The impact of algorithmic trading in a simulated asset market

Timothy Walsh, Bo Xiong, Christine Chung
Department of Computer Science

Purba Mukerji
Department of Economics

Connecticut College
{twalsh, bxiong, cchung, pmukerji}@conncoll.edu

Abstract. Algorithmic trading (AT) in asset markets has risen in importance over the past decade, revolutionizing the way market transactions are conducted. The extant empirical literature provides sometimes contradictory results on the impact of AT on market quality parameters such as liquidity and volatility. In this work we create a computer simulated asset market in order to make the effects of algorithmic trading more transparent to observers. Our market consists of human and algorithmic counterparts of traders that trade based on technical analysis, fundamental analysis and pairs trading strategies. The difference between human traders and algorithmic traders is in the speed of trading. Our preliminary results indicate that liquidity improves as the share of algorithmic traders increases, while no clear trend emerges for volatility. Pairs trading leads to an decrease in correlation between price of a stock and its fundamental value.

Introduction

Algorithmic trading (AT) in asset markets has risen in importance over the past decade, revolutionizing the way market transactions are conducted. While the underlying strategies used by algorithms to transact assets might often follow the same logic as human traders, two important differences set them apart. First, algorithmic trading can take into account much more information than a human trader alone could process in the time needed to make a given decision to transact. Second, while human traders may make transaction decisions based primarily on a particular rule, just as an algorithm would, in addition to this rule humans intuitively defer to an amorphous set of parameters, which taken together, we think of as human judgment.

It is harder to program all elements of human judgment into trading algorithms, therefore in practice firms use a mix of algorithms and human judgment (SEC report on flash crash). Even in fully automated transactions, where not only the actual timing and carrying out of the transaction, but even the decision to transact is driven by algorithm, there are pauses built into the system such that if there is any doubt about the quality of data or price or quantity of transactions in the marketplace, the system stops for human input to move forward. This was illustrated rather dramatically during the “flash crash” of May 2010 when many traders withdrew from the markets as their automated systems paused for human input due to unusual market fluctuations.

Empirical analysis of algorithmic trading (and a subset called high frequency trading), have estimated its impact on various aspects of market quality such as liquidity, price spreads, the extent of adverse selection, trade-related price discovery, and the volatility of asset prices. While some of the findings are at odds with each other, there is a clear consensus that the impact of algorithmic trading has significantly influenced all aspects of market quality listed above.

The contradictory results in the empirical literature may be the result of differences in methodologies, time periods and samples of assets analyzed. Even in the same work, there are often differences in the degree to which the impact of algorithmic trading is felt among different types of assets, and transactions. In addition, for each empirical finding there are different explanations that could produce the observed impact of algorithmic trading.

These factors indicate that it would be useful to go a step beyond measuring AT’s impact and attempt to unravel *how* this impact is actually created by the operation of different strategies in AT. Toward this ultimate goal, in this work, we create a computer simulation of the asset market. Our computer simulation gives us the capability to characterize the impact of AT on market quality and test its sensitivity to changing situations such as the volume of algorithmic versus human trading. Since proprietary trading algorithms are not public information, we make simplifying assumptions about the traders that allow us to approximate their behavior. Naturally the algorithms that actually operate in the marketplace have developed nuances that this exercise does not capture, however we expect that our results still provide a useful estimation of overarching and longer-term market outcomes.

We limit our model to two types of basic traders: fundamental analysts and technical analysts. Both types of traders can be either *human* or *algorithmic*. We simulate the behavior of simplified forms of these two types of traders. We then test the liquidity and volatility of the market as the percentage of algorithmic traders (of both types) increases. As a measure of liquidity, we use effective half-spread (as in Hendershott and Riordan, 2011), and as a measure of volatility we use variance in price (as in Zhang, 2010). Our results indicate that liquidity improves as the proportion of algorithmic traders increase, while volatility follows no clear trend. Finally, we ask what happens when algorithmic pairs traders, who trade two stocks based on the belief that their prices are correlated, are mistakenly “pairing” two unrelated stocks. Here we are motivated by the idea that, “an algo doesn't know or care why two assets are moving together; it merely is programmed to recognize that they are doing so” (Zweig, 2010). Our results indicate that if pairs traders begin trading unrelated stocks, their prices will begin to “artificially” or incorrectly move together, even as their underlying fundamentals diverge.

