Dashboard Confusion Revisited

Stephen Few, Perceptual Edge

Visual Business Intelligence Newsletter
March 2007

In March of 2004, three years ago exactly, my article titled "Dashboard Confusion" appeared in *Intelligent Enterprise* magazine. I wrote it because, at the time, I was concerned that the potential benefits of dashboards were being undermined by a great deal of confusion about what a dashboard was. The term "dashboard" needed a clear definition. In the article, I proposed a working definition in an attempt to reduce the confusion:

> *A dashboard is a visual display of the most important information needed to achieve one or more objectives; consolidated and arranged on a single screen so the information can be monitored at a glance.*

Although many people have embraced this definition over the years, the term dashboard continues to be used in reference to almost any type of screen-based display that combines more than a single chart, no matter what its purpose. This not only creates confusion, but it also makes it impossible to say anything useful about dashboards. For instance, you cannot say anything about how dashboards ought to be designed without first specifying the purpose of the display. Displays that are used for monitoring what's going on (dashboards) must be designed and must function quite differently than displays that are used to analyze data.

As a metaphor, the term dashboard lends itself naturally to a particular meaning. Dashboards in automobiles, and their close cousins, cockpits in airplanes, are used to *monitor* what's going on at a glance. They are not used for analysis. When applied to monitoring displays, the dashboard metaphor is useful and clear. Unfortunately, despite its merits, the metaphor has also opened the door to an interface that undermines the effectiveness of dashboards. Most vendors have spent more time making their displays look like automobile gauges and meters, and even traffic lights (a metaphorical leap), than they have on creating displays and functions that effectively support monitoring.

Multi-chart analytical displays have tremendous potential, but they are very different in design and function from dashboards. I would like to propose a unique name for them so we can discuss and promote them without confusion. I suggest that we call this a *Faceted Analytical Display*.

Introducing "Faceted Analytical Displays"

According to the *Shorter Oxford English Dictionary*, the adjective "faceted" means "having a form which has many faces or aspects." This perfectly describes the nature of these analytical displays, which combine several views of (or perspectives on) a common set of data to provide a rich display for analysis.

Once again, in an effort to promote clarity and the basis for fruitful discussion, I would like to propose a definition:

> *A "faceted analytical display" is a set of interactive charts (primarily graphs and tables) that simultaneously reside on a single screen, each of which presents a somewhat different view of a common dataset, and is used to analyze that information.*

Comparison is the fundamental activity of analysis. Something powerful happens when we are able to see data simultaneously from several perspectives. We are able to spot connections and discern relationships that would otherwise remain hidden. If we were to look at these same views of the data independently, one at a time, we would never gain the same insights, largely due to the limits of short-term memory. Once one of these views is no longer in front of our eyes, we remember little, if anything, of what we saw, so we cannot compare it to what we're looking at now. We can hold about four chunks of information at one time in short-term memory. One of the best ways to work around this limitation is to place everything that you wish to compare within eye span. Software that allows us to do this eliminates the need to hold information in short-term memory. An example of what I mean by a faceted analytical display appears on the next page, which I created using Spotfire's visual analysis software called DecisionSite.

"The greatest clarification that is needed today is a distinction between dashboards, which are used for monitoring what's going on, and displays that combine several charts on a screen for the purpose of analysis."

Here we are now in March of 2007, three years from the release of my original article, and I am compelled once again to plead for clarity to end the confusion. The greatest clarification that is needed today is a distinction between dashboards, which are used for monitoring what's going on, and displays that combine several charts on a screen for the purpose of analysis.

