One-way ANOVA

Comparing several means

Inference for the mean

- We know how to perform inference for 1 or two samples of quantitative data where μ is our parameter of interest.
  - Z vs. T inference

- Now we will compare one variable over more than two samples using a one-way Analysis of Variance (ANOVA) Test

Comparing several means

- The first step is to see if there is an overall difference between the groups:

  H_0: **All means μ_i are equal** $\Leftrightarrow \mu_1 = \mu_2 = \ldots = \mu_k$

  H_a: **Not all means μ_i are equal**

- This will require a new distribution, the F distribution

- If we find a difference, we will follow up to see which group (or groups) are different

Means of groups visually

If we have sampled from the same (or similar enough) population then the variance of the sample means should be close to the average sample variance

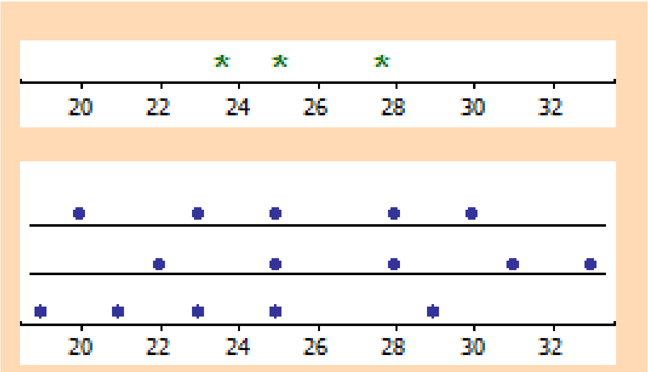

Equal Population means

If we have sampled from different populations then the variance of the sample means > the average sample variance

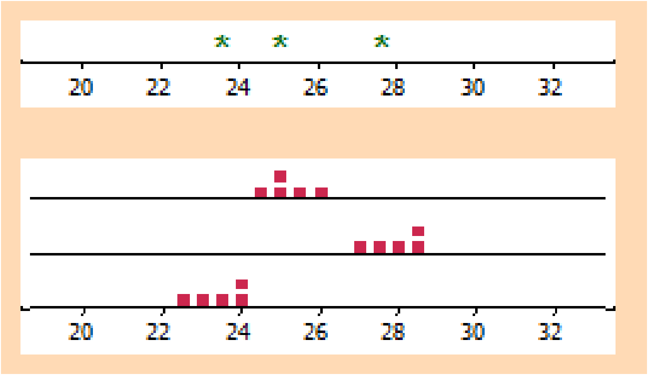

Unequal Population means

ANOVA Assumptions

Similar to Regression assumptions

1. Each sample must be independent SRSs.

2. Each population represented by the k samples must be Normally distributed. However, the test is robust to deviations from Normality for large-enough samples

3. The ANOVA F-test requires that **all k populations have the same standard deviation σ.**
   - There are inference tests for this, but they tend to be sensitive to deviations from the Normality assumption or require equal sample sizes.
   - Equal sample sizes make the ANOVA more robust to deviations from the equal σ rule.
   - **A simple and conservative approach**: The ANOVA F test is appropriate when the largest sample standard deviation is no more than about twice as large as the smallest sample standard deviation.

Analysis of Variance (ANOVA) tables

- The basic format of a one way ANOVA table is as follows:

Source	D.o.F	SS	MS	F	P-Val
Groups	DF_g	SSG	MSG	Test Statistic	?
Error	DF_E	SSE	MSE	–	–
Total	DF_T	SST	–	–	–

- Ultimately our goal is to find our F test Statistic and P–val to make our decision

Analysis of Variance (ANOVA) tables

- From there we begin filling out our table. The easiest place to start is D.o.F:

Source	D.o.F	SS	MS	F	P-Val
Groups	DF_g	SSG	MSG	Test Statistic	?
Error	DF_E	SSE	MSE	–	–
Total	DF_T	SST	–	–	–

- D.o.F. for Groups:

$$DF_g = \#\text{ of Groups} - 1 = k - 1$$

- Next we need the total D.o.F.

$$DF_T = n - 1$$

- From there:

$$DF_E = DF_T - DF_g = (n-1) - (k-1) = n - k$$

Analysis of Variance (ANOVA) tables

- We keep filling out our table moving right:

Source	D.o.F	SS	MS	F	P-Val
Groups	DF_g	SSG	MSG	Test Statistic	?
Error	DF_E	SSE	MSE	–	–
Total	DF_T	SST	–	–	–

- The next step deals with calculating the sum of squares:
  - SSG $= \sum_{i=1}^{k} n_k(\bar{x}_k - \bar{x}_{Overall})$
  - SSE $= \sum_{i=1}^{k}(n_i - 1)s_i^2$
  - SST$=$ SSG $+$ SSE

- The calculation for these gets complicated, but if you have 2/3 SSs you can easily find the other.

