Problem set 02. Statistics and probability theory

Exercise 1. Mean versus median. A small accounting firm pays each of its five clerks \$35,000, two junior accountants \$80,000 each, and the firm's owner \$320,000. What is the mean salary paid at this firm? How many of the employees earn less than the mean? What is the median salary?

Odp. 1
Mean = \$81,875
Median = \$35,000
Everyone but the owner earns less than the mean wage. In this example, the median value is more indicative of the salary distribution among the employees.

Exercise 2. How does the median change? The firm in the previous exercise gives no raises to the clerks and junior accountants, while the owner's take increases to \$455,000. How does this change affect the mean? How does it affect the median?

Odp. 2
Mean changes to \$98 750 while the median remains unchanged.

Exercise 3. Distributions for time spent studying. We asked the students in a large first-year college class how many minutes they studied on a typical weeknight. Here are the responses of random samples of 30 women and 30 men from the class:

Women					Men				
180	120	180	360	240	90	120	30	90	200
120	180	120	240	170	90	45	30	120	75
150	120	180	180	150	150	120	60	240	300
200	150	180	150	180	240	60	120	60	30
120	60	120	180	180	30	230	120	95	150
90	240	180	115	120	0	200	120	120	180

The most common methods for formal comparison of two groups use $\bar{x}$ and s to summarize the data. We wonder if this is appropriate here.
Make the stemplot for this data set.
(a) What kinds of distributions are best summarized by $\bar{x}$ and s? It isn't easy to decide whether small data sets with irregular distributions fit the criteria.
(b) Each set of study times appears to contain a high outlier. Are these points flagged as suspicious by the $1.5 \times$ IQR rule? How much does removing the

outlier change $\bar{x}$ and s for each group? The presence of outliers makes us reluctant to use the mean and standard deviation for these data unless we remove the outliers on the grounds that these students were exaggerating.

Odp. 3

Stemplots:

Women		**Men**	
6	0	0	0
9	0	3	0 0 0 0
11	5	4	5
12	0 0 0 0 0 0 0	6	0 0 0
15	0 0 0 0	7	5
17	0	9	0 0 0 5
18	0 0 0 0 0 0 0 0 0 0 0	12	0 0 0 0 0 0 0
20	0	15	0 0
24	0 0 0	18	0
36	0	20	0 0
		23	0
		24	0 0
		30	0

Five-number summary

	Women	**Men**
Minimum	60	0
Q1	120	60
Median	175	120
Q3	180	150
Maximum	360	300

Mean and standard deviation

	Women	**Men**
Mean($\overline{x}$)	165.16	117.17
Std. deviation (σ)	56.51	74.23

a) Usually, the mean and standard deviation are most appropriate for distributions that are approximately symmetric without any strong outliers. Strong outliers tend to shift these parameters, thus complicating the distributions. For skewed or distributions with outliers, the median and IQR are generally more robust and better reflect a typical value

and spread. Here, the stemplots we created earlier are very useful in identifying outliers, and since the distributions show skew and possible outliers, the median/IQR are safer summaries.

b) IQR=Q3-Q1

Women: $\text{IQR}_{Women} = 180 - 120 = 60$

Men: $\text{IQR}_{Men} = 150 - 60 = 90$

IQR Rule: An outlier is defined when a value is lower than Q1 – 1.5×IQR or higher than Q3 + 1.5×IQR

Women:

$1.5\times\text{IQR}_{Women}=90$

Q1 – 90 = 30 and Q3 + 90 = 270

So, 360 is an outlier ($>[\text{Q3} + 1.5\times\text{IQR}_{Women}]$)

Men:

$1.5\times\text{IQR}_{Men}=90$

Q1 – 135 =-75 and Q3 + 135= 285

So, 300 is an outlier ($>[\text{Q3} + 1.5\times\text{IQR}_{Men}]$)

By removing the outlier from each group:

	Women	**Men**
Mean($\overline{x}$)	158.45	110.86
Std. deviation (σ)	43.65	66.88

Exercise 4. Hummingbirds and flowers. Different varieties of the tropical flower Heliconia are fertilized by different species of hummingbirds. Over time, the lengths of the flowers and the form of the hummingbirds' beaks have evolved to match each other. Here are data on the lengths in millimeters of three varieties of these flowers on the island of Dominica:

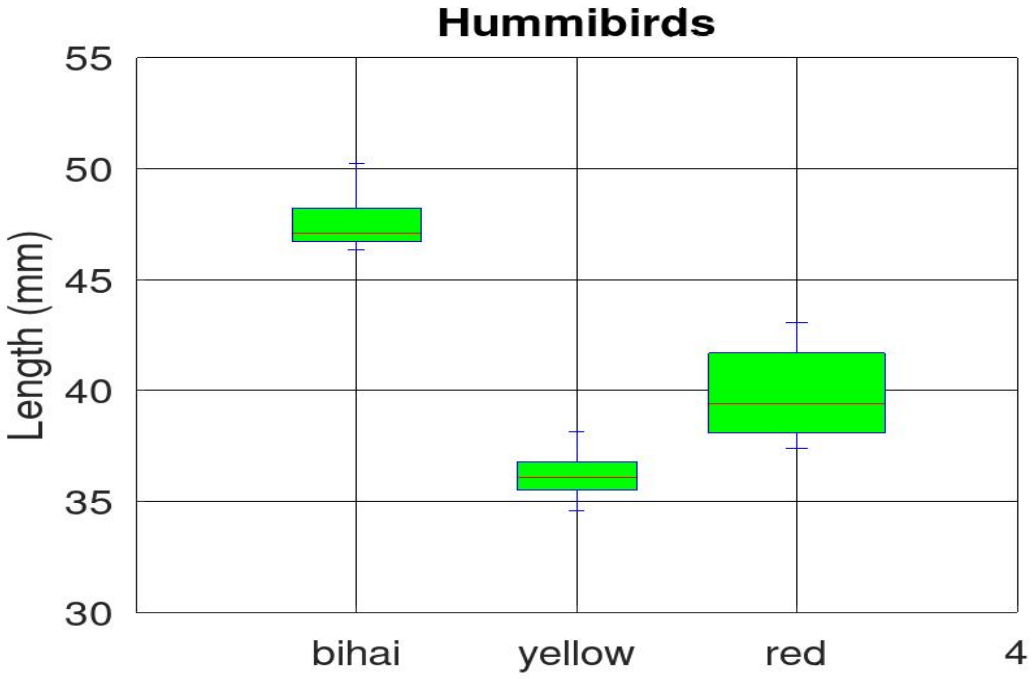
H. bihai							
47.12	46.75	46.81	47.12	46.67	47.43	46.44	46.64
48.07	48.34	48.15	50.26	50.12	46.34	46.94	48.36

H. caribaea red							
41.90	42.01	41.93	43.09	41.47	41.69	39.78	40.57
39.63	42.18	40.66	37.87	39.16	37.40	38.20	38.07
38.10	37.97	38.79	38.23	38.87	37.78	38.01	

H. caribaea yellow							
36.78	37.02	36.52	36.11	36.03	35.45	38.13	37.1
35.17	36.82	36.66	35.68	36.03	34.57	34.63	

Make boxplots to compare the three distributions. Report the five-number summaries along with your graph. What are the most important differences among the three varieties of flower?

Odp. 4



Five-number summary

	H. binai	**H. caribaea (red)**	**H. caribaea (yellow)**
Minimum	46.32	37.40	34.57
Q1	46.37	38.09	35.57
Median	47.12	39.16	36.11
Q3	48.20	41.58	36.80
Maximum	50.26	43.09	38.13

The three varieties are very well separated with no overlap between the medians. H. binai are longer and have narrow spacing. H. caribaea (red variety) is characterized by the widest spacing, whereas H. caribaea (yellow variety) is the shortest, has the narrowest spread, and the most symmetrical distribution. The boxplots are very useful here, as the three groups are very well separated, and the extraction of results for this data analysis is easy.

Exercise 5. A uniform distribution. If you ask a computer to generate "random numbers" between 0 and 1, you will get observations from a uniform distribution. Figure 1.37 graphs the density curve for a uniform distribution. Use areas under this density curve to answer the following questions.

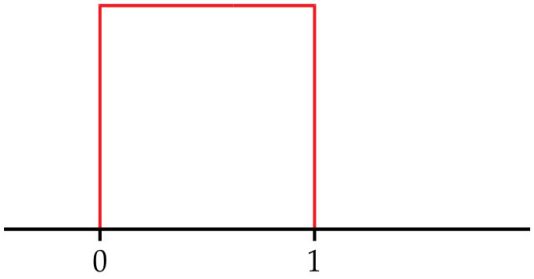


FIGURE 1.37 The density curve of a uniform distribution, for Exercise 1.108.

(a) Why is the total area under this curve equal to 1?
(b) What proportion of the observations lie below 0.35?
(c) What proportion of the observations lie between 0.35 and 0.65?

Odp. 5
a) Height: $= 1 \times 1 = 1$
b) $0.35 - 0 = 0.35$ (35%)
c) $0.65 - 0.35 = 0.30$ (30%)

Exercise 6. Use a different range for the uniform distribution. Many random number generators allow users to specify the range of the random numbers to be produced. Suppose that you specify that the outcomes are to be distributed uniformly between 0 and 4. Then the density curve of the outcomes has constant height between 0 and 4, and height 0 elsewhere.
(a) What is the height of the density curve between 0 and 4? Draw a graph of the density curve.

(b) Use your graph from (a) and the fact that areas under the curve are proportions of outcomes to find the proportion of outcomes that are less than 1.
(c) Find the proportion of outcomes that lie between 0.5 and 2.5.

Odp. 6

a) The area of the uniform distribution must always account for 1, so height $= \frac{1}{4} = 0.25$

b) $1 \times 0.25 = 0.25(25\%)$

c) $(2.5 - 0.5) \times 0.25 = 0.5(50\%)$

Exercise 7. Find the mean, the median, and the quartiles. What are the mean and the median of the uniform distribution in Figure 1.37? What are the quartiles?

