AARHUS UNIVERSITY

Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the post-print version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Arnoldi, J. (2016). Computer Algorithms, Market Manipulation and the Institutionalization of High Frequency Trading. *Theory, Culture & Society*, 33(1), 29-52. doi:10.1177/0263276414566642

Publication metadata

Title:	Computer Algorithms, Market Manipulation and the Institutionalization of High Frequency Trading
Author(s):	Arnoldi, J.
Journal:	<i>Theory, Culture & Society</i> , 33(1), 29-52
DOI/Link:	http://journals.sagepub.com/doi/10.1177/0263276414566642
Document version:	Accepted manuscript (post-print)

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

**COMPUTER ALGORITHMS, MARKET MANIPULATION AND THE
INSTITUTIONALISATION OF HIGH FREQUENCY TRADING**

Abstract

The article discusses the use of algorithmic models in finance (algo or high frequency trading). Algo trading is widespread but also somewhat controversial in modern financial markets. It is a form of automated trading technology, which critics among other things claim can lead to market manipulation. Drawing on three cases, this article shows that manipulation also can happen in the reverse way, meaning that human traders attempt to make algorithms 'make mistakes' by 'misleading' them. These attempts to manipulate are very simple and immediately transparent to humans. Nevertheless, financial regulators increasingly penalise such attempts to manipulate algos. The article explains this as an institutionalisation of algo trading, a trading practice which is vulnerable enough to need regulatory protection.

Introduction

Is it possible to fool an algorithm and if yes, does that constitute punishable manipulation? Such questions have emerged as issues surrounding the use of executing algorithms (algos) in financial trading, most often referred to as algo trading. This article investigates cases of manipulation not with, but of, algos, meaning that algos are being ‘cheated’. The article furthermore describes how such instances only become describable as manipulation because the interpretations of regulatory statutes have shifted, the effect of which is that algos are granted protection. Thus the significant change is not only the use of algos, nor the manipulation of algos, but also the regulatory changes made to protect algos.

There exist many types of algos and they are being used with different purposes. One important purpose is using algos to place orders at a high frequency - many of which are limit orders that will never be executed. This type of trading (which is central for this article’s discussion of algos) is generally referred to as high frequency trading but high frequency trading is an ambiguous term and the distinction between high frequency and algo trading is fuzzy (see below). This article will generally use the term ‘algo’ to refer to executing algos used for high speed order processing.

Algo trading generally and high frequency trading specifically have attracted considerable controversy. Firstly, there are fears over the risks that algo trading poses to the financial system, essentially because use of algos means leaving much decision making to the algos with little possibility for human supervision and intervention (Cliff and Northrop, 2010). This risk is often discussed with reference to the ‘flash

crash', a sudden and dramatic market fluctuation happening in May of 2010, which in all likelihood was triggered by (an) algo(s) (MacKenzie, 2011). Secondly, there are discussions about issues of fairness because high frequency trading and innovations surrounding it at least according to critics of the technology create two tier markets due to differences in access to market information (Haldane, 2011; Lewis, 2014). Thirdly, use of algos generally, and use of algos for high frequency trading specifically, has raised discussions about the possibility of market manipulation (Biais and Woolley, 2011).

This article takes the latter of these three themes as its point of departure but it stands the subject of manipulation on its head. The topic of the article is not market manipulation by means of algorithms but rather manipulation of algos. That manipulation happens when human traders make algos 'take actions' that are not in the 'interest' of the algos (or to be precise, those operating them), instead benefitting the manipulating humans. Algo trading potentially represents one of the most dramatic shifts of intellectual labour from humans to machines (Cliff et al., 2010).

While the manipulation of algos in this article is discussed in terms of regulation, such manipulation is but one particular issue of regulation raised by algo trading. For a review of other regulatory issues of algo trading, in a similar theoretical framing, and leading to a in many ways similar conclusion, see Lenglet (2011).

In the decade before the turn of the century several social theorists began describing forms of, if not non-social, then 'less-social' interaction. In the context of finance Castells highlighted information processing technologies as a crucial economic driver and talked pessimistically about automatons, which would lead to a dominance of

‘social morphology’ over social action (Castells, 2000: 16). Knorr-Cetina in a series of publications on finance discussed ‘sociality with objects’ (Knorr-Cetina, 1997) and advocated a ‘new post-social theory’ (Knorr-Cetina, 2000) to understand such sociality. Also, ANT-inspired scholars such as MacKenzie (2003, 2004, 2006) studied the performative role of technology and pricing models in finance. Obviously, the use of algos in financial trading evokes all of these theories. But there are also aspects of algo trading that need a more critical gaze than what is afforded by these theories. Indeed, current usage of algos arguable exposes a blind spot in, if not all, then at least Knorr Cetina’s work. On the one hand, the use of algos in trading in many ways represents the coming true of the descriptions (or predictions) of Knorr Cetina’s theory. But on the other, regulative measures towards curbing manipulation of algos touch on an area, which Knorr-Cetina did not cover in the original theory, namely that of codified rules. For Knorr Cetina, ‘post-social’ denotes human engagement with objects. These objects are ‘unfolding structures of absence’ (Knorr-Cetina, 2000: 11). It is this temporality or dynamism that gives the ‘objects’ enough being for there to be a basis for some form of social interaction with them, yet as said above that interaction remains relative ‘flat’ or ‘shallow’ (devoid of deeper meaning). One example Knorr-Cetina uses is the interaction of a trader with ‘the market’ (Knorr-Cetina and Brügger, 2002). The market is conceived by the trader to be some form of social entity although it could also simply be seen as a series of market rates projected on a screen. While Knorr Cetina’s theorizing may at the time it was presented have seemed slightly contrived, reality has, I argue, caught up with the theory, in some ways validating it, in other ways superseding it. Now there really are objects - computerised algos – acting in the financial markets. It is now no longer a vague notion of ‘the market’ to which being or otherness can be projected but these algos. It

is to those post-social actors that some degree of agency, and, significantly, also regulatory protection, is attributed and granted.

