Sampling Distributions and Estimation

PSY 5101: Advanced Statistics for Psychological and Behavioral Research 1

Sampling Distributions

- Pivotal subject: distributions of statistics. "Foundation…linchpin…important…crucial"
- You need sampling distributions to make inferences:
  - To get probabilities of statistics for decision making about parameters
  - To get information necessary to estimate parameters
- A distribution that could be formed by drawing all possible samples of a given size N from some population, computing the statistic for each sample, and arranging these statistics in a distribution
- Every statistic has a sampling distribution

Types of Distributions

- Population:
  - Distribution of all possible scores (Xs)
  - Usually large, unobtainable, and hypothetical
  - Has parameters μ and σ^2, the values of which are usually unknown
  - Unknown shape
  - We want to infer to one of the parameters or to the distribution itself

Types of Distributions

- Sample:
  - Distribution of the N scores that we actually have (Xs)
  - Usually a manageable size, already obtained, and real
  - Contained in what we will call our "real world"
  - Has known statistics like $\overline{X}$ and s^2
  - Known shape
  - We want to infer from one of the statistics to a parameter

Types of Distributions

- Sampling distribution:
  - Distribution of a statistic over all possible samples
    - Example: $\overline{X}$
  - Shows the variability of the statistic
  - Theoretical
  - Has parameters and usually a known shape
  - The bridge for the inference from the sample to the population (i.e., from the statistic to the parameter)
  - Where we get the probabilities of the statistic so we can make decisions about the parameter

Example

- IQ of deaf children example: Do deaf children have lower IQ scores than other children? Or are their scores the same as the average for the general population? If μ=100 and σ^2=225, then would $\overline{X}$ = 88.07 from a sample of N=59 deaf children be considered to be significantly lower than 100?
  - Note: IQ differences observed for deaf children appear to be due to the fact that their primary language (American Sign Language) is not English, so they score lower on the verbal part of the total IQ test
- What we are actually testing is the probability of obtaining a sample mean of 88.07 from 59 participants drawn from a population with a mean of 100

Types of Distributions

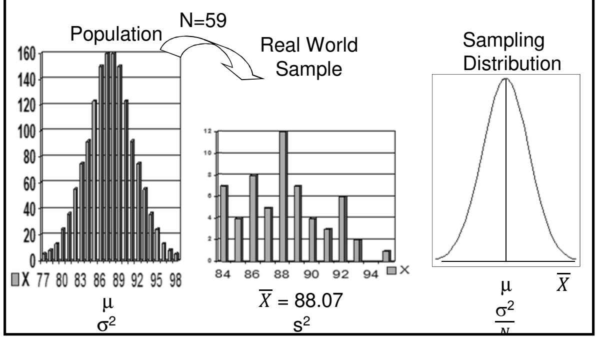

Sampling Distribution of $\overline{X}$

- The sampling distribution of $\overline{X}$
  - Has the purpose of any sampling distribution: to obtain probabilities
  - Has the definition of any sampling distribution: the distribution of a statistic
  - Has specific characteristics:
    - Mean: $\mu_{\overline{X}} = \mu$
    - Variance : $\sigma^2_{\bar{x}} = \frac{\sigma^2}{N}$
    - Standard error of the mean: $\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{N}}$
      - This is just the square root of the variance
    - Shape is normal if
      - Population is normal
      - N is large (more than about 30)
        - Central Limit Theorem

Sampling Distribution

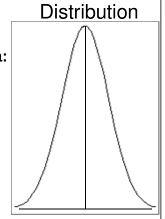

μ $\overline{X}$

$\frac{\sigma^2}{N}$

Sampling Distribution of $\overline{X}$: Use of $z_{\overline{X}}$

- IQ of deaf children:
  - What is the mean of this population distribution? Is it 100 as it is for the population of all IQ scores (μ=100 and σ^2=225)?
  - What is the probability of getting $\overline{X}$ = 88.07 or less if μ=100 (and σ^2=225)?
  - To get this probability, we need a new statistic, $z_{\overline{X}} = \frac{\overline{X} - \mu}{\sqrt{\frac{\sigma^2}{N}}}$
    - $z_{\overline{X}} = \frac{88.07 - 100}{\sqrt{\frac{225}{59}}}$ = -6.11
    - p($\overline{X} \leq$ 88.07) = p(z $\leq$ -6.11) < .00003

Sampling Distribution

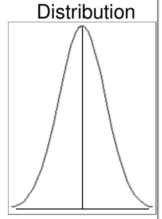

μ $\overline{X}$

$\frac{\sigma^2}{N}$

Use of $z_{\overline{X}}$

- IQ of deaf children:
  - So what does this look like and how does it help us decide about μ=100? Is the mean of the IQ of deaf children 100?
  - Because the probability of getting $\overline{X}$ = 88.07 or less if μ=100 is so small (less than .00003) we reject the idea that μ=100.
  - It is very unlikely to get the data that led to $\overline{X}$ = 88.07 from a population with μ=100.