Background

The literature in the area of algorithmic trading is vast relative to its recent vintage (for example see Biasis and Woolley, 2011, for background and literature survey). In this section we briefly discuss some studies that are most closely related to ours.

There are a few studies whose results directly inspire the present work. Hendershott and Riordan (2011) and Hendershott et. al. (2011) analyze NYSE electronic message traffic and conclude that as AT grows, liquidity improves. Zhang (2010) analyzes a sample of firms from the CRSP and the Thomson Reuteres Institutional Holdings database from 1985-2009 and concludes that high frequency trading (HFT) is positively correlated with market volatility. Brogaard (2010) analyzes data from the financial crisis in 2008-2009 and concludes that HFT reduces volatility.

Our study also relates to existing simulations in the context of asset markets. For example, Kearns et al (2010) study the profitability of HFT. Using stock market data from 2008, they simulate an “omniscient” HFT trader that effectively overestimates the maximum possible profit that could be earned from HFT traders that year. They conclude that the total profit derived from HFT trading is in general quite modest relative to the total trading volume. In contrast with their work, we study the impact of HFT on the market, not the profitability of HFT. Their work is also based on the complete history of every message from the NASDAQ stock exchange in 2008, so they are simulating additional HFT trading on top of the existing history of trades. Arthur et al. (1997) created a stock market simulation environment called the Santa Fe Artificial Stock Market. Their work is part of a line of work in “agent-based financial markets.” The goal this line of work is to simulate autonomous agents that learn and adapt their strategies over time. Therefore the focus is more on the performance of the agents and their strategies. In our work, we do not program adaptive learning agents. Instead our agents have set algorithmic trading strategies and our goal will be to see what kind of impact such strategies have on market outcomes.

Experimental Method

In this section we will start by describing the initial assumptions and settings of our simulation environment. We then describe the trader types and parameter settings for each. We then give the strategies that each of our traders uses, and discuss the procedure that each trader follows for deciding the volume of shares and the price at which to either sell or buy. Finally, we outline the tests we have run and provide our results.

Initial settings

This section serves as a discussion of all assumptions and initial settings in our simulation. In our simulation, technical traders use 20 day moving averages of a stock's price, where *price* refers to the dollar amount a stock is trading at. Because of this we must initially generate a stock price *history*. We first simulate a random price between \$5 and \$400. We use that number as the start price on the first day of the history. Then, for every day that follows, the price is decided by drawing from a normal distribution whose mean is the previous day's price. The final price in the price history is the *initial price* of the stock for the simulation. We also randomly generate the *fundamental value* of the stock, or how much the stock should actually be worth given the value of the company. The fundamental value of the stock on the first day of our simulation is a price uniformly chosen at random within ±8% of the initial price in the market. Then each day after that, the fundamental value is once again a random value within ±8% of the value the day before. A day in our simulation is considered 12 hours. The 20 day stock averages are the last 20 end-of-day prices averaged together.

We also implement a standard order book. Our order book is made up of two ordered lists, one for buy orders and one for sell orders. The sell order list is ordered from high to low and the buy ordered list is ordered from low to high. When an order is placed it gets put into one of these two lists, maintaining order.

The Traders

This section will serve to give a brief description of trader types as well as all parameters that are used for each type.

In our simulation we have 3 different trader strategies: *fundamental* analysis, *technical* analysis, and *pairs* trading. All fundamental traders have the same strategy whether they are algorithmic traders or humans. The same is true for technical traders. Pairs traders are only simulated as algorithmic traders. The difference between *human* traders and *algorithmic* traders is simply that algorithmic traders trade with higher frequency and lower latency than human traders. *Frequency* refers to how long it takes the trader to make a decision after the last decision. *Latency* refers to the time after a decision has been made and before the trade is executed. Therefore, the algorithmic traders can both make trade decisions more often and can execute their trades faster.

Each trader has a number of parameters. Human traders make their first decision in the fifth second of the simulation. They can execute their first trade at 10 seconds. Human trading frequency is between 95 and 105 seconds (after their last decision they wait a random value between 95 and 105 seconds before they make the next decision), and latency is between 35 and 45 seconds. Algorithmic traders are much faster. They make their first decision in the first second and can execute in the next second. Then latency and frequency are both one second so they can make a decision every second and always trade the second after that. These parameters are summarized in the table below.

