

Statistics 512: Applied Linear Models

Topic 8

Topic Overview

This topic will cover

- More on Multiple Comparisons / Confidence Intervals
- Two-way ANOVA with unequal numbers of observations in the cells
- Three-way ANOVA example

Estimation of Factor Level Means (19.8)

- Still assume n constant across cells here.
- Point Estimates are $\hat{\mu}_{i.} = \bar{Y}_{i.}$, $\hat{\mu}_{.j} = \bar{Y}_{.j}$, and $\hat{\mu}_{i,j} = \bar{Y}_{i,j}$.
- These have associated variances (estimate by plugging in MSE): $s^2\{\bar{Y}_{i.}\} = MSE/bn$, $s^2\{\bar{Y}_{.j}\} = MSE/an$, and $s^2\{\bar{Y}_{i,j}\} = MSE/n$.
- These may be used with t -critical values to form confidence intervals. The degrees of freedom are those associated with the MSE: $(n-1)ab$. It is not really appropriate to look at $\hat{\mu}_{i.}$ or $\hat{\mu}_{.j}$ when there is serious interaction.

Computation

- Means can be obtained from `proc means` in SAS
- MSE for the model can be obtained from SAS as well
- Construct the CI using these values and the appropriate critical value. The critical value can be from the t -distribution. Or it may be Tukey, Bonferroni, or Scheffe adjusted as is appropriate.

Contrasts

- We can look at contrasts of means (on the same factor) including multiple comparisons as we did in One-way ANOVA.
- *When there is no interaction*, for factor A the contrast $L = \sum c_i \mu_{i.}$ is estimated by $\hat{L} = \sum c_i \bar{Y}_{i.}$. An unbiased estimator of the variance is $s^2\{\hat{L}\} = \frac{MSE}{bn} \sum c_i^2$.
- Using a t -critical value (use error d.f.) we may construct a CI for the contrast.

- Factor B is analogous: the contrast $L = \sum c_j \mu_{.j}$ is estimated by $\hat{L} = \sum c_j \bar{Y}_{.j.}$. An unbiased estimator of the variance is $s^2\{\hat{L}\} = \frac{MSE}{an} \sum c_j^2$.
- If there is interaction, we may consider contrasts of the form $\hat{L} = \sum c_{i,j} \bar{Y}_{i,j.}$. For these, $s^2\{\hat{L}\} = \frac{MSE}{n} \sum c_{i,j}^2$ and CI's may be obtained using an appropriate critical value

Multiple Comparisons

- The multiple comparison procedures *with no interaction* are the same as for one-way ANOVA.
- Can use LSD, Tukey, Bonferroni, or Scheffe in SAS as appropriate in the `means` statement.

Example

- Recall the bread sales example (`knn1860.sas`)
- Shelf height (A) has 3 levels
- Shelf width (B) has 2 levels
- There are 2 observations at each level (total 12 observations)

Find a 95% CI for the mean sales using the middle wide shelf

So we want a 95% CI for $\mu_{2,2}$. From SAS we have the means output and also $MSE = 10.333$. We also had $\hat{\mu}_{2,2} = \bar{Y}_{2,2} = 69$. There are $(n-1)ab = 1(3)(2) = 6$ degrees of freedom, and if this is the only interval of interest we may use the t -distribution so that the critical value is 2.447. Hence the CI is given by $\hat{\mu}_{2,2} \pm 2.447 \left(\sqrt{MSE/n} \right) = 69 \pm 2.447 \left(\sqrt{10.333/2} \right) = (63.44, 74.56)$.

Find a 95% CI for the difference in sales between the middle shelf and the top shelf

Here we are averaging across width, looking at the contrast $\mu_{2.} - \mu_{1.}$. The marginal sample means were 67 and 44 respectively, so our point estimate is 23. The variance of the point estimate will be $s^2\{\hat{L}\} = \frac{MSE}{bn} \sum c_i^2 = \frac{10.33}{2 \times 2}(2) = 5.167$. In this case perhaps we are looking at all of the differences in means. It would then be appropriate to use a Tukey-adjusted critical value. There are 3 comparisons so the degrees of freedom will be 3 and 6. We have $q(0.95; 3, 6) = 4.34$ (from table B9) and so our critical value is $4.34/\sqrt{2} = 3.07$. The confidence interval is (16.02, 29.98).

Two-way ANOVA: Unbalanced Designs (Ch. 23)

From a data analysis point of view, the balanced design is the nicest. The “orthogonality” of this design makes it the most straightforward to analyze and understand. (Complete independence among factors.) However, there are times when equal sample sizes are not possible. Here are some of the reasons this might be the case (there may be others):

1. The experiment planned for a balanced design, but because of dropouts some data is unavailable (missing data).
2. The data were simply collected, not planned, and so the experimenter had no control over the number of observations in each “treatment”.
3. For reasons such as cost or ethics, it is not possible to examine all possible factor combinations. (Complicated designs such as these will be examined in Stat 514)
4. Some factor levels may be more important or more prevalent than others, and the experimenter wishes these to be more highly represented in the data.

