

Statistical Foundation for Business Analytics

Sumeet Gupta

Professor (Information Technology and Systems)

Indian Institute of Management Raipur

Outline

- Variation and Variance
- Summarizing Variance
- Statistical Inference: Random Sampling
- Confidence Interval and Hypothesis Testing
- Hypothesis Testing for Ordinal Data
- Hypothesis Testing for Nominal Data – Chi Square
- Testing for Differences among Groups – ANOVA
- Testing for Relationships – Correlation and Regression

Variation and Variance

Principle of Variation

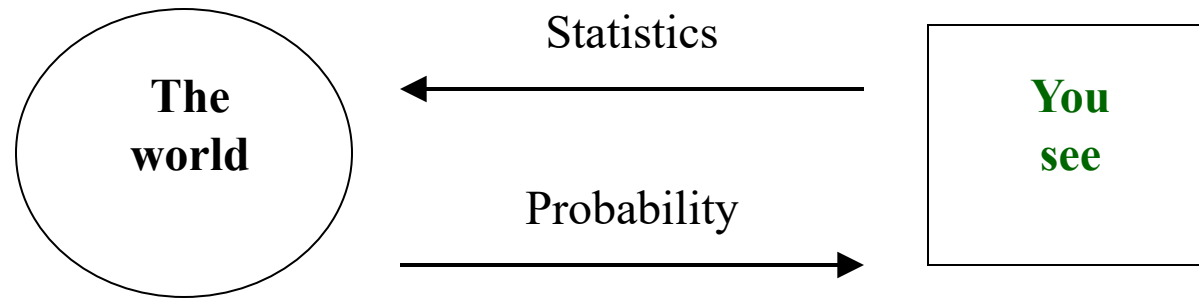
- Constant: Whose value does not change
 - E.g. Gravitational Constant for earth
- But most observations in life vary
 - E.g. Rising time of the sun; Amount we eat everyday; Sales; Production figures
- Can we capture this variation?
 - If the variation is small, our predictions can be fairly accurate
 - E.g. Rising time of the sun
 - But if the variation is large, our predictions remain inaccurate
 - E.g. Occurrence of an Earthquake

Data Science

- Capturing this variation and making predictions
 - Data → Estimate the Unknown
- Statistics: Art and Science of Collecting and Understanding DATA
 - DATA = Recorded Information (Sales, Productivity, Quality, Costs, Return, ...)
- Data Mining: Search for patterns in large data sets
 - Statistics: Prediction, Classification and Clustering
 - Computer Science: efficient algorithms (instructions) for collecting, maintaining, organizing, analyzing data
 - Optimization: calculations to achieve a goal

Probability

- “Inverse” of statistics



- **Statistics**: generalizes from data to the world
- **Probability**: “What if ...” Assuming you know how the world works, what data are you likely to see?
- Examples of probability:
 - Flip coin, stock market, future sales, IRS audit, ...
- Foundation for statistical inference

Types of Variables

- A Variable
 - The type of measurement being done
 - e.g., Sales volume, Cost, Productivity, Number of defects, ...
- Types
 - Quantitative variables (Ratio Vs Interval)
 - Qualitative (Nominal Vs Ordinal)
- Collection
 - Cross-Sectional (No meaningful sequence)
 - Over a period of time (Meaningful sequence – Time Series)

How Many Variables?

- Univariate data set: One variable measured for each elementary unit
 - e.g., Sales for the top 30 computer companies.
 - Can do: Typical summary, diversity, special features
- Bivariate data set: Two variables
 - e.g., Sales and # Employees for top 30 computer firms
 - Can also do: relationship, prediction
- Multivariate data set: Three or more variables
 - e.g., Sales, # Employees, Inventories, Profits, ...
 - Can also do: predict one from all other variables

Visualizing and Summarizing Variance


Visualizing and Summarizing Variance

- Quantitative Variables
 - Visual Representation: Histograms and Boxplots
 - Central Tendency: Mean, Median and Mode
 - Dispersion: Std Deviation, Variance, Range, Coefficient of variation
- Qualitative Variables
 - Visual Representation: Bar Charts, Pie Charts
 - Central Tendency: Median and Mode
 - Dispersion: No of Categories
- Time Series Data
 - Visual Representation: Time Series Chart
 - Descriptors: Trend, Seasonality and Cyclicity

Histograms

- A Picture of a list of numbers




Data
11 15
8 26
10 5
15



- BARS ARE HIGH when many elementary units fall within this range
- Shows typical value (center), dispersion (variability), distribution shape, outliers (if any)

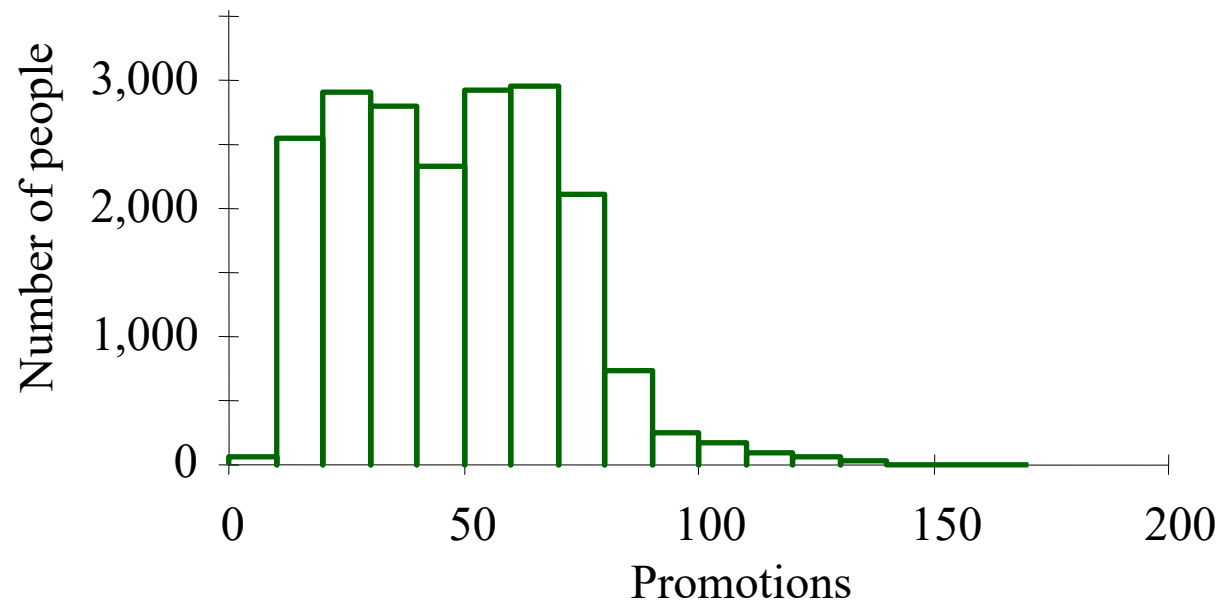
Histogram: Ideal Shapes

- Normal
 - Symmetric
 - Bell-Shaped
- Skewed
 - Not symmetric
 - Can cause trouble
 - Transform? Logarithm?
- Bimodal
 - Two clear groups
 - Find out why!
 - Analyze separately?



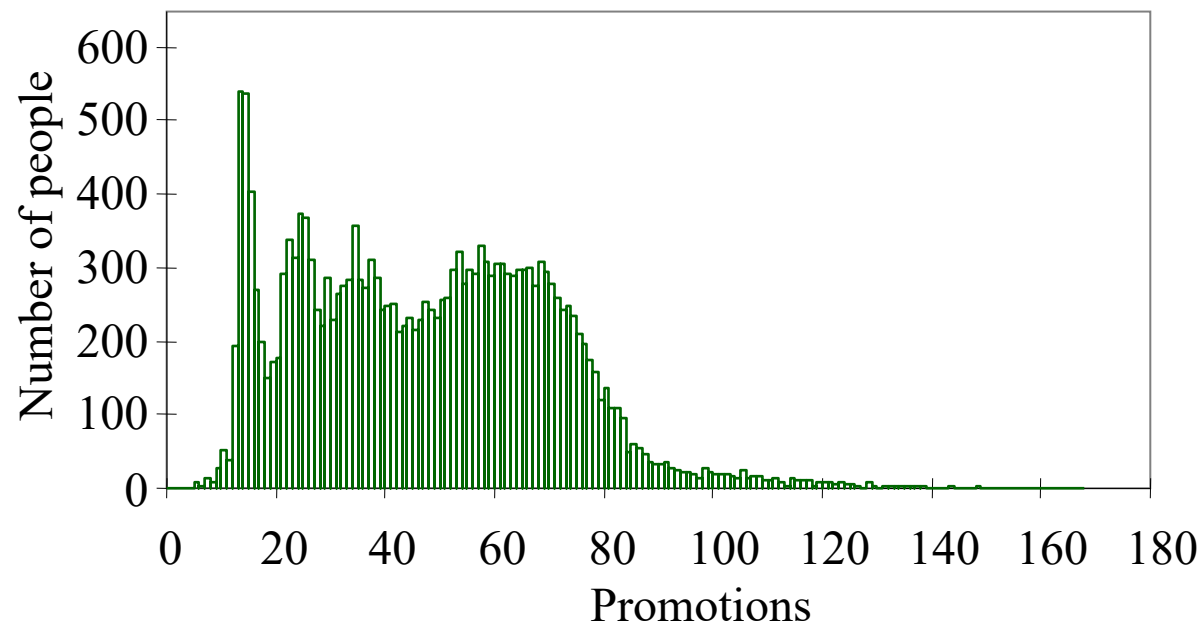
Histogram: Example

- **Promotions Dataset**
- Number of promotions received by 20,000 people in the donations database



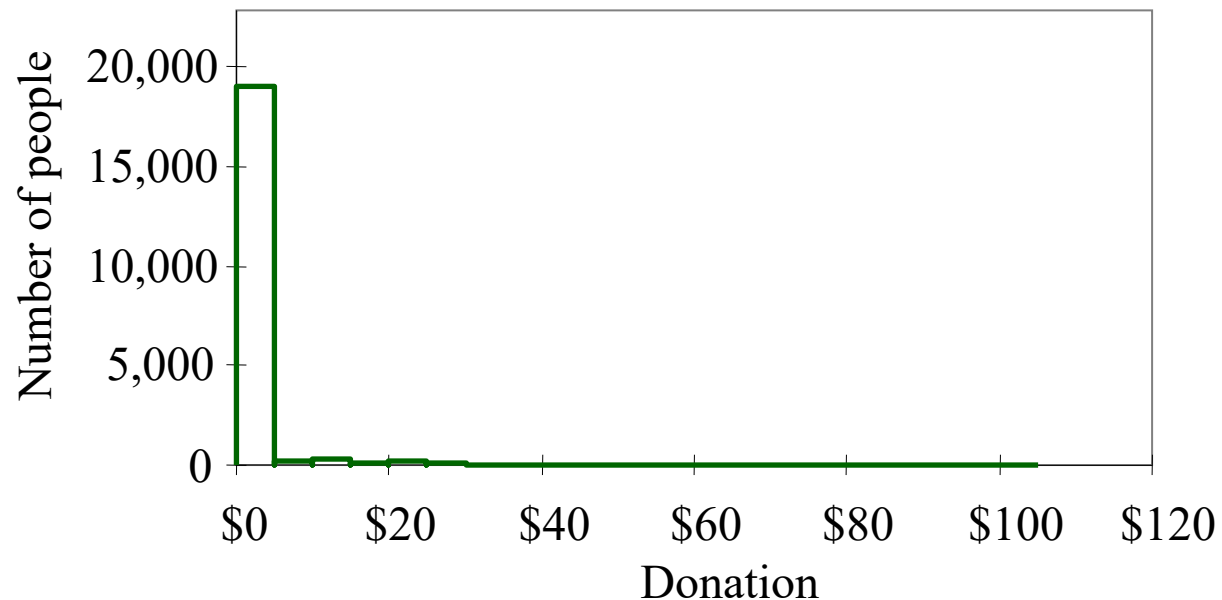
Histogram: Example

- **Promotions Dataset**
- Reduce bar width from 10 to 1 promotion
- With large data set, can see interesting structure
 - such as the peak at about 15 promotions



