

Information Visualization for Stock Market Ticks: Toward a New Trading Interface

by

Pasha Roberts

B.A., Economics, Russian Studies
College of William and Mary, 1984

Submitted to the MIT Sloan School of Management
in Partial Fulfillment of the Requirements for the Degree of

Master of Science

at the

Massachusetts Institute of Technology

February 2004

© 2003 Pasha Roberts. All rights reserved

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part.

Signature of Author

MIT Sloan School of Management
December 5, 2003

Certified by

Andrew W. Lo
Harris & Harris Group Professor
Thesis Supervisor

Accepted by

Margaret C. Andrews
Executive Director of the Sloan MBA Program

Information Visualization for Stock Market Ticks: Toward a New Trading Interface

by Pasha Roberts

Submitted to the MIT Sloan School of Management
on December 5, 2003 in Partial Fulfillment of the
Requirements for the Degree of

Master of Science in Management

Abstract

Ticks, the second-to-second trades and quotes of a market, might be considered the atoms of finance. They represent the basic, defining transactions that represent an asset in the market. Almost all financial concepts, such as returns or risk, are essentially abstractions from tick data. Like atoms, ticks form a truly massive dataset – millions per day, far too numerous to represent with traditional graphical methods.

Information-intensive disciplines such as medicine, bioinformatics, the earth sciences, and computational fluid dynamics have adopted modern visualization methods to manage and explore their respective mountains of data. Finance, however, has done little with the discipline of visualization to date.

This research presents the hypothesis that information visualization of tick data can improve human performance for intraday equity trading. To consider the hypothesis, the act of equity trading is broken into functional tasks, and the tasks mapped to information requirements. Using the TAQ historical dataset, the research evaluates new 2D and 3D information designs for tick data, and creates a visual language for equity trading.

The functional tasks and information designs are implemented in a visualization application, which provides an almost purely graphical trading interface to historical ticks. The user experience, trading performance, and analytical insight from this application are evaluated versus numeric methods. Based upon this experiment, the research concludes by exploring the usability, potential issues and future directions of trading and tick visualization in general.

Thesis Supervisor: Andrew W. Lo
Title: Harris & Harris Group Professor

Acknowledgements

Professor Andrew W. Lo initially focused my interest in financial visualization into ticks and the trading problem, with the challenge to create software that would enable “a reasonably intelligent 13-year old to be a successful trader.” His ongoing interest, insight and assistance as thesis advisor have been invaluable in the completion of this research.

I have greatly appreciated the opportunity to do much of my work at MIT’s Laboratory for Financial Engineering. Far more than a quiet place to work, the LFE was a broad and valuable resource of industry-oriented knowledge. Mike Epstein brought decades of experience straight from the great trading pits of the world, and kept me firmly in touch with the real, non-academic needs of traders. Svetlana Sussman was instrumental at holding the lab together; organizing dozens of seminars with exceptional speakers, and keeping everyone well supplied and fed.

Felice Frankel, research scientist at the MIT School of Science, was an exceptional source of energy and ideas. In particular I thank her for several clarifying and informative design concepts for the prototype visualization.

Mark Kritzman of Windham Capital Management Boston was kind enough to offer industry input and a sanity check for the thesis, plus some ideas about future validation.

On the home front, most of all, I thank my wife-to-be, Greta Kauffman, for her constant support and patience with my work on this project. Her interest and care have been a tremendous boost to me personally, and I sincerely look forward to “going long” on this relationship with her.

Table of Contents

1	Introduction.....	7
1.1	Statement of Problem and Methodology.....	8
1.2	Overview.....	8
2	Information Visualization - Motivation and Practice.....	9
2.1	Visualization for Presentation.....	10
2.2	Visualization for Analysis.....	11
2.3	Tactical Visualization.....	11
2.4	Potential Pitfalls of Visualization.....	12
2.4.1	Metacognitive Gap.....	12
2.4.2	Chartjunk and Other Design Faux Pas.....	13
2.4.3	Novelty Factor.....	13
2.4.4	Navigation Problems.....	13
2.4.5	Vigilance Fatigue.....	13
2.5	Visualization in Non-Financial Fields.....	14
3	Prior Financial Work.....	17
3.1	The Nature of Financial Data.....	17
3.2	Current Graphical Standards.....	18
3.3	Academic Work.....	19
3.4	Commercial Ventures.....	24
3.4.1	Generic Visualization Products.....	24
3.4.2	Financial Analysis Products.....	25
3.4.3	Professional Trading Products.....	28
3.4.4	Web-Based Consumer Products.....	31
3.4.5	Custom Development.....	32
4	A Functional View of the Trading Problem.....	36
4.1	Asset Search.....	36
4.2	Position/Timing Search.....	37
4.3	Execution Problem.....	38
4.4	Monitoring Problem.....	38
4.5	Trading Policy.....	39
5	A Trading Visualization System.....	40
5.1	TAQ Data.....	40
5.2	Mapping Functional Tasks to Views.....	41
5.3	Single-Asset Focus View.....	41
5.4	Multiple-Asset Context View.....	45
5.5	Trading Interface.....	46
5.6	Data Animation.....	51
5.7	Navigation Issues.....	51
5.8	Technical Implementation.....	52
6	Observations.....	53
6.1	Dominance of MSFT and INTC.....	53
6.2	Opening Volume.....	55

6.3	Propagation & Basket Selling.....	57
6.4	Price Change Behavior.....	59
6.5	Late Prints & Errors.....	62
7	A Day of Trading.....	63
7.1	Traders Log.....	63
7.2	Trading Results.....	72
8	Future Work.....	73
8.1	User Testing.....	73
8.2	Algorithmic Testing.....	73
8.3	Visual Behavioral Finance.....	73
8.4	Visualization Enhancements.....	74
8.4.1	Broader Context View.....	74
8.4.2	Integration of Quote Data.....	74
8.4.3	Market and Technical Indicator Integration.....	75
8.4.4	Sound.....	75
8.4.5	Arbitrage Formation.....	75
8.4.6	Graphical Order Placement.....	76
8.4.7	Intelligence Augmentation.....	77
8.5	Related Financial Applications.....	77
8.5.1	Trading Presentation.....	77
8.5.2	Compliance and Surveillance.....	77
8.5.3	Microstructure Research.....	78
8.5.4	Monte Carlo Analysis.....	78
9	Conclusions.....	79
10	References and Bibliography.....	80

List of Figures

Figure 1: Visualization Aids Ocean Current Mapping.....	14
Figure 2: MRI Studies Map Neural Pathways	15
Figure 3: Genomic and Proteomic Visualizations	15
Figure 4: Battlespace Visualization for Tactics and Strategy.....	16
Figure 5: Typical Price-And-Volume Chart Design.....	18
Figure 6: Financial Viewpoints Shows Three Intersecting Spreadsheets	20
Figure 7: A Self-Organizing Map Shows Distribution of Shapes.....	21
Figure 8: Columns and Worms Show Money Flow Over Time.....	22
Figure 9: StockVis Web-Enables Advanced Charting.....	23
Figure 10: Portfolio Impact Shows an Asset Landscape	25
Figure 11: A GSphere Shows Diversification.....	26
Figure 12: Market Topology Shows Complex Correlations	27
Figure 13: Heatmaps Show Price Changes.....	28
Figure 14: GL Genie Trades in 3D	29
Figure 15: vrTrader's Pioneering Display	30
Figure 16: Map of the Market Shows 500 Stocks at Once	31
Figure 17: The NYSE 3DTF Monitors Operations	33
Figure 18: JPMorgan Trading History Viewer	34
Figure 19: Unified Field Wins the Beauty Contest.....	35
Figure 20: Functional View of the Trading Problem	36
Figure 21: Data Mapping for Focus View	42
Figure 22: Color Map Setup Dialog	43
Figure 23: Close, Open, High, Low, Current Markers for Focus View	44
Figure 24: Merged Markers Are Still Visible	44
Figure 25: Early Multi-Asset Design.....	45
Figure 26: Data Mapping for Context View.....	46
Figure 27: Trade Entry Interface.....	47
Figure 28: Each Position Lot is Represented as a Cube	47
Figure 29: Aggregate Position is Represented as a Cone.....	48
Figure 30: Multiple Trade Lots – 3 Open, 1 Closed	49
Figure 31: Scaling into a Successful Trade	49
Figure 32: Context Screen Shows Multiple Positions.....	50
Figure 33: Animation Control Panel	51
Figure 34: MSFT and INTC Dominate Trade Volume	53
Figure 35: MSFT Consistently Shows More Volume on Single-Asset Charts	54
Figure 36: Extra Volume at Opening Bell	55
Figure 37: Context View of Open Shows High Volume	56
Figure 38: Searching for Price Propagation	57
Figure 39: JPM Gains \$0.20 in 30 Seconds	59
Figure 40: HON Loses \$0.09 in 45 Seconds.....	60
Figure 41: HON Gains \$0.09 Back in 15 Seconds.....	60
Figure 42: A Typical Late Print.....	62

1 Introduction

The earliest mathematicians thought about their problems entirely in geometric terms. Archimedes and the ancient Greeks created elegant and beautiful geometric proofs for many of today's advanced mathematical concepts. But without the symbolic notation of algebra, they were ultimately limited from the calculus and higher forms of abstraction.¹ With the change to symbolic notation, mathematical discovery exploded, but almost completely left its graphical roots. In the hundreds of years since the development of algebra and calculus, graphical reasoning has taken a far back seat to symbolic reasoning.

Over the last several decades the volume and complexity of information has avalanched, and computer-aided graphical capacity has advanced tremendously. As a result there has been a new interest in using computerized graphical techniques to present large, multidimensional results or to explore complex datasets. From this background was born the modern field of visualization.

Numerous scientific and quantitative fields have derived insight and value from the use of visualization, but finance has only scratched the surface of the discipline. This is a curious fact, since finance has an abundance of highly accessible, precisely measurable data, with immediate and direct rewards available to those who are able to derive meaning from these data quickly and accurately.

One of the ripest areas for financial visualization is the study of “ticks”, the second-by-second trades and quotes in a market. These transactions are essentially the atoms of finance, underlying all further abstractions of market behavior. This dataset is far too large in scope and scale to be shown using simple line-plot methods without aggregation.

¹ Maor, Eli. *E: The Story Of A Number*. Princeton: Princeton University Press, 1994. 44-46.

1.1 Statement of Problem and Methodology

This research presents the hypothesis that advanced visualization techniques can improve human performance for intraday equity trading. To consider the hypothesis, this research evaluates new information designs for equity ticks, and demonstrates a prototype application to embody these designs for the TAQ historical dataset. The resulting system provides an almost purely graphical trading interface, which is then evaluated in terms of user experience, trading performance, and analytical insight.

1.2 Overview

In the next chapter, information visualization will be described and motivated as a discipline, with examples from non-financial visualization applications. Prior academic and commercial works in financial visualization will be reviewed in Chapter 3. Chapter 4 will break the act of trading down into functional tasks, and map the tasks to existing information designs. The same tasks will then be fit into a more advanced visualization framework in Chapter 5, chronicling the development of a prototype application written by the author.

Chapter 6 will observe qualities of the underlying data using this new tool, and Chapter 7 will describe one day of trading using the system. Future work will be explored in Chapter 8, and finally in Chapter 9, conclusions will be drawn from the exercise.

2 Information Visualization - Motivation and Practice

Various definitions have emerged in the multi-disciplinary field of visualization. “Information Design” is the broadest, defined as “the art and science of preparing information so that it can be used by human beings with efficiency and effectiveness.”² This is broad enough to include anything from calligraphy to tables to advanced graphics. “Visualization” is defined more narrowly, as “the use of computer-supported, interactive, visual representations of data to amplify cognition.”³ Discussion of visualization typically involves formally mapping data elements to visual elements, such as mapping price to a Y-axis. Some definitions of visualization⁴ allow for non-visual forms of media, such as auditory, tactile, and other interfaces. Some practitioners have gone so far as to rename the practice “Multisensory Data Perceptualization” or “Interactive Perceptualization”.

As much of science deals with physical phenomena, “Scientific Visualization” typically refers to the visualization of physically based data. In this case, the three dimensions of the model’s physical space are almost always mapped directly to identical dimensions in a graphical representation. Examples abound, including cartography, weather information, MRI scans, molecular visualizations, and fluid dynamics plots. Cloud patterns or a patient’s brain already exist in space in three dimensions, which are scaled down to fit into a visualization’s view.

“Information Visualization” is the visualization of abstract data. With no natural mapping between data and graphical elements, the information designer is left with the considerable challenge to find a fit that can maximize understanding and usability. A common line plot of stock prices is an example of simple information visualization; time and price do not naturally exist in 3D space, but were mapped by the designer into X and Y dimensions, respectively.

For any information design, graphical or not, there are three possible goals: presentation, analysis or tactics. It is vital to be clear which is intended, and who the audience will be. The expository, even persuasive, goals of presentation often directly conflict with the open-ended,

² Horn, Robert E. “Information Design: Emergence of a New Profession”. *Information Design*. Ed. Robert Jacobson. Cambridge: MIT Press, 1999. 15.

³ Card, Stuart, et al., eds. *Readings in Information Visualization*. San Francisco: Morgan Kaufman, 1999. 6.

⁴ Foley, J., and B. Ribarsky. “Next-Generation Data Visualization Tools”. *Scientific Visualization: Advances and Challenges*. Ed. L. Rosenblum, et al. London: Academic Press, 1994.

exploratory goals of analysis, or the terse, time-critical demands of a tactical display. Each will be considered separately.

2.1 Visualization for Presentation

The goal of a presentation may be to educate, to entertain, to persuade, to sell, or for many other motivations. Even absent of an explicit agenda, a presentation still conveys meaning and conclusions – even seemingly raw data such as a phone book or train schedule.

The classic works in presentation graphics are by Jacques Bertin⁵ and Edward R. Tufte⁶. The theories and designs of these two men have crystallized and formalized the discipline, and have become to some degree synonymous with the practice of good presentation design.

In many cases, a simple textual table can be the best information design for a series of numbers to support a point.⁷ Graphical designs are exceedingly useful, however, to speed and enhance understanding of relationships, causation, scale, and quantity. Good images can be, as Tufte puts it, “quantitatively eloquent”.⁸

This eloquence is partially possible because the image acts as a repository for the data under consideration, relieving our “working memory”⁹ from having to remember an entire multi-dimensional series. This visual crutch lets the viewer remember and compare more data, more quickly than from rows and columns of numbers.

Certain graphical features are seen with “pre-attentive processing”¹⁰ so that they are understandable at a glance. These pre-attentive features are understood much faster than others because they are literally processed prior to conscious attention. As a result, well-designed pre-attentive graphical elements pop out in one’s view of a visualization.

Another source of this eloquence is the tremendous pattern-finding capacity of the human visual system. A point of meaning can be clearly supported by a graphic representation of underlying data, lined up to display the intended pattern. A well-crafted image makes the mental

⁵ Bertin, Jacques. *Semiology of Graphics: Diagrams, Networks, Maps*. Madison, Wis.: University of Wisconsin Press, 1983.

⁶ Tufte, Edward R.. *The Visual Display of Quantitative Information*. Cheshire, Conn.: Graphics Press, 1983.

Tufte, Edward R.. *Envisioning Information*. Cheshire, Conn.: Graphics Press, 1990.

Tufte, Edward R.. *Visual Explanations*. Cheshire, Conn.: Graphics Press, 1997.

⁷ Tufte, *Envisioning Information*, 104.

⁸ Tufte, *Visual Explanations*, 13.

⁹ Ware, Colin. *Information Visualization: Perception for Design*. San Diego: Academic Press, 2000. 373.

¹⁰ *Ibid.*, 163.

model of the presenter explicit, and adds credibility to the hypothesis. The viewer can make sense of a model, visually test alternative hypotheses, and validate that the dataset is reasonably free from errors.

2.2 Visualization for Analysis

All of the presentation factors above are operational for an analytical, exploratory visualization. The difference is in the goal – rather than presenting a known model or conclusion, the visualization is intended to help the user find one. More formally, the goal of such a system is to harness the human visual system to find and act upon patterns in the dataset, by allowing flexible manipulation and display of the data. John W. Tukey was the first to formalize this practice, in his *Exploratory Data Analysis*,¹¹ emphasizing statistical insight and analysis using graphics.

Deriving a mental model from raw data and acting upon it is typically understood as a cyclical process with feedback at many stages.¹² Ultimately, this type of process requires an information visualization system to allow the user to easily search and categorize the raw data, create multiple models and test hypotheses iteratively. The system may also enable the user to directly act upon the data, for example by changing simulation parameters.

2.3 Tactical Visualization

Tactics might be considered to be “seat of the pants” analysis. In situations such as military engagement or stock trading, there is no time to think, and the data are usually fuzzier than in a detached analysis. In these situations, the threat of physical or financial peril requires the user to act immediately – without leisurely research. Because information must flow to the user in real time, it is more urgent, more error prone, and technically more difficult to deliver.

Information designs for these environments must be simple enough to avoid misunderstanding during the fog of war, but powerful enough to provide insight. There is simply not time for many of the affordances available to classical analytic visualizations. A milder example of this concept applies to instrument displays for cars or aircraft – while a detailed

¹¹ Tukey, John W.. *Exploratory Data Analysis*. Reading, Mass.: Addison Welsley, 1977.

¹² Card, *Readings in Information Visualization*, 10-12.

analysis of a car's fuel utilization may be interesting, a driver typically has less than a second to check the fuel level and get the eyes back on the road.

There is little treatment of this category in the literature, outside of military work or industrial design. The object of this study, intraday equity trading, is clearly in this camp.

2.4 Potential Pitfalls of Visualization

2.4.1 Metacognitive Gap

It is a running joke in the cognitive sciences to compare the growth in computer power, with transistor density doubling every 24 months, to the growth in human computing power, with brain volume doubling every 3×10^7 months.¹³ Their point is that human intuitions about how people see, hear, and remember are embarrassingly inflated. We are only physically capable of seeing about five objects at a time, and can only observe one change at a time, but our intuition tells us that we can see everything at once.¹⁴ This lack of awareness is formally called the “metacognitive gap”, manifesting in numerous ways. One dramatic example is called “change blindness” where viewers are able to not “see” directly visible objects¹⁵— the classic example is of drivers colliding with plainly visible pedestrians.¹⁶

For the information designer, the net of this visual cognition research is twofold. First, it is vital to realize the limitations of human perception, avoiding designs that ask the impossible of users. Though this sounds obvious, it is in fact difficult to avoid with a large, multidimensional dataset. The second lesson is to create visualizations that support the ability of human cognition to use the scene as an external memory bank and processor. A coherent visualization with a predictable interface can extend working memory and amplify the cognitive ability of the user. A badly designed visualization can easily mislead or confuse, simply because human cognition isn't as powerful as most people believe.

¹³ Course Notes from “Design of Interactive Multimodal Media Systems”. Proceedings of *ACM Siggraph 2003*.

¹⁴ Rensink, Ronald A.. “Internal vs. External Information in Visual Perception”. Proceedings of *ACM Symposium on Smart Graphics*, June 11-13, 2002.

¹⁵ Simons, Daniel J, and Christopher J Chabris. “Gorillas in Our Midst: Sustained Inattentive Blindness for Dynamic Events”. *Perception*, 28 (1999), 1059-1074.

For examples, see http://viscog.beckman.uiuc.edu/djs_lab/demos.html

¹⁶ For examples, see <http://www.cs.ubc.ca/~rensink/flicker/index.html>

2.4.2 Chartjunk and Other Design Faux Pas

Edward Tufte coined the term “chartjunk”¹⁷ to describe redundant, decorative embellishments to information designs. Much has been said about this elsewhere, but it is important to underscore the degree to which good design is vital to the success of a visualization. Design problems can render a well-intentioned view into an ineffective or even misleading tool.

2.4.3 Novelty Factor

One frequently noted criticism of visualization, particularly of 3D formats, is the “novelty factor”. Interactive 3D views are still relatively rare outside of video games and movies, and there is legitimate concern that users will become fascinated with the technology rather than the content. This can in fact be a positive force, leading the user to spend more time with and offer more patience to a new system than they would normally.

The novelty factor is, however, a temporary fascination, after which the visualization must stand or fall based on its inherent content and usability. If it fails to deliver value or is frustrating to use, then users will go back to previous tools, good or bad.

