Algorithms at the Edge

David Leinweber
dleinweber@post.harvard.edu
626-398-9415

Draft article for Institutional Investor Alpha Magazine

Revised Jan 3, 2007

Spy vs. Spy

This article looks inside the meanest, smartest algorithms on the street, and the even meaner, smarter algorithms that will replace them. The cold war is history, but there's an arms race underway in algo trading. It started in the eighties, and shows no signs of slowing down.

Readers of this fine journal are also no doubt familiar with MAD Magazine's Spy vs. Spy cartoons. Running continuously since 1961, in the "Joke and Dagger" department, we see the spies, identical except for the color of their coats and hats, engage in an endless series of ever more elaborate schemes to gain an advantage.

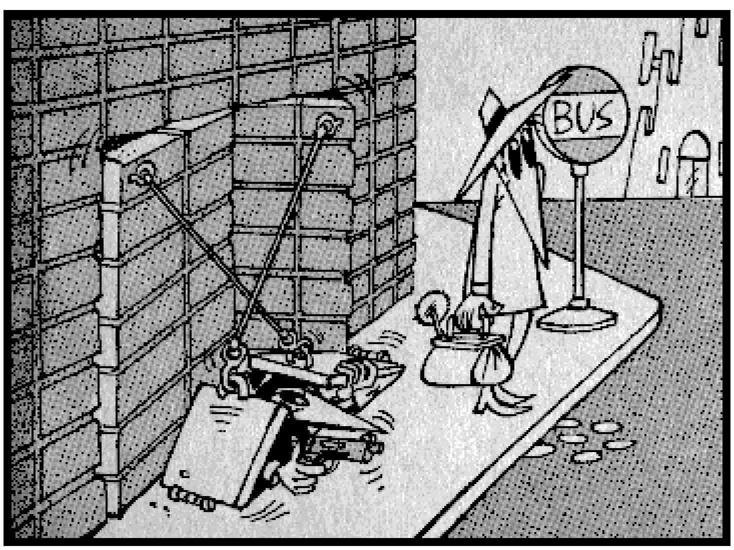

MAD's spies use an assortment of daggers, explosives, poisons, military hardware, and Rube Goldberg schemes in their war. The battle for supremacy in algorithmic execution uses an assortment of mathematics, programming, communications, computing hardware, and Rube Goldberg schemes.

It 's worthwhile to understand the simpler beginnings of electronic trading to better appreciate today's elaborate systems, and the more elaborate systems that will replace them. When market systems involved chalkboards, shouting, hand signals, and large paper limit order books, there was no possibility of using a computer to execute trades.

This changes in 1976, when the NYSE introduces the DOT (Designated Order Turnaround) systems, the first electronic executions system. It was designed to free specialists and traders from the nuisance of 100 share market orders. The NASDAQ market, started in 1971, used computers to display prices, but relied on telephones for actual transactions until 1983 with the introduction of the Computer Assisted Execution System, and the Small Order Execution System in 1984.

Simultaneous improvements in market data dissemination allowed computers to be used to access timely quote and trade streams. The specialists at the NYSE had a major technology upgrade in 1980, when the specialist posts themselves, which had not changed since the 1920s, were made electronic for the first time, dramatically reducing the latencies in trading. A recent study[1] of trading before and after that important upgrade found significant improvements in market quality, liquidity, and transaction costs.

Early electronic execution channels started out as a service for only the smallest orders. But the allowed sizes grew fast. Support for limit orders were added. DOT became SuperDOT. It and the NASDAQ systems accommodated larger and larger orders. Orders exceeding the size limits for automation could be routed to specialists and market makers, the others would be executed automatically, by the market center's computers.

This was algorithmic trading without algorithms, an early form of direct market access (DMA). The initial user interfaces were for one stock at a time. They resembled electronic versions of paper buy and sell slips. This became tedious, and soon execution capabilities for a list of names followed. Everyone was happy to be able to produce and screen these lists using their fancy Lotus 123 spreadsheets, which totaled everything up nicely to avoid costly errors.

We are only a step away from algorithmic trading. Programmers at the order origination end grew more capable and confident in their abilities to generate and monitor an ever larger number of small orders. Aha! Algo trading has snuck up on us.

Algos for Alpha

Early adopters of these ideas were not looking to minimize market impact, or match VWAP. They were looking to make a boatload of cash. Nunzio Tataglia, a Jesuit educated Ph.D. physicist with the vocabulary of a sailor started an automated trading group at Morgan Stanley in the mid 80s. He hired people like young Columbia computer science professor David Shaw. At first, a few papers about hooking Unix systems to market systems emerged. Then the former academics realized there was no alpha in publications. Shaw, of course, went on to found D.E. Shaw, one of the world's largest and most consistently successful quantitative hedge funds.

---

[1] The Impact of Trading Technology: Evidence from the 1980 Post Upgrades, Dec. 2006 working paper, D. Easley (Cornell) L. Hendershott (Berkeley) and T. Ramadori (Oxford) http://faculty.haas.berkeley.edu/hender/TradingTech.pdf

The early alpha seeking algo players were the first combatants in the Spy vs Spy algorithm wars. Pairs trading, a popular activity at the time, relied on statistical models. Finding stronger short term correlations than the next guy had big rewards. Escalation beyond pairs, to groups of related securities was inevitable. Many exploited parallel developments in futures markets to open the door to electronic index arbitrage trading.

Automated market making was a valuable early algorithm. In quiet normal markets buying low and selling high across the spread made money. Real market makers have obligations to maintain a two sided quote for their stocks, even in turbulent markets, which is often expensive. Electronic systems, without the obligations of market makers, were not only much faster at moving quotes, they could choose when not to make markets in any stock. David Whitcomb, founder of Automated Trading Desk[2], one of the pioneers in this, describes his firm's activity as "playing NASDAQ like a piano".

Faster data feeds and faster computation means you get to run ahead of the other kids in line. This was a time when the lags between one desktop datafeed, like Quotron, and another might be as much as 15 minutes in a busy market. All introduced significant delays relative to what was actually happening. Slow computers, sending information to slow humans over slow lines were easy marks for early algo warriors willing to spend to buy faster machinery, and smart enough to code the programs to use it. This aspect of the arms race continues unabated today.