Risk and *budget* are the two parameters that are the same for all trader types. *Risk* is a random value set for all traders between 1 and 10. This represents how much risk a trader is willing to take when trading stocks. This risk value for each trader is constant throughout a simulation. Budgets are the amount of money each trader starts with. The budget is set to $360. This initial budget value allows our simulation to have similar numbers of buy and sell orders. Naturally, the budget gets updated as traders execute trades.

The Strategies

To approximate the behavior of traders in real markets, we implemented simple trading strategies based on commonly understood behaviors of traders in the market.

Fundamental traders trade based on what they believe to be the fundamental value of a stock. At any given time, each fundamental trader is endowed with a random value that is within 30% of the actual fundamental value. This random value represents the trader's *perceived fundamental value* of the stock. If the perceived value is greater than the *current price* (the most recent trading price), the trader will buy because they think the stock is worth more than it is selling for. In the opposite case, where the current price is greater than the perceived fundamental value, the trader thinks the stock is trading at a price higher than what it is worth, so they will sell. Specific bid and sell prices and share volumes are decided on as described in the next section.

Technical traders look at the 20-day moving average price of the stock, which updates at every second based on the last 20 days worth of seconds. If the current price crosses from above this moving average to below the moving average, then they sell. When the current price crosses from below to above then they buy. This simple strategy is based on the description from Janssen et al.

Given two stocks whose prices move together, a *pairs trader* sells or buys one stock based on the movements of the other. Our goal is to determine the impact on the market when a pairs trader mistakenly believes two stocks are correlated, and begins trading them as though they are paired, even though they are unrelated. So our pairs trader will assume two unrelated stocks are positively and closely correlated. Then, if one stock moves a significant amount (±3%) since the trader checked last, the trader will buy/sell the other stock at an appropriate price and volume.

Buying and Selling

Once each trader in our simulation makes the decision to either buy or sell, they all share a similar framework for deciding how many shares of a stock to buy or sell and at what price. We use *trade price* to refer to the dollar value at which a trader wishes to buy or sell the stock, and *trade volume* to refer to the number of shares that the trader wants to buy or sell. This section will outline how these values are set for any given trader.

We begin by defining a quantity called *movement*:

$$movement = (p\_value - c\_price) / c\_price,$$

where *c_price* refers to the current price of the stock, and *p_value* refers to the trader's perceived value of the stock. Specifically, *p_value* is determined as follows for each trader type.

Trader Type	***p_value***
Technical	20-day moving average
Fundamental	perceived fundamental value
Pairs	max price over the past 40 seconds, if price is currently decreasing min price over the past 40 seconds, if price is currently increasing

Note that pairs traders are deciding how much of a stock to trade based on the *movement* value for a separate, "paired" stock. Loosely speaking, *movement* reflects how much a trader wants the stock. For instance, when the perceived fundamental value is far above the current price, the trader will see this either as a situation in which to make a lot of money, or one in which they must trade quickly before the market moves the price closer to its true, higher value.

For a buy order, volume is determined using the following formula.

$$Buy\ Volume = baseVol * movement * random * risk$$

where *baseVol* is an initial trade volume, which is set to 500, *movement* is as defined above, *risk* is a random integer between 1 and 10 (as described above), and *random* is a value drawn uniformly at random between 1 and 2. The *random* factor is to account for unpredictable elements that may change from one trade to the next, like investor mood.

When a trader buys, the price is set by the current price subtracted by an offset value, then added to *movement* multiplied by *risk* and *random*. I.e.,

$$Buy\ Price = current\ price - offset + movement * random * risk$$

where *offset* is 0.3, and everything else is as described above for volume. The purpose of *offset* is to reflect the fact that a trader wanting to buy will want to buy low. Note that if the trader is not very motivated to buy the stock (as reflected in *movement*) and if they are risk-averse, then they might bid below the current price. Of course, buy volume is also subject to the trader's budget. If their budget is less than the cost of the volume of shares they wish to buy, then the volume is automatically truncated as needed.

When a trader sells a stock, the volume is determined as follows.

$$Sell\ Volume = volume\ held * movement * random2 * risk2$$

where *volume held* is the number of shares of the stock that the trader currently owns, *movement* is as defined above, *random2* is a value between .85 and 1.15, and *risk2* is the trader's risk scaled to a value between .5 and 1.5. When a trader sells, the price is set by adding the current price to the offset value, which is once again 0.3, minus the *movement* term

$$Sell\ Price = current\ price + offset - movement * random2 * risk2$$

Experiments & Results

We first describe our test of how liquidity and volatility change in the simulated market as the proportion of algorithmic traders increases. We then describe our tests that explores how pairs traders, trading on two unrelated stocks, will affect the prices of these stocks relative to their fundamental values.

