JOHNSON
Cornell University

Johnson School Research Paper Series #08-2011

A Dysfunctional Role of High Frequency Trading in Electronic Markets

Robert Jarrow—Cornell University
Philip Protter—Columbia University

March 2011

A Dysfunctional Role of High Frequency Trading in Electronic Markets

Robert A. Jarrow*and Philip Protter†

March 8, 2011

Abstract

This paper shows that high frequency trading may play a dysfunctional role in financial markets. Contrary to arbitrageurs who make financial markets more efficient by taking advantage of and thereby eliminating mispricings, high frequency traders can create a mispricing that they unknowingly exploit to the disadvantage of ordinary investors. This mispricing is generated by the collective and independent actions of high frequency traders, coordinated via the observation of a common signal.

KEY WORDS: High frequency traders, algorithmic traders, electronic trading, arbitrage opportunities, martingale measures.

1 Introduction

Arbitrageurs are viewed positively by economists as serving a useful role in competitive financial markets. Arbitrageurs search for mispricings, and in exploiting them, they eliminate these mispricings and increase market efficiency. This is the ruthless Darwinian nature of competitive markets. The initial empirical literature on algorithmic and high frequency trading (see Brogaard (2010), Castura, Litzenberger, Gorelick, and Dwivedi (2010), Hasbrouck and Saar (2010), Hendershott, Jones and Menkveld (2008), Hendershott and Moulton (2010), Hendershott and Riordan (2009), Riordan and Storkenmaier (2009), and Stoll (2006)) supports the belief that high frequency (algorithmic) traders serve a similar function in electronic markets, making them more efficient. Preliminary evidence suggests that computer based trading reduces bid/ask spreads, increases market liquidity, and decreases market volatility. Yet, the verdict is

*Johnson Graduate School of Management, Cornell University, Ithaca, New York 14853 and Kamakura Corporation.

†Statistics Depatrment, Columbia University, New York, NY 10027; supported in part by NSF grant DMS 0604020

still out, both in the academic (see Zhang (2010) and Kirilenko and Kyle (2011)) and financial community[1].

Adding to this debate, we construct a model to show that high frequency traders may not increase the efficiency of electronic markets. In fact their trades can create increased volatility and mispricings (deviations from fundamental value) that they exploit to their advantage. These self-induced mispricings are generated by two market realities: (i) that demand curves for equity shares are downward sloping, and (ii) that there is a differential speed of transacting across traders. Combined, these two market conditions enable high frequency traders to create a trend in market prices that they exploit to the disadvantage of ordinary traders. The price trend is generated by their collective but independent actions, coordinated via the observation of common signals. The common signal could be the difference between the futures and forward prices of a stock index, wrongly believed to be an arbitrage opportunity, or trading based on electronic news generated in the financial press. Their speed advantage is captured by making the high frequency traders' strategies optional processes, instead of predictable processes. This technical distinction incorporates the economic advantage of speed in the execution of trades.[2]

From a regulatory policy perspective, this predatory aspect of high frequency trading should be excluded whenever possible. But, this is a difficult if not an impossible task. One cannot (and should not) prevent investors from trading based on common signals. However, it is the differential *speed advantage* of high frequency trading that causes the inequity.[3] To the extent that the speed advantage is generated by preferential treatment in the execution of market orders, these can be eliminated. To the extent that the speed advantage is due to financial resources, they will be impossible to remove.

To formalize our model, we use the tools of mathematical finance. We define an economy to be well-functioning if, given the relevant information, there exist a probability measure making the market price process a local martingale.[4] In a well functioning market, we can show that there exist no arbitrage opportunities for ordinary traders. In contrast, in such an economy, we show that the high frequency traders create abnormal profit opportunities for themselves that cannot be exploited by ordinary traders. These abnormal profit opportunities depend on their trading speed advantage. If one removes their speed advantage by making their trading strategies predictable instead of optional processes,

[1]See, e.g. *Wall Street Journal*, June 30, 2010, "Fast Traders Face Off with Big Investors over Gaming" by Scott Patterson.

[2]Predictable processes are adapted processes that can be obtained as limits of left continuous processes with right limits. Optional processes are adapted processes that are obtained as limits of right continuous processes with left limits. See Jarrow and Protter (2008) for an explanation of the economic relevance of these two types of processses.

