

Algo vs.

By David Leinweber

Algo

Algorithmic trading has changed investing forever. Today more than 90 percent of hedge funds use electronic trading algorithms — computer models designed to execute trades more easily, cheaply and, most important, anonymously. As David Leinweber writes, the battle for algorithmic trading has turned into a full-scale arms race. The winners, he predicts, will design algorithms able to probe, learn and adapt to the increasingly complex information available to them. Leinweber knows of what he speaks. For much of the past 30 years, he has been in the trenches, building or advising on some of the computerized quantitative trading and investment systems that make algorithmic trading possible.

Remember *Mad* magazine's "Spy vs. Spy" comic strip? Created by Cuban artist Antonio Prohias after he fled Castro's rule to go to the U.S., "Spy vs. Spy" has been running continuously in *Mad*'s "Joke and Dagger" department since 1961. The comic features two spies, identical except for the colors of their coats and hats, engaged in an endless series of elaborate schemes of Cold War one-upmanship to try to gain an advantage. *Mad*'s spies use an assortment of daggers, explosives, poisons, booby traps and Rube Goldberg-like machines to try to win their war.

The battle for supremacy in algorithmic trading is simi-

larly amusing. Traders use an assortment of mathematics, programming, communications, computing hardware and, yes, Goldbergesque schemes in hopes of gaining an edge in electronic trading. Like the longtime rivals in "Spy vs. Spy," who continue to fight well after the Cold War has ended, today's algorithmic traders are engaged in an arms race that shows no signs of slowing down.

It's helpful to understand the simpler beginnings of electronic trading to better appreciate today's elaborate systems and the ever more elaborate systems that will replace them. When markets involved chalkboards, shouting, hand signals and large, paper limit-order books, there

Algorithmic Trading

was no possibility of using a computer to execute trades.

That changed in 1976, when the New York Stock Exchange introduced for its members the designated order turnaround, or DOT, system, the first electronic execution system. It was designed to free specialists and traders from the nuisance of 100-share market orders. The Nasdaq Stock Market, which opened in 1971, used computers to display prices but relied on telephones to transact trades until the introduction of the computer assisted execution system, or CAES, in 1983 and the small order execution system, SOES, in 1984.

Simultaneous improvements in market data dissemination allowed computers to be used to access 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, dramatically reducing the latencies in trading. A 2006 study of trading before and after the upgrade, by professors David Easley of Cornell University, Terrence Hendershott of the University of California, Berkeley, and Tarun Ramadorai of the University of Oxford, found major improvements in the quality of executions.

Early electronic execution channels were for only the smallest market orders. But the permitted sizes grew fast. Support for limit orders was added. DOT became SuperDOT, and the tool was adapted for direct use by the buy side, first by the little guys — a joint venture between Richard Rosenblatt, founder and CEO of Rosenblatt Securities, and a technology provider, Davidge Data Systems (more on that later) — and then by the big brokers, who gave the product away for clearing business. SuperDOT and the automated Nasdaq systems accommodated ever-larger orders. Those exceeding the size limits for automation were routed to specialists and market makers.

This was algorithmic trading without algorithms, an early form of direct market access. The first user interfaces were designed for one stock at a time — 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 1-2-3 spreadsheets, which totaled everything up nicely to avoid costly errors.

Algorithmic trading was only one step away. As 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, it had snuck up on us.

Early adopters of these ideas were not looking to minimize market impact or match volume-weighted average price, or VWAP, the price at which the majority of a stock's trading occurs on a given day. They were looking to make a boatload of cash. Nunzio Tartaglia, a Jesuit-educated Ph.D. physicist, started an automated trading group at Morgan

Stanley in the mid-1980s and hired young Columbia University computer science professor David Shaw. At first, the group produced a few papers about hooking Unix computer operating systems to market systems. Then the former academics realized there was no alpha in publications. Shaw went on to found D.E. Shaw & Co., one of the biggest and most consistently successful quantitative hedge fund managers. Fischer Black's Quantitative Strategies Group at Goldman, Sachs & Co. was another algorithmic trading pioneer. Goldman's quants were perhaps the first to use computers for actual trading as well as identifying trades.

