



HAL
open science

What if the Screens Went Black? The Coming of Software Agents

Karin Knorr Cetina

► **To cite this version:**

Karin Knorr Cetina. What if the Screens Went Black? The Coming of Software Agents. Working Conference on Information Systems and Organizations (ISO), Dec 2016, Dublin, Ireland. pp.3-16, 10.1007/978-3-319-49733-4_1. hal-01619192

HAL Id: hal-01619192

<https://inria.hal.science/hal-01619192>

Submitted on 19 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

What If the Screens Went Black? The Coming of Software Agents

Karin Knorr Cetina

University of Chicago
knorr@uchicago.edu

Trading screens are not supposed to be black. In fact, when we see them on trading floors, on TV, or in media centres, they attract us with catching colours and blinking information. They project urgency, speed, and power – the power of big money, the power of winning and losing. When we are near them, we feel their heat. We want to give in to their considerable attraction. We want to be players of the game and part of the action.



Of course, the screens are only a medium. What beckons us is the market. All the beguiling signals come from a market that moves – prices, volumes, ratios, differentials. They move in a context and in response to it. The context, of course, also moves, but usually more slowly. How often do politicians make decisions? How often do tsunamis strike? The context is everything that happens in the world that may have an impact on the market. But that everything happens on a human schedule and is adapted to how we conduct our lives and how nature and the environment conduct theirs. What parts of the context impact a market? The question is ongoing and difficult. Should the stars and their movements have a presence onscreen? Astrologists may think so, but Reuters and Bloomberg don't, and they are two of the largest companies feeding material to the screens. Is Namibia a country that should be present in a global financial market? It has its weekend and peasant markets, to be sure, and some stocks, but its gross domestic product is shrinking and is little more than a rounding error on the trillion dollar economies that count. Trading screens cover the three major time zones and not every little nation and group that populates the globe.

Even with these restrictions, however, trading screens are rich in texture and information. They are not just surfaces. Behind every window, dozens of others can be opened and closed. Calculations can be called up and made on the spot. Information can be requested and is instantly available, at all times. Social relations carry on and jokes are made. And, of course, one trades through the screen, through the electronic channel. All trading in currency markets, for instance, is conducted electronically.

Screens are rich, and they are not only there for trading. They are the world for those who sit in front of them. The question posed at the beginning of this essay is larger than it may have initially seemed. When the screens go black, a world goes black. But what happens if that world goes black?

Twenty years from now, even in ten years, the answer could be “nothing at all”. The screens project the market, but in the future the market might carry on just as well without screens. The contextual information will be necessary and will remain available – available to programs that receive and decode it without displaying it onscreen. The scopic regime under which electronic financial markets operate could simply migrate fully underground. Why “scopic” and why “regime”? Financial screens are like scopes. A scope is an instrument for observation. A periscope, for example, protruding from the surface of the ocean, can give a submarine a 360-degree field of view – a view richer and deeper than that of the naked eye. A telescope, using lenses or mirrors, can “see” the objects that circle us on planet Earth – it brings close what’s behind the twinkle in the sky that the eye can see. The screens that project the market are not optical instruments, of course. They bring the activities of globally dispersed participants near and condense them onscreen, making them visually available to everyone connected to the information technologies that produce and transmit the screens’ content – in real time and in real speed. News about changes in the market’s context augment financial content: political activities (e.g. elections), announcements related to the economy (e.g. Fed forecasts and earnings calls), social events (e.g. riots), cultural trends (e.g. tourism), and natural disasters (e.g. floods). The observational system at the centre of a fully electronic financial market then uses its knowledge of market transactions and external events to perform instant analyses of the material. It combines new data with historical content and enriches both with financial indicators processed elsewhere and released to the screen. This extension in the scope of financial transactions to, in principle, all that is globally pertinent to the market, as well as to the analysis of these materials, is part of the notion of “scopic”. A scopic

regime then simply refers to a combination of technological media, software, and information that enables a global financial market. I call it a regime because scopic media rule the market, so to speak. Without these media, the market could not operate. Of course, the actors in the market could try to go back to telephone trading, or even to voice trading as it once existed in alleys and streets. But unless all computerized infrastructures break down, no one in the foreign exchange market, with its \$4 trillion daily turnover, could survive. A scopic regime forces actors to participate in the current venue if they want to be in the market – though it also allows and even encourages improvements, customization, and changes of venue. The point of the current venue is that it works by making the flow of the market visually available in augmented fashion for human actors to perceive, interpret, and act upon. It also streams the market, putting it in a visually running sequence. In other words, it continually updates the market as new information flows in and old information flows out (or scrolls down the screen).