[3]In a sense, this is similar to the intent of insider trading laws. Inside information is based on the "fundamental" price process. The law states that it is illegal to trade on inside information until it is released to the market. In our situation the "inside" information is not based on the "fundamental" price process but the "order flow" process. Both types of information affect prices.

[4]Of course, in probability theory, martingales have a long history of characterizing fairness in the winnings generated from gambling.

then these abnormal trading profits disappear.

An outline for the paper is as follows. Section 2 presents the basic idea underlying the model, which is presented in section 3. Section 4 presents the key theorem in the paper, while section 5 concludes.

2 The Basic Idea

Broadly speaking, there are two types of computer based or algorithmic traders. One set of algorithmic traders use computers to reduce execution or liquidity costs. The second set use the computer to automate trading based on mispricings or market signals. This second set of algorithmic traders often employ high frequency technology to trade quickly, before the mispricing disappears or the market signal is incorporated into the price. The paper concerns this second set of high frequency traders. For an analytical model considering the first class of traders see Cvitanic and Kirilenko (2010). For an empirical study of some effects of the second class of traders, see Kirilenko and Kyle (2011).

We consider a frictionless and competitive market. Hence, from the perspective of any trader, the market is perfect. There are no bid/ask spreads and markets are perfectly liquid. This price process is taken to be exogenous and all traders act like price takers believing that their trades do not change the price. We have two types of market participants. One type we call *ordinary traders*, e.g. pension funds and small investors. The second type we call *high frequency traders*. The ordinary and high frequency traders have different speeds of transacting, characterized by the dependence of their trading strategies on their different information sets.

The ordinary traders just observe the price process and their trading strategies are represented by predictable processes (limits of left continuous adapted processes, or even continuous adapted processes). In contrast, the high frequency traders see a common signal, which could be the realization of some market related event or a mispricing. The ordinary investors do not see this signal[5]. By construction, the high frequency traders can transact instantaneously based on this signal, before the signal is incorporated into the market price. This implies that their trading strategies need not be predictable, but just adapted processes (technically known as *optional* processes) with respect to the price process. This difference in trading strategies characterizes the speed advantage distinction between ordinary and high frequency traders in our model.

In addition, we assume that seeing the same signal, all high frequency traders do the same trade at the same time. Acting independently based on this signal, unknowingly but in unison, they collectively act like a large trader. This collective action has a quantity impact on the market price since demand curves are downward sloping. This correlated trading activity corresponds, in practice, to high frequency traders who use the same "alpha" generating trading strategies, e.g. index arbitrageurs or momentum traders. In contrast to the large trader

[5] In fact, they could be allowed to see this signal, but in this case, due to unspecified constraints, they cannot transact quickly enough based on this observation.

literature (see, e.g. Jarrow (1992), Bank and Baum (2004), Cvitanić and Ma (1996)), however, because the high frequency traders are price takers, they do not trade strategically anticipating the impact of their trade on the price.

We will show that under this structure, high frequency traders create abnormal profit opportunities that they exploit to the disadvantage of the ordinary traders. Intuitively, they "front-run" their own collective trades. Or, alternatively stated, in unison the high frequency traders create their own momentum, which generates profitable returns. The remainder of this paper formalizes this basic idea.

3 The Model

We assume a complete filtered probability space $(\Omega, P, \mathcal{F}, \mathbb{F})$ where P is the statistical probability measure. Trading takes place continuously in a frictionless and competitive market over the time horizon $t \in [0, \infty)$. Let S be the market price of a stock that is adapted to the given filtration. We assume that the stock pays no dividends. Also trading is a money market account paying the default free spot rate of interest. Without loss of generality, we also assume that the spot interest rate is zero.[6]

We let H denotes a possible trading strategy of an *ordinary trader*, assumed to be a predictable and admissible self-financing trading strategy in the money market account and stock. Recall that an admissible trading strategy is a trading strategy where the generated value process (expression (5)) below) is bounded below by an arbitrary constant, while a self financing trading strategy has no cash flows after it is initially constructed (see Jarrow and Protter (2008)).

We let X denote the collective admissible and self financing trading strategies of all the high frequency traders, assumed to be an adapted process. In addition, we assume that X is a semimartingale, hence càdlàg, with no continuous part, i.e. it is a pure jump process in the sense that its paths change only by jumps. This feature captures the notion that the signal generating the high frequency trades is a pure jump process, with unanticipated events.

