# Three blind men and an elephant

## The advantage of multiple views

Author: Stephen Few, Perceptual Edge

January 2012

There's a well known folktale about three blind men who encounter an elephant for the first time and attempt to learn about it independently, by touch alone. The experience of each man is unique because each touches a different part of the elephant. This ancient story can teach us something important today about business intelligence (BI). Here's the tale:

> *One day, three blind men happened to meet each other and gossiped a long time about many things. Suddenly one of them recalled, "I heard that an elephant is a queer animal. Too bad we're blind and can't see it."*
>
> *"Ah, yes, truly too bad we don't have the good fortune to see the strange animal," another one sighed.*
>
> *The third one, quite annoyed, joined in and said, "See? Forget it! Just to feel it would be great."*
>
> *"Well, that's true. If only there were some way of touching the elephant, we'd be able to know," they all agreed.*
>
> *It so happened that a merchant with a herd of elephants was passing, and overheard their conversation. "You fellows, do you really want to feel an elephant? Then follow me; I will show you," he said.*
>
> *The three men were surprised and happy. Taking one another's hand, they quickly formed a line and followed while the merchant led the way. Each one began to contemplate how he would feel the animal, and tried to figure how he would form an image.*
>
> *After reaching their destination, the merchant asked them to sit on the ground to wait. In a few minutes he led the first blind man to feel the elephant. With outstretched hand, he touched first the left foreleg and then the right. After that he felt the two legs from the top to the bottom, and with a beaming face, turned to say, "So, the queer animal is just like that." Then he slowly returned to the group.*
>
> *Thereupon the second blind man was led to the rear of the elephant. He touched the tail which wagged a few times, and he exclaimed with satisfaction, "Ha! Truly a queer animal! Truly odd! I know now. I know." He hurriedly stepped aside.*
>
> *The third blind man's turn came, and he touched the elephant's trunk which moved back and forth turning and twisting and he thought, "That's it! I've learned."*
>
> *The three blind men thanked the merchant and went their way. Each one was secretly excited over the experience and had a lot to say, yet all walked rapidly without saying a word.*
>
> *"Let's sit down and have a discussion about this queer animal," the second blind man said, breaking the silence.*
>
> *"A very good idea. Very good." the other two agreed for they also had this in mind. Without waiting for anyone to be properly seated, the second one blurted out, "This queer animal is like our straw fans swinging back and forth to give us a breeze. However, it's not so big or well made. The main portion is rather wispy."*

> *"No, no!" the first blind man shouted in disagreement. "This queer animal resembles two big trees without any branches."*
>
> *"You're both wrong." the third man replied. "This queer animal is similar to a snake; it's long and round, and very strong."*
>
> *How they argued! Each one insisted that he alone was correct. Of course, there was no conclusion for not one had thoroughly examined the whole elephant. How can anyone describe the whole until he has learned the total of the parts.*
>
> (Kuo, Louise and Kuo, Yuan-Hsi, Chinese Folk Tales, 1976, Celestial Arts: Millbrae, CA, pp. 83-85.)

If I were to retell this story today to teach a lesson about BI, I might call it "Three blind analysts and a data set." People struggle every day to make sense of information, stumbling blindly, touching only small parts of the whole, which results in narrow, fragmented, and often erroneous understanding. Conventional BI tools make it unnecessarily difficult to explore data from multiple perspectives, so analysts tend to pursue only a limited set of predetermined questions and do so myopically. It is simply too time consuming to explore the data thoroughly, allowing fresh discoveries to lead them to free-flowing and comprehensive exploration. When their tools support data exploration and analysis poorly, people tend to accept the first answer that comes to mind, because it's too much work to look further. Without the ability to examine data from multiple perspectives simultaneously, many of the meaningful relationships that exist in your data will remain hidden.

This problem is becoming only more acute as we deal with the massive data sets of the modern day: ecommerce, retail and web operations are only a few of the domains that have seen an explosion of data. The use of mobile technology promises information access but adds more potential for people to see only a part of any given problem.

We stumble blindly and understand only in part, often because we are disabled by ineffective tools. Tableau software offers a solution. Tableau makes it so easy to shift from one perspective to another while exploring and analyzing data that we, as analysts, are encouraged to pursue every question that arises during the process, almost as fast as we can think of them. Because we are not distracted by the mechanics of using the software or forced to go through time-consuming steps to get from one view of the data to another, we can become immersed in the data and the analytical process. We are able to spend our time thinking about the data, not wrestling with the software.

In this paper, I want to focus on the insights that emerge when software enables us to view a data set from several perspectives at the same time. I'll use the term "coordinated multi-view display" to describe a screen that contains multiple concurrent views of a common data set so that comparisons can be made and connections can be seen between them. Some call these dashboards, but this is a term that I reserve for displays that are used for performance monitoring, not data analysis. Don't worry about the name, however. Whatever you choose to call it is fine. What matters is that you have the means to expand your analytical reach by viewing data in this way.

“ People struggle every day to make sense of information, stumbling blindly, touching only small parts of the whole ”

“ which results in narrow, fragmented, and often erroneous understanding. ”

# Data visualization and the human brain

The human brain is amazing. If you are at all aware of the current research that focuses on the brain and how it works, you know that this field is now experiencing exponential growth and discoveries as scientists take advantage of new technologies such as fMRI (functional magnetic resonance imaging) scans to observe the brain in action. Despite how powerful computers have become, much that humans can do quite easily may never be possible for computers. The reverse is also true. Many of the things that computers do best are unapproachable by the human brain, which is why we rely on them for lightning fast calculations and are happy to let them do those procedural repetitive tasks that we find so boring.

The human brain is extraordinarily good at doing some things and surprisingly limited in other ways. For instance, on one hand our brains make us extremely good at recognizing visual patterns (something computers can't do well at all), while on the other hand they are able to attend to and remember relatively little of what we perceive. When we think about things, trying to make sense of them, information is temporarily stored to support this process in working memory.

Working memory is a bit like RAM (random access memory) in a computer in that it is limited and designed for temporary storage. Compared to that hard disk drive that is built into your computer or attached to it externally, RAM seems very limited, but compared to working memory in the human brain, RAM seems enormous. Only three or four chunks of visual information can be stored in working memory at any one time. Three or four chunks! Information that comes in through our eyes or that is retrieved from long-term memory in the moment of thought is extremely limited in capacity. If all three or four storage slots are occupied, you must let something go from memory to allow something new to come in. When you release information from working memory, it can take one of two possible routes on its way out: 1) it can be stored permanently in long-term memory by means of a rehearsal process that we call memorization, or 2) it can simply be forgotten.

Perhaps the most fundamental activity of data analysis is comparison. Edward Tufte once wrote that the essential question of data sensemaking is "Compared to what?" Individual facts mean nothing in isolation. Facts become meaningful when we compare them to one another and see them in context. To say that quarter-to-date revenue is $1,383,593 means little until you put it into context through one of more comparisons, such as by considering it in relation to the revenue target of 1.5 million dollars or to the amount of quarter-to-date revenue that was earned by this date last year.

To compare facts, you must hold them in working memory simultaneously. Because we can hold so little in working memory at any one time, however, to do analysis effectively, we must rely heavily on external aids to memory. This sounds like an ideal job for a computer. Even a piece of paper that you write notes on to keep track of information as you're analyzing data is an external memory aid that is quite powerful despite being low-tech. A computer running properly designed software, however, can augment our ability to retain, access, and think about information much better than pen and paper.

