

Financial Visualization Case Study: Correlating Financial Timeseries and Discrete Events to Support Investment Decisions

Eugene Sorenson¹, Richard Brath²

¹ Bloomberg L.P., ² Oculus Info Inc.

{esorenson1@bloomberg.net, richard.brath@oculusinfo.com}

Abstract

In financial timeseries visual data analysis, there are many potential data elements to display, including continuous quantitative data and discrete event data. Bloomberg has designed and implemented a scalable visual representation for the depiction of many discrete timestamped events in use by hundreds of thousands of financial markets experts. This visualization enables a single screen to visually organize a large volume of event data; to facilitate inference through visual alignment of related data; and to provide a workflow from the single point of access to a wide variety of detailed information. It allows an order of magnitude more event types to be associated with timeseries data.

Keywords---Timeseries visualization, financial visualization, multi-attribute glyphs.

1. Introduction

Participants in the capital markets are confronted with the challenge of considering many discrete pieces of information when analyzing a security and making an investment decision. For example, users of the Bloomberg Professional® service have access to tens of thousands of stocks, four million fixed income securities, millions of indices and economic data series, both historic and real-time and across many variables. The volume of transactions and information continues to grow. Furthermore, the variety of data includes financial news, analyst reports and audio/video segments from thousands of sources as well as mandatory filings, government economic data, management personnel, legal dockets, natural disasters, user-defined alerts and more. There is an on-going need to design and develop effective user interfaces to enhance users' abilities to exploit this data.

All this information is available interactively and a significant portion updates in real-time. Much of the value of this data depends on speed, and visualization can provide a key role to aid faster comprehension than traditional rows and columns of data, or simple charts and graphs. Bloomberg has a number of different visualization projects to improve the comprehension of complex

financial data and associated financial applications, such as derivative valuation, portfolio monitoring, sector rotation, etc.

This paper describes one significant visualization project in 2006-2007 combining the visual depiction of continuous timeseries data (e.g. stock prices) together with associated discrete event data (e.g. news, earnings releases, announcements) as multi-variate glyphs:



Fig. 1. A financial price timeseries with discrete events aligned along a common time axis provides visual access to possible causes of price movement

The prior technique displayed an alphanumeric label representing the event adjacent to a timeseries line. The technique originally worked well when the event selection was limited to earnings reports and corporate actions. However, as we focused on adding many more events and event types it was clear this technique would not scale beyond 36 unique item types (i.e. limited to alphanumeric characters) and could result in occlusion when there were many near simultaneous events:



Fig. 2. Alphabetic characters superimposed over a price chart indicate events. The technique will not scale to greater event density due to occlusion.

Going forward, a more general technique was required that would scale to depict more event types, greater volume of events, reduce occlusion and facilitate workflow.

2. User Needs

Line charts and other related high-frequency charts plotting price over time (e.g. candle charts) are very heavily used in financial services. A user can view a chart to validate an assumption (e.g. is the price still trending); look for potential behaviors and patterns (has the price re-tested a high and failed indicating a new potential downtrend); etc. When there is an unexpected anomaly (e.g. unexpected price drop) the immediate question is “why” and an explanation is sought. At this point, the viewer needs to understand what other events happened at that point in time, for example:

- Was there a significant news story?
- Did a research or credit analyst change a rating?
- Was there an unexpected result in newly released economic data?
- Did the company have a significant action, such as an acquisition or investor presentation?

More complex investment hypotheses emerge when considering event patterns over time. For example:

- Regular patterns between event occurrence and price movement, e.g. Is there a consistent price response each time this company makes a public announcement regarding earnings?
- Event sequences and price movement, e.g. If a particular analyst’s commentary on this stock has the potential to impact the price, are there any economic events that immediately precede and influence that analyst.

The product team collected user feedback and enhancement requests through informal interviews with users and from comments collected by the support desk. This user feedback included:

- Overlapping alphabetic characters were difficult to read.
- If characters overlapped the price line, critical price information could be obscured.
- Ability to access more detailed information, e.g. via tooltips or workflow.