When things migrate underground, something splits off the huddles of traders and screens of the trading floors we know. When it splits off and turns inside, so to speak, another world opens up underneath the level of human agency: that of data centres, programs, and electronic connections. This is the environment of algorithms, or as we may want to think of it, of software agents. Software agents operate underground already, but they only partly substitute for humans. Human traders are still on the floors and in exchanges, watching the market, and watching what their “tools”, the algorithms, do. Traders refer to algorithms as their tools. Having a tool suggests that you have control over it and that it serves you. In the foreign exchange market, for example, as judged by evidence I have from two top banks in worldwide trading volume, trades valued at over \$5 million (first bank) or over \$50 million (second bank) are executed by human traders. Those valued less than these amounts are left to algorithms, but human traders oversee the activity. The British Foresight report that investigated the future of computer trading just recently stated that: “Despite the increase in computer power, processing speed and sophistication of computer algorithms, the present day financial markets still involve large numbers of human traders”. And it continues: “There are good reasons to expect, for the next decade or so, the number of human participants in the market will remain significant” [1] (p. 32).

As the timeline and the way things are hedged in this formulation suggest, the presence of algorithms already is quite significant. Today, traders expect that when things migrate underground fully, perhaps after the next decade, some human traders and others are likely to remain on the floor

and continue their work. For current traders, it is hard to imagine a market without human participants, or without screens. Such an arrangement would not be possible, they tell me, since someone has to be there to watch and think and to develop more tools. Let's assume they are right and ask: what it would mean for our trading culture and sociology if the screens went "half black"? What would it mean if traders were still around but had limited roles to play, and underground agents had taken over plenty of trading functions? I want to turn to this shortly, but before I do so, here are some data and definitions, and an explanation of algorithmic trading.

Electronic-Automatic-Algorithmic: Some Data and Definitions

The Bank for International Settlement (BIS) described the rapid growth the foreign exchange market experienced between 2007 and 2010, when it was last measured, as the "the \$4 trillion question: what explains FX growth since the 2007 survey?" [2]. The foreign exchange market had grown during that time, the time of the financial crisis (!), from a daily turnover of \$3.21 trillion to approximately \$4 trillion – after an already unprecedented rise in activity by 69% between 2004 (\$1.420 trillion) and 2007. The \$4 trillion question can be answered, the BIS says, by a 20% surge in high-speed buying and selling by computer programs using algorithms.

The number of nonfinancial institutions, or retail traders, also increased during the period, making up almost 10% of the global market in 2010 [2]. These non-banks and day traders increasingly use algorithms offered to them through the trading platforms they employ. Several sources, including the BIS, estimated the total volume of automated trading based on algorithms in the foreign exchange market in 2010 to be about 25% [3] (p. 11). Traders interviewed in 2013 thought it had remained roughly around 30%. The volume varies between banks, as some banks make a stronger push towards automated trading using algorithms than others. Barclays, for example, was said to have only four spot traders left in foreign exchange, and they manage trades worth over \$50 million – everything else is done automatically through the "Algobook". One of the two largest Swiss banks joined the push towards algorithmic trading, and the other appeared to still ponder it in 2012 (interview data). As the BIS report noted, "FX dealing banks unsurprisingly do not welcome the compression of spreads" associated with the rise of high frequency trading, one of the main trading strategies using algorithms [3] (p. 9).¹ The total number of this sort of automated trading is considerably higher in the stock market. The Foresight Report estimated that

¹ "Spreads" here refer to the difference between bid and ask or buy and sell prices from which traders profit.

the percentage of trades that are automated can be as high as 60% in US equity trading, with estimates ranging from 30% to 60% for European equity markets. Equity markets are not global – equities can be purchased or sold globally, but they are organized regionally in exchanges – hence the differences between Europe and the United States [1] (p. 43).