Good visual analysis software, such as Tableau, can help us overcome the limitations of working memory in several ways. The goal is to enable as many enlightening comparisons as possible. Good software can support this goal by increasing each of the following:

1. The amount of information that we can compare (*more* data)
2. The range of information that we can compare (*more* dimensions)
3. The different view of information that we can compare (*more* perspectives)

Tableau software enriches our ability to compare data in each of the three ways just mentioned:

1. *More* data, by displaying data visually, which allows us to see and compare patterns and trends, and also allows us to chunk more data together into the limited number of storage slots in working memory.
2. *More* dimensions, by displaying data in a series of small graphs arranged as a visual crosstab, which allows multiple dimensions to be compared simultaneously, with less reliance on working memory.
3. *More* perspectives, by supporting coordinated multi-view displays to allow several perspectives to be rapidly examined and compared by quickly swapping them in and out of working memory.

## More data through visual encoding

Traditional BI relies mostly on tabular data displays. Tables are wonderful if you need to look up individual values, compare a single value to another, or know values precisely, but they do not display patterns and trends, nor do they highlight exceptions. This is a problem, because data analysis relies on our ability to spot and make sense of patterns, trends, and exceptions in data, and to compare entire series of values to one another. Take a look at the table in Figure 1 and compare it to the line graph in Figure 2 of the same data. Relying on the table to discern the ups and downs of sales through time and to compare the patterns of change from region to region would yield very little of the information that is immediately obvious in the graph. Visual representations give form to data, making patterns, trends, and exceptions easy to see.

| Market | January | February | March | April | May | June | July | August | September | October | November | December |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Central | 11,580 | 11,190 | 11,255 | 11,478 | 10,914 | 11,174 | 11,486 | 11,333 | 11,432 | 11,270 | 10,885 | 11,368 |
| East | 7,587 | 7,374 | 7,372 | 7,688 | 7,275 | 7,824 | 7,975 | 7,734 | 7,593 | 7,565 | 7,438 | 7,753 |
| South | 4,451 | 4,348 | 4,321 | 4,333 | 4,159 | 4,302 | 4,406 | 4,601 | 4,547 | 4,556 | 4,480 | 4,576 |
| West | 11,698 | 11,280 | 11,407 | 11,613 | 11,046 | 11,507 | 11,963 | 12,039 | 11,697 | 11,596 | 11,300 | 11,883 |

**Figure 1.**

*A simple table of monthly sales data, grouped by regional market.*

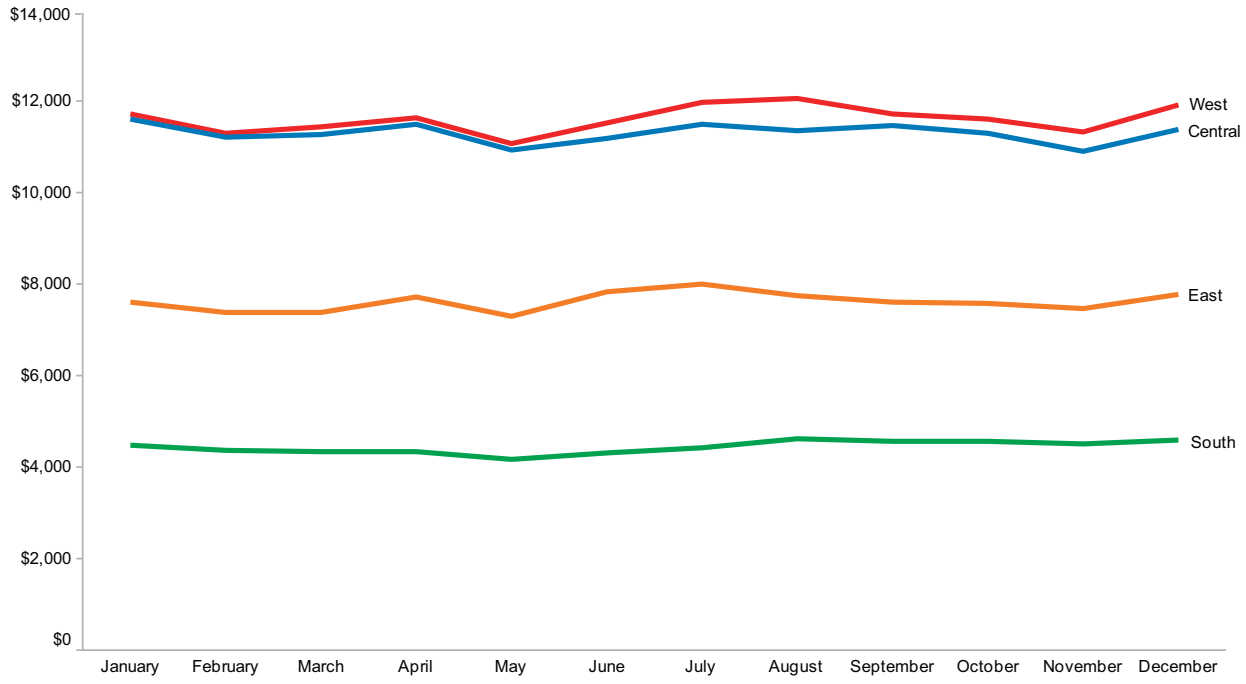


**Figure 2.**

*The same sales data that appears in Figure 1 is now displayed visually as a line graph.*

Another advantage of properly designed graphs over tables for analytical purposes is less obvious. If you needed to remember information in the table, you could hold only about four of the values (that is, four of the monthly sales numbers) in working memory at any one time. By relying on the graph, however, 12 values are combined into each of the four lines to form a pattern that you might be able to hold entirely as a single chunk in working memory. Because the values are represented as a simple visual shape, we are able to hold much more information simultaneously in memory.

## More dimensions through visual crosstabs

Tableau extends the benefits of data visualization further by organizing many graphs on the screen at the same time in the form of a visual crosstab. Figure 3 displays 24 small graphs and arranges them in the familiar crosstab fashion to present sales data across four different dimensions at once: products within product types by row, regions by column, and market size by the color of the line.