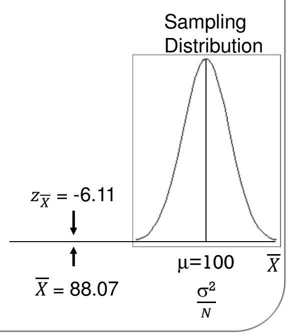

Other Sampling Distributions

- The sampling distribution of $\overline{X}$ is the first sampling distribution we learn but it is not the only one (all statistics have sampling distributions)
- All sampling distributions have in common:
  - Purpose: to obtain probabilities
  - Definition: the distribution of a statistic
- …but each sampling distribution has specific characteristics like mean, variance, and shape

Other Sampling Distributions

- Sampling distributions of s^{*2} and s^2:
  - Both have shapes that are positively skewed
  - The mean of s^{*2} is $\frac{N-1}{N}\sigma^2$, always smaller than σ^2
  - The mean of s^2 is σ^2

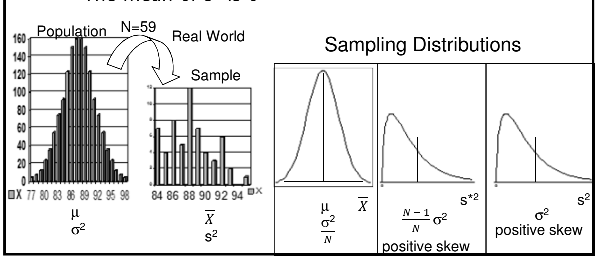

Other Sampling Distributions

- Sampling distributions of r, s*, and s
  - r: the mean is ρ (rho) if ρ=0, and the shape is symmetric but not normal
  - s* and s: neither has a mean equal to σ

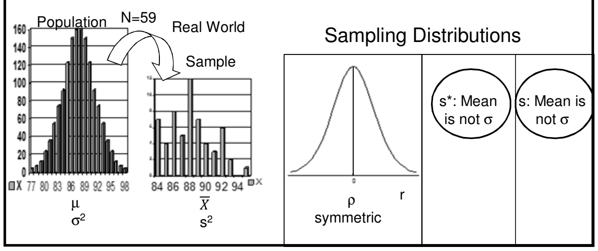

Estimation

- You need sampling distributions to make inferences:
  - To get probabilities of statistics for decision making about parameters
  - To get information necessary to estimate parameters
- Estimation is the calculation of an approximate value of a parameter
  - **Point estimation** is the use of a statistic as a single value (point) to estimate a parameter
  - Any statistic "can" be used to estimate any parameter
  - Some statistics are good (and logical) estimates of particular parameters such as using $\overline{X}$ as an estimate of μ
  - "Unbiased estimate" is one definition of "good estimate"

Estimation: Unbiased

- Unbiased estimate: A statistic is an unbiased estimate of a parameter if the mean of its sampling distribution is equal to the parameter: $\mu_{statistic}$ = desired parameter
- The following statistics are unbiased estimates of their corresponding parameters:
  - $\overline{X}$ is an unbiased estimate of μ because $\mu_{\overline{X}} = \mu$
  - s^2 is an unbiased estimate of σ^2 because $\mu_{s^2} = \sigma^2$
  - r is an unbiased estimate of ρ because $\mu_r = \rho$ if ρ=0
- Note that the statistic and parameter can change but the definition of unbiased is $\mu_{statistic}$ = desired parameter

Estimation: Unbiased

- The following statistics are not unbiased estimates of their corresponding parameters (each is a biased estimate):
  - s^{*2} is a biased estimate of σ^2 because $\mu_{s^{*2}} \neq \sigma^2$
  - s^* is a biased estimate of σ because $\mu_{s^*} \neq \sigma$
  - s is a biased estimate of σ because $\mu_s \neq \sigma$
- Note that the statistic and parameter can change but the definition of unbiased is $\mu_{statistic}$ = desired parameter
  - Always μ of the statistic with this μ being equal to the desired parameter

The Bias in s^{*2}

- The mean of the sampling distribution of s^{*2} is only a fraction of σ^2
  - s^{*2} is a biased estimate of σ^2 because it tends to be too small (not every value of s^{*2} will be too small…but on average it is an underestimate of σ^2)
- Here is the "ideal" estimate of population variance: $\sigma^2 = \frac{\sum(X-\mu)^2}{N}$
  - The problem is that we almost never know μ so we usually cannot use this formula

