

## Statistics and parameters

Tables, histograms and other charts are used to summarize large amounts of data. Often, an even more extreme summary is desirable.

- A number that summarizes *population* data is called a *parameter*.
- A number that summarizes *sample* data is called a *statistic*.

### Observations:

- Population parameters are (more or less) constant, while sample statistics depend on the particular sample chosen—i.e., sample statistics *vary with the sample*.
- Sample statistics are *known* because we can compute them from the (available) sample data, while population parameters are often *unknown*, because data for the entire population is often unavailable.
- One of the most common uses of statistics is to *estimate* (unknown) population parameters.

## Measures of *central tendency*

The most extreme way to summarize a list of numbers is with a single, *typical* value. The most common choices are the *mean* and *median*.

- The *mean* (*average*) of a set of numbers is the sum of all the values divided by the number of values in the set.
- The *median* of a set of number is the middle number, when the numbers are listed in increasing (or decreasing) order. The median splits the data into two equally sized sets—50% of the data lies below the median and 50% lies above.

(If the number of numbers in the set is *even*, then the median is the average of the two middle values.)

The mean and median different ways of describing the *center* of the data. Another statistic that is often used to describe the typical value is the *mode*, which is the *most frequently occurring* value in the data.

**Example.** Find the mean, median and mode of the following set of numbers:

$$\{12, 5, 6, 8, 12, 17, 7, 6, 14, 6, 5, 16\}.$$

- The *mean* (average).

$$\frac{12 + 5 + 6 + 8 + 12 + 17 + 7 + 6 + 14 + 6 + 5 + 16}{12} = \frac{114}{12} = 9.5.$$

- The *median*. Arrange the data in ascending order, and find the average of the middle two values in this case, since there are an even number of values:

$$5, 5, 6, 6, 6, 7, 8, 12, 12, 14, 16, 17 \longrightarrow \text{median} = \frac{7 + 8}{2} = 7.5.$$


- The *mode* is 6, because 6 occurs most frequently (three times).

## Comments:

- The mean is sensitive to *outliers*—extreme values in the data (much bigger or much smaller than most of the data). Big outliers pull the mean up and small outliers pull the mean down.
- The median gives a better sense of ‘middle’ when the data is skewed in one direction or the other.
- The mean is easier to use in mathematical formulas.
- Both the median and the mean leave out a lot of information. E.g., they tell us nothing about the *spread* of the data or where we might find ‘peaks’ (modes) in the distribution, etc.


The relative positions of the mean and median provides information about how the data is distributed...

In this histogram the mean is bigger than the median.