Meetings among different business groups indicated potential demand for many different tasks and types of events. For example:

- *Monitoring*: Traders were interested in monitoring their trades in relationship to price movement in order to better manage and assess their trading activity on a one day chart as well as the impact of significant news or corporate actions, such as earnings reports.
- *Exploration*: Economic analysts were interested in reviewing statistical agency announcements (e.g. GDP, unemployment, factory orders) and associated price movements to explore potential trends over a recent period such as a year or two.
- *Evaluation*: Sales professionals were interested in seeing their stock recommendations versus the actual market price over a period of a few months to assess the quality of their forecasts.

In addition to the immediate need regarding the representation, there were many additional constraints, e.g.:

- *Workflow*: The visualizations are intended to be part of a broader workflow. They must co-exist with other elements in the interface. Further, as the visuals represent large amounts of data, they need to aid navigation.
- *Small-space*: The visualizations need to be usable from full-screen down to a small size, including a mandatory 480 x 320 size; and even smaller as part within a user-configurable tiled-window desktop called Launchpad.
- *Clear, Immediate Comprehension*: All screens in the system must provide immediate useful information. It was not an option to require one or two clicks of the user to configure at start.
- *Fast Start-up*: Financial markets require speed throughout the application. Screens in professional systems must launch in a second, completely populated with data.
- *High-Performance*: The visualization needed to update at sub-second intervals, ideally it should be able to achieve 30 frames per second.
- *Low CPU*: A trader may have many different applications, all updating in real-time, running on a desktop with as many as 8 screens. The visualizations could not consume more than 5% CPU at peak-load (e.g. start-up and fast moving markets) on lowest common denominator hardware.
- *Rapid product cycle*: At Bloomberg, new information products are typically conceived, designed, implemented, tested and deployed in months if not weeks. The financial industry is competitive and fast-paced: therefore, the product processes are similarly rapid [1].

3. Background and Related Work

Visualization is not new to professional users in the financial markets. Bloomberg’s solution needed to accommodate the existing professional user community and respect techniques they were already familiar with.

3.1. Financial Visualization Literacy

Market professionals have long relied on high-frequency charts of time series data with time plotted horizontally and stock price (or other continuous variable) plotted vertically. This may be more simply described as a “price chart”. The financial industry’s focus on returns places greater demand on understanding a security’s current price in the context of its history. As a result there is far greater use and familiarity with time series charts versus other types of visual representations.

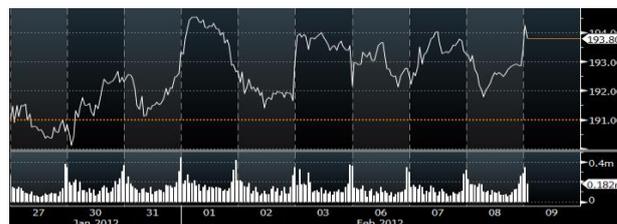


Fig. 3. Sample price chart over 10 days showing a

security's price in the upper panel and trade volume in the lower panel.

Some market professionals, such as technical analysts, utilize many unique variants of price charts, including open-high-low-close charts, candle stick charts, point and figure, renko, market picture charts, etc. Also, there is a long history of visual analytics with tight integration between charts and analytical functions, such as averages, ratios and spreads, with particular configurations known by acronyms such as MACD (moving average convergence/divergence), DMI (directional movement indicator), RSI (relative strength indicator), etc.



Fig. 4. Sample price chart over 6 months showing common technical studies, including moving average lines (yellow and green lines) and MACD lines (bottom panel).

Thus there are established conventions for timeseries representations within the professional financial markets. Within this community, all users are familiar with the basic conventions and smaller groups of users have acquired deep expertise in specialized charts and analyses.

3.2. Related Work

The previous technique to represent events was to display a single alphanumeric character as close as possible to the line in the chart, without obscuring the line (figure 2). Many other financial applications used the similar techniques:

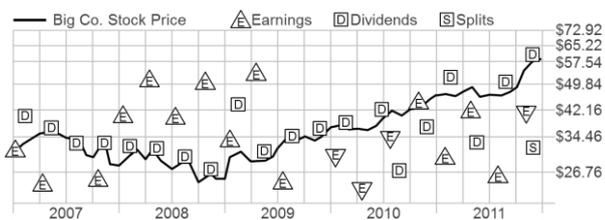


Fig. 5. Example chart showing a single character per event (e.g. E for Earnings). The seemingly random placement of the event icons throughout the plot area increases the sense of clutter and could reduce speed of comprehension of the timeseries price line. (wireframe drawing based on chart from wsj.com).