2.4.4 Navigation Problems

One area of particular risk is the issue of 3D navigation. With the exception of heavy video gamers or 3D professionals, most computer users are not comfortable with the mechanics of moving around in a 3D world on a 2D screen. Even with 3D input devices, many find it easy to get “lost in space” and become quickly frustrated. It is unfair to blame the user for a lack of discipline in memorizing, say, that shift-left-click will pan an object, or even for not knowing what panning is. It is the information designer’s job to make navigation feel “natural”. Since flying can hardly be considered natural for humans, effective 3D interfaces often severely constrain the user’s mobility, for example by adding a false sense of gravity, forcing the user to walk across a landscape.

2.4.5 Vigilance Fatigue

A phenomenon that can be easily confused with the novelty factor is “vigilance fatigue”. There is a known tendency, first observed in aircraft traffic controllers, and more recently in baggage screeners, for the performance of high-attention visual tasks to rapidly degrade after one

¹⁷ Tufte, *The Visual Display of Quantitative Information*, Chapter 5.

hour.¹⁸ Thus an excessively demanding new visualization may seem at first to be revolutionary, but will seem to be next to useless at the end of a full day. Vigilance fatigue is particularly strong in situations that ask the user to detect infrequent patterns – everything begins to look the same. This form of fatigue can be addressed directly, rather than writing off the entire visualization as a novelty.

2.5 Visualization in Non-Financial Fields

Visualization has been a significant tool in numerous other fields of scientific endeavor. In order to evaluate the applicability of the practice for finance, it will be useful to briefly overview the uses of visualization in these other fields. It will be seen that most of these early successes involve scientific visualization, which models physical, non-abstract data.

In physics, visualization is commonly used in the analysis of computational fluid dynamics, astronomy, and nuclear physics. Earth science, intrinsically related to 3D space, makes extensive use of the tool – especially in meteorology, ocean current analysis, and geology. The business of oil and gas exploration has derived great financial benefit from visualization, where virtual exploration is far less expensive than actual drilling.

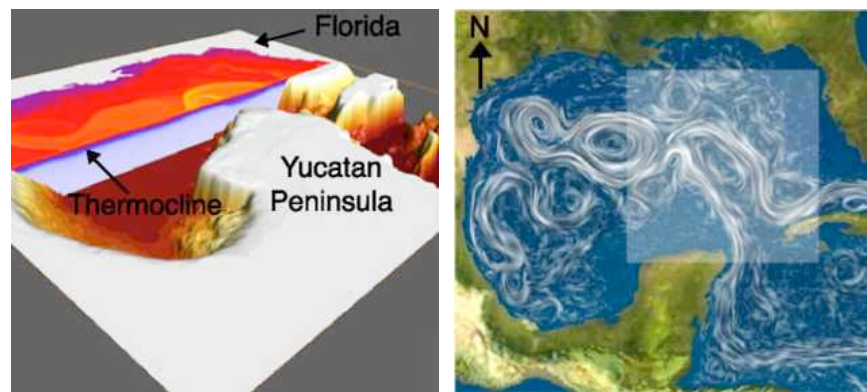


Figure 1: Visualization Aids Ocean Current Mapping¹⁹

¹⁸ Ware, *Information Visualization: Perception for Design*, 341-342.

¹⁹ Grant, Josh, et al. "Case Study: Visualizing Ocean Flow Vertical Motions using Lagrangian-Eulerian Time Surfaces". Proceedings of *IEEE Visualization 2002*.

Similarly, the practice of medicine has found many areas where visualization is simpler and more practical than surgery. Visualization is synonymous with the related practices of radiology, computed tomography (CT) and magnetic resonance imaging (MRI). In addition to life-saving diagnostic procedures, CT and MRI have been an invaluable aid to the study of brain structure and function.

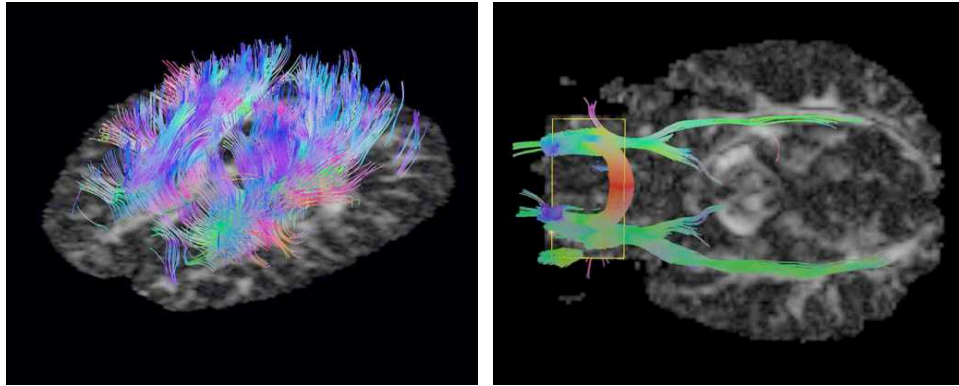


Figure 2: MRI Studies Map Neural Pathways²⁰

Smaller, but equally physical is the study of bioinformatics, especially genomics and proteomics. Visualization has enabled scientists to map and begin to understand the microscopic and complicated relationships of genes, proteins and life.

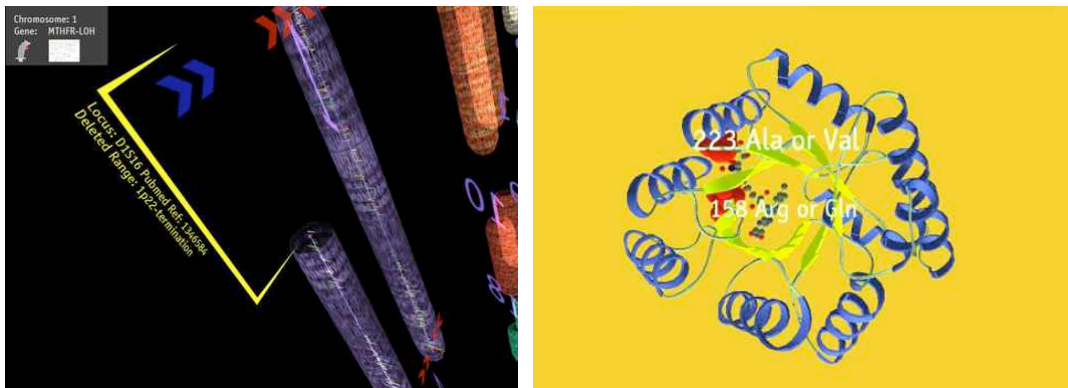


Figure 3: Genomic and Proteomic Visualizations²¹

²⁰ Zhukov, Leonid, and Alan Barr. "Oriented Tensor Reconstruction: Tracing Neural Pathways from Diffusion Tensor MRI". Proceedings of *IEEE Visualization 2002*.

²¹ Adams, Mark, et al. "Case Study: A Virtual Environment for Genomic Data Visualization". Proceedings of *IEEE Visualization 2002*.

The US military has been a leader in adopting this discipline for simulation, training, and battlespace visualization. This work is of particular interest because it often takes place in a tactical context, with an appreciation for real-time, incomplete data.



Figure 4: Battlespace Visualization for Tactics and Strategy²²

Lastly, visualization is starting to see regular use in the more abstract domain of data communication. Information designs and products are now available for visual network monitoring, debugging of cluster hardware, and particularly for web traffic and ecommerce interactions.

²² US Office of Naval Research, *The Dragon Battlefield Visualization System*.
<http://www.ait.nrl.navy.mil/vrlab/projects/Dragon/Dragon.html>

3 Prior Financial Work

3.1 The Nature of Financial Data

At its simplest form, financial data is a time series of price and volume values for a wide array of assets. From these prices are derived returns, and most other financial information. In mathematical terms, the time series of prices across many assets can be seen as a vector field, with an amplitude (volume) and velocity (return) at each time and price point. The realized set of actual price values is one integral path through the vector field, essentially a streamline. A Monte Carlo analysis will generate a wider range of hypothetical price paths through the same field. Another form of vector field comes from the analysis of payoffs from option or trading strategies, examining sensitivities to asset price, interest rate, time, and so on.

Most other concepts of financial data also revolve around the central concepts of time and price. Some may describe the state of a market or portfolio, showing inventory levels of assets, risk levels or the current market depth of bids and offers.

In this light, financial information is hardly unique – there are many information designs that are appropriate to time series, vector fields, and stocks of inventory.

3.2 Current Graphical Standards

The common two-dimensional line plot is clearly the workhorse of today's financial graphics world. Most financial graphs are a variation of Figure 5, showing time on the X-axis, price (or return) on the Y-axis, and volume on the Y-axis of a separate bar chart. It is an effective way to show three dimensions for a single asset, in one compact view. Minor modifications can add dimensions for high/low prices, moving averages, and similar metrics on the same timeline.

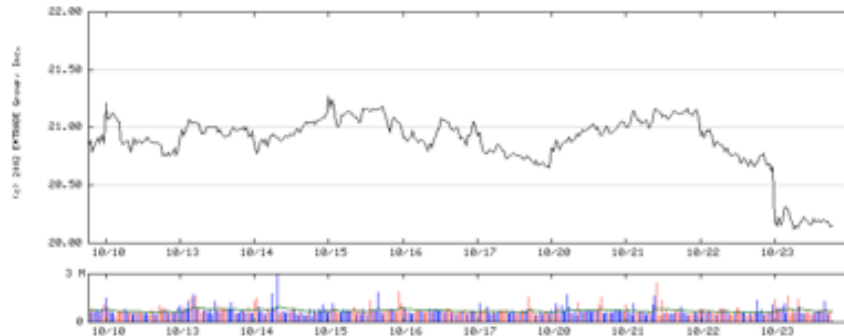


Figure 5: Typical Price-And-Volume Chart Design²³

This design is pervasive, and is found in everything from publications such as the *Wall Street Journal* to trading software, such as Bloomberg® or Tradestation®. It is notable that the discipline of Technical Analysis, a major and popular body of market action studies, is entirely based on the use of charts such as the above.²⁴ For low-dimensional time series information, this design is very good, and there is no reason to fix something that is not broken.

However, many financial problems produce high-dimensional information that cannot fit into this simple framework. The following sections will explore efforts that have attempted to solve this design problem.

²³ Courtesy of E*Trade, <http://clearstation.etrade.com>

²⁴ Edwards, Robert D., et al.. *Technical Analysis of Stock Trends*. Boca Raton, Fla.: St Lucie Press, 8th Edition, 2000. 9.

3.3 Academic Work

Visualization currently does not garner a great deal of attention in the financial literature, and likewise, finance is only a minor category in the visualization literature. Prior academic work in financial visualization can be divided into two categories. First, there is visualization-oriented research that incidentally uses financial data, exploring techniques or designs without serious regard for the underlying financial data. These are not uncommon, since financial data is numeric, easy to access, and has high intrinsic credibility. The designs from these are sometimes interesting, but the financial analysis is often flawed. Second, there is a limited amount of research that attempts to solve financial problems with visualization. Five of the latter efforts will be considered here.

Ben Schneiderman and Martin Wattenberg pioneered the application of treemaps to financial data, carefully dealing with layout algorithms that remain consistent across multiple iterations.²⁵ Their work has been commercialized as Smart Money's Map of the Market, and will be discussed in the commercial section below.

²⁵ Wattenberg, M. "Visualizing the Stock Market". Proceedings of *ACM CHI 99*, Extended Abstracts. 188-189.

In 1995 Lisa Strausfeld of the MIT Media Lab created an information environment named “Financial Viewpoints”. Though her effort is focused on creating an “embodied virtual space”, the result is an interesting financial analysis tool. To consider 7 mutual funds, with 50 data elements each, the visualization employs three movable planar spreadsheets, one on each axis.²⁶ To explore relationships in the data, the user moves the planes against each other, like a slide rule in 3D. The system also integrates simple bar charting and video feeds to support the data. Though useful and attractive, the structure of this view limits it from expressing relationships beyond three linear dimensions.²⁷

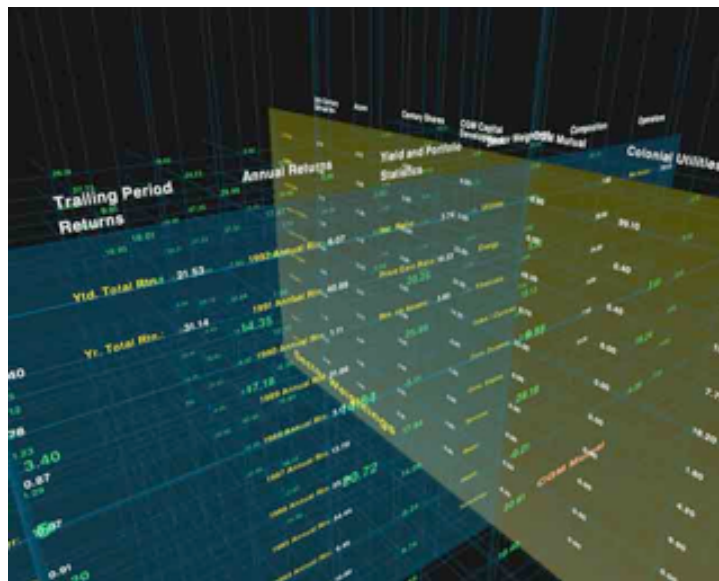


Figure 6: Financial Viewpoints Shows Three Intersecting Spreadsheets²⁸

²⁶ Strausfeld, Lisa. *Embodying Virtual Space to Enhance the Understanding of Information*. Masters Thesis, Massachusetts Institute of Technology, Program in Media Arts & Sciences, 1995. 17.

²⁷ Fry, Benjamin Jotham. *Organic Information Design*, Masters Thesis, Massachusetts Institute of Technology, Program in Media Arts & Sciences, 2000. 32.

²⁸ Strausfeld, *Embodying Virtual Space to Enhance the Understanding of Information*.

Kresimir Simunic from the Institute for Media Communication, Fraunhofer Gesellschaft in Germany, proposed a method of assisting Technical Analysis.²⁹ An index of assets will present a single price path, or shape, for a given time period – but within that aggregate shape, the component assets may display a wide variety of patterns. This work clusters similar chart shapes of component assets and maps them onto a self-organizing Kohonen feature map, an artificial neural network algorithm. The result shows the distribution of patterns within the aggregate, arranged topologically in a compact 2D view. The color brightness of each sub-chart represents the distance of its shape from the neighboring shape. This prototype also enabled zooming, so that an analyst could move in from a large index to sectors, or from a general shape to sub-patterns.

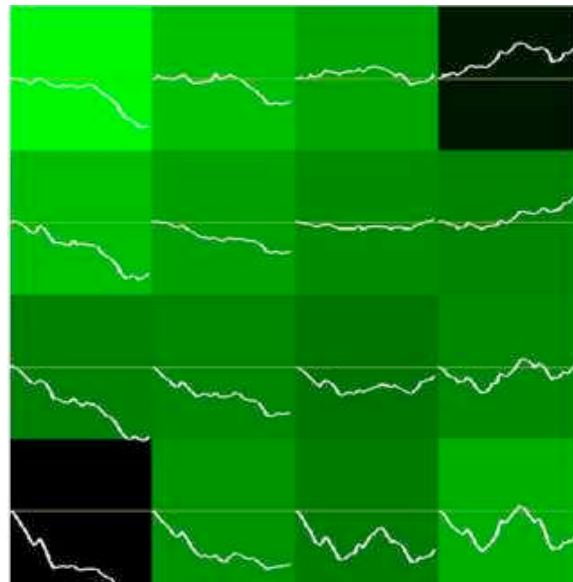


Figure 7: A Self-Organizing Map Shows Distribution of Shapes³⁰

²⁹ Simunic, Kresimir, “Visualization of Stock Market Charts”. Proceedings from *The 11th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision 2003*.

³⁰ Ibid.

Tim Dwyer and Peter Eades from the University of Sydney visualized share ownership between selected managers on the UK Stock market, using a proprietary dataset.³¹ The “columns and worms” metaphor maps individual assets to columns, varying in width with share price, and worm-like graph connectors to indicate fund transfers at each time period. This stratified graph maps time across 5 vertical levels over the 2-week period, allowing the user to move the blue plane through the time axis to clarify the view.

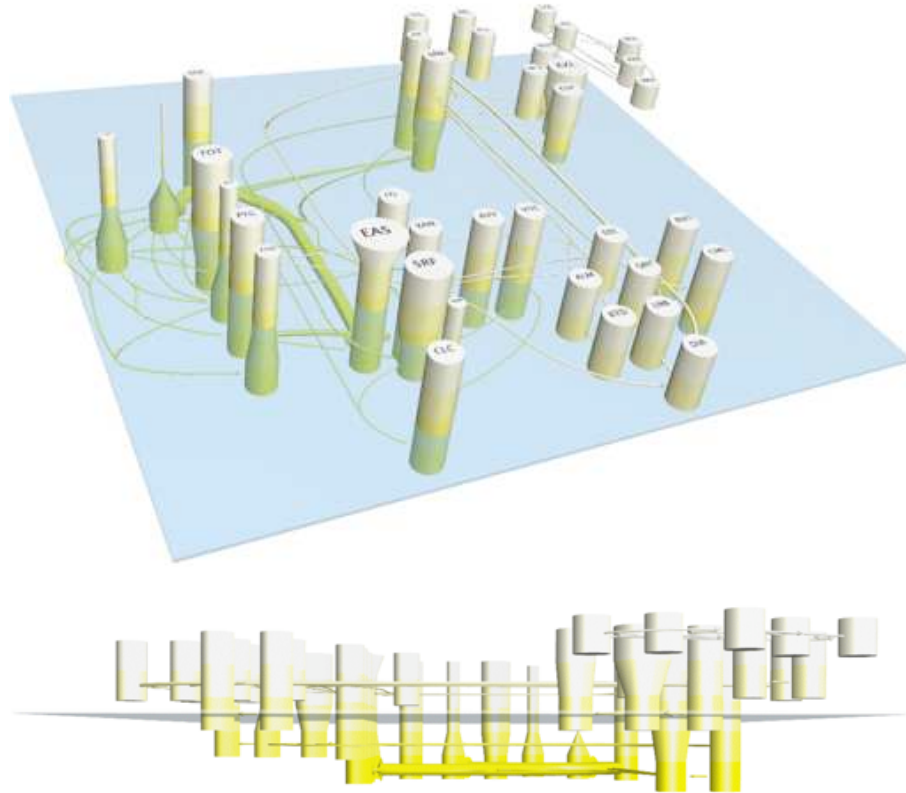


Figure 8: Columns and Worms Show Money Flow Over Time³²

³¹ Dwyer, T and P. Eades. “Visualising a Fund Manager Flow Graph with Columns and Worms”. Proceedings of *IEEE IV 2002*. 147-158.

³² Courtesy of T. Dwyer, University of Sydney.

Edward L. Parrish from UC Santa Cruz explored stock chart improvement within the constraints of web-enabled technology.³³ His StockVis system provides dynamic charts with the option to present multiple time periods and 3D charting. As a computer science thesis, the primary contribution of this work is to web-enable advanced charting, a difficult problem in itself.

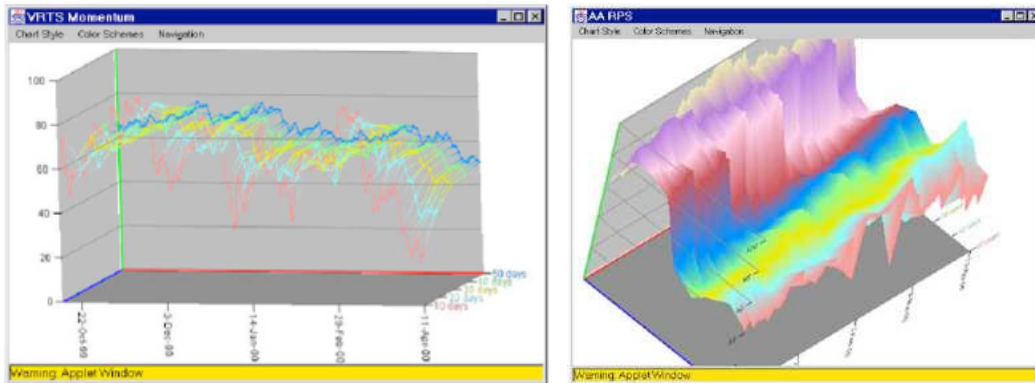


Figure 9: StockVis Web-Enables Advanced Charting³⁴

All of these views are designed to assist the process of financial analysis and can provide useful insight to traders, investors or academics. Most are retrospective, not real-time, and none yet address tactical interfaces for traders.

³³ Parrish, Edward. "StockVis: An Internet-Based System for Visualizing Stock Market Data". Master's Thesis, UC Santa Cruz, Department of Computer Science, 2000.

