Manipulation in a Simple Stock Market

Econophysics

Nicolò Gozzi

Fisica dei Sistemi Complessi
Università degli Studi di Torino

Manipulation in a Simple Stock Market

Econophysics

Nicolò Gozzi

Abstract

The aim of this work is to analyse the behavior of trading agents in a simple stock market in which benchmarks for risk exposure (in our case an index that estimates volatility) are manipulated. The idea comes from recent events: at the beginning of February 2018, large fluctuations of the Cboe Volatility Index (VIX), whose aim is to estimate the implied volatility of S&P500 stocks, caused a flash-crash in the US stock market. Furthermore, it is emerging now that the VIX (or the "fear gauge", as it is called) has been manipulated by a group of individuals who were trying to keep it low in order to speculate on volatility. We investigate also the role of Algorithmic Trading in a crisis situation: automatic reaction to a sudden fluctuation in the market can amplify or even trigger a crash. To build the model we use NetLogo, an open source agent-based programmable modeling environment.

Contents

Chapter 1

Introduction

1.1 Cboe Volatility Index

To understand what the VIX Index is, we can look directly at the definition given by the Chicago Board Options Exchange, the largest U.S. options exchange which calculates and publishes it:

> The Cboe Volatility Index (VIX) is a key measure of market expectations of near-term volatility conveyed by S&P500 stock index option prices. Since its introduction in 1993, the VIX Index has been considered by many to be the world's premier barometer of investor sentiment and market volatility. Several investors expressed interest in trading instruments related to the market's expectation of future volatility, and so VIX futures were introduced in 2004, and VIX options were introduced in 2006. Options and futures on volatility indexes are available for investors who wish to explore the use of instruments that might have the potential to diversify portfolios in times of market stress.

In other words, the VIX is the expected range of movement in the S&P500 index over the next period. It can be considered then as a measure of risk of a financial transaction: investors would take greater risks if the volatility of the market is expected to be low, otherwise they would be more prudent if it is expected to be high. However, recent events showed a serious bug in this system: it seems that through clever and complex techniques someone succeeded in manipulating the VIX without even exposing on the market. The purpose of this work is certainly not to explain how the VIX has been - or better, can be - manipulated, however, the main idea behind this complicated and obscure procedure is that by only submitting targeted orders for S&P500 options during pre-open auction period it is possible to drive up or down the VIX. A paper published in May 2017 by Griffin and Shams of the University of Texas at Austin [1], detailed "interesting" patterns in trading related to the VIX and tries to explain these methods.

But why someone would manipulate the VIX? As said above, options and futures on volatility are available since 2004. This means that investors with positions in derivatives markets could have interest in controlling the VIX, since basically they are "betting" on volatility. Of course market manipulation is illegal, and of course we are not saying that every single investor who is betting on VIX is also trying to manipulate it. However, in a period of low interest rate and high liquidity in the financial markets, many investors tried to implement "exotic" strategies to generate profit, and speculating on volatility is one of them.

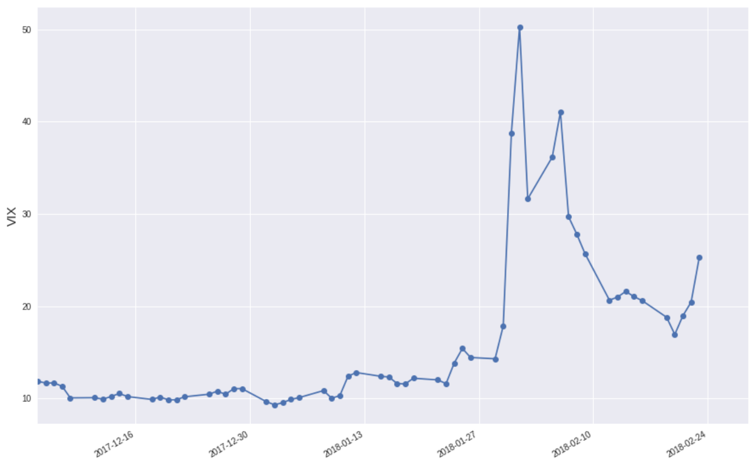

Figure 1.1: **Evolution of the VIX Index over the period Dec 2017 - Mar 2018**. At the beginning of February 2018 the index jumped suddenly from around 20 points to over 50.

