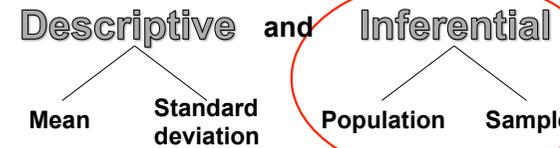


Distribution of sample means

Two types of statistics:



The relationship between Population and Samples is described in terms of Probability.

A jar with 100 marbles, 50 white and 50 black. You pick one marble. What is the probability that it is a black marble?

A jar with 100 marbles, 90 black and 10 white. You pick one marble. What is the probability that it is a black marble?

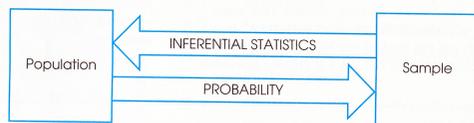
- Here a "jar with 100 marbles" is the population, and the "one marble" you pick is a sample.

In inferential statistics we do not know about the population, and must make decisions about the population on the basis of samples. So we make **inferences** about the population.

Suppose we have two jars with marbles, one has 50 black/50 white and the other has 90 black/10 white, but we do not know which is which. We pick 4 marbles.

There are two populations. Your sample is: 4 marbles.

Suppose all marbles are black; which jar (population) did we take the marbles out of?

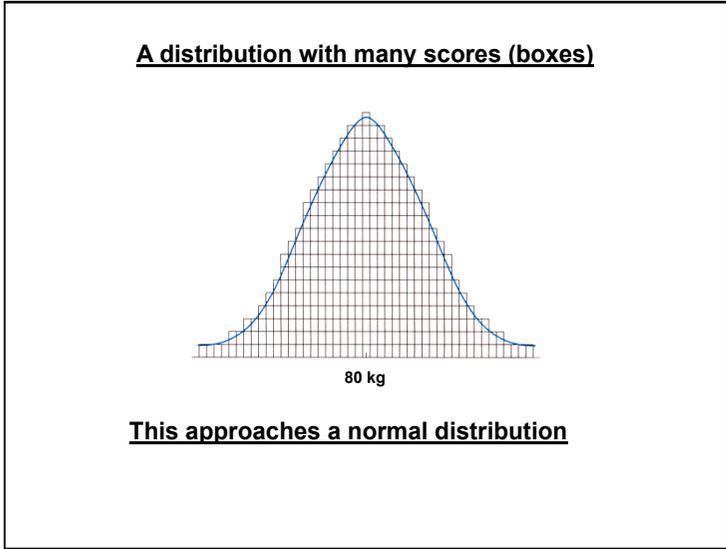
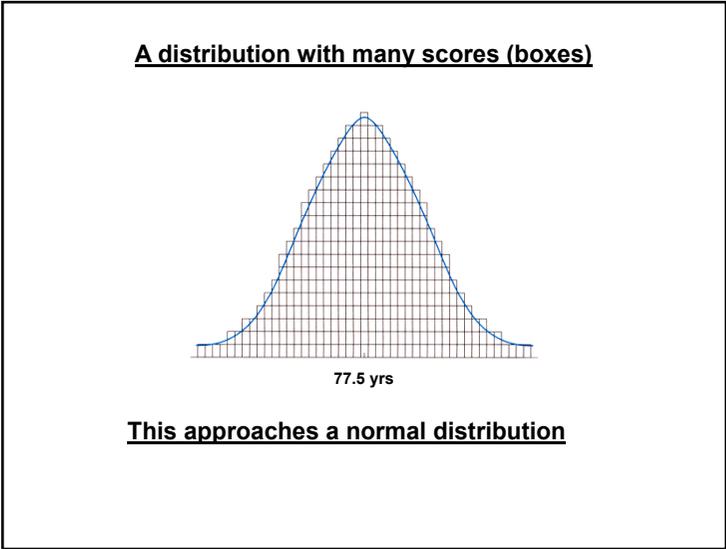
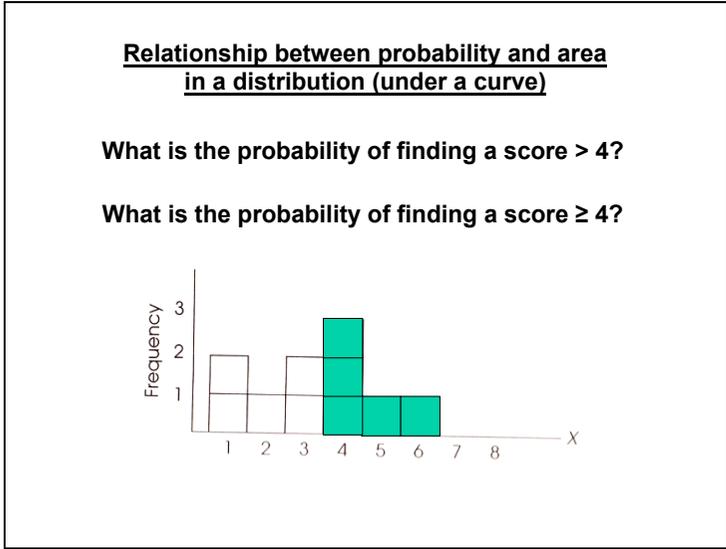
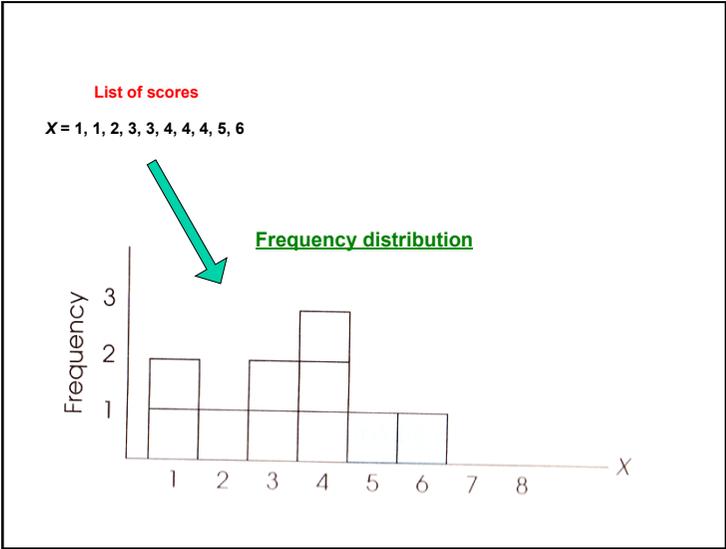