It is relatively easy to see how and why Knorr-Cetina, coming from a tradition of phenomenology and symbolic interactionism, has made this theoretical move. She basically identifies two features that according to phenomenology and symbolic interactionism are essentially social, namely temporal processing of (inter)subjective meaning and generalization of otherness, as things that are in play when traders talk, describe and interpret 'the market' as a social other. But social interaction is not just upheld by projections of meaningful otherness and tacit norms but also by formal, institutionalised, and enforced rules. Likely as a result of her theoretical background, which tends to downplay the importance of such formal norms and regulations in social interaction, Knorr-Cetina largely ignores formal regulatory rules in her theorising of post-social interaction. But if formal institutionalised rules are ignored, it also becomes impossible to ask whose interests are furthered by any given formal rules and how such rules are legitimized. Thus, the post-social becomes strangely depoliticized. As part of its conclusion, this article attempts to reintroduce a theoretical basis for a more political view. That entails asking which rules and regulations are created and enforced. And when answering that question, it should be remembered that such formal institutional frameworks exert power not just by banning and punishing certain practices but also by simultaneously protecting other practices (and technologies).

What is algo trading?

Before defining algo trading, it may be worthwhile to briefly place algorithms in their general socio-technological context. Algorithms are finite sets of rules, specifying operations that can solve problems (Kronsjö, 1987: 1). Algorithms date back to Euclid but their contemporary importance is due to algorithms being coded as software whereby the algorithms' operations can be executed by computers. That should make the importance of algorithms clear for all: Algorithms are used for data mining in marketing research, search engines, autofocus systems in cameras, washing machines – anything involving devices containing a computer chip. The implications of this prolific technology have been theorised by social and media theorists, of which several arguments echo also in the current article. Firstly, algo trading is typical of what Fuller and Goffey (2012) call 'gray media'; technologies that are 'black boxed' and recessive. When reading reports about the events in the world's financial markets few presumably think about these events as the results of computers executing numerous trades. And the power dynamics of the financial markets have been fundamentally changed by algos and trading strategies using algos, yet this has been largely unnoticed by the general public. Secondly, algorithms profoundly change traditional ideas about perception and action. For example, algos that scan for arbitrage possibilities make such opportunities accessible to human traders who otherwise would not be able to perceive of these opportunities. Uricchio (2011) has argued that the mediating effect of computer algorithms ultimately threatens the Cartesian distinction between subject and object. The subject's representation of an object is increasingly a product of other objects. This applies to non-executing trading algos (see below). But changes in the relationship between subject and object also applies to executing algos (again, see below), albeit here not in regard to perception but to action. Algos that execute, that act, have a performative power much stronger

than any information merely encoded in ordinary language (Cheney-Lippold, 2011; Mackenzie and Vurdubakis, 2011). Much of the following deals with such executing algos.

There are many different types of algorithms used in the financial sector. They can be used for different trading strategies and process different types of information. Non-executing algos are merely ‘seeing tools’ (Beunza and Stark, 2004, Pryke, 2010) by means of which human traders can reduce complexity and frame and process information in higher quantities and/or faster than what the human brain can do. Final decision making and action taking is however still left to humans. Executing (or automated) algos on the other hand are computer software directly linked to the exchange so that trading is done ‘automatically’ based on the parameters of the algo. In such cases, algos execute trades, that is, feed information directly into electronic trading platforms. Such automated trades are conducted without active involvement, and only under supervision, of humans. Hence algos are also referred to as robots in the industry although they remain largely virtual robots in the shape of computer code.

Clearly, one advantage executing algos offer is fast information processing. That speed can be used for different purposes. One important purpose is using algos to place orders at a high frequency - many of which are limit orders that never are executed. This type of trading (which is central for this article’s discussion of algos) is generally referred to as high frequency trading although high frequency trading is an ambiguous term. Reflecting this, the US Commodity and Futures Trading Commission established an expert group with the mandate of defining high frequency

trading. Tellingly, they are expecting only be able to define different (sub) types. However, the working definition proposed by the Commodity and Futures Trading Commission focuses on ‘the use of computer programs or algorithms for automated decision making where order initiation, generating, routing, and execution are determined by the system without human direction for each individual trade or order’, which are electronically linked to exchanges with ‘extraordinarily high-speed order submission/cancellation/modification systems’ (CFTC 2012).

The advantage of algos is a combination of processing power and speed. For example, algos are often used to conduct arbitrage, that is, exploit (often miniscule) price differences of the same or, more often, different securities that statistically or theoretically should correlate in value but where there may be temporarily uncorrelated prices. Arbitrage was a dominant trading strategy also before algo trading but algo trading makes it possible to exploit miniscule price disparities occurring in time frames down to microseconds. Given that arbitrageurs’ exploitation of such differences tends to even out the same differences, speed is a key competitive advantage for such strategies. Also, more complex forms of arbitrage consist of identifying price correlations (or rather, momentary lack of such correlations) between different assets, something which only can be assessed - indeed can only be identified - based on complex calculations and vast amounts of historical data on previous price fluctuations (Pole, 2007). That information processing can only be done by means of powerful computers and algorithms.

Algos can also be used to not exploit price differences between two similar assets but to exploit upcoming changes in the value of one asset. This type of strategy, which

clearly falls under the category of high frequency trading, is controversial because one way of determining price changes is by monitoring, and responding to, the order flow at exchanges before that information is disseminated to other exchanges and/or market participants. This is possible if algos can monitor and respond to orders placed in the exchange system mili or micro seconds before others. This practice has spawned considerable debate because it is somewhat similar to front running, a kind of market manipulation were a broker would use his or her knowledge about customers' orders to take (price-impacting) positions in the market to his or her own advantage prior to executing the customer's order (Lewis, 2014). High frequency trading comprises a range of different trading strategies all made possible by the use of automated algos and high-speed electronic exchange systems. To this should be added that at least most of the controversial high frequency strategies hinge on having access to exchange order data. It is this data that many (but not all) algos process and respond to. But as we shall see, it is also this information, which can be manipulated in order to trigger certain responses from algos.

High frequency trading is all about reducing latency, which is the key word of the industry: Receiving information without delay; being ahead of the rest of the market; being able to exploit market opportunities before others have spotted them (MacKenzie et al., 2012). Significantly, 'before' in this context means milli, and now increasingly, microseconds ahead of other market participants. Because the time differences are so minuscule, there is a considerable competition between high frequency traders to get the fastest hardware in order to stay ahead. This has recently meant not only investment in specialised hardware but also in a de facto merger of hard and software: The latest innovation is processing chips with algos embedded in

the chips' circuitry. There is thus no longer software being run on a hardware platform but a seamless connection between the two. This reduces data processing time slightly – it reduces latency.