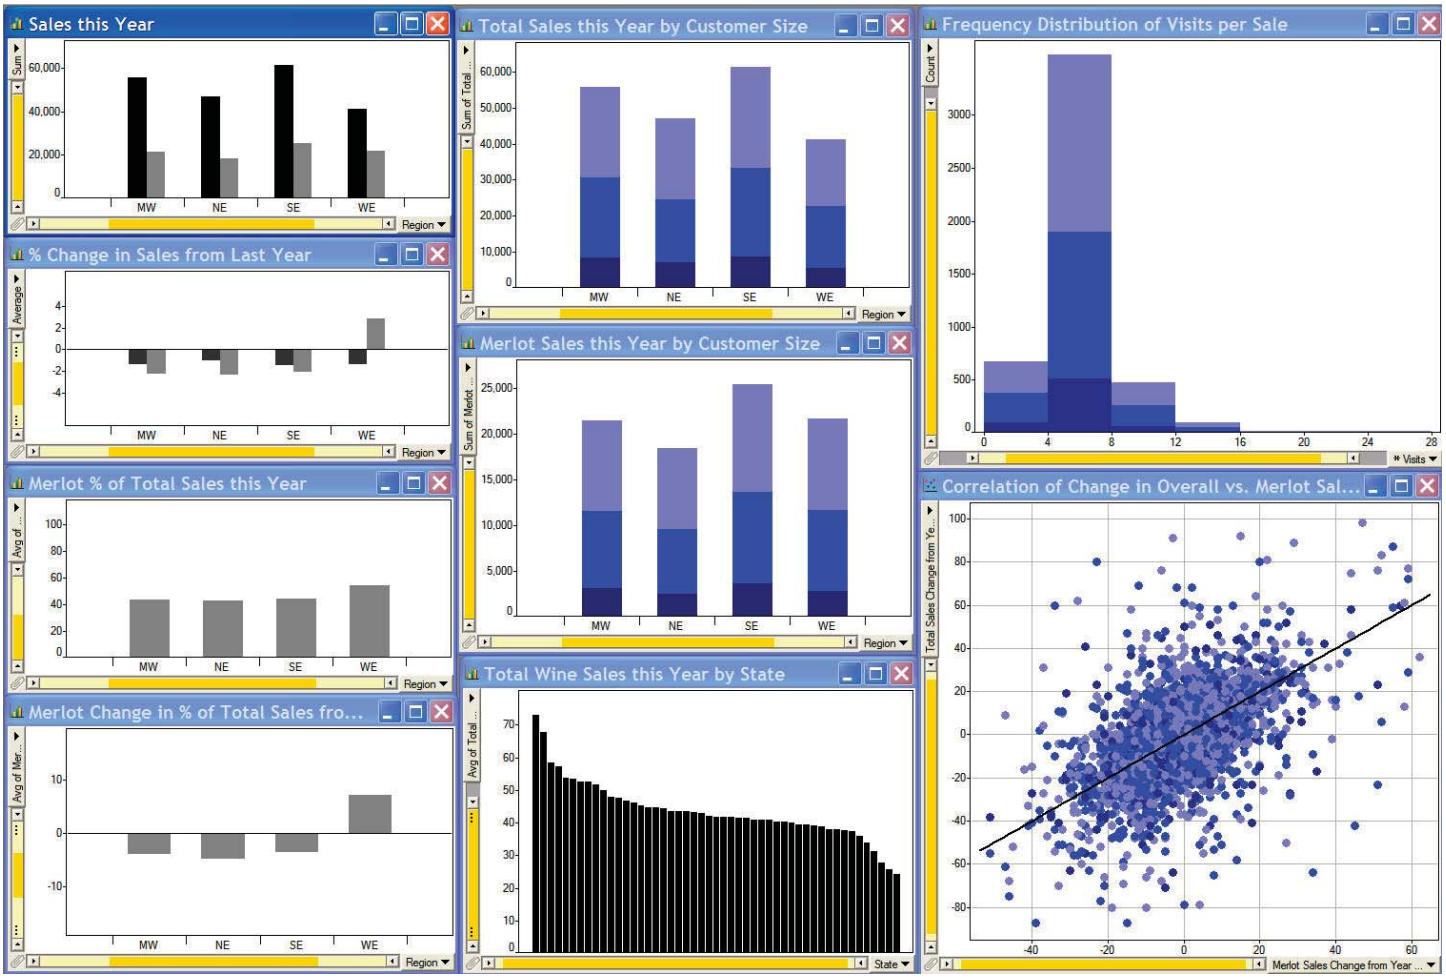

Figure 1: Example of a faceted analytical display created using Spotfire DecisionSite.

Each graph in this collection of nine reveals a different aspect of wine sales, based on a common set of data. Starting with the graph in the upper left-hand corner and working down each column before moving right to the next, here's what they display:

Column 1:

- Total wines sales (black bars) compared to Merlot sales (gray bars) in U.S. dollars for this year per region.
- Percentage change in sales between last year and this year—total (black bars) and Merlot (gray bars)—per region.
- Merlot's percentage of total wine sales per region.
- Change in Merlot's percentage of total sales from last year to this year per region.

