9 GENERAL GRAPH DESIGN

The visual nature of graphs requires a number of unique design practices. The volume and complexity of quantitative information that you can communicate with a single graph are astounding but only if you recognize and avoid poor design practices that would undermine your story.

Because of their visual nature, graphs tap into the incredible power of visual perception to communicate quantitative information. When the story that you wish to tell is contained in the data's patterns, trends, and exceptions; or when it depends on your audience's ability to compare entire series of values to one another (e.g., monthly domestic sales for the entire year compared to international sales), a graph will do the job best, but only if you avoid far-too-common design pitfalls.

We've already covered the aspects of quantitative communication that apply to both tables and graphs. None is more important to the design of graphs than the fundamental principle that was stated so eloquently by Edward Tufte: "Above all else show the data."[1] Quantitative stories reside in the facts, not in the containers that we use to present them. The general practice of highlighting the data and subduing all else is even more important in the design of graphs than in the design of tables. Tables are a bit more forgiving of visual design flaws because tables encode data through the use of verbal language (i.e., text), visually displayed. Graphs, in contrast, encode data as visual objects. These objects must be prominent, accurate, and clear.

1. Edward R. Tufte (2001) *The Visual Display of Quantitative Information,* Second Edition. Graphics Press, page 92.

Two fundamental principles of quantitative communication apply exclusively to graphs:

- Maintain visual correspondence to quantity.
- Avoid 3D.

Both principles are firmly rooted in practical concerns; you can wreak havoc on communication if you ignore these principles.

Maintain Visual Correspondence to Quantity

You can only use two attributes of visual perception to encode quantitative information in a way that can be easily and accurately interpreted: length and 2-D position. Quantitative values in graphs are either encoded visually as length in the form of bars or boxes or as 2-D position in the form of points and lines. Other visual attributes are either not perceived quantitatively at all (e.g., hue) or not well enough (e.g., 2-D area and color intensity) to justify their use for quantitative encoding when length and 2-D position are available.

A bar that is twice as long as another is perceived as having twice the quantitative value. Visual objects that encode quantitative values in graphs are interpreted by means of a scale line along the vertical or horizontal axis. When a bar that is twice as long as another corresponds to a value of two on the scale line, visual perception alone tells us that the value of the shorter bar is one, or very close to it. If the shorter bar actually corresponds to a value of 1.75 or 0.5, something is amiss.

People sometimes intentionally manipulate graphs to mask the truth contained in numbers. Darrell Huff, in his 1954 classic *How to Lie with Statistics*,[2] was one of the earliest to express this concern. Advertisements are notorious sources of deliberately misleading graphs, but deception is not confined to advertising. You'll be faced many times with the temptation to manipulate graphs to give your case more strength than it deserves based on the actual numbers. Given the understanding of visual design that you are developing by reading this book, you will be even better equipped to manipulate visual design to exaggerate or hide the truth. It's easy to rationalize little design manipulations here and there to shade the truth slightly for a just cause. Be aware, though, that this manipulation does not qualify as design for communication. The goal of design for communication is always to promote an accurate understanding of the truth.

2. Darrell Huff (1954) *How to Lie with Statistics*. W. W. Norton & Company.

Here's a simple illustration of the potential for deliberate misinformation:

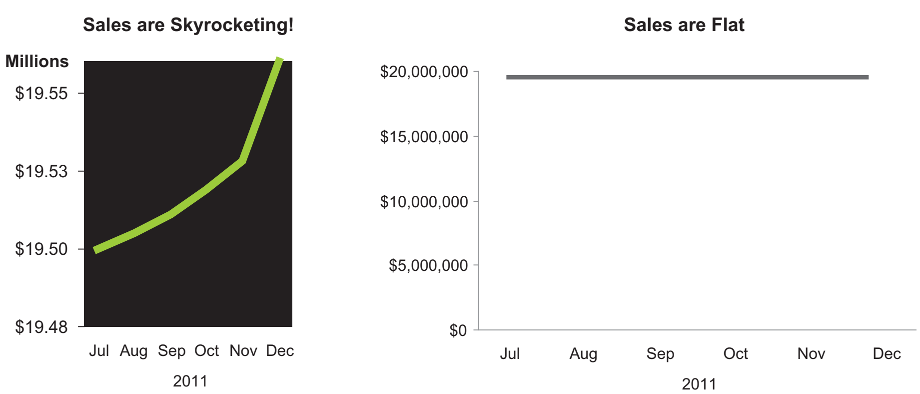

FIGURE 9.1 These two graphs display the same information in dramatically different ways, producing radically different messages.

The graph on the left has been deliberately manipulated to make an increase in sales from $19,500,000 in July to $19,560,000 in December, which is an increase of less than one-third of 1%, look like an increase of more than 200%. The graph on the right more accurately presents the data. Do you see the specific aspects of the graph on the left that were used to exaggerate the increase in sales? Take a moment to see how many you can find, and list them in the margin to the right.

• • • • • • •

Five design characteristics of the graph on the left give the false impression that sales have risen dramatically from July to December:

1. The scale on the Y axis does not start at zero. Rather, it starts at $19,475,000 and extends only to $19,560,000, thus making minor changes in sales appear extreme.

2. The plot area of the graph is taller than it is wide. This dramatically increases the slope of the line.
3. The line is green. The color green carries the meaning of growth and health in English-speaking cultures and dollars in the United States, so it reinforces the positive spin of the message. Also, placing a bright green line on a black background makes it pop with visual impact.
4. The highest value—the final value of $19,560,001—is set as the top of the scale. This gives the green line the appearance of extending right off the top of the graph.
5. Placing the boldfaced axis label Millions in the prominent upper left position near the title "Sales are Skyrocketing" suggests that they are increasing by millions.

This design certainly exaggerates the good news about sales, but I've seen much worse. Can you think of any additional design changes that could be made to further hide the truth?

• • • • • • •

Here's one that I've seen:

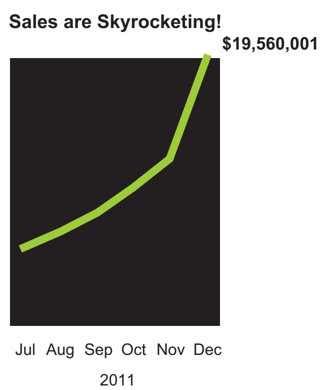

FIGURE 9.2 This is an extreme example of intentional deceit through graph design.

Notice the changes? Values along the Y axis have been removed, and only the final data point has been labeled. Without at least one more value on the scale, there is absolutely no way to know the extent of the increase. The single value of $19,560,001, combined with the characteristics we've already discussed, together suggest a huge rate of increase. By making the graph 3D and manipulating the angle, I could exaggerate the increase even more, which is done all the time.

Now, back to the principle that prompted our journey through the dark alleys of visual obfuscation. A quantity that is visually encoded in a graph should match the actual quantity that it represents. Two specific design practices will help you honor this correspondence:

- Make the distance between tick marks on a scale line correspond to the differences in the values that they represent.
- Generally include the value zero in your quantitative scale, and alert your readers when you don't unless you're confident that they won't be misled.

Correspondence to the Tick Marks

You should always keep the distance between tick marks on a scale line consistent with the difference in the quantitative values that they represent. Software that generates graphs for you based on specified sets of values automatically enforces this practice. If the tick marks represent the values 1, 2, 3, 4, and 5, they will be positioned an equal distance from one another. If you ever produce graphs without the aid of graphing software, you should be sure to honor this practice. Approaching this from the opposite perspective, if you have a set of tick marks that are positioned at equal distances from one another, the values that you use to label them should also represent equal numeric intervals. Never place a gap in the values, such as in consecutive tick marks labeled as 1, 2, 7, 8, and 9, even if there are no values in the graph that fall within the missing range. To do so would undermine the graph's visual integrity.

You may recognize that these tick marks would not be equidistant if you were using something other than a standard scale, such as a logarithmic scale. We'll look at the special qualities and uses of logarithmic scales a little later.

Even if you indicate a break in the quantitative scale where a section of values has been eliminated, your readers could still be easily misled.

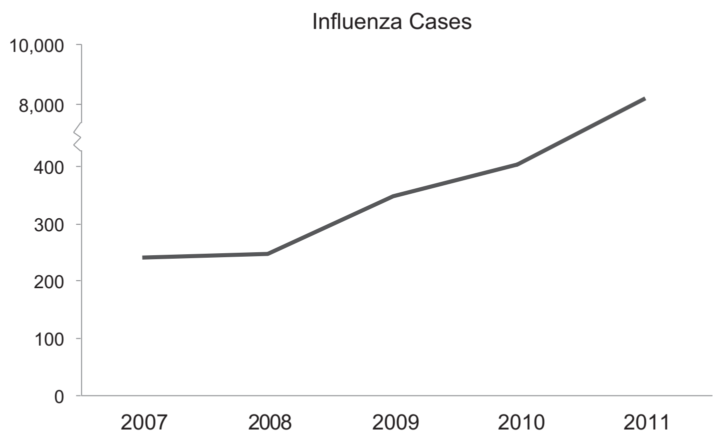

FIGURE 9.3 Scale breaks can be misleading.

Despite the fact that the scale starts at zero, the increase in influenza cases from 2010 to 2011 is underrepresented to a huge degree because of the scale break between 400 and 8,000 along the Y axis. Here's how the same values appear with a proper scale.

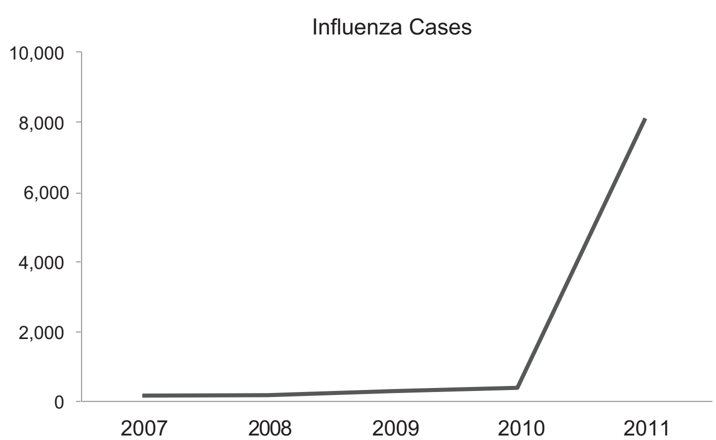

FIGURE 9.4 With a proper scale, the dramatic rise from 2010 to 2011 in influenza cases is striking.

Something is missing in this graph, however, that you might want your readers to see: the pattern of changes that occurred from 2007 through 2010. To accommodate the high number of cases in 2011, the graph's scale now forces all other values into a tiny space near the bottom, which makes the line appear almost flat during that period. How can we show the earlier pattern of change and yet still tell the more important story that the number of cases dramatically increased from 2010 to 2011? Can you think of a solution?

• • • • • • •

To tell this entire story, two graphs are needed, such as the following.

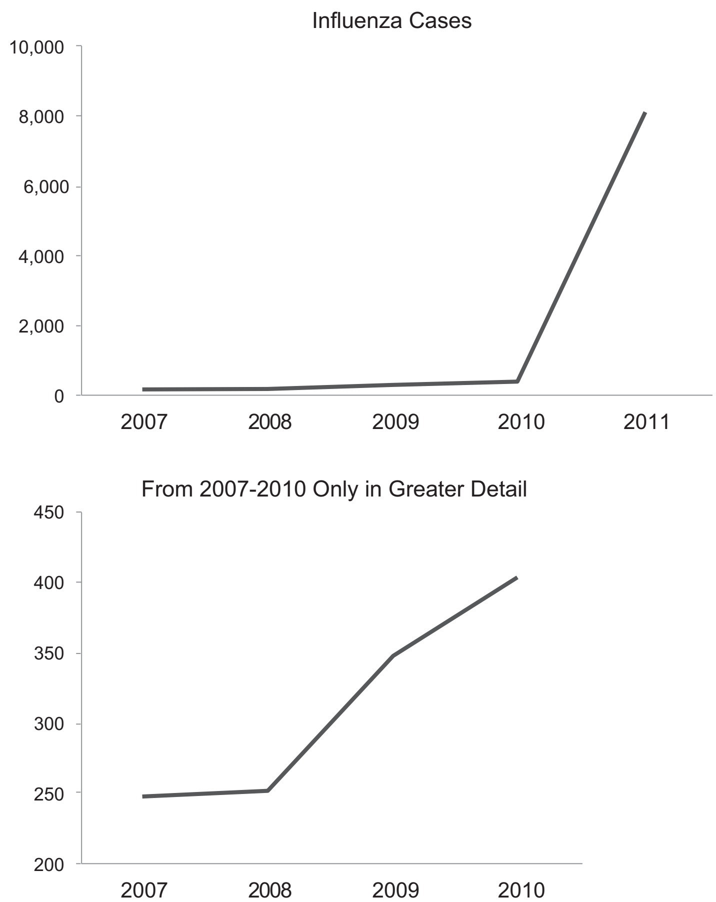

FIGURE 9.5 Two graphs are needed to tell all aspects of this story clearly.

It's important to know that quantitative stories can often only be told with more than one graph. Nothing is gained by attempting to squeeze into a single graph what can be more effectively presented in several.

Zero-Based Scales

When you set the bottom of your quantitative scale to a value greater than zero, differences in values will be exaggerated visually in the graph. Usually, you should avoid starting your graph with a value greater than zero, but when you need to provide a close look at small differences between large values, it's appropriate to do so. When you do so, alert your readers to the fact if you have any doubt that they'll notice. Perhaps you observed that the scale in the lower graph in *Figure 9.5* doesn't start at zero. Because the same information was already shown using a zero-based scale in the upper graph, the fact that the scale was adjusted in the lower graph would be hard to miss. If the sales manager of a company with the subtly rising sales that we examined in *Figures 9.1* and *9.2* wanted to examine that increase in great detail, however insignificant it might be as a percentage increase, the following graph would make this possible, but textual alerts similar to those shown in red might be needed.

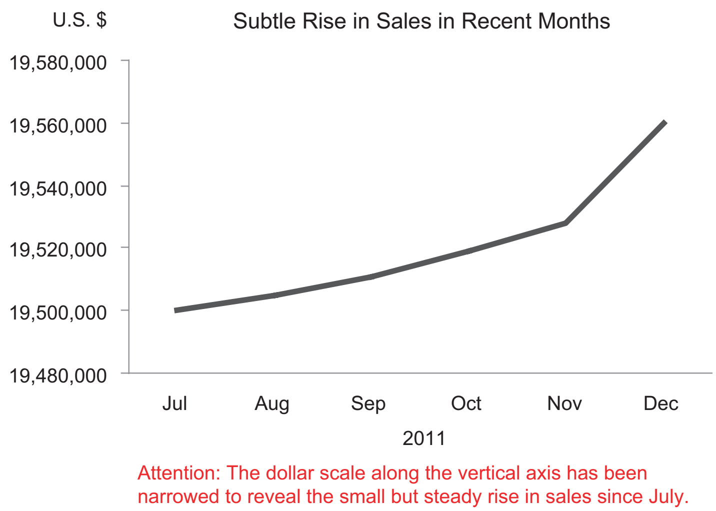

FIGURE 9.6 This is an example of an exception to the zero-based scale, illustrating how such an exception can be clearly noted to prevent misunderstanding.

Never eliminate zero from the quantitative scale when bars are used to encode the values, however. Why? Because a bar encodes quantitative value primarily through its length, and, without zero as the base, the length will not correspond to its value. In the following software ad, which I clipped from a magazine, how much greater is customer loyalty to MicroStrategy than Cognos Powerplay?

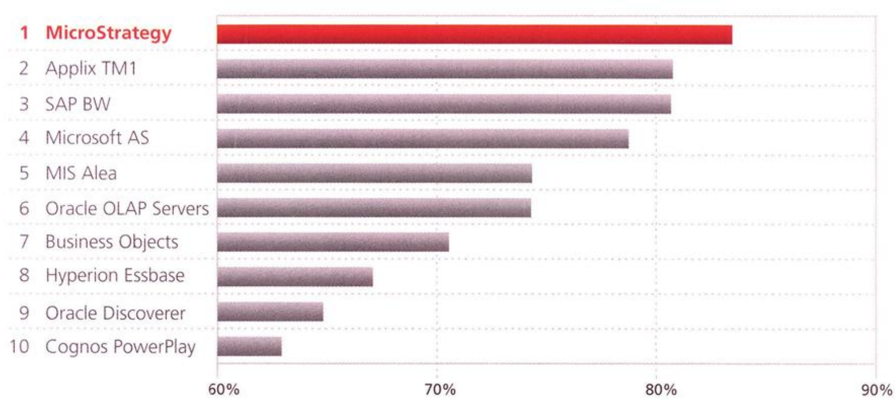

FIGURE 9.7 This graph misrepresents the values by starting the scale at 60%.

The MicroStrategy bar appears to be more than six times greater in length than the Cognos PowerPlay bar. The difference between the values, however, is about 83% versus 63%—quite a different story. Here's the same information, properly displayed:

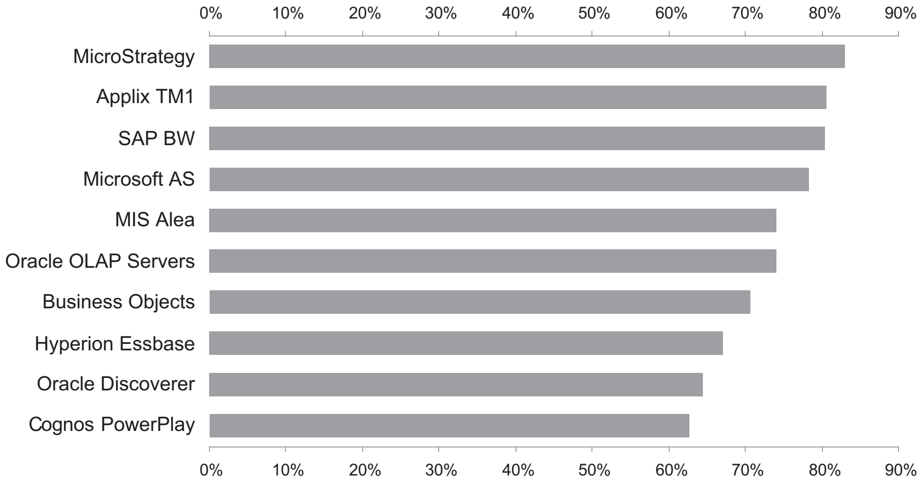

FIGURE 9.8 This graph displays the values in *Figure 9.7* properly.

When a graph represents both positive and negative numbers, zero will not mark the bottom of the scale, but it will still represent the base from which all values extend. The following two graphs contain the same set of positive and negative values. The graph on the right correctly displays zero as the base of its scale from which bars extend upwards for positive values and downwards for negative values, but the one on the left mistakenly sets the base to slightly below the lowest value, resulting in a confusing and misleading representation of the values.

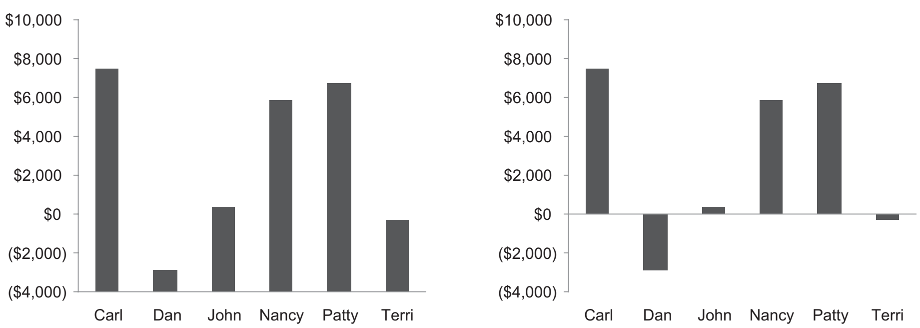

FIGURE 9.9 Both of these graphs display both positive and negative numbers. The graph on the right correctly sets zero as the base of its scale at the point where the X axis intersects the Y axis. The graph on the left incorrectly sets the base to slightly below the lowest value.

Avoid 3D

When 3D is used in graphs, it takes one of two possible forms:

- The addition of a third dimension of depth to objects (e.g., bars) that are used to encode quantitative values, without the addition of a third quantitative scale.
- The addition of a third dimension of depth to the overall graph with an associated quantitative scale (the Z axis).

Neither form is effective, but the reasons are entirely different.

Data Objects with 3-D Depth

We're using four objects to encode quantitative values in graphs: points, bars, lines, and boxes. The addition of depth to a value-encoding object does not affect the object's value. Add depth to a series of bars, and what do you have? Nothing more than bars that now occupy more space and are harder to tie to values along the scale. If you add depth to value-encoding points, like dots and squares, you get spheres and cubes that represent the same values as before, but now their depth makes it harder to align them accurately with the scale. 3-D versions of lines look like thick ribbons and suffer from the same problems.

Here are four variations of the same graph, three of which have 3-D effects added to the bars:

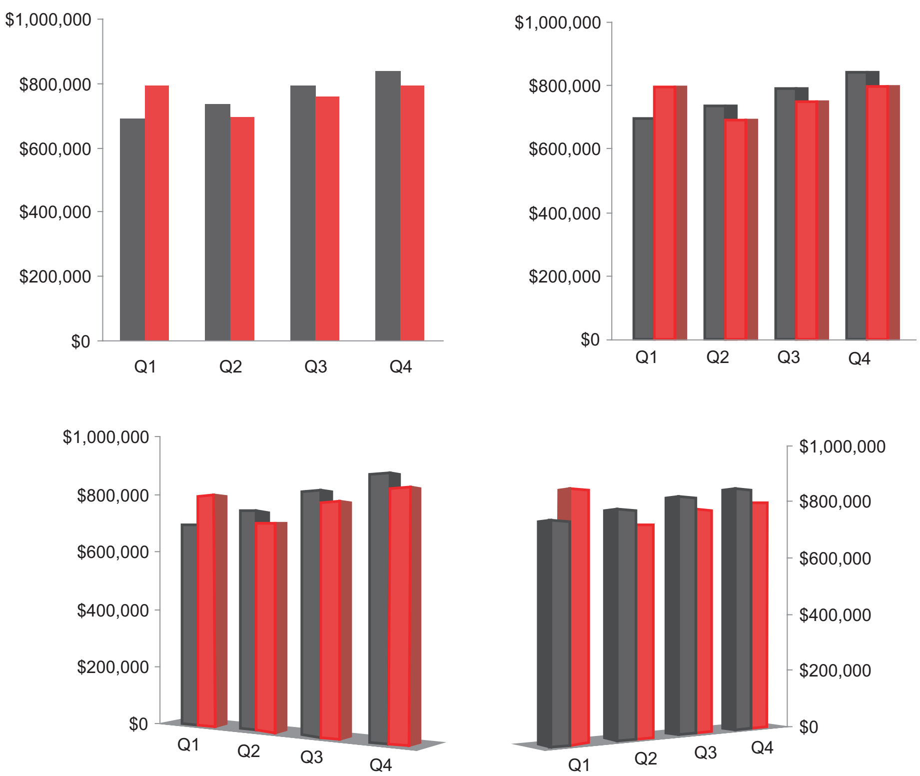

FIGURE 9.10 These four examples display the same values using bars in four different ways, three of which incorporate 3D.

Which graph is easiest to read? When shown all of these at once, the answer is obvious, isn't it? I was careful in the graphs above to keep the 3-D effects simple. If we take advantage of the many options that most software provides, we can bury truth in visual effects. In the following example, I've manipulated perspective and angles to make a steady increase in expenses from $100,000 to $121,000 look like a flat series of consistent values.

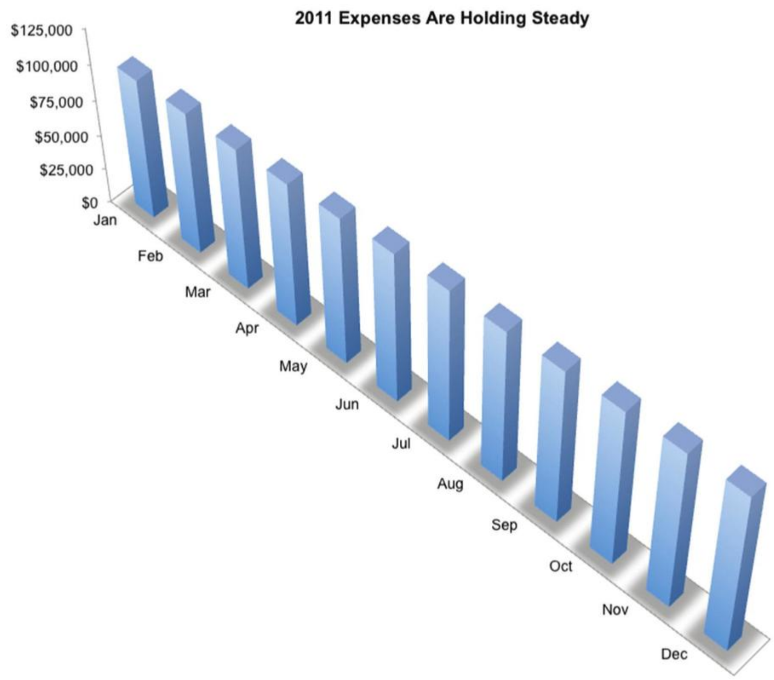

FIGURE 9.11 3-D effects are sometimes used to tell lies.

Most software makes it far too easy and tempting to add a third dimension to objects in graphs. This functionality is thrown in because people expect it, not because it's useful. It is far better to impress your readers with graphs they can easily understand and use than graphs that look like video games and are difficult to interpret.

Remember the data-ink ratio. The addition of 3D to value-encoding objects adds ink but not data. That is, it adds meaningless visual content that your readers must take in and process, resulting in nothing but wasted time and effort.