Odp. 7

Mechanical interpretation of the mean value:

1. Discrete distribution of mas m. Let $x_1, x_2, \ldots, x_n$ be a given dataset. We place the same portion of mass m/n at each point x_i. Then **the mass center** $\overline{x}$ is obtained from the following formula

$$m\overline{x} = \frac{m}{n}x_1 + \frac{m}{n}x_2 + \frac{m}{n}x_3 + \cdots + \frac{m}{n}x_n$$

Hence, we get

$$\overline{x}_n = \frac{1}{n}(x_1 + x_2 + x_n + \cdots + x_n).$$

2. Continuous mass distribution. Let a mass m be distributed over the interval $[a, b]$ according to a density function $f(x)$, $x \in [a, b]$:

the mass Δm in the any interval $[\alpha, \beta] \subset [a, b]$ is equal

$$\Delta m = m \int_{\alpha}^{\beta} f(z)dz.$$

We divide the interval $[a, b]$ into system of subintervals $[z_i, z_{i+1}]$ $i = 0, 1, 2, \ldots, n$: $a = z_0 \leq z_1 \leq_2\leq \cdots \leq z_n = b$. Then the masses Δm_i concentrated in the intervals $[z_i, z_{i+1}]$, $i = 0, 1, \ldots, n-1$ are equal to $\Delta m_i = m \cdot \int_{z_i}^{z_{i+1}} f(z)dz$.

Let $z_{i+\frac{1}{2}}$, $i = 0, 1, 2, \ldots, n-1$ denote the centers of the subintervals. We discretize the model placing the portions $\Delta m_i = m \cdot f(z_{i+\frac{1}{2}})\Delta z_i$ of mass m at each point $z_{i+\frac{1}{2}}$. Then the mass center $\overline{x}$ (in discretized model) is given by the following formula

$$m \cdot \overline{x} \cdot = \sum_{i=0}^{n-1} = m \cdot \sum_{i=0}^{n-1} z_{i+\frac{1}{2}} f(z_{i+\frac{1}{2}})\Delta z_i$$

$$\overline{x} = \sum_{i=0}^{n-1} z_{i+\frac{1}{2}} f(z_{i+\frac{1}{2}})\Delta z_i$$

As $n \to \infty$, the sum on the right-hand side goes to $\int_a^b zf(z)dz$. Hence in the limit case we get the following formula for the mean value

$$\overline{x} = \int_a^b zf(z)dz.$$

a) $\mu = \int_0^1 z \cdot 1dz = 0.5$, median=0.5

b) Q1 = 0.25, because $\int_0^{0.25} 1dz = 0.25$ (=area over interval $[0, 0.25]$)

Q3 = 0.75, because $\int_0^{0.75} 1dz = 0.75$ (=area over interval $[0, 0.75]$)

Note
A uniform distribution is always symmetric, so the mean, the median, and the quartiles are evenly spaced, and the values are the same.
MEDIAN AND MEAN OF A DENSITY CURVE
The median of a density curve is the equal-areas point, the point that divides the area under the curve in half.
The mean of a density curve is the balance point, at which the curve would balance if made of solid material.
The median and mean are the same for a symmetric density curve. They both lie at the center of the curve. The mean of a skewed curve is pulled away from the median in the direction of the long tail.

Exercise 8. Horse pregnancies are longer. Bigger animals tend to carry their young longer before birth. The length of horse pregnancies from conception to birth varies according to a roughly Normal distribution with mean

336 days and standard deviation 3 days. Use the 68-95-99.7 rule to answer the following questions.
(a) Almost all (99.7%) horse pregnancies fall in what range of lengths?
(b) What percent of horse pregnancies are longer than 339 days?

Odp. 8
We know that for a normal distribution:

$\sigma:\ 68\%,\ 2\sigma:\ 95\%,\ 3\sigma:\ 99.7\%$
and for this exercise, we have:

$\mu = 336$ days, $\sigma = 3$ days or it can be written as 336 ± 3 days

a) For 99.7% pregnancies:
$336 \pm 3\sigma$, so $336 - 3 \times 3 = 327$ and $336 + 3 \times 3 = 345$ so the range is $[327, 345]$
Thus, almost all horse pregnancies are between 327 and 345 days.

b) We want > 339 days:
339 is $\mu + \sigma$(68% of the population)
$100\% - 68\% = 32\%$, divided by 2 as we want only the positive part so 16%

Exercise 9. Park space and population. Below are data on park and open space in several U.S. cities with high population density. In this table, population is reported in thousands of people, and park and open space is called open space, with units of acres.