1.2 Algorithmic Trading and Flash Crashes

With Algorithmic Trading we refer to a method of executing financial transactions using automated algorithms accounting variable such as time, price and volume. This *black box trading* is considered a great opportunity to make profit, since it combines computational power to consolidated trading strategies and to modern machine learning models. However, behind this motivated confidence towards algorithmic trading can hide dangers and pitfalls.

An example of how automated trading strategies can turn from solution to problem is represented by the United States trillion-dollar stock market crash of May 6, 2010, known as the Flash Crash. On that occasion at 2:32 p.m. EDT, stock indexes such as S&P500, Dow Jones Industrial Average and Nasdaq Composite collapsed and rebounded within 36 minutes (Figure 1.2). Financial markets are known to be volatile, of course. But how a crash of this magnitude can occur and recover within 36 minutes apparently without no reason? The answer can sound a little bit creepy. As pointed out in the paper "Automation, Intermediation and The Flash Crash" [2] published in February 2018, the crash was mainly caused by a large sell order of 75,000 E-mini stock index futures contract generated by a high frequency trading software (HFT). Besides this, the massive presence of HFT algorithms in the market contributed to a cascade effect with dramatic consequences.

Events like this raised questions and concerns about the structure of electronic markets: should we introduce limitations to HFT? Should we introduce strict control in order to block immediately panic selling situations and so avoid crashes like the Flash Crash of May 2010? We have to keep well in mind this questions. In fact, the "VIX crash" of February 2018 was amplified by operations generated by HFT algorithms: it seems that the explosion of the bubble of the VIX, which probably was manipulated by speculators, caused massive sell orders especially generated by automated systems and thus the crash.

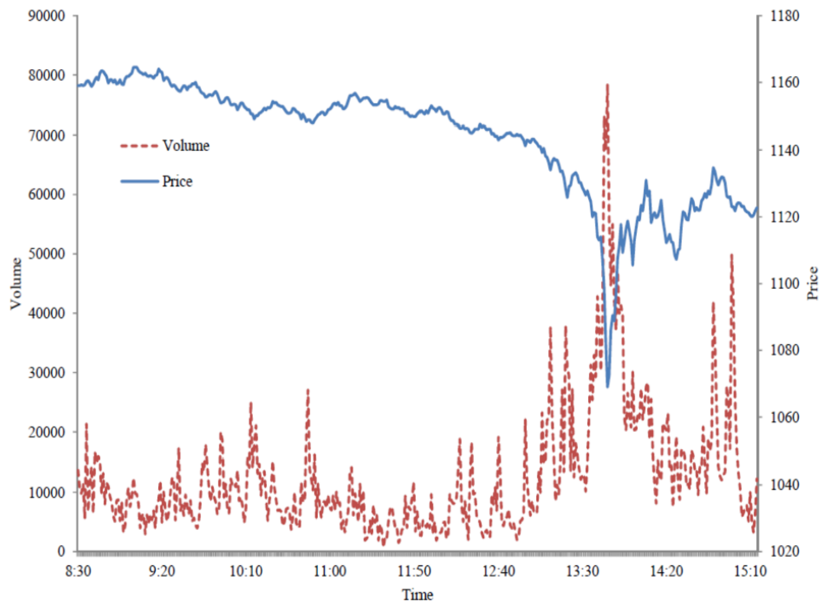

Figure 1.2: **Prices and trading volume of the E-mini S&P500 stock index futures contract during the Flash Crash**. Source: Kirilenko, Kyle, Samadi, and Tuzun (2017)

