The Role of Software and the Importance of Thoughtful Defaults

Stephen Few, Perceptual Edge
Visual Business Intelligence Newsletter
December 2007

I believe that every business ought to exist because there is a need for what it provides and that it is the responsibility of every business to serve these needs to the best of their ability. Too many software companies fail to serve the real needs of their customers. They strive to do only what they must to drive sales. Buyers of software are far too forgiving and they expect too little. Software that is hard to use, full of bugs, and barely does its job is what buyers have come to expect, because it is the norm, which they feel powerless to change. There are a few delightful exceptions, but we are still far from the tipping point that must be reached to significantly change the software industry for the better.

Software vendors should support the needs of their customers, but not necessarily their desires. Vendors ought to know the difference. They should understand what people really need and how to meet those needs most effectively, rather than address them superficially, which is what they often do instead. Software should routinely be designed to make it easy for people to do what works and difficult to do what doesn't. One important and powerful way that software products can do this is through well-designed defaults. I think that Barry Schwartz made this point eloquently in a recent *Scientific American Mind* article:

> *Why do defaults have such power? Some of it may come from inattention. Life is busy and complicated, and it is not possible to pay attention to everything. That is why most of us keep our cell phone plan whether or not it is the best one for us. Researching alternatives is time-consuming, and we do not want to be bothered. But laziness and inattention are not the sole reasons for the power of defaults. As University of California, San Diego, psychologist Craig R. M. McKenzie and his colleagues showed in a 2006 study, most people infer that the default is the recommended option.*

> *Given the power of defaults, policymakers could use them to nudge people in a direction that will enhance their well-being, something University of Chicago legal scholar Cass R. Sunstein and economist Richard Thaler call "libertarian paternalism." In this practice, leaders would choose defaults with an eye on people's stated or implied preferences (the "paternalistic" part) while allowing anyone to opt out (the "libertarian" element).*

> Barry Schwartz, "When Words Decide," *Scientific American Mind,* August/September 2007, pp. 40 and 41.

In regards to data visualization software in particular (the kind that uses graphs to display quantitative data), there are two basic ways that vendors can make the lives of users easier through default settings: 1) Automatically provide the right type of graph for the task at hand, and 2) automatically format the graph to represent the data most effectively.

"Software vendors should support the needs of their customers, but not necessarily their desires."

Defaults Settings for Selecting the Right Visualization

Most products do little, if anything, to help people select the kind of graph that will do the best job of displaying the data, based on the nature of the data and the task at hand. The best example of attempting to do this that I've seen so far exists in *Tableau* as one of their *Show Me!* features. You can select the data that you wish to examine simply by clicking them in the list of available data items and then click the *Show Me!* icon on the toolbar, and the data will automatically be displayed using the type of graph that is most often useful for examining data of the selected type. If you then wish to view the data in a different type of graph, you can click the *Show Me! Alternative* icon to open a window, which shows every type of graph that's available, but with the ones that wouldn't be useful for the data you're examining dimmed out as inactive. The folks at Tableau understand data visualization well enough to know the rules that determine which types of graphs are appropriate for particular types of data. For instance, if I am examining data that only includes a single quantitative measure, such as sales revenue, it wouldn't make sense to allow me to use a scatterplot, which is designed to compare two quantitative variables to see if and how they are correlated.

One of the simplest ways that software can help people select the right graph is to only support graph types that are useful, skipping those that are obscure or just plain silly. One of the best examples of what *not* to do can be found in *Excel* prior to the 2007 version, which forced people to choose from some pretty bad graphs, and was often partly to blame for people selecting ineffective and even downright misleading graphs. In *Excel 2007*, the list of basic graph types has thankfully been pared down considerably, and those types that are the least commonly useful are not directly exposed in the new ribbon interface, but accessed through an icon labeled *Other Charts.*
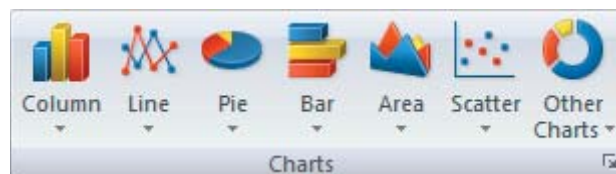
I applaud this improvement in the design of Excel, but bemoan the fact that now, once you've selected one of the common types of graphs, you must then make an additional selection from a long list of versions of that graph, many of which are awful, such as 3-D, cylindrical, conical, or pyramidical versions of a column graph (a.k.a., a vertical bar graph).

