Estimation:

variability

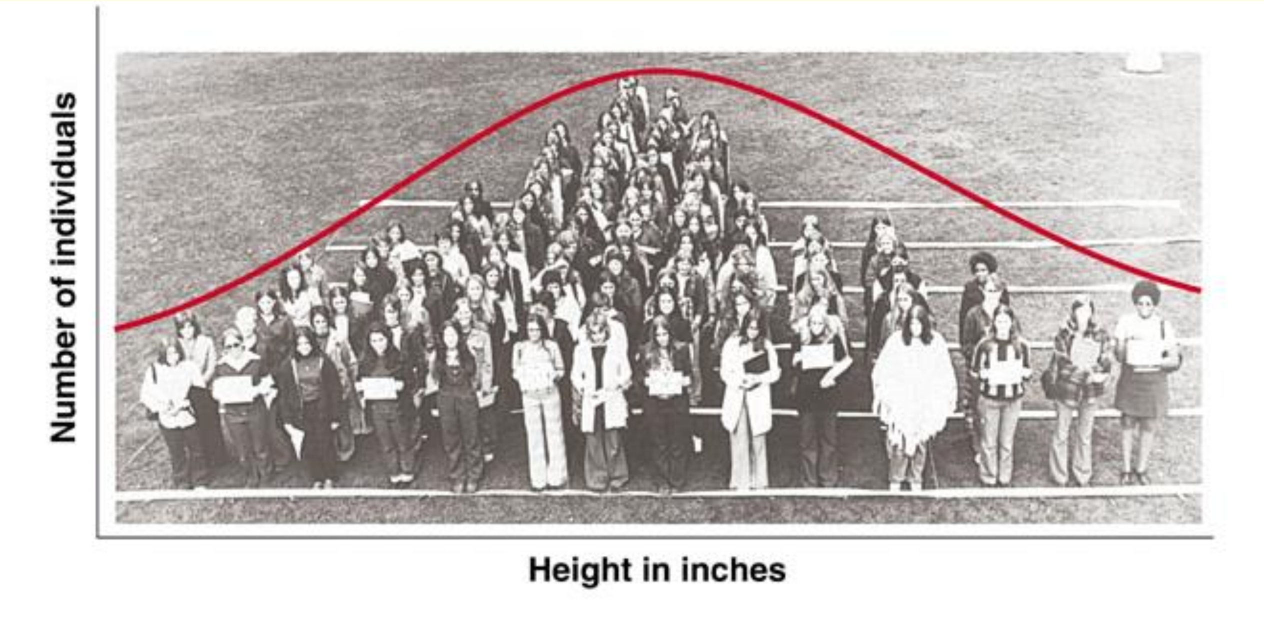

http://www.pelagicos.net/classes_biometry_fa16.htm

Next Step: Going Beyond the Data

A simple Example:

Roll two dice and add the value – 8 times

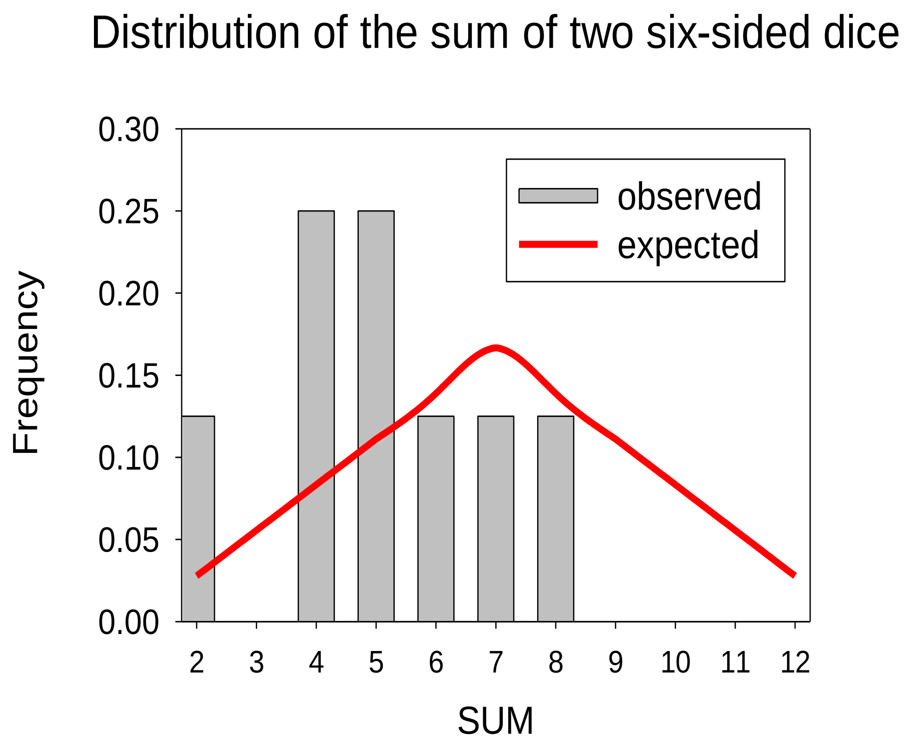

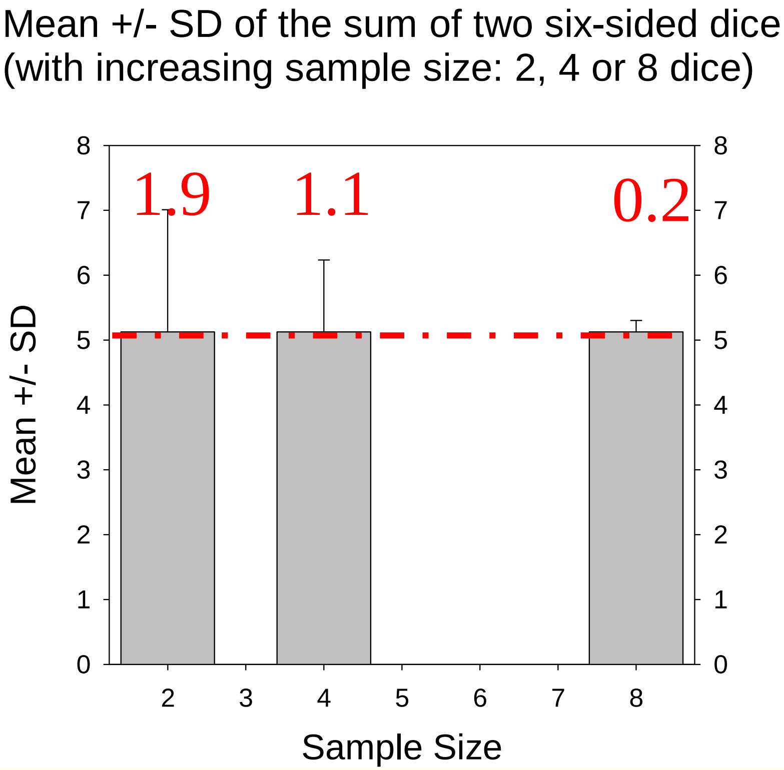

Mean = 5.1 +/- 1.9 (S.D.)

Next Step: Going Beyond the Data

Sampling allows us to guess about population parameters

However, different samples from the same population will differ… due to random variation (sampling)

Therefore, it is critical to assess how well any given sample represents the population

To do this, we use the Standard Error (S.E.)

S.E. : Standard Deviation / Sqrt (N)

$$S.E. = \frac{s}{\sqrt{N}}$$

Next Step: Going Beyond the Data

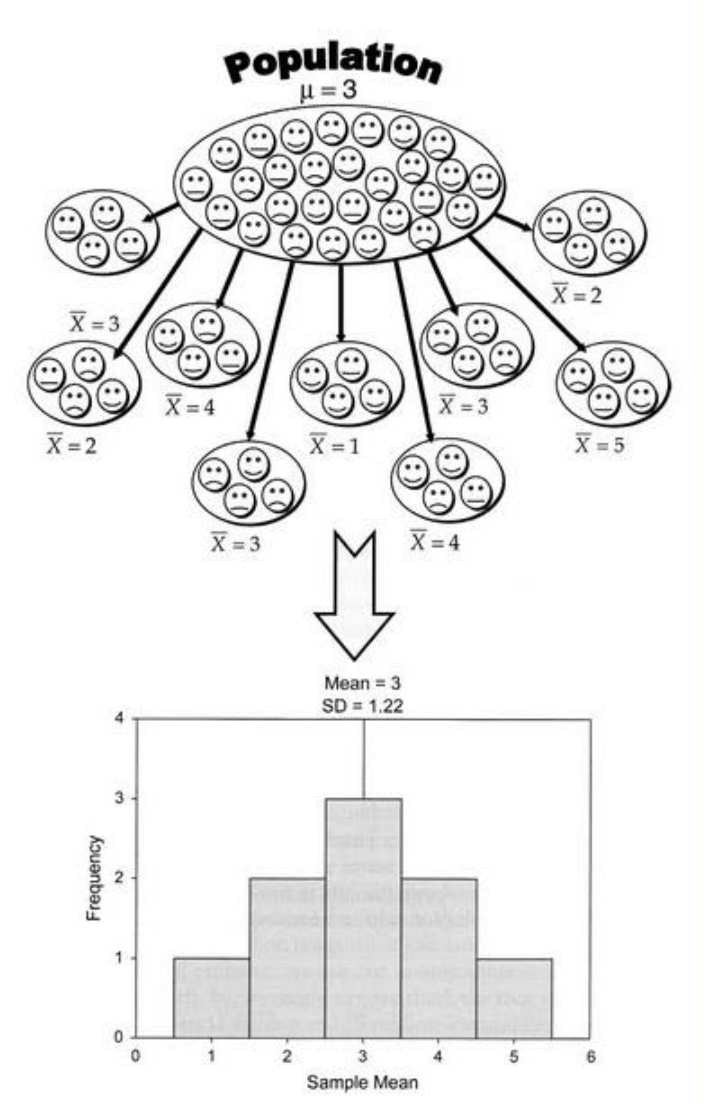

Resample the same population, by looking at small number of individuals at a time… 9 times

This illustrates sampling variation

We plot frequency distribution of the nine estimated means.

