Chapter 1

Discussion Problem Solutions

D1. Reasonable suggestions at this stage include: compare the average age of those laid off with the average age of those retained; compare the proportion of those, say, age 50 or older who were laid off with the proportion of those under age 50 who were laid off; make a graph that shows both the ages of those laid off and those not laid off. (If students suggest this option, you might take the opportunity to show them how to make a back-to-back stem-and-leaf plot.) Note: Students may do a lot or a little with this question. In some classes, students may offer only very general ideas. On the other hand, students who have learned to "take charge" in their mathematics classes may want to thoroughly explore data tables, bar graphs, and averages before continuing with the discussion. At any rate, all students should see the need for some method of summarizing and displaying the data in Display 1.1.
D2. Both groups, hourly and salaried, show similar patterns: older workers were more likely than younger ones to lose their jobs. From the dot plots alone, it is hard to decide which group, hourly or salaried, provides stronger evidence of age bias. The salaried workers are a larger group, although the difference in ages between laid off and retained may be more pronounced for the hourly workers. On balance, though, variability within each group makes it hard to judge the strength of the evidence from the dot plots. Useful as graphs are, they have limitations, and there are times when numerical summaries are helpful. Most of the time, a good analysis involves both graphical displays and appropriate numerical summaries.
D3. Statistical inference can provide an answer to this question, and Section 1.2 will sketch how this is done. For now, the purpose of the question is to involve students in thinking about what the question asks and why it is important. Note: Students may be worried that there are too few workers to make conclusions or that the differences are too small to make any conclusions. Try to get them to focus on the crucial question: Do the patterns look like the kind you would expect to occur just by chance if there was no age discrimination? No calculations need to be done; have students simply observe the pattern.
D4. Considerations that might explain how older workers could have been laid off disproportionately without being victims of age discrimination include: Older workers tended to hold jobs that had become obsolete; or, older workers tended not to have the up-to-date training needed to use Westvaco's computers or other technology.
D5. No, it is not likely. To get an average age of 58 or greater, you couldn't select any worker under age 55—and half the workers are under age 55.

D6. A small probability favors Martin, showing it is unlikely for a company to lay off, just by chance, workers as old as those laid off by Westvaco.
D7. The probability of getting any specified average age is quite small. For example, if Westvaco laid off the 35-, 38-, and 48-year-olds, the average age would be 40.3 and there certainly would have been no suspicion of age discrimination. However, if you look at Display 1.10, you see that the probability of getting an average age of exactly 40.3 is actually smaller than getting an average age of 58. Westvaco shouldn't be under suspicion because the average age of its laid-off workers was exactly 58, and getting exactly 58 just by chance is very unlikely; Westvaco should be under suspicion because the average age it got is in the extreme upper part of the distribution of all possibilities. That is, it is the location of age 58 in the distribution that is important, not the probability of getting that exact number.
D8. a. No, the probability is larger than 0.025. **b.** If the Martin probability had been 0.01, it would have met this requirement and Westvaco owes the court an explanation. If it had been 0.10, it would not have met the requirement and Westvaco would not be required to defend itself. Note about D8 and significance levels: Using any one choice of threshold for deciding statistical significance is somewhat arbitrary, even though fixed-level testing commonly uses a significance level of $\alpha = 0.05$. There is no best value (except in a strict decision-theoretic setting where you know the costs of Type I and Type II errors), but there is an unambiguous ordering: the smaller the probability, the stronger the evidence against the null hypothesis. Though students may legitimately disagree about how low a value they personally would require in order to rule out the chance model, they should all come to recognize that a lower significance level means stronger evidence is required.
D9. a. If the coin were fair, it would be very unlikely to get 19 heads or more in 20 flips. (The probability is $\approx$ 0.00002.) **b.** The chance model is that the coin is fair. Students should flip a fair (or assumed to be fair) coin 20 times and count the number of heads. The number of heads is the summary statistic. This process should be repeated as many times as you have time for in your class. You may display the distribution on a dot plot. A typical distribution for 500 repetitions is given here. (Each dot represents 10, or fewer, repetitions.) 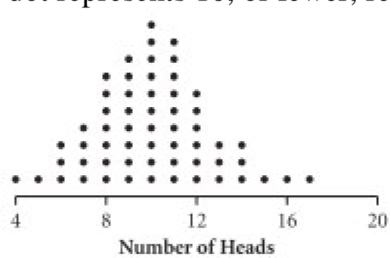   In this simulation of 500 repetitions, the largest number of heads ever to appear was 17. The probability of getting 19 or more heads is very close to 0. If the coin is fair, it is extremely unlikely to get 19 or 20 heads, so we conclude that the coin is not fair.

