Analysis of Variance

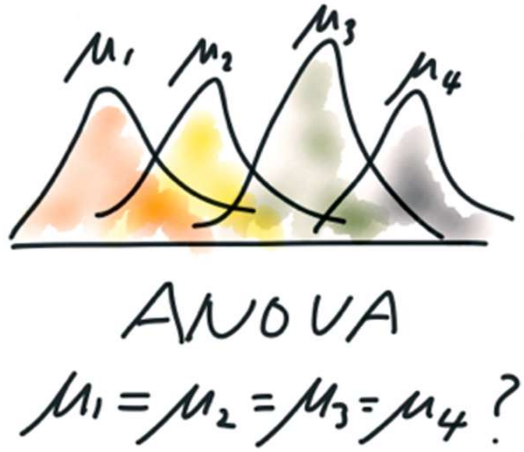

Comparing Multiple Groups: ANOVA

- Want to compare more than 2 independent groups?
- Use analysis of variance (ANOVA)
- Simultaneous examination: do groups differ from one another?
- Focus on one-way ANOVA

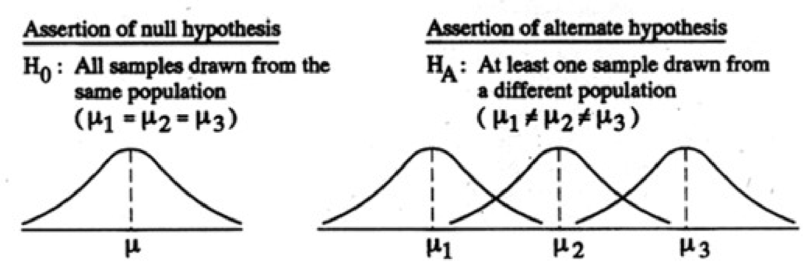

Testing Based on Variance Comparison

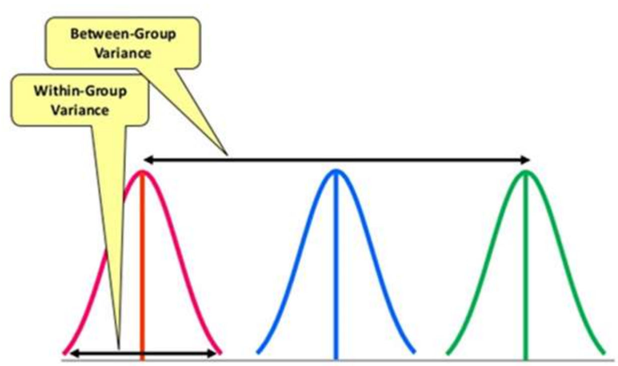

- http://rpsychologist.com/d3-one-way-anova
- http://digitalfirst.bfwpub.com/stats_applet/stats_applet_1_anova.html

Same with Overlapping Distributions

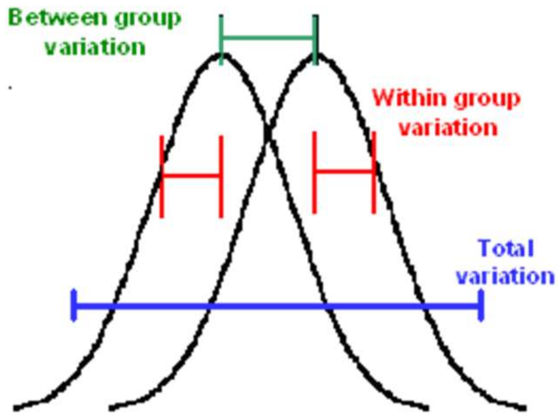

F Statistic

- Test statistic: F statistic
- If used for two groups → same result as t-test (for two groups, squaring the t value will give us F statistic)
- Called F-test in honor of Sir Ronald A. Fisher who developed the statistic in the 1920s

Analysis of Variance Table

Source	Sum of Squares (SS)	df	Mean SS	F
Between groups	$BSS=\Sigma(\bar{X}_{group} - \bar{X}_{grand})^2$	k-1	BSS/(k-1)	$\frac{BSS/(k-1)}{WSS/(N-k)}$
Within groups	$WSS=\Sigma(X - \bar{X}_{group})^2$	N-k	WSS/(N-k)	
Total	$TSS=\Sigma(X - \bar{X}_{grand})^2$	N-1	TSS/(N-1)	

- Total Mean SS = total variance of X
- BSS + WSS = TSS
- Applet: https://demonstrations.wolfram.com/VisualANOVA/

