HIGH FREQUENCY TRADING, HIDDEN ORDERS AND MARKET QUALITY IN EQUITIES

BY CHENG GAO

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Economics

Written under the direction of

Bruce Mizrach

and approved by

New Brunswick, New Jersey

October, 2015

ABSTRACT OF THE DISSERTATION

High Frequency Trading, Hidden Orders and Market Quality in Equities

by Cheng Gao

Dissertation Director: Bruce Mizrach

The first essay studies the roles of trading speed and hidden orders in limit order markets. We develop a model where liquidity suppliers differ in speed of revising their limit orders and have an option of hiding their orders. The model predicts that fast liquidity suppliers bear lower adverse selection risk and therefore submit orders with narrower bid-ask spreads. Slow liquidity suppliers may overcome their speed disadvantage by using hidden orders. We also provide empirical results that support the model. We find that non-high frequency trading firms account for 70% of liquidity provision in hidden executions, and hidden orders have significantly narrower spreads and lower adverse selection risk than visible orders. Our theoretical model and empirical findings suggest that high frequency technology and hidden orders are substitutes in reducing adverse selection risk.

The second essay investigates market quality breakdowns in equity markets. A breakdown occurs when an order book thins to the point where extreme price movements are observed. These are frequently reversed as the market learns that nothing fundamental has occurred. The daily average breakdown frequency from 1993-2011 is 0.64%, with averages in 2010-11 below this amount. Controlling for microstructure effects, breakdowns have fallen significantly since Reg NMS. Spikes in market correlation and high frequency trading (HFT) surges make breakdowns more likely. Exchange traded funds (ETFs) break down more often than non-ETFs. Both ETFs and HFT Granger cause market correlation. Breakdowns are predictable for up to two days.

The third essay analyzes HFT activity in equities during U.S. Treasury permanent open market (POMO) purchases by the Federal Reserve. We construct a model to study HFT quote and trade behavior when private information is released. We estimate that HFT firms reduce their inside quote participation by up to 8% during POMO auctions. HFT firms trade more aggressively, and they supply less passive liquidity to non-HFT firms. Market impact also rises during Treasury POMO. Aggressive HFT trading becomes more consistently profitable, and HFT firms earn a higher return per share. We also estimate that HFT firms earn profits of over $105 million during U.S. Treasury POMO events.

Acknowledgements

Many people have contributed to the production of this dissertation. I would like to express my gratitude to all people who have made this dissertation possible.

My sincerest thanks are to my advisor, Dr. Bruce Mizrach. I have been extremely fortunate to have an advisor who inspired me with scientific thinking and provided freedom for me to pursue research interests. I am very grateful for his guidance, encouragement and support during my Ph.D. studies. I want to thank Dr. Mizrach for everything he has done for me.

I am indebted to my doctoral committee members: Dr. Norman R. Swanson, Dr. John Landon-Lane and Dr. Xing Zhou, for their insight and suggestions on my dissertation. I appreciate their time and efforts for serving on my committee.

I also thank NASDAQ OMX for generously providing the high frequency data set and TotalView ITCH data for this research.

My thanks also go to the staffs in the Department of Economics for all kinds of support during my studies.

Furthermore, I would like to thank my friends and classmates for their help.

Finally, I want to express my eternal gratitude and love to my parents and my wife for their understanding and continuous support in the past years.

Dedication

For My Parents and My Wife

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The equity markets in the U.S. have changed dramatically over the past two decades. With the advancement of technologies for submitting and executing orders, the markets have evolved from primarily manual trading to highly automated algorithmic trading (AT) systems. Hendershott, Jones, and Menkveld (2011) define AT as the use of computer algorithms to make trading decisions and manage orders. High frequency trading (HFT) is a subset of AT and uses advanced, high-speed technology for order submission, cancellation, and execution in milliseconds. The Securities and Exchange Commission (SEC, 2010) describes HFT as "professional traders acting in a proprietary capacity that engage in strategies that generate a large number of trades on a daily basis." HFT has grown, since the adoption of the Regulation National Market System (Reg. NMS) in 2005, and now represents the majority of equity trading volume. HFT has received lots of attention in academic, regulatory, and public media, and there is an active debate on whether HFT is beneficial or harmful to market quality.

This thesis starts with studying the role of trading speed in reducing adverse selection risk. A substitutional tool to decrease the amount of adverse selection by using hidden orders is also examined in Chapter 2. Chapter 3 analyzes the trend of market quality in the recent decades from 1993-2011 and shows the benefits from AT on market quality. We then explore changes in market structure and market correlation that may have effect on market quality. The last chapter of the thesis investigates the effect of HFT activity in stressful market conditions. It provides a different portrait of HFT from their contribution to market quality in normal times. All these findings help to answer questions about the social value of HFT and may provide policy implications for regulation of HFT activities.

Chapter 2 contributes to the literature by jointly examining the roles of trading speed and hidden orders in limit order markets. We first develop a theoretical model by introducing the speed characteristics of market participants into the Glosten and Milgrom (1985) model. Liquidity suppliers in the model submit bid and offer quotes competitively to the limit order book, and may cancel their limit orders when market condition changes. When fast and slow traders compete to provide liquidity in the market, slow liquidity suppliers are in a worse situation because their limit orders have wider spreads and therefore a lower execution probability. By including hidden orders in the model, We find that slow liquidity suppliers may overcome their speed disadvantage by using hidden orders.

