Poison Pill or Satiating Lozenge?

Order flow toxicity and algorithmic market makers

by

Nicholas Garmulewicz

An essay submitted to the Department of Economics

in partial fulfillment of the requirements for

the degrees of Master of Arts

Queen's University

Kingston, Ontario, Canada

August 2015

Contents

1 Related Literature **4**

1.1 Non-Traditional Market Making: Algorithmic Trading 4

2 The Model: **5**

2.1 Probability of Informed Trading (PIN): Foundation of VPIN 5

2.2 Information Arrival and Clock Time vs. Event Time 10

2.3 VPIN: A new way forward 12

3 Data and Summary Statistics: **15**

3.1 Estimating order flow toxicity: VPIN 19

4 Results: VPIN and market characteristics in volume time **21**

5 Formal model of AT trade participation: **27**

6 Conclusion **31**

7 Appendix: Data and Matching Process **32**

Financial markets are now different in fundamental ways. In recent years information technology has increased in both its speed and the capacity leading to a revolution in how financial markets operate. Growth and development in IT have increased the speed at which traders receive information about the market, process it, and resubmit orders. As a result of technology revolutionizing how financial markets operate. a new breed of traders who rely solely on technology and automation to execute their orders has risen, algorithmic traders (ATs). ATs encompass a large variety of traders and strategies varying from non-traditional market making to the broadly cited 'cheetah traders' or high frequency trading (HFT) firms. However, improvements in technology has not only made trading faster but has fundamentally changed the way markets are structured.

Market participants who have speed advantages has always been a part of of financial markets. As early as the 1800s legend has it that Nathan Mayer Rothschild used racing pigeons to front run his competitors and trade on the news of Napoleons's defeat a full day ahead of His Majesties official messengers, Gray and Aspey (2004). This begs the question, what is different this time around? Following the insight of Easley et al. (2012a) the paradigm shift in financial markets is the way silicon traders (ATs) measure time.

With the technology to process information at incredibly fast speeds and respond by submitting large order volumes across multiple markets, chronological time fails to capture how ATs engage in sequential strategic trading. Rather, ATs design strategic trading strategies using event time, the natural way computers process information, Easley et al. (2012a). This paradigm shift has changed how ATs engage in trade, to how liquidity is provided, and to how efficient prices are discovered, O'Hara (2015). Using this new lens with which we can view financial markets, I analyze adverse selection and order flow toxicity measured by event time and analyze how AT trading activity responds.

The rise of high speed silicon traders has caused, academics, regulators, and market participants to worry about negative externalities from high speed trading. Such negative externalities are, front running, increased volatility, decreased liquidity, market manipulation, and adverse selection. While there is general, but not universal, agreement that AT market making enhances market quality by reducing spreads and enhancing informational efficiency (Jones (2013); Hendershott and Riordan (2013); Carrion (2013)), little work has been done on formally analyzing if ATs increase adverse selection and how ATs respond by changes in order flow toxicity.

Intuitively, adverse selection in the order flow is the “Natural tendency for passive orders to fill quickly when they should fill slowly and fill slowly when they should fill quickly.”, Sofianos (2008). That is, adverse selection can be thought of as informed traders taking advantage of uniformed traders by trading on their superior information of where future prices are moving. Viewed through the lens of a market maker (traditional or non traditional) order flow is regarded as toxic when it adversely selects market makers who may be unaware that they are providing liquidity at a loss, Easley et al. (2012a). The order arrival process, that is a proxy for new information coming to the market, makes trade volume informative for subsequent price moves in general and order flow toxicity in particular. Volume-synchronized Probability of Informed Trading (VPIN) is a metric for estimating order flow toxicity based on a process subordinated to volume arrival.

To study the relationship between adverse selection in the order flow and AT trading activity I use electronic limit order book data from the Deutscher Aktien Index (DAX) stocks (the 30 largest market capitalization stocks) that is traded on the Deutsche Boerse (DB). While direct identification of ATs is often not possible in other markets the DB uniquely identifies if an order was generated by an algorithm. Using this unique dataset I analyze the link between adverse selection in the order flow and AT trade participation.

Using Volume-synchronized Probability of Informed Trading (VPIN) metric to capture order flow toxicity and adverse selection, I find that AT trade activity is negatively correlated with adverse selection. By analyzing AT trade activity in event time the analysis is consistent with the way ATs view sequential strategic trading in the limit order book. Using updated methodology I further find that AT trade activity is more sensitive to market conditions when order flow toxicity, as measured by VPIN, is high. As a result, ATs closely monitor the market and when adverse selection risks are high they decrease their trade participation.

Section 1 gives further insight to the rise of algorithmic trading. Section 2 gives the methodology of the probability of informed trading metric that is the foundation of VPIN. Section 3 describes the data and summary statistics and Section 4 gives results for the link between VPIN and AT trade participation. Finally, Section 5 presents a formalization of AT trade participation and Section 6 concludes.

1 Related Literature

1.1 Non-Traditional Market Making: Algorithmic Trading

Algorithmic trading, while seemingly precise, actually describes a large and diverse set of activities and behaviours of market participants that interact with the market at extremely low latency and are completely automated[1]. There has been documentation in the literature that ATs pursue a wide array of strategies that range from market making, statistical arbitrage, and opportunistic trading, O'Hara (2015). Regardless of the type of strategies that ATs pursue, their presence is undeniable. According to SEC (2010) estimates of HFT typically exceeded 50% of total volume in U.S. listed equities and concludes that, "[b]y any measure, HFT is a dominant component of the current market structure and likely to affect nearly all aspects of its performance".

The flash crash of May 6, 2010 initially pushed HFTs and algorithmic trading into the public eye and to the attention of market regulators. During May 6th 2010, the S&P 500 fell by 10% within 15 minutes before recovering. Commentators initially criticized HFT firms in exacerbating the crisis, but the CFTC and SEC were able to identify most of the HFT firms that were active during the crisis were providing liquidity and stabilizing prices until they were eventually overwhelmed and had to liquidate their positions, Jones (2013). These findings are in line with the general consensus that low latency trading improves market quality. Brogaard et al. (2014) show that HFTs improve price discovery, Hasbrouck and Saar (2013), and Menkveld (2013) show that HFTs improve spreads in both certain and uncertain market conditions.

HFT market making differs from traditional market making in that it is often implemented across and within markets, making it akin to statistical arbitrage. Intuitively, non-traditional market making (AT) uses historical correlation patterns in price changes to move liquidity between securities or markets. For example, consider market making within a market. If statistically an upward price tick in stock A is generally followed by a similar upward price tick in stock B, then a high frequency market maker would want to sell stock A and buy stock B (essentially striving to buy low/ sell high). This involves submitting an order at the ask in

[1]Latency is the time it takes to send market data such as orders to an exchange or server. In other words, it is the time it takes to learn about and event (e.g change in a bid), generate a response, and have the exchange act on the response, Hasbrouck and Saar (2013). Latencies are typically measured on the millisecond scale (one thousandth of a second), microsecond scale (one millionth of a second), or nano-seconds (one billionth of a second). To put it in perspective it takes about 300ms for the human eye to blink.

stock A and at the bid in stock B. The process becomes far more complex when it goes across markets and as a result requires a close monitoring of market conditions and tight risk management practises.

While a small portion of AT traders use extremely low latency and opportunistic trading to front-run orders AT firms typically act as market makers, providing liquidity to position takers by placing passive orders at various levels across the electronic limit order book, Easley et al. (2012a). Rather than making profits on directional bets, AT market makers work to earn margins on large trading volume. For non-traditional algorithmic market makers to maintain profitability, they must have the ability to limit position risk, which is greatly affected by the ability to control adverse selection risk. ATs are more likely to trade in the direction of permanent price changes and trade against transitory price movement indicating that profitable AT market makers closely monitor the market for adverse selection risk, Brogaard et al. (2014). However, there has been little formal analysis on the link between changes in adverse selection and AT trade participation.

2 The Model:

2.1 Probability of Informed Trading (PIN): Foundation of VPIN

To evaluate the importance of adverse selection in the order flow, measuring the probability of information based trading (PIN) is one of many models that tries to capture the process of information arrival. PIN at its foundation is a measure to capture the asymmetry in the order flow between informed traders and uniformed traders. The model builds on the foundations of Easly and O'Hara (1987), and Easley and O'hara (1992) with the original formulation of PIN introduced by Easley et al. (1996). As a seminal market micro-structure model, PIN has been empirically tested in a wide array of subjects ranging from trading venue choice Grammig et al. (2001) to the effects of analyst coverage, Easley et al. (1998) . While PIN is not directly observable from the order flow its theoretical parameters can be estimated using a maximum likelihood approach (MLE). In a high frequency world numerical estimation of PIN has become increasingly difficult with convergence problems and accuracy of conventional trade classification algorithms that it relies upon, Easley et al. (2012a). Volume-synchronized Probability of Informed Trading (VPIN) is a response to some of these problems as it doesn't require numerical estimation of intermediate parameters and can be derived directly from the order flow.

At its core, PIN is a micro-structure model that views trading as a game between liquidity providers and traders (position takers) that is repeated over $i = 1, ..., I$. As a repeated game, how the investors information set is updated is a critical part of the model. At the beginning of each period nature chooses wether an information event occurs with an independent probability α and if an event doesn't occur with probability $(1-\alpha)$. The information event can be viewed as information that is relevant to the future value of the asset, i.e information that can effect company prospects if the security is an equity. Expanding upon the type of information that comes to market, we can differentiate if the news has a positive effect or a negative effect on the securty's valuation. We can specify that if an information event occurs, there is a $(1-\delta)$ probability of good news occurring and a δ probability that bad news occurs. If the information is 'good' news traders know that by the end of the trading period the asset will be worth $\bar{S}_i$. If the information is 'bad' news the valuation of the asset at the end of the period will be $\underline{S}_i$. Furthermore, to prevent any arbitrage possibilities, $\bar{S}_i > \underline{S}_i$.

