

Unit 1
Summarizing Data

Topics	1. Scales of Measurement	2
	2. Descriptives for Nominal and Ordinal Data	7
	a. The Bar Chart	9
	b. The Pie Chart	12
	3. Descriptives for Continuous Data	13
	a. The Histogram	14
	b. The Frequency Polygon	18
	c. The Cumulative Frequency Polygon	19
	d. Percentiles (Quantiles)	20
	e. Quantile Quantile Plot	22
	f. Stem and Leaf Diagram	23
	g. Box and Whisker Plot	24
	4. The Summation Notation.....	26
	5. Measures of Central Tendency	27
	a. The mode	29
	b. The mean	30
	c. The median	33
	6. Measures of Dispersion	35
	a. Variance	36
	b. Standard Deviation	37
	c. Median Absolute Deviation from Median	39
	d. Standard Deviation v Standard Error	40
	e. A Feel for Sampling Distributions	43
	f. The Coefficient of Variation	45
	g. The Range	46

1. Scales of Measurement

Data can be of different types, and it matters...

The distinction between **qualitative** versus **quantitative** is a familiar one. For example, to describe a flower as **pretty** is a **qualitative** assessment while to record a child's age as **11 years** is a **quantitative** measurement.

Consider that we can reasonably refer to the child's 22 year old cousin as being twice as old as the child whereas we cannot reasonably describe an orchid as being twice as pretty as a dandelion.

We encounter similar stumbling blocks in statistical work. Depending on the type of the variable, its **scale of measurement** type, some statistical methods are meaningful while others are not.

- **Nominal Scale:** Consists of a finite set of possible values or categories which have no particular order.

Example: Cause of Death

- Cancer
- Heart Attack
- Accident
- Other

Example: Gender

- Male
- Female

Example: Race/Ethnicity

- Black
- White
- Latino
- Other

Other Examples: Eye Color, Type of Car, University Attended, Occupation

- **Ordinal Scale:** Consists of a finite set of possible values or categories which DO have an order.

Example: Size of Container

- Small
- Medium
- Large

Example: Pain Level

- None
- Mild
- Moderate
- Severe

For analysis in the computer, both nominal and ordinal data might be stored using numbers rather than text.

Example: Race/Ethnicity

- 1 = Black
- 2 = White
- 3 = Latino
- 4 = Other

Nominal - The numbers have NO meaning
They are labels ONLY

Example: Pain Level

- 1 = None
- 2 = Mild
- 3 = Moderate
- 4 = Severe

Ordinal - LIMITED meaning
4 > 3 > 2 > 1 is all we know
apart from their utility as labels.

Discrete, Count Data: When we count the number of times some event or outcome occurs, the values we can observe are limited to whole numbers: 0, 1, 2, 3, ...

Examples:

Number of children a woman has had
Number of clinic visits made in one year

The numbers are meaningful
We can actually compute with these numbers.

Interval Scale: Interval data are generally measured on a continuum and differences between any two numbers on the scale are of known size.

Example: Temperature in °F on 4 successive days

Day:	A	B	C	D
Temp °F:	50	55	60	65

For these data, not only is day A with 50° cooler than day D with 65°, but it is 15° cooler. Also, day A is cooler than day B by the same amount that day C is cooler than day D (i.e., 5°).

An important property of interval scale data is that there is no true zero point. That is, the value "0" is arbitrary and doesn't reflect absence of the attribute. This is clear in the above example since a temperature of 0° doesn't mean there is no temperature.

Ratio Scale: Ratio scale values are also measured on a meaningful continuum. The distinction is that ratio data have a meaningful zero point.

Example: Weight in pounds of 6 individuals
136, 124, 148, 118, 125, 142

Note on meaningfulness of “ratio”-

Someone who weighs 142 pounds is two times as heavy as someone else who weighs 71 pounds. This is true even if weight had been measured in kilograms.

- Examples of attributes having meaningful zeros:
Height in centimeters
Age in years

- In the sections that follow, we will see that the possibilities for meaningful description (tables, charts, means, variances, etc) are lesser or greater depending on the scale of measurement.
- The following chart gives a sense of this idea.
- For example, we'll see that we can compute relative frequencies for a nominal random variable (eg. Hair color: e.g. "7% of the population has red hair") but we cannot make statements about cumulative relative frequency for a nominal random variable (eg. it would not make sense to say "35% of the population has hair color less than or equal to blonde")

Our data are values of a measured random variable				
Type	Discrete		Continuous	
	Nominal	Ordinal	Interval	Ratio
Descriptive Methods	Bar chart Pie chart - -	Bar chart Pie chart - -	- - Dot diagram Scatter plot (2 variables) Stem-Leaf Histogram Box Plot Quantile-Quantile Plot	- - Dot diagram Scatter plot (2 vars) Stem-Leaf Histogram Box Plot Quantile-Quantile Plot
Numerical Summaries	Frequency Relative Frequency Frequency	Frequency Relative Frequency Cumulative Frequency Frequency	- - - means, variances, percentiles	- - - means, variances, percentiles

Note – This table is an illustration only. It is not intended to be complete.

2. Descriptives for Nominal and Ordinal Data

Consider a study of 25 consecutive patients entering the general medical/surgical intensive care unit at a large urban hospital.

- For each patient the following data are collected:

<u>Variable</u>	<u>Code</u>
• Age (in years)	
• Type of Admission:	1= Emergency 0= Elective
• ICU Type:	1= Medical 2= Surgical 3= Cardiac 4= Other
• Systolic Blood Pressure (in millimeters of mercury)	
• Number of Days Spent in ICU	
• Vital Status at Hospital Discharge:	1= Dead 0= Alive

The actual data are provided on the following page.