As we will see, the reason for the unbalance may influence our interpretation of the results.

Data for two-way ANOVA

- Y , the response variable
- Factor A with levels $i = 1$ to a
- Factor B with levels $j = 1$ to b
- $Y_{i,j,k}$ is the k th observation for treatment (i, j) , $k = 1$ to $n_{i,j}$.
- Now (in Chapter 23) we do not have equal sample size (i.e. we have an *unbalanced design*) in each treatment combination. This causes complications in our analysis.

KNNL Example

- KNNL page 954 (`knn1954.sas`)
- Y is the change in growth rates for children after a treatment
- A is gender, $a = 2$ levels: male, female
- B is bone development, $b = 3$ levels: severely, moderately, or mildly depressed
- $n_{i,j} = 3, 2, 2, 1, 3, 3$ children in the groups

Read and check the data

```
data hormone;
infile 'h:\System\Desktop\CH22TA01.DAT';
input growth gender bone;
proc print data=hormone;
```

Obs	growth	gender	bone
1	1.4	1	1
2	2.4	1	1
3	2.2	1	1
4	2.1	1	2
5	1.7	1	2
6	0.7	1	3
7	1.1	1	3
8	2.4	2	1
9	2.5	2	2
10	1.8	2	2
11	2.0	2	2
12	0.5	2	3
13	0.9	2	3
14	1.3	2	3

Prepare the data for a plot

```
data hormone; set hormone;
  if (gender eq 1)*(bone eq 1) then gb='1_Msev ';
  if (gender eq 1)*(bone eq 2) then gb='2_Mmod ';
  if (gender eq 1)*(bone eq 3) then gb='3_Mmild';
  if (gender eq 2)*(bone eq 1) then gb='4_Fsev ';
  if (gender eq 2)*(bone eq 2) then gb='5_Fmod ';
  if (gender eq 2)*(bone eq 3) then gb='6_Fmild';
```

Plot the data

```
title1 'Plot of the data';
symbol1 v=circle i=none c=black;
proc gplot data=hormone;
  plot growth*gb;
```

Find the means

```
proc means data=hormone;
  output out=means mean=avgrowth;
  by gender bone;
```

Plot the means

```
title1 'Plot of the means';
symbol1 v='M' i=join c=black;
symbol2 v='F' i=join c=black;
proc gplot data=means;
    plot avgrowth*bone=gender;
```

Cell Means Model

$$Y_{i,j,k} = \mu_{i,j} + \epsilon_{i,j,k}$$

where $\mu_{i,j}$ is the theoretical mean or expected value of all observations in cell (i, j) .
the $\epsilon_{i,j,k}$ are iid $N(0, \sigma^2)$
 $Y_{i,j,k} \sim N(\mu_{i,j}, \sigma^2)$, independent

Estimates

Estimate $\mu_{i,j}$ by the mean of the observations in cell (i, j) , $\hat{\mu}_{i,j} = \bar{Y}_{i,j} = \frac{\sum_k Y_{i,j,k}}{n_{i,j}}$.

For each (i, j) combination, we can get an estimate of the variance $s_{i,j}^2 = \frac{\sum_k (Y_{i,j,k} - \bar{Y}_{i,j})^2}{n_{i,j} - 1}$, as long as $n_{i,j} \geq 2$.

We pool these to get an estimate of σ^2 .

Pooled Estimate of σ^2

In general we pool the $s_{i,j}^2$, using weights proportional to the df, $n_{i,j} - 1$.

The pooled estimate is $s^2 = \frac{\sum_{i,j} (n_{i,j} - 1) s_{i,j}^2}{\sum_{i,j} (n_{i,j} - 1)}$.

(Notice that if $n_{i,j} = 1$ we cannot calculate $s_{i,j}$, but its weight is zero anyway).

Run proc glm

```
proc glm data=hormone;
    class gender bone;
    model growth=gender|bone/solution;
    means gender*bone;
```

The syntax `gender|bone` is short for `gender bone gender*bone`. See SAS help on the “bar operator” for more information.

Parameter Estimates

The `solution` option on the `model` statement gives parameter estimates for the `glm` parameterization.

Parameter		Estimate	Standard Error	t Value	Pr > t
Intercept		0.900000000 B	0.23273733	3.87	0.0048
gender	1	-0.000000000 B	0.36799004	-0.00	1.0000
gender	2	0.000000000 B	.	.	.
bone	1	1.500000000 B	0.46547467	3.22	0.0122
bone	2	1.200000000 B	0.32914029	3.65	0.0065
bone	3	0.000000000 B	.	.	.
gender*bone	1 1	-0.400000000 B	0.59336610	-0.67	0.5192
gender*bone	1 2	-0.200000000 B	0.52041650	-0.38	0.7108
gender*bone	1 3	0.000000000 B	.	.	.
gender*bone	2 1	0.000000000 B	.	.	.
gender*bone	2 2	0.000000000 B	.	.	.
gender*bone	2 3	0.000000000 B	.	.	.