The mean of the means provides the best estimate of μ

But what about the S.D. of this distribution?

Next Step: Going Beyond the Data

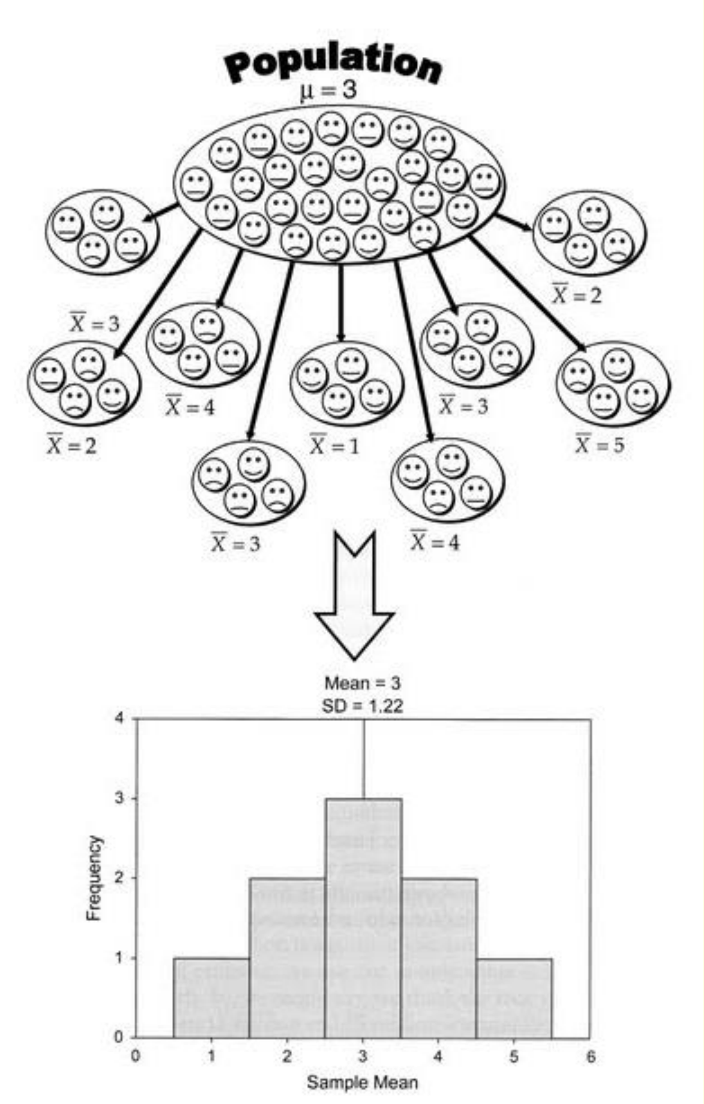

Basically, the S.E. is the S.D. of the means of the samples.

Rather than doing this exercise, statisticians have figured out that when $N \geq 30$, this distribution has the following properties:

mean = μ S.D. = S.E.

$$\text{S.E.} = \frac{s}{\sqrt{N}}$$

Confidence Intervals

Because different samples produce slightly different estimates, we can assess the accuracy of our estimates by calculating the boundaries within which we believe the true population parameter value lies.

The confidence interval (**CI**) indicates reliability of an estimate. It is the observed interval (i.e. calculated from observations), different from sample to sample, that frequently includes the parameter – if we repeat the experiment

How frequently the observed interval contains the parameter is determined by the **confidence level**.

Confidence Intervals

The standardized normal distribution (mean = 0 and a S.D. = 1) has same properties as any normal distribution:

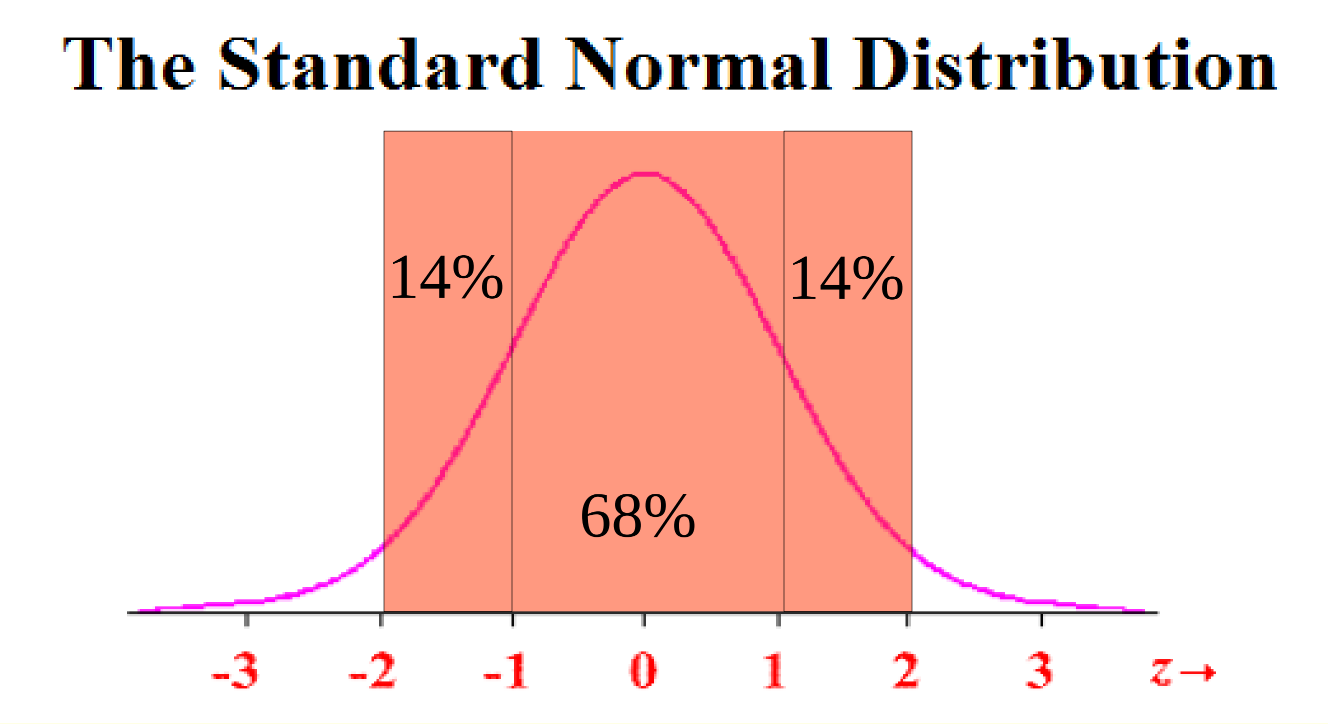

How did I get 14%?

96 - 68 = 28

28 / 2 = 14

Confidence Intervals

So, to contain 95% of the mass of the distribution, We must encompass from Z = -1.96 to Z = +1.96

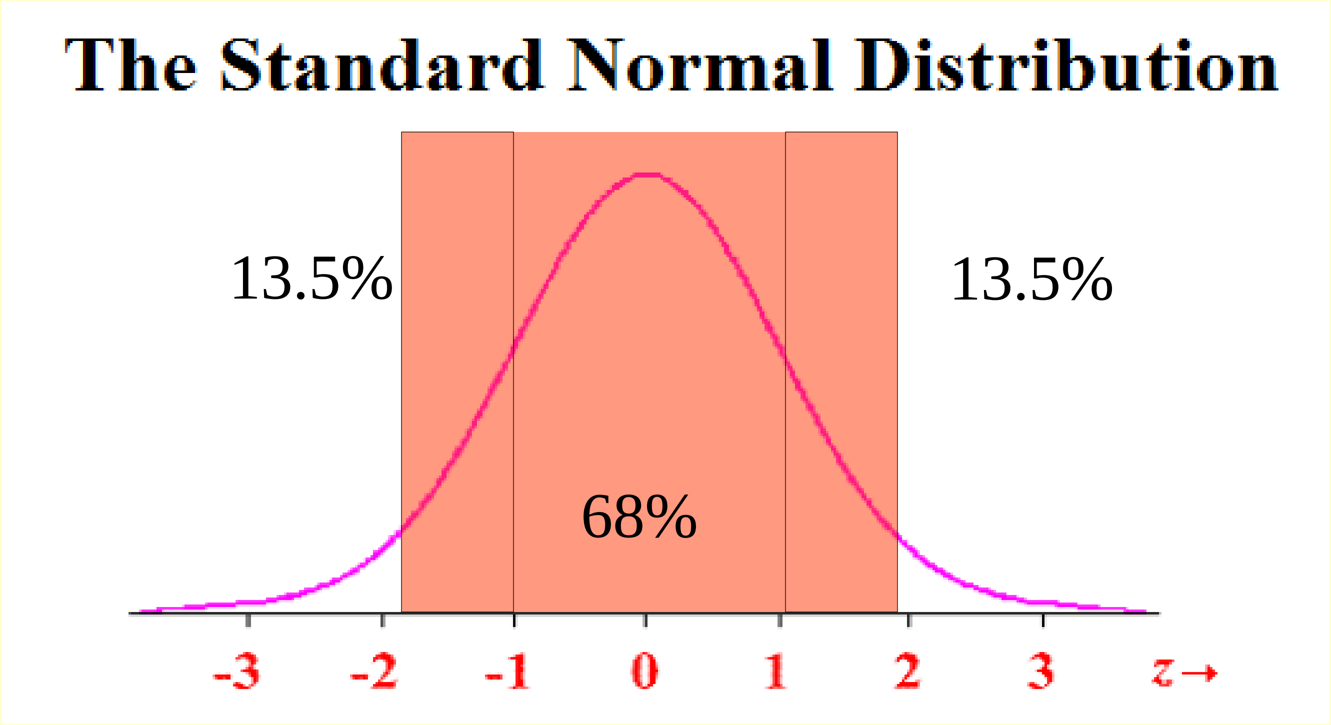

NOTE: Area under the curve is 95%

Z values from Table A1 in Field's Text Book

Confidence Intervals

A machine fills cups with margarine, and is adjusted so that the content of each cup is exactly 250 grams.