The next step in PIN is to model how heterogeneous traders come to the market to trade, i.e informed and uninformed traders. After an information event occurs or does not occur, trading for the period begins with traders arriving to the market throughout the trading period according to a Poisson process. After an information event occurs it is assumed that only a portion of traders will have information that they can then use to earn rents from uninformed traders. Rents are earned by informed traders who use their information about the assets ending period value. That is, informed traders will buy if they see good news, and sell if they see bad news. To capture the arrival rate of traders orders during period with an information event, orders from informed traders arrive at rate μ. Wether there is an information event (good or bad) or no new information coming to market uninformed traders arrive at a rate of ε_i where $i = [B, S]$ denotes uninformed buyers or sellers.

PINs strength lies in how it relates the observable market outcomes to the unobservable information and order processes that underly trading decisions. Intuitively, the model interprets the normal level of buys and sells in a stock as uninformed trade, and it uses that data set to identify the rate of uninformed order flow (ε). Therefore, abnormal order flow can be interpreted as informed trading and is used to then identify the number of informed traders coming to market (μ). Finally, the number of periods in which there is abnormal buy or sell volume is used to identify both the probabilities of information events occurring (α) and the type of event (δ).

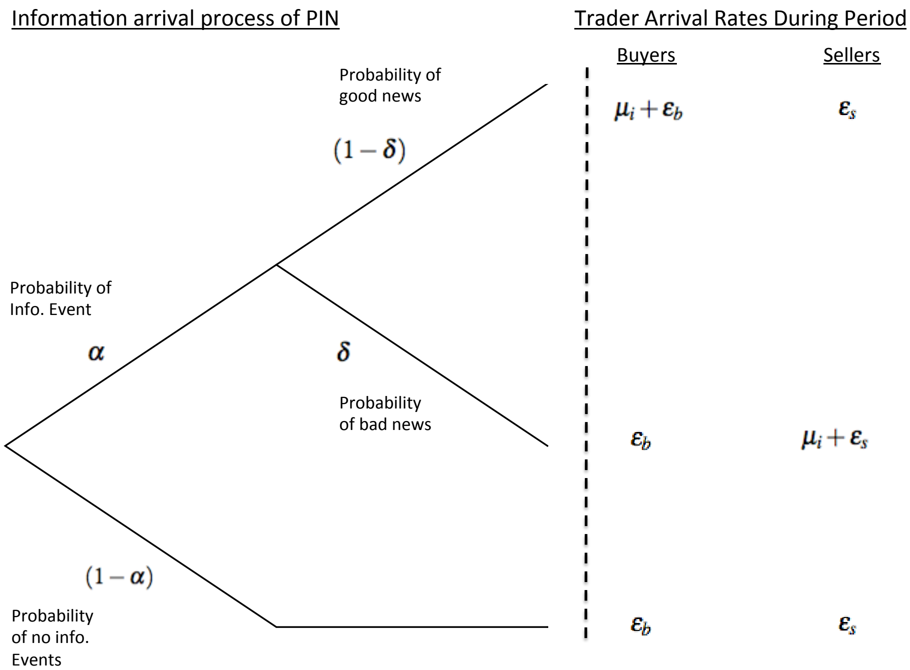

Figure 1: Information and order arrival process of PIN

The likelihood of observing k orders from informed traders in a day with an information event is,

$$\frac{\mu_i^k e^{-\mu_i}}{k!}$$

where e is the exponential function. A liquidity provider uses their knowledge of these unobservable parameters to determine the prices that they post at the bid and at the ask.

The bid ask spread in is partially generated because the liquidity provider does not know wether the counterparty to their trade is informed or not and the spread must cover the risk of adverse selection and inventory risk.

The bid/ask spread can be viewed as the difference in the expected value of the asset conditional on someone buying from the liquidity provider and the expected value of the asset conditional on someone selling to the liquidity provider, Easley et al. (2012a).

Liquidity providers observe trades and in the the PIN framework are modelled as if they use Bayes' rule to update their beliefs of the level of informed traders in the market, i.e order flow toxicity. Let,

$$P(t) = (P_n(t), P_b(t), P_g(t))$$

be the be the parameters related to beliefs about, no, bad, and good information events occurring. With the PIN framework viewing each trading period as a repeated game, beliefs about the information set at $i = 0$ is,

$$P(0) = ((1-\alpha), \alpha\delta, \alpha(1-\delta))$$

.

To determine the bid/ask spread at time t, the liquidity provider updates their information set conditional on the arrival of an order of the relevant type. Here arrival of orders to the market are proxies for new information coming to the market. Therefore, the expected value of the asset conditional on time t is,

$$E[S_i|t] = P_n(t)(\delta \underline{S}_i + (1-\delta)\bar{S}_i) + P_b(t)\underline{S}_i + P_g(t))\bar{S}_i$$

Using this characterization of the expected value of the asset, the bid is the expected value of the asset conditional on someone wanting to sell the asset to the liquidity provider,

$$B(t) = E[S_i|t] - \frac{\mu P_b(t)}{\varepsilon_b + \mu P_b(t)}(E[S_i|t] - \underline{S}_i)$$

.

Similarly the ask can be defined as the expected value of the asset conditional on someone wanting to buy from the liquidity provider.

$$A(t) = E[S_i|t] + \frac{\mu P_g(t)}{\varepsilon_s + \mu P_g(t)}(\bar{S}_i - E[S_i|t])$$

.

If there are no informed traders in the market ($\mu = 0$) then a trade carries no information and the bid and the ask are equal to their prior expected value of the asset ($E[S_i|t]$). At the other extreme, if the market is only composed of fully informed traders ($\varepsilon_s = \varepsilon_b = 0$) then $B(t) = \underline{S}_i$ and $A(t) = \bar{S}_i$, the final value of the asset if the information is bad or good respectively. In a fully informed market, traders won't trade and the market shuts down.

Formally, the bid/ask spread is,

$$\begin{aligned}
\Sigma(t) &= A(t) - B(t) \\
\Sigma(t) &= \frac{\mu P_g(t)}{\varepsilon_s + \mu P_g(t)}(\bar{S}_i - E[S_i|t]) + \frac{\mu P_b(t)}{\varepsilon_b + \mu P_b(t)}(E[S_i|t] - \underline{S}_i)
\end{aligned}$$

The first term in the second expression can be interpreted as the probability that a sell is an information based trade, and the second term can be interpreted as the expected loss to an informed seller. If the spread in the initial period, Σ, has the form where there is a symmetric amount of uninformed buyers and sellers ($\varepsilon_b = \varepsilon_s = \varepsilon$) and good and bad information arrives to the market with equal probabilities ($\delta = 1 - \delta$),

$$\begin{aligned}
\Sigma &= \frac{\mu\alpha(1-\delta)}{\varepsilon + \mu\alpha(1-\delta)}(\bar{S}_i - E[S_i]) + \frac{\mu\alpha\delta}{\varepsilon + \mu\alpha\delta}(E[S_i] - \underline{S}_i) \\
&= \frac{\mu\alpha\delta}{\varepsilon + \mu\alpha\delta}(\bar{S}_i - \underline{S}_i) \\
&= \frac{\alpha\mu}{2\varepsilon + \alpha\mu}(\bar{S}_i - \underline{S}_i)
\end{aligned}$$

From this expression of the bid ask spread we can see that a key component is the probability that an order is from an informed trader. The first term in this expression can be interpreted as the probability that an opening trade in a period is information based, also known as PIN.

$$PIN = \frac{\alpha\mu}{2\varepsilon + \alpha\mu}$$

PIN is a measure of the fraction of orders that come from informed trades ($\alpha\mu$) compared to the overall order flow ($2\varepsilon + \alpha\mu$).

2.2 Information Arrival and Clock Time vs. Event Time

When viewing equity markets we can think about information events as being events relevant to the future prospects of a company that directly impacts its current market value. If markets are fully efficient and security valuation reflects all available information, the value of a security should converge to its full information value as informed traders seek to profit on their information by trading, Easley et al. (2012a). Since market makers can take various positions in a stock depending on their holdings or the market characteristics, fluctuations in the future value of securities will greatly effect their profitability.

Market makers need to understanding what drives future price movements to stay in the market and provide liquidity. Since relevant information to conventional market makers will arrive at a humanly perceptible time, the decisions to engage in liquidity provision happen over the same frequency. However, in a high frequency market where AT traders can react to information at the millisecond level ATs still face the same problem as conventional market makers, albeit at a much smaller time interval.

In comparison to conventional market makers who typically hold inventory for hours or even days, an AT market maker who anticipates holding the stock for minutes is affected by information that influences its value over that interval, Easley et al. (2012a). The change in the reactionary time of ATs initially raised concerns about increased volatility by AT market making and the possibility of adversely selecting non-AT traders. However, recent literature has shown that the increased speed and reaction times of AT market making enhances market quality by reducing spreads and increasing informational efficiency and price discovery, Brogaard et al. (2014).

With ATs reacting at higher speeds the standard definition of an information event needs to be expanded upon. Information events that effect the profitability of AT market makers need not have a direct link to a companies future value but can include information that is relevant to the liquidity demand characteristics of the order book at a millisecond level. This means that information events may occur more frequently, and may have varying importance for the size of future price movements.

