Average weight of Eisenhower dollar: 23 grams

Average cost of dinner in Decatur: 23 dollars

Would it be more surprising to see

- A dinner that costs more than 27 dollars, or
- An Eisenhower dollar that weighs more than 27 grams?

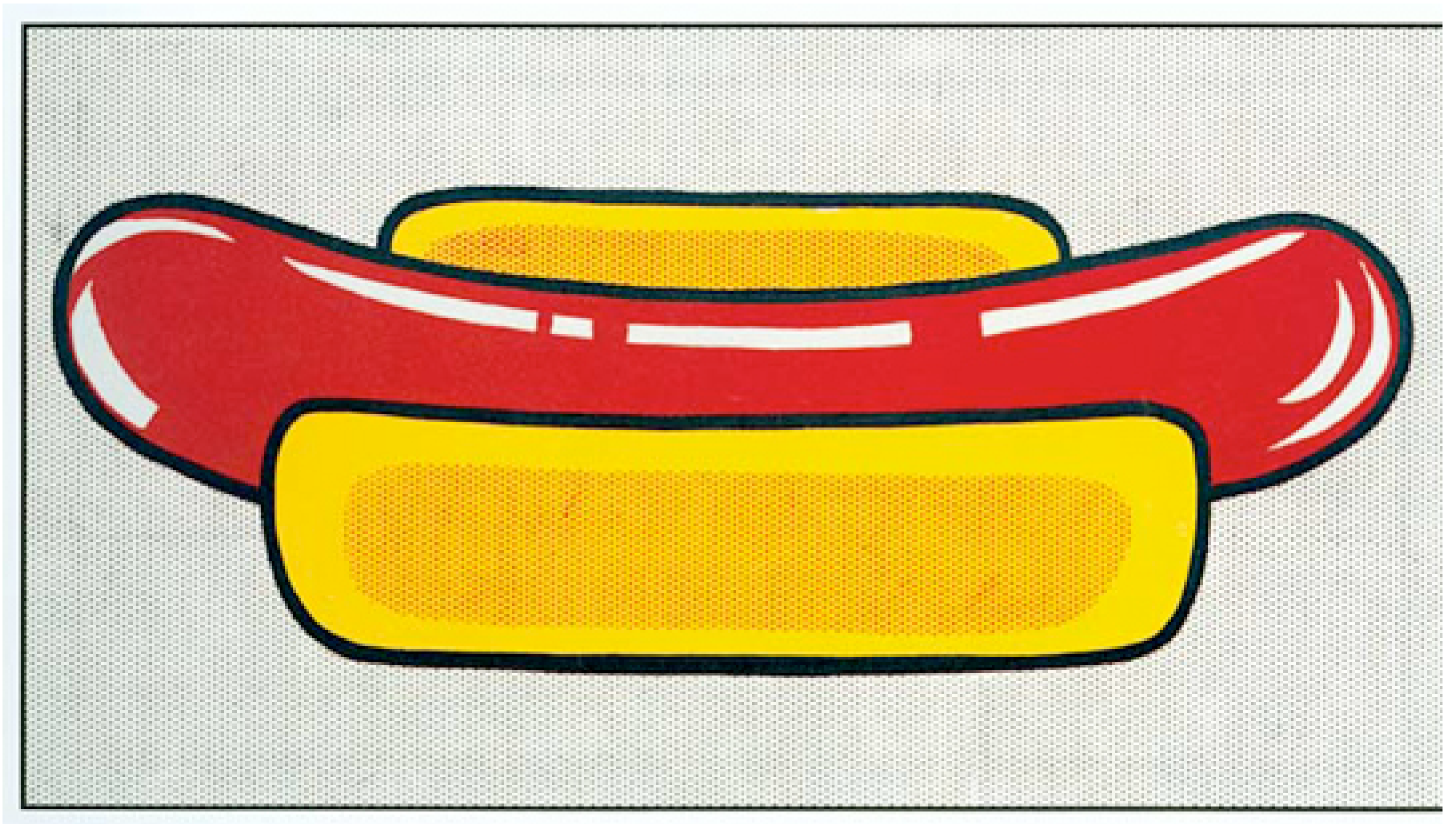

Description: Results of a laboratory analysis of calories of 54 major hot dog brands. Researchers for *Consumer Reports* analyzed three types of hot dog: beef, poultry, and meat (mostly pork and beef, but up to 15% poultry meat).