<u>ID</u>	<u>Age</u>	<u>Type_Adm</u>	<u>ICU_Type</u>	<u>SBP</u>	<u>ICU_LOS</u>	<u>Vit_Stat</u>
1	15	1	1	100	4	0
2	31	1	2	120	1	0
3	75	0	1	140	13	1
4	52	0	1	110	1	0
5	84	0	4	80	6	0
6	19	1	1	130	2	0
7	79	0	1	90	7	0
8	74	1	4	60	1	1
9	78	0	1	90	28	0
10	76	1	1	130	7	0
11	29	1	2	90	13	0
12	39	0	2	130	1	0
13	53	1	3	250	11	0
14	76	1	3	80	3	1
15	56	1	3	105	5	1
16	85	1	1	145	4	0
17	65	1	1	70	10	0
18	53	0	2	130	2	0
19	75	0	3	80	34	1
20	77	0	1	130	20	0
21	52	0	2	210	3	0
22	19	0	1	80	1	1
23	34	0	3	90	3	0
24	56	0	1	185	3	1
25	71	0	2	140	1	1

Discrete variables:

- Type_Adm
- ICU_Type
- Vit_Stat

Continuous variables:

- Age
- ICU_LOS
- SBP

Frequency Table, Relative Frequency Table

A tally of the possible outcomes, together with “how often” and “proportionately often” is called a frequency distribution.

- ◆ It is a sensible thing to do for discrete variables only.
- ◆ For the variable ICU_Type, the frequency distribution is the following:

ICU_Type	Frequency (“how often”)	Relative Frequency (“proportionately often”)
Medical	12	0.48
Surgical	6	0.24
Cardiac	5	0.20
Other	2	0.08
TOTAL	25	1.00

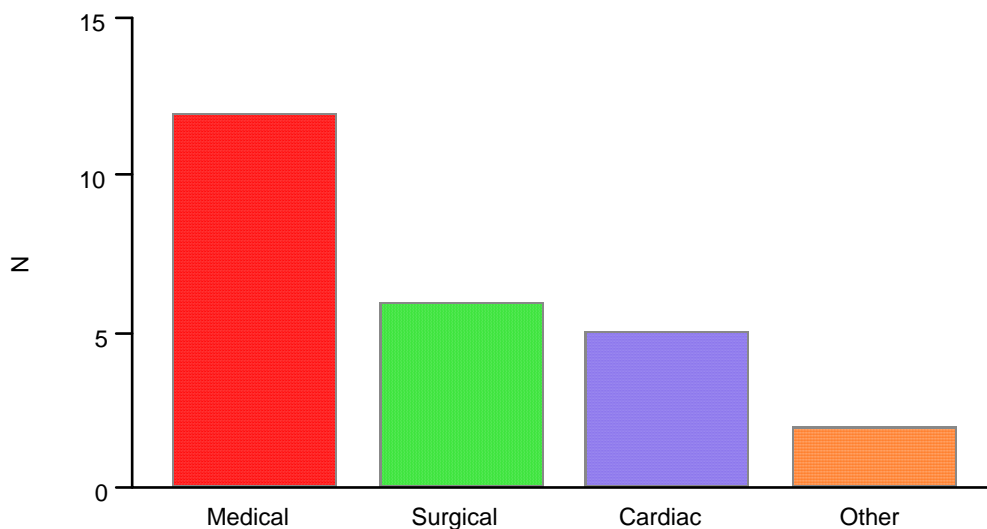
- ◆ This summary will be useful in constructing two graphical displays, the bar chart and the pie chart.

Bar Chart

On the horizontal are the possible outcomes. On the vertical is plotted either

- ◆ “how often” - frequency
- ◆ “how proportionately often” - relative frequency

Example - Bar chart for Type of ICU Patients (n=25)




Guidelines for Construction of Bar Charts

- Label both axes clearly
- Leave space between bars
- Leave space between the left-most bar and the vertical axis
- When possible, begin the vertical axis at 0
- All bars should be the same width

There's no reason why the bar chart can't be plotted horizontally instead of vertically. And, not surprisingly, if you change the choice of scale, you can communicate to the eye a very different message.

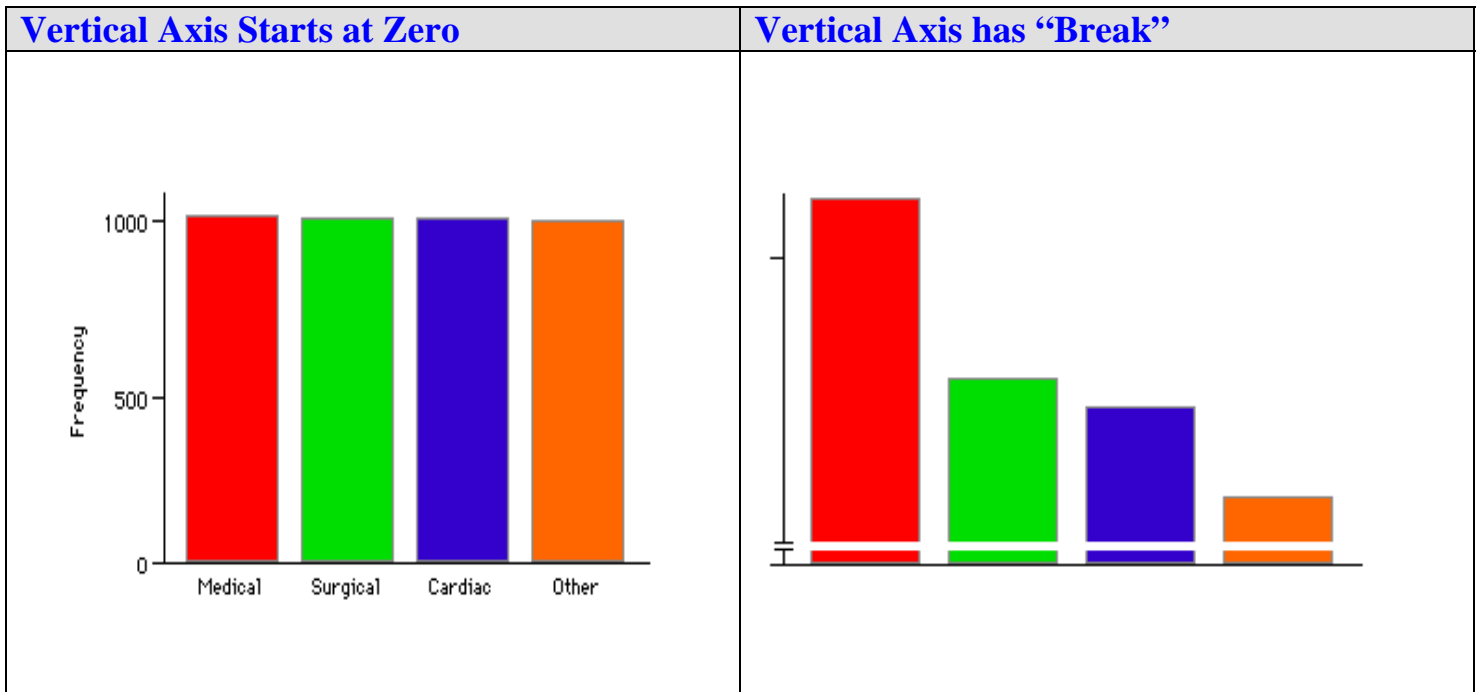
Example, continued - Consider two choices of scale for the vertical axis in the bar chart for ICU_Type:



It's difficult to know how to construct a bar chart when the frequencies are very high.

- Suppose the frequencies were 1012 medical, 1006 surgical, 1005 cardiac and 1002 other type.
- Is it better to have a vertical axis start at zero or to “break” the axis?

Example, continued -




Pie Chart

Instead of “stacks” rising up from the horizontal (bar chart), we could plot instead the shares of a pie.

Recalling that a circle has 360 degrees, that 50% means 180 degrees, 25% means 90 degrees, etc, we can identify “wedges” according to relative frequency

Relative Frequency	Size of Wedge, in degrees
0.50	50% of 360 = 180 degrees
0.25	25% of 360 = 90 degrees
p	$(p) \times (100\%) \times 360$ degrees

Example, continued -




3. Descriptives for Continuous Data

Tables and graphs for continuous data address many things. Among the most important are:

- What is typical (location)
- What is the scatter (dispersion)

The following are 3 scenarios for the patterns of values (eg – values of age) in a simple random sample (eg – a simple random sample of cholesterol values where the sample size is $n=25$)



The 3 patterns are quite different. The leftmost pattern is symmetric and bell shaped. The middle pattern has a tail to the right. The rightmost pattern is comprised of two symmetric bell shaped patterns that are separated in their location.

“Good” choices for summarizing location and dispersion are not always the same and depend on the pattern of scatter.

Histogram

For a continuous variable (e.g. – age), the frequency distribution of the individual ages is not so interesting.

Age	Frequency
15	1
19	2
29	1
31	1
34	1
39	1
52	2
53	2
45	2
65	1
71	1
74	1
75	2
76	2
77	1
78	1
79	1
84	1
85	1


We “see more” in frequencies of age values in “groupings”. Here, 10 year groupings make sense.

Age Interval	Frequency
10-19	3
20-29	1
30-39	3
40-49	0
50-59	6
60-69	1
70-79	9
80-89	2
TOTAL	25

A plot of this “grouped” frequency table gives us a better feel for the pattern of ages with respect to both location and scatter. This plot is called a [histogram](#)

- A **histogram** is a graphical summary of the pattern of values of a **continuous random variable**. More formally, it is a graphical summary of the frequency distribution
- It is **analogous** to the bar graph summary for the distribution of a discrete random variable

Figure 12: Histogram of the Age Distribution of ICU Patients



How to Construct a Histogram

Step 1: Choose the number of groupings (“class intervals”). Call this k .

- ◆ The choice is arbitrary.
- ◆ k too small over-summarizes. k too large under-summarizes.
- ◆ Sometimes, the choices of intervals are straightforward – eg 10 year intervals, 7 day intervals, 30 day intervals.
- ◆ Some text books provide formulae for k . Use these if you like, provided the resulting k makes sense.
- ◆ Example – For the age data, we use $k=8$ so that intervals are sensible 10 year spans.

Step 2: Rules for interval beginning and end values (“boundaries”)

- ◆ Boundaries should be such that each observation has exactly one “home”.
- ◆ Equal widths are not necessary. WARNING – If you choose to plot intervals that are of Unequal widths, take care to plot “area proportional to relative frequency”. This is explained in Step 3.

Step 3: In a histogram, always plot “Area Proportional to Relative Frequency”


- ◆ An example is the best explanation.

Age Interval	Frequency
10-29	4
30-39	3
40-49	0
50-59	6
60-69	1
70-79	9
80-89	2
TOTAL	25

- ◆ For the intervals 30-39, 40-49, 50-59, 60-69, 70-79, 80-89: the widths are all the same and span 10 units of age. Heights plotted are 3, 0, 6, 1, 9, and 2.
- ◆ But the interval 10-29 spans 20 units of age. A frequency of 4 over 20 units of age corresponds to a frequency of 2 over 10 units of age.
- ◆ For the interval 10-29, a height of 2 would be plotted.

The result is the histogram shown on page 15, shown again here:

Figure 12: Histogram of the Age Distribution of ICU Patients



:

Frequency Polygon

- ◆ The frequency polygon is an alternative to the histogram. It's not used often. Its companion, the cumulative frequency polygon (next page), is more commonly used.
- ◆ Both the histogram and frequency polygon are graphical summaries of the frequency distribution of a continuous random variable
- ◆ Whereas in a histogram ...
 - ◆ X-axis shows intervals of values
 - ◆ Y-axis shows bars of frequencies
- ◆ In a frequency Polygon:
 - ◆ X-axis shows midpoints of intervals of values
 - ◆ Y-axis shows dot instead of bars

Some guidelines:

- i. The graph title should be a complete description of the graph
- ii. Clearly label both the horizontal and vertical axes
- iii. Break axes when necessary
- iv. Use equal class widths
- v. Be neat and accurate

Example -

The following frequency polygon is similar in interpretation to a histogram.




Figure 1: Frequency polygon for the psychology test scores.

Source: <http://cnx.org/content/m11214/latest/>

Cumulative Frequency Polygon


As you might think, a **cumulative** frequency polygon depicts **accumulating** data.