We assume that the stock price process without high frequency traders is given by the process

$$dS_t = S_{t-}\sigma(S_{t-})dZ_t$$

where $\sigma(\cdot)$ is the volatility function, S_{t-} means the left limit of S_s as $s \uparrow t$, and Z is a semimartingale with respect to $\mathbb{F}$. This process can be interpreted as the stock's *fundamental value.*

In the presence of the high frequency traders, we assume that the price process takes the form

$$dS_t = S_{t-}\sigma(S_{t-})dZ_t + \eta(S_{t-})dX_t$$

[6]This is without loss of generality because one can normalize by the value of a traded money market account, see Jarrow and Protter (2008).

where $\eta(\cdot) > 0$ is the sensitivity coefficient of the high frequency traders' quantity impact on the price.

Note that when acting in concert, the high-frequency traders' trades change the market price, with a sensitivity coefficient equal to the function $\eta = \eta(s) > 0$. The component $\eta(S_{t-})dX_t$ can be interpreted as a deviation from fundamental value caused by the high frequency trader activity. This component captures the notion that the demand curve for the stock is downward sloping and it is consistent with the evidence in Zhang (2010) and Kirilenko and Kyle (2011) that the high frequency traders' activities move the market price away from fundamental value and increase the return's volatility.

This liquidity impact is analogous to the impact assumed in the large trader asset pricing literature (see Jarrow (1992), Bank and Baum (2004), Cvitanić and Ma (1996)). In contrast to this literature, however, high frequency traders do not trade strategically anticipating the impact of their trades on the price because they act like price takers. In addition, the large trader literature incorporates the large traders' actions into the drift coefficients and to a lesser extent, the volatility of the price process itself. An excellent example of this is Cvitanić and Ma (1996) where this is done through the use of forward-backward stochastic differential equations. In contrast, we isolate the impact of the actions of the high frequency traders through an additive term in order to differentiate the fundamental price from the market price. This is consistent with the price process construction previously used in Jarrow, Protter, Roch (2010) in the study of price bubbles.

This assumption is consistent with the idea that these high frequency traders see the same signal, which generates the same trade (buy or sell). They trade the instant the signal is observed, since they are high frequency traders. Hence, their trading strategy is only adapted to the signal process and it need not be a predictable process with respect to the filtration generated by S. This difference means that high frequency traders can trade in a more timely fashion with respect to the relevant information than can the ordinary traders. We emphasize, however, that the signal is only implicit in the trading strategies process. The trading strategies themselves, and not the signal, directly impact the market price.

To simplify the notation, and with a small loss of generality, we assume this sensitivity coefficient is constant and equal to one, effectively folding it into the strategy process X itself. That is, we assume for the remainder of the paper that $\eta \equiv 1$, i.e.

$$dS_t = S_{t-}\sigma(S_{t-})dZ_t + dX_t. \tag{1}$$

Depending upon the form of the volatility function, solutions exist for this stochastic differential equation. For example, if $\sigma(x) \equiv 1$, then we are in the case of a stochastic exponential where we even have a solution given explicitly by a formula:

$$S_t = \mathcal{E}_X(Z)_t = \mathcal{E}(Z)_t \left(\int_0^t \frac{1}{\mathcal{E}(Z)_s} d(X_s - [X, Z]_s) \right). \tag{2}$$

We need to impose some additional structure on the high frequency trading strategy process X to insure that the stock price does not go negative. In this regard, these trading strategies need to be bounded below in some manner. These bounds are consistent with, but more restrictive than, these trading strategies being admissible. There are two cases to consider: $\Delta X_s \Delta Z_s \neq 0$ and $\Delta X_s \Delta Z_s = 0$. If $\Delta X_s \Delta Z_s \neq 0$, we assume that

$$\Delta X_s \geq -\left(S_{s-}\sigma(S_{s-})\Delta Z_s\right). \tag{3}$$

And, if $\Delta X_s \Delta Z_s = 0$, we assume that

$$\Delta X_s \geq -S_{s-}. \tag{4}$$

There is one more issue: S can *a priori* become negative through the drift of X. To simplify the analysis, we assume that X has no drift. However, this assumption can be relaxed by imposing alternative restrictions. Finally we assume markets are incomplete, so we do not have uniqueness of a risk neutral measure. We circumvent this problem by assuming that *the market itself chooses a risk neutral measure with which it prices derivatives.* See Jacod and Protter (2010) or Jarrow, Protter, and Shimbo (2010) for the development of this idea. We summarize these assumptions by formalizing our standing hypotheses.