Liquidity

The results of our liquidity test, with 95% confidence intervals, can be seen in Figure 1. Each point on the graph represents the average over 50 runs, each of 30-days of trading. For the first point on the graph, our simulation has 1000 human technical traders and 1000 human fundamental traders, with no algorithmic traders. In each successive point of the graph, the proportion of algorithmic traders increases. The number of technical and fundamental traders is equal at all points.

As in Hendershott and Riordan (2011) we use effective half-spread as our measure of liquidity. As half-spread decreases, liquidity increases. For a given trade of a given stock, the effective half-spread is computed as follows.

Half-spread =sign(trade price-midpoint)/midpoint*

Where *trade price* is the sale price, *midpoint* is the price halfway between the transacting bid and ask prices, and *sign* is determined by how the trade was initiated. If a seller initiated the trade then *sign* = 1, if a buyer initiated the trade then *sign* = -1. We compute the half spread at 800 evenly-distributed points in time over the course of the 30-day run. At the end of each run, we compute the average half-spread over the 800 values. Each point on the curve in Figure 1 is thus the average of 50 average half spreads.

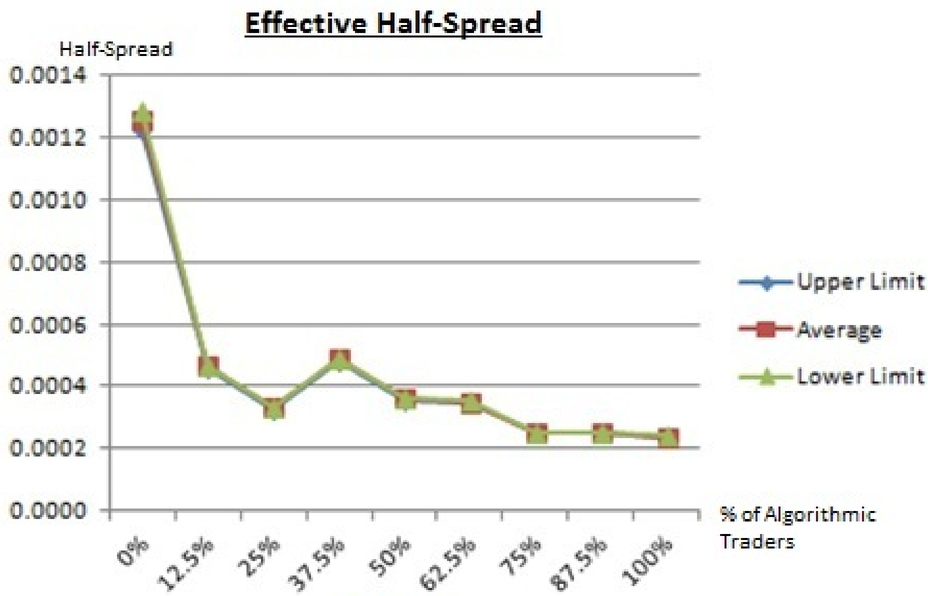

Figure 1

Volatilty

We define volatility, as in Zhang (2010), as variance in price. We compute this variance over the duration of each 30-day simulation period, using the 800 stored prices, similar to the calculation of average half-spread, described above. Our results can be seen below in Figure 2. For each point on the graph, we again average 50 runs, as with our liquidity test. The figure also shows the 95% confidence intervals.