F-Distribution: Shape

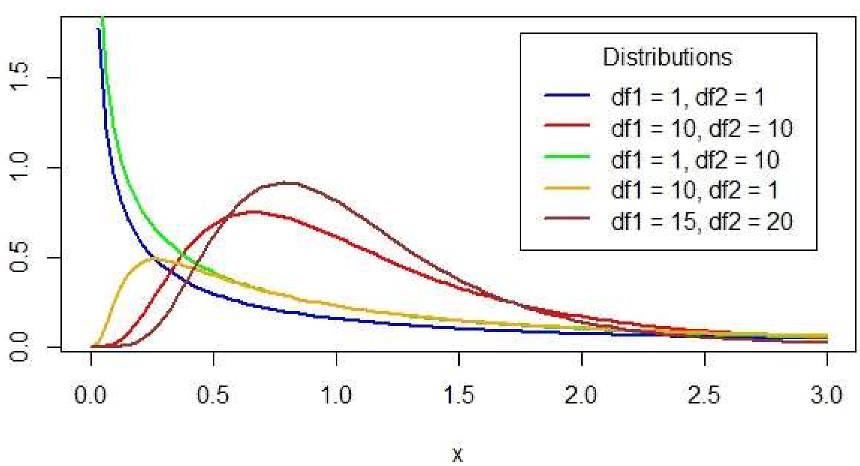

F-Distribution and Critical Region

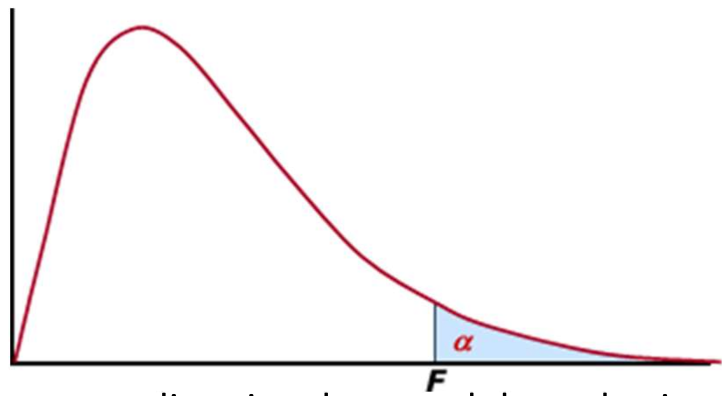

- Always non-directional research hypothesis
- Always one-tailed test!

ANOVA Step by Step

1. State your null and research hypotheses:

- Null: In the population, the means of all groups are the same.

  H0: $\mu_1 = \mu_2 = \mu_3$ [etc. if more groups]

- Research: In the population, the mean of at least one group differs from the others (always non-directional)

  H1: $\mu_1 \neq \mu_2 \neq \mu_3$ [etc. if more groups]

ANOVA Step by Step

2. Select the alpha level
3. Identify the test statistic: F statistic
4. Formula:

$$F = \frac{MS_{Between}}{MS_{Within}}$$

MS between = $\Sigma(\bar{X}_{group} - \bar{X}_{grand})^2 / (k-1)$

MS within = $\Sigma(X - \bar{X}_{group})^2 / (N-k)$

df1=k-1, df2=N-k

ANOVA Step by Step

5. Use Table B3 to find the critical value of F: based on three pieces of information: (1) df for numerator = k-1 (where k=number of groups); (2) df for denominator = N-k, (3) our selected alpha level.
6. Compare the computed and the critical value.
7. State your decision about H0: If your computed value is larger than the critical value → reject H0 in favor of H1. If your computed value is smaller than the critical value → fail to reject H0.

Example

Suppose three large groups of students were taught statistics using different pedagogical methods. We randomly sampled 4 students from each group and have the following scores on achievement test:

- Method 1: 71, 75, 65, 69
- Method 2: 90, 80, 86, 84
- Method 3: 72, 77, 76, 79

We want to know whether these three methods produce significantly different achievement results or not.

We want to use 99% confidence level.

Example: Step by Step

1. Our null hypothesis is that the three methods produce the same results; our research hypothesis is that at least one of the methods produces results distinct from others.