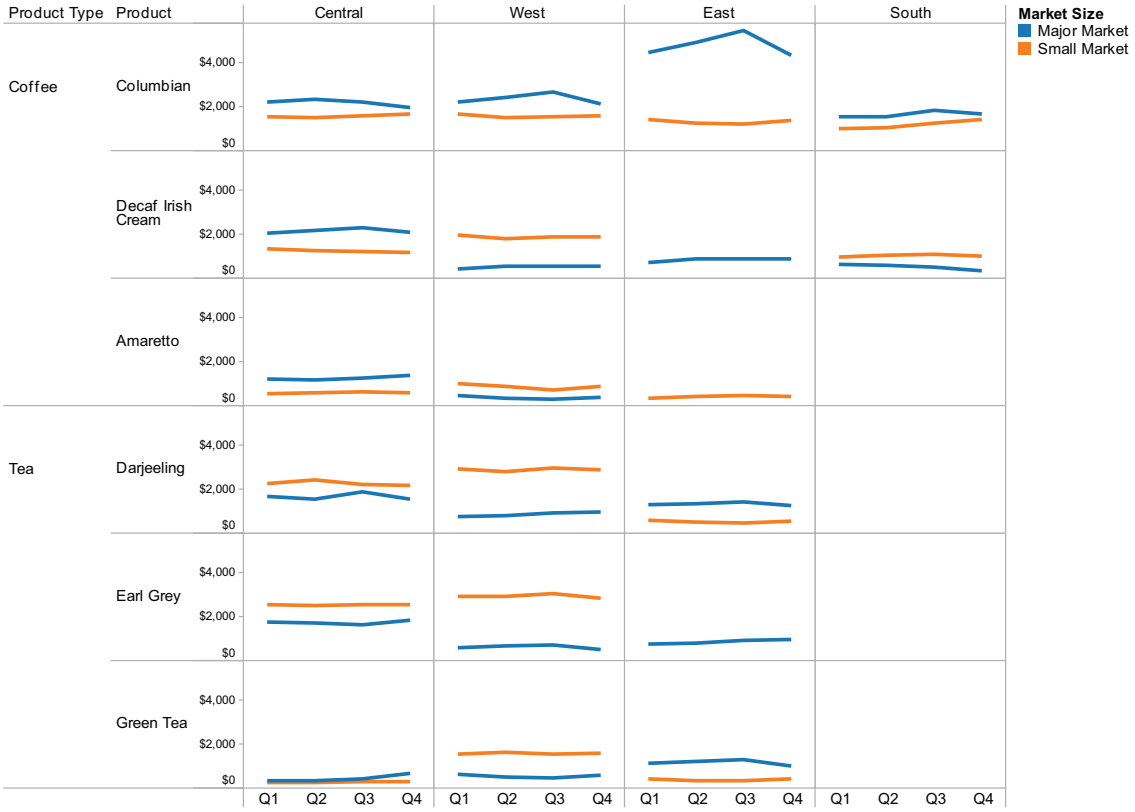


**Figure 3.**

*A visual crosstab.*

This approach makes a great deal of data available to our eyes and does so across several variables, thus expanding dimensionality well beyond traditional graphical displays.

## More perspectives through multi-view displays

To understand something, you must examine it from very angle and focus on various parts. Too much data analysis involves looking only for one thing in particular and doing so from a single perspective. Is revenue going up? According to this line graph, the answer is “yes,” end of story. Perhaps, however, you ought to look at expenses, profits, marketing campaigns, seasonality, composition of the sales force, new product introductions, and the competition to understand the richer story that your data has to tell.

Multifaceted data analysis like this is seldom done, not so much because its value isn’t obvious, but because the tools that analysts generally use discourage a thorough pursuit of the facts. It would be very difficult—perhaps impossible—to build a full picture of what’s going on in one’s mind if you were forced to examine each useful view of the data independently, separated by individual queries and the mechanics of constructing each new table or graph. This slow and fragmented process will resign you to the narrow and misleading world of the three blind men who encountered the elephant. You’ll never be able to stitch together a whole cloth from these

individual swatches, partly due to the limitations of working memory, but also because the cumbersome and time-consuming mechanics of the process will discourage any momentum that your inquiring mind tries to build. For your eyes to be opened and your mind to become enlightened, you need tools that can display the data from several perspectives simultaneously, allowing you to shift your gaze rapidly between those perspectives to spot connections and build a rich picture of what's going on.

## Multi-view displays

In their book Visual Statistics, Young, Valero-Mara, and Friendly explain why we must view data in multiple ways.

> *It is usually the case that a single display cannot show you everything that is interesting about a set of data. Even for the simplest data, we will often see different features of interest in different displays, if only because most graphics are designed to help us explore one (or at most a few) features. Thus, it is usually necessary to examine several different graphics in order to get the full picture of the problem at hand.*
>
> *(Visual Statistics: Seeing Data with Dynamic Interactive Graphics, Forrest W. Young, Pedro M. Valero-Mara, and Michael Friendly, John Wiley and Sons, Inc., 2006, p. 133)*

But why not use one view at a time in sequence, rather than placing several on the screen at once? When we want to compare information that's contained in multiple views, we must place them on the screen at the same time because of working memory's limitations. Good visual analysis software augments working memory by allowing us to place multiple charts on the screen at once that each view the data from a different perspective. This uses the screen as a means of external data storage that provides immediate access to the information. Anything that's in front of our eyes can be rapidly swapped in and out of working memory as needed. If we were forced to examine those same views sequentially, one at a time, we could only compare something that appears in the current view to three or four specific chunks of information at most that appeared in previous views. Analyzing data in this manner would take forever, resulting in fewer discoveries.

Multi-view displays can be used to solve various problems. The following five occasions are particularly useful and common:

1. The number of variables that we need to compare exceeds what can be effectively displayed in a single chart.

2. Our ability to examine the data would be undermined by visual clutter if displayed in a single chart.

3. We would be perceptually and cognitively overwhelmed using a single chart to perform the task at hand.

4. We must compare various levels of summary and detail (for example, the same year's worth of data expressed quarterly, monthly, and weekly).

5. We must examine the data from multiple perspectives that require different types of charts (line graphs, bar graphs, scatterplots, maps, tables, etc.).

To understand the strength of multi-view displays, you definitely need to see one. Figure 4 shows an example that I quickly threw together using the multi-chart functionality that Tableau calls a dashboard, but I prefer to call a coordinated multi-view display to distinguish this, which is used for interactive data explorations and analysis, from a multi-view display that is used to monitor what's going on, which is what I call a dashboard:

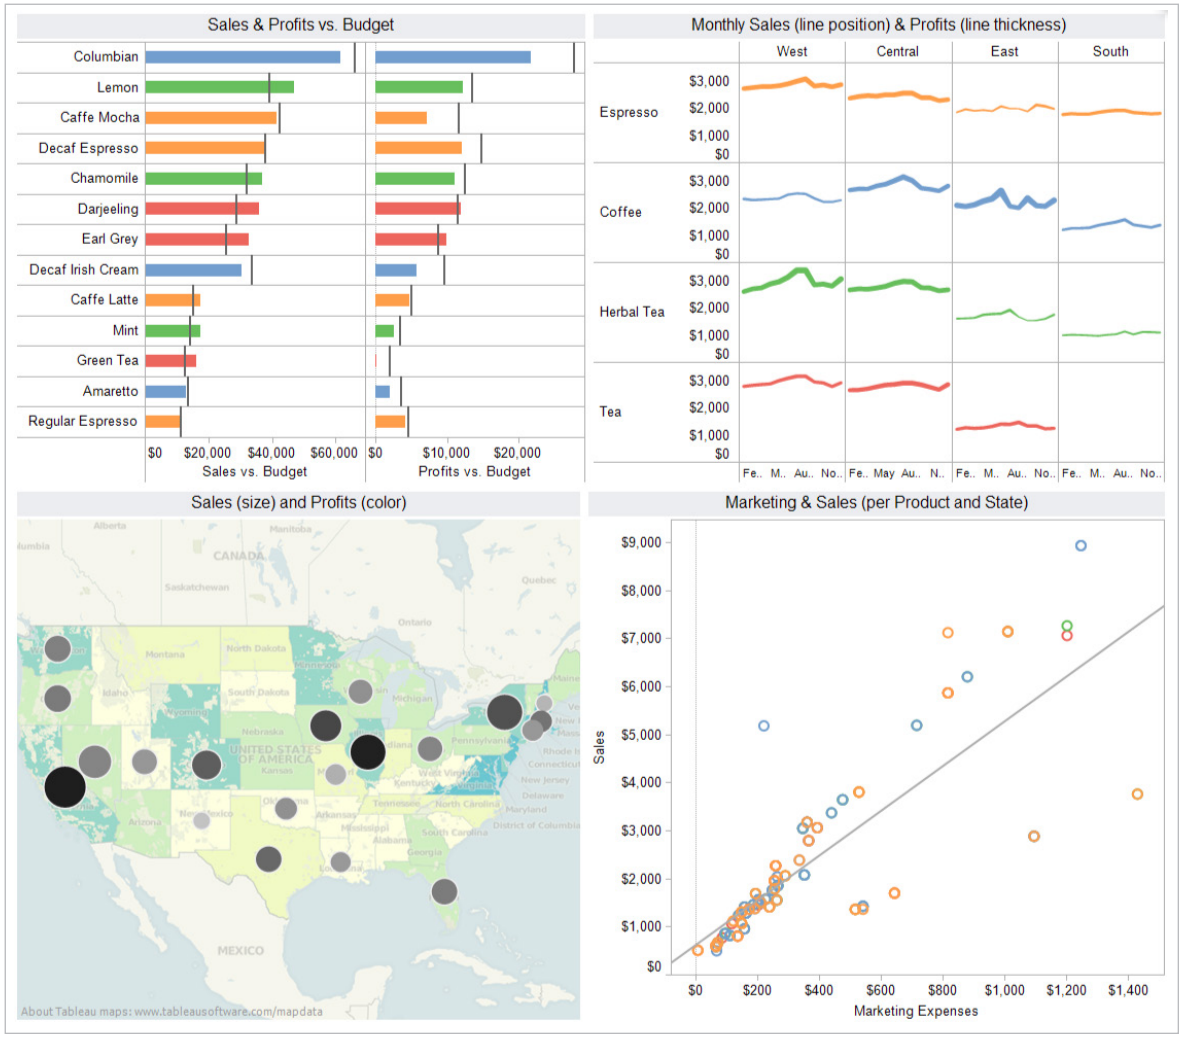


**Figure 4.**

*Example of a multi-view display.*

This one screen combines four views of a single data set; each focuses on a different aspect of the data. In the upper left corner you see two horizontal bar charts that work together to display sales vs. budget (the vertical gray lines represent budget amounts) on the left and profits vs. budget on the right. In the upper right corner you see a visual crosstab arrangement of line graphs that display monthly sales and profits by product types (the rows) and regional markets (the columns), sorted by total sales in descending order from top to bottom by product type and from left to right by region. Each line represents sales, based on vertical position, and profits, based on thickness. Even though we can't decode or compare profits precisely when encoded as line thickness, an approximate sense of profits combined in the same set of graphs with a precise

encoding of sales allows us to detect useful features, which we'll see in a moment. In the lower left corner you see both sales, profits, and per capita income by state on a map: bubble sizes represent sales, bubble colors represent profits (the darker the greater), and background colors ranging from yellow to blue represent per capita income. In the lower right corner you see a scatter plot that displays the correlation between marketing expenses and sales, grouped by product type using color, with a separate data point for each combination of state and product type. What we have here is a display that will allow us to explore a year's worth of sales data from several perspectives with little effort and no delay as we move rapidly from view to view.

We can discern several facts about sales and profits using this single display, even before beginning to interact with the data. Here are a few facts that I noticed with ease:

- Espresso sales were best overall, which I can see because the line graphs are sorted by product type in descending order of overall sales.
- Despite the fact that Espresso is the leading product type, the horizontal bar graphs show that the leading individual product in both sales and profits is Columbian Coffee, followed by Lemon Tea in the number two position. Espresso sales take the overall lead, however, because the third and fourth top-selling products both fall into the Espresso category (orange items).
- The bar graphs also reveal that, despite Columbian Coffee's overall lead, it is one of only two products that failed to meet its sales budgets (represented by the vertical gray line). The other is Decaf Irish Cream. Both are coffee products (blue items). Perhaps the person in charge of budgeting coffee sales doesn't handle budgeting as well as the other product managers. Of course, there are other possible explanations, so an explanation would require deeper digging. It's possible that the high sales performance of tea products (the red items) compared to budget might also be due to poor budgeting skills.
- Sales are highest in the West region, shown by the fact that it is in the first and thus highest-ranked column of line graphs, which is probably due to the fact that sales in the state of California lead the nation, as we can see by the large size of California's bubble on the map. Although the Central region does not perform as well as the West overall, it outperforms the West in coffee sales, which can also be seen in the line graph. Coffee sales increase in the Central region in the month of July, as they also do in the East, and in the South as well, but a month later. A similar summer peak in sales does not occur in the West for coffee, however, but it does for herbal tea.
- In general, the scatter plot reveals a positive correlation between marketing expenses and sales, and no single product type stands out as being better or worse than the others in this respect. The scatter plot also reveals a few outliers in the data—data points that seem far removed from the norm—and there appear to be two separate groups of data points that form linear positive correlations, revealed by the one linear series of points above the trend line and another below it.
- The fact that profits can be seen in the line graphs based on thickness in addition to sales causes some interesting patterns to emerge. For instance, notice how Coffee profits in the West appear to be considerably less than expected given the high sales, whereas in the East profits are higher than they are in the Central region, even though sales are lower. Notice also that Espresso profits in the East seem unexpectedly low.

“ It is usually the case that a single display cannot show you everything that is interesting about a set of data... ”

“ Thus, it is usually necessary to examine several different graphics in order to get the full picture of the problem at hand. ”

– Visual Statistics, Young, Valero-Mara and Friendly

The story that this display reveals, even without interacting with the data, is richer than this, but these are the facts that caught my eye noticed with little effort.

Despite their potential benefit, a simultaneous display of multiple views is not a panacea. Multi-view displays can sometimes tax the mind to a counter-productive degree. There is no simple rule of thumb such as “more than four simultaneous views are too many.” The point at which a multi-view display becomes too complex is determined by several factors, including the nature of the data, the number of variables and values in each view, the nature of the task, the expertise of the person examining the data, and last but perhaps most important, the design of the display. A single chart with few variables and meager data could overwhelm our brains if it is poorly designed, while a well-designed display that incorporates many charts could be perfectly usable.

Two general guidelines for the visual design of multi-view displays are especially helpful:

- Keep everything that isn’t data to a minimum.
- Maintain consistency among views wherever possible, except when doing so might suggest connections that don’t actually exist in the data.

When combining multiple views on a screen, because we’re adding more visual content for our eyes to process, it’s important to eliminate all non-data content that isn’t necessary to support the display of data in a meaningfully way. Axis lines on graphs are a common example of non-data content that meaningfully supports the display of data by clearly indicating the boundaries of the plot area where the data appears and by providing straight lines to which tick marks and labels can be attached, which tell us how to interpret the data. Wherever non-data content such as axis lines are included, they should be only visible enough to do their job effectively and never so visible that they compete with or distract from the data. Unnecessary and overly salient non-data content should be avoided even in a single chart, but in a multi-view display like this, the practice is critical. Visual analysis tools that keep non-data content to a minimum as a design default make it easy to follow this useful practice. This is important, because you don’t want to waste time arranging and formatting the display to reduce non-data content; you want to remain in the flow of analysis without distraction or interruption.

## Coordinated multi-view displays

Multi-view displays really come alive when you begin to interact with the data in a coordinated manner that affects all views simultaneously through a single interaction, such as by filtering or highlighting. Let’s look at filtering and highlighting independently to see how they can bring facts to light that might otherwise remain hidden in the shadows.

### Coordinated filtering

Filters in Tableau can function either locally, affecting a single view (that is, a single chart), or globally, simultaneously affecting all charts that appear on the screen. The ability to filter all views in the same way simultaneously can be especially revealing. Let’s look at a few examples, using the same multi-view display as above. We’ll start with something simple.

What happens if we remove California—the state with the highest sales—from the mix? When I unchecked California in the state filter to remove it from view, I was able to watch the charts change before my eyes, resulting in the version shown in Figure 5.