The empirical results provide support for the model. Using NASDAQ TotalView-ITCH and HFT data, We find that non-HFT firms account for nearly 70% of liquidity provision in hidden executions, and their hidden orders have significantly narrower spreads and lower adverse selection costs than visible orders. In contrast, HFT firms who are able to reduce the adverse selection risk by updating their limit orders quickly are less likely to hide orders. Our theoretical model and empirical findings suggest that high-speed technology and hidden orders play equivalent roles for liquidity suppliers to reduce adverse selection risk.

Traditional market quality measures, such as bid-ask spread, may not capture some aspects of today's market. In the third chapter, we develop a new market quality metric that is motivated by events like the Flash Crash on May 6, 2010 and the Twitter hack on April 23, 2013 where stock prices dropped sharply and rebounded. A stock is identified as having a market quality breakdown if its intraday prices experience (1) decline: the best bid prices fall 10% or more below the 09:35 price; (2) recovery: the low price on (09:35,15:55) must be 7.5% lower than the 15:55 price; and (3) non-fleeting: the low tick must be repeated at least once in a subsequent calendar second. Our measure relies on quotes rather than trades because breakdowns in market quality impede trading, and that the consolidated quotes provide the best real time portrait of the market. We focus on the best bid and offer from the listing exchange. In addition, we look at movements in the time frame from 09:35-15:55, given the fact that opening and closing procedures vary across exchanges and may not be comparable.

We find that the daily average breakdown frequency is 0.64% throughout the sample period from 1993-2011, an average of 44 stocks per day. The trend of market quality breakdowns shows one peak in 2000 and another one in 2008 during the financial crisis. After 2009 the breakdown frequency has returned to the level in the late 1990s. By 2011 it has fallen to 0.39%, half the rate in 1998 when humans provided the majority of quotes. Controlling for microstructure effects, breakdowns occur 41.87% less frequently since Reg. NMS was fully in place in 2007. Despite the Flash Crash, 2010 has the fewest breakdowns of any year since 2007. These findings show that market quality is not becoming worse more recently. Indeed the Reg. NMS market structures that have facilitated HFT have produced improvements in market quality. However, this does not mean that HFT activities are all benign to the market.

The academic literature on HFT's impact on market quality has focused on its liquidity provision, contribution to price discovery, influence on bid-ask spread and volatility. Menkveld (2013) analyzes the arrival of the Chi-X high frequency platform in Europe and concludes that HFT firms act as market makers in the new market. Brogaard, Hendershott and Riordan (BHR, 2014) find that HFT increases price efficiency through their marketable orders. Hasbrouck and Saar (2013) identify HFT activity by linked orders and suggest that HFT lowers short-term volatility and bid-ask spreads, and increases displayed depth.

Although HFT firms are passive liquidity providers in most trades, it does not indicate that they always act as market makers. A key part of market makers is that they have affirmative obligations to contribute liquidity not only in good times but more importantly in bad ones, and also not front-running by trading for their own accounts ahead of completing customers' orders. Non-AT systems have suffered similar limitations. Christie and Schultz (1994) provide evidence on the implicit collusion by NASDAQ market makers to maintain artificially wide bid-ask spreads. This price-fixing case leads to the introduction of the SEC Order Handling Rules. Ellis (2010) documents that NYSE specialists engaged in widespread misconduct between 1999 and 2003, including inter-position, front-running, and freezing. Their violations result in a subsequent NYSE reform and further regulations on securities markets. Although regulations that promoted HFT have narrowed spreads, it does not mean that regulation of specific HFT activities is unnecessary.

For this purpose, we provide evidence in Chapter 4 that HFT firms behave quite differently when the market is under uncertainty. Using the same NASDAQ high frequency data set as BHR, we are able to identify whether an HFT firm initiated or filled a trade. The data set also tells whether an HFT firm is providing liquidity at each tier of the order book. We study HFT activity in a period of potential market stress, the U.S. Treasury Permanent Open Market Operations (POMO) by the Federal Reserve starting in late 2008 as part of its quantitative easing program. We find that HFT firms scale back from inside quotes during POMO auctions. They also trade more aggressively and provide less passive liquidity to non-HFT market participants. HFT firms make more consistent profits in aggressive trading from non-HFT firms during POMO events than normal times, with profits per share rise 300%. They are profitable 88% of the time on aggressive trades and 100% of the time on passive trades. We estimate that HFT firms earn profits of over $105 million during U.S. Treasury POMO events. These findings indicate that studying the average effect of HFT firms may provide a misleading conclusion of their contribution to market quality.

Chapter 2

Strategies to Reduce Adverse Selection Risk: Are Speed and Hidden Orders Equivalent?

2.1 Introduction

High frequency trading (HFT) and market transparency are two crucial aspects in the regulation of equity market structure. HFT has grown since the adoption of the Regulation National Market System (Reg. NMS) in 2005, and now represents the majority of equity trading volume. HFT firms use high speed technology for order submission, cancellation, and execution. There is an active and ongoing debate on the impact of HFT on market quality. In a speech given in June 2014 the Securities and Exchange Commission (SEC) Chairman Mary Jo White discussed the importance of addressing the concern about HFT and promoting fairness for investors. [1] Pre-trade opacity is another concern for regulators and policy makers. Many exchanges allow traders to submit hidden orders that are not visible to the market. The impact of hidden liquidity remains an open question. Studies on HFT and hidden orders are two isolated strands in the academic literature. This paper contributes to the literature by analyzing the roles of trading speed and hidden orders *jointly*.

