Stockholm School of Economics
Department of Finance
Thesis in Finance
Spring 2012

Mean Machines? – High-Frequency Trading and Market Quality in the Swedish Financial Market

---

Abstract: The role and effects of high-frequency trading have been heavily discussed over the past few years. However, the actual effects of this type of trading in Sweden have not been properly investigated. By using a previously unmatched dataset, with minute observations, we analyze the effects of high-frequency trading on the quality of the Swedish financial market. To do this we define three parameters of market quality: volatility, liquidity and market efficiency. The research design is based on a difference-in-difference approach on the structural change of NASDAQ OMX Stockholm, the implementation of the INET trading platform, and variance ratio tests. We find that high-frequency trading overall has positive effects on market quality in Sweden; while volatility is somewhat reduced and liquidity is improved on the margin the market efficiency remains unaffected. We suggest that the risks with high-frequency trading are acceptable and that the general fear and aversion towards this type of trading needs to be reconsidered.

Keywords: High-frequency trading, Market quality, Volatility, Liquidity, Market efficiency

Authors: Carl-Emil Lindholm[♥] and Erik Petersson[♣]

Tutor: Assistant Professor Laurent Bach

[♥]21148@student.hhs.se, [♣]21135@student.hhs.se

Acknowledgements: We would like to thank our tutor Assistant Professor Laurent Bach for his guidance and support throughout the course of the work with this thesis. We are also grateful for the data and assistance provided by Avanza Bank AB, especially Elias Lindholm. Finally, we appreciate the thoughts and insights provided by the business professionals we met with and interviewed.

Contents

List of Tables

List of Figures

I. Introduction

High-frequency trading[1] is a topic widely debated in Swedish, as well as international, media today. However, going back a few years hardly anyone outside the field of finance knew anything about the matter. In December 2005 a Swedish business newspaper, Dagens Industri, published an article about the existence of "black boxes" on the financial markets. The article described a phenomenon where financial institutions used computers programmed with algorithms to exploit trading opportunities by monitoring price movements in the market (Munkhammar, 2005). This kind of strategies can be seen as the predecessor to HFT strategies which takes advantage of high-speed connections with the exchanges and algorithms to get in and out of trades within seconds. It was not until August 2009, to the best of the authors' knowledge, that HFT was mentioned in the broader Swedish media when Dagens Nyheter published an article about HFT and the debate that flash orders had created in the U.S. market (Levander, 2009). Since then there have been numerous articles about the matter and much of the debate is focused on the price manipulation and increased volatility that certain HFT strategies have been accused of.

Just recently (January, 2012), Swedish investment bank Pareto Öhman was fined SEK 0.5M for their algorithmic trading; the company's explanation was a faulty algorithm (Rex, 2012). Several important actors in the finance industry have strong opinions on the subject, where traditional investors claim that the quality of the market has been negatively affected by HFT. The CEO of NASDAQ OMX Stockholm[2], Jens Henriksson, on the other hand claims that HFT has not had an impact on the volatility on the Swedish stock exchange. Instead, he argues, much of the negative attention regarding HFT stems from the general turbulence we have seen in the stock market over the past few years. Others defend HFT by pointing to the fact that it is much more prevalent in other financial markets and that the Swedish financial market would be at a disadvantage if it was not allowed (Neurath, 2011).

In the spring of 2012 The Swedish Financial Supervisory Authority (Finansinspektionen) published a qualitative investigation of HFT in Sweden, based on two surveys, where they argue that the risks with HFT and algorithmic trading are acceptable. However, they recognize that the Swedish financial markets have changed, to some extent for the worse, but that this is due to other reasons than HFT. They also note a general concern for market manipulation that cannot be overlooked (Finansinspektionen, 2012). Further, the minister of the Swedish financial markets, Peter Norman, has defended HFT to some extent, claiming that the main criticizers are those

[1] From now on we refer to high-frequency trading and trades as HFT.
[2] From now on we refer to this as OMX Stockholm or the Stockholm stock exchange.

that have lost the opportunity to make easy money, e.g. day traders. The revision of the Markets in Financial Instruments Directive (MiFID), MiFID II, also directly address HFT and algorithmic trading by suggesting that safeguards[3] should be put in place for these trading methods. The safeguards include "the requirement for all algorithmic traders to become properly regulated, provide appropriate liquidity and rules to prevent them from adding to volatility by moving in and out of markets" (European Commission, 2011). Finally, in the U.S. HFT was greatly criticized following the events on May 6th, 2010, later known as The Flash Crash, when the equity markets plunged about six percentage points in a mere five minutes, only to recover most of the decline in the following twenty minutes.

The advent of the debate is logical and to be expected, especially considering the rapid rise seen in the number of trades performed through HFT. Today we see financial institutions that were nearly invisible in the Nordic exchanges a few years ago, constitute up to five percent of the total volume traded on the exchange. Furthermore, HFT has gone from being virtually non-existent to represent about 20 percent of the total trading on the Stockholm exchange in a relatively short period of time (Hertzberg, 2012) – in the more mature U.S. market it has been claimed that this number is as high as 60 percent (Psomadelis and Powell, 2011). What is most interesting with the current discussion on HFT is the overall negative view presented in the general business media. Most of the articles are based on statements by large investors, traders or others with a conservative view on how the equity market should function. However, despite the negative view presented in the media very few concrete arguments have been made as to in which way the HFT is bad for the financial markets. The arguments mostly rest upon a notion that the markets are more volatile today, without proving that this is in fact the case or that there is any relationship between the existence of HFT and increased market volatility.

The fact that HFT actors operate at very high speeds make the topic challenging, to be able to infer anything about the impact it might have on financial markets high-frequency data needs to be analyzed; both the unavailability of such data and its magnitude makes for a difficult task. However, with a previously unmatched dataset provided by Avanza Bank AB, we show that there are good reasons to question the negative view presented above. By using high-frequency data and employing a difference-in-difference framework, on a period in which a structural change important to HFT was made at the Stockholm stock exchange, we show that the volatility has actually decreased due to HFT. Compared to the pre-event period our tests suggest that HFT might have decreased the volatility with as much as 15-30% of the pre-event standard deviation of volatility. Albeit not as conclusive, our analysis also suggests an improvement of liquidity in

[3] In fact NASDAQ OMX already has both static and dynamic volatility guards, also known as circuit breakers, in place that halts the trading for about one minute if stock prices show abnormal deviations.

terms of a tighter bid-ask spread due to HFT. We also investigate the effect HFT has had on the market efficiency, as measured by the serial correlation in returns, but these results are inconclusive. The results lead us to conclude that HFT is not detrimental to market quality, but rather seems to improve it.

A. Purpose and Limitations

Given the contemporaneity of the subject and the strong debate in media, a study of the impact HFT has had on the Swedish stock market is much needed. Unlike the American market, where HFT has been studied to some extent already, there is a gap of knowledge concerning this subject in the academic literature of the Swedish financial market. Therefore, this paper aims to uncover what effect HFT has had on the quality of the Swedish financial market. Mainly, we want to answer the question:

What is the effect of high-frequency trading on the quality of the Swedish financial market?

Also, since HFT, at least in the case of Sweden, is a relatively recent phenomenon we delimit ourselves to examining a period from 2006 until today.

To measure the quality of the market we will consider the price process followed by Swedish stocks and evaluate trends in the liquidity and volatility. By doing this we hope to bring some hard facts to a debate that thus far has been rather one-sided and until now mostly based on opinions. Hence, we hope to contribute to a more informed debate about an issue that is likely to come under closer scrutiny in the years to come, especially considering the development of MiFID II.

B. HFT Definition

Before moving further, it is important to define what is meant by HFT in this paper; particularly important is the difference between algorithmic trading and HFT. Algorithmic trading does not depend on the low latency requirements most HFT strategies do. They were developed to be used by mainly buy-side participants to minimize the market impact of their trades. Instead of, as historically, using a broker-dealer to manage execution of large positions algorithms can be used to submit and manage orders and make certain trading decisions. They determine timing, quantities, routing and prices for the orders to optimize the outcome of the operation (Hendershott et al., 2011).

While HFT is a subset of algorithmic trading and is indeed performed using algorithms it is not used to manage predetermined decisions but analyze data and make trading decisions autonomously (Chlistalla, 2011). These algorithms produce a very large number of orders which are sent to the marketplace at very high speeds, and the executions are measured in microseconds. Further, they earn small profits on each trade but since they perform a huge amount of trades the profits become sizeable (Zhang, 2010). They analyze vast amounts of data and make trading decisions based on opportunities that exist for only parts of seconds, i.e. they do not depend on human interaction to decide whether or not to enter a trade. The low latency requirement previously mentioned has had the implication that most HFT-algorithms are operating from servers physically located at the market places' matching engines (so called co-location), in a pursuit to cut off microseconds in the data transmission process. The positions held by these algorithms are generally closed within very short time horizons (seconds or minutes) and they most often close the day flat, i.e. with no significant positions in any securities or direction. Finally, HFT-algorithms are often used in the market making process and this is often the base activity for HFT firms (Hertzberg, 2012).

C. Previous Research

Research directly addressing HFT and its impact on financial markets is still not that extensive, especially in the Nordic region, but it is continuously growing. Due to difficulties in identifying HFT many of them use proxies of different kinds. The results differ somewhat depending on method, market and time-horizon as well as definition of market quality; some papers suggest a positive impact and others a negative. Nonetheless, the results are mainly in favor of HFT. The most important papers directly addressing HFT and their results are discussed below; some papers addressing algorithmic trading deemed relevant are also presented.

The most commonly tested characteristics of market quality in relation to HFT are liquidity, volatility and price discovery, i.e. the price formation process where buyers and sellers interact to find the equilibrium price. Brogaard (2010) tests these parameters using Hasbrouck's measures of price discovery, natural experiments and hypothetical price paths and suggests that HFT adds to the price discovery process and that it might dampen intraday volatility. However, even though HFT provides the best bid and offer quotes for a substantial part of the trades it does not provide the same liquidity-depth as non-HFT actors. Volatility in exchange rate markets was investigated by Chaboud et al. (2009) using a range of methods including an instrumental variable approach, with the fraction of trading floors able to use algorithmic trading as instrument. Their results are in line with Brogaard (2010) as they do not find a causal relationship between HFT and volatility. However, their results indicate that non-HFT adds more to the variance in exchange rate returns than HFT, i.e. non-HFT traders are driving the price-discovery. Riordan and Storkenmaier (2011) studies the impact of reduced latency, i.e. the amount of time it takes to submit and receive feedback on an order, on liquidity and price discovery using a natural experiment (the upgrade of the trading system on Deutsche Boerse) in an electronic limit order market. They show that price discovery increases post upgrade, indicating that prices are more informative. Even though algorithmic trading effects must not per se be valid for HFT, the study by Hendershott et al. (2011) is interesting since they look at the above market quality characteristics and show strong evidence that liquidity and the informativeness of quotes are increased by algorithmic trading.

In addition to the aforementioned characteristics used to test market quality some papers also study bid-ask spreads, which represents one dimension of liquidity. Hasbrouck and Saar (2010) look at the activity in what they call the "millisecond environment", i.e. order-level NASDAQ data, and conclude that short-term volatility, spreads and the displayed depth in limit order books are improved. These results are confirmed by Castura et al. (2010); they show that US equity markets have become more efficient in terms of tighter spreads, less mean reversion of stock prices and greater liquidity on the inside over the past several years. Furthermore, by using variance ratio tests they investigate the market efficiency. Even though they do not show a causal relationship they demonstrate that exchanges that moved earlier to automation saw earlier improvements these factors. Additionally, Riordan and Storkenmaier (2011) find an increase in liquidity due to lower adverse selection costs, i.e. costs that arise due to asymmetric information between buyers and sellers, by using different bid-ask spread measures, e.g. quoted and effective spreads. However, Hendershott and Moulton (2009) investigates the implementation of the Hybrid market at the New York Stock Exchange and find that the increase in execution speed (from about 10 to about 1 second) raised the spreads due to increased adverse selection, but lead to more information being incorporated into prices, i.e. more efficient prices. Jovanovic and Menkveld (2011) use a theoretical model to show that when investors arrive at different times in a limit-order market the presence of HFT can raise welfare due to elimination of adverse selection because they post competitive price quotes in the market. Hence, their results add to the overall opinion that HFT, in most cases, seems to reduce adverse selection in limit-order markets.

Even though the majority of papers looking at volatility and price discovery find improvements in financial market quality when HFT is present there are a number of papers stating the contrary. Zhang (2010) classifies investors into categories to estimate HFT and finds a positive correlation between HFT and stock price volatility, even after controlling for a number of volatility drivers. Further, he finds evidence that the market overreacts to fundamental news when HFT is present and hence obstructs the price discovery process. Jarrow and Protter (2011) finds results in line with this when they create a model without bid-ask spreads and with perfect liquidity. They show that market volatility increases and that HFT actors generate abnormal profits at the expense of non-HFT actors. Cvitanic and Kirilenko (2010) look at transaction prices in relation to HFT by modeling a limit-order market with low frequency traders and then add HFT and find that transaction prices are affected. Their results, like those of Jarrow and Protter (2011), indicate that HFT-algorithms make profit by "sniping" out human orders not at the front of the order book. They also show that it might be the case that the faster humans submit their orders, the more they are "exploited" by algorithms.

Finally, there are a number of papers investigating the Flash Crash and the role HFT played in those events. Kirilenko et al. (2011) use audit-trail data and find that while HFT did not trigger the crash their responses did indeed exacerbate the market volatility. They argue that the high trading volumes by HFT can be mistaken for liquidity, and when HFT-algorithms rebalance their portfolios they compete for liquidity and by that contribute to increased volatility. The liquidity issue is further investigated by Easley et al. (2010); they suggest that the Flash Crash should be seen as a liquidity event arising from structural changes due to HFT. When order flow toxicity, i.e. expected losses from trading with better informed counterparties, increases HFT market makers might reduce their risk by reducing or even liquidating their positions. This might have devastating consequences for other market participants in terms of market illiquidity.

Overall, the current literature directly addressing and investigating HFT mostly shows improvements in a number of market quality characteristics. However, there are a number of papers showing market obstructions created by HFT, but they are currently in minority. Finally, as mentioned in the introduction there is a vast amount of non-academic articles discussing the potential detrimental market effects introduced by HFT; these are not presented here since they are not based on academic research.

II. Research Hypotheses

As the purpose of this thesis is to investigate effects of HFT on the quality of the Swedish financial market we need a number of parameters as tools for the analysis. Inspired by the previous research within the subject, were we mainly draw on the paper by Castura et al. (2010), and general market quality discussions we have chosen a set of parameters we consider most relevant. The first parameter, and perhaps most important for market participants, is volatility. Although it has been scrutinized by several previous papers it is still deemed a very important market quality characteristic; large price-swings might create large costs for all participants. Another parameter important to market participants is the market liquidity, which has also been profoundly investigated, but is still interesting due to the recent rise of HFT. Liquidity is imperative for participants to determine the timing of their trade executions and to reduce the price impact of their trades. While these parameters have been previously investigated we argue that our contribution is the combination of them on a previously unmatched dataset. The final parameter is the overall market efficiency in the fashion of Fama (1965) and Samuelson (1965). This has not been thoroughly investigated in previous HFT-papers and might represent a significant cost to larger investors. Theories and facts concerning these parameters, as well as our hypotheses regarding them, are presented in what follows.

A. Volatility

Stock return volatility is one of the most common measures of financial risk and is the tendency for prices to unexpectedly change (Harris, 2003). It is an important issue for both investors, whom require higher risk premiums for stocks with higher volatility, and managers, whose companies face higher costs of capital if their stocks have higher volatilities; Bushee and Noe (2000) argue that these are reasons to be concerned about volatility. However, Ozenbas et al. (2002) argue that there are both good and bad volatility in that some (good) volatility reflects news incorporating into the prices. They further suggest that it is important for market structure to control bad volatility, which occur due to the arrival of trade orders to the market.

Volatility is also central to finance in general, in such areas as asset pricing, portfolio allocation and risk management (Andersen et al., 2002) and Bessembinder (1998) suggests that market, i.e. listing exchange, characteristics have implications for stock price volatility. Hence, there are several reasons for why volatility is an important parameter in the field of market quality. Despite its importance, the consensus regarding volatility measurements are somewhat ambiguous; but Andersen et al. (2002) provide a number of different ways to measure and interpret volatility, including ex-post (realized) return variability, model volatility procedures (parametric and non-parametric) and ex-ante (implied) expected volatility.

The effect of HFT on market volatility is not obvious. While the previous research within the field shows mostly improvements in that volatility is reduced Zhang (2010) discuss a number of different effects, both increasing and reducing volatility. The market-making activity performed by HFT-algorithms might reduce volatility due to the fact that they supply liquidity in the market and thereby reduce the price impact from large block traders. On the other hand, he argues that the interaction between autonomous HFT-algorithms and fundamental investors might increase volatility. The high trading volumes generated by these algorithms does not necessarily have to indicate higher liquidity since they might withdraw from the market in times of high uncertainty. When fundamental investors execute large orders automatically, often using volume as a proxy for liquidity, they can potentially cause large price movements, e.g. the Flash Crash of May 2010. HFT-algorithms can also detect and front-run large orders by institutional investors, thereby pushing prices either up or down and hence increase volatility. This behavior is made possible due to co-location, where HFT firms place their servers close to the exchanges' matching engines to gain low data transfer latency. Finally, due to the statistical properties, e.g. short-term correlation trading, of many HFT-algorithms they might generate price momentum and thus increase market volatility.

The importance of volatility for market participants together with the ambiguity around the effects due to HFT makes it interesting in the context and is therefore the first parameter used in our analysis. Finally, our hypothesis regarding volatility is based mainly on the previous research on this parameter. Even though there are papers showing results of increased volatility we consider the arguments and obtained results for decreased market volatility stronger. Hence, our first hypothesis states the following:

Hypothesis 1: Market volatility has decreased as a result of high-frequency trading.

B. Liquidity

Liquidity is another important characteristic for financial markets; in fact some argue that for a well-functioning market liquidity is the uttermost important characteristic (e.g. Harris, 2003). It is generally defined as the ability for market participants to trade the amount they want at the time they prefer, without having a too large effect on price (Castura et al., 2010). If participants do not have to depend on others to execute their trades when they desire, based on their current needs, the markets become more efficient and funds can flow more freely. However, according to Harris (2003) there is widespread confusion regarding the definition, since it might mean different things to different participants. There are also several dimensions of liquidity, further obstructing any uniformly accepted definition. Is it the ability to trade quickly, to trade at any time or any size preferable? Or is it all of them?

Since there are several definitions of liquidity there are naturally a number of different measures that can be used, each focusing on different dimensions of liquidity. Aldridge (2010) lists the following measures to assess market liquidity: tightness of the bid-ask spread; market depth (size of all limit orders posted at the current market price); market resilience (how quickly the price reverts to the equilibrium level); price sensitivity to block transactions; and an illiquidity ratio (relative price change to quantity traded). This means that there are several ways to determine potential liquidity effects due to HFT.