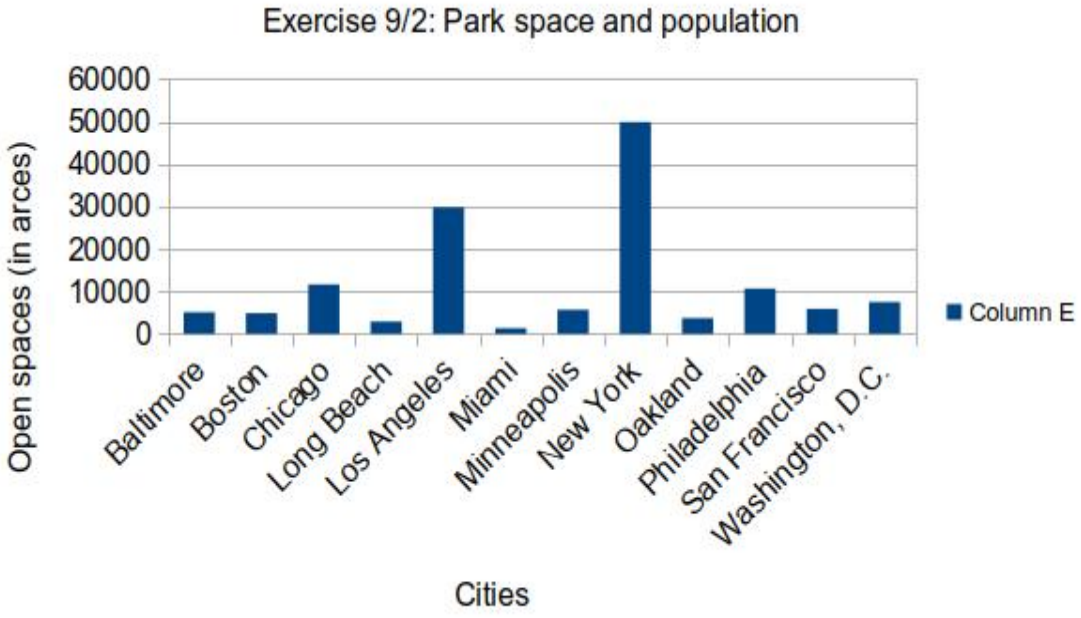
City	Population	Open space
Baltimore	651	5,091
Boston	589	4,865
Chicago	2,896	11,645
Long Beach	462	2,887
Los Angeles	3,695	29,801
Miami	362	1,329
Minneapolis	383	5,694
New York	8,008	49,854
Oakland	399	3,712
Philadelphia	1,518	10,685
San Francisco	777	5,916
Washington, D.C.	572	7,504

(a) Make a bar graph for population. Describe what you see in the graph.
(b) Do the same for open space.

(c) For each city, divide the open space by population. This gives rates: acres of open space per thousand residents.
(d) Make a bar graph of the rates.
(e) Redo the bar graph that you made in part (d) by ordering the cities by their open space to population rate.
(f) Which of the two bar graphs in (d) and (e) do you prefer? Give reasons for your answer.

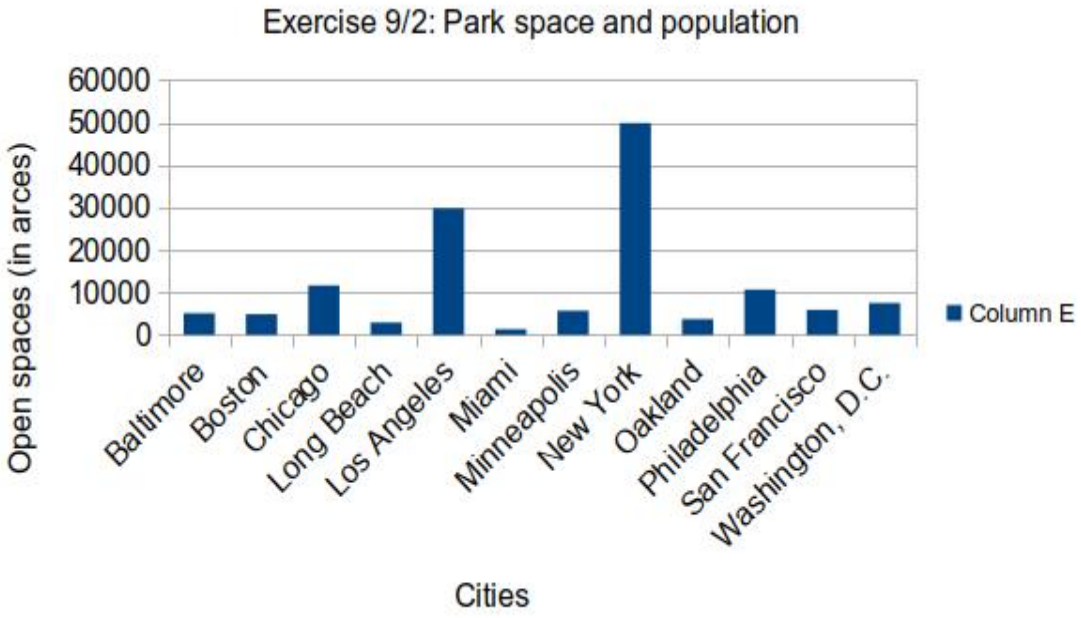
Odp. 9

a) For the population graph



New York has the highest population, followed by Los Angeles and Chicago third.