Automated trading is electronic trading that depends on one or more algorithms at some stage of the process [3] (pp. 3, 5). Electronic trading that is not automated uses computers to trade, exchange information, chat, and everything else scopic media allow one to do through a screen. Electronic trading is, then, manual and “high touch”: a trader hits a key to trade what’s on offer on an electronic broker system, which displays the bid and ask prices and volumes onscreen, or a trader sends a message to deal directly with a counterparty using an onscreen conversational dealing system. Electronic trading also can be “medium” or “low” (human) touch, and only then is it algorithmic (and electronic) trading; more or less, the work of human traders – scanning prices and executing trades – is performed by algo-tools. When does it become what one might call “algo” touch? There are various types of trading based on algorithms; for example, they include the vast area of algorithmic execution, in which the computer program is responsible for executing a trading order. This sort of algorithmic trading “is the bullet, not the finger on the trigger”, since the algorithm isn’t making trading decisions [4] (pp. 5–7, 12–15). But algorithms that make trading decisions are increasingly common, and are based on models and incoming information. For example, they can read the news and interpret current events, and in response they can execute a trade and hedge against it. When the algorithm has its finger on the trigger, we can extend the current parlance and describe it as “algo” touch. Humans still outperform algorithms in analysing the semantic information carried in human-readable data streams, which range from written stories to audio and video sequences and tweets on social media websites. But news analysis performed by algorithm is a significant focus of research, with the goal of enabling computers not only to understand the numerical information of market prices, but also to understand non-numerical information. As this research advances, we will approach the “algo” touch mode, in which no human trader will be involved in trades, but in which software agents will be essential. Of course, the human touch in this case will just shift to other professions: for instance, to engineers and programmers that develop the software agents.

Is a software agent a professional? An algorithm is simply a set of instructions for accomplishing a task, according to various definitions. It may be a formula, or a set of rules, a set

of steps, or on the most general level an approach to solving a problem. The latter definition sees an algorithm more as an idea behind a set of instructions. There are likely to be many possible solutions to any problem, and for that reason there can be many different algorithms trying to accomplish the same thing. Even when the problem seems as simple as finding an entry in a phone book, there are different ways, as Durbin explains, of solving it: we can start from the beginning and flip pages until we find the name, or we can guess where in the book names with the first letter of the entry might appear and flip back and forth till we find it, or we could possibly work with an index, if one exists.² On the programming level, an algorithm is a computer code. When you write a program to filter the light in a certain way so as to create a specific photographic effect, you have created an algorithm. The use of algorithms in trading is not new. The Electronic Broker System (EBS), developed by a consortium of banks in the early 1990s, is a combination of programs that on the most basic level helps prospective buyers and sellers find one another. It does this by sorting incoming trading requests according to the best bid and ask prices, putting them in a sequence, by distinguishing and adding up volumes, by supplying the next price once a price–volume offer has been consumed, and so on. To do this, it needs criteria and information and the capability to sort, compare, archive, etc. An order-executing algorithm, for example, may simply perform order slicing; that is, it may split a large order into smaller segments of “child orders” that are put on the market (or sent to the exchange) sequentially, perhaps every hour. The goal is to reduce the impact one large order can have on the market price of a given security. A human trader likely would initiate this order by placing it in a more sophisticated way, for example, by placing individual child orders on the market at irregular times. A trader might also consider price movements before executing a portion of the order, and so on. Algorithms evolved at least in part by becoming more sophisticated by imitating experienced traders. For example, they learned to include more randomization – to elude market participants able to identify the linear actions of a tool that trades according to a fixed schedule. Then they became responsive – they learned to execute orders in response to live market volume – instead of using historical volumes to determine when to execute a child order. And then, according to Johnson, they started to base their responses on particular types of analysis or conditions, evolving within categories of algorithms for a while, before a new generation of algorithms took hold [4] (p. 14). The first generation of algorithms learned from humans – it is the one that evolved from a linear fixed execution schedule to responding to market

² I take the example from [5].

conditions in order to reduce market impact. The second generation learned from theories and models as it became more price and cost sensitive – it used transaction cost analysis to estimate and reduce transaction costs. That is when algorithms began to move beyond trying to imitate the behaviour of experienced traders. The third generation became reflexive – algorithms learned to examine and use data and venues that the market itself provided. For example, they learned to examine order books, which became more widely available as markets transitioned to electronic trading platforms, and to use more than one execution system, which also became available – e.g. electronic crossing networks (ECNs) and alternative trading systems (ATSs). This generational mobility will continue naturally, one assumes, though it may also be pushed forward by system-internal dynamics. Algorithms can be copied, which can eliminate their advantage. As a consequence, one looks to identify new advantages. The most sought after of these in the last ten years has been time advantage, which is used by high-frequency trading (HFT).