The most efficient way of gaining microseconds however is through proximity to, or co-location with, the servers of the exchanges as even the most sophisticated cables enhance latency. Hence co-location has become a standard service offered by the world's security exchanges. Trading firms can rent spaces adjacent to (meaning in the same building if not the same room as, or not right next to) an exchange's server. By doing that they avoid long cables between their own computers and the exchange server, which reduces latency. At the same time exchanges constantly upgrade their networks to offer more bandwidth.

Proximity trading, just like high frequency trading, raises concerns over a two-tier market with some participants getting information before others (MacKenzie et al., 2012). Andrew Haldane of the Bank of England has for example issued a stark warning about such two tiers with different access to market information (2011). So-called flash trading, a practice which has been offered by four major American exchanges including NYSE and NASDAQ (Nanvar and Harris, 2011), has attracted even more controversy and also lead the Securities and Exchange Commission (SEC) to propose that it be outlawed. Flash trading means ask and bid prices being offered - against payment of a fee to the exchange – to market participants 30 milliseconds (0.03 second) before they are routed to the official electronic platform and become visible for all market participants (and see the mentioning of front running above).

High frequency trading is thus a term used with reference to different trading strategies all using automated algos. As can be seen from the above, a crucial question concerning executing algos is what information they respond to, that is, what they base their ‘decisions’ on. By and large, algos are primarily responding to changes in the market, primarily (changes in) data on order flow. However, there are also algos that respond to news media reports on markets or other key sources of information. Such attempts have also spread into attempts to create so-called semantic algos. Here the aim is to create information interpreting models on a larger scale. It seems however clear that most algos currently are ‘tuned to’ changes in market volume or similar (Biais and Woolley, 2011). In what follows, I will generally use the terms ‘algos’ and ‘algo trading’ to refer to executing algos used for order processing based on order flow information. Most of these usages would generally be seen also as high frequency trading. I however only use that term in the contexts of specific persons or firms using automated algos monitoring order flow data running on machines co-located with exchange servers. The three cases and the discussion deal with automated algos monitoring order data flow but it cannot be determined if they are running on co-located machines. Thus the executing algos in the case studies may incorporate strategies with a focus other than latency reduction.

Algo trading generally improves the liquidity of the market and reduces the ‘spread’ (the difference between current bid/buy and ask/sell prices). But that also means that algo trading becomes a necessity for (those having the role of) market makers (liquidity providers) when operating in such high liquidity markets where the spread is small. The spread is the margin derived from buying at the bid and selling at the ask (the bid is therefore always lower than the ask). The ability to execute quickly reduces

the risk of the market having moved against the market maker in the time between having bought and being able to sell. Reducing that risk in turn means the margin does not have to be so big. As a result there is a huge competitive pressure towards high frequency trading. Exactly how much of trading on modern exchanges is high frequency trading is difficult to determine but most estimates are in the 50-70 % range for the major exchanges in the US and 30-50% for European exchanges. The impact of algo trading is clearly visible in that the average size of trades has decreased markedly while the number of orders has increased. There is some uncertainty about the profitability of algo trading (Kearns et al., 2010) but the general consensus is that algo trading creates a flow of individually small but, thanks to the high frequency, collectively significant profits, often with little risk-taking involved.

Three cases of layering and spoofing

Case study data

This article is based on three sources of data. The first and most important source is the electronic order book of a major European security exchange, from which two short time sequences are presented containing what has been identified by exchange overseers as manipulation by, but not least of, algos. In addition to these two small batches of trading data, a third case provides another example of the same type of manipulation as the second case, but this time from the US and on a larger scale with a much longer time frame. In addition, unstructured interviews were conducted with one high frequency trading trader (interviewee 3) who has executive experience working for one of the, if not the, leading high frequency trading firm in the US, two governmental regulators (of which one, interviewee 2, was interviewed twice over a

18 month span), and an official at a security exchange, interviewee 1 (who was also interviewed twice over a 18 month span). The order book data was provided by one of the interviewees. I was from the beginning attempting to generate data on new information inequalities created by high frequency trading. Therefore the specific subject of manipulation of models was a research topic that emerged from the interviews rather than being a preconceived research theme. After its emergence, however, this new topic came to structure later phases of the research.

As mentioned, the first two case studies are two small sequences from the order book. Case study data normally is of such a magnitude, and also analysed in such great depth, that they become larger narratives, often detailing developments over time (Marrelli, 2007). Here, only the third case possesses these traits. However, case study data can also be short spanned or ephemeral events such as one of ‘Goffman’s smiles’ (Thomas, 2011: 514). Leaving the selection of case material to an interviewee as was the case here obviously has potential pitfalls. However, the researcher not having selected case material may also eliminate the researcher’s personal biases. The present case material furthermore fulfills several key criteria for case study method. Firstly, discrete events are often used for case studies even though such events normally occur on a different time scale than what is the case in the first two cases where we are dealing with timeframes of maximum 19 centi seconds. Case material data is, secondly, well suited for falsification of theories (Flyvbjerg, 2006). In this study, the cases serve the purpose of showing that market manipulation not only is done with algo trading, which is a common assumption (Biais and Woolley, 2011), but that algos themselves are manipulated. Thirdly, although case study methodology is not by default linked to grounded theory, case studies often function well when no

hypotheses were formed prior to conducting research. In the current research, the overall research question was whether new trading technologies create new information asymmetries and unfairness in markets. Manipulation of algos was as mentioned above not initially an issue. Fourthly, case studies are not primarily concerned with issues of representativeness; indeed, they may be analysis of singular outliers (Thomas, 2011). That said, while the present cases may contain singular instances of manipulation, together with the interviews they present relatively strong evidence that the described manipulations are not isolated events. In line with case study methodology however, this article is not trying to answer whether they are often occurring and hence ‘large’ problems - here the aim is to demonstrate that such cases of manipulation exist and to draw out the implications of that existence.