³⁴ Ibid

3.4 Commercial Ventures

Advanced visualization software is still very much a niche market, characterized by small companies and a high level of per-client customized work. There is no “killer visualization application” in any industry, despite a fair bit of hype in the early 1990’s.³⁵ Most commercial visualization applications focus on the technology’s early adopter industries: medicine, biotechnology and oil exploration.

3.4.1 Generic Visualization Products

A small number of vendors produce generalized information visualization tools, which are applicable to the financial world. Many of these tools are part of calculation software packages, such as Matlab, Mathematica, SAS, and the like. These are capable of producing very acceptable 2D and 3D presentation graphics, but are not particularly flexible or interactive for exploratory analysis or complete, stand-alone systems.

Spotfire, Inc.,³⁶ a local start-up, has created a leading analytical/visualization workbench with higher-end visualization capability, but has focused entirely on non-financial markets as a matter of priority and market knowledge.³⁷

There also exist a number of visualization programming libraries, such as VTK by Kitware,³⁸ or PV-WAVE by Visual Numerics.³⁹ These libraries are not stand-alone applications, as is Spotfire. They are building blocks for a programmer building a custom application, and will be discussed later, as the research prototype uses one of these toolkits.

³⁵ “Grokking the Infoviz”. *The Economist* 367:8329 (6/21/2003).

³⁶ Spotfire, Inc. <http://www.spotfire.com>

³⁷ Craig Hattabaugh, Spotfire Director of Sales, Interview by author, 10/30/03.

³⁸ Kitware, Inc. <http://www.kitware.com>

³⁹ Visual Numerics, Inc. <http://www.vni.com>

3.4.2 Financial Analysis Products

A small number of financial analysis products do integrate advanced visualization into their offering, primarily for portfolio or risk management.

High Tower Software's Portfolio Impact⁴⁰ platform includes a navigable 3D asset landscape; intended for brokers, fund managers, portfolio analysts, and exchanges. This "Cybergrid" operates in real time and can be layered hierarchically.



Figure 10: Portfolio Impact Shows an Asset Landscape⁴¹

⁴⁰ High Tower Software, Inc. <http://www.portfolioimpact.com/>

⁴¹ Courtesy of High Tower Software, Inc.

Gravity Investments has created an interesting visualization for portfolio diversification. Their GSphere⁴² represents an ideally diversified portfolio as a wireframe sphere, into which is placed a polygon model of an actual portfolio. The customizable axes are set to represent dimensions for diversification, and assets are plotted into that space.

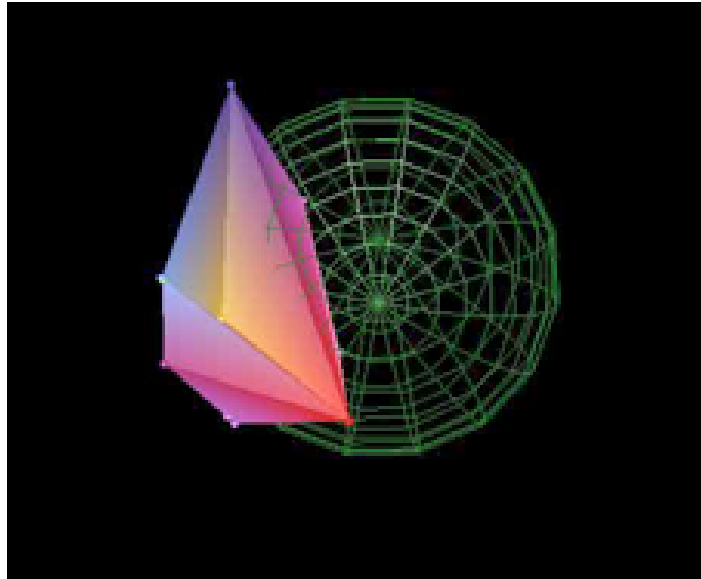


Figure 11: A GSphere Shows Diversification⁴³

An analyst is directed to (1) make the polygon shape as big as possible, and (2) as symmetrical, or sphere-like as possible.⁴⁴ The system interfaces with their optimization software, so the analyst is not left out on a limb for how to accomplish the above. One clear issue with this software is that since diversification is used as a proxy for risk, the analyst with a near-sphere may be left with the false impression of safety. Emotionally, the graphics may convey an untrue sense of concreteness and security to the portfolio's actual future volatility.

Also, a simple 2D bar chart could show this diversification for more than three factors, or a single sum-of-squares metric could measure the whole problem – but the view would be far less interesting.

⁴² Gravity Investments, Inc. <http://www.gravityinvestments.com/portfoliopics.htm>

⁴³ Courtesy of Gravity Investments, Inc.

⁴⁴ Gravity Investments Inc. <http://www.gravityinvestments.com/faq.htm>

Market Topology,⁴⁵ also known as Impactopia, arranges assets into tree graphs, based on correlation and performance. The tree naturally grows into geographical and sector branches. These fractal-like 2D and 3D patterns can be navigated and animated by the analyst seeking to understand complex relationships between assets.⁴⁶

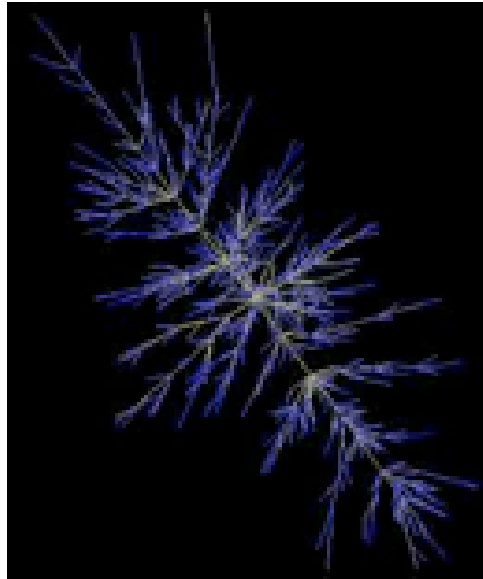


Figure 12: Market Topology Shows Complex Correlations⁴⁷

⁴⁵ Market Topology, Inc. <http://www.market-topology.com/>

⁴⁶ Inf@Vis! “Impactopia”. http://www.infovis.net/E-zine/2002/num_99.htm

⁴⁷ Courtesy of Market Topology, Inc.

3.4.3 Professional Trading Products

All popular trading platforms include the essential line plot representations of asset prices and volumes. There are some add-on packages to these, such as RINA Systems “3D SmartView” enabling traders to graphically evaluate the stability of various trading strategies.⁴⁸

NeoVision Hypersystems, Inc. claims to be the market leader in “tactical visualization”⁴⁹ with their HeatMap product line. These are simple but effective 2D charts, using color to convey price or risk movement for many simultaneous assets. They cover many asset types, integrate with many financial data feeds and are well positioned to sell to Wall Street. The displays are not particularly information-rich, nor attractive, but they are an important step forward from the numeric ticker.

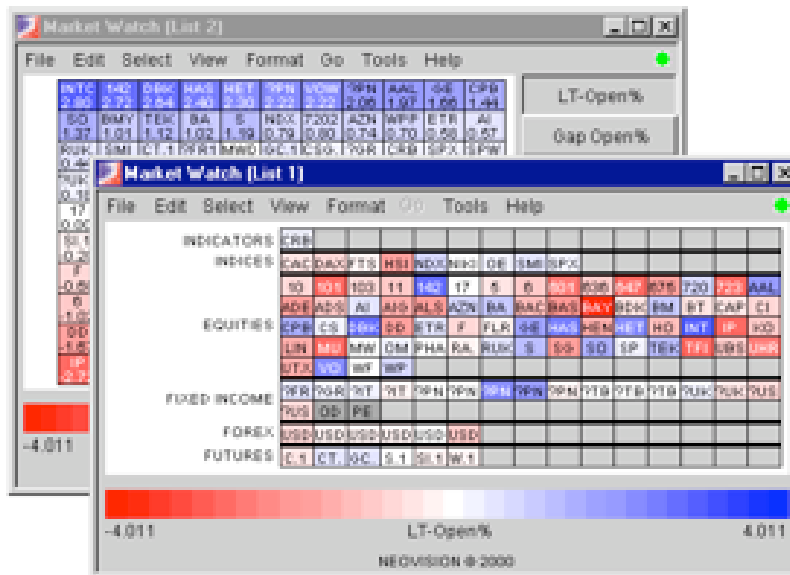


Figure 13: Heatmaps Show Price Changes⁵⁰

⁴⁸ RINA Systems, Inc. <http://www.rinafinancial.com/>

⁴⁹ NeoVision HyperSystems, Inc. <http://www.neovision.com/products/strategy/>

⁵⁰ Courtesy NeoVision HyperSystems, Inc.

GL Trade,⁵¹ a large French firm, markets a range of innovative trading tools. One of their products is a truly 3D trading platform, named GL Genie. This system has two main views, one of a portfolio, and another for a single asset. In the trading view, a line plot of the asset price recedes into the distance alongside a histogram of volume or quotes at each price.



Figure 14: GL Genie Trades in 3D⁵²

Interestingly, it is possible to trade directly from this interface. The trader can see his bid or ask represented directly in the quote histogram, and can drag the quote from one price to another if so desired. Limits to the trade can also be directly manipulated. It is an interesting paradigm from a large vendor, gained from the 2001 acquisition of a small London-based company, 4D Trading.

Hotspot FX⁵³ presents a web-based view of market quotes for specific foreign currencies. Its view is essentially a histogram of bids and offers, showing market depth far more clearly than a textual screen.

⁵¹ GL Trade SA. <http://www.gltrade.com>

⁵² Courtesy of GL Trade SA

⁵³ Hotspot FX, Inc. <http://www.hotspotfx.com>

In 1994, DIVE Laboratories created an early flash in this market with vrTrader, a trading application that was an offshoot of DIVE's military VR technology.⁵⁴ The product has been discontinued after a series of acquisitions – the current owner, Viewpoint Corporation⁵⁵ focuses entirely on 3D product marketing tools for the web.

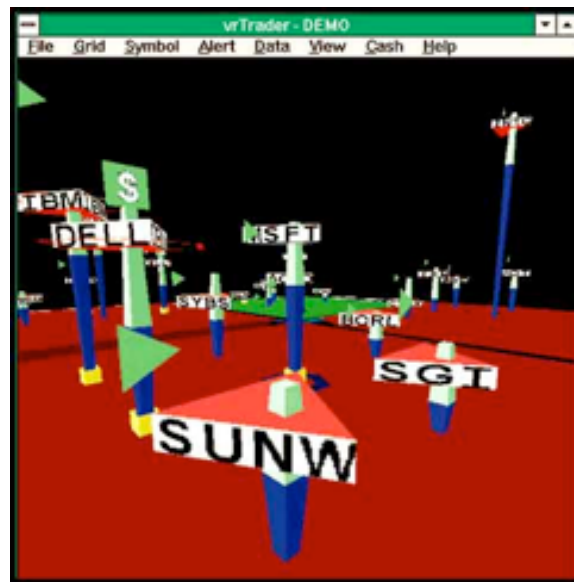


Figure 15: vrTrader's Pioneering Display⁵⁶

There are a number of marginal, smaller entries into the trading visualization market. Vibrant 3D sells "Market Helmsman", a tool to create interactive 3D scatter charts with customizable axes.⁵⁷ 3DStockCharts⁵⁸ sold monthly subscriptions to a small java-based application to visualize quotes and price action. However, its parent company is now in bankruptcy and this service may or may not be continued by another party.

In the UK, Clem Chambers made a splash by transitioning from computer games to finance, prototyping game-like 3D market representations, including whimsical views such as a fish tank.⁵⁹ However, his company⁶⁰ now only offers traditional quotes and stock picks.

⁵⁴ Baum, David, "Assets in Wonderland". *Byte Magazine*, July 1995. <http://www.byte.com/art/9507/sec8/art3.htm>.

⁵⁵ Viewpoint, Corporation. <http://www.viewpoint.com/>

⁵⁶ Courtesy of DIVE Laboratories

⁵⁷ Vibrant 3D, Inc. <http://www.vibrant3d.com/helmsman.htm>

⁵⁸ 3DStockCharts, Inc. <http://www.3dstockcharts.com>

⁵⁹ Fisher, Adam. "Playing the Market in 3-D". *Wired Magazine*, 8.06 (June 2000).

⁶⁰ ADVFN, Inc. <http://www.advfn.com>

From beginnings such as these, other companies may yet emerge in this space to sell to the trading community. There are now early glimmers of progress from firms such as Fractal Edge⁶¹ or ImagineXD⁶² – but for the moment, they remain small and under funded.

Sadly, many of these firms still use flashy cyber-lingo to name their companies and products. This kind of naming is problematic in a more mature market, and only evokes the failed hype surrounding VR in the early 1990's.⁶³

3.4.4 Web-Based Consumer Products

The most well-known and visible financial visualization is the “Map of the Market” by Smart Money. Utilizing the “Tree Map” view developed by Ben Shneiderman of UMD,⁶⁴ and Dow Jones, the Map of the Market is able to show the relative returns and market capitalization of over 500 assets in one compact 2D view.⁶⁵ Smart Money has enjoyed success in licensing this technology to online publications, such as the Wall Street Journal.

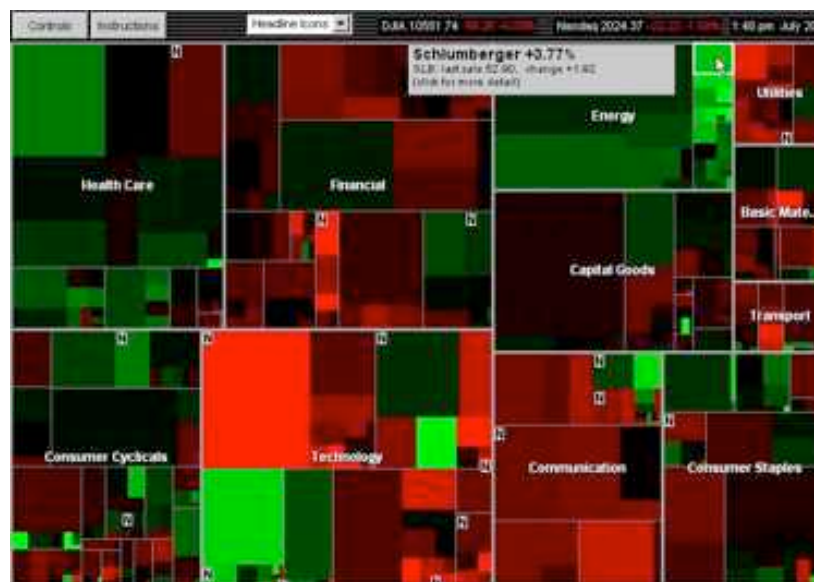


Figure 16: Map of the Market Shows 500 Stocks at Once⁶⁶

⁶¹ Fractal Edge Ltd. <http://www.fractaledge.com>

⁶² ImagineXD, Inc. <http://www.imaginexd.com/>

⁶³ The early 1990's VR hype was perhaps epitomized by *Lawnmower Man*, a 1992 movie directed by Brett Leonard, written by Stephen King.

⁶⁴ Ben Schneiderman <http://www.cs.umd.edu/~ben/>

⁶⁵ Schneiderman, Ben, and Martin Wattenberg. “Ordered Treemap Layouts”. Proceedings from *IEEE Visualization 2001*.

⁶⁶ Courtesy of Smart Money, Inc.

Smart Money has adapted the map in many ways to show mutual funds, risk and other concepts. It is a strong design, giving the user an immediate gist of the market, especially compared to the experience of viewing a multi-columned spreadsheet with these 500 assets.⁶⁷

While the view provides a good overview of the market, it is not a tactical tool for traders. As compelling as it is, the map is primarily intended to draw users and ultimately subscribers to the web site, rather than to be a heavily used tool.⁶⁸

Two important limitations apply to web-based consumer applications. First, the level of financial detail that is required or comprehensible by the average consumer is relatively simple. It is not useful to visualize option sensitivities to the average investor. Second, the web is a very limited delivery technology – for the immediate future, bandwidth and browser limitations will be a significant challenge for this market.

3.4.5 Custom Development

The most innovative financial visualizations remain in the province of custom development for financial firms. The largest market visualization in the world is certainly the Nasdaq tower in Times Square, New York City. This tower and the television studio's video wall underneath were a custom solution by Visible Insights/AdvizorSolutions.⁶⁹ Clearly they were designed for presentation and mindshare, rather than analysis – but they are an impressive and constant supply of real-time market information.

Visible Decisions, founded by William Wright, was one of the pioneers in financial visualization. Visible Insights acquired the firm in 1999,⁷⁰ and shifted the firm focus almost entirely to ecommerce visualization. They continue to produce custom financial visualizations, from their AdvizorSolutions business unit in partnership with Wright's new company, Oculus.⁷¹

⁶⁷ Inf@Vis! "Visualizing Finance". http://www.infovis.net/E-zine/2002/num_77.htm

⁶⁸ Martin Wattenberg, Interview by author at IEEE Visualization 2001.

⁶⁹ AdvizorSolutions, Inc. http://www.advizorsolutions.com/customer_successes/financial_services.asp

⁷⁰ Chevreau, Jonathan. "Toronto Firm Pioneers 'Financial Animation'" *The Financial Post*, January 4, 1994.

⁷¹ Oculus, Inc. <http://www.oculusinfo.com>

As if in reply to Nasdaq, NYSE has sponsored one of the most ambitious information visualization projects to date – the 3D Trading Floor. The 3DTF is not intended to be a trading interface, but a tool for the real-time management of floor operations. The visualization is unique in that it was designed by architects, and actually follows the floor pattern of the physical exchange. The 3DTF came online in March 1999, and came to fulfill a significant marketing role, projecting a brand and the technological commitment of the exchange.⁷² The next generation of the system is expected to move away from the busy, floor-based design and into cleaner, more abstract representations of the business process.⁷³



Figure 17: The NYSE 3DTF Monitors Operations⁷⁴

⁷² Petroff, Gregory. “The Virtual Stock Exchange: Lessons Learned from the NYSE 3DTF”. Proceedings from *CODATA Workshop*, 10-11 July 2002.

⁷³ Ibid.

⁷⁴ Courtesy NYSE

J.P. Morgan had the foresight to create a formal Visualization Group, supporting a wide range of trading and research operations. At its height, the group created a number of 2D and 3D visualizations to analyze complex options, risk, and trading behavior.⁷⁵ Some have left JP Morgan to found new custom development firms, such as Digital Image Design Incorporated⁷⁶ or Interactiview.⁷⁷



Figure 18: JPMorgan Trading History Viewer⁷⁸

⁷⁵ Night Flight Studios. <http://www.nightflightstudios.com/JPMorgan/viscvr2.htm>

⁷⁶ Digital Image Design, Inc. <http://www.didi.com>

⁷⁷ Interactiview, Inc. <http://www.interactiview.com/>

⁷⁸ Courtesy of JP Morgan

Unified Field is better known for design-oriented views than for visualization tools, but they deserve notice for their beautiful work. Their work for Goldman Sachs is typical, projecting beautiful views of real-time and historical data onto 15-foot screens.⁷⁹ They are not, however, as rigorous as the above visualizations – for example, the attractive waves in ribbon view below represent price changes by day of week,⁸⁰ which is next to useless as a quantitative measure.

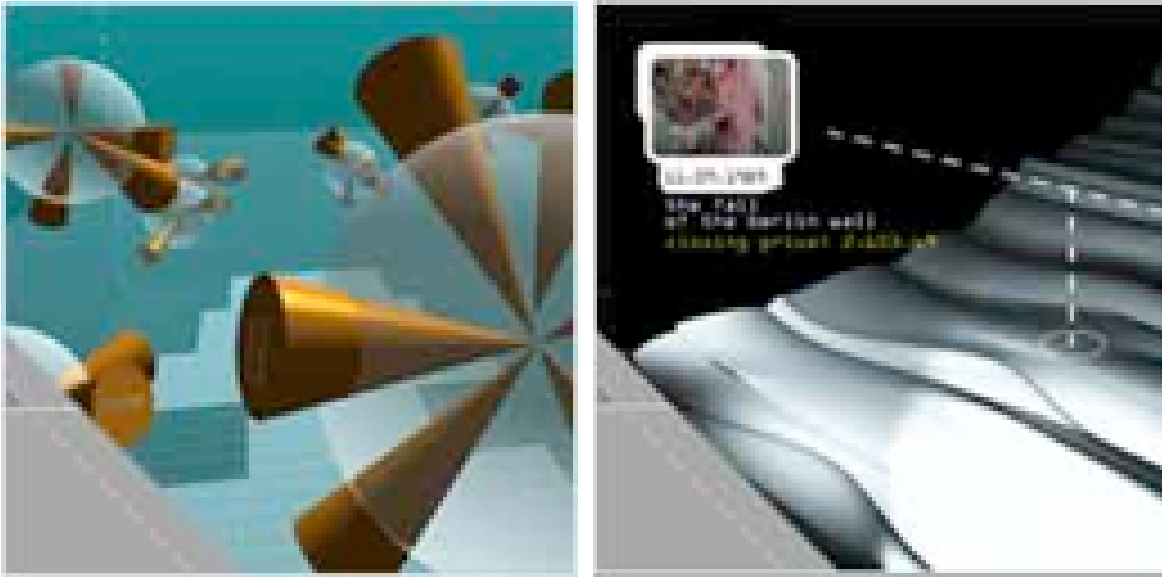


Figure 19: Unified Field Wins the Beauty Contest⁸¹

⁷⁹ Unified Field, Inc. http://unifiedfield.com/4d/case_studies/case.htm

⁸⁰ Eli Kuslansky, Unified Field, Inc. Interview by author at ACM Siggraph 2001.