Algos for the buy-side

It didn't take long to notice that these new electronic trading techniques (they weren't called algorithms then) had something to offer to the buy-side. Financial journals for academics and practitioners offered a stream of opinion, theory and analysis of transaction costs. Firms like Wayne Wagner's Plexus Group made persuasive well-supported arguments about the importance of transaction costs. Pension plan sponsors, sitting at the top of the financial food chain, were convinced in large numbers.

Index managers did not have to be convinced. With no alpha considerations in the picture at all, they observed that it was possible to run a lousy index fund, or a particularly good one. The difference was the cost of implementation, the cost of trading. Passive managers, on their way to becoming the trillion dollar behemoths of today, are high value clients to brokers

In addition to giving their largest customers what they want, brokers had an additional incentive to adopt electronic trading. The end of fixed equity commissions spawned new competitive pressures. Electronic trading had the potential to cut costs dramatically, while improving quality of service.

The largest firms developed their own electronic order entry systems. Others bought from niche vendors. One of these was Davidge Data, headquartered in a loft on White Street, in the meat district, a short walk from the financial center. Nick Davidge had many clients to support, and used bicycles to dispatch service people, mostly himself. In the course of this activity, he made the important discovery that girls's bikes on New York streets have a half-life ten times longer than boy's bikes.

[2] They have an appropriately snazzy website: www.atdesk.com

From Order Pad to Algos

The first direct access tools from the sell-side were single stock electronic order pads, followed shortly by lists. The next step is breaking those orders down into pieces small enough to execute electronically, and spreading them out over the course of the trading day. Innovative systems like ITG's Quantex allowed traders without large software staffs to use and define analytics and rules to control electronic trading[3]. This looks like what we consider to be algorithmic trading today.

The big news in algorithmic trading in the late 80s was that you could do it at all. The first algo strategies were based on simple rules, like "send this order out in ten equal waves, spaced equally from open to close". But these are predictable and easy to game by manipulating the price on a thin name with a limit order placed just before the arrival of the next wave, getting bagged in classic Spy vs spy style. There was little or no mathematical underpinning, just rules of thumb and educated guesses.

A Scientific Approach: Mathematics, Behavior, and Discovery

The obvious shortcomings of these simple strategies have motivated several generations of mathematically based algorithms, using increasing levels of mathematical and econometric sophistication to include models of market impact, risk, order books, and the actions of other traders. The idea of an efficient frontier of trade path strategies and the use of optimization establish a firm foundation.

Markets have become even more fragmented and complex, with less information conveyed by the BBO and the book, creating a need to exploit new order types and to access "dark liquidity". This has given rise to a generation of "behavioral" algorithms, probing for liquidity, driven by procedural logic and stimulus-response principles and as well as mathematical models. Note that the word "behavioral" here is not referring to behavioral finance, it used in the robotics and computer science sense of designing programs by describing what you want them to do, but not exactly how to do it.

Probe, Learn, Adapt, Cooperate

Algorithms need to probe, learn, and adapt. They need to make effective use of analytic tools, and how recognize their limitations. Algos at the edge seek to exploit information beyond the traditional data from feeds and fills, including news, pre-news, and other forms of market color found on the web.

There has been an explosion of progress in tools for processing text. Think, for example, of Google . This is another instance of a technology that used to require a pack of nerds becoming accessible commercially. Reuters and Dow Jones unveiled new offerings to facilitate this in 2006.

When it comes to millisecond scale "cancel and replace" decisions, algorithms rule. No human can react as fast. The combination of quantitative methods and "artificial

[3] A Little AI Goes a Long Way on Wall Street, Leinweber & Beinart, JPM 96

intelligence" methods is increasingly effective. But how can human traders work with algorithms, using ***"intelligence amplification***" to form a partnership that enhances the skill of both? Finding the proper mix of human and machine skills is a challenge for traders.

Gary Kasparov (the world chess champ who lost to Deep Blue in 1997) suggested that chess tournaments be open to human-machine teams. You can see at least part of Kasparov's job in that situation is to keep an eye on the machine's decisions, just in case it misses some of his insights. Carrying this analogy to trading a little further, imagine if the game was not tournament chess, which allows up to seven hours for the game, but blitz chess, which allows a total of three minutes. Given Moore's law and all that, it's not long before the computer that beat Kasparov with seven hours could beat him with three minutes. Many facets of trading are more like blitz than tournament play,

Job Insecurity for Traders

There no shortage of future paycheck anxiety among traders. Numbers have been dropping. Big time industry mavens like IBM publish a report called "The Trader is Dead, Long Live the Trader"[4] Finextra is an excellent free fount of global market technology information. In May, they ran a headline "City (of London) trading jobs to fall by 90% as banks take up algorithmic technology" [5] Even the Economist magazine, in a story headlined "The march of the robo-traders" observes "programs that buy and sell shares are becoming ever more sophisticated. Might they replace human traders?" [6]

This kind of press makes for anxious career planning. The traders who exit the trade will be replaced by algorithms, the traders who stay will be the ones who play well with machines

Understanding algorithms is a survival skill for traders. Someone will be making the decisions that the algorithms can't make (yet). The history in this essay so far skipped over some important details. Here is a Cliff's Notes for AlgoTrading 101, annotated along the way with features from algorithms at the edge.

You are what you eat.

Algorithms feed on market data, and their sophistication has grown with the scope, timeliness and accuracy of that information. Feeding market data to computers was a black art.

Algorithms have ***sensors*** and ***effectors***, analogous to the "eyes and wheels/arms" of real robots. In between the sensors and effectors, there is a computer program that provides ***control.***

The sensors are the feed of market transparency information, quotes, trades, order books, indications of interest. The effectors are order placement, cancel and modify interfaces. They result in an additional sensor stream of execution information. Control comes from a program based on a combination of market models, rules, and procedures.