Case one

Figure 1 is a spreadsheet with a sequence of orders from the electronic order book. As mentioned above, algo trading makes it possible to place a large number of limit orders, which only are executed if market prices reach the specified limit - meaning that many such limit orders never are executed. Nevertheless, other algos may respond to the limit orders. The first case concerns a trader (A) who has the intention of driving up the price of a security. A has placed several sell limit orders already in the order book at the asking price 113.25 (order 779) and 113.24 (order nos. 503, 898 and 980). These three orders are at the best asking price. Following this, A makes a large limit buy order at 113.25, which is three ticks (minimum price increments) over the current best bid. This order however means that A’s own smaller sell orders are automatically cancelled (as A cannot trade with A self). The order at 113.25 is in this

regard of no consequence but because the other three were at best asking price, the quantity of contracts offered at best asking price is reduced. As a result, the quantity of contracts at the best ask price is reduced to 343.

Insert fig. 1 around here

Another part of the order is executed but more than 50% (1732 lots) remain in the order book as a non-executed bid. This now constitutes the best bid price (at time 13:23:56.05). At the same time, the balance between bids and asks, which is a signal for other algos in the markets, has also changed. As other algos respond, bid price has increased two ticks and the spread decreased with one tick. As other market participants – read, algos – follow troupe, the price goes up

Case two

The second case concerns a trader, B, who intends to sell 800 lots of a security. In order to drive up demand and hence price, B also places a string of buy limit orders low enough that they will not be executed. Figure 2 is a screen shot from the electronic order book. B's orders are marked with a curly bracket B in Figure 2 with B's 800 lot sell order at the top of the bracket and B's limit buy orders below.

Insert figure 2 around here

While not being executed, these limit (buy) orders nevertheless trigger other algos, which respond with higher buy orders, which raises the prices and allows B to sell the 800 lot (see figure 3.) Algos here respond to changes in volume of (potential) demand, which serves B's purpose.

Insert figure 3 around here

Preliminary comments

The first thing that stands out is the time frame of the events. In the first case the time frame is 19 centi seconds, in the second the time frame is 6 centi seconds. These are relatively long periods by high frequency trading standards but nevertheless show that we are dealing with algos responding to (an) algo(s). And we see especially in the second case that algos respond almost instantly to market changes. Secondly, we see that algos respond extremely quickly to minor changes in the market, in the case of the second data sequence changes in the volume of orders. The attempts to manipulate the market are hence on a micro level both in terms of time frame (centi seconds) and market changes (a few ticks). It follows from this, thirdly, that the manipulation takes place within the (electronic) market, by means of placing limit orders that are never supposed to be executed. This is in stark contrast to some (but not all – see below) of the traditional forms of manipulation, which often centre on spreading misleading

information external to, but with relevance for, the markets. Such manipulation (for example spreading false rumours) is a process that both takes much more than centi seconds and which will induce much larger effects than an increase of two ticks or the spread narrowing one tick – it is in other words macro manipulation which stands in contrast to the micro manipulation seen in the cases. And as mentioned before, micro manipulation of algos has to rely on information transmitted within the exchange (because this is the information which algos process) while traditional manipulation could be done either through collusion or misinformation within the exchange or through information exchanged externally.

Case three

On December 18, 2012, the United States Security and Exchange Commission (SEC) published the proceedings on a case against the Canadian trading firm Biremis and the firm's co-founders and owners Peter Beck and Charles Kim. The broker-dealer registration of Biremis was revoked and Beck and Kim were both fined USD 250,000 and barred from trading on U.S. securities markets. The case against Biremis is extensive. Essentially, the SEC found that Biremis organised a large number of day traders (in excess of 4000) around the world through an intermediary firm named Opal Stone (financed by a fund controlled by Beck). A 'certain' (SEC, 2012: 8) number of the day traders engaged in spoofing and layering. SEC's ruling finds Biremis ultimately responsible for this conduct. The manipulation done by the day traders was happening over a long period of time. Generally, the manipulation done was layering similar to case two. SEC provides one specific example, concerning layering of the stock of Colonial Properties Trust listed at the NYSE:

That day, as of 11:00:58 a.m., the National Best Offer (“inside ask”) for CLP was \$19.19 and the National Best Bid (“inside bid”) was \$19.15. At that moment, the trader began placing orders to sell CLP shares. Over the next 31 seconds, the trader submitted 28 consecutive orders to NASDAQ (OUCH), NASDAQ (RASH), a dark pool operated by a broker-dealer, and NYSE, each to sell 100 shares of CLP at prices successively decreasing from \$19.21 to \$19.13. In addition, the trader periodically interspersed the 100-share orders with seven other sell orders submitted to NASDAQ (RASH), each for 14,000 shares at several cents above the prevailing inside ask. These 14,000-share “pressure orders” falsely signalled to the marketplace that – in addition to the 100-share sell orders near the prevailing inside ask – there was additional and substantial interest in selling the stock at price levels above the prevailing inside ask. When viewed in combination with other displayed orders, the “pressure orders” created the appearance of a liquidity imbalance, that is, a substantial difference between the quantities of shares demanded for purchase and the quantities offered for sale.

(SEC, 2012: 9)

One minute after the day trader had made the last sell order, he then placed buy orders at the now deflated price of USD 19.12. After that order was executed, he cancelled all the other sell orders, which were the ‘signals’ of liquidity imbalance that had driven the price down. The seller then repeats the layering, only in reverse:

Less than a second after submitting his last cancellation, the trader switched sides, conducting the manipulation in reverse. At 11:01:42 a.m., with the inside bid for CLP at \$19.14 and the inside ask at \$19.17, the trader submitted 57 buy orders, each for

100 shares, at prices that successively rose from \$19.08 to \$19.19 over the course of 47 seconds.

(SEC, 2012: 10)

These tricks are then repeated again several times during the day.

The third case shares many traits with the first two. The time frame here is however significantly longer, as the day traders filled in the orders manually. The most advanced technique they used was in fact a keyboard shortcut that enabled them to delete the limit orders quickly once they had moved the price. Nevertheless, we are still dealing with short time frames, and especially fast market responses to the limit orders. Also, we still are dealing with micromanipulation. The net gain from the two instances of layering was USD 266. Clearly, these manipulative actions have to be conducted continuously for it to be possible to accumulate significant gains.