Definition 1 (Standing Hypotheses) *For the remainder of the paper we will assume:*

(a) Z *and* X *are both semimartingales;*

(b) X *changes only by jumps, and moreover* $[X, Z] = 0$*, which implies that* X *and* Z *have no common jumps;*

(c) The equation $dS_t = S_{t-}\sigma(S_{t-})dZ_t$ *has a unique solution given* $S_0 > 0$*;*

(d) The process X *represents the trading strategy of the high frequency traders, and it can be both positive and negative;*

(e) In the presence of high frequency traders, the price process S *changes to follow the evolution of equation* (1)*;*

(f) A trading strategy H *of an ordinary trader is assumed to be predictable, i.e. the limit of a process which is left continuous with right limits, and adapted; for clarity we denote it as* H_{t-} *at time* $t \geq 0$*;*

(g) The notation X^* *denotes the trading strategy of the high frequency traders, where* $X^*_s = X_{s-}$ *for ordinary (not high frequency) trades in the market, but* $X^*_s = X_s$ *(and not* X_{s-}*) when trades are used in a high frequency manner and operate against the process* X *itself, since* X *is known to the high frequency traders, it being their own strategy.*

Given this structure, we can now write down the different traders' value processes generated by their respective self financing and admissible trading strategies.

Lemma 2 (Traders' Value Processes) *Assume the Standing Hypotheses. We have three cases to consider:*

1. *Ordinary traders,* not *in the presence of high frequency traders:*

$$V_H(t) = H_0 + \int_0^t H_{s-} dS_s = H_0 + \int_0^t H_{s-} S_{s-} \sigma(S_{s-}) dZ_s \tag{5}$$

2. *Ordinary traders, in the presence of high frequency traders:*

$$V_H(t) = H_0 + \int_0^t H_{s-} dS_s = H_0 + \int_0^t H_{s-} S_{s-} \sigma(S_{s-}) dZ_s + \int_0^t H_{s-} dX_s. \tag{6}$$

3. *High frequency traders:*

$$\begin{aligned} V_X(t) &= X_0 + \int_0^t X_s^* dS_s \\ &= X_0 + \int_0^t X_{s-} S_{s-} \sigma(S_{s-}) dZ_s + \int_0^t X_s dX_s \\ &= X_0 + \int_0^t X_{s-} S_{s-} \sigma(S_{s-}) dZ_s + \frac{1}{2}(X_t^2 + \sum_{s \le t} \Delta X_s^2). \end{aligned} \tag{7}$$

Note that in both cases the stochastic integral with respect to Z is well defined.

Proof. The use of X^* is explained in the Standing Hypotheses. Only the third equation needs a proof. Consider the integral $\int_0^t X_s dX_s$. By integration by parts, we know that

$$\int_0^t X_{s-} dX_s = \frac{1}{2}(X_t^2 - [X, X]_t).$$

Since $\int_0^t X_s dX_s = \int_0^t (X_{s-} + \Delta X_s) dX_s$, we have that

$$\begin{aligned} \int_0^t X_s dX_s &= \frac{1}{2}(X_t^2 - [X, X]_t) + \sum_{s \le t} \Delta X_s^2 \\ &= \frac{1}{2}(X_t^2 - [X, X]^c + \sum_{s \le t} \Delta X_s^2) \end{aligned}$$

Noting that $[X, X]^c = 0$ a.s. by part (b) of the standing hypotheses (Definition 1) and hence $[X, X]_t = \sum_{s \le t} \Delta X_s^2$ completes the proof. ■

To continue, we need some additional regularity conditions on the market price process to guarantee that the market is well-functioning in the absence of high frequency traders. In particular, we add the following assumption:

Assumption 3 (Well-functioning Markets) *There exists an equivalent probability measure Q which makes the two equations*

$$dS_t = S_{t-}\sigma(S_{t-})dZ_t \tag{8}$$

$$dS_t = S_{t-}\sigma(S_{t-})dZ_t + dX_t \tag{9}$$

into local martingales.

This assumption implies that the economy is well-functioning in the sense that there are no arbitrage opportunities for ordinary traders (proven in the next section).