These constraints are (as we have seen before)

- Last level of each main effect is zero
- Interaction terms with *a* or *b* are zero

These can be rearranged to get the cell means in the usual way.

ANOVA Summary

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	5	4.47428571	0.89485714	5.51	0.0172
Error	8	1.30000000	0.16250000		
Corrected Total	13	5.77428571			

Note DF and SS add as usual.

Type I and Type III SS

In our previous ANOVA example, Type I and Type III SS were identical. This was true because the fact that the sample sizes were all the same made the variables completely orthogonal.

When sample sizes are unequal, the SS do not break down in the usual way. The various SS that we can calculate will not necessarily add up to the SSM.

SAS actually has four types of SS (I, II, III, IV) it can calculate. It does `ss1` and `ss3` by default but you can also ask for `ss2` and `ss4`.

We will focus on Type I and Type III in ANOVA.

Type I

Recall that Type I SS refer to the difference in SS when variables are added sequentially in the model, i.e. $SS(A)$, $SS(B|A)$, $SS(A \times B|A, B)$. Type I weights each observation equally, with the result that the treatments are weighted in proportion to their $n_{i,j}$.

Type II

Recall from regression that Type II SS referred to the difference in SSM when a variable is included last in the model or not (i.e., $SS(A|B, A \times B)$, $SS(B|A, A \times B)$, $SS(A \times B|A, B)$). Type II also weights each observation equally, with the result that the treatments are weighted in proportion to their $n_{i,j}$.

Type III

The ANOVA Type III SS are similar to the Type II SS, in that the other variables are assumed to already be in the model (this variable included last). Type III SS adjust for the cells having different $n_{i,j}$, by weighting each treatment equally, so that the observations are weighted differently. Therefore, when the sample sizes are unequal, Type III SS are more informative about the treatments than Type I. The Type III SS are calculated using regression with indicator variables to do the ANOVA, and to calculate the SSM for the full and reduced models. In Sections 23.2-3, KNNL are discussing Type III SS (they don't call them that; the type numbers are a SAS convention).

Type IV

Type IV SS are like Type III, except that Type IV additionally take into account possibly empty cells ($n_{i,j} = 0$). If there are empty cells, then Type IV SS are preferred. See KNNL Section 23.4 about empty cells.

Output Type I

Source	DF	Type I SS	Mean Square	F Value	Pr > F
bone	2	4.30628571	2.15314286	13.25	0.0029
gender	1	0.09257143	0.09257143	0.57	0.4720
bone*gender	2	0.07542857	0.03771429	0.23	0.7980

$$SSG + SSB + SSGB = 4.47429 = SSM$$

Output Type III

Source	DF	Type III SS	Mean Square	F Value	Pr > F
bone	2	4.18971429	2.09485714	12.89	0.0031
gender	1	0.12000000	0.12000000	0.74	0.4152
bone*gender	2	0.07542857	0.03771429	0.23	0.7980

$$SSG + SSB + SSGB = 4.38514 \neq SSM$$

Type I vs Type III

- SS for Type I add up to model SS .
- SS for Type III do not necessarily add to SSM .
- Type I and Type III are the same for the interaction because it is the last term in the model, but the Type I and Type III analysis for the main effects are not necessarily the same.
- Different hypotheses are being examined with the two types.
- Most people prefer the Type III analysis.
- This can be misleading if the sample sizes differ greatly.
- Contrasts can provide some insight by showing us what is actually being calculated.

Using Contrasts to illustrate exactly what is being calculated with Type I and Type III SS

It would not be necessary to construct these contrasts in a typical analysis. But for illustration purposes, we are going to construct specific contrasts in terms of the cell means / factor effects parameters and show that they come out to the Type I and Type III SS , which should help you understand Type I/III SS better.

Contrast for $A \times B$

- This is the same for Type I and Type III.
- Null hypothesis is that the profiles are parallel; see plot for interpretation: the difference between the factor levels for *bone* is the same whether gender is 1 or 2.
- $H_0 : \mu_{1,2} - \mu_{1,1} = \mu_{2,2} - \mu_{2,1}$ and $\mu_{1,3} - \mu_{1,2} = \mu_{2,3} - \mu_{2,2}$
- Written with contrasts this is: $H_0 : L_1 = \mu_{1,1} - \mu_{1,2} - \mu_{2,1} + \mu_{2,2} = 0$ and $L_2 = \mu_{1,2} - \mu_{1,3} - \mu_{2,2} + \mu_{2,3} = 0$
- In terms of the factor effects parameters these are (μ 's, α 's and β 's cancel):

$$\begin{aligned}L_1 &= (\alpha\beta)_{1,1} - (\alpha\beta)_{1,2} - (\alpha\beta)_{2,1} + (\alpha\beta)_{2,2} \\L_2 &= (\alpha\beta)_{1,2} - (\alpha\beta)_{1,3} - (\alpha\beta)_{2,2} + (\alpha\beta)_{2,3}\end{aligned}$$

- Recall that SAS interpretes the coefficients in the contrast in terms of the factor effects parameters.