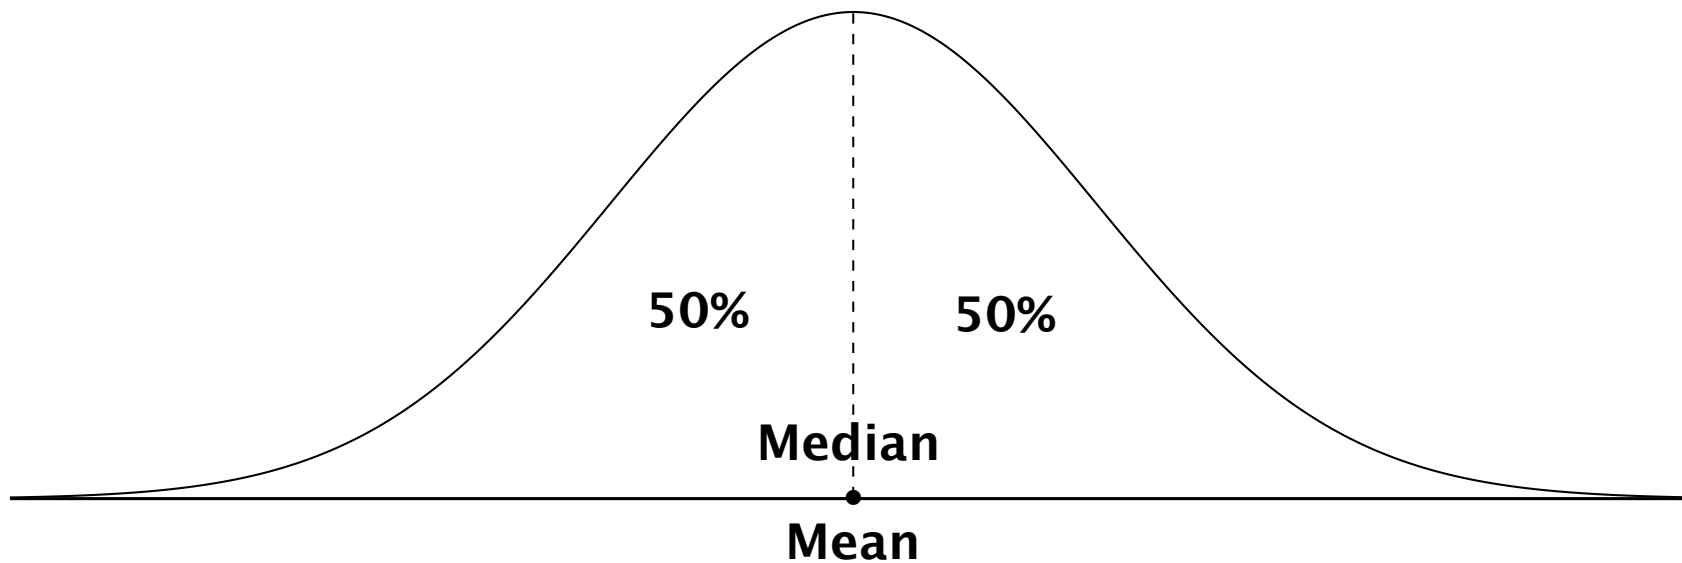
This is an indication that there are large outliers — the histogram has a longer tail on the right. We say that the data is *skewed to the right*.

In this histogram the mean is smaller than the median.




This is an indication that there are small outliers — the histogram has a longer tail on the left. We say that the data is *skewed to the left*.

If the mean and median are (more or less) equal, then the tails of the distribution are (more or less) the same, and the data has a (more or less) symmetric distribution around the mean/median, as depicted below.