⁸¹ Courtesy of Unified Field, Inc.

4 A Functional View of the Trading Problem

The goal of this research is to apply and evaluate visualization for the act of trading equities. Therefore, it is important to clearly delineate the work of a trader. At an abstract level, the work can be divided into four main blocks: Asset Search, Timing Search, Execution, and Monitoring. Each of these functions is guided by, and feeds back to, a trading policy. Though in real life the process is not as neat as in this abstraction, each function will be considered separately below:

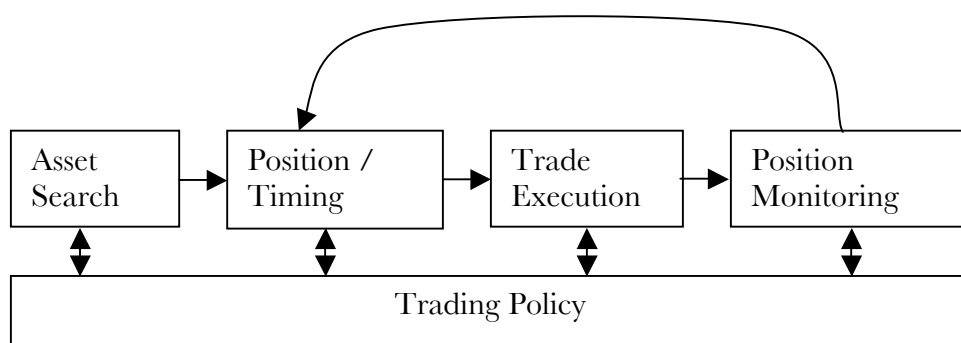


Figure 20: Functional View of the Trading Problem

4.1 Asset Search

For speculative traders, the most obvious and yet elusive task is to identify a trading opportunity, typically a liquid asset with a seemingly predictable price path. Non-speculative traders have the related goal of finding an appropriate hedging or indexing asset. From a universe of thousands of equities, bonds, derivatives, and other instruments, the trader must find one opportunity. Since one opportunity may involve a combination of assets, the potential range of positions is very large indeed – and there is no shortage of advice, good and bad, about trading opportunities.

An experienced trader will typically closely follow a short list of individual assets to limit this search problem. Software can assist the trader in this idiosyncratic process, by searching and visualizing asset characteristics, and by accessing news, rumors, fundamental and macro

information. Charting software is especially useful to technical traders, enabling them to plot trendlines and other indicators to predict future price paths.

4.2 Position/Timing Search

Once a trader has formed a belief about the price path of a particular asset, the task shifts to making the trade. The trader must choose the instrument (usually the underlying asset or a derivative), a position size, and the right time to get into the position. Timing is especially crucial for short-term trades, often taking advantage of rumor-based price changes or intraday volatility.

The goal is to make money, rather than to be right – so controlling risk is paramount. Using stops and limit orders, the trade can be structured to control loss without limiting upside. It is also common to scale into a trade, committing more capital to it once the price path is clearer. Some markets, such as Nasdaq, allow participants to view the book of bids and offers on a name – many traders find this Level 2 data to be vital to understanding the depth of the market and when to act.

Software has a clearer and more vital role in this stage, since real-time trading moves quickly in many simultaneous dimensions. Such a system can be considered a tactical aide, presenting information and alternatives to the trader. Ideally, the system could be set up to recommend a trader's pre-determined strategies, to minimize impulsive, emotional responses.

For historical and practical reasons, traders think and speak about these assets in terms of raw prices, rather than returns or other abstract concepts. Any display of trading data that does not respect this predilection for pennies is doomed to failure, given the prevailing mental model.

4.3 Execution Problem

Once the trader decides to pull the trigger on an opportunity, it is vital to execute the trades quickly and efficiently. Unless the trader is physically on the floor of the NYSE, this is not always as simple as it sounds, particularly for large blocks of shares. Although inefficiencies are measured in seconds or pennies, a bad execution can erase any edge the trader had in an otherwise perfect opportunity.

By spanning multiple exchanges, software is able to help find the best execution for a trade. Current systems are less integrated than one would imagine – even now many traders have multiple terminals for multiple systems to accomplish this task. Vendors such as Lava Trading, Inc.⁸² are able to offer this integration as a major selling point.

4.4 Monitoring Problem

An individual trader might have a dozen or more positions simultaneously open, and a trading firm could have hundreds of open interests. It is an attention-demanding task to keep track of all of these, much less to trade each one effectively. For each monitored item, the trader must decide when and whether to end the position or to add to it. The trader must evaluate risk and act by changing stops or hedges. A trading manager has the additional responsibility of managing the firm's overall risk, as well as watching for unusual trader behavior.

Once the decision has been made to end, change, or hedge a position, the trader's focus returns to the timing and execution problems. Interestingly, traders generally do all of this by watching the numbers change on their screens, and sometimes by flipping between line charts. Some trading packages can be programmed with simple heuristics to automate some of these tasks or to draw attention to important events. The only use of color is to highlight the name and/or price in red if it is down for the day, and green if it is up for the day.

The monitoring problem is open for much improvement. New systems should provide an edge in watching more open positions with more intelligence. Though it is not hard to provide a better real-time interface than a table of numbers, the numbers are not to be underestimated.

⁸² Lava Trading, Inc. <http://www.lavatrading.com>

4.5 Trading Policy

Although the above cycle is messy and fast-paced in real life, there is a method to the madness. Successful traders have a constantly evolving set of policies that govern their behavior through the four functions above. The policy is informed by their performance, as their policy drives future trades. For example, a firm or trader may set their maximum position size to 10% of their trading assets, or they may choose to never place an order without a stop.

An ideal trading system would be able to implement and remind the trader of these policies, and even to learn from past performance. In the heat of battle, traders are very prone to errors rooted in human emotions, and need a plan to carry them through temporarily irrational impulses. This factor is especially true for a larger trading firm with centralized policies, where many traders must be trained, evaluated, and contained from causing extensive damage.

5 A Trading Visualization System

This research posits the hypothesis that information visualization methods can improve the trading functions listed above. In order to partially test this hypothesis, the author built a prototype trading system, leveraging advanced visualization techniques from non-financial disciplines. By creating an experimental task-oriented design, it becomes possible to directly evaluate the look, feel, and performance characteristics of trading in a graphical environment.

While better designs may yet be found, this application is a significant step towards a visualization-enabled trading tool. To purify evaluation against a numeric-based system, this experimental system intentionally displays no numbers on the main screens. Full numeric results are available on a pullout drawer for validation purposes.

5.1 TAQ Data

As a prototype application, it was infeasible to connect to a live stock data feed. In addition to excessive costs, live feeds are not active 24 hours a day and do not produce replicable results. Instead, historical tick data were loaded from the Trade And Quote (TAQ) database, hosted by Wharton Research Data Services (WRDS).⁸³ The NYSE publishes the TAQ database, which contains second-by-second trade and quote transaction data for every asset listed on the NYSE, AMEX, NMS, and some small cap issues.

Two weeks of trading (March 4-15, 2002) for the Dow 30 were chosen as a test dataset. This typical 10-day range had over 2.15 million trades for 30 stocks, or an average of 7,167 trades per day per asset.⁸⁴ The dataset also contained over 3.5 million quotes, which were not used. The 83 MB of trade data were fed into a SQL database for rapid retrieval and analysis. The use of SQL was a significant benefit and is highly recommended for similar financial research.

Each trade record contains the symbol, trade time, price, size, various condition and error codes, and the exchange on which the trade occurred.⁸⁵ Each day in the dataset contains over 215,000 points with 4+ dimensions – a very large visualization dataset, by modern standards.

⁸³ Wharton Research Data Services <http://wrds.wharton.upenn.edu/>

⁸⁴ As will be seen, MSFT and INTC are responsible for half of the transactions in the database.

⁸⁵ New York Stock Exchange, Inc. “CT Binary File”. *TAQ User's Guide*. Version 3.31. 17.

5.2 Mapping Functional Tasks to Views

This prototype focuses on improving three of the functional tasks listed above: the position/timing search, position monitoring, and policy formation/feedback. The asset search problem is wide-ranging and idiosyncratic – other datasets and types of visualization may be well suited to this search, but not a visualization of the tick dataset. The execution problem does not need visualization at all; it needs the kind of fast, intelligent systems integration such as is provided by Lava.

The prototype implements the classic focus+context paradigm. A 2D context view shows all assets, and a navigable 3D focus view shows just one asset in detail. Simulated trades can be placed against any asset from the focus view.

A trader engaged in the position/timing search would open the focus view of the targeted asset, watch the ticks in motion and pull the trigger when training and intuition demand action. Position monitoring can be continued from the same view – multiple windows can support as many individual asset views as a screen will allow. More compact position monitoring can be performed from the context view, showing each asset and all open positions. Variations of the context view could display additional management parameters, such as risk and hedging. A trading manager could use a master context view to see firm-wide positions in real time, as well as the actions of individual traders in these assets.

This prototype would easily accept additional features in regards to the policy function, in both views. Policy could be embodied by position size suggestions, suggestions to increase or end a position, and the like. Since these policies do not yet have purely visual analogies, this prototype did not implement them.

5.3 Single-Asset Focus View

Since tick data is abstract, there is no natural mapping to form, space, or color. However, the time series line plot has a powerful presence in the mind of Wall Street. It is more useful to extend this line plot analogy than to fight it – however, the design can handle far more dimensionality when it is placed in an animated, navigable 3D environment.

Both views in this prototype share the same mapping design, a time series “ribbon”. It is as if the single-asset line plot were made of a long strip of paper, coursing through space. Formally, the mapping is as follows:

- X Axis: Time in Seconds
- Y Axis: Width of Ribbon = Log_{10} of Volume
- Z Axis: Price in Dollars
- Color of Ribbon: Tick-by-tick price change (green = up, grey = zero, red = down)

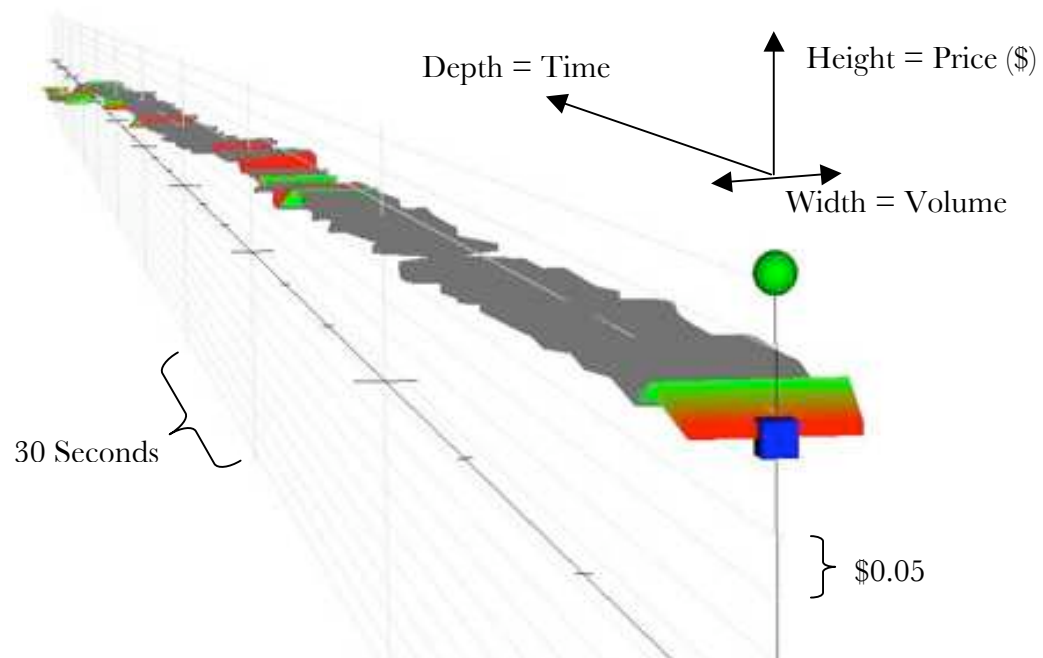


Figure 21: Data Mapping for Focus View⁸⁶

The color of the ribbon follows the financial convention of red for loss, green for gain. The colors and price thresholds are configurable via the following color map editor, to allow for experimentation and the accommodation of color blindness. The realized color follows a S-ramp algorithm between the grey zero value and the target color at threshold. Personal tastes vary, but for many this visualization is clearer on screen with a black background, and for print with a white background.

⁸⁶ TAQ data for MSFT on 3/4/02 at 10:01:39

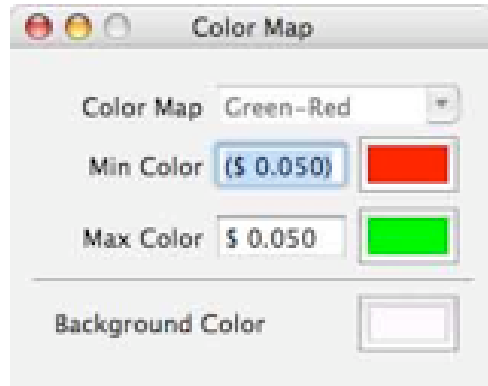


Figure 22: Color Map Setup Dialog

Note that the focus view seen from the side is just like a traditional price line plot, across the “graph paper” of the grey axes. Each horizontal axis represents five cents, and each vertical axis represents thirty seconds. The base of each axis is black – the time base axis is always at the current time, and the price base axis is the previous day’s close price. Thus, the ribbon above the base axis in Figure 21 clearly shows a stock that is up for the day.

On the price axis are five price markers – green and red balls indicate the day’s high and low, a black and yellow diamonds show the day’s previous close and open, and a blue cube shows the current price.

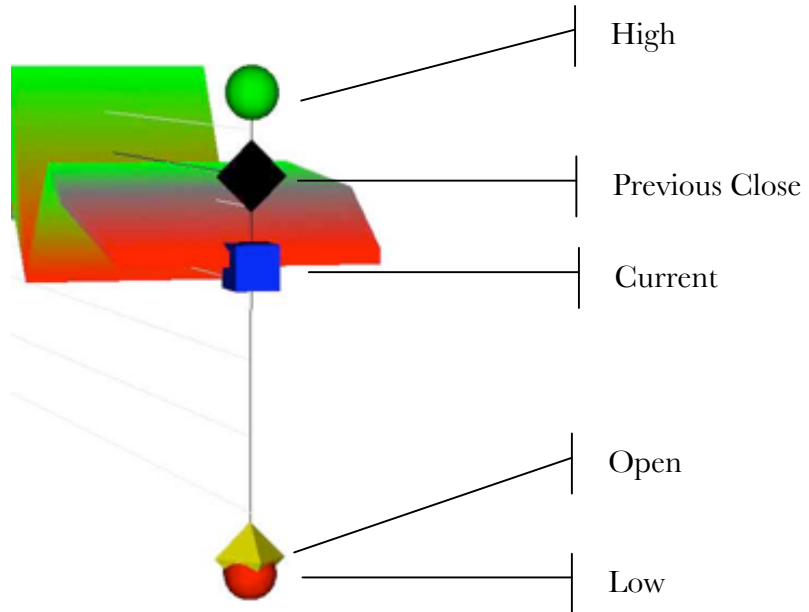


Figure 23: Close, Open, High, Low, Current Markers for Focus View⁸⁷

The diamonds, spheres, and cube are sized such that they can visibly merge, if for example the current price is also the day’s low.

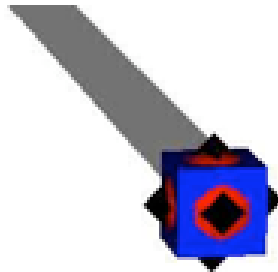


Figure 24: Merged Markers Are Still Visible⁸⁸

In order to reinforce the visualization, key graphical elements are unobtrusively redundant – for example, the top and bottom price grids coincide with the green and red balls, and the

⁸⁷ TAQ data for MSFT on 3/4/02 at 9:31:55

⁸⁸ TAQ data for IBM on 3/4/02 at 9:30:00

colored tick price changes are reflected by height changes. These supporting graphical elements also make the ribbon metaphor portable to other views.

When the animation is running, new data points continually appear at the price axis. Older data points shift down the time axis, creating the illusion that a ribbon is being printed from the price axis. Long-time market watchers may notice this nod to this product's predecessor, the stock ticker.

5.4 Multiple-Asset Context View

Initial designs placed all asset ribbons side by side in 3D. While this was quite visually interesting, it didn't work as a visualization. As can be seen below, asset ribbons would occlude each other and make a specific reading impossible. Also, normalization was difficult between differently priced assets – traders will insist on seeing their assets represented by raw price, rather than even simple abstractions such as return. Finally, the 3D view would not scale – though 30 stocks are somewhat visible, the design would have no hope of showing 300 assets.

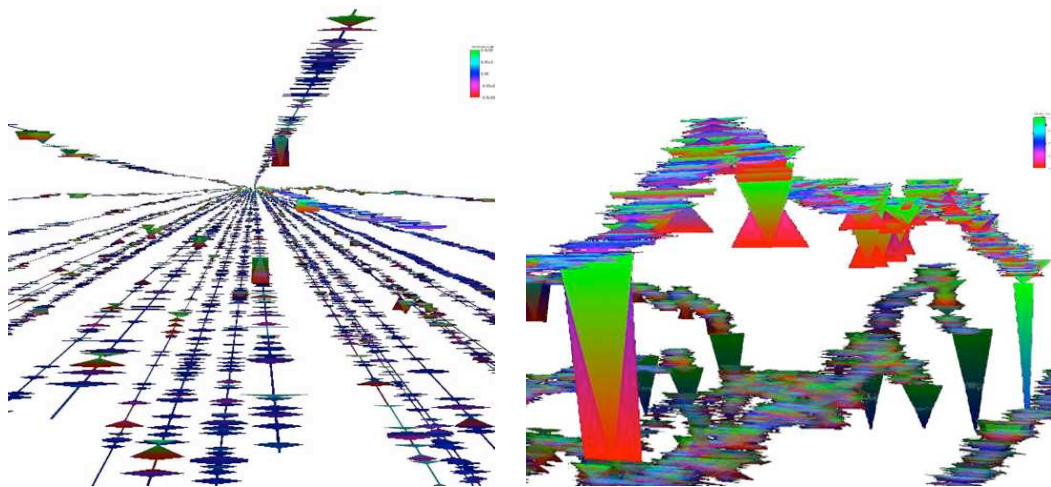


Figure 25: Early Multi-Asset Design

Other design experiments included creating a single landscape-like surface for all assets, but this suffers from even more occlusion and false signals. Since the order of assets is somewhat arbitrary (by sector, then alphabetical) there is no logical reason to graphically link neighboring assets' returns.

By viewing the same asset ribbons from the top, in parallel projection, many of the above problems are fixed. The raw price dimension is lost, but price changes remain visible as color. The result is a terse, scalable display that can show volume and price changes for many assets at once.