This is actually useful in that it lets us see **percentiles** (eg, median)

Its plot utilizes cumulative frequency information (right hand columns below in blue)

Age Interval	Frequency	Relative Frequency	Cumulative through Interval	
			Frequency	Relative Frequency
10-19	3	12	3	12
20-29	1	4	4	16
30-39	3	12	7	28
40-49	0	0	7	28
50-59	6	24	13	52
60-69	1	4	14	56
70-79	9	36	23	92
80-89	2	8	25	100
TOTAL	25	100		

Notice that it is the **ENDPOINT** of the interval that is plotted on the horizontal. This makes sense inasmuch as we are keeping track of the **ACCUMULATION** of frequencies to the **end of the interval**.



Percentiles (Quantiles)

Suppose that 50% of a cohort survived at least 4 years.

This also means that 50% survived at most 4 years.

We say 4 years is the median.

The median is also called the 50th percentile, or the 50th quantile. We write $P_{50} = 4$ years.

Similarly we could speak of other percentiles:


P_{25} : 25% of the sample values are less than or equal to this value.

P_{75} : 75% of the sample values are less than or equal to this value.

P_0 : The minimum.

P_{100} : The maximum.

It is possible to estimate the values of percentiles from a cumulative frequency polygon.



Example – Consider $P_{10} = 18$. This allows us to say: “10% of the sample is age ≤ 18 ” or “The 10th percentile of age in this sample is 18 years”.

How to Determine the Value of the pth Percentile in a Data Set

Consider the following sample of n=40 data values

0	1	1	3	17	32	35	44	48	86
87	103	112	121	123	130	131	149	164	167
173	173	198	208	210	222	227	234	245	250
253	256	266	277	284	289	290	313	477	491

Step 1:

Order the data from smallest to largest

Step 2:

Compute $L = n \left[\frac{p}{100} \right]$ where

n = size of sample (eg; n=40 here)
p = desired percentile (eg p=25th)

L is NOT a whole number

Step 3:

Change L to next whole number.

Pth percentile = Lth ordered value in the data set.

L is a whole number

Step 3:

Pth percentile = average of the Lth and (L+1)st ordered value in the data set.

Solution for 25th Percentile

$L = (40)(25/100) = 10 \rightarrow$ 25th percentile is average of 10th and 11th $\rightarrow = (1/2)(86 + 87) = 86.5$

Quantile-Quantile (QQ) and Percentile-Percentile (PP) Plots

- We might ask: *“Is the distribution of my data normal?”*
 QQ and PP plots are useful when we want to compare the percentiles of our data with the percentiles of some reference distribution (eg- reference is normal)

- **QQ Plot:** **X= quantile in sample** **Y=quantile in reference**


Desired Quantile	10th	20th	...	99 th
X = Value in Sample of Data				
Y = Value in Reference Distribution				

- **PP Plot:** **X= percentile rank in sample** **Y= percentile rank in reference**

Data value	##	##	...	##
X = Rank (percentile) in Sample of Data				
Y = Rank (percentile) in Reference Distribution				

- **What to look for:** A straight line suggests that the two distributions are the same

Example – Normal Quantile-Quantile Plot of Age for data on page 8



Note on the Y-axis values: These have been standardized (more on this later)


Stem and Leaf Diagram

- A stem and leaf diagram is a “quick and dirty” histogram
- It is also a quick easy and easy way to sort data.

Example - Stem and Leaf Plot of Age of 25 ICU Patients:

- The **stems** are to the left of the vertical. In this example, each value of stem represents a multiple of 10.
 - *Step 1:* List your stems
- The **leaves** are to the right of the vertical.

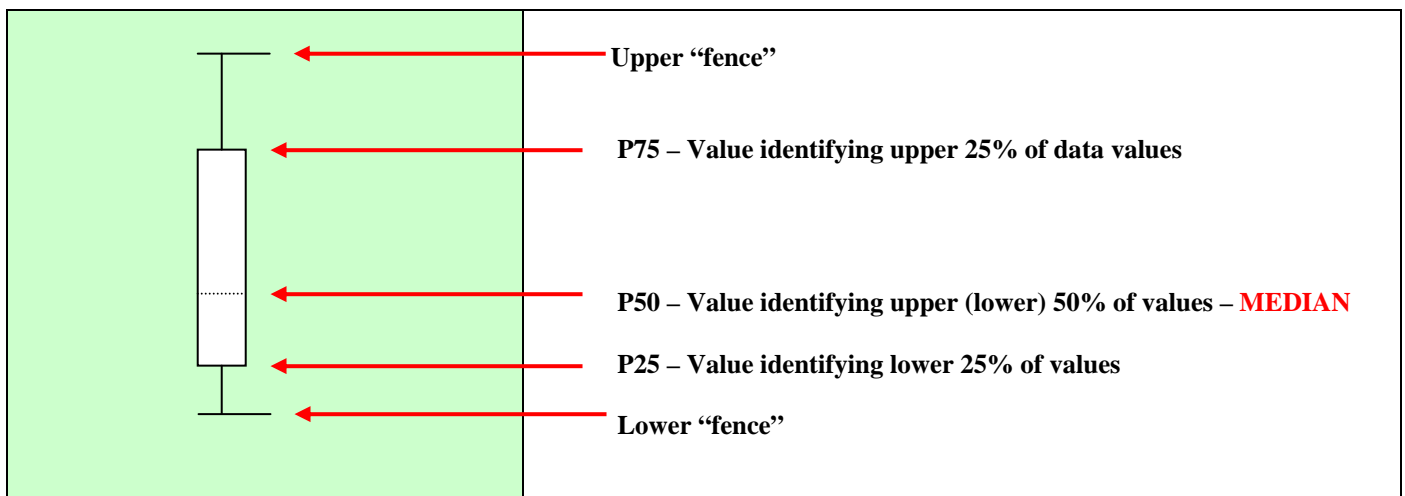
Step 2: “Plot” your leaves.