Graphs with 3-D Depth

A third dimension of depth may be added to an entire graph through the use of a third axis, conventionally called the Z axis. The Z axis may be used for either a categorical or a quantitative scale. A categorical scale along the Z axis allows you to add another set of categorical items that extend back along the axis, accompanied by related rows of quantitative values. A quantitative scale along the third axis can display a third quantitative variable in a scatter plot. In theory, this is a valid way to include more information in a graph. In practice, with rare exceptions, it is simply too hard to read. Simulating 3-D space on a 2-D surface works nicely for paintings or technical illustrations but almost never for graphs.

A few examples will vividly illustrate this point. Let's start with the same data that we examined earlier as the dark gray bars in *Figure 9.10*.

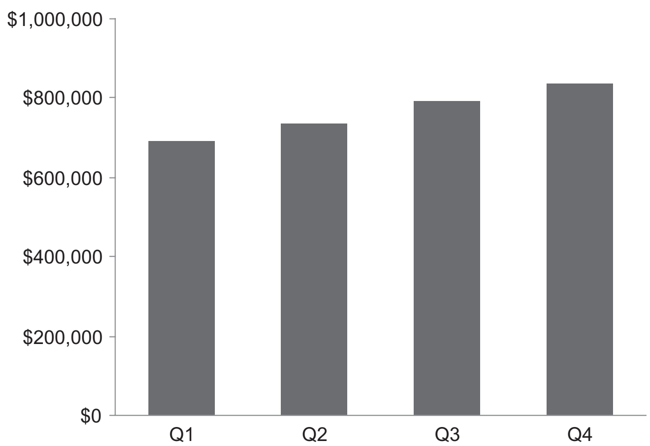

FIGURE 9.12 This is an example of a simple 2-D graph.

So far we have a very simple 2-D graph. Now let's say that we want to display these quarterly bookings by the four sales regions of North, East, South, and West. To do so, we could encode each region as a different hue and keep the graph 2D, as follows:

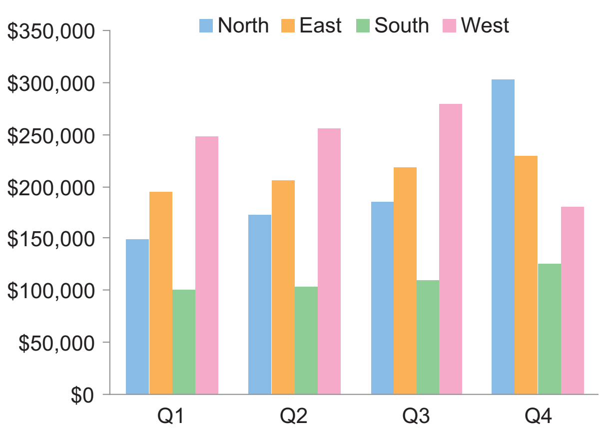

FIGURE 9.13 This 2-D graph has two sets of categorical items: quarters along the X axis and sales regions encoded as different hues.

This is still fairly easy to read. Rather than using hue to encode the four sales channels, we could instead add a Z axis to the graph, making it 3D, and display the sales channels along that axis.

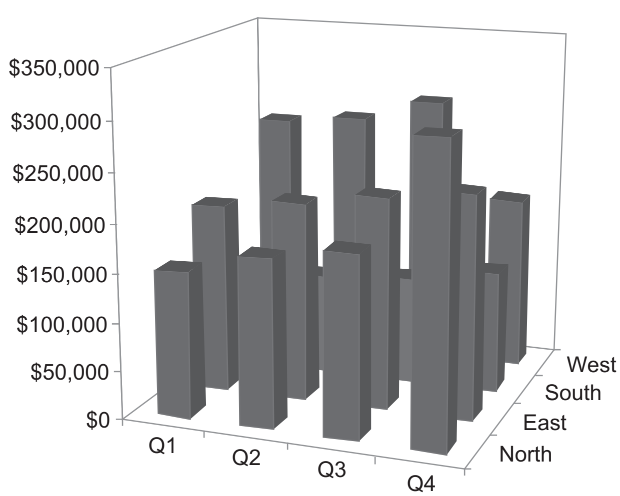

FIGURE 9.14 This is a 3-D graph, with sales in dollars along one axis, quarters along another, and sales regions along a third.

This is a very simple example of a 3-D graph with two categorical scales (quarters and regions) on one quantitative scale (dollars). What do you think? Does it work? Examine it for a moment, attempting to read and compare its values. Try to describe what makes this graph difficult to read.

• • • • • • •

When a third dimension is added to a graph, adjustments are usually made to the graph automatically by software—tilting, rotating, and adding perspective—to make its data more visible. A 3-D display like this is called an *axonometric projection*. The previous example was tilted down 15 degrees, rotated clockwise 20 degrees, and given 30 degrees of perspective. These variables can be altered in an effort to make the graph easier to read. Even though the graph has been tilted and rotated in an attempt to make the rows of bars more visible, some bars will always remain partially or entirely hidden. Also, it's nearly impossible to line the bars up with values along the quantitative scale.

Software that generates 3-D graphs often includes grid lines on the walls in an effort to make the quantitative values easier to align with the scales lines. Here's the same graph as before with the addition of these features, along with black borders around the bars to more clearly delineate them.

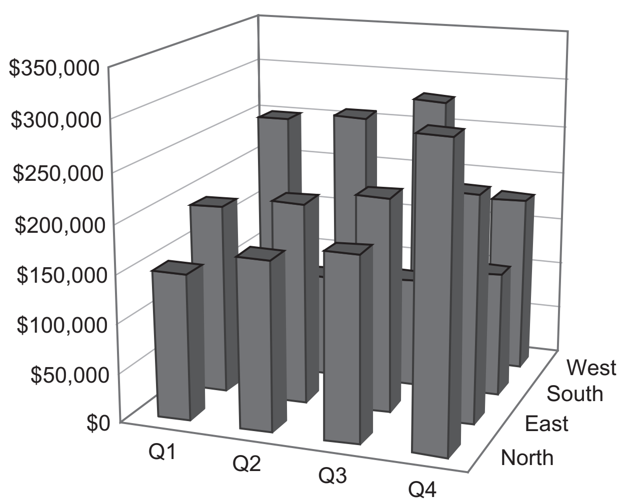

FIGURE 9.15 This is a 3-D graph that has been enhanced in an effort to make the values easier to read through the use of grid lines on the walls and borders around the bars.

Even though this is a fairly simple graph, these enhancements still don't solve the problems. Software vendors sometimes argue that this problem can be solved by rotating the graph to see bars that are hidden. The problem with this approach is that any new perspective will reveal some bars and hide others, never allowing us to see and compare all the values at once, which is a key benefit of graphs. Changing from the use of bars to lines to encode the data doesn't fix the problem either, as you can see in the following example:

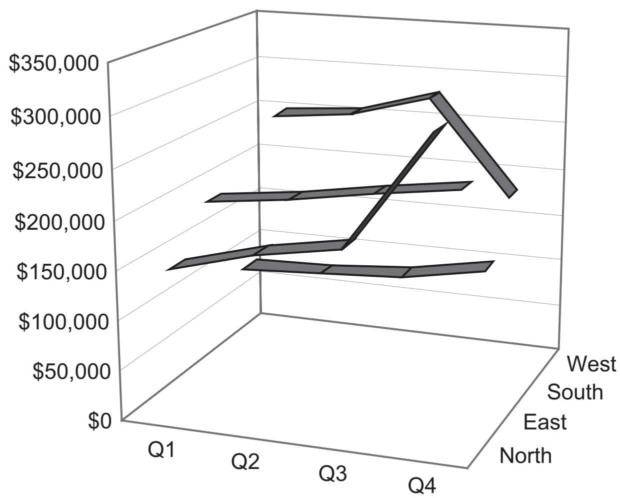

FIGURE 9.16 This graph displays the same data as *Figure 9.15* but this time using lines to encode the values.

Which of the lines in this graph represents the south region? When I ask this question in classes, fewer than 50% of my students answer correctly. The lowest of the four lines represents the south region, but this isn't at all obvious, is it? Support components called *drop lines* were invented to help us locate data objects in relation to scales along axes, especially in 3-D graphs, but they clutter the graph and reduce its interpretation to a slow series of look-ups.

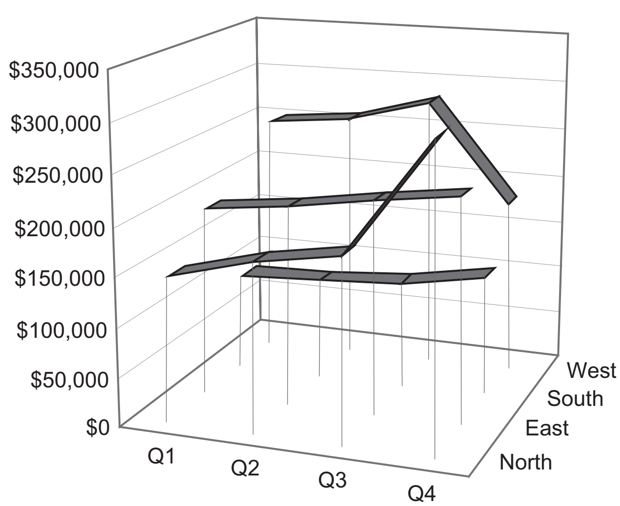

FIGURE 9.17 This graph is the same as the one in *Figure 9.16* with the addition of drop lines.

So far we've only examined the association of a categorical scale with the third axis. The problems don't get any better when the third axis is used for a quantitative scale. Imagine a scatter plot that correlates employee salaries in dollars along one axis, tenure on the job in years along another axis, and level of education in years along the final axis. It's too difficult to tell where the points are positioned along the third axis.

3-D renderings of quantitative information rarely work. Don't sacrifice effective communication through the use of 3-D fluff. Even when you are driven by a sincere desire to give your readers more information by using a third dimension, there are better ways to realize these good intentions. One effective technique is to use multiple related graphs in a series, which we'll explore in Chapter 11, *Displaying Many Variables at Once*.

Summary at a Glance

- Encode quantities to correspond accurately to the visual scale.
  - Keep the distance between tick marks on a scale line consistent with the difference in the quantitative values that they represent.
  - In most cases include the value zero in your quantitative scale, and alert your readers when you don't. Always start the quantitative scale at zero when you use bars to encode the values.
- Avoid 3-D displays of quantitative data.

10 COMPONENT-LEVEL GRAPH DESIGN

Several visual and textual components work together in graphs to present quantitative information. If these components are out of balance or misused, the story suffers. For each component to serve its purpose, you must understand its role and the design practices that enable it to fulfill its role effectively.

Graphs are constructed from components. Most but not all components represent data. Data components can be divided into two groups: primary components (points, bars, lines, and boxes) and components that serve secondary roles (scales, trend lines, tick marks, and so on). This chapter is organized into the following sections:

- Primary data component design
- Secondary data component design
- Non-data component design

Before jumping in, let's get our terminology straight. Here's a diagram that shows most of the terms used to describe graph components in this chapter:

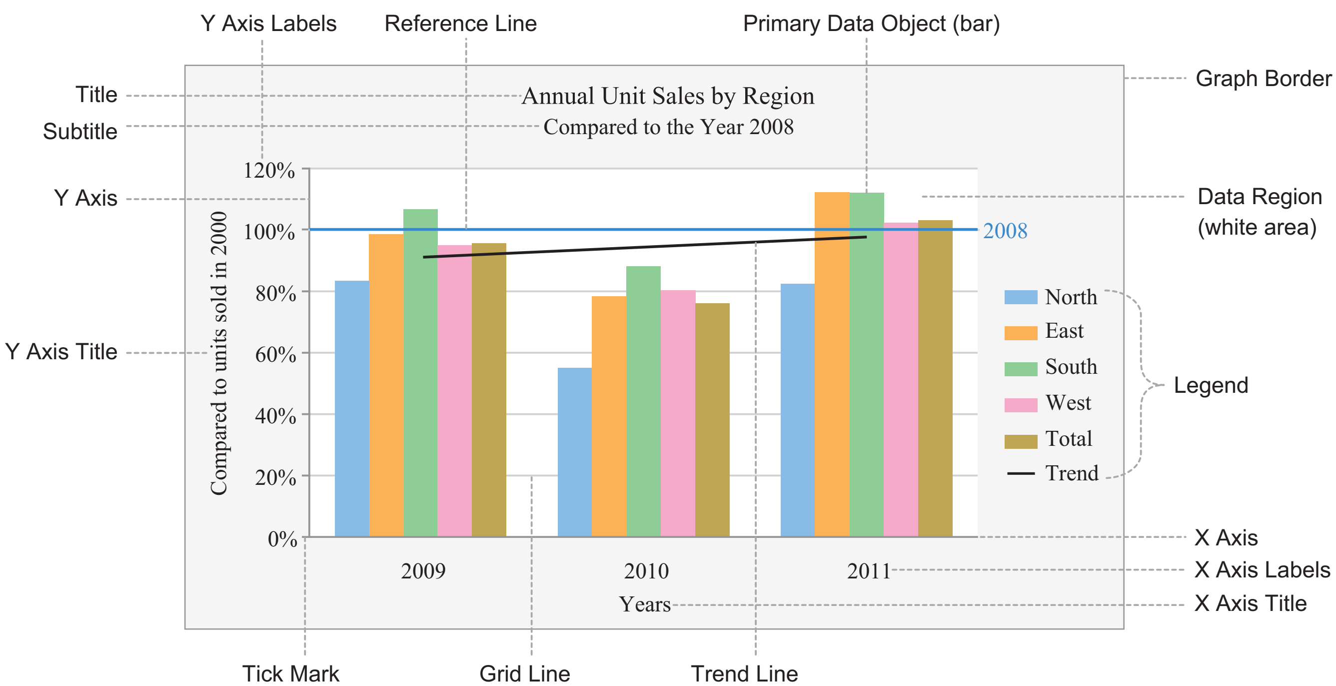

FIGURE 10.1 This figure is designed only to label the parts of a graph, not to illustrate best practices in graph design.

Primary Data Component Design

Graphs that are properly designed can be decoded as an act of visual perception, tapping into our powerful ability to detect patterns in quantitative information that is displayed visually. Quantitative values are encoded in graphs as points, bars, lines, or boxes. To assign precise values to these objects, most graphs use quantitative scales along one or more axes, labeled with numbers and tick marks. Categorical items are primarily encoded as visual attributes of the points, bars, lines, or boxes (e.g., color or location in 2-D space). Text labels are used to assign categorical items to these attributes, either along the axes or by means of legends. Just as in tables, text may also be used in graphs to introduce the display and provide explanations, reinforce and highlight particular items, sequence elements, recommend responses in the form of decisions and actions, and pose important questions.

Because we've already spent time in Chapter 6, *Fundamental Variations of Graphs* examining how points, bars, lines, and boxes encode values, we'll move on now to 1) the finer details of their design, 2) ways they can be combined for specific purposes, and 3) a few data components that we haven't examined so far, such as tick marks and reference lines.

Points

When you encode quantitative values as points (dots, squares, triangles, and so on), all design practices address a single primary objective: keep points easy to perceive. Let's walk through a few examples of graphs that illustrate various perceptual problems that can plague points.

This first example uses different point shapes to distinguish two sets of values. However, it's harder than it should be to see the two groups as distinct. Take a moment to describe, in the margin to the right, the problem as you see it:

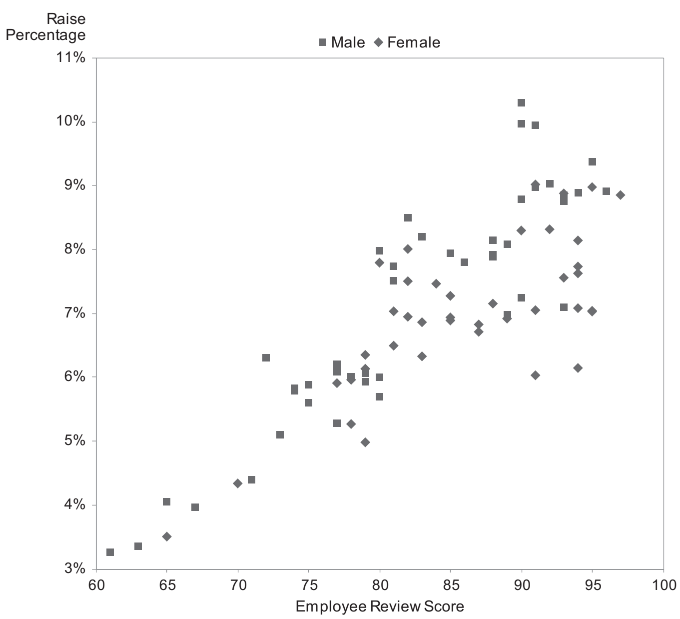

FIGURE 10.2 This is an example of points that suffer from a particular perceptual problem.

The problem here is that the points are so small that their distinct orientations—vertically oriented squares for the men versus squares rotated 45° (i.e., diamonds) for the women—cannot be distinguished easily. We can remedy this lack of perceptual distinction between the two sets in a number of ways. Take a minute and try to identify one or two ways to fix this problem.

• • • • • • •

One simple solution involves enlarging the points. Nothing is different about the following example except that the points are bigger:

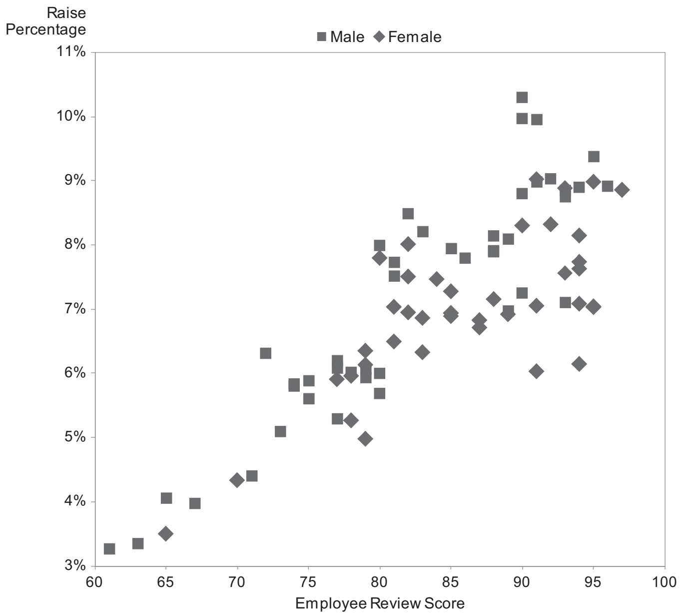

FIGURE 10.3 This example demonstrates that sets of points can be made more distinct by enlarging them.

Enlarging the points certainly helps, but another minor change will help even more: remove their fill color. Differences between simple shapes, or in this case between a single shape that has been oriented in two ways (the diamonds are merely squares rotated 45 degrees), can be more easily distinguished when the fill color has been removed, leaving only their borders, as you can see in this next example:

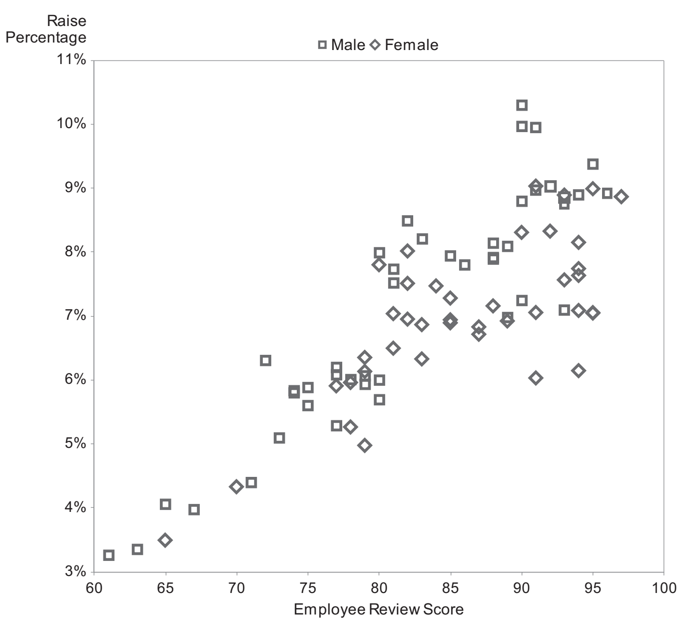

FIGURE 10.4 This example demonstrates that sets of points can be made more distinct by removing fill color.

If you want a method that doesn't involve enlarging the points, select shapes that are more visually distinct. An example appears on the next page.

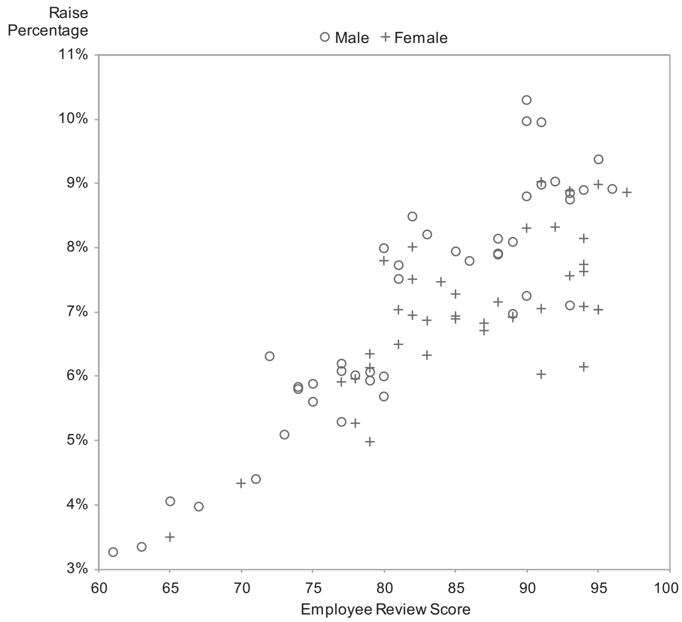

FIGURE 10.5 This example demonstrates that sets of points can be made more distinct from one another by using shapes that are very different.

Some shapes that are easiest to distinguish from one another are the following:

All of these shapes are now available in Excel. These symbols are easier to distinguish if you keep those with interiors (circle, square, and triangle) empty of fill color.

So far we've used different point shapes to distinguish groups, but there is a preattentive visual attribute that we can distinguish better than shapes. What do you think it is?

• • • • • • •

The answer is hue. In the next example, the two groups of points can be easily distinguished even though they're fairly small.

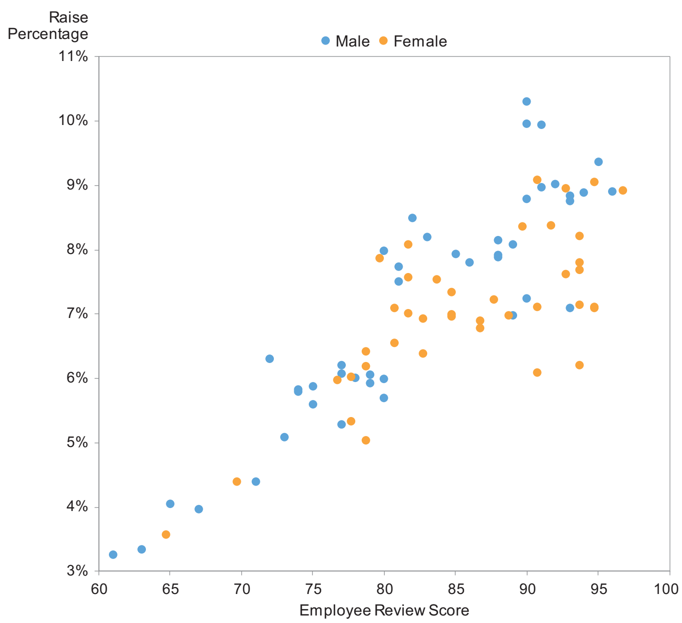

FIGURE 10.6 This example demonstrates that sets of points can be made more distinct from one another by using different hues.

Hues work best for distinguishing groups of points. Stick with hue for this purpose unless you can't because you've already used hue differences for another purpose in the graph, or you're concerned that the graph will be printed in black and white, resulting in the loss of hue distinctions.

Let's move on to another perceptual problem that can plague points, especially in scatter plots that show a large number of points. Take a look at the following example and describe in the right margin the problem that you detect:

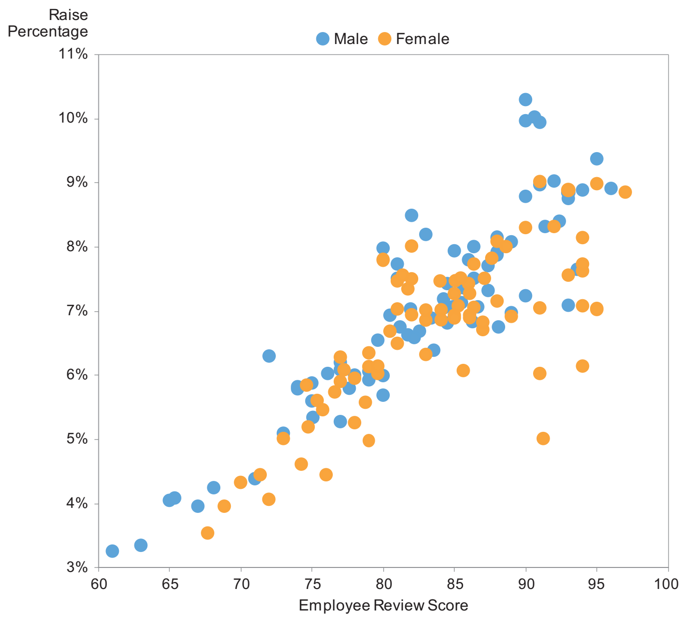

FIGURE 10.7 This example shows points that suffer from a visibility problem.

• • • • • • •

This scatter plot contains many more points than the previous examples. The problem that it illustrates is that points can overlap, potentially causing some to be obscured completely. This problem is called *over-plotting*. When points are hidden or overlapped by other points to the degree that they can't be distinguished, incomplete and inaccurate communication results.