Summarizing, here's the logic: Assume the model is correct that your friend's coin is fair.

Question: Is the actual data (19 heads in 20 flips) easy to get from this model?
Answer: No, it's almost impossible to get 19 heads in 20 flips of a fair coin. Conclusion: Your friend's coin is not fair. Emphasize to your students that you can estimate the chance of getting 19 or more heads from a fair coin, but you cannot estimate the chance that your friend's coin is fair. If you assume the model is correct (the coin is fair), you can estimate the probability you need by simulation. Notice, however, that it doesn't work to assume that outcome (19 heads) and then use simulation to estimate the chance of a fair coin.

Practice Problem Solutions

P1. This plot is rather striking; the older hourly workers were far more likely to be laid off in Rounds 1 through 3 than were the younger workers.

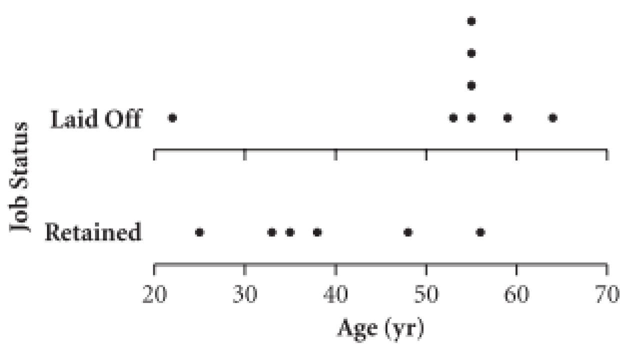

P2. B.

P3. Overall, the patterns in the two displays are strikingly similar. Specifically, both graphs show that older workers were especially likely to be targeted for layoff in the earlier rounds, and that younger workers were more likely to be targeted in Rounds 4 and 5. Both displays also show that most of the layoffs were planned in the first two rounds. This sequence of layoffs for hourly workers provides stronger evidence for Martin's case. In the first three rounds, seven of the eight people laid off were over 50, while only one of the six retained in these three rounds was over 50.

P4. **a.** The average age is now 52.67. On the dot plot, this is the column that is just a bit more than half-way between 50 and 55. That is, it is the tallest column in the group of seven columns of dots *between* 50 and 55. The number of repetitions out of 200 that gave an average age of 52.67 or larger is 37. Thus, the estimated probability of getting an average age of 52.67 or larger is 37 out of 200 or about 0.185.

b. There is no evidence of age discrimination because an average age of 52.67 or larger is relatively easy to get just by chance.

P5. **a.** The ten workers laid off were ages 22, 33, 35, 53, 55, 55, 55, 55, 59, and 64. Their average age was 48.6 years.

b. The chance model is that 10 workers were selected at random for layoff from 14 workers with ages 22, 25, 33, 35, 38, 48, 53, 55, 55, 55, 55, 56, 59, and 64. Write these ages on 14 cards of equal size. Place the cards in a bag or box, mix the cards up by shaking the bag, and then draw 10 of the cards at random. Calculate and record the average age of the ten picked cards. Then, replace the cards, mix, and repeat many times, recording the average age each time. Make a dot plot showing the distribution of these average ages and compute the proportion of times you get an average age of at least 48.6.

c. 45 of the 200 dots are at or above 48.6, for a proportion of 22.5%.
d. No, this evidence does not help Martin's case. Martin wants to show that the company systematically targeted older workers for layoff. These results show that we would expect an average age of 48.6 or more, purely by chance, about 22.5% of the time, which is not a rare occurrence.

