Candidate Number: 124376
Word Count: 2377

INT14: Agent Based Simulations of High Frequency Trading in Financial Markets

Supervisor: Dr A. Gerig

Abstract

Financial markets are subject to a growing proportion of trades being executed by autonomous algorithms. This is often called High Frequency Trading (HFT). In this paper I create a simple model to study the impact of HFT on market participants. I find that such strategies improve liquidity through allowing more transactions to take place, without an adverse effect on pricing or volatility - HFT increases the likelihood of less competitive orders being filled. This is an important result in light of current debates regarding changes to financial market regulation, and how increased automation is changing markets for the average investor.

Introduction

Financial markets have evolved considerably over the last 20 years. All major exchanges have switched to electronic trading platforms from the open outcry "pits" of the past [1], such that it is now possible for orders to be sent to the market and executed without any direct human involvement. This has resulted in the development of automated trading software that analyses market data in real time, placing orders based on heuristic strategies. These programs can buy and sell on a much shorter timescale than traditional traders, hence their name – High Frequency Traders (HFT). Such traders seek to make small profits per transaction through exploiting short-term price fluctuations, without assuming overnight positions.

My study focuses on one particular effect linked to HFT – the synchronising of price responses. "In an empirical analysis, Gerig [2]demonstrates that price movements on the NASDAQ have become increasingly *synchronised* over the past 10 years, and this is specifically due to HFT. Figure 1 exhibits this effect. I model this effect of connecting markets in order to ascertain the impact on investors and traditional traders operating different strategies.

As a physicist, I feel it important to clarify my assumptions and approximations to make the model as simple as possible whilst still being relevant. In a treatment similar to the statistical mechanics of gases, I wish to ignore specific properties of the majority of traders – their risk profile, specific strategy or motivation. Traders submit orders at random, based on a Zero-Intelligence model [3, 4]. This treatment strips out the idiosyncrasies of individuals' behaviour and assumes only local interactions are of significance – traders are only interested in meeting their own specific price expectations. To consider the HFT's effect, I

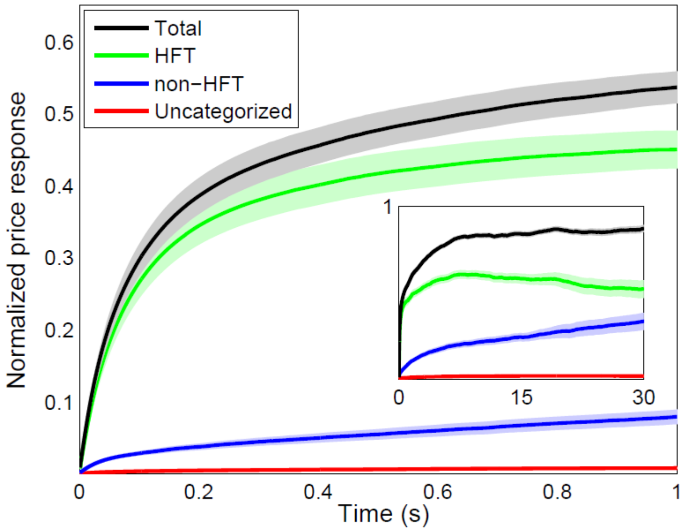

FIG 1: Normalised price response of stock *i* due to stock *j≠i,* for 40 US stocks, decomposed into the amount due to HFT activity (green), non-HFT activity (blue) and uncategorised activity (red). Standard errors of the sample means are indicated in the shaded colour. Taken from A. Gerig – *"High Frequency Trading: What is it Good for?"* (Working Paper)

run 2 identical markets simultaneously and allow the HFT to trade between them when transactions on individual exchanges cannot occur. I show that the markets with HFT increase the likelihood that a party entering the market can transact. In the regime where HFT makes profit, we see that traders pay a premium for transactions – sellers receive less and buyers pay more, on average. However, in the limiting case of a competitive landscape for HFT, buying and selling prices converge and the average price is no different from the control market. This satisfies the important concept of liquidity, which is a key goal of financial markets and an important measure of their success– agents are able to transact when they want to, and at a fair price. Volatility, measure by the standard deviation of transaction prices, is also a relevant observable. I do not report a significant impact on volatility.

