



Lecture 3

Sampling distributions. Counts, Proportions, and sample mean.

- **Statistical Inference:** Uses data and summary statistics (mean, variances, proportions, slopes) to draw ***conclusions*** about a population or process.
- **Statistic:** Any random variable measured from a random sample or in a random experiment.
- **Sampling distribution of a statistic:** shows how a statistic varies in repeated measurements of an experiment. The probability distribution of a statistic is called its sampling distribution.
- **Population distribution of a statistic:** distribution of values for all members of the population. Unknown, but estimable using laws of statistics.

Sampling Distribution for Counts and Proportions:

- In a survey of 2500 engineers, 600 of them say they would consider working as a consultant. Let X = the number who would work as consultants.
- X is a count:
- Sample Proportion of people who would work as consultants:

Distinguish count from sample proportion, they have *different* distributions.

Binomial Distribution for Sample Counts:

- Distribution of the count, X , of successes in a binomial setting with parameters n and p
- n = number of observations
- p = P (Success) on any one observation
- X can take values from 0 to n
- **Notation:** $X \sim \text{Bin}(n, p)$
- **Setting:**
 1. Fixed number of n observations
 2. All observations are independent of each other
 3. Each observation falls into one of two categories: Success or Failure
 4. P (Success) = P (S) = p

EXAMPLES (Bin or not Bin)

- Toss a fair coin 10 times and count the number X of heads. What about a biased coin?
- Deal 10 cards from a shuffled deck of 52. X is the number of spades. Suggestions??
- Number of girls born among first 100 children in a (large) hospital this year.
- Number of girls born in this hospital so far this year.

Finding Binomial Probabilities

- You can use R to evaluate probabilities:
 - » `pbinom(4,size=10,prob=0.15)` (calculates $P(\text{Bin}(10,0.15) \leq 4)$)
 - » `[1] 0.990126`

Example

Your job is to examine light bulbs on an assembly line. You are interested in finding the probability of getting a defective light bulb, after examining 10 light bulbs.

- Let X = number of defective light bulbs
 - $P(\text{defective}) = .15$
 - $N = 10$
1. Is this a binomial set up?
 2. What is the probability that you get at most 2 defective light bulbs?
 3. What is the probability that the number of defective light bulbs you find is greater than eight?
 4. What is the probability that you find between 3 and 5 defective light bulbs?