In a limit order market, buyers and sellers submit orders into a limit order book. A buy limit order specifies the highest price the trader is willing to pay for a given amount of the security. A sell limit order specifies the lowest price the trader is willing to accept

[1] Mary Jo White, “Enhancing Our Equity Market Structure,” http://www.sec.gov/News/Speech/Detail/Speech/1370542004312.

for a given amount. Buy and sell limit orders are called bids and asks respectively. The difference between the best bid and the best ask price is called the bid-ask spread. Incoming orders are matched directly with orders that are already present in the limit order book. A trade occurs if a newly entered order cross the spread. The trade initiators are often called liquidity demanders, and traders on the passive side of trades are called liquidity suppliers. Copeland and Galai (1983) note that buy and sell limit orders are respectively put and call options for liquidity demanders. When new information arrives at the market, liquidity suppliers face the risk of being picked off by liquidity demanders if they are delayed to cancel or update their orders. In this situation limit orders may get unfavorable executions, which is often referred to as the adverse selection problem. In this paper, we study the effects of HFT and hidden orders on adverse selection.

HFT is indeed a general term that includes a variety of trading strategies. The SEC (2010) discusses four types of HFT strategies: passive market making, arbitrage, structural, and directional trading. Biais and Foucault (2014) note the HFT heterogeneity and classify HFT strategies into five categories ranging from market making to manipulation. Despite differences among HFT strategies, what's in common is that they all rely on advanced high speed technology. Speed plays an important role for participants in equity markets. Liquidity suppliers who submit bid and ask quotes to the limit order book actively monitor the market and update their limit orders accordingly. The use of high speed technology enables them to incorporate real-time information quickly into their quotes and therefore be better protected against unfavorable executions. As fast liquidity suppliers face lower risk of being adversely selected, they are generally willing to narrow the bid-ask spread and provide more liquidity. Both theoretical and empirical studies suggest that HFT liquidity supplying activities generally improve market quality. Jovanovic and Menkveld (2012) propose a model in which high frequency traders may reduce adverse selection costs by acting as middlemen in limit order markets.They also find empirical evidence from the Dutch equity markets that the effective bid-ask spread declined by 15% and the adverse selection cost fell by 23%

on the entry of a large, passive HFT firm. Hasbrouck and Saar (2013) find that increased low-latency activity improves liquidity by lowering bid-ask spreads and increasing displayed depth. Malinova, Park, and Riordan (2013) find that bid-ask spreads and adverse selection costs rise in Canadian equity markets after an increase in regulatory fees that affected HFT market making activities. Hagströmer and Norden (2013) find that the majority of HFT volume at NASDAQ OMX Stockholm (NOMX-St) comes from market making activities and HFT market making reduces short-term volatility. Brogaard, Hagströmer, Norden, and Riordan (2014) find that there is an increase in the share of liquidity provision by traders who upgrade colocation subscriptions at NOMX-St and a decrease in their adverse selection costs.

Speed also matters for liquidity demanders because those who have fast technology may access news quickly and trade before others. Martinez and Roşu (2013) model HFT participants as informed traders who observe news stream and trade quickly. They find that HFT generates trading volume and volatility and decreases liquidity. Foucault, Hombert, and Roşu (2013) argue that the ability of HFT firms to receive news faster creates additional information asymmetry and thus reduce liquidity. Biais, Foucault, and Moinas (2013) suggest that differences in speed increase adverse selection costs and thus HFT generates negative externalities. Carrion (2013) and Brogaard, Hendershott, and Riordan (2014) analyze a data set from NASDAQ that identifies HFT firms and find that HFT firms act as aggressive liquidity demanders in more than 50% of their trading activities. Brogaard, Hendershott, and Riordan (2013) use the 2008 short sale ban as an exogenous shock and find that HFT firms decrease liquidity and increase volatility. Gao and Mizrach (2013) find that HFT firms decrease their market making activity and increase their aggressive trades during Federal Reserve Treasury purchases. Hirschey (2013) notes that HFT firms anticipate the order flow from non-HFT investors and their aggressive trades are highly correlated with future returns. Breckenfelder (2013) finds that competition among HFT firms is associated with an increase in their liquidity demanding activities that lead to a decline in liquidity and a rise in short-term volatility.

Therefore, the net effect of speed improvement on market quality relies on its relative contribution to liquidity suppliers and demanders. In this paper, we first develop a theoretical model by introducing the speed characteristics of market participants into the Glosten and Milgrom (1985) model. Liquidity suppliers in our model submit bid and ask quotes competitively to the limit order book, and may choose to cancel their limit orders when market condition changes. We assume that there are two types of liquidity demanders in the market: informed traders and noise traders. Informed traders know the fundamental value of the security and want to trade quickly using high speed technology. Noise traders trade for exogenous reasons other than the fundamental and do not rely on high speed. Whether liquidity suppliers are able to cancel their limit orders depends on their relative speed to liquidity demanders. We find that fast liquidity suppliers bear lower risk of being adversely selected and therefore submit limit orders with narrow bid-ask spreads. In contrast, slow liquidity suppliers are less likely to cancel their limit orders on news arrival and charge wide bid-ask spreads. The model predicts that adverse selection costs fall and spreads narrow when the relative speed of liquidity suppliers increases.