Yet, while the machine is supposed to fill every cup with exactly 250 grams, the content added to individual cups shows some variation, and can be considered a random variable X.

This variation is assumed to be normally distributed around the desired mean of 250 grams, with a given standard deviation of 2.5 grams.

Confidence Intervals

To determine if the machine is calibrated properly, a sample of $n = 25$ cups of margarine is chosen at random and the cups are weighed.

The resulting measured masses of margarine are $X_1, \ldots, X_{25}$, a random sample from X.

$$\hat{\mu} = \bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i. \qquad \bar{x} = \frac{1}{25}\sum_{i=1}^{25} x_i = 250.2 \text{ grams.}$$

Note: this is the result from one experiment

Confidence Intervals

If we do this experiment many times, each time we get a slightly different result.

We can determine the endpoints of the distribution of expected replicated experiment results by considering that the sample mean from a normally distributed sample is also normally distributed, with the same expectation μ, and a standard error of:

$$\frac{\sigma}{\sqrt{n}} = 0.5 \text{ grams}$$

Confidence Intervals

$$P(-z \leq Z \leq z) = 1 - \alpha = 0.95.$$

$$0.95 = 1 - \alpha = P(-z \leq Z \leq z) = P\left(-1.96 \leq \frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \leq 1.96\right).$$

$$= P\left(\bar{X} - 1.96\frac{\sigma}{\sqrt{n}} \leq \mu \leq \bar{X} + 1.96\frac{\sigma}{\sqrt{n}}\right)$$

$$Lower\ endpoint = \bar{X} - 1.96\frac{\sigma}{\sqrt{n}}, \qquad Upper\ endpoint = \bar{X} + 1.96\frac{\sigma}{\sqrt{n}}.$$

$$0.95 = P\left(\bar{X} - 1.96 \times 0.5 \leq \mu \leq \bar{X} + 1.96 \times 0.5\right)$$

$$= P\left(\bar{X} - 0.98 \leq \mu \leq \bar{X} + 0.98\right).$$

$$(\bar{x} - 0.98; \bar{x} + 0.98) = (250.2 - 0.98; 250.2 + 0.98) = (249.22; 251.18).$$

Confidence Intervals

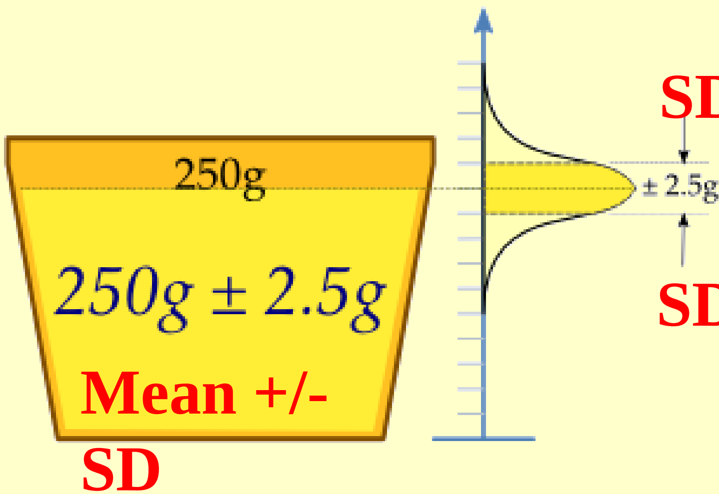

When we sample this distribution (n = 25 samples) and calculate the mean, its 95% C.I. lies between:

the lower endpoint: 249.22 g

and

the upper endpoint: 251.18 g.

Confidence intervals consist of a range of values (interval) that estimate an unknown population parameter.

The level of confidence of the confidence interval indicates the probability that the confidence range captures this true population parameter given a distribution of samples.

Point Estimates

Point Estimate: Involves the use of sampling data to calculate a single value, which serves as a "best guess" or "best estimate" of an unknown population parameter (for example: u, σ)

For example: $\bar{X}$ estimates μ

Yet, if we recalculate any point estimate many times, we get slightly different answers every time.

Why? Because of sampling (random noise in the data).

Confidence Intervals

Confidence Interval (CI) is the range of values (envelope) that includes the estimated parameter, with a probability determined by the **confidence level** (usually 95%, but not necessarily).

