NOT RECOMMENDED

Fig. 2.11 Energy Data

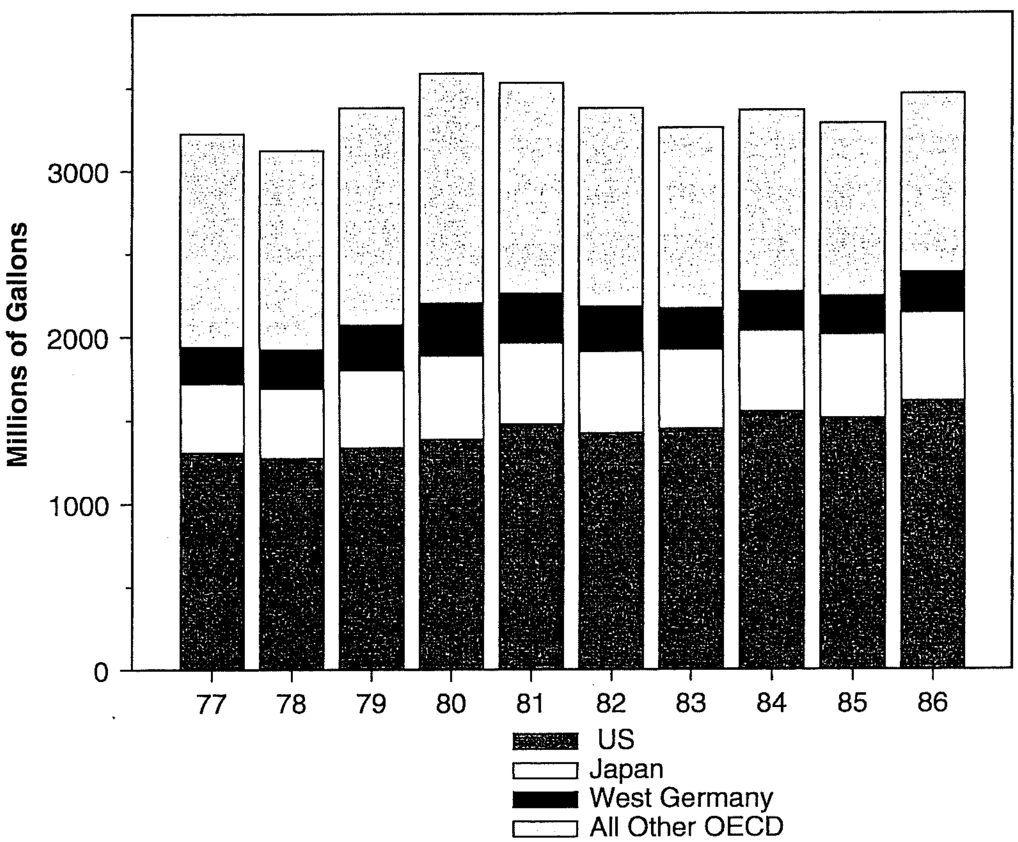

2.3 BAR CHARTS: STACKED AND GROUPED

Another common graph form is a *stacked bar chart*. Figure 2.11 shows petroleum stocks from 1977 to 1986 in millions of gallons for the United States, Japan, West Germany, and all other countries of the Organisation for Economic Co-operation and Development (U.S. Dept. Energy, 1986). You probably read the values for the United States and the totals quite accurately. Study the chart and see what you can discover about the other countries.

Fig. 2.12 Energy Data: All Other OECD

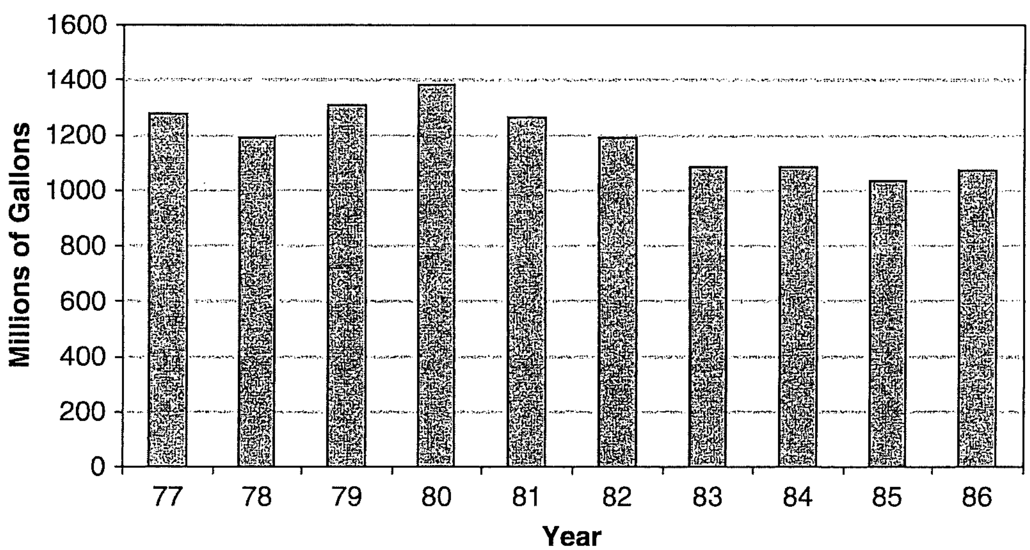

Did you notice in Figure 2.11 that the values for "all other OECD" generally tend to decrease over time? You probably didn't. As we shall see in Chapter 3, it is very difficult to judge lengths that do not have a common baseline.

Fig. 2.13 Energy Data: Grouped Bar Chart

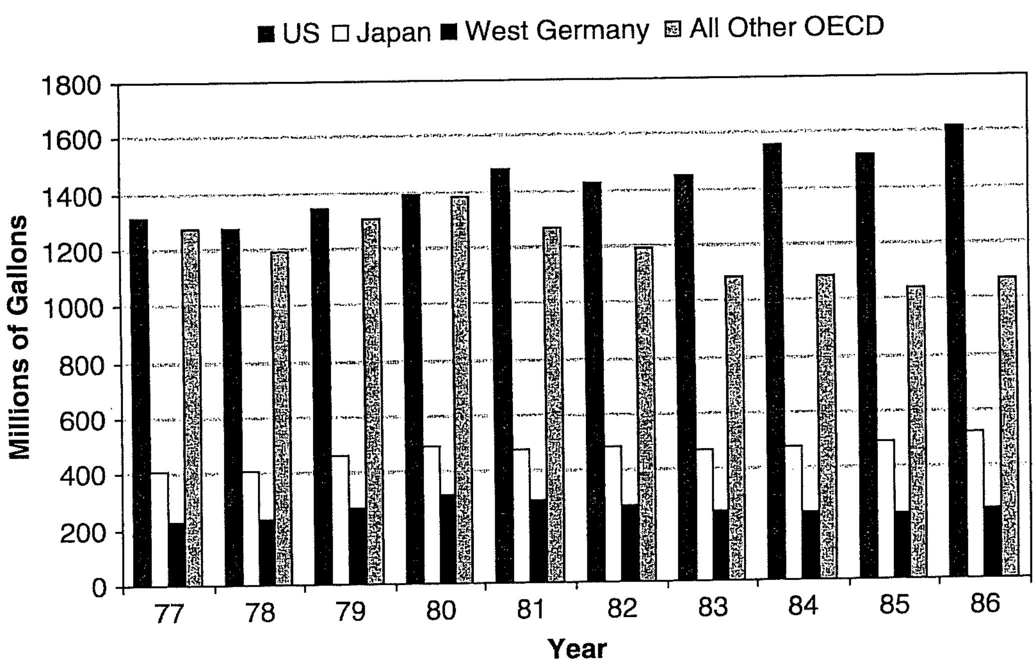

The bars in *grouped bar charts* do have a common baseline. Iowever, a grouped bar chart such as Figure 2.13 becomes lifficult to read with even a few groups. It is difficult to ollow the trend for a given group such as Japan because the lata for the other groups fall between the consecutive values or Japan. Reordering the shadings helps to make the groups listinguishable. The pattern of the "all other OECD" group is ertainly clearer than in the stacked bar chart. However, trellis lisplays, which are discussed in Chapter 5, are far clearer than s a grouped bar chart.

Fig. 2.14 Playfair's Balance-of-Trade Data

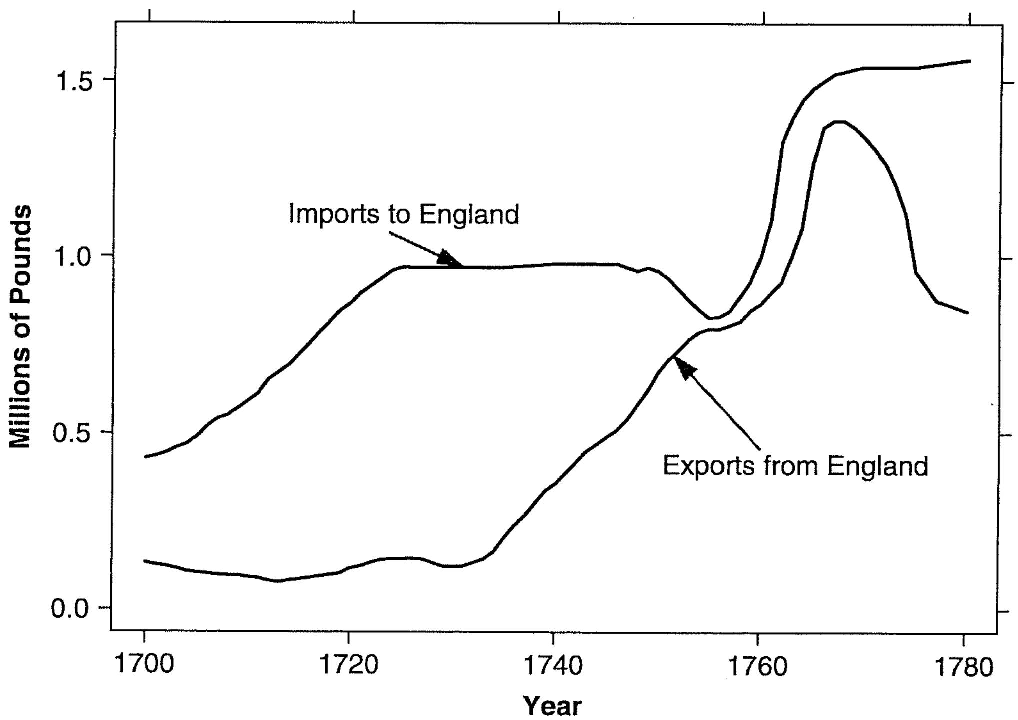

2.4 DIFFERENCE BETWEEN CURVES

Most of the graph forms that have been used until recently were introduced by William Playfair in the late eighteenth and early nineteenth centuries. Figure 2.14 uses Playfair's data (Playfair, 1786) to show exports from England and imports to England in trade with the East Indies. We're interested in the balance of trade, which is the difference between exports and imports. We see that the difference is about 0.4 minus 0.2 or 0.2 in 1700, and then it increases for awhile. I'd like you to continue sketching the difference.