Most quantitative data analysis and presentation can be performed using five basic types of graphs:

- bar graph (both vertical and horizontal bars)
- line graph (with or without data points along the line)
- dot plot
- scatterplot
- box plot (both vertical and horizontal boxes, along with variations for displaying stock data)

There are times when other types are needed, such heatmaps and treemaps, which can be quite useful when used properly, but the five types of graphs listed above do the job quite well in most circumstances. General purpose visual analysis and presentation software ought to make it easy to access these five types, but might place any others out of close reach, perhaps accessible through an additional click.

"One of the simplest ways that software can help people select the right graph is to only support graph types that are useful, skipping those that are obscure or just plain silly."

Here are a few simple guidelines that the software can follow to help people select the appropriate type of graph from the available list:

- Use a line graph if the selected data includes an interval scale (see the note about interval scales on the next page).
- Use a bar graph to feature data along a nominal or ordinal scale.
- When a bar graph is selected, use horizontal bars if the labels along the categorical scale are either long or many, otherwise use vertical bars.
- If the quantitative scale of a bar graph is narrowed such that zero is no longer the baseline value, suggest that the graph be changed to a dot plot.
- Use a line graph whenever the selected data includes time (years, quarters, months, etc.).
- Use a scatterplot if two quantitative variables have been selected.
- Use a box plot if the selected data consist of many quantitative values belonging to a single variable, which are divided into two or more categorical subsets (such as a separate subset per region, department, or product). In other words, a box plot should be selected if it can be assumed that the graph's purpose is to display how the values in each subset are distributed across the quantitative range from lowest to highest, such that those distributions can be easily compared.

I have intentionally stopped far short of providing a full set of guidelines. I have also refrained from arranging these guidelines in the form of a decision tree, which you would need to implement graph selection intelligence that is comparable to Tableau's "Show Me!" functionality. You'll need to make the effort to fill in the gaps on your own. As vendors who develop the software, you are, after all, supposed to be experts.

Interval Scales

When labels along the graph's X or Y axis identify *what* is being measured, rather than the amount of the measurement itself, those labels compose what's called a *categorical scale* (that is, as opposed to a quantitative scale). It is useful to differentiate three types of categorical scales:

Nominal	The labels along the scale refer to items in a category that have no particular logical order and no close connection with one another. A list of departments in a company is a common example. The departments are clearly separate from one another and can be ordered in any way.
Ordinal	The labels along the scale refer to items in a category that have no close connection to one another, but they do have a logical order. For example, the list "small," "medium," and "large" has a logical order. Another common example would be a list of salespeople, when they are arranged in order of performance, thereby adding the following meanings to the scale: "top salesperson," "second place salesperson," and so on.
Interval	The labels along the scale refer to items that have a logical order and are closely connected to one another. They are intimately connected one to the next because they represent contiguous intervals—ranges of quantitative values—along a larger range of quantitative values. For example, in a graph that counts the number of employees in a company by age from 20 to 80, if age has been divided into equal intervals of 10 years each (20-29, 30-39, etc.), what you have is an interval scale. The most common interval scale that appears in graphs is time, divided into intervals such as years, quarters, months, or days. Time is an interval scale because it is linear and each item is intimately related. For instance, the year always moves from June to July.