We then empirically examine the effect of trading speed on adverse selection costs and spreads using NASDAQ TotalView-ITCH data. Our initial sample consists of the 120 stocks that are studied by Carrion (2013) and Brogaard, Hendershott, and Riordan (2014). We excluded delisted stocks during our sample period from 2008-2013. The final sample contains 108 stocks that are stratified by market capitalization and listing exchanges. We find that the speeds of both liquidity suppliers and demanders on NASDAQ have increased significantly over the sample period. We measure the relative speed of liquidity suppliers to demanders by the ratio of the 5th percentile life of executed orders to that of cancelled orders. Controlling for volatility, trading volume, price, and market capitalization, we find that quoted and effective spreads narrow and adverse selection costs fall when the relative speed of liquidity suppliers to demanders increases. We also find that HFT participation in either the liquidity supplying or demanding side has been rising from 2008-2013. The share of HFT cancellations increases from 19% in 2008 to nearly 30% in 2013. The proportion of HFT aggressive executions rises from 17% in 2008 to more than 23% in 2013.

Our model also predicts that limit orders of slow liquidity suppliers are less likely to be executed in the presence of fast liquidity suppliers. Slow liquidity suppliers face the choice of exiting the market or using some other strategy to compete with fast liquidity suppliers. Many stock exchanges allow traders to hide part or all of the quantity of their limit orders. We find that slow liquidity suppliers may use hidden orders as a tool to overcome their speed disadvantage.

Empirical studies show that hidden orders account for a large proportion of liquidity on exchanges. Aitken, Berkman, and Mak (2001) report that undisclosed orders represent about 28% of trading volume on the Australian Stock Exchange. Pardo and Pascual (2006) find that 18% of trades in the Spanish Stock Exchange involve hidden order executions. Bessembinder, Panayides, and Venkataraman (2009) study hidden liquidity on Euronext Paris and find that hidden orders account for 44% of volume. Frey and Sandås (2012) analyze DAX-30 stocks on Deutsche Börse's electronic trading platform Xetra and find that the average share of iceberg orders is 9% of all non-marketable orders and 16% of overall volume. Yao (2012) finds that about 11% of trades on NASDAQ are executed against hidden orders in 2010 and 2011.

We study the sample of 108 stocks during the period from 2008-2013 using the ITCH data. Over the sample period, 13% of executions on NASDAQ involve hidden orders. The proportion of hidden executions is 17% in 2008 and then falls below 11% during the period from 2009-2011. The frequency of hidden executions gradually rises in more recent years. Hidden executions occur nearly 15% of time in 2013.

The academic literature provides some possible explanations for the use of hidden orders. Aitken, Berkman, and Mak (2001) as well as Pardo and Pascual (2006) suggest that the use of hidden orders is mainly a way for liquidity suppliers to reduce the option value of limit orders. Buti and Rindi (2013) present a theoretical model where traders who have no private information on asset values may rationally use reserve orders to reduce adverse selection costs. Tuttle (2006) studies hidden order usage in the SuperSOES system and finds that hidden liquidity reduces the adverse selection risk for limit order submitters. De Winne and D'Hondt (2007) also find empirical evidence that traders use hidden orders to manage picking off risk. An alternative explanation for hidden order usage comes from the information exposure problem of displayed orders. Harris (1997) argues that order exposure could reveal trader's intentions which other traders may respond to. He discusses that other traders may either withdraw from trading or front-run limit order traders. Therefore, hidden orders provide a way for traders to reduce the informational impact.

The relationship between hidden orders and market quality remains an open question. Moinas (2010) develops a theoretical model in which informed traders use hidden limit orders to conceal their private information and trading intentions. Her model predicts that the introduction of hidden orders improve market liquidity and total welfare. However, Anand and Weaver (2004) find that market quality neither improved nor deteriorated after the Toronto Stock Exchange abolished hidden limit orders in 1996. They also find that the re-introduction of hidden orders in 2002 had no effect on spreads or visible depth. Other studies suggest that the strategic use of hidden orders may reduce market liquidity and harm price discovery. Foley, Malinova, and Park (2013) find that the introduction of hidden orders on the Toronto Stock Exchange in 2011 resulted in wider quoted spreads and higher trading costs. Cebiroğlu, Hautsch, and Horst (2014) suggest that hidden orders generate trading frictions and thus reduce price efficiency and lead to excess volatility.

We model the use of hidden orders by liquidity suppliers with heterogeneous speeds. When fast and slow traders compete to provide liquidity in the market, slow liquidity suppliers are in a worse situation because their limit orders have wider spreads and therefore lower execution probabilities. Our model suggests that slow liquidity suppliers may overcome their speed disadvantage by using hidden orders. Hidden orders reduce the risk of being picked off by informed traders and therefore enable slow liquidity suppliers to submit quotes that are competitive with fast liquidity suppliers. Given that hidden orders also reduce the chance of trading with uninformed noise traders, fast liquidity suppliers have less incentive to hide their orders than slow ones. Our empirical results provide support for the model. We find that hidden orders have significantly lower adverse selection costs and narrower spreads than visible orders, and non-HFT firms account for 70% of liquidity provision in hidden executions. HFT liquidity suppliers who are able to reduce the adverse selection risk by updating their limit orders quickly are less likely to hide limit orders. Our results suggest that hidden orders and speed technology are substitutes to manage the adverse selection risk.