Exercise Solutions

E1. a. There were 6 hourly workers under age 50, and 3 were laid off. So 50% of the hourly workers under age 50 were laid off. 50% were not laid off.

b. 10 of the hourly workers were laid off. 3 were under age 50, so 3/10 = 0.3 of the hourly workers who were laid off were under age 50. 70% were age 50 or over.

c. You should compute the proportion of hourly workers under age 50 who were laid off and the proportion of those age 50 or over who were laid off. 50% (3 out of 6) of those under age 50 and 87.5% (7 out of 8) of those ages 50 and over were laid off. Thus, a disproportionately high proportion of workers age 50 and over were laid off.

d. For the salaried workers, 37.5% of those under age 50 were laid off and 60% of those age 50 and older were laid off. For both hourly and salaried workers, the proportion of workers age 50 and over who were laid off is disproportionately high, but the difference is more extreme in the case of the hourly workers and so appears to provide more evidence for Martin.

However, note that the number of hourly workers in each age category is very small. Thus, a small change in the number laid off in each category would make a large change in the difference in the proportions laid off. For example, suppose Westvaco had laid off one more worker under age 50 instead of one of the workers age 50 or over. Then the difference in the proportions would be 4 out of 6 or 67% of those under age 50 were laid off and 6 out of 8 or 75% of the workers age 50 or over were laid off, now, there is not a very big difference.

E2. a. 14 out of 27, or about 52% of the workers age 40 and over were laid off. 14 out of 18, or about 78% of the workers who were laid off were age 40 or over.

b. 4 out of 9, or about 44% of the workers under age 40 were laid off. So, about 56% of the workers under age 40 were not laid off.

c. Since your goal is to determine whether workers age 40 and over are treated differently from workers under age 40, you would compute and compare the proportion of workers age 40 and over who were laid off to the proportion of workers under 40 who were laid off. Forty-four percent of workers under age 40 were laid off, but 52% of workers age 40 and over were laid off. This favor's Martin's case but not a lot because 8% isn't a very large difference considering the relatively small numbers of people.

d. For this table of the salaried workers, 44% of those under 40 were laid off and 52% of those age 40 and over were laid off. From the table of the salaried workers, 6/16 (37.5%) of those under 50 were laid off and 12/20 (60%) of those age 50 or over were laid off. There is a greater discrepancy between the proportions using age 50 as the cutoff, so Martin's lawyer would want to use age 50 as the cut-off age.

E3. a.

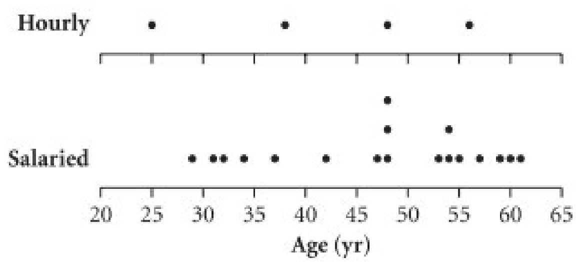

Only four hourly workers kept their jobs, making comparison difficult. However, the distribution of salaried workers who kept their jobs generally falls to the right of the distribution of hourly workers who kept their jobs. Indeed, the average age of the hourly workers who kept their jobs was 41.75, whereas it was 47.17 for the salaried workers.

b. We cannot conclude this from the dot plots in part a alone. We would have to take into consideration the age distributions for hourly and salaried workers *before* the layoffs began, as well as after. All we can say is that the salaried workers who kept their jobs tended to be older than the hourly workers who kept their jobs.

E4. a.

	Laid Off	Retained	Total
Hourly	10	4	**14**
Salaried	18	18	**36**
Total	**28**	**22**	**50**

b. To see if hourly workers and salaried workers are treated differently, compute and compare the proportions of each type of worker who were laid off. For hourly workers, 10 out of 14, or 71.4%, were laid off. For salaried workers, 18 of 36, or 50%, were laid off.

c. From these proportions, it is clear that hourly workers were more likely than salaried workers to lose their jobs.