With the development of HFT on financial markets the opinion is that trading volumes have increased substantially (Kirilenko et al., 2011). Even though many market participants use this as a proxy for liquidity it does not have to be the case that overall liquidity has increased due to HFT. While HFT-algorithms generally trade large volumes they do it in smaller portions. Hence, there might be available liquidity at every price-increment but the total size might be smaller than it used to be, i.e. smaller market depth (Sandhagen, 2012). However, several papers suggest a tightening of the bid-ask spread indicating reduced adverse selection and increased liquidity in line with Aldridge (2010). Further, the speed of HFT-algorithms might increase the market resilience, which, as mentioned, is also important for market liquidity. Overall, there are several indicators pointing towards beneficial effects of HFT, but the vague definition of liquidity generates some doubts about the true effects.

Liquidity is the second parameter we aim to use to use in our analysis. Both the academic literature and market participants agree that ample liquidity is very important for healthy financial markets. Hence, together with volatility and market efficiency we consider it to be a necessary part of any market quality analysis. The second hypothesis is in line with most previous research on the parameter. Even though there are ambiguities regarding the definitions we agree with the general conception that liquidity should be improved as a result of HFT. The hypothesis is therefore stated in the following way:

Hypothesis 2: Market liquidity has improved as a result of high-frequency trading.

C. Market Efficiency

The overall efficiency of financial markets is important for all market participants. First and foremost, an efficient market is fundamental to have a "fair market" for investors, in that inside information should not create investment opportunities. Further, efficient financial markets are important in that they ease the transfer of funds from lenders to borrowers so that they are used in a way that is most socially useful.

The most important theory in the area of market efficiency is the Efficient Market Hypothesis (EMH). Developed by Fama and Samuelson independently during the 1960's, the EMH has had a large impact on how the financial market has been viewed in academic literature over the past five decades. Three forms of market efficiency have been put forth since then: the strong, semi-strong and weak form of market efficiency. In the weak form of efficiency it should be impossible to earn abnormal returns by using previous price data; in the semi-strong form all publicly available information is assumed to be incorporated in the price of assets; in its strongest form all information should be incorporated and hence it should be impossible to earn abnormal returns even with private information.

Samuelson (1965) showed that in an informationally efficient market, where the price reflects all available information, price changes are impossible to forecast. One statistical implication of this is that stock prices should follow a random walk, i.e. that the expected return of a security, conditional on the information of previous realized returns, is always equal to the unconditional expectation:

$$E(\tilde{r}_t | r_{t-1}, r_{t-2} \dots) = E(\tilde{r}_t) \text{ for all } r_{t-1}, r_{t-2} \dots.$$

Since the advent of the EMH several papers has aimed to test whether or not the random walk property of stock prices hold. Fama (1970) finds support for this property when testing the different forms of market efficiency, observing little serial correlation in the sample. However, LeRoy (1973) show that if the required rate of return is not exogenously given, as in Samuelsson (1965), theoretical proof of the martingale property of stock prices cannot be derived. Nevertheless, he concludes that the martingale property works well as an approximation of how efficient capital markets behave. Although it would be hard to argue the case that capital markets are fully efficient and even harder to proof such statement, the degree of market efficiency over time serves as an interesting topic in a time when the markets are going through a major change.

Whether HFT makes markets more or less efficient is rather hard to predict ex-ante. The increased speed of trading introduced by HFT-algorithms might increase market efficiency in that prices reflect their fundamental values faster. Furthermore, some HFT strategies make sure that the law of one price holds, by keeping stock prices at different marketplaces the same by exploiting any discrepancies. However, it is argued that there are some deviant HFT strategies that try to exploit other market participants by some degree of price manipulation, or front-running, which could cause market disruptions and make them less efficient.

Given the fact that market efficiency is an important characteristic of market quality and the lack of research in this area in relation to HFT we consider this an interesting parameter in this context. While the effects, given our reasoning above, is somewhat difficult to predict we expect, in line with previous papers in the area, that the positive effects are stronger and that market efficiency therefore has increased as an effect of HFT. Hence, our final hypothesis is:

Hypothesis 3: Market efficiency has improved as an effect of high-frequency trading.

III. Method

A. Data

The data used in this paper is based on a unique dataset provided by Avanza Bank AB. It covers all stocks on the Stockholm Stock Exchange's Large and Mid Cap lists between January 1st, 2006 and February 22nd, 2012. Even though the data is provided by Avanza it is not restricted to their trades but covers trades by all members on the Stockholm stock exchange. Due to the nature of the dataset a slight survivorship bias might be present. The reason is that our dataset contains all stocks currently listed on NASDAQ OMX Stockholm, and hence does not contain stocks that have been delisted (independent of the reason for a delisting). However, we believe that the effect of this survivorship bias will be small due to the research design used in this paper. Companies that have been listed during the time frame of our dataset have been included in the dataset from that point in time. The classification of firms into Large and Mid Cap is based on the date, February 8th, 2010, at which INET (the new trading system at OMX Stockholm) was implemented. From there on we do not allow securities to move between the lists, even if it has changed after or before the introduction of INET. The reason for this approach is mainly due to the setup of our difference-in-difference (DD) test, which will be described below.

The dataset includes actual trade data on the following variables for each minute: highest and lowest price; opening and closing price; value-weighted average price (VWAP); and the number of stocks traded; this result in 784 110 observations for each security (that has been listed throughout the whole period). It is important to note that since the data is only based on actual trades and not outstanding bids and asks, there are minutes for which there is not any trading activity in a security. This is a large issue for some companies that are not frequently traded, which generally means either the A-shares of a company or some of the smaller Mid Cap companies. We replace such missing observations with the last traded price to facilitate the variance-ratio test described below. We argue that when no trade has taken place in a security the best indicator of the current price is the most recent price at which the security was traded. Some securities are therefore showing an almost constant price, only changing a few times a day. This could skew our results and hence we limit our dataset to only include securities for which we have more than 30 000 observations over the whole sample period, i.e. more than about 20 trades per day. The final dataset therefore includes 68 and 53 securities from the Large and Mid Cap lists, respectively.

The dataset provided by Avanza is complemented with data from Thomson Reuters Datastream (Datastream) for the same period and sample of securities. We retrieve daily bid-ask quotes and total number of outstanding shares. Finally, we also retrieve data indicating corporate actions taken by a company, e.g. stock splits. This is used to account for such actions so that our results are not skewed.

We present descriptive statistics for a sub-sample of the dataset, i.e. the event period surrounding the INET implementation, and the entire dataset in Tables I and II, respectively.

Table I: Descriptive Statistics – DD parameters

This table presents the mean, median and standard deviation of the four measures used in the difference-in-difference regressions in this paper, i.e. the Standard Deviation, Trading Range, Stock Turnover and Quoted Spread; for definitions see the Section III, subsection C and D. The measures are reported for Mid and Large Cap, respectively, for a period of six months before and after the INET implementation. We also report these statistics for the entire event period, i.e. from August 10th, 2009 to August 9th, 2010. All statistics are in percentages.

	Before INET		After INET		Total Event Period	
	Mid Cap	**Large Cap**	**Mid Cap**	**Large Cap**	**Mid Cap**	**Large Cap**
(in %)	**Hourly Standard Deviation**					
Mean	0.1202	0.0928	0.1269	0.0816	0.1235	0.0872
Median	0.1046	0.0817	0.1051	0.0680	0.1048	0.0747
Std. Dev.	0.0906	0.0579	0.1053	0.0624	0.0983	0.0605
N	59 201	75 956	59 466	76 296	118 667	152 252
(in %)	**Trading Range (minute average per hour)**					
Mean	0.0555	0.0483	0.0642	0.0482	0.0599	0.0482
Median	0.0185	0.0368	0.0216	0.0364	0.0200	0.0366
Std. Dev.	0.1088	0.0570	0.1442	0.0573	0.1278	0.0572
N	59 201	75 956	59 466	76 296	118 667	152 252
(in %)	**Hourly Stock Turnover (% of outstanding shares)**					
Mean	0.0366	0.0506	0.0343	0.0539	0.0354	0.0523
Median	0.0151	0.0278	0.0134	0.0281	0.0142	0.0279
Std. Dev.	0.0928	0.0876	0.0872	0.0965	0.0900	0.0922
N	59 201	76 296	59 466	75 956	118 667	152 252
(in %)	**Daily Quoted Spread (% of mid-point quote)**					
Mean	0.8437	0.5773	1.4221	0.7438	1.1360	0.6611
Median	0.5380	0.2719	0.8130	0.2608	0.6667	0.2670
Std. Dev.	1.0361	1.0735	1.8310	1.3697	1.5194	1.2343
N	8 351	9 610	8 532	9 731	16 883	19 341

Data source: Avanza Bank AB and Datastream

Table II: Descriptive Statistics – Full Sample

This table presents the mean, median and standard deviation of minute VWAP returns (in basis points) and the number of shares traded per minute for Large Cap and Mid Cap, respectively. The VWAP returns are adjusted for corporate actions. Missing prices has been replaced by the last observed price. The sample covers a period from January 1st, 2006 and February 22nd, 2012.

	Large Cap		Mid Cap	
	VWAP Return (‰)	**Shares Traded**	**VWAP Return (‰)**	**Shares Traded**
Mean	0.0105	5 939	0.0114	750
Median	0.0020	3 391	0.0000	347
Std. Dev.	4.0643	13 253	3.5114	2 609
N	783 599	783 599	783 599	783 599

Data source: Avanza Bank AB

B. High-Frequency Trading in Sweden

1. Structural Changes of the Stockholm Stock Exchange

Before discussing the structural changes in Sweden, some background of trading systems in general and the Stockholm stock exchange in particular might be useful. The central part of stock exchanges is the trading platform which links the market participants together so that they can trade with each other; one can even say that the trading platform is the exchange itself. While many exchanges historically used the open outcry system on the trading floor, where information was transferred using shouting and hand signals, the trading today mostly takes place in automated electronic trading systems. In Sweden the automated trading system SAX (Stockholm Automated Exchange) was implemented on June 2nd, 1989 and this meant that the members no longer needed representatives at the physical exchange. In 1999 this system was replaced by SAXESS, which could handle different types of securities such as equities, derivatives and bonds. Even though the base structure remained the same the SAXESS system was upgraded a number of times, which among other things resulted in reduced latency (Gaudy, 2012). In 2008 OMX was acquired by the American stock exchange NASDAQ and a so called "technology roadmap" was introduced for the integration of the technical systems (Malmqvist, 2009); this lead to the implementation of the INET trading system on February 8th, 2010. The system was rolled out in NASDAQ OMX's exchanges in the Nordics and Baltics and this trading system shift was one of the biggest infrastructural changes in the history of these equities markets (NASDAQ OMX, 2010b).[4] INET had been used for some time on NASDAQ's other exchanges both in the US and UK. It can handle one million messages (information sent to the exchange to buy, sell or cancel orders) per second at speeds below 250 microseconds and is thereby the fastest trading system in the world (NASDAQ OMX, 2010b). The upgrade from the old SAXESS trading system was rather significant since the order speed was reduced from 2.5 milliseconds, i.e. a reduction of the factor 10. Further, the number of orders that the new system can handle has increased, from 1,500 to 7,000 per second. Another new feature in INET is the "order routing"-function, which enables members to push an order to other marketplaces to receive the best possible price in the market.

The technical structure of the INET system is very different from SAXESS and hence the change had a large impact on all members (Gaudy, 2012). However, while the differences were mostly on the infrastructural level the new system allows for significantly reduced latency, as mentioned above, and increased stability. These structural changes therefore become especially important for trading strategies highly dependent on speed, i.e. HFT. The implementation of the INET platform therefore significantly increased the attractiveness of using HFT-algorithms on the Nordic exchanges. The tick sizes for the stocks on the Large Caps were reduced in connection with the INET implementation; this is also advantageous for HFT, since such strategies are based on opportunities that are only small fractions of the share prices. It is therefore reasonable to assume that the structural change of the Stockholm stock exchange, due to the implementation of INET, significantly affected the share of HFT in the market. Even though the share might had increased over time previous to February 2010, it is reasonable to assume that this event, due to its beneficial HFT-features, had an effect on the trend and that many HFT players entered the Nordic markets generally, and Stockholm specifically, after the implementation.

To confirm our conjectures regarding the effects of INET on HFT we met with one of the technical departments on NASDAQ OMX Nordic. They played an important part in the implementation of INET and are also responsible for the co-location services used by many HFT firms. Richard Gaudy (2012), Managing Director at this department, argues that INET had a large impact on the share of HFT on the Stockholm stock exchange and once it was launched a number of firms entered the market. The old SAXESS system was not as attractive for such trading; the speed and technical structure of this system did not really meet their requirements. Additionally, the number of co-location customers increased significantly following the implementation of INET (Gaudy, 2012). Since these services are mainly used by firms employing

[4] The Nordic exchanges are Copenhagen, Helsinki, Iceland and Stockholm and the Baltics are Riga, Tallinn and Vilnius.

HFT-algorithms this further indicates the impact of INET. Even though co-location has been available since 2004, back then it was called proximity, it was not until INET and HFT that there was any wider interest in such services. Overall, we argue that these structural changes on the Stockholm stock exchange was both important and beneficial for HFT and that it had a significant effect on the share of the trading performed using such strategies.

2. High-Frequency Trading Proxy

To evaluate effects of HFT and obtain causality in the analysis we need a measure of the share of HFT in the market. However, access to and availability of such data is very limited and is therefore unfortunately not included in our dataset. This is due to a number of issues such as a strong desire for HFT players to keep their actions as anonymous as possible – i.e. by trading through other direct-connected participants, so called direct market access. Also, the enormous amount of data that needs to be analyzed, since one would need to use tick data, is problematic. The amount of raw data needed to obtain a measure of HFT on OMX Stockholm is estimated to be 10 to 20 TB for the time period used in this thesis (Hertzberg, 2012). Hence, such an analysis is outside the scope of this thesis.

As we do not possess an actual measure of the HFT share we need a proxy as is the procedure in many other papers in the area. The proxies used in these papers are of different kinds and qualities. The proxy we are using is not an actual measure per se, but rather it builds on the distinct characteristics of the securities on the different lists on OMX Stockholm. The Stockholm stock exchange is divided in three different segments depending on the size of the companies: Large Cap (share values over EUR 1 billion); Mid Cap (share values between EUR 150 million and EUR 1 billion); and Small Cap (share values up to EUR 150 million). These segments have rather different characteristics in terms of trading, as the volume traded and available liquidity on Large Cap is significantly higher than for Mid and Small Cap. The share of the turnover in terms of number of trades on the Large Cap on Stockholm stock exchange (excluding OTC) in January 2010[5] was 82%, compared to 11% and 6% for Mid and Small Cap, respectively. The same pattern is shown in the average daily statistics for the same month and year; where Large Cap stands for 64% of the number of shares traded and Mid and Small Cap stands for 10% and 25%, respectively.[6] Finally, Large Cap has 91% of the total turnover in SEK for this month (NASDAQ OMX, 2010a).

---

[5] Statistics for January 2010 is used due to the date of the INET implementation and the Difference-in-difference method. Details of this procedure are described below in the volatility section.

[6] The last percentage point in these statistics is made up of the shares traded on the external list for foreign companies. Also, the reason for Small Cap having a higher share of shares traded is their relatively low price.

Many market participants use trading volume as a proxy for liquidity; in the light of the above statistics we can hence assume that the liquidity is significantly higher on the Large Cap than on the other lists. This is also confirmed by market participants, representatives for the stock exchange and regulators.[7] Due to the structure of HFT strategies, e.g. need for liquidity and volume,[8] it is therefore reasonable to assume that practically all such trading takes place on the Large Cap. Hence, the segmenting of OMX Stockholm generates a natural proxy for HFT, where we are assuming that such trading is taking place on Large Cap stocks and not on Mid and Small Cap stocks. We are aware that this assumption is not perfect and that there are in fact some HFT on both Mid and Small Cap. However, the assumption made is that the vast majority of this trading is performed on the Large Cap. Further, this proxy assumption is sufficient due to the structures of the methods used in this thesis. The details of these methods used to examine HFT effects on market quality are described below, but are in essence based on analyses of differences in the parameters between the two groups of stocks.

Since this is a central part of the thesis we have performed a number of interviews to confirm the proxy assumption. The interviewees are either HFT players themselves or have significant knowledge of the market due to their positions as representatives for the Stockholm stock exchange (NASDAQ OMX Nordic) or market regulators (Finansinspektionen). They are all confirming the above reasoning that Small and Mid Cap characteristics are not enough for HFT and that it is therefore performed mostly on shares on the Large Cap. The amount of trading performed on the other lists is just not enough to be interesting for HFT-algorithms. The low volumes and liquidity pose a significant risk for such algorithms that requires ability to enter and exit positions whenever needed. Not being able to act as desired for fractions of seconds might mean success or total failure for these algorithms. However, normal "low frequency" algorithms might be able to trade on these lists since they do not have the same urgent needs to enter and exit positions.

To further confirm our proxy assumptions we have investigated the market shares of some of the most well-known actors within the HFT-field in the Nordic region and Sweden. However, due to the strong secrecy surrounding these companies the available data is rather scarce. One such firm is Citadel Securities, which perform a number of different high-frequency strategies including market making. In January 2010 they had a market share of the number of trades on NASDAQ OMX Nordic of 4.1%. Their market share in total turnover was 3.9%, making them the 8th largest player on these exchanges at the time. Citadel performed about 98.6% of their trades and 99.6% of the total turnover on the Large Caps of the Nordic exchanges. Another,

[7] This topic was extensively discussed during our interviews with Julander, Gaudy and Sandhagen (2012).
[8] Additional details of how HFT strategies are working are described closer in the introduction.

smaller, algorithmic player called International Algorithmic Trading, with a turnover market share of 0.38%, has about the same distribution of trades; 94.5% of their trades in January 2010 were done on the Large Caps (NASDAQ OMX, 2010a). The growth of HFT over the past years has led to the entrance of a number of new players within the area on the Nordic exchanges, e.g. Spire Europe, Getco Europe and Virtu Financial whom all seems to work within the area (Sandhagen, 2012).[9] In January 2012 they had a market share of total turnover on the Nordic exchanges of 3.86%, 3.09% and 2.29%, respectively. Hence, they all belong to the 20 largest members of the Nordic exchanges. The trading patterns of these companies further confirms our assumptions as their shares of trades performed on the Large Caps are 97.4%, 97.0% and 97.4% for Spire, Getco and Virtu, respectively (NASDAQ OMX, 2012).

In conclusion, the HFT share assumption we are using in this thesis is supported by both exchange-wide and company specific statistics; it is further supported by business professionals with significant knowledge of both markets generally and HFT specifically. We are therefore confident that the Large and Mid Cap lists works well as a proxy and can be used to test effects of HFT in Sweden. There are two reasons for excluding the Small Cap securities from the analysis. Firstly, the trading in these securities is too low and hence prices become artificially stable in our dataset. Secondly, these shares have characteristics that are poorly matched with those of the Large Cap securities; hence, these securities would serve as a poor benchmark.

C. Volatility

The first parameter used to evaluate market quality is volatility. There are a number of different methods to determine stock volatility, including implied and realized measures as well as model- and non-model dependent measures. We follow the recent tendency in financial research to use model-free measurements of actual return variation, i.e. ARCH and GARCH models will not be used. The structure of our dataset allows for different ways to measure the variation of stock returns. While the typical definition of market volatility is the historical standard deviation of stock returns we are also using the intra-minute price variation to measure volatility. Hence, two different measures of market volatility will be used in the analysis, inspired by the ones used in Hendershott and Moulton (2009). The reason for using two measures is that we want to capture different effects in the price movements and thereby looking at all potential effects of HFT.