Formulation:

Lower: Mean – (Z score * SE)

Upper: Mean + (Z score * SE)

Margerine Cup Example:

(Repeat experiment 50 times)
n = 25 cups each time

Results of the First Experiment:

95% CI:

From first experiment:
249.22 to 251.18

Lower: 250.0 – (1.96 * 0.5)

Upper: 250.0 + (1.96 * 0.5)

Confidence Intervals

By standardizing, the distribution
(so it has a mean = 0 and a S.D. = 1)
we get a random variable (not dependent on μ):

Z Score:

$$Z = \frac{\bar{X} - \mu}{\sigma / \sqrt{n}} = \frac{\bar{X} - \mu}{0.5}$$

Thus, it is possible to find two values, independent of μ, between which Z lies with a given probability. This is the measure of how confident we want to be. We take $1 - \alpha = 0.95$.

Confidence Intervals

Developing Predictions with Built-in Margin of Error:

In Grand Forks, thousands of people, prepared for the flood by building sandbag dikes.

These dikes were based on a 49-foot estimate of flooding by the National Weather Service.

What went Wrong ?

Statistical Inference & Reliability

Confidence intervals built around point estimates to capture the variability inherent in the data and in the estimation method.

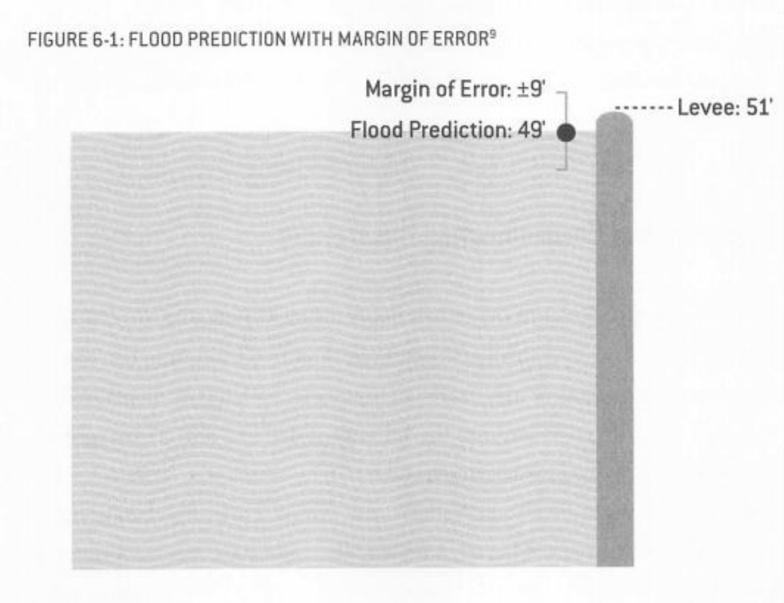

To show reliability of an estimate, confidence intervals are reported along with the point estimate: 49 +/- 9 C.I.

A point estimate is a single value of a population parameter.

A confidence interval specifies a range within which the parameter is estimated to lie (e.g., envelope of outcomes).

Confidence Intervals

The river crested at 54 feet in Grand Forks.