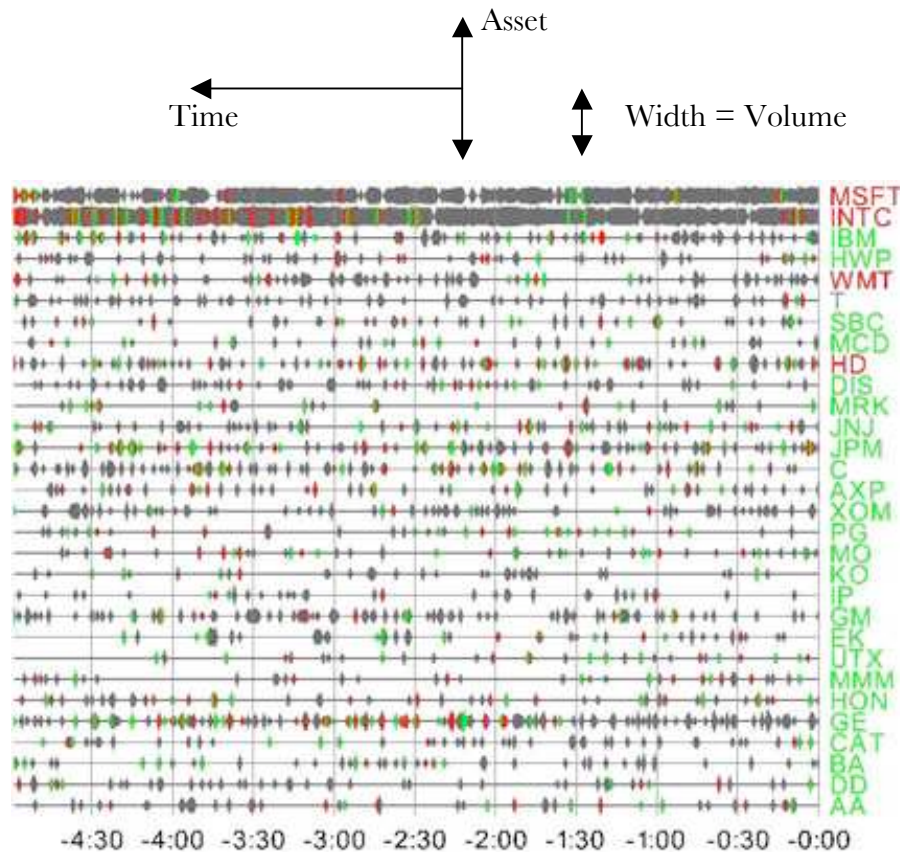


Figure 26: Data Mapping for Context View⁸⁹

As in many financial displays, the color of the symbol text is green if the name is up for the day, and red if it is down.

5.5 Trading Interface

A trade dialog is initiated from a button on the single-asset view. It will be an interesting piece of future work to define a completely graphical method of trade entry that provides the same transactional clarity and precision as this simple dialog – there are complex human interface issues associated with graphically defining a precise order size, much less limit prices or stop loss prices. Since large sums of money ride on each trade, this interface designer chose to

⁸⁹ TAQ data for Dow 30 on 3/4/02 at 9:56:15

err on the side of simplicity and clarity. Since the purpose of this prototype is to explore visualization concepts, the trading system did not implement certain real-world features, such as transaction costs, stops, or only allowing short sales on an uptick.



Figure 27: Trade Entry Interface

Once each order is entered, a cube representing the position appears next to the single-asset ribbon. The cube has its base set at the position's entry price, and the height of the cube represents the price gain or loss in the position, reinforced by a green or red color. A short sale is identically configured, but the gain or loss is inverse to the price.

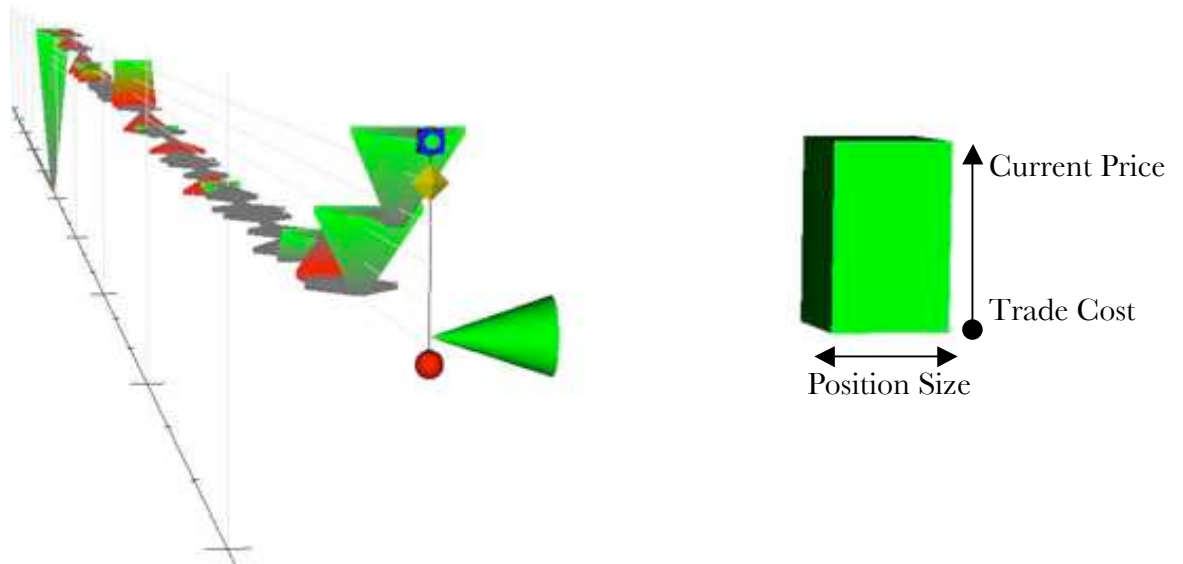


Figure 28: Each Position Lot is Represented as a Cube⁹⁰

⁹⁰ TAQ data for IBM on 3/4/02 at 9:32:41

A cone, attached directly to the price axis, represents the aggregate position. The horizontal length of the cone is proportional to the total position size, and the cone's axis, or point, is located at the volume-weighted average cost of the aggregate position. As before, the color indicates whether the position is up or down. The result is a simple indicator for the aggregate position, which can be seen quickly in relation to the other price icons.

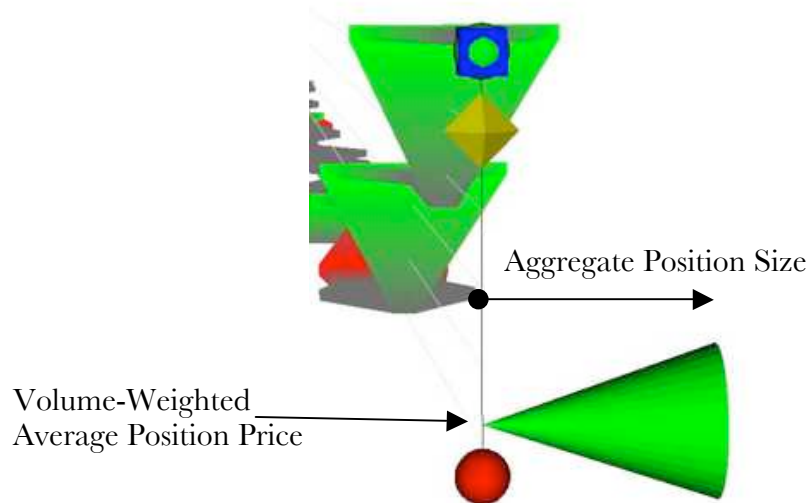


Figure 29: Aggregate Position is Represented as a Cone⁹¹

⁹¹ TAQ data for IBM on 3/4/02 at 9:32:41

When a trade is closed, the cube becomes semi-transparent and fades back into the timeline. Multiple trades are represented as adjacent cubes, that is, separate lots.

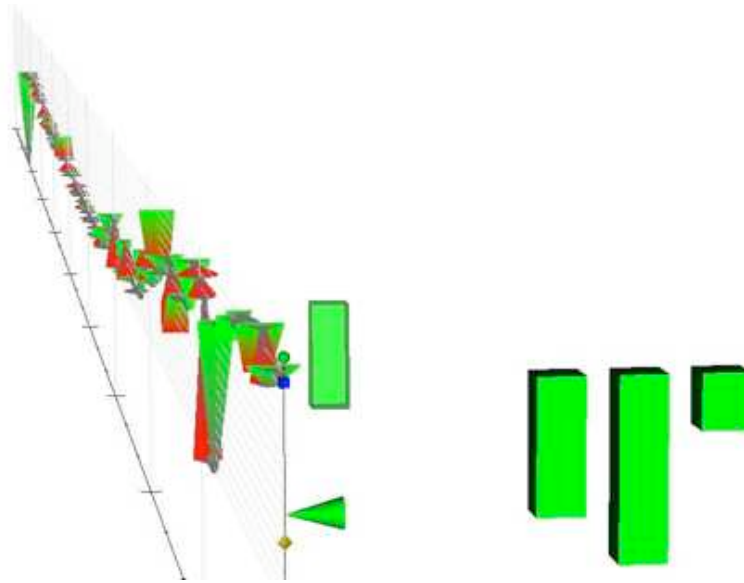


Figure 30: Multiple Trade Lots – 3 Open, 1 Closed⁹²

To date, specific lot information has been of interest primarily to tax accountants. The intent of showing this optional detail is to visually depict actual trading strategies in action, providing feedback to the individual's or firm's trading policy. For example, the common practice of scaling into a trade is clearly visible in the following visualization:

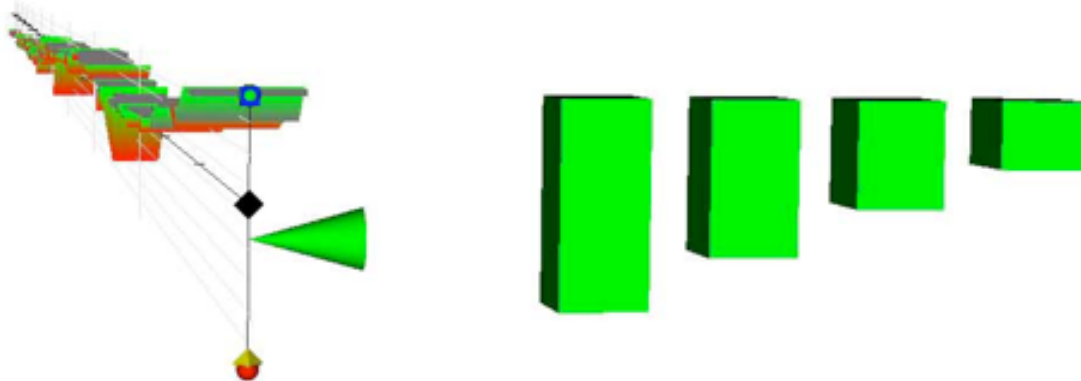


Figure 31: Scaling into a Successful Trade⁹³

⁹² TAQ data for IBM on 3/4/02 at 9:34:12

On the multi-asset context view, the aggregated position appears as a rectangle to the right of the symbol name, showing the size and profitability of open trades in that name. This view represents an aggregation of all open trades, providing the trader with a single monitoring post.

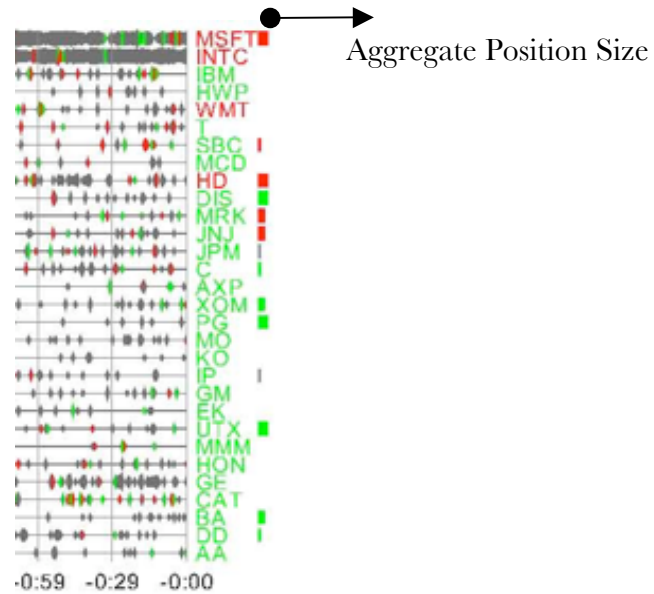


Figure 32: Context Screen Shows Multiple Positions⁹⁴

To provide numeric backup, the application also includes a traditional blotter screen, showing aggregated positions or individual lots.

⁹³ TAQ data for MSFT on 3/4/02 at 9:33:42

⁹⁴ TAQ data for Dow 30 on 3/4/02 at 9:46:06

5.6 Data Animation

In order to emulate a live data feed, the program creates a stream of historical data from the SQL database. This detail enables a number of conveniences that would not be possible in a live application – time can be stopped, slowed down, or sped up.⁹⁵ It is made possible by a technical innovation, whereby a moving window of data is animated within the visualization pipeline. This is a dramatic shift from typical visualization systems, in which an entire static dataset is loaded and manipulated. Data animation enables this system to visualize a much larger dataset with far more flexibility than would otherwise be imaginable, even on a laptop computer.

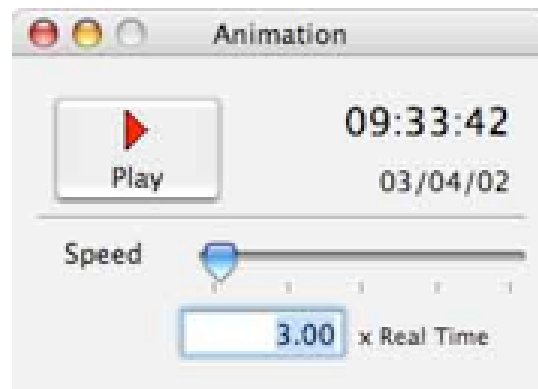


Figure 33: Animation Control Panel

5.7 Navigation Issues

Navigation in three-dimensional space has historically been a difficult problem for many users,⁹⁶ and this prototype's single-asset focus view is no exception. The issue is somewhat simplified by a side effect of data animation, whereby the data are continually refreshed at the same X and Y origin. Of further assistance, the virtual camera has been programmed to follow the asset up and down the Z-axis as it changes price. But this camera-following algorithm is not perfect, and the user must occasionally intervene to get an ideal view. Motion is also useful for many users as a perceptual aid, gaining better depth perception and seeing around occlusions.⁹⁷ Moving the camera or viewpoint in a view is essential for many viewers to gain a sense of presence or control.

⁹⁵ It is technically possible to reverse time as well, but this introduces cheating possibilities for the trading simulation.

⁹⁶ Strausfeld, *Embodying Virtual Space to Enhance the Understanding of Information*.

⁹⁷ Ware, *Information Visualization: Perception for Design*, 282-4.

User interaction uses the default “trackball camera” style, where a combination of mouse and keyboard entries moves the virtual camera through the scene. To a novice user, it is still far too easy to get confused and stuck in an undesirable viewpoint – the system should ultimately turn off freestyle navigation and provide a more controlled user interface to control the user’s viewpoint.

5.8 Technical Implementation

Initial concepts for this application considered a full-screen game-style environment. Game engines, such as Quake or Unreal,⁹⁸ were seriously considered and ultimately rejected. The game approach is interesting for a simplified, single-asset world, but falls apart under more complex demands. Game engines typically expect to manage a single landscape, which did not make sense for this data set. The multiple views required for context and multiple asset focus views are strained or impossible in a full-screen engine. It is also a well-known secret that game engine performance is achieved through a wide range of graphical and computational tricks. As has been noted, the human eye is easy to fool, and is often fooled intentionally in the name of game performance. The use of these shortcuts could undermine the accuracy of the financial visualization and lead to an improper trade.

Another approach considered was to program directly to OpenGL or a lower-level graphics API. This quickly proved to be a costly exercise, yielding highly specific, rigid results.

The best approach found was to leverage a scientific visualization library. Ultimately the graphics library chosen was the Visualization Toolkit (VTK), an open-source, multi-platform scientific visualization library.⁹⁹ VTK is heavily used in medical and computational fluid dynamics visualization applications. It has a full-featured graphics pipeline, extensive support for large datasets, and a good reputation for data integrity. Since VTK is commonly used with vector field visualization, it was particularly well suited to the problem at hand.

This prototype application was written for the Apple OS X platform using a combination of Objective-C and C++, consisting of over 10,000 lines of code.

⁹⁸ Epic Games, Inc. <http://www.unreal.com>

⁹⁹ VTK is published by Kitware, Inc. <http://www.vtk.org>

6 Observations

With these new eyes on the market, what can be seen? Do these views help choose trade timing? Do the views help to monitor many simultaneous positions? Several observations emerge upon visual examination of the dataset:

6.1 Dominance of MSFT and INTC

One persistent feature of the context view, on all days, at all times, is the thick volume of MSFT (Microsoft) and INTC (Intel). While all of these Dow 30 stocks are considered to be the most liquid in the world, two technology stocks vastly exceed all of the others. This heavy volume is visible as the consistently thick horizontal lines for these assets, compared to the thin and variable lines for all others.

Looking at the numbers directly, it can be seen that MSFT and INTC comprise 46% of the transactions, 28% of the share volume, and 26% of the dollar volume during this 10-day period.

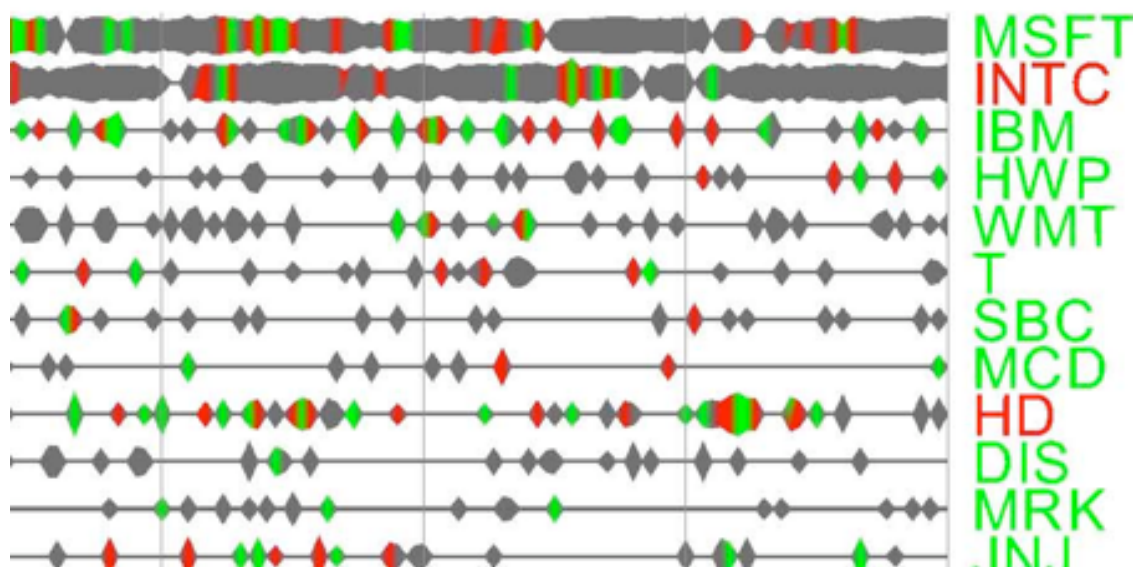


Figure 34: MSFT and INTC Dominate Trade Volume¹⁰⁰

¹⁰⁰ TAQ data for Dow 30 on 3/4/02 at 9:38:25

This volume is also clear upon comparing single-asset charts:

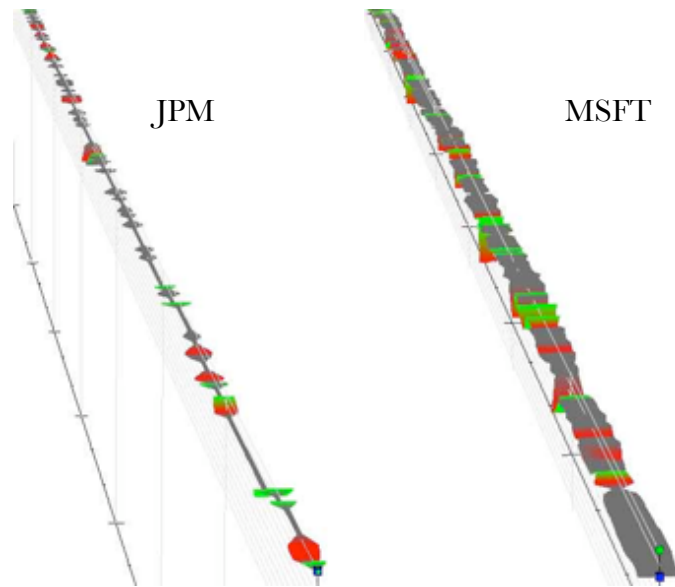


Figure 35: MSFT Consistently Shows More Volume on Single-Asset Charts¹⁰¹

Although a statistical analysis or trading floor experience can yield the same insight, these simple views are able to show the pattern immediately. Line thickness is a prime example of graphic elements that are pre-attentively processed,¹⁰² so it can be legitimately claimed that this pattern is literally discovered before you think about it.