First we note that the two processes, both labelled S, are not the same in Assumption (3). Second we note that under Assumption 3 we see that Z itself is a local martingale under Q, because

$$dZ_t = \frac{1}{S_{t-}\sigma(S_{t-})}dS_t$$

so that Z can be expressed as the stochastic integral with respect to a local martingale, and hence is itself a local martingale, since the integrand $\frac{1}{S_{t-}\sigma(S_{t-})}$ is left continuous and hence locally bounded. (Note that we are using here that $S_{t-}\sigma(S_{t-})$ is strictly positive; i.e., it is never zero.) Since we now know that Z is a local martingale under Q, the integral on the right side of equation (9) is a local martingale, and thus we must have that X *itself is a local martingale under* Q. To convince the reader that such an X exists, and that there are many such, we provide next a simple example of a Lévy process martingale with this property (Example 4), the construction of which given here being partially inspired by a similar one in the book of von Weizsäcker and Winkler (1990).

Example 4 *We let X be a "purely discontinuous" martingale, a compensated sum of jumps. An example of such a martingale X which is not in L^2 consists of the following: let $N^i, i = 1, 2, 3, \ldots$ be i.i.d. Poisson processes with common parameter λ. Let $(U^i)_n$ be sequences of i.i.d. random variables, all independent from all of the N^i, with $E(U^i) = \mu$ and $E\{(U^i)^2\} = \infty$. Let $(T^i_n)_{n\geq 1}$ be the jump times of N^i and let*

$$Y^i_t = \sum_{n=1}^{\infty} U^i_n 1_{\{t\geq T^i_n\}} - \mu\lambda t$$

so that each Y^i is a compound Poisson process, with the same distribution as $Y^j, j \neq i$, but independent from it. We then define

$$X_t = \sum_{i=1}^{\infty} \frac{1}{i^2} Y^i_t \tag{10}$$

Since each U_n^i is in L^1 we have that

$$\begin{aligned} E(|X_t|) &\leq \sum_{i=1}^{\infty} \frac{1}{i^2} E(|Y_t^i|) \\ &\leq \sum_{i=1}^{\infty} \frac{1}{i^2} \left(\sum_{n=1}^{\infty} \left\{ E(|U_n^i|) 1_{\{t \geq T_n^i\}} \right\} + \mu\lambda t \right) \\ &< \infty \end{aligned}$$

and that X is a convergent sum of a countable number of independent martingales, and is therefore itself a martingale. Note that it has an infinite number of jumps on a compact time interval such as $[0,t]$, and since $U^i \notin L^2$, also $X \notin L^2$ and we have the stronger statement that $[X,X]$ is not in L^1; that is, $E([X,X]_t) = \infty$.

We can make this example even more interesting, by again letting $N^j, j = 1,2,3,\ldots$ be i.i.d. Poisson processes with common parameter λ^*. Let $(V^j)_m$ be sequences of i.i.d. random variables, all independent from all of the N^i, with $E(V^j) = \nu$ and $E\{(V^i)^2\} = \infty$. Then if we set

$$R_t^j = \sum_{m=1}^{\infty} V_m^j 1_{\{t \geq T_m^j\}} - \nu\lambda^* t$$

we have that

$$X_t^* = \sum_{j=1}^{\infty} \frac{1}{j^2} R_t^j \tag{11}$$

Combining (10) *and* (11) *we have that*

$$X_t' = X_t - X_t^* \tag{12}$$

is again a martingale in L^1 but with $E([X',X']_t) = \infty$ and moreover it is an example of a Lévy process.

Finally, and most importantly, we note that if $\mu\lambda = \nu\lambda^*$ then the drifts cancel, and X' has no drift. *This example did not have jumps from Z, but one could modify it to include jumps of Z as well. In this case we would not have $[X,Z] = 0$, but rather $[X,Z]_t = \sum_{s \leq t} \Delta X_s \Delta Z_s 1_{\{\Delta X_s \Delta Z_s \neq 0\}}$.*

4 The Result

This section proves the key theorem in the paper which consists of four interrelated results. The general theory states that an absence of arbitrage is equivalent to the existence of an equivalent probability measure that turns the price process into a sigma martingale, or (better) a local martingale if (for example) the price process is bounded below. See for example [6],[7], or [8]. In the next theorem, we see that under one situation (the important one) we cannot have the existence of such an equivalent probability measure for the 3-vector process; this does not necessarily mean, however, that we have arbitrage, since we have left the precise framework of the theorem within the general theory. Nevertheless the result is indicative of a problem from the standpoint of the absence of arbitrage; the third process cannot be transformed into a local martingale, and is instead a strict submartingale. This implies, at a minimum, the existence of a statistical arbitrage, if not a pure arbitrage.