Column 2:

- Total wines sales this year per customer size (darkest blue represents big customers, medium blue represents medium-sized customers, and light blue represents small customers) by region.
- Merlot sales this year per customer size by region.
- Total wines sales this year by state.

Column 3:

- Frequency distribution of sales based on the number of sales visits that led to each sale (0-3, 4-7, 8-11, etc.).
- Scatterplot showing the correlation of the percentage change in total wine sales from last year to this year (vertical axis) to the percentage change in Merlot wine sales from last year to this year (horizontal axis). Each data point represents a particular customer.

Seen together, these graphs help us see connections (interesting interactions) between various aspects of sales. Here are a few features that caught my attention:

- Merlot sales in the west region represent around 55% of total sales, the greatest share of any region, but this is mostly due to a dramatic increase in Merlot sales in the west. Despite this strong showing in the west, Merlot sales decreased from the first year to the next more than total wine sales in all regions except the west.
- Even though the west came in last in overall wine sales this year, they came in second in Merlot sales.
- Most sales resulted from between four and seven visits to a customer and the number of visits that seems to be required for sales is fairly consistent between big, medium, and small customers.
- While there is a positive correlation between changes in total wines sales vs. Merlot sales from last year to this year (that is, as total wines sales increased or decreased, Merlot sales tended to do the same), there are a number of exceptions to this pattern, especially in the lower right corner where a disproportionate number of Merlot sales increases appear relative to decreases in overall sales.

Many of these observations might have remained invisible using more traditional business intelligence methods of analysis, and certainly would have been much more difficult and time-consuming to uncover.

Important Faceted Analysis Functionality

The power of these rich analytical displays is dramatically extended when software provides two specific ways to interact with the data:

- Global filters
- Brushing

Global filters provide the means to remove unwanted data from the entire display. The wine sales data in the example above segments customers into “big,” “medium,” and “small” customers. An example of a filtering activity would be to remove medium and small customers from the display so you can focus without distraction on the big customers only (illustrated in Figure 2). Another example could involve removing all wine orders that were for less than $500. Filters in some analysis tools can only be used to affect individual charts. Global filters, however, work across the entire collection of charts, keeping the data in sync across the entire display. When we examine data from multiple perspectives simultaneously, we almost always want to filter the data in a consistent manner across all of the charts. The easier the filter controls are to use and the faster they can be set, the more powerfully they support the free flow of analysis.

One powerful means of interacting with the data, called “brushing,” extends our ability to see connections in the data by highlighting selected items across all of the views. Imagine that while examining the wine sales data, we become interested in a particular subset of the data, such as those Merlot sales that decreased since last year even though overall wines sales increased, as shown in the upper left quadrant of the scatter plot in Figure 3 at the bottom of this page. We don’t want to filter out everything else from the display, because we want to compare this subset of data to the data as a whole. We could brush this subset of data, which involves highlighting these values through a simple action such as drawing a rectangle around them.

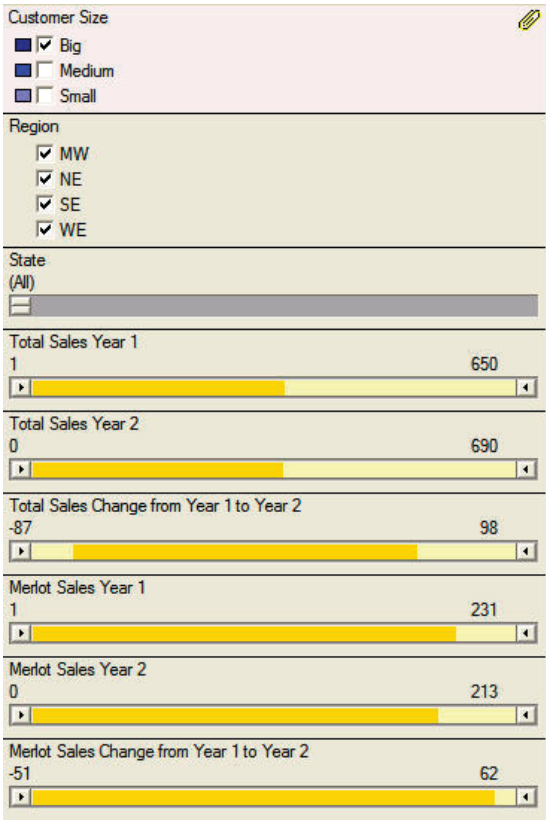

Figure 2: Example of global filters, which use simple mechanisms such as checkboxes and sliders.