$A \times B$ contrast statement

```
contrast 'gender*bone Type I and III'
  gender*bone 1 -1 0 -1 1 0,
  gender*bone 0 1 -1 0 -1 1;
```

Type III Contrast for *gender*

- Null hypothesis is that the average for males and females is the same. For Type III each treatment mean has the same weight regardless of the sample size, so some observations are weighted more heavily than others.

$$H_0 : \mu_{1,1} + \mu_{1,2} + \mu_{1,3} = \mu_{2,1} + \mu_{2,2} + \mu_{2,3},$$

i.e.,

$$H_0 : L = 0, \text{ where } L = \mu_{1,1} + \mu_{1,2} + \mu_{1,3} - \mu_{2,1} - \mu_{2,2} - \mu_{2,3}$$

- In the hypothesis, all cell means are weighted equally (ignore different sample sizes).

Write L in terms of the factor effects:

$$\begin{array}{rcl} \mu_{1,1} & = & \mu + \alpha_1 + \beta_1 + (\alpha\beta)_{1,1} \\ \mu_{1,2} & = & \mu + \alpha_1 + \beta_2 + (\alpha\beta)_{1,2} \\ \mu_{1,3} & = & \mu + \alpha_1 + \beta_3 + (\alpha\beta)_{1,3} \\ -\mu_{2,1} & = & -(\mu + \alpha_2 + \beta_1 + (\alpha\beta)_{2,1}) \\ -\mu_{2,2} & = & -(\mu + \alpha_2 + \beta_2 + (\alpha\beta)_{2,2}) \\ -\mu_{2,3} & = & -(\mu + \alpha_2 + \beta_3 + (\alpha\beta)_{2,3}) \\ \hline L & = & 3\alpha_1 - 3\alpha_2 + (\alpha\beta)_{1,1} + (\alpha\beta)_{1,2} + (\alpha\beta)_{1,3} - (\alpha\beta)_{2,1} - (\alpha\beta)_{2,2} - (\alpha\beta)_{2,3} \end{array}$$

Contrast statement: gender Type III

```
contrast 'gender Type III'
  gender 3 -3
  gender*bone 1 1 1 -1 -1 -1;
```

Type I Contrast for *gender*

- Null hypothesis is that the average for males and females is the same.
- For Type I each treatment is weighted by its sample size because each observation is weighted equally.
- Note that the data are actually balanced for gender; that is, there are the same number of males ($3 + 2 + 2$) as females ($1 + 3 + 3$) so we can work with the sum instead of the averages. (Or we could just divide everything by 7.)

$$\begin{aligned} H_0 : \quad & \frac{3\mu_{1,1}+2\mu_{1,2}+2\mu_{1,3}}{7} = \frac{\mu_{2,1}+3\mu_{2,2}+3\mu_{2,3}}{7}; \text{ i.e.,} \\ H_0 : \quad & L = 3\mu_{1,1} + 2\mu_{1,2} + 2\mu_{1,3} - (\mu_{2,1} + 3\mu_{2,2} + 3\mu_{2,3}) = 0 \end{aligned}$$

Write L in terms of the factor effects:

$$\begin{array}{rcl} 3\mu_{1,1} & = & 3(\mu + \alpha_1 + \beta_1 + (\alpha\beta)_{1,1}) \\ 2\mu_{1,2} & = & 2(\mu + \alpha_1 + \beta_2 + (\alpha\beta)_{1,2}) \\ 2\mu_{1,3} & = & 2(\mu + \alpha_1 + \beta_3 + (\alpha\beta)_{1,3}) \\ -\mu_{2,1} & = & -(\mu + \alpha_2 + \beta_1 + (\alpha\beta)_{2,1}) \\ -3\mu_{2,2} & = & -3(\mu + \alpha_2 + \beta_2 + (\alpha\beta)_{2,2}) \\ -3\mu_{2,3} & = & -3(\mu + \alpha_2 + \beta_3 + (\alpha\beta)_{2,3}) \\ \hline L & = & (7\alpha_1 - 7\alpha_2) + (2\beta_1 - \beta_2 - \beta_3) + 3(\alpha\beta)_{1,1} + \\ & & 2(\alpha\beta)_{1,2} + 2(\alpha\beta)_{1,3} - (\alpha\beta)_{2,1} - 3(\alpha\beta)_{2,2} - 3(\alpha\beta)_{2,3} \end{array}$$

Contrast statement: gender Type I

```
contrast 'gender Type I'
gender 7 -7
bone 2 -1 -1
gender*bone 3 2 2 -1 -3 -3;
```

We could do the same thing for *bone* (work out the details yourself). The answers are as follows.

