Probability and Samples

The Distribution of Sample Means
Chapter 7

Class Outline – 7-16-08

- Distribution of Sample Means – Chap. 7
- Questions

Chap. 7 – Probability & Samples

- The z-scores and probability that we dealt with on Monday only looked at a single score
  - (A z-score representing a single score and the probability of obtaining a single score from a population)
- Most research uses samples containing more than one score
- We will now look at z-scores and probability in terms of larger samples
  - Transforming a sample mean into a z-score

Samples

- Remember, a sample provides an incomplete picture of the population from which it was drawn
- The statistics computed for a sample will not be the same as the population it is drawn from
- This difference or error between a sample and the population is called ***sampling error***

Samples

- If you take different samples from a population, they will not be the same
- Example on board (sample from population, each sample will not be the same)
- Some samples represent the population well, others do not. How do we know whether the sample is accurately describing the population?

Samples

- **Distribution of sample means** - the collection of sample means for *all the possible random samples* of a particular size (n) that can be obtained from a population.
  - Distribution of *means*, not raw scores
- Collect every possible sample – can then compute probabilities for obtaining a particular sample

Samples

- **Sampling Distribution** – a distribution of *statistics* obtained by selecting all possible sample of a specific size from a population
  - Could be other statistics also
- The distribution of sample means is an example of a sampling distribution

Constructing a Sampling Distribution

- We can collect every possible random sample of size *n* from a population of scores
- We then get the mean for each sample and plot the means in a frequency distribution histogram

Population

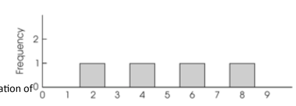

Figure 7-1 (p. 155)
Frequency distribution histogram for a population of four scores: 2, 4, 6, 8.

All Possible Random Samples

Table 7.1 (p. 156)
All the possible samples of $n = 2$ scores that can be obtained from the population presented in Figure 7.1. Notice that the table lists random samples. This requires sampling with replacement, so it is possible to select the same score twice.

	Scores		Sample mean
Sample	First	Second	(M)
1	2	2	2
2	2	4	3
3	2	6	4
4	2	8	5
5	4	2	3
6	4	4	4
7	4	6	5
8	4	8	6
9	6	2	4
10	6	4	5
11	6	6	6
12	6	8	7
13	8	2	5
14	8	4	6
15	8	6	7
16	8	8	8

Figure 7-2 (p. 156)
The distribution of sample means for *n* = 2. The distribution shows the 16 sample means from Table 7.1.

Distribution of Sample Means

- Characteristics:

1. Sample means tend to pile up around the population mean
2. Distribution of sample means is approximately a normal distribution
3. The larger the sample size, the closer the sample means should be to the population mean (large samples are more representative than small samples).

Back to the example…

- We can use the distribution of sample means to answer probability questions about sample means
  - $p\ (M > 5)$?

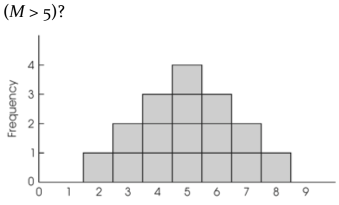

Reality

- Impossible to collect every possible random sample of any n from population N = 100, let alone N = 1,000!
- Fortunately, there is a mathematical solution

The Central Limit Theorem

- The central limit theorem shows what a distribution would look like if you took every possible sample, calculated every sample mean, and then constructed the distribution of sample means
- **Central limit theorem** – for any population with mean μ and standard deviation σ, the distribution of sample means for the sample size n will have a mean of μ and a standard deviation of $\sigma/\sqrt{n}$ and will approach a normal distribution as n approaches infinity

The Central Limit Theorum

- Importance:
  - Describes *any* population
  - Approaches normality quickly (n = 30)

The Central Limit Theorum

- The distribution of sample means will be normal if:
  - The population from which the samples are selected is a normal distribution

  AND/OR

  - The number of scores (n) in each sample is relatively large, around 30 or more (even if the samples were taken from a skewed distribution)

Expected value of *M*

- The mean of the distribution of sample means is equal to the population mean and is called the **expected value of *M***
- We *expect* the mean of the sample means to be equal to the population mean