When points overlap somewhat, but none are entirely hidden, this can sometimes be remedied by enlarging the graph, decreasing the size of the points, or a combination of the two. These steps will reduce the overlap. If the problem persists, another technique can be used that usually fixes the problem. Knowing what you now know about visual perception, look again at the last example and see whether you can come up with any other way to remedy the problems associated with the overlapping points.

• • • • • • •

Any luck? Your ability to recognize opportunities to improve visual perception by manipulating visual attributes should be sharpening, so you may have detected that the solid nature of the points, the fact that they are filled with color, reduces their discrete visibility where overlapping occurs. Let's see what happens when we remove the fill colors, leaving only the outlines of the dots:

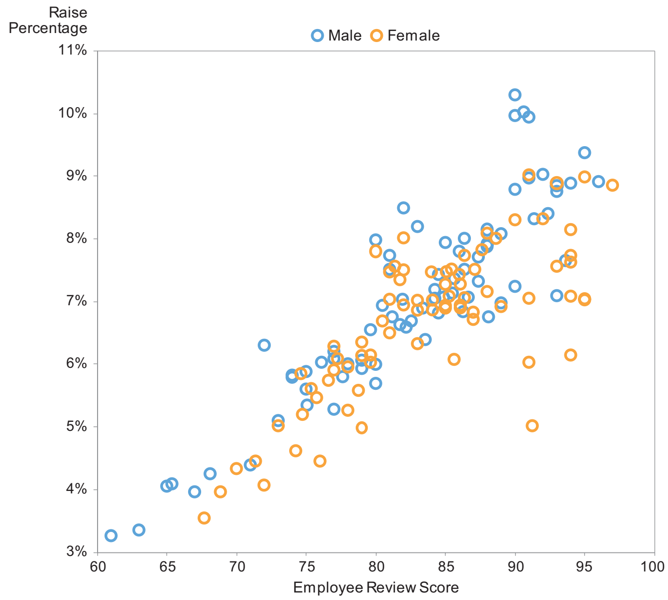

FIGURE 10.8 This example demonstrates that the problem of overlapping data points can be reduced by removing the fill color from the points.

This makes a big difference. When the points have transparent interiors, we can more readily see when points overlap. This same technique works with a variety of point shapes, not just dots.

Bars

Several characteristics of bars deserve attention. We'll consider each of the following:

- Orientation
- Proximity
- Fills
- Borders
- Base value

ORIENTATION

In this context, orientation refers to the direction that the bars run, either horizontally from left to right or vertically from bottom to top. Each orientation has advantages in particular circumstances.

Horizontal bars are the best choice when it would be difficult to get categorical labels to fit side by side under vertical bars. In the following example, even though the names of these sales representative names are not abnormally long, placing them side by side to label vertical bars results in an awkwardly wide graph:

FIGURE 10.9 This graph illustrates the problem with vertical bars when their categorical labels are long.

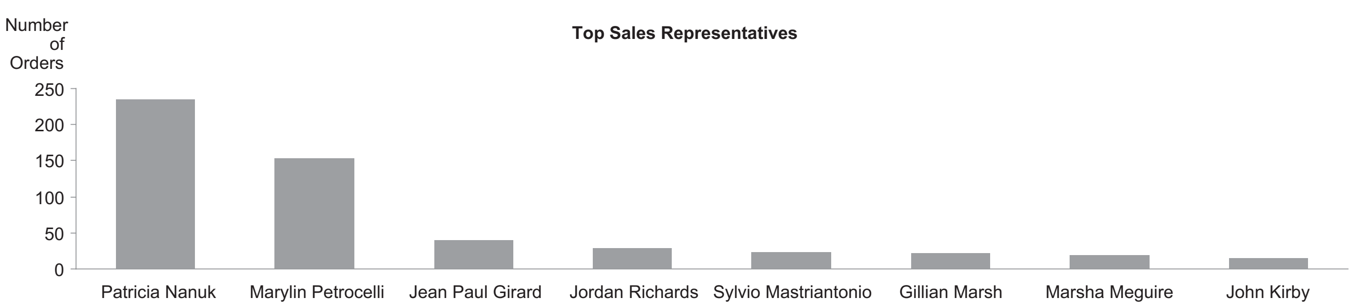

Perhaps you could make this work, but what if the graph displayed twenty sales representatives rather than eight?

You might be thinking that this problem could be solved by orienting the names at an angle or vertically, as in the following example.

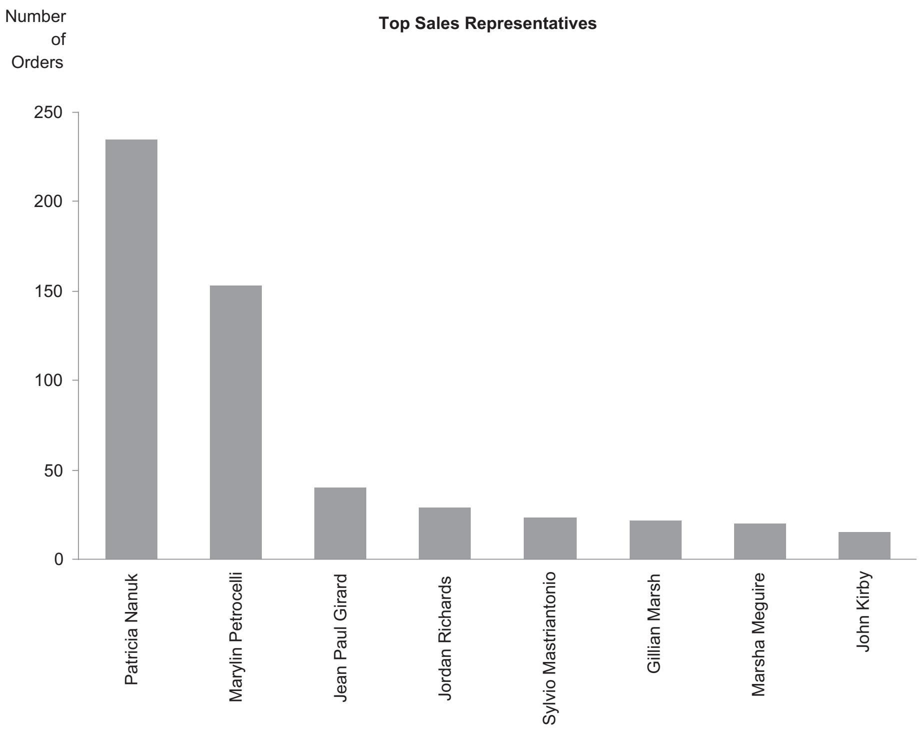

FIGURE 10.10 This graph illustrates the problem that results when we associate long categorical labels with vertical bars by orienting them vertically.

This solves the horizontal space problem but makes the names hard to read. When you must use vertical bars and the categorical labels are too long to fit side by side, opt for an upward-sloping angle of 45° or less rather than orienting the labels vertically. This is much easier to read, as in the following example:

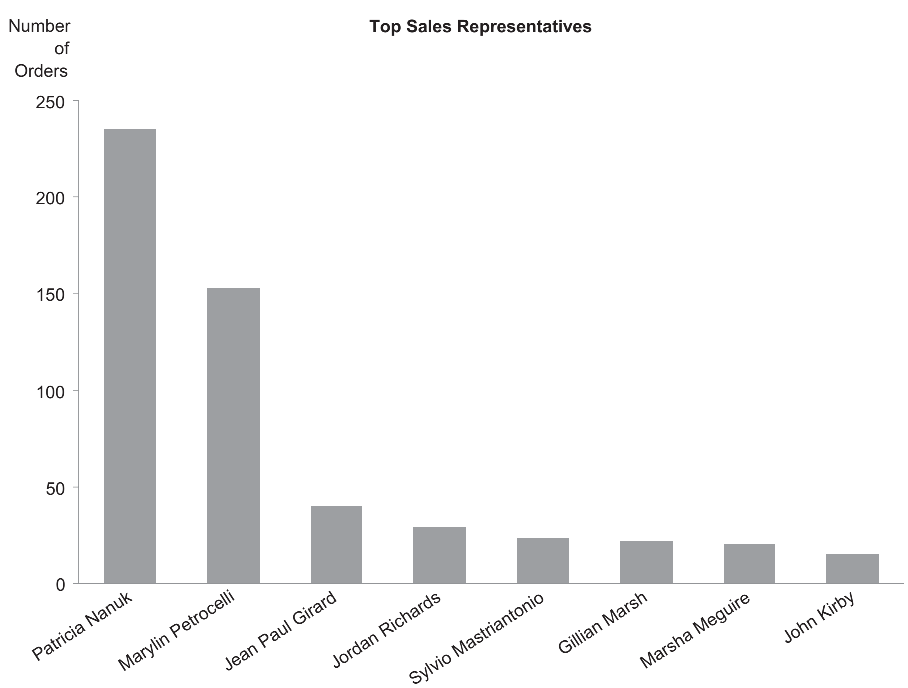

FIGURE 10.11 This graph orients the categorical labels at an angle, rather than vertically, to make them more legible. These are oriented at a 33° angle.

It is almost always better, however, to solve the problem with horizontal rather than vertical bars if you can. Notice how well this solution works.

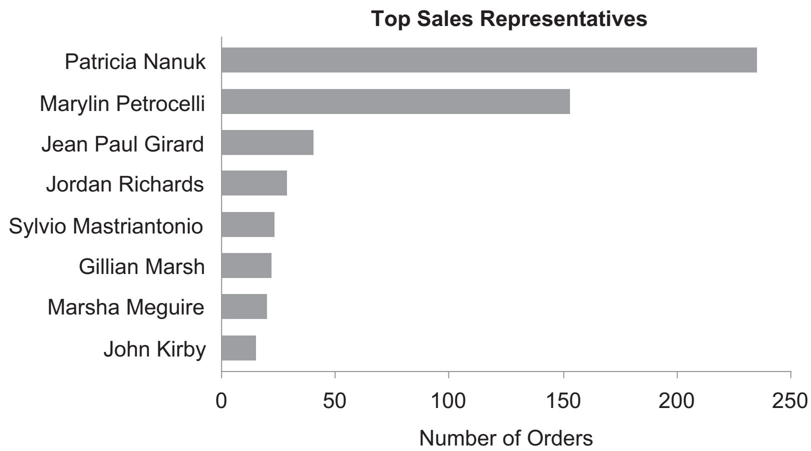

FIGURE 10.12 Horizontal bars handle long categorical labels nicely.

Long labels aren't a problem at all when you use horizontal bars. Even when labels aren't especially long, if you need to show a large number of bars in a graph, orienting them horizontally usually works best. The one exception is when you're displaying change through time (years, months, days, etc.) in a bar graph; you should almost always place time on the X axis so that it runs horizontally from left to right, which will result in vertical bars.

PROXIMITY

Now let's consider how closely bars should be placed to one another. You should always maintain a balance between the width of the bars themselves and the white space that separates them. In other words, rather than thinking of the space between the bars in absolute terms, think of it in terms of a ratio of the width of the bars to the width of the white space between them. Here are several examples of graphs with different ratios of bar width to white space. Take a look at the next five figures and determine which seems to work best:

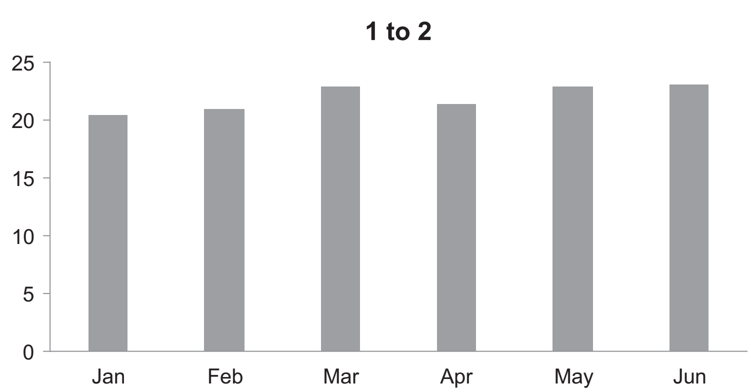

FIGURE 10.13 This is an example of a 1-to-2 ratio of bar width to intervening white space.

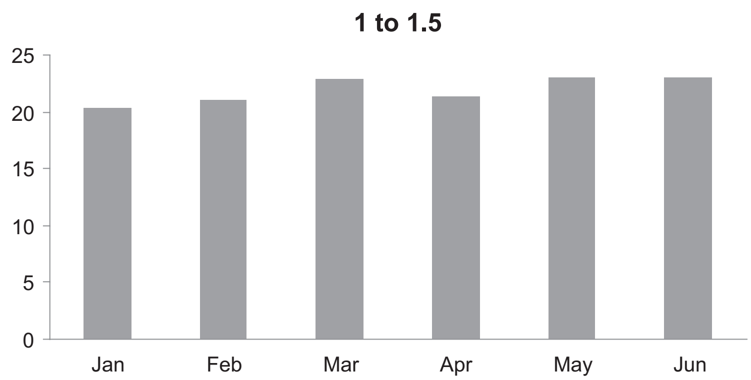

FIGURE 10.14 This is an example of a 1-to-1.5 ratio of bar width to intervening white space.

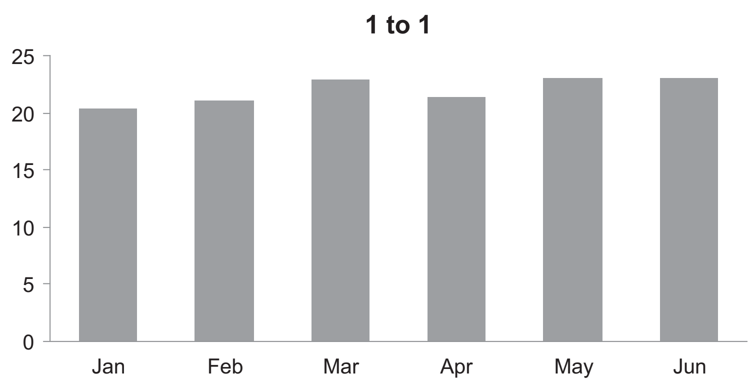

FIGURE 10.15 This is an example of a 1-to-1 ratio of bar width to intervening white space.

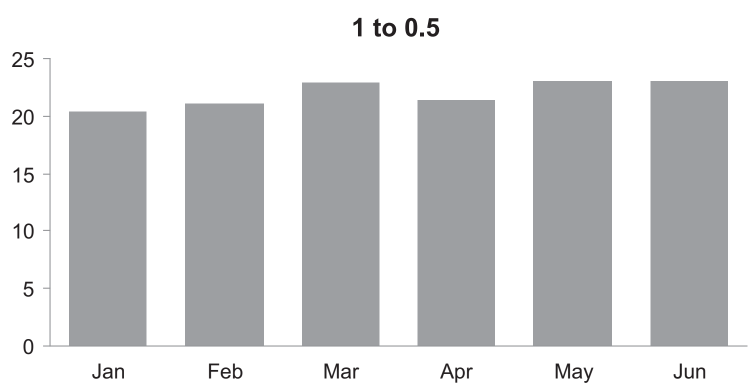

FIGURE 10.16 This is an example of a 1-to-0.5 ratio of bar width to intervening white space.

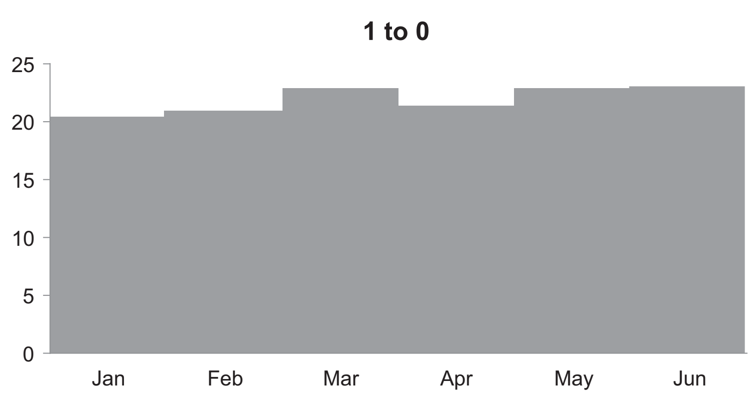

FIGURE 10.17 This is an example of a 1-to-0 ratio of bar width to intervening white space.

I'm not aware of any research that suggests which of these ratios works best. Personally, I prefer ordinarily to stick within the range extending from a ratio of

1:1.5 to 1:0.5 and lean toward a ratio of 1:1 as ideal. Larger ratios produce too much white space. At the other extreme, as in the 1:0 ratio example, the bars cease to be discrete, suggesting a continuous range of values that is only appropriate along an interval scale or when displaying a part-to-whole relationship. At this extreme, the unique ability of bars to display individual values as discrete is almost entirely lost.

The primary occasion when space between bars isn't needed is when several bars correspond to each categorical item labeled along the axis. In the examples below, four bars, one per region, appear side by side in each quarter.

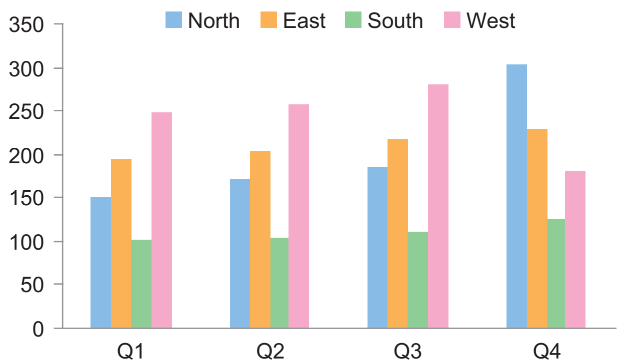

FIGURE 10.18 This graph illustrates when it is appropriate to place bars side by side without white space in between.

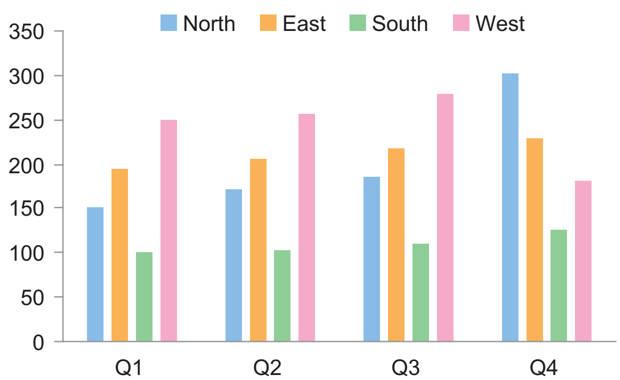

FIGURE 10.19 This graph shows that it's unnecessary to insert white space between bars that are associated with the same categorical item along the axis.

Allowing the bars within each quarter to touch, as they do in the upper example, reinforces the fact that they belong together as part of a group.

What about overlapping the bars? You could go further than removing white space between bars and actually overlap them. Examine the next example to see how this looks:

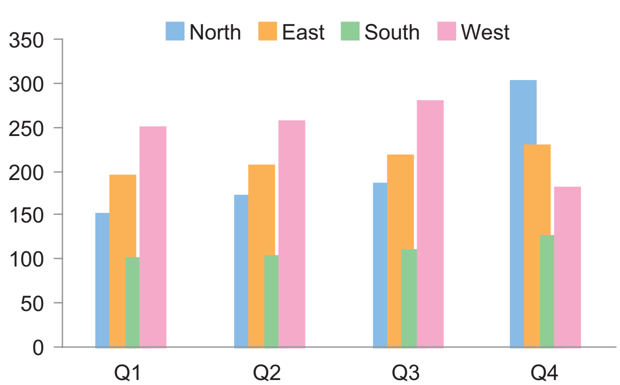

FIGURE 10.20 This is an example of overlapping bars, which illustrates the visual problems that overlapping creates.

When bars overlap in this manner, they look a bit like a jigsaw puzzle. The pink bars appear more dominant than the others because they occupy the front position. Some bars, like the orange ones, take on odd shapes. Overall, the graph is visually confusing, and thus hard to read, showing clearly why it's best to avoid overlapping bars.

FILLS

The use of fill colors or patterns in bars should follow the general design practices that we've already examined:

- Avoid the use of fill patterns (e.g., horizontal, vertical, or diagonal lines) because they create disorienting visual effects.
- Use fill colors that are clearly distinct.
- Use fill colors that are fairly balanced in intensity for data sets that are equal in importance.
- Use fill colors that are more intense than others when you wish to highlight particular values above the others.

This last practice in this list is a useful way to direct your readers' attention to particular data without being overbearing, such as by giving the south region in the following example a darker, more saturated color.

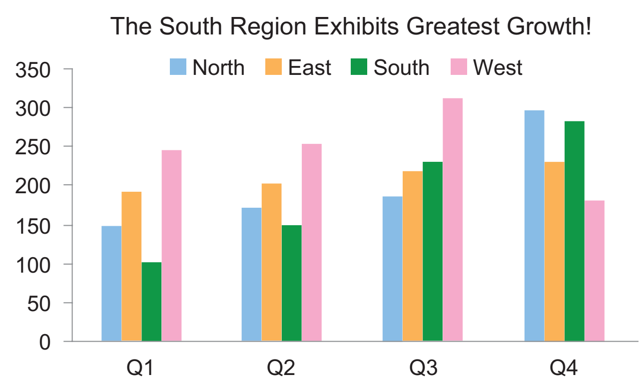

FIGURE 10.21 This is an example of using an intense fill color, in this case dark green, to highlight a particular set of values.

We should add one more practice to this list: use only one fill color per set of related values. This practice hasn't been followed in the graph on the left below:

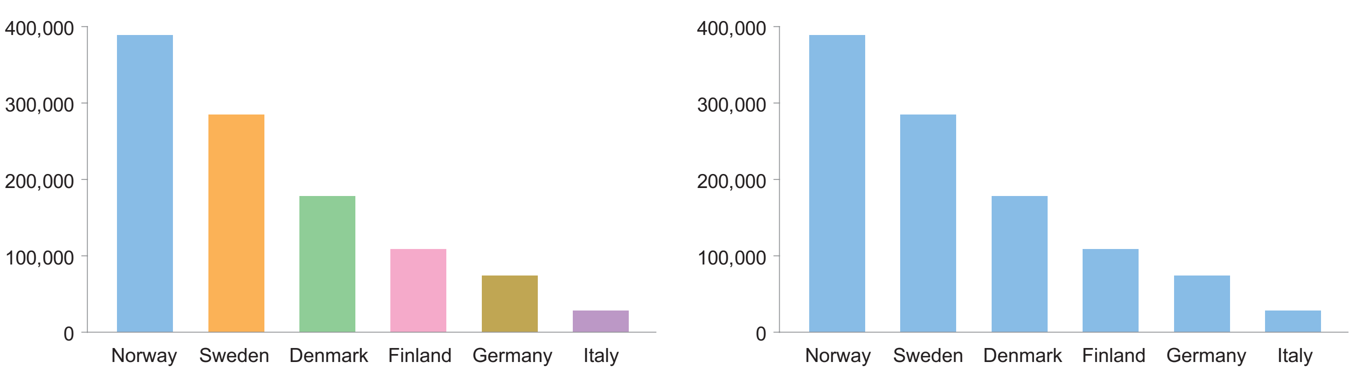

FIGURE 10.22 For a single series of values, the bars on the left differ in color and those on the right are the same.

Which one of these graphs does a better job of encouraging your eyes to compare the values? The one of the right, correct? The sameness of the color causes you to see the bars as a group, which in fact they are, and your eyes are therefore encouraged to compare them to one another. The one on the left does the opposite. Differences in color discourage us somewhat from comparing the bars. Given the fact that each bar is identified with labels along the X axis, what do the colors in the graph on the left mean? Nothing. A good general principle in graph design is to never vary anything—color, size, position, angle, etc.—unless the differences mean something. Meaningless variation unnecessarily complicates the display and sends your readers on a search for meanings that don't exist.

BORDERS

A border around a bar is only visible if the border's color is different from the fill color of the bar. The use of bar borders usually adds a visual component to the graph without adding information. Here are examples of some of the possible variations:

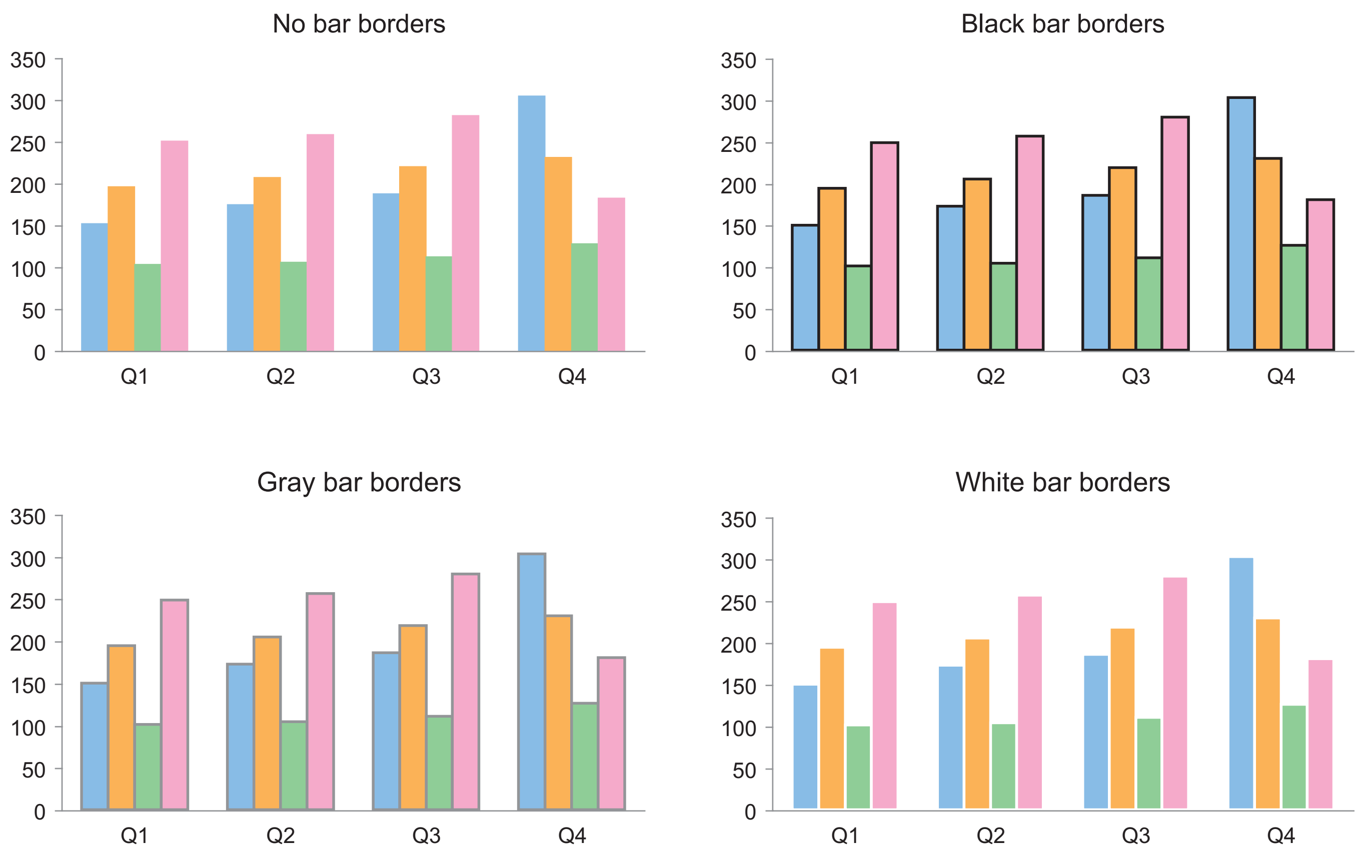

FIGURE 10.23 These examples show variations in the use of borders around bars.