HFT is another subset of automated trading. It uses the famous processing power of computers to monitor and analyse thousands of variables to optimize trades. The latency periods exchanges achieved in 2010 were below 10 milliseconds – a real blink of an eye takes 300 to 400 milliseconds. HFT's first trademark characteristic is the use of extraordinary high speed together with sophisticated algorithms, a mix that is hardly beatable by a human being. IBM's Deep Blue began to overtake human chess masters in the late 1990s;³ HFT trading took off only around the middle of the 2000s and reached volumes accounting for more than half of stock trading in the United States in 2010. There are signs that this expansion is ending, given an increase in competition, the high costs of HFT, and new regulatory restrictions. HFT firms may co-locate their computers at exchanges to improve the speed of computer linkages, and they use data feeds offered for purchase to gain processing advantages over others in the HFT universe. One example that recently has caught the attention of US regulators is the sale of early financial data by exchanges and information provider firms like Thompson Reuters and Bloomberg. Thompson Reuters, for example, offered the release of the University of Michigan's consumer confidence index, a market-moving economic indicator, two seconds early to paying clients, and the company in turn paid the university "at least \$1million a year" to distribute the data early to its customers [6]. NASDAQ and the Chicago Mercantile Exchange promise to deliver NASDAQ data to Chicago customers two

³ For a brief history of the respective competitions see the Wikipedia entry: http://en.wikipedia.org/wiki/Computer_chess (accessed 16 July 2013).

milliseconds faster than they are otherwise available using microwave transmission, at a cost of \$20,000 per month to subscribers. Such costs of the race to receive information first, together with the costs of shaving additional milliseconds off trade times through the co-location of servers with exchanges and through building new data connections (e.g. from Chicago to New York), make it harder to generate profits, particularly during market downturns. Traditional investors, like mutual funds, have embraced the high-speed industry's automated strategies and moved some of their business away from exchanges that are popular with high-speed traders. "Profits from high-speed trading in American stocks", the *New York Times* reported in October 2012, "are on track to be . . . down 35 percent from last year and 74 percent lower than the peak of about \$4.9 billion in 2009", citing the estimates of a brokerage firm [7]. "Wells Fargo and JPMorgan Chase each earned more in the last quarter than the high-speed trading industry will earn this year", it added. HFT firms also were accounting for "a declining percentage of a shrinking pool of stock trading, from 61 percent three years ago to 51 percent now", according to the Tabb Group, a data firm. The report also mentioned that firms have been cutting staff, or even have shut down [7].

HFT has several other features having to do with the short time horizon of trading. For example, traders hold positions only for very short periods of time, they may submit large numbers of orders only to cancel them soon thereafter, and they end the trading day "flat", without carrying significant positions (inventories of an investment instrument) overnight.⁴ Observers also note that HFT firms implement typical trading strategies such as classical arbitrage, which exploits the difference between actual market prices and prices implied by "no arbitrage" conditions; latency arbitrage, which exploits the short time lag between the moment a trade occurs, the moment an HFT player detects the resulting change in the security's price, and the moment an update is made to the price quoted by market making traders;⁵ liquidity redistributing strategies, which detect imbalances in order books and pricing discrepancies across trading platforms; and "complex event processing", which exploits properties of currency pairs (e.g. correlations among pairs of currencies and other assets) to find profit opportunities. There are also – to return now to the more general level of automated trading – different styles of investing by which one can classify the use of trading with algorithms. If HFT is one style, statistical arbitrage is another. This is a term used for trading that spans longer time frames than HFT, but statistical arbitrage seeks to exploit mis-pricings

⁴ See the Foresight report [1] (pp. 42–43) for the summary on which my discussion is based.

⁵ Market makers provide liquidity to the market, meaning they are at least in principle obligated to buy or sell when a request is made. They attempt to profit from the difference between the buying and selling price (the bid–ask spread).

between correlated assets in a similar way. Quantitative trading based on proprietary models whose actual mechanisms are not disclosed is a third style. Some authors, like Johnson, reserve the term algorithmic trading for the execution of orders [4] (pp. 3–5).

Transformations

What little I have said about types of trading with algorithms and some of the issues this raises may at least suggest one thing – if the screens go black, or more realistically, if they go just half black, and the world that emerged with digitalization since the 1980s fades into history, then another world will rise in its place, one full of actors, business models, attempts at advantage taking, and sophisticated trading strategies. And while there are important differences between the former and the latter, markets will not end, and in some cases, as in the case of the Foreign Exchange Market, they may stealthily become larger and richer rather than smaller. The new world will be curated and monitored by humans. What’s the difference between a market curator and a trader? Or more broadly speaking, what is the sociology of the new world with its understructure of software agents that work, make decisions, and determine investment success?