The first measure of stock volatility is the one-minute trading range, which is the highest minus the lowest price during the minute divided by the closing price of the minute. This

[9] The quest for secrecy again makes it difficult to determine the companies exact operations, i.e. which markets they operate in and what kind of trading they do.

measure is then averaged for each stock and hour of the trading day (*i* is firm, *t* is hour and *n* is minute):

$$TR_{i,t} = \frac{\sum_{n=1}^{60} \frac{P(high)_{i,n} - P(low)_{i,n}}{P(close)_{i,n}}}{60}$$

Hendershott and Moulton (2009) uses a five-minute interval and daily averages, but we argue that one-minute intervals and hour averages better captures any potential intra-day effects caused by HFT. The second measure is based on the quote return volatility in Hendershott and Moulton (2009) and is the hourly standard deviation of minute stock returns:

$$SD_{i,t} = \sqrt{\frac{\sum_{n=1}^{60}\left(r_{i,n} - \overline{r_{i,n}}\right)^2}{(60-1)}}$$

The minute stock returns are computed in the following standard procedure:

$$r_n = \frac{P_n}{P_{n-1}} - 1$$

While they use the mid-quote returns we are instead using the minute-average price (VWAP) return. The VWAP return shares the reduced effect of bid-ask bounces, i.e. price jumps created by the fact that the last trade could be instigated from the bid or ask side, with the mid-quote price return. Hence, the outcome of our procedure should be rather similar. The trading range measure captures the most extreme price movement during the minute, including noise from bid-ask bounces and high-frequency price jumps, and the quote return volatility is the more commonly used measure of volatility not as affected by such noise.

To evaluate any potential effects of HFT on the volatility measures described above we intend to use a DD approach. This method is commonly applied in evaluation research to study effects of public interventions. Hendershott and Moulton (2009) use it to evaluate effects of a change in execution time for market orders on the New York Stock Exchange by comparing the difference between NYSE and NASDAQ stocks before and after the change. Hendershott et al. (2011) are also using this method to examine the effects of algorithmic trading on liquidity by using a change in market structure (the implementation of automatic quote dissemination in 2003) that increased algorithmic trading as an exogenous instrument. Zhang (2010) goes even further and uses both DD and difference-in-difference-in-difference approaches to examine the impact of HFT on stock volatility. He uses a main sample period and an estimation period to compare differences in volatility over time. Hence, the use of the DD approach in the research field of HFT is somewhat established. The combination of the previous use of the method, with the fact that we have identified a structural change in the Swedish financial market, leads us to conclude that the DD approach is the most appropriate one to assess the impact of HFT on the volatility in the Swedish financial market.

The DD method uses panel data in such a way that causal inference might be identified. This method is an effective technique used to evaluate the effect of a treatment on a specific variable by using a treatment and control group, but as Abadie (2005) states: it is based on rather strong assumptions. The average outcome of the investigated variable of the two groups does not have to be the same per se, but the difference in absence of treatment has to stay about the same before and after the event for the method to be applicable. In other words, the method rests on the crucial assumption that this variable for the treatment and control groups would have followed parallel paths over time; this is known as the Common Trends Assumption or CTA (Angrist and Pischke, 2009). A second important assumption underlying this method according to Bach (2012) is called the No Anticipation Assumption (NAA) and states that the effect of the event must not be realized until the event has occurred. If this is not true evaluation of the effects become difficult, if not impossible.

The framework used to obtain the DD estimator is explained by Angrist and Pischke (2009) and it is in fact a version of a fixed effects estimation using aggregate data. They show that the heart of the DD method is an additive structure for potential outcomes for the control group. This structure is used to build a model where the effect of treatment can be identified by using the change in outcome for the non-treated and treated groups and the aforementioned assumption about common trends. The basic DD model is the following:

$$Y_{it} = \gamma_i + \lambda_t + \delta D_{it} + \varepsilon_{it}$$

Where Y_{it} is the outcome of the observed variable for firm *i* in time *t* and γ_i and λ_t represents firm and time specific (fixed) effects. D_{it} is a firm-time dummy indicating treated firms in the post-treatment period and the structure of the error term is such that $E(\varepsilon_{it}|i,t) = 0$. This makes δ the estimator of the DD effect and hence the causal effect of interest.[10] The DD effect is possible to estimate using sample estimates of the population means in line with the model above, but we can also use a regression framework for this. The regression model used to estimate the above equation is the following:

$$Y_{it} = \alpha + \gamma D_i + \lambda D_t + \delta(D_i * D_t) + \varepsilon_{it}$$

Where D_i and D_t are firm and time dummies for treated firms in the post-treatment period. Again δ is the coefficient of interest as it measures the effect of treatment on the treated.

[10] Note that we only describe the basic setup of the DD framework in this thesis, for a closer description and derivation please see Angrist and Pischke (2009) or Abadie (2005).

As mentioned previously the DD method rests, together with assumptions regarding trends in outcomes, upon the identification of an event (or treatment) for a specific group. As we discussed above the structural change by the implementation of the INET trading platform had important implications for HFT. The argument that there was a noticeable impact in the share of HFT in the market due to this event is therefore reasonable. Consequently, the INET implementation serves as a suitable event in a DD approach to investigate HFT. The discussion above on the HFT proxy lead us to conclude that the Large Cap securities could serve as a treatment group, while the Mid Cap securities is the control group. This result in the specific DD regression-model we are using in this thesis:

$$Vol_{it} = \alpha + \gamma L_i + \lambda Post_t + \delta(L_i * Post_t) + \beta X_{it} + \varepsilon_{it}$$

Where Vol_{it} is the volatility measure, according to the specifications above, for firm *i* in time *t*, L_i is a dummy indicating whether the firm is listed on Large Cap or not and $Post_t$ is a time-dummy indicating observations after the implementation of INET (February 8th, 2010). The two dummies are interacted to capture the treatment effect. Finally, X_{it} represents the controls used in our model, namely firm-fixed and time-fixed effects; the time fixed effects are captured on a weekly level.

The DD tests will be performed on a sub-sample of the larger dataset to isolate the effects of the INET implementation. Hendershott and Moulton (2009) use data for about eight months surrounding their event, but to make sure that we capture both immediate and potential delayed effects of the event we use data for about six months before and after the implementation, from August 10th, 2009 to August 9th, 2010. Furthermore, since return series tend to exhibit heteroskedasticity we cluster our standard errors on a firm level to account for this. Since we have two different specifications of volatility we will run the same model on both to capture any potential effects of HFT.

D. Liquidity

The second parameter of market quality is liquidity; as is the case with volatility this parameter can be measured in a number of different ways. The most common measure mentioned in the majority of papers discussing liquidity, and also mentioned by Aldridge (2010) as one of the best estimators of liquidity, is the bid-ask spread. The specific bid-ask spread measure we are using is called the quoted spread (QS) and is expressed as a proportion of the quote midpoint. This specific measure has been used in various papers looking at bid-ask spreads, e.g. Stoll (1989) and Lee et al. (1993), and also papers looking more specifically at liquidity; Kale and Loon (2011) use it as one of their parameters of liquidity when assessing the impact of market power on stock liquidity. The quoted spread for each stock and day is computed in the following way (*i* is firm and *t* is day):

$$QS_{i,t} = \frac{Ask_{i,t} - Bid_{i,t}}{(Ask + Bid)_{i,t}/2}$$

The quoted spreads are measured on a daily frequency, instead of hourly as with the volatility measures, due to the absence of bid and ask prices in the higher-frequency dataset. To further establish any potential effects of HFT on liquidity we use an additional measure of liquidity. This second measure is the hourly number of shares traded in proportion to the total number of shares outstanding, called stock turnover. As mentioned in the theory section many market participants, especially HFT actors, use different measures of volume as a proxy for liquidity. While it might be a rough proxy we argue that the stock turnover could be used as a measure of the ability for market participants to trade at a time preferable to them. Hence, a higher turnover is beneficial for participants that want to be able enter and exit positions as easy as possible. The stock turnover is computed using the following formula (*i* is firm, *t* is hour and *n* is minute):

$$ST_{i,t} = \frac{\sum_{n=1}^{60} ST_n}{SO_{i,t}}$$

Where ST_n is the number of shares traded per minute and SO_{it} is shares outstanding for each firm and hour.

These measures of liquidity will be evaluated using the same method as for the volatility, i.e. a DD approach. The event and regression model used are the same but for the dependent variable, which is now the liquidity measures instead of volatility. Hence, the model used is the following:

$$Liq_{it} = \alpha + \gamma L_i + \lambda Post_t + \delta(L_i * Post_t) + \beta X_{it} + \varepsilon_{it}$$

Where Liq_{it} is the liquidity measure, using aforementioned definitions, and the other variables are defined as per above.

E. Market Efficiency

The third and final parameter of market quality we identified is market efficiency. Several tests of the efficient market and the random walk hypotheses have been developed since their introduction. The one used in this paper is a version of a test developed by Lo and MacKinlay (1988). They tested the random walk hypothesis for weekly returns in the U.S. market by comparing variance estimators computed from data sampled at different intervals. Following their model[11], denote by P_t the stock price at time *t* and define $X_t \equiv \ln(P_t)$ as the lognormal

[11] We just give a short review of the Lo and MacKinlay variance-ratio test and focus on a subset of the statistics presented relevant for this paper. For full derivations of their results please refer to Lo and MacKinlay (1988).

price process – what follows is a short description, using their terminology, of the variance-ratio test under the assumption of homoscedasticity:

$$H: \epsilon_t \; i.i.d. N(0, \sigma_o)$$

$$X_t = \mu + X_{t-1} + \epsilon_t$$

Where μ is a drift parameter and ϵ_t is a random disturbance term. The price change process looks as follows:

$$dX(t) = \mu \, dt + \sigma_o dW(t)$$

Where $dW(t)$ denotes a standard Wiener process. One important implication of the random walk (X_t) is that the variance of its increments is linear in the time interval – hence, the variance of $X_t - X_{t-1}$ is half of that of $X_t - X_{t-2}$. It is by using this property, that Lo and MacKinlay develop a formal test of the random walk model. Consider a sample of $nq + 1$ observations, $X_0, X_1 \dots X_{nq}$, where $q > 1$ and define the following estimators:

$$\hat{\mu} \equiv \frac{1}{nq} \sum_{k=1}^{nq} (X_k - X_{k-1}) = \frac{1}{nq} (X_{nq} - X_0)$$

The estimator of the sample variance using every observation is:

$$\hat{\sigma}_a^2 \equiv \frac{1}{nq} \sum_{k=1}^{nq} (X_k - X_{k-1} - \hat{\mu})^2$$

Whereas the estimator of the sample variance using every qth observation is given by:

$$\hat{\sigma}_b^2 \equiv \frac{1}{nq} \sum_{k=1}^{n} \left(X_{qk} - X_{qk-q} - q\hat{\mu}\right)^2$$

With these estimators the following test statistic, the variance ratio minus one, can be defined:

$$J_r(q) \equiv \frac{\hat{\sigma}_b^2(q)}{\hat{\sigma}_a^2} - 1$$

Under assumption of homoscedasticity and independent Gaussian increments the following asymptotic distributions are given for $J_r(q)$:

$$\sqrt{nq} \, J_r(q) \sim N(0{,}2(q-1))$$

Hence, by using this statistic we can test if the price process in fact follows a random walk. However, using only every qth observations puts a strict constraint on a finite sample, therefore Lo and MacKinlay refine the statistics by using overlapping differences and define the following estimator of σ_o^2:

$$\hat{\sigma}_c^2 \equiv \frac{1}{nq^2} \sum_{k=q}^{nq} \left(X_k - X_{k-q} - q\hat{\mu}\right)^2$$

This, instead of using n observations, uses $nq - q + 1$ observations. Using this estimator in the variance-ratio test gives the following test statistics for the random walk model:

$$M_r(q) \equiv \frac{\hat{\sigma}_c^2(q)}{\hat{\sigma}_a^2} - 1$$

Another refinement made is to use unbiased variance estimators when calculating the M-statistic. More interesting though is the intuition for the M-statistic that they develop, using $q = 2$, i.e. using two period increments for estimating $\hat{\sigma}_c^2(q)$, they show that:

$$M_r(2) \backsimeq \hat{\rho}(1) + \frac{1}{4n\hat{\sigma}_a^2}[(X_1 - X_0 - \hat{\mu})^2 + (X_{2n} - X_{2n-1} - \hat{\mu})^2] \backsimeq \hat{\rho}(1)$$

Hence, for the case of $q = 2$ the $M_r(q)$-statistic is approximately equal to the first order auto-correlation coefficient estimator. Therefore, by using the variance-ratio test not only the correctness of the random walk hypothesis is tested, but the statistics themselves gives an intuitive interpretation of how the market is behaving. Put simply, one can detect if there is a momentum effect (positive autocorrelation) or mean-reversion effect (negative autocorrelation) present in the price process of a security. For higher orders of aggregation values (q), the $M_r(q)$-statistics are (approximately) a combination of the first $q - 1$ autocorrelation coefficient estimators with declining weights. Since much literature has found that most return series contains heteroscedasticity this paper will make use of the heteroscedasticity robust variance ratios in Lo and MacKinlay's article. Since the principle and intuition behind the heteroscedasticity robust variance-ratios is the same as for the homoscedastic ones, the full derivation is not presented in this paper.

Further improvement on the methodology of the variance ratio test was made by Chow and Denning (1991), who account for the fact that for the random walk hypothesis to hold, the variance-ratios need to equal one for all the aggregation parameters (q) considered. Hence, when considering several different aggregation values, a more appropriate test method would be to test the variance-ratios jointly. In their paper they propose the following test statistics:

$$Z^*(q) = \max_{1 \le i \le m} |Z(q_i)|$$

Where q_i corresponds to a set of aggregation parameters such that $2 \le q_i \le \frac{N}{2}$ for all $i = 1,2 \dots m$, where N is the sample size and the $Z(q_i)$ is the variance-ratio statistic as defined as in Lo and MacKinlay. The asymptotic critical value for this statistic is given by the studentized maximum modulus (SMM) distribution, $SMM(\alpha; m; \infty)$. The asymptotic SMM critical value can then be calculated from a standard normal distribution by:

$$SMM(\alpha; m; \infty) = Z_{\alpha^+/2} \text{ where } \alpha = 1 - (1 - \alpha)^{\frac{1}{m}}.$$

By running simulations they show that the probability of type I errors, i.e. rejecting the null hypothesis when it is actually true, is much larger when failing to control for the joint test size. By using their methodology instead, this risk is reduced significantly.

To a large extent the usage of variance-ratios to analyze the efficiency in the stock market after the introduction of HFT is based on an article by Castura et al. (2010). The main difference, with respect to methodology, between our paper and theirs, is that we have a specific date after which there is reason to believe that the level of HFT is increasing, i.e. the implementation of INET. Also, since we have argued that HFT is more or less only present on the Large Cap we can let Mid Cap-listed securities serve as a benchmark in our analysis.

Since Lo and MacKinlay's variance-ratios is a measure of the serial correlation present in the price process of a security we can observe the efficiency of the market by studying how close to unity this ratio is over a selected time period (as done in Castura et al., 2010). As mentioned, a variance-ratio larger than one would be a sign of a momentum effect, and contradictory a ratio smaller than one would be a sign of a mean reversal effect, in the price process of the security. Therefore, in the analysis that follows, we will present the average variance-ratio over each quarter for the securities with a supposedly higher amount of HFT and compare this to the evolution of the variance-ratios of less HFT intense securities – by doing this any effect that HFT has had on the efficiency of the Stockholm stock exchange can be uncovered. The variance ratios will be plotted for several different values of q to account for different orders of autocorrelations. The different intervals (aggregation values, q) considered is 2 minutes/1 minute; 5 minutes/1 minute; 10 minutes/1 minute; and 4 minutes/2 minutes – preferably one would like to consider even higher frequencies than minutes as well, due to the very high speed at which HFT take place. However, due to data availability higher frequencies will not be considered in this paper.

As a more formal test of the efficiency individual Chow-Denning tests will be performed on all the stocks individually in the sample for each quarter of the examined period. The Chow-Denning test will use variance-ratios calculated with the same aggregation parameters as mentioned above. By doing this one can see if the relative efficiency of the stock market has changed, i.e. if there is any trend, especially surrounding the implementation of INET in February 2010, were more (or less) stock's price processes reject the random walk hypothesis.

As noted in Castura et al. (2010), one potential downside of using high-frequency data is that micro-structural effects are expected to be present. For instance, the bid-ask bounce skews results as appearing mean reverting; these effects are decreasing in the sampling rate which means that they are not as prevalent in our sample, albeit still present. However, this is mitigated to some extent by using the VWAP, as mentioned above. Furthermore, assuming that the mean reverting effect is relatively constant over time it does not affect our results drastically, since what is most interesting is the evolution throughout the sample period, between Mid and Large Cap, and not the variance-ratios per se.

IV. Results and Analysis

A. Volatility

1. Method robustness

Before interpreting and analyzing the results of the DD approach we need to make sure that the results are valid, i.e. that the underlying assumptions of the model are fulfilled. As described in the method section the two underlying assumptions for this method are the Common Trends Assumption (CTA) and the No Anticipation Assumption (NAA). In our setting the NAA is fulfilled by the construction of the event; the implementation of INET was a structural change of the market taking place at a specific date. Hence, the system was not available for trading on the Nordic exchanges before the implementation. This means that the NAA is fulfilled for all variables used in the analysis, including the liquidity variables. The procedure to evaluate the credibility of the CTA is somewhat different as one has to make individual evaluations for each variable. The general practice is to plot the trends of the variables of interest for the treatment and control groups and thereby determine whether they are similar or not (Angrist and Pischke, 2009).

Figures VII and VIII in the appendix shows the two volatility parameters we are using, the one-minute trading range and the standard deviation of stock returns; in these we observe rather similar trends between Large and Mid Cap for the six-month period before the event. Even though the trends in volatility are not identical one can easily argue that they have common characteristics and developments; the CTA is therefore considered to be fulfilled for these parameters. Further, the regression model we estimate includes time- and firm-fixed effects as control variables, and to confirm the validity of the results we have run the regression models with different frequencies for the time-fixed effects. Altering the model specification gives very small effects on the results, the coefficient changes slightly with kept significance. We therefore argue that the results are robust of simple arbitrary model specifications. Finally, as have been discussed previously it is reasonable to believe that the event had an impact on the share of HFT in the market. All this in combination gives for a result of the market volatility effects of HFT in Sweden that ought to be credible.

2. Volatility effects

The first definition of market volatility was the one-minute trading range, which captures the most extreme price movements. Hendershott and Moulton (2009) argue that such a measure captures the high-frequency transitory volatility, i.e. excessive short-term volatility, as well as the microstructural noise. Hence, this measure should in a very effective way capture any effects that HFT-algorithms working in the market might have had. Even though these algorithms work at a much higher frequency than the minute their effects on prices and thereby trading range volatility will definitely be visible, since the highest and lowest prices during the course of the minute are used. To further elaborate on the effects of HFT we defined another measure of market quality, namely the hourly standard deviation of minute stock returns. This measure is not as affected by the bid-ask bounces as the trading range and is also the measure commonly referred to in discussions of market volatility, even though the frequency of stock returns might differ. Further, this measure might be more familiar to common and institutional investors than the trading range is.