Among other things, we see that the minimum age in this sample is 15 years and the maximum age is 85 years. We can also see that most of the subjects in this sample are in their seventies.

Box and Whisker Plot


- The **box and whisker plot** (also called **box plot**) is a wonderful schematic summary of the distribution of values in a data set.
- It shows you a number of features, including: extremes, 25th and 75th percentile, median and sometimes the mean.
- Side-by-side box and whisker plots are a wonderful way to compare multiple distributions.
- **Definition**
 - The **central box** has P_{25} and P_{75} for its limits. It spans the middle half of the data.
 - The **line within the box** identifies the median, P_{50} . Sometimes, an asterisk within the box is shown. It is the **mean**. The lines coming out of the box are called the “whiskers”. The ends of these “whiskers” are called “fences”.
 - **Upper “fence”** = The largest value that is less than or equal to $P_{75} + 1.5*(P_{75} - P_{25})$.
 - **Lower “fence”** = The smallest value that is greater than or equal to $P_{25} - 1.5*(P_{75} - P_{25})$.



Note on the multiplier 1.5


This is a convenient multiplier if we are interested in the comparability of the distribution of our sample values to a normal (Gaussian) distribution in the following way. If the data are from a normal distribution, then 95% of the data values will fall within the range defined by the lower and upper fences.

Example, continued - Box plot summary of the distribution of the 25 ages.



Note: min=15, P₂₅=39, P₅₀=53, P₇₅=76, max=85, and $1.5*(P_{75}-P_{25})=55.5$

Example – The following is an example of side-by-side box and whisker plots. They are plots of values of duration of military service for 5 groups defined by age.



- Duration of military service increases with age.
- With age, the variability in duration of military service is greater.
- The individual circles represent extreme values.

4. The Summation Notation

The summation notation is nothing more than a secretarial convenience. We use it to avoid having to write out long expressions.

For example,

Instead of writing $x_1 + x_2 + x_3 + x_4 + x_5$,
 We write $\sum_{i=1}^5 x_i$

Another example –

Instead of writing out the product of five terms $x_1 * x_2 * x_3 * x_4 * x_5$,
 We write $\prod_{i=1}^5 x_i$
 This is actually an example of the [product notation](#)

The summation notation

Σ The Greek symbol sigma says “add up some items”

$\sum_{\text{STARTING HERE}}$ Below the sigma symbol is the starting point


$\overset{\text{END}}{\Sigma}$ Up top is the ending point

5. Measures of Central Tendency

Previously we noted that among the most important tools of description are that address

- What is typical (location)
- What is the scatter (dispersion)

Recall - “Good” choices for summarizing location and dispersion are **not** always the same and **depend on the pattern of scatter**.



Mode. The mode is the most frequently occurring value. It is not influenced by extreme values. Often, it is not a good summary of the majority of the data.

Mean. The mean is the arithmetic average of the values. It is sensitive to extreme values.

$$\text{Mean} = \frac{\text{sum of values}}{\text{sample size}} = \frac{\sum (\text{values})}{n}$$

Median. The median is the middle value when the sample size is odd. For samples of even sample size, it is the average of the two middle values. It is not influenced by extreme values.

We consider each one in a bit more detail ...

Mode

Mode. The mode is the most frequently occurring value. It is not influenced by extreme values. Often, it is not a good summary of the majority of the data.

Example

- Data are: 1, 2, 3, 4, 4, 4, 4, 5, 5, 6
- Mode is 4

Example

- Data are: 1, 2, 2, 2, 3, 4, 5, 5, 5, 6, 6, 8
- There are two modes – value 2 and value 5
- This distribution is said to be “bi-modal”

Modal Class

- For grouped data, it may be possible to speak of a modal class
- The modal class is the class with the largest frequency

Example

Interval of Values (“Class”)	Frequency, f
31-40	1
41-50	2
51-60	5
61-70	15
71-80	25
81-90	20
91-100	12

- The modal class is the interval of values 71-80

Mean

Mean. The mean is the arithmetic average of the values. It is sensitive to extreme values.

$$\text{Mean} = \frac{\text{sum of values}}{\text{sample size}} = \frac{\sum (\text{values})}{n}$$


Calculation of a “mean” or “average” is familiar; e.g. -

grade point average
mean annual rainfall
average weight of a catch of fish
average family size for a region

A closer look using summation notation

- Suppose data are: 90, 80, 95, 85, 65
- sample mean = $\frac{90+80+95+85+65}{5} = \frac{415}{5} = 83$
- sample size, $n = 5$
- $x_1 = 90, x_2 = 80, x_3 = 95, x_4 = 85, x_5 = 65$
- \bar{X} = sample mean
- $\bar{X} = \frac{\sum_{i=1}^5 x_i}{n} = \frac{x_1 + x_2 + x_3 + x_4 + x_5}{5} = \frac{90 + 80 + 95 + 85 + 65}{5} = 83$

The mean can be thought of as a “balancing point”, “center of gravity”



- By “balance”, it is meant that the sum of the departures from the mean to the left balance out the sum of the departures from the mean to the right.
- In this example, sample mean $\bar{X} = 83$
- **TIP!!** Often, the value of the sample mean is not one that is actually observed

Skewness

When the data are skewed, the mean is “dragged” in the direction of the skewness

Negative Skewness (Left tail)	Positive Skewness (Right tail)
<p>Mean is dragged left</p>	<p>Mean is dragged right</p>

The Weighted Mean

- In a weighted mean, separate outcomes have separate influences.
- The influence attached to an outcome is the weight.
- Familiar is the calculation of a course grade as a weighted average of scores on separate outcomes.
- **Example**

Outcome, X_i is	Weight (Influence) w_i is
Homework #1 = $X_1 = 96$	$w_1 = 20$ (20% of grade)
Homework #2 = $X_2 = 84$	$w_2 = 20$
Homework #3 = $X_3 = 65$	$w_3 = 20$
Midterm Exam = $X_4 = 73$	$w_4 = 10$
Final Exam = $X_5 = 94$	$w_5 = 30$

$$\begin{aligned}
 \text{weighted mean} &= \frac{\sum(\text{weight})(\text{score})}{\sum(\text{weight})} \\
 &= \frac{\sum_{i=1}^n (w_i)(X_i)}{\sum(w_i)} \\
 &= \frac{(20)(96) + (20)(84) + (20)(65) + (10)(73) + (30)(94)}{(20) + (20) + (20) + (10) + (30)} \\
 &= 84.5
 \end{aligned}$$

Note – The use of the weighted mean is often used to estimate the mean in a sample of data that have been summarized in a frequency table. The values used are the interval midpoints. The weights used are the interval frequencies.