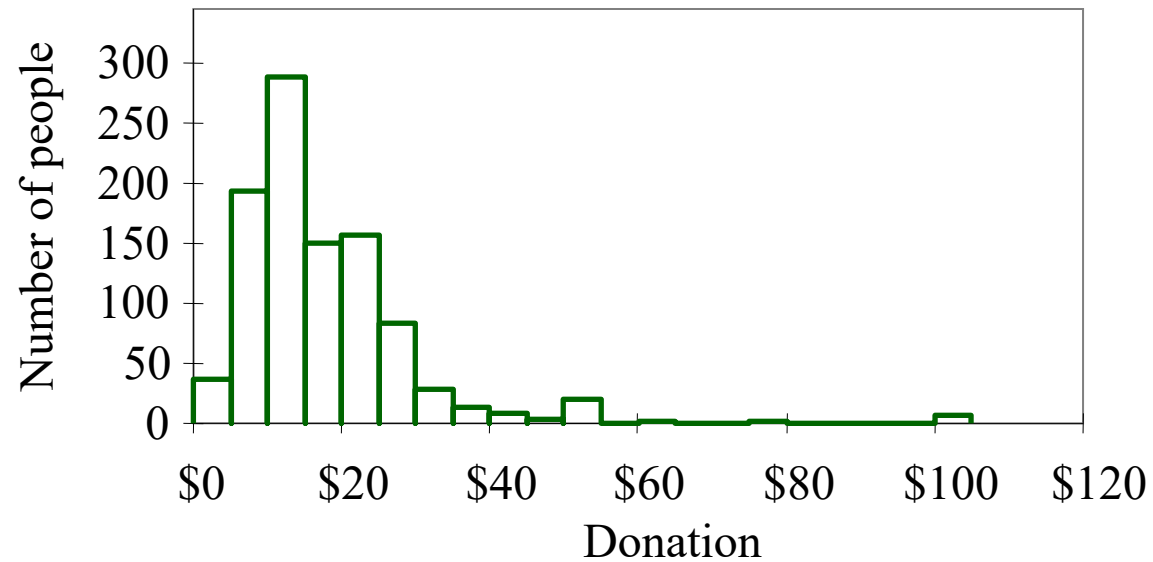
Histogram: Example

- **Donations Dataset**
- Size of donation received in response to mailing
- Note: many donations of \$0 among these 20,000
 - Difficult to see anything else! (six donated \$100)




Histogram: Example

- **Donations Dataset**
- Keep only the 989 who donated (eliminate \$0)
 - to see detail among those who made a gift
- Can now see the distribution of the gift amounts



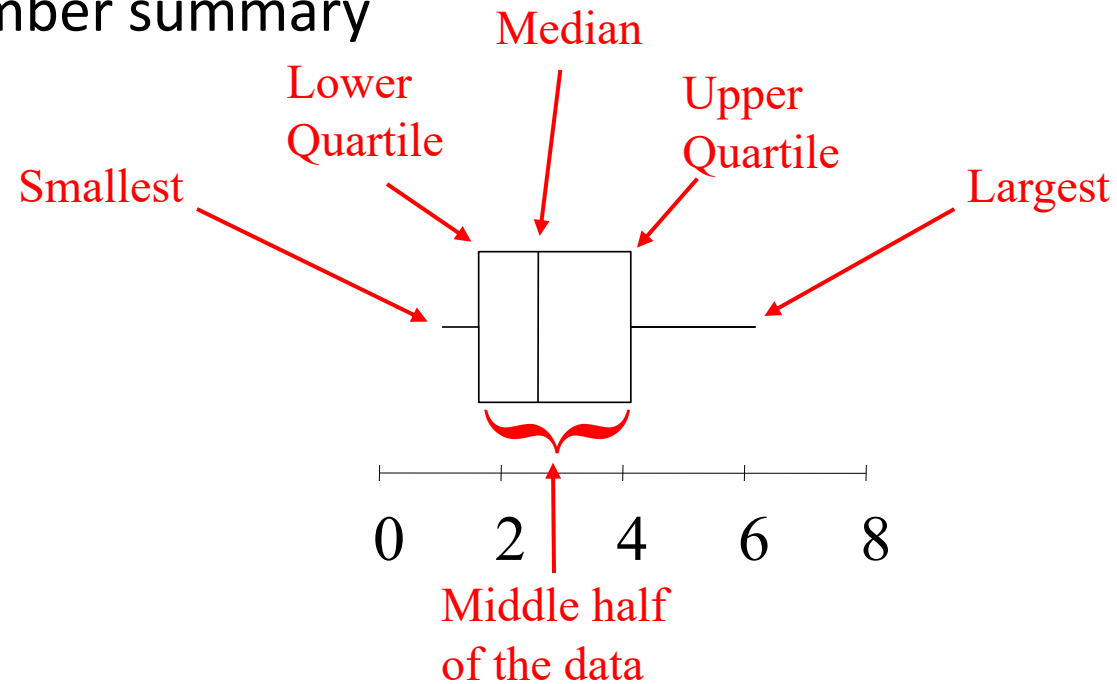
Histogram: Example

- **Donations Dataset**
- With so much data (989 people)
 - we can use smaller bars to see more details
- Note the “spikes” at \$5, 10, 15, 20, 25, and 50



Box Plot

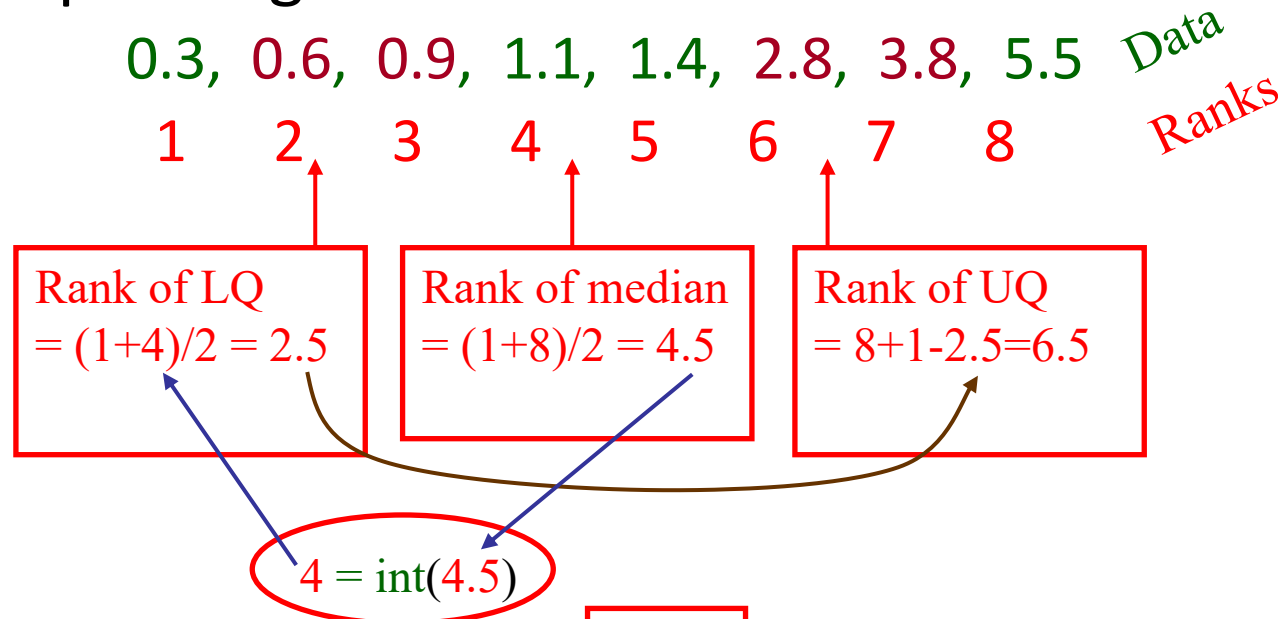
- Displays five-number summary



- Less detail than histogram
 - Easier to compare many groups

Box Plot: Example

- Spending rank ordered from smallest to largest



- LQ is $(0.6+0.9)/2 = 0.75$
- UQ is $(2.8+3.8)/2 = 3.3$

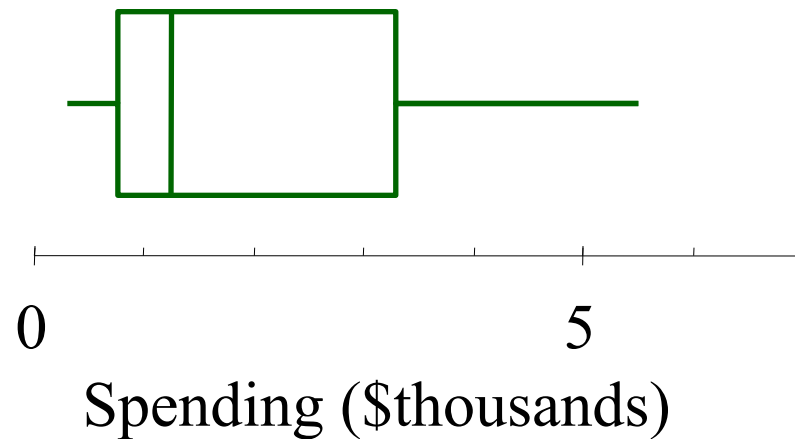
Box Plot: Example

- Five-number summary

0.3, 0.75, 1.25, 3.3, 5.5

Smallest, LQ, Median, UQ, Largest


- Box plot



- Shows some skewness (lack of symmetry)

Box Plot: Identifying Outliers


- Outliers are defined as observations, if any, either:
 - More than $UQ + 1.5(UQ - LQ)$, or
 - Less than $LQ - 1.5(UQ - LQ)$
- Outliers are far from the center of the distribution
 - and may be interesting as special cases




Box Plot: Example - CEO Compensation

- CEO compensation in technology companies
 - Detailed box plot identifies outliers
 - and identifies the most extreme non-outliers,
 - gives more detail than the (ordinary) box plot

Detailed Box Plot

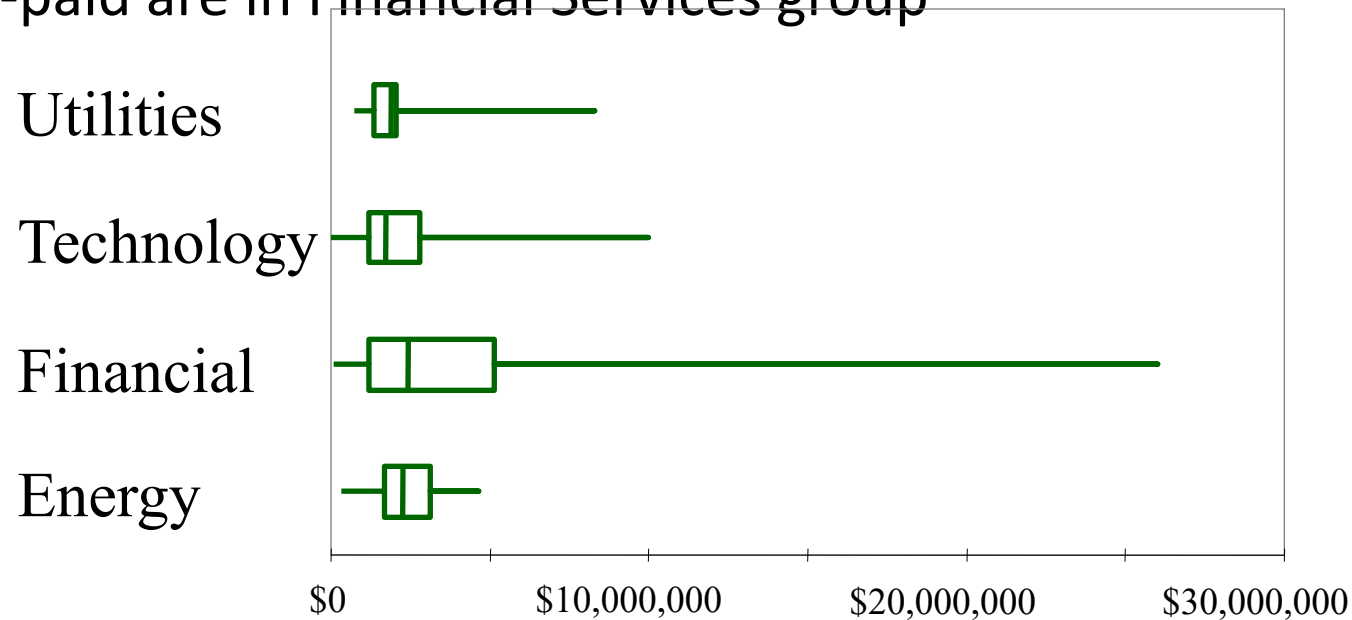


Box Plot




Box Plot: Example - CEO Compensation

- Box plots to compare firms within industry groups
 - Utilities group generally shows lower compensation
 - Highest-paid are in Financial Services group