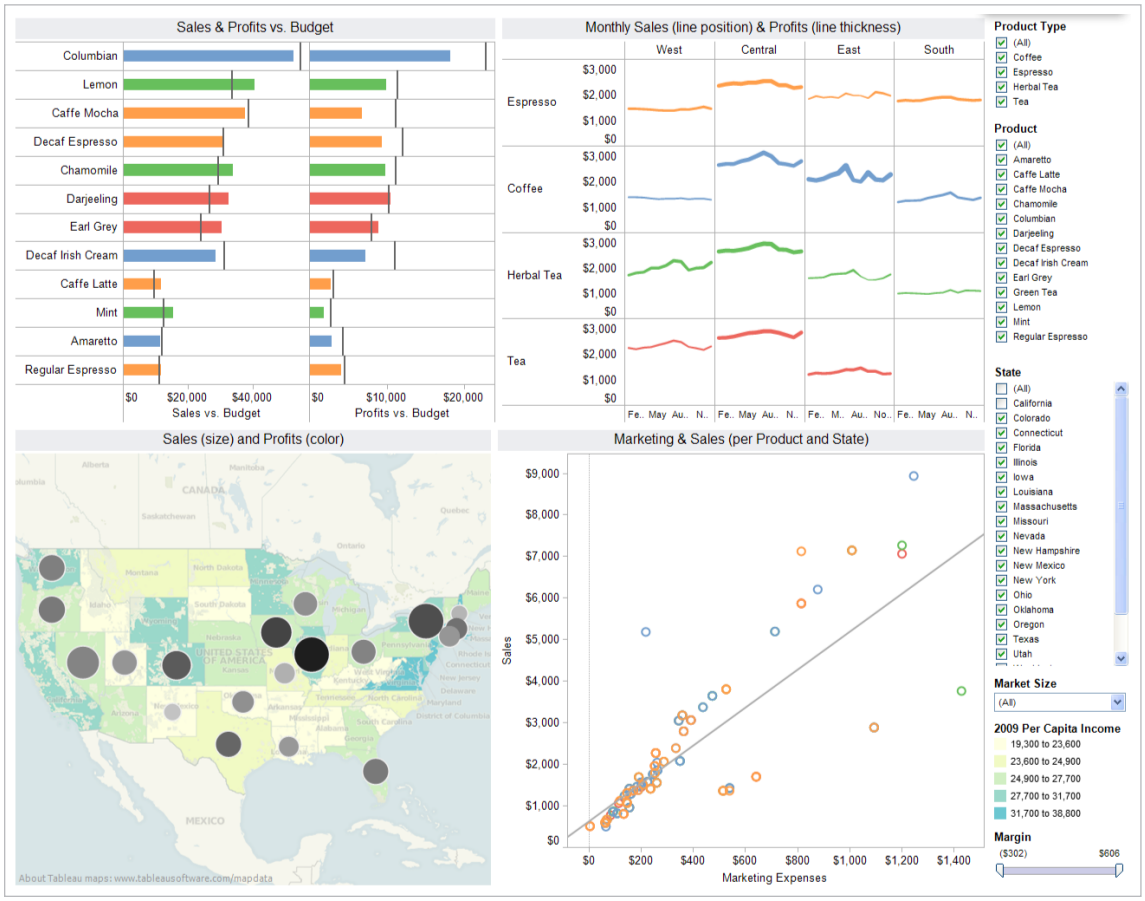


**Figure 5.**

*The state of California has been filtered out of this coordinated multi-view display.*

The first thing I noticed was that, without California, the relative order of product sales was rearranged. I intentionally set the bar graphs to not re-sort automatically when filters are applied. This make it easy to spot when the leftmost set of bars (sales vs. budget) are no longer in order by size, which indicates a change in ranking. To see the new order of products ranked by sales with California excluded, I could simply click the sort icon on the toolbar.

Let's put California back in the mix and see if we can learn anything about Major Markets vs. Small Markets by filtering out the latter. Figure 6 show data for Major Markets only:

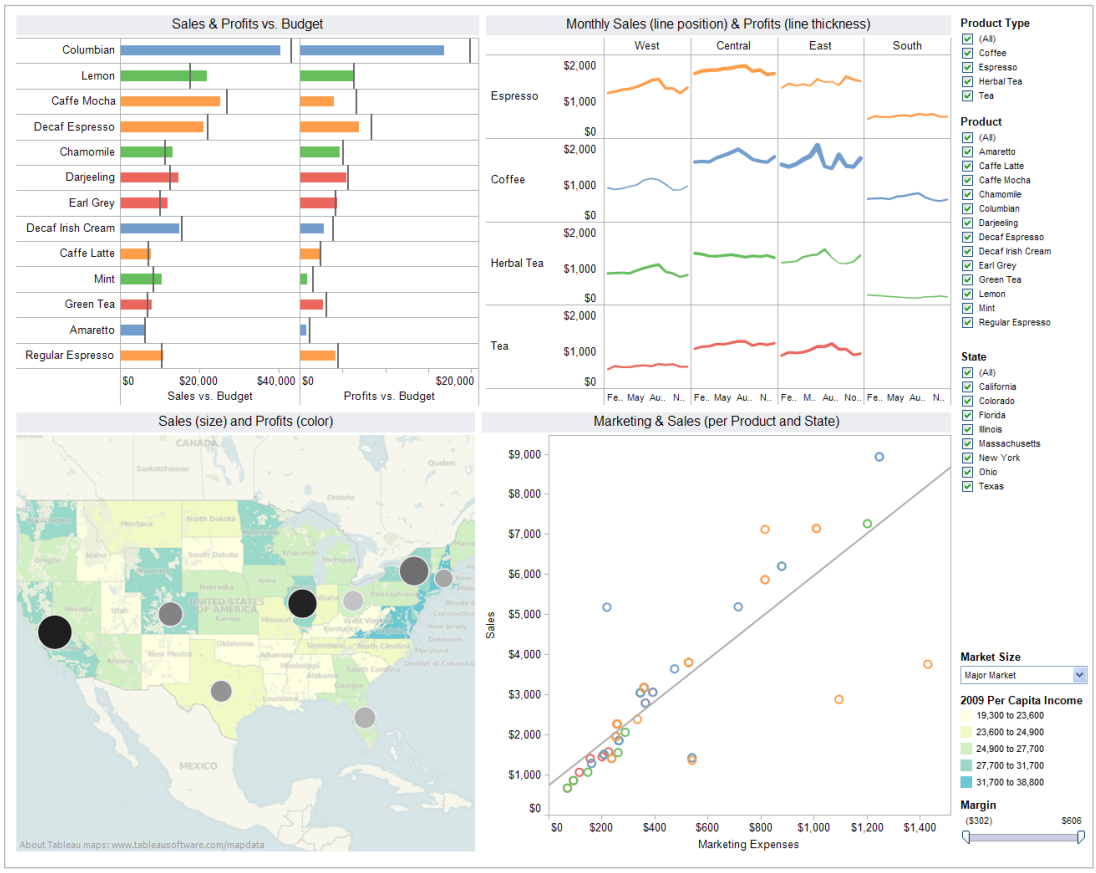


**Figure 6.**

*Small markets have been filtered out of this display.*

Notice that the map now shows bubbles for states in Major Markets only. Once again, we can see that the order of products ranked by sales in the bar graphs is different when Small Markets are excluded, and something else has changed as well. No long are the profits earned by Decaf Irish Cream, Green Tea, and Amaretto below budget, but now Mint Tea, which was doing fine overall before, is failing to meet its budget.