One of the more important aspects of high-frequency data is that trades are not equally spaced in terms of time. By taking trades as a proxy for new information coming to market, we can draw the link between between uneven arrival of trades and the uneven arrival of information. This is not a new phenomenon in equity markets with Mandelbrot and Taylor (1967) noting that this was true in equity trading during 1960s. In other words, the arrival of new information to the market triggers waves of decisions by market participants and market makers that then translates into volume bursts as traders either adjust their holdings or inventory, Easley et al. (2012a). However, in the context of high frequency trade data this can have serious impact on the analysis of order flow. If information relevant to different securities arrives at different time throughout the day there will then be a distinct intraday volume seasonality that can lead to volatility clustering and non-stationary volume distributions.

The more relevant a piece of news is, the more volume it attracts. The link between information and volume forms the foundation for changing from a time clock to an event clock measured by trade volume. A volume clocks basis is drawing a sample every time the market exchanges a constant volume (volume bucketing). The sampling of a constant volume attempts to mimic the arrival of information that is equal and comparable across volume buckets. If a particular piece of news generates twice as much volume as another such piece of news, twice as many observations are drawn creating double the weight in the observation, Easley et al. (2008).

Changing from clock time to volume time accomplishes two goals. First it allows for each observation to have comparable information contained in it regardless of clock time taken to exchange that amount volume. When viewed through the framework of PIN it provides the link to measure PIN at an intraday frequency and measure the level of informed traders. But more importantly when analyzing a high frequency market and AT market making activity, measuring order flow by event time is more consistent with how ATs view the market. Computers by their construction operate on an event based clock rather than a chronological clock, i.e a cycle, Easley et al. (2012b). Therefore, it is perfectly natural for an AT trader to formulate a strategy based on an event time as all trading strategies are automated. One of many examples would be that they turn their portfolio every time a fixed number of shares are traded.

The second goal in moving to a volume clock is the improvement in the statistical properties of measurements derived from the order flow. First, by measuring in event time intra-day seasonal effects are removed from the order flow. Secondly, volume bursts due to new information cause the distributions of of volatility, volume, and returns to be leptokurtic and skewed. By measuring market variables by event time, normality and the assumption of IID is partially recovered, Easley et al. (2012b). The guiding principal is that the process measured by an event clock is more likely to be covariance stationary than when measured by chronological time, Hasbrouck (2006). Finally, sampling according to a volume clock addresses data entry errors such as incorrectly assigned time stamps to trades or in the data used here, an error in data merging and manipulation.

2.3 VPIN: A new way forward

To address the new problems faced in a high frequency environment Easley et al. (2012a) define a new measure of adverse selection and probability of informed trading called *Volume-synchronized Probability of Informed Trading* (VPIN). While VPIN has its roots in the sequential trade model of PIN it breaks away from using chronological time and uses a volume clock with an expanded interpretation of information events. To measure adverse selection in the order flow, VPIN uses a measure of order imbalance over an equal volume sample. By using volume time VPIN allows one to divide the trading session into periods of equal information content allowing trade imbalances to have a meaningful economic impact on liquidity providers.

Easley et al. (2008) provide the fundamental link between PIN and VPIN. By allowing for a dynamic microstructure where the arrival rates of informed and uninformed trades vary by time, similar to a GARCH (Bollerslev (1986)) specification on volatility, visible trade data can approximate unobservable PIN parameters. Specifically, for a particular period τ (e.g equal amount of volume, minutes, days) the expected trade imbalance approximately equals the level of informed traders, and the expectation of the total number of trades equals the total number of traders in the market. Formally,

$$PIN = \frac{\alpha\mu}{2\varepsilon + \alpha\mu} \approx \frac{E[V_\tau^S - V_\tau^B]}{E[V_\tau^S + V_\tau^B]} = VPIN$$

where V_τ^S and V_τ^B are the volume of buys and sells over period τ.

There are two key innovations from PIN to VPIN: broadening the definition of an information event, and the use of an event clock measured by equal transactions of volume. First, PIN relies on fundamental information events that relate to the true underlying value of the stock. Informed traders use that information to land on the right side of the order imbalance. VPIN information events include the concept of order flow toxicity that is an encompassing measure that includes a component of informed traders. Adverse selection may come from factors relating to trading in the overall market to specific liquidity characteristics over a certain time period. For example, if the issuing firm of a stock releases information about unexpected future revenue, this information will generally effect the stock price and be relevant to a market maker who holds inventory.

By using a volume clock VPIN more accurately measures levels of intraday order flow toxicity. The PIN model collects daily order imbalances under the assumption of independence and that all information is impounded into the stock at the end of the day. In comparison VPIN calculates order imbalances every time the market exchanges a constant level of volume. Since trade volume is a proxy for information coming to market, volume bucketing mimics the arrival of information of a comparable relevance. As previously noted, by using a volume clock for estimation, VPIN reduces the impact of volatility clustering.

To estimate VPIN, trades need to be signed as buy or sell so the expected trade imbalances can be calculated. Easley et al. (2012a) argue that in a high frequency market, standard itemized classification models such as the Lee-Ready algorithm (Lee and Ready (1991)) are problematic and can lead to misclassification errors. One problem is that reporting conventions in markets could treat orders differently depending if they are buy or a sell. For example the New York Stock Exchange (NYSE) reports one trade if a large sell block crosses multiple buy orders, and reports multiple trades if a large buy order is crossed with multiple sell orders. To remedy trade misclassification problems that can lead to biases in VPIN estimation, Easley et al. (2012a) develop a 'bulk-volume' trade classification. Trades are first aggregated into small time bars (or volume bars) and then a standardized price change is estimated for that time bar. The aggregated trade volume over the small time bar is then signed as buy or sell using a normal distribution and the standardized price change. Each time bar (with signed aggregated trade volume) is then used to fill each of the equal sized volume buckets.

In the original PIN model, order imbalances are observed in terms of the number of buys and sells with the size of the trade not taken into account, Abad and Yagüe (2012). In comparison, with a 'bulk-volume' order classification VPIN takes into account trade size by treating each reported trade as if it were an aggregation of trades of one unit size. For example, if there are 3 one minute time bars with a signed and aggregated trade volume of 1,100, 550, and 800 and a standardized volume bucket of 1,200, to fill the first volume bucket the first entire aggregated time bar would be used and then 100 out of the second time bar. Using bulk volume order classification explicitly puts 'trade intensity' into the analysis.

However, Chakrabarty et al. (2012), and Andersen and Bondarenko (2014) find that standard tick transaction classifications rules (TTR) outperform bulk volume (BV) classifications. Using a large sample of stocks traded on NASDAQ's INET platform Chakrabarty et al. (2012) find that TTRs outperform BV classifications in a large cross-section of stock sizes and for time and volume bars of all sizes. Further, Chakrabarty et al. (2012) find that BV accuracy is adversely effected by episodes of market instability such as high volatility, and the presence of hidden liquidity. To further complicate the matter, Easley et al. (2014) reply to these criticisms by restating their claim of the inherent difficulties of estimating the sign of trades due to the new high frequency microstructure environment[2].

Finally it is important to note that the VPIN estimation procedure produces a series of estimations of order flow toxicity for a time period and doesn't require the estimation of unobservable parameters. In comparison, PIN requires a maximum likelihood estimation of unobservable parameters and provides a single estimate of informed trading for the period (day, month, year), Abad and Yagüe (2012). However, as VPIN allows the capturing of adverse selection risk variations on an intraday level, changes in the raw VPIN measure are not

[2]Both parties in this debate claim that the other have misinterpreted the others results and claims of accuracy. Rather than taking a stand on wether one method is more appropriate than another, I calculate VPIN using a bulk volume classification as initially proposed by, Easley et al. (2012a). Given that I don't explicitly use VPIN to try and measure or predict extreme market turbulence, and rather use VPIN as a measure of adverse selection that effects AT market makers trade participation, I feel that I can skirt this imbroglio.

informative. To see if there is a significant change in the level of order flow toxicity the suggested approach is to compare the raw VPIN measure to its empirical CDF. With the focus of my analysis on how order flow toxicity effects AT trade participation I only highlight the central moments of VPIN and cross-sectional correlations with AT trade participation and other market characteristics.

3 Data and Summary Statistics:

To examine the link between Algorithmic trading and adverse selection in the order flow, as measured by *Volume Synchronized Probability of Informed Trading* (VPIN), I use German stock market data for a cross-section of 30 stocks that have AT activity uniquely identified. The data used is the merging of the Deutsche Boerse (DB) Automated Trading Program (ATP) system order data and Reuters DataScope Tick History data provided by the Securities Industry Research Centre of Asia-Pacific (SIRCA). The SIRCA data contains two separate datasets, trades and orderbook updates. There are 30 stocks in the sample that are constituents of the DAX index between Jan. 1 2008 and Jan. 18 2008 giving a total of 13 trading days.

To generate a single dataset that uniquely identifies trades, orders, and cancelations by Trader Type, algorithmic trader (AT) or everybody else (EE), I recreate the order causing the observed book changes. To identify trades, that look like cancels at the best bid/ask in the recreated message stream of the order book, I matched the SIRCA trade file with the reconstructed message book. Finally, to generate the complete data set I merged the generated events (insert, cancel, trade) with the DB-provided AT order data set. The final dataset contains all orders submitted by ATs and EE for the first 10 levels on the bid and the ask side. See the Appendix for more information.

To better capture how differences in adverse selection and market characteristics vary with AT activity I broke the cross-section of 30 DAX stock into 3 AT activity deciles, low, medium, and high. To determine the cutoff for AT activity deciles I first calculated the VPIN metric for each stock using a volume clock that is calibrated to have 1/50th of average daily trading volume in each volume bucket. Using the volume clock I then calculated the daily average percentage of AT trade activity in each volume bucket. This identification strategy allows me to see the difference between stocks that have high AT trade participation and low AT trade participation based in event time, i.e every time an equal amount of information comes to market. The cross-section of stocks are then ranked from lowest to highest AT trade activity and then split into 3 equally weighted deciles.