Box Plot: Example - CEO Compensation

- Detailed box plots (with outliers and most extreme non-outliers named)



Central Tendency: Average or Mean

- Add the data, divide by n or N (the number of elementary units)

“X-bar”

$$\bar{X} = \frac{X_1 + X_2 + \dots + X_n}{n}$$

Sample average

“mu”

$$\mu = \frac{X_1 + X_2 + \dots + X_N}{N}$$

Population average

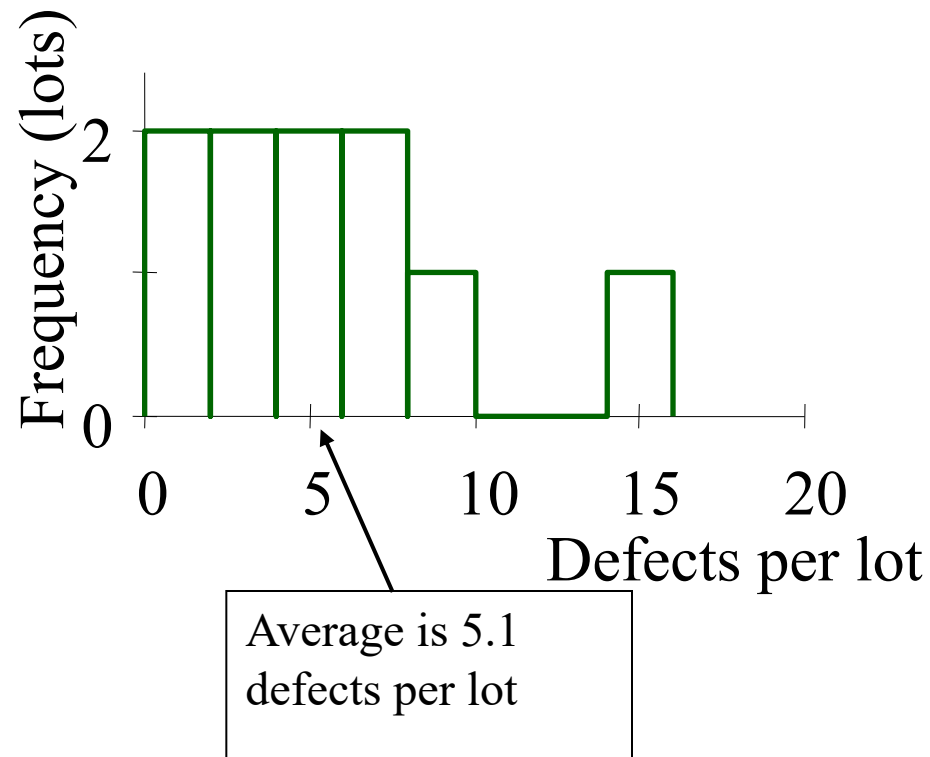
- Divides total equally. The only such summary
- A representative, central number (if data set is approximately normal)
- Summation notation
 - Σ is capital Greek sigma

$$\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$$

$$\mu = \frac{1}{N} \sum_{i=1}^N X_i$$

Example: Number of Defects

- Defects measured for each of 10 production lots
4, 1, 3, 7, 3, 0, 7, 14, 5, 9



Central Tendency: Median

- A representative, central number
 - If data set *has* a center
- Less sensitive to outliers than the average
- For skewed data, represents the “typical case” better than the average does
 - e.g., incomes
 - Average income for a country equally divides the total, which may include some very high incomes
 - Median income chooses the middle person (half earn less, half earn more), giving less influence to high incomes (if any)

Central Tendency: Median

- Also summarizes the data
- The **middle** one
 - Put data in order
 - Pick middle one (or average middle two if n is even)
 - Median (9, 4, 5) = Median(4, 5, 9) = 5
 - Median (9, 4, 5, 7) = Median (4, 5, 7, 9) = $(5+7)/2 = 6$
- Rank of the median is $(1+n)/2$
 - If $n=3$, rank is $(1+3)/2 = 2$
 - If $n=4$, rank is $(1+4)/2 = 2.5$ (so average 2nd and 3rd)
 - If $n=262$, rank is $(1+262)/2 = 131.5$

Central Tendency: Median Example

- Customers plan to spend (\$thousands)

3.8, 1.4, 0.3, 0.6, 2.8, 5.5, 0.9, 1.1

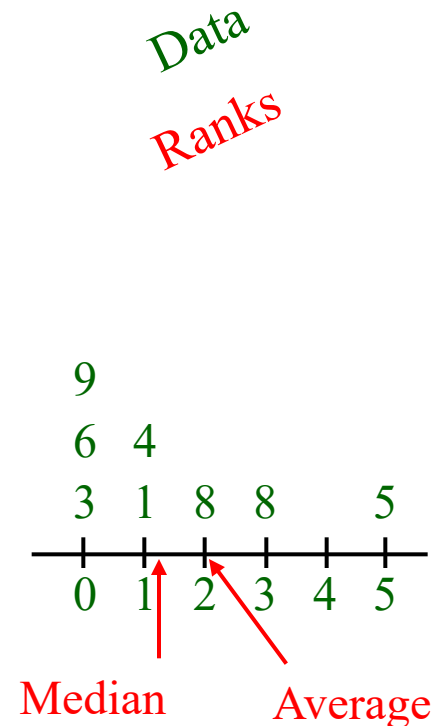
- Rank ordered from smallest to largest

0.3, 0.6, 0.9, 1.1, 1.4, 2.8, 3.8, 5.5

1 2 3 4 5 6 7 8

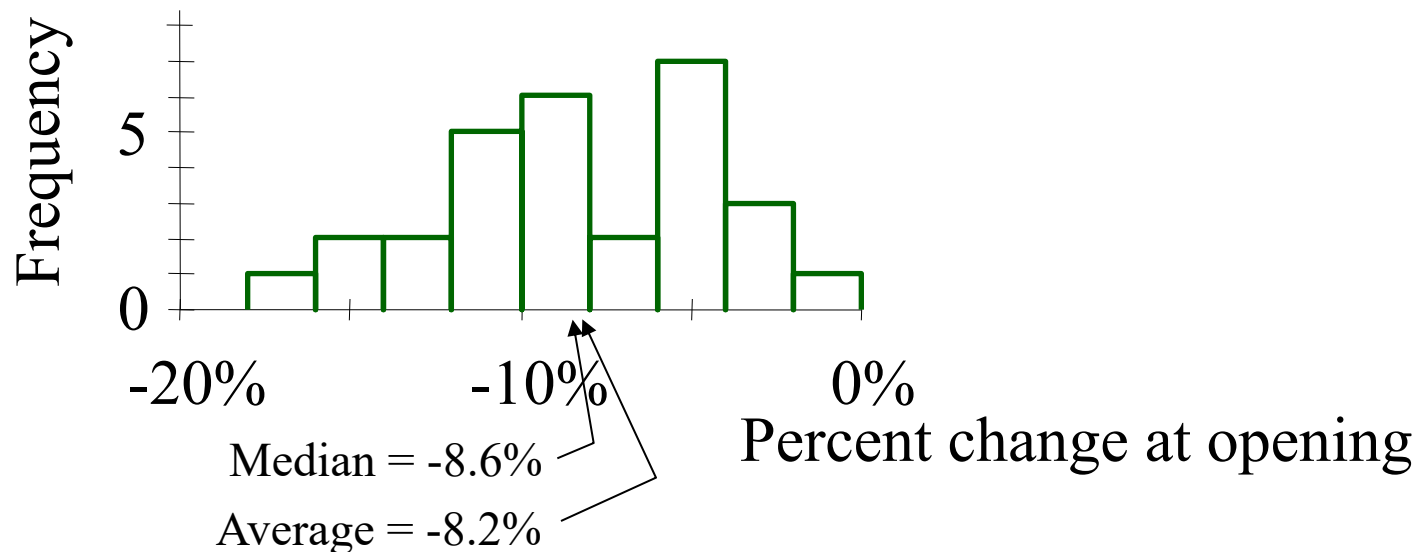
Rank of median
 $= (1+8)/2 = 4.5$

- Median is $(1.1+1.4)/2 = 1.25$
 - Smaller than the average, 2.05
 - Due to slight skewness?




Central Tendency: Median Example

- **Example: The Crash of 1987**
- Dow-Jones Industrials, stock-price changes as each stock began trading that fateful morning
- Fairly normal
- Mean and median are similar





Central Tendency: Median Example

- **Personal income of 100 people**
- Average is higher than median due to skewness



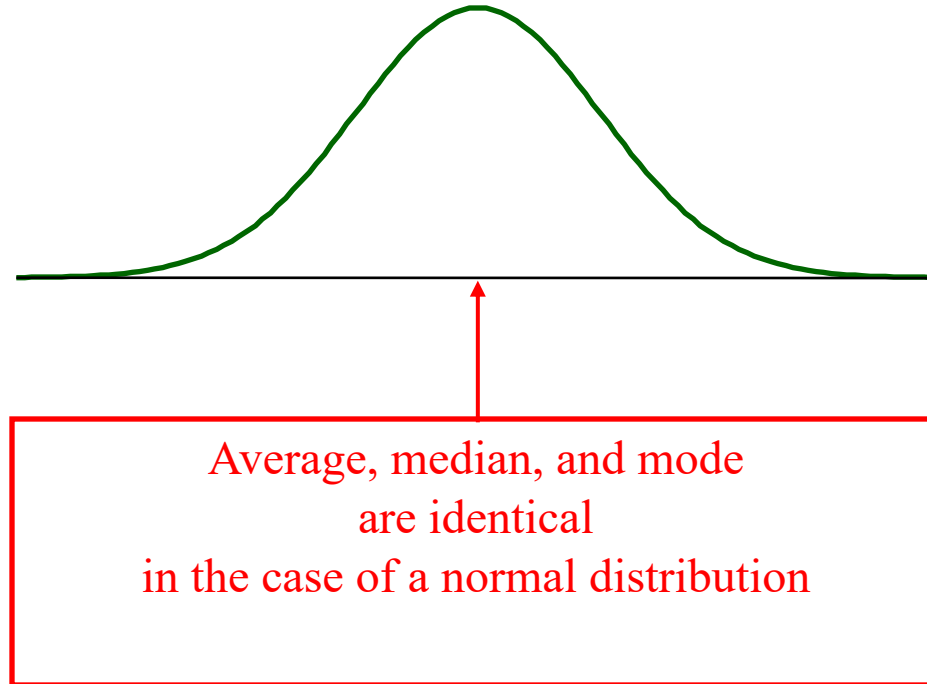
Central Tendency: Mode

- Also summarizes the data
- Most common data value
 - Middle of tallest histogram bar
- Problems:
 - Depends on how you draw histogram (bin width)
 - Might be more than one mode (two tallest bars)
- Good if most data values are “correct”
- Good for nominal data (e.g., elections)




Normal Distribution

- Average, median, and mode are **identical**
 - If the data come from a normal distribution



Skewed Distribution

- Average, median, and mode are **different**
 - The few large (or small) values influence the mean more than the median
 - The highest point is not in the center



Which summary to use?

- Average
 - Best for normal data
 - Preserves totals
- Median
 - Good for skewed data or data with outliers, provided you do not need to preserve or estimate total amounts
- Mode
 - Best for categories (nominal data).
 - The mode is the only summary computable for nominal data!