This chapter is organized as follows. We develop a theoretical model in Section 4.3 to study the roles of speed and hidden orders in reducing adverse selection risk. Section 4.4 describes the NASDAQ HFT data set and the TotalView-ITCH data. We present in Section 2.4 empirical results on speed improvement over time and its impact on adverse selection and spreads. We then analyze hidden orders in Section 2.5 and show that non-HFT firms use hidden orders to compete with HFT firms in liquidity provision. Section 4.7 concludes.

2.2 The Model

We introduce the speed characteristics of liquidity suppliers into the Glosten and Milgrom (1985) model. Liquidity suppliers may use high speed technology to monitor the market and revise their limit orders after submission. In this section, we first develop a baseline model with homogeneous liquidity suppliers. We then study the case in which liquidity suppliers have different speeds. Finally, we extend this framework by including hidden orders in the model.

2.2.1 Primitives

Consider a limit order book for a risky security with its fundamental value V_t that follows

$$V_{t+1} = V_t + \varepsilon_{t+1}, \tag{2.1}$$

where ε_{t+1} is the value change due to information arrival and can take three values: $\varepsilon_{t+1} = +\sigma > 0$ if the news is positive, $\varepsilon_{t+1} = -\sigma$ if negative, and $\varepsilon_{t+1} = 0$ if no news occurs. We assume that information about the fundamental arrives with probability γ and it could be positive or negative with equal probability. With the complementary probability $1 - \gamma$ no information arrives.

There are two types of agents who demand liquidity in the market: noise traders and informed traders. Noise traders submit market orders to trade for liquidity reasons that are not related to information about the security value. A noise trader could be a buyer or a seller with equal probability. Informed traders know the fundamental value of the security. They trade immediately upon information arrival and maximize their profits by trading up to the fundamental value using marketable limit orders.

Liquidity in the limit order book is provided by risk-neutral traders who place bid and ask quotes competitively. Liquidity suppliers in our model differ from traditional market makers in the aspect that they do not have obligations to maintain liquidity. In other words, they want to cancel limit orders as quickly as possible when the adverse selection risk is high. The success of order cancellation relies on their investment in high speed technology.

Liquidity suppliers in our model are not informed of the fundamental value. However, they may infer the informedness of an incoming order by analyzing order flow information. We assume that with probability α they make correct inference on the identity of an arriving liquidity demander. With probability $1 - \alpha$ they are wrong. In this setting α is large when liquidity suppliers have superior speed to monitor the market at high frequency and process order flow information quickly. Liquidity suppliers are perfectly informed of order flow if $\alpha = 1$, and absolutely uncertain about the informedness of an incoming order if $\alpha = 1/2$. Liquidity suppliers would choose to cancel their limit orders when they infer that an incoming order is from a informed trader, and decide not to cancel when an order is from a noise trader.

The use of high speed technology also provides liquidity suppliers with more chance to withdraw their liquidity when they anticipate a high risk of adverse selection. It is important to note that whether a liquidity supplier succeeds in cancelling a limit order depends on her relative speed to liquidity demanders. We assume that speed is not a primary concern for noise traders so that a liquidity supplier is always able to cancel her existing limit order when a noise trader arrives. In contrast, informed traders also invest in speed technology to ensure immediate execution. Therefore, on the arrival of a informed trader, a liquidity supplier is faster and able to cancel her limit order with probability λ, but being picked off with probability $1 - \lambda$.

The timing for traders arriving at the limit order market has the following sequence. In stage 1 liquidity suppliers submit limit orders. In stage 2 either a noise trader or a informed trader arrives, and liquidity suppliers choose whether to cancel existing limit orders based on their inference on the identity of the liquidity demander. If a liquidity supplier decides to cancel a limit order, the success of the request depends on her relative speed to the liquidity demander. The value of the security is then realized after stage 2.The trading game repeats from stage 1. The tree presented in Figure 1 describes the trading process.

[INSERT Figure 1 HERE]

The occurrence of an information event is selected by nature. If an event occurs, informed traders buy on positive news and sell on negative news. If no event occurs noise traders arrive to buy or sell with equal probability. Liquidity suppliers are not informed of information events and make inference on the identity of liquidity demanders. They choose to cancel limit orders when inferring the arrival of informed traders and able to do so only if they are faster. Liquidity suppliers would decide not to cancel limit orders when inferring that noise traders arrive.

To simplify the analysis we assume that each order submitted by traders is for one unit of the security. Due to symmetry it suffices to analyze only one side of the limit order market. Without loss of generality we study the decision problem of liquidity suppliers on the ask side of the book while liquidity demanders submit buy orders.

2.2.2 Order submission strategy of liquidity suppliers

Suppose a liquidity supplier places a ask quote at A_t and define $s_t = A_t - V_t$ as the half-spread on the ask side. With probability $(1-\gamma)/2$ a market buy order from a noise trader arrives at the order book. In this case the liquidity supplier recognizes the noise trader with probability α and chooses to stay in the market to earn a profit of s_t. With probability $1-\alpha$ the liquidity supplier cancels the ask quote due to her wrong belief that a informed trader arrives.