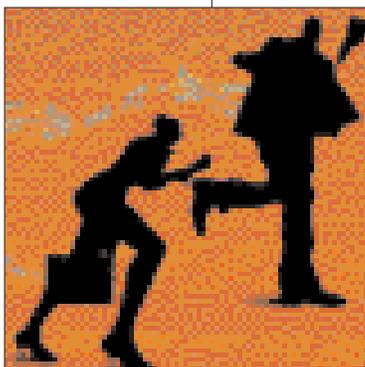
The early alpha seekers were the first combatants in the algo wars. Pairs trading, popular at the time, relied on statistical models to find the relationship between the price movements of stocks. Finding stronger short-term correlations than the next guy delivered big rewards. Escalation beyond pairs, to groups of related securities, was inevitable. Parallel developments in futures markets opened 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 was easy money. Real market makers have obligations to maintain two-sided quotes for their stocks, even in turbulent markets, and this is often expensive. Electronic systems, without the obligations of market makers, not only are much faster at moving quotes, but they can choose when not to make markets in a stock. David Whitcomb, founder of Automated Trading Desk, another algo pioneer, describes his firm's activity as "playing Nasdaq like a piano." There were other piano players: Morgan Stanley transformed its trading desk into an automated market-making system. Along with firms such as Getco and Tradebot Systems, they came to dominate the inside quote and liquidity in the largest names today. Joseph Gawronski, president of Rosenblatt Securities at the NYSE, says that the algo wars have brought a massive change in market structure.

Faster data feeds and faster computation let you run ahead of the other kids in line. In the early 1990s the lag between one desktop data feed and another might be as long as 15 minutes. The path from market event to screen event had significant delays. Slow computers, sending information to slow humans over slow lines, were easy marks for early algo warriors willing to buy faster machinery and smart enough to code the programs to use it. This aspect of the arms race continues unabated today.

Before long the industry noticed that these new electronic trading techniques had something to offer to the buy side. Financial journals offered a stream of opinion, theory and analysis of transaction costs. Firms such as Wayne Wagner's Plexus Group made well-supported arguments about the high cost of transactions. Pension plan sponsors, sitting atop the financial food chain, were persuaded in large numbers.

Index managers did not have to be persuaded. With no alpha considerations in the picture, they observed that it was possible to run either a lousy index fund or a particu-



larly good one — the difference was the cost of trading. Those passive managers, on their way to becoming trillion-dollar behemoths, were high-value clients to brokerages.

In addition to giving their high-value clients what they wanted, brokerage houses had another incentive to adopt electronic trading. The demise of fixed equity commissions had spawned new competitive pressures. Electronic trading had the potential to cut costs dramatically while improving quality of service.

The biggest firms developed their own electronic order-entry systems. Others bought from niche vendors. One of these was Davidge Data Systems, headquartered in a loft near the meatpacking district in New York City, not far from Wall Street. Nick Davidge had many clients to support, and he used bicycles to dispatch service people, including himself.

The first direct-access tools from the sell side were single-stock electronic order pads, followed shortly by lists. By this point, the sell side was looking for a way to break orders into pieces small enough to execute electronically and spread them out in time. Innovative systems such as Investment Technology Group's QuantEx allowed traders without large software staffs to use and define analytics and rules to control electronic trading. The result was what we consider to be algorithmic trading today.

THE BIG NEWS IN ALGORITHMIC trading in the late 1980s 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 strategies were 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, bagging your rivals in classic “Spy vs. Spy” style. There was little or no mathematical underpinning, just rules of thumb and educated guesses.

The obvious shortcomings of these simple strategies inspired several generations of mathematically based algorithms that used 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 conceptual foundation analogous to the efficient frontier in portfolio theory.

Markets have become even more fragmented and complex, with less information conveyed by the best bid and offer, or BBO, and the book — creating a need to exploit new order types and to access “dark liquidity.” This has given rise to behavior-based algorithms that probe for liquidity, driven by procedural logic and stimulus-response principles as well as mathematical models.

Algorithms need to probe, learn and adapt. They need to make effective use of analytic tools and learn how to recognize their limitations. Algos at the edge seek to exploit information beyond the traditional data, including news, prenews and other forms of market color found on the Web. There has been an explosion of progress in tools

for processing text. Think, for instance, of Google.