Which Summary? (continued)

- Average requires quantitative data (numbers)
- Median works with quantitative or ordinal
- Mode works with quantitative, ordinal, or nominal

	<u>Quantitative</u>	<u>Ordinal</u>	<u>Nominal</u>
Average	Yes	-	-
Median	Yes	Yes	-
Mode	Yes	Yes	Yes

Dispersion or Variability

- Also known as **dispersion, spread, uncertainty, diversity, risk**
- Example data: **2, 2, 2, 2, 2, 2, 2**
 - Variability = 0
- Example data: **1, 3, 2, 2, 1, 2, 3**
 - How much variability?
 - Look at how far each data value is from average $X = 2$: —
 - Deviations from average are **-1, 1, 0, 0, -1, 0, 1**
 - Variability should be between **0** and **1**

Dispersion: Examples

- Stock market, daily change, is uncertain
 - Not the same, day after day!
- Risk of a business venture
 - There are potential rewards, but possible losses
- Uncertain payoffs and risk aversion
 - Which would you rather have
 - \$1,000,000 for sure
 - \$0 or \$2,000,000, each outcome equally likely
 - Both have same average! (\$1,000,000)
 - Most would prefer the choice with less uncertainty

Dispersion: Standard Deviation S

- Measures variability by answering:
 - “Approximately **how far from average** are the data values?” (same measurement units as the data)
 - The square root of the average squared deviation
 - (dividing by $n-1$ instead of n for a sample)

- For a sample

$$S = \sqrt{\frac{(X_1 - \bar{X})^2 + (X_2 - \bar{X})^2 + \dots + (X_n - \bar{X})^2}{n-1}}$$

- For a population

“sigma”

$$\sigma = \sqrt{\frac{(X_1 - \mu)^2 + (X_2 - \mu)^2 + \dots + (X_N - \mu)^2}{N}}$$

Dispersion Example: Spending

- Customers plan to spend (\$thousands)

3.8, 1.4, 0.3, 0.6, 2.8, 5.5, 0.9, 1.1

- Average is 2.05. Sum of squared deviations is

$$(3.8-2.05)^2+(1.4-2.05)^2+\dots+(1.1-2.05)^2 = 23.34$$


- Divide by $8-1=7$ and take square root:

$$\sqrt{\frac{23.34}{7}} = \sqrt{3.334286} = 1.83 = \text{Standard deviation}$$

- Customers plan to spend about 1.83 (thousand, i.e., \$1,830) more or less than the average, 2.05.
 - Some plan to spend more, others less than average

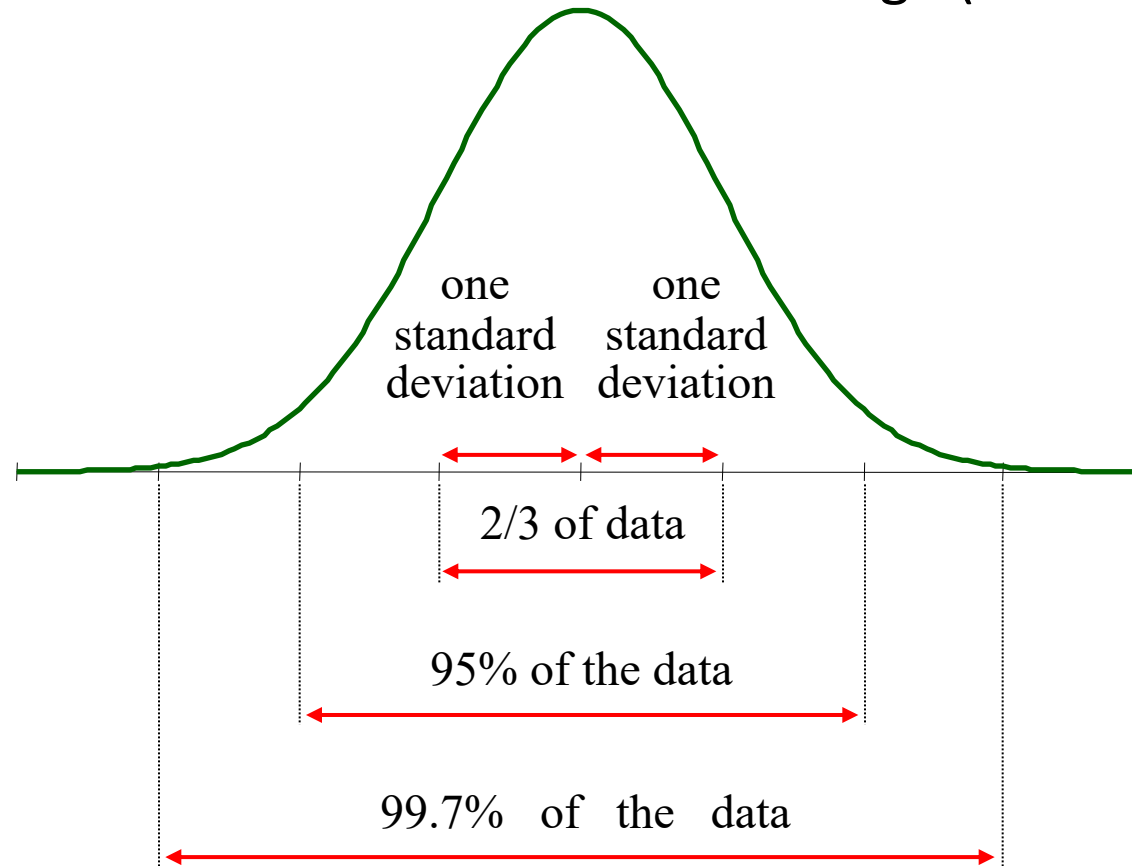
Dispersion Example: Spending

- On the histogram
 - Average is located near the center of the distribution
 - Standard deviation is a *distance* away from the average
 - Standard deviation is the *typical distance* from average



Dispersion in Normal Distribution

- For a **normal distribution only**
 - **2/3** of data within one standard deviation of the average (either above or below)
 - **95%** for 2 std. devs.
 - **99.7%** for 3




Dispersion in Skewed Distribution

- No simple rule for percentages within one, two, three standard deviations of the average
- Standard deviation retains its interpretation as the standard measure of
of

Typically how far the observations are from average

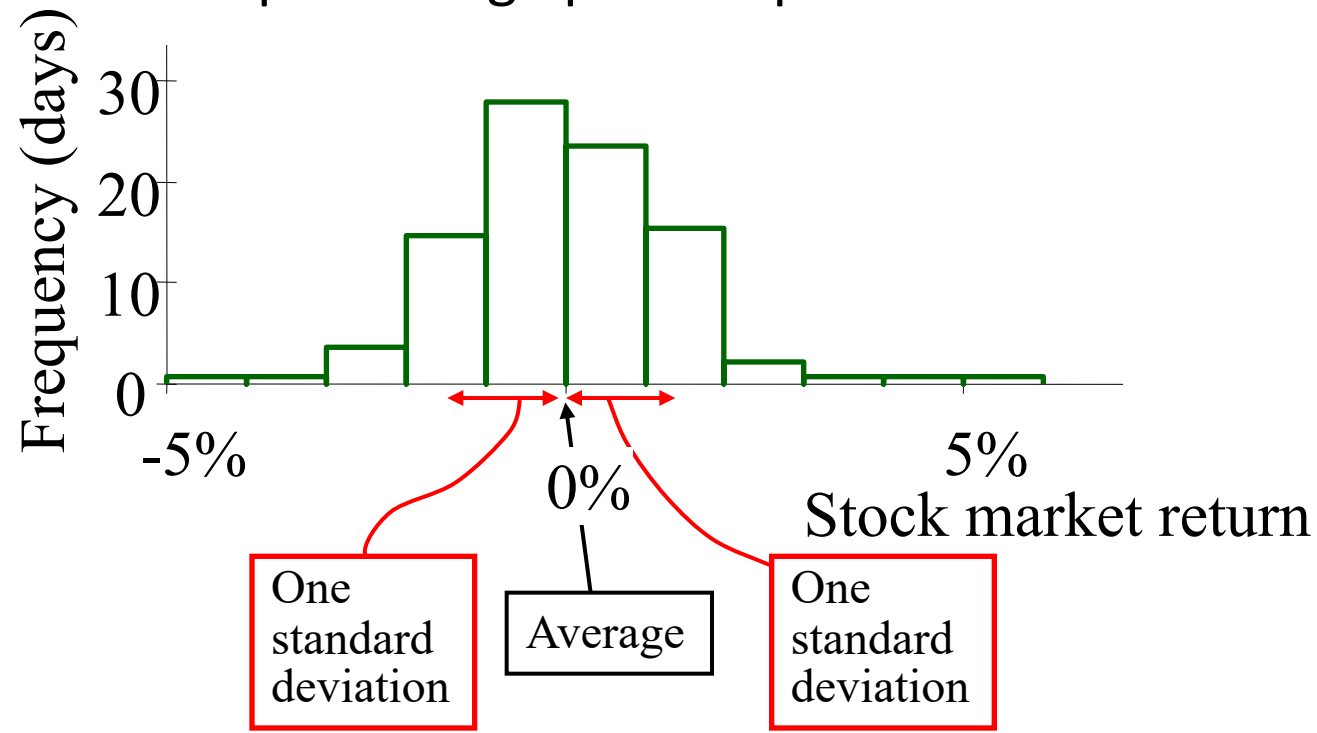
Dispersion Example: Quality Control Charts

- Control limits are often set at **3 standard deviations** from the average
- If the process is normally distributed, then
 - Over the long run, observations will stay within the control limits **99.7%** of the time
- If the process goes out of control, you will know



Dispersion in Example: The Stock Market

- Daily stock market returns, S&P500 index, first half of 2001. Standard deviation is **1.43%**
 - Average daily percent change: **-0.03%**
 - Typical day: about 1.5 percentage points up or down

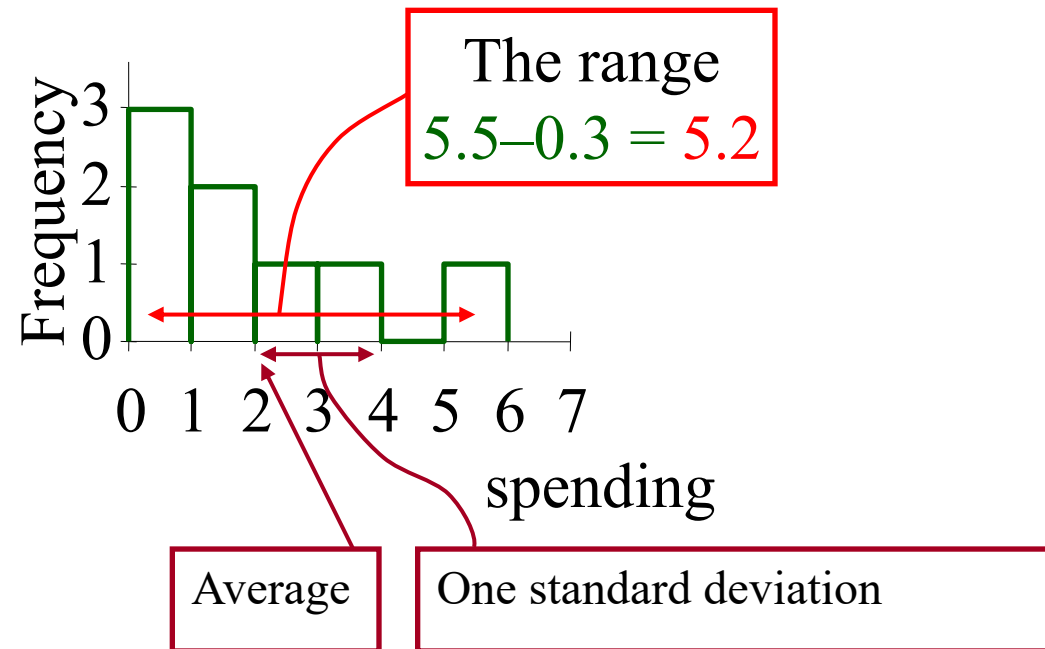


Dispersion: The Range

- The difference: **Largest – Smallest**
- Good features
 - Easy and fast to compute
 - Describe the data
 - Check the data: Is the range too big to be reasonable?
- Problem
 - Very sensitive to just two data values
 - Compare to standard deviation, which combines all data values

Dispersion Example: Spending

- \$Thousands: 3.8, 1.4, 0.3, 0.6, 2.8, 5.5, 0.9, 1.1
- The range is 5.2
 - larger than the standard deviation, 1.83



Dispersion: Coefficient of Variation

- A **relative** measure of variability
- The ratio: Standard deviation divided by average
 - For a sample: S/X
 - For a population: σ/μ
- No measurement units. A pure number. Answers:
 - “Typically, in percentage terms, how far are data values from average?”
- Useful for comparing situations of different sizes
 - To see how variability compares *after adjusting for size*

Statistical Inference: Random Sampling

Random Sampling


- The basis for statistical inference about a population based on a sample
 - Example: Build restaurant in neighborhood?
- Population: the collection of items you want to understand
 N items. Example: all people in neighborhood
- Sample: a smaller collection of population units
 n items. Example: 100 neighborhood residents who agree to be interviewed
 - Which 100 residents?
 - How to select?