Table 1 gives the breakdown of trades and messages generated by ATs and EE for the full sample and the three AT activity deciles. Looking at the full sample in Panel A and B of Table 1 ATs generate close to 59% of the trades and over 76% of messages. The ATs significantly higher share of limit orders is in line with Hendershott and Riordan (2013) who use the same dataset and find that ATs submit more orders than they execute either by submitting uncompetitive orders away from the best prices or by cancelling and replacing orders close to the best bid or ask. Looking at the difference between Decile 1 (low AT activity) and Decile 3 (high AT activity), ATs share of both trades increases marginally but their share of messages actually decreases. This indicates that looking at the mean trade and message participation using clock time (daily average) clouds how ATs actually participate in response to equivalent information set updates.

Table 1: Trade participation and order book message generation by trader type for 30 DAX index stocks from Jan. 1 2008 to Jan. 18 2008. Panel A gives the percentage of Algorithmic Traders (AT) trades versus everybody else (EE). Panel B gives the percentage of AT nonmarketable limit order participation versus EE. AT activity deciles are calculated as the percentage of AT trade activity in 1/50 of average daily volume buckets. Decile 1 is low AT activity, Decile 2 is medium AT activity, and Decile 3 is high AT activity. Each decile is composed of 10 equally weighted stocks.

Panel A: Trades

	AT	EE	Total
Decile 1	57.51(%)	42.49(%)	100(%)
Decile 2	61.07(%)	38.93(%)	100(%)
Decile 3	58.06(%)	41.94(%)	100(%)
Total	58.85(%)	41.15(%)	100(%)

Panel B: Messages

	AT	EE	Total
Decile 1	76.75(%)	23.25(%)	100(%)
Decile 2	78.17(%)	21.83(%)	100(%)
Decile 3	72.89(%)	27.11(%)	100(%)
Total	76.33(%)	23.67(%)	100(%)

Table 2 describes the 30 stocks in the DAX index and each AT activity decile. The mean of the descriptive variables are calculated daily for each stock over the sample period (30 stocks for 13 trading days giving 390 observations). Table 2 reports the mean, standard deviation, maximum, and minimum of these 390 observations. Panel B through Panel D reports the same statistics for the respective AT activity deciles (10 stocks 13 trading days giving 130 observations).

Looking at the full sample size and the 3 AT activity deciles, the DAX stocks are actively traded. With a full sample daily mean of 7007 trades and an average trade size of 925 shares, the securities in the sample are quite liquid. The quoted spread is the spread between best bid and best ask and is calculated when the trade occurs. An average quoted spread of 3.62 bp is comparable to other large and liquid securities markets, Hendershott and Riordan (2013).

Looking at the difference between Panel B: Decile 1 and Panel D: Decile 3, on average low AT activity stocks have a higher price, very little difference in average trade size, and have a higher average quoted spread and effective spread. The higher AT activity in the lower average spread stocks is consistent with previous results of ATs being more sensitive to market conditions compared to non-AT traders, Hendershott and Riordan (2013). Since trades allow market participants to change their holdings and market makers to adjust their inventory, large liquidity demanding orders placed in periods of low liquidity can have a significant price impact and signal to uniformed traders the arrival of new information. Therefore, breaking up a large order into multiple smaller orders and submitting them conditional on market conditions can improve the quality of trade executions while simultaneously camouflaging trading intentions , Barclay and Warner (1993). Therefore, ATs who monitor market conditions much closer than human traders (EE) are more likely to submit orders that have a greater sensitivity to market conditions such as the bid ask spread and order flow toxicity.

Table 2: Summary statistics of average daily trade data for full sample and three AT activity deciles of the 30 constituents of the DAX index between Jan. 1 2008 and Jan. 18 2008. The data set combines Deutsche Boerse (DB) Automated Trading Program (ATP) system order data and SIRCA trade, quote, and order data. Variables are averaged per stock per day (390 observations), and the mean, standard deviation, maximum, and minimum of these averages are reported. AT activity deciles are calculated as the percentage of AT trade activity in 1/50 of average daily volume buckets. Decile 1 is low AT activity, Decile 2 is medium AT activity, and Decile 3 is high AT activity. Each decile is composed of 10 equally weighted stocks.

Panel A: Full Sample

Variable	N	Mean	Std. Dev	Max	Min
Trades per day	390	7007	4144	26861	1570
Price (euros)	390	67.85	42.28	155.17	6.44
Size of trade	390	925	1087	6929	184
Quoted spread (bp)	390	3.62	1.87	9.99	1.07
Effective spread (bp)	390	1.79	0.94	5.07	0.52

Panel B: Decile 1 (Low AT activity)

Variable	Obs	Mean	Std. Dev.	Max	Min
Trades per day	130	9301	4897	26861	2616
Price (euros)	130	77.20	41.44	147.28	14.89
Size of trade	130	1099	1448	6929	213
Quoted spread (bp)	130	3.99	2.17	9.99	1.07
Effective spread (bp)	130	2.00	1.10	5.07	0.52

Panel C: Decile 2 (Medium AT activity)

Variable	Obs	Mean	Std. Dev.	Max	Min
Trades per day	130	7218	3434	17861	1734
Price (euros)	130	86.76	43.41	155.17	21.64
Size of trade	130	516	308	1911	184
Quoted spread (bp)	130	4.03	1.38	8.86	1.78
Effective spread (bp)	130	1.99	0.69	4.33	0.82

Panel D: Decile 3 (High AT activity)

Variable	Obs	Mean	Std. Dev.	Max	Min
Trades per day	130	4502	2084	11076	1570
Price (euros)	130	39.59	23.15	92.92	6.44
Size of trade	130	1160	1056	5262	184
Quoted spread (bp)	130	2.85	1.73	8.64	1.12
Effective spread (bp)	130	1.39	0.85	4.04	0.52

3.1 Estimating order flow toxicity: VPIN

To measure levels of order flow toxicity I estimate VPIN for each of the 30 DAX stocks that make up the sample. The VPIN metric encapsulates both levels of adverse selection and unfavourable liquidity demand characteristics, both which are critical to liquidity providers that are monitoring their trading risk. Following the methodology of Easley et al. (2012a) VPIN estimation uses the following information, time of trade, price, and volume exchanged. Using the volume clock that VPIN is estimated over I further calculate order book characteristics to examine what drives AT trading activity.

The first step in estimating VPIN begins with aggregating trades in time bars. Following Easley et al. (2012a) that posit initially aggregating by time bars leads to better identification of of buy and sell volume, and thus better order flow toxicity measurement, I aggregate trades into 1 minute time bars. Over the time bar I aggregate volume of all trades that occur during the 1 minute interval and then calculate the price change for this period. The standardized price change between the beginning and the end of the interval is used to then calculate the percentage of buy and sell volume. Aggregation into 1 minute time bars mitigates the effects of order splitting that leads to problems in standard order classification algorithms. Using the standardized price change allows for bulk volume classification in probabilistic terms.

Calculating the volume buckets, which forms the foundation of the volume clock, is the next step in calculating VPIN. The implicit goal behind volume bucketing is to spilt the order flow into intervals of homogeneous information content. Volume bucket size is calculated as 1/50th of daily average volume (in shares traded) as in Easley et al. (2012a). While the size of the volume bucket is a key component of calculating VPIN, different sizes of volume buckets cause order flow toxicity to be viewed at different levels. Abad and Yagüe (2012) show that by estimating VPIN using one volume bucket (average daily trading volume) that VPIN estimates closely resemble PIN[3]. Buckets are filled by adding trade volume for consecutive time bars until 1/50th of daily average volume is completed. If the volume in the last time bar is in excess of what is required to complete the volume bucket the excess trade volume spills over into the next bucket. In general multiple time bars generate a volume bucket, however, if there is a large enough trade volume in one time

[3] Abad and Yagüe (2012) posit that the size of the volume bucket is the more relevant variable of VPIN metric procedure. They argue that Easley et al. (2012a) choice of a volume bucket size of 1/50th of daily trading volume is arbitrary. They show that by varying the size of the volume buckets what VPIN captures may be different components of the adverse selection risk faced by HF liquidity providers. However, it seems unclear what kind of toxicity is measured when a lower number of buckets is employed.

bar then a bar (or a portion of one) could fill one volume bucket.

Once the volume bucket is completed the time bar volume is classified as buy or sell initiated under probabilistic terms. Easley et al. (2012a) show that classifying bulk volume by multiplying the volume in the time bars by the normal distribution evaluated at the the standardized price change can improve the accuracy of trade classification in a high frequency environment where order splitting and high speed cancelations are common. Let,

$$
\begin{aligned}
V_\tau^B &= \sum_{i=t(\tau-1)+1}^{t(\tau)} V_i \cdot Z\left(\frac{P_i - P_{i-1}}{\sigma_{\Delta P}}\right) \\
V_\tau^S &= \sum_{i=t(\tau-1)+1}^{t(\tau)} V_i \cdot \left[1 - Z\left(\frac{P_i - P_{i-1}}{\sigma_{\Delta P}}\right)\right] = V - V_\tau^B
\end{aligned}
$$

where V_τ^B is buy initiated volume in volume bucket τ and $t(\tau)$ is the index of the last time bar included in the τ volume bucket. Z is the cumulative distribution function (CDF) of the standard normal distribution, and $\sigma_{\Delta P}$ is the estimated standard deviation of the price changes between time bars over the whole sample. I used a rolling sample of 50 volume buckets to calculate the sample standard deviation. If there is a positive price change over the volume bucket, then volume is weighted more towards buys then sells, and vice versa with a negative price change. In other words, the higher the absolute price change, the higher the order imbalance between buy and sell initiated trade volume. Order imbalance (OI) is simply the absolute difference between buy and sell initiated volume in each volume bucket[4].