Fig. 2.15 Playfair's Balance-of-Trade Data: Imports Minus Exports

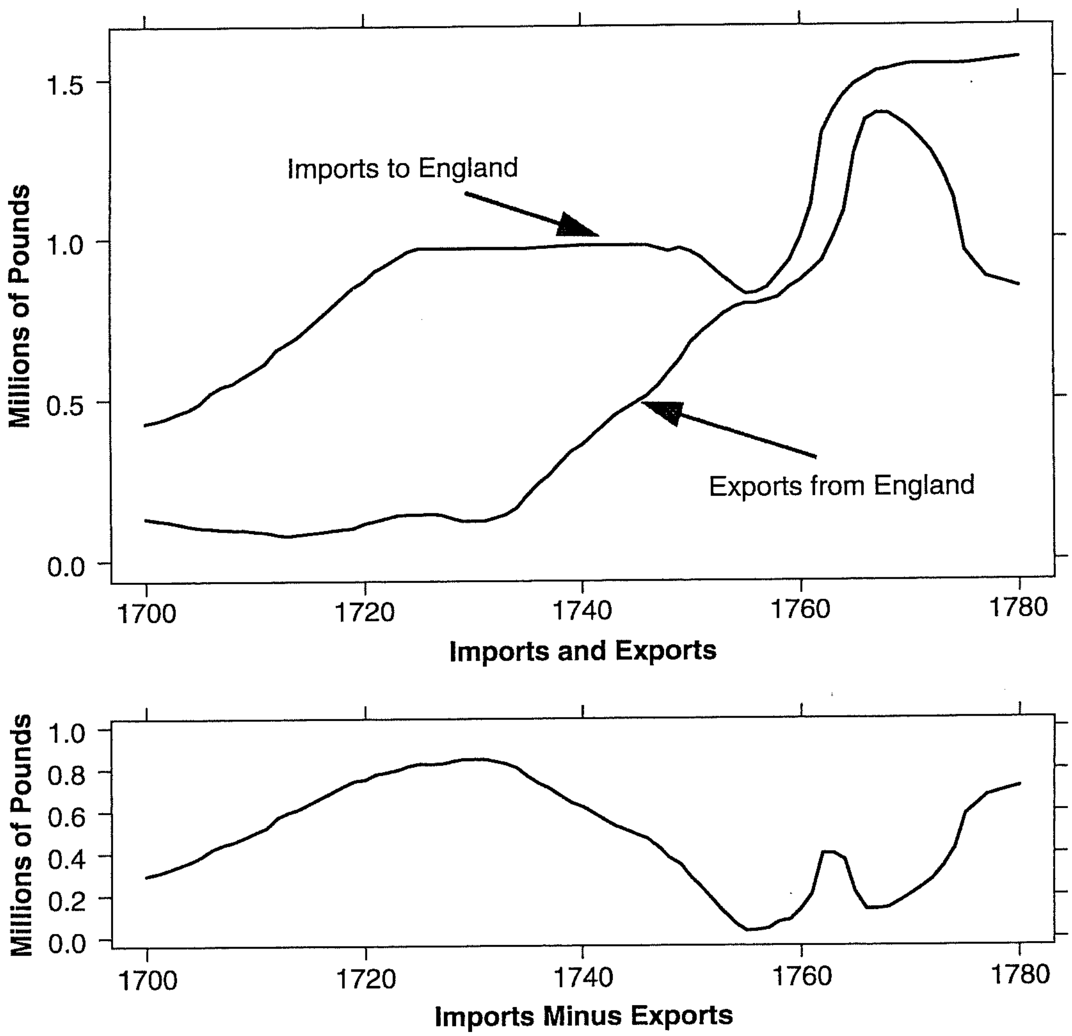

Did you notice the hump after 1760? We miss it because our eyes look at the closest point rather than the vertical distance.

It is important to remember to plot the variable of interest. If interested in the balance of trade, plot the difference rather than just the imports and exports. If we have *before* and *after* data and are interested in improvement, we plot the improvement, not just the *before* and *after* data.

Fig. 2.16 Difference between Curves

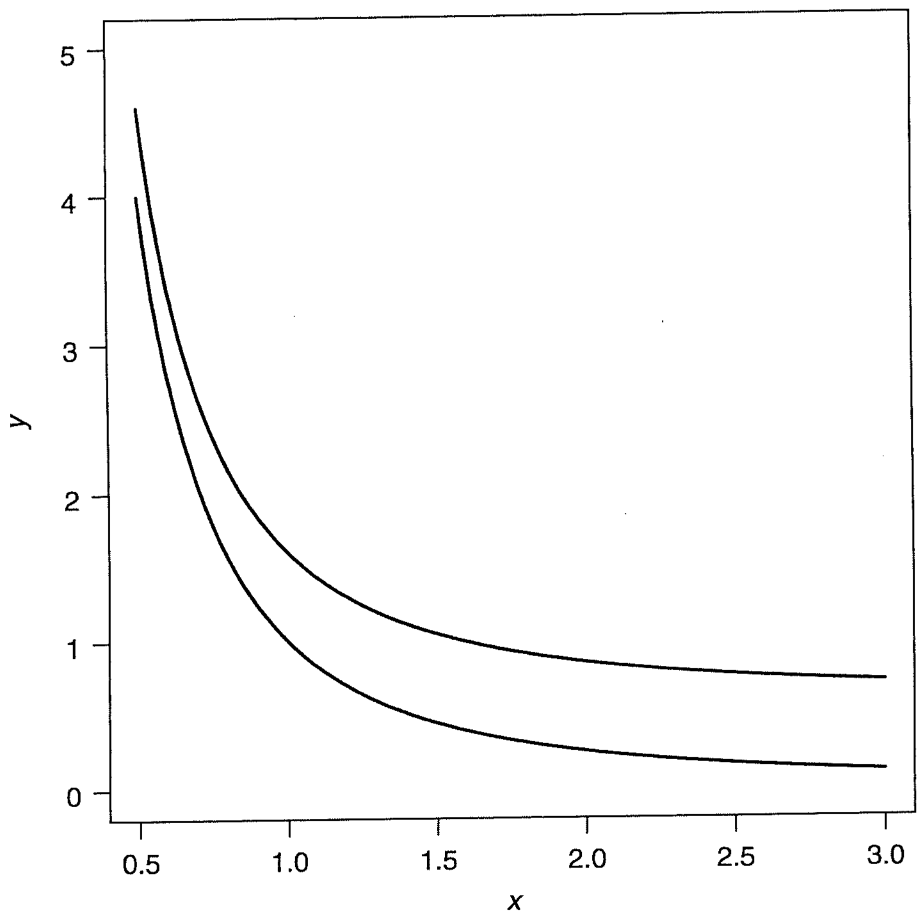

Look at the two curves in Figure 2.16. For what values of x are they closest together and also, farthest apart?

Fig. 2.17 Ownership of Government Securities

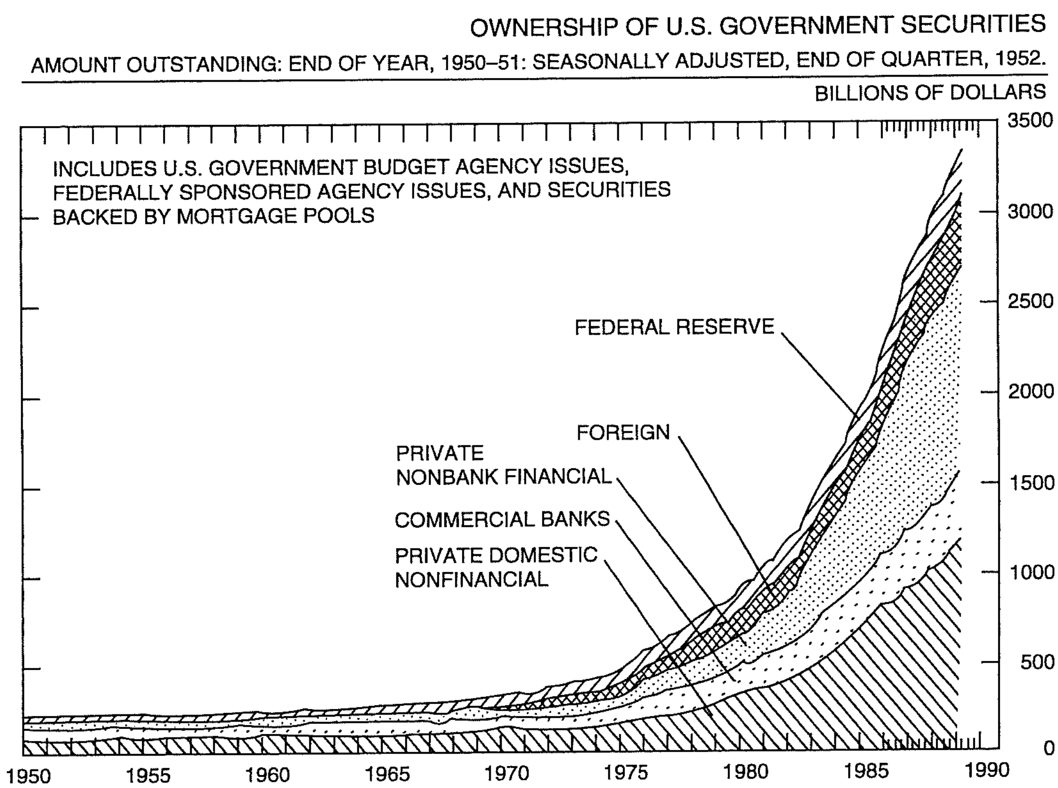

The last question appeared to be easy, but actually the wo curves differ by a constant amount. The curves plotted re $y_1 = 1/x^2$ and $y_2 = y_1 + 0.6$, so that one curve is always xactly 0.6 higher than the other.

The last few charts have taught me not to trust my judgment vhen viewing charts such as the ownership of U.S. govern-nent securities [Board of Governors of the Federal Reserve ystem (U.S.), 1989]. If interested in how a specific group, ay commercial banks, changed over time, I would perform he subtraction and plot that group over time as we did with he exports and imports in the Playfair example.

3 Human Perception and Our Ability to Decode Graphs

In Chapter 2 we saw that some common graphs do not communicate numerical information effectively. We also discovered other graphs that clearly communicate the information and the patterns of the data. In this chapter we examine briefly the tasks required to decode the information in a graph. Cleveland and McGill (1984) ran experiments to determine which of these tasks we do most accurately. This knowledge helps us to understand why some graphs work and others don't. We first list in alphabetical order 10 judgments we make when decoding quantitative information from graphs, describe each briefly, then order them by our ability to perform them accurately. In some cases the descriptions of two tasks with similar properties (e.g., area and volume) appear together.