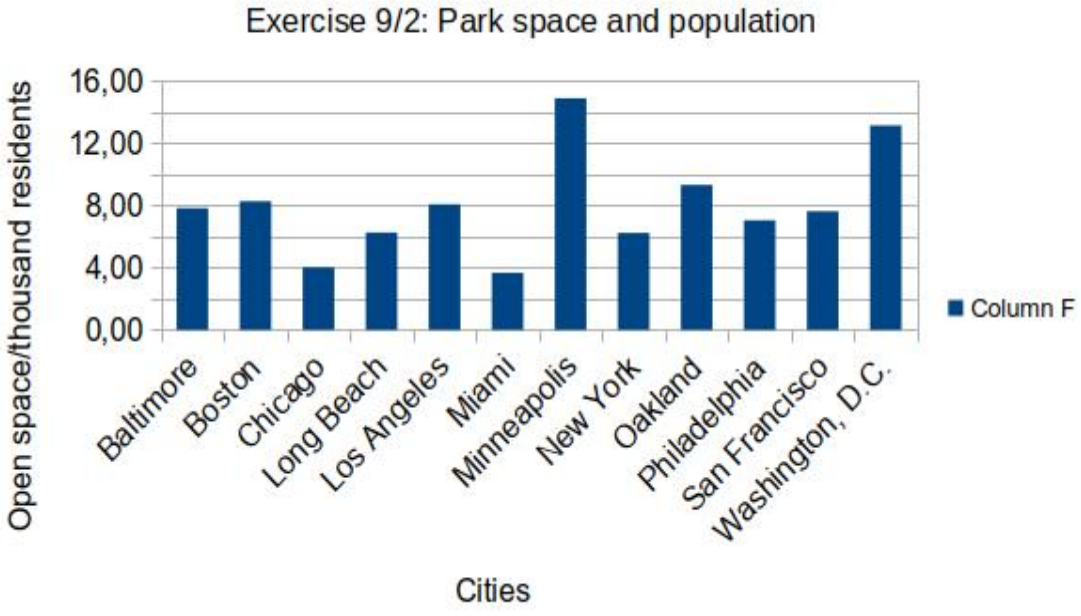
b) For the open space graph



The same trend as with the population graph can be observed.

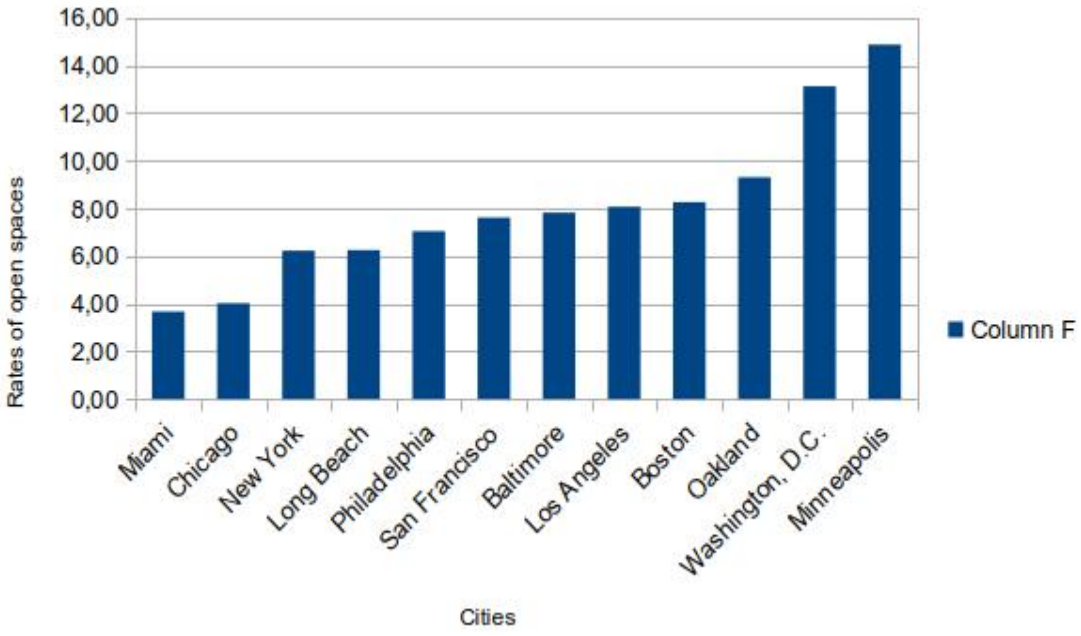
c) We divide open space by population

d) Bar graph of acres of open space per thousand residents



This is the true value we are interested in. Now, we see that Minneapolis has the highest open space per thousand residents, followed by Washington D.C. The three largest cities from before have dropped in the rankings.

e) Bar graph ordered from highest to lowest open space per thousand residents.



f) The last graph makes comparisons and rankings among the U.S. states much easier, where we can immediately define which is larger and draw conclusions between the states. However, by ordering the bar graph in alphabetical order, we can directly pinpoint the U.S. state we are interested in. Overall, it depends on what we are actually looking for.

Exercise 10. Compare two Normal curves. We use the fact that the distribution of scores for the 76,531 students who took the exam was approximately N(572, 51). These students were classified in a variety of ways, and summary statistics were reported for these different subgroups. When classified by gender, the scores for the women are approximately N(579, 49), and the scores for the men are approximately N(565, 55). Figure 1.43 gives the Normal density curves for these two distributions. Here is a possible description of these data: women score about 14 points higher than men on the ISTEP English/language arts exam. Critically evaluate this statement and

then write your own summary based on the distributions displayed in Figure 1.43.

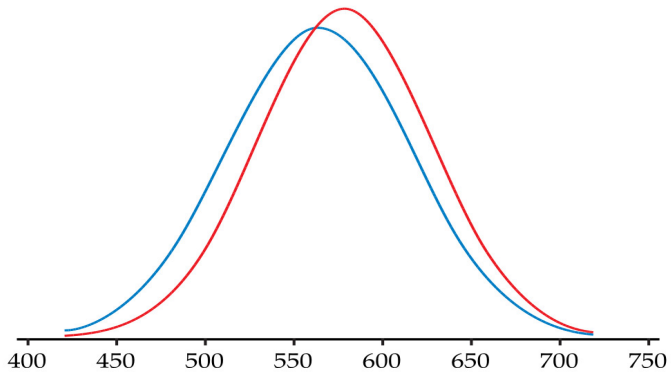


FIGURE 1.43 Normal density curves for ISTEP scores of women and men, for Example 1.53.