Other than for highlighting, borders around bars are only useful when bar fill colors do not sufficiently stand out against the background (e.g., light yellow bars against a white background). If you must use light-colored bars for some reason, the use of subtle borders (e.g., gray rather than black) creates sufficient separation between the bar and the background to make the bars stand out.

Just like fill colors, borders may be intentionally introduced to make particular values or sets of values stand out from the others. In the following example, attention is clearly drawn to a particular bar through the use of a border that is absent from the other bars:

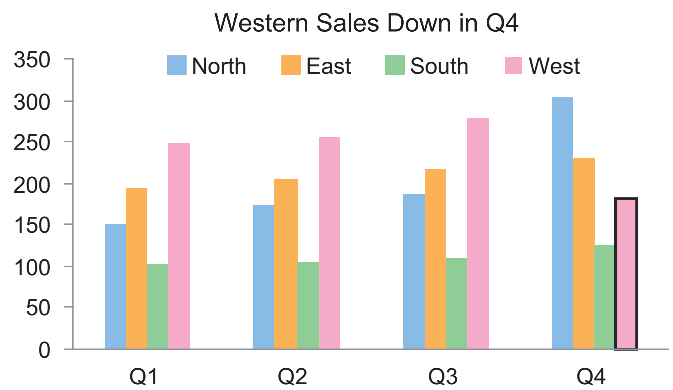

FIGURE 10.24 This example uses a border around a bar to highlight a particular value

BASE VALUE

A bar consists of two ends: the one that marks the value, called the *endpoint*, and the one that forms its beginning, called the *base*. The purpose of this section is to consider where bars should begin along a quantitative scale. Bars should begin at zero and extend from there. This is not the same as saying that bars should always begin at the bottom or left edge of the graph. The axis that the bars rest on does not intersect the other axis at its lowest quantitative value when both positive and negative values are included. In the following example, the X axis intersects the Y axis near the middle:

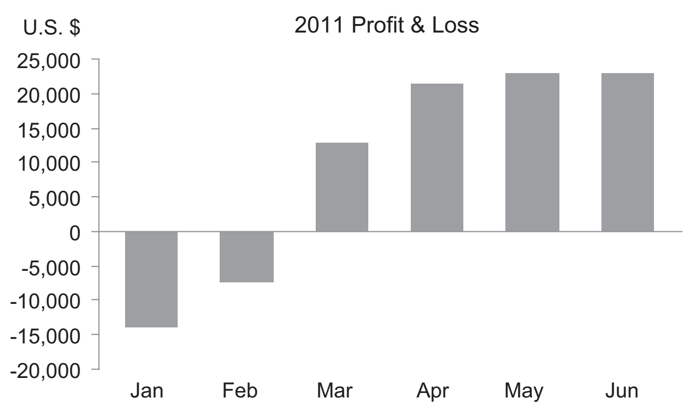

FIGURE 10.25 This is an example of a graph with an X axis that is not positioned at the end of the Y axis.

Whether bars extend upward or downward, to the right or to the left, they should start at zero. Otherwise, their lengths will not correspond to the values they encode.

Lines

In contrast to bars, which emphasize individual values, lines emphasize continuity and flow from one value to the next. Consequently, lines are particularly good at displaying values that change through time, featuring the overall shape of that change. With rare exceptions, lines should only be used to connect values along an interval scale (i.e., to display a time series or a distribution). In this section we'll examine ways to make lines easy to distinguish from one another and circumstances when points should be included on lines to mark the location of values.

DISTINGUISHING LINES

When a graph contains multiple sets of values, each encoded as a line, you must take care to make them visually distinct. Lines that look too much alike are hard to trace as they cross one another, as in the following example:

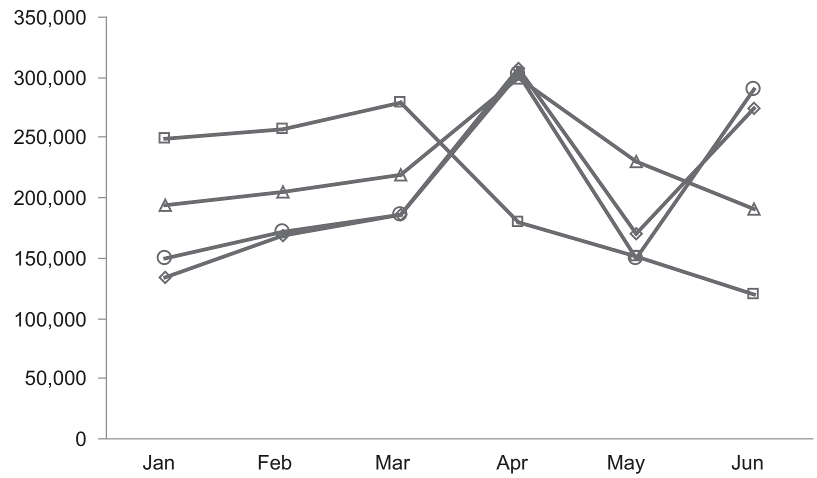

FIGURE 10.26 This graph shows lines that are not clearly distinct from one another.

In this example, the only visual differences that distinguish the lines are the distinct shapes of the points that mark the values. This distinction isn't sufficient, however, especially where lines intersect. Shapes that mark individual data points don't solve this problem because they only appear intermittently along the length of the line. Unless you want to simultaneously display the individual values as well as the overall shape of the data, you wouldn't bother to mark the individual values with points anyway, so you definitely shouldn't get into the habit of distinguishing lines by symbol shapes alone. Given what you know about the attributes of visual perception, how could you make the lines sufficiently distinct?

• • • • • • •

The most effective method uses differences in hue or color intensity. Here are two new versions of the same graph without distinct point shapes but with distinct colors instead:

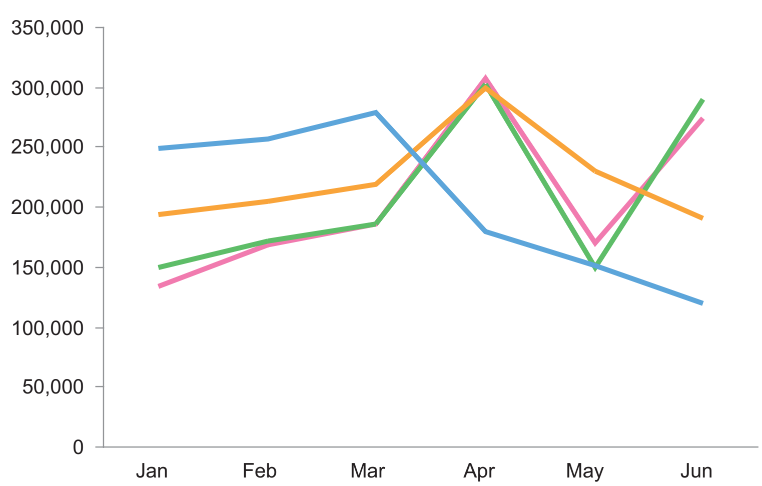

FIGURE 10.27 This graph shows hues used to make lines clearly distinct.

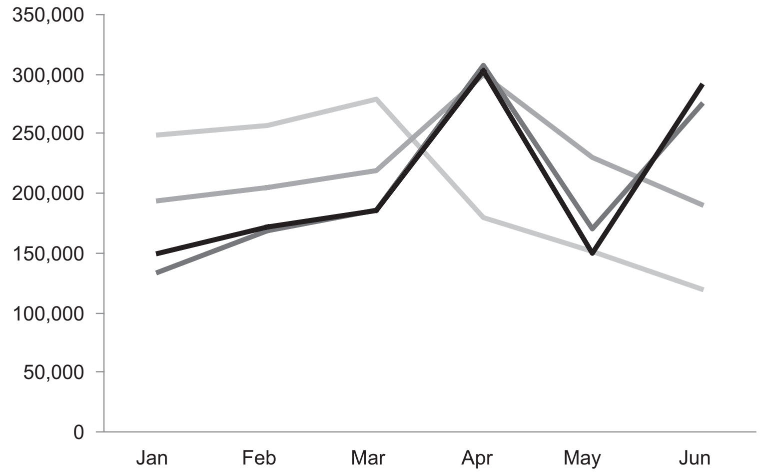

FIGURE 10.28 This graph shows different color intensities used to make lines distinct.

Hues work more effectively than color intensities, but if the graph isn't printed in color, differing shades of gray are more likely to remain distinct. One disadvantage of this approach is that the darker lines stand out more than those that are lighter, which suggests that the darker lines are more important that others. Sometimes this is an asset because it provides a way to display a ranking relationship. In the case above I've used the order of dark to light lines to rank the lines according to their values in the final month of June.

Although lines of significantly different shades of gray will usually look different when printed on a black and white printer, they will not necessarily look different or even be visible when copied on a black and white photocopier. When a gray line is scanned by a black and white photocopier, there is a tendency for light gray lines to become white and dark gray lines to become black. If you know your graph will be copied on a black and white photocopier, you might have to use black lines with one of the inferior encoding methods, such as differing symbol shapes or line styles.

Another option that's available for distinguishing lines, which is also useful when hues can't be used, is variations in line style, as in the following example:

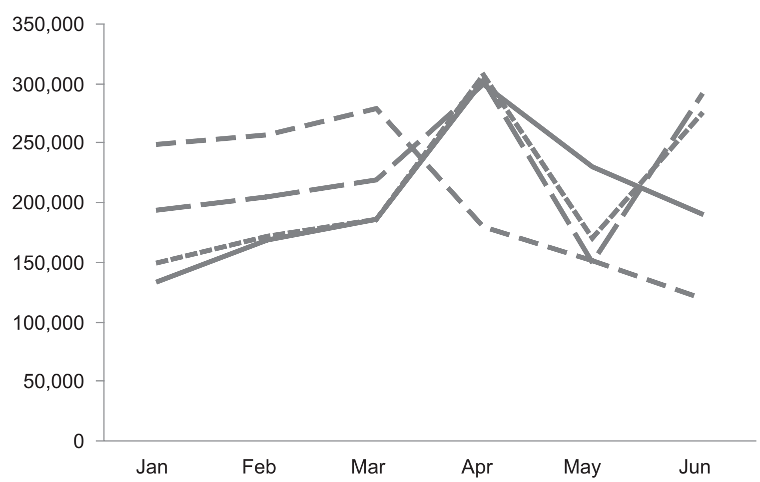

FIGURE 10.29 This graph shows line styles used in an attempt to make lines clearly distinct.

One significant drawback of this approach, as you can see, is that dashed styles break up the flow of the line, producing a jagged appearance that is not inviting to the eye. For this reason it is best to avoid varying line styles unless you can't use hue or color intensity.

POINTS OR NO POINTS

You have perhaps noticed that I don't always include data points to mark values along the line. The stories the lines tell in graphs are usually contained in the patterns that the lines form, not in precise values, so points aren't necessary. Nevertheless, there are times when it is useful to include points. The primary case is when the graph contains multiple lines that people will use to compare values at a particular point in time. For example, in the next figure, it would be difficult to know precisely where along the three lines you should pinpoint to compare values for the month of June.

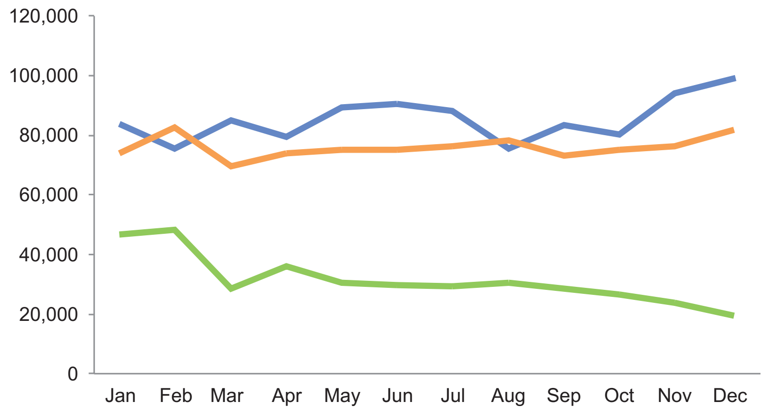

FIGURE 10.30 Without points, it is sometimes difficult to compare values on different lines at the same point in time.

With the addition of points, this task becomes easier because the points tell us precisely where to make the comparison, as you can see in the example below:

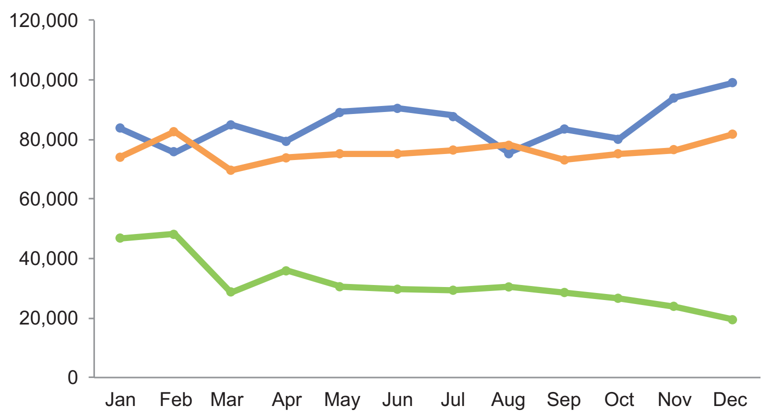

FIGURE 10.31 The addition of points makes it easier to compare values on different lines at the same point in time.

Notice that it isn't necessary to make the points large to do the job. In fact, it works best to keep them small to avoid clutter. Another way to solve this problem is to include light vertical grid lines in the graph to mark the position of each month, rather than points, as illustrated below:

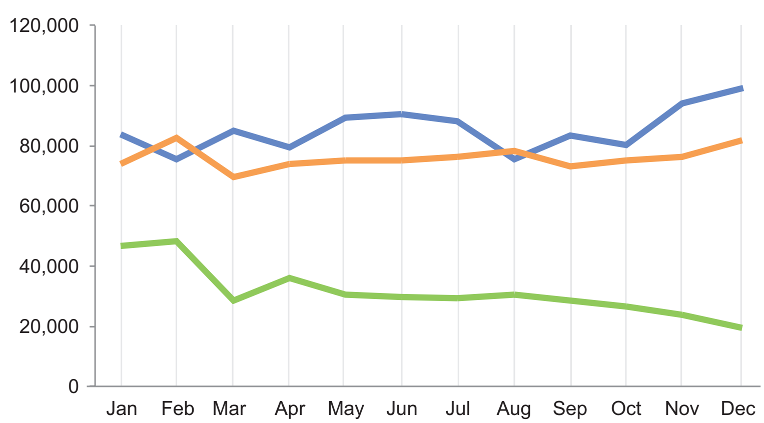

FIGURE 10.32 The addition of light vertical grid lines makes it easier to compare values on different lines at the same point in time.

Boxes

Most of the design practices that apply to bars apply to boxes as well because of the similarities between them.

- Boxes may be oriented either vertically or horizontally.
- Space between boxes that exceeds the width of a box is wasteful, making comparisons more difficult because of excessive distance between them, except when they're used in time series displays where the extra space between them is useful for showing the pattern of change through time. (I'll illustrate this in a moment.)
- When you have multiple sets of boxes in a graph, differing their fill colors works best to distinguish them.
- Borders around boxes aren't necessary unless the box's fill color does not stand out sufficiently against the background.

The one bar design practice that doesn't apply to boxes concerns the base. Bars must begin at a common base of zero on the quantitative scale, but boxes are different because they don't each display a single value as bars do. Instead, boxes show the distribution of a full set of values from low to high.

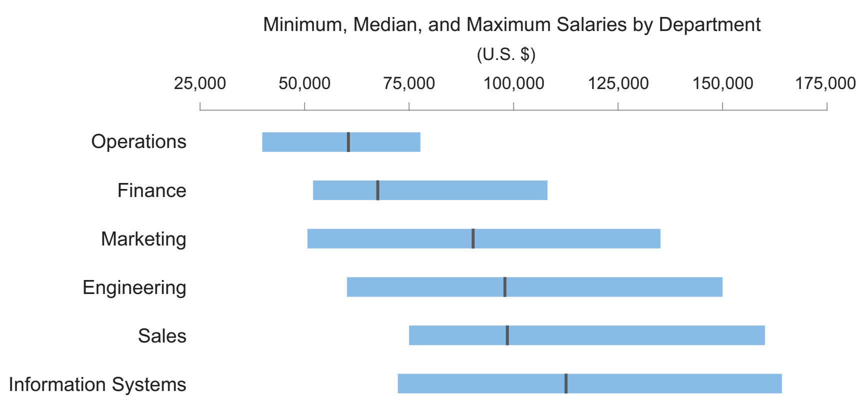

FIGURE 10.33 Unlike bars, boxes don't require that the quantitative scale begin at zero.

In this example, because the categorical labels are fairly long, I chose to use horizontal boxes, which allowed me to keep the boxes close together for easy comparison.

The two examples below both show the same series of distributions over time. Although the version on the left makes it slightly easier to compare individual distributions to one another by placing the boxes in close proximity, the increased width of the graph on the right allows it to do a better job of showing the pattern of how the mean selling price changed through time.

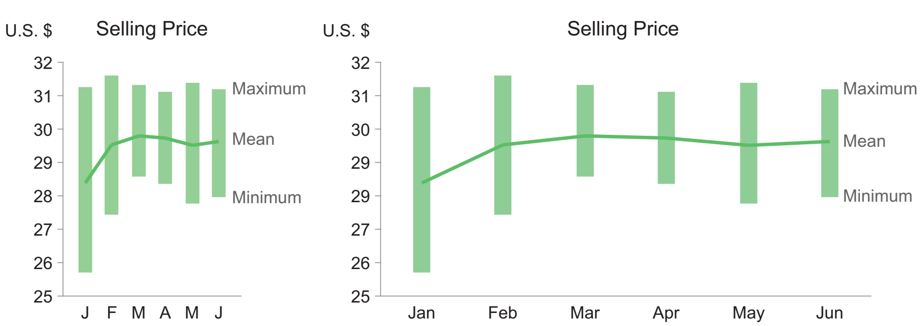

FIGURE 10.34 These box plots are designed to show how distributions change through time.

These two graphs were created in Excel using a line graph with high-low lines at each value, which is described in Appendix E, *Constructing Box Plots in Excel*.

As with most time-series displays, it usually works best to make the graph wider than it is tall to feature the pattern of change without exaggeration. Except when you are displaying distributions over time, keep boxes fairly close to one another so they can be easily compared.

Useful Combinations of Points, Bars, Lines, and Boxes

Points, bars, lines, and boxes are ordinarily used independently in graphs. They shouldn't be combined arbitrarily, but there are occasions when they collaborate in meaningful and effective ways. In the following example, the graph on the top combines bars and a line arbitrarily, which results in a solution that works less well than the one underneath:

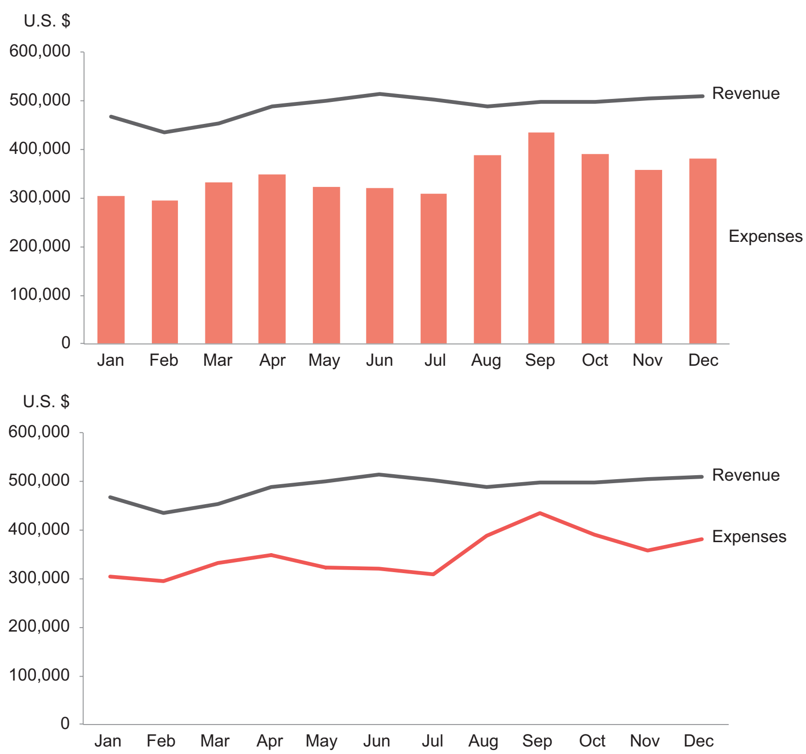

FIGURE 10.35 The top graph combines bars and a line arbitrarily whereas the one on the bottom uses two lines, which reveals patterns and supports comparisons better.

When displaying time-series data, we would only use bars if we wanted to feature values individually rather than as an entire series. For what possible reason would we want to feature the pattern of changing revenues by using a line but feature monthly expenses individually by using bars as shown in the upper graph above? Doing so not only makes it harder to see the pattern of changing revenues but also harder to see how the pattern of change in revenue compares to that of expenses. Choices between points, bars, lines, and boxes should never be arbitrary. We should always choose the means of display that can tell the story best. There are three occasions in particular when it is useful to combine different value-encoding objects:

- Combining boxes and lines to display distributions through time
- Combining bars and lines for ranked contributions to the whole
- Combining bars and points for uncluttered comparisons

BOXES AND LINES FOR DISTRIBUTIONS THROUGH TIME

The first of these combinations we've already seen. The final box plot example, *Figure 10.34*, used boxes to show individual distributions and a line to show how their center values (which in this case are means) changed through time. This combination was not arbitrary; the boxes and lines took advantage of the strength of each form of display—boxes for distributions and a line for the pattern of change through time—to tell a story that featured both.

BARS AND LINES FOR RANKED CONTRIBUTIONS TO THE WHOLE

While studying the distribution of wealth in Italy, a social scientist named Vilfredo Pareto discovered that approximately 20% of the population owned 80% of the wealth. From there, he further observed that these proportions describe many other aspects of society as well, which led him to propose the 80-20 rule. A special kind of graph that can be used to show how large amounts of something are associated with small portions of a population, called the *Pareto chart*, was named in his honor. Pareto charts combine ranking and part-to-whole relationships to tell the story of how the biggest parts of something combine to dominate the whole. In the following example, different channels through which people discover an organization are displayed in a Pareto chart.

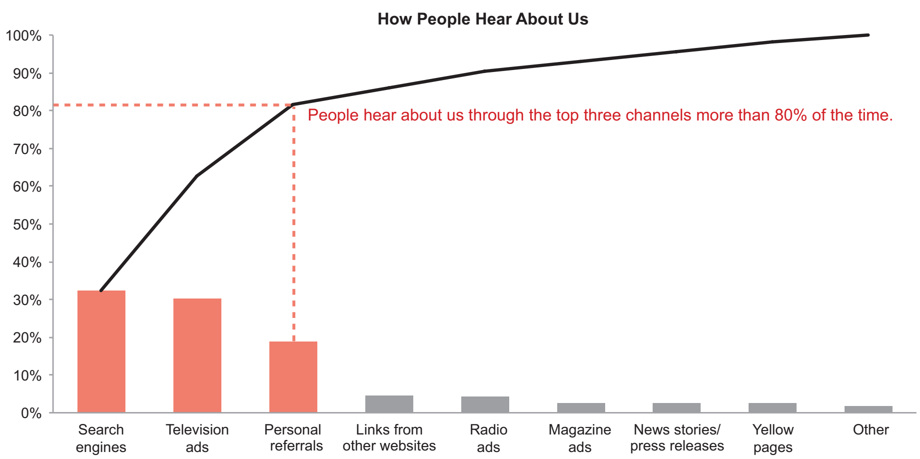

FIGURE 10.36 This Pareto chart presents, in ranked order, the channels through which people discover an organization.

A Pareto chart combines bars and a line in a useful way. The bars display a ranking relationship between parts of the whole from greatest to least, in this case channels through which people discover the organization. Bars make it possible to easily compare the contribution of individual channels. A line accumulates the same series of values to show the percentage of the whole contributed by subsets of values from high to low. A line makes sense in this case because there is an intimate connection from one value to the next as they combine to form the whole. This particular example was designed to tell the story that most people discover the organization through the three top channels. The clear message is that that the organization could increase its visibility by doing what it can to improve these three channels.