¹⁰¹ TAQ data for JPM, MSFT on 3/4/02 at 9:37:52

¹⁰² Ware, *Information Visualization: Perception for Design*, 165.

6.2 Opening Volume

The single-asset focus view shows a characteristic green or red triangle on market open, showing the first trades of the day, above or below the previous day's close price. This shape and color pattern would be expected, but the graphs universally show unusually high volume at the opening. Volume is mapped to width, using a log scale – the discrepancy is visible as an overhang.

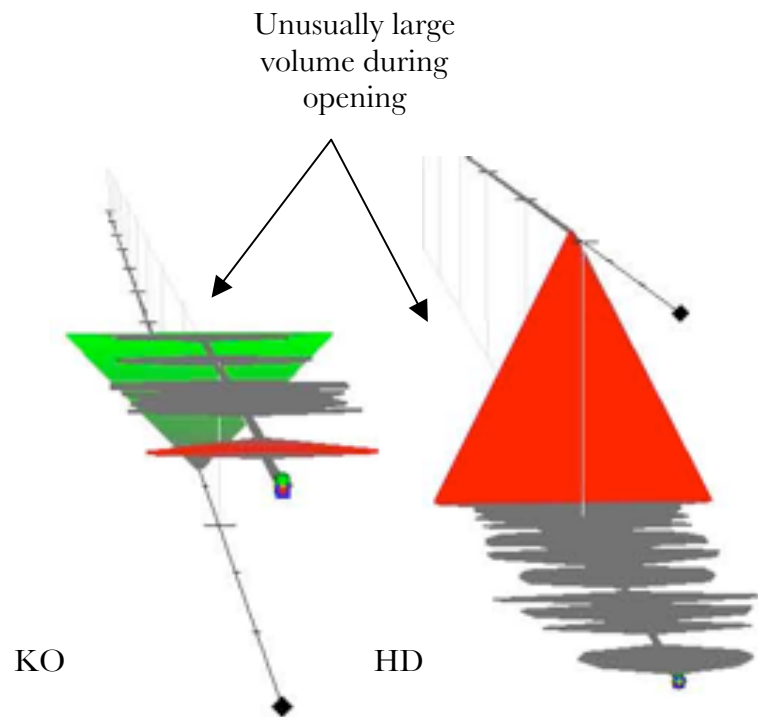


Figure 36: Extra Volume at Opening Bell¹⁰³

¹⁰³ TAQ data for KO on 3/4/02 at 9:30:55, HD on 3/4/02 at 9:36:03

The pattern is also visible in the context view, with thick volume at market open for most assets:

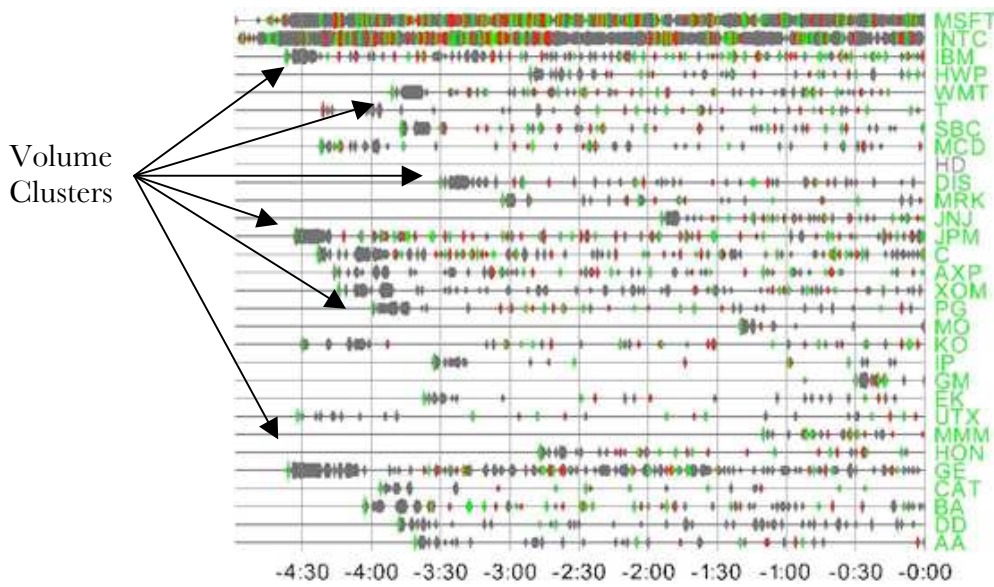


Figure 37: Context View of Open Shows High Volume¹⁰⁴

The reason for this pattern is that the NYSE operates as a “call market” at opening bell. In such a market, the specialist looks at the entire book of orders, including those from the previous evening, and chooses a single price that will execute the most orders, in shares. This special pricing manifests the large volumes seen in the above charts.

This phenomenon could be explored more deeply by mapping the ticks’ sale condition codes to ribbon color, instead of price change. This way, the ribbon color would clearly indicate special circumstances for clear correlation with price and volume action.

¹⁰⁴ TAQ data for Dow 30 on 3/4/02 at 9:34:41

6.3 Propagation & Basket Selling

It was hoped that skillful market visualization would be able to show patterns of price change propagation between assets. Because of basket selling, there should be clear correlations between sales of these Dow 30 equities. Further, there are news-related trends through the day that should move investments between sectors. By grouping the equities by market sectors, it would be possible to see gains and losses by sector.

In the context view below, there are hard-to-find, apparent patterns of gains or loss. For example, MSFT, INTC and IBM show several joint up- and downticks, as circled. The larger circles indicate other short-term correlations in volume and tick direction.

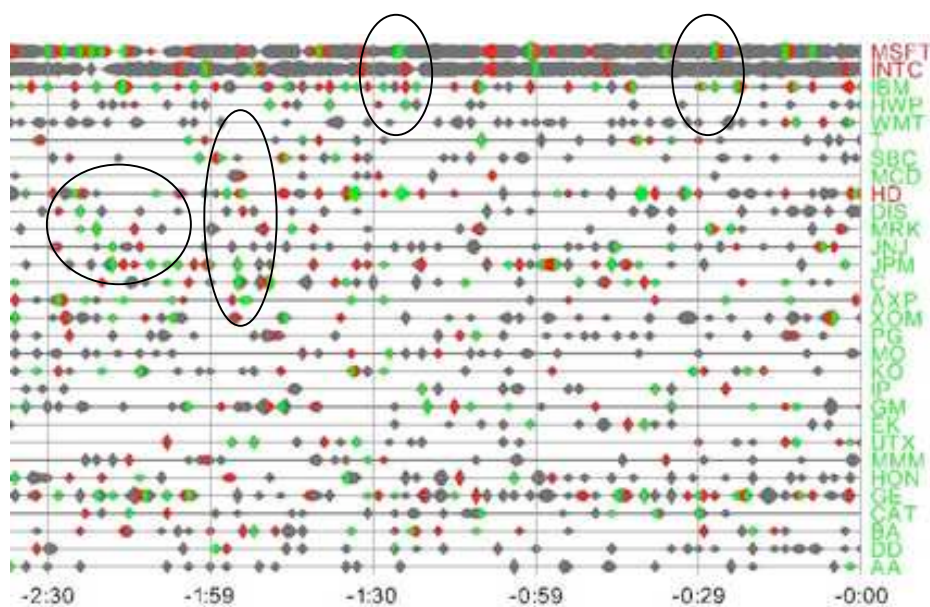


Figure 38: Searching for Price Propagation¹⁰⁵

Of course, these patterns could also be random - but the common practice of basket selling should create short-term correlations. The primary structural reason they are so difficult to find is inconsistent execution time. Though a basket may go to the floor to sell all at once, each asset will be sold separately over many seconds, muddying the waters.

¹⁰⁵ TAQ data for Dow 30 on 3/4/02 at 9:41:49

Visually, this view is also very hard to scan for such patterns. Once the view is in motion, the task becomes almost impossible. Changing the color map thresholds or background color does little to improve pattern recognition. Worse, after a short while the human mind becomes open to suggestion, imagining patterns where there are none. The climatic scene of *Poltergeist*, where a small girl is drawn into a static-filled TV screen, comes to mind after searching for propagation patterns in this view for too long.

The human visual system is not good at this scanning task, and tires quickly. This inability is closely related to change blindness and vigilance fatigue, and has been well documented in cognitive studies.¹⁰⁶ In order to improve at the task of finding price propagation, this system would need to abandon displaying raw data, and move to a more abstract view of derivative data. Also central to this problem is the ordering of assets – in order for any pattern to become clear, the correlated assets must be near each other. But there is no natural ordering for the assets – sectors are used here, but correlations may occur across other dimensions. A self-organizing asset list may resolve this issue, in which case individual assets would become difficult to follow over time.

¹⁰⁶ Rensink, “Internal vs. External Information in Visual Perception”.

6.4 Price Change Behavior

As the equities are bid up and down, their visualizations show characteristic price-seeking patterns far more vividly than line plots or raw numbers. The price action in motion appears to be “fighting” for a new level, feeling as if it is telling a story. This display is simply reflecting the human auction, a dialogue on the floor of the NYSE and on electronic exchanges. A trader’s ability to detect strength in these movements could be a profitable trading edge. In color and motion, the quantitative becomes qualitative.

For example, this chart shows JPM gaining significant share value, \$0.20, in only thirty seconds. After this rally, the price stays at the new level for many minutes. A trader’s edge would come from identifying the strength at the beginning of this pattern, and buying. Interestingly, the trader’s sale would become one of the ticks in the developing pattern.

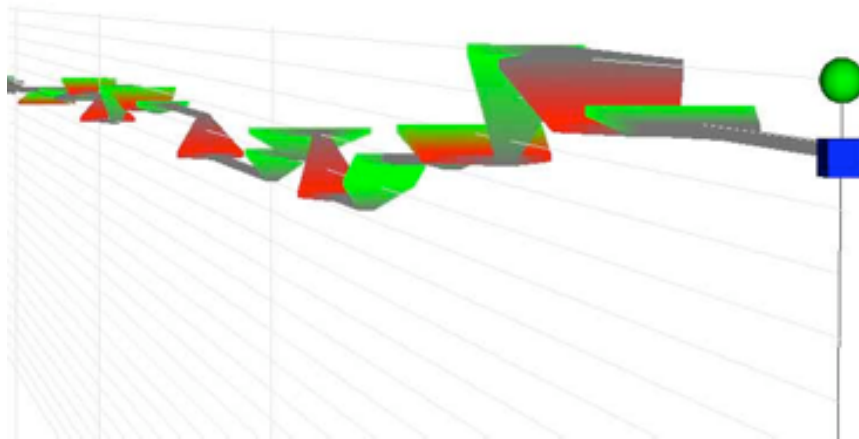


Figure 39: JPM Gains \$0.20 in 30 Seconds¹⁰⁷

¹⁰⁷ TAQ data for JPM on 3/4/02 at 10:01:43

Which kinds of price change patterns have persistent results, and which fail to stick? The following example shows price exploration for HON, where four downticks to a lower price are refuted by trades at the higher, original price. Finally, the fifth downtick sticks and ongoing volume proceeds at the new price, \$0.09 lower.

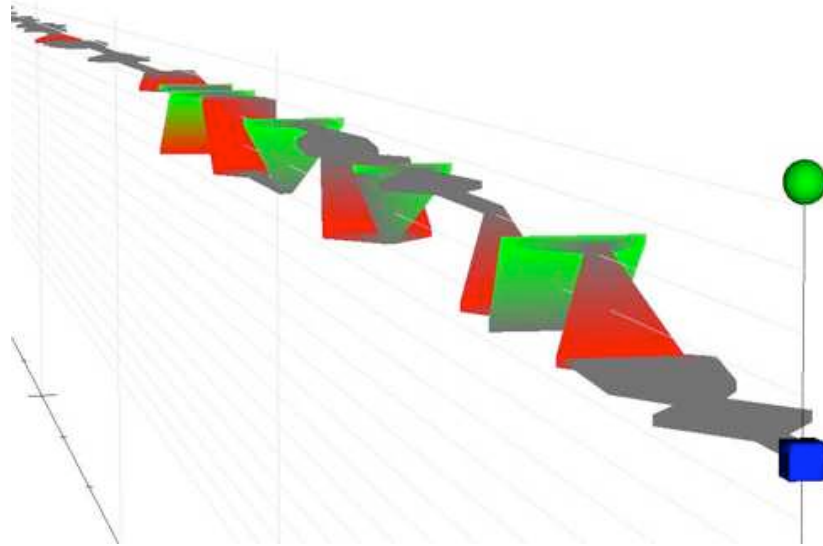


Figure 40: HON Loses \$0.09 in 45 Seconds¹⁰⁸

Six minutes later, the HON price is bid back up, successfully pushed down a few cents, and then strong bids resume at the new price, back up \$0.09. As with JPM above, there is a small, one-tick overreaction in the price adjustment.

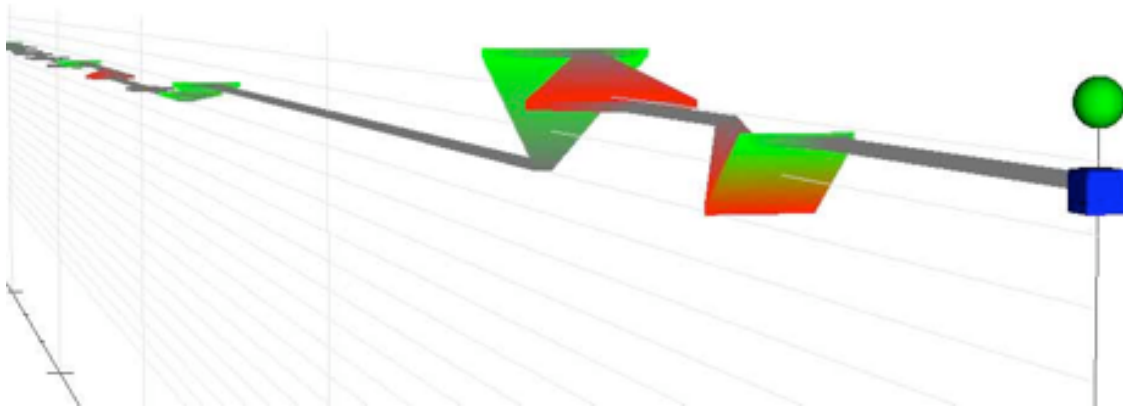


Figure 41: HON Gains \$0.09 Back in 15 Seconds¹⁰⁹

¹⁰⁸ TAQ data for HON on 3/4/02 at 10:08:52

These stories are not definitive – they only suggest the kind of visual reasoning that may be possible for a full-time trader watching these views. One might choose to spice up these stories with structural issues, rumors about the asset, or a simple battle between red vs. green. Alternatively, a pure economist might classify this dialogue as a fantasy to the tune of Brownian motion.

These visualizations could be augmented with icons to indicate news events, traders' annotations, or technical indicators, such as trend lines or moving averages. However, since this is a tactical interface, it is important to keep the view clean and free of extraneous pixels.

6.5 Late Prints & Errors

The application normally screens out ticks marked in TAQ as errors or duplicates. However, in the flow of normal price action, anomalies are still occasionally observed. Many of these are so-called “late prints” or “sold sales”, transactions that occurred earlier in the day but are applied to the tape with a later time.

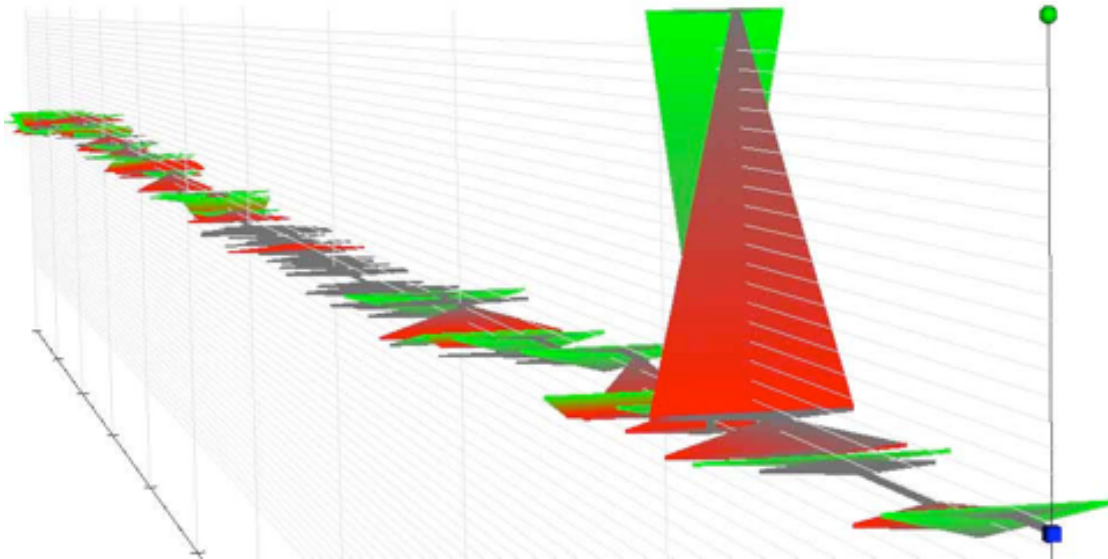


Figure 42: A Typical Late Print¹¹⁰

In the example above, the abnormal trade in IBM was not marked in the TAQ database as an error or late print, though it probably should have been. Modifying the query to intentionally include error-marked data simply causes more of these blips to show up. Some errors are dramatic, but most are small enough to be lost in the deluge of correct trades.