Let's put the Small Market states back in the mix and use filters to examine Margin. Notice in Figure 7 what happened when I filtered out all but sales with margins of $400 and greater:

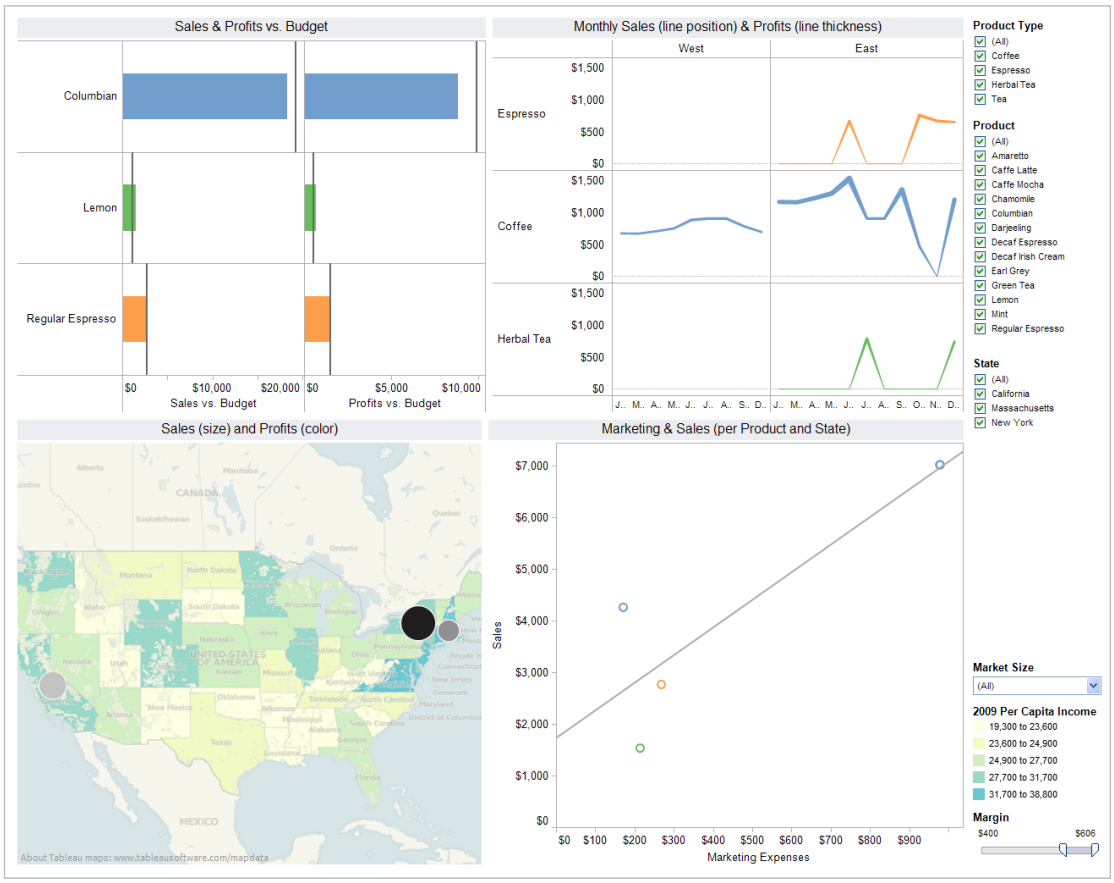


**Figure 7.**

*All but sales with the highest margins have been filtered out of this display.*

We can see in the bar graphs that only three products remain: Columbia Coffee, Lemon Tea, and Regular Espresso. We see on the map that only three states remain, and that New York now outperforms California in sales and profits. We see in the scatter plot that Coffee products are being sold with high margins in two states, which I've annotated to reveal as Columbian Coffee in Massachusetts and California. Regular teas—products in red, which no longer appear—have been eliminated altogether. What's perhaps most eye-catching, however, is what's happened to the line graphs: the Central and South regions have been eliminated, only Coffee products remain in the West, and in the East region Espresso and Herbal Tea products only appear during the latter half of the year.

Had we been relying on individual charts viewed independently from one another rather than a richer multi-view display, most of these facts that we've discovered quickly and easily might have remained unknown.

All of these filtering examples so far have illustrated the use of global filters, which allow us to remove the same subset of data simultaneously from all of the charts. As an alternative to this approach, let me show you something that you can do with a multi-view display that involves local filters. In Figure 8, notice that the same visual crosstab of line graphs appears twice: one above and one below. Both present monthly sales by Product Type and Region. The sales and margin filters in the top right affect only the top view and those below affect only the bottom view.

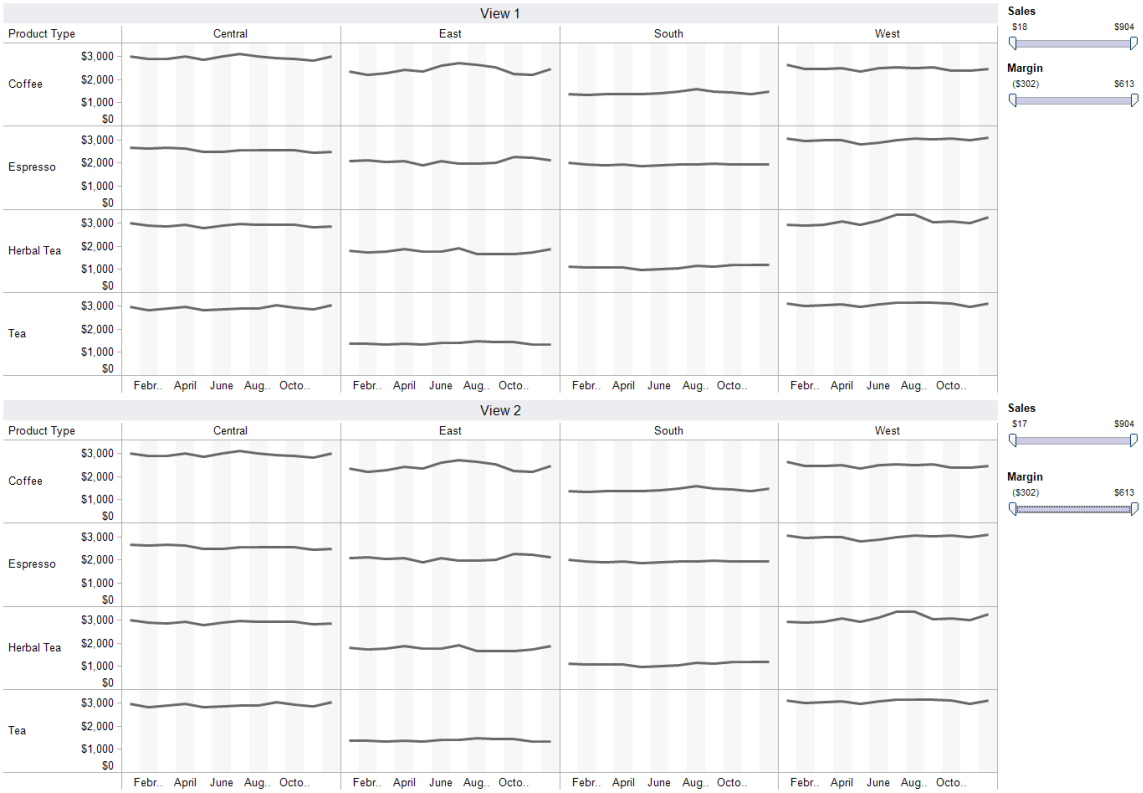


**Figure 8.**

*This display shows the same view of the data twice, but allows separate filters to be applied to each.*

A display of this type allows us to compare different subsets of sales by filtering each of the graphs differently. For instance, look at what happen in Figure 9 when I left the upper view alone, without applying any filters, but filtered out all but sales with margins of $100 and greater in the lower view:

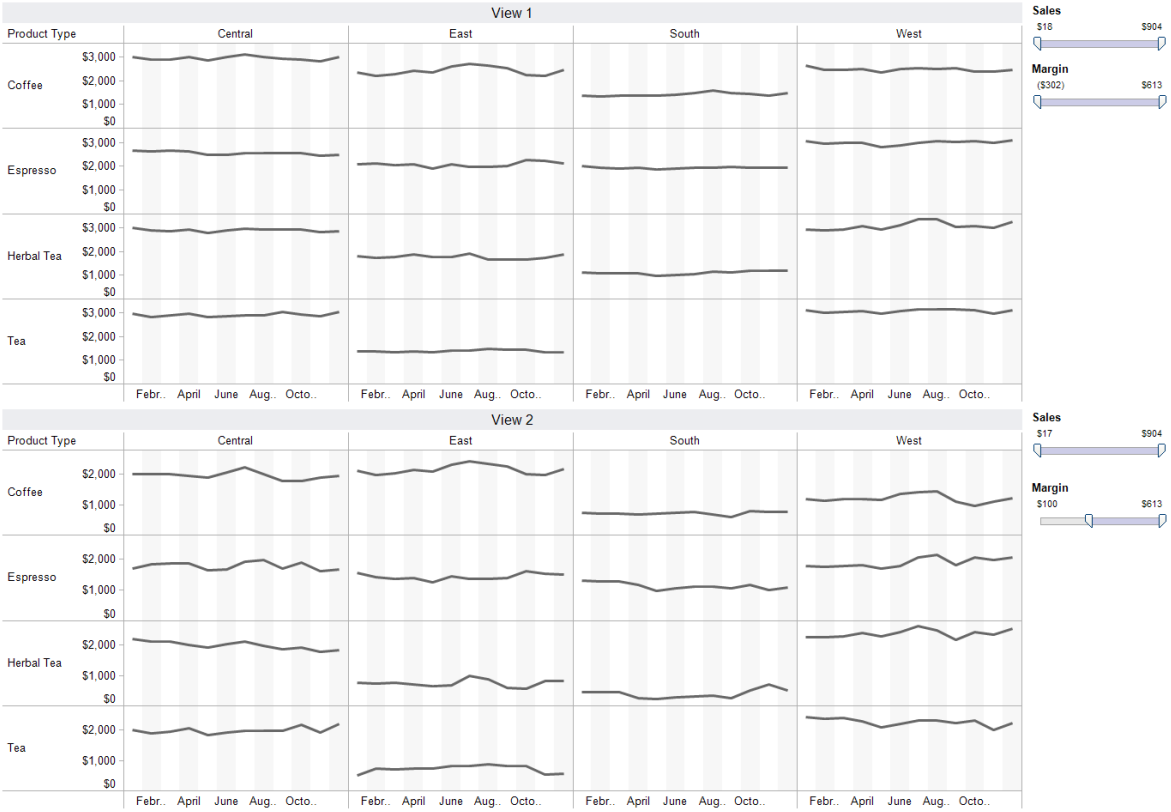


**Figure 9.**

*Margins less than 100 have been filtered out of the lower visual crosstab (view 2).*

By displaying these two versions of the data simultaneously within eye span, we can make subtle comparisons with ease. For instance, we don't need to look beyond the leftmost graph in the top row of each view to see that the elimination of sales with low margins has caused the peak in Coffee sales in the Central region that occurred in April to disappear completely.

## Coordinated highlighting

Filtering is not the only way to focus on a specific subset of data, and at times it's not the best way. A potential downside of filtering is the fact that we can easily lose a sense of the whole when we restrict our view to particular parts. Sometimes it's necessary to focus on particular parts without losing a sense of the whole. We can do this by highlighting those parts or by dimming everything else, without removing any data from view. When we're using multi-view displays, coordinated highlighting causes selected items to be featured in all charts automatically and simultaneously. In Figure 10, a multi-view display of healthcare costs, I selected the "Hypertension" bar in the upper-left chart, which caused data related to hypertension to be automatically highlighted in all of the charts.

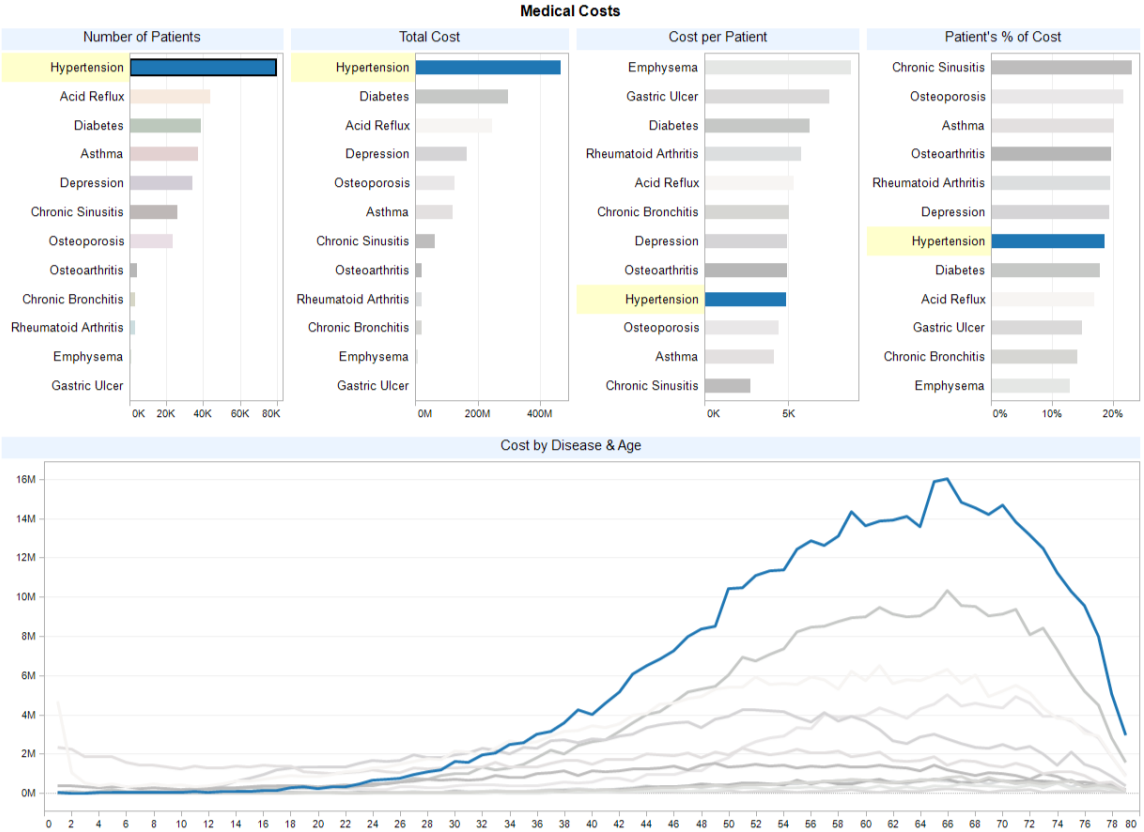


**Figure 10.**

*Costs of treating hypertension have been highlighted in all of these charts.*

Had I filtered out information related to everything but hypertension, the ability to see how much greater the number of patients treated are for hypertension, along with the costs, compared to other conditions, and how much greater it affects people in the second half of life than any other condition, would have been missed.

In Figure 11, the same set of 2009 house sales in Seattle are displayed from two distinct but complementary perspectives—a scatter plot that features house sizes and prices and a map to feature their locations. In both views, the sizes of the circles represent when the houses were built (from 1940 through 2009): small circles are old houses and the largest are the newest.