(*Consumer Reports*, June 1986, pp. 366-367)

BEEF: 

POULTRY: 

MEAT: 

Do different types of hot dogs have different numbers of calories?

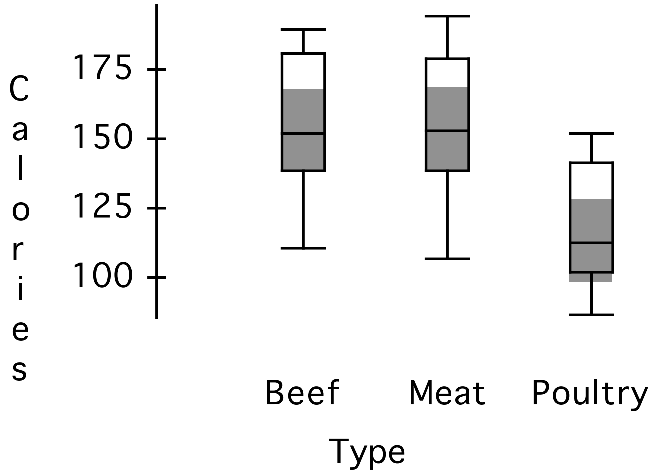

Group	Count	Mean	Median	Std Dev
Beef	20	156.85	152.5	22.642
Meat	17	158.706	153	25.236
Poultry	17	118.765	113	22.551

Multiple comparisons

(Comparing means, for example, among three or more groups)

Two steps:

1. Overall test: Is there good evidence of any difference among the parameters?

2. Follow-up analysis: Decide which parameters differ, and how large the differences are.

We'll talk about only overall test, and only one method: one-way ANOVA.

Analysis of Variance For Calories by Type

Source	df	Sums of Squares	Mean Square	F-ratio	Prob
Typ	2	17692.2	8846.10	16.074	≤ 0.0001
Error	51	28067.1	550.336		
Total	53	45759.3			

ANOVA F Test

H_0: all populations have the same mean
H_a: not all populations have the same mean

To test the null hypothesis, calculate the F statistic

$$\text{Roughly, } F = \frac{\text{variation among sample means}}{\text{variation among individuals in same sample}}$$

When H_0 is true, the F statistic has the $F(I-1, N-1)$ distribution.
When H_a is true, the F statistic tends to be large.
We reject H_0 if the F statistic is sufficiently large.