FIGURE 1.43 Normal density curves for ISTEP scores of women and men, for Example 1.53.

Odp. 10

Women: 579 ± 49
Men: 565 ± 55

The statement that women score about 14 points higher than men on the exam is partially true. Indeed, judging from the mean values, the difference is 14, but the two graphs are overlapping and we can see that some men have scored higher points than some women. Thus, the term is a bit exaggerated and it should be corrected to women score higher than men on average.

Exercise 11. Biological clocks. Many plants and animals have "biological clocks" that coordinate activities with the time of day. When researchers looked at the length of the biological cycle in the plant Arabidopsis by measuring leaf movements, they found that the length of the cycle is not always 24 hours. Further study discovered that cycle length changes systematically with north-south location.

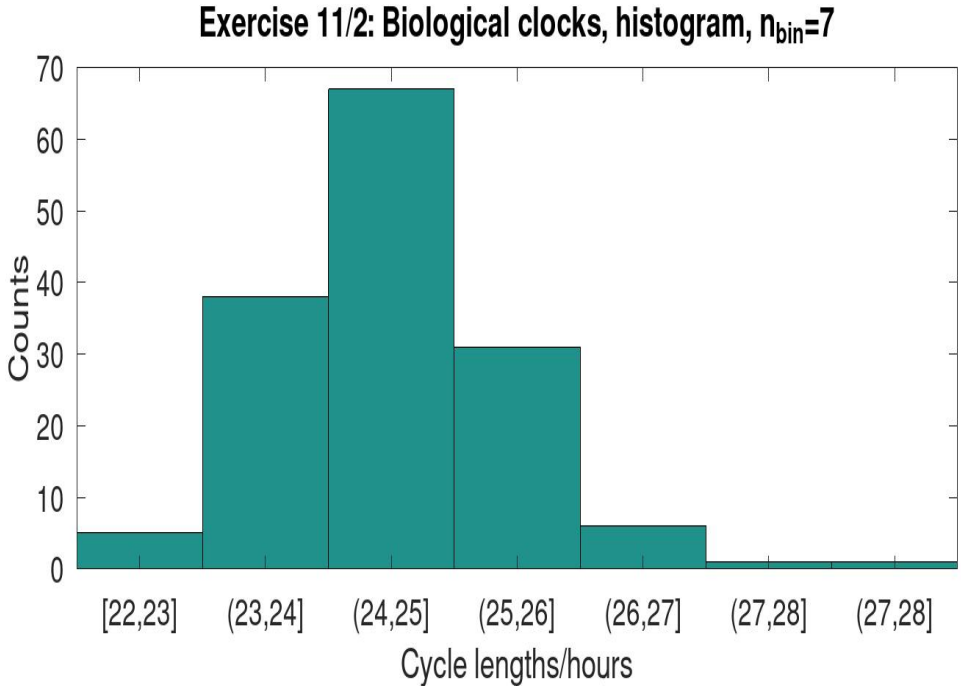
Biological clock cycle lengths for a plant species in different locations

23.89	23.72	23.74	24.35	25.05	24.56	23.69	22.33	23.79	22.12
25.39	23.08	25.64	23.98	25.84	25.46	24.37	24.13	24.40	24.74
24.44	24.82	23.56	24.96	24.21	23.85	24.57	23.44	23.64	24.23
24.01	24.58	25.57	23.73	24.11	23.21	25.08	24.03	24.62	23.51
23.21	23.41	23.69	22.97	24.65	24.65	24.29	23.89	25.08	23.89
24.95	23.09	23.21	24.66	23.88	25.33	24.38	24.68	25.34	25.22
23.45	23.39	25.43	23.16	23.95	23.25	24.72	24.89	24.88	24.71
23.58	25.98	24.28	24.25	23.16	24.19	27.22	23.77	26.21	24.33
24.34	24.89	24.32	24.14	24.00	23.48	25.81	24.99	24.18	22.73
24.18	23.95	24.48	23.89	24.24	24.96	24.58	24.29	24.31	23.64
23.87	23.68	24.87	23.00	23.48	24.26	23.34	25.11	24.69	24.97
24.64	24.49	23.61	24.07	26.60	24.91	24.76	25.09	26.56	25.13
24.81	25.63	25.63	24.69	24.41	23.79	22.88	22.00	23.33	25.12
24.00	24.31	23.03	24.51	28.55	22.96	23.61	24.72	24.04	25.18
24.30	24.22	24.39	24.73	24.68	24.14	24.57	24.42	25.62	

Table 1.11 contains cycle lengths for 149 locations around the world. Describe the distribution of cycle lengths with a histogram and numerical summaries. In particular, how much variation is there among locations?

Odp. 11

Total (n):	149
Five-number summary	
Minimum	22.00
Q1	23.74
Median	24.31
Q3	24.82
Maximum	28.55



The distribution is slightly right-skewed because of two outliers, mainly. The distribution is tight around the median of 24.31 with a very small spread.