Pareto charts can be easily created in Excel simply by associating a column chart (vertical bars) with the series of individual values and a line chart with the series of cumulative values.

BARS AND POINTS FOR UNCLUTTERED COMPARISONS

Another combination that is often handy uses bars to feature a primary set of values and points to display one or more sets of values that you want to compare with the primary set. A good example involves a comparison of actual expenses to budgeted expenses, by department:

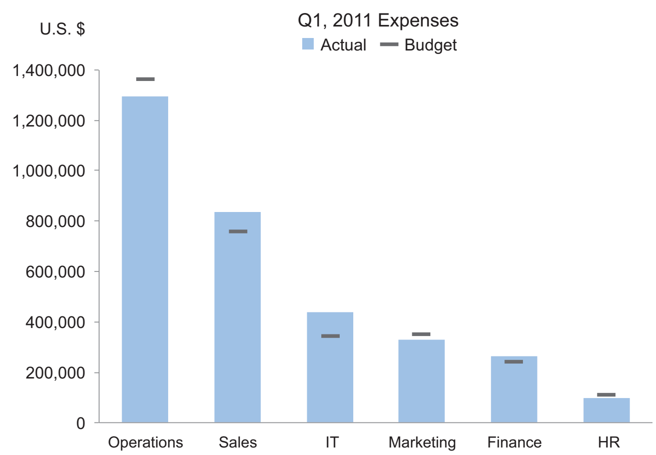

FIGURE 10.37 This graph uses bars to display a primary set of values and points to display comparative values.

In this case, actual expenses are featured, and the budget amounts are there merely to provide context in the form of a comparison that allows the viewer to assess performance. Rather than using bars for both sets of values, using bars for actual expenses only and points for budget values causes actual expenses to stand out clearly as the more important of the two. Another benefit of this design is that it is less cluttered and requires less space than a graph with two sets of side-by-side bars. Graphs of this type can be used for any comparison that features a single set of values but also provides other values for the purpose of comparison.

Graphs of this type can be created in Excel by using a column chart (vertical bars) for the primary set of values and a line chart for the other set. Display points along the line and then turn the line off, so only the points remain visible.

Secondary Data Component Design

In addition to the primary data objects that are used to encode values in graphs (points, bars, lines, and boxes), several other components of graphs also provide information. In this section, we'll look at ways to design each of the following components for maximum effect:

- Trend lines
- Reference lines
- Annotations
- Scales
- Tick marks
- Grid lines
- Legends

Trend Lines

A trend line describes the overall nature of a set of values, usually in a graph that displays correlations or time series. As such, trend lines do not represent actual values but instead summarize their essence. In this section we'll examine the role of trend lines that display time series and correlations.

TREND LINES IN TIME SERIES

The following time-series display includes a trend line to show the overall trend of change during a particular 12-month period:

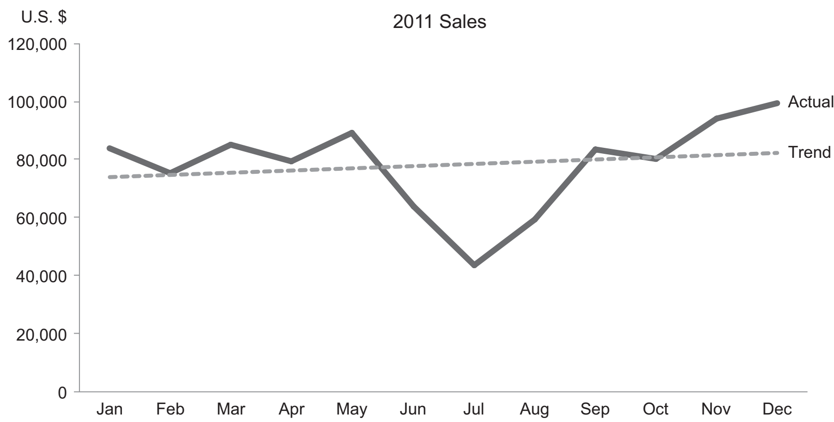

FIGURE 10.38 This graph shows a trend line that represents the overall pattern of sales during the year.

A trend line of this type is also called straight line of best fit or a linear trend line. It is an attempt to draw a straight line from left to right through the vertical center of the full set of values, keeping the line as close as possible to all values without curving it. You should use trend lines of this type in time-series displays with caution because they can be misleading. A trend line is based only on the values that you've included in the graph. Slightly changing the time period that you're showing might have a significant impact on the trend line, as you can see in this example where adding one more month to the beginning of the line graph has caused the trend line to reverse direction.

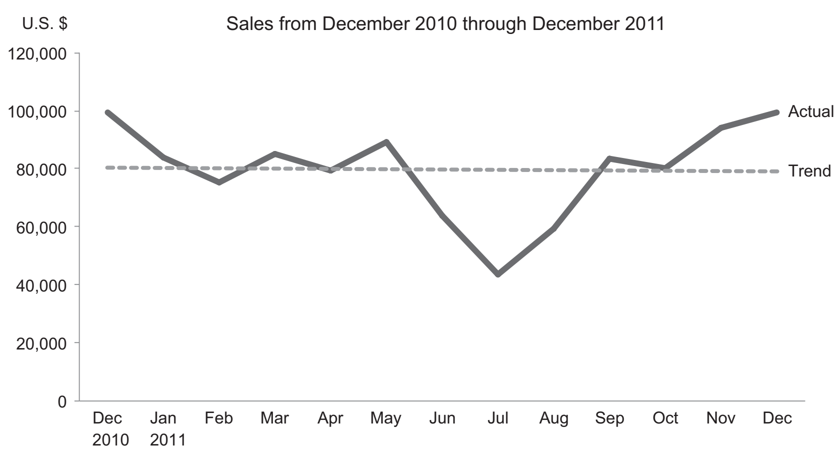

FIGURE 10.39 This trend that sloped upward in *Figure 10.38* now slopes downward simply because the month of December, 2010 was added.

Because of this potential for misrepresentation, it often makes sense to show the overall trend of change through time by using a *moving average* (a.k.a. *running average*). This approach softens the jaggedness that often occurs in change through time by replacing actual values with the average of several consecutive values. For example, a five-month moving average of monthly sales would show for each month the average of that month's sales and the four previous months.

Here's the same sales data as above, but in this case the trend is shown as a five-month moving average.

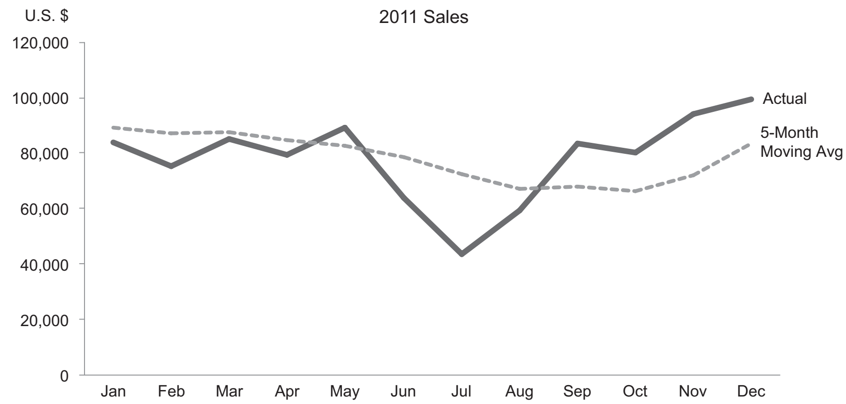

FIGURE 10.40 In this graph the overall trend is represented by a moving average.

Trend lines should always stand out as different so they aren't mistaken for actual values. In the three previous examples, I made the trend line lighter and dashed, to alert readers to its difference in meaning. It certainly isn't necessary to use dashes for a trend line although this often works well. In all examples so far, I've deemphasized the trend line relative to the line of actual values to make the latter stand out as more important. If the trend line is more important to the story that you're trying to tell in the graph, you should do the opposite to feature the trend above the actual values.

TREND LINES IN CORRELATIONS

Trend lines are often included in scatter plots to show the overall pattern of correlation. Here's a simple example:

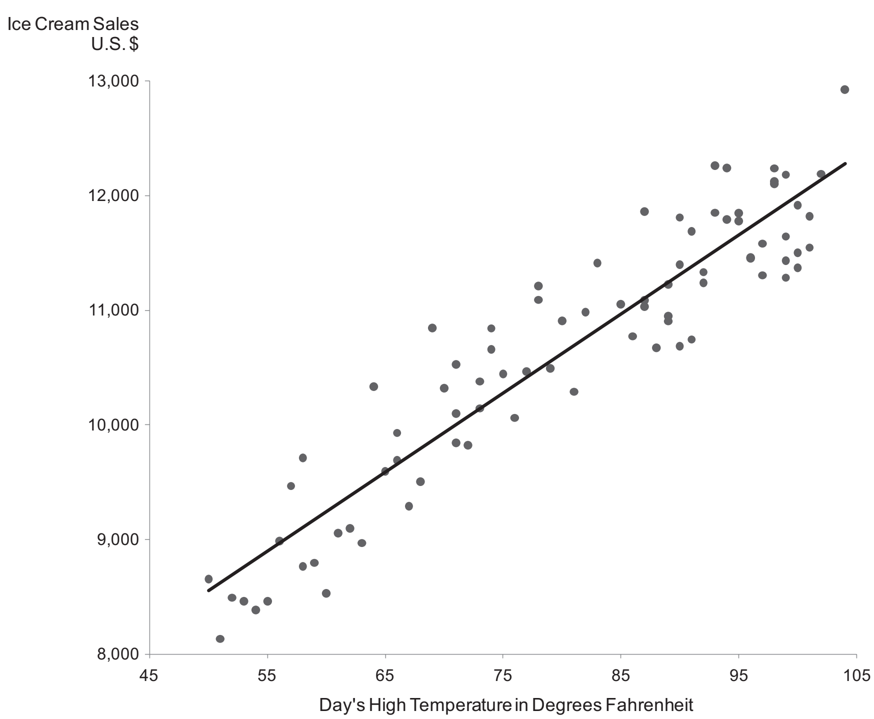

FIGURE 10.41 This scatter plot includes a trend line that highlights a correlation.

Trend lines do a great job of revealing the forest when your readers might otherwise get lost in the trees. Especially when you're dealing with a large number of erratic values, trend lines reveal what might be hard to see otherwise.

As with time-series trends, the overall patterns of correlations are frequently not linear. To do its job, a trend line should summarize the overall shape of the data. In the following scatter plot, the pattern formed by the values cannot be summarized by a straight line.

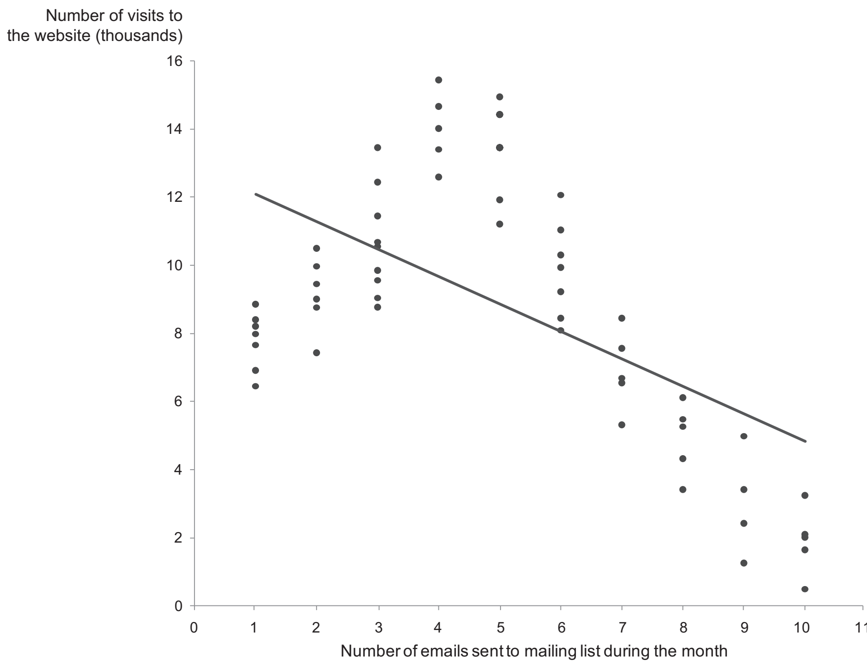

FIGURE 10.42 The trend line in this scatterplot does not match the pattern in the data.

In this case, beginning from the left the values slope upwards to a peak at 4 on the X axis and then turn downwards from there. A curved line such as the one shown below is needed to summarize the overall shape of the data.

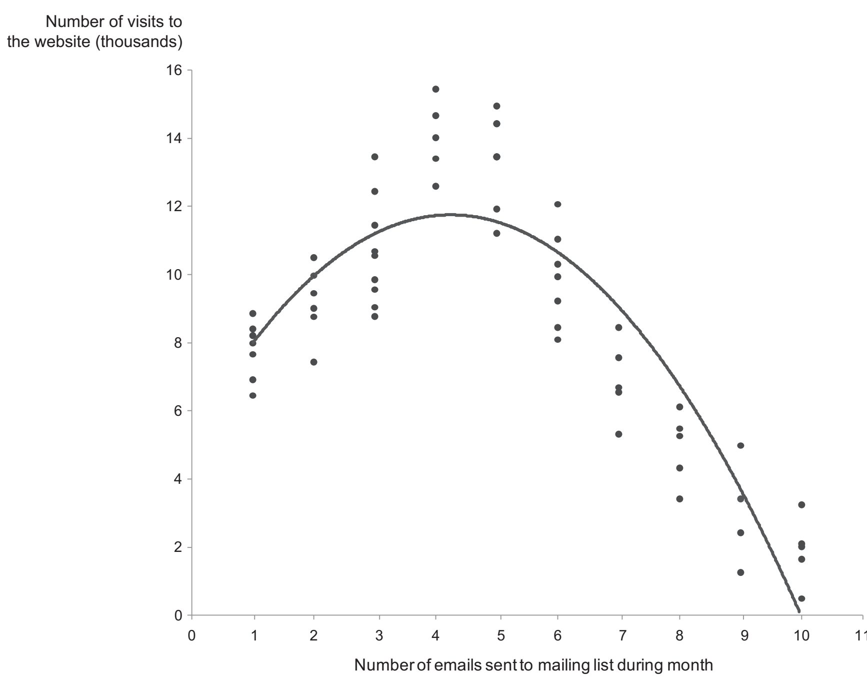

FIGURE 10.43 The trend line that describes this correlation curves to fit the data.

The story of this data set seems to be that people on the mailing list respond positively to emails by visiting the website but only up to a certain point—four emails in a month—and every additional email beyond that produces a decrease in website visitors.

The primary design practice to keep in mind for trend lines is the use of visual attributes such as hue, color intensity, line patterns, and line thicknesses to either highlight or subdue the lines, as needed.

Reference Lines

Reference lines are primarily used to provide context for the information in the form of a meaningful comparison. If you're displaying a year's worth of monthly sales revenues, you might want to include reference lines to show the previous year's lowest and highest monthly revenues, illustrated below.

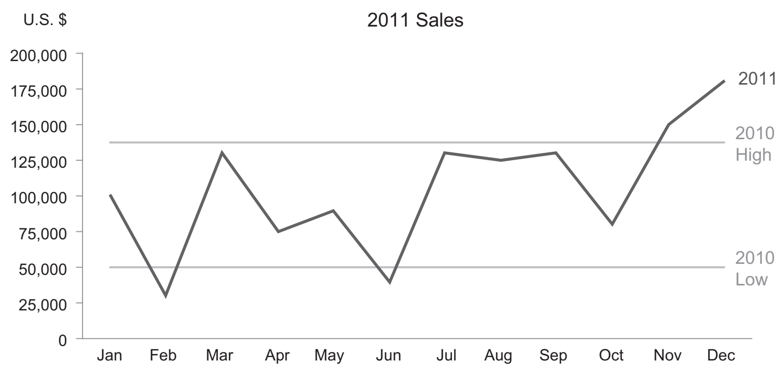

FIGURE 10.44 This graph includes reference lines to mark the highest and lowest monthly revenues during the previous year.

Reference lines are especially useful for displaying a measure of the norm, making it easy to see how values deviate from that norm. Averages (e.g., the mean or median) and standard deviations work well with reference lines. Here's an example of the same monthly sales values as in the graph above, but this time a reference line allows the viewer to compare them to average monthly sales:

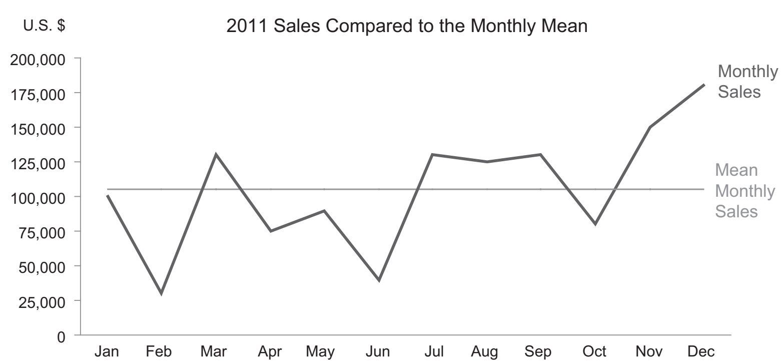

FIGURE 10.45 This graph shows a reference line that represents mean monthly sales to which each month's sales can be compared.

The next example is similar, but this time three reference lines are included to provide additional measures of the norm: one displays the mean, and two display one standard deviation above and below the mean.

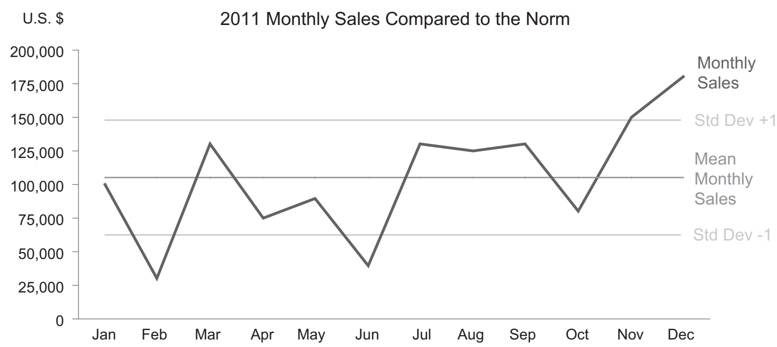

FIGURE 10.46 This graph shows reference lines that represent multiple measures of the norm.

Sometimes it works even better to mark entire reference regions. For instance, rather than using reference lines to mark standard deviations in the previous example, the entire range within one standard deviation above and below the mean could be displayed using fill color, such as in the example below.

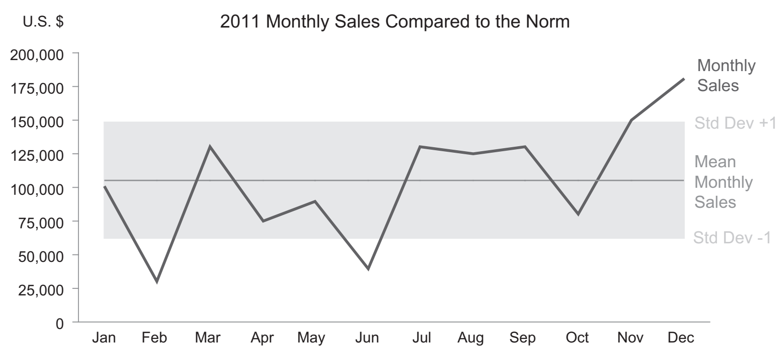

FIGURE 10.47 This graph includes reference lines in the form of a gray rectangle that marks a region of normal activity, in this case one standard deviation above and below the mean monthly sales.

If the purpose is to make values that fall outside the range of normal stand out, this does the job wonderfully, making it easy to distinguish between the normal range and everything else.

In addition to reference lines that mark particular quantitative thresholds, lines can also be used to separate categorical items into meaningful groups. In the following example, a line is used to separate products into two groups, hardware and software.

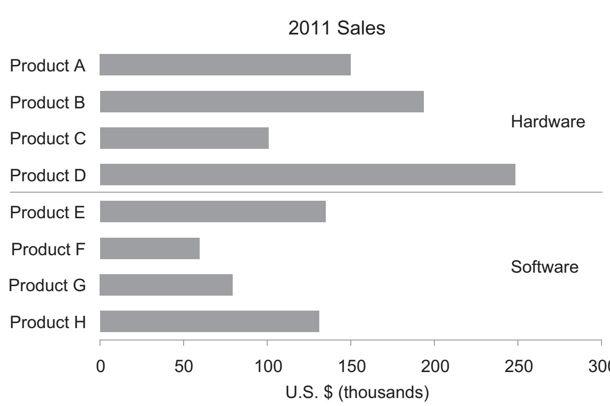

FIGURE 10.48 This graph shows a reference line that is used to mark a break point in a series of categorical subdivisions, in this case dividing products into two groups: hardware and software.

Notice that this example includes text annotations to label the two groups. We'll look more closely at annotations in a moment.

Reference lines and lines that are used to separate items into groups can be formatted in a number of ways, using various hues, styles (solid or dashed), and thicknesses, depending on the role that the lines play. As with trend lines, when you place a reference line in a line graph, you should be careful to make it look different from lines that encode the primary data. In addition, you must decide how much the line needs to stand out to do its job and then format it accordingly. In the previous example, to separate the bars into two groups, one for hardware and one for software, the line did not need to stand out very much, so it is relatively thin and light gray. If you used a reference line to mark the threshold of acceptable performance and the primary point of your story was to highlight values that exceed that threshold, you would make the line stand out to a degree that catches the reader's attention.

Annotations

Text appears in graphs in the title, the titles of and labels along axes, and labels in legends. It can also be used to annotate a graph in various ways. When words are needed to clarify meaning or to make a point, they belong in or near the graph. Regardless of its purpose or form, text should be tightly integrated with the graph but in a way that does not distract from important graphical elements. By convention, titles generally appear just above the data region of the graph, and notes generally appear just below the graph. As long as they don't disrupt perception of the visual information in the data region, notes may be placed in the data region, especially when they apply to particular data. Notes that are instructional or interpretive in nature usually don't need to appear in the data region and can be placed adjacent to the graph (below, above, or to either side), close enough that they are not overlooked. Here's an example of effective text placement:

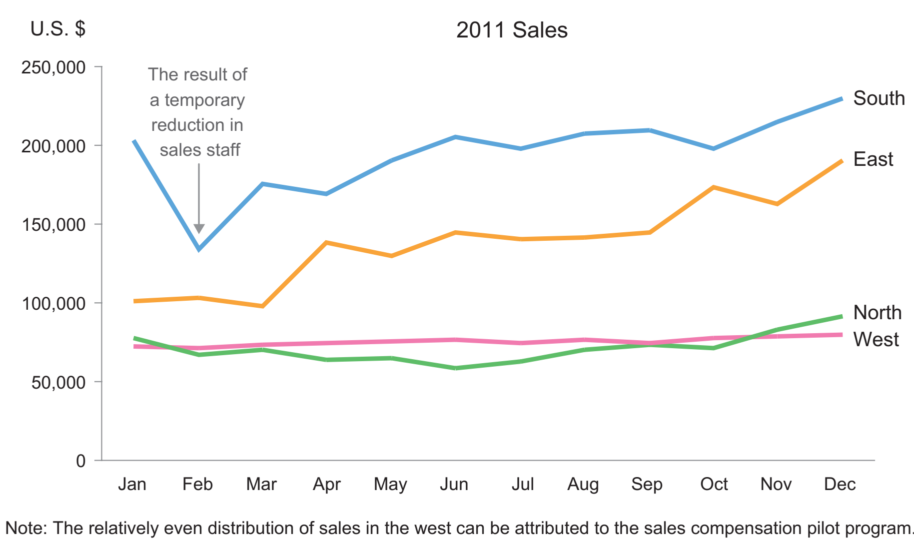

FIGURE 10.49 This graph Illustrates how annotations can be located near the point of reference without interfering with the data.

In the following example, annotations point out relevant events:

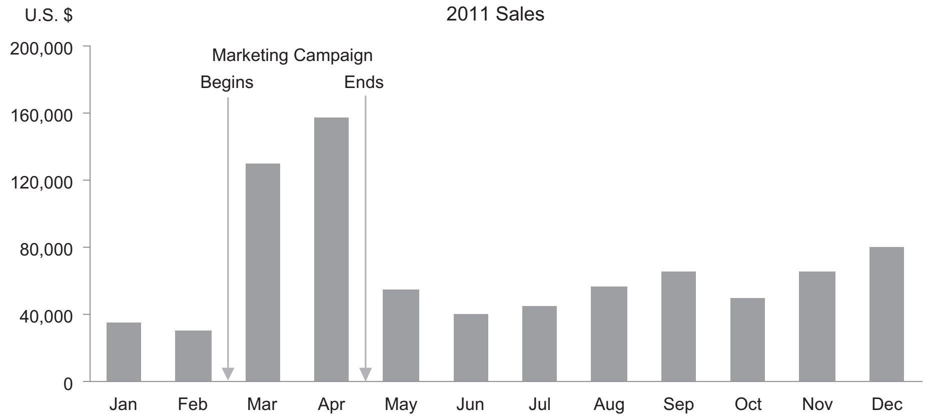

FIGURE 10.50 This graph shows annotations that identify significant events in time.

Another annotation that is often useful involves labeling values that deserve to be featured. In the next example, the highest values along each line have been labeled to draw attention to them.