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 intelligence methods is increasingly effective. But how best can human traders work with algorithms, using intelligence amplification to form a partnership that enhances the skills of both? Finding the proper mix of human and machine skills is a challenge for traders. “Humans definitely cannot react faster, but they can react smarter in many instances,” Rosenblatt Securities president Gawronski observes. “One thing algos do extremely well is allow for one to reflect what one anticipates they would want to do if a certain set of circumstances occurred.”

As Gawronski explains, a human trader reacts in the true sense to new information and changes his plan based on that new information. An algorithm works differently, being forced to anticipate what will occur and then having a set plan for dealing with those circumstances if they do in fact occur. “In an auto-execution, millisecond world,” he says, “one has no choice but to use algos and play the anticipation game, as trading will go on without your participation if you simply try to react.”

Garry Kasparov, the world chess champion who lost to Deep Blue in 1997, suggested that chess tournaments be open to human-machine teams. 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. Applying this analogy to trading, imagine if the game were not tournament chess, which allows up to seven hours for a game, but blitz chess, which allows each player just a few minutes per game. Given Moore's Law, it wouldn't be long before the computer that beat Kasparov with seven hours could beat him with three minutes. Many facets of trading are more like blitz chess than tournament play.

THERE IS NO SHORTAGE OF paycheck anxiety among traders. Their numbers have been dropping. Specialist firms have been cutting staff by 30 to 50 percent, says Gawronski. “Algos are being employed to do some of the routine heavy lifting of market making,” he explains. Last spring industry maven IBM Business Consulting Services published a report titled “The Trader Is Dead, Long Live the Trader!” Like global warming, the changes in trading are a reality that can't be ignored.

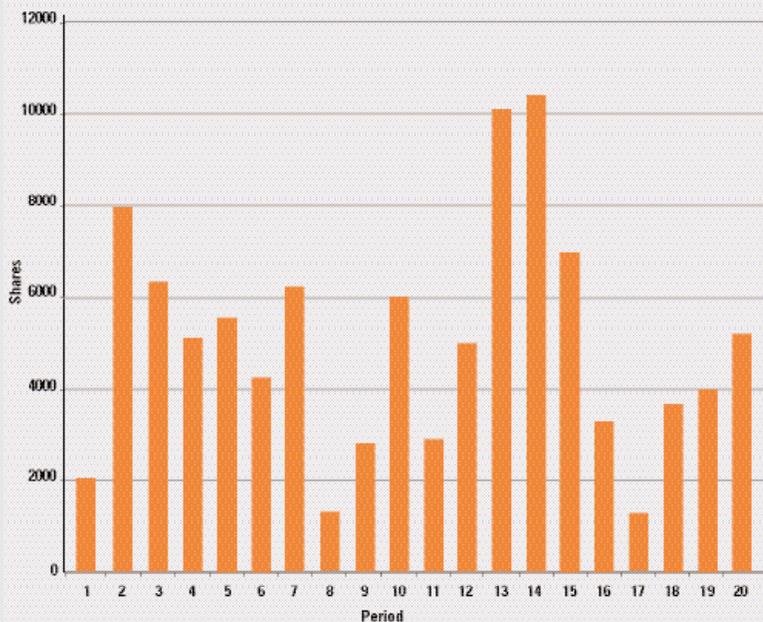
The traders who survive will be the ones who play well with machines. Understanding algorithms is critical.

“Early adopters of these ideas were not looking to minimize market impact or match VWAP. They were looking to make a boatload of cash.”

The Academic Argument

Algorithmic trading got a boost in 1998 when Massachusetts Institute of Technology professors Dimitris Bertsimas and Andrew Lo showed that properly parceling out equity trades in smaller share packages over a given time period minimizes the expected execution cost.

Best-Execution Strategy



Information Component



The figures above, reproduced from Bertsimas and Lo's 1998 paper, "Optimal Control of Execution Costs," show the relationship between a best-execution trading strategy and the information component. The top figure presents the number of shares traded in each time period; the bottom figure has the corresponding information component. As the illustration demonstrates, trading is strongly driven by the information component.

Algorithms have sensors and effectors, analogous to the eyes and motors of robots. In between the sensors and effectors, there is a computer program that provides control.

Sensors include the data feed of market information, quotes, trades, order books and indications of interest. Algos feed on market data, and their sophistication grows with the data's scope, timeliness and accuracy.