Exercise 12. Product preference. Product preference depends in part on the age, income, and gender of the consumer. A market researcher selects a large sample of potential car buyers. For each consumer, she records gender, age, household income, and automobile preference. Which of these variables are categorical and which are quantitative?

Odp. 12

Gender: Categorical
Age: Quantitative
Income: Quantitative
Preference: Categorical

Exercise 13. Distance-learning courses. The 222 students enrolled in distance-learning courses offered by a college ranged from 18 to 64 years of age. The mode of their ages was 19. The median age was 31. Explain how this can happen.

Odp. 13

Mode: The most frequent value
Median: The "middle" value
This happens because many students are 19 years old but fewer than half (because of the median).

Exercise 14. Norms for reading scores. Raw scores on behavioral tests are often transformed for easier comparison. A test of reading ability has mean 75 and standard deviation 10 when given to third-graders. Sixth-graders have mean score 82 and standard deviation 11 on the same test. To provide separate "norms" for each grade, we want scores in each grade to have mean 100 and standard deviation 20.
(a) What linear transformation will change third-grade scores x into new scores $x_{new} = a + bx$ that have the desired mean and standard deviation? (Use $b > 0$ to preserve the order of the scores.)
(b) Do the same for the sixth-grade scores.
(c) David is a third-grade student who scores 78 on the test. Find David's transformed score. Nancy is a sixth-grade student who scores 78. What is her transformed score? Who scores higher within his or her grade?
(d) Suppose that the distribution of scores in each grade is Normal. Then both sets of transformed scores have the $N(100, 20)$ distribution. What percent of third-graders have scores less than 78? What percent of sixth-graders have scores less than 78?

Odp. 14

3rd graders: 75 ± 10
6th graders: 82 ± 11
We need 100 ± 20
Equation: $x_{new} = a + bx$
a) For 3rd graders:

$b = \frac{\sigma_{new}}{\sigma_{old}} = \frac{20}{10} \Leftrightarrow b = 2$

$a = \mu_{new} - b\mu_{old} = 100 - 2 \times 75 \Leftrightarrow a = -50$

$x_{new} = -50 + 2x$

b) Similarly for 6th graders: $b = 1.818$, $a = -49.08$

$x_{new} = -49.08 + 1.818x$

c) David (3rd grader)

$x_{new} = -50 + 2 \times 78 \Leftrightarrow x_{new} = 106$

Nancy (6th grader)

$x_{new} = -49.08 + 1.8 \times 78 \Leftrightarrow x_{new} = 92.72$

d) Note: Because the transformation of the distributions is linear, using both new and old distributions will give the same result.

3rd graders

$z_{3rd} = \frac{x-\mu}{\sigma} = \frac{106-100}{20} \Leftrightarrow z_{3rd} = 0.3$

6th graders

Similarly:

$z_{6th} = \frac{x-\mu}{\sigma} = \frac{92.72-100}{20} \Leftrightarrow z_{6th} = -0.364$

By checking tables of the standard normal distribution we find out that the percentage for scores below 78 are:
For 3rd graders: 61.8%
For 3rd graders: 35.8%

Exercise 15. Damage caused by tornadoes and amount of oil recovered from some 64 wells. The average damage caused by tornadoes in the states

State	Damage ($ millions)	State	Damage ($ millions)	State	Damage ($ millions)
Alabama	51.88	Louisiana	27.75	Ohio	44.36
Alaska	0.00	Maine	0.53	Oklahoma	81.94
Arizona	3.47	Maryland	2.33	Oregon	5.52
Arkansas	40.96	Massachusetts	4.42	Pennsylvania	17.11
California	3.68	Michigan	29.88	Puerto Rico	0.05
Colorado	4.62	Minnesota	84.84	Rhode Island	0.09
Connecticut	2.26	Mississippi	43.62	South Carolina	17.19
Delaware	0.27	Missouri	68.93	South Dakota	10.64
Florida	37.32	Montana	2.27	Tennessee	23.47
Georgia	51.68	Nebraska	30.26	Texas	88.60
Hawaii	0.34	Nevada	0.10	Utah	3.57
Idaho	0.26	New Hampshire	0.66	Vermont	0.24
Illinois	62.94	New Jersey	2.94	Virginia	7.42
Indiana	53.13	New Mexico	1.49	Washington	2.37
Iowa	49.51	New York	15.73	West Virginia	2.14
Kansas	49.28	North Carolina	14.90	Wisconsin	31.33
Kentucky	24.84	North Dakota	14.69	Wyoming	1.78