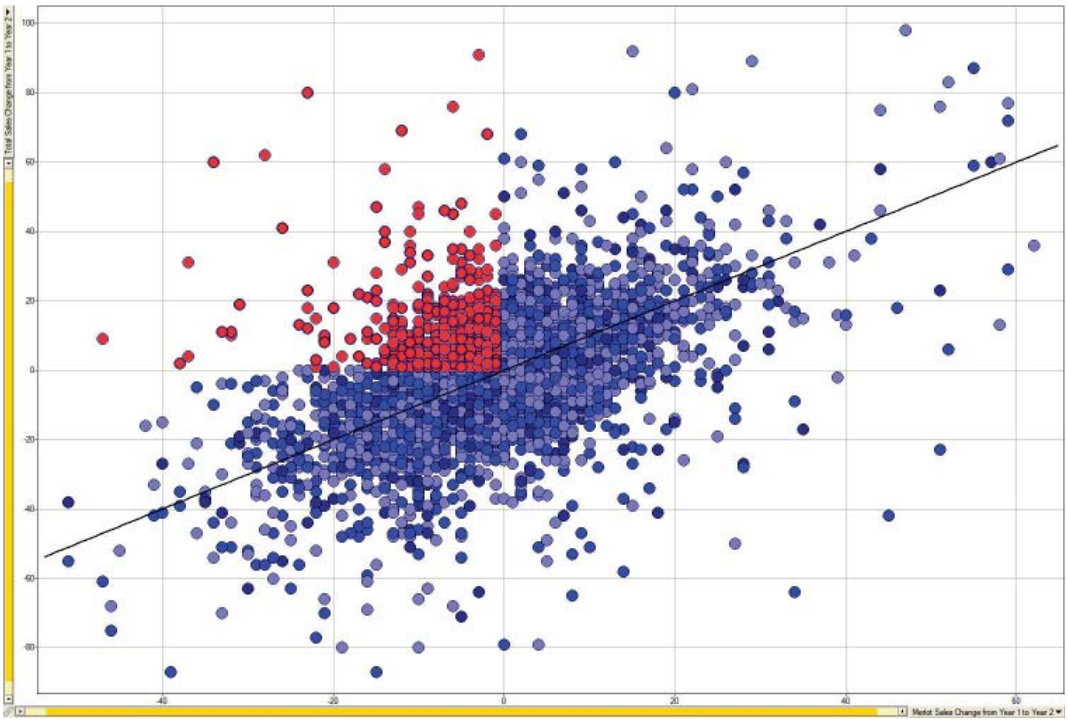

Figure 3: All values above zero on the Y-axis and below zero on the X-axis have been highlighted.

Figure 4: Brushing caused a subset of data that was highlighted in one chart to automatically appear highlighted in all charts. This example was created using Spotfire DecisionSite.

This action isn't very useful when viewing the scatterplot alone, but look what happened throughout the display when these values were brushed, shown above in Figure 4.

One of the first things that I discovered by brushing the data in this way was that these decreases in Merlot sales that were out of sync with corresponding increases in overall wine sales occurred to a slightly less degree in the west (especially seen in upper graph in the left column and the middle graph in the center column). By looking at each of the bars in the "Total Wine Sales this Year by State graph" (bottom middle) that lack red highlights altogether, I quickly learned that this situation did not occur at all in Delaware (the state with the greatest wine sales), Maine, Montana, New Mexico, and Washington D.C. I could also easily see states in which this occurred to a disproportionately high level, such as in Vermont (represented by the last bar to the right), where the ratio of the red section of the bar to its overall height indicates that at least 40% of wine orders in Vermont fell into this category.