A word about List of scores, Frequency tables and Frequency distributions

List of scores
X = 1, 1, 2, 3, 3, 4, 4, 4, 5, 6



Frequency Table	
X	F
6	1
5	1
4	3
3	2
2	1
1	2

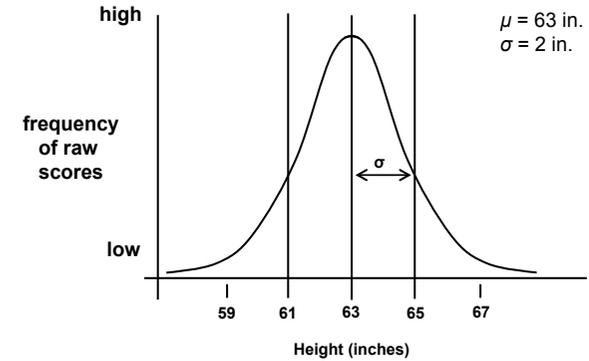


In **Inferential Statistics** we make inferences about **populations** on the basis of **samples**.

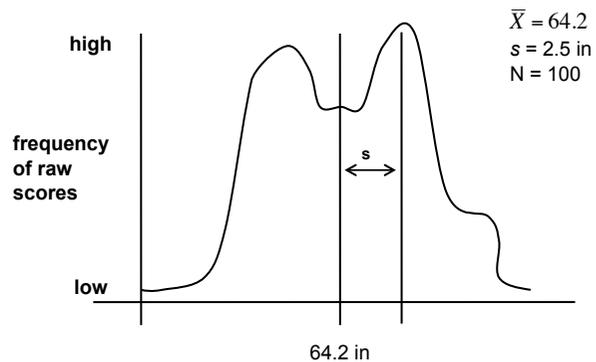
We therefore need to know what relationship exists between samples and populations.

A population of scores has:
mean (μ),
standard deviation (σ) and
shape (e.g., normal distribution)