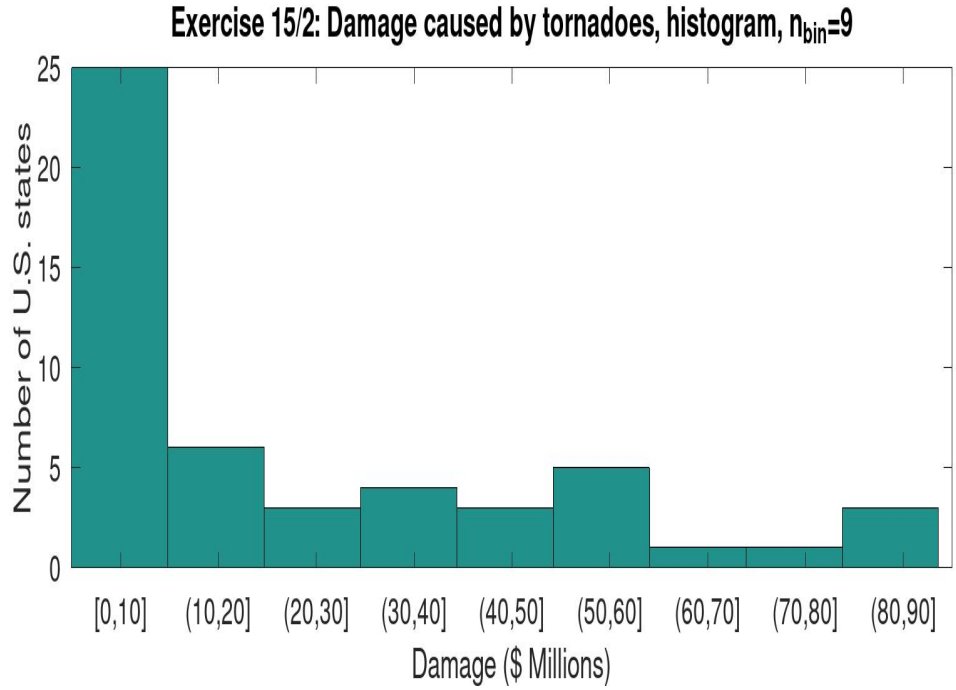
Here are the estimated total amounts of oil recovered from 64 wells in the Devonian Richmond Dolomite area of the Michigan basin, in thousands of barrels:

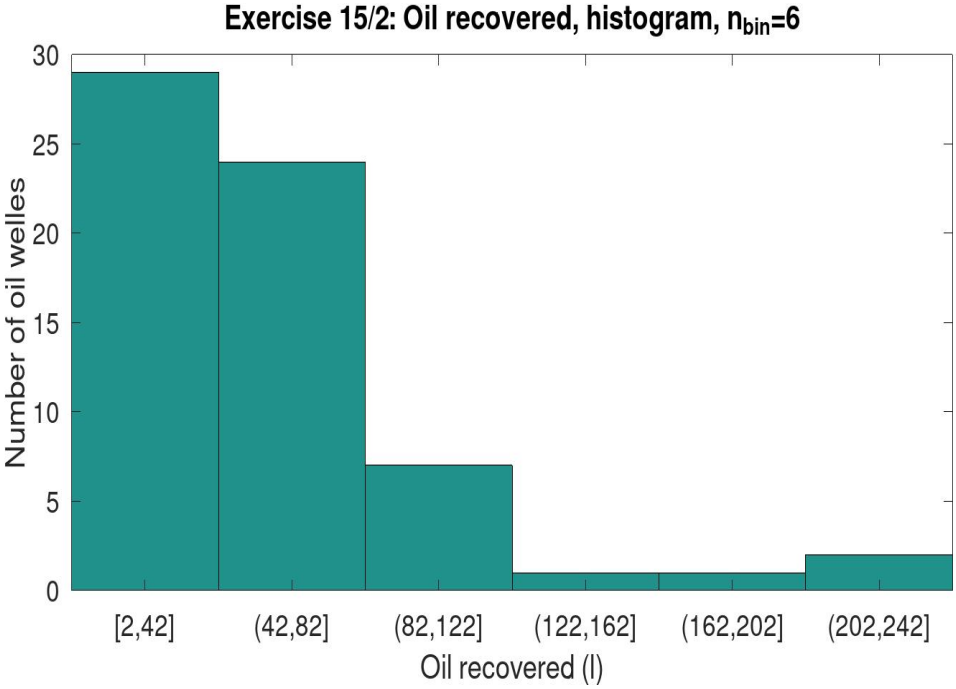
21.7	53.2	46.4	42.7	50.4	97.7	103.1	51.9
43.4	69.5	156.5	34.6	37.9	12.9	2.5	31.4
79.5	26.9	18.5	14.7	32.9	196.0	24.9	118.2
82.2	35.1	47.6	54.2	63.1	69.8	57.4	65.6
56.4	49.4	44.9	34.6	92.2	37.0	58.8	21.3
36.6	64.9	14.8	17.6	29.1	61.4	38.6	32.5
12.0	28.3	204.9	44.5	10.3	37.7	33.7	81.1
12.1	20.1	30.5	7.1	10.1	18.0	3.0	2.0

Make histograms of those distributions and show that they both are right-skewed.

Odp. 15

Looking at both histograms, we see that most data is centered towards lower values with longer tails to the right (less counts in higher values)


Exercise 15/2: Damage caused by tornadoes, histogram, n_{bin}=9
Number of U.S. states
0
5
10
15
20
25
[0,10]
(10,20]
(20,30]
(30,40]
(40,50]
(50,60]
(60,70]
(70,80]
(80,90]
Damage ($ Millions)


Exercise 15/2: Oil recovered, histogram, n_{bin}=6
Number of oil welles
0
5
10
15
20
25
30
[2,42]
(42,82]
(82,122]
(122,162]
(162,202]
(202,242]
Oil recovered (l)