

## Assumptions of the ANOVA F-test:

- Again, most assumptions involve the  $\varepsilon_{ij}$ 's (the error terms).
  - (1) The model is correctly specified.
  - (2) The  $\varepsilon_{ij}$ 's are normally distributed.
  - (3) The  $\varepsilon_{ij}$ 's have mean zero and a common variance,  $\sigma^2$ .
  - (4) The  $\varepsilon_{ij}$ 's are independent across observations.
- With multiple populations, detection of violations of these assumptions requires examining the residuals rather than the  $Y$ -values themselves.
- An estimate of  $\varepsilon_{ij}$  is:
- Hence the residual for data value  $Y_{ij}$  is:
- We can check for non-normality or outliers using residual plots (and normal Q-Q plots) from the computer.
- Checking the equal-variance assumption may be done with a formal test:  
 $H_0: \sigma_1^2 = \sigma_2^2 = \dots = \sigma_t^2$   
 $H_a: \text{at least two variances are not equal}$

- **The Levene test is a formal test for unequal variances that is robust to the normality assumption.**
- **It performs the ANOVA F-test on the absolute residuals from the sample data.**

**Example pictures:**

### **Remedies to Stabilize Variances**

- **If the variances appear unequal across populations, using transformed values of the response may remedy this. (Such transformations can also help with violations of the normality assumption.)**
- **The drawback is that interpretations of results may be less convenient.**

## **Suggested transformations:**

- **If the standard deviations of the groups increase proportionally with the group means, try:  $Y_{ij}^* = \log(Y_{ij})$**
- **If the variances of the groups increase proportionally with the group means, try:  $Y_{ij}^* = \sqrt{Y_{ij}}$**
- **If the responses are proportions (or percentages), try:  $Y_{ij}^* = \arcsin(\sqrt{Y_{ij}})$**
- **If none of these work, may need to use a nonparametric procedure (e.g., Kruskal-Wallis test).**

## **Making Specific Comparisons Among Means**

- **If our F-test rejects  $H_0$  and finds there are significant differences among the population means, we typically want more specific answers:**

**(1) Is the mean response at a specified level superior to (or different from) the mean response at other levels?**

**(2) Is there some natural grouping or separation among the factor level mean responses?**

- **Question (1) involves a “pre-planned” comparison and is tested using a contrast.**

- **Question (2) is a “post-hoc” comparison and is tested via a “Post-Hoc Multiple Comparisons” procedure.**

## Contrasts

- A contrast is a linear combination of the population means whose coefficients add up to zero.

Example ( $t = 4$ ):

- Often a contrast is used to test some meaningful question about the mean responses.

Example (Rice data): Is the mean of variety 4 different from the mean of the other three varieties?

We are testing:

What is the appropriate contrast?

Now we test:

We can estimate  $L$  by:

Under  $H_0$ , and with balanced data, the variance of a contrast

is:

- Also, when the data come from normal populations,  $\hat{L}$  is normally distributed.
- Replacing  $\sigma^2$  by its estimate MSW:

**For balanced data:**

- To test  $H_0: L = 0$ , we compare  $t^*$  to the appropriate critical value in the t-distribution with  $t(n - 1)$  d.f.
- Our software will perform these tests even if the data are unbalanced.

**Example:**

- **Note:** When testing multiple contrasts, the specified  $\alpha$  (=  $P\{\text{Type I error}\}$ ) applies to each test individually, not to the series of tests collectively.

## Post Hoc Multiple Comparisons

- When we specify a significance level  $\alpha$ , we want to limit  $P\{\text{Type I error}\}$ .
- What if we are doing many simultaneous tests?
- Example: We have  $\mu_1, \mu_2, \dots, \mu_t$ . We want to compare all pairs of population means.
- Comparisonwise error rate: The probability of a Type I error on each comparison.
- Experimentwise error rate: The probability that the simultaneous testing results in at least one Type I error.
- We only do post hoc multiple comparisons if the overall F-test indicates a difference among population means.
- If so, our question is: Exactly which means are different?
- We test:
- The Fisher LSD procedure performs a t-test for each pair of means (using a common estimate of  $\sigma^2$ , MSW).
- The Fisher LSD procedure declares  $\mu_i$  and  $\mu_j$  significantly different if:

- **Problem: Fisher LSD only controls the comparisonwise error rate.**
- **The experimentwise error rate may be much larger than our specified  $\alpha$ .**
  
- **Tukey's Procedure controls the experimentwise error rate to be only equal to  $\alpha$ .**
  
- **Tukey procedure declares  $\mu_i$  and  $\mu_j$  significantly different if:**
  
- **$q_\alpha(t, df)$  is a critical value based on the studentized range of sample means:**
  
- **Tukey critical values are listed in Table A.7.**
  
- **Note:  $q_\alpha(t, df)$  is larger than**
  
- **Tukey procedure will declare a significant difference between two means \_\_\_\_\_ often than Fisher LSD.**
  
- **Tukey procedure will have \_\_\_\_\_ experimentwise error rate, but Tukey will have \_\_\_\_\_ power than Fisher LSD.**
  
- **Tukey procedure is a \_\_\_\_\_ conservative test than Fisher LSD.**

## Some Specialized Multiple Comparison Procedures

- **Duncan multiple-range test**: An adjustment to Tukey's procedure that reduces its conservatism.
- **Dunnett's test**: For comparing several treatments to a "control".
- **Scheffe's procedure**: For testing "all possible contrasts" rather than just all possible pairs of means.

**Notes**: ● **When appropriate**, preplanned comparisons are considered superior to post hoc comparisons (more power).

- Tukey's procedure can produce **simultaneous CIs** for all pairwise differences in means.

**Example:**



## Random Effects Model

- Recall our ANOVA model:
- If the  $t$  levels of our factor are the only levels of interest to us, then  $\tau_1, \tau_2, \dots, \tau_t$  are called fixed effects.
- If the  $t$  levels represent a random selection from a large population of levels, then  $\tau_1, \tau_2, \dots, \tau_t$  are called random effects.

**Example:** From a population of teachers, we randomly select 6 teachers and observe the standardized test scores for their students. Is there significant variation in student test score among the population of teachers?

- If  $\tau_1, \tau_2, \dots, \tau_t$  are random variables, the F-test no longer tests:

**Instead, we test:**

**Question of interest:** Is there significant variation among the effects for the different levels in the population?

- For the one-way ANOVA, the test statistic is exactly the same,  $F^* = MSB / MSW$ , for the random effects model as for the fixed effects model.