Effectors are order-entry components, including instructions to cancel or modify. 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.

You are what you eat, so a basic algo war tactic is to improve the timeliness, scope and accuracy of market data. Anyone using more than one data service notices lags from one to another, and they all lag the event. Companies like Wombat Financial Software will sell you the docking adapter to sidle right up to the Securities Industry Automation Corp. original mother ship, where the price and volume data of stock sales get consolidated, so you don't have to rely on data vendors like Reuters and Bloomberg.

In the algo wars, as in real wars, it's a good idea to control your communications and avoid those slow satellite links. Communication satellites are in geosynchronous equatorial orbits 22,240 miles above the equator. Light travels at 186,000 miles per second, so a satellite hop takes at least 250 milliseconds, long enough for a crowd to get ahead of you.

You can 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 Stock Exchange, Euronext, Tokyo Stock Exchange, Globex and a growing list of other market centers. Co-location, as it is called, can do wonders for latencies in execution. As the algo wars proceed, brokerages willing to commit capital will be able to offer zero-latency executions. Zero is sure to appeal to 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 and its own clientele (say, brokerages or the buy side). In many of these systems, order size is hidden. Finding liquidity may require being in multiple systems for a period of time. This can create a risk of overexecuting unless very conservative rules are followed. Larger firms, willing to risk some capital by incurring the risk of overbuying (or overselling) can allow their clients to make use of more-aggressive trading tactics. Algos at the edge combine analytic tools with expert rules and procedures to profit from the complexity of multiple execution systems.

Future algos will have uniform access to a mix of securities and derivatives. This opens a door to improve patient execution of large "transitions," or program trades, by controlling risk. Nearly all current algo trades occur over the course of a single day. There is no fundamental reason for this. Without a one-day rule, future algos will better serve

institutions by using patient transition-trading to make sizable adjustments in their portfolios. Full-service brokerages will be able to offer customized short-term derivatives for controlling risk exposures along the paths of longer trades.

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

THE EARLIEST ALGORITHMS USED the “keep it simple, stupid” strategy of splitting orders into N parts, every $1/N$ of a trading day: For example, an order for 10,000 shares would be sent out as ten orders for 1,000 shares, at ten times spaced equally over the trading day. This signaling made it easy for traders on the other side to spot these algorithms and pick them off.

The next round of the algo wars was to get algorithms to be less naive, to hide trades by randomizing times and sizes. This worked very well, according to a 2004 study by Quantitative Services Group of actual institutional trades. It showed a reduction in cost from 26 basis points per trade to 2 basis points. But randomization can make some stupid decisions — placing small orders at the open and close, not reflecting urgency or tolerance for risk, missing transient opportunities in liquidity.

In 1998 professors Dimitris Bertsimas and Andrew Lo of the Massachusetts Institute of Technology co-authored one of the first academic papers on scientific approaches to trading, “Optimal Control of Execution Costs.” They start with an analysis of the merits of mindless naive strategies, asking in what sort of environment would such strategies be optimal. This turns out to be 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 examine how optimal trading strategies depend on it. Their findings, some of which are reproduced in the figures at left, show that trading is strongly driven by the rather abstract information variable. Determining the information variable is not easy and could include anything from conducting a microlevel empirical analysis to listening for rumors on a bus.

Modeling market impact and information was a significant advance. The next step was to incorporate the idea of risk aversion and the distinction between passive and alpha-seeking trades. In their groundbreaking 2000 paper, “Optimal Execution of Portfolio Transactions,” Robert Almgren and Neil Chriss introduced the idea of using liquidity-adjusted value at risk as a metric for trading strategies. The research of the two former professors, who now work at Banc of America Securities and SAC Capital Management, respectively, has been widely adopted in today’s algorithmic systems. Here is what they did, as explained in the paper’s abstract:

“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 by minimizing Value at Risk.”

Almgren and Chriss show how to optimize trading strategies to create an efficient frontier for trade execution depending on the risk tolerance of the trader. For risk-averse traders, accelerating the execution speed will reduce risk but at the cost of higher market impact. A trader with short-term alpha would use such a strategy to reduce opportunity costs. More aggressive traders, who likes risk, would slow down the execution to get more risk, also incurring more market impact cost.