A informed trader buys in the occurrence of a positive news with probability $\gamma/2$. The liquidity supplier chooses to cancel the ask quote with probability α and not to cancel with probability $1-\alpha$ given her inference on the order flow information. In the case that the liquidity supplier chooses to cancel her limit order, she is not able do so with probability $1-\lambda$ when slower than the informed trader. Therefore, the expected profit for the liquidity supplier is

$$\begin{aligned}\pi_t &= \frac{1-\gamma}{2}[\alpha s_t + (1-\alpha)\cdot 0] + \frac{\gamma}{2}[\alpha(\lambda\cdot 0 + (1-\lambda)(s_t-\sigma)) + (1-\alpha)(s_t-\sigma)] \\ &= \frac{1}{2}[(1-\gamma)\alpha s_t + \gamma(1-\alpha\lambda)(s_t-\sigma)]. \qquad (2.2)\end{aligned}$$

This shows that the probability of being adversely selected for a liquidity supplier is $\gamma(1-\alpha\lambda)/2$, which is a decreasing function of α and λ. This is consistent with the intuition that the adverse selection risk is lower if liquidity suppliers are better informed of order flow or faster to cancel their limit orders. The equivalent role of α and λ in the expression suggests that order flow information and cancellation speed have the same effect on reducing the adverse selection risk for liquidity suppliers.

Let $s_t^*(\alpha, \lambda)$ be the competitive half-spread given that liquidity suppliers are informed of order flow with probability α and faster than informed traders with probability λ. $s_t^*(\alpha, \lambda)$ is obtained when the zero expected profit condition is satisfied, i.e. $\pi_t = 0$.

$$s_t^*(\alpha, \lambda) = \frac{\gamma(1-\alpha\lambda)}{(1-\gamma)\alpha + \gamma(1-\alpha\lambda)}\sigma. \tag{2.3}$$

It is straightforward to show that $\partial s_t^*(\alpha, \lambda)/\partial\alpha < 0$ and $\partial s_t^*(\alpha, \lambda)/\partial\lambda < 0$, that is, the competitive spread decreases in α and λ. Other things being equal, the bid-ask spread will become narrower when liquidity suppliers are better informed of order flow. Similarly, the spread will be smaller when liquidity suppliers are faster and more likely to cancel limit orders on the arrival of informed traders. This result is consistent with the finding by Hendershott, Jones, and Menkveld (2011) that algorithmic trading narrows spreads.

The result in (2.3) also suggests that different combinations of α and λ may result in the same competitive spread. For example, when $\alpha = 1$ and $\lambda = 0$, liquidity suppliers are able to identify noise and informed traders perfectly but cannot cancel outstanding limit orders successfully when informed traders arrive. In this case the competitive half-spread is $s_t^*(1, 0) = \gamma\sigma$. In another case of $\alpha = 1/2$ and $\lambda = 1$, when liquidity suppliers are always able to cancel limit orders when recognizing informed traders but completely uninformed of order flow, the competitive half-spread is the same, i.e. $s_t^*(1/2, 1) = \gamma\sigma$. In general, when $\lambda = 1/\alpha - 1$, the competitive half-spread of liquidity suppliers reduces to $s_t^*(\alpha, \lambda) = \gamma\sigma$ that is exactly the same as the result from the Glosten and Milgrom (1985) model.

2.2.3 Liquidity suppliers with heterogeneous speeds

The competitive half-spread s_t^* in (2.3) is derived when liquidity suppliers are all identical with respect to cancellation speed and order flow information. We then consider the case that technology innovation enables a group of liquidity suppliers to have greater advantage than others. Given that order flow information is equivalently effective with cancellation speed to reduce the risk of adverse selection, we assume that all liquidity suppliers have the same capability of processing order flow information but only differ in cancellation speed. In other words, α is the same for all liquidity suppliers but λ is not. Let λ_F be the probability that fast liquidity suppliers are able to cancel limit orders on the arrival of informed traders, and λ_S be the probability for slow liquidity suppliers with $\lambda_S < \lambda_F$.

From (2.3) we know that the competitive half-spread is narrower for liquidity suppliers who have higher speed. Since the existence of slow liquidity suppliers does not affect the expected profit of fast ones, the competitive half-spread for fast liquidity suppliers would be

$$s_F^* = s_t^*(\alpha, \lambda_F) = \frac{\gamma(1-\alpha\lambda_F)}{(1-\gamma)\alpha + \gamma(1-\alpha\lambda_F)}\sigma. \tag{2.4}$$

In contrast, the order submission strategy of slow liquidity suppliers may change in the presence of fast ones. Slow liquidity suppliers have opportunities to trade with noise or informed traders only when fast liquidity suppliers cancel their outstanding limit orders. In specific, the probability that slow liquidity suppliers are contacted by noise trader falls to $(1-\gamma)(1-\alpha)/2$, and the probability of being contacted by informed traders decreases to $\gamma\alpha\lambda_F/2$. Therefore, the expected profit for slow liquidity suppliers is given as

$$\pi_S = \frac{1}{2}[(1-\gamma)(1-\alpha)\alpha s_t + \gamma\alpha\lambda_F(1-\alpha\lambda_S)(s_t - \sigma)]. \tag{2.5}$$

The competitive half-spread for slow liquidity suppliers is then obtained by setting $\pi_S =$

0.

$$s_S^* = \frac{\gamma\lambda_F(1-\alpha\lambda_S)}{(1-\gamma)(1-\alpha)+\gamma\lambda_F(1-\alpha\lambda_S)}\sigma. \tag{2.6}$$