Another approach was popularized by the stock charts in Google Finance that launched in early 2006 [2]:

- Each event can have a unique alphanumeric character. The character is attached via a leader line to the chart, as if a flag.

- Details for each event can be elaborated in a side panel, referring to each character.
- In areas of density, the flag can be stacked to indicate high density of events. Zooming in will show each flag individually.



Fig. 6. Example chart similar to financial charts from Google Finance, showing alphabetic flags in the chart, and the corresponding news item in the adjacent panel on the right (wireframe drawing based on image from [2]).

The Google approach was considered very innovative but possibly had limitations that may not have been acceptable to the professional user:

- The approach requires a unique event label per each event. This approach was either limited to 36 individual alphanumeric labels or required more complex labels which would increase marker size and potentially increase perceived clutter.
- The approach does not easily differentiate between types of events. Google has addressed this somewhat in later versions by creating different shapes and colors of markers for event types such as dividends and stock splits.
- The approach requires the viewer to shift gaze back and forth from the chart to the list, scanning the chart and scanning the list with each shift, implying a potential perceptual inefficiency.

This last point is similar to Larkin and Simon's [3] conclusions that (1) visually grouping information together that is used together avoids the search effort and (2) removes the need to match symbolic labels, leading to reductions in working memory and search effort.

Another common approach to organize and display related information in a singular interface is the portal. A portal simply puts all the different information together across a series of separate view tiles with no explicit linkages across views. Finding related data at a particular time point requires the user to exhaustively search through all the portal tiles (often more tiles than can be arranged in a single display screen) and perform mental arithmetic to evaluate whether the timestamp associated with a piece of information is in relevant timeframe or not before evaluating the information.

4. Design, Implementation and Feedback

The project went through an initial design phase with feedback primarily from internal stakeholders and was followed by an implementation phase.

4.1. Requirements and Design

The approach consisted first of assessing all the potential event data items that needed to be addressed. Ini-

tially, 4 different categories of events were required with approximately 25 different unique types of events. The design approach also needed to be flexible for future extensibility, on the assumption that users would widely adopt the newer design. (This was not a given).

In assessing the requirements, we considered that our design approach needed to fulfill Jacques Bertin's [4] three visualization functions:

1. **Recording information** i.e. a visual representation that acts as an inventory, where everything can be collected in one place, e.g. a road map.
2. **Processing information** i.e. an environment for visual analysis facilitating visual inferences through perceptual mechanisms, e.g. a scatterplot with color-coded points; and interactive techniques such as filtering.
3. **Communicating information** i.e. a visual to convey a message, e.g. an infographic or narrative visualization.

We felt fulfilling Bertin's three functions required first and foremost a successful execution of the inventory of all timeseries information in a single graphic. In effect, the goal was to create the graphical equivalent of a roadmap for timeseries data. The visualization needed to create a frame of reference where any kind of event data could be placed and a user would know what the event was and what other events were associated with it of the same type or at the same time.

In addition to Bertin's functions, we required a visual solution that differentiated many event types in a compact size independent of language to meet the needs of a global user community.

Early conceptual sketches focused on two different design approaches:

A. Icons close to the line. Event icons are placed along the price line in the chart, consistent with the then current approach. In addition, leaderlines and logic are used to lessen issues with occlusion and address data density. To differentiate between event categories and event types, glyphs would need to be designed to differentiate between categories and between types within a category.

B. Markers in time aligned horizontal panels. Simple event markers are located in parallel panels immediately below the price chart, and aligned to a common time axis. In this case, the icons are simply markers, as labels indicating the categories of events can be located within each horizontal panel.

C. Other approaches. Some other visualization approaches were considered that did not include a horizontal linear time axis. However, it became a requirement to use a familiar horizontal linear time axis as the professional users were extremely familiar with horizontal time oriented charts.

Approach B, horizontal time-aligned panels, was appealing, as it seemed highly scalable and fit with the objective of creating the timeseries roadmap.

Approach A, icons along the line, was also appealing in that users were already familiar with the technique

and it worked well when events were few in number (e.g. 5-25 events).