Creating More Effective Graphs, by Naomi B. Robbins

3.1 ELEMENTARY GRAPHICAL PERCEPTION TASKS

Angle

Area

Color hue

Color saturation

Density

Length

Position along a common scale

Position along identical, nonaligned scales

Slope

Volume

Fig. 3.1 Angle Judgments

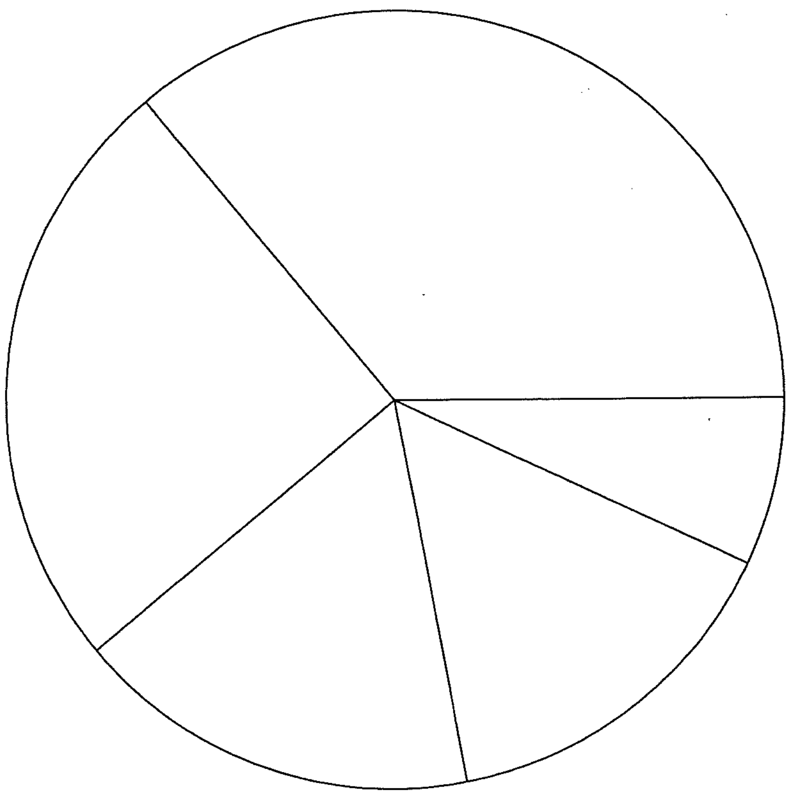

We make angle judgments when we read a pie chart, but we on't judge angles very well. These judgments are biased[1]; ve underestimate *acute angles* (angles less than 90°) and verestimate *obtuse angles* (angles greater than 90°). Also, ngles with *horizontal bisectors* (when the line dividing the ngle in two is horizontal) appear larger than angles with ertical bisectors.

Biased means that we consistently under- or overestimate the true value.

Fig. 3.2 Area and Volume Judgments

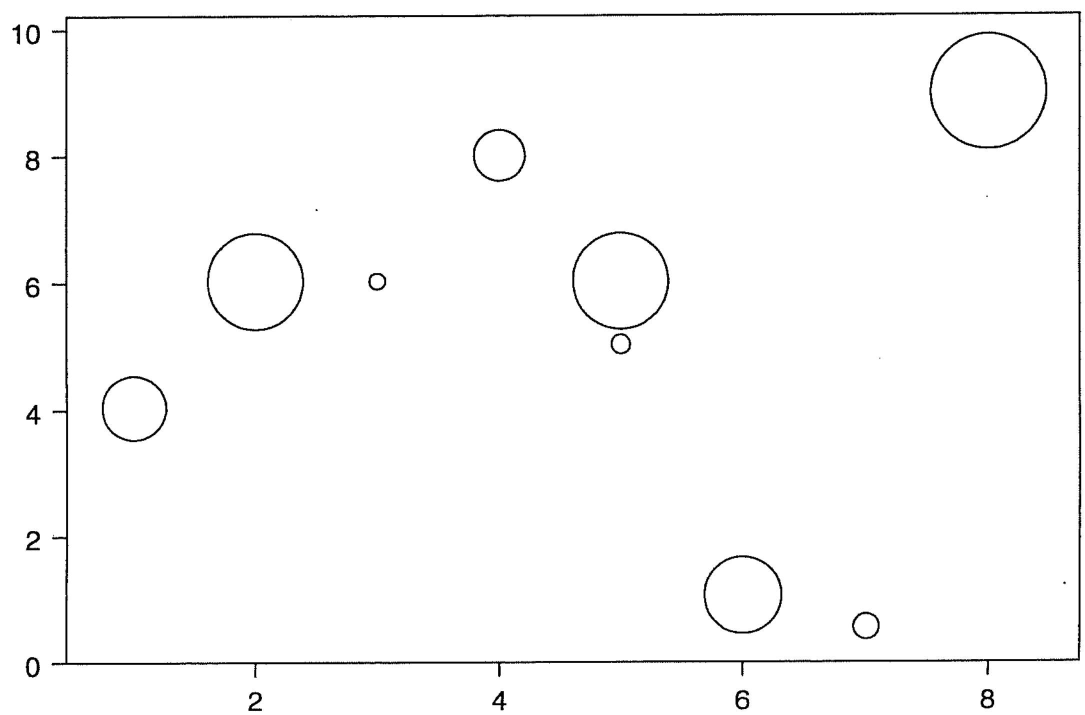

The circles on Figure 3.2 show three variables by the horizontal position of the center of the circle, the vertical position, and the area of the circle. For example, the horizontal axis could be the months that you have held a security, the vertical could be the price you paid in thousands of dollars, and the areas of the circles could represent your gain. These charts are often called *bubble plots.*

Area judgments are also biased. They are much less accurate than length and position judgments. Volume judgments are even more biased. Stevens (1975) presents the following law: Let x be the magnitude of an attribute of an object, such as its length or area. According to *Stevens' law*, the perceived scale is proportional to x^{β}, where β has been determined by experimentation to range generally from 0.9 to 1.1 for length, 0.6 to 0.9 for area, and 0.5 to 0.8 for volume.

When $\beta = 1$, $x^{\beta} = x$, and when $\beta < 1$, $x^{\beta} < x$. Since the β for area is less than 1, we perceive areas to be smaller than they really are. This bias is more pronounced with volumes. The range of beta for lengths includes 1, so we perceive lengths more accurately than areas or volumes.

Fig. 3.3 Color Hue, Saturation, and Density

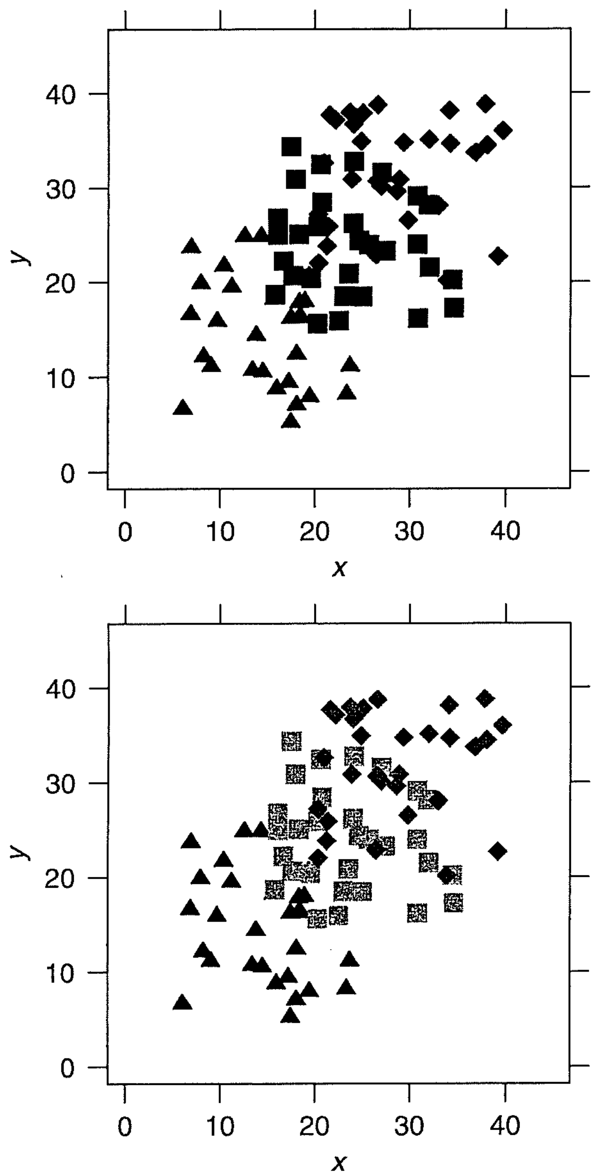

Color coding (*hue*[2]) is very effective for distinguishing data rom various groups. For example, suppose that the triangles 1 Figure 3.3 represent data for England; the squares, France; nd the diamonds, Italy. It is difficult visually to separate the hree groups in the top plot. Varying density or saturation an also be used to distinguish groups of data. Using different ensities in the bottom plot makes this task easier, even hough we are limited to shades of gray. If we use just one ue, we can rank by saturation or density to show levels of a uantitative variable. Weather maps often vary shades of red nd blue to show temperature. But color coding with different ues does not work well for showing numerical information ince we don't perceive an ordering to red, green, blue, and ther hues.

Hue is the technical term for what we call *color* (red, yellow, blue, etc.). *aturation* refers to the intensity of the color. As saturation increases, the olor becomes purer; as saturation decreases, the color becomes more ray. *Density* refers to the shading or amount of black.

Fig. 3.4 Length Judgments

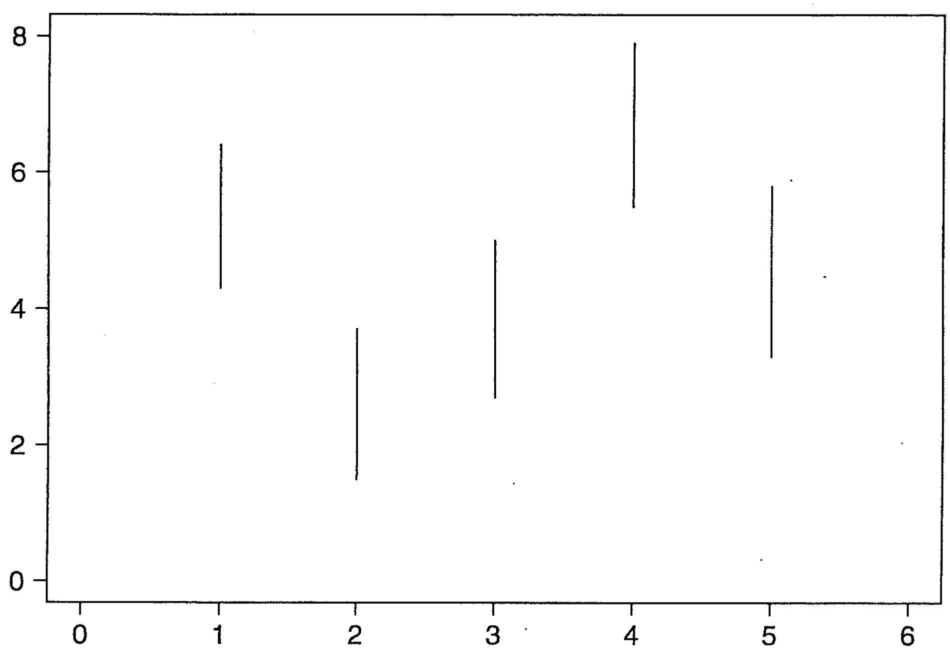

We all know what length means, but can you order the engths of the five line segments in Figure 3.4? We learned rom Stevens' law that we judge lengths more accurately than ıreas or volumes, but judging lengths is still not easy.