These results are useful in that they help ascertain HFT's role in the modern marketplace. HFT is already the dominant form of trading in global equity markets, and with the advent of electronic trading in the bond markets [5] it can only be anticipated that automated strategies will spill over. As such, understanding HFT and its impact is an increasingly relevant undertaking.

Model

The model simulates the *continuous double auction*, a market structure common to most, if not all modern financial exchanges. Traders may submit *bids* and *offers* to buy and sell respectively, that represent the best price that they are willing to transact at. If prices *cross* - a bid meets or exceeds a previous offer, or the converse – a transaction takes place at the earlier listed price. If an incoming order is unable to fill any existing orders, it is placed in the *limit order book.* This consists of 2 lists, the *bid book* and the *ask book* – which contain the previously unfilled orders. At each time step, a random order is generated with equal probability to buy or sell, with a price following a uniform probability distribution between 1 and 200. Trades take place in units of one "share" at a time, as in Gode and Sunder's original model. We wish to

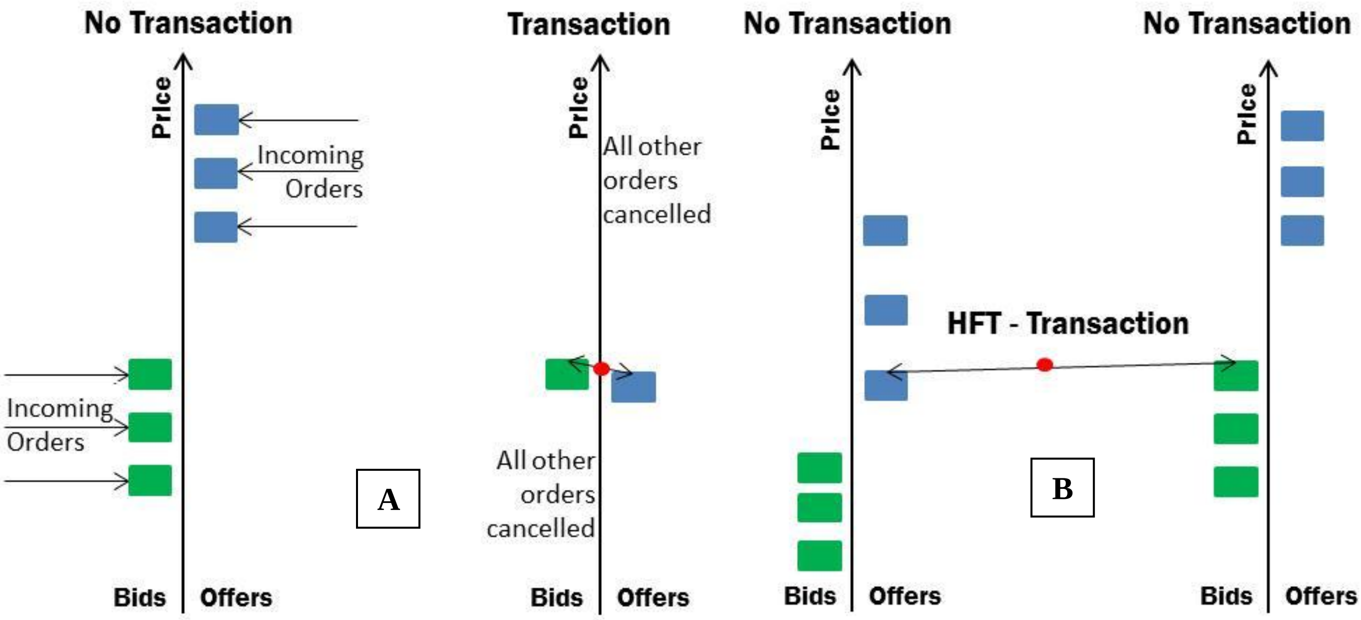

FIG. 2: A diagram of the order book in both scenarios modelled. (**A**) exhibits the market without HFT, at time steps with and without a transaction occurring. Note that transactions can only occur if a bid exceeds an offer. (**B**) shows the connecting effect of the HFT. The order book is cleared entirely after each transaction.

evaluate the effect on a large number of market participants, therefore it is more relevant to ignore volume of individual transactions and focus on price and total number of transactions. When a transaction takes place, I clear all existing orders, so as to emphasise the importance of liquidity – market participants need to be able to trade when they want to, at a fair price. If orders are not met at the earliest opportunity, then they are cancelled. Figure 1 illustrates this diagrammatically.