Variability

- For our distribution of sample means, we describe the variability using the **standard error of *M*** instead of the standard deviation
  - Standard error of *M* = standard deviation for distribution of sample means
- Standard error measures the standard amount of difference between *M* and μ that is reasonable to expect simply by chance
  - (not all samples will have *M* = μ)
- The standard error is represent by σ_M

Standard Error

- Standard error tells you how well a sample mean estimates the population mean from which it was drawn
- The **law of large numbers** states that the larger a sample size is, the closer the sample mean will be to the population mean
  - (this means that the larger our sample size, the smaller our standard error)
- The formula for standard error uses both the standard deviation and the sample size
  - Standard error = $\sigma_M = \frac{\sigma}{\sqrt{n}}$

Standard Error

- Sampling Error - A sample will typically not provide a perfectly accurate representation of the population from which it was drawn
- Although they occur less frequently, samples drawn that have means that fall in the extremes (either tail) are not at all representative of the population
- Standard error provides a way to measure the average or standard distance between a sample mean and the population mean (sampling error)

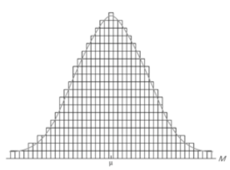

Figure 7-5 (p. 165)
An example of a typical distribution of sample means. Each of the small boxes represents the mean obtained for one sample.

3 Different Distributions

- Population (*N*)
  - μ (mean), σ (standard deviation)
- Sample (*n*)
  - *M* (mean), *s* (standard deviation)
- Distribution of Sample Means
  - All possible samples of *n*
  - μ (expected mean), σ_M (standard error)

Probability & Distribution of Sample Means

- The distribution of sample means is used to find the probability associated with any specific sample
  - Example: SAT Scores - what is probability of drawing a sample (n = 25) with M > 540?

Z-scores and Sample Means

- We can find a z-score for a sample mean just as we did for an individual score
  - A z-score of +2.0 tells us our sample mean is much larger than we would expect
- The formula for finding the z-score of a sample mean is slightly different:
  - $z = \frac{M - \mu}{\sigma_M}$
- With the z-score of our sample mean, we can look at the unit normal table to determine the probability of drawing that sample

Practice calculating sample Zs

- Population μ = 80, σ = 12
- What is the Z score for the following sample means?
  1. M = 84 for a sample of n = 9 scores
  2. M = 74 for a sample of n = 16 scores
  3. M = 81 for a sample of n = 36 scores

  - If we can calculate z scores, then we can determine probability

Example (*M* to *p*)

- Normal population, μ = 90, σ = 10
- Select sample of *n* = 25 scores
- What is the probability that the sample will have a mean greater than 94?
- *p* (*M* < 91)?

Example (*p* to *M*)

- SAT scores again
  - μ = 500, σ = 100
- What is the range of sample means that can be expected 80% of the time (*n* = 25)?

Inferential Statistics

- Treatment and control groups – do they differ?
- Using distribution of sample means, we can determine what would be expected for control sample
- If treatment sample falls outside expectations, can *infer* that treatment had an effect

Example

- A population forms a normal distribution with $\mu = 50$, $\sigma = 10$
- A sample of $n = 25$ scores is selected and given a treatment ($M = 55$)
- Is the sample typical of the population, or extreme?
- Or, Did the treatment have an effect?

Example

- An automobile manufacturer claims that a new model will average $\mu = 45$ miles per gallon (normal distribution), with $\sigma = 2$.
- You test a sample of $n = 4$ cars, and find $M = 42$ miles per gallon.
- Is the manufacturer lying/ wrong?
  - Is the sample mean likely to be obtained from the population? Or rare?
  - What is the range of sample means that we can expect 95% of the time?

Example

- In the general population, average reaction time to an obstacle in the road is $\mu = 600$ ms, $\sigma = 70$ ms.
- We "treat" a sample of $n = 4$ drivers to 3 tequila shots prior to testing.
  - $M = 630$ ms
- Did alcohol affect reaction time?
- What if our sample was $n = 25$?