Horizontal time-aligned panels (approach B) are not new to this community. Analytic functions of continuous data based on mathematical algorithms used to detrend the data such as MACD, DMI and RSI are typically displayed in adjacent time-aligned panels immediately below the key price timeseries. What is different is the placement of discrete data, i.e. event data, in multiple adjacent aligned panels. This was not common in the financial timeseries analysis community. Yet, there is historic precedent for placing discrete events along a common time axis with discrete horizontal bands, as seen in various old and recent charts e.g. [5,6,7,8]:

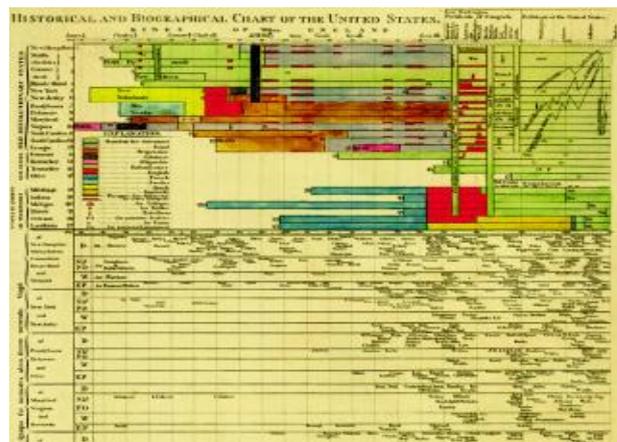


Fig. 7. Example showing bands of continuous data above (bars and lines) and discrete events below (text) aligned to a common horizontal time axis [5].

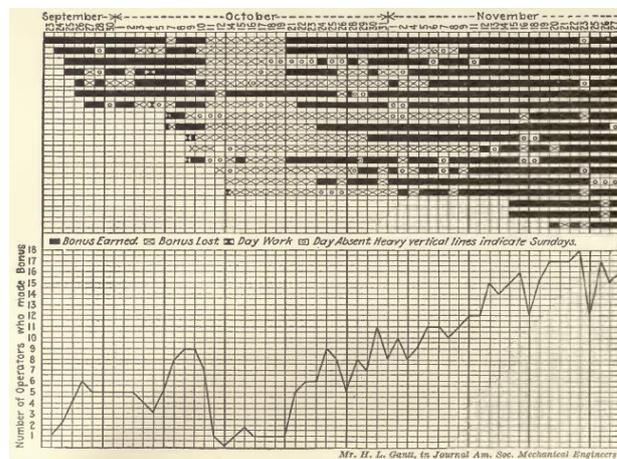


Fig 8. Example showing continuous data below and discrete categorical events above, all aligned to a common time axis shown at the top of the chart [6].

Rather than choose only one or the other, both ideas were blended together on the assumption that some users would prefer to have events located close to the price line when events were few (and as occurred in their existing interface); but there would inevitably be situations where the main chart would be overwhelmed with too many events interfering with the users' required need to visually parse the price line for patterns. The resulting design allowed users to move items between locations with either drag and drop or via a traditional dialog box.

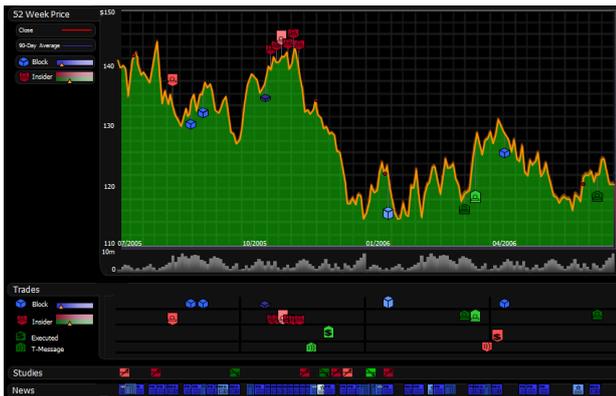


Fig. 9. Early design for the depiction of events along price line in chart and in separate panels below.

In order to maintain a consistent representation of events in both the chart and the aligned panels, a single design was required. As these representations need to be compact, a multivariate icon approach was used. The glyphs needed to:

- Differentiate between categories, such as news or stock trade.
- Differentiate within a category, such as a popular news story or an earnings release.
- Indicate potential price impact, for example, an insider selling stock in a company could be perceived as a negative influence on the stock price.