- H0: $\mu_1 = \mu_2 = \mu_3$ H1: $\mu_1 \neq \mu_2 \neq \mu_3$

We use non-directional test (always for ANOVA)

2. We choose α= .01 (1-.99=.01).

3. We will use F statistic.

Calculating Means

- $\bar{X}_{group1}$ =(71+75+65+69)/4=70
- $\bar{X}_{group2}$=(90+80+86+84)/4=85
- $\bar{X}_{group3}$=(72+77+76+79)/4=76
- $\bar{X}_{grand}$ =(70*4+85*4+76*4)/12=77

Calculations

Group	X	$\bar{X}_{group}$	$\bar{X}_{grand}$	X-$\bar{X}_{group}$	(X-$\bar{X}_{group})^2$	X-$\bar{X}_{grand}$	(X-$\bar{X}_{grand})^2$	$\bar{X}_{group}$ -$\bar{X}_{grand}$	($\bar{X}_{group}$-$\bar{X}_{grand})^2$
1	71	70	77	1	1	-6	36	-7	49
1	75	70	77	5	25	-2	4	-7	49
1	65	70	77	-5	25	-12	144	-7	49
1	69	70	77	-1	1	-8	64	-7	49
2	90	85	77	5	25	13	169	8	64
2	80	85	77	-5	25	3	9	8	64
2	86	85	77	1	1	9	81	8	64
2	84	85	77	-1	1	7	49	8	64
3	72	76	77	-4	16	-5	25	-1	1
3	77	76	77	1	1	0	0	-1	1
3	76	76	77	0	0	-1	1	-1	1
3	79	76	77	3	9	2	4	-1	1
Σ				0	**130**	**0**	**586**	0	**456**

ANOVA Table

- $df_{between}$ = k-1 = 3-1 = 2
- df_{within}= 12-3 = 9
- df_{total}= 12-1 = 11

Source	SS	df	Mean SS	F
Between groups	456	2	228	15.8
Within groups	130	9	14.44	
Total	586	11		

- Check if math is ok: 456 + 130 = 586; 2+9=11

Example Step by Step

4. Computed test statistic F = 15.8

5. Critical value: Use table B3: df for numerator=2, df for denominator =9 → critical value = 8.02

6. Test statistic > critical value (15.8>8.02)

7. We reject the null hypothesis and conclude that there are statistically significant differences among these means.

Conclusion: at least one of these three methods of teaching produces significantly different results from the others.

- We can also report our finding as F(2, 9) = 15.8, p<.01

Post-Hoc Comparisons (follow-up pairwise tests)

- Interested in pinpointing which groups specifically differ from each other?
- Can make comparisons using two group methods (t-tests for independent samples)
- This is called post-hoc analyses

Inflated Alpha

- Making so many comparisons → risk of Type I error (alpha) is inflated
- E.g., 3 groups → could compare group 1 to group 2, group 2 to group 3, and group 1 to group 3 → 3 comparisons in total
- Need to adjust our test to the total number of comparisons