Analysis of Variance (ANOVA) tables

- Once we have our SS column calculated we then scale by the D.o.F. to find the MS

Source	D.o.F	SS	MS	F	P-Val
Groups	DF_g	SSG	MSG	Test Statistic	?
Error	DF_E	SSE	MSE	–	–
Total	DF_T	SST	–	–	–

- $MSG = SSG/DF_g$

- $MSE = SSE/DF_E$

Analysis of Variance (ANOVA) tables

- Finally we find our F Test Statistic by looking at the ratio of the MS terms

Source	D.o.F	SS	MS	F	P-Val
Groups	DF_g	SSG	MSG	Test Statistic	?
Error	DF_E	SSE	MSE	–	–
Total	DF_T	SST	–	–	–

- F Test Stat:

$$F = \frac{\text{MSG}}{\text{MSE}}$$

- We then go to the F table w/ both D.o.F. to find a p-val and make a decision

The F Distribution

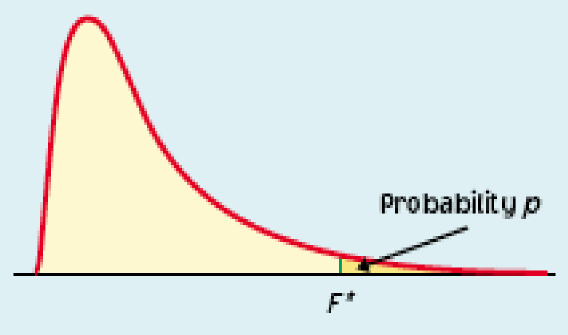

- The F distribution is asymmetrical and has two distinct degrees of freedom. One for the numerator (DF_g) and denominator (DF_E)
- In words the F test statistic looks like this:

$$F = \frac{\text{variation among the sample means}}{\text{variation among individuals in the samples}}$$

- If the variability of means is:
  - Smaller than variability within samples then F tends to be small.
  - Larger than variability within samples then F tends to be large.

ANOVA Table calculation Example

- Suppose we were comparing $k=6$ groups and there were $n=10$ observations

- Given:

ANOVA Table					
Source	D.o.F	SS	MS	F	p-value
Groups		0.17820			
Error				–	–
total		0.18665	–	–	–

- Fill out the rest of the ANOVA table:

F Table Example

- Full table should be:

ANOVA Table					
Source	D.o.F	SS	MS	F	p-value
Model	5	0.1782	0.03564	16.87100592	0.008595
Error	4	0.00845	0.0021125	–	–
total	9	0.18665	–	–	–

- From the Table:
  - Between 0.01 and 0.001

- In Excel:

0.008565	=F.DIST.RT(16.87,6-1,10-6)