Binomial Mean and Standard Deviation

$$\mu_x = np$$

$$\sigma_x^2 = np(1 - p)$$

$$\sigma_x = \sqrt{np(1 - p)}$$

Example: Find the mean and standard deviation of the previous problems

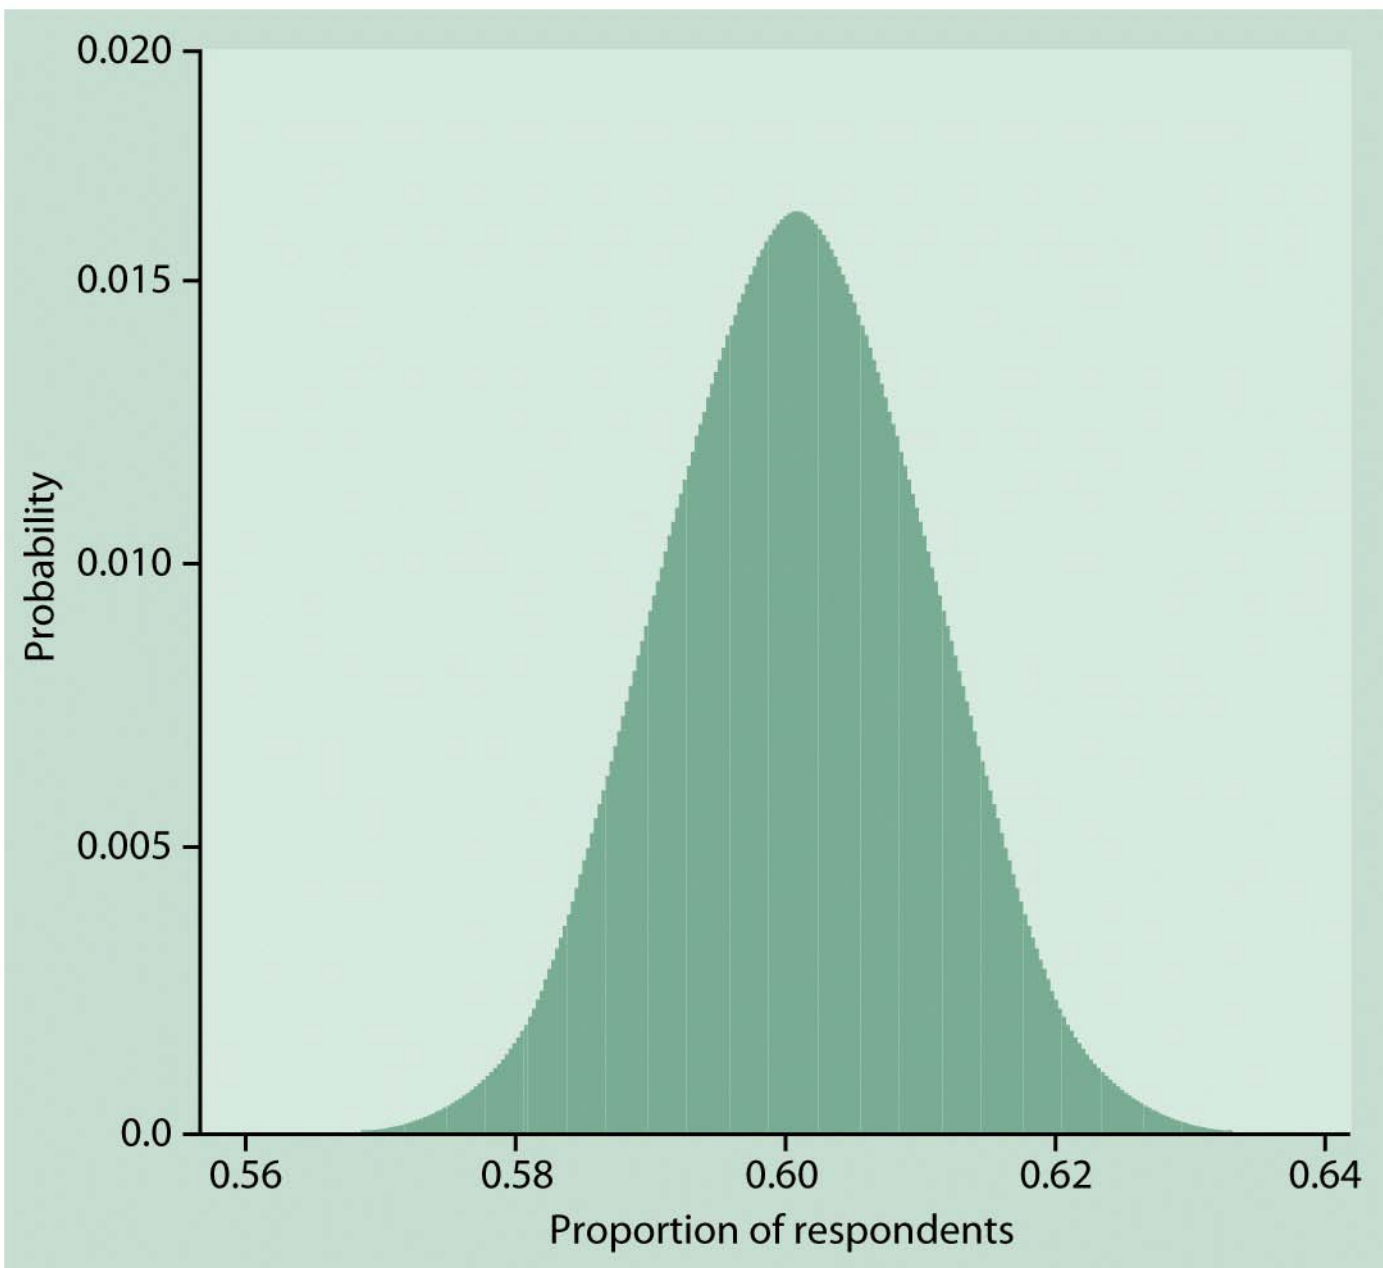
Sample Proportions

- Let X be a count of successes in $n =$ total number of observations in the data set.
- Then the sample proportion:

$$\hat{p} = \frac{X}{n}$$

– NOTE!!!!

- We know that X is distributed as a Binomial, however \hat{p} is NOT distributed as a Binomial.



Normal approximation for counts and proportions

- If X is $B(n,p)$, $np \geq 10$ and $n(1-p) \geq 10$ then:

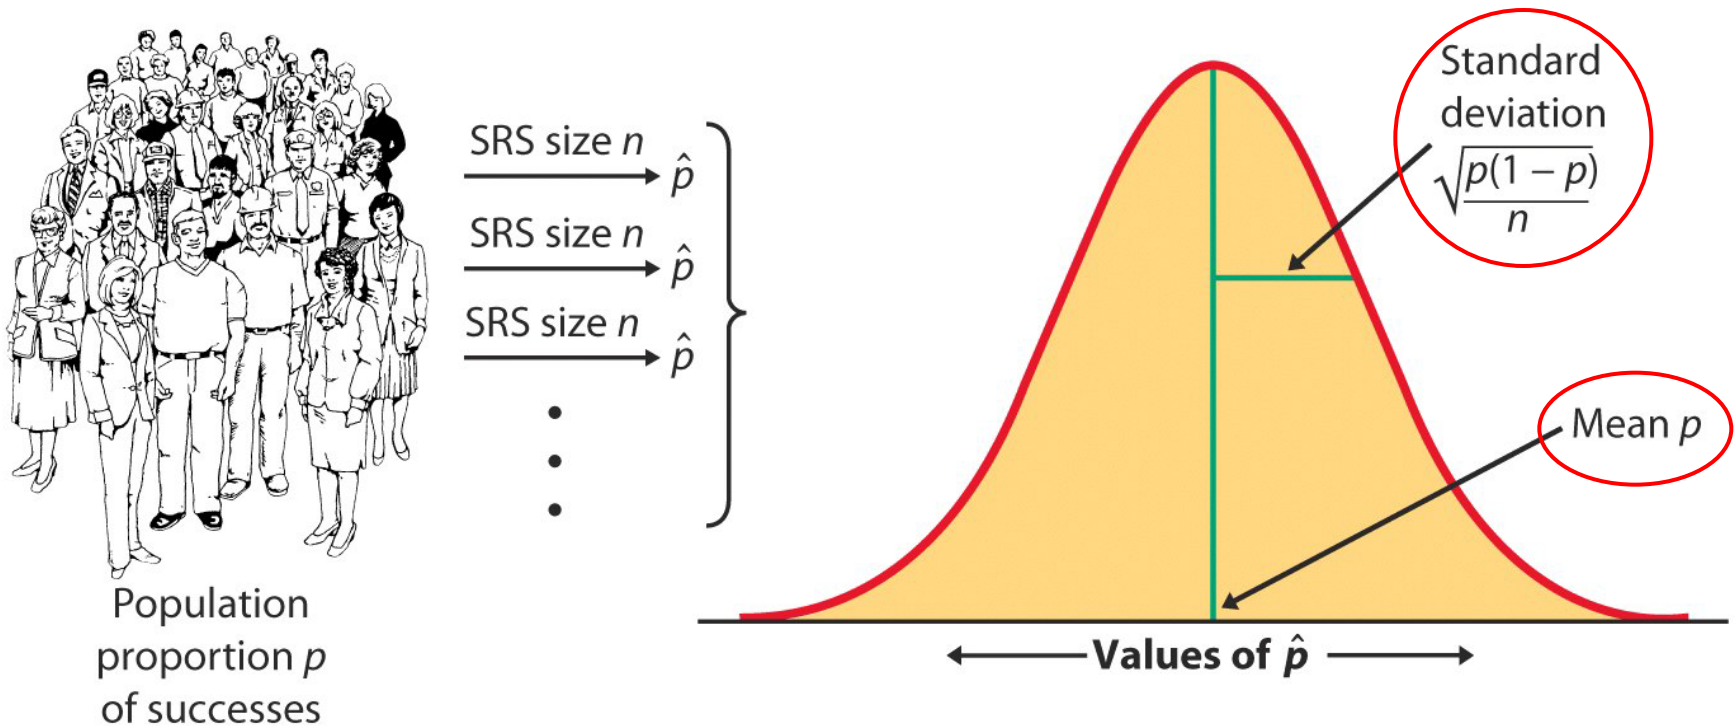
X is approximately $N(np, \sqrt{np(1-p)})$

\hat{p} is approximately $N(p, \sqrt{\frac{p(1-p)}{n}})$

Sampling distribution of \hat{p}

The sampling distribution of \hat{p} s never exactly normal. But as the sample size increases, the sampling distribution of \hat{p} becomes approximately normal.

The normal approximation is most accurate for any fixed n when p is close to 0.5, and least accurate when p is near 0 or near 1.



Example:

- In a survey 2500 engineers are asked if they would consider working as consultants. Suppose that 60% of the engineers would work as consultants. When we actually do the experiment 1375 say they would work as consultants

Find the mean and standard deviation of \hat{p} .

What is the probability that the percent of to be consultants in the sample is less than .58?

Between .59 and .61?

The continuity correction:

- **Example:** According to a market research firm 52% of all residential telephone numbers in Los Angeles are unlisted. A telemarketing company uses random digit dialing equipment that dials residential numbers at random regardless of whether they are listed or not. The firm calls 500 numbers in L.A.
 1. What is the exact distribution of the number X of unlisted numbers that are called?
 2. Use a suitable approximation to calculate the probability that at least half the numbers are unlisted.

The continuity correction(cont.):

- In the previous problem if we compute the probability that exactly 250 people had unlisted numbers using the normal approximation we would have find this probability equals zero.
- That is obviously not right because this number has to have some probability (small but still not zero).
- The problem comes from the fact that we use a continuous distribution (Normal Distribution) to approximate a discrete one (Binomial Distribution).
- So to improve the approximation we use a correction:
 - Whenever we compute a probability involving a count we will move the interval we compute 0.5 as to include or exclude the endpoints of the interval depending on the type of interval (closed or open) we compute in the problem.
 - Then we use the normal approximation to compute the probability of this new interval.

- Example: In the previous problem find:

$$P(X \geq 250)$$

$$P(X > 250)$$

$$P(X = 250)$$

$$P(X < 250)$$

$$P(248 < X < 251)$$

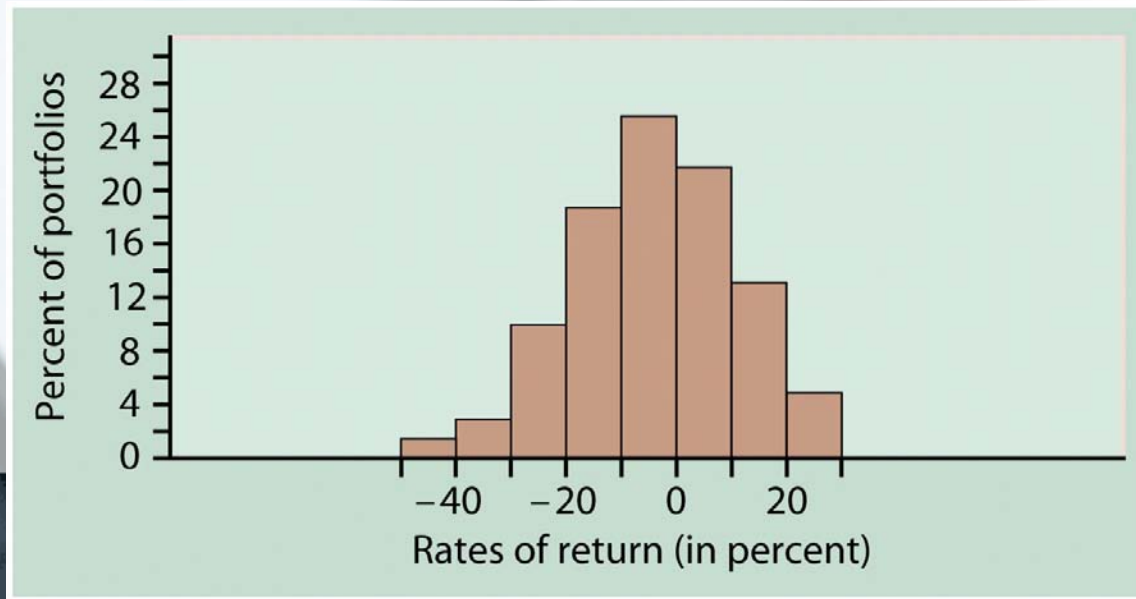
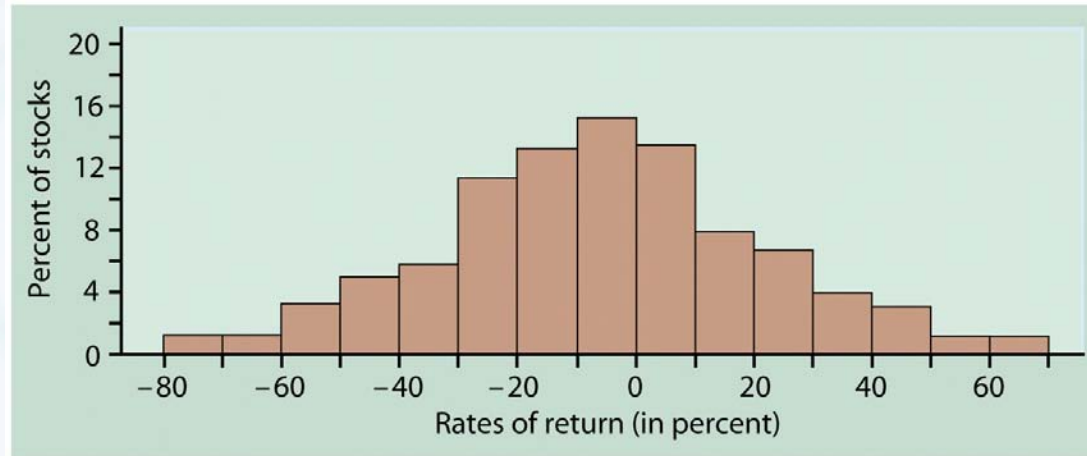
$$P(248 < X \leq 251)$$

$$P(248 \leq X \leq 251)$$

Sampling distribution of the sample mean

- **Distribution of the center and spread**
- Setup:
- Draw a SRS (simple random sample) of size n from a population.
- Measure some variable X (i.e. income)
- Data: n random variables, $X_1, X_2, X_3 \dots X_n$, where X_i is a measurement on 1 individual (i.e. income of 1 individual in the sample)
- Since the individuals are randomly chosen, the X_i 's can be considered to be independent

Example: Distribution of individual stocks (up) vs. distribution of mutual funds (down)