To detect a difference in length between two line seg- nents, we need a fixed percentage increase in the length. For example, if one line is 99 inches and the other is 100 nches, it will be much more difficult to distinguish this 1-inch lifference than if one line is 1 inch and the other is 2 inches, even though the absolute differences are the same. By the way, line 1 is 2.1 units, line 2 is 2.2, line 3 is 2.3, line 4 is 2.4, ınd line 5 is 2.5 units.

Fig. 3.5 Position along a Common Scale

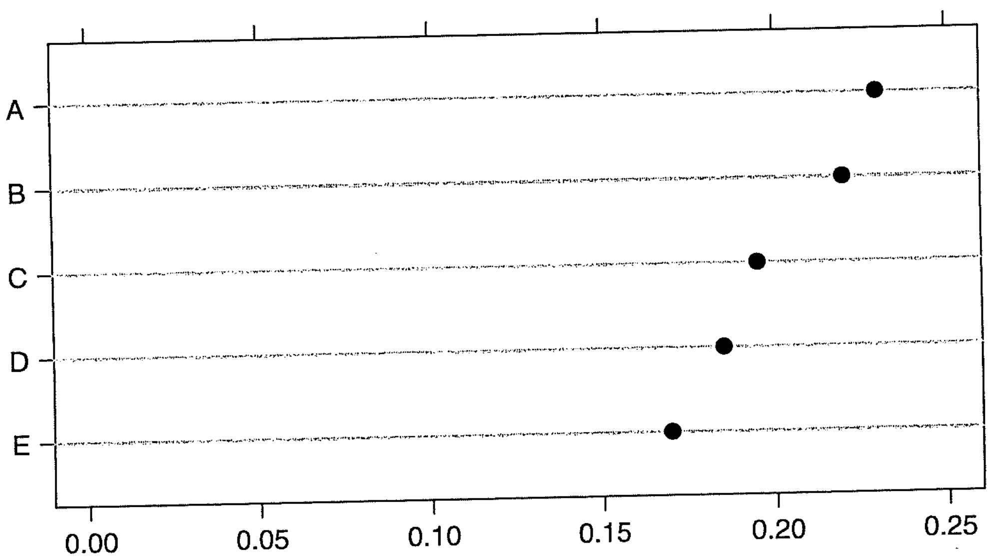

The dot plot shown in Figure 3.5 allows us to make judgments of positions along the common horizontal scale. Experiments by Cleveland and McGill (1984) have shown that this is the most accurate of the elementary graphical tasks. The dot plot was designed to take advantage of the knowledge gained from these experiments on perception and decoding information from graphs.

Fig. 3.6 Position along Identical, Nonaligned Scales

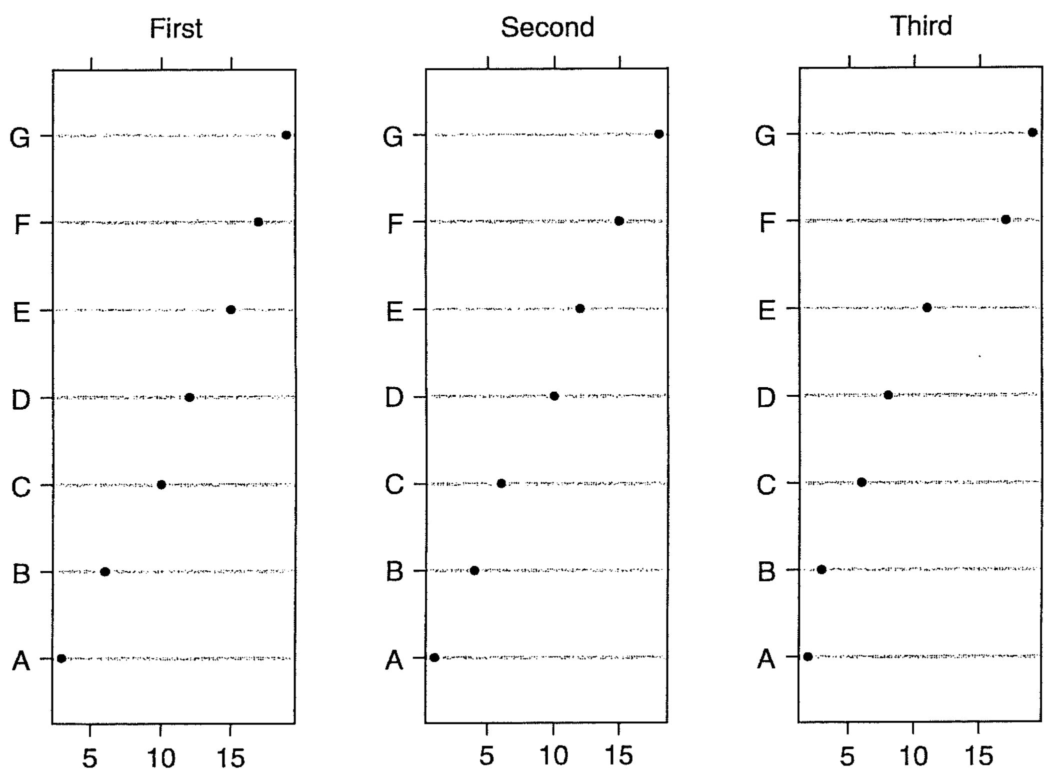

Note that the horizontal scales in Figure 3.6 are the same for the three cases shown. Within the same panel we judge position along a common scale. To compare values on separate panels, we compare positions along identical but nonaligned scales. We make these judgments very accurately.

Multipanel displays are extremely useful when we have more than two variables. Each panel shows two variables for one value of the third variable. For example, if the third variable is countries and we have data for England, France, and Italy, there would be one panel for each country. Multipanel displays are discussed in Chapter 5.

Fig. 3.7 Slope Judgments

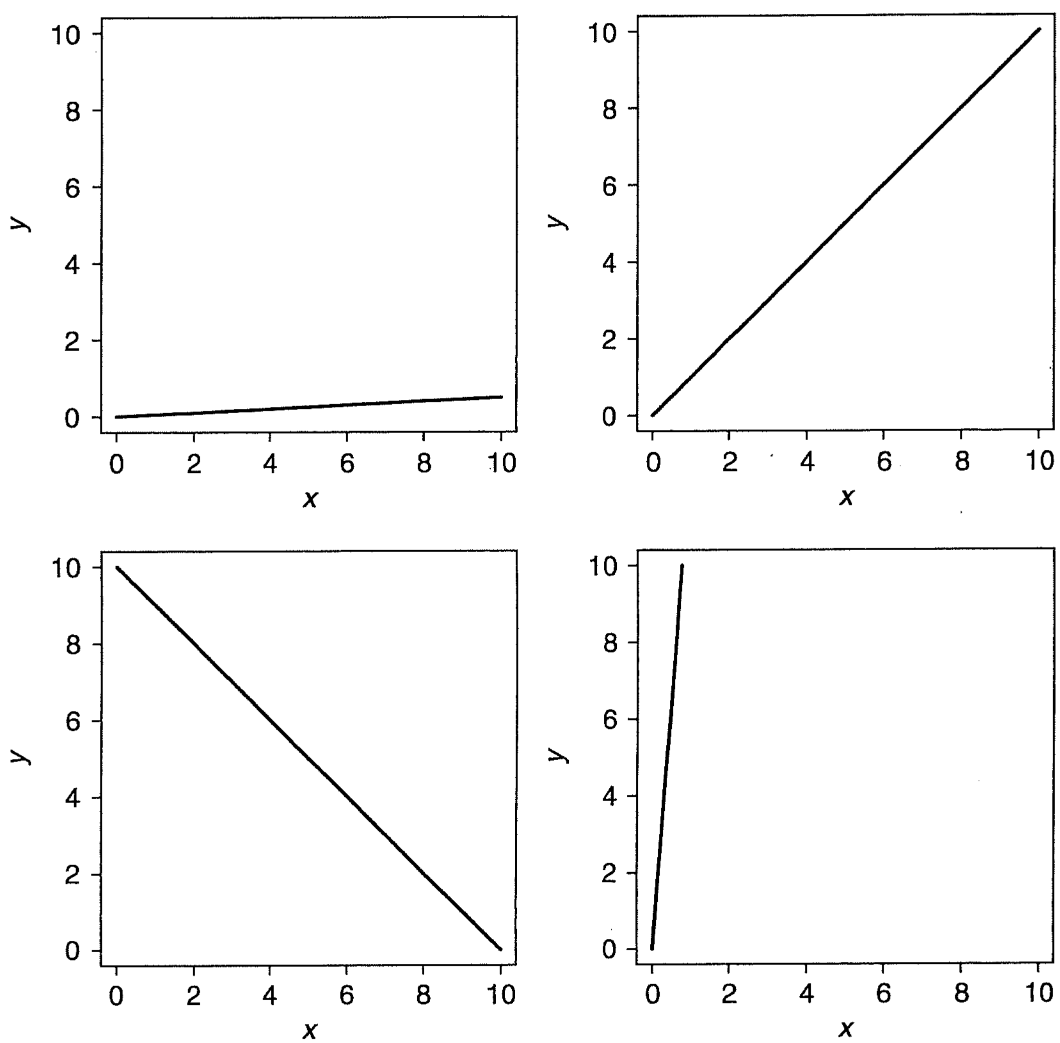

Readers make angle judgments to determine slopes, but we don't judge angles very accurately. The accuracy of judgments of slopes of line segments depends on the angle with the horizontal. Poor accuracy results from angles close to 90°. We judge angles near 45° most accurately. Included in Chapter 7 is a technique called *banking to* 45° that will tell you how to make use of this knowledge.