The HFT interaction is modelled by running 2 identical exchanges simultaneously. If transactions are unable to take place on either exchange in a given time step, the HFT is permitted to transact between the 2 markets. They are able to do this in 2 ways: maximising profit, and with no profit at all. In the first case, they transaction occurs at the best ask price in one market, and the best bid, with the HFT taking the difference. In the second regime, the transaction occurs half way between the bid-ask spread, giving the best price to both parties as the buyer pays less and the seller receives more than they had initially specified.

Consider a transaction between the 2 *markets*
Let the bid price on market 1 *be denoted by* b_1 *and the ask price on market* 2 *be denoted by* a_2*, where* $b_1 \geq a_2$

If the HFT is permitted to make a profit, then the transaction occurs at b_1 *on market* 1 *and* a_2 *on market* 2*, generating a profit of:*

$$b_1 - a_2 \qquad (1)$$

Assuming the HFT is forbidden to make a profit, the transaction then takes place at ,

$$\left(\frac{b_1 + a_2}{2}\right) \qquad (2)$$

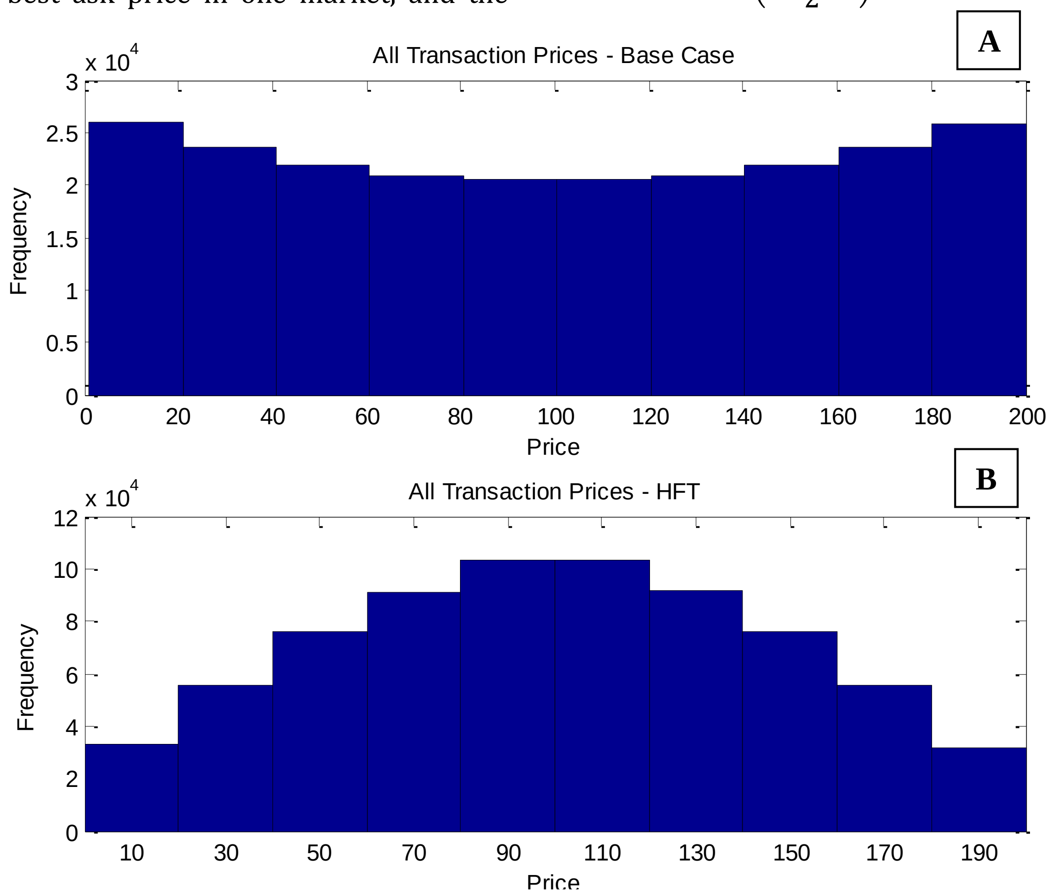

FIG. 3: Distribution of transaction prices across the entire simulation. (**A**) shows the market without HFT; (**B**) shows the market with HFT, including transactions over both exchanges. Note that without HFT the average price does not occur as often as prices in the extremes.