Arguably, micro manipulation equals high frequency manipulation.

A final commonality is the most important one. In all three cases, the manipulation is, at least seen in retrospect, extremely transparent. Thus the question is who falls for such simple tricks? In regard to the third case, for whom could the limit orders, deviating significantly from the best ask price, create a perception of liquidity imbalance? The answer is of course algos. SEC states that:

Thus, this strategy creates a false picture of the pricing of, and/or demand for, a stock in order to induce, or trick, other market participants, often those using algorithmic trading platforms, to execute against the layering trader's bona fide order.

(SEC, 2012: 8)

Twice later in the proceedings, SEC (2012: 10-11) emphasises the use of high volume limit orders, which are deemed to have specifically targeted algos. SEC state that ‘the pressure orders created the appearance of a liquidity imbalance. Market participants that use trading algorithms often program their algorithms to react to such market signals.’ (SEC, 2012: 10). The interesting aspect is of course that these attempts to manipulate would be immediately transparent for humans. One mentioned piece of evidence against Biremis mentioned by SEC (2012: 12) was an internal memo describing day traders’ strategy to “make money from those stupid programmed orders’ (another term for orders submitted by algos). Only algos are ‘stupid’ enough to fall for such transparent tricks.

The reinterpretation of manipulation

The manipulation of ‘naïve’ algos (naïve is not used to imply that algos possess any kind of consciousness) is, I argue, the novel issue in all three cases. The three reported cases of manipulation can only occur if the market counterparts are algos - no human would fall for such simple tricks. This is something to which security exchanges have had to adapt to in the recent years, simply because it raises the question who is liable. Is it the manipulating traders or the traders who leave executions to algos? The following is from an interview with the head of the surveillance section of a major exchange:

The main point is - this is the tendency that we can see - given the fact that you have, at the end of the day, a machine which is acting, you lose the kind of human intention to look at things and be aware, 'this could be an issue', because you have a dumb machine more or less, and the machine reacts not like humans - let's face it like this - and it is from all points of view an issue in the market that some specialists on the other side found out that you can just manipulate machines much easier than human beings. Try to react in the market, for example, putting order in the market and so on. And what we have found out is that the actors in the market who try to manipulate those things, they have a new target - the machine.

[...] What we just can realize in the market is that high frequency trading, algo trading, is more the target of individuals trying to enable or to manipulate algo trading to do something, and not the way that algo trading is responsible for new things in the market, let's say, this be based on market manipulation or something. I think that is a clear fact that we can find just in the last 3-4 years.

[...] what we have realised in the last three years or four years [...] those who are acting as high frequency or algo traders are more the target.

While the cases may seem to be straightforward instances of manipulation, many of these, including the first two cases, are surrounded by ambiguity. Indeed, they can only be seen as manipulation if the definition of manipulation is shifted, however slightly. Reflecting this, the first case was after a review by the financial surveillance unit of the exchange deemed to be in a grey zone, which could not warrant criminal persecution or sanctions from the exchange. The second case was considered manipulation and did lead to sanctions from the exchange. The third case obviously

lead to sanctioning but it should be remembered that this was a case of large-scale organized and well coordinated manipulation running over a long period of time. The offenses committed by Biremis was not only that the firm facilitated layering but that they did not respond to warnings from the authorities and that they did not report suspicious behaviour, thereby violating a variety of financial and other legislation including the Patriot Act (containing rules attempting to hinder money laundering). Furthermore, Peter Beck had been previously fined by SEC. And rather bizarrely, another Beck owned company, Swift Trade, was fined for manipulation on the London Stock Exchange in 2011. The manipulation there was very similar to that of Biremis.

There are at least three reasons for the ambiguity. Before going into detail with these, it may be worth reminding ourselves that even codes for good market conduct that today seems self-evident are relatively recent constructs. For example, insider trading today arguably constitutes the most familiar and clear-cut transgression of such codes, it is broadly condemned, and heavily sanctioned. Yet, it was only after the 1929 crash that American authorities outlawed insider trading and only in the 1980s that insider trading became criminalized to the degree it is today, something which, it has been argued, happened as part of a political process driven by a host of different motives, not necessarily all economic (Joo, 2007).

Returning to the specifics of algo manipulation, the first source of ambiguity is that there exist a host of algo-based high frequency strategies that aim at ‘sniffing’ out hidden liquidity in the market. This is generally considered a legitimate strategy and is of a nature where many small (limit) orders are placed in order to get an idea of

where larger orders (often placed by executing algos) are hidden. Placing many limit orders, the main bulk of which never are executed, has become a common phenomenon due to high frequency trading. Trading firms having a 150:1 order-to-trade ratio is normal in today markets. This may raise suspicion, as an abundance of limit orders is the main indicator of layering happening. Exchanges therefore may inquire why that ratio is so high (interview data). And new regulation has been put into place to curb excessive limit order to trade ratios. But often high ratios are linked - or at least the traders are able to link – to explanations other than intentions to manipulate. An example from the same exchange as that of case one and two illustrates this: The order-to-trade ratio of an exchange member suddenly changed dramatically from 1:1 to 50:1. This brought the exchange to inquire why. The answer was that the trading firm had adopted a new strategy, which involved trading “at high frequency a series of intraday spreads between the current observed volatility and the theoretical mean value” (answer supplied by exchange member after inquiry by surveillance unit).

A second reason for why manipulation can be hard to define and establish is that motive and responsibility can easily disappear in the actor-networks of algo trading (Lenglet, 2011). If a regulator such as the one quoted above believes that it is algos that – to use the words of another interviewee – are ‘the victims’, proving intent (as opposed to negligence or mistake) is difficult when dealing with actor-networks consisting of, on the human side, programmers, hardware designers, traders, risk managers and on the non-human side algos and hardware.

Interviewee 2: This is sort of a dislocation of knowledge and responsibility.

You have a bunch of people who are in fact more or less involved in the operation of the machine, but it's difficult to find out one individual who is the really responsible person. This can be difficult in some situations.

In other words, the trader operating/overseeing an algo used to place orders in a way that constitutes layering might not be the person who designed the algorithm. Indeed the design of the algo might have been a joint collaboration between different persons from different firms.