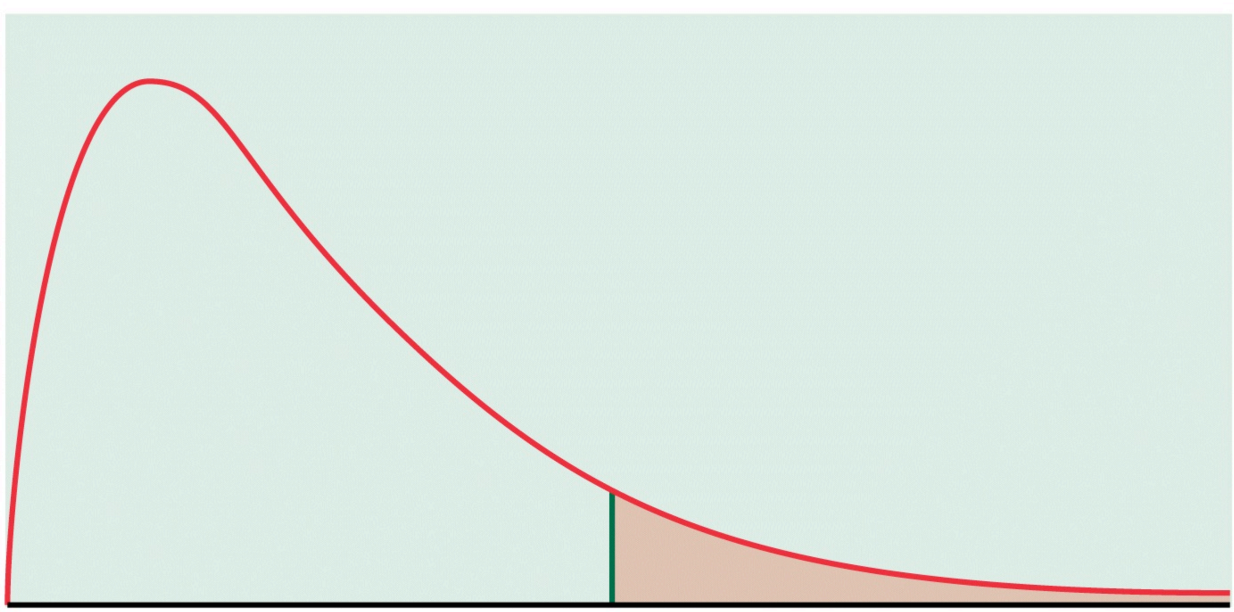

ANOVA Assumptions

1. We have I **independent SRS's**, one from each of the I populations

2. The *i*th population has a **Normal distribution** with unknown mean μ_i.

3. All of the populations have the **same standard deviation σ**, whose value is unknown.

ANOVA Assumptions

1. We have I **independent SRS's**, one from each of the I populations

- We're used to this. Garbage in, garbage out.

2. The *i*th population has a **Normal distribution** with unknown mean μ_i.

- We're used to this. Central Limit Theorem helps, but LOOK AT THE DATA.

3. All of the populations have the **same standard deviation σ**, whose value is unknown.

- ???

Rule of Thumb

The results of the ANOVA F test are approximately correct when the largest sample standard deviation is no more than twice as large as the smallest sample standard deviation.

Rule of Thumb

The results of the ANOVA F test are approximately correct when the largest sample standard deviation is no more than twice as large as the smallest sample standard deviation.

EXAMPLE:

Group	Count	Mean	Median	Std Dev
Beef	20	156.85	152.5	22.642
Meat	17	158.706	153	25.236
Poultry	17	118.765	113	22.551

(Also, try to make all samples about the same size, and no sample too small.)