3.2 ORDERED ELEMENTARY TASKS

Page 47 shows the elementary tasks for decoding quantitative information in alphabetical order. The following list shows the same tasks in order of our ability to perform them accurately:

1. Position along a common scale
2. Position along identical, nonaligned scales
3. Length
4. Angle - slope
5. Area
6. Volume
7. Color hue - color saturation - density

Fig. 3.8 Detection

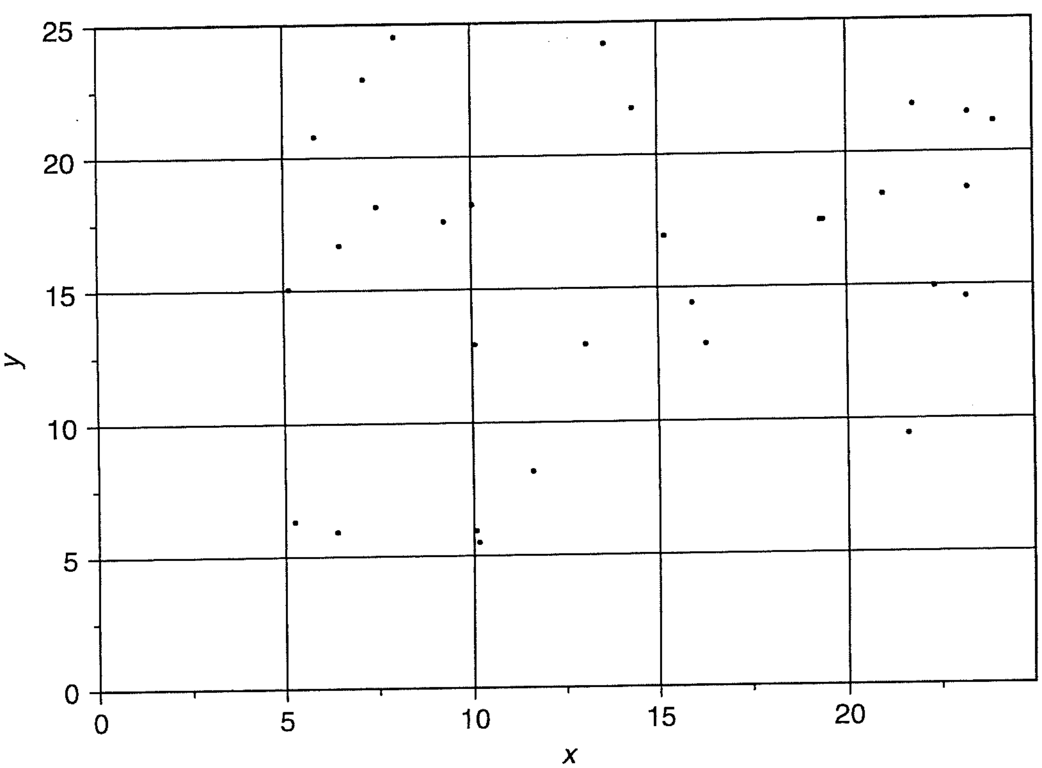

3.3 ROLE OF DISTANCE AND DETECTION

Distance and detection also play a role in our ability to decode information from graphs. The closer together objects are, the easier it is to judge attributes that compare them. As distance between objects increases, accuracy of judgment decreases. It is certainly easier to judge the difference in lengths of two bars if they are next to one another than if they are pages apart.

Before we can perform any of the elementary tasks, we must be able to detect the data. We often cannot do so if data points overlap one another; are hidden in the axes, tick marks, or grid lines; or are too light to see. Figure 3.8 illustrates some of these problems. Additional examples of data hidden by other graphical elements are provided in Chapter 6.

SUMMARY

Creating a more effective graph involves choosing a graphical construction in which the visual decoding uses tasks as high as possible on the ordered list of elementary graphical tasks while balancing this ordering with consideration of distance and detection.

Fig. 4.1 State Areas: Strip Plot

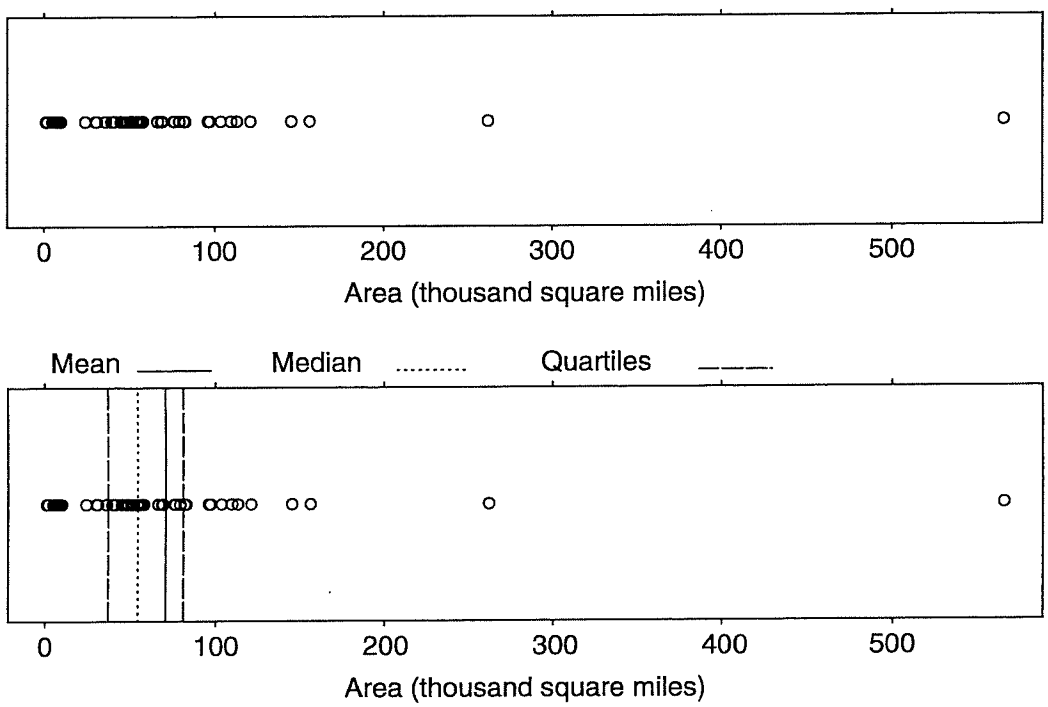

.1 DISTRIBUTION OF ONE VARIABLE

.1.1 Strip Plots

strip plot shows the distribution of data points along a umerical axis; it is also called a *one-dimensional scatterplot*, *ne-dimensional data distribution graph*, or *point graph*. he top strip plot in Figure 4.1 shows the distribution of the reas of the 50 states of the United States. It shows clearly the ınge of the areas and where most of the values lie, but not ıuch more. The bottom strip plot includes some summary :atistics; the mean is shown as a solid line, the median as dotted line, and the 25th and 75th percentiles as dashed nes. Strip plots are sometimes used in the margins of two-imensional displays to show the distribution of each variable :parately.

Fig. 4.2 State Areas: Dot Plot

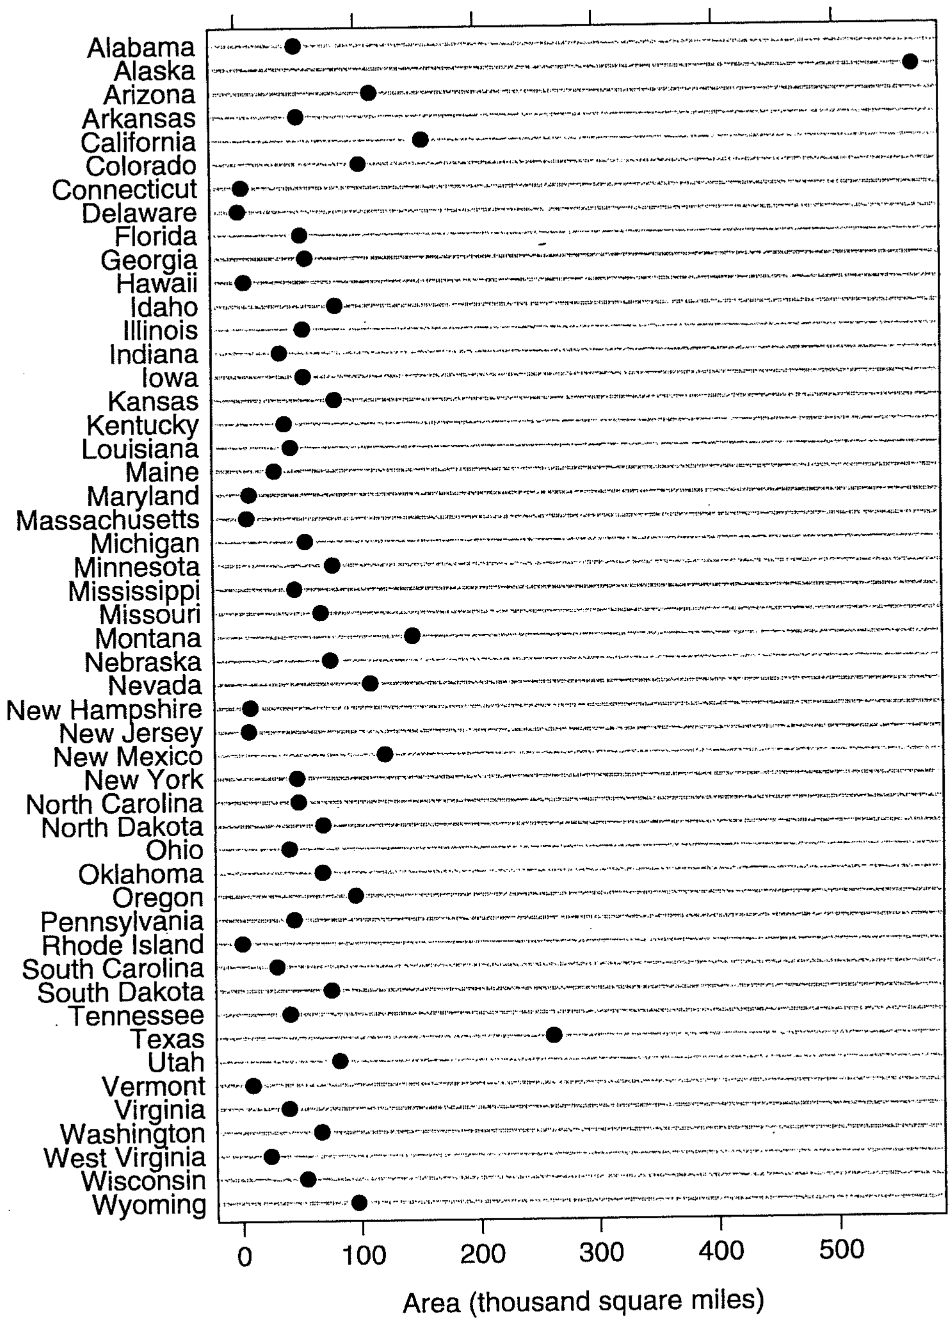

4.1.2 Dot Plots

Figure 4.2 shows the areas of the 50 U.S. states plotted using a dot plot. Dot plots were introduced by Cleveland (1984) after extensive experimentation on human perception and our ability to decode graphical information. Since the judgments the reader makes when decoding the information are based on position along the common horizontal scale, these plots display data effectively. Since it would be very difficult to fit the names of the states on the horizontal axis, dot plots place them on the vertical axis, and the quantitative variable, area in thousands of square miles, on the horizontal axis. Notice how much more informative the dot plot is than the strip plot. Although the data appear clearly, we can still improve this chart.