Defaults for Graph Formatting

When a graph is selected and it first appears on the screen, it ought to already be designed as effectively as possible without requiring extra steps by the user. This means that its default appearance ought to correspond to graph design best practices, based on an understanding of the purpose for which that particular graph is most often used. Of course there will be times when something other than the default appearance might be needed, but this should be the exception, rather than the rule. In other words, when determining the formatting defaults for each type of graph that they supply, a software vendor should ask, "Under normal circumstances, what components should be included and how should they look?" A good example is whether grid lines should be included in a bar graph as a default, and if so, what color should they be, how dark should they be, and so on.

"When a graph is selected and first appears on the screen, it ought to already be designed as effectively as possible without requiring extra steps by the user."
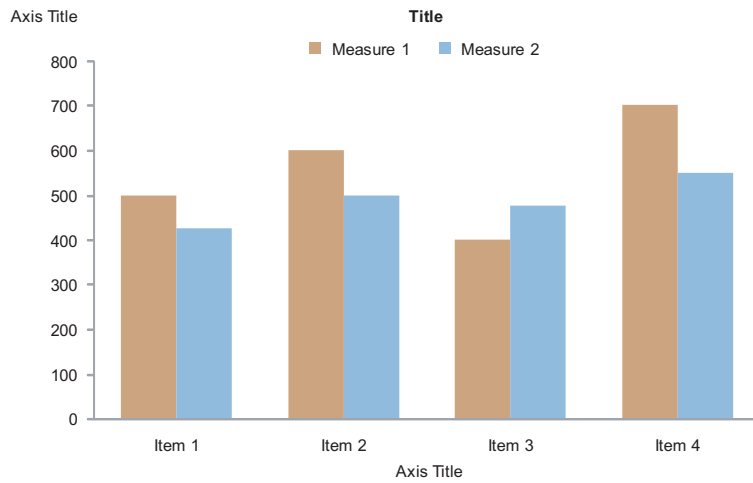
Here are a few of the formatting defaults that I find useful:

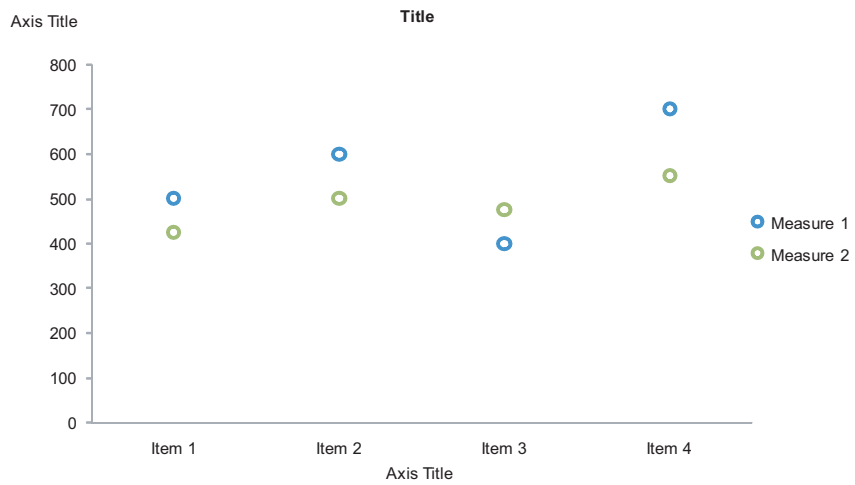| Component	Suggested Formatting Defaults
Entire graph area	• No fill color in the background of the graph • No border around the graph as a whole
Plot area	• No fill color in the background of the plot area • No border around the plot area • No grid lines. Allow the user to add them on those occasions when they are needed. When they are added by the user, make them just visible enough to do the job by using a light gray color and a thin stroke weight. • For scatterplots, set the length of the X-axis equal to the length of the Y-axis, resulting in a plot area that is square in shape.
Axes and scales	• Visually subdue the axes and tick marks by rendering them as a medium-saturated gray with a medium stroke weight. • Place tick marks along the outside of the axis rather than in the plot area. • Only include tick marks along quantitative scales (not categorical scales). • Whenever bars are used to encode the data, always begin the quantitative scale at zero. Negative numbers may be included in the range, but bars should always begin at a baseline of zero. This is because one of the ways that bars encode quantitative values is through their length (independent of the scale), but if the scale does not begin at zero, when bars are compared their differences are perceived inaccurately. • If the user then narrows the scale such that zero is no longer included, ask if she would like to switch to a dot plot rather than a bar graph. • Orient labels along both the X and Y axes horizontally.