EXAMPLE:
HEIGHT OF ALL ADULT WOMEN IN ENGLAND



Sample of 100 adult English women

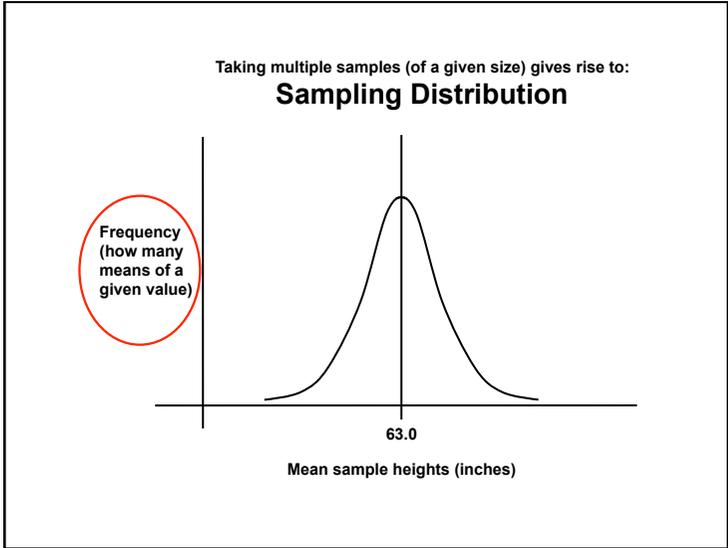


If we take repeated samples, each sample has a mean height (\bar{X}), a standard deviation (s), and a shape.

$\bar{X}_1 = 61$	$s_1 = 3.1$
$\bar{X}_2 = 69$	$s_2 = 2.9$
$\bar{X}_3 = 60.5$	$s_3 = 2.5$
.	.
.	.
.	.

Due to random fluctuations, each sample mean and SD is *different* from other samples, and from the parent population.

Fortunately, these differences are *predictable*. So we can use samples to make inferences about populations.



What are the properties of the distribution of Sample Means (Sampling Distribution)?

(a) The mean of the sample means is the *same* as the mean of the population of raw scores:

$$\mu_{\bar{X}} = \mu$$

(b) The sampling distribution of the mean also has a standard deviation. The standard deviation of the sample means is called the "standard error". It is **NOT** the same as the standard deviation of the population (of raw scores).

$$\sigma_{\bar{X}} \neq \sigma$$

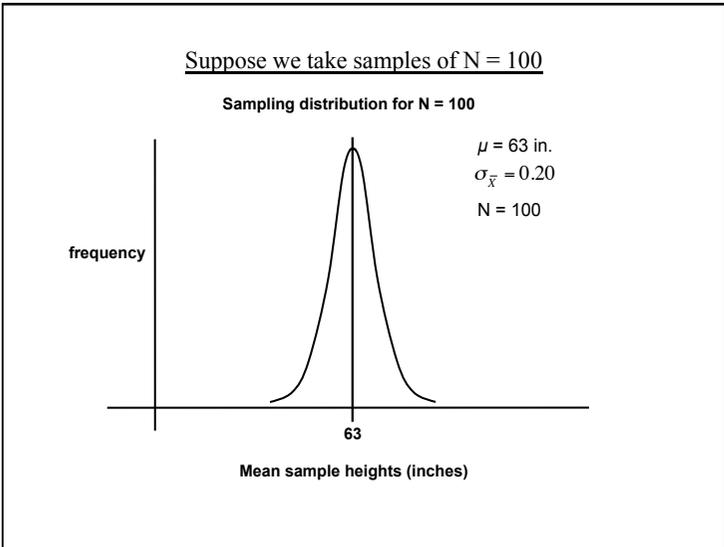
The standard error is smaller than the population SD:

$$\sigma_{\bar{X}} = \frac{\sigma}{\sqrt{N}}$$

The bigger the sample size, the *smaller* the standard error.

i.e., variation between sample means *decreases* as sample size *increases*.

(This is because sample mean based on a large sample reduces the influence of any extreme raw scores in a sample.)



How do we calculate the standard deviation of a sampling distribution (known as the standard error)? For $N = 100$

$$\sigma_{\bar{x}} = \frac{2}{\sqrt{100}} = \frac{2}{10} = 0.20$$

Suppose $N = 16$, instead of 100

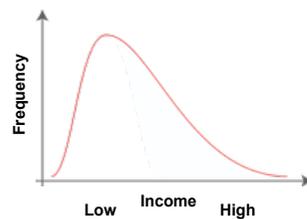
$$\sigma_{\bar{x}} = \frac{2}{\sqrt{16}} = \frac{2}{4} = 0.50$$

(c) -The distribution of sample means is normally distributed if the population of scores is normally distributed.

- And, for $N \geq 30$ the distribution of sample means is normally distributed *no matter what the shape of the original distribution of raw scores is!*

Example: Annual income of Americans.

Many people in the lower and medium income bracket; very few are ultra rich. (So distribution is NOT normal.)



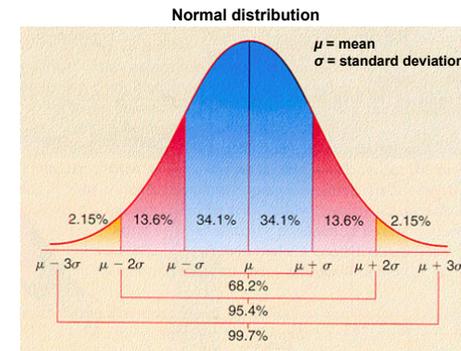
If we take many samples of size $N > 30$. The resulting sampling distribution of the mean will be normal.