The Median

Median. The median is the middle value when the sample size is odd. For samples of even sample size, it is the average of the two middle values. It is not influenced by extreme values.

If the sample size n is ODD	median = $\frac{n+1}{2}$ th largest value
If the sample size n is EVEN	median = average of $\left(\left[\frac{n}{2}\right]th, \left[\frac{n+2}{2}\right]th\right)$ values

Example

- Data, from smallest to largest, are: 1, 1, 2, 3, 7, 8, 11, 12, 14, 19, 20
- The sample size, $n=11$
- Median is the $\frac{n+1}{2}$ th largest = $\frac{12}{2} = 6$ th largest value
- Thus, median value is = 8
- Five values are smaller than 8; five values are larger.

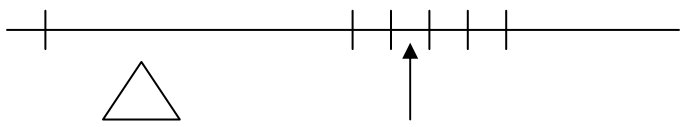
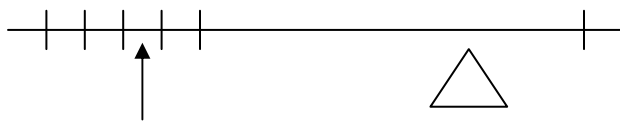
Example

- Data, from smallest to largest, are: 2, 5, 5, 6, 7, 10, 15, 21, 22, 23, 23, 25
- The sample size, $n=12$
- Median = average of $\frac{n}{2}$ th largest, $\frac{n+2}{2}$ = average of 6th and 7th largest values
- Thus, median value is = average (10, 15) = 12.5


Skewed Data – When the data are skewed the median is a better description of the majority than the mean

Example

- Data are: 14, 89, 93, 95, 96
- Skewness is reflected in the outlying low value of 14
- The sample mean is 77.4
- The median is 93

Negative Skewness (Left tail)	Positive Skewness (Right tail)
MEAN < Median	MEAN > Median
 <p data-bbox="121 1711 357 1753">Mean is dragged left</p> <p data-bbox="454 1711 568 1753">MEDIAN</p>	 <p data-bbox="933 1711 1047 1753">MEDIAN</p> <p data-bbox="1193 1711 1461 1753">Mean is dragged right</p>

6. Measures of Dispersion



There are choices for describing *dispersion*, too. As before, a “good” choice will depend on the pattern of scatter.

Variance

A summary measure of the variability of individuals in nature

Population Variance. The population variance is written as σ^2 and is a summary measure of the squares of individual departures from the mean *in a population*. For a population that is finite and of size N, there exists a formula for population variance. This formula makes use of the mean of the population which is represented as μ .

$$\sigma^2 = \frac{\sum_{i=1}^N (X_i - \mu)^2}{N}$$

Note – The population variance can be interpreted as the average of the individual squared deviations from the mean. It can also be interpreted as the expected squared deviation from the mean.

Sample variance A sample variance is written as S^2 and is a summary measure of the squares of individual departures from the sample mean in a sample. For a simple random sample of size n (recall – we use the notation “N” when we speak of the size of a finite population and we use the notation “n” when we speak of the size of a sample)

$$S^2 = \frac{\sum_{i=1}^n (X_i - \bar{X})^2}{(n-1)}$$

Note – Comparison of a finite population variance with a sample variance allows us to appreciate that the sample variance is a good guess of the population variance. The formulae are very similar. Compared to the formula for the population variance, the formula for the sample variance has the following modifications: (1) N is replaced by n (2) μ is replaced by \bar{X} and (3) the divisor N is replaced by the divisor (n-1) . Thus, provided n is reasonably large, the sample variance can be interpreted (almost anyway) as an average squared deviation

Standard Deviation

Standard Deviation. The population standard deviation (σ) and a sample standard deviation (S or SD) are the square roots of σ^2 and S^2 . As such, they are additional choices for summarizing variability. The advantage of the square root operation is that the resulting summary has the same scale as the original values.

$$\text{Sample Standard Deviation (S or SD)} = \sqrt{\frac{\sum (X - \bar{X})^2}{n - 1}}$$

Disparity between individual and average ...	$(X - \bar{X})$
Disparity between individual and average ...	$(X - \bar{X})^2$
The average of these ...	$\frac{\sum (X - \bar{X})^2}{n}$
The sample variance S^2 is an “almost” average	$\frac{\sum (X - \bar{X})^2}{n - 1}$
The related measure S (or SD) returns measure of dispersion to original scale of observation ...	$\sqrt{\frac{\sum (X - \bar{X})^2}{n - 1}}$

Example of Sample Variance (S^2) and Standard Deviation (S) Calculation –

Consider the following sample of survival times (X) of n=11 patients after heart transplant surgery. Interest is to calculate the sample variance and standard deviation.

- ◆ Patients are identified numerically, from 1 to 11.
- ◆ The survival time for the “ith” patient is represented as X_i for $i= 1, \dots, 11$.