Chapter 2

The Model

2.1 Continuous Double Auction Market

In our model we consider a single stock market in which agents behave more or less randomly. At each time step each agent is asked about what it wants to do: it can either buy or sell (or even neither of the two). This is substantially the reason why we call the model *Continuous Double Auction Market* (CDA) [3]. As said above, only one stock is traded between agents. Its price is basically driven up or down by the requests of the agents. In fact each agent has a variable *price*, that records the price at which the agent is willing to buy or sell, according to its decision. This variable has a stochastic evolution since at every time step it is corrected of a gaussian factor sampled from $\mathcal{N}(0{,}100)$. Then, decisions of the agents (buy or sell) and prices related to these decisions are recorded in two "books": like in a real electronic market, buy orders are collected and ordered from the highest to the lowest price in a list, while sell orders are collected in a separated list and ordered from the lowest to the highest price. If the first two orders in the two lists match (in other words, if the first sell price is equal or lower to the first buy price) the transaction is executed and the two orders are cancelled from lists. So, to sum up, if there are more agents buying in the market, the price of the unique stock will naturally tend to rise, since higher buy prices in the book become more probable; otherwise, if there are more agents selling, the price will go the opposite way.

We still have two points to discuss in order to describe the CDA Market. The first one is that, in order to prevent prices to become negative, we facilitate the opening of buy positions when the price is "too" low. This is not an absurd assumption. In fact, in real world situations, it is quite reasonable to think that when the price of a stock or of a commodity is low, there will be someone disposed to buy it (especially if he or she thinks that the price will rise in the future). The second point is regarding the buy orders and the sell orders book. In our model we can either decide to clean the two books after each trading period or we can decide not to do this. For our intentions this is not a key point, so, if not specified otherwise, we will clean lists after each period.

2.2 Volatility Index

We suppose there exists a method to keep low the volatility index. So we fix the initial value of the index to 15 and we let it oscillate in a range $\pm 2\%$. Then, with a certain probability (that we can control), the index can jump suddenly to high values. This is not that far from what happened at the beginning of February 2018, when the VIX spiked over 50 points. Another parameter related to the volatility index is the length of the crisis. If not specified differently, it will be set at 30 ticks.

2.3 Agents

We consider two breeds of agents: *random traders* and *algorithmic traders*. **Random traders** behave randomly. At every time step each random trader:

- according to the probability of inaction (a parameter that we can control and we leave at 0.2) decides either to pass or to open a position;
- if it has decided to open a position, it evaluates either to go buy or sell.

The decision between buy and sell is taken substantially tossing a coin: each agent has a *threshold* which represents the probability of opening a buy position and is set to 0.5. This threshold is modified when:

- the price of the unique stock is too low, in this case the threshold is set to 0.8, in order to encourage the opening of buy positions and so avoid negative prices;
- the volatility index is low, in this case traders are more confident so to the threshold is added a small bonus of $+0.01$.

Algorithmic traders do not behave completely randomly. When the volatility index is low they act exactly like random traders. However, when the index spikes over 50 points they start to sell with probability 1. This behaviour is similar to the one observed in the Flash Crash: sudden change in benchmark parameters like volatility can cause an immediate reaction in automated trading systems with the danger of cascade effects and contagion to the whole market.

2.4 Simulation

As said above, we realized the simulation in NetLogo (while data has been analysed and plotted with Matplotlib, a graphic Python library). Through the control panel we can manage the parameters of the simulation:

- the *lengthSimulation* slider controls the length of the simulation (expressed in ticks);
- the *bubbleProbability* slider controls the probability of the explosion of the bubble of the volatility index. In other words it represents the probability with which the volatility index spikes over 50 points;

Figure 2.1: **Netlogo control panel.**

- the *bubbleLength* slider controls the length of the bubble, or in other words the time that needs for the volatility index to reach again low values;
- the *passLevel* slider controls the probability of inaction of all agents;
- the *nRandomTraders* slider controls the total number of trading agents (it is called nRandomTraders beacuse initially all agents have this breed);
- the *nAlgoTraders* slider controls the number of RandomTraders that are turned into Algorithmic Traders;
- the *floorActing* switch controls the possibility to manipulate or not the threshold of agents when price is low in order to avoid negative prices;
- the *cleanTheLogs* switch controls the possibility to clean or not the *buy* and *sell* books at the end of each cycle;
- the *setup* button prepare the simulation initializing agents and parameters;
- the *go* button starts the simulation.