[4] IBM Business Consulting Services, Trader is dead, long live the trader. 2006
[5] http://www.finextra.com/fullstory.asp?id=15955
[6] http://www.economist.com/displaystory.cfm?story_id=E1_QPGRNTQ

A central theme in the Algorithm vs Algorithm battle is to get an edge in time, scope and accuracy on market data. Anyone using more than one data service notices lags from one to another, and they all lag the event. There is disintermediation going on in these businesses as well. Companies like Wombat[7] will sell you the docking adapter to sidle right up to the SIAC mother ship[8], just like Reuters and Bloomberg and the rest. Being able to do this disintermediates the data vendors, at least for the raw material.

In the algo wars, as in real wars, it's a good idea to control your communications, avoiding those slow satellite links. Communication satellites are in GEO (Geosynchronous Equatorial) orbits, 22,240 miles over the equator. Light travels at 186,000 miles/second, so a satellite hop takes at least a quarter of a second. Hardly noticeable while talking, but that 250 milliseconds is long enough for a crowd to be ahead of you in US equity and derivative markets.

So many markets, so little time.

You can also rent a parking space for your execution computer right next to the market center computers, eliminating communication latency. This service is now offered by the NYSE, NASDAQ, London, Euronext, TSE, Globex, and a growing list of others <verify>.

This co-location at market centers can do wonders for latencies in execution. As the algo wars proceed, we will see tools that combine fast broad market access with access to proprietary execution channels. Brokers willing to commit capital will be able to offer zero latency executions. Zero is low enough for lots of fast trading strategies that are subject to the vagaries of execution. Watch out, here comes a mob of new hedge funds.

Algos at the edge see a thousand points of light, each with its own Alternative Trading System (ATS), its own clientele (brokers allowed or buyside only), and character (some have large size over days, some have cyclicality). The nature of many of these systems is that size is hidden, and participation may require being in the system for a period of time. This can create a risk of over executing unless the most conservative rules are followed. Larger firms, willing to risk some capital by incurring the risk of over buying (or selling) will be able to allow their clients to make use of more aggressive trading tactics.

In areas like managing and monitoring multiple variegated execution systems, algos on the edge combine analytic tools with expert rules and procedures to profit from the complexity.

Patient trading, Transitions, Trade path risk control

[7] http://www.wombatfs.com Wombat Financial Software is a big arms dealer in the algo wars. The Adnan Khashoggi of low latency connections.

[8] The Securities Industry Automation Corporation, is the place where the consolidated tape gets consolidated. Once, the mother of all market data. It has a long history as the market data arm of the New York exchanges, back to the New York Quotation Company, formed in 1889. The current firm was formed in 1972. NYSE Group acquired the part it didn't own in 2006.

Algos at the edge have a unified way to access a mix of asset classes and derivatives. This opens a door to automated patient execution of large transitions. Nearly all current algo trades occur over the course of a single day. There's no fundamental reason for this. Without a one day rule, future algos will serve institutions using patient transition trading to make sizable adjustments in their portfolios.

Risk is the reason algo trades are nearly all completed in one day now. Full service brokers will able to offer customized short term derivatives for controlling risk exposures along the path of a longer trade. Algos at the edge can act like skilled transition traders.

The Naïve Strategy

The well wired trader has spared no effort or expense in obtaining the finest kind in data and market access of all flavors. What to do with it? Do math.

The earliest algorithms used the KISS strategy of splitting up orders in N parts, every 1/N of a trading day, eg, an order for ten thousand shares would be sent out as ten orders for one thousand shares, at ten times spaced equally over the trading day. This sort of signaling made it easy for traders on the other side to spot these algorithms, and pick them off. Later, smarter authors, who we will dicsuss soon, called this "The Naïve Strategy", which is a kind assessment.

So the next round of Algorithm vs Algorithm/People was to get the algorithms to be less naïve, to hide, by randomizing times and sizes. But randomization made some stupid decisions – placing small orders at the open and close, not reflecting urgency or tolerance for risk, missing transient opportunities in liquidity.

Mix in market impact

Dimitris Bertsimas and Andew Lo[9] at MIT co-authored one of the first academic papers on more sophisiticated approaches to algorithmic trading, "Optimal Control of Execution Costs", in 1998. It may be a look deeper than Algos 101, but motivated readers with some math will get an excellent compact description of their work in the abstract below:

> *"We derive dynamic optimal trading strategies that minimize the expected cost of trading a large block of equity over a fixed time horizon. Specifically, given a fixed block S of shares to be executed within a fixed finite number of periods T, and given a price-impact function that yields the execution price of an individual trade as a function of the shares traded and market conditions, we obtain the optimal sequence of trades as a function of market conditions---closed-form expressions in some cases---that minimizes the expected cost of executing S within*

---

[9] Dimitris Bertsimas and Andrew W. Lo , "OPTIMAL CONTROL OF EXECUTION COSTS" *Journal of Financial Markets* 1(1998), p. 1-50. http://web.mit.edu/alo/www/Papers/bertlo98.html

T periods. Our analysis is extended to the portfolio case in which price impact across stocks can have an important effect on the total cost of trading a portfolio.

They start with a mathematical analysis of the merits of mindless naïve strategies, by asking what sort of environment it would take for them to actually be 'optimal' which turns out to an unrealistically simple world. They then model a more realistic world, where the trading strategy incorporates ideas of market impact, and an "information variable", and how optimal trading strategies depend on it. One of their illustrations is below, it shows the number of shares traded in each period on the top half, and the information variable for that period below.