Finally, to calculate the VPIN metric I chose a sample of 50 volume buckets as the length to calculate VPIN from. Following the link that Easley et al. (2008) established,

$$
VPIN = \frac{\alpha\mu}{2\varepsilon + \alpha\mu} \approx \frac{E[V_\tau^S - V_\tau^B]}{E[V_\tau^S + V_\tau^B]} = \frac{\sum_\tau^n |V_\tau^S - V_\tau^B|}{n \cdot V}
$$

where VPIN is simply the average order imbalances in the sample. The sum of the order imbalances over

[4] Andersen and Bondarenko (2014) show that signed order imbalances actually contain relevant information in measuring real time market stress indicators.

the sample is a proxy for the expected trade imbalance and the sample size times the volume bucket size is the expected number of trades. The VPIN metric is updated over the entire stock by a rolling window process over each volume bucket. With a given sample size of 50, the first VPIN measure is captured over the first 50 volume buckets, but once bucket #51 is filled the new VPIN observation is calculated using bucket #2 to #51. The VPIN metric is updated in volume time for two reasons. First, volume time mimics the arrival speed of information to the market. With trade volume used as a proxy for new information, this goal is achieved. Second, volume time allows for updating of VPIN with the arrival of comparable amounts of information, Easley et al. (2012a).

4 Results: VPIN and market characteristics in volume time

Table 3 shows the estimated VPIN and descriptive market characteristic variables using a volume clock for 30 stocks in the DAX index. VPIN is calculated using a a rolling sample length of 50 volume buckets, price volatility is the variance in prices over the volume bucket, absolute return is calculated as the absolute volume bucket net return, AT percentage is calculated as the percentage of AT trade participation of the corresponding volume bucket, depth is the average sum of the depth placed at the best bid and best ask, and quoted spread and effective spread are defined as the average respective spreads over the volume bucket. The mean, standard deviation, maximum, and minimum for each of these variables for each stock is then calculated. Table 3 reports the cross-sectional mean of the average summary statistics for each stock for the full sample and three AT activity deciles.

Looking at the full sample summary statistics the mean cross-sectional VPIN measure alone is not that informative of the level of order flow toxicity. The cross-sectional average VPIN is 0.272 indicating that on average across all stocks 27.2% of the trades come from informed traders[5]. The liquidity variables, depth, quoted spread, and effective spread all indicate that the average cross-sectional volume buckets are quite liquid. As seen before in Table 2 spreads narrow in the high AT activity decile compared to the low AT

[5]Raw VPIN measures on their own are not informative of the level of order flow toxicity. Meaningful comparison of VPIN to its empirical CDF can show how toxic the order flow is to the VPIN distribution. For example, Easley et al. (2012a) find that prior to the Flash Crash of 2010 the estimated VPIN measure for the E-mini S&P 500 futures series using one-minute bars, from January 1, 2008, to August 15, 2011, for a bucket size 1/50th of average daily volume buckets per day, and a sample length of fifty buckets, the CDF of VPIN was above 0.9 for two hours prior to the crash.

activity decile.

By comparing VPIN across AT activity deciles we can see that there is a significant difference between low AT activity and high AT activity. From Table 3 Panel B and Panel D, average VPIN is higher (0.303) in the low AT activity decile compared to the high AT activity decile (0.228). Using a t-test, and under the assumption of independent samples, I test for the null hypothesis of equivalence of means. Panel E of Table 3 shows that the difference in means for VPIN is statistically significant except for the difference between Decile 1 and Decile 2.

To further see the difference in estimated VPIN across low AT participation deciles and high AT participation deciles, Figure 2 shows the empirical CDFs of the VPIN measures for each AT decile. As can be clearly seen the low AT activity decile contains a greater percentage of higher VPIN values than the high AT activity decile. While a individual raw VPIN value would only be informative by comparing it to its respective empirical CDF, Figure 2 is informative by showing that lower AT participation is associated with higher order flow toxicity.

Interestingly, looking at Table 3 when comparing across AT participation deciles, only price volatility shows a difference in mean estimates. The low AT activity decile is associated with a high level of VPIN but also a high level of average price volatility over each volume bucket. This result is inline with Easley et al. (2012a), Andersen and Bondarenko (2014), Bethel et al. (2011), Wu et al. (2013), that all conclude that volatility is correlated with VPIN. While there is a significant dispute in the literature on the direction of causality (if any) between VPIN and volatility, I don't take a stance and rather look at how AT activity responds to both VPIN levels and factors such as price volatility, absolute return, depth, and bid ask spread.

Table 3: Descriptive statistics for VPIN measure and explanatory market variables over volume clock for full sample and 3 AT activity deciles. VPIN is the order flow toxicity measure calculated over 50 volume buckets and a 1 min time bar. $VPIN \approx \frac{\sum_{\tau=1}^{n}|V_{\tau}^{S}-V_{\tau}^{B}|}{nV}$, n=50 and V equals one-fiftieth of average daily volume. Order imbalance is signed using bulk classification of Easley et al. (2012). Other variables are calculated over the same volume clock (τ). Price volatility is the volatility of prices over each volume bucket (τ), Abs. return is the absolute return $\left(\left|\frac{p_{\tau}}{p_{\tau-1}}-1\right|\right)$ over volume bucket (τ), AT percentage is the proportion of trades that are executed by AT traders over volume bucket τ, Depth is the sum of the orders at the best bid and ask, Quoted spread (reported in bp) is the spread between the best ask and best bid when the trade is reported, and Effective spread (reported in bp) is the difference between the price and the midpoint. The mean, standard deviation, maximum, and minimum for the variables is calculated for each stock individually and then averaged for each stock. The cross-sectional mean of the descriptive statistics is reported for the full sample and 3 AT activity deciles.

Panel A: Full Sample

	VPIN	Price Volatility	Abs. Return (%)	AT percentage	Depth	Q. Spread	E. Spread
Mean	0.272	0.009	0.204	0.689	7942.173	3.670	1.492
Stdev.	0.098	0.013	0.195	0.100	3836.750	1.760	17.014
Max.	0.475	0.157	1.861	0.988	35346.824	23.981	347.834
Min.	0.053	0.000	0.000	0.276	1352.772	0.860	0.520

Panel B: Decile 1 (Low AT activity)

	VPIN	Price Volatility	Abs. Return (%)	AT percentage	Depth	Q. Spread	E. Spread
Mean	0.303	0.013	0.222	0.649	12706.332	4.040	1.577
Stdev.	0.100	0.017	0.209	0.103	5976.187	1.789	18.497
Max.	0.498	0.180	1.986	0.988	49739.979	21.838	373.394
Min.	0.060	0.000	0.000	0.268	1917.637	0.916	0.53

Panel C: Decile 2 (Medium AT activity)

	VPIN	Price Volatility	Abs. Return (%)	AT percentage	Depth	Q. Spread	E. Spread
Mean	0.284	0.011	0.183	0.682	2517.818	4.049	1.662
Stdev.	0.102	0.014	0.175	0.099	1325.543	1.981	19.514
Max.	0.499	0.153	1.631	0.975	16211.889	26.915	404.384
Min.	0.058	0.000	0.000	0.243	294.591	0.877	0.811

Panel D: Decile 3 (High AT activity)

	VPIN	Price Volatility	Abs. Return (%)	AT percentage	Depth	Q. Spread	E. Spread
Mean	0.228	0.004	0.206	0.738	8602.369	2.921	1.236
Stdev.	0.092	0.008	0.201	0.098	4208.519	1.509	13.031
Max.	0.428	0.139	1.966	1.000	40088.603	23.189	265.725
Min.	0.042	0.000	0.000	0.316	1846.087	0.787	0.631

Panel E: t-test for equivalence of means

	VPIN	Price Volatility	Abs. Return (%)	AT percentage	Depth	Q. Spread	E. Spread
Diff. 1 - 2	0.85	0.54	1.134	-3.44^{***}	1.08	-0.01	-0.20
Diff. 2 - 3	3.08^{***}	2.36^{**}	-1.41	-1.92^{*}	-1.44	1.59	0.85
Diff. 1 - 3	3.13^{***}	2.31^{**}	0.47	-2.90^{**}	0.40	1.24	0.71

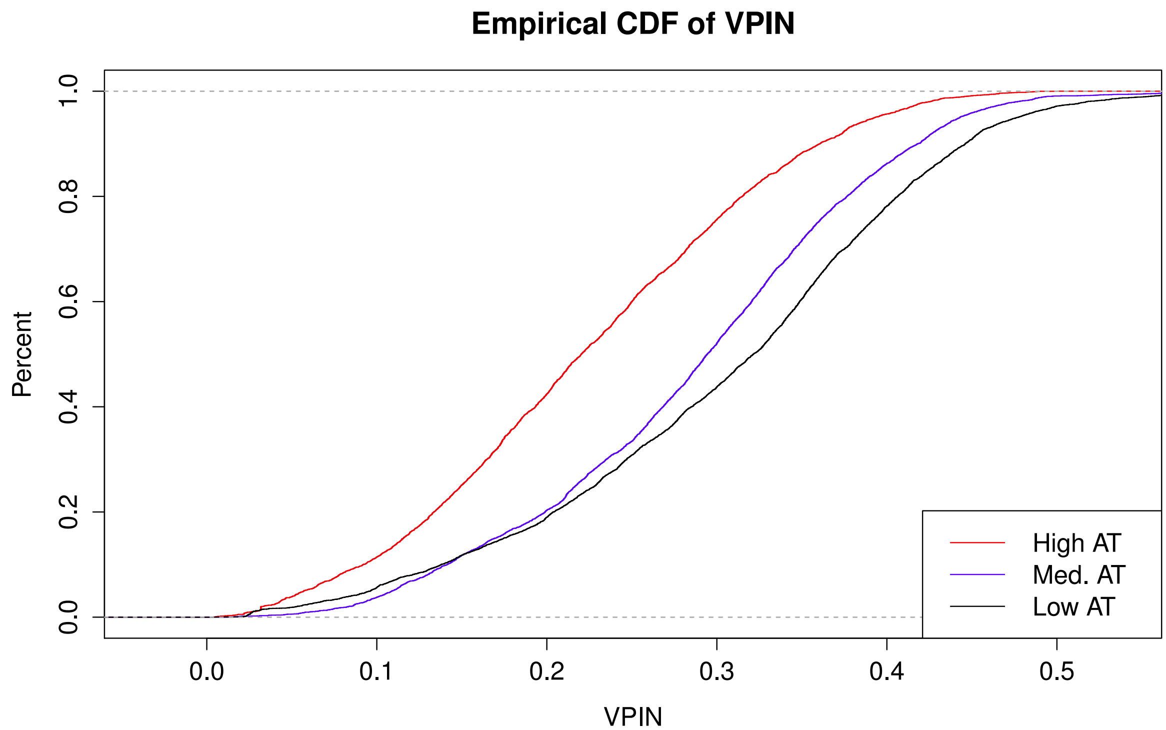

Figure 2: Empirical CDFs of VPIN measure for 3 AT activity deciles. AT activity deciles are ranked by average AT participation over volume buckets that are 1/50th of average daily trading volume.