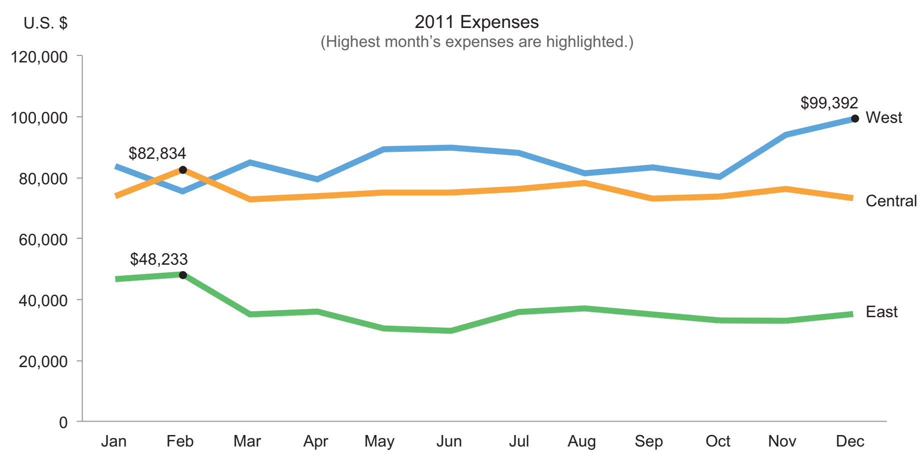

FIGURE 10.51 Particular values have been featured by labeling them.

As with reference lines, annotations should be tightly integrated in a way that complements other information without distracting from it. The prominence of annotations should be determined by their relative importance to the story that you're trying to tell. When placing annotations in the data region of a graph, it is often useful to make them inconspicuous to prevent the appearance of clutter. Reducing the intensity of text from black to gray can often accomplish this because it makes the text lighter than other data elements and thus prevents distraction, as illustrated in the example on the following page:

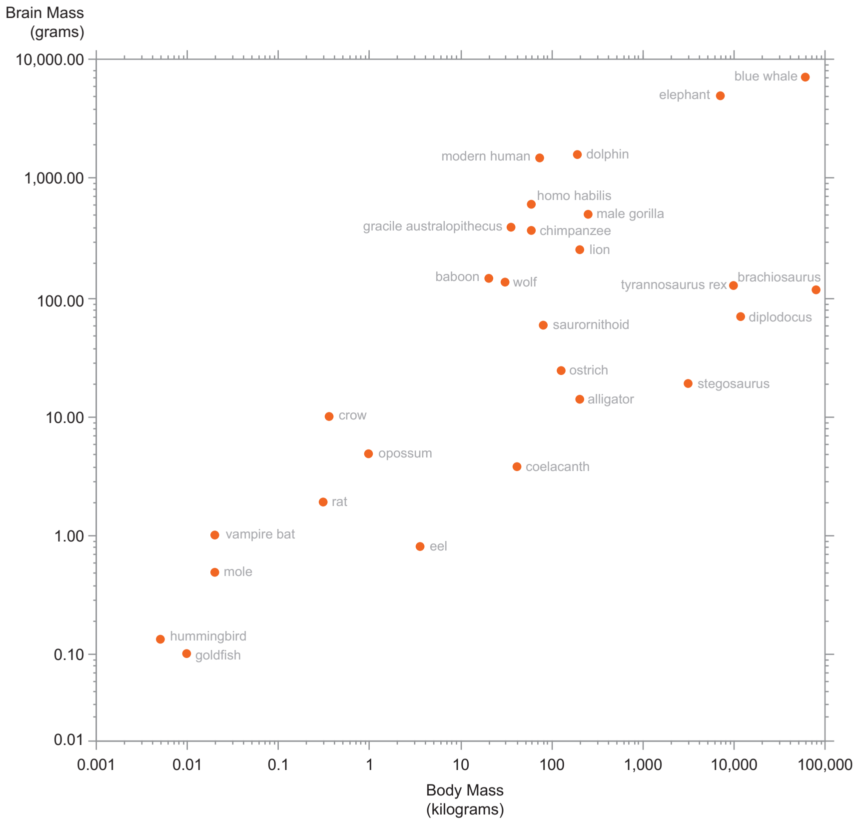

FIGURE 10.52 Values can be labeled in a way that does not distract from the values themselves. A version of this graph originally appeared in Carl Sagan's book *The Dragons of Eden: Speculations on the Evolution of Human Intelligence* (1977), Random House.

If your primary goal is to show the pattern of correlation between body mass and brain mass but you also want to allow readers to identify specific animals, this design works well. The labels can be easily read but can also be ignored when focusing on the data points because of the hue and color intensity differences between the labels and the data points.

Scales

Scales and axes are intimately related. A scale, regardless of the form it takes in the world at large (e.g., a ruler or a weight scale), measures the value of something. In graphs, scales are laid out along axes. They divide axes into increments of equal lengths and assign quantitative measures or categorical labels to those increments. Quantitative scales provide a means of assigning specific numeric values to the data objects that encode those values, based on the location of the objects along the scale line. Categorical scales associate categorical items with data objects in the same manner. The tiny lines that intersect scale lines to mark the increments are called *tick marks*; we'll cover them in the next section. Without scale lines on the axes, graphs could not tell their stories sufficiently.

Quantitative scales come in a variety of types, but graphs rarely require more than the following two:

- Linear scale
- Logarithmic scale

Along a *linear scale*—the type used in most graphs—the quantitative interval from one tick mark to the next is always the same. If the scale starts at 0, and the next tick mark as you go up the scale has a value of 10, then you know that the next tick mark after that will have a value of 20, and so on. To calculate the value of each tick mark, you simply add the value of a single interval to the value of the previous tick mark.

A *logarithmic* (abbreviated as *log*) *scale* works differently. Each value along a logarithmic scale increases by a particular rate rather than a particular amount. Another way of saying this is that each value is calculated by multiplying the previous value by a particular number, called the base. Here's an example of a log scale with a base of 10:

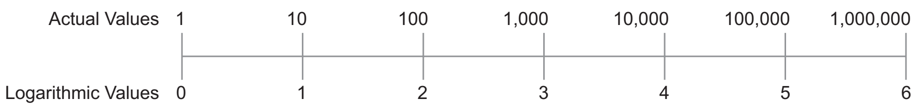

FIGURE 10.53 This illustration shows a log scale with a base of 10.

You might have noticed when looking at the scatter plot that correlated body mass to brain mass in *Figure 10.52* that the scales along both axes looked like this log base 10 scale.

What's the base 10 logarithm of 10,000? The answer is 4, since 10^4 (i.e., 10 to a power of 4, or 10 × 10 × 10 × 10) equals 10,000. The standard way of writing this is $\log_{10}(10{,}000) = 4$. Here's a concise definition:

> The logarithm of a number is the mathematical power to which another number, called the base, must be raised to equal that number.

In a log scale, then, actual values from one interval to the next are equal to the previous value multiplied by the base. For example, in *Figure 10.53*, the actual value 1,000 is equal to the previous value, 100, multiplied by the base of 10.

If you couldn't define it before, now you can, but what good is it? When are log scales useful in graphs? Take a look at the following scatter plot, which displays the correlation between income and lifespan for various countries. The lifespan scale on the Y axis is linear, ranging from an average lifespan per person of 40 to 85 years, but the income scale in international dollars on the X axis is logarithmic, ranging from an average annual income of $100 to $100,000 dollars per year.

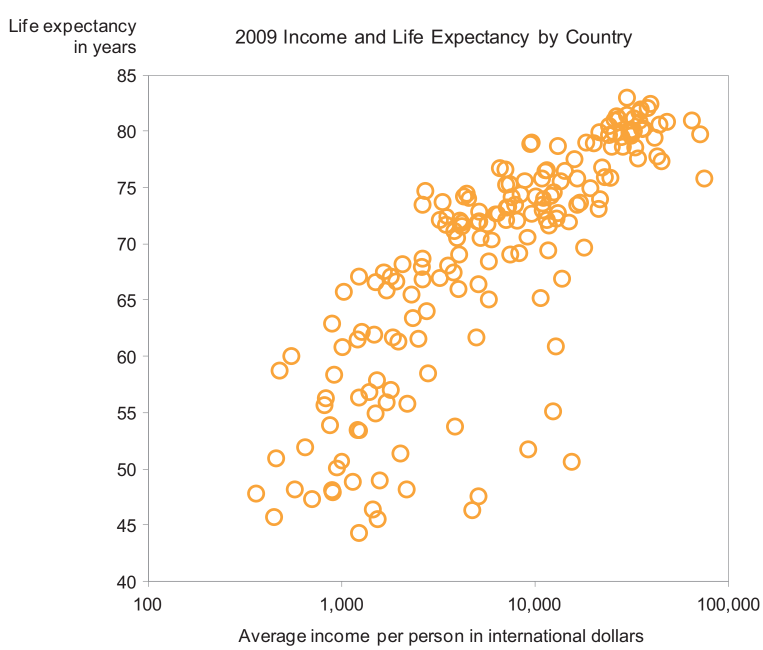

FIGURE 10.54 This scatter plot has a log scale on the X axis.

Why wouldn't this graph work well with a linear scale on the X axis? Of the 179 countries in this scatter plot, the average per-person incomes in 16 are less than $1,000, and in 76 they are less than $5,000. In a scatter plot with a linear scale, the data would look like this:

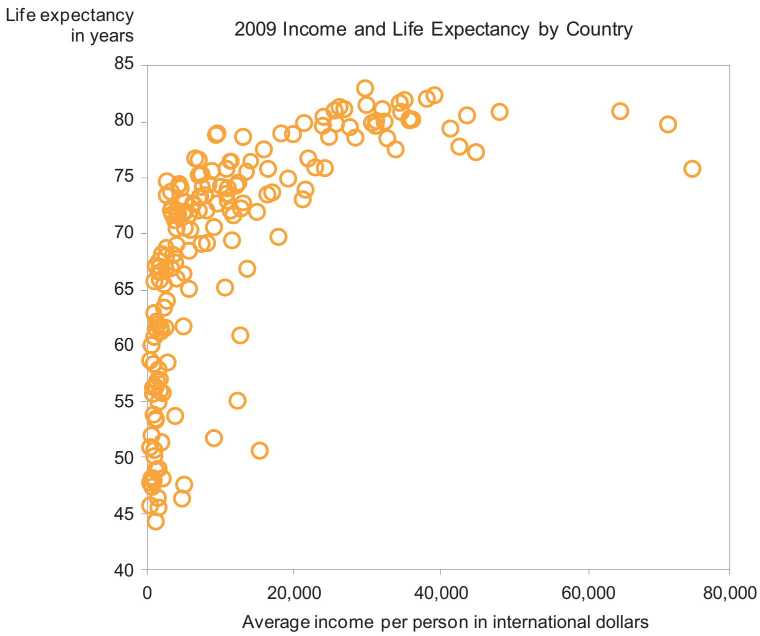

FIGURE 10.55 This scatter plot has a linear scale on the X axis.

For many purposes, this would work fine, but if you wanted to differentiate countries with low incomes, you wouldn't want them buried in a clutter of overlapping points. A log scale is a handy alternative on these occasions, but only if readers don't need to compare values based on distances between them. Here's another example of this situation:

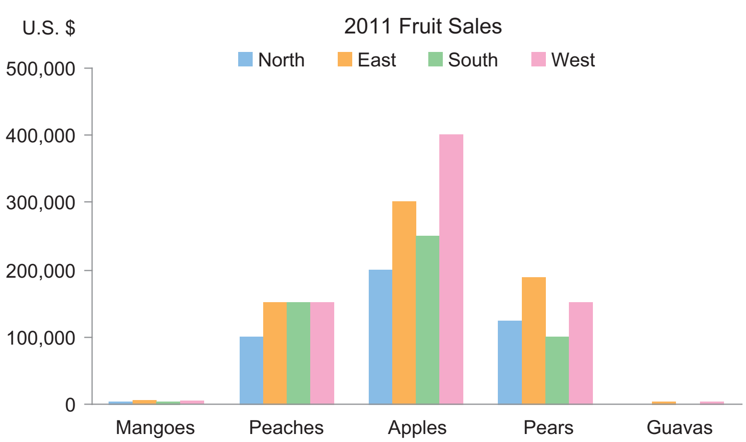

FIGURE 10.56 Several values on this graph are difficult to read because they are extremely small compared to others.

The huge difference between sales of apples at the high end and guavas and mangoes at the low end makes it difficult to interpret the low values along this linear scale. You could try to resolve this by making the graph much taller, but the result would look ridiculous. Log scales can provide an easy solution to this problem. Here are the same values, this time displayed as a dot plot with a log scale:

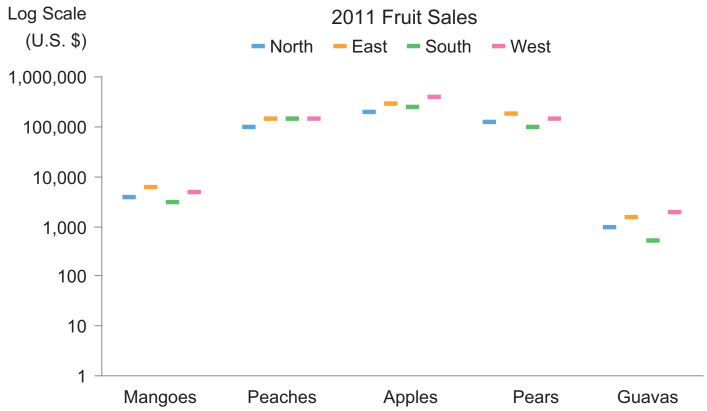

FIGURE 10.57 This graph uses a log scale to allow disparate values to be viewed together.

I substituted points for bars in this example because using the lengths of bars to represent logarithmic rather than linear values can lead to a great deal of confusion. It's hard to overcome the natural tendency to perceive a bar that is twice the length of another as representing twice the value of the other bar. Even when points are used, people can still be easily confused by a log scale, which we'll consider more carefully in a moment. For now, notice that the huge gap along the scale between the different sets of values has been dramatically reduced so that we can now see all of them more clearly and compare them more easily. Notice, however, that some care has been taken in this example to alert the reader to the fact that a log scale is being used. This is important, because readers would misinterpret differences in sales between various fruits, such as guavas and apples, if they failed to take the log scale into account and adjust their perception of differences accordingly.

With a linear scale, you know that if one value is twice as high as another, assuming the scale begins at zero, one value is twice the other. This is clearly not the case with log scales, but there is a unit of measure that will yield the same difference in value between equal differences in position along a log scale. Care to make a guess? Take a moment to think about it.

• • • • • • •

The same distance anywhere along a log scale equals the same *percentage* or *ratio*. In the next example, I've removed two of the regions, leaving only the north and east, and have added short black vertical lines to feature the distance between sales in the north and east for each fruit.

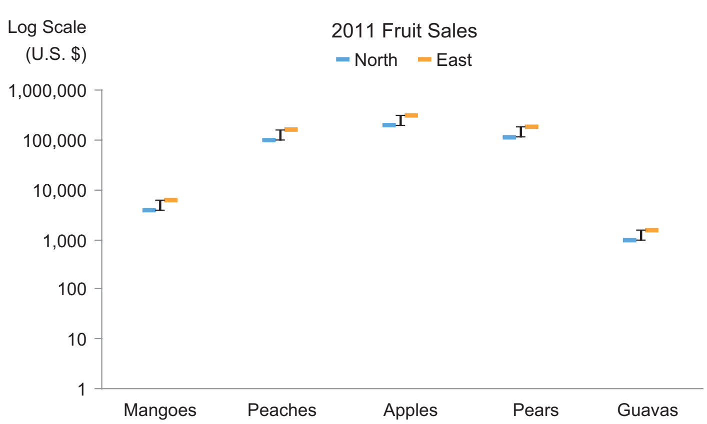

FIGURE 10.58 This example shows that, on logarithmic scales, equal distances correspond to equal percentages.

When I created this data set, I made sales values in the east precisely 50% greater than sales in the north for each fruit. For instance, the comparison of apple sales between the north and the east is $200,000 to $300,000, and for apples and guavas it is $1,000 to $1,500; in both cases a difference of 50%. The fact that the percentage difference between sales in the north and east for all fruits is the same is something that we couldn't see if the graph had a linear scale.

This feature of a log scale can be put to use, can't it? Whenever you want to compare differences in values as a ratio or percentage, log scales can do the job nicely. They function in this way because, whereas equal distances along a linear scale are equal in value, equal distances along a log scale are equal in ratio. In *Figure 10.58* on the previous page, beginning with the value of 1 at the bottom of the scale, each successive value is precisely 10 times the previous value. This is especially useful when displaying time series if you want to make it possible for your readers to compare rates of change through time. In the next example, the line graph has a linear scale, which makes it difficult to compare rates of change in sales from quarter to quarter among fruits and nuts.

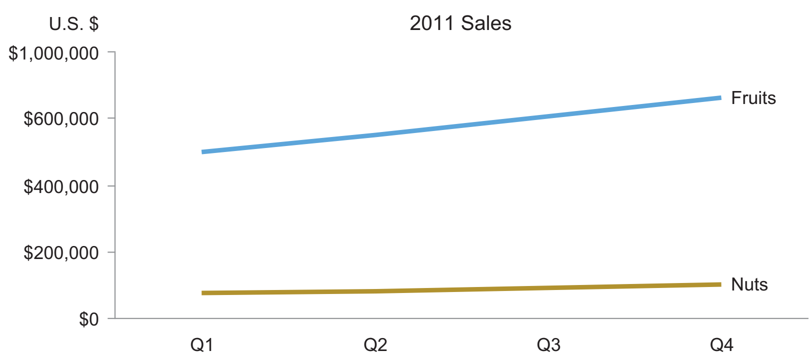

FIGURE 10.59 In this graph, time-series information is displayed using a linear scale. With a linear scale, the rate of change through time is difficult to discern.

In this next example, the same values are displayed using a log scale:

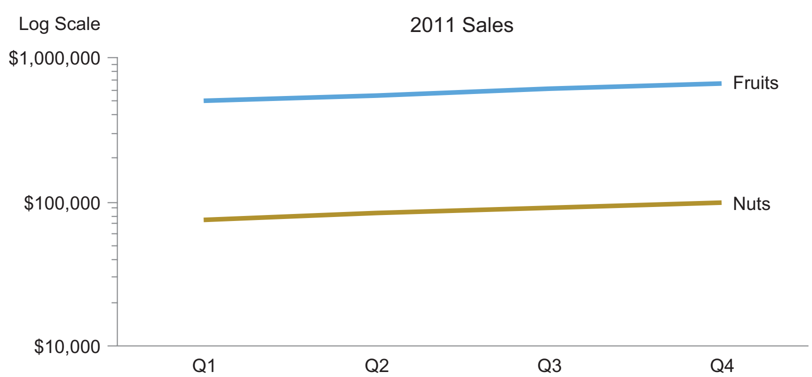

FIGURE 10.60 In this graph, the same time-series information shown in *Figure 10.59* above is displayed using a log scale to make rates of change easy to compare.

Notice that the slopes of the two lines are precisely the same. When you use lines to encode time-series values along a log scale, equal slopes indicate equal rates of change. In fact, sales of fruits and nuts both increased by exactly 10% each quarter. When the rate of change (rather than change in quantity) is the heart of your message, a log scale will serve you well. Just be sure to explain how they work to your readers because it's not at all obvious.

Despite their usefulness, the down side of log scales is the fact that they are not visually intuitive. Equal distances along the scale correspond to equal log values (ratios) but not equal linear values (amounts). Note that the minor tick marks (i.e., the smaller ones that aren't labeled) along the log scale in the last figure are not equally spaced. Each minor tick mark represents an additional increment of the value represented by the labeled tick mark below it. The first minor tick mark above the one labeled $10,000 represents $20,000 (i.e., $10,000 + $10,000). The one above that represents $30,000, and so on up the scale until you reach $100,000; above $100,000, each minor tick mark represents an additional $100,000. By displaying the minor tick marks, you can give your readers a visual clue that the scale is not linear, but this isn't enough if they aren't accustomed to log scales. If you wish to use log scales in the ways that I've described, and your readers don't already understand how to read them, you must take the time to provide a little education. You can head off inevitable confusion by including notes of explanation next to the graph where they can't be missed.

One more point to understand about log scales is that they don't have to use a base of 10. The base can be any number, but a base of 10 is the most common. In fact, some software restricts you to a base of 10. This is unfortunate, for log scales work in exactly the same way and offer exactly the same advantages with any base, and sometimes base 10 is too big to provide the best spread of values in a graph.

Base 2 is every bit as useful as base 10. With a base of 2, every major tick mark on the scale is exactly twice the actual value of the one before it (2, 4, 8, 16, 32, 64, 128, etc.). A base-2 log scale makes it easy to interpret rates of change in terms of doubling. With a time-series display in a line graph, a line that ascends upwards during an interval of time (e.g., month or year) from one major tick mark to the next indicates that the value doubled during that period. The other advantage of a base-2 log scale is that if the values you are graphing are not spread across a broad range, this scale allows you to display greater detail than a log 10 scale that jumps from 10,000 to 100,000 to 1,000,000 in single bounds.

Tick Marks

The information contained in a scale line consists solely of the tick marks, which establish the positions of values along the scale, and the labels, which assign numbers to the tick marks. The line on which the tick marks reside encodes no information itself. We'll explore the following questions about tick marks:

- How visible should tick marks be?
- Where should tick marks appear on an axis?
- When can you eliminate tick marks?
- When should you use minor tick marks?
- How many tick marks should you use?
- Which values should you indicate with tick marks?

HOW VISIBLE SHOULD TICK MARKS BE?

Tick marks and labels provide critical information because they allow us to interpret the data values in a graph, but they are not the values themselves. Consequently, they should be visually subdued in comparison to the actual data values in the graph but prominent enough to be read easily. Thin, light gray lines work well as a default.

WHERE SHOULD TICK MARKS APPEAR ON AN AXIS?

You have three options for the location of tick marks: the inner side of the axis, the outer side of the axis, or across the axis. The following example illustrates these options:

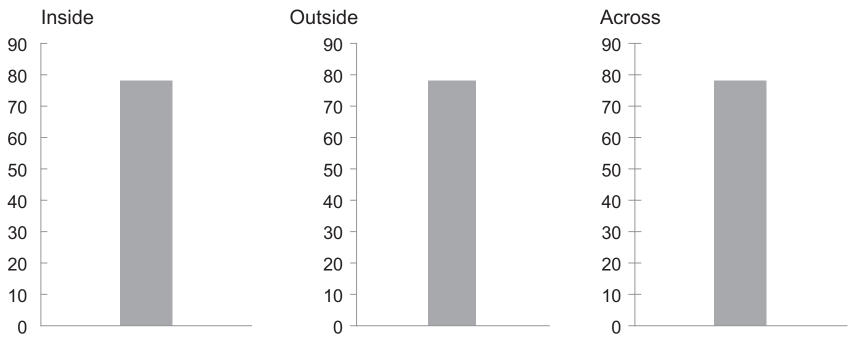

FIGURE 10.61 These examples show the three options for placing tick marks on an axis: inside, outside, and across.

As long as the tick marks are visually subdued in relation to the data values, any of these options will work, but I prefer placing them on the outside because this leaves the data region of the graph completely free of all but the values themselves, providing a clean, uncluttered backdrop.

WHEN CAN YOU ELIMINATE TICK MARKS?

Tick marks are always needed for quantitative scales, but what about categorical scales? In all but one case, tick marks are superfluous on categorical scales. The categorical labels themselves identify the locations of the values clearly enough without assistance from tick marks. Take a look at the two bar graphs below, one with categorical tick marks and one without:

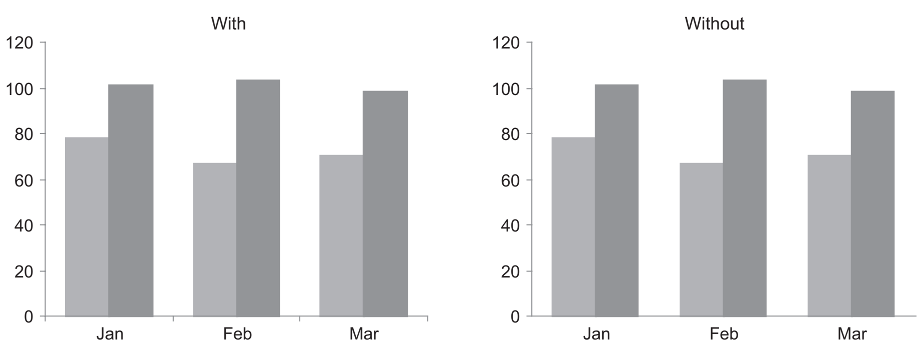

FIGURE 10.62 Tick marks appear along the categorical scale on the left but not on the right.

As you can see, tick marks along the categorical scale of a bar graph aren't needed. The same is true of graphs that encode values using points or boxes. Even though tick marks are small, if you eliminate them when they aren't needed, you'll produce cleaner, less cluttered graphs.

The one case in which tick marks can be slightly useful on a categorical scale is in line graphs that don't already mark the positions of values along the lines (e.g., by using points or grid lines). But only use tick marks when knowing the precise location of values is necessary. In the following example, notice how difficult it is to spot precisely where along the lines the value for February is positioned because the lines are relatively straight without an angle to mark the position:

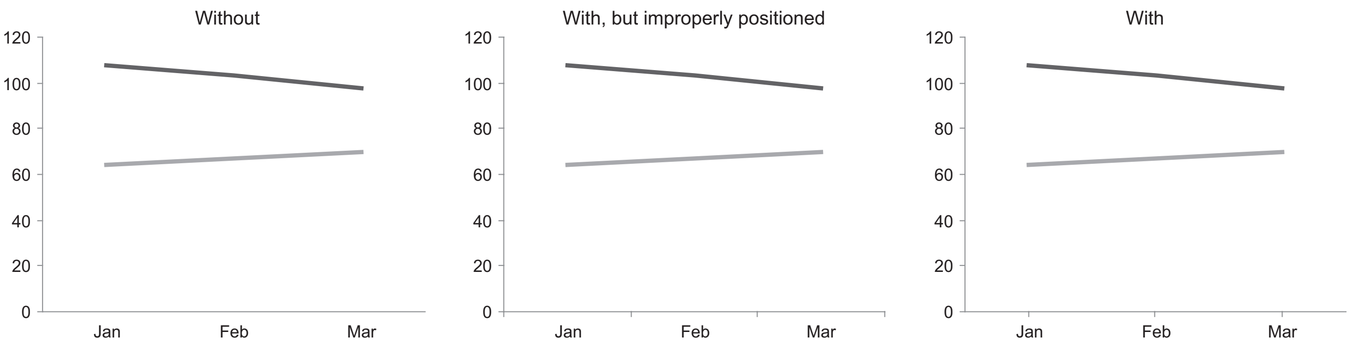

FIGURE 10.63 Three lines graphs: one without tick marks on the categorical scale, one with them but in the wrong position, and one with them in the correct position.