Random Sampling

- Provides a foundation for statistical inference
- A random sample must satisfy:
 1. Each population unit must have an equal chance of being selected
 - This helps assure representation, because all units in the population are equally accessible
 2. Units must be selected independently of one another
 - This guarantees that each item to be selected will bring new, independent information
- Properties
 - Sample is representative of population (on average)
 - Statistical inference will use randomness of the sample

Random Samples of 5 from 36

- On 6 by 6 grid
 - 5 squares shaded in, selected at random, one at a time
 - Without replacement
 - Note that adjacent (touching) squares can be selected
 - To exclude them would make selection non-independent
 - Perhaps you see systematic patterns
 - They are not there by design, but by coincidence
 - But a “checkerboard” pattern would very likely *not* be random



Random Sampling

- Sampling Without Replacement
 - No unit may appear more than once in the sample
 - After a unit is selected, it is removed from the population before another population unit is drawn
- Sampling With Replacement
 - A unit can be represented more than once in the sample
 - After a unit is selected, it is “replaced” back in the population before another unit is drawn
- Census
 - A sample that consists of the *entire population*
 - Often too expensive
 - Even when affordable, may not be cost-efficient

Statistic and Parameter

- Sample Statistic

- Any number computed from sample data
 - A random variable. Known
 - Example: Average weekly food expenditures for 100 sampled residents
 - Random? Yes! Due to randomness of sample selection

- Population Parameter

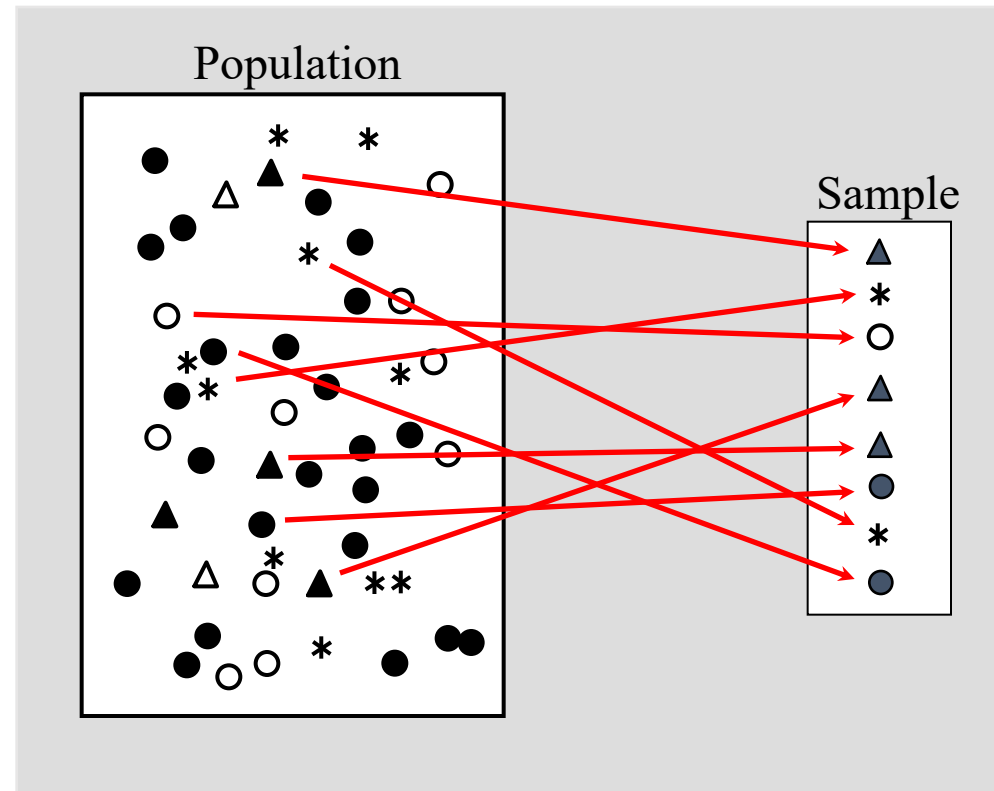
- Any number computed for the entire population
 - A fixed number. Unknown
 - Example: mean weekly food expenditures for all 77,386 residents
 - Do we ever know this? NO!
 - But we estimate it (with error)

Estimator and Estimate

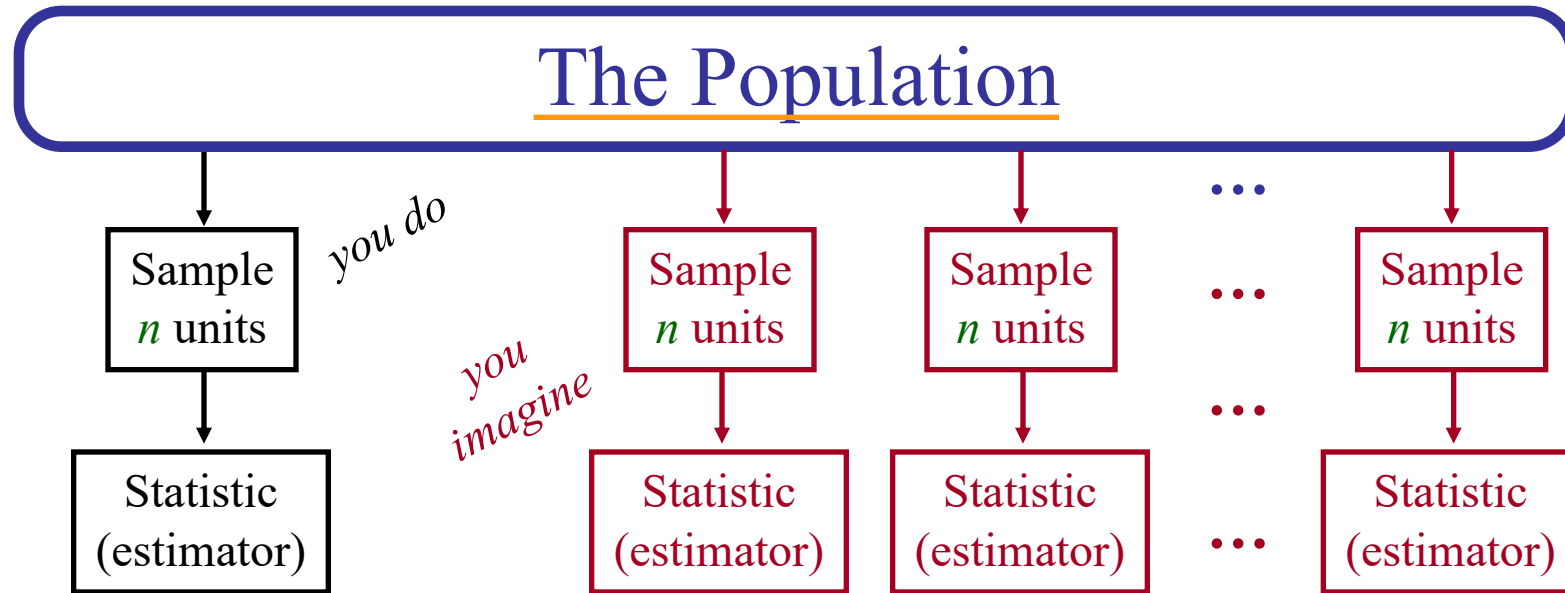
- Estimator
 - A sample statistic used to guess a population parameter
 - Example: Sample average for 100 selected residents is an estimator of the population mean of all 77,386 residents
- Estimate [WRONG! Estimators are usually wrong. Often useful anyway]
 - The actual number computed from the data
 - Example: Rs 33.91 is an estimate of neighborhood weekly food expenditures per person
- Estimation error
 - Estimator minus population parameter. Unknown
 - Example: $33.91 - 35.69 = -1.78$

Visualizing the Sampling Process

- Here is a fairly (but not perfectly) representative sample
 - Note that neither open triangle was selected



Visualizing the Sampling Process



- The process:
Sample n , compute statistic
is the same as the process:
Choose 1 from sampling distribution


A histogram of *these* imagined values represents the sampling distribution of *this* statistic

Central Limit Theorem


- Helps you find probabilities for an *average* \bar{X} of n independent individuals by giving you
 - The mean $\mu_{\bar{X}} = \mu_X = \mu$
 - The standard deviation $\sigma_{\bar{X}} = \sigma / \sqrt{n}$
 - The right to use the normal probability tables:
 - If n is large, then the average \bar{X} is approximately normal, even if individuals are skewed
- Works for totals also by giving you
 - The mean $\mu_{total} = n\mu$
 - The standard deviation $\sigma_{total} = \sigma\sqrt{n}$
 - The right to use the normal probability tables for *total*

Central Limit Theorem (continued)


- Individuals in population
 - Highly non-normal distribution
 - Mean μ , standard deviation σ
- Averages of $n = 3$ individuals
 - Non-normal, but less so
 - Same mean μ
 - Lower std. deviation:
- Averages of $n = 10$ individuals
 - Close to normal
 - Same mean μ
 - Lower std. deviation



$$\sigma_{\bar{X}} = \sigma / \sqrt{3}$$



$$\sigma_{\bar{X}} = \sigma / \sqrt{10}$$



Central Limit Theorem

- What good to us is $\sigma_{\bar{X}} = \sigma / \sqrt{n}$?
 - Why bother choosing a larger sample to estimate μ ?
 - Even sampling just $n = 1$ (or a few) we have an *unbiased* estimator that is, on average, equal to μ
 - But the error can be very large!
 - When we sample $n = 100$ or $n = 1,000$, what we gain for our trouble is the LOWER ERROR $\sigma_{\bar{X}} = \sigma / \sqrt{n}$ which says that our estimator, \bar{X} , is closer to the unknown population mean μ
 - It tells us how the variability of the sample average \bar{X} is related to the variability of individuals in the population
 - This will be useful soon in defining the *standard error*

Standard Error

- Definition: the **Estimated Standard Deviation** (of the sampling distribution) of a statistic

- Standard Error of the Average

$$S_{\bar{X}} = \frac{S}{\sqrt{n}}$$

- Indicates approximately how far the sample average \bar{X} is from the population mean μ
- Advantage: can be computed using sample data

- Standard Deviation of the Average

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

- Also indicates approximately how far the sample average \bar{X} is from the population mean μ
- Problem: cannot be computed without population parameters

Confidence Interval and Hypothesis Testing

Overview

- Confidence Interval
 - Computed from *data*
 - Has a *known* probability of including the *unknown* population parameter being estimated
- Statistical Inference
 - An exact probability statement about the population, based on sample data
- Confidence Level
 - The probability of including the population parameter within the confidence interval
 - **95%** is the usual standard. Also: **99%**, **99.9%**, **90%**

Approximate Confidence Interval

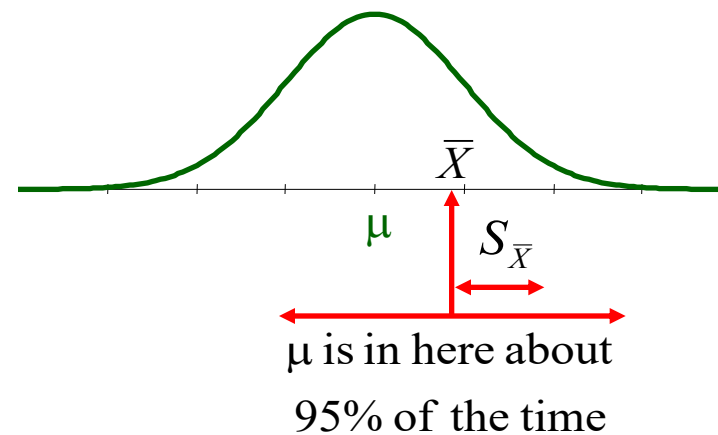
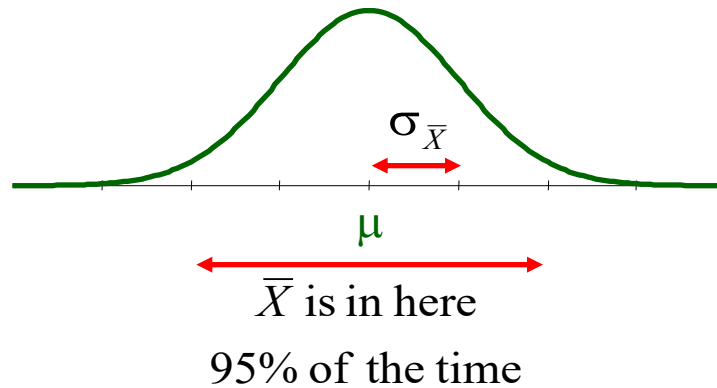
- Estimator ± 2 (Standard error of estimator)
 - We are 95% sure the population parameter is in this interval
- For population mean:

$$\bar{X} \pm 2S_{\bar{X}}$$

- Because an estimator (if normally distributed) has a 95% chance of being within **two** of its standard deviations from its mean

Why does it work?