Chapter 3

Results

3.1 Simulations and Analysis of Market Trends

3.1.1 50 Algorithmic Traders

We start considering 50 Algorithmic Traders out of 100 total trading agents. We notice that, when the value of the volatility index is low, the market is generally *bullish* (Figure 3.1). To see better the effect of Algorithmic Traders on the market, we focus our attention on the bubble explosion. It is interesting to notice that the reaction to the spike of the volatility index is immediate and the price is driven down towards the *floorActing* level rapidly. Furthermore when the value of the volatility index is high, the length of the selling queue stabilizes around the number of Algorithmic Traders (Figure 3.2).

3.1.2 20 Algorithmic Traders

Now we consider 20 Algorithmic Traders out of 100 total trading agents. Again, when the volatility is low, since agents are more confident, the trend of the market is generally *bullish* (Figure 3.3). If we focus on the explosion of the volatility bubble, we notice again that the reaction is immediate and drives down radipdly the price towards the *floorActing* level (Figure 3.4).

3.1.3 5 Algorithmic Traders

Finally we consider only 5 Algorithmic Traders out of 100 total trading agents. Quite surprisingly, it seems that, even if their effect is limited, they still influence the market (Figure 3.6). In fact, if we focus on the bubble part, we see that the trend is *bearish*, even though the price does not reach *floorActing* level.

Figure 3.1: **50 Algorithmic Traders**: OHLC chart of price, Volatility Index evolution and Length of the selling queue.

Figure 3.2: **50 Algorithmic Traders - Focus on Bubble**: OHLC chart of price, Volatility Index evolution and Length of the selling queue.

Figure 3.3: **20 Algorithmic Traders**: OHLC chart of price, Volatility Index evolution and Length of the selling queue.

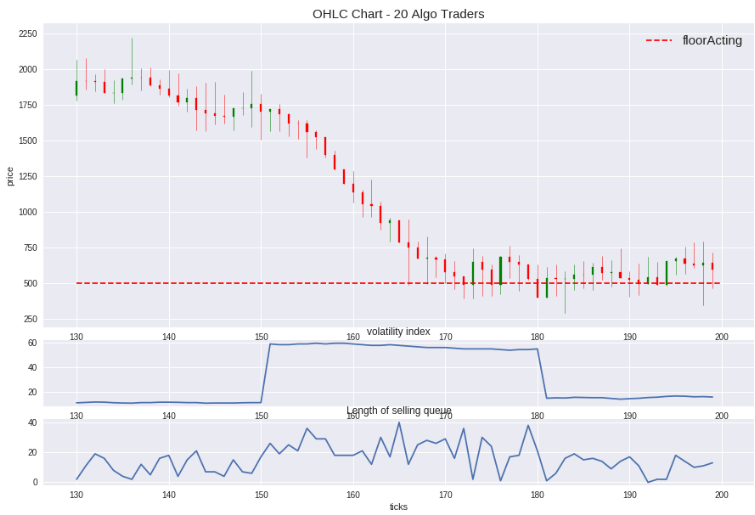

Figure 3.4: **20 Algorithmic Traders - Focus on Bubble**: OHLC chart of price, Volatility Index evolution and Length of the selling queue.

Figure 3.5: **5 Algorithmic Traders**: OHLC chart of price, Volatility Index evolution and Length of the selling queue.