Data objects (points, lines, bars, and boxes)	• No 3-D effects
	• No borders around bars or boxes
	• Set the gap width between bars or boxes, or between related sets of bars or boxes, to 100%. In other words, the space between individual or grouped bars or boxes should be equal to the width of a single bar or box.
	• No data points to mark the values along a line.
	• When data points are shown along a line, make them the same color as the line and make the diameter of the point only slightly larger than the width of the line.
	• Use a dot (circle) as the shape for data points. This applies to scatterplots, dot plots, and line graphs when data points mark values along the line.
	• In dot plots, show the outline of the dot only, without fill color, to make it easier to see the dots when they overlap.
	• Provide a default color palette for data objects (data points, lines, bars, and boxes) that consists of two basic sets of colors: a large set of soft colors (medium saturated colors) that are used for most items, and a smaller set of dark and bright (fully saturated) colors that can be used to highlight particular data.
Titles	• Include a graph title, X-axis title, and Y-axis title.
	• Center the graph's title just above the plot area and orient the text horizontally.
	• Place the Y-axis title above the axis and orient the text horizontally.
	• Center the X-axis title just below the X-axis labels and orient the text horizontally.
Legends	• No legend if only a single data set is displayed in the graph.
	• When a legend is displayed, do not include a border around it.
	• For vertical bar graphs and vertical box plots, place the legend above the plot area with the labels arranged horizontally (side-by-side) in the same order as the bars or boxes that they represent.
	• For horizontal bar graphs and horizontal box plots, place the legend to the right of the plot area with the labels arranged vertically (one on top of the other) in the same order as the bars or boxes that they represent.
	• For all other graphs, place the legend to the right of the plot area with the labels arranged vertically.