- \bar{X} is within $2S_{\bar{X}}$ of its mean μ about 95% of the time
 - Because $S_{\bar{X}}$ is an estimator of $\sigma_{\bar{X}}$ (the standard deviation of the sampling distribution of $\sigma_{\bar{X}}$)



- This also says that μ is within $2S_{\bar{X}}$ of \bar{X} about 95% of the time

Confidence Interval for the Mean

- We are 95% sure that the *unknown* population mean μ is between

$$\bar{X} - tS_{\bar{X}} \quad \text{and} \quad \bar{X} + tS_{\bar{X}}$$

- where t is from the t table
- For smaller sample sizes ($n = 40$ or less)
 - t is larger than 1.96 (approximately 2)
 - Because we used the **Standard Error** ($S_{\bar{X}}$) as an estimator in place of the **Population Standard Deviation** ($\sigma_{\bar{X}}$) of the sample average

Assumptions

- Assumptions needed for validity of the Confidence Interval
 1. Data are a **RANDOM SAMPLE** from the population of interest
 - (So that the sample can tell you about the population)
 2. The sample average \bar{X} is approximately **NORMAL**
 - *Either* the data are normal (check the histogram)
 - *Or* the central limit theorem applies:
 - Large enough sample size n , distribution not too skewed
 - (So that the t table is technically appropriate)

Interpretation

- Interpreting a Confidence Interval
 - Population mean μ is fixed and unknown
 - Confidence interval is random and known
 - The probability is 0.95 that μ is between

$$\bar{X} - tS_{\bar{X}} \quad \text{and} \quad \bar{X} + tS_{\bar{X}}$$

- We are 95% sure that (for example) μ is between
57.86 and 78.74
- Lifetime average: about 95% of confidence intervals included μ .
 - You may never know which ones!

Hypothesis Testing

- Deciding between two possibilities based on data
 - e.g., “Is it real? Or is it just coincidence?”
- Hypothesis: a statement about the population
 - e.g., More than 30% of customers recognize our product
 - e.g., You will win the election
 - e.g., Strategy Z will make you rich in the stock market
- Note: a hypothesis is either **TRUE** or **FALSE**
 - Even with data, you may never know for sure, because of *randomness*

Example: Dishwasher Detergent

- From a box of Cascade: A hypothesis (italics added)
 - “Individual packages of Cascade may weigh slightly more or less than the marked weight due to normal variations incurred with high speed packaging machines, but *each day’s production of Cascade will average slightly above the marked weight*”
- This hypothesis is either *true* or *false*
 - We do not know which
 - The package claims that it is *true*
 - We could test it, e.g., by weighing a sample of boxes

Example: Pure Randomness

- Coin tossing: Probability $1/2$ for Heads or Tails
 - 2 tosses: Probability $1/4$ for each of HH, HT, TH, TT
 - 3 tosses: Probability $1/8$ for each of HHH, HHT, HTH, HTT, THH, THT, TTH, TTT
 - HHH is no less likely than any other *particular* sequence
 - 10 tosses: probability $1/1,024$ of **HHHHHHHHHH**
- If you toss **1,000** times, it is not surprising to find ten Heads in a row *somewhere* in the sequence
- But if you toss only **10** times and find ten Heads
 - **You will rightly be suspicious!!!!**
 - Reject the hypothesis that this is an ordinary coin?

Null and Research Hypothesis

Null Hypothesis H_0	Research Hypothesis H_1
The Default . Accept unless disproven	Has burden of proof . Requires convincing evidence
Often specific	Often general
Often randomness	Often “ your theory ”
A, B are independent (no connection)	A, B are dependent (related)
Ad has no effect on purchase	Ad works
$\mu = \mu_0$ Population mean equals reference value	$\mu \neq \mu_0$ Population mean does not equal reference value
Long-run mean oven temperature, μ , equals the desired setting $\mu_0=325^0$	Long-run mean oven temperature, μ , does not equal the desired setting $\mu_0=325^0$

Interpretation

- If you **reject H_0** and **accept H_1**
 - H_0 could **not** reasonably have produced the data
 - Either
 - H_1 is true, or
 - H_0 is true, but you made a **TYPE I ERROR**
 - Happens **5%** of the time when H_0 is true
 - A **strong conclusion**
 - A **significant** result
- You have earned a “license to explain” the observed difference

Reject H_0 when
 H_0 is really true

Interpretation (continued)

- If you **accept H_0**
 - H_0 could reasonably have produced the data
 - Either
 - H_0 is true, or
 - H_1 is true, but you made a **TYPE II ERROR**
 - Difficult to control
 - This error is possible, and is very likely if μ_0 is close to μ
 - A **weak conclusion**
 - **Not** a significant result
 - Little or nothing to explain
 - The observed difference might just be random

Accept H_0 when
 H_1 is really true

Errors in Hypothesis Testing

		Your Decision	
		Accept Null Hypothesis H_0	Accept Research Hypothesis H_1
The Truth	Null Hypothesis H_0	Yay! Correct Decision	Whoops! Type I Error [level 0.05]
	Research Hypothesis H_1	Whoops! Type II Error [not easily controlled]	Yay! Correct Decision

p -Values

- The smallest test level that is significant
 - Often provided by computer analysis
 - e.g., $p = 0.0297$ ←
- Tells the strength of the evidence against H_0
 - Small p value says data unlikely to come from H_0
 - Reject H_0 if p is small enough
 - Not significant ($p > 0.05$)
 - Significant ($p < 0.05$)
 - Highly significant ($p < 0.01$)
 - Very highly significant ($p < 0.001$)
 - What if $p = 0.374$? *Not significant* because $p > 0.05$

Significant, because
it is less than 0.05

Says that if H_0 were true,
your data would
not be surprising

Hypothesis Testing for Ordinal data

Nonparametric Methods

- Statistical Procedures for Hypothesis Testing that ***do Not Require a Normal Distribution***
 - Because they are based on ***Counts*** or ***Ranks***
 - A ***Random Sample*** is still required
- The Nonparametric Approach Based on ***Counts***
 - Count the number of times some event occurs
 - Use the binomial distribution to decide whether this count is reasonable or not under the null hypothesis
- The Nonparametric Approach Based on ***Ranks***
 - Replace each data value with its rank (***1, 2, 3, ...***)
 - Use formulas and tables created for testing ranks

Advantages and Disadvantages

- Advantages of Nonparametric Testing

- No need to assume normality
- Avoids problems of transformation (e.g., interpretation)
- Can be used with ordinal data
 - Because ranks can be found
- Can be much more efficient than parametric methods when distributions *are not* normal

- Disadvantage of Nonparametric Testing

- Less statistically efficient than parametric methods when distributions *are* normal
 - Often, this loss of efficiency is slight

Testing the Median

- Testing the Median against a Known Reference Value
 - **Without** assuming a normal distribution
 - Note: The number of sample data values below a continuous population's median follows a binomial distribution where $p = 0.5$ and n is the sample size
 - The **Sign Test**
 1. Find the **modified sample size m** , the number of data values **different from** the reference value θ_0
 2. Find the limits in the table for this modified sample size
 3. Count how many data values fall below the reference value
 4. Significant if the count (step 3) is outside the limits (step 2)

Sign Test: Hypotheses, Assumption

- Sign Test for the Median (for a Continuous Population Distribution)
 - $H_0: \theta = \theta_0$ and $H_1: \theta \neq \theta_0$
 - where θ is the unknown population median and θ_0 is the (known) reference value being tested
- Sign Test for the Median (in General)
 - H_0 : The probability of being above θ_0 **equals** the probability of being below θ_0 in the population
 - H_1 : These probabilities are **not equal**
 - where θ_0 is the (known) reference value being tested
- Assumption required:
 - The data set is a random sample from the population

Example: Family Income

- Comparing Local to National Family Income
 - Survey median is \$70,547, based on $n = 25$ families
 - National median is \$27,735
 - This is the reference value θ_0
 - Performing the sign test
 - Modified sample size is $m = 25$, since all sampled families have incomes different from the reference value
 - Limits from the table are 8 and 17 (for testing at the 5% level with $m = 25$)
 - There are 6 families with income below the reference value
 - Since 6 falls outside the limits (from 8 to 17),

Median family income in the community is significantly higher than the national median

Example: Family Income

Income of Sampled families	Is = 27,735	Is > 27,735	Is < 27,735
\$39,465.00	0	1	0
\$80,806.00	0	1	0
\$267,525.00	0	1	0
\$163,819.00	0	1	0
\$58,525.00	0	1	0
\$25,479.00	0	0	1
\$29,341.00	0	1	0
\$96,270.00	0	1	0
\$85,421.00	0	1	0
\$56,240.00	0	1	0
\$14,706.00	0	0	1
\$54,348.00	0	1	0
\$7,081.00	0	0	1
\$137,414.00	0	1	0
\$16,477.00	0	0	1
\$5,921.00	0	0	1
\$187,445.00	0	1	0
\$83,414.00	0	1	0
\$36,346.00	0	1	0
\$19,605.00	0	0	1
\$156,681.00	0	1	0
\$138,933.00	0	1	0
\$70,547.00	0	1	0
\$81,802.00	0	1	0
\$78,464.00	0	1	0
\$70,547.00	0	19	6

Testing for Nominal data

Chi-Squared Tests

- Hypothesis Tests for Qualitative Data

- Categories instead of numbers
- Based on counts
 - the number of sampled items falling into each category
- Chi-squared statistic
 - Measures the difference between **Actual** counts and **Expected** counts (as expected under the null hypothesis H_0)

$$\text{Chi - squared statistic} = \text{Sum of } \frac{(\text{Observed Count} - \text{Expected Count})^2}{\text{Expected Count}} = \sum \frac{(O_i - E_i)^2}{E_i}$$

where the sum extends over all categories or combinations of categories

- **Significant** if the chi-squared statistic is large enough

Summarizing Qualitative Data

- Use Counts and Percentages, for Example:
 - How many (of people sampled) prefer your product?
 - What percentage of sophisticated product users said they would like to purchase an upgrade?
 - What percentage of products produced today are both
 - designed for export
 - *and* imperfect

Testing Population Percentages

- Testing if Population Percentages are Equal to Known Reference Values
 - The chi-squared test for equality of percentages
 - Could a table of observed counts have reasonably come from a population with known percentages (the reference values)?
 - **The data**: A table indicating the **count** for each category for a single qualitative variable
 - **The hypotheses**:
 - H_0 : The population percentages are **equal** to a set of known, fixed reference values
 - H_1 : The population percentages are **not equal** to a set of known, fixed reference values

Testing Percentages (continued)

- ***The expected counts:***
 - For each category, multiply the population reference proportion by the sample size, n
- ***The assumptions:***
 1. Data set is a random sample from the population of interest
 2. At least 5 counts are expected in each category
- ***The chi-squared statistic:***

$$\text{Sum of } \frac{(\text{Observed Count} - \text{Expected Count})^2}{\text{Expected Count}} = \sum \frac{(O_i - E_i)^2}{E_i}$$