Bone **Type III**

$$\begin{aligned} H_0 : \quad & \frac{\mu_{1,1}+\mu_{2,1}}{2} = \frac{\mu_{1,2}+\mu_{2,2}}{2} = \frac{\mu_{1,3}+\mu_{2,3}}{2} \\ L_1 = \quad & \mu_{1,1} - \mu_{1,2} + \mu_{2,1} - \mu_{2,2} = 2\beta_1 - 2\beta_2 + \alpha\beta_{1,1} - \alpha\beta_{1,2} + \alpha\beta_{2,1} - \alpha\beta_{2,2} \\ L_2 = \quad & \mu_{1,2} - \mu_{1,3} + \mu_{2,2} - \mu_{2,3} = 2\beta_2 - 2\beta_3 + \alpha\beta_{1,2} - \alpha\beta_{1,3} + \alpha\beta_{2,2} - \alpha\beta_{2,3} \end{aligned}$$

$H_0 : L_1 = 0$ and $L_2 = 0$.

```
contrast 'bone Type III'
bone 2 -2 0
gender*bone 1 -1 0 1 -1 0,
bone 0 2 -2
gender*bone 0 1 -1 0 1 -1;
```

Bone **Type I**

$$\begin{aligned} H_0 : \quad & \frac{3\mu_{1,1}+\mu_{2,1}}{4} = \frac{2\mu_{1,2}+3\mu_{2,2}}{5} = \frac{2\mu_{1,3}+3\mu_{2,3}}{5} \\ L_1 = \quad & 15\mu_{1,1} - 8\mu_{1,2} + 5\mu_{2,1} - 12\mu_{2,2} \\ & = 7\alpha_1 - 7\alpha_2 + 20\beta_1 - 20\beta_2 + 15\alpha\beta_{1,1} - 8\alpha\beta_{1,2} + 5\alpha\beta_{2,1} - 12\alpha\beta_{2,2} \\ L_2 = \quad & 2\mu_{1,2} - 2\mu_{1,3} + 3\mu_{2,2} - 3\mu_{2,3} \\ & = 5\beta_2 - 5\beta_3 + 2\alpha\beta_{1,2} - 2\alpha\beta_{1,3} + 3\alpha\beta_{2,2} - 3\alpha\beta_{2,3} \end{aligned}$$

```
contrast 'bone Type I'
  gender 7 -7
  bone 20 -20 0
  gender*bone 15 -8 0 5 -12 0,
  bone 0 5 -5
  gender*bone 0 2 -2 0 3 -3;
```

Contrast output

Contrast	DF	Contrast SS	F Value	Pr > F
gender Type III	1	0.12000000	0.74	0.4152
gender Type I	1	0.00285714	0.02	0.8978
bone Type III	2	4.18971429	12.89	0.0031
bone Type I	2	4.30628571	13.25	0.0029
gender*bone Type I/III	2	0.07542857	0.23	0.7980

Only *bone* is significant. Notice that the contrast *SS* match the appropriate Type I or III *SS* above. So these contrasts help clarify exactly what hypothesis is being tested by each *SS*.

Source	DF	Type I SS	Mean Square	F Value	Pr > F
gender	1	0.00285714	0.00285714	0.02	0.8978
bone	2	4.39600000	2.19800000	13.53	0.0027
gender*bone	2	0.07542857	0.03771429	0.23	0.7980

Source	DF	Type III SS	Mean Square	F Value	Pr > F
gender	1	0.12000000	0.12000000	0.74	0.4152
bone	2	4.18971429	2.09485714	12.89	0.0031
gender*bone	2	0.07542857	0.03771429	0.23	0.7980

Type I *bone* does not *exactly* match the contrast but it is close (4.396 vs. 4.306). This is because *bone* is second in the model. With *bone* listed first, the Type I *SS* for *bone* exactly matches the contrast.

Source	DF	Type I SS	Mean Square	F Value	Pr > F
bone	2	4.30628571	2.15314286	13.25	0.0029
gender	1	0.09257143	0.09257143	0.57	0.4720
bone*gender	2	0.07542857	0.03771429	0.23	0.7980

Remember: in a typical analysis, you would not do all these **contrast** statements. These only served the purpose of illustrating what the Type I and III *SS* actually mean.