¹¹⁰ TAQ data for IBM on 3/4/02 at 10:12:02

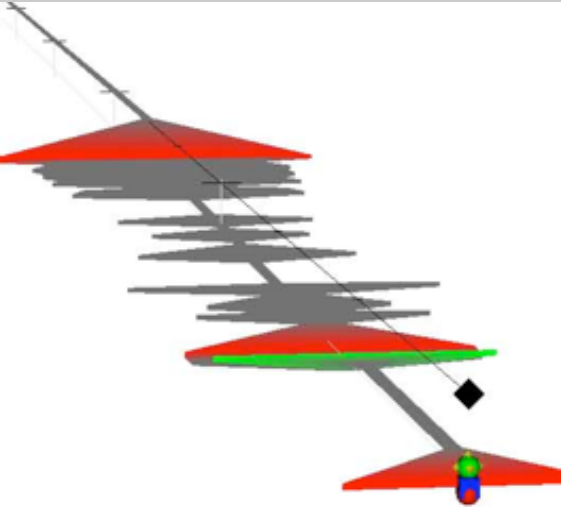
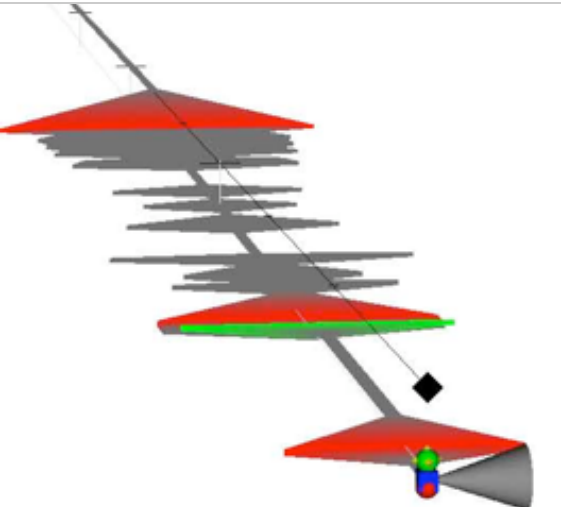
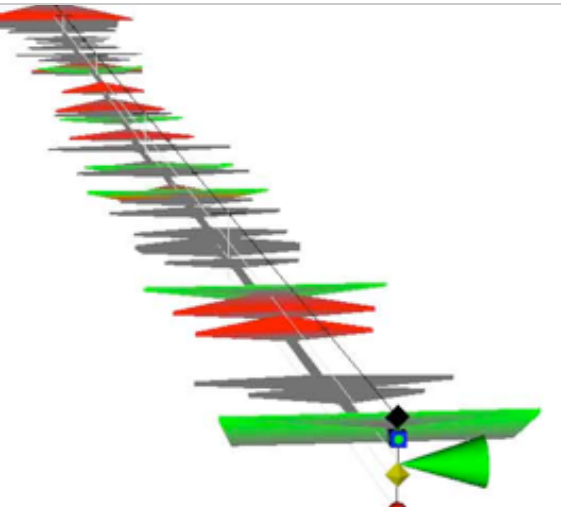
7 A Day of Trading

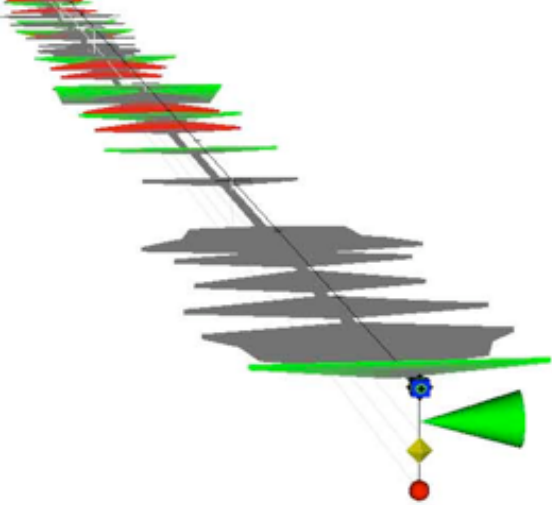
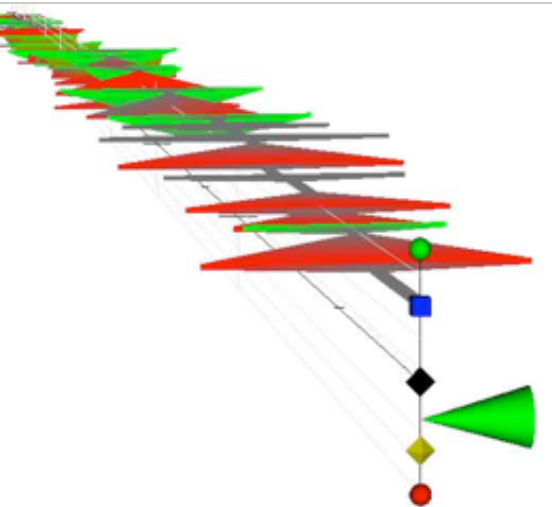
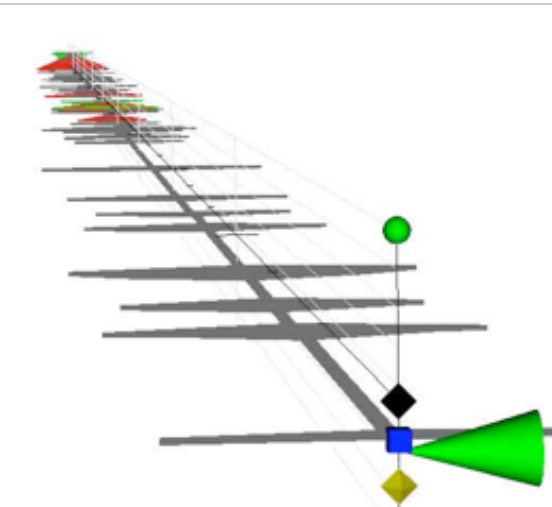
The simple presence of a working “Trade” button on the interface radically changes the meaning and intent of this visualization. The potential to interact with the data, much less to make money from it, transforms the tool from a detached laboratory tool into an opportunity-seeking engine. The act of trading turns very much into a video game. This is a mixed blessing – while well-informed, visceral engagement with the market is good, the visual stimuli may over-stimulate already volatile emotional trading traps.

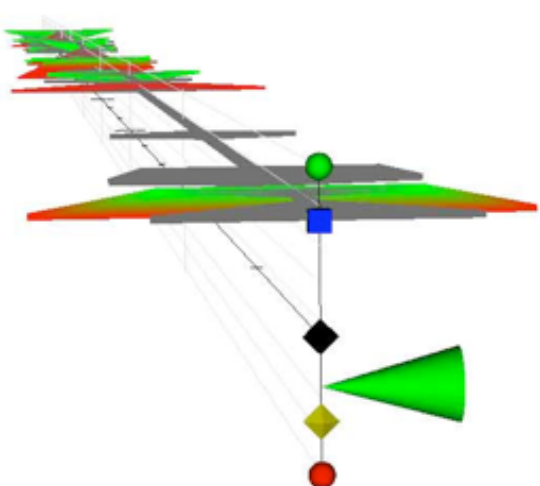
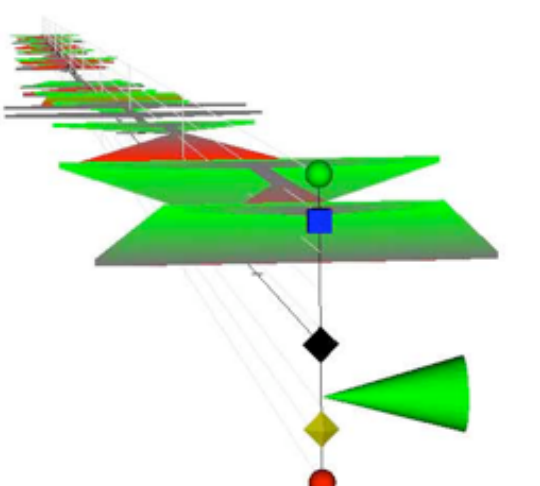
In order to explore this dynamic, the author sat down to trade one asset with this prototype for one day. AXP was chosen as the target equity based on past liquidity and volatility for this short trading period. Only the graphical views would be used, ignoring the numeric prices. The price behavior for this date, 3/5/02, was not known before trading. The only “cheating” behavior employed was to occasionally stop the clock in order to take notes, or speeding up the clock to get through long flat periods. The following log chronicles the experience.

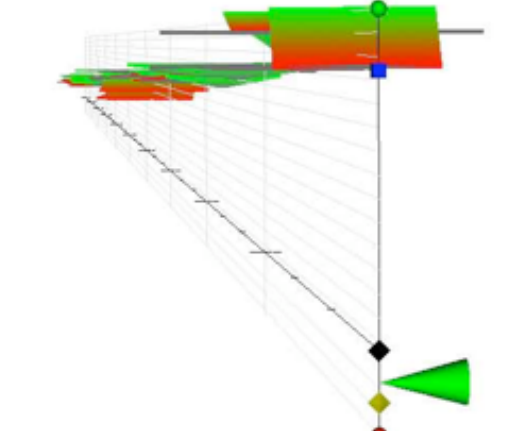
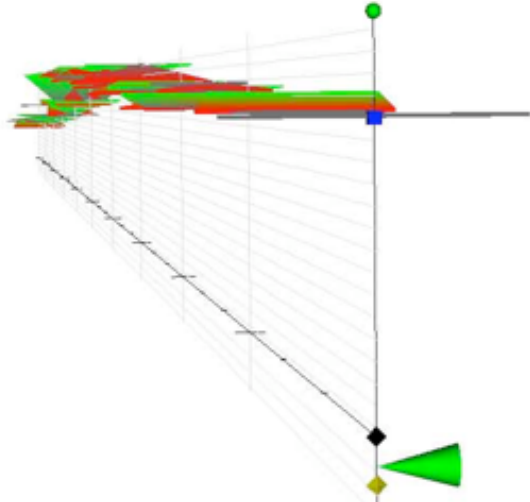
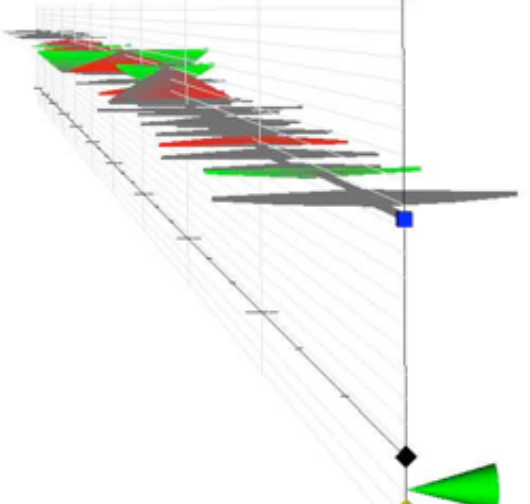
7.1 Traders Log

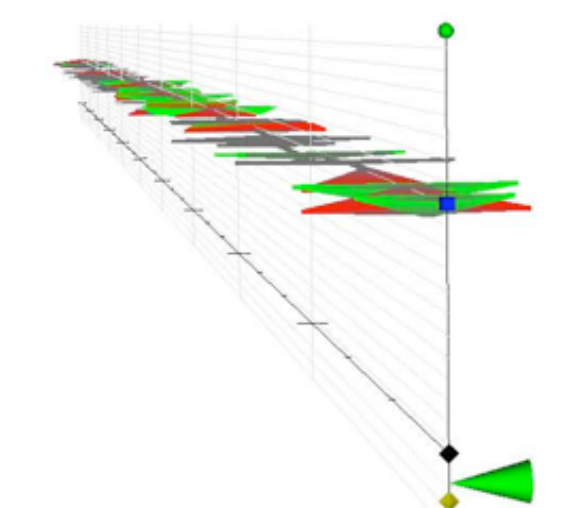
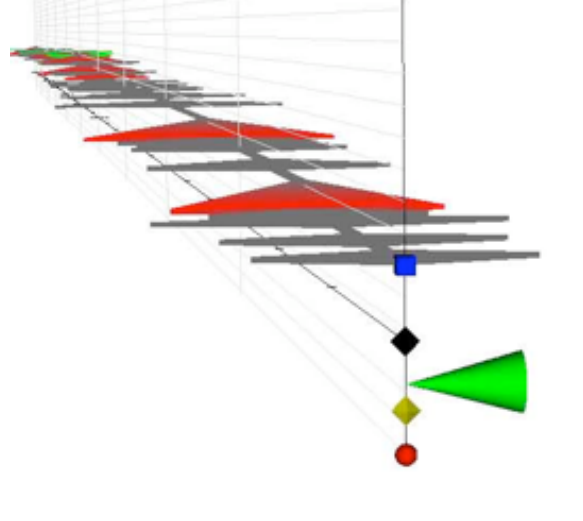
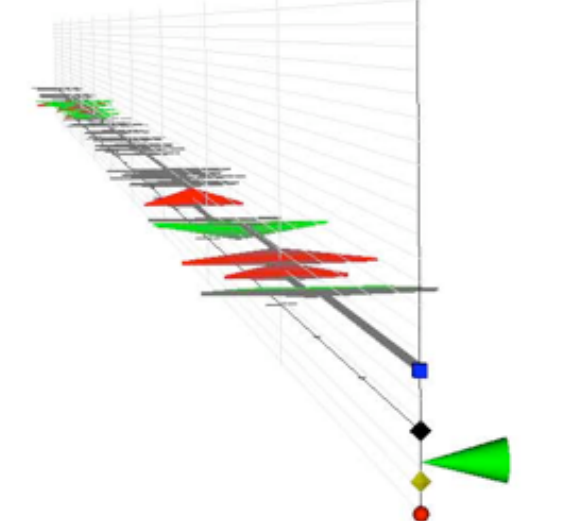
3/5/02 9:30:00	A large single-asset window was set up and positioned to see prices, but not position lots.	
3/5/02 9:30:17	AXP opens slightly down, soon after market open. Total Position: 0	

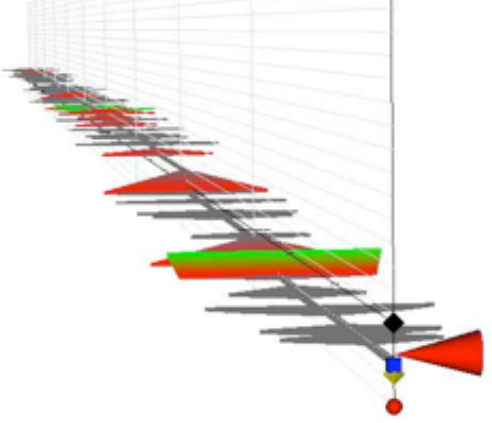
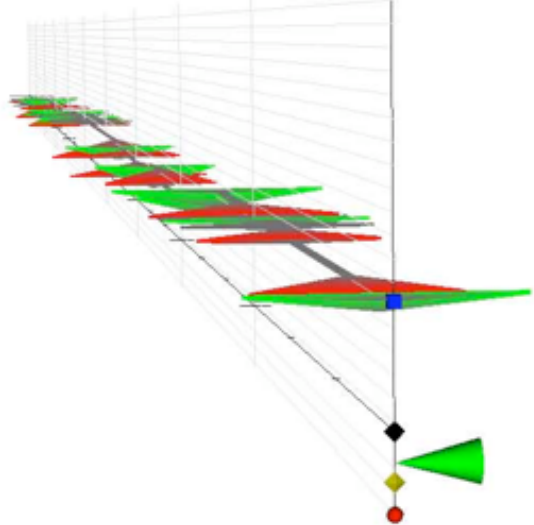
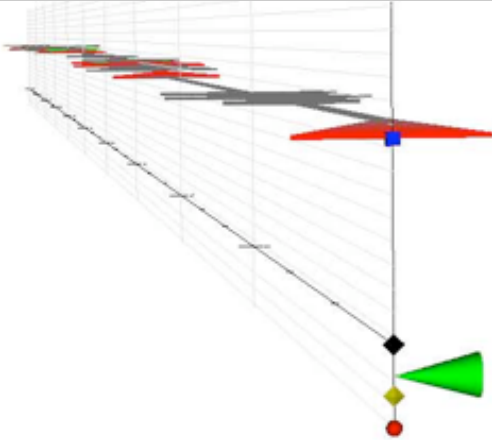
<p>3/5/02 9:31:01</p> <p>After a couple of down ticks, the trader starts to believe that it will bounce back.</p> <p>Total Position: 0</p>		
<p>3/5/02 9:31:03</p> <p>Ready to bet on this pattern by going long. Will scale into the trade more if it keeps up the momentum. Will stop out of the trade if it falls two notches below open.</p> <p>Trade 1: Long 1000 shares.</p> <p>Total Position: 1000</p>		
<p>3/5/02 9:32:26</p> <p>Big uptick with strong volume, but still below the closing price. Will scale in more long shares.</p> <p>Trade 2: Long 1000 shares.</p> <p>Total Position: 2000</p>		

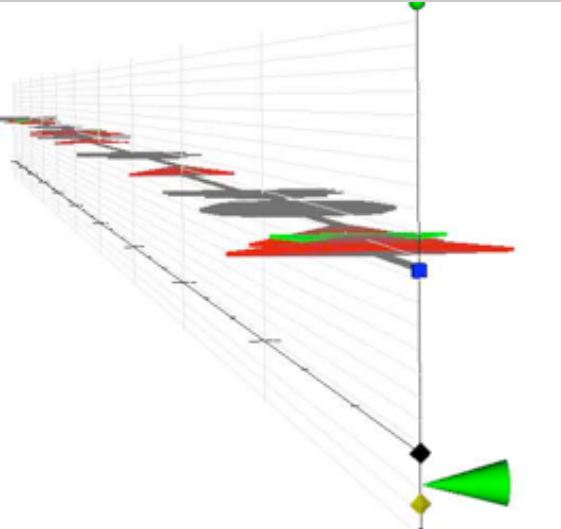
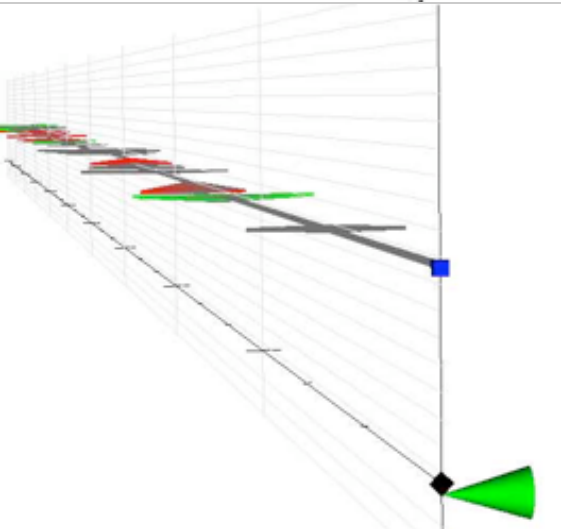
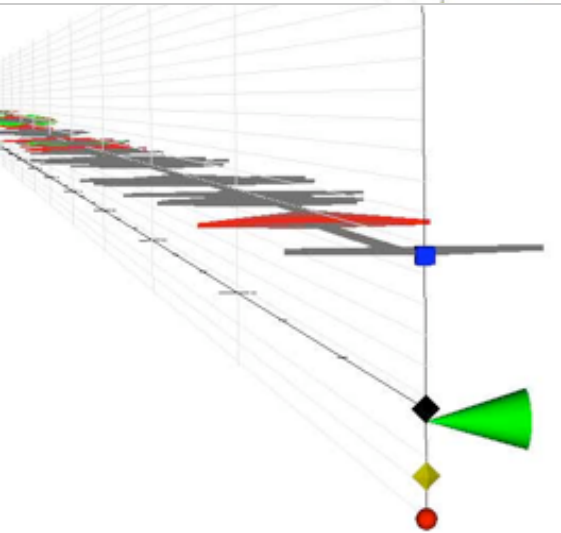
<p>3/5/02 9:33:27</p>	<p>Price has edged up to previous close. Scale in more long, though this might be too high. After this trade, sit back and let it run. Will stop out if it goes below open.</p> <p>Trade 3: Long 1000 shares</p> <p>Total Position: 3000</p>	 <p>The chart shows a price line moving upwards from left to right. A green triangle on the right indicates a trade execution. Below the price line, a vertical line has a blue circle at the top, a yellow diamond in the middle, and a red circle at the bottom. A green cone is attached to the bottom of this vertical line.</p>
<p>3/5/02 9:38:07</p>	<p>Price continues to rise, a bit volatile at this new level. Hold steady, making money.</p> <p>Total Position: 3000</p>	 <p>The chart shows the price line continuing to rise but with more frequent and larger price fluctuations. A green triangle on the right indicates a trade execution. Below the price line, a vertical line has a green circle at the top, a blue square in the middle, a black diamond in the middle, a yellow diamond in the middle, and a red circle at the bottom. A green cone is attached to the bottom of this vertical line.</p>
<p>3/5/02 9:44:23</p>	<p>After rally, price drops down below previous close, just above average position cost. Still optimistic, willing to bet that there is more upward volatility than down. Buy more while it is low.</p> <p>Trade 4: Long 1000 shares</p> <p>Total position: 4000</p>	 <p>The chart shows the price line dropping downwards from left to right. A green triangle on the right indicates a trade execution. Below the price line, a vertical line has a green circle at the top, a black diamond in the middle, a blue square in the middle, a yellow diamond in the middle, and a red circle at the bottom. A green cone is attached to the bottom of this vertical line.</p>

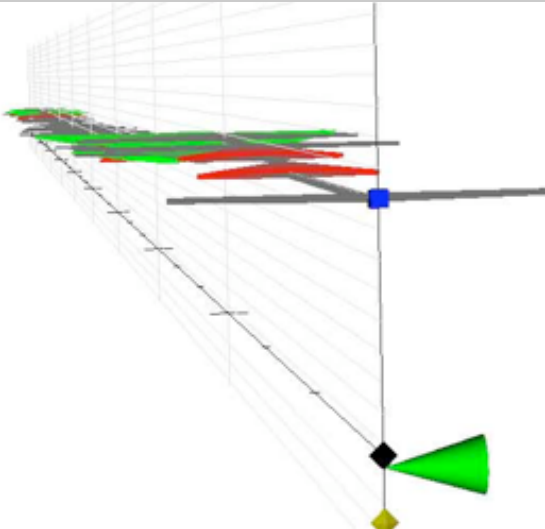
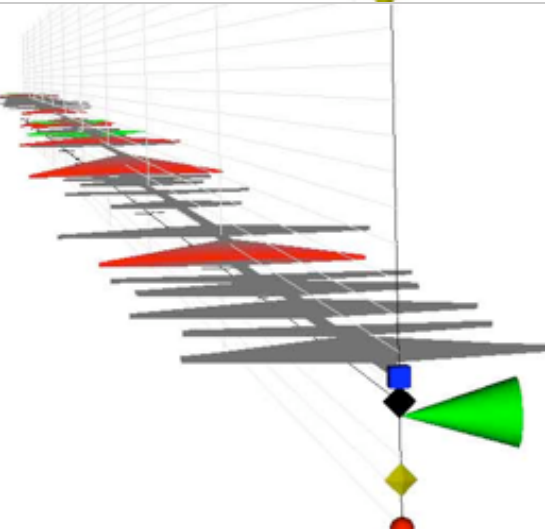
<p>3/5/02 9:52:27</p>	<p>Very happy to see price go back up and stay flat at new level. Will hold steady unless price dips down to average cost again, or stop out if it reaches below open.</p> <p>Total Position: 4000</p>	
<p>3/5/02 9:59:53</p>	<p>Price takes a sudden, sharp drop down below average cost – strategy unchanged, choose to buy more shares.</p> <p>Trade 5: Long 1000 shares</p> <p>Total Position: 5000</p>	<p>(no snapshot)</p>
<p>3/5/02 10:03:09</p>	<p>Price bounces back up quickly, affirming belief in upward skew.</p> <p>Total Position: 5000</p>	