The solution was a series of pictographic glyphs (e.g. [9]) as opposed to abstract glyphs (e.g. [10,11,12,13,14]), each glyph indicating multiple data attributes:

- **Broad outline shape** indicates the category, e.g. square for news, chevron for trades.
- **Icon** inside the shape to indicate the specific type, e.g. big letter E for earnings release, a pictograph of a human figure for insider trades.
- **Color** of the broad shape indicates impact, e.g. red for a stock sell, green for a buy, blue for neutral.
- **Size** was considered, but not used.

4.2. Implementation and Examples

As design and implementation progressed, internal product reviews by different product areas revealed strong positive response to the design with many suggestions for additional event types, including events for commodities, messaging, automated patterns, rating agency announcements, etc. Internal design critiques provided incremental feedback to refine the design, e.g. glyph sizing based on data attributes was deferred as there were unresolved design issues and sizing could be added in later. Design critiques and reviews by subject matter experts highlighted issues with co-incident events, a need to contract multiple tracks together, leader lines, labels, magnification, workflow, event location, etc.

The first version was targeted for users interested in individual stocks. It was made available as one button on a toolbar of seven buttons immediately above the chart. Within twelve weeks there were thousands of users that had organically discovered the new functionality, were using it many times per day, and usage was increasing.

As intended, the visual alignment of events allowed for quick decoding of significant events in relationship to stock price movements. Some examples:

4.2.1. Single Price Move and Corresponding Event. The stock Motorola Mobility (MMI) had a huge leap in stock price on August 15, 2011. Visually scanning vertically at this date reveals many news stories (blue 1,2,5 icons, e.g. [A]), an acquisition (light blue A [L]), and many analyst reports (white diamonds and blue down arrow [H]). The catalyst for all the activity is likely the acquisition event. This can be confirmed via tooltips and drill-throughs, which show that Google announced their intent to acquire Motorola Mobility for \$10B.



Fig. 10. Events associated with Motorola Mobility in 2011. The sharp price increase in mid-August is heavily covered in news stories (blue icons above the line), the trigger event being the acquisition (vertically aligned A icon) of Motorola Mobility by Google.

4.2.2. Multiple Price Drops and Corresponding Events. Research In Motion (RIM/BB) stock price has dropped throughout 2011. In particular there are a number of sharp drops. Three sharp drops correspond to earnings events (green and red E icons, [E], directly along the price line), i.e. corporate filings of quarterly earnings results, which in these three instances, the market reacted negatively towards. There is one more sharp drop in price on April 28. Scanning vertically reveals a GU (guidance update [GU]), a number of analyst revisions (various H and S down arrows and white diamonds [H][S]) and a negative GDP report (red flag in lowest panel [G]). Figure 1 provides a large image. In this case, the trigger event was likely the guidance update (GU [GU]), which is a corporate announcement forecasting future earnings, again resulting in negative market response. This alignment can be confirmed with a mouse over interaction which overlays a visual leader line connecting the icon to the price line and additional details in a tooltip (fig. 11).



Fig. 11. Events associated with RIM in 2011. The stock has a number of sharp drops, each of which

now more than 200 unique event types available, showing that the original design goal of efficiency and extensibility has been achieved and showing that there is demand from various user communities for additional data events not originally considered.

The approach led to several benefits for end-users:

Information Inventory: The original concept provides “a place for everything and everything in its place.” Similar to a roadmap, it provides a repeatable, consistent framework for organizing information. Like a roadmap, it provides ways to understand different classes of local information through the use of data dense glyphs. And, like interactive maps, it provides means to add and remove different groups of data from the display in an expected manner.

Visual Correlation. By visually organizing data vertically along a common time axis, the user can visually infer relationships at points in time by scanning vertically. For example, for a given price anomaly, the viewer can scan vertically to see events occurring at the same time and infer a possible relationship, such as a trigger event. By organizing groups of events into panels, patterns within one type of event can be inferred by scanning horizontally. Horizontally aligned labels reduce the requirement to memorize event types and the need to rely on tooltips. The use of connotative coloring of icons aids visual screening for particular events (e.g. looking for red events associated with a price drop).

Data Density. The original need was to increase the data density. Previously, only a few dozen events could be clearly depicted along a timeline. The new technique provides compact icons that can be densely packed and is routinely used to show effectively anywhere from a few events to a few hundred events.