- ***The degrees of freedom:*** Number of categories minus 1
- ***The test result: Significant*** if the chi-squared statistic is larger than the critical value from the table

Independence

- Two Qualitative Variables are *Independent* if:
 - knowledge about the value (i.e., category) of one variable does not help you predict the other variable
- For Example: Background Sales Information
 - The two variables are
 - where the customer lives, and
 - the customer's favorite product
 - Independence would say that
 - customers tend to have the same pattern (distribution) of favorite products, regardless of where they live, and
 - where customers live is not associated with which product is their favorite

Testing for Association

- The chi-squared test for independence
 - **The data:** A table indicating the counts for each combination of categories for two qualitative variables
 - **The hypotheses:**
 - H_0 : The two variables are **independent** of one another
 - H_1 : The two variables are **associated**; they are **not independent**
 - **The expected table:**

$$\text{Expected count} = \frac{\left(\begin{array}{c} \text{Count for category} \\ \text{for one variable} \end{array} \right) \left(\begin{array}{c} \text{Count for category} \\ \text{for other variable} \end{array} \right)}{n}$$

- Tells you what the counts *would have been*, on average, if the variables were independent and there were no randomness

Testing Association (continued)

- **The assumptions:**
 1. Data set is a random sample from the population of interest
 2. At least 5 counts are expected in each combination of categories

- **The chi-squared statistic:**

•
$$\text{Sum of } \frac{(\text{Observed Count} - \text{Expected Count})^2}{\text{Expected Count}} = \sum \frac{(O_i - E_i)^2}{E_i}$$

where the sum extends over all combinations of categories

- **The degrees of freedom:**

$$\left(\begin{array}{c} \text{Number of categories} \\ \text{for first variable} \end{array} - 1 \right) \left(\begin{array}{c} \text{Number of categories} \\ \text{for second variable} \end{array} - 1 \right)$$

- **The test result: Significant** if the chi-squared statistic is larger than the critical value from the table

Example: Market Segmentation

- Data: Rowing Machine Purchases
 - Which model rowing machine was purchased?
 - Basic, Designer, or Complete
 - Which type of customer purchased it?
 - Practical or Impulsive

Observed Counts			
	Practical	Impulsive	Total
Basic	22	25	47
Designer	13	88	101
Complete	54	19	73
Total	89	132	221

Example (continued)

- Overall Percentages

- Divide each count by the overall total, 221

- e.g., 10.0% = 22/221 were practical customers who purchased basic machines

- Note: impulsive customers bought most designer machines, while practical customers bought most complete machines

Overall Percentages			
	Practical	Impulsive	Total
Basic	10.0%	11.3%	21.3%
Designer	5.9%	39.8%	45.7%
Complete	24.4%	8.6%	33.0%
Total	40.3%	59.7%	100.0%

Example (continued)

- Percentages by Model

- Divide each count by the total for its model type

- e.g., $22/47 = 46.8\%$ of the 47 basic machines were purchased by practical customers
- Note: Practical customers make up 40.3% of all customers, but represent 74.0% of complete-machine purchasers

	Practical	Impulsive	Total
Basic	46.8%	53.2%	100.0%
Designer	12.9%	87.1%	100.0%
Complete	74.0%	26.0%	100.0%
Total	40.3%	59.7%	100.0%

Example (continued)

- Percentages by Customer Type

- Divide each count by the total for its customer type

- e.g., $22/89 = 24.7\%$ of the 89 practical customers purchased basic machines

- Note: Of all machines, 33.0% are complete; looking only at practical-customer purchases, 60.7% are complete

	Practical	Impulsive	Total
Basic	24.7%	18.9%	21.3%
Designer	14.6%	66.7%	45.7%
Complete	60.7%	14.4%	33.0%
Total	100.0%	100.0%	100.0%

Example (continued)

- Expected Counts
 - Multiply row total by column total, then divide by the overall total of 221
 - e.g., if there were no association between type of customer and type of machine, we would have expected to find $89 \times 47 / 221 = 18.93$ basic machines purchased by practical customers

	Practical	Impulsive	Total
Basic	18.93	28.07	47.00
Designer	40.67	60.33	101.00
Complete	29.40	43.60	73.00
Total	89.00	132.00	221.00

Example (continued)

- Chi-Squared Statistic (do not use the *total* row or column)

66.8

- Degrees of freedom

$$(\text{Rows} - 1)(\text{Column} - 1) = (3 - 1)(2 - 1) = 2$$

- Result

- If there were no association, such a large chi-squared value would be highly unlikely

The association between customer type and model purchased is very highly significant ($p < 0.001$)

Testing for Differences among Groups - *ANOVA*

ANOVA

- **Analysis of Variance**
 - A general framework for statistical hypothesis testing
 - Based on comparing sources of variability
 - Uses an ***F test*** based on an ***F statistic*** (a ratio of two variance measures) to perform each hypothesis test
 - One-Way ANOVA
 - To test whether or not the averages from several different situations are different from one another
 - Generalizes the ***t*** test to more than two samples

Comparing Many Samples

- Use Box Plots to Explore the Data
 - Do box plots look reasonable?
 - Do centers appear different from one box plot to another? (This is what one-way ANOVA will test)
 - Is the variability reasonably constant from one box plot to another?
 - Important because ANOVA will assume that these variabilities are equal in the population
- Use the F Test to see if the averages show any significant differences

One-Way ANOVA: Data Set

- k Groups, Sample Sizes n_1, n_2, \dots, n_k

Grand average is

$$\bar{X} = (n_1\bar{X}_1 + n_2\bar{X}_2 + \dots + n_k\bar{X}_k) / (n_1 + n_2 + \dots + n_k)$$

- Data Set

	Sample 1	Sample 2	...	Sample k
	$X_{1,1}$	$X_{2,1}$...	$X_{k,1}$
	$X_{1,2}$	$X_{2,2}$...	$X_{k,2}$
	•	•	•	•
	•	•	•	•
	•	•	•	•
	X_{1,n_1}	X_{2,n_2}	...	X_{k,n_k}
Average	\bar{X}_1	\bar{X}_2	...	\bar{X}_k
Standard deviation	S_1	S_2	...	S_k
Sample size	n_1	n_2	...	n_k

One-Way ANOVA: Background

- Sources of Variation
 - Between-sample variability (one sample to another)
 - Within-sample variability (inside each sample)
- Assumptions
 - Data set is k random samples from k populations
 - Each population is normal and std. devs. are identical
- Hypotheses
 - $H_0: \mu_1 = \mu_2 = \dots = \mu_k$ (the means are equal)
 - $H_1: \mu_i \neq \mu_j$ for at least one pair of populations
(the means are **not** all equal)

One-Way ANOVA: Computations

- Between-Sample Variability ($k - 1$ degrees of freedom)

$$\frac{n_1(\bar{X}_1 - \bar{X})^2 + n_2(\bar{X}_2 - \bar{X})^2 + \dots + n_k(\bar{X}_k - \bar{X})^2}{k - 1} = \frac{1}{k - 1} \sum_{i=1}^k n_i (\bar{X}_i - \bar{X})^2$$

- Within-Sample Variability ($n - k$ degrees of freedom)

$$\frac{(n_1 - 1)S_1^2 + (n_2 - 1)S_2^2 + \dots + (n_k - 1)S_k^2}{n - k} = \frac{1}{n - k} \sum_{i=1}^k (n_i - 1)S_i^2$$

- F Statistic ($k - 1$ and $n - k$ degrees of freedom)

$$F = \frac{\text{Between sample variability}}{\text{Within sample variability}}$$

One-Way ANOVA: Results

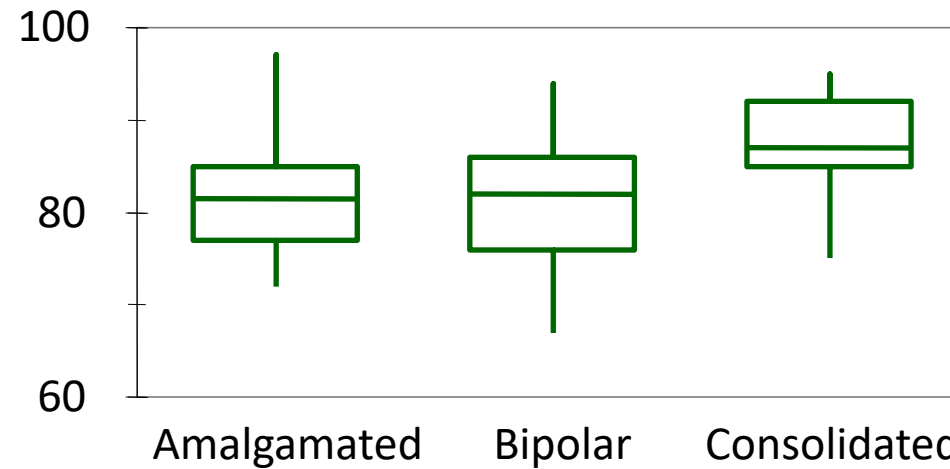
- If F Statistic is **Less** than the critical table value (So that $p > 0.05$)
 - Accept null hypothesis H_0 as a reasonable possibility
 - Do **not** accept the research hypothesis H_1
 - The sample averages are **not significantly different** from one another
 - The observed differences among sample averages could reasonably be due to random chance alone
 - The result is **not statistically significant**
 - (All the above statements are equivalent to each other)

One-Way ANOVA: Results (cont'd)

- If F Statistic is **More** than the critical table value (So that $p < 0.05$)
 - Accept the research hypothesis H_1
 - **Reject** the null hypothesis H_0
 - The sample averages are **significantly different** from one another
 - The observed differences among sample averages could **not** reasonably be due to random chance alone
 - The result is **statistically significant**
 - (All the above statements are equivalent to each other)

Example: Quality Scores

- Box Plots Show some Possible Differences



- Consolidated's scores appear higher, on average, than the others
 - Is this significant?
 - Or, if suppliers were of equal quality, could we reasonably see this large a difference due to random chance alone?

Example (continued)

- The Apparent Differences are ***highly significant***
 - $F = 5.897$ with 2 and 55 degrees of freedom
 - $p = 0.0048$ is less than 0.01
 - Differences as large as these could not reasonably be due to random chance alone
- But *which* Suppliers are Different?

Least-Significant-Difference Test

- t tests for the average difference between pairs of samples
 - To see which pairs are different
 - Can **only** be used if the F test is **significant**
 - If F test is **not** significant, then you already know that **no** pairs are significantly different
 - Based on
 - The average difference for pair i, j
 - The standard error of this average difference
 - The number of degrees of freedom ($n - k$)

$$\sqrt{(\text{within sample variability}) \left(\frac{1}{n_i} + \frac{1}{n_j} \right)}$$

Example: t Tests

- Amalgamated (82.056) and Bipolar (80.667)
 - Average difference is -1.389 , Std. error = 2.170
 - Confidence interval is from -5.64 to 2.86 , $t = -0.640$
 - **Not significant**
- Amalgamated (82.056) and Consolidated (87.684)
 - Average difference is 5.628 , Std. error = 2.222
 - Confidence interval is from 1.27 to 9.98 , $t = 2.533$
 - **Significant**
- Bipolar (80.667) and Consolidated (87.684)
 - Average difference is 7.017 , Std. error = 2.139
 - Confidence interval is from 2.82 to 11.21 , $t = 3.281$
 - **Significant**

Two-Way ANOVA

- Two Factors Define the Situations
 - e.g., First Factor is **Shift** (Day, Night, or Swing)
 - e.g., Second Factor is **Supplier** (A, B, or C)
 - Production Quality is measured
- Main Effect For Each Factor
 - Are there differences from one **Shift** to another?
 - Are there differences from one **Supplier** to another?
- Interaction
 - Are there particular combinations of **Shift** and **Supplier** that work better or worse than expected for the overall abilities of that **Shift** and that **Supplier**?