Fig. 4.3 State Areas: Ordered by Size

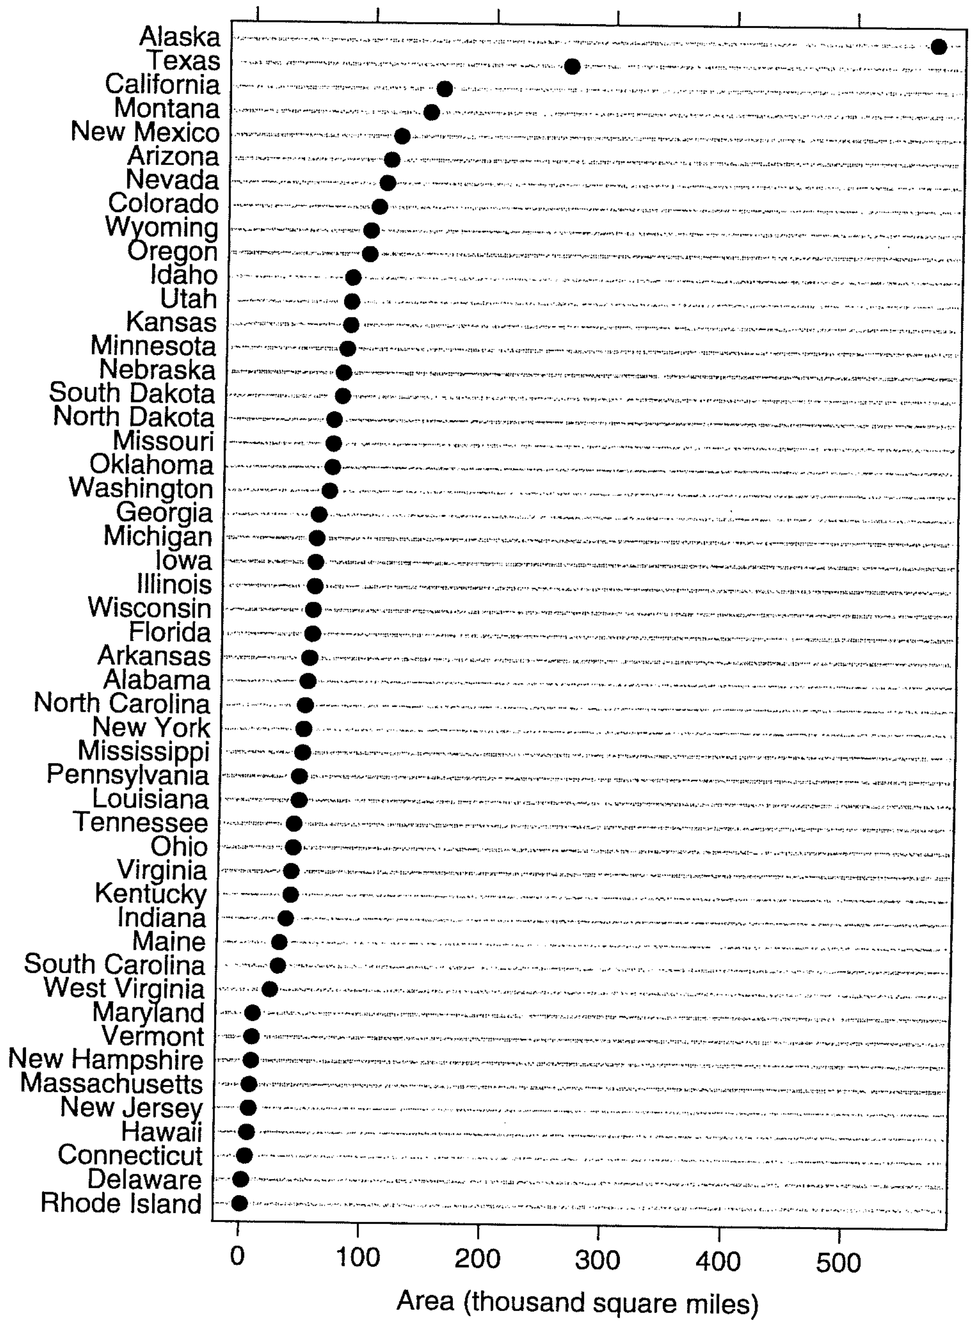

The states appear alphabetically in Figure 4.2. In Figure 4.3 hey are listed in order of size. This presentation is much nore informative. It is much easier to answer such questions ıs: "How many states are smaller than Indiana?" or "What is he median size of a state?" This figure could be improved ;ven more. Notice that Texas and Alaska are so much bigger han the other states that most of the data are on the left side ıf the chart. Figure 4.4 shows how to handle this problem, ıut it is not the best choice for every situation. As in any form ıf communication, we must know our audience and tailor vhat we say to be appropriate for that audience, the readers ıf the chart.

Fig. 4.4 State Areas: Logarithm with Base 2

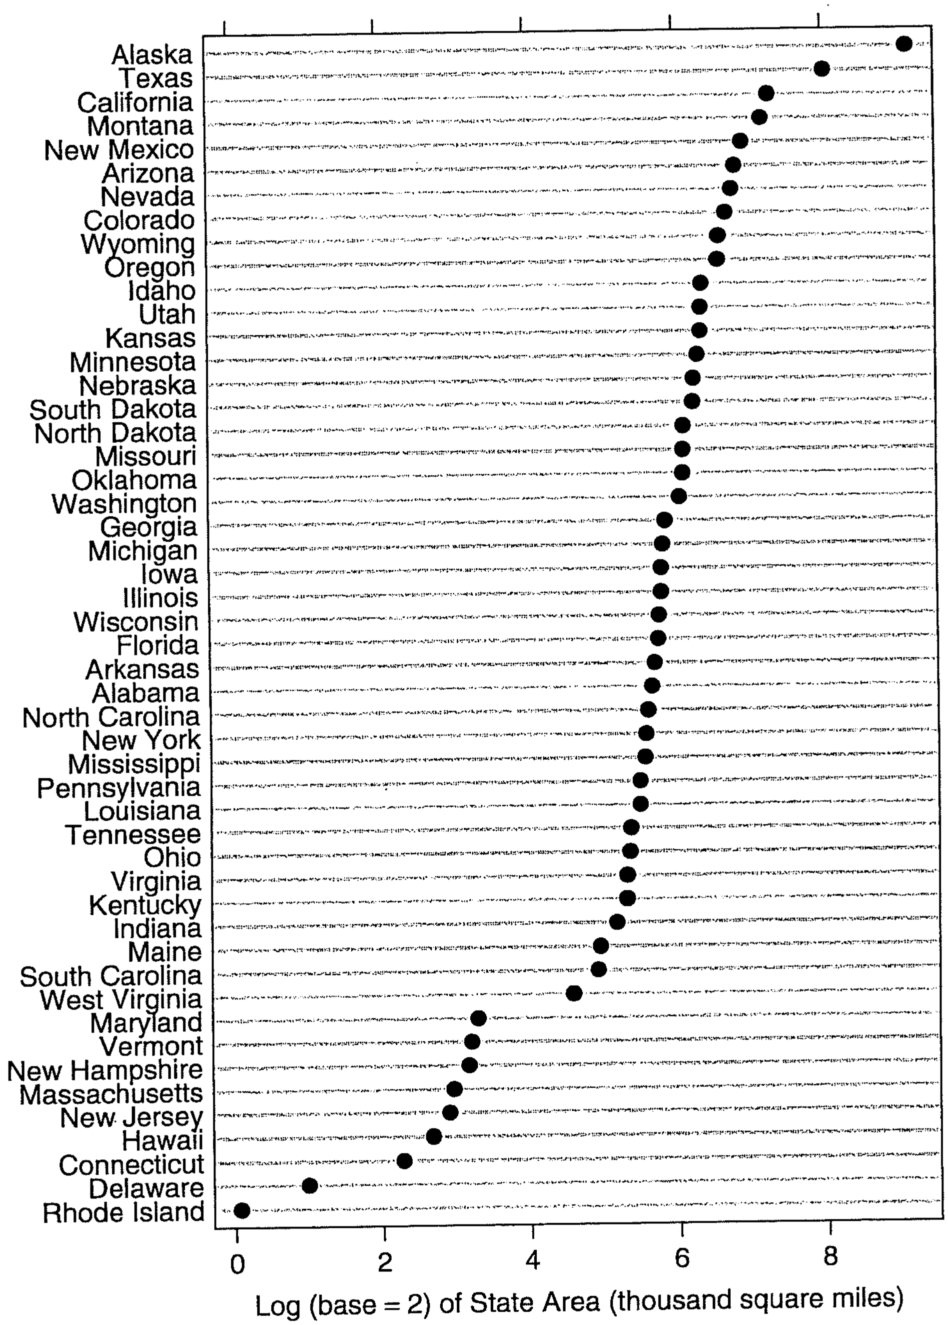

A logarithmic scale makes it possible to plot values with)o wide a range for a linear scale. You have probably seen raphs plotted on a logarithmic scale, with the axis labeled 10, 00, 1000, and so on, in financial reports. Let's review what 'e mean by *logarithms*. If $10^2 = 100$, then $\log_{10}(100) = 2$. $= \log_b x$ means that b is raised to the exponent y in order) get x.

Bases other than 10 are also useful. Ten is useful when the ata range over several orders of magnitude. A base of 2 is seful for plots when we want to spread the data over a smaller inge. Figure 4.4 shows the state areas on a logarithmic scale 'ith a base of 2. This allows us to see details that were not as .ear on the linear scale: for example, the large jump in size om Maryland to West Virginia.