<replace with II drawn version – highlight or color high information trades – bars 2 3 7 13-15 – show in both upper and lower bars. I'll write a caption>

18 D. Bertsimas, A.W. Lo/Journal of Financial Markets 1 (1998) 1–50

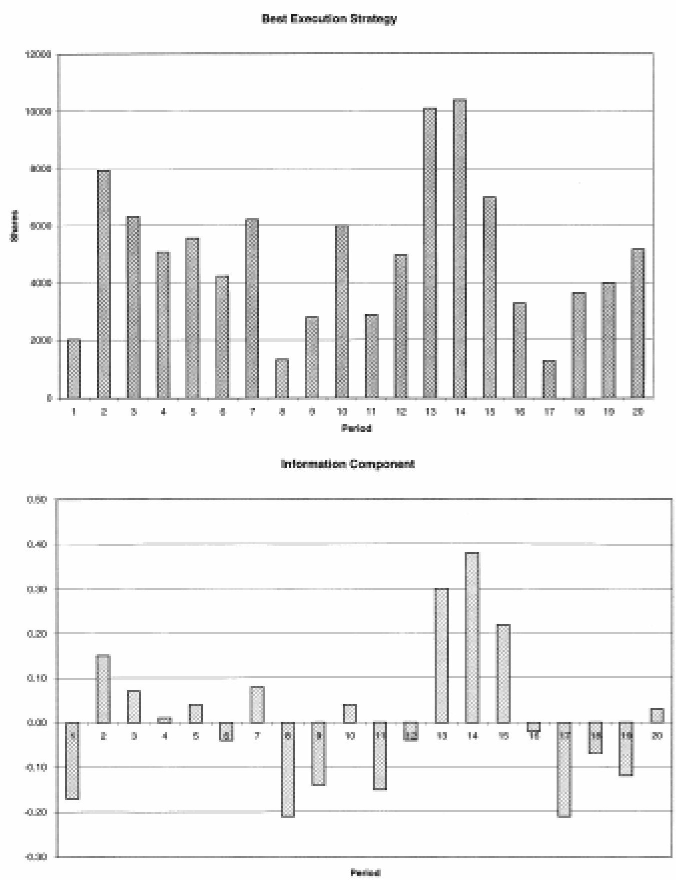

Fig. 2. Best-execution strategy and information-based component for Table 2.

Notice that the trading is strongly driven by the information component. This information variable is quite abstract, exactly how this is determined could include anything from micro-level empirical analysis to rumors heard on the bus.

Mix in risk

Modeling market impact and information was a significant advance. But there is room for improvement. The next step in mathematical magic was to incorporate the idea risk aversion the difference between passive and alpha-seeking trades, which traders are paid to handle very differently.

Robert Almgren and Neil Chriss (now with Bank of America Securities and SAC,respectively) made the next major step in their 2000 paper "Optimal Execution of Portfolio Transactions"[10]. It explicitly included the risk aversion of traders, and introduced the idea of liquidity-adjusted value at risk as a metric for trading strategies. OK, let's call this Algos 201, but again, the authors do a fine job explaining in more details for the mathematically inclined.

> *"We consider the execution of portfolio transactions with the aim of minimizing a combination of volatility risk and transaction costs arising from permanent and temporary market impact. For a simple linear cost model, we explicitly construct the **efficient frontier** in the space of time-dependent liquidation strategies, which have minimum expected cost for a given level of uncertainty. We may then select optimal strategies either by minimizing a quadratic utility function, or minimizing Value at Risk. The latter choice leads to the concept of Liquidity-adjausted Value at Risk, or L-VaR, that explicitly considers the tradeoff between volatility risk and liquidation costs."*

The illustration below is worth many words in understanding this trading model. The intuitively appealing implications of their results for traders of different motivations are seen in the illustration below. The risk averse patient trader C takes his time doing the trade. The more motivated trader A moves faster, willing to take her chances on a larger market impact to get it done.

<replace with II version. Use colors for A B C, label A-risk taker, trades fast, B- risk neutral, C- risk averse, trades slow>

December 2000 *Almgren/Chriss: Optimal Execution* 18

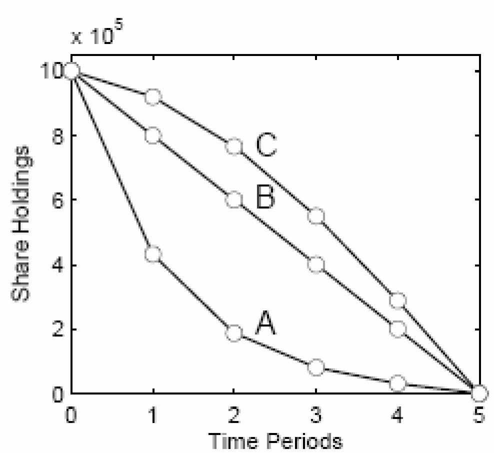

Figure 2: *Optimal trajectories.* The trajectories corresponding to the points shown in Figure 1. (A) $\lambda = 2 \times 10^{-6}$, (B) $\lambda = 0$, (C) $\lambda = -2 \times 10^{-7}$.

[10] http://www.courant.nyu.edu/~almgren/papers/optliq.pdf

Mathematical models of markets can become very elaborate. Game theoretic approaches to other market participants, human and machine, in the spirit of the "Beautiful Mind" ideas of John Nash, bring another level of insight.

<maybe more on this>

Known Unknowns and Unknown Unknowns

Algren and Chriss close with an important point about the limitations of all model driven strategies. As part of the Algos 201 track, here is what they say about connecting algorithms to real world events:

> *"Finally, we note that any optimal execution strategy is vulnerable to unanticipated events. If such an event occurs during the course of trading and causes a material shift in the parameters of the price dynamics, then indeed a shift in the optimal trading strategy must also occur. However, if one makes the simplifying assumption that all events are either 'scheduled' or 'unanticipated', then one concludes that optimal execution is always a game of static trading punctuated by shifts in trading strategy that adapt to material changes in price dynamics. If the shifts are caused by events that are known ahead of time, then optimal execution benefits from precise knowledge of the possible outcomes of the event. If not, then the best approach is to be actively 'watching' the market for such changes, and react swiftly should they occur. One approximate way to include such completely unexpected uncertainty into our model is to artificially raise the value of the volatiltity parameter."*

This is a version of the wisdom of two time Secretary of Defense Donald Rumsfeld that there are "known unknowns and unknown unknowns".

Examples of known unknowns are

- Scheduled announcements that affect particular stocks, like earnings or conference call
- Announcements that affect groups of stocks, like housing starts
- Announcements that affect broad markets, like macroeconomic data and interest rates

There are many sources of this type of information. Thompson's StreetEvents[11] has a wide selection. Econoday[12] is a calendar book for information on the later two types of event above. It was found in every trading room since the fifties is also now a web service. Some algorithms use this information, some don't. Guess which ones are better.