And what happens to what I called the scopic regime when algorithmic traders become widely used? The first answer is that algorithmic traders and their effects have to be monitored and observed. Algorithms move underground, or rather, into the belly of computer networks, where computer programs do their work. They are software agents that traders don’t get to know face to face on trading floors or during happy hour at a bar. What they do will be known through the textual descriptions created for them, but in the market we will know them mainly through their effects. Traders may not trade anymore below a certain volume – employing their algo-tools for that purpose – but they surely will continue to monitor the market, the price and volume movements that result from human and algorithmic trading, and they will continue to monitor outcomes of software agents’ activities. Automated trading systems will follow their logic regardless of the outcome, as Mary Schapiro, then chair of the Securities and Exchange Commission, told a congressional committee in the wake of the flash crash of 2010.⁶ Human monitoring is needed to prevent this from happening, in addition to the circuit breakers that have been implemented in

⁶ As reported in [8]. See also [9].

exchanges to automatically halt trading when prices of stocks fluctuate by more than a certain percentage.

These may appear to be minor changes in trading practice, but they reflect a tectonic shift on a macro level that is less benevolent for those they affect. The most drastic change we witness today occurs on a professional level, and its most obvious effect is the historical replacement of one elite profession by another, in which the losers are human touch traders, and in exchanges, where the losers are voice brokers. The winners for now are professionals with a higher education and one or more degrees in the relevant disciplines: software engineering, mathematics, statistics, physics (physicists tend to be trained in both and know some informatics), and finance. The institutional winners are speed-focused trading firms, hedge funds, or proprietary desks and funds internal to banks, in which some of the new professionals work. To some degree, and in some areas, the displacement involved a power struggle; traditionally powerful and dominant exchanges, for example the New York Stock Exchange (NYSE), were slow to adapt to the new realities, and in the process lost volume and status. The NYSE, whose roots go back more than two hundred years, once held a preeminent market position at the apex of the world financial system, but it now accounts for only 22.5% of all stock trading. The replaced category of actors subsequently has moved into asset management, private trading groups and firms, or has been retrained. For example, there has been an exodus in recent years of traders from (investment) banks to proprietary trading firms and hedge funds, which have greatly increased in number. “Hundreds of talented traders” have left as banks shrank their trading operations in anticipation of the regulatory overhaul, the Volcker Plan, which would bar banks from such proprietary trading. “They are armed with skill sets”, this report continues, “that are in demand at hedge funds. Some are launching their own funds – a total of 549 new hedge funds were launched during the first six months of this year (2012) . . . while others are joining established ones . . . Rather than wait for banks to close their ‘prop desks,’ many star traders have gone into business for themselves over the last couple of years” [10]. The demand illustrated here may help explain the lack of a power struggle in some areas, but this is only part of the story. Many former traders, including those with technical and organizational skills, simply lost their jobs, and it is unclear where they went. The new skill sets expected of human participants in hedge funds often involve a combination of quantitative skills and programming knowledge as well as trade design and fund management skills. Here is the wording of a recent ad accessed on the web in May 2014:

A Hedge fund in Greenwich is seeking a Quantitative Research Analyst to join their Global Currency Portfolio Management team. This is a great opportunity to work on the cutting edge of quantitative investment & portfolio strategies for FX. Responsibilities will involve: 1) building trading strategies, tools & applications for global currency trading, 2) portfolio construction/optimization, and 3) FX/Currency exposure management. Applicants should have: 1+ years' experience working with economic and financial databases: very strong programming skills in Java, C++, SQL and Excel: a Ph.D. graduate or Ph.D. candidate from a Top School in Computer Science, Math or Financial Engineering: and superior communication skills. Currency strategy experience is strongly preferred.⁷

What about the non-human participants, our algorithmic agents? Their preferred qualifications are less mixed, as algorithms tend to be specialized. What they can do technically within their specialized functions already has been illustrated in the previous section. But how should these functions be executed, or, in other words, what properties do we look for in a trading algorithm? We won't find job descriptions for software agents on the web, but, luckily for us, computer science itself talks about desirable properties of algorithms and how these properties should be determined. Developing a solution to some problem typically involves at least four steps: "(1) designing an algorithm or step-by-step procedure for solving the problem, (2) analyzing the correctness and efficiency of the procedure, (3) implementing the procedure in some programming language, and (4) testing the implementation".⁸ Analysing the qualification of an algorithm in terms of its efficiency in the present context means, first and foremost, measuring its speed in executing tasks. Speed depends on several factors, for example:

- the size of the input ("searching through a list of length 1,000 takes longer than searching through a list of length 10")
- the algorithm type ("Unordered-Linear-Search is inherently slower than Binary-Search")
- the programming language used ("interpreted languages such as Basic are typically slower than compiled languages such as C++")

⁷ See www.simplyhired.com/job/foreign-exchange-trading-quantitative-analyst-ph-job/greenwich-based-hedge-fund-finance-industry/z6bw2hxuqg (accessed 12 May 2014).