<p>3/5/02 10:08:27</p>	<p>Impressive rally. Hold steady. Little concern now about stopping out of trade. Instead, wondering how and when to lock in these gains.</p> <p>Total Position: 5000</p>	
<p>3/5/02 10:11:35</p>	<p>Price settles at new level, with over \$3000 in gains for the day. Decide to sell 2000 shares if price drops below midpoint.</p> <p>Total Position: 5000</p>	
<p>3/5/02 10:28:04</p>	<p>Price has dropped close to that midpoint. Decide to ride it out.</p>	

<p>3/5/02 10:41:34</p> <p>Price remains choppy, but essentially flat near midpoint. Will stay the course.</p> <p>Total Position: 5000</p>		
<p>3/5/02 10:57:23</p> <p>Price edged down below selling point, without selling. Disappointed at lack of discipline, but still profitable.</p> <p>Total Position: 5000</p>		
<p>3/5/02 11:10:06</p> <p>Flat at new level. Resigned to new prices.</p> <p>Total Position: 5000</p>		

<p>3/5/02 11:26:04</p>	<p>Sudden price decrease throws the portfolio into a loss for the first time. Price is still above stop, so staying the course. Still believe in upward skew of this price distribution.</p> <p>Total Position: 5000</p>	
<p>3/5/02 11:52:15</p>	<p>A small rally just before lunch places the portfolio back in the green.</p> <p>Total Position: 5000</p>	
<p>3/5/02 12:58:36</p>	<p>Flat through lunch. Not sure if it is going anywhere. Will sell 3000 at next peak.</p> <p>Total Position: 5000</p>	

<p>3/5/02 12:59:36</p>	<p>Continues flat.</p> <p>Total Position: 5000</p>	
<p>3/5/02 13:06:52</p>	<p>Still flat with thin volume. Decide to sell 3000 shares here to partially lock in gains. Note that average cost rises with these shares sold.</p> <p>Trade 6: Sell 3000 shares</p> <p>Total Position: 2000</p>	
<p>3/5/02 13:21:29</p>	<p>Edging slowly downwards.</p> <p>Total Position: 2000</p>	

<p>3/5/02 14:21:28</p>	<p>Rode through another period of loss, and back up to a comfortable level. Considering selling the rest here, but decide to wait.</p> <p>Total Position: 2000</p>	
<p>3/5/02 14:55:20</p>	<p>Fell again – don't like seeing it this low, this close to market close. Will sell on next high.</p> <p>Total Position: 2000</p>	
<p>3/5/02 15:10:37</p>	<p>Sell the rest on a small rally.</p> <p>Trade 7: Sell 2000 shares</p> <p>Total Position: 0</p>	<p>(no snapshot)</p>

<p>3/5/02 15:15:48</p>	<p>Price fell to a new low shortly afterwards. This would have triggered the stops if sale had not already been made.</p> <p>Total Position: 0</p>	
<p>3/5/02 16:00:02</p>	<p>AXP ended the day slightly above the previous close, a deceptively small change for a volatile day.</p> <p>Total Position: 0</p>	

7.2 Trading Results

In one day and 7 trades, the trader made \$1,820. If he could have identified the early-morning peak, he would have made \$3,250 in less than an hour.

The important question is to discern the effects that the view had on trading behavior: Did the view provide a good sense of future price direction? Did the view encourage irrational speculation? Was it riskier because it was more fun? Did it make the trader think more about the position than the price action? Did the view encourage too many trades? Are traders able to learn from this system more quickly than a numerical system?

8 Future Work

As with all research or software, there is much to be done. The following are conceptual, technical, and financial future directions for academic and commercial work.

8.1 User Testing

A clear next step in this research is more user testing. A simple test would be to compare trading behavior and performance between users of numeric vs. graphical systems, using the same historical data set. The same test could in fact be done with the same user to see if he even recognizes that he is working with the same dataset.

8.2 Algorithmic Testing

As has been noted earlier, Technical Trading is a discipline that has derived trading rules from a related visualization, the line plot. By extending these rules or discovering new rules for trading under this new visual paradigm, it would be possible to compare trading by visual guidance versus by more traditional, non-visual rules. This kind of back-testing has been long been done informally within the Technical Trading community. Recently, this subject has started to command formal, academic attention with Lo's analysis of the technical "head and shoulders" and "double-bottom" patterns¹¹¹.

The logic would imply that if visual rules are profitable, then visual tools could be useful. Of course, it would also imply the utility of automated pattern-finding – a subject that will be touched again under "Intelligence Augmentation".

8.3 Visual Behavioral Finance

Most behavioral finance research to date has measured behavior in response to numeric or behavioral stimuli. Visual input is perceived and processed differently than textual or verbal input,¹¹² and this perceptual difference may cause different emotional or behavioral responses.

This is a vital subtlety to the designer of a trading interface, who is trying to avoid negative trader stimuli. The ultimate goal would be to design a trading system that shows all aspects of

¹¹¹Lo, Andrew, et al. "Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation" *Journal of Finance*. 55(2000), 1705-1765.

¹¹² Ware, *Information Visualization: Perception for Design*, 26.

the trading problem in behaviorally correct terms. For example, red may be an inappropriate color to represent loss, and certain shapes may ultimately prove to be more behaviorally correct than others.

These low-level signals may or may not be overpowered by the context of trading, where there are clear consequences, regardless of color or shape. Learned behavior may ultimately bring traders back to perverse, risk-seeking behavior – but the emotional power of these low-level visual signals is poorly understood and merits more study.

8.4 Visualization Enhancements

8.4.1 Broader Context View

The prototype visualization’s context view is not broad enough. A simple solution to this problem would be to add a third view, showing the price and volume action in the context of days or weeks. However, this approach could be unwieldy and cause confusion.

An experimental but potentially more fruitful approach could be to fit the current views with a log or other nonlinear scale for time. In the context view, it would be possible to use one or more “table lenses”¹¹³ to magnify areas of interest.

8.4.2 Integration of Quote Data

This prototype focused on trade data, ignoring quotes. While the TAQ dataset could have been used to build a semi-transparent envelope of bid and ask prices around the price ribbon, this approach was discarded as too cluttered. More interesting is the route taken by GL Genie, with a histogram of quote (or trade) density at each price level. The depth of the book would clearly be of great use to a trader.

This would require an algorithm to create a list of open bids/offers from TAQ’s historical stream, or a live feed of level 2 data for a real-time application. Quote data would also enable the use of derivative statistics, such as bid-ask spread, price pressure, and the like.

¹¹³ Rao, Ramana, and Stuart K. Card. “The Table Lens: Merging Graphical and Symbolic Representations in an Interactive Focus+Context Visualization for Tabular Information”. Card, Stuart, et al., eds. *Readings in Information Visualization*. San Francisco: Morgan Kauffman, 1999. 343-349.

8.4.3 Market and Technical Indicator Integration

As noted above, a clear next step for this prototype is to cleanly add indicators for the market index and technical metrics to the asset focus view. With an eye to keeping the system clean and configurable, it is possible to layer more graphical elements onto the time axis or around the price ribbon.

8.4.4 Sound

Any trader from a live exchange, such as the NYSE, will quickly emphasize the importance of sound. The noise of the crowd and machines serve as a vital environmental cue about general market action.

The NYSE opened the “Blue Room” in 1969, its first major expansion in almost half a decade, with state-of-the-art soundproofing. Soon afterward, it became apparent that traders couldn’t trade effectively in these quiet rooms. The auditory input was so important that the exchange wired the floor sound into the new offices.¹¹⁴

Removed from the floor, an upstairs trader still has options today to hear floor action. Multiple services broadcast sound from the major exchange floors over the web, and companies such as MarketSound¹¹⁵ have attempted to synthesize price and crowd data for specific assets.

Trading applications need to use environmental sound more effectively. Some platforms have programmable event-based noises, but do little to convey the background of emotion in the market. It would be interesting to integrate this visual model with recent advances in 3D sound.

8.4.5 Arbitrage Formation

This research has focused on the display of a single asset, rather than multiple-asset opportunities. Instead, a more complex arbitrage opportunity could be displayed in real time with multiple ribbons on one axis, highlighting any relative mispricing. For example, such an application might simultaneously track a company’s stocks, corporate bonds, and government bonds to monitor the size of a capital structure arbitrage opportunity.

¹¹⁴ Mike Epstein, Interview with author.

¹¹⁵ MarketSound, Inc. <http://www.marketsound.com>

8.4.6 Graphical Order Placement

In this prototype, trades were placed using a traditional textual interface, in order to assure precision and clarity. If the same clarity were possible with a completely graphical interface, then trading could become a completely graphical task.

To start such a trade, the user might click on the asset ribbon, and the system would present him with a movable bar to indicate the desired position size, long or short. Dragging the bar up or down would visually create a cube in the trading space, interactively sized. Limit orders and stops could be drawn on an adjacent target price palette, or even directly on the asset's price axis.

It would even be possible for the trader to draw a more complex pricing/payoff profile, which the system could replicate using a combination of options, bonds and equity. With computer assistance, there is no reason to constrain the trader's instruments to simple equities. Such a system could take a more active role in automating position size and risk management boundaries.

Video games have arguably created the most popular, pervasive, and successful virtual tactical interfaces, so perhaps it would be useful to borrow visual elements from them. The trader's assets and risk could be displayed on a screen overlay. A trader could create a set of favorite, memorized "moves" to be applied in different situations. Like a soldier in the game of Doom[®], the trader could choose their favorite "weapon" and "shoot" as soon as they see an opportunity.

In such a scenario, the trader may not even know about the price, size, or even the name of his positions. Such a system would move the practice of trading even further into a virtual, symbolic discipline, rewarding traders more for their pattern-finding capacity, discipline, and mental speed, rather than any particular mathematical or financial acuity. Of course, this model would severely punish those who do not think before they shoot, just like in a video game or real life. It is an interesting scenario to consider, given the video-game immersion of today's youth, who are tomorrow's traders.

8.4.7 Intelligence Augmentation

It is interesting to consider the place of heuristics and artificial intelligence (AI) in the context of visualization. Discussion of this issue quickly turns to a contentious debate within the AI community, on the role of “Intelligence Augmentation”, or IA. A traditional AI-centric viewpoint would apply machine learning of some variety upon the raw data, to directly find and report patterns. The counter argument can be summarized as “IA>AI”, stating that a symbiotic human-computer team operating together is inherently a more powerful model than pure-computer model. IA considers the computer to be an extension of the user, much like a prosthetic device.¹¹⁶ Much IA discussion centers on wearable computers and software agents, but analysis and tactically oriented visualization clearly fits into the framework as well.

8.5 Related Financial Applications

8.5.1 Trading Presentation

Rather than using the previous visualizations as a tactical or analytical tool, these same representations may also be useful for presentation purposes. Static or animated, these images clearly and directly show traders’ choices in a moving price path.

They may be useful to convey strategies and mental models for trader training, policy review, or performance assessment. Such representations may also prove useful in litigation support, clarifying complicated trader actions to a lay jury or judge. Recent lawsuits about mutual fund “market timing” may find this kind of visualization useful in making a case.

8.5.2 Compliance and Surveillance

Similarly, a large trading firm, exchange, or a regulatory agency could visualize real-time trading, searching for illegal or excessively risky behavior. Notably missing from the TAQ dataset is the identity of the trading counterparties – this is not public information. However, in a compliance, surveillance, or litigation context such information would become available. This would enable a visualization to sketch a trader’s path through a sea of assets, making an overall strategy clear.

¹¹⁶ Maes, Patti. Class notes for Reading Seminar on Intelligence Augmentation. Massachusetts Institute of Technology, Program in Media Arts & Sciences. <http://ia97.www.media.mit.edu/courses/ia97/>

8.5.3 Microstructure Research

An important area of research today is in the area of market microstructure. This type of visualization is a natural way for researchers to see ticks in action. Freed of tactical constraints, a variety of analytical tools can be applied to the ticks, depending on the nature of the research.

8.5.4 Monte Carlo Analysis

As mentioned at the outset, asset returns can be considered to be a vector field, and actual realized prices a set of integral values, a streamline. Monte Carlo analysis typically considers thousands of possible price paths to discover characteristics of the underlying assets. In simple cases, distribution statistics of the simulated final asset price are all that is needed from a Monte Carlo analysis.

However, more complex instruments or analyses may require understanding of the complete time series, rather than just the final price. Volume visualization in particular can be a powerful tool for exploring these results, in combination with the more traditional time series methods used in this project.

9 Conclusions

This research has proposed a functional view of equity trading, and applied an advanced visualization prototype to these tasks. Visualization, even at this prototype level, does add value to the trade timing problem, and particularly to the monitoring problem. It also shows to be useful for the formation, evaluation, and communication of trading policy.

Temporal and equity patterns in volume are visible in the dataset. Erroneous ticks, such as mislabeled late prints, stand out clearly. Price action is clear, with a vivid presence that line charts do not project. The give and take of prices as they are bid up and down are visible in an animated, even dramatic fashion. Trading positions are clearly connected to the price action, and monitored in real time.

The act of equity trading is deeply bound to time and price – as a result, trading visualizations cannot effectively stray far from these two literal dimensions. In this way, information design for visualizing ticks is an exercise in creative constraint, compared to visualization for risk or portfolio visualization.

This is the beginning of an attempt to find a visual language for trading. These designs must be refined, clarified, and tested. The technology must be solidified and accelerated to handle institution-sized portfolios. These are a simple matter of creativity and hard work.

A more interesting aspect of this research, however, is in discovering and managing the ways that this kind of visualization affects human emotions and behavior. These views clearly show the market more vividly than ever before – it will be an subtle and interesting design challenge to channel this emotion into greater trading performance, while minimizing irrationally risky trader behavior.

10 References and Bibliography

- “Grokking the Infoviz”. *The Economist* 367:8329 (6/21/2003).
- Adams, Mark, et al. “Case Study: A Virtual Environment for Genomic Data Visualization”. Proceedings of *IEEE Visualization 2002*.
- Albers, Josef. *Interaction of Color*. New Haven, CT: Yale University Press, 1975.
- Arnheim, Rudolf. *Visual Thinking*. Berkeley, CA: University of California Press, 1969.
- Baum, David, “Assets in Wonderland”. *Byte Magazine*, July 1995.
<http://www.byte.com/art/9507/sec8/art3.htm>
- Bertin, Jacques. *Semiology of Graphics: Diagrams, Networks, Maps*. Madison, Wis.: University of Wisconsin Press, 1983.
- Card, Stuart, et al., eds. *Readings in Information Visualization: Using Vision to Think*. San Francisco: Morgan Kaufman, 1999.
- Chevreau, Jonathan. “Toronto Firm Pioneers ‘Financial Animation’” *The Financial Post*, January 4, 1994.
- Cleveland, William S.. *The Elements of Graphing Data*. Monterey, CA: Wadsworth Advanced Books and Software, 1985.
- Cleveland, William S.. *Visualizing Data*. Murray Hill, NJ: AT&T Bell Laboratories, 1993.
- Cooper, Alan. *About Face 2.0*. Indianapolis: Wiley, 2003.
- Course Notes from “Design of Interactive Multimodal Media Systems”. Proceedings of *ACM Siggraph 2003*.
- Dwyer, T and P. Eades. “Visualising a Fund Manager Flow Graph with Columns and Worms”. Proceedings of *IEEE IV 2002*.
- Edwards, Robert D., et al.. *Technical Analysis of Stock Trends*. Boca Raton, Fla.: St Lucie Press, 8th Edition, 2000.
- Fisher, Adam. “Playing the Market in 3-D”. *Wired Magazine*, 8.06 (June 2000).
- Foley, J., and B. Ribarsky. “Next-Generation Data Visualization Tools”. *Scientific Visualization: Advances and Challenges*. Ed. L. Rosenblum, et al. London: Academic Press, 1994.

- Fry, Benjamin Jotham. *Organic Information Design*, Masters Thesis, Massachusetts Institute of Technology, Program in Media Arts & Sciences, 2000.
- Grant, Josh, et al. “Case Study: Visualizing Ocean Flow Vertical Motions using Lagrangian-Eulerian Time Surfaces”. Proceedings of *IEEE Visualization 2002*.
- Horn, Robert E. “Information Design: Emergence of a New Profession”. *Information Design*. Ed. Robert Jacobson. Cambridge: MIT Press, 1999.
- Inf@Vis! “Impactopia”. http://www.infovis.net/E-zine/2002/num_99.htm
- Inf@Vis! “Visualizing Finance”. http://www.infovis.net/E-zine/2002/num_77.htm
- Lo, Andrew, et al. “Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation” *Journal of Finance*. 55(2000), 1705-1765.
- Maes, Patti. Class notes for Reading Seminar on Intelligence Augmentation. Massachusetts Institute of Technology, Program in Media Arts & Sciences.
<http://ia97.www.media.mit.edu/courses/ia97/>
- Maor, Eli. *E: The Story Of A Number*. Princeton: Princeton University Press, 1994.
- Margolin, Victor, and Richard Buchanan. *The Idea of Design*. Cambridge, MA: MIT Press, 2000.
- McCloud, Scott. *Understanding Comics: The Invisible Art*. Northampton, MA: Kitchen Sink Press, 1993.
- Mitchell, William J.. *City of Bits: Space, Place, and the Infobahn*. Cambridge, MA: MIT Press, 2000.
- New York Stock Exchange, Inc. *TAQ User’s Guide*. Version 3.31.
- Norman, Donald A.. *The Design of Everyday Things*. New York: Doubleday, 1990.
- Parrish, Edward. “StockVis: An Internet-Based System for Visualizing Stock Market Data”. Master’s Thesis, UC Santa Cruz, Department of Computer Science, 2000.
- Petroff, Gregory. “The Virtual Stock Exchange: Lessons Learned from the NYSE 3DTF”. Proceedings from *CODATA Workshop*, 10-11 July 2002.
- Rao, Ramana, and Stuart K. Card. “The Table Lens: Merging Graphical and Symbolic Representations in an Interactive Focus+Context Visualization for Tabular Information”. Card, Stuart, et al., eds. *Readings in Information Visualization*. San Francisco: Morgan Kauffman, 1999.
- Rensink, Ronald A.. “Internal vs. External Information in Visual Perception”. Proceedings of *ACM Symposium on Smart Graphics*, June 11-13, 2002.

- Schneiderman, Ben, and Martin Wattenberg. "Ordered Treemap Layouts". Proceedings from *IEEE Visualization 2001*.
- Simons, Daniel J, and Christopher J Chabris. "Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events". *Perception*, 28 (1999), 1059-1074.
- Simunic, Kresimir, "Visualization of Stock Market Charts". Proceedings from *The 11th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision 2003*.
- Spence, Robert. *Information Visualization*. New York: Addison Wesley, 2001.
- Strausfeld, Lisa. *Embodying Virtual Space to Enhance the Understanding of Information*. Masters Thesis, Massachusetts Institute of Technology, Program in Media Arts & Sciences, 1995.
- Tufte, Edward R.. *Envisioning Information*. Cheshire, CT: Graphics Press, 1990.
- Tufte, Edward R.. *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press, 1983.
- Tufte, Edward R.. *Visual Explanations*. Cheshire, CT: Graphics Press, 1997.
- Tukey, John W.. *Exploratory Data Analysis*. Reading, Mass.: Addison Welsley, 1977.
- Ware, Colin. *Information Visualization: Perception for Design*. San Diego: Academic Press, 2000.
- Wattenberg, M. "Visualizing the Stock Market". Proceedings of *ACM CHI 99*, Extended Abstracts. 188-189.
- Wright, William. "Seeing Investment Strategy". *Global Investment Management*, March, 1998.
- Wright, William. and J. O'Neil. "Visualization Value for Northern Trust Investors". *Data Mining for Financial Services Conference*, IBC, Chicago, 2000.
- Zhukov, Leonid, and Alan Barr. "Oriented Tensor Reconstruction: Tracing Neural Pathways from Diffusion Tensor MRI". Proceedings of *IEEE Visualization 2002*.