Table F example

$df_{num} = k - 1$

$df_{den} = N - k$

TABLE F *F* distribution critical values

Degrees of freedom in the denominator	*p*	1	2	3	4	5	6	7	8
		Degrees of freedom in the numerator							
1	.100	39.86	49.50	53.59	55.83	57.24	58.20	58.91	59.44
	.050	161.45	199.50	215.71	224.58	230.16	233.99	236.77	238.88
	.025	647.79	799.50	864.16	899.58	921.85	937.11	948.22	956.66
	.010	4052.20	4999.50	5403.40	5624.60	5763.60	5859	5928.40	5981.10
	.001	405284.00	500000.00	540379.00	562500.00	576405.00	585937.00	592873.00	598144.00
2	.100	8.53	9.00	9.16	9.24	9.29	9.33	.35	9.37
	.050	18.51	19.00	19.16	19.25	19.30	19.33	19.35	19.37
	.025	38.51	39.00	39.17	39.25	39.30	39.33	39.36	39.37
	.010	98.50	99.00	99.17	99.25	99.30	99.33	99.36	99.37
	.001	998.50	999.00	999.17	999.25	999.30	999.33	999.36	999.37
3	.100	5.54	5.46	5.39	5.34	5.31	5.28	5.27	5.25
	.050	10.13	9.55	9.28	9.12	9.01	8.94	8.89	8.85
	.025	17.44	16.04	15.44	15.10	14.88	14.73	14.62	14.54
	.010	34.12	30.82	29.46	28.71	28.24	27.91	27.67	27.49
	.001	167.03	148.50	141.11	137.10	134.58	132.85	131.58	130.62
4	.100	4.54	4.32	4.19	4.11	4.05	4.01	3.98	3.95
	.050	7.71	6.94	6.59	6.39	6.26	6.16	6.09	6.04
	.025	12.22	10.65	9.98	9.60	9.36	9.20	9.07	8.98
	.010	21.20	18.00	16.69	15.98	15.52	15.21	14.98	14.80
	.001	74.14	61.25	56.18	53.44	51.71	50.53	49.66	49.00
5	.100	4.06	3.78	3.62	3.52	3.45	3.40	3.37	3.34
	.050	6.61	5.79	5.41	5.19	5.05	4.95	4.88	4.82
	.025	10.01	8.43	7.76	7.39	7.15	6.98	6.85	6.76
	.010	16.26	13.27	12.06	11.39	10.97	10.67	10.46	10.29
	.001	47.18	37.12	33.20	31.09	29.75	28.83	28.16	27.65
6	.100	3.78	3.46	3.29	3.18	3.11	3.05	3.01	2.98
	.050	5.99	5.14	4.76	4.53	4.39	4.28	4.21	4.15
	.025	8.81	7.26	6.60	6.23	5.99	5.82	5.70	5.60
	.010	13.75	10.92	9.78	9.15	8.75	8.47	8.26	8.10
	.001	35.51	27.00	23.70	21.92	20.80	20.03	19.46	19.03

ANOVA Example 1

- Male lab rats were randomly assigned to 3 groups differing in access to food for 40 days. The rats were weighed (in grams) at the end.
  - Group 1: access to chow only.
  - Group 2: chow plus access to cafeteria food restricted to 1 hour every day.
  - Group 3: chow plus extended access to cafeteria food (all day long every day).
- Does the type of food access influence the body weight of male lab rats significantly?

Chow	Restricted	Extended
516	546	564
547	599	611
546	612	625
564	627	644
577	629	660
570	638	679
582	645	687
594	635	688
597	660	706
599	676	714
606	695	738
606	687	733
624	693	744
623	713	780
641	726	794
655	736	
667		
690		
703		

ANOVA Example 1

- Hypotheses:

$$H_0: \mu_{chow} = \mu_{restricted} = \mu_{extended}$$

H_a: H_0 **is not true** *(at least one mean μ is different)*

- First we must examine the data to check our assumptions:
  - We can assume good sampling techniques
  - Look at Boxplots for deviations from normality, outliers, equal variances
  - Check summary Statistics

Descriptive Statistics: Chow, Restricted, Extended

Statistics

Variable	N	Mean	StDev
Chow	19	605.6	49.6
Restricted	16	657.3	50.7
Extended	15	691.1	63.4

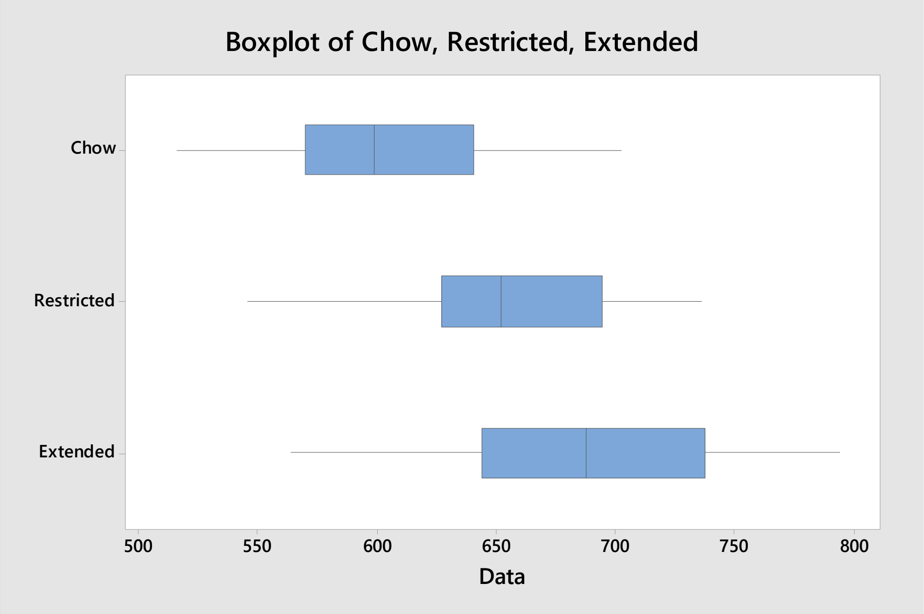

Minitab ANOVA Output

- Stat -> ANOVA -> One-Way
  - Default options are fine for our purpose:
- This means we Reject the Null.
- In our context, this means there are significantly different responses for different treatments.
- To see which ones, we must do a follow-up examination

One-way ANOVA: Chow, Restricted, Extended

Method

Null hypothesis	All means are equal
Alternative hypothesis	Not all means are equal
Significance level	α = 0.05

Equal variances were assumed for the analysis.