We begin the analysis by comparing the development over the event window for the two volatility measures. To do this we plot the difference between the weekly average of the measures for Mid and Large Cap, respectively, i.e. $\overline{VOL}_{MID} - \overline{VOL}_{LARGE}$. We see from Figures I and II below that the INET implementation seems to have had an effect on both measures. The difference for the trading range measure is somewhat higher after the event; even though it is not overly observable. However, there is an interesting spike in the difference just after the event where the trading range for Mid Cap increased significantly, while Large Cap remained on lower levels. For the standard deviation measure we see that the difference is increasing after the event; this increase is more pronounced than for the trading range measure.

The observations from the figures below are further analyzed using the DD framework, where the variable of interest is the interaction variable; it measures the effect of being a treated firm after the event, i.e. Large Cap companies after the INET implementation. Since we defined two volatility measures, to capture different effects, we also have two variables of interest. Our results for the trading range measure are presented in Table III in the appendix and they show a statistically significant decrease (at the 1% level) in the trading range for Large Cap companies after the implementation of INET. The coefficient on the interaction variable is -0.0088, i.e. the trading range would have been 0.0088 percentage units higher for Large Cap companies in absence of the treatment (HFT). This implies that the effect of HFT is a decrease in the trading range, as the treatment produces a deviation from the common trend illustrated by Mid Cap companies.

Figure I: Difference of Weekly Average Trading Ranges

This figure presents the difference, (in percentage points) between the weekly average trading range for Mid and Large Cap, i.e. $\overline{TR}_{MID} - \overline{TR}_{LARGE}$, including the 95% confidence interval, over the period August 10th, 2009 to August 9th, 2010. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

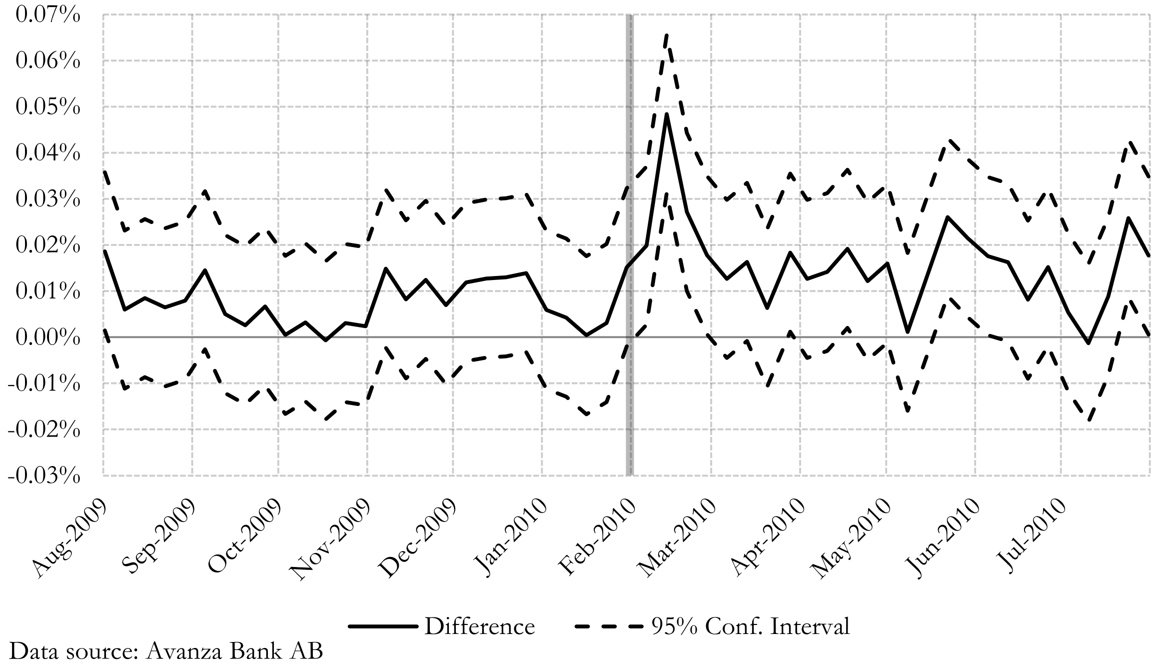

Data source: Avanza Bank AB

Figure II: Difference of Weekly Average Hourly Standard Deviations

This figure presents the difference (in percentage points) between the weekly average hourly standard deviation for Mid and Large Cap, i.e. $\overline{SD}_{MID} - \overline{SD}_{LARGE}$ including the 95% confidence interval, over the period August 10th, 2009 to August 9th, 2010. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

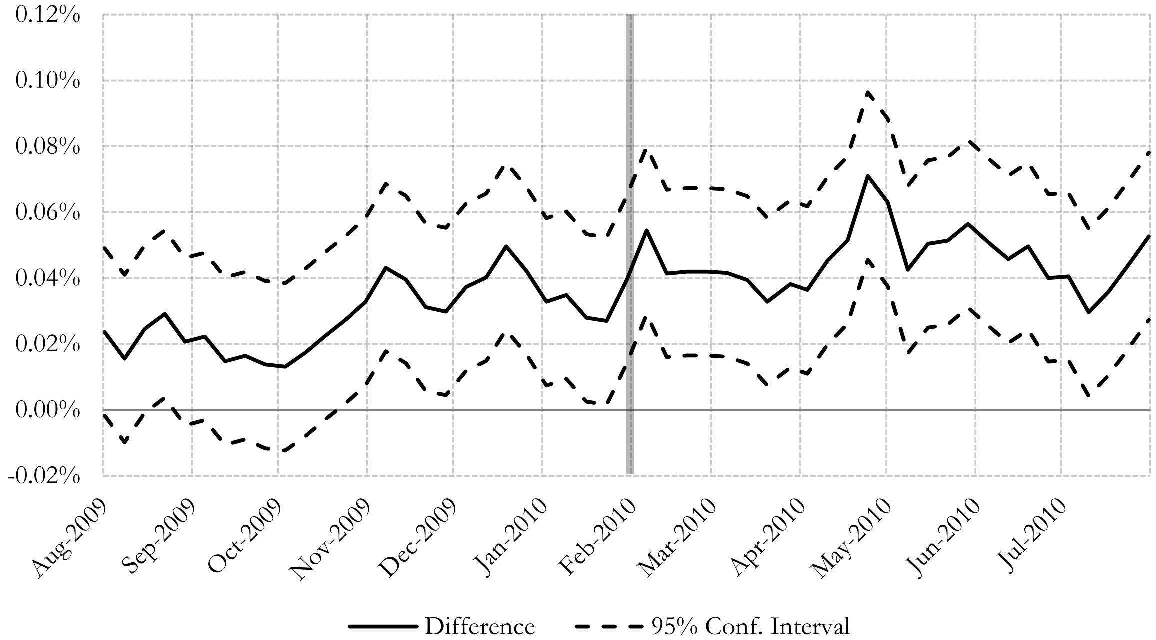

Data source: Avanza Bank AB

Although not a very large effect, we still argue that the economic implications of the result is somewhat significant as the average trading range for Large Cap companies before INET was 0.0483%, with a standard deviation of 0.0570%. This implies that the magnitude of the decrease is about 15% of the pre-event standard deviation. The post-event average and standard deviation of the Large Cap trading range, 0.0482% and 0.0573%, are about the same as the pre-event but it is the Mid Cap increase in the respective statistics that generates the results; this is due to the common trends assumption implicit in the DD framework. We therefore argue that HFT seems to have beneficial effects on the trading range. Hence, this measure, which the academic field considers includes microstructural features of the market that very well may be affected by HFT, has not deteriorated due to HFT. In support of our hypothesis this definition of market volatility has in fact improved as an effect of HFT.

Our DD results for the standard deviation measure are presented in Table III and the implications are about the same as for the trading range. The standard deviation for Large Cap relative to Mid Cap companies after INET has decreased with statistical significance at the 1% level, i.e. the DD estimate is statistically significant at this level. This implies that HFT has an effect also on this measure of market volatility. The economic significance of this result is somewhat higher as the coefficient on the interaction variable is -0.0179; the standard deviation has therefore decreased by 0.0179 percentage units as an effect of HFT. The average hourly standard deviation for Large Cap before the event was 0.0928% and this measure has a standard deviation of 0.0579%; the magnitude of the effect is therefore about 30% of the pre-event standard deviation, and clearly of economic significance.

Overall, we see that the tests of both volatility measures show a significant and negative coefficient on the interaction variable. This is a very interesting result and also in line with our hypothesis; it means that the volatility, independent of the definition, of Large Cap companies has decreased relative to Mid Cap companies after the implementation of INET. This implies that HFT has positive effects on market quality in that it reduces volatility. While the most important result is that HFT does not create volatility in the Swedish market, it is also interesting to see that when HFT is introduced on a large scale in the market the volatility is actually somewhat reduced. Even though the economic significance of these results is not too remarkable we still argue that HFT is beneficial for the quality of the market, at least in terms of volatility. The implication of these results is that exchanges potentially can reduce volatility, to the benefit of members and investors, by facilitating the use of HFT. This implication is further confirmed by Chaboud et al. (2009) as they show that non-HFT trades add more to the variance in returns than HFT trades. This is especially interesting in the light of the recent debate regarding the detrimental effects of HFT, where strong forces in the financial markets have intensely opposed the increasing share of HFT. While they argue that markets become overly sensitive, with substantial risks for disruptions and volatility shooting through the roof, when shares are traded by HFT-algorithms we suggest the complete opposite.

To be able to infer what our results implies for market volatility it is important to understand the general state of the economy at the time of our DD test. Clearly, if the results are obtained in a time where the market economy is very unstable, as for example during the recent global financial crisis, these might have less significance during what we would call normal market conditions. Considering Figure IX in appendix, we observe some turbulence during the period. However, considering the large increases in volatility during the financial crisis, in 2008 and the dot-com bubble in the beginning of the 2000's, the recent volatility pattern is more similar to that of the period 2003-2006. Hence, at a first glance this appears to be a period that is a least somewhat representative of normal market conditions.

There are reasons to be careful to define the period as normal though. The main reason is that the global financial crisis was followed by what we today call the European sovereign debt crisis, during which several Eurozone countries have had problem paying off their debt. Although this crisis was just starting in 2010 it had an impact on the market during our test period. The very high returns that we can see during the period are a reaction to one of the many rescue packages that has been presented during this crisis. Also, during this period, Swedish auto manufacturer Saab entered into the financial difficulties which later made the company bankrupt; although not one of Sweden's largest companies this clearly had an effect on the stability in the Swedish market. However, we argue that the period can be described as one with fairly normal market conditions. While this implies that our results ought to be valid in normal markets it does not have to be the case for periods of turmoil. In such conditions HFT-algorithms might actually intensify market volatility as showed by Kirilenko et al. (2011) in their study of the Flash Crash of 2010.

Our results are in line with most previous research that has shown reductions in volatility due to HFT. However, it is interesting that the volatility measures we are using, which were inspired by a previous paper by Hendershott and Moulton (2009), shows contradictory results to what they find. They show increases in their measures in connection with the introduction of the Hybrid market at NYSE, which expanded electronic execution and thereby increased the speed within the market. Even though their analysis examines effects of speed in the market it does not investigate HFT per se. The Hybrid reduced the execution time from 10 seconds to less than one second which is a significant change indeed, but still well above the microsecond environment where HFT-algorithms are working. Further, general trading patterns and market characteristics are slightly different in the US than in Sweden.

There might be several reasons to why the market volatility has decreased; Zhang (2010) discusses theories concerning the market-making activities by HFT-algorithms. By continuously posting and also instantly providing liquidity in the market these algorithms might reduce the price impact due to trades, i.e. bad volatility; this is especially true for larger trades. This trade matching is done at higher speeds and efficiency than the older manual matching, which might lead to smaller price impacts and lower volatility. Due to their speed advantage the HFT-algorithms might also post liquidity, i.e. order matching, at the inside of the bid-ask spread and the price jumps for these trades are consequently smaller. The market volatility can therefore decrease as an effect of lower price jumps. Further, Hendershott and Moulton (2009) discuss the possibility that increasing the execution speed might affect the rate and efficiency at which information is incorporated into prices. Hence, the price discovery process might be more efficient as the interaction between buyers and sellers to determine the price is faster and smoother; this is due to the continuous and instantaneous market-making by HFT-algorithms. This in turn might lead to fewer and smaller price jumps when new information arrives; the volatility thereby becomes lower. Our results are therefore also in line with Brogaard (2010), who suggest that HFT reduce the intraday volatility by adding to the price discovery process. Finally, Chaboud et al. (2009) do not find a causal relationship between algorithmic trading and increased volatility, more specifically exchange rate volatility. Our findings support their conclusions as they actually suggest that if anything more algorithmic trading[12] is associated with lower volatility. Even though our results show relatively small volatility effects we are still suggesting the same inferences for HFT and market quality.

Overall, we suggest that the main reason behind the decreasing volatility is the efficient market-making provided by HFT. There is a lot of such market-making in the Swedish market as the major HFT players are all employing this more or less as their base activity (Julander, 2012 and Sandhagen, 2012). We have argued that there is a strong case to be made that such activities is beneficial for volatility. However, given our measure of HFT, the Large Cap proxy, we cannot establish the exact magnitude of the effect on volatility of increasing HFT, other than that in our sample an increased HFT-share has decreased volatility by 0.0088 and 0.0179 percentage units for the trading range and hourly standard deviation, respectively. While the magnitude of these effects, compared to the statistical properties of the measures, is not extraordinary we argue that

[12] While the authors perform an empirical analysis of algorithmic trading on exchange rates we argue that the results are valid for HFT. This is because their definition of algorithmic trading includes a high-frequency consideration and their tests are done at the second frequency.

these results support our first hypothesis and is a valid argument opposing the general negative opinion of these autonomous trading methods.

B. Liquidity

1. Method robustness

The effect of HFT on liquidity is examined using the same method as for the volatility, i.e. the DD framework. This means that the procedure to determine the underlying assumptions is the same. As the event we are using is also the same we can determine that the NAA is fulfilled by the previous arguments, i.e. the construction of the event. However, for the CTA we need to scrutinize the trends of the liquidity variables for the treatment and control group.

As is the case for the trends of the volatility measures Figures X and XI show rather similar trends for the liquidity measures for both Large and Mid Cap. The average quoted bid-ask spread, plotted in weekly averages, for Large and Mid Cap follow the same, relatively stable, trend over time. The trends for the turnover measure are less stable over time, but still show similar characteristics. Even though the magnitude of the changes is not the same they follow similar patterns to a large extent. Overall, we therefore argue that the crucial CTA is confirmed for both measures and the DD approach can be used for the analysis. Finally, the fixed effects in the regression model were examined also for the liquidity measures. The results remain the same while varying the model specification in terms of different fixed effects; this means that the results are robust to the model specification.

2. Liquidity effects

The hypothesis we stipulated regarding liquidity proposed that market liquidity has improved as an effect of HFT. The opinion that HFT has increased trading volumes as well as a relatively general consensus within the academic world that bid-ask spreads has tightening indicates a liquidity improvement and also a support of our hypothesis. The price discovery process mentioned in the context of volatility might also have an effect on liquidity as increased execution speed means that prices can more quickly revert back to the equilibrium price, i.e. the market flexibility is improved.

The definition of market liquidity is somewhat problematic due to the wide range of measures available. However, inspired by Aldridge (2010) we defined the first measure as the quoted bid-ask spread; it measures the width of the bid-ask spread relative to the quote midpoint. It has been argued that this is one of the best estimators of liquidity as it indicates the cost of an instantaneous reversal of the trading position (Aldridge, 2010); a lower spread therefore indicates better liquidity. The spread could also be argued to measure the level of adverse selection in the market; a narrower spread reduces the effect of asymmetric information. The second liquidity measure was the stock turnover, since many HFT actors commonly use the trading volume as a proxy for liquidity. While the bid-ask spread measures the cost of an instantaneous reversal of a trading position the stock turnover could be argued to measure the ability to enter or exit that position whenever necessary.

Proceeding as in the analysis of volatility we begin by comparing the development of the two measures over the event window. The difference between the weekly average of the measures for Mid and Large Cap, respectively, are plotted in the Figures III and IV below. In the first figure we clearly observe an increase in the difference for the quoted spread measure after the event. However, for the turnover measure we cannot identify any changes in the difference in relation to the event. The INET implementation therefore seems to have affected only the first measure and not the second. We further notice an interesting increase in the difference in the quoted spread measure around the event date; the quoted spread for all stock increased in the wake of the event with Mid Cap stocks increasing the most. We believe that this was a short-term reaction due to uncertainties amongst members of the new system. This is also the opinion of NASDAQ OMX, a somewhat lower trading activity than normal, mainly from algorithmic traders, in the first few days after the implementation indicates that market participants wanted to test the new technology and then gradually increase the trading (Pinter, 2012). These actions of uncertainty had an impact on the spread for stocks on both Large and Mid Cap.

The results for the quoted spread from the DD regression are presented in Table IV in the appendix; they show a statistically significant (at the 1% level) decrease in the bid-ask spread for Large relative to Mid Cap after the implementation of INET. The implication is that the presence of HFT on the market reduces the bid-ask spread, indicating that liquidity is improved by HFT. The result is not only statistically significant, but has also an economic significance since the coefficient on the interaction variable is -0.34. This means the spread for Large Cap has decreased 0.34 percentage units as an effect of HFT, even when controlling for the general trend after the event. The effect of being a treated company after the event, i.e. the effect of HFT, is a reduction in the spread of 0.34 percentage units relative to the expected trend. When contrasting this to the average quoted spread of Large Cap before the event of 0.58% and the standard deviation of 1.07%, we immediately realize the economic importance of the result; the effect is about 30% of this pre-event standard deviation. While the total average spread has increased after the event (the coefficient on the *post* variable is 1.21%), which to a large extent is an effect of the post-INET spread disturbance, the effect of HFT is a reduction in the Large Cap spread by as much as 0.34 percentage units (relative to what would have been the level without HFT).