The third reason for ambiguity is the most fundamental and most important one. Any regulation is based on some interpretation of what reasonably can be said to have tricked another into a false belief about the market, 'inducing the purchase or sale of such security by' (SEC, 2012: 3) that other. The SEC ruling in case three shows that the bar for what can reasonably be said to be misleading 'somebody' have had to be lowered out of consideration for algos that respond to very simple cues. But where exactly shall the bar then be placed? What can reasonably be expected of an algo? That question is not made less complicated by the fact that the answer must balance also the trading strategies such as the 'liquidity sniffing' just mentioned above, which might be rendered illegal if the bar is set too low.

Cases two and three indicate that in spite of the just mentioned ambiguities, the interpretation of the regulatory statutes have changed in order to accommodate algos. The SEC ruling is the clearest indicator of this but importantly it is not only SEC who has changed its regulative practices in recent years to protect algos. The following is

from the head of surveillance of a major European security exchange who describes a similar shift in Europe:

Interviewee 1: Five years ago, everybody would say: where's the problem? Why are you programming these stupid algos? It is your fault. But now it is market manipulation.

Thus, what five or ten years ago was so transparent that no one could claim to have been manipulated now is an offense by a manipulating party. It is worth noting that the changes in the definition of what constitutes manipulation just mentioned by Interviewee 1 are not formal changes of regulation - meaning that regulative or legal statutes are revised - but rather changes in the interpretations of said statutes. The same is visible in case three. The SEC verdict is made with reference to already existing legal code defining manipulation.

Transparent manipulation of course also raises the question why some traders in the first place take the risk of granting dumb machines the responsibility for trading large sums, a question which also was raised in an interview:

Interviewer: But given the assumption that people are generally not stupid, why would the people who are responsible for that machine not notice that (attempt to manipulate)?

Interviewee 1: But this is too expensive to some, sometimes. Because one reason to just change human beings into machines is the question of cost. This

is a very simple issue, from my point of view, and a simple example, but in all of the cases we have just focused on this manipulation of machines based on wrong impressions, order situations and so on in the market.

Leaving trading to ‘naïve’ algos may hence be a choice of economic necessity for high frequency traders. It nevertheless seems clear, but is hard to verify, that more sophisticated algos easily could be created and run. Crucially, the premium on that would be latency. The current technological developments strongly indicate that strategic priority is granted to speed rather than sophisticated data processing. Crudely put, algos get faster but not smarter. And when such fast, but not smart, algos dominate the market, the rules of the game must be changed accordingly.

Conclusion – the institutional underpinnings of computer code.

So far we have investigated an industry where, as a result of technological development, attempts are made by market participant to trick algos to act not to their own benefit or to the benefits of their owners and controllers, but to the benefit of the first mentioned market participants. Algo-human interaction constitutes a type of post-social interaction (Luca et al., 2011) with humans interacting and competing with, and trying to outsmart, computerised algos and vice versa. Seen from one angle, there arguably is little news in that; market participants have always tried to fool each other and such attempts have always run up against, and often transgressed, the boundaries of legality. That it now happens to machines where it before happened to humans would, by that line of reasoning, constitute only a minor change. The real significance, I have suggested, is the fact that the interpretations of regulatory code

defining market manipulation are changing. The machines are not only taking the place of actors, I have argued, they at the same time have been granted regulatory protection through changed interpretations of what constitutes manipulation.

That change in formal rules, in regulatory code, is something not foreseen or indeed included in Knorr-Cetina's theory of post-social interaction, which was developed based on observations of work in the financial industry. Referring back to the blind spot in Knorr-Cetina's theory alluded to in the introduction, we are now in a position to see that algo trading indeed is post-social, but this post-social order is, like any social order, upheld not only by humans attributing some degree of otherness to non-human objects and by co-temporality involving both human and non-humans, but is also upheld by regulatory institutions. Thus, the central argument of this article is that the changes in definitions of market manipulation described here – however tacit – are changes towards a regulatory order that can enable algo trading as an institutionalised practice. Algo trading (or at least several important forms of algo trading), it is suggested, would be too vulnerable a trading practice if the regulatory system was not supporting it.

Talking about institutions and institutionalised practices here is useful for at least two reasons. Firstly, organizational sociology has demonstrated how regulative institutions such as financial regulations are analytically distinct from, but in social reality intertwined with, other types of social rules and cognitive frames that in total help maintain stable forms of social interaction (Powell & DiMaggio, 1991; Scott, 2008). Thus, institutional theory creates a wider focal field including regulatory institutions but without excluding cognitive framings and other elements theorized by

Knorr-Cetina. Secondly, institutional theory also opens up for questions concerning how such frameworks are legitimised by certain bodies of knowledge and furthered by specific stakeholders.

That latter question is beyond the scope of this article. I would however like to indicate some likely entry points for an analysis of that question. Also, in closing, I would like to address very briefly another question, namely whether the relationship between post-social technologies and regulative institutions is relevant only for algo trading and finance or whether it applies also elsewhere.

As mentioned in the introduction, algo-trading is contested, viewed with skepticism both from outside but also contested within the financial industry itself with the financial press reporting heated debates and even fist fights (Metha, 2009). The recent book by Michael Lewis on high frequency trading (2014) purveys a strong sense of moral outrage towards high frequency trading where market participants find the markets dissolving on the screen as soon as they place their orders because algos react to the orders being placed before exchanges can process them. Lewis's book also indicates the dilemma these market participants face. Seeing market dynamics change, it is difficult for trading firms not to enter the arms race for speed. And finally, Lewis portrays the great power of the big Wall Street banks, most of which apparently support algo trading. Exchanges are for their part under considerable competitive pressure due to deregulation and electronic trading technologies, both of which threaten the exchange's status as the central locations for trade. The exchanges however find themselves with one valuable product due to the technological

development, namely proximity trading, and therefore the exchanges may for that reason want to protect and further algo trading.