A certain stock may have an inherently lower or higher level of adverse selection so by grouping stocks into deciles I may lose information by looking across decile averages. To provide a robustness check to the already stated results, Figure 3 plots the mean of VPIN and the descriptive market variables for each stock sorted according to AT intensity (low to high). As we can see in Figure 3 there is a clear downward trend in VPIN and price volatility as we move from low AT activity stocks to high AT activity stocks.

To better understand how each of the variables vary with each other in the cross-section, Table 4 gives the cross-sectional correlation coefficients for all 7 variables. In confirmation of previous evidence, VPIN is negatively correlated with AT activity and positively correlated with price volatility and absolute return. Table 4 also shows that VPIN is negatively correlated with depth, and positively correlated with spreads, indicating that when VPIN is high the market is more illiquid. Interestingly AT activity is negatively correlated with price volatility and absolute return indicating that ATs are more active in stable market conditions are favourable and liquid. These results are consistent with Hendershott and Riordan (2013) that find depth is negatively related to AT initiated trades, there is a negative relationship between AT initiation and spreads, and that AT trade initiation and lagged volatility are inversely related. These facts cumulate to provide no evidence to support a hypothesis that ATs exacerbate volatility or adverse selection.

While not taking a specific stance on the drivers of VPIN, the fact that AT activity is negatively correlated with VPIN is important. The negative correlation between AT activity and VPIN is notable for two reasons. First, there exists a link between AT activity and order flow toxicity giving evidence that ATs acting as as market makers may be less likely to engage in trade when adverse selection risk is high. Second, it gives the launch pad to formally look at what drives AT trade participation, i.e what factors do ATs monitor when deciding to engage in liquidity provision.

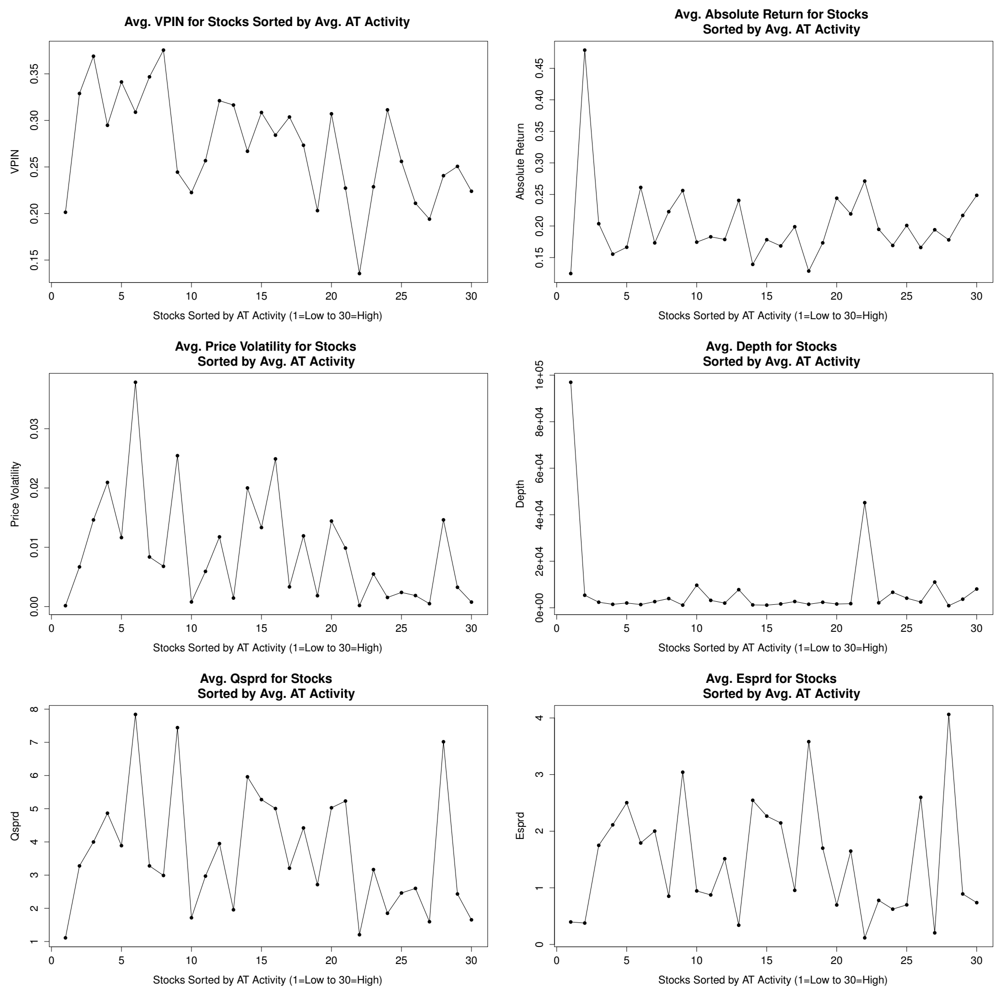

Figure 3: Cross-sectional daily average of VPIN, absolute return, price volatility, depth, quoted spread, and effective spread for 30 DAX stocks sorted according to average AT activity in volume time (low to high).

Table 4: Cross-sectional pearson correlation coefficients of average VPIN and explanatory market variables for 30 DAX stocks. VPIN is the order flow toxicity measure calculated over 50 volume buckets and a 1 min time bar. $VPIN \approx \frac{\Sigma_{\tau=1}^{n}|V_{\tau}^{S}-V_{\tau}^{B}|}{nV}$, n=50 and V equals one-fiftieth of average daily volume. Order imbalance is signed using bulk classification of Easley et al. (2012). Other variables are calculated over the same volume clock (τ). Price volatility is the volatility of prices over each volume bucket (τ), Abs. return is the absolute return $\left(\left|\frac{p_{\tau}}{p_{\tau-1}}-1\right|\right)$ over volume bucket (τ), AT percentage is the proportion of trades that are executed by AT traders over volume bucket τ, Depth is the sum of the orders at the best bid and ask, Quoted spread (reported in bp) is the spread between the best ask and best bid when the trade is reported, and Effective spread (reported in bp) is the difference between the price and the midpoint. The mean is taken of each variable for each stock to calculate the cross-sectional correlations.

	VPIN	Price Volatility	Abs. Return	AT (%)	Depth	Qsprd	Esprd
VPIN	1.000	·	·	·	·	·	·
Price Volatility	0.356	1.000	·	·	·	·	·
Abs. Return	0.121	0.011	1.000	·	·	·	·
AT (%)	-0.289	-0.235	-0.019	1.000	·	·	·
Depth	-0.425	-0.336	-0.103	-0.199	1.000	·	·
Qsprd	0.288	0.906	0.017	-0.192	-0.444	1.000	·
Esprd	0.126	0.568	-0.375	-0.068	-0.376	0.716	1.000

5 Formal model of AT trade participation:

To understand the factors that drive levels of AT trade participation and liquidity provision (under the assumption that ATs act as non-traditional market makers) there must be a more formal framework to tease out causality. Following Hendershott and Riordan (2013) that use the same dataset and the correlations already established, I propose a simple framework where AT trade participation is driven by multiple market factors, among them VPIN. Establishing a framework to view the the link between order flow toxicity and AT trade participation allows for the formalization of the results in the previous section.

It has been shown that ATs consume liquidity when it is cheap and provide liquidity when it is expensive, likely reducing the volatility of liquidity, Hendershott and Riordan (2013). Furthermore, AT liquidity providers closely monitor the market and respond more quickly to changes in market conditions as they have the technology to do so and tight risk management practises in place to reduce the chances of being adversely selected. By monitoring the state of the order book that includes factors such as order flow toxicity, price volatility, and liquidity factors, ATs have a competitive advantage over EEs to enter and exit out of orders that have gone stale or have the potential to be adversely selected[6].