The Bias in s^{*2}

- Using $\overline{X}$ as an estimate of μ leads us to:

  $s^{*2} = \frac{\sum(X-\overline{X})^2}{N}$
  - Remember that $\overline{X}$ minimizes the sum of the squared deviations such that $\sum(X-\overline{X})^2$ is as small as it can be and smaller than $\sum(X-\mu)^2$
  - The least-squares property of $\overline{X}$ makes it so that the estimate of σ^2 is too small when $\overline{X}$ is used in place of μ

The Bias in s^{*2}

- How much of an underestimate is s^{*2}?
  - Consider the means of the sampling distributions of the numerators of "ideal" formula and s^{*2} formula
  - Mean of sampling distribution of $\sum(X-\mu)^2$ = $N\sigma^2$
  - Mean of sampling distribution of $\sum(X-\overline{X})^2$ = $N\sigma^2$-σ^2
    - The numerator of the formula for s^{*2} is too small by σ^2
    - Rewriting $N\sigma^2$-σ^2 as $\sigma^2(N-1)$ offers a solution to the problem
    - Mean of sampling distribution of $\sum(X-\overline{X})^2 = \sigma^2(N-1)$
    - Because $\sum(X-\overline{X})^2$ has an average value of $\sigma^2(N-1)$, all you have to do is divide this numerator by N-1 instead of N and it will be an unbiased estimator of σ^2
    - Mean of sampling distribution of $\frac{\sum(X-\overline{X})^2}{N-1} = \sigma^2$

Bias in X_{50}

- $\overline{X}$ is an unbiased estimate of μ
  - $\overline{X}$ for each sample will not always be equal to μ but there will not be any systematic error
  - That is, $\overline{X}$ may sometimes be larger than μ and may sometimes be smaller than μ
  - However, if you repeatedly sample from the population, then the average $\overline{X}$ would equal μ
- X_{50} is only an unbiased estimate of μ when the population is symmetric
  - Because X_{50} would be the same as $\overline{X}$ in that case
- It makes sense that the sample mean would be the best estimate of the population mean

Interval Estimation

- So far, we have focused on point estimation (estimating a parameter with a statistic)…now we will focus on interval estimation
- Interval estimation involves obtaining an interval of potential values for a parameter
- Example: $\overline{X}$ = 88.07 is the point estimate of μ for the IQ of deaf children…but we cannot say that the population mean of the IQ of deaf children is exactly 88.07 because there is error in the estimate of $\overline{X}$

Interval Estimation and Error

- Standard error of the mean is a measure of the amount of error in $\overline{X}$ as an estimate of μ
  - Standard error of the mean = $\frac{\sigma}{\sqrt{N}}$
    - Notice that the standard error of the mean is the standard deviation of the sampling distribution of $\overline{X}$…that is, it is the square root of the variance of the sampling distribution of $\overline{X}$ which is $\frac{\sigma^2}{N}$
    - Large standard error means that there is a lot of variability between the means of different samples…so the mean of our sample may not be representative of the population
      - To calculate standard error for a sample, we substitute s for σ: $\frac{s}{\sqrt{N}}$
- We use the standard error of the mean to place bounds on $\overline{X}$ in such a way that we are confident that they include the true value of μ a certain percentage of the times μ could be estimated

Confidence Interval

- Confidence interval is a range of values that is constructed such that a certain percentage of the time (usually 95% or 99%) the true value of the population mean will fall within these limits
  - If we collected 100 samples, calculated the mean for each sample, and then calculated a confidence interval for that mean…then for 95 of those samples the confidence interval would contain the true value of the mean in the population

Calculating the Confidence Interval

- 95% of z-scores fall between -1.96 and +1.96
- If our sample means are normally distributed with a mean of 0 and a standard error of 1, then the limits of our 95% confidence interval would be -1.96 and +1.96
- We could convert our raw scores to z-scores but we usually use the following equations instead (to maintain the original metric)
  - Lower boundary of CI = $\overline{X}$ - (1.96*Standard Error)
  - Upper boundary of CI = $\overline{X}$ + (1.96*Standard Error)
  - Sample mean is always at the center of the CI

Confidence Intervals for μ

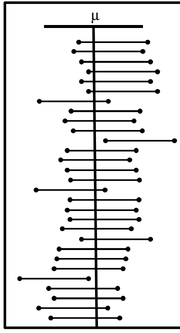

- Confidence intervals for μ each based on a different sample of size N
- Vertical line represents the constant fixed value of μ whereas the horizontal lines represent different confidence intervals for μ
- Most interval brackets include μ but some do not
- Note that the intervals vary but μ does not
- The location of each interval is a function of $\overline{X}$, the length of each interval is fixed for a given N, σ^2, and degree of confidence