Figure III: Difference of Weekly Average Quoted Spreads

This figure presents the difference (in percentage points) between the weekly average quoted spread for Mid and Large Cap, i.e. $\overline{QS}_{MID} - \overline{QS}_{LARGE}$, including the 95% confidence interval, over the period August 10th, 2009 to August 9th, 2010. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

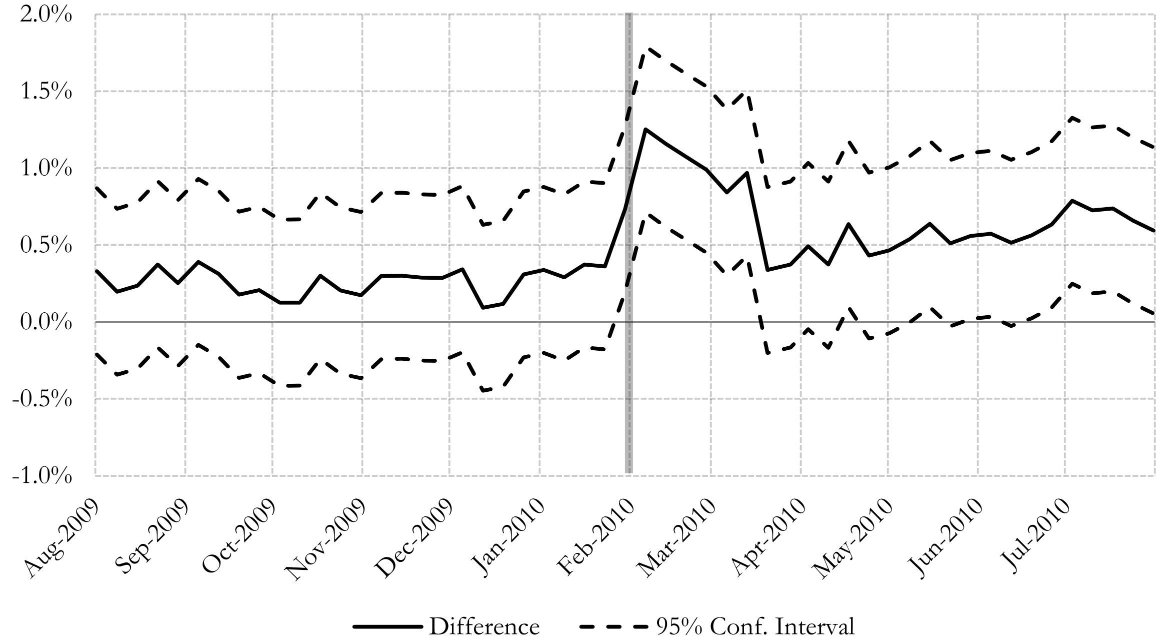

Data source: Datastream

Figure IV: Difference of Weekly Average Share Turnovers

This figure presents the difference (in percentage points) between the weekly average share turnover for Large and Mid Cap, i.e. $\overline{ST}_{LARGE} - \overline{ST}_{MID}$, including the 95% confidence interval, over the period August 10th, 2009 to August 9th, 2010. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

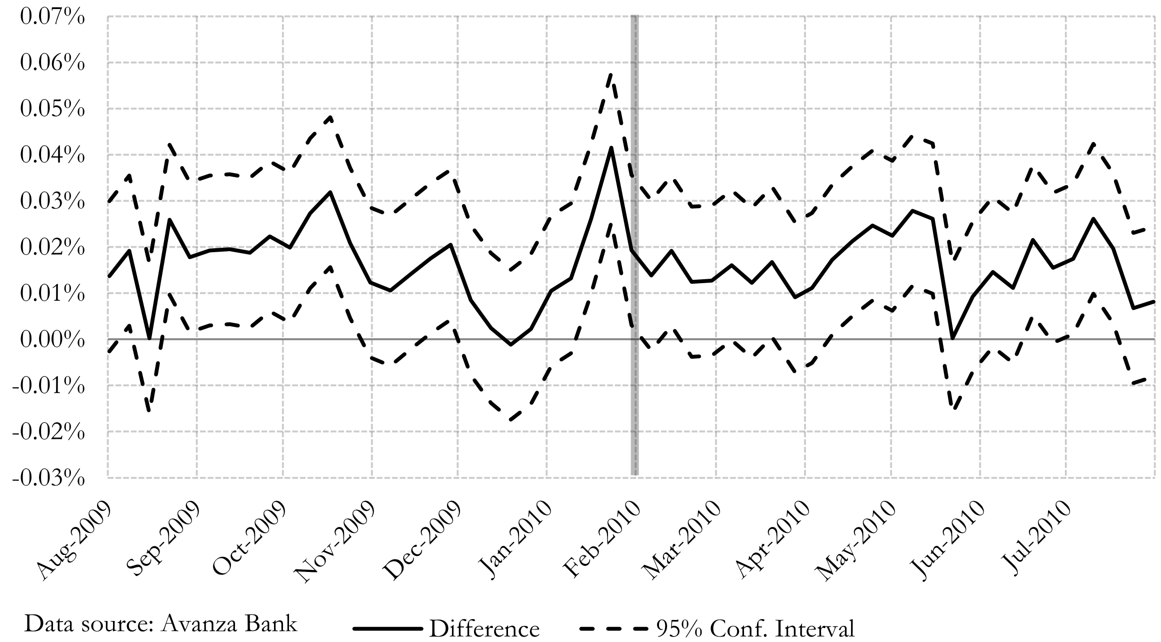

Data source: Avanza Bank AB and Datastream

Since the bid-ask spread is generally considered an appropriate measure of liquidity, and a lower spread indicates a better liquidity, this result supports our hypothesis regarding the beneficial effects of HFT on liquidity.

The results for the stock turnover measure are presented in Table IV and they to some extent contradict the results above. We do not find a significant effect of HFT on stock turnover as the coefficient on the interaction variable is not significant on any acceptable significance levels (the p-value is 78%). While this means that HFT has not improved this measure of liquidity, it also implies that it has not reduced it. Nonetheless, it does not support our hypothesis regarding improved liquidity of HFT. Since the effect of this measure is non-significant the economical interpretations become somewhat redundant; the coefficient is somewhat negative, but close to zero with a large standard error implying absence of effects. The results are still interesting, and somewhat surprising, as there is a general opinion that the trading performed by HFT-algorithms increases trading volume; our results show indications that the stock turnover (in Sweden) is in fact not affected by HFT.

The results of the liquidity analysis are somewhat ambiguous as we have shown significant improvements in the first measure and a non-existent effect on the second. Hence, we have one test that supports the hypothesis and one that does not. While the absence of significance for the turnover measure does not contradict the hypothesis per se we cannot prove its support. Harris (2003) discusses the widespread confusion around the liquidity definitions and this might influence our results, since the two measures we are using reflect dimensions of liquidity that are very different. However, based on the results we argue that the effect of HFT on liquidity is either positive or has no effect; we find no indications of a negative liquidity impact by HFT. It is further interesting to see that HFT affects one dimension of liquidity but not the other. Nonetheless, this means that we have weak support for our hypothesis. Hasbrouck and Saar (2010) show similar results, including improvements in spreads, in their study on order-level NASDAQ data. Even though Castura et al. (2010) fail to prove a causal relationship to HFT they use an extensive tick-size dataset to show tighter spreads as well as greater liquidity on the inside. Inside liquidity, i.e. the volumes available at the best bid and ask price in every period, is also investigated by Brogaard (2010), although he show that even though HFT generally provides the best bid and ask prices, the depth is not the same as for non-HFT. According to Sandhagen (2012) this is also the case in Sweden; HFT-algorithms continuously post bid and ask orders on the inside of the spread, thereby reducing the spread, but they do it in smaller volumes than what a regular trader would use. However, if the algorithms' orders are filled they might instantly provide new liquidity. This means that even though the visible liquidity might be lower, the actual available liquidity might be a lot higher.

During our sub-sample period, where we perform the DD test, the tick size for Large Cap companies was reduced by NASDAQ OMX. This was done in two steps where tick sizes for companies in OMXS30 (the 30 largest companies on the exchange) was reduced on October 26th, 2009 and the rest of the Large Cap on June 7th, 2010. There are a number of empirical studies showing that a reduced tick size might reduce bid-ask spreads (see for example Harris (1994) and Goldstein and Kavajecz (2000)). This implies that the tick size reductions might have an impact on our DD results for the bid-ask spread, even though one change is before and one after the DD event. To control for this we have performed an additional DD test for the quoted spread measure on another sub-sample excluding the periods before and after these changes. Thereby we only include a time period where tick-sizes are constant for all stocks and this would allow for an estimation of the effect of HFT. However, this extended analysis shows no significant impact on the previously obtained results and the effect of HFT on the bid-ask spread remains on the same magnitude and significance.

Even though our results are not uniformly showing a liquidity improvement we still argue that they indicate positive effects on liquidity by HFT. The quoted spread measure, which is considered the better measure of liquidity (Aldridge, 2010), is improved and the turnover measure is simply unaffected. We believe that the reasons for these results are mainly those identified by Brogaard (2010), i.e. that HFT-algorithms commonly provide the best bid and ask prices at the inside of the spread; Sandhagen (2012) is of the same opinion. This would in turn lead to a narrower spread when HFT is present. Additionally, Jovanovic and Menkveld (2011) showed, although using a theoretical model, that the presence of HFT can reduce adverse selection and thereby the bid-ask spread. Riordan and Storkenmaier (2011) also show reduced adverse selection as a result of reduced latency. Hence, our results of a narrower bid-ask spread due to HFT is supported by several research papers. However, the absence of effects on the stock turnover is somewhat puzzling; HFT-algorithms generally trade large volumes and an increased share of HFT should therefore increase the trading volumes. Among others Kirilenko et al. (2011) argue that HFT adds to the trading volumes, but continues the argument by stating that this does not always have to imply increased liquidity as such algorithms compete for it when they rebalance their portfolios. We do not observe these effects in our data as our results show no effects of HFT on stock turnover and thereby trading volumes. The reasons for this might be that the expansion of HFT has had an effect on the trading by other participants, such as day traders. Their strategies might not be profitable these days as a result of HFT-algorithms "stealing" their opportunities. Hence, such traders may be pushed out of the market and their trading might be taken over by HFT leaving the total turnover unchanged. Further, some larger institutional investors are of the opinion that their trades are abused by HFT since they move the market against them (Demos and Grant, 2011). This has had the effect that they move their trades off the market to dark pools of liquidity, i.e. over-the-counter trading volume unavailable to the public, where HFT is not present. This means that even if HFT increases the trading volumes the flight to these dark pools might have a counteracting effect.

To conclude, the analysis shows ambiguous support for the hypothesis that HFT improves market liquidity. HFT seems to reduce the bid-ask spread but has no effect on the stock turnover, i.e. trading volumes. We still argue that HFT has important and beneficial implications for liquidity. However, an important measure for many investors is the visible liquidity, both at the inside and the full depth. According to Sandhagen (2012) and Julander (2012) there is a general opinion that these volumes have decreased compared to historical levels; this might be due to both HFT and additional price increments by lower tick sizes. We argue that the actual liquidity in the market has in fact not decreased; HFT-algorithms instantaneously post new orders if their previous ones are filled, they simply keep their liquidity off the radar. These theories are supported by our results of an unaffected stock turnover; this further means that investors still can trade their desired volumes whenever they want. While this is mostly true one must recognize the potential dangers with such "off radar" liquidity and market-making by HFT. They can withdraw from the market in time of turmoil, and thereby leaving the ordinary investors without any possibility to exit (or enter) desired positions. The Flash Crash of 2010 is a perfect example of such an event, where markets became illiquid due to the reactions of HFT market-makers (Easley et al., 2010).

C. Market Efficiency

To analyze the effect of HFT on market efficiency on the Stockholm stock exchange several variance ratio tests was performed. Considering Figure V, an upward trend for the average of the variance ratios of the Large Cap listed firms is evident during the examined period. This pattern is confirmed when considering the average variance ratios calculated with other aggregation parameters, i.e. higher orders of serial autocorrelation (see Figures XII to XIV in appendix). What can be noticed is that until the end of 2008, Large Cap stocks have variance ratios lower than one on average, which implies that the price process is mean reverting, i.e. returns above the average return tend to be followed by returns that are below the average. The same holds true for stocks on the Mid Cap, when higher orders of autocorrelation than the first are considered. After 2008 the average variance ratio implies a basically efficient price process for the Large Cap stocks during about one and a half year. However, after the implementation of INET (first quarter 2010), the average variance ratio implies a price process inefficiency again; going forward from this date the inefficiency takes the form of a momentum effect, i.e. returns above average tend to be followed by returns above average and vice versa. The pattern is the same when considering higher levels of serial autocorrelation, but the effect appears to become prevalent somewhat later (during the second and third quarter of 2010, respectively).

Figure V: Average Variance Ratios, 2/1 minutes

This figure presents the average variance ratio for stocks listed on the Large and Mid Cap, respectively, over each quarter of the period 2006Q1 to 2012Q1. The variance ratios are computed by using 2 minutes over 1 minute VWAP prices for the Large Cap and Mid Cap stocks, respectively. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

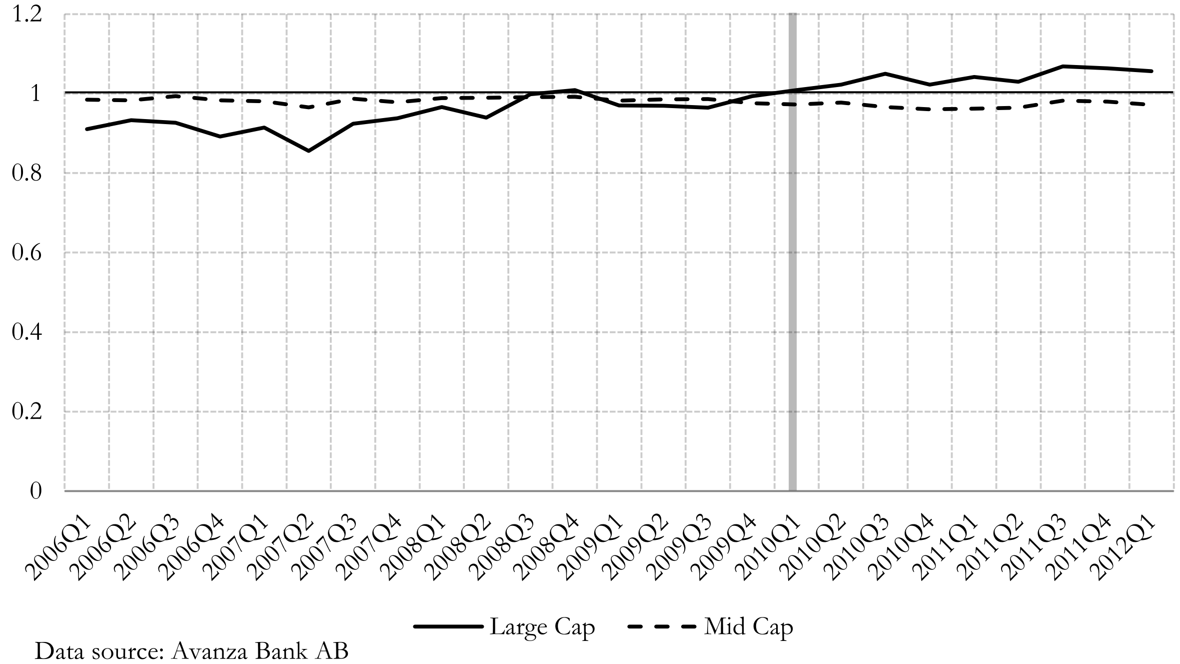

Data source: Avanza Bank AB

Something else that is striking is the apparent efficiency in the price process of the Mid Cap listed stocks when one considers only the first order autocorrelation, as presented in Figure I. Clearly, we would expect the price process to be more efficient for Large Cap listed stocks since these tend to be traded more often, covered by more analysts and also less susceptible to insider trading. However, when considering the higher orders of autocorrelation (see Figures XII to XIV) the average variance ratios of Mid Cap listed stocks seem to suggest inefficiencies in the price process for these stocks as well. Although the extent of inefficiency is less than that of the Large Cap listed stocks, at least before 2010, the Mid Cap listed stocks have average variance ratios consistent with a mean reverting price process. While the amount of mean reversion is less for Mid Cap stocks in earlier periods it seems to be higher than the corresponding momentum effect of the Large Cap stocks in the period after the first quarter of 2010. Put simply, the variance ratios of Large Cap listed stocks deviates less from one than those of the Mid Cap listed stocks going forward from 2010. However, we stress again that this is only true when looking at higher orders of autocorrelation, i.e. using larger aggregation parameters in the variance ratio tests. As mentioned earlier, the first order autocorrelation, as suggested by the average variance ratios for 2 minutes over 1 minute, clearly suggests that Mid Cap listed stocks follows a more efficient price process.

To be able to deduce anything regarding the impact of HFT on the price process in the equity market we need to compare the trends that we see for the average variance ratios of the Mid- and Large Cap stocks, respectively. If HFT has had an impact on the efficiency in the market, as measured by the random walk properties of the price process, we should be able to see a difference in the trends of the average variance ratios for Mid- and Large Cap listed stocks surrounding the implementation of INET. However, since the average variance ratios for the Large Cap listed stocks have been increasing ever since 2006, whereas a fairly stable ratio is seen for the Mid Cap listed stocks, it would be difficult to argue that this has anything to do with HFT. What we see is just two trends that are constant, relative to each other, throughout the examined time period, which would suggest that market efficiency is unaffected by HFT.

It is interesting to see the type of inefficiency prevalent in the Large Cap listed stocks price process though, which has gone from being mean reversal to a momentum effect, especially considering the fact that the moment effect is found shortly after or in connection with the introduction of INET. Similar results were found by Zhang (2010), who showed that HFT is positively correlated with price momentum. He argues that this phenomenon could be explained by the occurrence of large orders by fundamental investors that together with the presence of HFT might attract other investors, such as momentum traders. One could argue that the introduction of algorithmic trading has created an illusion of positive autocorrelation; since it means making small trades throughout the day, to not induce large price changes in an instrument. Thereby it slowly drives a price upwards or downwards. However, this argument fails to take into account that algorithmic trading is by no means a recent phenomenon. Hence, it would be hard to argue why we would see a move towards a positive autocorrelation in recent years due to algorithmic trading. HFT being truly feasible only in recent years, on the other hand, could be a potential explanation for this recent increase in autocorrelation; although it does not explain the fact that we see an increasing trend already in 2007.

Our results are to a large extent in line with those of Castura et al. (2010); they also find an increasing trend in the average variance ratios for the stocks in the Russell 1000 and Russell 2000 lists over the period 2006 to 2010. On some points the results differ though, for example they observe unanimous results with respect to the fact that the smaller firms average variance ratios are lower, i.e. further from unity, than the larger ones. We find results that are more mixed, e.g. for the 2 minutes over 1 minute average variance ratios the Mid Cap listed firms are very close to one over the entire period. Also, we observe variance ratios for the Mid Cap listed stocks that seemingly follows a more efficient price process in the earlier periods, while Castura et al. (2010) observe a clearly lower efficiency, with average variance ratios of 0.6-0.8, in the Russell 2000 firms over their entire examined period. However, taking into account the fact that their sample only covers a period until the second quarter of 2010, the results are very similar. In line with our results they observe variance ratios (for the larger firms) very close to one by the first half of 2010, whereas the efficiency is lower in the smaller stocks in this period. The differences are not that surprising though, besides examining the phenomenon on an entirely different market they also have access to data with second-by-second observations. Their data is of a higher frequency than ours, which allows them to make variance ratio tests on even higher frequencies than those performed in this paper. However, the comparisons made in this paper are with their variance ratio of 10 minutes of over 10 seconds, which could be said to be somewhat similar to our test of 10 minutes over 1 minute. And for these tests the results are, as mentioned, to a large extent consistent.

Until now we have mainly given an interpretation of the average variance ratios. The benefit of this approach is that it tells us something about the general trend. However, without considering the statistical significance of these results very little can be said about their reliability. When performing a variance ratio test, what is being tested is whether the price process follows a random walk or not. Hence, any rejection of our hypothesis is a rejection of the efficiency of that particular stock's price process. Further, rejections on a wider scale would thus be a rejection of the efficiency of that market. Hence, we extend our variance ratio test to see if the market today, when HFT is present, is less efficient compared to the period before HFT. As explained in the method section the Chow-Denning test-statistic is used for this analysis, since several aggregation parameters are considered, instead of a standard variance ratio test. Since the variance ratios are tested on a five percent significance level we would expect our sample to contain about five percent of false positives, i.e. rejections of the market hypothesis when it actually holds true.