⁸ For this and the following list and its explanation, see [11].

- the quality of the implementation (“good, tight code can be much faster than poor, sloppy code”)
- the speed of the computer executing the code.

In analysing the efficiency of an algorithm, Aslam and Fell [11] continue, one typically focuses on the speed of the algorithm as a function of the size of the input on which it is run, and one determines the number of program steps or some count of other computer operations as a function of the input size – the actual time, however, still also depends on the programming language used, the quality of the code produced, and the gigahertz (the speed) of the computer. In other words, an algorithm needs to live up to very particular and measurable criteria. Speed is just one of these criteria, though its importance in trading cannot be overestimated. Other criteria include precision (the more accurate the solution is, the better), memory (the less required, the better), optimization (some instructions, when repeated, may lead to learning and greater success in meeting a particular condition), and so on. If this sounds different from how we “size up” humans, even when they compete over speed, when they run a race for instance, then that’s because it is. Algorithms are not expected to have the same properties that humans do. Although they started out more or less imitating human behaviour in trading and are still advertised as being capable of doing exactly that (see [12, 13]) algorithms do think differently from “us”. “An algorithm is a well-ordered collection of unambiguous and effectively computable operations that when executed produces a result and halts in a finite amount of time“ [14] (p. 9) – I know of no account according to which a human being would fit this description.

Let me summarize, then, some of the differences between “us” and the algorithms that traders confront in the market, before ending with a brief outlook. The qualities of “them”, in contrast to those belonging to “us”, are taken from what humans dealing with algorithms attribute to them.⁹ Table 1 offers an overview.

⁹ These include my own interview sources and writings on algorithms by insiders in algorithmic trading. See for example [15].

Table 1. Different kinds of “subjectivities”

Human Actors	Software Agents
Emotion	Unemotional
Interpretation (taking role of other)	Scientific
Discretionary judgements	Dumb
Discipline	Seeing and reading
Self-regulation	Regulation via risk assessment
Research in delayed time	Research in real time
Division of labour	Division of labour
	Seductive (fast and cheap)

The first difference between humans and algorithms is that humans have emotions and that algorithms don't. Neurophysiologists think emotions are needed for intelligent human behaviour, but traders, taking their cue from the trading lore and culture, often maintain that emotions must be curbed, and that they can be controlled by “discipline”, a second characteristic and skill humans have.¹⁰ Trading culture recommends that traders be “disciplined”. A trader should take a view of how a stock or currency or other financial instrument will do within the relevant time period, and stick with it. The trader should liquidate a position when a price exceeds a pre-set limit, or when it gets too low, rather than wait and see if the investment will recover. One of the relevant stories that circulates among investors is about a famous trader who would leave the trading floor and go to the movies whenever markets were in turmoil – in order to maintain discipline and not react in impromptu ways. Discipline is recommended and taught to novices working on the floor to rein in their emotional responses. Algorithms, of course, don't need to have discipline, because they are disciplined by nature – programmed to do certain things and not others, and programmed to follow rules.

A third difference is that human beings are communicative actors using natural language rather than formalisms. Natural language words mostly have no fixed, unambiguous meaning. Meanings are pragmatic. They depend on context and the way words are used. Think of the fact that most dictionary entries have more than one listed meaning, and then consider the various ways

¹⁰ For an example of a neurophysiologist who thinks emotions are needed for intelligent human behaviour, see [16].