Fig. 4.5 State Areas: Top Axis Labeled with Original Scale

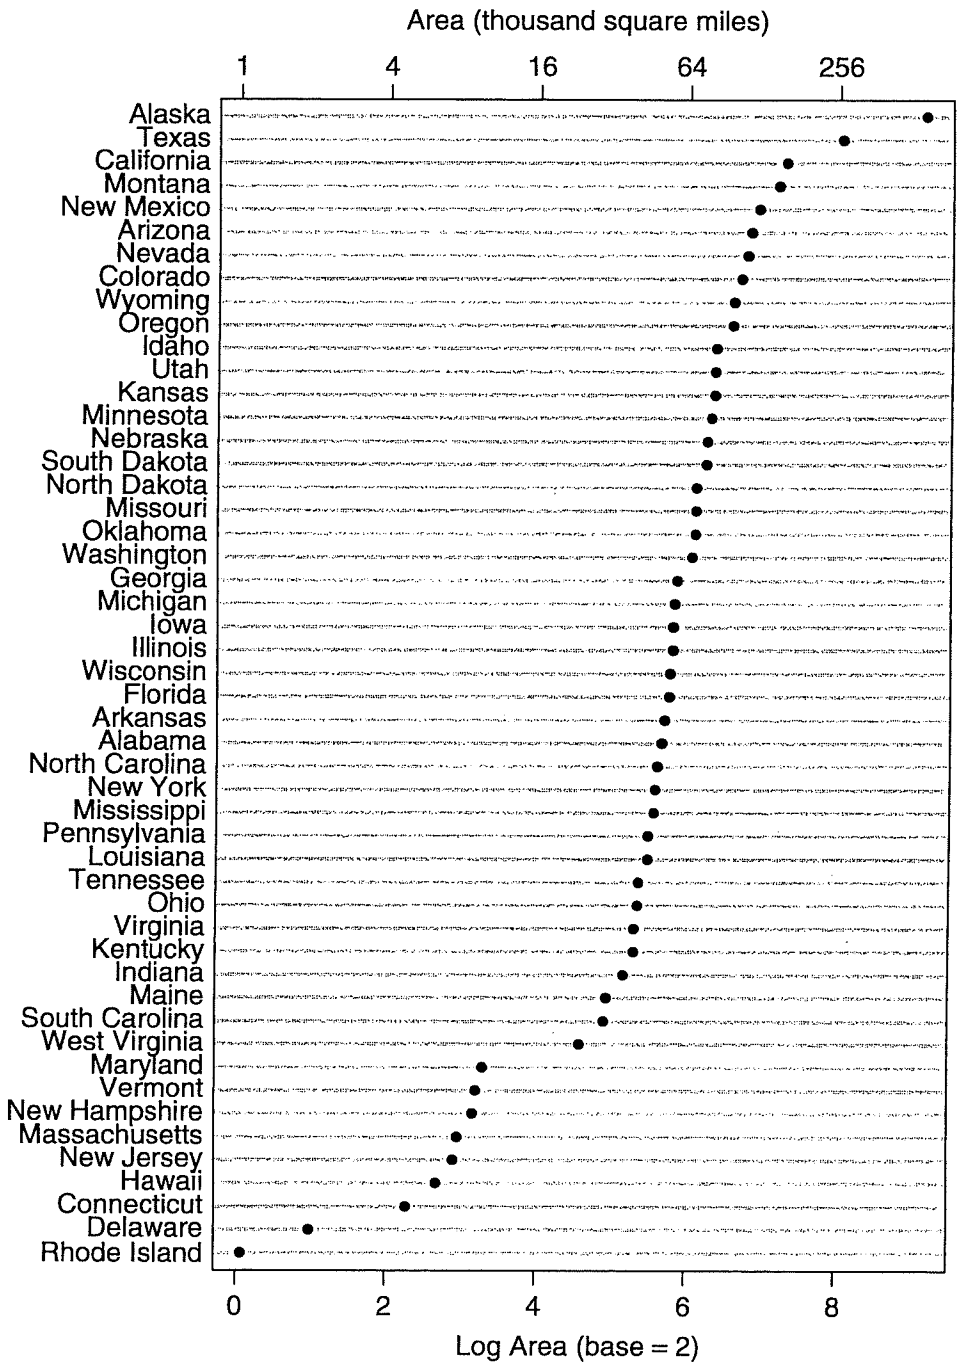

Logarithmic scales are useful for understanding multi-
ılicative factors. From Figure 4.5 we see that the state of
Vashington is approximately $2^6 \times 1000$, or 64,000 square
ıiles, whereas Texas is approximately $2^8 \times 1000$, or 256,000
quare miles. That tells us that Texas is the size of Washington
imes 2^2, or about four times the size of Washington, since
$^8 = 2^6 \times 2^2$.

The horizontal axis is labeled with the logarithmic scale.
: is useful also to label the figure with the original scale, to
ıake it easier to understand. We do that on the top axis of
igure 4.5. Note the connection between the labels on the
ottom and top; 2 raised to the value of the bottom label gives
ıe value of the top label. More information on logarithmic
cales is given in Chapter 7.

Fig. 4.6 Families Exhibition: Histogram

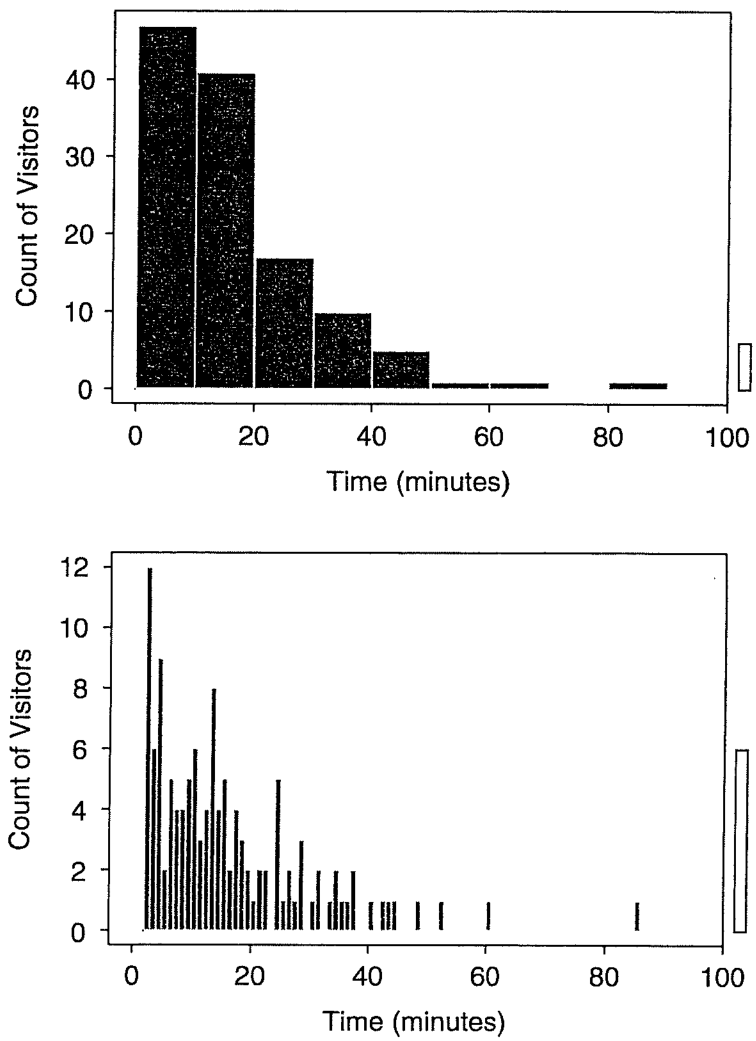

4.1.3 Histograms

Histograms show the distribution of a set of data. Serrell (1998) examined the number of minutes that visitors spent ıt museum exhibitions. To draw a histogram, the data are grouped into bins or intervals. For example, in the top chart n Figure 4.6, which shows the time that visitors spent at an exhibition named "Families," all times up to and including 10 ninutes are in the first bin, times from 10 to 20 minutes are n the second bin, and so on. Then the count of the number n each bin or the percent of the total in each bin is plotted. There is a trade-off between showing detail or showing a etter overall picture. The top figure shows the shape of he distribution more clearly, and the bottom figure shows nore detail. Histograms do a reasonable job of showing the hape of one data set but are not very useful for comparing listributions.

Note that the scales of the two histograms are not the ame. Since there are more bins in the bottom figure, there re fewer visitors in each bin. To help visualize this difference ı scales, the rectangles on the right both have a height of six isitors.

Fig. 4.7 Avoid Misleading Histograms

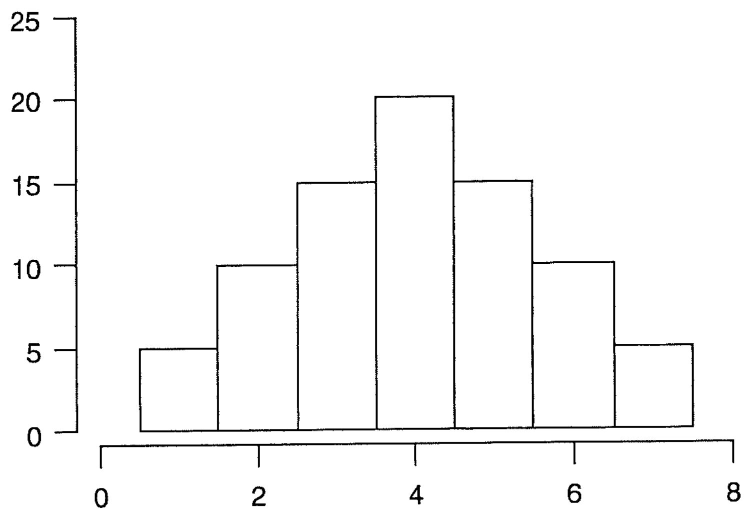

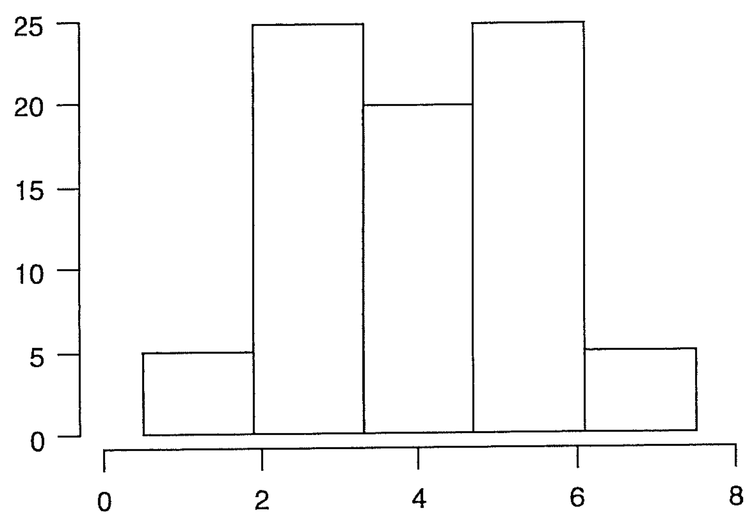

Some software packages for drawing histograms allow the ser to make an unreasonable choice of the number of bars to se. For example, suppose that there are seven possible data alues and that all are integers. The value "1" appears five mes, "2" appears 10 times, and the others appear as shown n the top chart in Figure 4.7.

The x axis of the top chart goes from 0.5 to 7.5 with a range f 7, so that each bar contains exactly one of the integers. [owever, as the bottom chart shows, if the user requests five ars, an unreasonable number for these data, each bar has a ridth of 1.4, which is 7 divided by 5. The first bar goes from .5 to 1.9, including the five occurrences of 1. The second bar oes from 1.9 to 3.3, including the 10 occurrences of 2 and 15 ccurrences of 3. This creates a very misleading impression of ıe shape of the distribution, as shown in the bottom chart.