Mathematical models of markets can become very elaborate. Game theory approaches to other market participants, human and machine, in the spirit of the *Beautiful Mind* ideas of John Nash, can bring a further level of insight. But Almgren and Chriss remind us about the limitations of all model-driven strategies: “Any optimal execution strategy is vulnerable to unanticipated events,” they write.

“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.”

This comment from Almgren and Chriss is the academic version of the wisdom of former secretary of Defense Donald Rumsfeld, who said there are “known unknowns and unknown unknowns.” In the investment world, known unknowns include scheduled announcements like earnings or conference calls that affect particular stocks; those like housing starts that affect groups of stocks; and those like macroeconomic data and interest rates that affect broad markets.

There are many sources of this information. Thomson StreetEvents offers a wide selection of potentially market-moving corporate information found in Securities and Exchange Commission filings and quarterly earnings calls. Econoday, a calendar book of upcoming economic data releases and U.S. Treasury announcements, long found in trading rooms, is now a Web service. Some algorithms use this information, some don’t. Guess which ones are better.

Unknown unknowns include news, discussion, rumor, market color, agency actions and research results. Computers are pretty good at finding this kind of thing — often, too good. Determining when an “unknown unknown” will change the trading strategy is a place where humans working with machines have an edge over either working alone. The microstructure tactics based on these cost-minimizing trading models are also deployed in VWAP and similar applications. These anticipate volume

and try to participate throughout the day (or given time period), optimizing to those volume and price targets.

Models are not markets. Even the most elegant models are abstractions of true markets. The real thing is a rapidly changing mélange of market fragments, continuous and call markets, electronic communication networks, innovative matching systems, indications and dark-liquidity pools.

IN THE DAYS WHEN STOCKS were measured in eighths and orders were being modified by people, the inside quote conveyed actionable information. Now the average life of a limit order is measured in milliseconds, and the quote is a fast-moving target. With decimalization, the old total size at the inside spread was distributed over six or more price levels, and the best bid and offer conveyed much less information. So ECNs and the exchanges exposed more of the book. But just when you could see the book, the algo battlefield shifted with dark liquidity, hidden pre-programmed orders to execute when others are filled, anonymous indications and matching systems. These take the liquidity that back in the day would have been 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 to tell us what to do. Algorithms at the edge can use models as a basis for action,

but they have a wider vocabulary of rule and procedural tools to execute across all market segments. As markets change, people will need to monitor and adjust algo and electronic strategies. Markets change rapidly, so humans will be important here.

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

Robots have done well in complex dynamic environments. Looking into how these robots “think” is looking at the future of algorithms. Looking at how humans and physical robots interact reveals how humans and trading robots will coexist.

There are always multiple approaches to robotic tasks. Structuring and coordinating these approaches is the goal of multiagent 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.** The world in general and markets in particular are not static. Things change;

information is incomplete. A reactive agent 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, with wide error bands. This is particularly true for equities. Financial models never capture every aspect of market participants’ motivations.

- **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 and reinforcement.** Performance measurement is natural for trading agents. For alpha-seeking algos, metrics like the Sharpe ratio fit. Pure execution algorithms use implementation cost or VWAP shortfall.

- **Layered behaviors.** Agents should have default behaviors that complete their tasks and avoid errors. Basic behavior is at the lower layers; more sophisticated behavior is above.

Some of these agents will be programs; some will be people. We can call these people “the employed traders of 2015.” They will be operating in a world where electronic equity execution is rapidly becoming a commodity and the buy side is able to bypass the sell side to access markets directly. Investment firms, both the bulge-bracket outfits and the wanna-bes, will be driven to develop full-service electronic interfaces to accommodate complex, multi-asset-class, leveraged trades.

Two recent prognostications for markets in 2015 are remarkably similar. One is a study by McLean, Virginia-based management and technology consulting firm BearingPoint, titled “Shifting from Defense to Offense: A Model for the 21st Century Capital Markets Firm.” The other is “Profiting Today by Positioning for Tomorrow: A Field Guide to the Financial Markets of 2015,” by IBM Global Business Services.

Both reports describe a shift from a product paradigm to a risk paradigm. Predictions include a willingness by both investors and sell-side firms that act as principals to commit capital in innovative ways and increased trading interest in risk classes over individual securities. They forecast an increasingly risk-centric view of trading, driven by the demands of complex alpha-seeking strategies.