[11] https://www010.streetevents.com/loginAll.asp

[12] http://www.econoday.com/

Unknown unknowns include news, discussion, rumor, market color, agency actions, research results. Computers are pretty good at finding this kind of thing. Too good, there is a vast amount of chaff, to be kind.

This is another front in the Algo vs Algo war. Figuring out what to do with news, and pre-news events. Determining when an ‘unknown unknown” that will change the trading strategy is a place where humans with machines have an edge on either one working alone.

What’s all this math?
How does it relate to the algos we know? VWAP and the boys?

The microstructure tactics based on these trading models are often deployed in VWAP and TWAP. These anticipate volume, and try to participate throughout the day (or time period), optimizing to that target.

Models aren’t Markets

We’ve convinced ourselves that even the most elegant models are abstractions of real markets. The real thing is a rapidly changing mélange of market fragments, continuous and call markets, ECNs and innovative matching systems, indications, dark liquidity and an explosion of activity by algorithms.

Limit order books just aren’t what they used to be. In days of eighths, and orders being modified by people, the inside quote conveyed actionable information. Now the average life of a limit order is measured in milliseconds, so the quote is a fast moving target. With decimalization, the old total size at the inside spread was distributed over 6 or more price levels, so the BBO conveyed much less information. ECNs and the exchanges exposed more of the book. Just when you could see the book, the algo battlefield shifts with dark liquidity, hidden pre-programmed orders to execute when others are filled, anonymous indication and matching systems. These take the liquidity that, back in the day, would be in the light (visible in the open book), and conceal it in the dark of less transparent markets and real-time programs.

Here we need to look at the control part of algorithms. With models, we can write formulas and make plans to tell us what to do. Edge algorithms can use models as a basis for action, but they have a wide and current vocabulary of rule and process tools to execute on all market segments.

As long as markets change, people will need to monitor and adjust algo and electronic strategies. Markets change so rapidly that humans will be important here.

Robots, RoboTraders, and Traders

Often, the best model of something is the thing itself. This has been a key concept in robotics. Building a robot that explores a model of Mars is different than building one that explores Mars.

An ever growing collection of impressive robots have done well in complex dynamic environments. Doing bomb disposal, surgery, open terrain driving, and floor washing well, with varying degrees of autonomy. [13]

Looking at how these robots think is looking at the future of algorithms. Looking at how humans and physical robots interact is a look at how humans and trading robots will coexist.

So what are these robots thinking?

There are always multiple approaches to robotic tasks, the computer science literature about structuring and coordinating these approaches, is called multi-agent systems[14]. Agents are a new approach to controlling complex systems. The agents are programs that cooperate, coordinate, and negotiate with each other. The list of key features of multiagent systems reads like a description of key features of algorithmic trading:

- **Embedded in the real world.** If a program's environment is always the same, and it is of no great consequence if it fails, then that program can just run blindly. But the real world in general and markets in particular are not static. Things change, information is incomplete. Everything is dynamic.

  A reactive agent maintains an ongoing interaction with the markets, and with information relevant to actions in the markets, and responds to events rapidly enough for the response to be useful.

- **Partial, imperfect models.** Models of financial market behavior never have the precision of engineering models. They are statistical in nature, and have wide error bands. This is particularly true for equities, less so for securities that have structural or arbitrage relations. Financial models are partial, they don't capture every aspect of market participant's motivations, and they are always imperfect.

- **Varied outcomes likely.** Simple games like tic-tac-toe can be modeled exactly. One action always leads to another. This is clearly not the case in trading.

- **Performance feedback & reinforcement**. Performance measurement is natural for trading agents. For alpha seeking algos, risk adjusted reward metrics like the Sharpe ratio are natural. For execution algos, implementation shortfall or other transaction cost measurements are appropriate. For simple systems, giving credit to reinforce successful behavior is basically keeping score. When systems become more complex, e.g., those involving cooperative decisions by humans and machines, the credit allocation procedures become more elaborate.

---

[13] An excellent series of photo articles on robotics – Forbes, Sep 04 '06 Elizabeth Corcoran

[14] An Introduction to MultiAgent Systems, Michael Woolridge, Wiley, 2005 http://www.csc.liv.ac.uk/~mjw/pubs/imas/

-
- **Layered behaviors.** This is the idea that agents should have default behaviors that complete their tasks, and avoid errors. Basic behavior is at the lower layers, more sophisticated behavior above. For example, a trading algorithm that included consideration of news to exploit short term reversals could fall back to a participation rule when there was no news.

Some of these agents will be programs, some will be people. We can call these people "the employed traders of 2015".

Multiple Markets, Multiple Asset Classes, Multiple Agents

Electronic equity executions are rapidly becoming commodities, and disintermediation allows the buy-side to bypass the sell-side to access markets directly. Full service investment firms, bulge brackets and wannabes, will be driven to develop full service electronic interfaces to accommodate complex, multi-asset class, leveraged trades.

Markets 2015

Two recent reports on markets in 2015 are remarkably similar. One is from BearingPoint, "Shifting from Defense to Offense: A Model for the 21st Century Capital Markets Firm"[15] They describe a shift from a product paradigm to a risk paradigm. Prognostications here include: an increase in the complexity of derivative and structured products driven by the demands of alpha seeking strategies; some products will require willingness to commit capital in innovative ways; and increased trading interest in risk classes, over individual securities.

Future markets study contestant number two is "Profiting today by positioning for tomorrow: A field guide to the financial markets of 2015", from IBM Global Business Services. [16]

There is a remarkable similarity in views of the two Markets 2015 studies on the shift to a risk centric view of trading. The IBM report opines, "As the industry matures, many traditional activities wil come under increasing pressure and new value engines will emerge. Activities under pressure are unnecessary bundles and transaction businesses. Value engines will be risk assumption and risk mitigation.