This is due to the Central Limit Theorem.

Given the distribution is normal, we use properties of normal distribution to do interesting things.

Various proportions of scores fall within certain limits of the mean (i.e. 68% fall within +/- 1 SD on either side of the mean; 95% fall within +/- 2 SD, etc.).

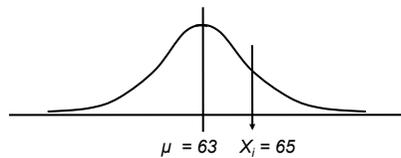
These proportions are best understood in terms of z-scores and the Unit Normal Table.



Calculating a z-score tells where in the distribution a particular mean lies. And from that you can decide how many means are above and below that particular mean.

Quick reminder about z-scores (using the height example)

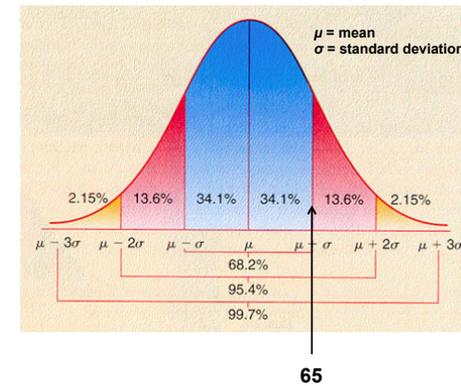
Calculating the distance of a given **score** from the population mean:



$\mu = 63$ in.
 $\sigma = 2$ in.

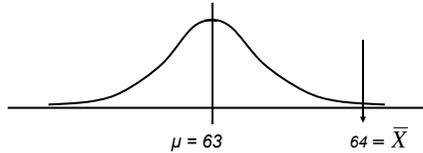
Calculating z-score

$$z_x = \frac{X_i - \mu}{\sigma} = \frac{65 - 63}{2} = \frac{2}{2} = +1$$



34.1% scores lie between 63 and 65

Calculating the distance of a given **sample mean** from the population mean:



We can do the same with sample means:

- (a) we obtain a particular sample mean;
- (b) we can represent this in terms of how *different* it is from the mean of its parent population.

$$z_{\bar{X}} = \frac{\bar{X} - \mu}{\frac{\sigma}{\sqrt{N}}} = \frac{64 - 63}{\frac{2}{\sqrt{100}}} = \frac{1}{\frac{2}{10}} = +5.00$$

If we obtain a sample mean that is *much* higher or lower than the population mean, there are two possible reasons:

- (a) our sample mean is a rare "fluke" (a quirk of sampling variation);
- (b) our sample has *not* come from the population we thought it did, but from some other, *different*, population.

The greater the difference between the sample and population means, the more plausible (b) becomes.

Take another example:

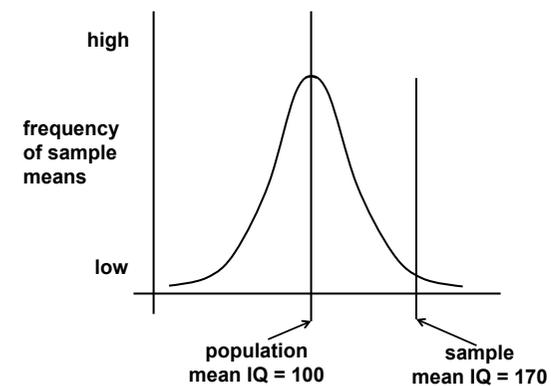
The human population IQ mean is 100.
A random sample of people has a mean IQ of 170.

There are two explanations:

(a) the sample is a fluke: by chance our random sample contained a large number of highly intelligent people.

(b) the sample does not come from the population we thought it did: our sample was actually from a different population - e.g., aliens masquerading as humans. Or, more likely, it was taken from the Mensa society members.

Relationship between population mean and sample mean:



This logic can be extended to the *difference between two samples from the same population*:

A common experimental design

We compare two groups of people:

- One group get the "wolfman" drug (Experimental group).
- The second group get a placebo (Control group).

At the start of the experiment, they are two samples from the same population ("humans").

At the end of the experiment, are they:
(a) still two samples from the *same* population (i.e., still two samples of "humans" – i.e. our experimental treatment has left them unchanged).

OR

(b) now samples from two *different* populations - one from the "population of humans" and one from the "population of wolfmen"?