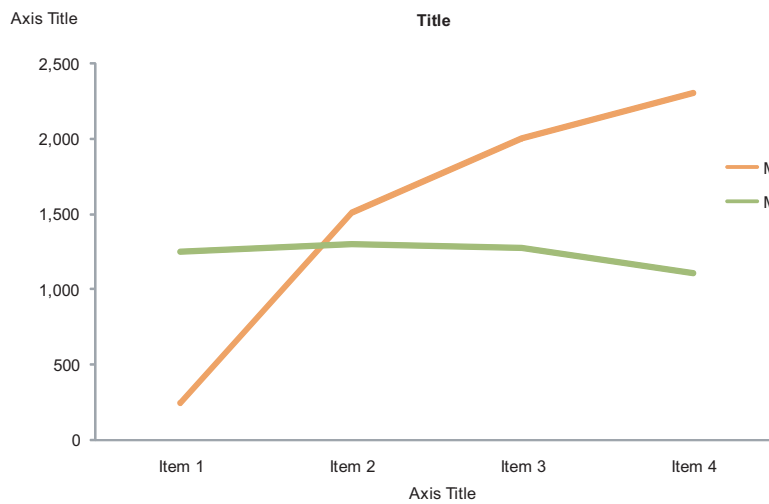Here are examples of effective default formatting for each of the five most useful types of graph:
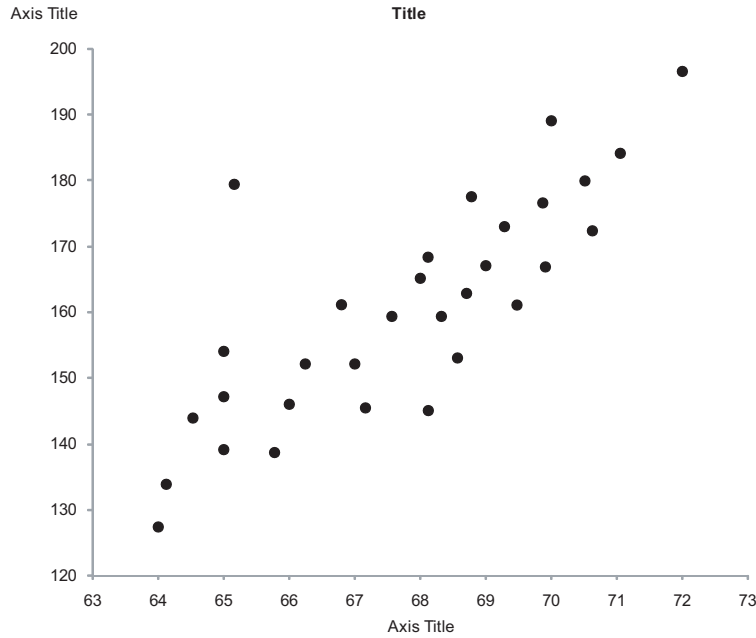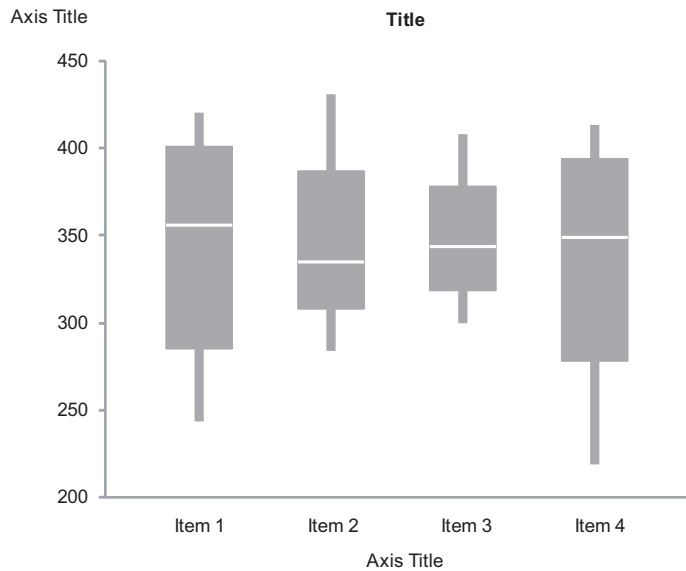
Bar Graph

Axis Title

Title

■ Measure 1 ■ Measure 2

Axis Title

Dot Plot

Axis Title

Title

○ Measure 1
○ Measure 2

Axis Title

Line Graph

Axis Title

Title

— Measure 1
— Measure 2

Axis Title

Scatterplot

Axis Title

Title

Axis Title

Box Plot

Axis Title

Title

Item 1   Item 2   Item 3   Item 4

Axis Title

Good defaults not only help by setting people's feet on the right path, but also by saving them time, supporting best practices from the beginning that would otherwise take several steps to implement. When shopping for new software, add good defaults to the criteria that you use to evaluate products: "Do the products defaults demonstrate that the vendor knows what it's doing and cares about its customers?" Good vendors take the time to provide good defaults. Although most of the new features that are routinely crammed into software product releases are rarely used, good defaults improve our lives every time we use the product.

About the Author

Stephen Few has worked for over 20 years as an IT innovator, consultant, and teacher. Today, as Principal of the consultancy Perceptual Edge, Stephen focuses on data visualization for analyzing and communicating quantitative business information. He provides training and consulting services, writes the monthly *Visual Business Intelligence Newsletter*, speaks frequently at conferences, and teaches in the MBA program at the University of California, Berkeley. He is the author of two books: *Show Me the Numbers: Designing Tables and Graphs to Enlighten* and *Information Dashboard Design: The Effective Visual Communication of Data*. You can learn more about Stephen's work and access an entire library of articles at www.perceptualedge.com. Between articles, you can read Stephen's thoughts on the industry in his blog.