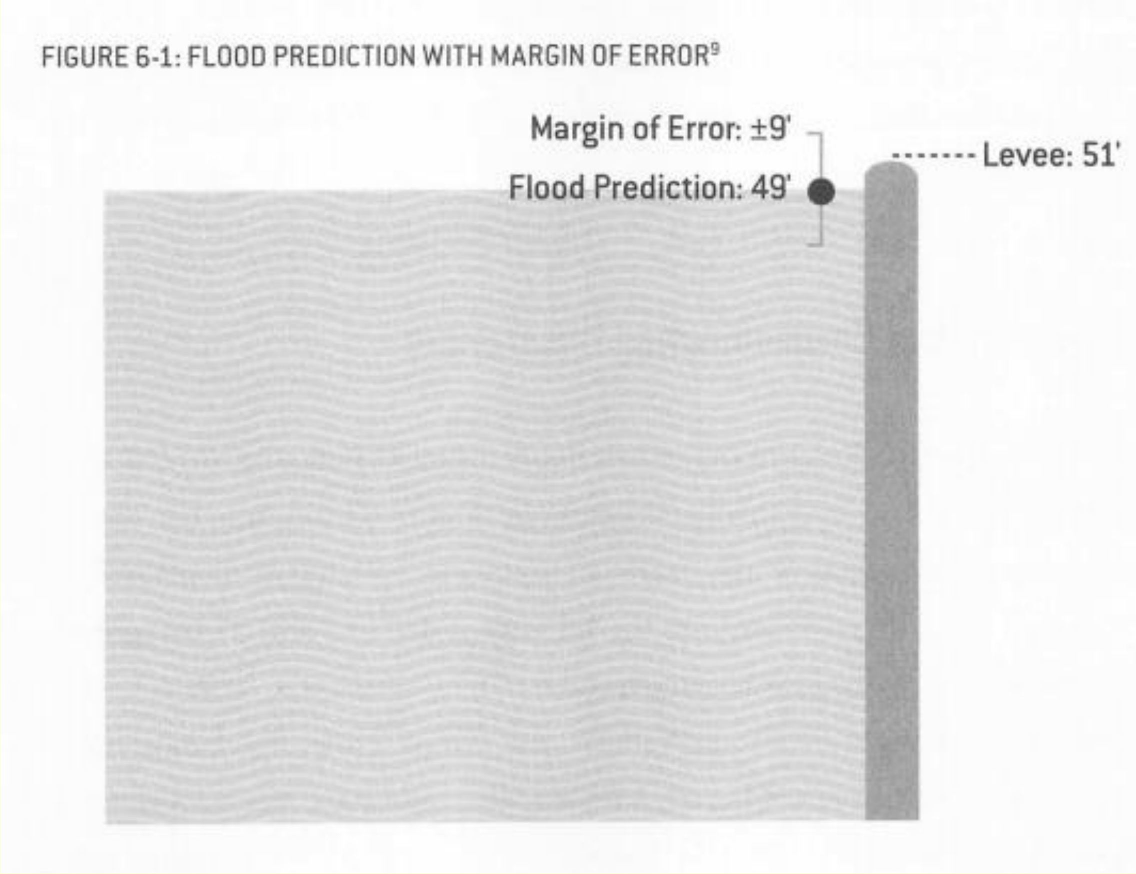

Confidence interval = +/- 9.0 (from mean)

UB = mean + (Z * SE)

58 = 49 + (1.96 * SE)

9 = 1.96 * SE

SE = 9 / 1.96 = 4.59

Z Score of 54 feet = (54 – 49) / (4.59) = 5 / 4.59 = 1.09

Remember: Confidence Intervals Decrease with Increasing Sample Size

What is my age: 37.95 years (SD = 2.45)

C.I. Formulation: Mean +/- (Z score * SE)

Mean +/- (1.96 * SE)

S.E. = S.D. / sqrt (n)

n	mean	std	sqrt(n)	se	ci
10	40.20	3.02	3.16	0.96	1.87
15	39.47	2.67	3.87	0.69	1.35
20	40.03	2.69	4.47	0.60	1.18

Confidence Intervals – Many Tests

Formulation = 95% confidence intervals

Lower bound: Mean – (1.96 * SE)
Upper bound: Mean + (1.96 * SE)

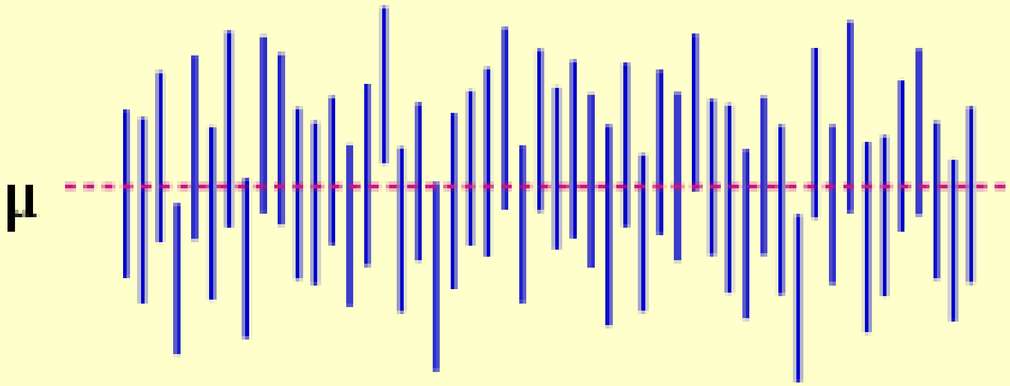

By definition: 95% of the confidence intervals (from different experiments) will overlap the real parameter μ

Interpreting Confidence Intervals

The (**CI**) is the interval that includes the estimated parameter, with a probability determined by **confidence level** (usually 95%).