The IBM study included a survey of 271 financial firms, buy- and sell side and processors. When asked which technological innovations would have the greatest impact on their ability to achieve their ten year goals, five of the top six responses relate to future electronic trading: risk measurement systems, client analytics, multi-asset class platforms, connectivity, and software to enable smarter trading

---

[15] http://www.bearingpoint.com , peter.horowitz@bearingpoint.com.
[16] Sarah Diamond
http://www.sia.com/ops2006/pdf/SarahDiamond.pdf

Trade Path Risk Control Algorithms

How will these trends be reflected in algorithmic trading systems? If the shifts described occur as predicted, we can anticipate that clients will want to control trade path risk, and sell-side firms will want to accommodate them.

Controlling risk exposures during the course of a complex trade using custom derivatives plays to one of the strengths (and profit generators) for large firms.

Agents will have to be able to price these derivatives, using quantitative measures, and the firms risk profile.

Playing well with robots and algorithms

People will have to find their place in these multi-asset risk-mitigated fragmented algorithm-infested markets of 2015 and beyond. It's interesting to look at how people are working with other algorithms in real robots.

Some of the real robots work largely on their own. They have stimulus/response rules and internal representations of their tasks. There are two million iRobot Roomba vacuums sucking up dirt solo.

The Mars Rovers, Spirit and Opportunity, the Energizer Bunnies of Space, are given a lot of control over their actions. They have an autonomous vision based mobility system.[17] , illustrated below. Humans set the goal, the rover takes care of the rest.

Other robots are on an extremely short leash. The iRobot Packbot Explosive Ordanance Disposal robot[18] comes with a substantial remote control. These are impressive items to see, and worth a visit to the website. This is made by the same company that makes the Roomba vacuum cleaners.

Robot surgeons, like the Da Vinci Surgical Robots[19] , are on the shortest possible leash Every move is controlled by a human surgeon. This is relly a teleoperated system, with very little autonomy, other than safety stops. The medical version of a single stock electronic order pad. Wait a few years and we'll be seeing algorithmic surgeons.

Just like being there

These robots and the people they work with have a great advantage in being able to see what they are doing using cameras, Well armored ones for PackBot, and little tiny ones in

[17] http://marsrovers.nasa.gov/technology/is_autonomous_mobility-02.html
If you are looking for quality internet entertainment, check the surprising video there. Those guys at JPL are such a bunch of cut-ups.

[18] http://www.irobot.com/sp.cfm?pageid=138
[19] http://www.ohioheartsurgery.com/robot.htm

tubes for Dr. DaVinci. Force feedback and texture sensors let you feel what it's like to be there.

In the real world of bombs and gall bladders, looking around is great way to work with robots. But where are the cameras for trading?

The Big Picture for markets

How can traders get the equivalent of a robot camera view into the markets? The employed trader of the future will have learned to amplify his intelligence by working shrewdly with computers. How will this look?

Ideas about how to do this have evolved from simple to sublime, as we saw with algorithms. Human access to market data has moved from ticker tapes to green screens to windowed graphics. Progress, no doubt, but not of the scale seen in other fields, like meteorology and molecular biology, where visual tools have truly created new insights. The reason for this is that, unlike weather and molecules, markets don't have a natural physical representation to use as a model for the visual representation. They are abstract entities.

Research labs in academia and innovative firms have been working on better visualizations for abstractions like markets for a long time.

Here's part of NYSE's Marketrac display, mapped to the physical floor. This rarely if ever used for trading,. It is more of a WWW style world's fair display.