Best Practices of Faceted Visual Analysis

Just placing multiple charts on a screen alone does not produce analytical discoveries such as those we've just seen. In other words, it is possible to produce a faceted analytical display that doesn't work very well. One typical problem that plagues other forms of data visualization as well, such as dashboards, is that the effectiveness of faceted analytical displays can be undermined by gratuitous visual effects, decorative display media, and poor use of colors. On the next page, Figure 5 exhibits some of these problems.

"The effectiveness of faceted analytical displays can be undermined by gratuitous visual effects, decorative display media, and poor use of colors."

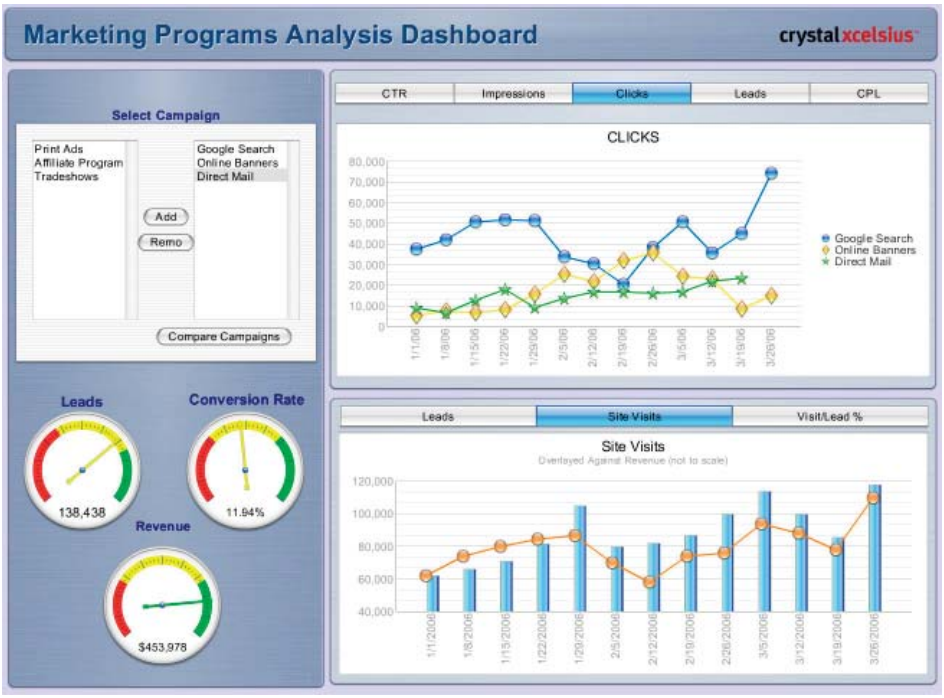

Figure 5: Poorly designed display media and distracting visual effects undermine the effectiveness of this display.

In this example, the gauges on the left use a great deal of space to provide far too little information. The graphs on the right are cluttered with large, flashy data points to mark the values along the lines, which make it difficult to see the shape of change through time. The bottom right graph is completely indecipherable. Which represents the site visits: the line or the bars? What does it mean that revenue is "not to scale"? Also, the fact that the two graphs on the right are not aligned so that the same points in time match up with one another makes is almost impossible to compare clicks (above) to site visits (below).

Perhaps the greatest problem in this example, however, is the fact that it restricts filtering to one choice only—the selection of marketing campaigns—which is controlled by using the list of campaigns that appears in the upper left corner. (By the way, don't you just love the "Remo" button?) Data analysis requires flexibility to filter the data in a variety of ways. For instance, it would be useful to filter out all but those campaign responses that generated leads.

Also, this example severely restricts the comparisons that can be made to one measure at a time in each of the two graphs. For instance, we could compare leads to clicks, or impressions to site visits, but we could never compare leads, clicks, impressions, and site visits, except as separate steps limited to one pair of measures at a time.

This same theme of hobbled analysis is repeated in the next example also (see Figure 6).

