared =  0.7691
   Total |  23211.7681    68  341.349531               Root MSE      =   8.878

------------------------------------------------------------------------------
     SBP        Coef.   Std. Err.       t     P>|t|       [95% Conf. Interval]
------------------------------------------------------------------------------
_cons         96.77353   3.620854     26.727   0.000       89.54426    104.0028
age            .956058   .0715298     13.366   0.000       .8132441    1.098872
gender
       1      13.51345   2.169318      6.229   0.000       9.182273    17.84464
       2    (dropped)
------------------------------------------------------------------------------
```

Comments:

1. The partial sums of squares $SS(\text{AGE}\,|\,\text{GENDER}) = 14080.56$, $SS(\text{GENDER}\,|\,\text{AGE}) = 3058.52$. These lead to Type III *F* statistics *F*(AGE|GENDER)=178.65 and *F*(GENDER|AGE)=38.80, both of which are statistically significant.
2. The estimate of the common slope between males and females is given in the output of the regression model as $\hat{\beta}_1 = 0.9561$ which slightly less than 0.9833, the unadjusted estimate of the slope from the simple linear regression of SBP on age.
3. We saw previously that the raw (unadjusted) mean SBP for females and males were $\overline{Y}_{\text{F}} = 139.86$ and $\overline{Y}_{\text{M}} = 155.15$ respectively. The *adjusted* (for the effect of age) mean SBP for females and males are respectively $\overline{Y}_{\text{Fadj.}} = \overline{Y}_{\text{F}} - \hat{\beta}_1\left(\overline{X}_{\text{F.}} - \overline{X}_{..}\right) = 139.86 - 0.9561(45.069 - 46.144) = 140.89$ and $\overline{Y}_{\text{Madj.}} = \overline{Y}_{\text{M}} - \hat{\beta}_1\left(\overline{X}_{\text{M.}} - \overline{X}_{..}\right) = 155.15 - 0.9561(46.925 - 46.144) = 154.40$. Thus, adjustment for the age of the two groups produced a slight adjustment in the mean SBP in each group.

If we want to investigate the presence of interaction between age and gender we use the following model:

```
. anova SBP age gender age*gender, continuous(age)

                           Number of obs =      69     R-squared     =  0.7759
                           Root MSE      = 8.94551     Adj R-squared =  0.7656

                  Source |  Partial SS    df       MS           F     Prob > F
              -----------+----------------------------------------------------
                   Model |  18010.3287     3   6003.4429      75.02     0.0000
                         |
                     age |   13857.691     1   13857.691     173.17     0.0000
                  gender |  273.443297     1  273.443297       3.42     0.0691
              age*gender |  .549356192     1  .549356192       0.01     0.9342
                         |
                Residual |   5201.4394    65  80.0221447
              -----------+----------------------------------------------------
                   Total |  23211.7681    68  341.349531
```

To obtain the regression estimates for `age` we add the option `regress` in the previous `anova` command as follows:

```
. reg

  Source |       SS       df       MS                  Number of obs =      69
---------+------------------------------               F(  3,    65) =   75.02
   Model |  18010.3287     3   6003.4429               Prob > F      =  0.0000
Residual |   5201.4394    65  80.0221447               R-squared     =  0.7759
---------+------------------------------               Adj R-squared =  0.7656
   Total |  23211.7681    68  341.349531               Root MSE      =  8.9455