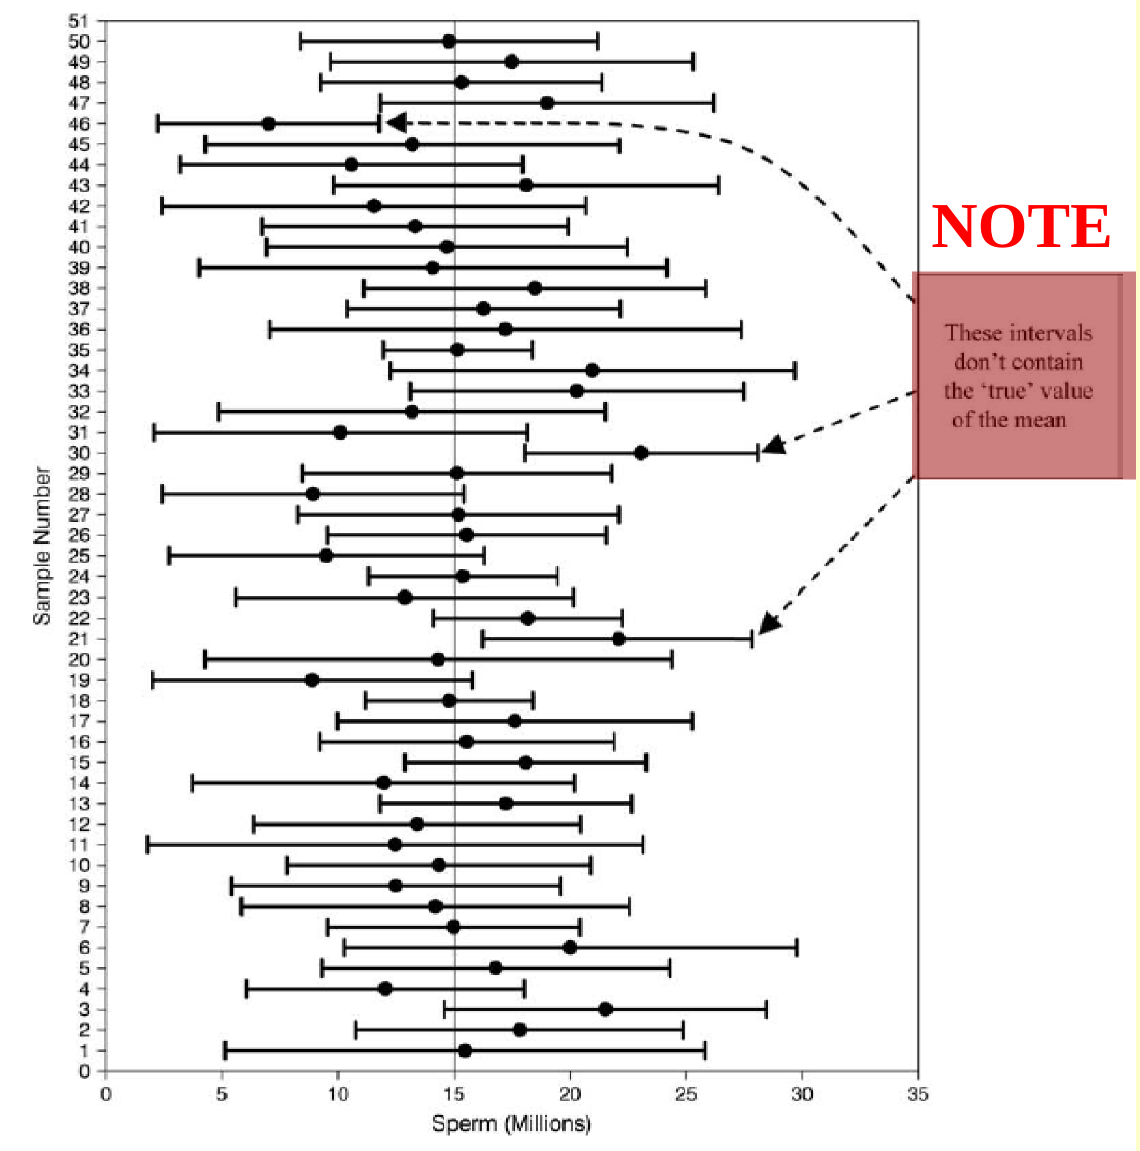

Interpreting Confidence Intervals

Case 1.
Two samples indistinguishable. They are from same population

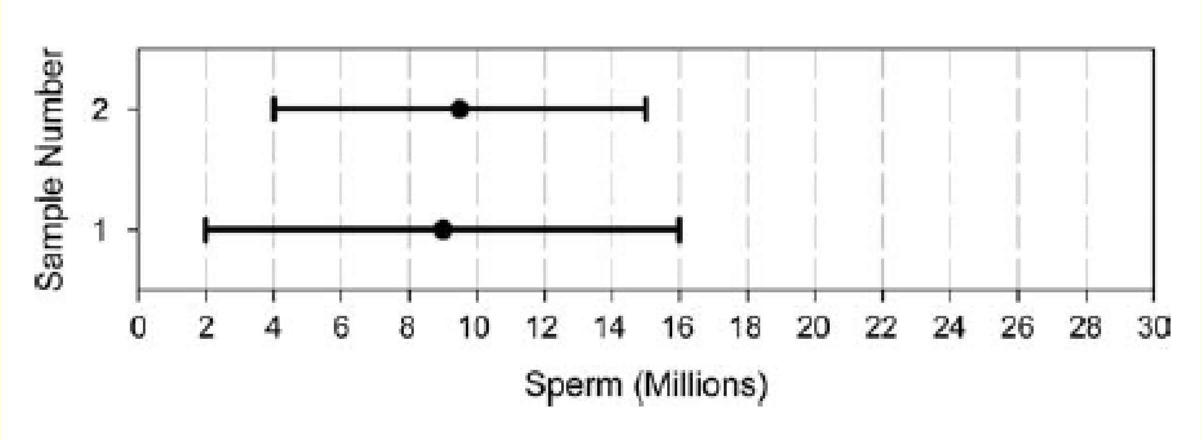

Case 2.
Two samples different. They are not from same population