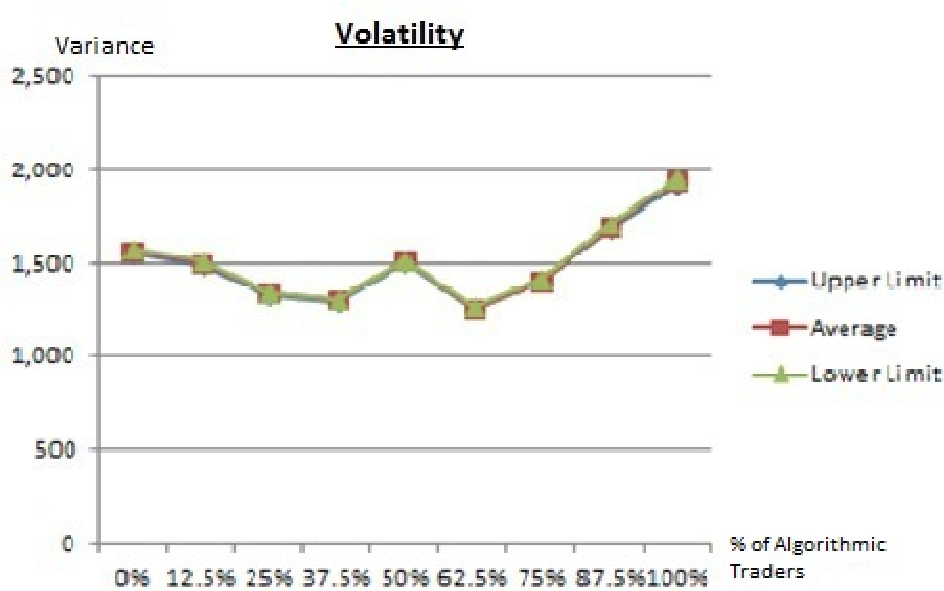

Figure 2

Again, for each successive point on the curve in Figure 2, we increase the proportion of algorithmic traders, while always maintaining the same number of technical traders and fundamental traders, and always having a total of 2000 traders.

Pairs trading

In our pairs trading test we simulate two unrelated (un-"paired") stocks. We introduce both fundamental analysts, as well as pairs traders who are trading as though they expect the two stocks' values to be positively correlated. We measure the correlation of the prices and the correlation of the values of the stocks as the proportion of pairs traders increases. To measure correlation, we use the linear correlation coefficient (CC) computed on the 800 stored prices and values for each run mentioned in earlier tests. Our goal is to determine whether pairs traders that are mistakenly trading stocks that do not move together will make the prices of the stocks move away from their true fundamental values.

We run these tests 10 times each over 30-day periods. The results are shown below in Figure 3. Each point on each line represents the average correlation coefficient value of the 10 runs.

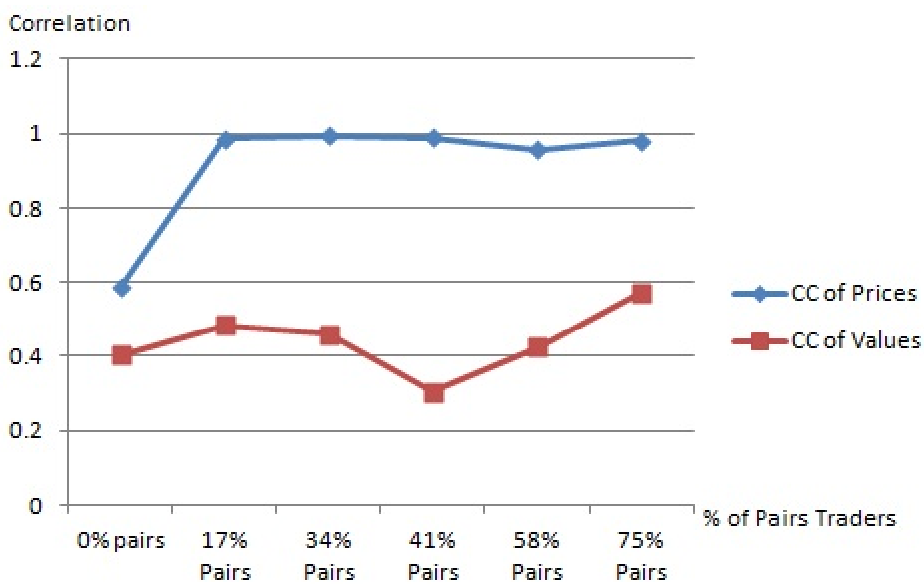

Figure 3

For each successive point on the curve in Figure 3, we increase the proportion of pairs traders to fundamental traders, while maintaining a total of 200 traders.

We also measure correlation between the prices and the values of each stocks. All other test parameters remain the same, except we run these tests 15 times each over 30-day periods. Many individual runs show a downward trend in correlation between price and value, which indicates that the pairs traders are causing the prices to move away from their true values. See, for example, the graph of the correlation between price and value of one stock over a single 30-day run in Figure 4.