I included the middle example of a line graph with tick marks that are incorrectly positioned because this is a common default of software, including Excel. Even when they are correctly positioned, as on the right, tick marks are sometimes far from the lines, which makes it only slightly easier to determine the exact position of values. For this reason, if the precise location of values along the lines is really needed, it usually works best to eliminate the tick marks and either replace them with light vertical grid lines or to include points to mark values on the lines.

WHEN SHOULD YOU USE MINOR TICK MARKS?

What about minor tick marks, the ones that can be included between the labeled tick marks to indicate finer increments along the scale line? Minor tick marks suggest a level of quantitative precision that graphs just aren't meant to provide. Minor tick marks are only useful on log scales, not for the purpose of greater precision, but to alert readers that a linear scale isn't being used.

HOW MANY TICK MARKS SHOULD YOU USE?

Another consideration is how many tick marks you ought to include on the scale line. Always aim for a balance between having too many, resulting in clutter, and too few, which would make it hard for your readers to determine the values of data objects that fall between them. Here are examples of both extremes:

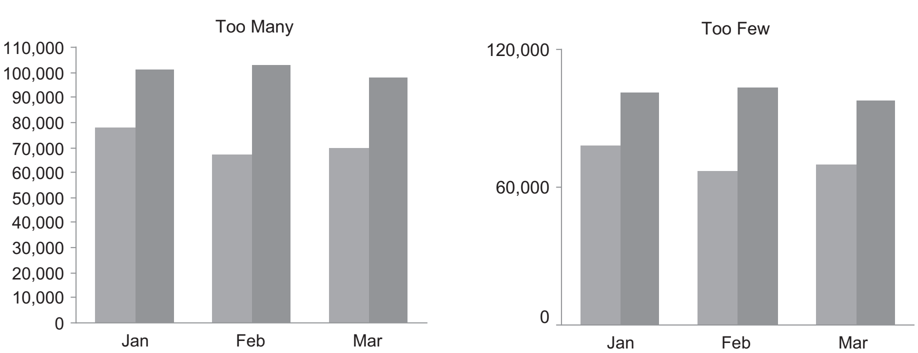

FIGURE 10.64 These are examples of graphs with too many and too few tick marks.

There is no exact number that works best in all circumstances, and the size of the graph is a factor that you must consider: the longer the scale line, the more tick marks it should contain. The following example shows a more appropriate number of tick marks than the two examples above, for the given scale length.

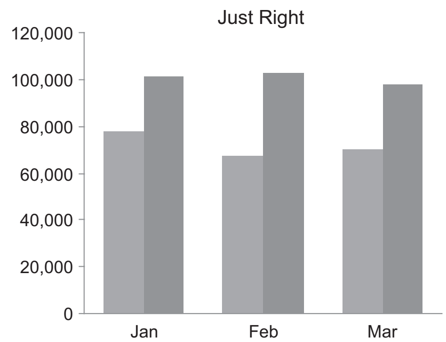

FIGURE 10.65 This graph shows an appropriate number of tick marks.

WHICH VALUES SHOULD YOU MARK WITH TICK MARKS?

Our final consideration regarding tick marks is the choice of values that they designate. It generally works best to use nice round numbers to which your readers can easily relate. Avoid irregular values like the following:

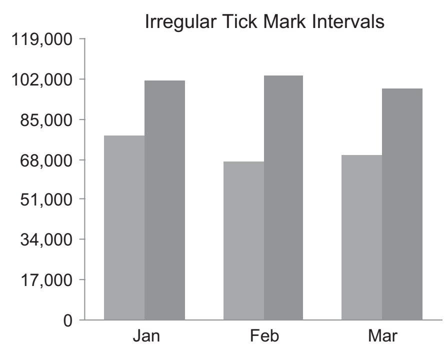

FIGURE 10.66 This graph shows tick marks that label values at peculiar increments, making them difficult to interpret.

If you use numbers like those above, you'll make your readers work extra hard and lead them to wonder why the increments are so peculiar, searching for meaning where none exists.

Grid Lines

From the early days of quantitative graphs, one of the primary purposes of grid lines was to make it easier to draw graphs. The lines were preprinted on graph paper as guides to assist the placing and sizing of objects. Now that almost all graphs are computer generated, this purpose is obsolete. Today, grid lines can be used to assist graph perception in three ways:

- Ease look-up of values
- Ease comparison of values
- Ease perception and comparison of localized patterns

In all cases, grid lines are support components, assisting in the visual perception of the data values themselves. Thus, grid lines should always be visually subdued. Thin, light lines (e.g., light gray) usually work best. Dark or heavy grid lines, such as those in the example below, should always be avoided because they disrupt perception of the data and produce annoying clutter.

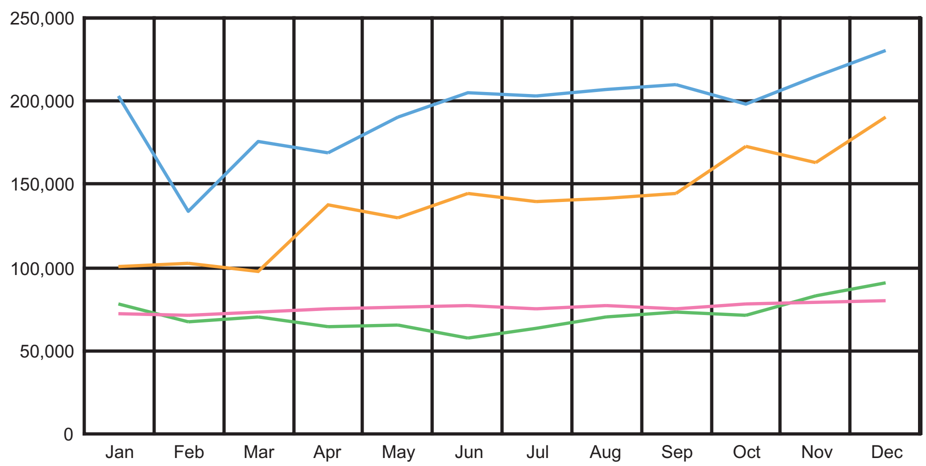

FIGURE 10.67 This graph includes dark, heavy grid lines, which you should always avoid.

Notice how much easier it is to focus on the lines of data in the example below.

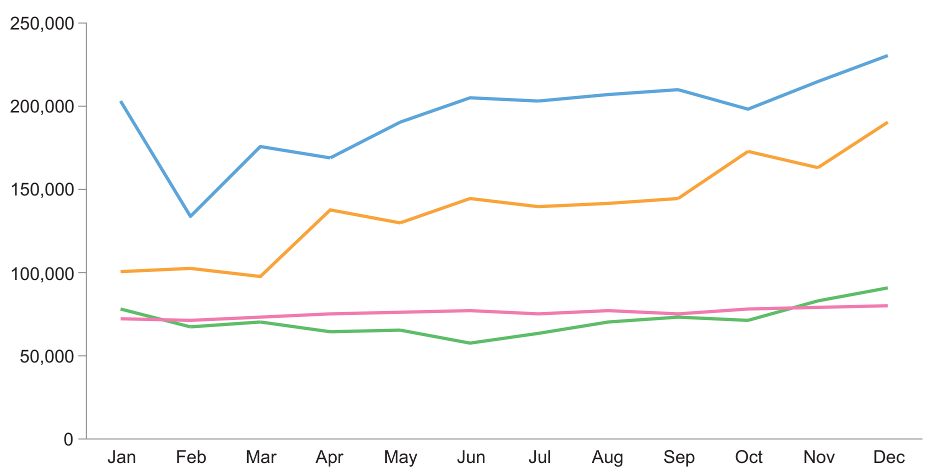

FIGURE 10.68 Without the dark grid lines, it is much easier to focus on the data.

ENHANCE THE LOOK-UP OF VALUES

The most common but least necessary use of grid lines is to assist readers in looking up values. We established previously that tables work much better than graphs for looking up precise values. Because graphs are used primarily to present the shape of the data, precise values usually aren't necessary. Nevertheless, grid lines do help when tick marks alone do not provide the level of precision that your readers would find useful for interpreting data values in graphs. This situation is most common with graphs that are especially wide, which makes it difficult to assign values to data objects that reside far from the scale line. In the following example, subtle grid lines provide all the look-up assistance that is necessary.

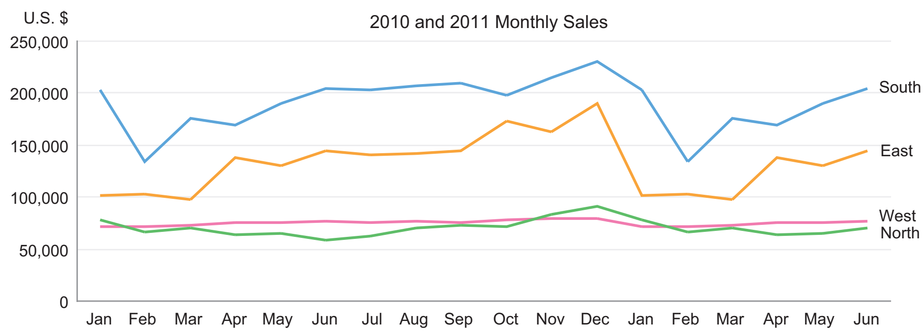

FIGURE 10.69 This graph uses subtle grid lines to enhance look-up.

ENHANCE THE COMPARISON OF VALUES

When perception of subtle differences is necessary, grid lines help by dividing the data region into smaller units, which makes otherwise subtle differences in size and location obvious. Grid lines allow us to compare differences in the size and position of objects relative to a much smaller space than the entire data region of the graph, which significantly increases the percentage differences between objects, making these differences easier to see. The following examples illustrate this point:

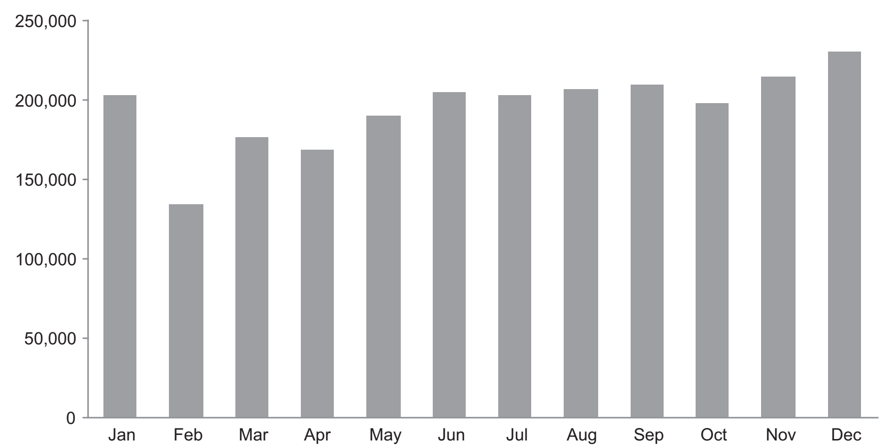

FIGURE 10.70 This graph demonstrates how difficult it is to perceive differences in the size and location of data objects when those differences are relatively small.

It's difficult to assess differences between bars that are close in length, such as January, June, and August. However, with the addition of light horizontal grid lines in the example below, these subtle differences become much easier to discern:

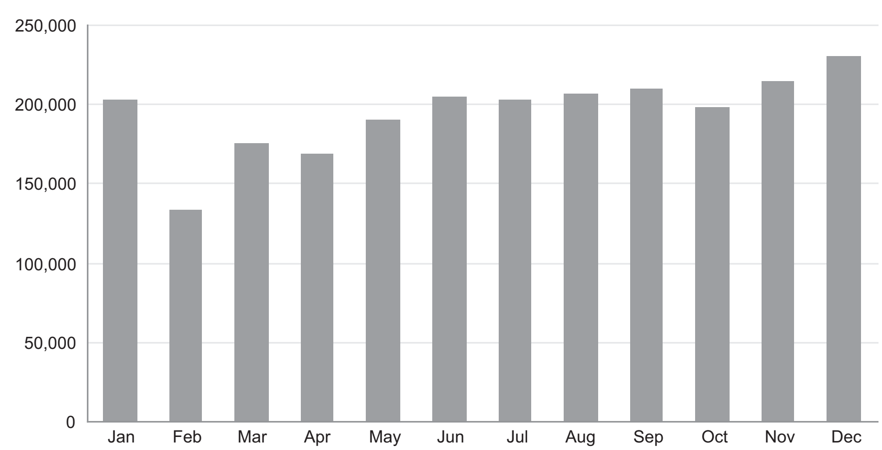

FIGURE 10.71 This graph uses grid lines to make it easier to perceive differences in the lengths of the bars.

ENHANCE PERCEPTION AND COMPARISON OF LOCALIZED PATTERNS

Graphs do a marvelous job of revealing the overall shape of an entire collection of values. And if you need to enhance perception of values in a subsection of the graph, grid lines can provide useful assistance. Imagine that you're examining the correlation between people's heights and salaries, which you've broken into two regional groups: West and East. For the moment, you want to compare values in the West to those in the East that fall from 65 to 70 inches in height and from $50,000 to $60,000. The grid lines in the graphs below would make this easy, but you'd struggle without them.

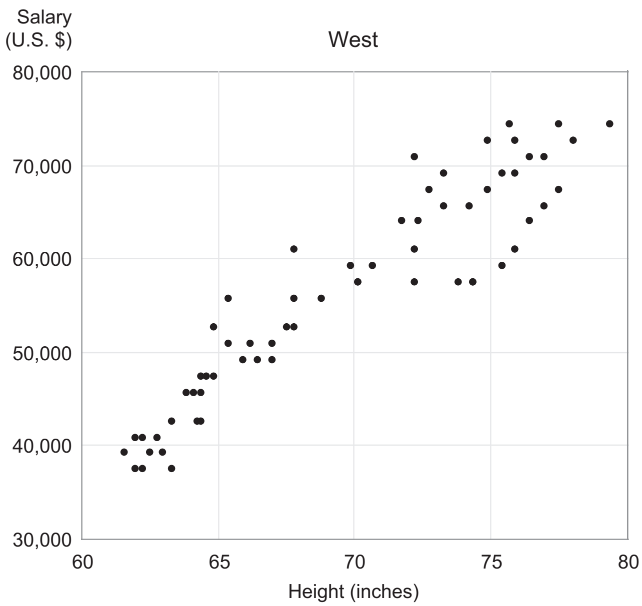

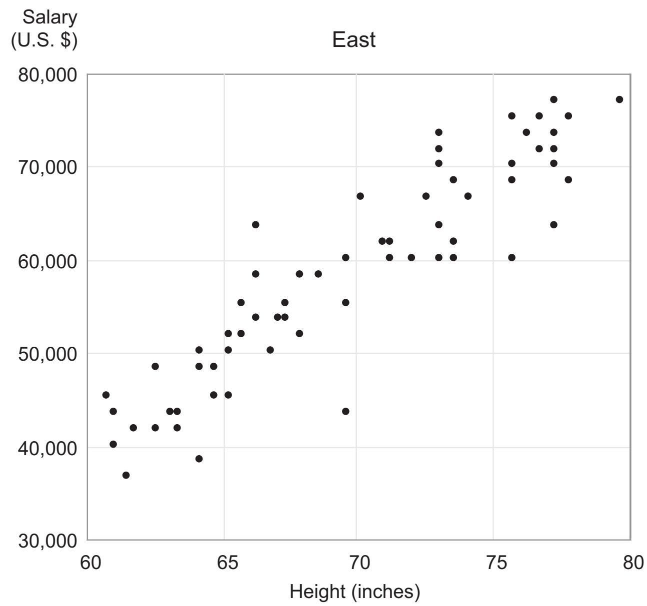

FIGURE 10.72 It would be difficult to isolate and compare regions of data in these scatterplots without the grid lines.

Lines that don't correspond to every value along scale (technically not grid lines) can also be used in a scatter plot to delineate meaningful sections of data. For instance, it is sometimes useful to divide a scatter plot into quadrants, as in the following example:

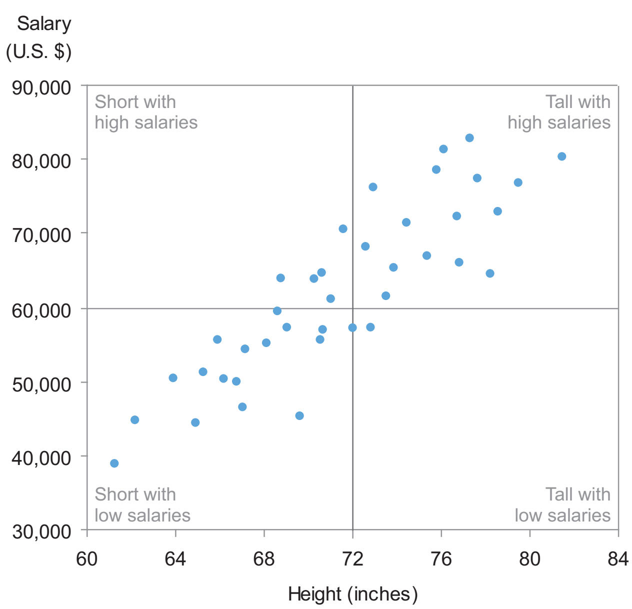

FIGURE 10.73 This scatter plot uses two intersecting lines to evenly divide the data region into quadrants.

The use of lines to divide scatter plots into quadrants can bring features of the data to light that might otherwise be missed. The previous example clearly reveals, with its nearly empty lower right quadrant, that tall employees rarely receive low salaries.

Legends

Legends are the tiny tables that associate categorical items with the visual attributes that encode them (e.g., hue, color intensity, or point shape). We'll look briefly at the following issues:

- When can you eliminate legends?
- Where should legends appear on the graph?
- How visible should legends be?
- Should legends have borders?
- How should you arrange labels in legends?

WHEN CAN YOU ELIMINATE LEGENDS?

Whenever categorical items are encoded in a graph, they must be labeled. If they appear on a categorical scale along an axis, they are labeled along the axis. However, if they are encoded using some other visual attribute, such as hue, they must be labeled in some other way. Legends are the conventional way to label these categorical items, but one circumstance lends itself to an alternate form of labeling. The graph below gives a hint. Take a look at it and see whether you can determine a different way to label the lines.

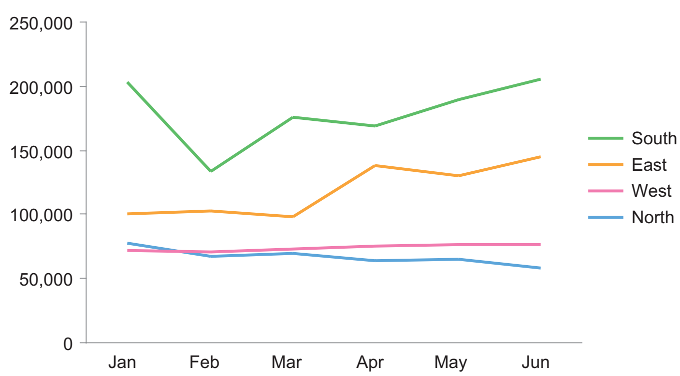

FIGURE 10.74 This graph includes a legend that could be replaced using a different means of labeling the categorical items of South, East, West, and North.

It's fairly easy to recognize in this example that the legend could be replaced by labeling the lines directly as you've seen me do on several occasions. Here's the improved version:

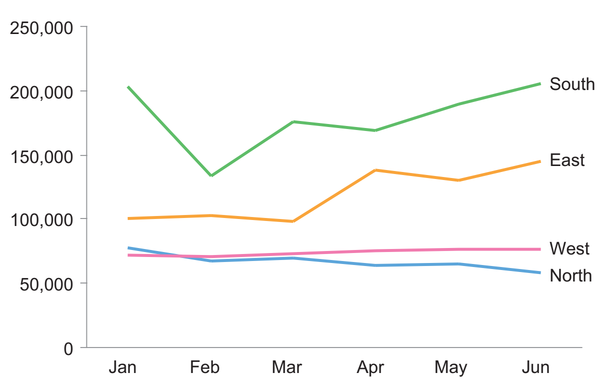

FIGURE 10.75 In this graph categorical items are labeled directly, next to the lines that encode them, rather than separately in a legend.

If you are producing graphs using software that doesn't offer the option of placing labels right next to the data sets, you can often type the labels in yourself.

The answer to the question "When can you eliminate legends?" is "Whenever the data items that need labels are grouped together so that you can place a label right next to each set." When lines are used to encode sets of values, each line acts as a grouping mechanism, making it easy to place a label right next to the line. However, as you can see in the next example, bars don't work this way:

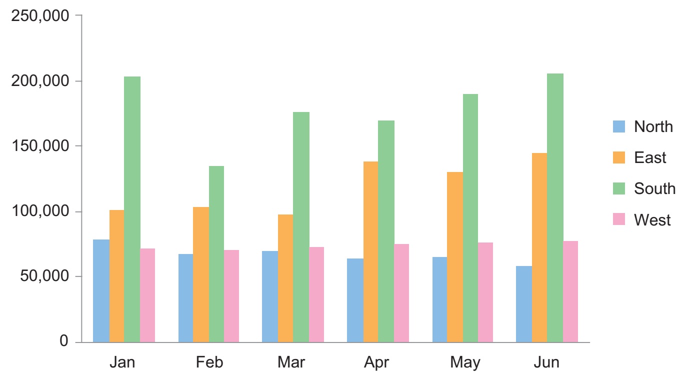

FIGURE 10.76 This is a graph that must use a legend to label the categorical items because the values are encoded as bars.

You certainly wouldn't want to repeat the same labels over and over again for each bar.

WHERE SHOULD YOU PLACE LEGENDS?

Legends are generally placed outside of the data region, but this needn't be a rigid rule. If you can place the legend inside the data region without getting in the way of the data values or interfering with perception of their shape, there is no reason why you shouldn't put it there. In fact, the closer the legend is to the data values, the easier it is to read the graph. Legends may be placed anywhere they fit and don't interfere with other more important components of a graph.

HOW VISIBLE SHOULD LEGENDS BE?

Legends provide the means to interpret categorical data, which is critical information, but they are not the actual data themselves. Therefore, although legends should be clearly visible in the graph, they should be somewhat less prominent than the actual data. You want your readers' eyes to be drawn predominantly to the data region of the graph.

SHOULD LEGENDS HAVE BORDERS?

You may have noticed that few of the legends that appear in examples in this book have borders around them. This is an intentional omission. Based on your understanding of visual perception, why do you think this is? Look at the two graphs below to see the difference:

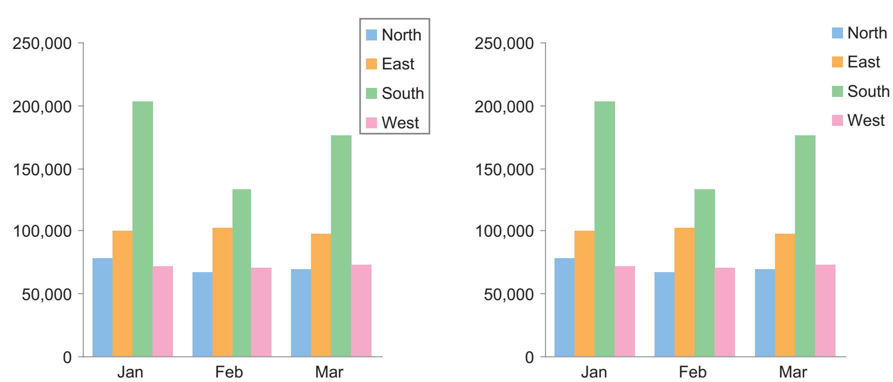

FIGURE 10.77 These two graphs are the same except that in one a border surrounds the legend and in the other the legend has no border.

First, the border around the legend does not add any meaning to the graph, and the legend doesn't need a border in order to be clearly distinct from other parts of the graph. Second, enclosing the legend with a border draws undesirable attention to it. The legend plays only a supporting role, so it isn't where you want your readers' eyes to be drawn. And finally, the border serves as a container, perceptually separating the legend information from rest of the graph.

When a legend is placed near another graph component and these two components must appear distinct from one another, something is needed to visually distinguish the legend from that other component; in this case, a border is useful. Be sure to make the border light, just visible enough to set the legend apart.

HOW SHOULD YOU ARRANGE LABELS IN LEGENDS?

You are probably accustomed to seeing the labels in legends arranged vertically rather than horizontally, but either arrangement can work fine. It really depends on the space you're working with and how the data is encoded. In the following example, the legend is arranged horizontally below the title, which prevents the graph from being wider than necessary.

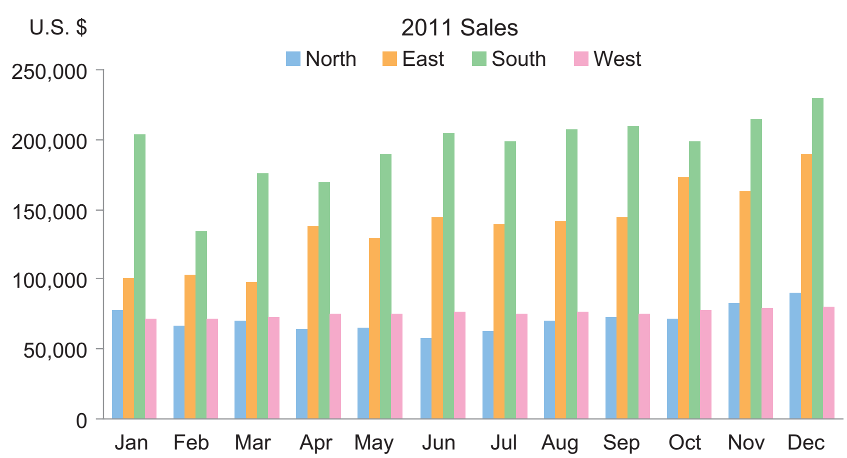

FIGURE 10.78 In this graph the labels in the legend are arranged horizontally.