Workflow. Furthermore, the technique provides an effective workflow that provides a single consistent context (the visualization) and details as needed (through both tooltips and a click-through pop-up). While not a “focus + context” display (e.g. [17]), the approach enables the users to easily move back and forth from a common context to detail data. This enhances discoverability of functionality associated with any event type.

Adoption. The initial view of the new system did not require extra cognitive effort by the users. The icons utilized the same characters and colors for events that the user were familiar with. Also the events were located adjacent to the price line, with the addition of subtle leader lines to aid the viewer’s ability to discriminate the time of the event in relation to the price line. Through a revised dialog box for managing events, viewers were exposed to the new option of aligned panels. The new functionality was steadily adopted by users as they discovered the additional functionality in their current workflow and did not require any specialized training or documentation.

Finally, another important result is that this creation of a unique, efficient, easy-to-use visualization provides unique benefits only available at Bloomberg. This provides a competitive business advantage in the growth of new customers and retention of existing customers.

6. Future Work

The technique has scaled to 200 unique event types and there remain on-going challenges (e.g. designing unique, quickly understood icons) and new challenges (e.g. browsing, finding and managing event types of interest). Other applications areas for the visualization of event data are also under consideration.

The strategy and design objectives used here to achieve Bertin’s functions through organizing an inventory of information in a roadmap layout to support analytic processes is being used to solve other information challenges at Bloomberg.

References

- [1] M. Bloomberg. *Bloomberg by Bloomberg*. Wiley, New York. 2001.
- [2] M. Arrington. “Google Finance – Lots of Flash & Ajax,” *TechCrunch*, <http://techcrunch.com/2006/03/21/google-finance-lots-of-flash-ajax> 2006.
- [3] J. Larkin, H. Simon. “Why a Diagram is (Sometimes) Worth Ten Thousand Words.” In *Cognitive Science*, vol. 11, no. 1, pp. 65-99, Jan. 1987.
- [4] J. Bertin. *Semiologie Graphique*, Gauthier-Villars, Paris 1967.
- [5] A. Grafton, D. Rosenberg. *Cartographies of Time: A History of the Timeline*. Princeton Architectural Press. 2010.
- [6] W. C. Brinton. *Graphic Methods for Presenting Facts*. Engineering Magazine Company, New York, 1919.
- [7] S. B. Cousins and M. G. Kahn, “The visual display of temporal information,” *Artificial Intelligence in Medicine*, vol. 3, pp. 341-357, 1991.
- [8] C. Plaisant, et al., “LifeLines: Visualizing Personal Histories,” in ACM CHI’96 Conference, 1996, pp. 221-227.
- [9] S. Rose. *Social Stratification in the United States*. The New Press. 2007.
- [10] A. Barr. Superquadrics and angle-preserving transformations. *IEEE Computer Graphics and Applications* 1, 1 (Jan. 1981), 11–23.
- [11] H. Chernoff. The use of faces to represent points in k-dimensional space graphically. *J. American Statistical Association* 68, 342 (1973), pp. 361–368.
- [12] R. Rohrer, D. Ebert, J. Sibert. The shape of Shakespeare: visualizing text using implicit surfaces. *Proceedings Visualization 98* (Research Triangle Park, NC). Los Alamitos, CA: IEEE Computer Society Press, 1998; 121 – 129.
- [13] F. Post, T. Walsum, F. Post. Iconic techniques for feature visualization. *Proceedings Visualization 95* (Atlanta, GA). Los Alamitos, CA: IEEE Computer Society Press, 1995; 288 – 295.
- [14] M. Ward. A Taxonomy of glyph placement strategies for multidimensional data visualization. *Journal of Information Visualization* 1 (3-4). 2004. pp. 194-210.
- [15] R. Brath. “The Many Dimensions of Shape.” *IV09 Opening Keynote*, www.oculusinfo.com/expertise.html
- [16] R. Brath. “Multiple Shape Attributes in Information Visualization: Guidance from Prior Art and Experiments.” *Proc. Information Visualization 2010 (IV10)*, 14th Int’l Conf. pp. 433-438, 2010.
- [17] B. Spence and M. Apperley. “A Professional’s Interface Using the Bifocal Display.” In: *Proceedings of the 1981 Office Automation Conference 1981*. pp. 313-315, 1981.