------------------------------------------------------------------------------
     SBP        Coef.   Std. Err.        t     P>|t|       [95% Conf. Interval]
------------------------------------------------------------------------------
_cons         97.07708   5.170455     18.775   0.000       86.75097    107.4032
age           .9493225   .1086412      8.738   0.000        .732351    1.166294
gender
       1      12.96144   7.011725      1.849   0.069      -1.041936    26.96482
       2     (dropped)
age*gender
       1      .0120301   .1451933      0.083   0.934      -.2779409    .3020011
       2     (dropped)
```

Comments:

1. By our assumptions ($\sum_{i=1}^{a}\tau_i = 0$ and $\sum_{i=1}^{a}\gamma_i = 0$), only one degree of freedom is associated with the effect τ (`gender`) and interaction γ_i `age*gender`. Their levels are *deviations* from overall quantities μ and β respectively. STATA defines the level with the largest coded value defines as the reference level, (female -- gender=1) so, $\hat{\beta}_F = 0.9493 = \hat{\beta}$ (i.e., $\hat{\gamma}_F = 0$), $\hat{\beta}_{oF} = 97.007 = \hat{\mu}$, (i.e., $\hat{\tau}_F = 0$), $\hat{\beta}_M = 0.9613 = 0.9493 + 0.0120 = \hat{\beta} + \hat{\gamma}_M$ $\hat{\beta}_{oM} = 109.968 = 97.007 + (12.961) = \hat{\mu} + \hat{\tau}_{ÁM}$
2. A crucial point is that in this statistical model the *t* test is not equivalent to the Type III *F* test and the estimates are dependent on the coding of the variable `gender`. Here male=0 and female=1. However, had we coded male=1 and female=0, then $\hat{\beta} = 0.9613 = \hat{\beta}_M$, while $\hat{\beta}_F = 0.9613 - 0.0120 = 0.9493$. A variation of this model that produces *t* tests exactly equivalent to the partial *F* tests will be considered during the next lecture.
3. Since interaction is not significant, the assumption of parallelism holds. We must return to the previous analysis without interaction (presented in the regression output and the figure above).

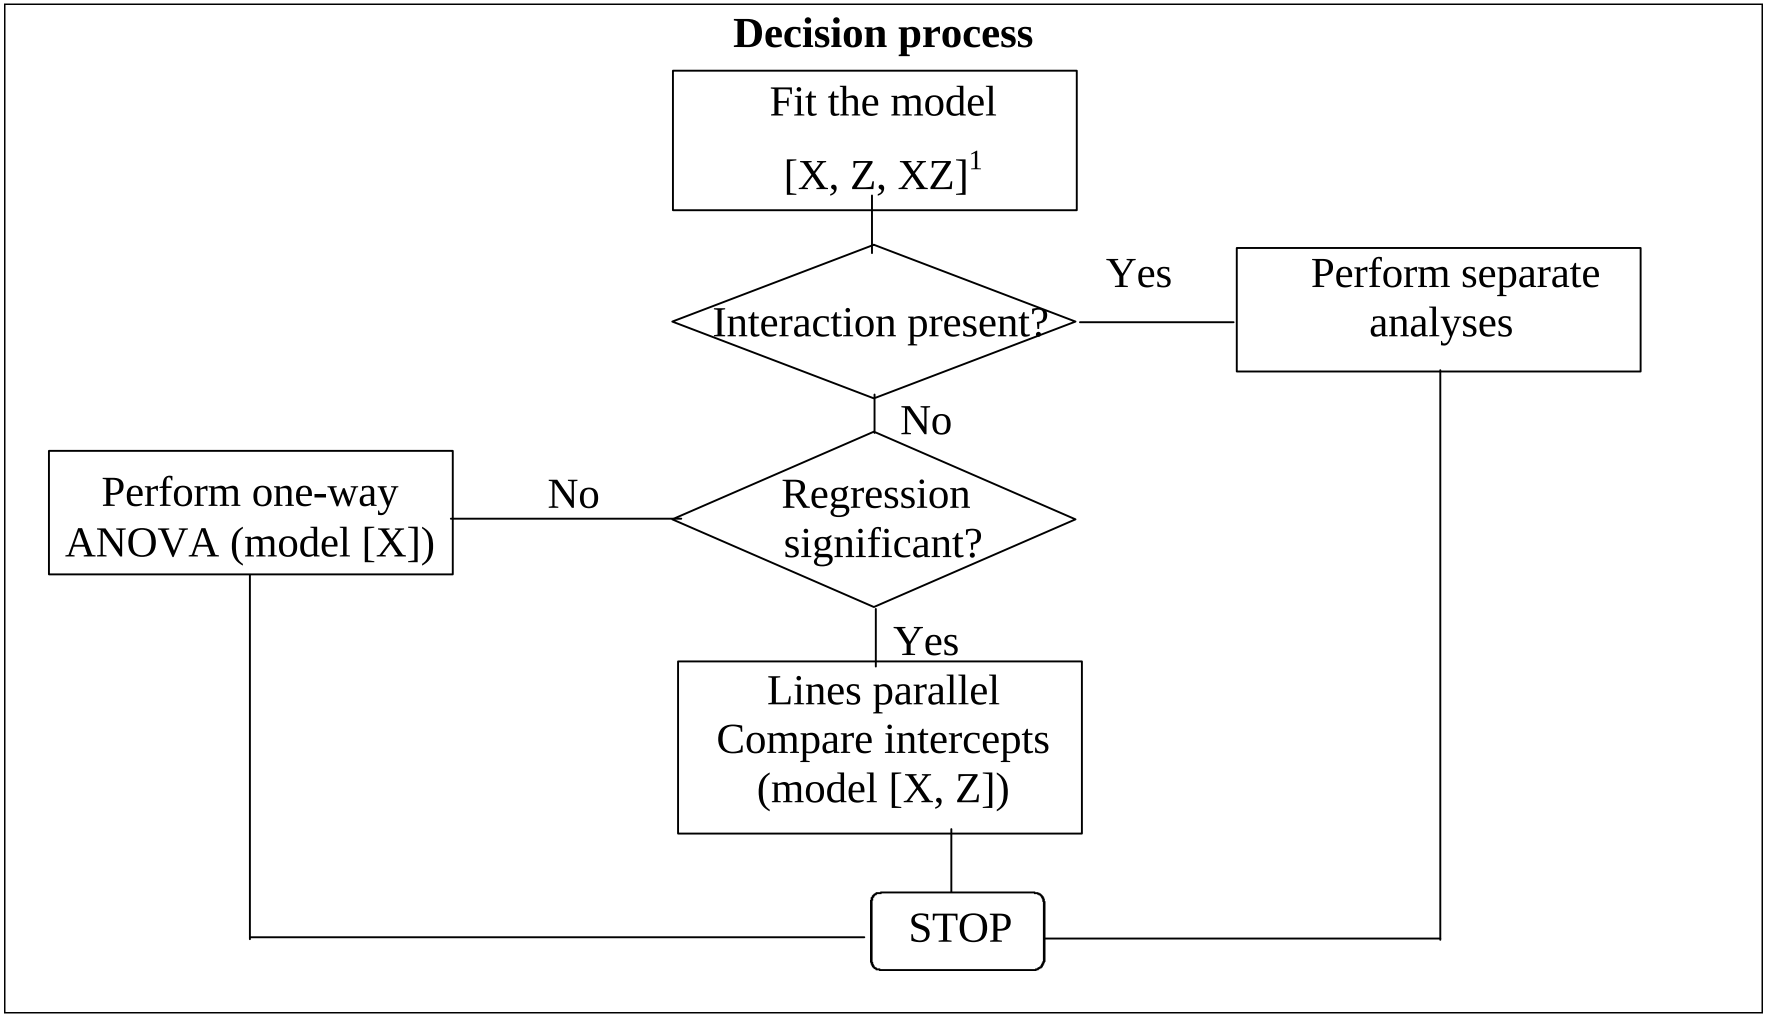

[1] This is the full model containing the interaction term XZ, the main effect X and the covariate Z