We see that s_S^* could be different the competitive half-spread for slow liquidity suppliers in the absence of fast ones, $s_t^*(\alpha, \lambda_S)$. Suppose that $s_t^*(\alpha, \lambda_S) = \gamma\sigma$, so we have $\lambda_S = (1-\alpha)/\alpha$ and $\lambda_F > (1-\alpha)/\alpha$. In this case, $s_S^* > s_t^*(\alpha, \lambda_S)$, which suggests that slow liquidity suppliers will widen their spreads when they are not able to adopt more advanced speed technology to compete with fast liquidity suppliers. In particular, the case of $\alpha = 1$ in which both fast and slow liquidity suppliers are perfectly informed of order flow information results in $s_S^* = \sigma$, which means that slow liquidity suppliers submit competitive ask quotes at $V_{t-1} + \sigma$ in the presence of fast ones. Although slow liquidity suppliers will not suffer any loss at this price, their orders indeed have little execution probability. This suggests that slow liquidity suppliers may be driven out of the market by fast ones.

2.2.4 Model with hidden orders

We now extend the above framework to allow liquidity suppliers to hide their limit orders in the book. Hidden orders are not visible to other market participants but have lower time priority to be executed than the displayed orders at the same price. By choosing hidden orders instead of displayed ones, liquidity suppliers reduce the risk of being picked off by informed traders upon the arrival of news. We assume that with probability θ a hidden limit order is executed by informed traders even if there is no visible limit order in the book. Under this assumption liquidity suppliers reduce the probability of trading with informed trader by $1-\theta$ if submitting hidden orders. However, when placing hidden orders liquidity suppliers also have a lower chance to trade with noise traders to earn profits. We assume that noise traders are less likely to submit market orders in the absence of visible orders, and with probability η they still choose to trade. Noise traders generally trade for exogenous reasons and are less sensitive to the visibility of orders than informed traders. Therefore we assume in the model that $\theta < \eta$.

Liquidity suppliers have an option to choose between visible and hidden orders. There are five possible equilibria in which both fast and slow liquidity suppliers have zero expected profits. We present them in Table 2.1.

[INSERT Table 2.1 HERE]

Visible orders submitted by slow liquidity suppliers have wider competitive spreads than those by fast liquidity suppliers and therefore have lower execution probability. Hidden orders provide a tool for slow liquidity suppliers to reduce the adverse selection risk so that they may improve their quote prices. By using hidden orders slow liquidity suppliers are able to compete with fast ones at the top of the queue. As a result, the optimal strategy for slow liquidity suppliers is to use hidden orders.

Fast liquidity suppliers are able to manage the adverse selection risk by updating their limit orders quickly. They are less likely to submit hidden orders because it reduces their chance to trade with noise traders to earn profits. Therefore, fast liquidity suppliers will choose visible orders as their optimal strategy to attract high trading volume.

2.2.5 Empirical predictions

The theoretical model suggests that the speed of liquidity suppliers would affect their order placement strategies in limit order markets. It has several empirical predictions that we summarize as follows.

H1: Adverse selection costs fall when the relative speed of liquidity suppliers to demanders increases.

H2: Spreads narrow when liquidity suppliers improve their relative speed.

H3: Hidden orders are used more frequently by slow liquidity suppliers than fast ones.

H4: Hidden orders bear lower adverse selection risk than visible orders.

H5: Hidden orders have narrower spreads than visible orders.

We then use NASDAQ data to test these hypotheses.

2.3 Data

The main data sets used in this paper are NASDAQ HFT database and TotalView-ITCH data. In addition, we obtain the number of outstanding shares from the Center for Research in Security Prices (CRSP) to calculate firms' market capitalization.

2.3.1 HFT database

The NASDAQ HFT database provides information about trades and quotes on 120 randomly selected stocks with half of the firms listed on NASDAQ and the other half on New York Stock Exchange (NYSE). The sample consists of three market capitalization groups: large, mid, and small. Each group contains 40 stocks. Stocks in the large-cap group are selected from the largest market capitalization stocks. The mid-cap group is composed of stocks around the 1000th largest stocks in the Russell 3000, and the small-cap group consists of stocks stocks around the 2000th largest stocks in the Russell 3000. We list in Table 2.2 the ticker symbols in the HFT database.

[INSERT Table 2.2 HERE]

The trade data set contains all trades on NASDAQ during regular trading hours from January 2008 to December 2009, plus the week of February 22-26, 2010. For each trade, the data set contains the following fields: ticker symbol of stock, date, timestamp, shares, price, buy/sell indicator, and type. The timestamps are in milliseconds. The buy/sell indicator tells whether the trade is initiated by a buyer or a seller. The type variable captures whether traders are HFT or non-HFT firms at both sides of a trade. Accordingly, trades are classified into four types: HH, HN, NH, and NN, where the first letter refers to the liquidity demander and the second to the liquidity supplier. NASDAQ identifies 26 HFT firms based on the trading styles of firms. As noted by Brogaard, Hendershott and Riordan (2014), it is possible that some HFT activities are not identified in this data set. For example, HFT firms that route their orders through large integrated firms, such as

Goldman Sachs and Morgan Stanley, cannot be clearly identified and therefore are excluded. Despite the limitation, this is the only data set that directly classify HFT activities in the U.S. equity markets.