We analyse the latter case, where transaction prices are governed by (2) because such is the competitive landscape in which HFTs operate, we assume that their margins are very small per trade. The HFT matches orders within the same time step, to replicate the speed at which they are able to transact. Automated orders can be submitted many times per second [6], far quicker than a human trader is able to process information, therefore this approximation is valid. Speed is in fact such an important factor that trading companies have taken to *colocation,* placing their servers in the same building as stock exchanges. Wissner-Gross and Freer discuss the implications of trading frequency along with the limiting effect of light propagation delays in "Relativistic Statistical Arbitrage" [7].

Results

The key factors that are of interest to retail and traditional investors can be summarised in 3 questions. Can I transact? At what price? How does the price fluctuate? I therefore analyse these 3 observables in the form of: the probability that a submitted order will result in a transaction, average price and volatility (standard deviation of the price series), comparing the market with and without HFT. I ran each simulation over 10000 time steps, 100 times, and recorded average values of the relevant observables. Values for the HFT case are aggregated over the 2 separate markets. The system is totally symmetrical so both markets are equivalent. The results are represented graphically in Figure 4 and in tabular form in Tables A and B. Figure 4(A) depicts the average price at which a transaction is executed, in which I observe no change in the HFT regime. In the base case, we observe an average value of 100.4 over 100 iterations of 1000 "tick" trading days, with a standard deviation of 1.2. In the HFT case, I record an average price of 100.5 with a standard deviation of 0.7. Figure 4(B) shows the average volatility in both cases. The market with HFT shows a drop in volatility from 60.1 to 47.3. Figure 4(C) shows that the probability that a given order will result in a completed transaction is significantly greater in the HFT case, going from 0.2250 to 0.3581, and the final Figure 4(D) confirms this, with a marked increase in the number of transactions in a trading period. It is also worth noting that there is no difference between the average prices that buyers and sellers trade at, due to our assumption that HFT can make no profit.

Discussion

I find that the model reproduces several empirical findings that have yet to be explained:
(1) Transaction prices are more accurate when HFT is present, i.e., they are closer to the equilibrium value [8, 9].
(2) Liquidity is increased with HFT. [9 also, 10]
(3) Volatility is reduced with HFT [10 also].

I consider first the price series. Order prices are generated uniformly at integer values between 1 and 200 – traders are assumed to have "zero-intelligence". Therefore we expect the equilibrium price to be the average of this distribution – 100.5. Data from both simulations converge to this value, within the 2 standard error range that defines a 95% confidence interval. The overall distribution of transaction prices across each model is depicted in figure 3. Figure 3(A) shows the base case. Note that more transactions occur at the edges of the price range. This is because transactions take place at a specific order price in the original Zero-intelligence model. Therefore transactions are more likely to take place when extreme orders are placed – low offers and high bids. With HFT, transactions can be facilitated between bid and ask values. Figure 3(B) illustrates this well – the distribution of transactions has a single peak around the mean price. We can therefore assert that HFT increases the likelihood that a trader can transact closer to the equilibrium price, which in this case is the fair value of the security. Without HFT, the market still values the security at 100.5 on average, but that does not mean that a trader can transact at that price.

We assess precisely the movement of the price with time by examining the volatility. Volatility is the standard deviation of the prices. Brogaard, along with Hasbrouk and Saar [9, 10] have conducted empirical studies demonstrating that HFT reduces intraday volatility. My simulation supports this claim for the same reason discussed before. The matching of extreme bids and asks across the 2 markets and providing competitive prices to both means that prices are more tightly distributed around the mean.

Another key consideration is Liquidity. An asset is liquid if "if it is more certainly realisable at short notice without loss" [11] Liquidity can be defined quantitatively in a number of ways [12]. However, my model accounts for the requirement of short notice, as orders are cancelled if they do not result in a transaction, and when they do transact, the price must satisfy the reservation price initially generated. As a result, my measure of liquidity is the number of transactions that take place per simulation, or the probability that an order transacts. The fact that orders are more likely to be filled in the HFT case shows that HFT provides liquidity to the market. Brogaard [8] finds that in the short term, HFT provides liquidity as the overall volatility of the market increases. My results support this study. Intuitively this makes sense, because as I explained previously, the principle of connecting markets at a mutually beneficial price facilitates transactions closer to equilibrium. HFT takes orders at the boundaries of the price range and transacts closer to the mean.