E5. a.

Round	Number Laid Off	40 or Older	Percentage 40 or Older
1	11	9	82%
2	9	8	89%
3	3	2	67%
4	4	2	50%
5	1	0	0%

b. Most of the layoffs came early. Of the 28 who were chosen for layoff, 20 were identified in the first two rounds. Only 8 were chosen for layoff in the last three rounds altogether.

Early rounds hit the older workers harder; later rounds tended to have higher percentages of younger workers. In the first two rounds, the percentage was very high among those over 40 who were laid off, at 85% (17 of 20). In the later rounds, the percentage for those over 40 dropped to 67% in Round 3 (2 of 3), 50% in Round 4 (2 of 4), and 0% (0 of 1) in Round 5.

The pattern is consistent with what you would expect to see if the department head who planned the layoffs was trying to cut costs by laying off the older, more experienced, and thus possibly more expensive workers first.

E6. One reasonable way to approach such an analysis is to follow the analysis of age, using year of hire as the variable. (This is equivalent to working with a new variable, *seniority,* defined as 1991 – *hire year.*) This dot plot compares the years of hire for those laid off and those retained.

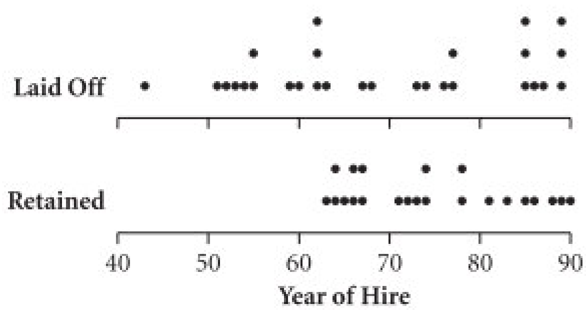

One striking feature is that all 12 of those hired in 1962 or earlier lost their jobs. Of those hired in 1985 or later, 8 of 13, or 62%, lost their jobs. Of those hired in the middle years, 1963–1984, only 8 of 25, or 32%, lost their jobs.

b. The information seems to suggest that "last hired, first fired" was not the policy, except for the most recently hired (1985–1990).

Year of Hire	Laid Off	Retained	Total	% Laid Off
62 or before	12	0	12	100.0%
63–74	5	13	18	27.8%
75–84	3	4	7	42.9%
85–90	8	5	13	61.5%
Total	**28**	**22**	**50**	**56.0%**

E7. a. The cases are the five individual players and the variables are games, at bats, hits, home runs, stolen bases, and batting average. Brock's batting average is 0.293. Hamilton's number at bats is 6276 or 6277. Raines' number of hits is about 2608.

b. The cases are the seven individual stocks and the variables are closing price on 10/28, closing price on 10/30, change (in price), percentage change (in price), and volume (of trading). Coca-Cola change is –8 5/8, Coca-Cola percentage change is –6.30. Eastman Kodak closing price on 10/30 is 170. General Motors closing price on 10/28 is 47 1/2. General Motors percentage change is –15.79. US Steel change is –12. US Steel percentage change is –6.45.

E8. Possible responses include income; proportion of people living below some defined poverty level; employment rates; health data such as life expectancy, infant mortality rate, incidence of diseases; social data such as crime rates, educational opportunities, average family size, waste disposal.

E9. a. Simulations will vary. For example, a student might draw the following 10 ages: 22, 25, 33, 48, 53, 55, 55, 55, 56, and 59. The summary statistic would then be 7 because 7 of the 10 are age 40 or older.