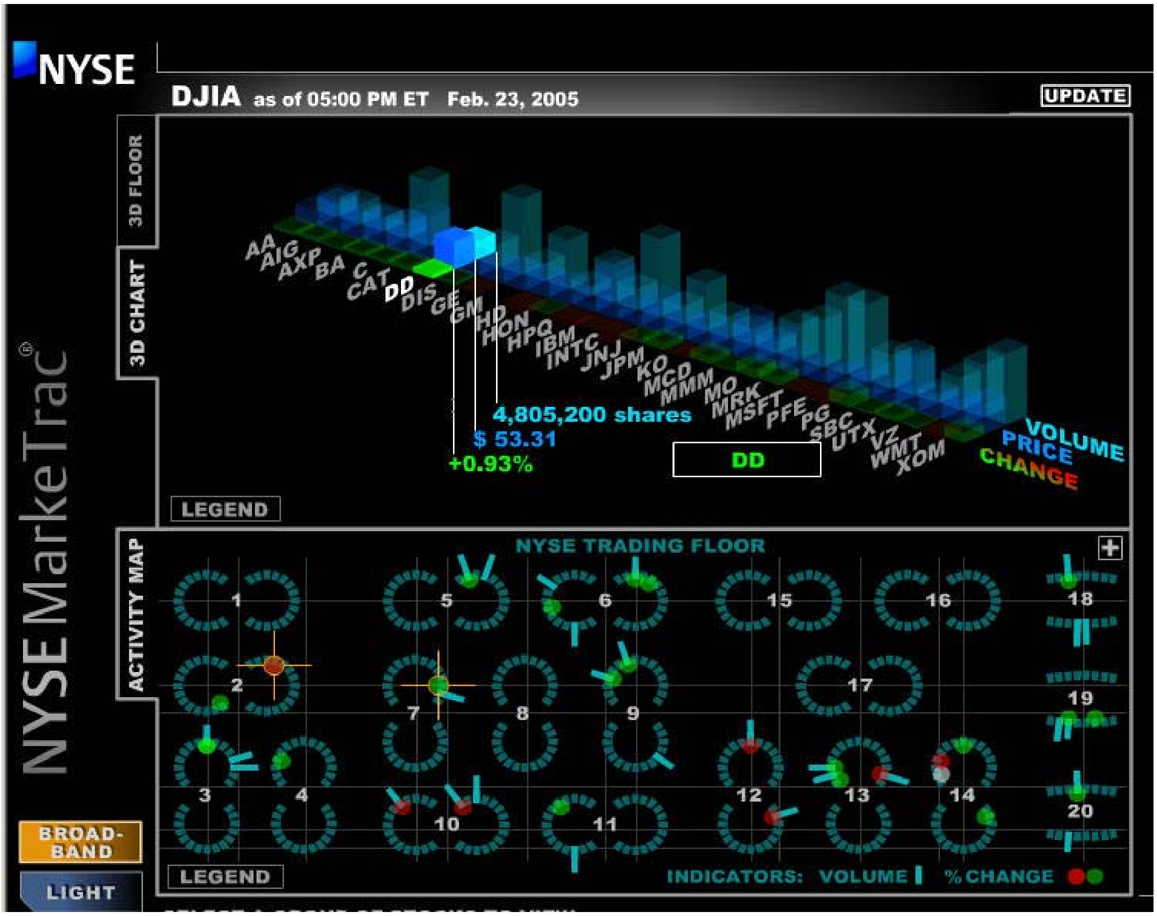

High Information Bandwidth Market Visualization

Seeing the market in the relevant dimensions is a necessary first step to trading effectively, and to intelligent automation of algorithmic trading beyond "one size fits all" strategies. Here is a current example of a high end viewer for market transparency data, trades, quotes, and the visible part of the limit order book. For some purposes, including the refinement and monitoring of algorithms, this kind of information is incomplete. Activity by algorithms make the display of the book an incomprehensible blur, or a dubious snapshot. Rates of cancels and replacement of limit orders are informative here. There are many real time trading analytics that can be used to decide when to help the algorithms out, by "changing the information variable".

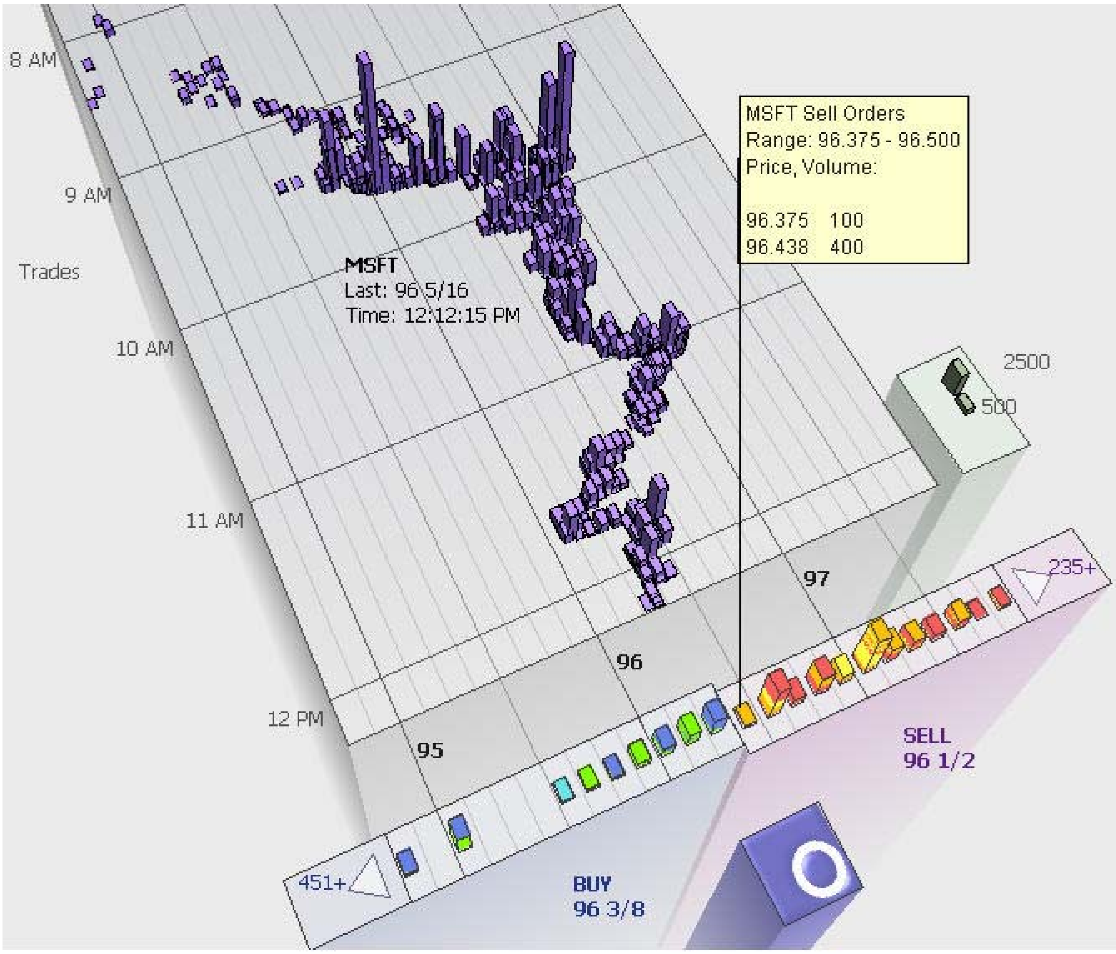

Here is a snapshot of a dynamic view of a market with limit order book transparency from Oculus Information's Visible Marketplace (c. 2004 Oculus Information). Layers on the book show grouped sizes at each price. As with all of these graphics, the dynamics and drill-down-to-details ability are content-rich aspects that don't translate well to the printed page. [20]

<Oculus is glad to have us use this. Contact is Bill Wright, bill.wright@oculusinfo.com>

If you are skillful and lucky, these market analytics will help you catch events before they are over, some of the time. Often, they are contemporaneous, or too late to exploit. News and other textual information have the same characteristic. Sometimes they lag market events, sometimes they lead.

The best free market visualization is the map of the market, from SmartMoney[21]. The sizes of the rectangles correspond to volume, or capitalization (or whatever you want), and there are many overlays available.

Interpreting this sort of complicated stuff to decide when and how to steer the algo ship is where people make the decisions that the algorithms can't make.

[20] An animated version can be seen at www.oculusinfo.com

[21] http://www.smartmoney.com/marketmap/

Agents for news and "pre-news"

The Map of the Market display combines measures of return (color), trading (size) and news. It can be interesting and useful, but for news what you are seeing is just the tip of the news iceberg. At least one headline daily for nearly all of the companies in the S&P 500 is almost inevitable. Certain types of news, e.g. earnings related, are more important than others. Market tools to deal intelligently withy news, for both humans and algorithms, are still evolving

Many news events lag the market, some lead it. An excellent (and free) place to see this live is the new Google Finance site. They combine excellent market graphics with an overlay of news stories. The example below shows news events A,B, and C, lagging the market for Accentia Biopharmaceuticals, which gained 57% on Oct 19, 2006.