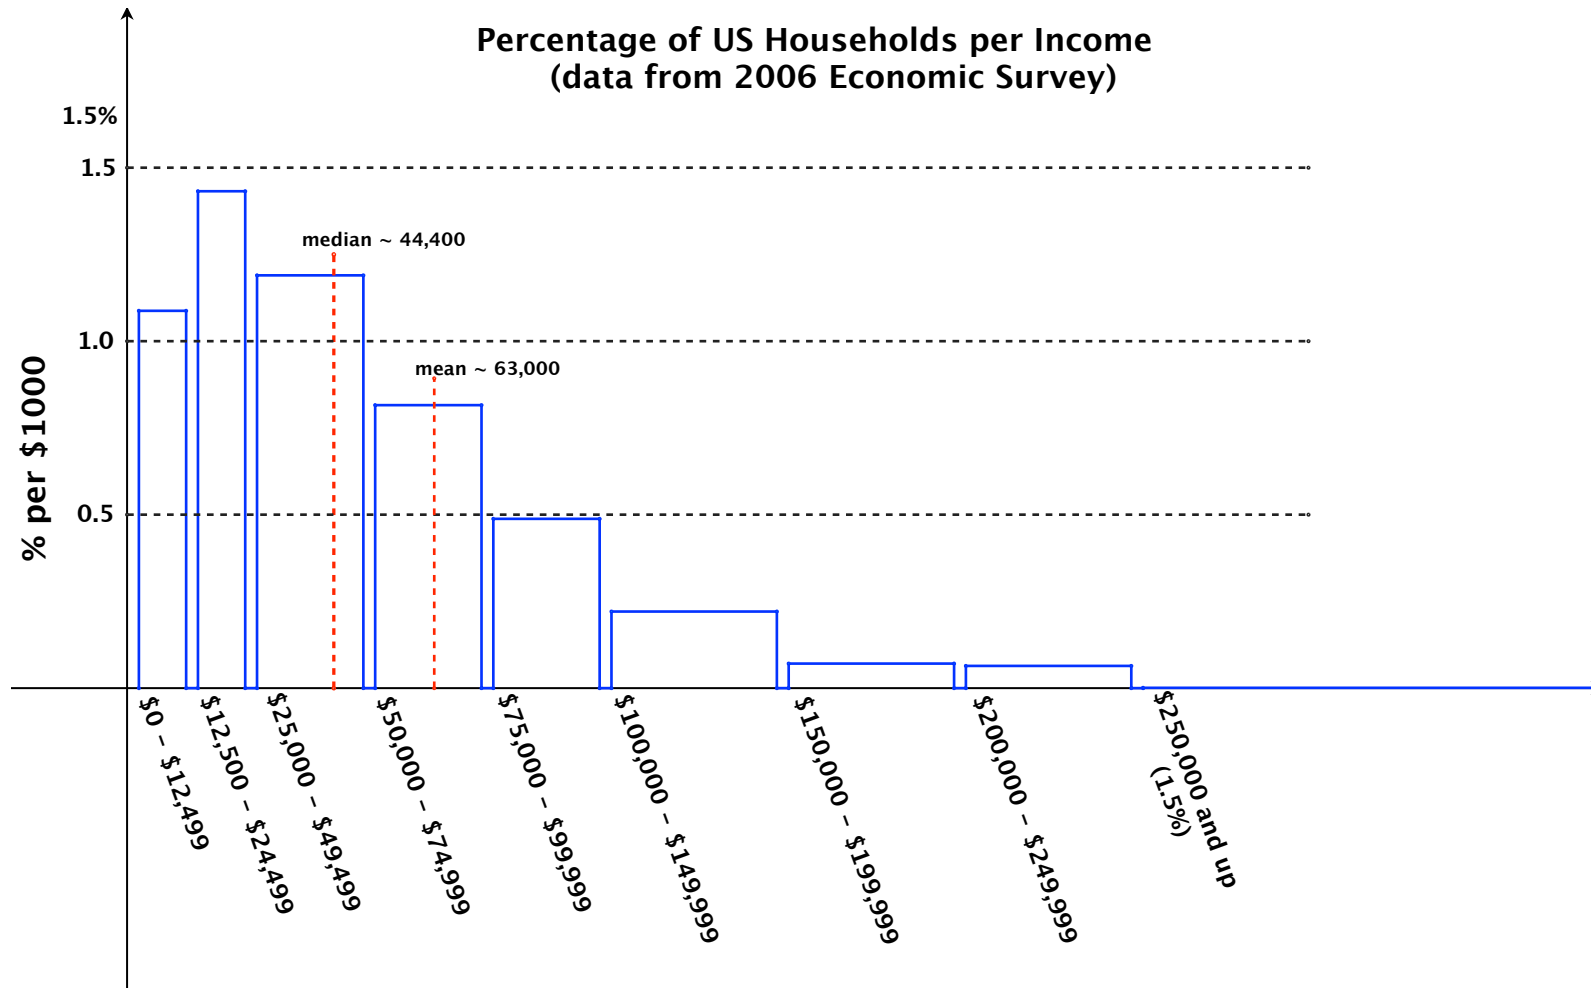


**Example:** The histogram below describes the distribution of household income in the United States. The data comes from the Wikipedia's summary of the *2006 Economic Survey*.