Using a volume clock that is measured by 1/50th of daily average trade volume I model AT trade participation in volume buckets as,

$$ln(AT_{i,\tau}) = \alpha_i + \beta_1 ln(VPIN_{i,\tau-1}) + \beta_2 ln(\sigma^2_{p,i,\tau}) + \beta_3 ln(|r_{i,\tau}|) + \beta_4 ln(Depth_{i,\tau}) + \beta_5 ln(Spread_{Quoted,i,\tau}) + \varepsilon_{i,\tau}$$

where τ is the volume bucket defined for stock i as 1/50th of average daily trading volume, $VPIN$ is the estimated order flow toxicity measure, $\sigma^2_{p,i,\tau}$ is the estimated price volatility over volume bucket i, $|r_{i,\tau}|$ is the absolute return over τ, $Depth$ is the average sum of the best bid and ask depth over τ, and $Spread_{Quoted,i,\tau}$

[6] One issue that complicates the identification strategy of AT trade participation is that extremely low latency traders (HFTs) have been shown to improve liquidity and drive down transaction costs, Jones (2013). This leads to a small simultaneity bias in both AT trade participation and bid ask spreads. However, with the difference in mean spreads between the high AT participation decile and low AT participation decile being not statistically different from zero, my data precludes me from running a system simultaneous of equations to determine a joint estimation of AT trade participation and spreads.

is the average quoted spread at time of trades[7]. To provide another angle on the relationship between AT trade participation and market characteristics I sorted the stocks according to average AT trade participation (low to high).

The log-log model of AT trade participation is estimated for each stock individually by OLS and the estimated coefficients and standard errors are presented in Figure 4. Two results are immediately clear by looking at Figure 4. First, the magnitude of the effect from a change in one of the explanatory variables on AT trade participation is higher for stocks where AT trade participation is low compared to where AT participation is high. This effect while appearing contradictory is actually consistent with the hypothesis that ATs closely monitoring the market. With ATs having a lower trade participation rate in stocks that have a high level of order flow toxicity, high depth, and large spread, a change in one of these explanatory variables will cause ATs to quickly leave the market. In other words, AT market makers are less sensitive to changes in order book characteristics like VPIN when average values are low. For example, in low AT activity stocks average VPIN is higher than in high AT activity stocks, therefore, one could posit that some AT firms have a higher risk appetite for order flow toxicity. However, a marginal increase in these 'riskier' stocks may cause a larger effect on AT participation as the adverse selection risks are higher.

The second key result is that a marginal increase in VPIN effects AT participation in a heterogeneous manner depending on whether average AT trade participation is high or low. For stocks where average AT trade participation is on average low, an increase in in order flow toxicity causes a decrease in AT trade participation. For example, a 1% increase in VPIN causes a 0.04% decrease in AT trade participation. On the other end of the spectrum for stocks where average AT trade participation is high a marginal increase in VPIN causes an increase in AT trade participation.

The difference in AT trade participation with a change in VPIN between low AT activity stocks and the high AT activity stocks can be partially explained by bid-ask spreads and different levels of order book monitoring. Drawing on the two stage make/take liquidity cycles of Foucault et al. (2013) where an order from a liquidity supplier first narrows the spread by offering a better price. Then in the second part of the cycle, a liquidity demander monitors the market and reacts to the narrow spread by initiating a trade. The

[7] I omitted including effective spread in the model as quoted spread adequately captures liquidity characteristics effective spread represents. Effective spread was also omitted due to the fact that it was not statistically different from zero when the models were estimated.

trade causes the spread to widen and the cycle repeats. Viewed through this framework, monitoring of the limit order book and the resulting liquidity cycles are manifestations of search frictions for investors seeking gains from trade, Hendershott and Riordan (2013). Smaller bid ask spreads can then be interpreted as smaller search costs that result in greater competition among traders due to lower bargaining frictions. As Decile 3 (high AT participation) also corresponds with a lower spreads we can argue that there is greater competition among AT traders in the high activity stocks. With increased competition among liquidity providers in liquid stocks, a marginal increase in order flow toxicity can cause AT market markers to maintain their market share or even increase it.

An alternative explanation for the increase in AT trade participation with an marginal increase in VPIN, is that ATs may increase their initiation of trades and profit from information gained by superior limit order book monitoring. Hendershott and Riordan (2013) find that AT trade initiation following lagged increases in futures of the DAX cause AT buy market orders more likely. When DAX futures drop they find the same conclusion that ATs increase their sell market orders. If futures prices lead the underlying stock prices, then the trades initiated by ATs impose adverse selection costs on the nonmarketable limit orders they execute against. This increase in informed trading will be contained in VPIN and will then be reflected in an increase in AT trade participation.

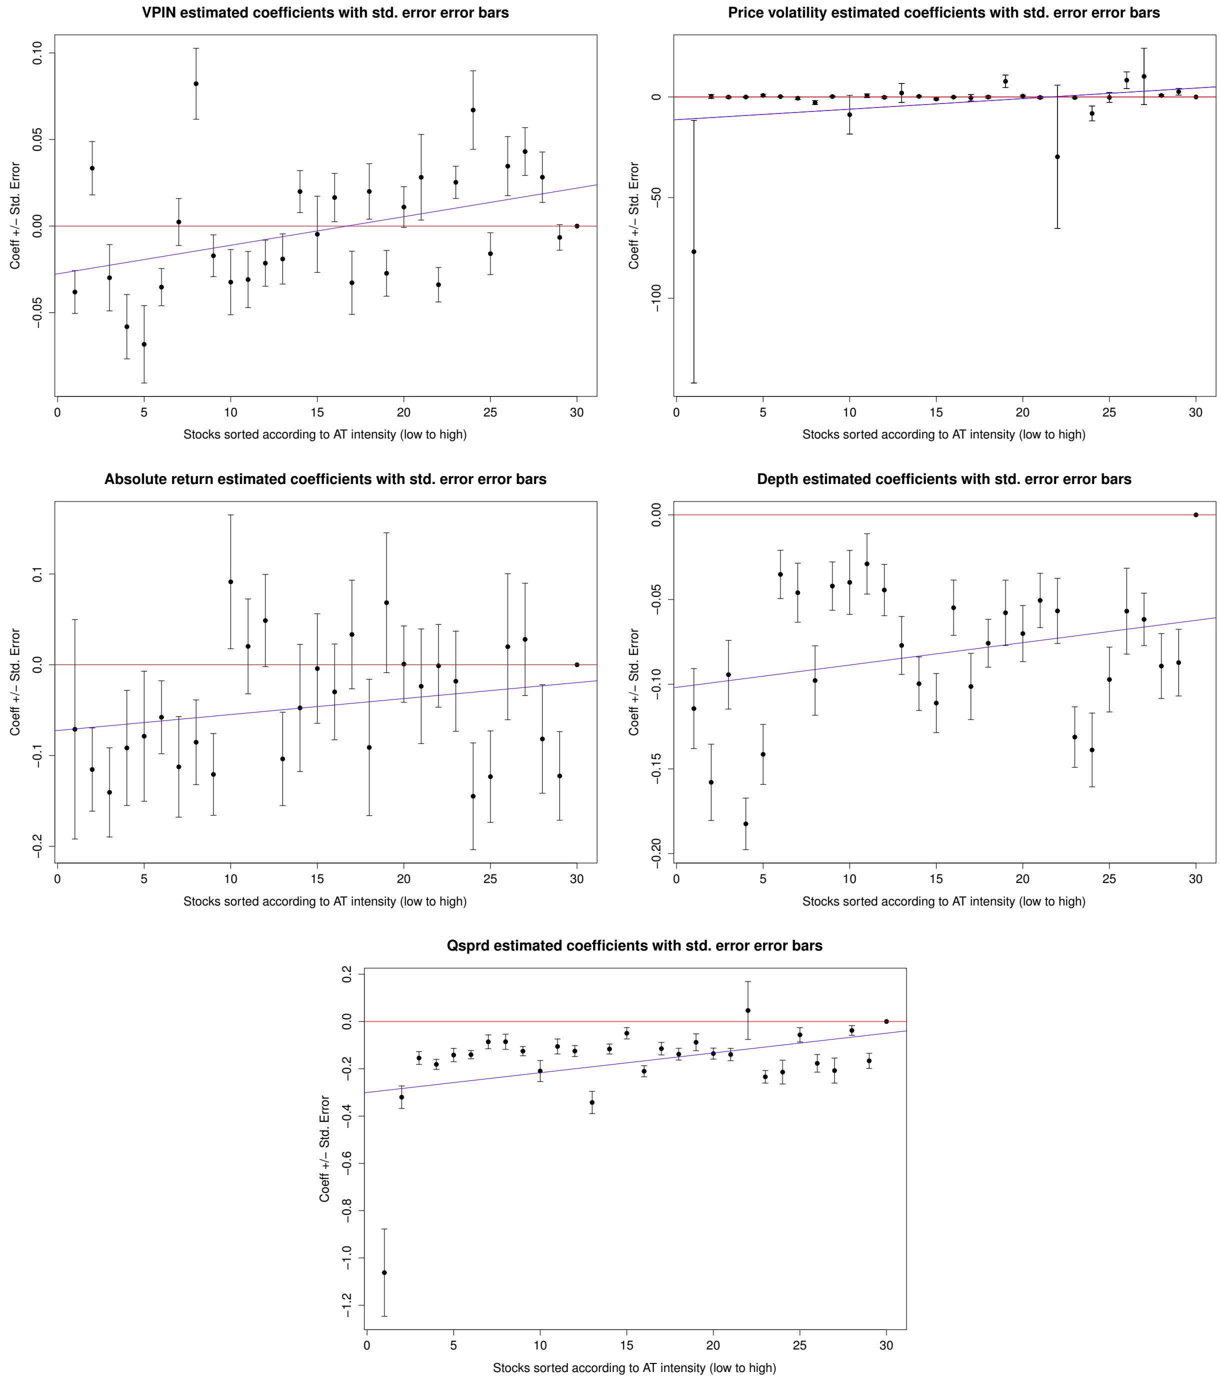

Figure 4: Estimated coefficients and standard errors from AT trade participation model. Model estimated is: $ln(AT_{i,\tau}) = \alpha_i + \beta_1 ln(VPIN_{i,\tau-1}) + \beta_2 ln(\sigma^2_{p,i,\tau}) + \beta_3 ln(|r_{i,\tau}|) + \beta_4 ln(Depth_{i,\tau}) + \beta_5 ln(Spread_{Quoted,i,\tau}) + \varepsilon_{i,\tau}$ where $i = 1, ..., 30$ stocks sorted according to AT trading intensity. Blue lines are additional cross-sectional trend lines.