individual words can be used towards different ends. Indeed, users may give a word nearly any meaning they want when they speak metaphorically or poetically. Natural language utterances and writings require interpretation, a skill humans learn by talking, textual exegesis, and training. For example, humans grow up “taking the role of the other” in Mead’s famous phrase. They learn by putting themselves in others’ positions and seeing the world (and themselves) from other points of view. That helps with decoding and understanding others’ intentions. Humans effortlessly and without explicit awareness pay attention to gestures, reading them as additional information. And they use all of their senses and not just their vision or prefrontal cortex to “understand” what they see. As they grow into a particular language and culture, they acquire a huge repertoire of implicit knowledge that they can bring to bear on familiar and unfamiliar situations. Based on their understanding and experience, human traders, for instance, usually have control over their trading decisions, using their own judgement. They can make discretionary judgements. They are not bound by the data and calculations they receive from analysts and sources on their screens, and they usually are not bound by the view their employer, say a global bank, has of the market. Algorithms don’t make discretionary judgements, and they are programmed to respond to data. Things get murkier, however, when we ask whether algorithms have interpretive skills. The answer is that algorithms currently are learning to read and decode texts in order to interpret incoming news items correctly, but, as of today, they are still making substantial errors, for example, by taking messages literally and without taking into account circumstances and contextual clues. One recent example is a hacked Twitter message sent out by the news agency Associated Press on 23 April 2013: “Breaking (news): Two Explosions in the White House and Barack Obama is Injured”. The message was wrong, but it caused the Dow to plunge by 140 points, the dollar to yen exchange rate to fall, and a downturn in bond yields, all within seconds. “No human believed the story. Only the computers react to something that serious disseminated in such a way”, a Wall Street trading director said about the market’s reaction. Human traders, seeing the selloff, checked the tweet against other, more reliable news sources and concluded the tweet was incorrect, and stopped the market plunge (Associated Press also quickly discovered the error).¹¹

Human beings have a moral sense. They can regulate their own behaviour, and they tend to continually and effortlessly monitor rule compliance in others. The phenomenon that Wall Street

¹¹ For a fuller report, see the following article: www.dailymail.co.uk/news/article-2313652/AP-Twitter-hackers-break-news-White-House-explosions-injured-Obama.html#ixzz2dxXoCW3t (accessed 18 August 2013).

does not conform to or even understand the moral expectations of Main Street is no contradiction to this assertion. Wall Street, too, has its morals, but its moral rules have little to do with what the external world thinks and wants. Wall Street morality is part of a law of trading practice, a *lex mercatoria* enforced by traders during trading. It's a law oriented to managing orderly behaviour among insiders without real concern for outsiders. Internal self-regulation is needed in markets in which there are no centralized national exchanges. Currency markets are a case in point. As over-the-counter markets, they are not subject to the standards and disclosure requirements of exchanges, which are much more regulated. Like derivatives markets and financial markets generally, currency markets lobby for self-regulation, and governments have set them free – partly, perhaps, in response to lobbying, but surely also to prevent the problems fixed exchange rate regimes tend to incur. Algorithms can be programmed to stop doing what they do under specified circumstances, but they clearly don't have the "moral compass" we attribute to humans. In fact, algorithms tend to be controlled by risk management techniques rather than by a moral sense and ethical considerations.

There are more differences between algorithms and humans, as the trading side sees it, and one is that while we consider ourselves intelligent, and undoubtedly are, algorithms are intelligent and simultaneously dumb. They are capable of making informed decisions, but they will make them "scientifically", strictly on the basis of data they receive and the analysis they conduct or that other algorithms conduct for them. The downside of this is that they cannot think "outside the box", outside what they are programmed and optimized to do. Algorithms can learn, for example, by going through iterative cycles of runs, within a given framework, but they appear to be less adept at responding circumspectly and cleverly to new situations. They can also be imitated and "out-programmed" once a party interested in doing so understands what they are designed to do. Humans, of course, have major drawbacks, too. One is their limited speed, the dimension along which algorithms always will outperform human beings. Their speed allows algorithms to conduct research practically instantly, for instance, while human analysts or traders need a substantially longer period of time for similar results. Both humans and algorithms tend to specialize in certain tasks and enjoy a division of labour, though algorithms surely always were more specialized than humans. This difference may become more pronounced in the future, as the ad cited above suggests. One of the seductive features of algorithmic trading is, of course, that once an algorithm has been developed its continued operation costs less than the labour of a human trader. This is a major driving force for algorithmic trading.