Analytical Strategy

First examine interactions

Some options when the interaction is significant and important:

- Interpret the plot of means (interaction plot)

- Run A at each level of B and/or B at each level of A
- Run as a one-way with ab levels
- Use contrasts

Some options when the interaction is not significant:

- Use contrasts for main effects
- Rerun without the interaction
- Use a multiple comparison procedure for the main effects

Example without interaction

```
proc glm data=hormone;
  class gender bone;
  model growth=gender bone/solution;
  means gender bone/ tukey lines;
```

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	3	4.39885714	1.46628571	10.66	0.0019
Error	10	1.37542857	0.13754286		
Corrected Total	13	5.77428571			

R-Square	Coeff Var	Root MSE	growth Mean
0.761801	22.57456	0.370868	1.642857

Source	DF	Type I SS	Mean Square	F Value	Pr > F
gender	1	0.00285714	0.00285714	0.02	0.8883
bone	2	4.39600000	2.19800000	15.98	0.0008

Source	DF	Type III SS	Mean Square	F Value	Pr > F
gender	1	0.09257143	0.09257143	0.67	0.4311
bone	2	4.39600000	2.19800000	15.98	0.0008

Parameter	Estimate	Standard Error	t Value	Pr > t
Intercept	0.968571429 B	0.18572796	5.22	0.0004
gender 1	-0.171428571 B	0.20896028	-0.82	0.4311
gender 2	0.000000000 B	.	.	.
bone 1	1.260000000 B	0.25931289	4.86	0.0007
bone 2	1.120000000 B	0.23455733	4.77	0.0008
bone 3	0.000000000 B	.	.	.

Tukey Comparisons

	Mean	N	bone
A	2.1000	4	1
A			
A	2.0200	5	2
B	0.9000	5	3

Multiple Comparisons in an Unbalanced Setting

- Standard Errors for similar comparisons will now be DIFFERENT (e.g. if we look at all the differences of the form $\mu_i - \mu_{i'}$, their variances will not be the same).
- See pages 961 and 962 for the various formulas. They now have $n_{i,j}$'s all over the place.
- Everything else (formation of CI's, use of Multiple Comparison critical values, etc) still applies.

Three-way ANOVA

Data for three-way ANOVA

- Y , the response variable
- Factor A with levels $i = 1$ to a
- Factor B with levels $j = 1$ to b
- Factor C with levels $k = 1$ to c
- $Y_{i,j,k,\ell}$ is the ℓ th observation in cell (i, j, k) , $\ell = 1$ to $n_{i,j,k}$
- A balanced design has $n_{i,j,k} = n$

Cell Means Model

$$Y_{i,j,k,\ell} = \mu_{i,j,k} + \epsilon_{i,j,k,\ell}$$

- $\mu_{i,j,k}$ is the theoretical mean or expected value of all observations in cell (i, j, k) .
- $\epsilon_{i,j,k,\ell} \sim^{iid} N(0, \sigma^2)$
- $Y_{i,j,k,\ell} \sim N(\mu_{i,j,k}, \sigma^2)$ are independent

Estimates

- Estimate $\mu_{i,j,k}$ by the mean of the observations in cell (i, j, k) , $\bar{Y}_{i,j,k} = \frac{1}{n} \sum_{\ell} Y_{i,j,k,\ell}$.
- For each (i, j, k) combination, we can get an estimate of the variance $\sigma_{i,j,k}^2$:

$$s_{i,j,k}^2 = \frac{\sum_{\ell} (Y_{i,j,k,\ell} - \bar{Y}_{i,j,k})^2}{n_{i,j,k} - 1}.$$

- Combine these to get an estimate of σ^2 , since we assume they are all equal. In general we pool the $s_{i,j,k}^2$, using weights proportional to the df , $n_{i,j,k} - 1$. The pooled estimate is obtained using weights proportional to degrees of freedom as usual:

$$s^2 = \frac{\sum_{i,j,k} (n_{i,j,k} - 1) s_{i,j,k}^2}{\sum_{i,j,k} (n_{i,j,k} - 1)} = \frac{\sum_{i,j,k} (n_{i,j,k} - 1) s_{i,j,k}^2}{n_T - abc} = MSE$$

Factor Effects Model

$$Y_{i,j,k,\ell} = \mu + \alpha_i + \beta_j + \gamma_k + (\alpha\beta)_{i,j} + (\alpha\gamma)_{i,k} + (\beta\gamma)_{j,k} + (\alpha\beta\gamma)_{i,j,k} + \epsilon_{i,j,k,\ell}$$

- μ is the overall (grand) mean
- $\alpha_i, \beta_j, \gamma_k$ are the main effects of factors A, B , and C
- $(\alpha\beta)_{i,j}, (\alpha\gamma)_{i,k}, (\beta\gamma)_{j,k}$ are the two-way (first order) interactions
- $(\alpha\beta\gamma)_{i,j,k}$ is the three-way (second-order) interaction
- An extension of the usual constraints applies.

ANOVA table

Sources of *model* variation include

- the three main effects
- the three two-way interactions
- the (one) three-way interaction.

With balanced data the SS and df add to the model SS and df .

Always have $Model + Error = Total$.

Each effect is tested by an F -statistic with MSE in the denominator.