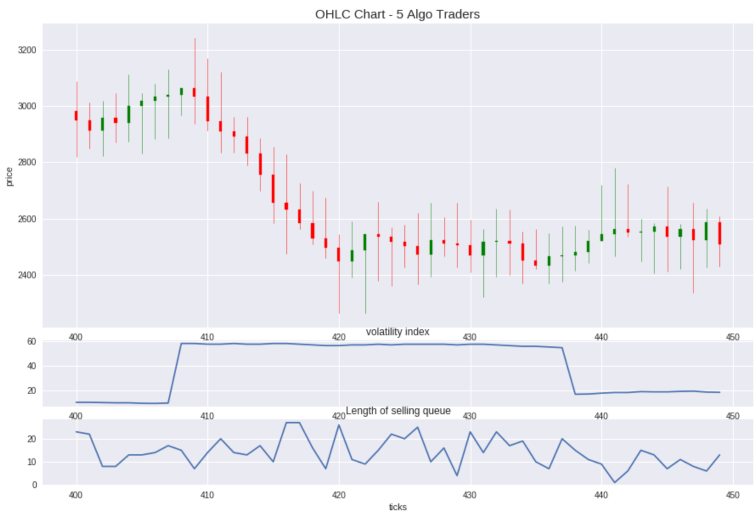

Figure 3.6: **5 Algorithmic Traders - Focus on Bubble**: OHLC chart of price, Volatility Index evolution and Length of the selling queue.

3.2 Impact of Algorithmic Trading

Now we want to quantify the effect of Algorithmic Trading on our simple market. In order to give a measure of the impact we define the quantity:

$$\Delta = startPrice \text{ - } endPrice$$

Where *startPrice* is the price when the volatility index jumps to high values and *endPrice* is the price when the index returns to low values. Then, if Δ is positive, it means that the price downed during the bubble explosion, if it is negative it means that the price rised.

We perform different simulations in which at every 100 ticks the volatility index jumps over 50 points and remains high for 30 ticks. Each time we record the value of Δ. We simulate this for 5, 20, 35, 50 and 65 Algorithmic Traders and we collect 100 values of Δ for each simulation. Then we evaluate mean and standard deviation so that at end we have 5 points and each one represents the mean Δ with the related error for 5, 20, 35, 50 and 65 Algorithmic Traders. From Figure 3.7 we see that all Δ are positive, so actually also the presence of only 5 Algorithmic affects the market and cause a deflection in the price. Furthermore, as it reasonable, the effect is increasing with the number of Algorithmic Traders. However, the mean Δ from 20 Algorithmic Traders on remains, with good approximation, constant. In other words the impact of 20 Algorithmic Traders it is more or less the same of 65 Algorithmic Traders.