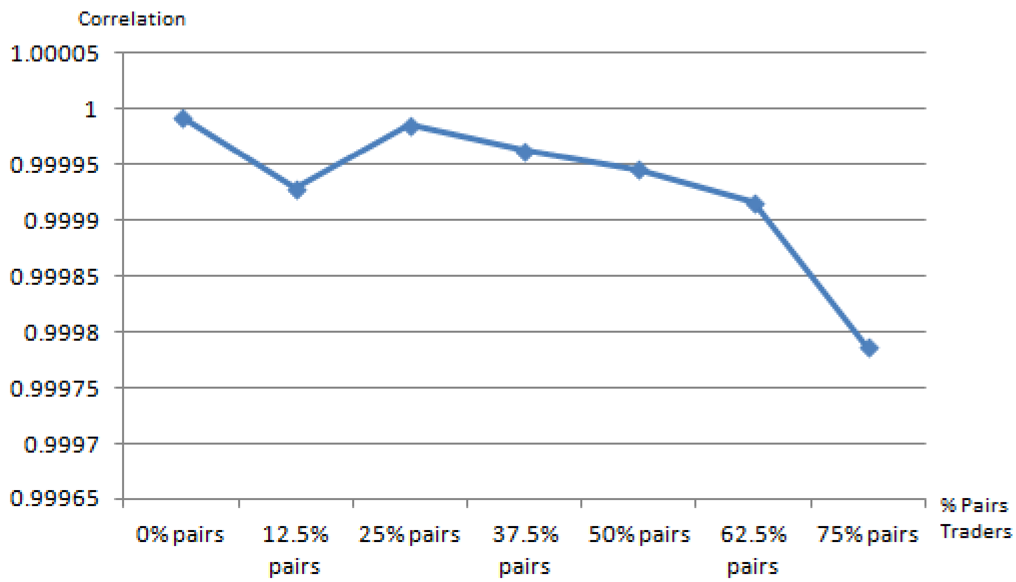

Figure 4

However when averaging all 15 runs together, the trend is less clearly defined. The average correlation between price and value of a single stock over all 15 runs is shown in Figure 5.

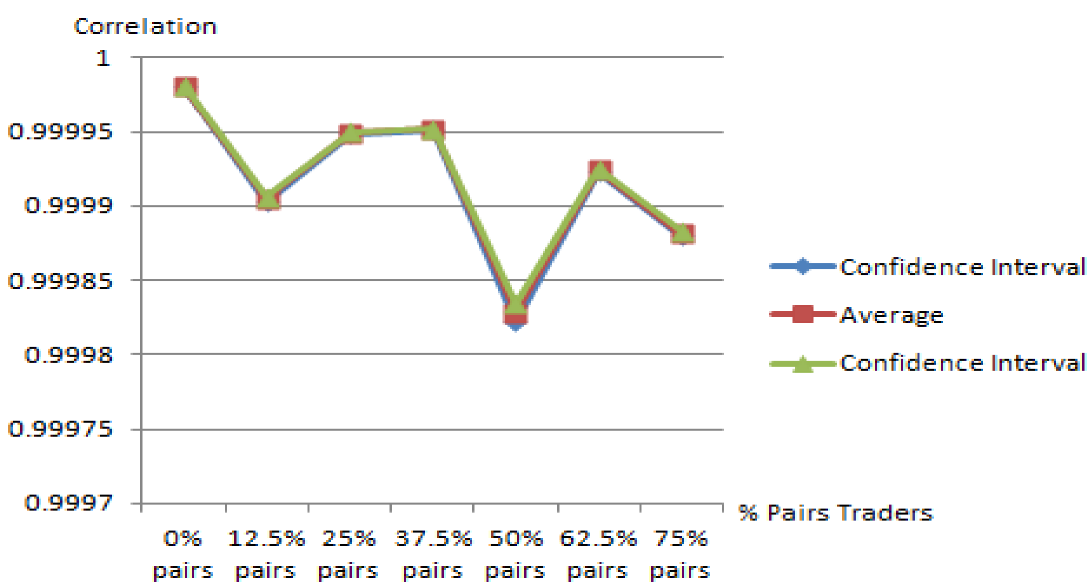

Figure 5

To take a closer look at what unfolds over the course of a single run, we present Figure 6, which shows what the market looks like when no pairs traders are present. In this particular run we control the fundamental values rather than randomly generating them. The first stock's fundamental value is the red line, increasing, then decreasing step function.. The second stock's fundamental value line is the flat green line. The figure shows that when there are no pairs traders, the prices stay near their corresponding fundamental values. However, in Figure 7 we show the same test with 50% pairs traders. Now the prices are far from their fundamental values and are highly correlated with one another.