Notice that all of the stories shown on Google occur after the big move. Either this is insider trading on a massive scale, or there is something missing. Don't call the SEC just yet, something is missing. There was a 9AM press release, not picked up by the mainstream media. It said, in part, ".., a new study shopwed that [Accentia's]… drug… can induce complete long-lasting remissions". News doesn't get much better than that for pharma companies.

This example illustrates the disintermediation of news, where those people who were looking at raw "pre-news" found the golden nugget before the main steam media. Similar nuggets are there to be found on thousands of sites, from news to legal to corporate to blogs. Looking at all of them all the time is a job for agents

For the pharmaceutical example here and similar situations, the approach to automation can be inferred by looking at the text in more depth. It reports on a clinical trial in glowing terms, "complete long term remissions", "80% achieved immune response", "complete tumor remission".

You definitely want to know when similar news comes from similar labs working for similar clinics testing hundreds of drugs for hundreds of companies around the world.

This is an example of persistent search. You can be persistent and do it yourself, early and often, or automate the process, where the machine finds and you evaluate.

Blogs and other forms of social media, are a new source of investment information. There are many items of anecdotal evidence to support the idea that bloggers sometimes have valuable information. Monitor110[22], a specialized Web 2.0 financial information firm, harvests blog content, and tracks reputation and relevance, tickerizes. Human analysts particpate in this process, along the lines of agent/human cooperation described previously.

Information inefficiencies we have seen so far are more common for smaller cap stocks. But they are not limited to small cap names, as seen in the report about MSFT below

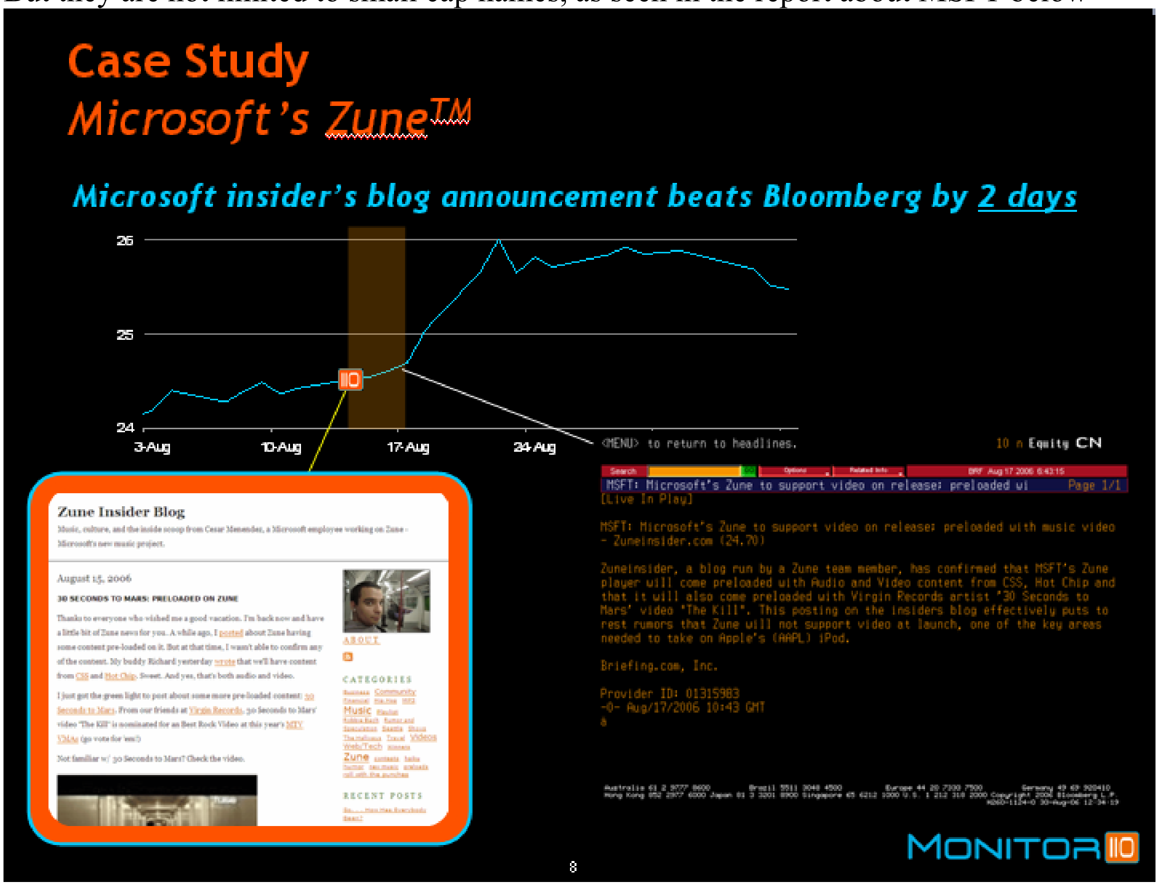

Innovative products will have features to facilitate this kind of selective approach to news, processed and in raw form. No dominant commercial paradigm has emerged. There is a great deal of research on gathering, aggregating characterizing, and filtering text. Some interesting start ups in this area are funded by the In-Q-Tel, the CIA's venture capital firm.[23]

Algorithms at the Edge

[22] www.monitor110.com the author is an advisor the this company.

[23] http://www.inqtel.org/

Algos will be pushed in all directions that will improve their performance. Mathematical models will improve. Adaptive probing strategies will adapt and probe. Latencies will go to zero and information will go to the sky. The minute to minute market games people used to play are now millisecond to millisecond games for computers. Understanding information and how to work with algorithms to profit from it are good skills for participants in the new fully wired financial markets.

You may be working with state of the art human computer interfaces, blazing trails in cyberspace to boundless wealth. But as my wife and other fans of stuff like Project Runway know, the important question is "what will you wear?"

In the web sleuth spirit of this article, I looked for images matching 'sensors effectors control human" and found this

Black is still good. Headphones big. Hair small. You will look very fine.

####