Conclusion

In this paper I consider the effect of High Frequency Trading in a simulated market. With the premise that HFT activity connects markets, I find that prices are more rational, volatility is reduced, and the market displays increased liquidity. This result is particularly relevant as it tells us that connected markets help investors. Rapid flow of price information helps market participants trade at a fair price, and reduces the opportunities for informed traders to benefit unfairly. I close with a quote by George Sauter of Vanguard, a large fund management company in his statement that *"high-frequency traders provide liquidity and 'knit' together our increasingly fragmented marketplace, resulting in tighter spreads that benefit all investors."* [13].

References

[1] Smith, R. 2010 *Is high-frequency trading inducing changes in market microstructure and dynamics?* Working Paper, Bouchet-Franklin Institute, Rochester, New York

[2] Gerig, A. 2012, *High Frequency Trading: What is it Good for?* Working Paper, University of Oxford

[3] Gode, D. K. and Sunder, S. 1993 *Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality, Journal of Political Economy 101, 1*

[4] Farmer, J. D., Patelli, and P., Zovko, I. I., *The Predictive Power of Zero Intelligence in Financial Markets, 2005 Proceeding of the National Academy of Sciences of the United States of America vol. 102, no. 6*

[5] Demos, T. March 2012 *Bond traders switch to electronic processing, The Financial Times, 28/03/2012* *http://www.ft.com/cms/s/0/9037a65e-78e6-11e1-88c5-00144feab49a.html#axzz1rBN4i1KM*

[6] Nanex, *Quote and Trade Rates 2004-2010* *http://www.nanex.net/FlashCrash/QuoteRates.html*

[7] Wissner-Gross, A. D. and Freer, C. E. 2010 *Relativistic Statistical Arbitrage* PHYSICAL REVIEW E 82, 056104

[8] Hendershott, T Jones, C. M. Menkveld A. J. 2011 JOURNAL OF FINANCE 66, 1

[9] Brogaard, J. 2012 *High Frequency Trading and Volatility,* Working Paper, University of Washington *http://papers.ssrn.com/sol3/papers.cfm?abstract id=1641387*

[10] Hasbrouk, J. and Saar, G. 2010 *Low-Latency Trading,* Working paper, New York University http://ssrn.com/abstract=1695460

[11] Keynes, J. M. 1930 *A Treatise on Money, vol. 2*

[12] Gabrielsen, A, Marzo, M., and Zagaglia, P. 2011 *Measuring market liquidity: an introductory survey,* Working Paper, University of Bologna Department of Economics

[13] 2012 *The Economist Special report: Financial innovation* http://www.economist.com/node/21547988

Tables: Simulation Data. **Table A** – Average data to compare both cases. **Table B** – Data from each market separately in the HFT simulation

Table A

Parameter (Average taken from 100 runs of 10000 iterations)	**Base Case**	**HFT (average over both markets)**
Price SD	100.5 1.38	100.5 0.71
Volatility SD	60.1 0.6	47.4 0.4
Volume SD	2255 33	3580 6
Probability of Transaction	0.226	0.358

Table B

Parameter (Average taken from 100 runs of 10000 iterations)	**Market 1**	**Market 2**
Price SD	100.5 0.86	100.5 0.78
Volatility SD	47.4 0.4	47.4 0.4
Volume SD	3588 6	3571 10
Probability of Transaction	0.359	0.357
Buyer price SD	100.5 1.0	100.5 0.85
Seller Price SD	100.5 0.97	100.5 0.98
HFT Trades SD	1899 29	1899 29

Figure 4: Comparison plots of Average Transaction Price, Volatility, Transaction Probability and Volume (number of trades) for both cases. Note that the introduction of HFT has no discernible effect on price, but statistically significant reduction of Volatility, along with an increase in the number of trades and the likelihood of a given order being filled. Error bars denote 95% confidence intervals.