Our Postsocial Future

Algorithms think and act quite differently from us, as we on the human side see it. Notions such as “thinking”, “intelligence”, and others mentioned previously in this essay, such as “seeing”, “reading”, and “acting”, suggest an agent. Are algorithms “tools”, a notion my interviewees assigned to them, or are they “software agents”, a term that denotes agency? The (financial) literature on algorithmic trading occasionally draws an analogy between algorithms and robots, and this analogy points to software agents. Herein lies another potential transition we confront – the transition from algos to agents, from what is merely a tool and instrument for human traders to traders that act on their own but are algorithmic in nature. The latter have to be reckoned with as strategic parties in human interactions, the former will do their work mainly by tooling us up. They are our “aides” with only instrumental rights and obligations. When we call these tools software agents, we elevate them to a higher status and suggest that they, too, are autonomous entities of a sort that can sense the environment and respond to it on the basis of specified goals. One way to draw the line between tools and software agents is to revert back to the distinction between algorithms that are the “bullet” and those that are “the finger on the trigger” – between those that merely execute trading decisions and those that make their own “decisions” on the basis of programmed analyses and models. One prediction one can make is that both types of algorithms will exist alongside each other in the future. The increasing volume of trading conducted by software agents poses interesting social science questions which will need to be explored. How do we have to imagine the interaction, on a microsociological level, between a software agent and a human agent? Software agents lack the “deep” play of humans – their generative ways of playing meanings, emotions, and existence. Relationships with these agents will not have the same quality as social relations. But we are likely to project some of our social relational expectations and fantasies onto them, and they may engage and attract us. I have used the term postsocial for object relations of this kind and have suggested that such unfolding, information-rich entities that lack completeness of being (like a market) may lure us into such relations [17]. The algorithm-based world that human traders and financial markets already inhabit, for instance, surely is a postsocial world. It’s trivial to say that algorithms lack soul. But it’s perhaps less trivial to find out what our postsocial relations with algorithms will look like in the future. When algos do the “touching”, as I envisaged in the beginning of this essay, when they trade according to decisions they themselves

make, they may also “touch” us – as competitors and antagonists, as counterparties, as objects of attachment.



References

1. The Government Office for Science: *Foresight: The Future of Computer Trading in Financial Markets. Final Project Report*. London (2012)
2. King, M.R., Rime, D.: The \$4 Trillion Question: What Explains FX Growth Since the 2007 Survey? *BIS Quarterly Review*. 27–41 (2010)
3. BIS (Bank for International Settlement): *High-Frequency Trading in the Foreign Exchange Market* (2011). Available at: www.bis.org/publ/mktc05.pdf (accessed 15 February 2012)
4. Johnson, B.: *Algorithmic Trading & DMA: An Introduction to Direct Access Trading Strategies*. 4Myeloma Press, London (2010)
5. Durbin, M.: *All About High Frequency Trading*. McGraw Hill, New York (2010)
6. See Mullins, B., Rothfeld, M., McGinty, T., Strasburg, J.: Traders Pay for an Early Peek at Key Data. *Wall Street Journal*, 12 June 2013. Available at: <http://online.wsj.com/article/SB10001424127887324682204578515963191421602.html> (accessed 15 June 2013)

7. See Popper, N.: High Speed Trading No Longer Hurling Forward. *New York Times*, 14 October 2012. Available at: www.nytimes.com/2012/10/15/business/with-profits-dropping-high-speed-trading-cools-down.html?pagewanted=all&_r=0 (accessed 16 October 2012)
8. Salmon, F., Stokes, J.: Algorithms Take Control of Wall Street. *Wired*, 27 December 2010. Available at: www.wired.com/magazine/2010/12/ff_ai_flashtrading/ (accessed 20 February 2013)
9. SEC (U.S. Securities and Exchange Commission) and the (CFTC) Commodity Futures Trading Commission: Findings Regarding the Market Events of May 6, 2010. Available at: www.sec.gov/news/studies/2010/marketevents-report.pdf (accessed 17 February 2011)
10. Shari, M.: The Volcker Generation. *Barron's*, 22 November 2012. Available at: <http://online.barrons.com/article/SB50001424052748703961304578129162892716642.html> (accessed 2 January 2013)
11. Aslam, J.A., Fell, H.: Analysis of Algorithms: Running Time. CSU200–Discrete Structures. Available at: www.ccs.neu.edu/course/csu200/05F/handouts/rt.pdf (accessed 28 August 2013)
12. Perlberg, S.: Now You Can Do Algorithmic Trading From Your Couch. *Business Insider*, 18 June 2013. Available at: www.businessinsider.com/now-you-can-do-algorithmic-trading-2013-6 (accessed 16 July 2013)
13. Cendrowski, S.: Quant Trading Comes to Main Street. *CNN Money*, 18 June 2013. Available at: <http://finance.fortune.cnn.com/2013/06/18/quant-trading-comes-to-main-street> (accessed 16 July 2013)
14. Schneider, G.M., Gersting, J.: *An Invitation to Computer Science*. West Group, New York (1995)
15. Aldridge, I.: *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*. John Wiley & Sons, Hoboken, NJ (2010)
16. Damasio, A.: *Descartes' Error: Emotion, Reason, and the Human Brain*. Vintage, London (1994)
17. Cetina Knorr, K.D., Bruegger, U.: Traders' Engagement with Markets: A Postsocial Relationship. *Theory, Culture & Society*. 19, 161–185 (2002)