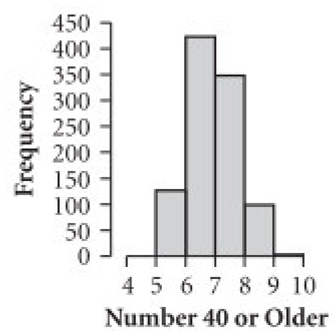

b. 24 of the 50 samples (48%) contained 7 or more people of age 40 or older.

c. It is quite likely to randomly draw ten ages from the fourteen and get 7, 8, or 9 who are aged 40 or older. In fact, nearly half of the samples in the dot plot contained 7 or more people of age 40 or older. Such a result is thus quite likely and provides no evidence for Martin.

NOTE: ***Choosing an Estimator***

A general rule in statistics is that you should not throw away information that might be relevant. Only using whether the person is 40 or older throws away information—the person's exact age. For example, suppose Westvaco had laid off everyone over 60 and kept everyone under 60. That would be age discrimination, but might not show up in an under/over age 40 analysis since all those aged 40–59 were kept.

Here is a specific example of that principle: In the Martin case, there were 10 hourly workers involved in the second round of layoffs; 4 of these workers were under 40. Three were chosen to be laid off; all 3 were 40 or older. To see whether this result could reasonably be due just to chance, consider drawing 3 tickets at random from a box with 10 tickets, 4 of which say "Under 40," and the 6 of which say "40 or older." A natural summary is the number of tickets out of 3 that say "40 or older." Simulating this process a large number of times shows that the chance that all 3 tickets say "40 or older" is about 1 / 6. (In fact, it is possible to count the possible outcomes. There are $\binom{10}{3} = 120$ possible outcomes, of which $\binom{6}{3} = 20$ have all 3 tickets saying "40 or older," so 1 / 6 is the exact probability.)

This probability is large enough that we cannot reasonably rule out the possibility that such data could arise by chance alone. Thus, this conclusion differs from the one conclusion on the actual ages, which was analyzed in an earlier activity. Because the analysis based on the actual ages uses more of the available information, it is more informative and trustworthy. The analysis based on the actual ages (the mean) uses more of the available information. On the other hand, it might be easier to convince a judge or jury by using 40 or over because that is the federal protected class.

E10. a. Write the ages of the 10 hourly workers on identical cards, mix the cards, and draw four at random. Then find the average age of those four workers. This chance model reflects a situation where chance alone decides who is laid off. Repeat this process many times, displaying the average ages on a dot plot. Decide whether it is reasonably likely to get an average age of 57.25 or greater by chance alone.

Alternatively, a calculator can be used to do the simulation in a variety of ways. For example, the ages of the workers can be numbered 1 through 10, and a random integer generator can be used to select the four numbers. On the TI-84, randInt(1,10,4) will generate a list of four integers from 1 to 10, but since it allows repeats, lists in which a number appears more than once will need to be ignored and another list generated until there is no repeat. Then students can find the average age of the workers who correspond to the four random integers.

b. Simulation results will vary.

c. From this simulation an average age of 57.25 or more occurred three times, or 1.5% of the time. In other words, getting an average age at least this high would happen less than 2 times out of every 100 if the selections were done randomly. This is strong evidence in Martin's favor. (The theoretical probability of getting an average age of 57.25 or more is only $4/210 \approx 0.019$. That is, getting an average age this high would happen less than 2 times out of every 100, if selections were done randomly.)
E11. C. 45 of the 200 samples, or 22.5%, had an average weight of 15 ounces or less. Since this could happen a little more than once out of every five times if the loaves are selected randomly, she should probably give her baker the benefit of the doubt.
E12. B. "If it is July 4,it is very unlikely to be snowing in Kansas. Therefore, this probably isn't July 4." The hypothesis is that it is July 4, so we begin by assuming that. We then examine the data (it's snowing in Kansas), and because it's unlikely to be snowing on July 4 in Kansas, we reject the hypothesis that it is July 4.
E13. a. Only four pair give an average age of 59.5 or older: 25 33 35 38 48 **55** 55 55 56 **64** 25 33 35 38 48 55 **55** 55 56 **64** 25 33 35 38 48 55 55 **55** 56 **64** 25 33 35 38 48 55 55 55 **56 64** **b.** Thus, the probability of getting an average of 59.5 or more is $4/45 \approx 0.09$. **c.** The evidence of age bias is somewhat weak.
E14. a. There are six sets of three that give an average of 58 or more: 25 33 35 38 48 **55 55** 55 56 **64** 25 33 35 38 48 **55** 55 **55** 56 **64** 25 33 35 38 48 **55** 55 55 **56 64** 25 33 35 38 48 55 **55 55** 56 **64** 25 33 35 38 48 55 **55** 55 **56 64** 25 33 35 38 48 55 55 **55 56 64** **b.** The probability of getting an average age of 58 or more by choosing three ages at random is $6/120 = 0.05$. **c.** Answers will vary. The probabilities differ because the result from the display was from a simulation and so is an estimate of the exact answer computed in part b. Unless the number of repetitions is quite large, there can be a considerable difference.