Other ANOVA Designs

- Three-Way ANOVA
 - Three factors
 - Main effect of each factor
 - Two-way interactions
 - Three-way interaction
- Analysis of Covariance (ANCOVA)
 - Combines regression analysis with ANOVA
 - Are slopes different in several regression situations?
 - Are group means different, after adjusting for a variable?
- Multivariate ANOVA (MANOVA)
 - More than one quantitative response variable

Testing for Relationships
Correlation and Regression

Bivariate Data: Relationships

Examples of relationships:

- Sales and earnings
- Cost and number produced
- Microsoft and the stock market
- Effort and results

- **Scatterplot**

- A picture to explore the relationship in bivariate data

- **Correlation r**

- Measures strength of the relationship (from -1 to 1)


- **Regression**

- Predicting one variable from the other

Interpreting Correlation


- $r = 1$

- A perfect straight line tilting *up* to the right




- $r = 0$

- No overall tilt
- No relationship?




- $r = -1$

- A perfect straight line tilting *down* to the right




Example: Internet Site Ratings

- Time Spent vs. Internet Pages Viewed
 - Two measures of the abilities of 25 Internet sites
 - At the top right are eBay, Yahoo!, and MSN
 - Correlation is $r = 0.964$
 - Very strong positive association (since r is close to 1)
 - Linear relationship
 - Straight line with scatter
 - Increasing relationship
 - Tilts up and to the right




Example: Merger Deals

- Dollars vs. Deals
 - For mergers and acquisitions by investment bankers
 - 244 deals worth \$756 billion by Goldman Sachs
 - Correlation is $r = 0.419$
 - Positive association
 - Linear relationship
 - Straight line with scatter
 - Increasing relationship
 - Tilts up and to the right



Example: Mortgage Rates & Fees

- Interest Rate vs. Loan Fee
 - For mortgages
 - If the interest rate is lower, does the bank make it up with a higher loan fee?
 - Correlation is $r = -0.890$
 - Strong negative association
 - Linear relationship
 - Straight line with scatter
 - Decreasing relationship
 - Tilts **down** and to the right



Example: The Stock Market

- Today's vs. Yesterday's Percent Change

- Is there momentum?


- If the market was up yesterday, is it more likely to be up today? Or is each day's performance independent?

- Correlation is $r = 0.11$

- A weak relationship?

- No relationship?

- Tilt is neither up nor down



Example: Stock Options

- Call Price vs. Strike Price

- For stock options


- “Call Price” is the price of the option contract to buy stock at the “Strike Price”
 - The right to buy at a lower strike price has more value

- A nonlinear relationship

- Not a straight line:
A curved relationship


- Correlation $r = -0.895$

- A negative relationship:
Higher strike price goes
with lower call price



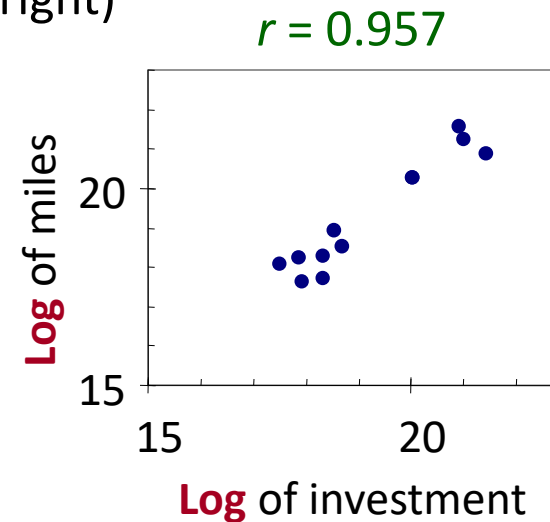
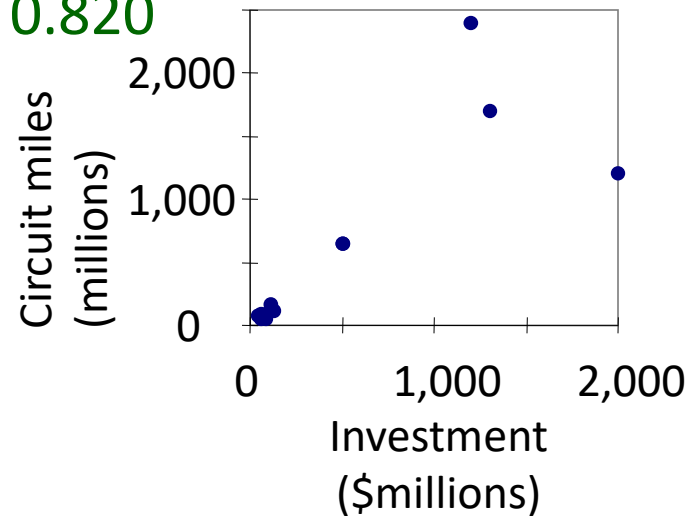
Example: Maximizing Yield

- Output Yield vs. Temperature
 - For an industrial process
 - With a “best” optimal temperature setting
 - A nonlinear relationship
 - Not a straight line:
 - A curved relationship
 - Correlation $r = -0.0155$
 - r suggests no relationship
 - But relationship is strong
 - It tilts neither
 - up nor down




Example: Telecommunications

- Circuit Miles vs. Investment (lower left)
 - For telecommunications firms
 - A relationship with unequal variability
 - More vertical variation at the right than at the left
 - Variability is stabilized by taking logarithms (lower right)
 - Correlation $r = 0.820$



Example: Bond Coupon and Price



- Price vs. Coupon Payment
 - For trading in the bond market
 - Bonds paying a higher coupon generally cost more
 - Two clusters are visible
 - Ordinary bonds (value is from coupon)
 - Inflation-indexed bonds (payout rises with inflation)
 - Correlation $r = 0.950$
 - for all bonds
 - Correlation $r = 0.994$
 - Ordinary bonds only



Example: Cost and Quantity

- Cost vs. Number Produced


- For a production facility
 - It usually costs more to produce more
- An outlier is visible
 - A disaster (a fire at the factory)
 - High cost, but few produced



Example: Salary and Experience

- Salary vs. Years Experience
 - For $n = 6$ employees
 - Linear (straight line) relationship
 - Increasing relationship
 - higher salary generally goes with higher experience
 - Correlation $r = 0.8667$


<u>Experience</u>	<u>Salary</u>
15	30
10	35
20	55
5	22
15	40
5	27



Mary earns \$55,000 per year, and has 20 years of experience

The Least-Squares Line $Y=a+bX$


- Summarizes bivariate data: Predicts Y from X
 - with smallest errors (in vertical direction, for Y axis)
 - Intercept is 15.32 salary (at 0 years of experience)
 - Slope is 1.673 salary (for each additional year of experience, on average)



Predicted Values and Residuals

- Predicted Value comes from Least-Squares Line
 - For example, Mary (with 20 years of experience) has predicted salary $15.32 + 1.673(20) = 48.8$
 - So does *anyone* with 20 years of experience
- Residual is actual Y minus predicted Y
 - Mary's residual is $55 - 48.8 = 6.2$
 - She earns about \$6,200 more than the predicted salary for a person with 20 years of experience
 - A person who earns *less* than predicted will have a negative residual

Predicted and Residual (continued)



Standard Error of Estimate

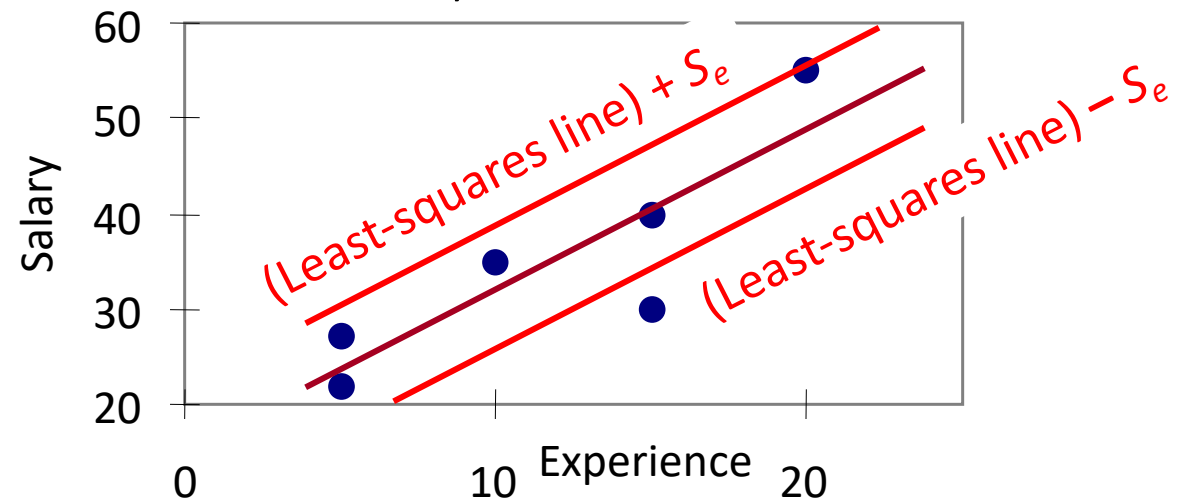
- $$S_e = S_Y \sqrt{(1 - r^2) \frac{n - 1}{n - 2}}$$
- Approximate size of prediction errors (residuals)
Actual Y minus predicted Y : $Y - [a + bX]$
- Example (Salary vs. Experience)

$$S_e = 11.686 \sqrt{(1 - 0.8667^2) \frac{6 - 1}{6 - 2}} = 6.52$$

Predicted salaries are about 6.52 (i.e., \$6,520) away from actual salaries

S_e (continued)

- Interpretation: similar to standard deviation
- Can move Least-Squares Line up and down by S_e
 - About 68% of the data are within one “standard error of estimate” of the least-squares line
 - (For a bivariate normal distribution)



Regression and Prediction Error

- Predicting Y as Y (not using regression)
 - Errors are approximately $S_Y = 11.686$
- Predicting Y as $a+bX$ (using regression)
 - Errors are approximately $S_e = 6.52$
 - Errors are **smaller** when regression is used!
 - This is often the *true payoff* for using regression
- Coefficient of Determination R^2
 - Tells what percent of the variability (variance) of Y is explained by X
 - Example: $R^2 = 0.8667^2 = 0.751$
 - Experience explains 75.1% of the variation in salaries

Linear Model


- Linear Model for the Population
 - The foundation for statistical inference in regression
 - Observed Y is a straight line, plus randomness

$$Y = \alpha + \beta X + \varepsilon$$




Randomness of individuals

Population relationship, on average



Why Statistical Inference?

- Because there can *seem* to be a relationship
 - when, in fact, the population is just random
- Samples of size $n = 10$
 - from a **population** with **no relationship** (correlation 0)
 - **Sample** correlations are not zero!
 - Due to the randomness of sampling



Standard Error of the Slope

- $$S_b = \frac{S_e}{S_X \sqrt{n-1}}$$

- Approximately how far the observed slope b is from the population slope β

- Example (Salary vs. Experience)

$$S_b = \frac{6.52}{6.06\sqrt{6-1}} = 0.48$$

- Observed slope, $b = 1.673$, is about 0.48 away from the unknown slope of the population

Statistical Inference

- Confidence Interval for the Slope

$b \pm tS_b$ where t has $n-2$ degrees of freedom

- Hypothesis Test

- Is β different from $\beta_0 = 0$?
- Is the regression significant?
- Are X and Y significantly related?

- YES

- If 0 is **not** in the confidence interval
 - Or if $|t_{statistic}| = |b/S_b| > t_{table}$

- NO

- Otherwise

*Same test
for all three
questions!*

Example (Salary and Experience)

- 95% confidence interval for population slope β
use $t = 2.776$ from t table, $6 - 2 = 4$ degrees of freedom

$$1.673 \pm 2.776 \times 0.48$$

- From 0.34 to 3.00
 - We are 95% sure that the **population slope** is somewhere between 0.34 and 3.00 (\$thousand per year of experience)
- Hypothesis test result
 - **Significant**
 - because 0 is not in the confidence interval

Experience and Salary are significantly related

Regression Can Be Misleading

- Linear Model May Be Wrong
 - Nonlinear? Unequal variability? Clustering?
- Predicting Intervention from Experience is Hard
 - Relationship may become different if you intervene
- Intercept May Not Be Meaningful
 - if there are no data near $X = 0$
- Explaining Y from X vs. Explaining X from Y
 - Use care in selecting the Y variable to be predicted
- Is there a hidden “Third Factor”?
 - Use it to predict better with multiple regression?