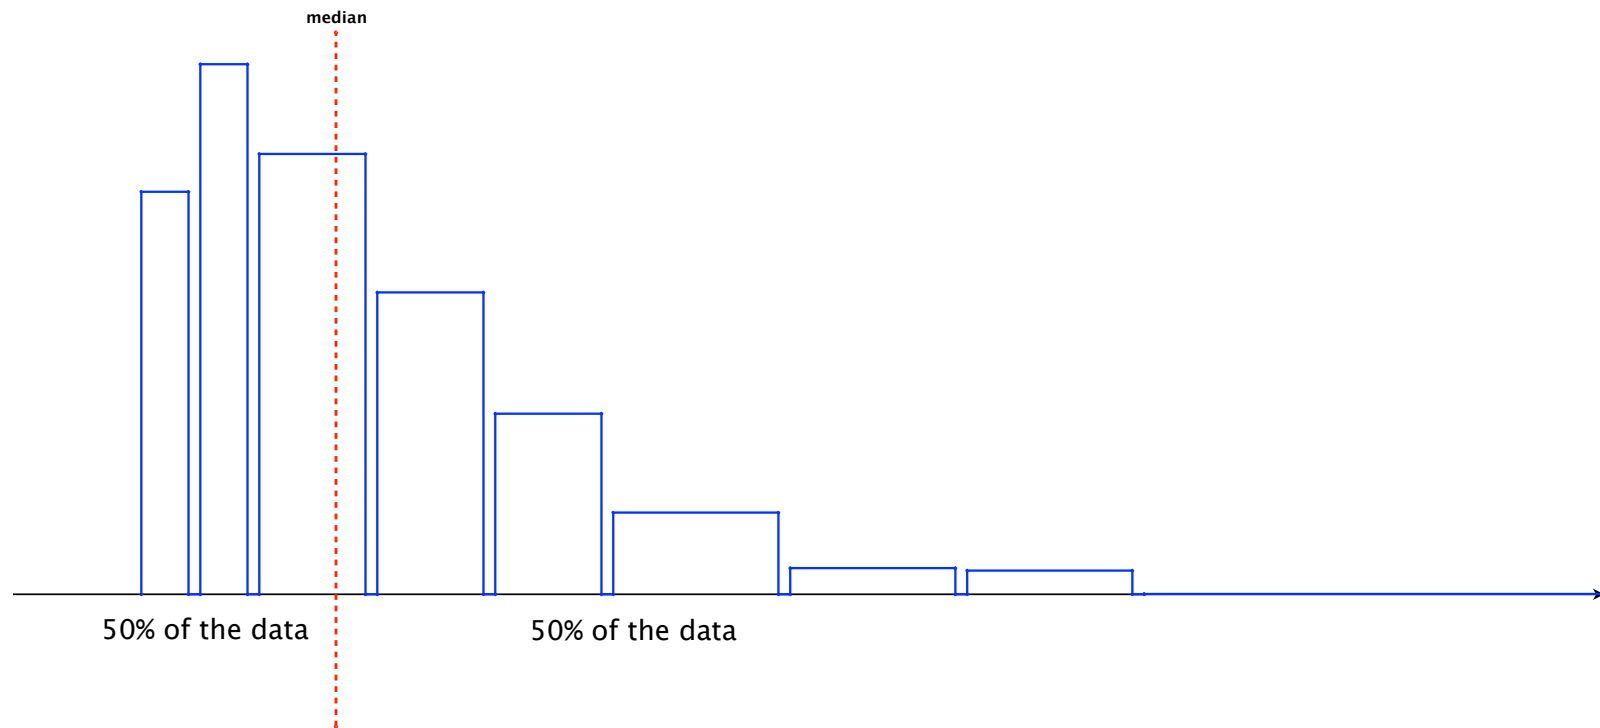


The distribution of household income is skewed to the right. The median household income in the survey was about \$44,400. The mean income was about \$63,000 (my estimate).




The mean and median describe the middle of the data in somewhat different ways:

- The median divides the histogram into two halves of equal area: it divides the data into two equal halves.



The mean and median describe the middle of the data in somewhat different ways:

- The mean is the ‘balance point’ of the data:




Averages and medians give a snapshot of a set of data. If the data comprises more than one variable, we can divide the data into categories with respect to one variable, and study the average/median of another variable in each category separately.

This allows researchers to discern relationships between different variables.

**Example:** The following graph comes from the 2005 American Community Survey of the US Census Bureau. It plots median household income by state.

Figure 1.  
**Median Household Income in the Past 12 Months With 90-Percent Confidence Intervals by State: 2005**



Source: U.S. Census Bureau, 2005 American Community Survey.