Patient Identifier, “i”	Survival (days), X_i	Mean for sample, \bar{X}	Deviation , $(X_i - \bar{X})$	Squared deviation $(X_i - \bar{X})^2$
1	135	161	-26	676
2	43	161	-118	13924
3	379	161	218	47524
4	32	161	-129	16641
5	47	161	-114	12996
6	228	161	67	4489
7	562	161	401	160801
8	49	161	-112	12544
9	59	161	-102	10404
10	147	161	-14	196
11	90	161	-71	5041
TOTAL	1771		0	285236

◆ $\sum_{i=1}^{11} X_i = 1771$ days

◆ Sample mean is $\bar{X} = \frac{1771}{11} = 161$ days

◆ Sample variance is $s^2 = \frac{\sum_{i=1}^{11} (X_i - \bar{X})^2}{n - 1} = \frac{285236}{10} = 28523.6$ days²

◆ Sample standard deviation is $s = \sqrt{s^2} = \sqrt{28523.6} = 168.89$ days

Median Absolute Deviation About the Median (MADM)

Median Absolute Deviation about the Median (MADM) - Another measure of variability is helpful when we wish to describe scatter among data that is skewed.

Recall that the median is a good measure of location for skewed data because it is not sensitive to extreme values.

Distances are measured about the median, not the mean.

We compute deviations rather than squared differences.

Thus

Median Absolute Deviation about the Median (MADM)

$$\text{MADM} = \text{median of } [|X_i - \text{median of } \{X_1, \dots, X_n\} |]$$

Example.

Original data: { 0.7, 1.6, 2.2, 3.2, 9.8 }

Median = 2.2

X_i	$ X_i - \text{median} $
0.7	1.5
1.6	0.6
2.2	0.0
3.2	1.0
9.8	7.6

$$\text{MADM} = \text{median } \{ 0.0, 0.6, 1.0, 1.5, 7.6 \} = \mathbf{1.0}$$

Standard Deviation (S or SD) versus Standard Error (SE)

So... what is the distinction between the standard deviation (SD or S) and the standard error (SE)?

the standard deviation (SD or S) addresses questions about variability of individuals in nature (imagine a collection of individuals), whereas

the standard error (SE) addresses questions about the variability of a summary statistic among many replications of your study (imagine a collection of values of a sample statistic such as the sample mean that is obtained by repeating your whole study over and over again)

The distinction makes use of the idea of **sampling distributions**. The following is an illustration. Sampling distributions are discussed further on page 42

Example

Suppose you conduct a study that involves obtaining a simple random sample of size $n=11$. Suppose further that, from this one sample, you calculate the sample mean (*note – you might have calculated other sample statistics, too, such as the median or sample variance*). Now imagine replicating the entire study 5000 times. You would then have 5,000 sample means, each based on a sample of size $n=11$.

If instead of replicating your study 5000 times, the study were replicated infinitely many times, the resulting collection of infinitely many sample means has a name: the **sampling distribution of $\bar{X}_{n=11}$** . Notice the subscript “ $n=11$ ”. This is a reminder to us that the particular study design that we have replicated infinitely many times calls for drawing a sample of size $n=11$ each time.

Why do we care at all?

We care because, often, we’re interested in knowing if the results of our one study conduct are similar to what would be obtained if someone else were to repeat it!!

Distinction between Standard Deviation and Standard Error of the Mean .

We're often interested in the (theoretical) behavior of the **sample mean** \bar{X}_n from one replication of our study to the next.

So, whereas, the typical variability among individual values can be described using the standard deviation (SD).

The typical variability of the sample mean from one replication of the study is described using the standard error (SE) of the mean:

$$SE(\bar{X}) = \frac{SD}{\sqrt{n}}$$

Note – A limitation of the SE is that it is a function of both the natural variation (SD in the numerator) and the study design (n in the denominator). *More on this later!*

Example, continued

Previously, we summarized the results of one study that enrolled $n=11$ patients after heart transplant surgery. For that one study, we obtained an average survival time of $\bar{X} = 161$ days.

What happens if we repeat the study? What will our next \bar{X} be? Will it be close? How different will it be? We care about this question because it pertains to the generalizability of our study findings.

The behavior of \bar{X} from one replication of the study to the next replication of the study is referred to as the sampling distribution of \bar{X} .