E15. a.

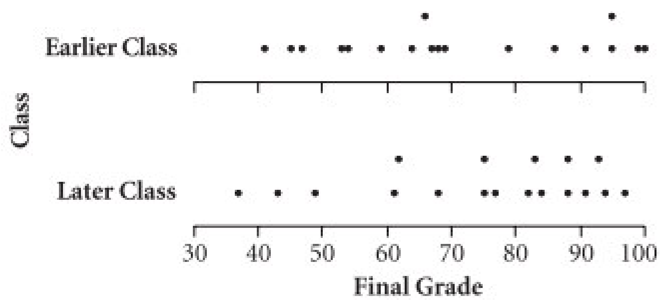

b. There does not appear to be a great difference between the sections. The mean score of students in the earlier class (71) is only slightly lower than that of students in the later class (74.8), and the range and variability of the scores are very similar, so it appears that the difference between class averages could reasonably be attributed to chance. The teacher need not look for an explanation.

E16. a.

		Pass		
		Fail	Pass	Total
Class	Earlier Class	6	12	18
	Later Class	3	15	18
	Total	9	27	36

b. 12 out of 18 or about 67% of the students in the earlier class passed. 12 out of 27 or about 44% of the students who passed were in the earlier class. 27 out of 36 or 75% of the students passed overall.

c. To decide if a disproportionate number of passing students are enrolled in the earlier class, you should compute the proportion of students in the earlier class who passed and the proportion of students in the later class who passed. 12/18, or 66.7%, of the students in the earlier class passed. 15/18, or 83.3%, of the students in the later class passed. Answers will vary depending on whether the students believe such a difference is likely to occur due to chance alone. (However, with only 18 students in each class, it turns out that a difference of 16.7% would not be a rare occurrence.)

E17. B is the best choice of those given, but the question remains whether such a difference is a reasonably likely outcome due to chance alone.

E18. The cases are the eight individual planets. The variables are the radius of the planet, the number of moons for the planet, and the average surface temperature of the planet.

E19. a. ii.. Martin's case would benefit from showing that it's pretty unusual to get 12 or more older workers if workers had been selected completely at random for layoff. That would mean that Westvaco could be asked for an explanation of these unusual results.

b-c. Results will vary.

d. This is 13% of the trials. A random selection process would generate a result at least this extreme about once every 7 or 8 trials. This is not a small enough percentage to cast serious doubt on a claim that those laid off were chosen by chance.

E20. a. The mean of the workers actually laid off is 49. There are 14 dots representing average ages of 49 and greater. That's 14 out of 600, or 0.023. The estimate for the probability that the average age of the five workers laid off is 49 or greater is about 2.3%.

b. Because the probability of getting a mean age of 49 or higher due to chance alone is small, Eastbanko would have some explaining to do.

E21. a. ${}_{14}C_{10} = \binom{14}{10} = 1001$ ways.

b. They could have laid off 5, 6, 7, 8, or 9 older workers.

c. $\binom{9}{7}\cdot\binom{5}{3} = 36\cdot 10 = 360$, $\binom{9}{8}\cdot\binom{5}{2} = 9\cdot 10 = 90$, $\binom{9}{9}\cdot\binom{5}{1} = 1\cdot 5 = 5$

d. $\frac{360+90+5}{1001} \approx .4545$. There is a 45.45% chance of getting seven or more workers age 40 and over if 10 workers are selected randomly from the 14 hourly workers.