Theorem 5 (Abnormal Profit Opportunities)

1. *There are no arbitrage opportunities for the ordinary traders.*
2. *The high frequency traders earn abnormal trading profits.*
3. *There exists no equivalent probability measure making both the ordinary and high frequency traders' value processes local martingales.*
4. *If the high frequency traders strategies are predictable processes, then their abnormal trading profits are removed.*

Proof. By Assumption 3, there exists an equivalent probability measure Q such that equations (8) and (9) are local martingales. By the standard Delbaen-Schachermayer theory (1994), for the ordinary traders, no free lunch with vanishing risk (NFLVR) is satisfied for both of these equations. This implies that there is no arbitrage for ordinary traders, since they are limited to those two pricing equations for their strategies, corresponding to equations (5) and (6). This proves result 1.

Under the Q of Assumption 3, expression (7) shows that the high frequency traders' value process is a submartingale, due to the last term in this expression. This proves result 2.

Result 4 follows by recognizing that if one uses $X_s^* = X_{s-}$ in expression (7), the resulting value process is a Q local martingale.

Now, for result 3. To contradict this result, we need to find an equivalent probability measure, call it Q^*, such that under Q^* we have the three equations (5), (6) and (7) are local martingales simultaneously. Using the same argument as used in Remark 3, if such a Q^* exists then Z is a local martingale, and also X is a local martingale. If the extra terms were $\int_0^t X_{s-} dX_s$ this would be trivial, but it is not: it is $\int_0^t X_s dX_s$, and the stochastic integral has an optional integrand, not a predictable one. Therefore because of Lemma 2 we are reduced to considering the process Y given by

$$Y_t = \frac{1}{2}(X_t^2 + [X,X]_t) \tag{13}$$

and Y needs to be a local martingale for the system of all three equations. But Y has the property that $Y_0 = 0$ and $Y_t \geq 0$ for all $t \geq 0$. Therefore if Y were to be a martingale or local martingale under any probability measure $Q^\star$ it would of necessity have the property that $Y_t = 0$ for all t, a.s. That is, $Y \equiv 0$ a.s., which is a contradiction as long as $X \not\equiv 0$. ■

Looking back to expressions (5) and (7) we can understand why the first two results in this theorem are true. For the ordinary traders, their value processes are martingales because they can only trade *after* the information revealed by the high frequency traders' trades is known. In contrast, the high frequency traders's trades *cause* the price movement. Hence, the value process has a positive component (the last term) because when the high frequency traders buy, the price rises; and, when they sell, the price falls. It is their trade that causes the price movement, generating self-fulfilling profits. This is akin to market manipulation generated by large traders, except that in this case the high frequency traders' profits are unknowingly generated via a coordinating mechanism - trading based on the same market signals - instead of strategic trading where they anticipate the impact of their trades on the price.

The next two results are complementary. The first shows that we cannot exclude the possibility that the high frequency traders' abnormal profits are in fact arbitrage opportunities. This will depend on the specifics of the market price process and the high frequency traders' strategies. The last result clarifies the role of speed in the high frequency traders' strategies. If the high frequency traders' strategies are predictable, and not optional strategies, then their abnormal trading profits disappear. This insight explains the high frequency traders' race for reduced execution speed by locating computers closer and closer to the exchange trading floor.[7]

Unfortunately, these abnormal trading profits generated by the high frequency traders are at the expense of the ordinary traders. As such, in this respect, high frequency trading introduces a market inefficiency. To the extent that this high frequency trading advantage is due to preferential treatment in the execution of market orders, it should eliminated by regulation. To the extent that the speed of execution is due to financial resources, it may be impossible to remove.