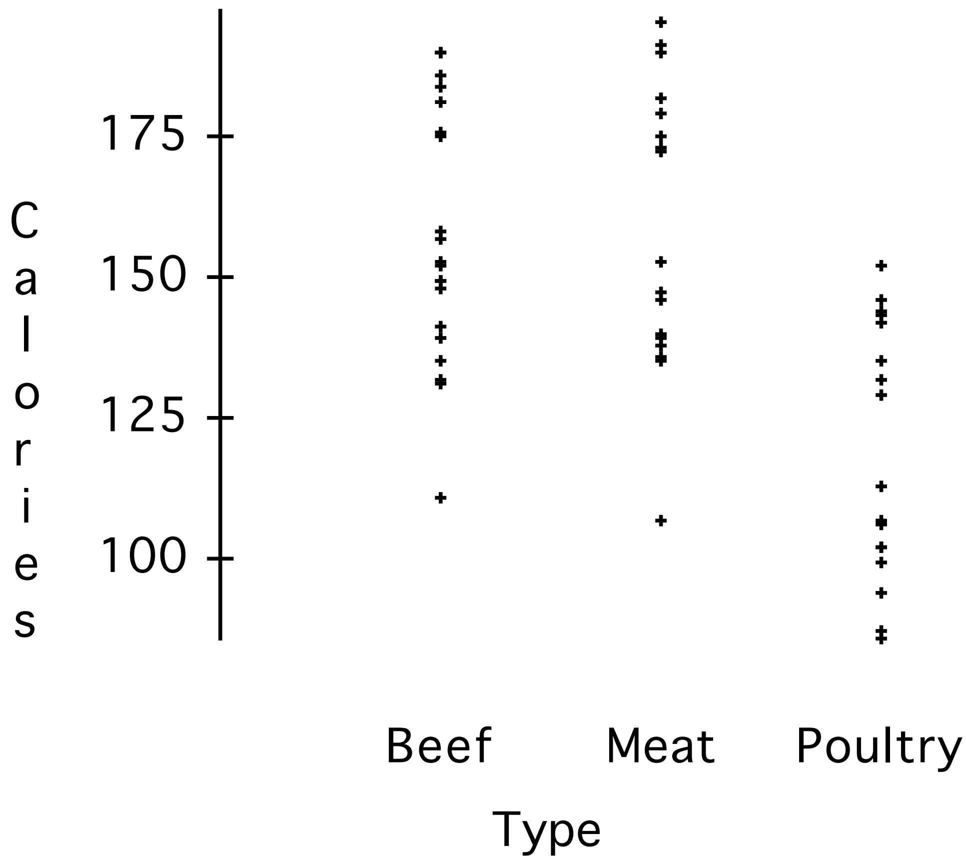

Look roughly normal.

What is the F statistic?

Roughly: $F = \dfrac{\text{variation among sample means}}{\text{variation among individuals in same sample}}$

Precisely: Assume that the *i*th population has the $N(\mu_I, \sigma)$ distribution, with sample size n_i, sample mean $\overline{x}_i$, and sample standard deviation s_i.

Then $F = \dfrac{MSG}{MSE}$,

where MSG (mean square group) is

$$MSG = \frac{n_1(\overline{x_1} - \overline{x})^2 + n_2(\overline{x_2} - \overline{x})^2 + \ldots + n_I(\overline{x_I} - \overline{x})^2}{I - 1}$$

and MSE (mean square error – "error" means "chance variation") is

$$MSE = \frac{(n_1 - 1)s_1^{\,2} + (n_2 - 1)s_2^{\,2} + \ldots + (n_I - 1)s_I^{\,2}}{N - I}$$

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

where MSG (mean square group) is

$$MSG = \frac{n_1(\overline{x_1} - \overline{x})^2 + n_2(\overline{x_2} - \overline{x})^2 + \ldots + n_I(\overline{x_I} - \overline{x})^2}{I - 1}$$

and MSE (mean square error – "error" means "chance variation) is

$$MSE = \frac{(n_1 - 1)s_1^{\,2} + (n_2 - 1)s_2^{\,2} + \ldots + (n_I - 1)s_I^{\,2}}{N - I}$$

Analysis of Variance For Calories by Type

Source	df	Sums of Squares	Mean Square	F-ratio	Prob
Typ	2	17692.2	8846.10	16.074	≤ 0.0001
Error	51	28067.1	550.336		
Total	53	45759.3			

Different tests for multiple comparisons

χ^2 test – compare proportions in different categories (categorical variables only)

one-way ANOVA F test – compare means (numerical variable) among categories

two-way ANOVA F test – same idea, but two KINDS of categories. (For example, Type of Meat and Kosher/Nonkosher.)

We won't talk about two-way (or three-way, etc.) ANOVA – similar, but more complicated.

Which test?

1. Do different breeds of dogs have different lifespans?

2. Does ethnicity have an effect on blood type?

3. Does gender affect SAT scores?

4. Do gender and hair color affect SAT scores?

5. Does the color of a team’s uniform affect its win-loss record?