If we consider Figure VI, where a variance ratio of 2 minutes over 1 minute is tested, something that stands out is the fact that the results overall suggest a fairly inefficient market (note: the lower the fraction the better the market). Furthermore, just as the average variance ratios suggested, Mid Cap listed stocks have a price process more efficient than that of the Large Cap listed ones throughout the entire period. This is contradictory both to what Castura et al.

(2010) finds for the American market and to what we would expect, based on the fact that Large Cap is under closer scrutiny than Mid Cap. Instead we see that more than 50% of the stocks are inefficient in most of the quarters and that as many as 60-70% are inefficient in the more recent years. This cannot be immediately tied to HFT though, since the trend is even worse for the Mid Cap listed stocks. Over the past two years the fraction of statistically inefficient Mid Cap stocks have come very close to the high levels of the Large Cap listed firms. Hence, the inefficiency in the market seems to be due to other reasons than HFT.

Figure VI: Fractions of Significantly Inefficient Stocks

This figure presents the fraction of stocks that have a 2 minutes over 1 minute Variance Ratio that is significantly different from one, i.e. considered inefficient, on a five percent significance level. This is done for stock listed on Mid and Large Cap, respectively, for each quarter over the period 2006Q1 to 2012Q1. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

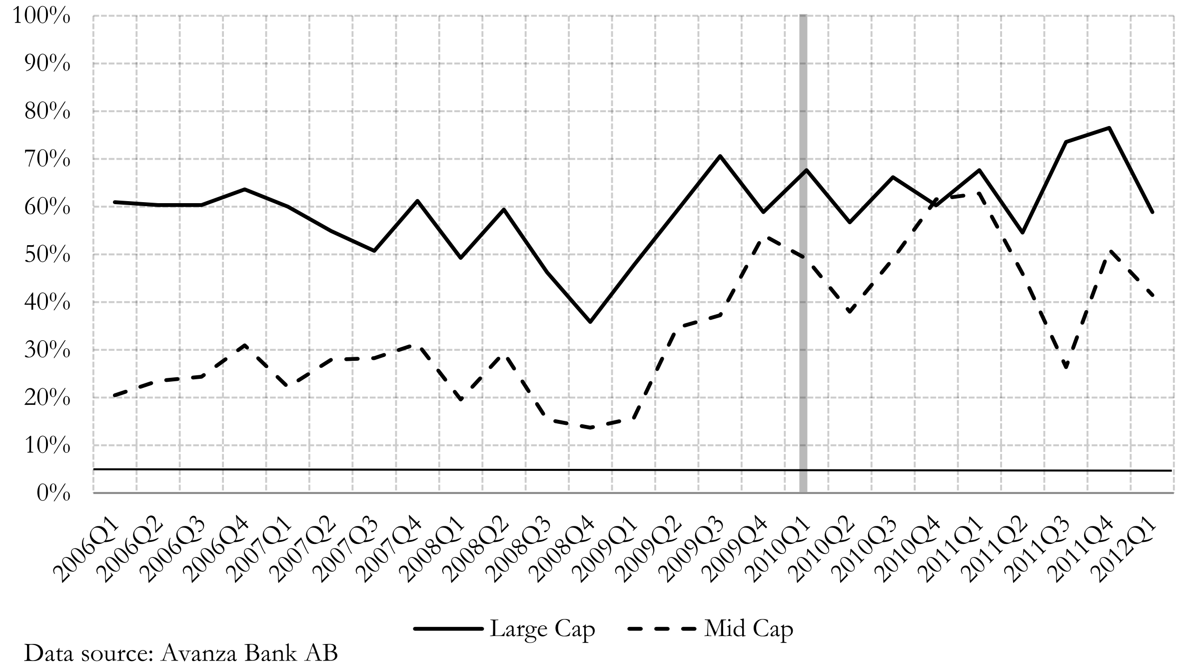

Data source: Avanza Bank AB

When higher orders of autocorrelations are considered the results suggest a relative deterioration for the Mid Cap listed stocks. Consider Figures XV to XVII in appendix, for all the higher order of autocorrelation the trend is that the fraction of Large Cap stocks following an inefficient price process is fairly high and stable. Simultaneously, the Mid Cap stocks have seen a fairly large increase in the fraction of inefficient stocks over the past two or three years. One could argue that the Large Cap would have seen an equal rise in inefficiency if not for the presence HFT. For the variance ratios of 4 minutes over 2 minutes, Large Cap listed stocks have actually seen a substantial decrease in the number of inefficient stocks in the past years, down to about 30%.

To conclude, the disparity in the results is somewhat large, but this comes as no surprise since this is exactly what Castura et al. (2010) finds for their higher sampling intervals, where they do not suggest any clear unanimous effect of HFT on market efficiency. Another potential reason for the absence of effects could be the already high level of inefficiency in the Swedish financial market. It might be that in a small market, where some market imperfections are likely to be present, the effect of increased HFT alone is too small to be noticeable. However, in a more developed market, like the one in the U.S., increased HFT might have an effect on the margin. Overall, our results and analysis suggest that HFT has not had a significant impact on market efficiency. Hence, we reject our hypothesis that HFT has improved market efficiency in Sweden.

D. Analysis Issues

The research design, including both variance ratio tests and the DD framework, was chosen to answer the overall research question, using the available dataset, in a comprehensive and exhaustive manner. However, a couple of caveats are still in order. Regarding the variance ratio tests we might have obtained other results if the frequency of the data was even higher; the current frequency might fail to capture the microstructural efficiency effects of HFT on market efficiency. Further, the fact that we use actual trade prices and not bid and ask quotes could have the effect of making prices look more stable than they actually are. If quoted prices change but there are no trades our dataset fails to capture the effects. While these issues are true also for the DD tests there might be additional concerns in relation to the event we are using, i.e. the INET implementation. Even though the case we have for a significant effect on the share of HFT in the market is strong, there is a risk that the effect is too small to be captured. In such cases the effect measured by the DD approach is related to something else than HFT. However, we believe that this is a minor risk, considering our arguments presented in the methodology section. Finally, even though the results regarding liquidity show either improvements or absence of effects it might be that these measures are not the most important for investors and other market actors. The difficulties in identifying and obtaining comprehensive liquidity measures might therefore reduce the implications and importance of these results. Nevertheless, we argue that our results and their implications are valid in general and have important consequences for the view on HFT and its effects on market quality in Sweden.

V. Conclusion

We set out to answer the question: *What is the effect of high-frequency trading on the quality of the Swedish financial market?* To be able do this we identify three parameters that ought to capture different dimensions of market quality: volatility, liquidity and market efficiency. By using a DD framework, on a structural change of the Stockholm stock exchange, we test different measures of volatility and liquidity with somewhat mixed results. Regarding volatility, both measures point towards a significant improvement due to HFT, with a magnitude of about 15-30% of the ex-ante volatility standard deviation depending on which measure we use. This leads us to conclude that HFT is in fact beneficial to the Swedish market in terms of volatility. We present several potential explanations for this, the most plausible one being the market-making activities by HFT firms. The results for our liquidity measures are not as conclusive. The quoted spread shows a significant improvement due to HFT, with a decrease of about 0.34 percentage units, whereas no such improvement is seen for the stock turnover measure. We suggest that the improved quoted spread is mainly due to the fact that HFT-algorithms quote bid and ask prices at the inside of the spread, as suggested by several other papers in the field. Even though the results for market liquidity are not as consistent as for volatility we argue that the overall effect of HFT on liquidity is at least not harmful. Further, by examining high-frequency data we show that, unlike results found in studies performed on U.S. data, there is no real support to conclude that HFT has improved the efficiency in the Swedish stock market. This is not to say that the effect has been detrimental since the market is about as efficient now as before the introduction of INET. However, whereas Large Cap listed stocks have a tendency for mean reversal on average in the earlier parts of our sample, they exhibits a momentum effect on average after the implementation of INET. This could possibly be explained by the larger share of HFT in Sweden today, since the same development is not seen for Mid Cap listed securities.

Overall, our results in general show positive effects of HFT on market quality in Sweden; three out of the five measures point to an improved quality and two measures show no significant effects in either direction. These results have important implications for both market participants and investors; the presence of HFT seems to both reduce volatility and increase liquidity. Further, the general fear and aversion towards HFT in Sweden needs to be reconsidered in the light of our results. Nonetheless, we recognize that there are risks associated with extraordinary events, such as the Flash Crash. However, considering the fact that the Swedish stock market already has volatility guards in place to protect investors of such occurrences this risk ought to be considered very small. Furthermore, some have argued that price manipulation is a reason alone to ban HFT from Swedish financial venues; this argument fails to realize all the benefits associated with HFT in a cross section. Additionally, price manipulation is not a phenomenon unique to HFT; unscrupulous investors have used such methods for a long time. While this is not an excuse, it shows the irrationality of punishing a certain type of investors based on the actions of a few. Hence, the question rather becomes what exchanges, such as OMX Stockholm, can do to become more attractive for HFT, e.g. lowering entry barriers by providing information for back-testing of algorithms. However, they still need to make sure that they have control systems in place that are sufficient to detect and penalize any price manipulators, may it be due to traditional or high-frequency trading. Finally, even though we are of the opinion that HFT is not harmful for financial markets and in agreement with Finansinspektionen think that the risks with HFT are acceptable we encourage the development of regulations within the area, e.g. MiFID II. However, while they should address potential risks of HFT, such as flash crashes and "off radar" liquidity, they need to recognize the beneficial effects of having such trading in the market.

A. Future research

Our results are obtained using a proxy for the share of HFT in the market, which restricts the methods that we are able to use. It further limits us to mainly concentrate our analysis around the structural change of the Stockholm stock exchange by the implementation of INET. A natural extension of our study is therefore to obtain an actual measure of the HFT share; this would allow for closer modeling of the market quality parameters in relation to HFT. While this is possible using the messaging system in the matching engine it places enormous demands on both computational as well as programming capacity. Further, obtaining a dataset with even higher frequency of the data might allow for an even deeper analysis of the microstructural effects of HFT on the parameters. It would also be interesting to extend the timeframe of the variance ratio analysis performed by Castura et al. (2010) in the U.S. market. This would allow for a further comparison of the development in the U.S. market, to determine whether the trends are the same as in Sweden. Finally, even though our parameters are chosen to exhaustively reflect market quality the analysis might be extended by including other dimensions of market quality, especially regarding liquidity. For example, one could use the exchange members' order books to obtain a measure of available liquidity in the market.

VI. References

Abadie, Alberto, 2005, Semiparametric Difference-in-Differences Estimators, The Review of Economic Studies, Vol. 72, pp. 1-19.

Aldridge, Irene, 2010, High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems (John Wiley & Sons).

Andersen, Torben G., Bollerslev, Tim, Diebold, Francis X., 2002, Parametric and Nonparametric Volatility Measurement, NBER Working Paper No. 279.

Angrist, Joshua D., and Pischke, Jörn-Steffen, 2009, Mostly Harmless Econometrics: an Empiricist's Companion (Princeton University Press).

Bach, Laurent, 2012, Lecture: Panel Data, Empirical Methods in Corporate Finance, Stockholm School of Economics.

Bessembinder, Hendrik, 1998, Trading Costs and Return Volatility: Evidence From Exchange Listings, Unpublished Manuscript, Arizona State University.

Brogaard, Jonathan A., 2010, High Frequency Trading and its Impact on Market Quality, Working Paper, Northwestern University, Kellogg School of Management.

Bushee, Brian J., and Noe, Christopher F., 2000, Corporate Disclosure Practices, Institutional Investors, and Stock Return Volatility, Journal of Accounting Research, Vol. 38, pp. 171-202.

Castura, Jeff, Litzenberger, Robert, Gorelick, Richard, and Dwivedi, Yogesh, 2010, Market Efficiency and Microstructure Evolution in U.S. Equity Markets: A High-Frequency Perspective, RGM Advisors, LLC.

Chaboud, Alain, Chiquoine, Benjamin, Hjalmarsson, Erik and Vega, Clara, 2009, Rise of the machines: Algorithmic trading in the foreign exchange market, Working Paper, Federal Reserve Board.

Chlistalla, Michael, 2011, High Frequency Trading: Better than its reputation?, Deutsche Bank Research, Research Briefing.

Chow, Viktor K., and Denning, Karen C., 1993, A Simple Multiple Variance Ratio Test, Journal of Econometrics, Vol. 58, pp. 385-401.

Cvitanic, Jaksa, and Kirilenko, Andrei, 2010, High Frequency Traders and Asset Prices, Working Paper, Caltech.

Demos, Telis, and Grant, Jeremy, 2011, Institutional Investors Air HFT Concerns, Financial Times, September 12th.

Easley, David, de Prado, Marcos M. López, and O'Hara, Maureen, 2010, The Microstructure of the 'Flash Crash': Flow Toxicity, Liquidity Crashes and the Probability of Informed Trading, Working Paper, Cornell University.

European Commission, 2012, New Rules for More Efficient, Resilient and Transparent Financial Markets in Europe, Press Release, Reference: IP/11/1219.

Fama, Eugene F., 1965, The Behavior of Stock-Market Prices, The Journal of Business, Vol. 38, No. 1, pp. 34-105.

Fama, Eugene F., 1970, Efficient Capital Markets: A Review of Theory and Empirical Work, The Journal of Finance, Vol. 25, No. 2, pp. 383-417.

Finansinspektionen, 2012, Investigation Into High Frequency and Algorithmic Trading, Finansinspektionen Report, Ref. 11-10857.

Goldstein, Michael A., Kavajecz, Kenneth A., 2000, Eights, Sixteenths, and Market Depth: Changes in Tick Size and Liquidity Provision on NYSE, Journal of Financial Economics, Vol. 56, No. 1, pp. 125-149.

Harris, Larry, 2003, Trading and Exchanges: Market Microstructure for Practitioners (Oxford University Press).

Harris, Lawrence E., 1994, Minimum Price Variation, Discrete Bid-Ask Spreads, and Quotation Sizes, The Review of Financial Studies, Vol. 7, No. 1, pp. 149-178.

Hasbrouck, Joel, and Saar, Gideon, 2010, Low-Latency Trading, Working Paper, Cornell University.

Hendershott, Terrence and Moulton, Pamela C., 2009, Speed and Stock Market Quality: The NYSE's Hybrid, Haas School of Business, University of California, Berkeley and Fordham Graduate School of Business.

Hendershott, Terrence, Jones, Charles M., and Menkveld, Albert J., 2011, Does Algorithmic Trading Improve Liquidity?, The Journal of Finance Vol. 66, No 1.

Jarrow, Robert, and Protter, Philip, 2011, A Dysfunctional Role of High-Frequency Trading in Electronic Markets, Johnson School Research Paper Series #08-2011.

Jovanovic, Boyan, and Menkveld, Albert J., 2011, Middlemen in Limit Order Markets, Working Paper, New York University.

Kale, Jayant R., and Loon, Yee Cheng, 2011, Product Market Power and Stock Market Liquidity, Journal of Financial Markets, Vol. 14, No. 2, pp. 376-410.

Kirilenko, Andrei A., Kyle, Albert S., Samadi, Mehrdad and Tuzun, Tugkan, 2011, The Flash Crash: The Impact of High Frequency Trading on an Electronic Market, Working Paper, Commodity Futures Trading Commission.

Lee, Charles M. C., Mucklow, Belinda, and Ready, Mark J., 1993, Spreads, Depths, and the Impact of Earning Information: An Intraday Analysis, The Review of Financial Studies, Vol. 6, No. 2, pp. 345-374.

LeRoy, Stephen, 1973, Risk Aversion and the Martingale Property of Stock Returns, International Economic Review, Vol. 14, pp. 436-446.

Levander, Margareta, 2009, Super-Computers Rules the Exchanges (Superdatorer Styr Börserna), Dagens Nyheter, August 30th.

Lo, Andrew W., and MacKinlay, A. Craig., 1988, Stock Market Prices Do Not Follow Random Walks: Evidence From a Simple Specification Test, The Review of Financial Studies, Vol. 1, No. 1, pp. 41-66.

Malmqvist, Mattias, 2009, The Pressure is Rising for NASDAQ OMX (Nu Ökar Trycket på NASDAQ OMX), Computer Sweden, June 8th.

Munkhammar, Viktor, 2005, Black Boxes Brings Down the SEK (Svarta Lådor Sänker Kronan), Dagens Industri, December 29th.

NASDAQ OMX, 2010a, Monthly Report - Total Equity Trading January 2010, News Release February 3rd.

NASDAQ OMX, 2010b, NASDAQ OMX Launches INET Trading System Across Its Seven Markets in the Nordics and Baltics, Press Release February 8th.

NASDAQ OMX, 2012, Monthly Report - Total Equity Trading January 2012, News Release February 9th.

Neurath, Carolina, 2011, They are Defending the Computers (De Försvarar Datorerna), Svenska Dagbladet, August 24th.

Ozenbas, Deniz, Schwartz, Robert A., and Wood, Robert A., 2002, Volatility in US and European Equity Markets: An Assessment of Market Quality, International Finance, Vol. 5, No. 3, pp. 437-461.

Psomadelis, Will, and Baden Powell, Stuart, 2011, High Frequency Trading - Credible Research Tells The Real Story, Schroders Investment Management, Special Report.

Rex, Martin, 2012, Repeated Negligence Expensive for Pareto Öhman (Upprepat Slarv Dyrt för Pareto Öhman), Dagens Industri, January 17th.

Riordan, Ryan, and Storkenmaier, Andreas, 2011, Latency, Liquidity and Price Discovery, Working Paper, School of Economics and Business Engineering, Karlsruhe Institute of Technology, Germany.

Samuelson, Paul A., 1965, Proof That Properly Anticipated Prices Fluctuate Randomly, Industrial Management Review, 6:2.

Stoll, Hans R., 1989, Inferring the Components of the Bid-Ask Spread: Theory and Empirical Tests, The Journal of Finance, Vol. 44, No. 1, pp. 115-134.

Zhang, X. Frank, 2010, High-Frequency Trading, Stock Volatility, and Price Discovery, Working Paper, Yale University.

A. Interviews and Seminars

Gaudy, Richard, Managing Director CoLo & Access Services (EMEA), NASDAQ OMX, 2012, Interviewed by Carl-Emil Lindholm and Erik Petersson, April 11th.

Hertzberg, Björn, Head of Economic & Statistical Research (Nordics), NASDAQ OMX, 2012, High-Frequency Trading at NASDAQ OMX, Seminar Presentation, January 19th.

Julander, Claes-Henrik, CEO Pan Capital, 2012, Interviewed by Carl-Emil Lindholm and Erik Petersson, April 12th.

Pinter, Stefan, Sales Director, NASDAQ OMX, 2012, Interviewed by Carl-Emil Lindholm and Erik Petersson, May 8th.

Sandhagen, Ludvig, Senior Investigator Finansinspektionen, 2012, Interviewed by Carl-Emil Lindholm and Erik Petersson, February 2nd.