5 Conclusion

An open question in the financial literature is whether high frequency trading improves or impedes market efficiency. The existing empirical literature, although mixed, supports the conclusion that it improves market efficiency by reducing bid/ask spreads and market volatility while making markets more liquid. In contrast, we provide a model with no bid/ask spreads and perfect liquidity, yet the introduction of high frequency trades both increases market volatility and generates abnormal profit opportunities for the high frequency traders at the expense of the ordinary traders. An open and important research question motivated by this paper is whether our model of the price process provides a good approximation to actual market prices. We conjecture that it does, and we look forward to the resolution of this conjecture in subsequent research.

[7]See New York Times, January 1, 2011, "The new speed of money, reshaping markets," Graham Bowley; (http://www.datacenterknowledge.com/archives/2010/12/14/speed-of-light-constrains-high-speed-traders), December 14, 2010; Rich Miller.

References

[1] P. Bank and D. Baum, 2004, Hedging and Portfolio Optimization in Illiquid Financial Markets with a Large Trader, *Mathematical Finance*, 14 (1), 1 - 18.

[2] J. Brogaard, 2010, High frequency trading and its impact on market quality, working paper, Northwestern University.

[3] J. Castura, R. Litzenberger, R. Gorelick, and Y. Dwivedi, 2010, Market efficiency and microstructure evolution in U.S. equity markets: a high frequency perspective, working paper, RGM Advisors, LLC.

[4] J. Cvitanić and A. Kirilenko, 2010, High frequency traders and asset prices, working paper, Caltech.

[5] J. Cvitanić and J. Ma, 1996, Hedging options for a large investor and forward-backward SDEs, *Ann. Applied Proba.*, **6**, 370-398.

[6] F. Delbaen and W. Schachermayer, 1994, A general version of the fundamental theorem of asset pricing, *Mathemathische Annalen*, 300:463–520.

[7] F. Delbaen and W. Schachermayer, 1998, The fundamental theorem of asset pricing for unbounded stochastic processes, *Mathemathische Annalen*, 312:215-250.

[8] F. Delbaen and W. Schachermayer, 2006, *The Mathematics of Arbitrage*, Springer-Verlag, Heidelberg.

[9] J. Hasbrouck and G. Saar, 2010, Low-latency trading, working paper, Cornell University.

[10] T. Hendershott, C. Jones and A. Menkveld, 2008, Does algorithmic trading improve liquidity?, forthcoming, *Journal of Finance*.

[11] T. Hendershott and P. Moulton, 2010, Automation, speed and stock market quality: The NYSE's hybrid, working paper, U. of California, Berkeley.

[12] T. Hendershott and R. Riordan, 2009, Algorithmic trading and information, working paper, U. of California, Berkeley.

[13] J. Jacod and P. Protter, 2010, Risk Neutral Compatibility with Option Prices, *Finance and Stochastics*, **14**, 285-314.

[14] R. Jarrow, 1992, Market Manipulation, Bubbles, Corners and Short Squeezes, *Journal of Financial and Quantitative Analysis*, September, 311 - 336.

[15] R. Jarrow and P. Protter, 2008, An introduction to financial asset pricing, *Handbooks in OR & MS*, eds. Birge and Linetsky, vol 15, Elsevier.

[16] R. Jarrow, P. Protter, and K. Shimbo, 2010, Asset Price Bubbles in an Incomplete Market, *Mathematical Finance*, **20**, 145-185.

[17] R. Jarrow, P. Protter and A. Roch, 2010, A liquidity based model for asset price bubbles, working paper, Cornell University.

[18] A. Kirilenko, A. Kyle, M. Samadi, T. Tuzun, 2011, The Flash Crash: The Impact of High Frequency Trading on an Electronic Market (January 12, 2011). Available at SSRN: http://ssrn.com/abstract=1686004

[19] P. Protter, *Stochastic Integration and Differential Equations, Second Edition, Version 2.1*, 2005, Springer-Verlag, Heidelberg.

[20] R. Riordan and A. Storkenmaier, 2009, Latency, liquidity and price discovery, working paper, School of Economics and Business Engineering, Karlsruhe Institute of Technology, Germany.

[21] H. Stoll, 2006, Electronic trading in stock markets, *Journal of Economic Perspectives*, 20 (1), Winter, 153 - 174.

[22] H. von Weizsäcker and G. Winkler, 1990, *Stochastic Integrals*, Vieweg & Sohn, Braunschweig.

[23] X. F. Zhang, 2010, High frequency trading, stock volatility and price discovery, working paper, Yale University.