Factor Information

Factor	Levels	Values
Factor	3	Chow, Restricted, Extended

Analysis of Variance

Source	DF	Adj SS	Adj MS	F-Value	P-Value
Factor	2	63400	31700	10.71	0.000
Error	47	139174	2961		
Total	49	202573			

Model Summary

S	R-sq	R-sq(adj)	R-sq(pred)
54.4164	31.30%	28.37%	22.06%

Means

Factor	N	Mean	StDev	95% CI
Chow	19	605.6	49.6	(580.5, 630.7)
Restricted	16	657.3	50.7	(629.9, 684.7)
Extended	15	691.1	63.4	(662.9, 719.4)

Pooled StDev = 54.4164

ANOVA Follow-up

- It is clear that the Chow group is significantly different from extended, and most likely from Restricted.

- We can run follow up (pooled) t tests to find out for sure:

μ_1: mean of Restricted
μ_2: mean of Chow
Difference: $\mu_1 - \mu_2$

T-Value	DF	P-Value
3.03	31	0.005

μ_1: mean of Restricted
μ_2: mean of Extended
Difference: $\mu_1 - \mu_2$

T-Value	DF	P-Value
-1.63	26	0.114

Means

Factor	N	Mean	StDev	95% CI
Chow	19	605.6	49.6	(580.5, 630.7)
Restricted	16	657.3	50.7	(629.9, 684.7)
Extended	15	691.1	63.4	(662.9, 719.4)

Pooled StDev = 54.4164

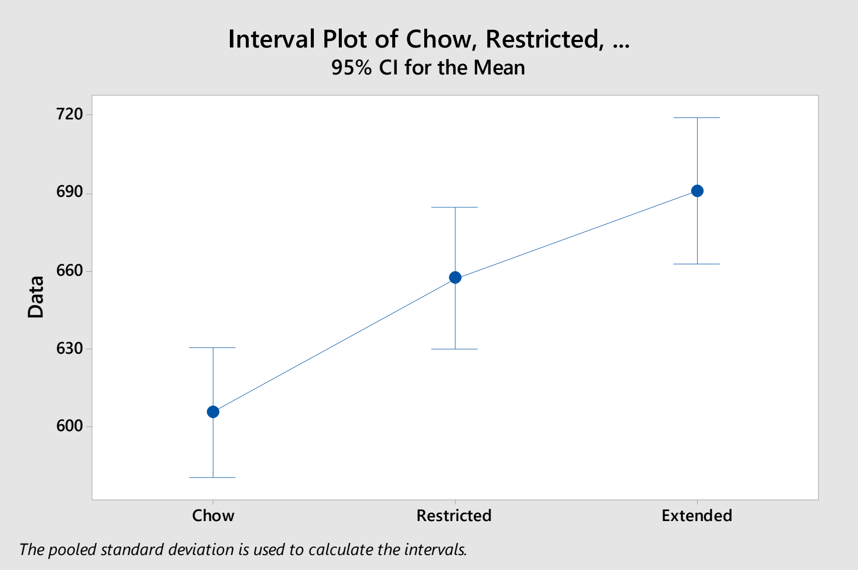

ANOVA Example 2

- The results of a study looking at the effect of treatments of brain tumors can be found in Canvas. We have 4 Groups:
  1. Irradiation
  2. Cannabinoids
  3. Irradiation & Cannabinoids
  4. Untreated

- Our hypotheses would be:

H_0: **Means are the Same or:** $\mu_I = \mu_C = \mu_{I\&R} = \mu_U$

H_a: H_0 **is not true** *(at least one mean μ is different)*

Examining the Data

- We can assume good sampling techniques
- Look at Boxplots for deviations from normality, outliers, equal variances
- Check summary Statistics

- We technically do not meet ANOVA assumptions here, however, there is clearly a difference

Statistics

Variable	N	Mean	StDev
Irradiation	4	54.3	28.2
Cannabis	5	26.00	16.63
Irradiation+cannabis	6	6.00	2.76
Untreated	7	48.29	24.80