E22. There are still 120 possible subsets of three, but there are many more ways to get an average as large as the observed average of 55, so the probability is substantially higher than in the actual case, and therefore the evidence of age bias is much weaker. (There are in fact 10 possible ways to get an average age of 55, so the exact probability is 10/120 ≈ 0.083 or about 8.3%.)

E23. a. 7/1000 = 0.007

b. Dear Society Members,

As you know, I've been suspicious ever since the two winners of last week's drawing were Filbert's cousins. To see how unlikely such an event would be, I made my own 50 raffle tickets, just like those in last week's drawing, and marked four of them to represent the four that belonged to Filbert's cousins. I then mixed the tickets up well and drew two of them. I did this 1,000 times. I got two tickets representing those belonging to Filbert's cousins only 7 times out of the 1,000 drawings. This gives me an estimated probability that both tickets would belong to Filbert's cousins of 0.007, which is a really small chance. This leaves us with the following possible explanations:

- The tickets weren't well mixed. Does anyone know if, for example, Filbert's cousins bought the last four tickets and these ended up on top? Or maybe they bought the first four tickets, which were on the bottom?

- A really unlikely event occurred.
- I don't like to use the word "cheating," but that is another possible explanation.

c. Even the most unlikely event is *possible*, so, yes, it's possible. In fact, the event that both winning tickets were held by Newman's cousins happened seven times in the simulation.

d. There are $\binom{4}{2} = 6$ ways to draw two winning tickets both held by Newman's cousins. There are $\binom{50}{2}$ or 1,225 ways to draw two tickets from the 50. So the probability is $\frac{6}{1,225} \approx 0.005$.

Concept Review Solutions

C1. B. The plot does not tell how many part-time and full-time students do not receive financial aid. There may be far more part-time students who do not receive financial aid than full-time students who do not receive financial aid. In this case, choice B would not be correct.

C2. E. Only choice E compares the proportion of hourly workers laid off with the proportion of salaried workers laid off.

C3. B. The total number of applicants admitted was 25, and of those, the number of females was 8. So, the proportion of admitted applicants who were female was 8 out of 25, or 32%. (This is much lower than the proportion of admitted applicants who were male, 68%, but more males applied for admission than females.)

C4. A. The actual proportion admitted, 32%, is near the middle of this distribution (which is around 29%), so women are admitted at a rate slightly higher than that of men, but the difference easily could be accounted for by chance alone.

C5. **a.** The data for compound B include more of the larger values, but there is a lot of overlap in the values. Technically, it can be argued both ways; on the one hand, compound B can be viewed as being better because it has more of the higher values, while it is reasonable to also say that there is no obvious winner.

b.

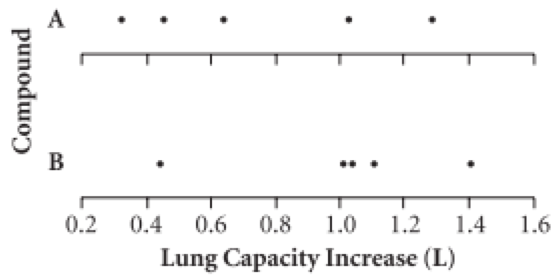

The points for compound B tend to fall more on the right side of the plot, so it appears that compound B gives larger measurements than compound A.

c. Compound A mean is 0.746; compound B mean is 1.002, for a difference of 0.256. Because the compound B mean exceeds the compound A mean, it might appear as though this is sufficient evidence to claim that compound B is better at opening the lungs than compound A. However, the difference is not very large and is not statistically significant, as illustrated in C6.

C6. a. Differences will vary depending on the samples chosen.

b. Mean(compound B) – Mean(compound A) = 0.256. This observed difference was exceeded 11 times out of the 50, or 0.22 of the time.

c. This high proportion points to a conclusion that the difference of 0.256 could well be the result of chance alone.