Analytical Strategy

First examine interactions

Some options when one or more interactions are significant

- Interpret the plot of means
- Run analyses for each level of one factor, eg run $A|B$ by C (`lsmeans` with slice option)
- Run as a one-way with abc levels
- Define a composite factor by combining two factors, eg AB with ab levels
- Use contrasts

Some options when no interactions are significant

- Use contrasts
- Rerun without the interactions
- Use a multiple comparison procedure for the main effects

KNNL Example

- KNNL page 1018 (`knnl1018.sas`)
- Y is exercise tolerance, minutes until fatigue on a bicycle test
- A is gender, $a = 2$ levels: male = 1, female = 2
- B is percent body fat, $b = 2$ levels: low = 1, high = 2
- C is smoking history, $c = 2$ levels: light = 1, heavy = 2
- $n = 3$ persons aged 25-35 per (i, j, k) cell

Read and check the data

```
data exercise;  
  infile 'h:\System\Desktop\CH23TA04.DAT';  
  input extol gender fat smoke;
```

Define variable for a plot

This is just to set a unique identifier for each treatment. There are other ways to do this.

```
data exercise;
set exercise;
gfs = 100*gender + 10*fat + smoke;
proc print data=exercise;
```

Obs	extol	gender	fat	smoke	gfs
1	24.1	1	1	1	111
2	29.2	1	1	1	111
3	24.6	1	1	1	111
4	20.0	2	1	1	211
5	21.9	2	1	1	211
6	17.6	2	1	1	211
7	14.6	1	2	1	121
8	15.3	1	2	1	121
9	12.3	1	2	1	121
10	16.1	2	2	1	221
11	9.3	2	2	1	221
12	10.8	2	2	1	221
13	17.6	1	1	2	112
14	18.8	1	1	2	112
15	23.2	1	1	2	112
16	14.8	2	1	2	212
17	10.3	2	1	2	212
18	11.3	2	1	2	212
19	14.9	1	2	2	122
20	20.4	1	2	2	122
21	12.8	1	2	2	122
22	10.1	2	2	2	222
23	14.4	2	2	2	222
24	6.1	2	2	2	222

Plot the data

```
proc sort data=exercise;
  by gender fat smoke;
title1 'Plot of the data';
symbol1 v=circle i=none c=black;
proc gplot data=exercise;
  plot extol*gfs/ haxis = 111 112 121 122 211 212 221 222;
```

Find the means

```
proc means data=exercise;
  output out=exer2 mean=avextol;
  by gender fat smoke;
```


Make a two-variable combination of *fat* and *smoke*

This is helpful for plotting.

```
data exer2;
    set exer2;
    fs = fat*10 + smoke;
proc print data=exer2;
```

Obs	gender	fat	smoke	avextol	fs
1	1	1	1	25.9667	11
2	1	1	2	19.8667	12
3	1	2	1	14.0667	21
4	1	2	2	16.0333	22
5	2	1	1	19.8333	11
6	2	1	2	12.1333	12
7	2	2	1	12.0667	21
8	2	2	2	10.2000	22

Plot the means

```
proc sort data=exer2; by fs;
title1 'Plot of the means';
symbol1 v='M' i=join c=black;
symbol2 v='F' i=join c=black;
proc gplot data=exer2;
    plot avextol*fs=gender / haxis = 11 12 21 22;
```

From this plot it appears that *gender* probably doesn't interact too much with the other variables.

Note: Interaction plots in the 3-variable model take the form of putting 2-factor combinations on the X-axis with separate lines for the third factor.

```
proc glm data=exercise;
    class gender fat smoke;
    model extol=gender|fat|smoke / solution;
    means gender*fat*smoke;
```

Recall that `gender|fat|smoke` is short for `gender fat smoke gender*fat gender*smoke fat*smoke gender*fat*smoke`.

The GLM Procedure

Class Level Information		
Class	Levels	Values
gender	2	1 2
fat	2	1 2
smoke	2	1 2

Number of observations 24

Dependent Variable: extol

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	7	588.5829167	84.0832738	9.01	0.0002
Error	16	149.3666667	9.3354167		
Corrected Total	23	737.9495833			

R-Square 0.797592
 Coeff Var 18.77833
 Root MSE 3.055391
 extol Mean 16.27083

Source	DF	Type I SS	Mean Square	F Value	Pr > F
gender	1	176.5837500	176.5837500	18.92	0.0005
fat	1	242.5704167	242.5704167	25.98	0.0001
gender*fat	1	13.6504167	13.6504167	1.46	0.2441
smoke	1	70.3837500	70.3837500	7.54	0.0144
gender*smoke	1	11.0704167	11.0704167	1.19	0.2923
fat*smoke	1	72.4537500	72.4537500	7.76	0.0132
gender*fat*smoke	1	1.8704167	1.8704167	0.20	0.6604

All main effects are significant. *Gender* and *fat* appear to have bigger effects than *smoke*. The two-way interaction between *fat* and *smoke* is also significant.

SAS Parameter Estimates

Solution option on the model statement gives parameter estimates for the glm parameterization.

These are as we have seen before; any main effect or interaction with a subscript of *a*, *b*, or *c* is zero.

These can be used to reproduce the cell means in the usual way.