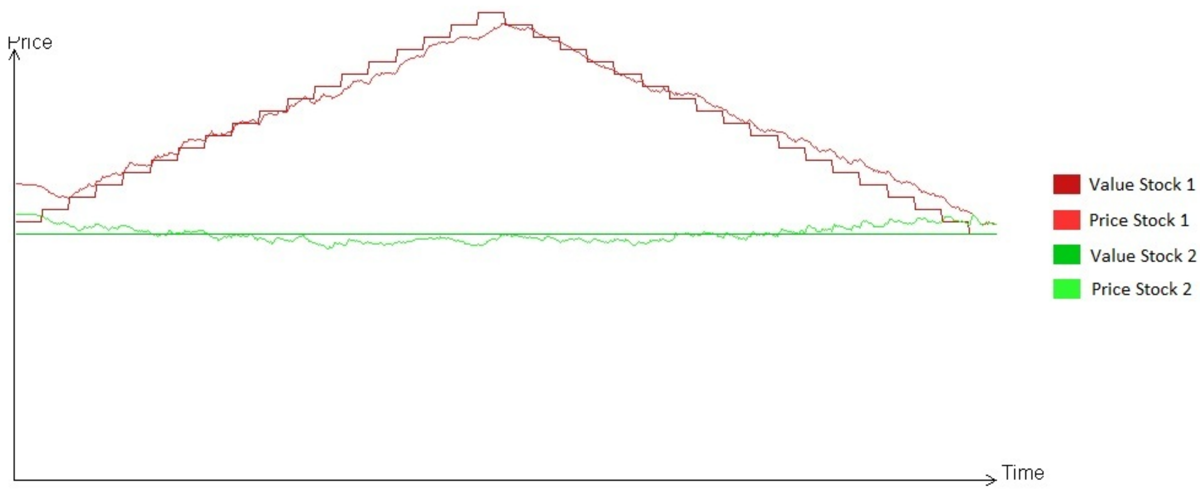

Figure 6

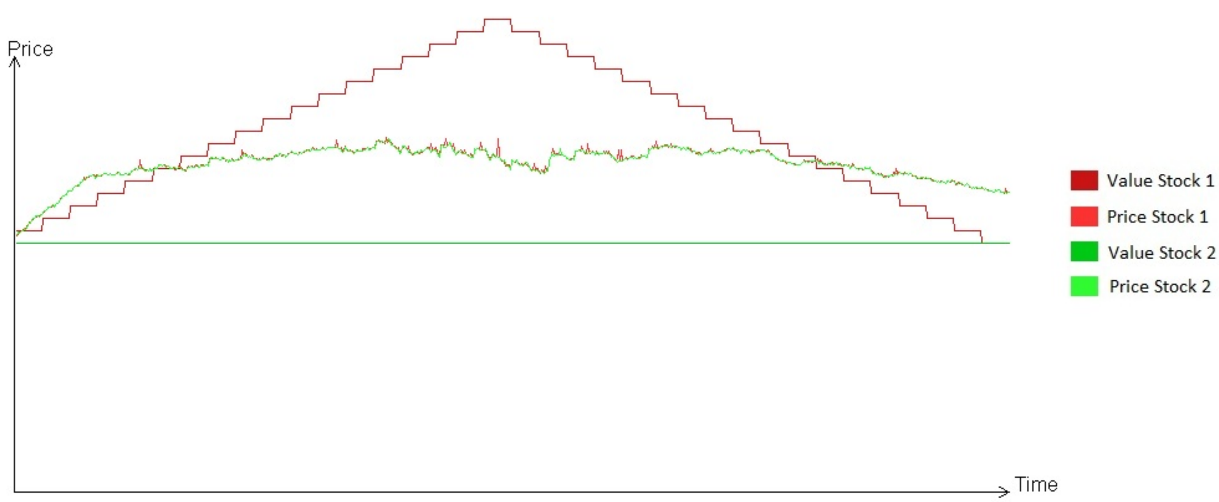

Figure 7

Conclusion

We created a simulation of an asset market with traders, both human and algorithmic, using very basic strategies of technical and fundamental analysis. We tested the impact of algorithmic trading on market quality measures like liquidity and volatility and conclude that liquidity improves as algorithmic trading increases, while the results for volatility are inconclusive. These results reflects the general consensus in the existing empirical literature regarding AT's impact on liquidity, as well as the lack of concensus about AT's impact on volatility. We also implemented a test with algorithimic pairs traders and conclude that if pairs traders mistakenly assume to unrelated stocks are correlated, the prices of those stocks with become decoupled from their fundamental values and will indeed become artificially correlated with one another.