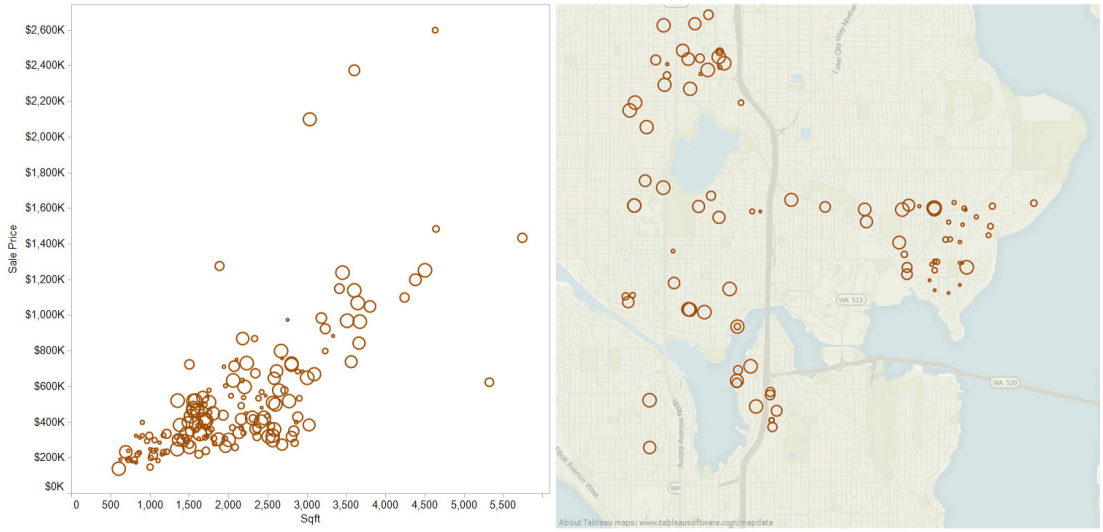


**Figure 11.**

*This multi-view display includes scatter plot and a map.*

Because these two views are linked, we can interact with them in a coordinated manner. For instance, we can select the five houses with sale prices greater than $1,400,000 in the scatter plot to simultaneously highlight them in both charts to see where they're located without losing awareness of the overall context (that is, the entire set of houses), as shown below.

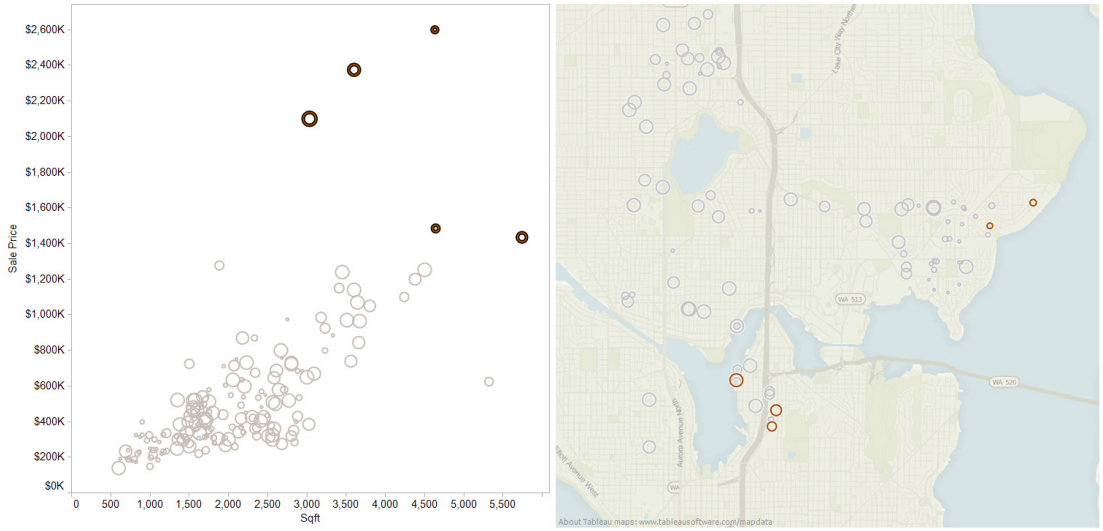


**Figure 12.**

*Selecting houses that sold for more than $1,400,000 in the scatter plot caused them to be highlighted in both charts.*

We can see that these houses are located in two geographically areas. What would not be obvious had we filtered out all houses except the five with highest sales prices is how much more expensive these houses were than most and the fact that all but one are located nearby at least one other house that sold.

Another question that we could ask is how similar the houses were to one another in size and price, which sold in the area near the highway below the water. I selected those houses in the map to highlight them in both charts, shown in Figure 13.

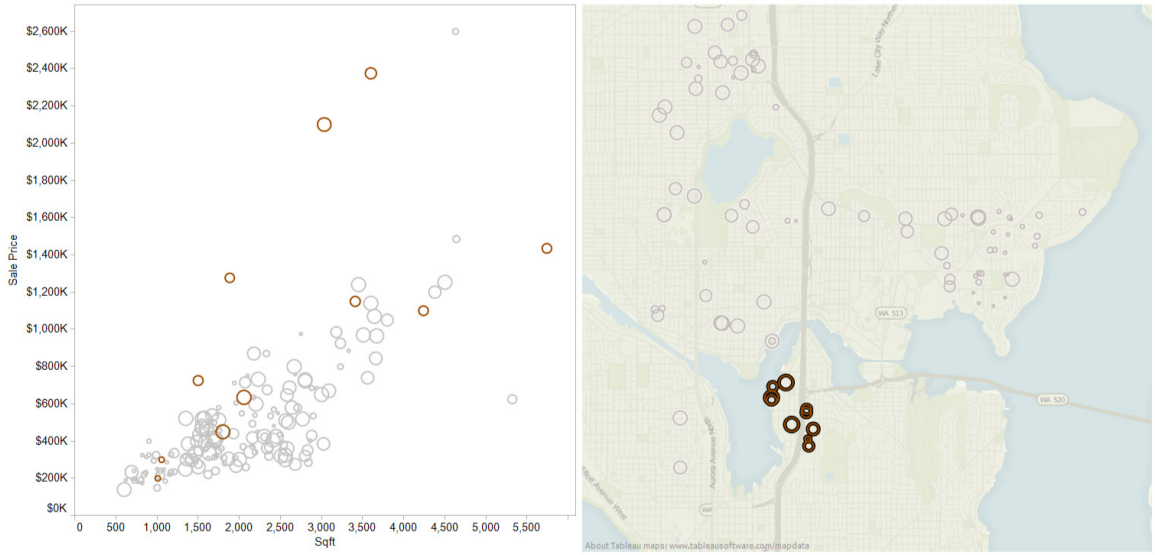


**Figure 13.**

*Selecting the houses in a particular geographical region caused them to be highlighted in both charts.*

What is now instantly obvious in the scatter plot is the fact that those houses were significantly different from one another in size and price, spanning almost the entire range across which the full set of houses are distributed. Had we filtered out all houses but these, we might have missed this fact.

# Final thoughts

What I hope is becoming obvious with these examples is that using coordinated multi-view displays can lead to insights that would be difficult to unearth otherwise. The real proof, however, will become clear only when you try them out on your own, using your own data. Every new fact that you discover could be one that provides real benefits to your organization.

# About the author

Stephen Few has worked for nearly 30 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 speaks and teaches internationally, writes the monthly Visual Business Intelligence Newsletter, and is the author of three popular books: Show Me the Numbers: Designing Tables and Graphs to Enlighten and Information Dashboard Design: The Effective Visual Communication of Data, and Now You See It: Simple Visualization Techniques for Quantitative Analysis.

# About Tableau

Tableau Software helps people see and understand data. Ranked by Gartner in 2011 as the world's fastest growing business intelligence company, Tableau helps individuals quickly and easily analyze, visualize and share information. With more than 6,500 customers worldwide of all sizes and across industries, Tableau is used by individuals throughout an organization, in an office and on-the-go. See the impact Tableau can have on your data by downloading the free trial at *www.tableausoftware.com/trial*.