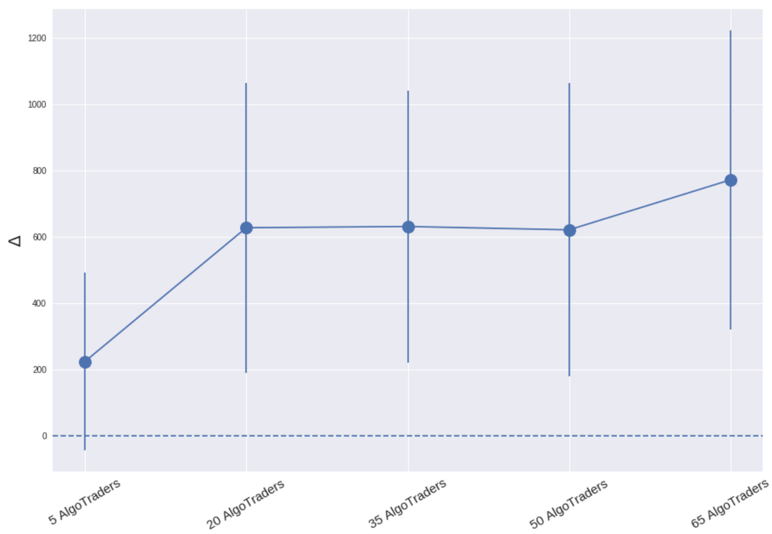

Figure 3.7: **Mean Δ**

3.3 Market Supervision

Now we suppose there exists an external agent that supervises the market. In particular, it is capable of detecting anomalous situations of panic selling and it tries to block them or at least to mitigate their effects. Technically we repeat the simulations of the previous section with the additional fact that after 10 ticks from the explosion of the bubble shorting becomes prohibited. This means that an agent can not open a sell position without having any stock. This is not far from what happens in the real world, where authorities, like the Consob in Italy, supervise operations in the stock market and manage irregularities. The time lag of 10 ticks is introduced mainly because it is quite reasonable to think that a certain time is needed to detect atypical situations and to intervene. We expect to see an attenuation of the effects caused by the presence of Algorithmic Traders, since they are more likely to be blocked by the external authority. This is actually what happens, as we can see in Figure 3.8: for 20, 35, 50 and 65 Algorithmic Traders mean Δ are smaller (this means that drops in prices are smaller) while for 5 Algorithmic Traders effects are completely vanished (in this case Δ is negative, this means that the price tend to rise also when the volatility index is high).

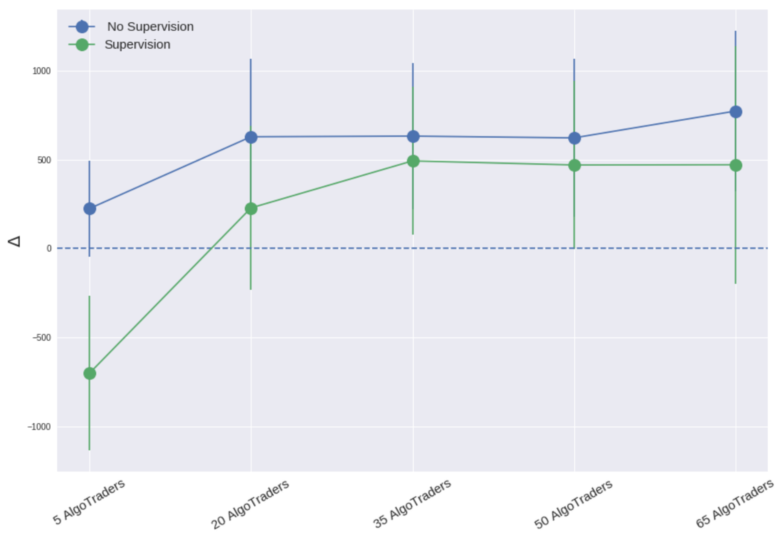

Figure 3.8: **Mean Δ with and without Supervision**

Chapter 4

Conclusions

Our goal was to investigate the effects on the market of manipulation of benchmark parameters for risk exposure like volatility. Through an agent-based approach we managed in reproducing the stock market environment with relatively simple assumptions. Then we have shown that, even if only a small fraction of the market is represented by automated trading systems, sudden change can cause reaction and cascade effects that may lead to crashes. Furthermore, we saw that the impact of Algorithmic Traders saturate quite rapidly: in other words, the effect caused by 20 Algorithmic Traders out of 100 total agents is already big enough to cause a crash and it is comparable to the effect caused by 65 Algorithmic Traders. Finally, we have shown how market supervision can mitigate these effects, in particular detecting and blocking risky operations like shorting.

Bibliography

[1] Griffin, John M. and Shams, Amin, Manipulation in the VIX? (May 23, 2017). Available at SSRN: https://ssrn.com/abstract=2972979 or http://dx.doi.org/10.2139/ssrn.2972979

[2] Kirilenko, Andrei A. and Kyle, Albert S. and Samadi, Mehrdad and Tuzun, Tugkan, Automation, Intermediation and the Flash Crash (February 6, 2018). Journal of Investment Management, Forthcoming. Available at SSRN: https://ssrn.com/abstract=3119363 or http://dx.doi.org/10.2139/ssrn.3119363

[3] P. Terna, Continuous double auction basic model (2018). Available at: http://terna.to.it/econophysics18/NetLogo_examples/CDA_basic_model.nlogo