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 — of large firms. Agents will have to be able to price these derivatives using quantitative measures and their firms’ risk profiles.

People will have to find their places in the multiasset, risk-mitigated, fragmented, algorithm-infested markets of 2015 and beyond. It is informative to ask how people today work with other algorithms, such as physical robots.

Some of today’s real robots work largely on their own.

“Algos at the edge seek to exploit information beyond the traditional data, including news, prenews and other forms of market color found on the Web.”

They have stimulus-response rules and internal representations of their tasks. There are 2 million iRobot Roomba vacuums sucking up dirt without human assistance.

The Mars Rovers, Spirit and Opportunity — the Energizer Bunnies of space — have a lot of control over their actions. Each has an autonomous mobility system. Humans set the goals; the rover takes care of the rest.

Other robots are kept on extremely short leashes. The iRobot PackBot Explosive Ordnance Disposal robot comes with a substantial remote control to manipulate the machine's electro-optical infrared thermal camera and other cool features. The U.S. military uses PackBots, made by the same company that brought us the Roomba, to identify and dispose of explosive devices in Iraq and Afghanistan.

Robot surgeons, like the Da Vinci Surgical System robots, are on the shortest possible leash. A human surgeon controls the robot's every move. This is really a teleoperated system, with very little autonomy other than safety stops.

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 tiny ones in tubes for Dr. da Vinci. Force feedback and texture sensors let users feel what it is like to be there. In the real world of bombs and gallbladders, looking around is a great way to work with robots.

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. Ideas about how to do this have evolved from the simple to the 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, such as 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.

These modern visualizations use techniques beyond the usual "picture on a screen" static displays. Many are three-dimensional, interactive and dynamic. The NYSE's MarkeTrac, which is available live on the Web, combines a stylized view of the trading floor with updated displays of market activity. Even better is Oculus Info's Web-based visible marketplace, which provides a three-dimensional view of how stocks trade over the course of a day, showing the relation between such things as rate of cancels and replacement of limit orders. The ability to drill down using visible marketplace or MarkeTrac can help humans turn the flood of market data into useful information and catch events before they are over.

The salient feature of the relationship between news and markets is that many news events lag the market, but some lead it. Textual and news systems like Google and its automated cousins help traders find the kind of unanticipated events that modify algorithmic strategies. Google

Finance, which combines excellent market graphics with an overlay of news stories, on occasion will pick up a press release not carried by the mainstream media. The growing disintermediation of news creates opportunities for traders to mine such golden nuggets of raw prenews. A trader interested in pharmaceuticals stocks, for example, would want to follow press releases from clinics around the world testing hundreds of drugs for hundreds of companies. This is an example of persistent search. Human traders can be persistent and do it themselves, early and often, or they can automate the process, using machines to find news for them to 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. Although innovative algorithmic systems undoubtedly will facilitate the use of news in processed and raw forms, no dominant commercial paradigm has yet emerged. There is a great deal of research on gathering, aggregating, characterizing and filtering text, some of which is being done by start-ups funded by In-Q-Tel, the venture capital arm of the U.S. Central Intelligence Agency.

Algorithms are 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. Traders who learn to work with algos at the edge will be throwing orders at the market faster than ever. Those who don't will suffer the same fate as the two heroes of "Spy vs. Spy," locked in a never-ending battle that neither one can win.

David Leinweber is an independent consultant on financial technology at Leinweber & Co. in Pasadena, California, and scientific adviser to New York-based Monitor110, a specialized Web 2.0 financial information firm. He founded two financial technology firms — one in algorithmic trading, the other to extract alpha from language on the Web. For seven years he managed \$6 billion in quantitative equities at First Quadrant in Pasadena. He was a visiting scholar at the California Institute of Technology and an information scientist at Rand Corp. He holds undergraduate degrees from the Massachusetts Institute of Technology and a Ph.D. in applied mathematics from Harvard University.

Leinweber would like to thank Robert Almgren, Jay Dweck, Joseph Gawronski, Scott Harrison, Bill Harts, Eli Ladopoulos, Richard Lindsey, Andrew Lo, Richard Rosenblatt and George Sofianos for their helpful comments on this article. Please contact the author at dleinweber@post.harvard.edu if you would like an expanded electronic version of this article with Web links, footnotes and pictures.