On the most general level the great promise of algo-trading is arguable to improve a factor held by neoclassical economic theory to be fundamental to market efficiency namely liquidity. Algo-trading has continuously been claimed to improve liquidity in the markets both because it can process endless amount of orders quickly and because the role of liquidity provider is broadened so that it no longer is the monopoly of a few designated market makers (Interviewee 3). In fact, what unites neoclassical economic thought on market liquidity and algo trading is speed. At its most basic, liquidity signifies the ability to sell any asset quickly, without paying a premium for that speed. That speed is of course in turn hinging on the markets' ability to exchange information to and between sellers and buyers and their ability to respond to that information quickly. Another aspect of liquidity is resting on the ability of the market to quickly regain equilibrium once large and rate-changing trades are made. Again, high speed automated trading promises to help improve that ability, something that is furthered by the ability of algo trading to slice orders into smaller one that will not disturb the market as much.

The use of executing algos is a practice that enables some controversial trading strategies where speed is of immense importance. That speed is the center of controversy because it potentially creates unfairness in the market. The speed is however also a weakness of algos because they may be triggered into responding in ways not to the advantage of the algos' controllers. This article has presented evidence that the regulators are changing their definitions of market manipulation,

banning such triggering and thus in effect protecting algos. It is in this regard noteworthy that these changes in definitions seemingly are being made so that they do not affect different practices of liquidity sniffing often associated with high frequency trading. Thus, the changing of the definitions of manipulation should be seen as a regulatory institutionalisation of algo trading generally and high frequency trading practices specifically. That emerging regulatory institutional support re-invokes the argument that any economic order, also the most laissez faire, needs a supporting institutional architecture (Fligstein, 2001). And it reminds us that which architectures ultimately are created is less a question of market forces and more a question of political forces than theories of markets often hold. Algos being more vulnerable than what is commonly thought (as suggested in this article) would only reinforce these basic insights.

And finally, it may seem worthwhile to ask if the relations between computerized technology and institutional support applies only to finance or also elsewhere? Two strands of scholarly work document interplays between computer codes and legal and other regulatory codes also in realms outside of finance. The first strand is work on how existing (often legal) concepts such as ‘private’ or ‘authorship’ have been extended into the digital realm – no doubt best exemplified by the extension of (intellectual) property rights into cyberspace (Boyle, 1992, Street, 2003). A second, partially overlapping, strand of literature has argued that digital innovations generally, and software code specifically, are codes also in the sense of being able to shape human conduct. Thus, computer code is seen as performative, in that it creates specific types of social (but not necessarily human) beings (Introna, 2011, Mackenzie

and Vurdubakis, 2011) or as governmental (Thrift and French, 2002), because it shapes human conduct.

As was the case with algos, both of these strands describe shifts that are political and contested. They, just as was the case for algo trading and financial regulation, describe regulatory shifts more discursive than formal, more based on changing interpretations of legal code than formal changes of same code. Algos may of course have less of an impact on human conduct than some other computerized representations. But as Introna (2011) remarks, any translation, in this case an extension of agency to the algo, is also an unpredictable transformation, which may have unintended consequences. No doubt the attempts to fool algos have been such an unintended consequence, the response to which has been regulatory changes. Last but perhaps most importantly, the literature points to an interdependent development of computer and legal code. Lessig (2006) uses a helpful distinction between East and West Coast code, referring to the legal codes created in Washington D.C. and the computer codes created in Silicon Valley. Lessig has a keen eye on the conflicts and commensurabilities between these two sources of code. This article has tried (albeit on another scale) to show yet another arena where these two types of code are intersecting, supporting each other, and together reshaping society.

References

- Arnoldi, J. and Borch, C. (2007) 'Market Crowds between Imitation and Control,' *Theory, Culture and Society* 24(8):164-180.
- Beunza, D. and Stark, D. (2004) 'Tools of the trade: The socio-technology of arbitrage in a Wall Street trading room,' *Industrial and Corporate Change* 13(2):369-400.
- Beunza, D. and Stark, D. (2012) 'From dissonance to resonance: cognitive interdependence in quantitative finance,' *Economy and Society* 41(3):383-417.
- Biais, B. and Woolley, P. (2011) 'High Frequency Trading.' Unpublished manuscript.
- Boyle, J. (1992) 'A Theory of Law and Information: Copyright, Spleens, Blackmail, and Insider Trading,' *California Law Review* 80(6):1413-1540.
- Castells, M. (2000) 'Materials for an exploratory theory of the Network Society,' *British Journal of Sociology* 51(1):5-24.
- CFCT (2012) www.cftc.gov/ucm/groups/public/@aboutcftc/documents/file/hftdefinitionletter111711.pdf. (accessed January 2014).
- Chaboud, A., Chiquoine, B., Hjalmarsson, E. and Vega, C. (2009) *Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market*. New York: The Federal Reserve.
- Cliff, D., Brown, D. and Treleaven, P. (2010) *Technology trends in the financial markets: A 2020 vision*. London: The Government Office for Science.
- Cliff, D. and Northrop, L. (2010) *The global financial markets: an ultra-large-scale systems perspective*. London: The Government Office for Science.

- Fligstein, N. (2001) *The Architecture of Markets*. Princeton: Princeton University Press.
- Flyvbjerg, B. (2006) 'Five Misunderstandings About Case-Study Research,' *Qualitative Inquiry* 12(2):219-245.
- Fuller, M. and A. Goffey (2012) *Evil Media*. Cambridge: MIT Press.
- Haldane, A. (2011) 'The Race to Zero.' Speech published by Bank of England.
- Introna, L. D. (2011) 'The Enframing of Code: Agency, Originality and the Plagiarist,' *Theory, Culture & Society* 28(6):113-141.
- Joo, T. W. (2007) 'Legislation and Legitimation: Congress and Insider Trading in the 1980s,' *Indiana Law Journal* 82(1):575-622.
- Kearns, M., Kulesza, A. and Nevmyvaka, Y. (2010) 'Empirical Limitations on High Frequency Trading Profitability.' Unpublished paper.
- Keynes, J. M. (1936) *General Theory of Employment, Interest, and Money*. New York: Harcourt Brace.
- Knorr-Cetina, K. (1997) 'Sociality with Objects,' *Theory Culture & Society* 14(4):1-30.
- Knorr-Cetina, K. (2000) 'Postsocial Theory' in Ritzer, G., and Smart, B. (eds.) *Handbook of Social Theory*, London: Sage.
- Knorr-Cetina, K. and Brügger, U. (2002) 'Traders' Engagement with Markets : A Postsocial Relationship,' *Theory, Culture & Society* 19(5-6):161-186.
- Kronsjö, L. (1987). *Algorithms - Their Complexity and Efficiency*. Chichester: Wiley.
- Lenglet, M. (2011) 'Conflicting Codes and Codings : How Algorithmic Trading Is Reshaping Financial Regulation ' *Theory, Culture & Society* 28(6):44-66.
- Lessig, L. (2006) *Code - version 2.0*. New York: Basic Books.