## Notation:

- The *population* mean (a parameter) is denoted by the Greek letter  $\mu$  ('mu'). If there are several variables being studied, we put a subscript on the  $\mu$  to tell us which variable it pertains to. For example, if we have data for population height ( $h$ ) and population weight ( $w$ ), the mean height would be denoted by  $\mu_h$  and the mean weight by  $\mu_w$ .
- The mean of a set of *sample* data (a statistic) is denoted by putting a bar over the variable. E.g., if  $\{h_1, h_2, h_3, \dots, h_n\}$  is a sample of heights, then the average of this sample would be denoted by  $\bar{h}$ .
- The median is usually denoted by  $m$  or  $M$ , and sometimes by  $Q_2$  (more on this later).
- We can use *summation notation* to simplify the writing of (long) sums:

$$h_1 + h_2 + h_3 + \dots + h_n = \sum_{j=1}^n h_j = \sum h_j.$$

For example we can write:

$$\begin{aligned}\bar{h} &= \frac{h_1 + h_2 + \cdots + h_n}{n} \\ &= \frac{1}{n} (h_1 + h_2 + \cdots + h_n) = \frac{1}{n} \sum h_j.\end{aligned}$$

**Comment:** The point of summation notation is to simplify expressions that involve sums with many terms, or in some cases, an unspecified number of terms. All the usual rules/properties of addition continue to hold. In particular

$$(i) \quad \sum (h_j \pm w_j) = \sum h_j \pm \sum w_j$$

$$(ii) \quad \sum (a \cdot h_j) = a \cdot \left( \sum h_j \right)$$

*and*

$$(iii) \quad \sum c = n \cdot c \quad (\text{here } n \text{ is the number of constant terms}).$$

## Measuring the spread of the data

The mean and median describe the middle of the data. To get a better sense of how the data is *distributed*, statisticians also use ‘*measures of dispersion*’.

- The *range* is the distance between the smallest and largest values in the data.
- The *interquartile range* is the distance between the value separating the bottom 25% of the data from the rest and the value separating the top 25% of the data from the rest. In other words, it is the *range* of the middle 50% of the data.

**Example:** In the histogram describing household income distribution, about 25% of all households have incomes below \$22,400 and about 25% of all households have incomes above \$78,000, so the interquartile range is  $78,000 - 22,400 = 65,600$ .



**The standard deviation:** The standard deviation of a set of numbers is *something like* the average distance of the numbers from their mean. Technically, it is a little more complicated than that.

If  $x_1, x_2, x_3, \dots, x_n$  are numbers and  $\bar{x}$  is their mean, then one candidate for measuring spread is the *average deviation* from the mean:

$$\frac{(x_1 - \bar{x}) + (x_2 - \bar{x}) + \dots + (x_n - \bar{x})}{n} = \frac{1}{n} \sum (x_j - \bar{x}).$$

**Potential problem:** positive terms and negative terms in the sum can cancel each other out... How much cancellation?

$$\begin{aligned} \frac{1}{n} \sum (x_j - \bar{x}) &= \frac{1}{n} \sum x_j - \frac{1}{n} \sum \bar{x} \\ &= \bar{x} - \frac{\overbrace{\bar{x} + \bar{x} + \dots + \bar{x}}^n}{n} \\ &= \bar{x} - \frac{\cancel{n} \cdot \bar{x}}{\cancel{n}} = \bar{x} - \bar{x} = 0 \end{aligned}$$

*Complete cancellation!*

Instead, statisticians use the *standard deviation*, which is given by

$$SD_x = \sqrt{\frac{1}{n} \sum (x_j - \bar{x})^2}.$$

In words, the SD is the *root of the mean of the squared deviations of the numbers from their mean*.

(\*) Squaring the deviations fixes the cancellation problem...

(\*) ... but exaggerates both very small deviations (making them smaller) and very large deviations (making them bigger)...

(\*) ... and also *changes the scale* (e.g., from inches to squared inches).

(\*) Taking the square root of the average squared deviation fixes both of these problems.