6 Conclusion

This paper studies the link between order flow toxicity and algorithmic trading activity in an electronic limit order books where Algorithmic Traders (ATs) are exactly identified. I find that by using a volume clock measure, AT trade participation in a constant measure of trade volume is negatively correlated with AT activity. Furthermore, AT trade activity is more sensitive to market conditions when order flow toxicity, as measured by VPIN, is high. ATs closely monitor the market and when adverse selection risks are high they decrease their trade participation. These results are consistent with previous work that show ATs have a positive impact on market conditions such as a reduction in spreads and volatility of liquidity (Jones (2013); Hendershott and Riordan (2013); Carrion (2013) ; Menkveld (2013)).

My results can have important implications for academics, regulators, and market participants. By showing that there is link between AT trade participation and changes in order flow toxicity, I establish two insights. First, it adds another piece to the puzzle about the effects of algorithmic trading and reinforces previous work that shows AT market making does not lead to increases in volatility and market turbulence. Secondly, I show the merits of using an event clock to measure AT activity and order flow toxicity. With a volume clock more consistent with how AT algorithms operate in the market, it only seems logical to follow suit and study the behaviour of AT market makers using the same outlook as they do. Furthermore, the VPIN measure used to capture order flow toxicity behaves as a proxy for adverse selection should. Moving forward, this leads me to conclude that VPIN is a useful metric for analyzing order flow toxicity and adverse selection.

By showing that VPIN is one of many important drivers for AT trade participation there is one important takeaway for non-AT market participants, VPIN captures order flow toxicity in a meaningful way. ATs are more sensitive to changes in market conditions and have tighter risk management practises than typical market makers so tracking order flow toxicity is important to maintain profitability. If traditional market makers and non-AT market participants want to compete for liquidity provision, understanding and tracking order flow toxicity should be an important part of maintaining market share for traditional market makers or in the very least regaining an equal foothold with silicon traders.

7 Appendix: Data and Matching Process

In Dec. 2007 the DB introduced its Automated Trading Program (ATP) to increase the volume of automated trading on Xetra. By offering fee rebates, the DB was implicitly subsidizing investment in AT technologies. To qualify for the ATP an electronic system must determine the price, quantity, and submission time for orders. In addition, the DB ATP agreement required that: i) the electronic system must generate buy and sell orders independently using a specific program and data; ii) the generated orders must be channeled directly into the Xetra system; and iii) the exchange fees or the fees charged by the ATP member to its clients must be directly considered by the electronic system when determining the order parameters.

The Xetra trading system is the electronic trading system operated by the DB and handles more than 97% of German equities trading by euro volume in DAX stocks, Hendershott and Riordan (2013). The DB is a publicly traded company that also operates the Eurex derivatives trading platform and the Clearstream European clearing and settlement system. DB admits participants that want to trade on Xetra based on regulations set and monitored by German and European financial regulators. After being admitted, participants can only connect electronically to Xetra; floor trading is operated separately, with no interaction between the 2 trading segments.

Liquidity in DAX stocks is provided by public limit orders displayed in the order book of each stock. Orders execute automatically when an incoming market or marketable limit order crosses with an outstanding non-marketable limit order. Order execution preference is determined using price-time-display priorities. Three types of orders are permitted: limit, market, and iceberg orders. Iceberg orders are orders that display only a portion of the total size of an order. Iceberg orders sacrifice time priority on the nondisplayed portion. Pre-trade transparency includes the 10 best bid and ask prices and quantities but not the identity of the submitting participant. Trade price and size are disseminated immediately to all participants. The tick size for most stocks is 1 euro cent, with the exception of 2 stocks that trade in 1/10ths of a cent[8].

To create the final data set of trades and orders, I use 3 separate data sources: AT order data from DB, public order book data from SIRCA, and public transactions data from SIRCA. Because SIRCA time stamps reflect

[8]Deutsche Telekom AG and Infineon AG have trade prices below 15 euros. Stocks with prices lower than 15 euros have a tick size of 1/10th of a cent.

routing delays between DB and Thompson- Reuters, the SIRCA data sets are subject to time lags relative to the AT system order. The time stamp or SIRCA order book data is lagged by up to 250 milliseconds (ms). The SIRCA transactions data set is lagged by up to 500 ms.

I generate order from successive order book updates as in Biais et al. (1995). Every generated order from the order book updates can only take on two values, addition of liquidity or removal of liquidity. To identify order that are trades I match the generated orders to the publicly available SIRCA trade data. To match the SIRCA trades to the orders I matched by:

- Symbol
- Date
- Price
- Size
- Timestamp
- Order type (insert or delete)

As Xtera permits iceberg orders and would not be visible in the order generation from visible order book updates, the SIRCA trade data was interleaved with the generated orders and the duplicates were deleted. Then to identify which orders and trades were generated by ATs, I matched the combined trade and order data with the AT DB data. As there is time lag between the SIRCA orderbook and the AT DB data I allowed for a time window of 500 ms when matching. Roughly 87% of all algorithmic orders are matched in the public data.

References

Abad, D. and Yagüe, J. (2012). From pin to vpin: An introduction to order flow toxicity. *The Spanish Review of Financial Economics*, 10(2):74–83.

Andersen, T. G. and Bondarenko, O. (2014). Vpin and the flash crash. *Journal of Financial Markets*, 17:1–46.

Barclay, M. J. and Warner, J. B. (1993). Stealth trading and volatility: Which trades move prices? *Journal of Financial Economics*, 34(3):281–305.

Bethel, E., Leinweber, D., Rübel, O., and Wu, K. (2011). Federal market information technology in the post flash crash era: roles for supercomputing. In *Proceedings of the fourth workshop on High performance computational finance*, pages 23–30. ACM.

Biais, B., Hillion, P., and Spatt, C. (1995). An empirical analysis of the limit order book and the order flow in the paris bourse. *Journal of Finance*, pages 1655–1689.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3):307–327.

Brogaard, J., Hendershott, T., and Riordan, R. (2014). High-frequency trading and price discovery. *Review of Financial Studies*, 27(8):2267–2306.

Carrion, A. (2013). Very fast money: High-frequency trading on the nasdaq. *Journal of Financial Markets*, 16(4):680–711.

Chakrabarty, B., Pascual, R., and Shkilko, A. (2012). Trade classification algorithms: A horse race between the bulk-based and the tick-based rules. *Available at SSRN 2182819*.

Easley, D., de Prado, M. M. L., and O'Hara, M. (2012a). Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies*, page hhs053.

Easley, D., de Prado, M. M. L., and O'Hara, M. (2012b). The volume clock: Insights into the high-frequency paradigm (digest summary). *Journal of Portfolio Management*, 39(1):19–29.

Easley, D., de Prado, M. M. L., and O'Hara, M. (2014). Vpin and the flash crash: A rejoinder. *Journal of Financial Markets*, 17:47–52.

Easley, D., Engle, R. F., O'Hara, M., and Wu, L. (2008). Time-varying arrival rates of informed and uninformed trades. *Journal of Financial Econometrics*, 6(2):171–207.

Easley, D., Kiefer, N. M., O'hara, M., and Paperman, J. B. (1996). Liquidity, information, and infrequently traded stocks. *The Journal of Finance*, 51(4):1405–1436.

Easley, D. and O'hara, M. (1992). Time and the process of security price adjustment. *The Journal of finance*, 47(2):577–605.

Easley, D., O'Hara, M., and Paperman, J. (1998). Financial analysts and information-based trade. *Journal of Financial Markets*, 1(2):175–201.

Easly, D. and O'Hara, M. (1987). Prices, trade size and information in security markets. *Journal of Financial Economics*, 33(1):173–199.

Foucault, T., Kadan, O., and Kandel, E. (2013). Liquidity cycles and make/take fees in electronic markets. *The Journal of Finance*, 68(1):299–341.

Grammig, J., Schiereck, D., and Theissen, E. (2001). Knowing me, knowing you: Trader anonymity and informed trading in parallel markets. *Journal of Financial Markets*, 4(4):385–412.

Gray, V. and Aspey, M. (2004). Rothschild, nathan mayer (1777–1836). *Oxford Dictionary of National Biography*.

Hasbrouck, J. (2006). *Empirical market microstructure: The institutions, economics, and econometrics of securities trading*. Oxford University Press.

Hasbrouck, J. and Saar, G. (2013). Low-latency trading. *Journal of Financial Markets*, 16(4):646–679.

Hendershott, T. and Riordan, R. (2013). Algorithmic trading and the market for liquidity. *Journal of Financial and Quantitative Analysis*, 48(04):1001–1024.

Jones, C. M. (2013). What do we know about high-frequency trading? *Columbia Business School Research Paper*, (13-11).

Lee, C. and Ready, M. J. (1991). Inferring trade direction from intraday data. *The Journal of Finance*, 46(2):733–746.

Mandelbrot, B. and Taylor, H. M. (1967). On the distribution of stock price differences. *Operations research*, 15(6):1057–1062.

Menkveld, A. J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4):712–740.

O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics*.

SEC (2010). Securities exchange act release. Technical Report No. 34-61358 75 FR 3594, 3606, Securities Exchange Commision.

Sofianos, D. J. . G. (2008). Passive orders and natural adverse selection. Street Smart 33, Goldman Sachs.

Wu, K., Bethel, W., Gu, M., Leinweber, D., and Ruebel, O. (2013). A big data approach to analyzing market volatility (june 5, 2013). *Available at SSRN 2274991*.