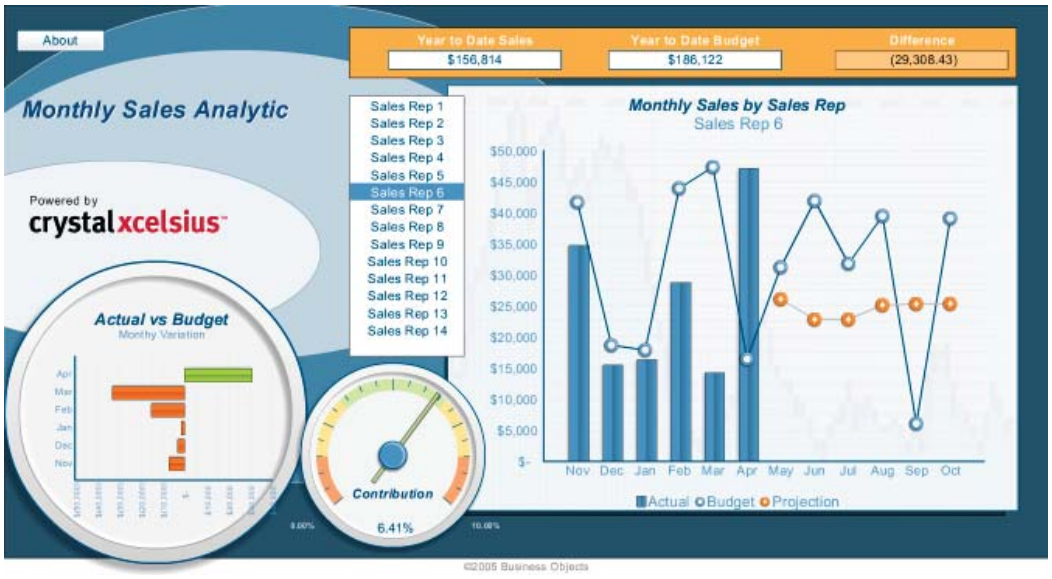

Figure 6: This display restricts what we can see at one time too severely, and also combines charts that are not synchronized.

The ability to only see one sales representative at a time, which entirely eliminates the ability to compare them, hardly lends itself to analysis. This display also suffers in another significant way, which is worth mentioning. Two of the charts—the bar graph on the left and the combination bar and line graph on the right—share a common period of time (November of one year through April of the next), but this synchronization is not obvious and the graphs are not arranged in a way that makes it possible to view deviation from budget (the graph on the left) in relation to the actual and budget values that appear in the other graph. The Actual vs. Budget graph displays the months in a horizontal arrangement from bottom to top, which is not intuitive. Time should almost always be arranged from left to right in chronological order. Also, the fact that some data sets are encoded as bars and some as lines is completely arbitrary. It would be much easier to compare the shape of change through time formed by each of these measures if lines were used to encode them all.

With the exception of silly display media (such as the gauges) and annoying visual effects, which come standard with this particular product, the problems that I've highlighted in these examples are problems of design. These displays were simply not designed effectively for analysis.

Data analysis—except for routine analytical procedures that are always done in the same exact manner—requires the ability to add, remove, and change views of the data on the fly, following the trail of analytical pursuit as quickly as questions arise. Any software tools that requires more than a few seconds for the analyst to add another chart to the screen or to modify an existing chart will not work for ad-hoc analysis. One of the ways that a tool such as Crystal Xcelsius from Business Objects fails to qualify as a serious analytical tool is the excessive work that is involved in creating the displays. For this reason, it functions more like a developer's tool for building analytical applications, rather than one that's designed for ad-hoc, free-flowing analysis.

I only know of a few visual analysis products that effectively support faceted analytical displays. In addition to Spotfire, which I used to create the good examples in this article, Advizor AnalystX does this well, and various SAS products also do the job but require extra skill and effort. Beginning in April with the next release of its visual analysis software, Tableau will complement its existing visual crosstab displays, which support rich multidimensional views of data, with the addition of faceted analytical displays. This will extend the power of this already fine product in ways that take it to the next logical level of its evolution.

Final Word

The primary points I hope you take from this article are the following four:

- Dashboards and multi-chart displays that are used for analysis are quite different.
- *Faceted analytical displays* provide a powerful means to see relationships in data.
- Global filters and brushing are powerful features that should be supported by any good faceted analytical display.
- No matter how great your tool, faceted analytical displays only work if you design them properly.

If you're an analyst and you don't already have a tool that allows you to prepare a rich visual feast from your data in the form of faceted analytical displays, I suggest that you find one and make it a part of your business intelligence toolset. It's time to move beyond the days when we were satisfied with a pivot table.

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.