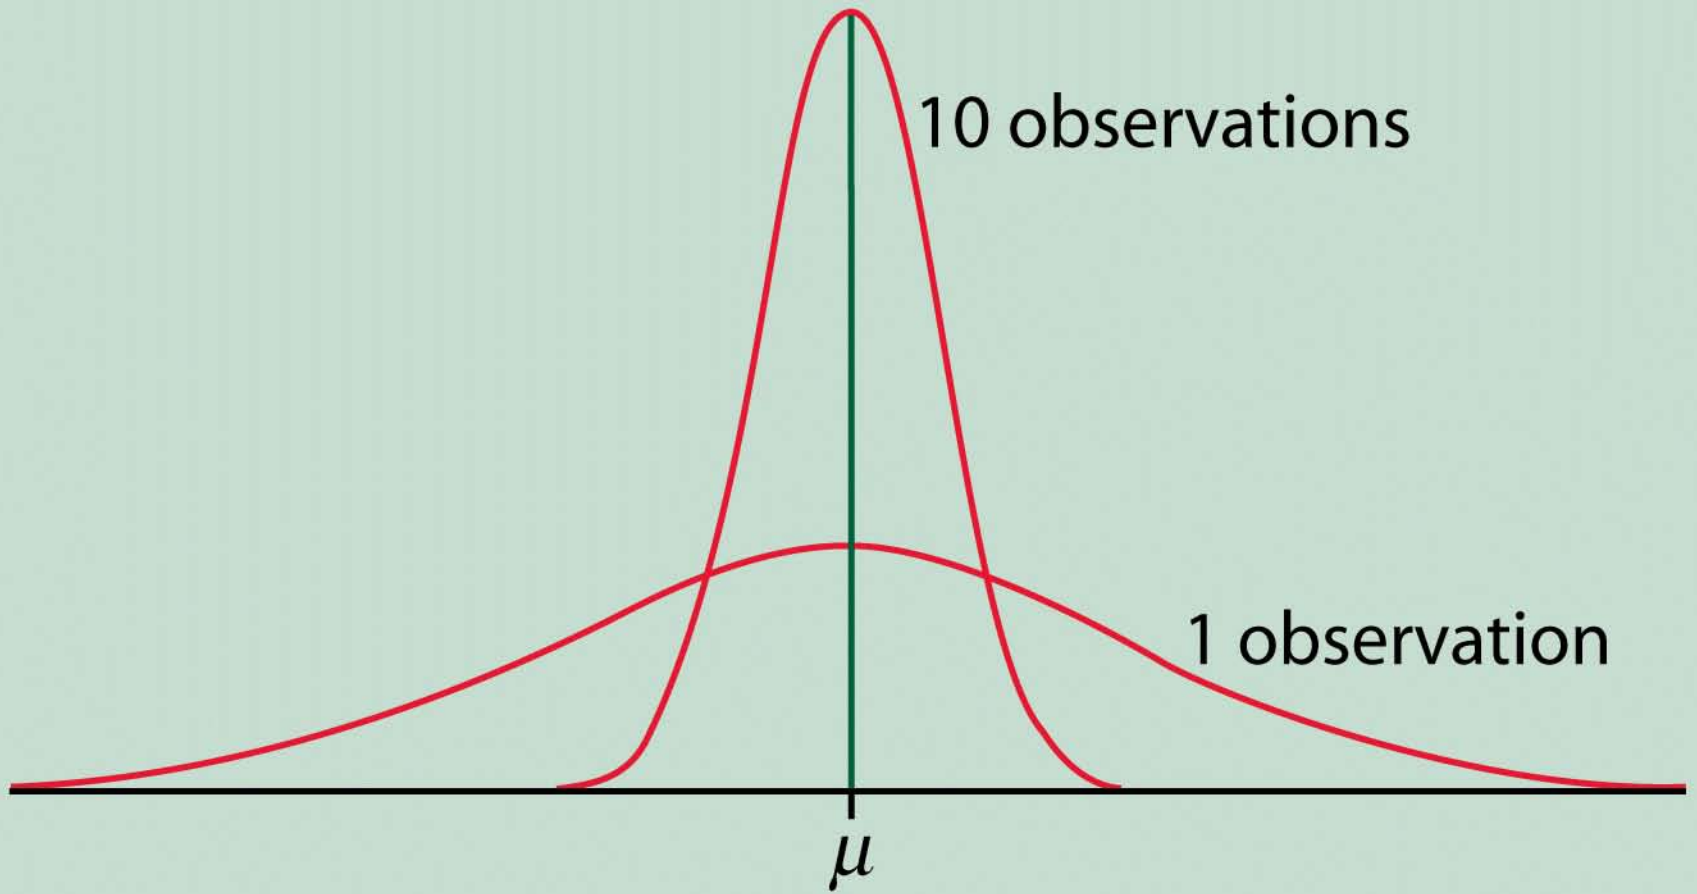


Sample mean: $\bar{X} = \frac{X_1 + X_2 + X_3 + \dots + X_n}{n}$

- Let \bar{X} be the mean of an SRS (simple random sample) of size n from a population with mean μ and standard deviation σ . The mean and standard deviation of \bar{X} are:

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

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



Central Limit Theorem:

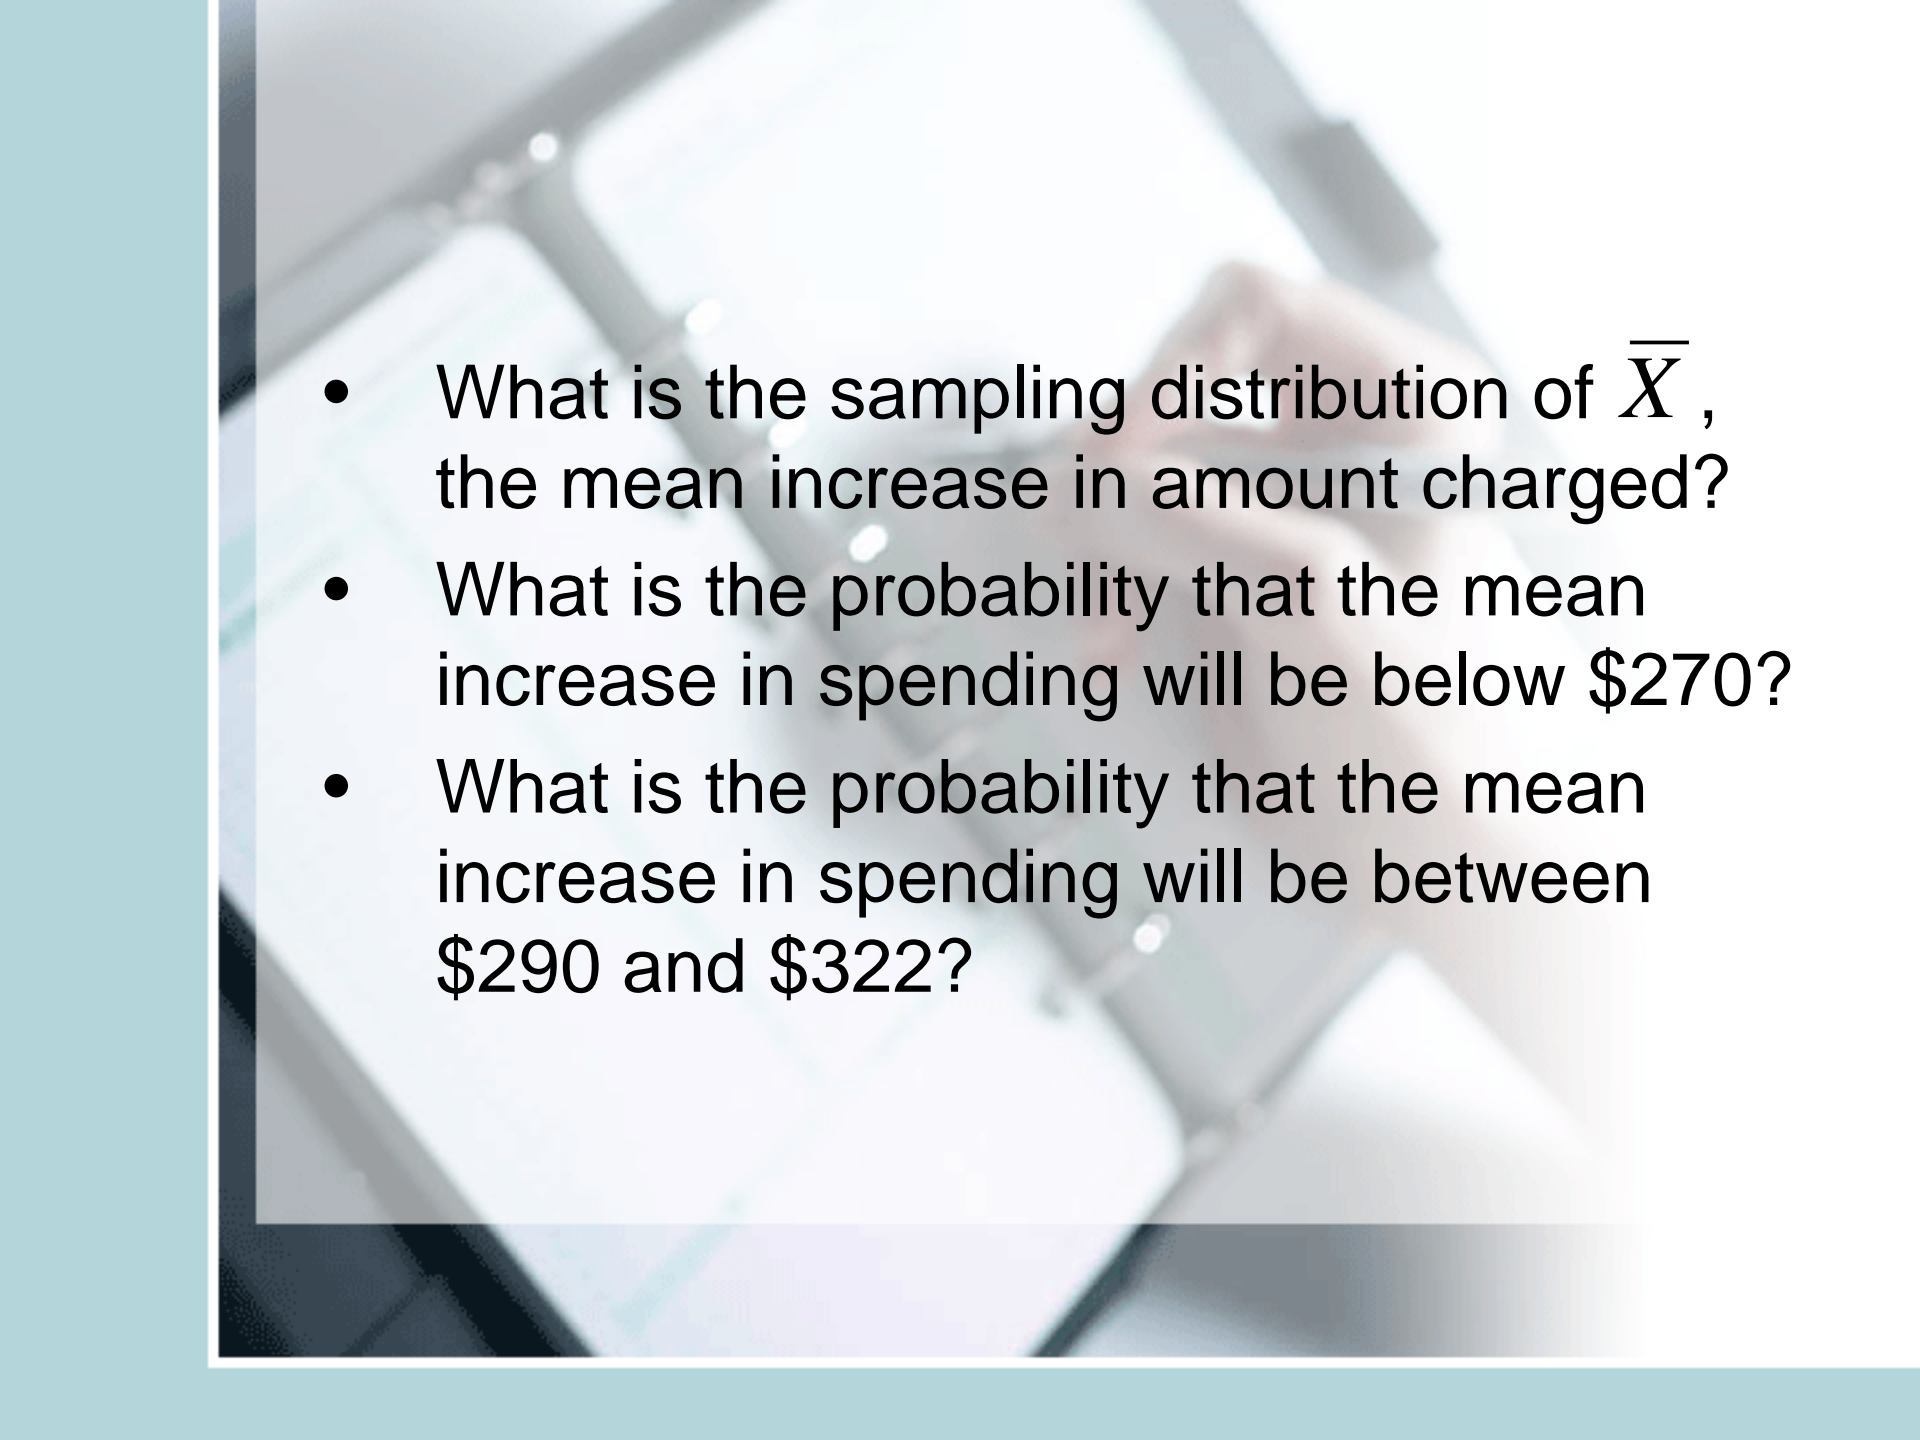
- Draw a SRS of size n (n large) from any population with mean μ and standard deviation σ . The sampling distribution of the sample mean is approximately normal:

$$\bar{X} \sim N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$$

- **Important special case:** If the *population is normal* then the sample mean has exactly the normal distribution: $N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$

Example:

- A bank conducts an experiment to determine whether dropping their annual credit card fee will increase the amount charged on the credit card. The offer is made to a SRS of 200 customers. The bank then compares the amount the customers charged on their cards this year, to the amount charged next year. A mean increase of \$308 with a standard deviation of \$108 was found.

- 
- A hand holding a pen over a document with a grid pattern, likely a ledger or account book. The background is a light blue and white grid.
- What is the sampling distribution of \bar{X} , the mean increase in amount charged?
 - What is the probability that the mean increase in spending will be below \$270?
 - What is the probability that the mean increase in spending will be between \$290 and \$322?

Example:

- The number of accidents per week at a hazardous intersection varies with mean 2.2 and standard deviation 1.4.
- What is the distribution of \bar{X} , the mean number of accidents in one year, (52 weeks)?
- What is the probability that \bar{X} is less than 2?
- What is the probability that there are fewer than 100 accidents in a year?

Example:

- The weight of eggs produced by a certain breed of hen is Normally distributed with mean 65 grams and standard deviation 5 grams. Let cartons of such eggs be considered to be SRSs of size 12. What is the probability that the weight of a carton falls between 750 grams and 825 grams?


Practical note

- ❑ Large samples are not always attainable.
 - ❑ Sometimes the cost, difficulty, or preciousness of what is studied drastically limits any possible sample size.
 - ❑ Blood samples/biopsies: No more than a handful of repetitions acceptable. Often, we even make do with just one.
 - ❑ Opinion polls have a limited sample size due to time and cost of operation. During election times, though, sample sizes are increased for better accuracy.
- ❑ Not all variables are normally distributed.
 - ❑ Income, for example, is typically strongly skewed.
 - ❑ Is \bar{x} still a good estimator of μ then?


The central limit theorem

Central Limit Theorem: When randomly sampling from any population with mean μ and standard deviation σ , **when n is large enough**, the sampling distribution of \bar{x} is approximately normal: $\sim N(\mu, \sigma/\sqrt{n})$.


Population with strongly skewed distribution




Sampling distribution of \bar{x} for $n = 2$ observations



Sampling distribution of \bar{x} for $n = 10$ observations



Sampling distribution of \bar{x} for $n = 25$ observations





Income distribution

Let's consider the very large database of individual incomes from the Bureau of Labor Statistics as our population. It is strongly right skewed.

- We take 1000 SRSs of 100 incomes, calculate the sample mean for each, and make a histogram of these 1000 means.
- We also take 1000 SRSs of 25 incomes, calculate the sample mean for each, and make a histogram of these 1000 means.