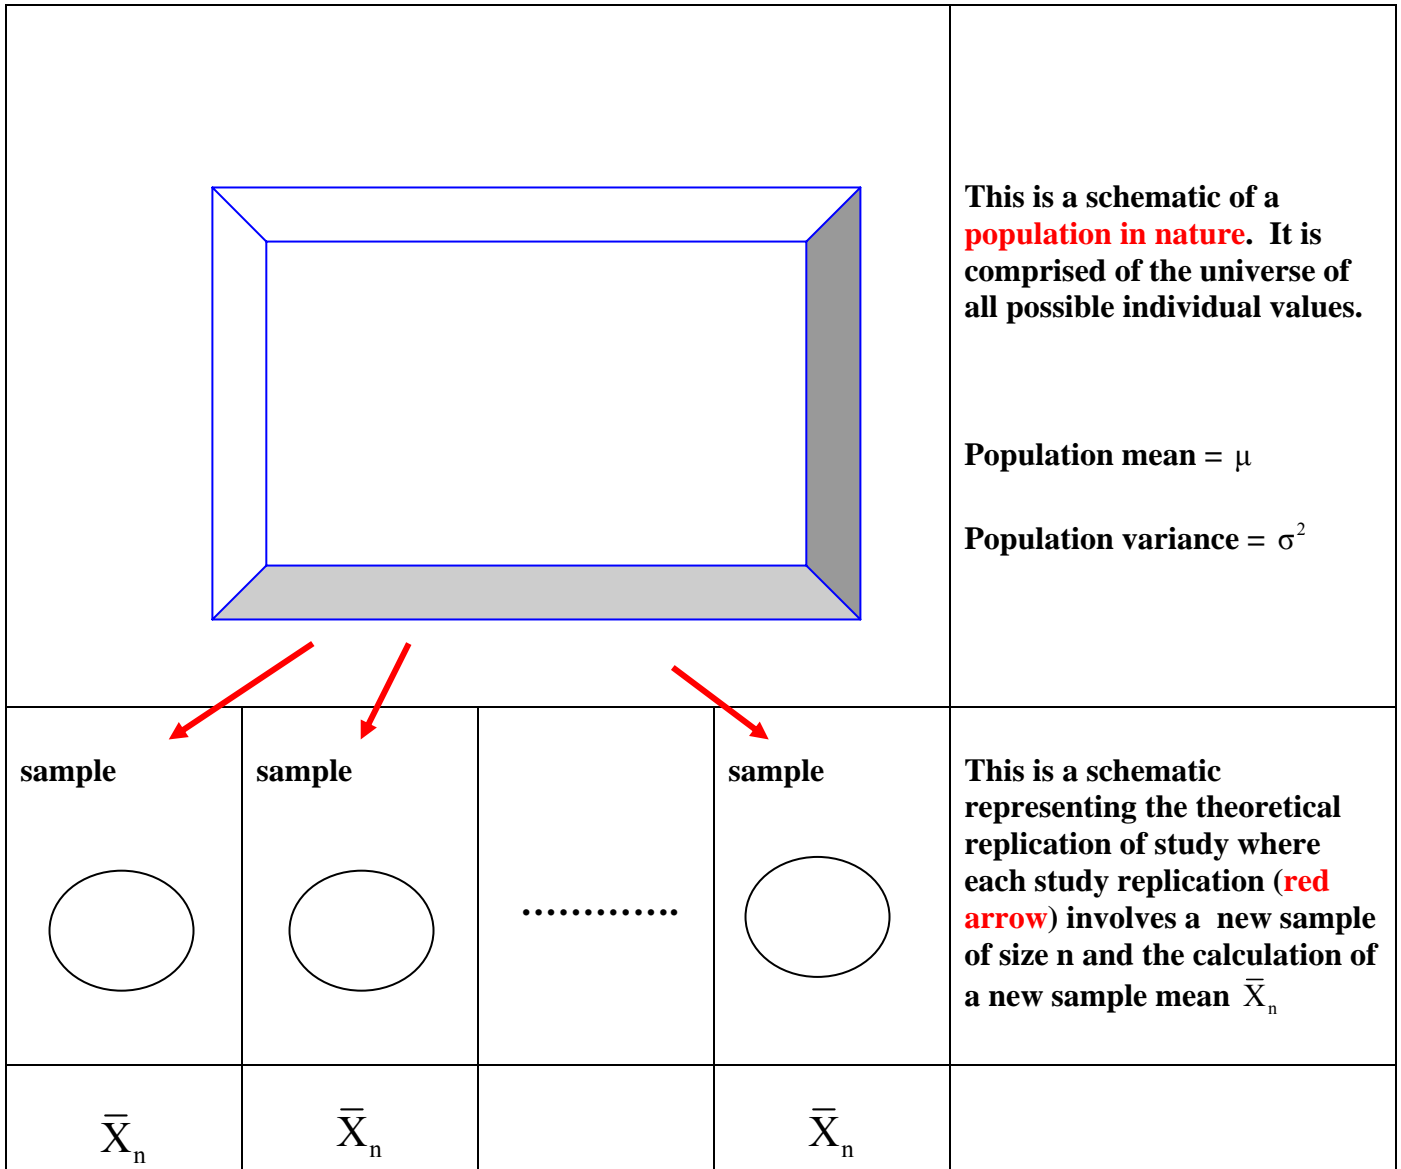
(We could just as well have asked about the behavior of the median from one replication to the next (sampling distribution of the median) or the behavior of the SD from one replication to the next (sampling distribution of SD).)

Thus, interest is in a measure of the “noise” that accompanies $\bar{X} = 161$ days. The measure we use is the standard error measure. This is denoted SE. For this example, in the heart transplant study


$$SE(\bar{X}) = \frac{SD}{\sqrt{n}} = \frac{168.89}{\sqrt{11}} = 50.9$$

We interpret this to mean that a similarly conducted study might produce an average survival time that is near 161 days, give or take 50.9 days.

A Feel for Sampling Distributions



Collecting together all the \bar{X}_n ...



This is a schematic of the resulting **sampling distribution of \bar{X}_n** . It is the collection of all possible sample means.

Sampling distribution has mean = $\mu_{\bar{X}_n}$

Sampling distribution has variance = $\sigma_{\bar{X}_n}^2$

Standard Deviation

- Describes variation in values of *individuals*.
- In the population of *individuals*: σ
- Our “guess” is S

Standard Error

- Describes variation in values of a *statistic* from one conduct of study to the next.
- Often, it is the variation in the *sample mean* that interests us.
- In the population of all possible *sample means* (“sampling distribution of mean”):
 σ / \sqrt{n}
- Our “guess” is S / \sqrt{n}

The Coefficient of Variation

The **coefficient of variation** is the ratio of the standard deviation to the mean of a distribution.

- It is a measure of the spread of the distribution relative to the mean of the distribution
- In the population, coefficient of variation is denoted ξ and is defined

$$\xi = \frac{\sigma}{\mu}$$

- The coefficient of variation ξ can be estimated from a sample. Using the hat notation to indicate “guess”. It is also denoted CV

$$cv = \hat{\xi} = \frac{S}{\bar{X}}$$

Example – “Cholesterol is more variable than systolic blood pressure”

	S	\bar{X}	$cv = \hat{\xi} = s/\bar{x}$
Systolic Blood Pressure	15 mm	130 mm	.115
Cholesterol	40 mg/dl	200 mg/dl	.200

Example – “Diastolic is relatively more variable than systolic blood pressure”

	S	\bar{X}	$cv = \hat{\xi} = s/\bar{x}$
Systolic Blood Pressure	15 mm	130 mm	.115
Diastolic Blood Pressure	8 mm	60 mm	.133

The Range

The range is the difference between the largest and smallest values in a data set.

- It is a quick measure of scatter but not a very good one.
- Calculation utilizes only two of the available observations.
- As n increases, the range can only increase. Thus, the range is sensitive to sample size.
- The range is an unstable measure of scatter compared to alternative summaries of scatter (e.g. S or MADM)
- HOWEVER – when the sample size is very small, it may be a better measure of scatter than the standard deviation S .

Example –

- Data values are 5, 9, 12, 16, 23, 34, 37, 42
- $\text{range} = 42 - 5 = 37$