This legend's horizontal arrangement of labels and position above the data region save horizontal space in the graph. The labels are also arranged in the order of the corresponding bars, which improves the graph's ease of use.

HOW MANY ITEMS CAN A LEGEND CONTAIN?

There is no conceptual limit to the number of categorical items you can include in a single graph, but there is definitely a practical limit. Too many visually distinct categorical items cannot be decoded without laborious effort because, as you learned in Chapter 5, *Visual Perception and Graphical Communication*, we can only hold the meanings of a few distinctions in working memory at one time. We've all seen graphs like the following:

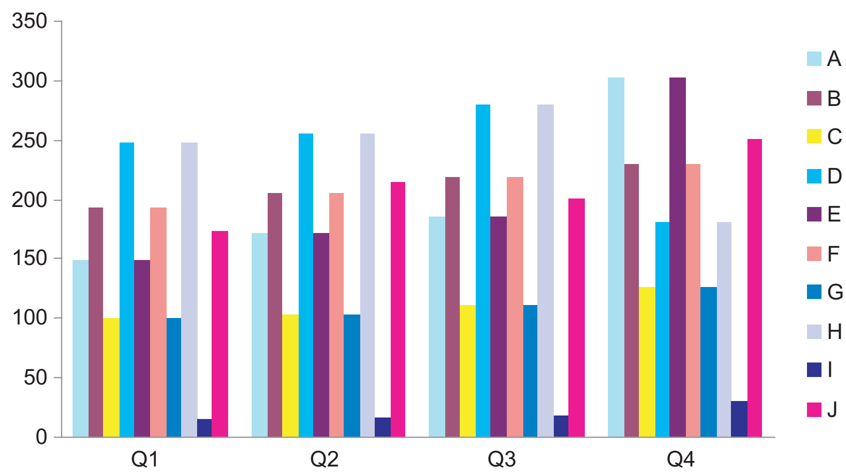

FIGURE 10.79 This graph has too many categorical subdivisions.

In fact, many of us have created similar graphical monstrosities. Not only is it impossible to keep the meaning of each color in mind, the graph also appears horribly cluttered. Clutter is visually exhausting, discouraging efforts to read the graph.

Non-Data Component Design

Axis lines are the only non-data components that are routinely useful in graphs. Axes give graphs dimension and serve as a container for the data. A single axis, either vertical or horizontal, produces a 1-D graph. Two perpendicular axes, one vertical and one horizontal, produce a 2-D graph, which is by far the most common type. The space that is defined by the axes where the data values are plotted (i.e., where the points, lines, or bars reside) is called the data region or the plot area.

Axes

When you design a graph, you need to ask two questions regarding the axes:

- Should the graph include one, two, three, or four axis lines?
- What ratio of the lengths of the X and Y axes works best?

SHOULD THE GRAPH INCLUDE ONE, TWO, THREE, OR FOUR AXIS LINES?

2-D graphs usually include one vertical axis on the left and one horizontal axis along the bottom. Under most circumstances, this convention is all that's needed. It is sometimes useful, however, to change the position of the axes. In the example that appears on the next page, you will see conventional axis placement as well as the alternatives:

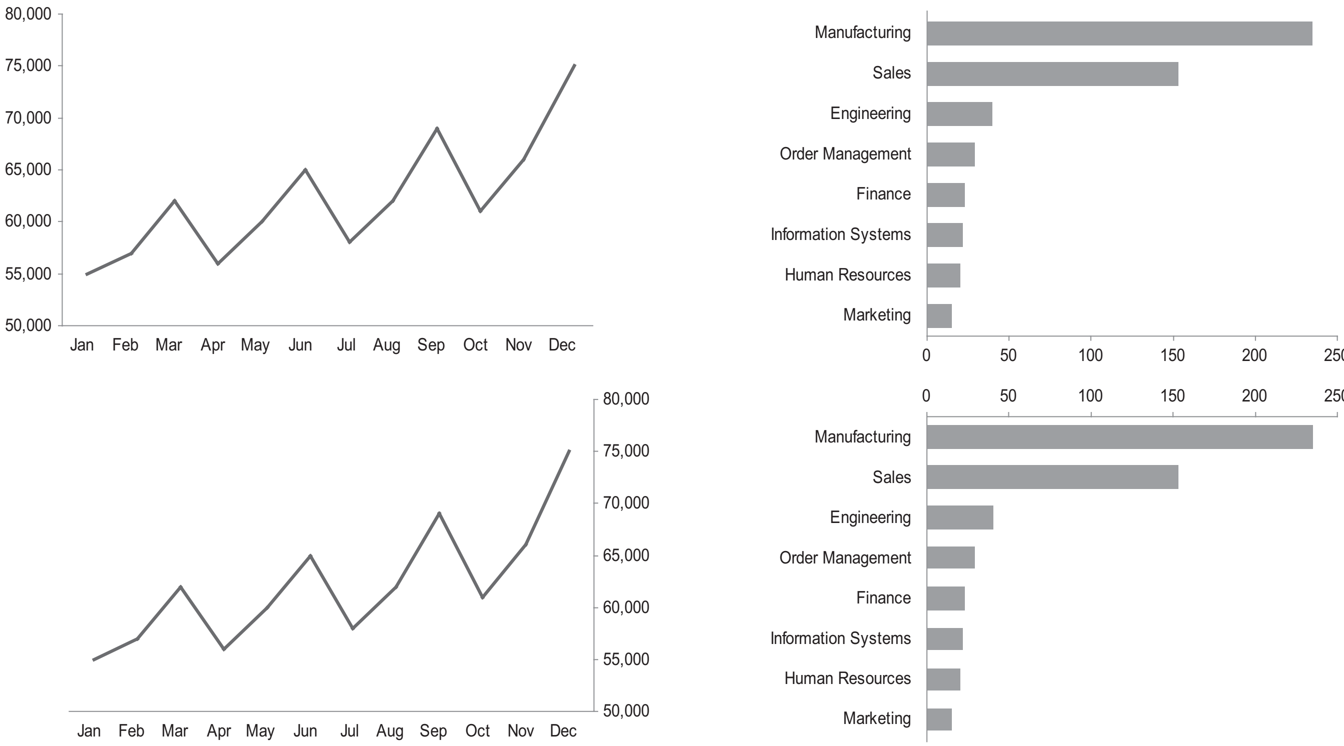

FIGURE 10.80 The upper graphs position the axes and scales conventionally and those below show the alternatives.

Looking at the lower line graph, can you think of an occasion when placing the Y axis on the right would make sense?

• • • • • • •

One obvious benefit is the fact that it is now easier to decode values along the line that appear on the right because the scale is closer to them. Did you also notice how strongly your eyes are drawn to the right half of the graph by the fact that the Y axis and scale are there. One occasion when this switch is useful is when you want to feature values on the right (e.g., the months in the final quarter of the year) more than those on the left. The mere presence of visual content on the right that is missing on the left draws the viewer's attention to the right. Switching the position of the axis from left to right or from bottom to top is a useful way to direct your reader's attention to the opposite region of the graph without adding anything to the graph to accomplish this.

Regardless of position, are there times when more or less than two axis lines are useful alternatives? In the following illustration, the conventional configuration of two axis lines is shown along with three alternatives:

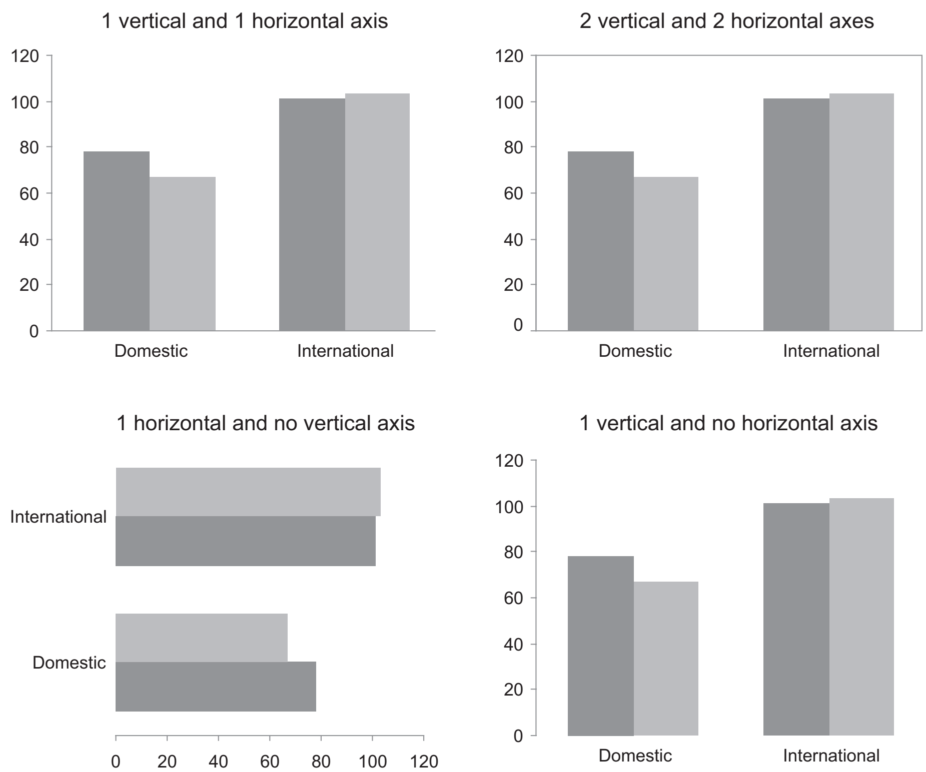

FIGURE 10.81 These graphs illustrate various axis configurations.

Two perpendicular axes tracing only two sides of the rectangular data region sufficiently define the space in most cases. Including two pairs of vertical and two pairs of horizontal axes to form a complete border around the data region is mostly useful when the data region must be clearly separated from surrounding content because the surroundings would otherwise compete too forcefully for attention. An axis may be left off without adverse affect when it hosts a categorical scale, and the values are encoded as horizontal bars. Because bars begin at the axis, their edges trace the line that the axis would otherwise display, which adequately delineates the data region. This works fine for horizontal bars, but I find that, without a base, vertical bars appear to float in space, as in the graph on the bottom right.

WHAT RATIO OF THE LENGTHS OF THE X AND Y AXES WORKS BEST?

The ratio of the length of the horizontal axis to the length of the vertical axis, or stated differently, the ratio of the data region's width to its height, is called the *aspect ratio*. It is calculated as width divided by height. The aspect ratio of a graph's data region greatly influences perception of the data. On the next page are a few examples of aspect ratios that vary from 2 to 0.5.

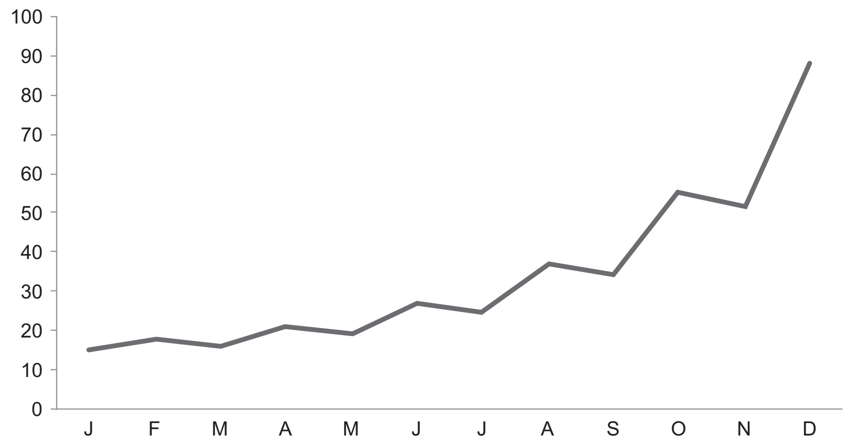

FIGURE 10.82 This graph has an aspect ratio of 2 to 1, or 2.

FIGURE 10.83 This graph has an aspect ratio of 1.5 to 1, or 1.5.

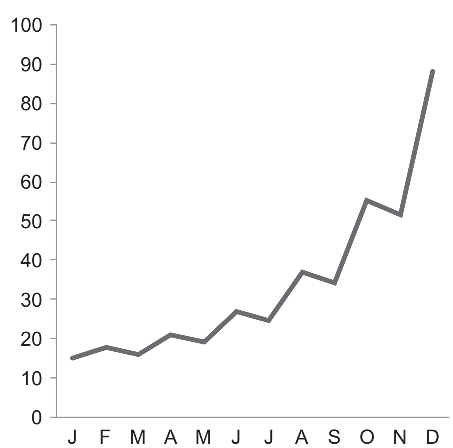

FIGURE 10.84 This graph has an aspect ratio of 1 to 1, or 1.

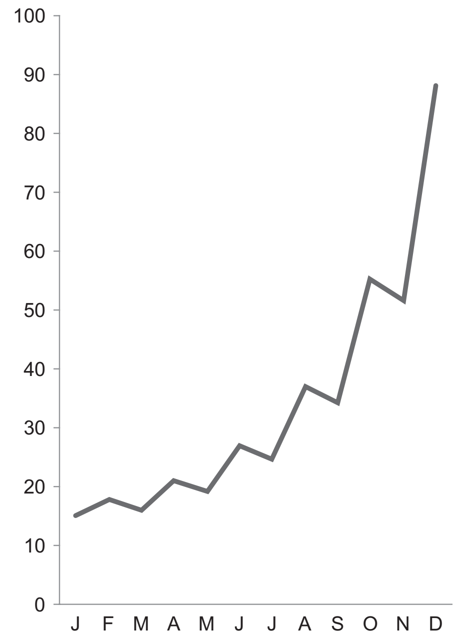

FIGURE 10.85 This graph has an aspect ratio of 1 to 1.5, or 0.67.

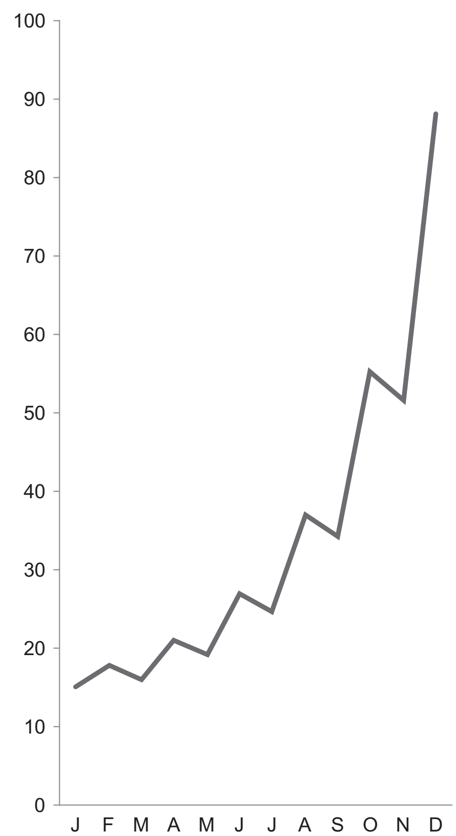

FIGURE 10.86 This graph has an aspect ratio of 1 to 2, or 0.5.

All of these graphs contain the same data; only their aspect ratios differ. When lines are used to encode values, as the aspect ratio increases, so does the appearance of the degree of change. A degree of change that already looks impressive with an aspect ratio of 2.0 looks like a blast-off at Cape Canaveral with an aspect ratio of 0.5. Both are accurate, both display precisely the same numbers, but they certainly differ in how they would be perceived.

There is no single aspect ratio that is always best. There are, however, two design practices that you should keep in mind. The first is that you should never manipulate the aspect ratio to intentionally exaggerate or downplay the degree of change. If your graphs usually appear wider than they are tall, suddenly making one taller than it is wide to convince your readers that sales are going through the roof would qualify as manipulation. The other general practice is to stick to the convention of making your graphs that display time series wider than they are tall. Emphasizing the horizontal rather than the vertical generally makes time series a bit easier to read and more in line with what people are accustomed to seeing.

Data Region

Apart from determining the data region's aspect ratio, we must design it to be a clean background in front of which the data can be easily seen. Readers' eyes should be drawn to the data. This is accomplished by making the points, bars, lines, or boxes visually prominent, not by energizing the background with vibrant color, a perception-skewing color gradient, or a silly image.

Data objects will stand out well against a light background unless they are excessively light themselves. A white data region is usually the best background, but there are times when other light colors, such as light gray or yellow, are useful, when a little extra is needed to draw readers' eyes to the data region because other content on the page or screen competes for attention. On such occasions subtle light fill color usually does the trick. Here are a few examples of fill colors that work:

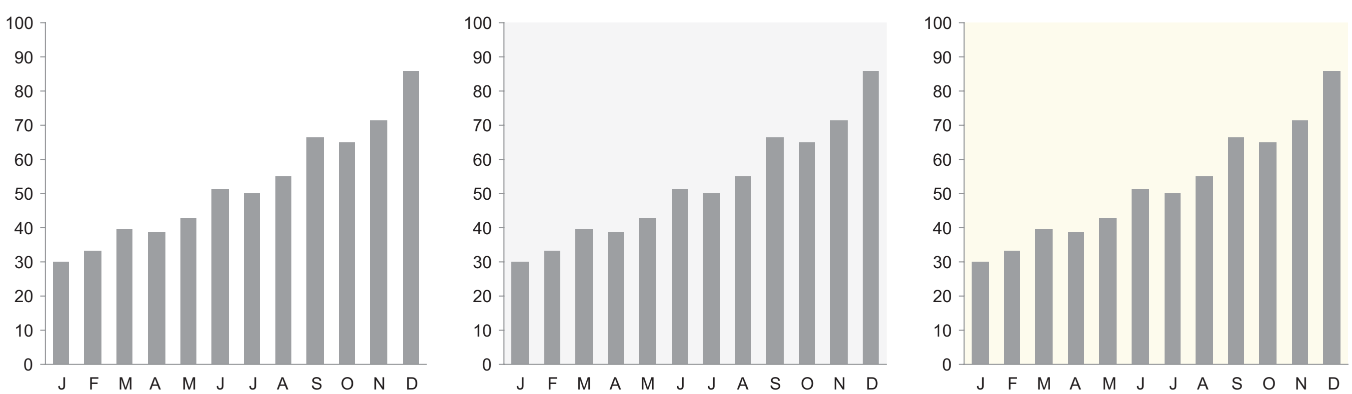

FIGURE 10.87 These examples show data region fill colors that provide an appropriate background for the data.

Another way to subtly highlight the data region is to use fill color for the graph except in the data region. By leaving the fill color of the data region white and coloring the surrounding areas of the graph with a light color, the white of the data region stands out in contrast, as in the following examples:

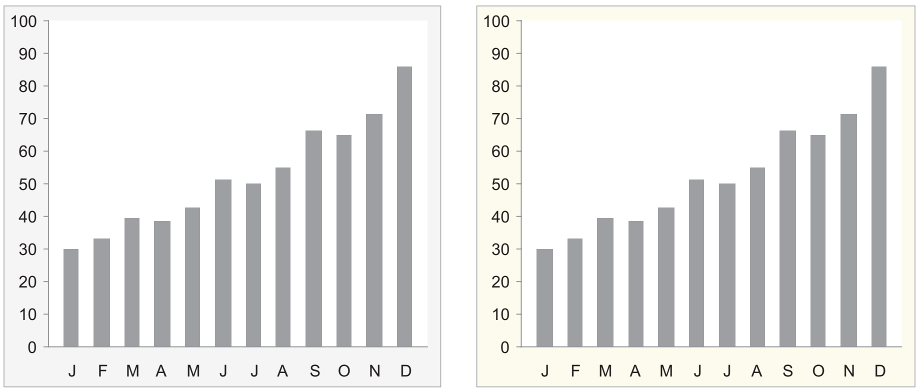

FIGURE 10.88 These graphs highlight the data region by using a subtle fill color for the surrounding areas.

When the data region needs an extra visual boost, the other visual attributes of the graph, especially the colors of the data objects, determine which of the two methods will work best, a fill color for the data region or for the surrounding areas of the graph. For instance, if your graph contains several sets of bars, and the color of one is light, the use of a subtle fill color in the data region might help them stand out more clearly than they would against white.

Gradients of color have their appeal, I suppose, but when they appear in the background of a graph, they skew perception of the data. Notice in the following example how the lines change in appearance depending on the background color.

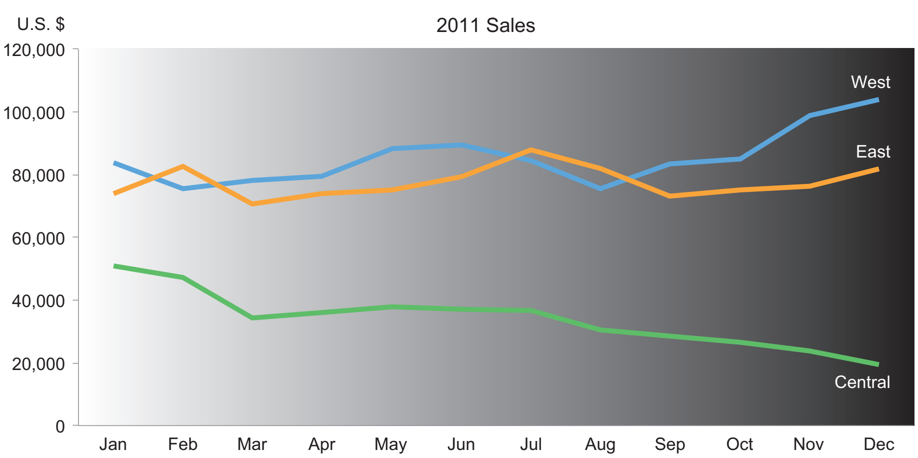

FIGURE 10.89 Color gradients in the background of the data region skew perception of the data.

Avoid gratuitous gradients and by all means resist the temptation to decorate you graphs with photographs or clip art.

FIGURE 10.90 Images in the background make it difficult to focus on the data.

An example like this might work fine in a magazine ad if you care little about readers' perception of the data, but they don't belong in the background if your objective is to tell a story contained in data. I'm not saying that photos and other images are never useful when presenting data. An appropriate photo can make a story real to people and touch their hearts in powerful way. Using a photo in conjunction with a graph in a way that complements the data, such as by placing the photo and the graph near one another, can tell a richer story without undermining the integrity of the data.

Summary at a Glance

Component	Practices

Points	• When sets of points cannot be clearly distinguished, correct by:
	  ◦ Enlarging the points
	  ◦ Selecting objects that are more visually distinct
	• When points overlap such that some are obscured, correct by:
	  ◦ Enlarging the graph and/or reducing the size of the points
	  ◦ Removing the fill colors

Bars	• Use horizontal bars when their categorical labels bars won't fit side by side.
	• Never use horizontal bars for time-series values.
	• Proximity
	  ◦ Set the width of white space separating bars that are labeled along the axis equal to the width of the bars, plus or minus 50%.
	  ◦ Do not include white space between bars that are differentiated by a legend.
	  ◦ Do not overlap bars.

Component	Practices
Bars *(continued)*	• Fills ◦ Avoid the use of fill patterns (e.g., horizontal, vertical, or diagonal lines). ◦ Use fill colors that are clearly distinct. ◦ Use fill colors that are fairly balanced in intensity for data sets that are equal in importance. ◦ Use fill colors that are more intense than others to highlight particular values. • Only place borders around bars when one of the two following conditions exists: ◦ The fill color of the bars is not distinct against its background, in which case you can use a subtle border (e.g., gray). ◦ You wish to highlight one or more bars compared to the rest. • Always start bars at a baseline of zero.
Lines	• Distinguish lines using different hues whenever possible. • Include points on lines only when values for the same point in time on different lines must be precisely compared.
Boxes	• Follow the principles for bar design, except when box plots are connected with a line to show change through time, which might require greater distance between boxes.
Combinations	• Use boxes and lines for distributions through time. • Use bars and lines in the form of Pareto charts for featuring the contribution of the largest portions of the whole. • Use bars and points for uncluttered comparisons.
Trend Lines	• In most cases, use moving averages rather than straight lines of best fit to show the overall nature of change through time. • Only use linear trend lines (straight lines of best fit) in a scatter plot when the shape of the data is linear rather than curved.
Reference Lines	• Use reference lines to mark meaningful thresholds and regions, especially for measures of the norm.
Annotations	• Use text to feature and comment on values directly when doing so is important to the story.
Log Scales	• Use log scales to reduce the visual difference between quantitative data sets with significantly different values so they can be clearly displayed together. • Use log scales to compare differences in value as percentages.
Tick Marks	• Mute tick marks in comparison to the data objects. • Use tick marks with quantitative scales but not with categorical scales, except in line graphs when slightly more precision is needed. • Aim for a balance between including so many tick marks that the scale looks cluttered and using so few that your readers have difficulty determining the values of data objects that fall between them. • Avoid using tick marks to denote values at odd intervals.

Component	**Practices**
Grid Lines	• Thin, light grid lines may be used in graphs for the following purposes: ◦ Ease look-up of values ◦ Ease comparison of values ◦ Ease perception and comparison of localized patterns
Legends	• Use legends for categorical labels when the labels are not associated with a categorical scale along an axis and cannot be directly associated with the data objects. • Place legends as close as possible to objects they label without interfering with other data. • Render legends less prominent than the data objects they label. • Use borders around legends only when necessary to separate legends from other information.
Axes	• Don't manipulate the aspect ratio to distort perception of the values.
Data Regions	• Keep the background clean and light.