Which histogram corresponds to the samples of size 100? 25?



How large a sample size?

It depends on the population distribution. More observations are required if the population distribution is far from normal.

- ▣ A sample size of 25 is generally enough to obtain a normal sampling distribution from a strong skewness or even mild outliers.
- ▣ A sample size of 40 will typically be good enough to overcome extreme skewness and outliers.

In many cases, $n = 25$ isn't a huge sample. Thus, even for strange population distributions we can assume a normal sampling distribution of the mean and work with it to solve problems.

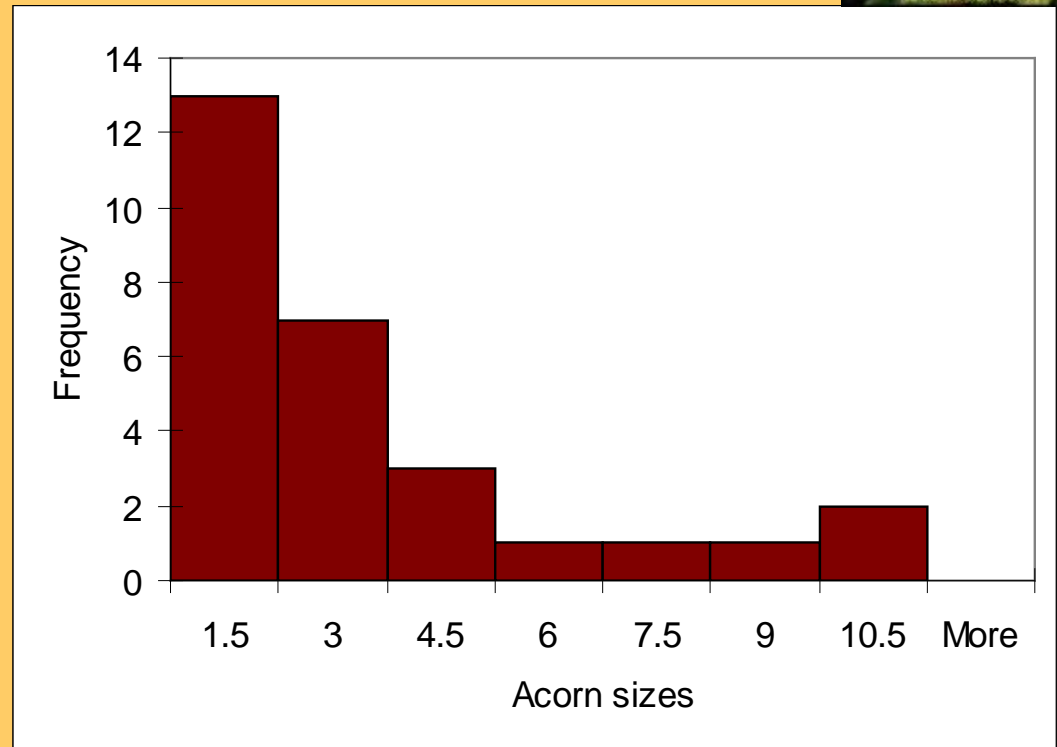
Sampling distributions



Atlantic acorn sizes (in cm^3)

— sample of 28 acorns:

- Describe the histogram.
What do you assume for the population distribution?



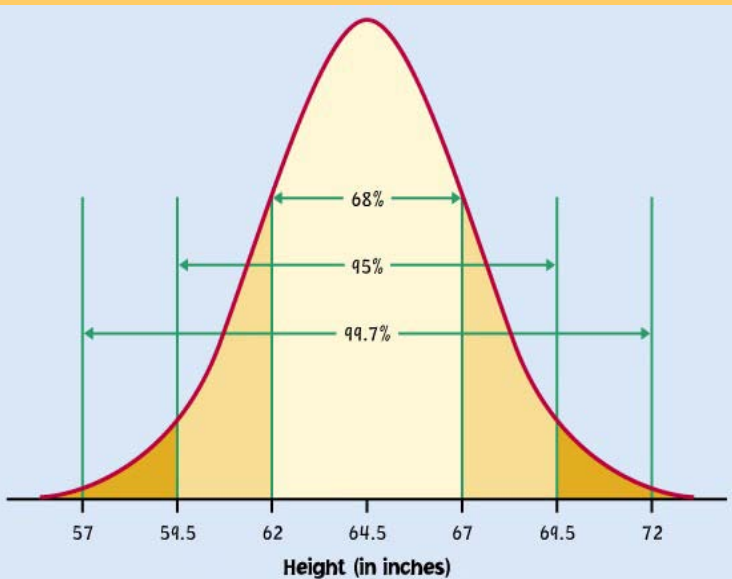
- What would be the shape of the sampling distribution of the mean:
 - For samples of size 5?
 - For samples of size 15?
 - For samples of size 50?

Further properties

Any linear combination of independent random variables is also normally distributed.

More generally, the central limit theorem is valid as long as we are sampling many small random events, even if the events have different distributions (as long as no one random event dominates the others).

Why is this cool? It explains why the normal distribution is so common.



Example: Height seems to be determined by a large number of genetic and environmental factors, like nutrition. The “individuals” are genes and environmental factors. Your height is a mean.