- Lewis, M. (2014) *Flash Boys: A Wall Street Revolt*. New York: W. W. Norton & Company
- Luca, M. D., Szostek, C., Cartlidge, J. and Cliff, D. (2011) *Studies of Interactions Between Human Traders and Algorithmic Trading Systems*. London: The Government Office of Science.
- Mackenzie, A. and Vurdubakis, T. (2011) 'Codes and Codings in Crisis: Signification, Performativity and Excess,' *Theory, Culture & Society* 28(6):3-23.
- MacKenzie, D. (2003) 'Long-Term Capital Management and the sociology of arbitrage,' *Economy and Society* 32(3):349-380.
- MacKenzie, D. (2004) 'The big, bad wolf and the rational markets: portfolio insurance, the 1987 crash and the performativity of economics,' *Economy and Society* 33(3):303-334.
- MacKenzie, D. (2006) *An Engine, Not a Camera: How Financial Models Shape Markets*. Cambridge: MIT Press.
- MacKenzie, D. (2011) 'How to Make Money in Microseconds,' *London Review of Books* 33(10):16-18.
- MacKenzie, D., Beunza, D., Millo, Y. and Pardo-Guerra, J. P. (2012) 'Drilling Through the Allegheny Mountains: Liquidity, Materiality and High-Frequency Trading.' <http://eprints.lse.ac.uk/44908>.
- Marrelli, A. F. (2007) 'Collecting Data Through Case Studies,' *Performance Improvement* 46(7):39-44.
- Metha, N. (2009) Equities industry clashes over flash and step-up orders. July issue: *The Traders Magazine*.
- Nanvar, E. and Harris, L. (2011) 'The Economics of Flash Orders and Trading.'

- Pole, A. (2007) *Statistical Arbitrage: Algorithmic trading insights and Techniques*. Hoboken: Wiley.
- Powell, W. W. & Dimaggio, P. J. (1991) *The New Institutionalism in Organizational Analysis*. Chicago: Chicago University Press.
- Pryke, M. (2010) 'Money's eyes: the visual preparation of financial markets,' *Economy and Society* 39(4):427-459.
- SEC (2012) 'Securities Exchange act of 1934.' in Administrative proceeding File No. 3-15136. New York: SEC.
- Scott, R. W. (2008) *Institutions and Organizations*. Los Angeles: Sage.
- Street, P. (2003) 'Stabilizing flows in the legal field: illusions of permanence, intellectual property rights and the transnationalization of law,' *Global Networks* 3(1):7-28.
- Thomas, G. (2011) 'A Typology for the Case Study in Social Science Following a Review of Definition, Discourse, and Structure,' *Qualitative Enquiry* 17(6):511-21.
- Thrift, N. and French, S. (2002) 'The Automatic Production of Space,' *Transactions of the Institute of British Geographers* 27(3):309-335.
- Uricchio, W. (2011) 'The algorithmic turn: photosynth, augmented reality and the changing implications of the image,' *Visual Studies* 26(1): 25-35.
- Zigrand, J.-P., Cliff, D. and Hendershott, T. (2011) *Financial stability and computer based trading*, London: The Government Office of Science.

List of figures (attached):

Figure 1

Figure 2

Figure 3

Figure one: copy of orders from exchange

Time	Product	Bid Quant	Best Bid	Best Ask	Ask Quant	Maint. Code	B/S	Order Limit	Order no.
13:23:55.90		561	113.22	113.24	1000				
13:23:55.96		561	113.22	113.24	1015				
13:23:55.98		546	113.22	113.24	1015				
13:23:56.05						D	S	500 113.25	779
13:23:56.05						D	S	22 113.24	503
13:23:56.05						D	S	500 113.24	898
13:23:56.05						D	S	150 113.24	980
13:23:56.05		546	113.22	113.24	343				
13:23:56.05						A	B	3000 113.25	165
13:23:56.05						P	B	343 113.25	165
13:23:56.05						P	B	925 113.25	165
13:23:56.05		1732	113.25	113.26	1033				
13:23:56.06						P	B	10 113.25	165
13:23:56.06		1722	113.25	113.26	1023				
13:23:56.09						P	B	18 113.25	165

Figure two: Screenshot from electronic order book system

Timestamp 10:14:03.59									
Phase Trading									
Ind Price									
round lots									
Order Book									
Body	AccQty	Surplus	Limit	Surplus	AccQty	AskQty	MatchNo	Member	OnBehalf
01	01	01	241	54681	54681	17001			
01	01	01	23,991	37681	37681	2501			
01	01	01	23,951	35181	35181	2001			
01	01	01	23,911	33181	33181	2461			
01	01	01	23,851	30721	30721	5001			
01	01	01	23,611	25721	25721	3801			
01	01	01	23,591	21921	21921	8921			
01	01	01	23,551	13001	13001	3001			
01	01	01	23,531	10001	10001	2001			
01	01	01	23,51	8001	8001	8001			
***	1	1001	1001	23,471	01	01			
1	1001	2001	2001	23,461	01	01			
21	11001	13001	13001	23,451	01	01			
1	1001	14001	14001	23,441	01	01			
21	11001	25001	25001	23,431	01	01			
1	1001	26001	26001	23,421	01	01			
21	16001	42001	42001	23,411	01	01			
1	1001	43001	43001	23,4	01	01			
21	16001	59001	59001	23,391	01	01			
1	1001	60001	60001	23,381	01	01			
1	1001	61001	61001	23,371	01	01			
Quantity Limit Act EXR TrR OT MatchNo Member									
Add	Buy	1,000.00	23.450	A	-	-	L		

[illegible]