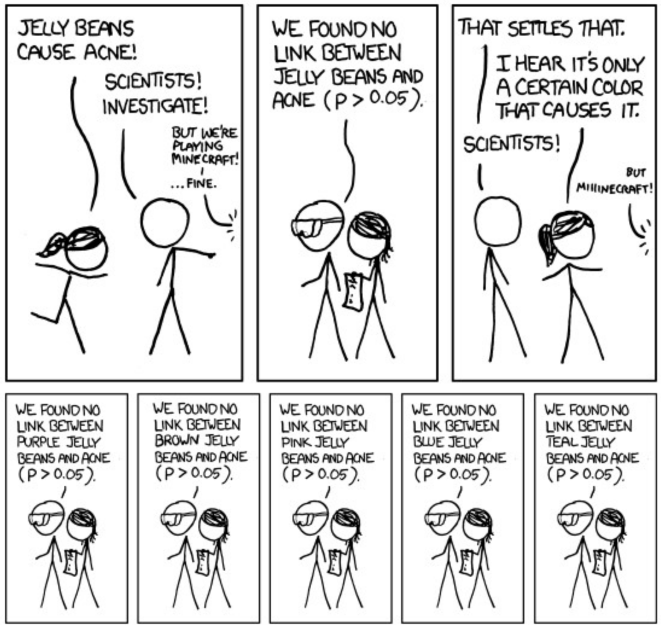

P-hacking Continues

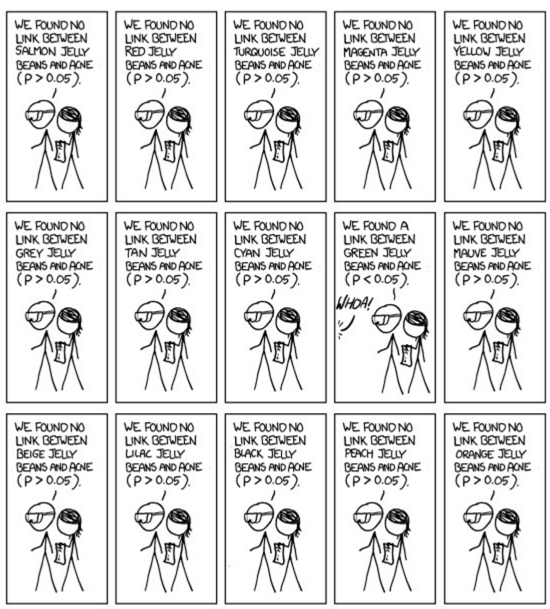

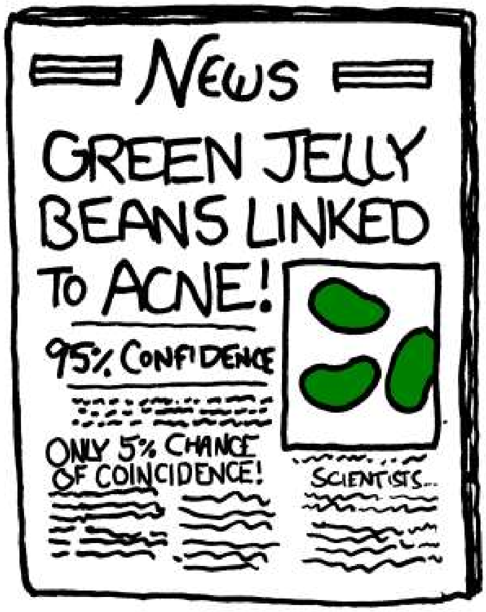

Post-Hoc Comparisons: Bonferroni Correction

- Use an adjustment called Bonferroni correction
- To adjust, we have two choices:
  1. Multiply p-values by the number of comparisons (Stata does that automatically)
  2. Divide alpha by the number of comparisons
- So either p should be larger or alpha should be smaller → more difficult to find that $p < \alpha$
- To calculate the number of comparisons for any number of groups: $k*(k-1)/2$
- 4 groups: $k = 4$ → $4*(4-1)/2 = 6$ comparisons

Age at First Childbirth and Social Class

- Does the average age when people have their first child differ by social class? (Want 99% confidence)
- H0: There are no social class differences in average age when people have their first child.
- H1: Average age when people have their first child varies by social class.
- Four social class groups, therefore:

H0: $\mu_1=\mu_2=\mu_3=\mu_4$ H1: $\mu_1\neq\mu_2\neq\mu_3\neq\mu_4$

ANOVA in Stata

```
. oneway agekdbrn class, means standard obs bonferroni

SUBJECTIVE |
     CLASS |  Summary of R'S AGE WHEN 1ST CHILD
IDENTIFICAT |              BORN
       ION |        Mean   Std. Dev.       Obs.
------------+-----------------------------------
  LOWER CLA |   21.798742   4.4718244        159
  WORKING C |   23.304207   5.2547438        618
  MIDDLE CL |   25.328836    5.822962        593
  UPPER CLA |   26.291667   6.1642702         48
------------+-----------------------------------
      Total |   24.083216   5.5985544       1418
                        Analysis of Variance
    Source              SS         df      MS            F     Prob > F
------------------------------------------------------------------------
Between groups      2359.01816      3   786.339387     26.44     0.0000
 Within groups      42055.1624   1414   29.7419819
------------------------------------------------------------------------
    Total           44414.1805   1417   31.3438112

Bartlett's test for equal variances:  chi2(3) =  19.3962  Prob>chi2 = 0.000
```

Posthoc Comparisons Portion (Bonferroni)

```
Comparison of R'S AGE WHEN 1ST CHILD BORN by SUBJECTIVE
CLASS IDENTIFICATION
                                    (Bonferroni)
Row Mean-|
Col Mean |    LOWER CL    WORKING     MIDDLE C
---------+--------------------------------
WORKING  |     1.50546
         |       0.012
         |
MIDDLE C |     3.53009    2.02463
         |       0.000      0.000
         |
UPPER CL |     4.49292    2.98746     .96283
         |       0.000      0.002      1.000
```

Conclusions

- We reject the null hypothesis of no social class differences in average age at first birth
- We are 99% confident that there are differences by class
- 4 out of 6 pairs of groups are significantly different ($p<.01$) from each other: higher class status → later childbirth
- But two exceptions:
  - lower and working class people seem to have their first children at similar ages
  - middle class and upper class people seem to have their first children at similar ages

Bivariate Bar Graphs for 3+ Groups

```
graph bar agekdbrn, over(class) blabel(bar) ytitle("Mean Age at First Childbirth")
```