Parameter		Estimate		Standard Error	t Value	Pr > t
Intercept		10.2	B	1.76403105	5.78	<.0001
gender	1	5.83333333	B	2.49471664	2.34	0.0327
gender	2	0.0	B	.	.	.
fat	1	1.93333333	B	2.49471664	0.77	0.4497
fat	2	0.0	B	.	.	.
gender*fat	1 1	1.9	B	3.52806211	0.54	0.5976
gender*fat	1 2	0.0	B	.	.	.
gender*fat	2 1	0.0	B	.	.	.
gender*fat	2 2	0.0	B	.	.	.
smoke	1	1.86666667	B	2.49471664	0.75	0.4652
smoke	2	0.0	B	.	.	.
gender*smoke	1 1	-3.83333333	B	3.52806211	-1.09	0.2933
gender*smoke	1 2	0.0	B	.	.	.
gender*smoke	2 1	0.0	B	.	.	.
gender*smoke	2 2	0.0	B	.	.	.
fat*smoke	1 1	5.83333333	B	3.52806211	1.65	0.1177
fat*smoke	1 2	0.0	B	.	.	.

fat*smoke	2 1	0.0	B	.	.	.
fat*smoke	2 2	0.0	B	.	.	.
gender*fat*smoke	1 1 1	2.23333333	B	4.98943328	0.45	0.6604
gender*fat*smoke	1 1 2	0.0	B	.	.	.
gender*fat*smoke	1 2 1	0.0	B	.	.	.
gender*fat*smoke	1 2 2	0.0	B	.	.	.
gender*fat*smoke	2 1 1	0.0	B	.	.	.
gender*fat*smoke	2 1 2	0.0	B	.	.	.
gender*fat*smoke	2 2 1	0.0	B	.	.	.
gender*fat*smoke	2 2 2	0.0	B	.	.	.

We can get the zero-sum constraints in the usual way (see the file `knn11018.sas` for the code).

Obs	gender	fat	smoke	mu	alpha	beta	gamma	alphabeta	alphagamma	betagamma	abc
1	1	1	1	16.2708	2.7125	3.17917	1.7125	0.75417	-0.67917	1.7375	0.27917
4	1	1	2	16.2708	2.7125	3.17917	-1.7125	0.75417	0.67917	-1.7375	-0.27917
7	1	2	1	16.2708	2.7125	-3.17917	1.7125	-0.75417	-0.67917	-1.7375	-0.27917
10	1	2	2	16.2708	2.7125	-3.17917	-1.7125	-0.75417	0.67917	1.7375	0.27917
13	2	1	1	16.2708	-2.7125	3.17917	1.7125	-0.75417	0.67917	1.7375	-0.27917
16	2	1	2	16.2708	-2.7125	3.17917	-1.7125	-0.75417	-0.67917	-1.7375	0.27917
19	2	2	1	16.2708	-2.7125	-3.17917	1.7125	0.75417	0.67917	-1.7375	0.27917
22	2	2	2	16.2708	-2.7125	-3.17917	-1.7125	0.75417	-0.67917	1.7375	-0.27917

Notice from the parameter estimates that $\beta\gamma$ is about the same size as γ . This makes it pretty hard to interpret the main effect of *smoke*.

```

title1 'Mean over gender vs smoke';
symbol1 v=L i=join;
symbol2 v=H i=join;
proc gplot data=BCdat;
plot muBC*smoke=fat;

```

Looking at this plot, it appears that smoking decreases tolerance for those of low body fat, but makes almost no difference for those at the high body fat.

Example Approach

Since there appears to be a *fat* by *smoke* interaction, let's run a two-way ANOVA (no additional interaction) using the *fat* \times *smoke* variable and *gender*. This will consider the four *fs* categories separately.

We will also use the interaction plot to describe the interaction.

```

proc glm data=exercise;
class gender fs;
model extol=gender fs;
means gender fs/tukey;

```

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	4	561.9916667	140.4979167	15.17	<.0001
Error	19	175.9579167	9.2609430		
Corrected Total	23	737.9495833			

Source	DF	Type I SS	Mean Square	F Value	Pr > F
gender	1	176.5837500	176.5837500	19.07	0.0003
fs	3	385.4079167	128.4693056	13.87	<.0001

Notice that the SS for *gender* is the same as before. Also, the SS now shown for *fs* is the sum of the SS for *fat*, *smoke*, and *fat* \times *smoke* in the original model. The SS for the remaining interaction terms has now been incorporated into the error term. SSE has gone up, but MSE has actually gone down a little.

Different means for *gender*

	Mean	N	gender
A	18.983	12	1
B	13.558	12	2

(Well, we knew that since *gender* was significant)

Tukey comparisons for *fs*

	Mean	N	fs
A	22.900	6	11
B	16.000	6	12
B	13.117	6	22
B	13.067	6	21

Category *fs* = 1 is the low body fat and light smoking history group. The other three groups were not significantly different from each other.