VII. Appendix

A. Tables

Table III: DD Regression Results: Volatility

This table presents the results of the standard OLS difference-in-difference regressions with our two volatility measures, Trading Range and Standard Deviation, as the dependent variable. The regression accounts for serial correlation within each security by using robust standard errors. Trading Range is the average over an hour of the difference between the highest and lowest observed price within a minute. Standard deviation is the hourly standard deviation of minute stock returns. Time fixed effects are controlled for on a weekly level.

	Volatility Measure	
	Trading Range	**Standard Deviation**
Large × Post	-0.0088***	-0.0179***
	(0.0032)	(0.0033)
Large	-0.0024	-0.0570***
	(0.0016)	(0.0017)
Post	0.0026	0.0006
	(0.0038)	(0.0043)
Constant	0.0472***	0.1244***
	(0.0025)	(0.0029)
Time Fixed Effects	YES	YES
Firm Fixed Effects	YES	YES
R^2	0.0492	0.1753
N	270919	270919

Standard errors are reported in parenthesis. ***, **, * denote significance at the 1, 5 and 10 percent levels, respectively.

Data source: Avanza Bank AB

Table IV: DD Regression Results: Liquidity

This table presents the results of the standard OLS difference-in-difference regressions with our two liquidity measures, Share Turnover and Quoted Spread, as the dependent variable. The regression accounts for serial correlation within each security by using robust standard errors. Share Turnover is the sum over an hour of the amount of shares traded related to the total amount of outstanding shares. Quoted spread is the daily closing difference between the bid and ask prices over the midpoint price. Time fixed effects are controlled for on a weekly level.

	Liquidity Measure	
	Share Turnover	**Quoted Spread**
Large × Post	-0.0010	-0.3373***
	(0.0036)	(0.0493)
Large	-0.0039**	-7.3998***
	(0.0018)	(0.0270)
Post	-0.0117**	-0.1665**
	(0.0051)	(0.0699)
Constant	0.0188***	7.3483***
	(0.0031)	(0.0432)
Time Fixed Effects	YES	YES
Firm Fixed Effects	YES	YES
R^2	0.2542	0.6205
N	270919	36224

Standard errors are reported in parenthesis. ***, **, * denote significance at the 1, 5 and 10 percent levels, respectively.

Data source: Avanza Bank AB and Datastream

Table V: Variance Ratio and Chow-Denning Test Results, Large Cap

This table presents the results from all the Variance Ratio Tests and Chow-Denning tests performed on the Large Cap stocks in our dataset. The number of observations varies mainly due to the fact that some firms are listed after 2006Q1 or corporate actions making the firm ineligible for the test in that particular period.

		Average Variance Ratios				Fraction of 5% significant Chow Denning tests			
Period	**N**	**(2/1)**	**(5/1)**	**10/1**	**(4/2)**	**(2/1)**	**(5/1)**	**10/1**	**(4/2)**
2006Q1	62	0.91	0.78	0.69	0.88	61%	61%	64%	64%
2006Q2	58	0.93	0.84	0.77	0.92	60%	59%	60%	57%
2006Q3	63	0.93	0.81	0.72	0.89	60%	68%	75%	73%
2006Q4	66	0.89	0.76	0.67	0.87	64%	67%	68%	71%
2007Q1	65	0.91	0.77	0.68	0.87	60%	63%	66%	57%
2007Q2	51	0.86	0.70	0.61	0.84	55%	59%	63%	65%
2007Q3	67	0.92	0.82	0.76	0.90	51%	52%	55%	48%
2007Q4	67	0.94	0.84	0.77	0.92	61%	70%	72%	67%
2008Q1	67	0.97	0.90	0.84	0.96	49%	54%	61%	49%
2008Q2	64	0.94	0.84	0.76	0.92	59%	63%	69%	69%
2008Q3	67	1.00	0.95	0.90	0.97	46%	48%	52%	39%
2008Q4	67	1.01	0.99	0.95	0.99	36%	34%	43%	27%
2009Q1	65	0.97	0.92	0.87	0.96	48%	49%	58%	55%
2009Q2	66	0.97	0.92	0.87	0.96	59%	62%	64%	59%
2009Q3	68	0.96	0.89	0.82	0.94	71%	72%	76%	69%
2009Q4	68	0.99	0.95	0.89	0.97	59%	66%	74%	54%
2010Q1	68	1.01	0.97	0.91	0.97	68%	69%	72%	49%
2010Q2	67	1.02	1.00	0.95	0.98	57%	46%	52%	22%
2010Q3	68	1.05	1.06	1.03	1.01	66%	65%	65%	26%
2010Q4	68	1.02	1.01	0.97	0.99	60%	50%	50%	31%
2011Q1	68	1.04	1.04	1.00	1.00	68%	62%	65%	38%
2011Q2	66	1.03	1.03	1.00	0.99	55%	50%	45%	23%
2011Q3	68	1.07	1.10	1.08	1.01	74%	74%	72%	37%
2011Q4	68	1.06	1.08	1.03	1.01	76%	78%	78%	43%
2012Q1	68	1.06	1.08	1.06	1.02	59%	54%	54%	28%

Data source: Avanza Bank AB

Table VI: Variance Ratio and Chow-Denning Test Results, Mid Cap

This table presents the results from all the Variance Ratio Tests and Chow-Denning tests performed on the Mid Cap stocks in our dataset. The number of observations varies mainly due to the fact that some firms are listed after 2006Q1 or corporate actions making the firm ineligible for the test in that particular period.

		Average Variance Ratios				Fraction of 5% significant Chow Denning tests			
Period	**N**	**(2/1)**	**(5/1)**	**10/1**	**(4/2)**	**(2/1)**	**(5/1)**	**10/1**	**(4/2)**
2006Q1	39	0.98	0.93	0.87	0.96	21%	33%	44%	38%
2006Q2	34	0.98	0.95	0.90	0.97	24%	26%	29%	26%
2006Q3	41	0.99	0.97	0.93	0.98	24%	27%	29%	27%
2006Q4	42	0.98	0.94	0.87	0.96	31%	48%	67%	55%
2007Q1	45	0.98	0.92	0.86	0.96	22%	29%	42%	33%
2007Q2	43	0.97	0.90	0.82	0.96	28%	42%	49%	42%
2007Q3	46	0.99	0.94	0.89	0.97	28%	30%	39%	35%
2007Q4	48	0.98	0.91	0.83	0.95	31%	44%	52%	48%
2008Q1	51	0.99	0.96	0.91	0.98	20%	20%	33%	31%
2008Q2	51	0.99	0.94	0.88	0.97	29%	31%	39%	37%
2008Q3	52	0.99	0.96	0.93	0.99	15%	13%	21%	15%
2008Q4	51	0.99	0.96	0.91	0.98	14%	20%	27%	25%
2009Q1	51	0.98	0.94	0.89	0.97	16%	35%	43%	35%
2009Q2	52	0.99	0.95	0.89	0.97	35%	44%	46%	40%
2009Q3	51	0.99	0.94	0.87	0.97	37%	49%	61%	47%
2009Q4	50	0.98	0.91	0.83	0.95	54%	64%	78%	68%
2010Q1	53	0.97	0.90	0.82	0.95	49%	70%	81%	72%
2010Q2	50	0.98	0.92	0.85	0.95	38%	50%	60%	50%
2010Q3	53	0.97	0.90	0.82	0.94	49%	58%	75%	72%
2010Q4	52	0.96	0.87	0.77	0.93	62%	79%	87%	83%
2011Q1	51	0.96	0.88	0.79	0.93	63%	78%	86%	80%
2011Q2	50	0.96	0.89	0.81	0.94	46%	60%	70%	66%
2011Q3	53	0.98	0.93	0.87	0.96	26%	30%	42%	38%
2011Q4	51	0.98	0.93	0.87	0.96	51%	55%	57%	53%
2012Q1	53	0.97	0.91	0.84	0.95	42%	43%	49%	51%

Data source: Avanza Bank AB

B. Figures

Figure VII: Weekly Average Trading Ranges

This figure presents the weekly average trading range (in percentage points) for Large and Mid Cap over the period 10 August, 2009 to 5 February, 2010, i.e. the 6 months period before the INET implementation.

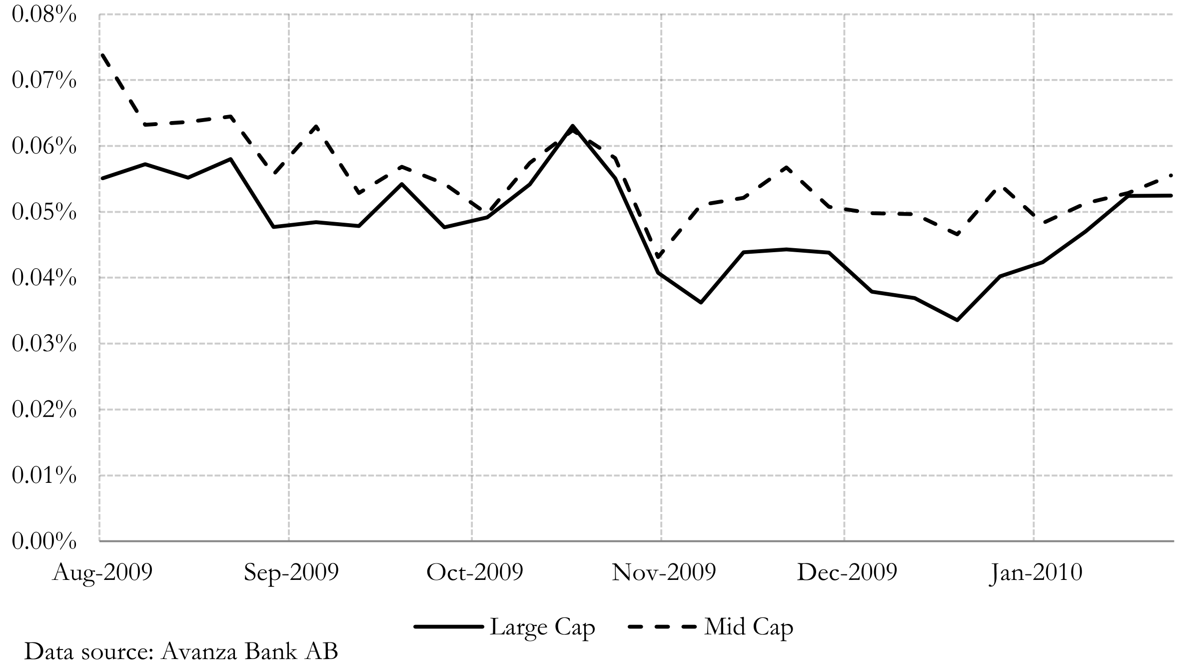

Data source: Avanza Bank AB

Figure VIII: Weekly Average Hourly Standard Deviations

This figure presents the weekly average hourly standard deviation (in percentage points) for Large and Mid Cap over the period 10 August, 2009 to 5 February, 2010, i.e. the 6 months period before the INET implementation.

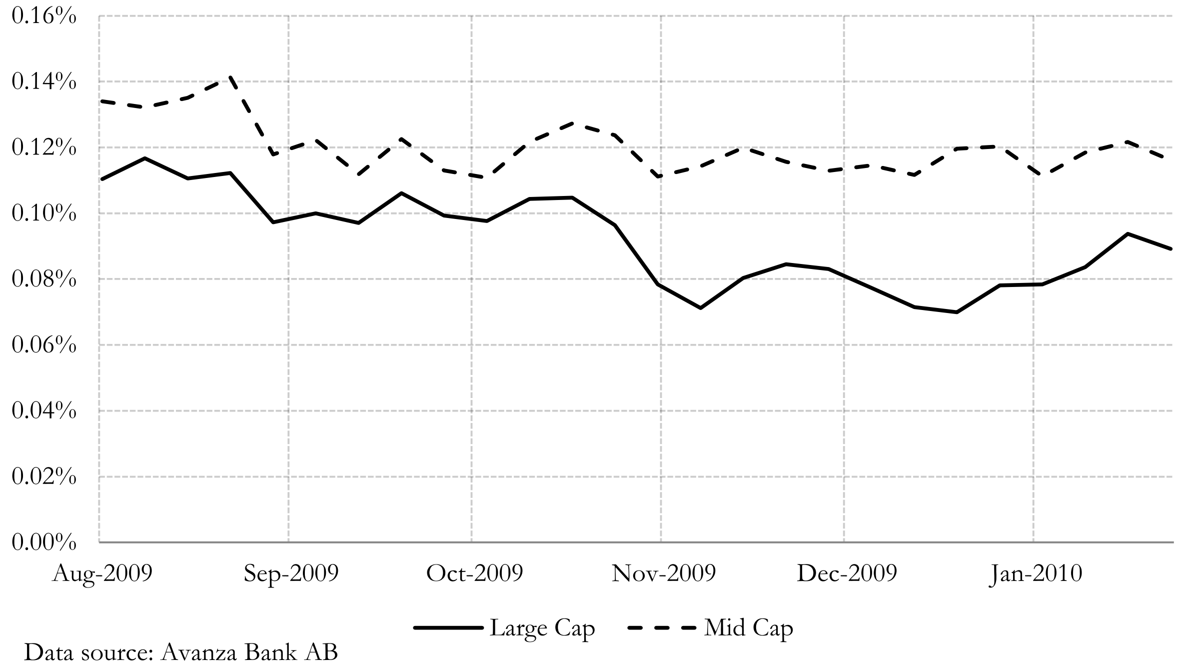

Data source: Avanza Bank AB

Figure IX: OMXS30 Return Series

This figure shows the daily realized returns of the OMXS30 index, i.e. the 30 most traded companies on the Stockholm stock exchange from January 2nd, 2002 to April 26th, 2012. The shaded area represents the event period for our DD regressions, i.e. August 10th, 2009 to August 9th, 2010.

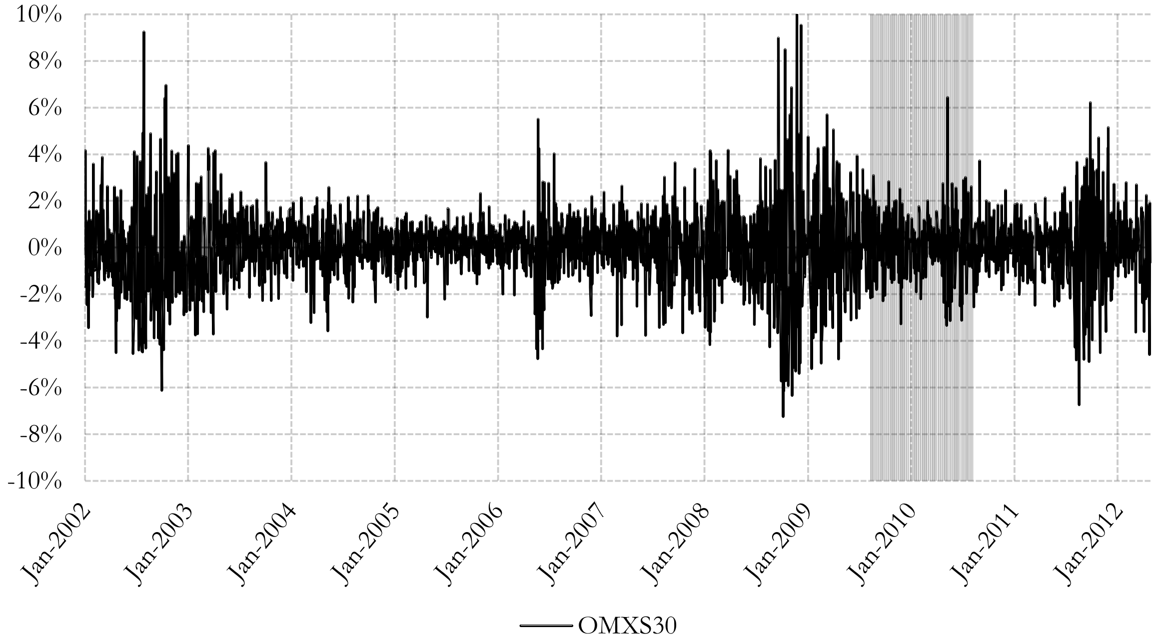

Data source: NASDAQ OMX

Figure X: Weekly Average Quoted Spreads

This figure presents the weekly average quoted spread (in percentage points) for Large and Mid Cap over the period 10 August, 2009 to 5 February, 2010, i.e. the 6 months period before the INET implementation.

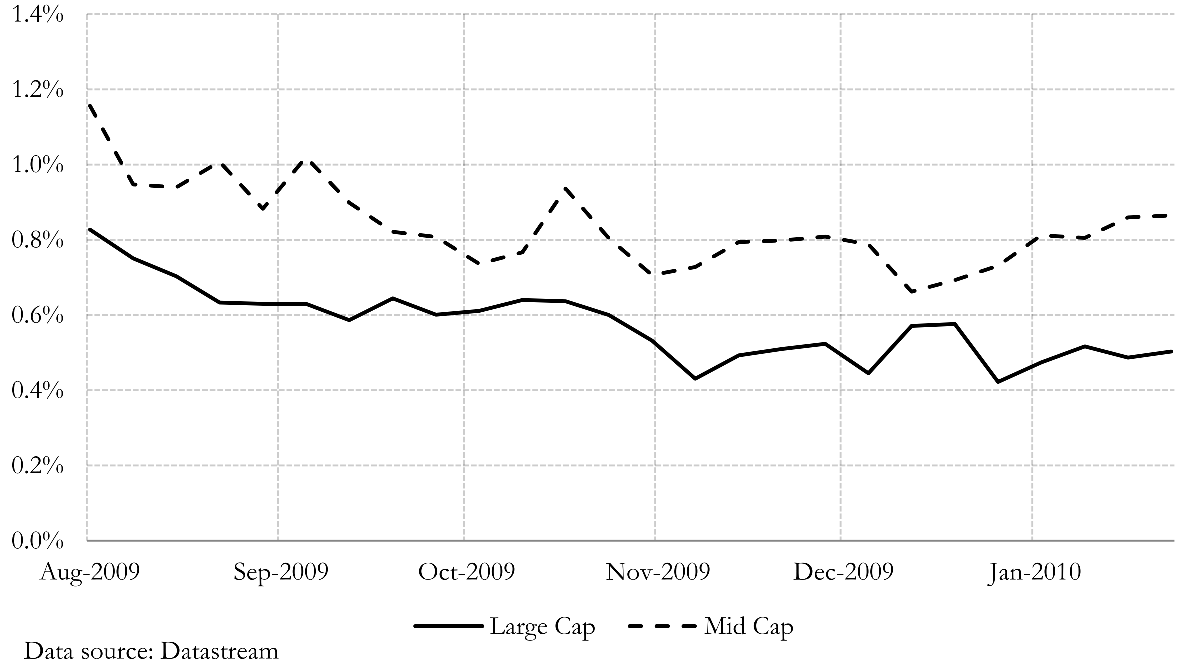

Data source: Datastream

Figure XI: Weekly Average Share Turnovers

This figure presents the weekly average share turnover (in percentage points) for Large and Mid Cap over the period 10 August, 2009 to 5 February, 2010, i.e. the 6 months period before the INET implementation.