References

1. Arthur B., Holland J. H., LeBaron, B., Palmer,R., and Tayler, P. (1997). Asset Pricing under Endogenous Expectations in an Artificial Stock Market. *The Economy as an Evolving Complex System, II*, edited by W. Brian Arthur, Steven Durlauf, and David Lane, Vol. XXVII, SFI Studies in the Sciences of Complexity, Addison-Wesley.

2. Biais, Bruno, and Woolley, Paul (2011). High Frequency Trading. *Working Paper*.

3. Bruno Biais & Pierre-Olivier Weill (2009). Liquidity Shocks and Order Book Dynamics. *NBER Working Papers* 15009.

4. Back, K. and Baruch, S. *Forthcoming*. Working Orders in Limit-Order markets and Floor Exchanges. *Journal of Finance.*

5. Brogaard, Jonathan A. (2010). High frequency trading and its impact on market quality. *Working Paper, Northwestern University*.

6. Cvitanic, Jaksa and Kirilenko, Andrei A., High Frequency Traders and Asset Prices (2010). Available at SSRN: http://ssrn.com/abstract=1569067.

7. Domowitz, Ian, and Yegerman, H. (2005). The Cost of Algorithmic Trading: A First Look at Comparative Performance. *Algorithmic Trading: Precision, Control, Execution.*

8. Gsell, Markus (2009). Algorithmic Activity on Xetra. *Journal of Trading*Bloch,

9. Gsell, Markus, and Gomber, Peter (2008). Algorithmic trading versus human traders – Do they behave different in securities markets? *Working paper, Goethe University*, Frankfurt am Main.

10. Hasbrouck, Joel and Saar, Gideon (2011). Low-Latency Trading. *Johnson School Research Paper Series* No. 35-2010. Available at SSRN: http://ssrn.com/abstract=1695460.

11. Hendershott, Terence, Jones, Charles and Menkveld Albert (2011). Does Algorithmic Trading Improve Liquidity? *The Journal of Finance*, LXVI(1), 1-34.

12. Hendershott, Terrence, and Riordan, R. (2011). Algorithmic trading and information. *Working paper, University of California, Berkeley*.

13. Janssen, Cory, Langager, Chad, and Murphy, Casey. Technical Analysis: Introduction. Online article, found at http://www.investopedia.com/university/technical/default.asp#axzz1YV3t5qY5

14. Kearns, Michael, Kulesza, Alex and Nevmyvaka, Yuriy (2010). Empirical Limitations on High Frequency Trading Profitability. Available at SSRN: http://ssrn.com/abstract=1678758

15. Kearns, Michael and Ortiz, Luis (2003). The Penn-Lehman Automated Trading Project. *IEEE Intelligent Systems*.

16. Menkveld, Albert J., High Frequency Trading and The New-Market Makers (August 14, 2011). ; EFA 2011 Paper. Available at SSRN: http://ssrn.com/abstract=1722924

17. Mukerji, P, 2009. Ready for Capital Account Convertibility? *Journal of International Money and Finance*, 28(6), 1006-1021.

18. Mukerji, P., and K. Saeed, 2011. Likely Causes of the US Housing Market Crisis: A System Dynamics Investigation. *Under review*

19. Riordan, Ryan, and Storkenmaier, Andreas (2011). Optical Illusions: The effects of exchange system latency on liquidity. *Working Paper, University of Karlsruhe, Germany*.

20. U.S. Commodities Futures Trading Commission, and U.S. Securities and Exchange Commission (2010). Findings Regarding the Market Events of May 6, 2010. (Washigton D.C.)

21. Zhang, Frank (2010). High-Frequency Trading, Stock Volatility, and Price Discovery. Available at SSRN: http://ssrn.com/abstract=1691679

22. Zweig, Jason (2010). The Intelligent Investor: Why your Stock Portfolio is Acting like a Commodity Basket. *The Wall Street Journal,* November 20, 2010.