Fig. 4.8 Judith Leyster Exhibition: Strip Plot

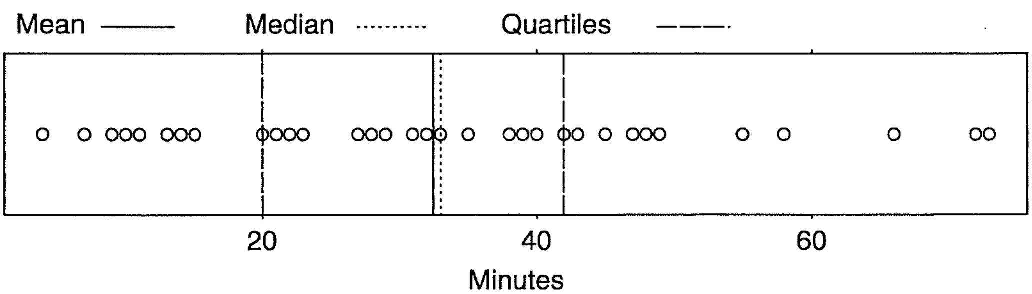

ʟ.1.4 Jittering

ʼhe strip plot in Figure 4.8 shows the number of minutes ʼisitors spent at an exhibition of the artist Judith Leyster. ʼhe data were collected to the nearest minute. There are 49 ›bservations:

4	7	7	9	10	10	11
11	13	14	15	15	20	21
22	22	23	27	27	28	28
29	31	32	33	33	35	38
38	39	40	40	40	40	42
42	42	43	45	47	48	48
49	49	55	58	66	72	73

ʼlotice that there are many repeat values, so that some plotting ;ymbols overlap one another.

Fig. 4.9 Judith Leyster: Jittered Strip Plot

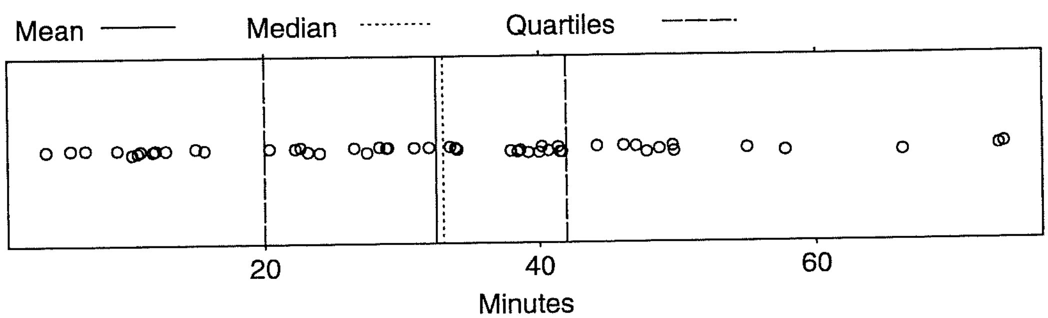

To make the data points distinguishable, we have added ındom noise to the data before plotting Figure 4.9. This echnique, called *jittering*, moves the data points a small, ındom amount from their original positions so that they o longer overlap. Many software packages allow you to tter your data. If yours does not, you can generate random umbers with a small variance or spread to achieve this effect.

There are a number of other solutions to this problem of verlapping data points, which appear on page 165 and in leveland (1994).

Fig. 4.20 Carbon Dioxide Data: Month Plot

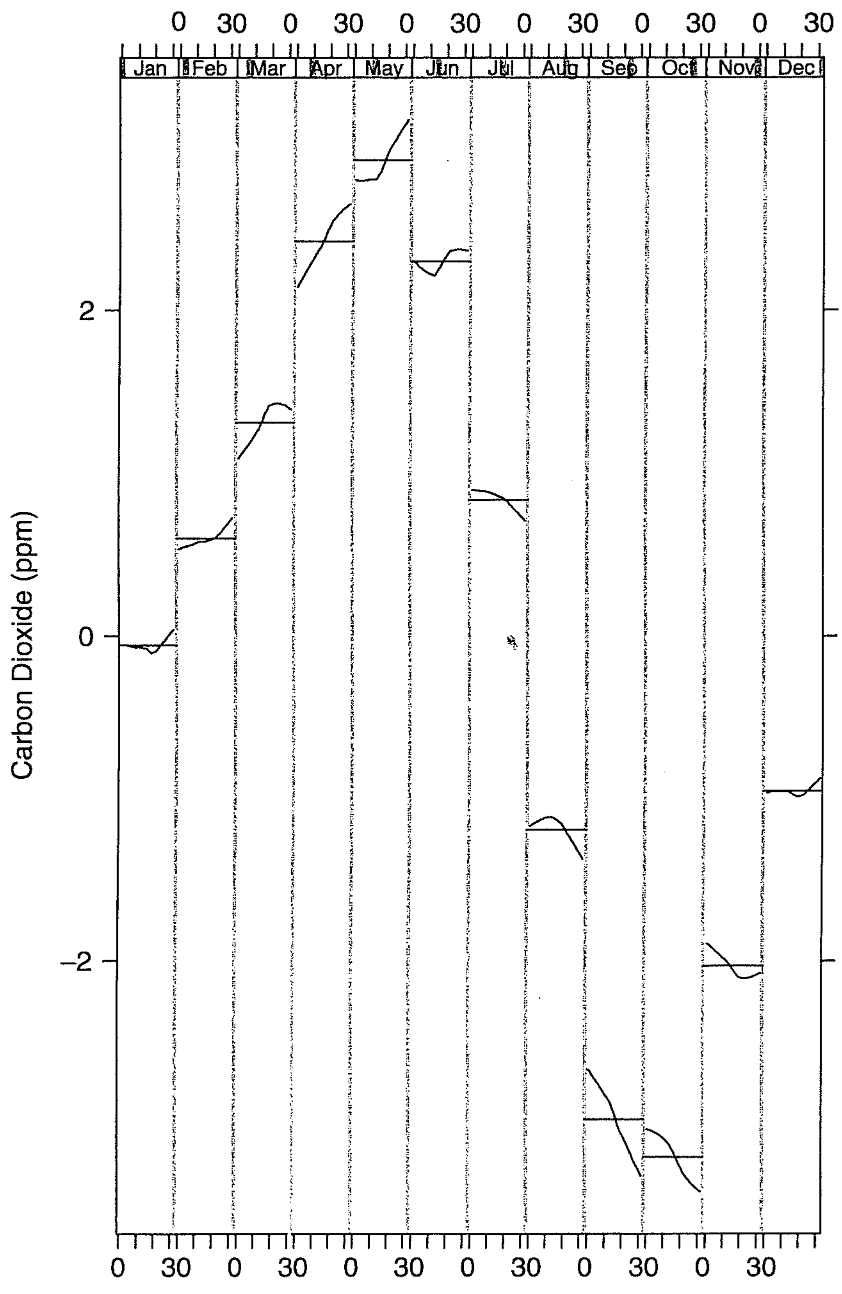

Figure 4.20 shows an alternative presentation of a month plot; it displays the seasonal component of the carbon dioxide data (Cleveland and Terpenning, 1982). This shows clearly that the annual cycle reaches a maximum level in May and a minimum level in October. It also clearly shows that the levels of carbon dioxide in the spring are increasing, whereas those in the fall are decreasing, so that the range of the cycles is increasing. The data set includes measurements of carbon dioxide for over 30 years (from 1959 through 1990).

Fig. 5.4 Barley Data

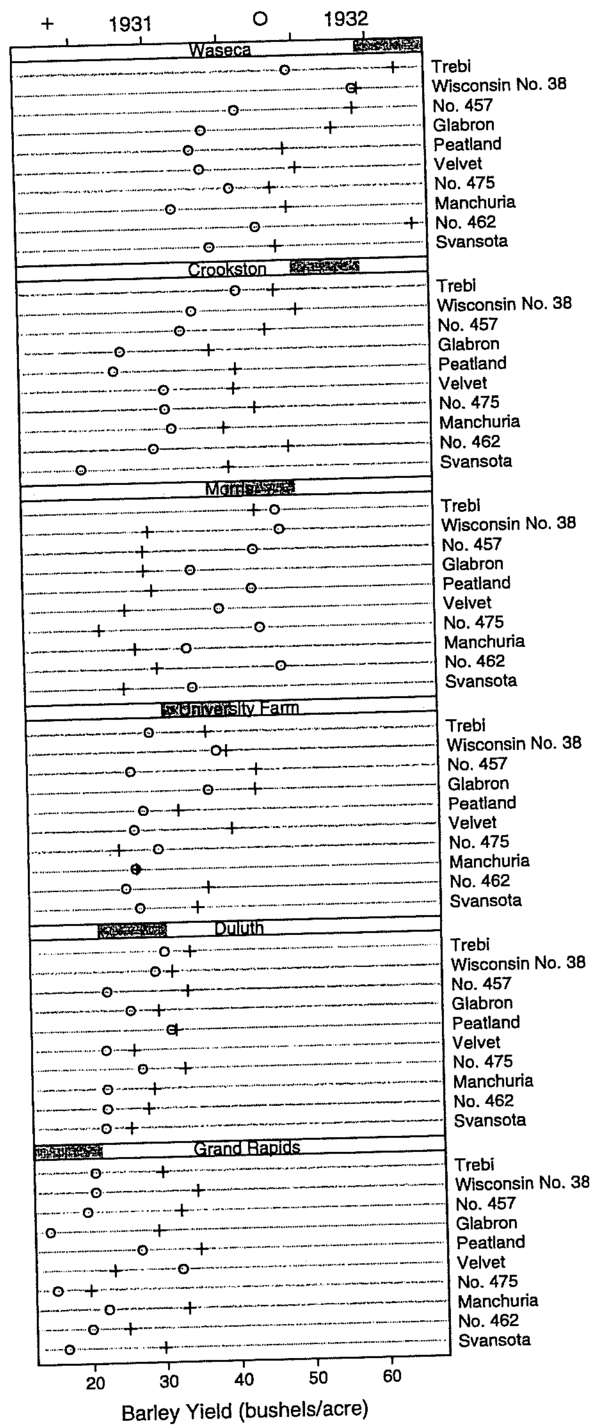

;.2 MORE THAN THREE VARIABLES

;.2.1 Superposed Data Sets

Many early statistical experiments were run on agricultural lata. One of these studied the yield of barley. Ten varieties of oarley were planted in six sites in Minnesota in 1931 and 1932. Yields were obtained for each combination of site, year, and variety. Therefore, there were 120 data points ($10 \times 6 \times 2$). Figure 5.4 shows these observations. Each panel contains a lot plot of the yield for each variety and each year for a specific site. Study this figure and comment on an interesting aspect of the data.