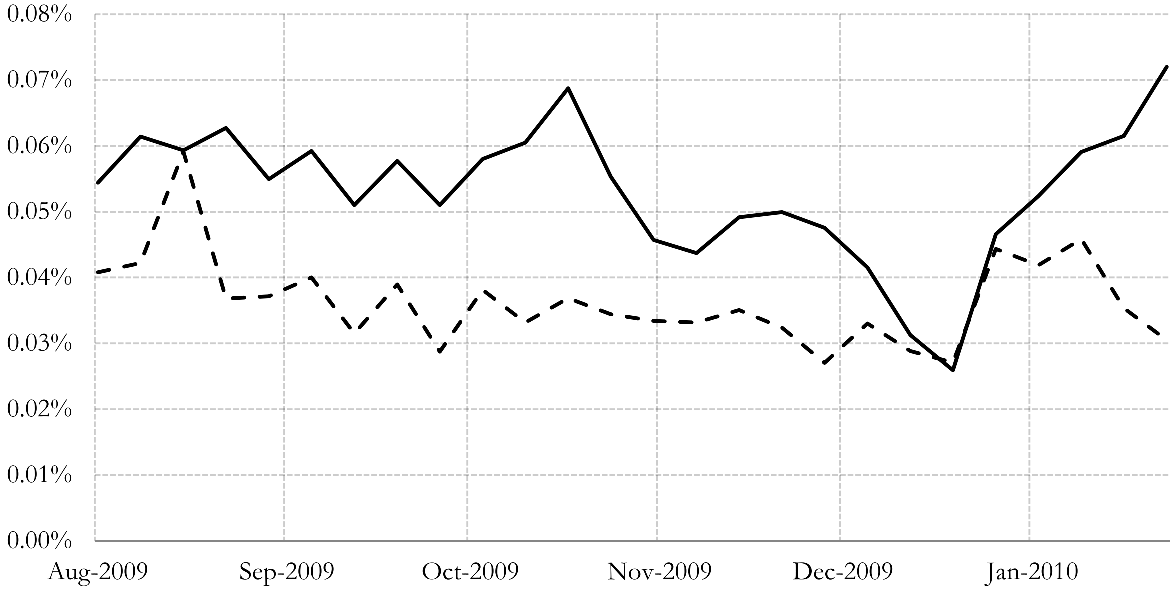

Data source: Avanza Bank AB and Datastream

Figure XII: Average Variance Ratios, 5/1 minutes

This figure presents the average variance ratio for stocks listed on the Large and Mid Cap, respectively, over each quarter of the period 2006Q1 to 2012Q1. The variance ratios are computed by using 5 minutes over 1 minute VWAP prices for Large and Mid Cap stocks, respectively. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

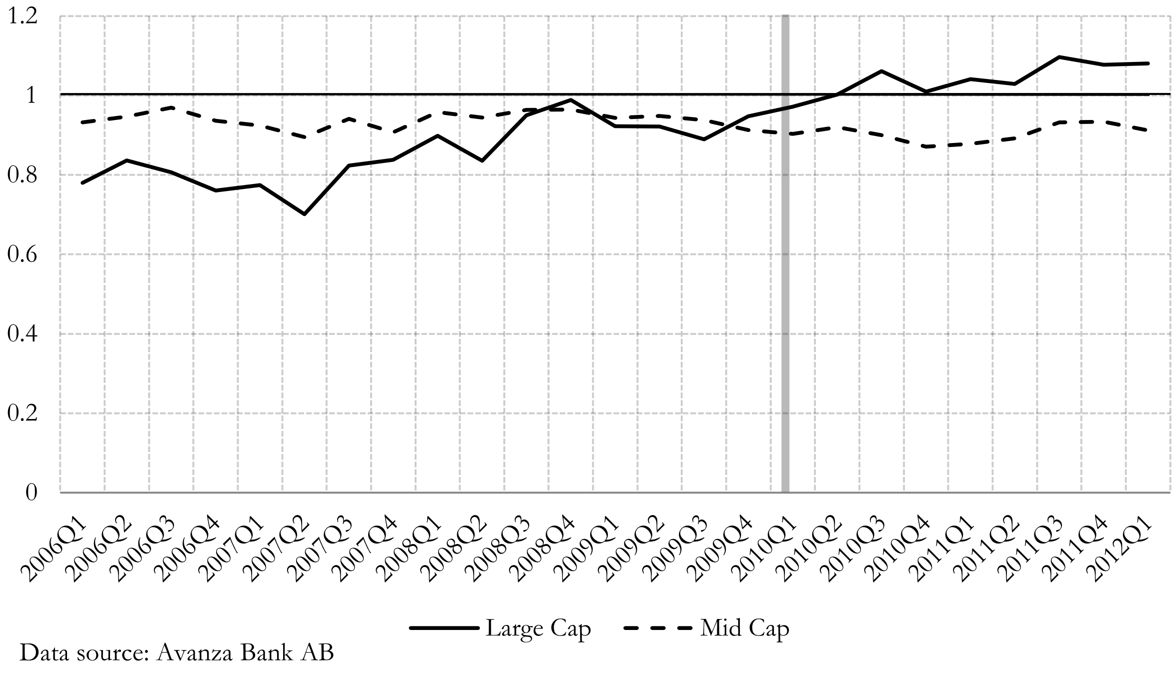

Data source: Avanza Bank AB

Figure XIII: Average Variance Ratios, 10/1 minutes

This figure presents the average variance ratio for stocks listed on the Large and Mid Cap, respectively, over each quarter of the period 2006Q1 to 2012Q1. The variance ratios are computed by using 10 minutes over 1 minute VWAP prices for Large and Mid Cap stocks, respectively. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

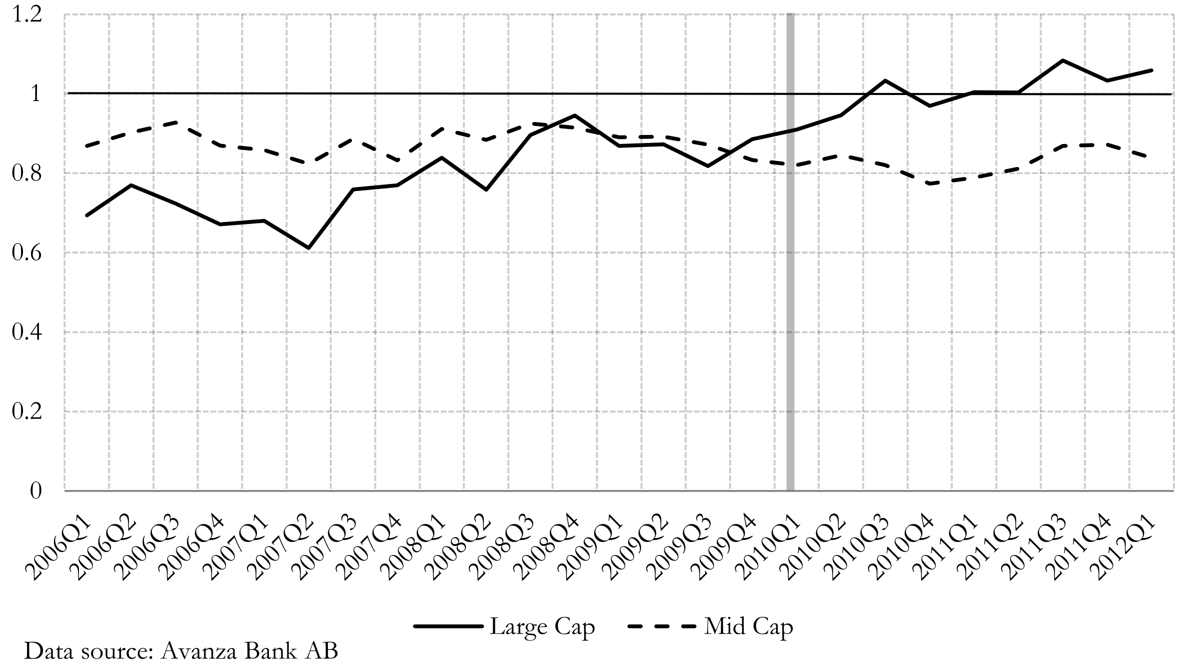

Data source: Avanza Bank AB

Figure XIV: Average Variance Ratios, 4/2 minutes

This figure presents the average variance ratio for stocks listed on the Large and Mid Cap, respectively, over each quarter of the period 2006Q1 to 2012Q1. The variance ratios are computed by using 4 minutes over 2 minutes VWAP prices for Large and Mid Cap stocks, respectively. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

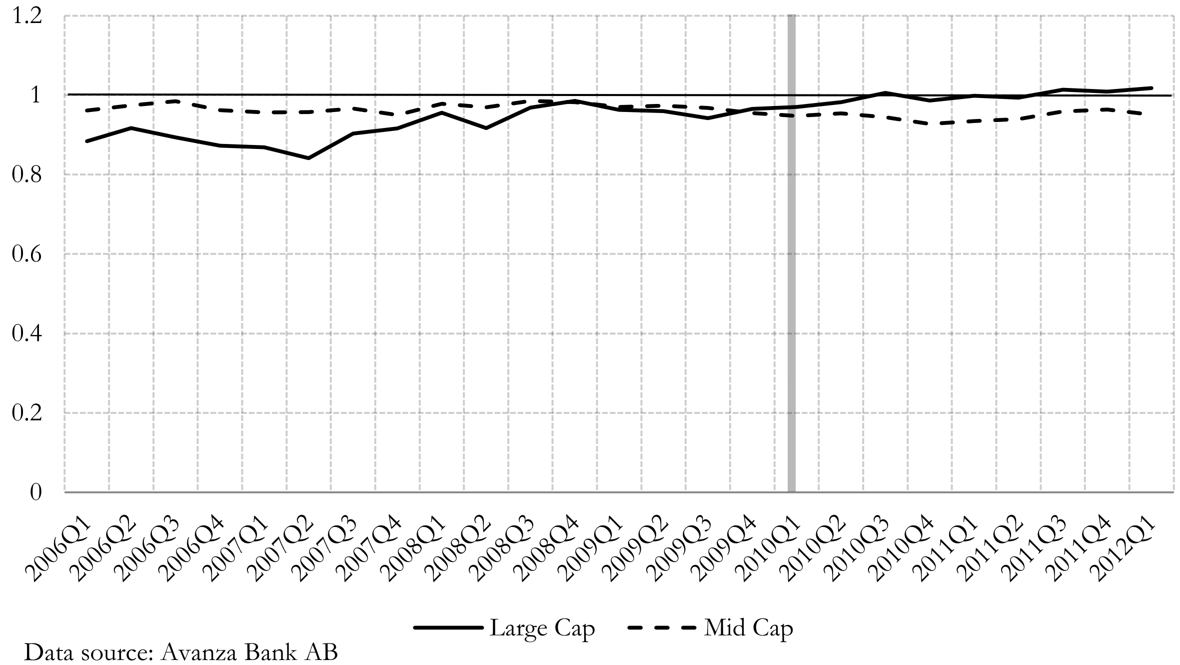

Data source: Avanza Bank AB

Figure XV: Fractions of Significantly Inefficient Stocks, Chow-Denning 5 minutes

This figure presents the fraction of stocks that are significant at the five percent level for the Chow-Denning tests at five minute sampling, i.e. considered inefficient, for Large and Mid Cap listed stocks, respectively, for each quarter over the period 2006Q1 to 2012Q1. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

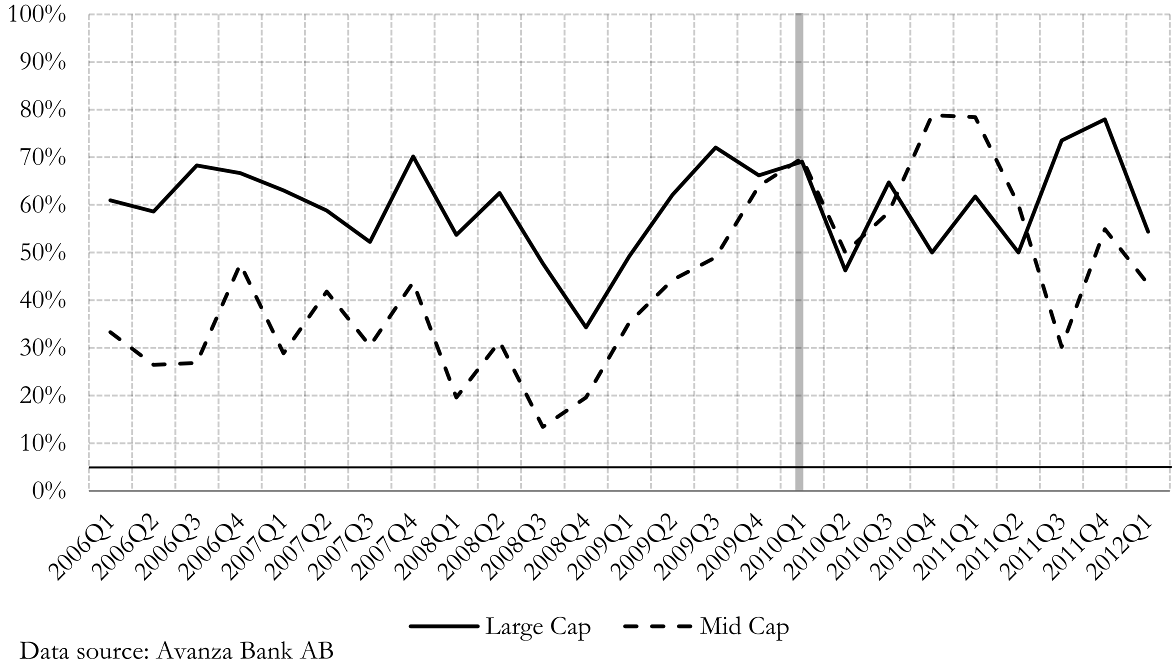

Data source: Avanza Bank AB

Figure XVI: Fractions of Significantly Inefficient Stocks, Chow-Denning 10 minutes

This figure presents the fraction of stocks that are significant at the five percent level for the Chow-Denning tests at ten minute sampling, i.e. considered inefficient, for Large and Mid Cap listed stocks, respectively, for each quarter over the period 2006Q1 to 2012Q1. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

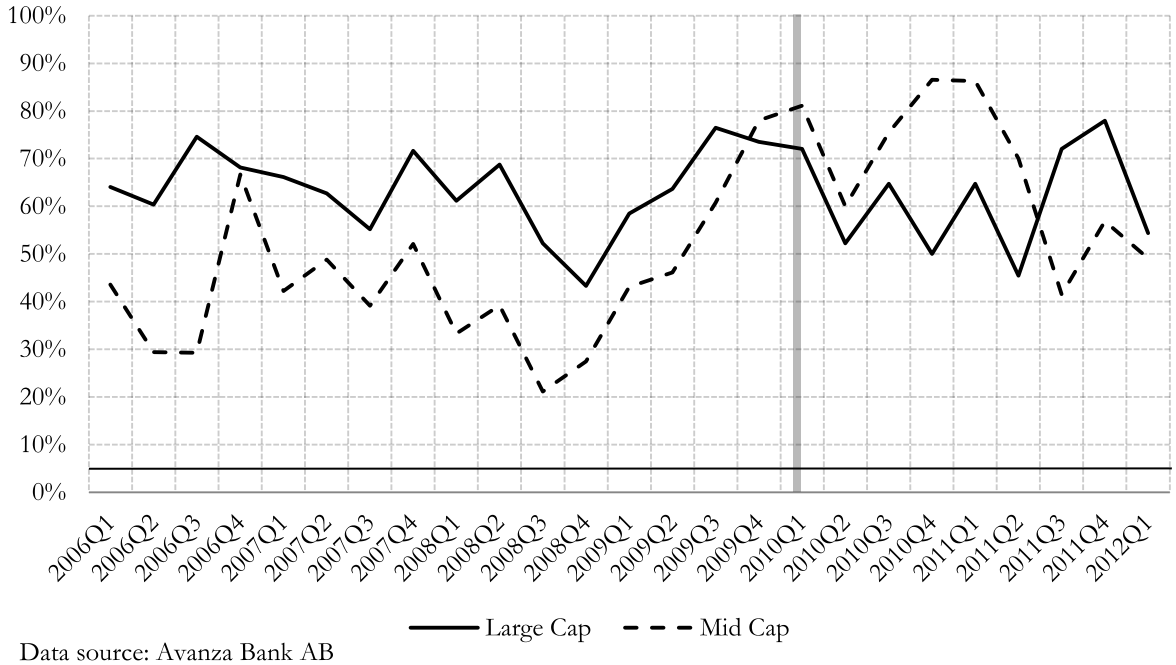

Data source: Avanza Bank AB

Figure XVII: Fractions of Significantly Inefficient Stocks, Chow-Denning 4/2 minutes

This figure presents the fraction of stocks that have a 4 minutes over 2 minute Variance Ratio that is significantly different from one on a five percent significance level, i.e. considered inefficient, for Large and Mid Cap listed stocks, respectively, for each quarter over the period 2006Q1 to 2012Q1. The vertical line marks the date of the INET implementation, i.e. February 8th, 2010.

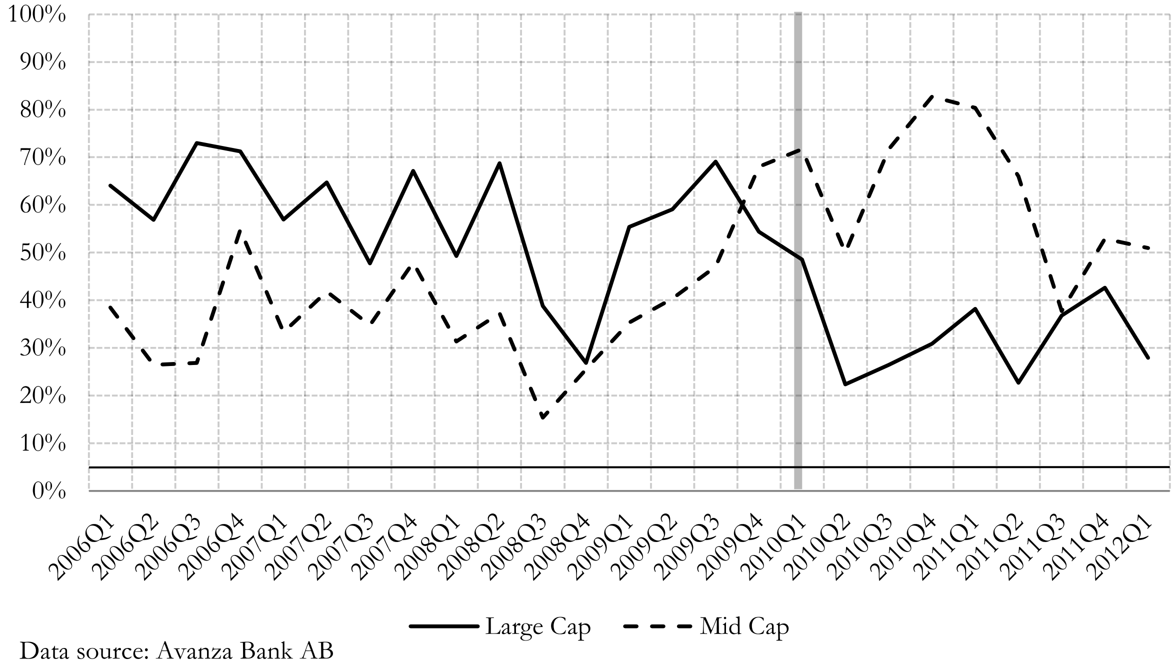

Data source: Avanza Bank AB