2.3.2 ITCH

TotalView-ITCH is a direct data feed that NASDAQ provides to market participants. It consists of sequenced messages for order addition, cancellation, replacement, and execution for securities traded on NASDAQ. The order level data enables us to construct the limit order book with all price tiers at any given time. The timestamps in the ITCH data are in milliseconds before August 10, 2009, and in nanoseconds afterward.

Messages in ITCH are marked by different symbols for their types. An "A" or an "F" message indicates that a new order has been added into the book. The difference between the two types "A" and "F" is that the market participant identifier (MPID) is associated with the entered order for "F" messages while "A" messages are anonymous. Each add order message has a unique order reference number, so that we may track the life of each order when it is removed from the book.

A limit order leaves the book if it get cancelled, executed, or replaced by another order. We see a "D" message when the entirety of an order is cancelled so there is no remaining shares available. An "X" message indicates that an order is partially deleted. When an order on the book is executed there is an "E" message. An order may be executed in whole or in several parts, so it is possible to see multiple execution messages for the same order reference number. An order is occasionally executed at a different price from the original quote, and these executions are designated with a "C". A "U" message indicates than an order is cancelled and replaced by another one.

The ITCH data also provides a trade message "P" for a match involving a hidden order. This message tells the price and size of a hidden execution and whether the original hidden order is a buy or a sell order.

The ITCH sample covers the period from 2008-2013. We analyze the 120 stocks that are selected for the HFT database. Two stocks, BZ and MAKO, do not exist at the beginning of the sample period. 10 stocks are delisted during our sample period: ABD, BARE, BW, CHTT, GENZ, KNOL, KTII, MFB, PPD, and RVI. To maintain a balanced panel, we exclude these stocks. The final sample consists of 108 stocks with half listed on NASDAQ and the other half on NYSE. Among the 108 stocks there are 39 large-cap, 38 mid-cap, and 31 small-cap stocks.

To retrieve the HFT status in hidden and visible executions separately, we match each trade in the NASDAQ HFT database with the corresponding one in the ITCH data. This enables us to analyze the proportion and characteristics of hidden liquidity supplied by HFT and non-HFT firms respectively.

2.4 Speed

Speed matters for both liquidity suppliers and demanders in the limit order market. For liquidity suppliers investment on speed technology enables them to reduce the latency of their messages being processed by the exchange. As their existing limit orders are more likely to be cancelled when new information arrives, liquidity suppliers get better protection against the risk of being picked off. However, informed liquidity demanders who have higher speed are able to pick off stale limit orders more quickly. Therefore, the overall effect of speed improvement on market quality depends on its relative contribution to market participants who try to avoid being adversely selected with respect to those who trade against stale quotes. In this section we first present the trend of speed improvement over our sample period from 2008-2013. We then study the impact of speed on adverse selection and market liquidity.

2.4.1 Summary statistics

We cannot directly observe the trading speed of market participants. The ITCH data tracks each limit order by a unique reference number so that we are able to determine the life of an order when it leaves the book. The order life provides a proxy for the latency of accessing market by traders. A shorter order life indicates a higher trading speed. We use the 5th percentile life of cancelled orders τ^C as a proxy to measure the latency of liquidity suppliers, and the 5th percentile life of executed orders τ^E as a proxy to measure the latency of liquidity demanders. [2] We plot in Figure 2.2 the 5th percentile order lives across all sampled stocks from 2008-2013.

[INSERT Figure 2.2 HERE]

Across all stocks in the sample, the 5th percentile life of cancelled orders has fallen significantly from 1 millisecond in 2008 to 0.05 milliseconds in 2013. There is a remarkable drop in cancellation latency from 2009-2011. This is consistent with the finding by Gai, Yao, and Ye (2014) of NASDAQ system upgrade and traders' speed improvement in 2010. The cancellation latency stays relatively stable in 2012 and 2013, the last two years of our sample.

The 5th percentile life of executed orders shows a similar decreasing trend with cancelled orders during the period from 2008-2013. It falls significantly from 1 millisecond in 2008 to 0.07 milliseconds in 2013. The execution latency is lower than the cancellation latency in 2009 and 2010, but becomes higher afterward.

The falling trend of trading latency over the period is associated with the rise of HFT. HFT firms make extensive infrastructure investments to increase their speed on order submission, cancellation and execution. Following Hasbrouck and Saar (2013), Hasbrouck (2013), and Gai, Yao, and Ye (2014), we define a high frequency message as any order chain with less than 50 millisecond link. In specific, if a limit order is cancelled or executed in

[2] The results are robust to using the minimal and the 1st percentile of order lives.

less than 50 milliseconds after it is added to the order book we attribute this activity to an HFT firm. Only HFT firms are able to operate with this latency that is more than 20 times faster than the reaction of a human trader.

We first look at HFT participation in the liquidity supplying side. We measure the share of HFT cancellations as the number of cancelled orders with less than 50 millisecond life divided by the total number of cancelled orders. Table 2.3 summarizes the proportion of HFT cancellations.

[INSERT Table 2.3 HERE]

Panel A shows the summary statistics for HFT cancellations with orders pooled across all stocks on each day. We report the results for each year during the period from 2008-2013. HFT liquidity suppliers account for an average of 19% of cancellations in 2008. The proportion falls a little to 16% in 2009 and starts to rise afterward. In 2012 and 2013 nearly 30% of cancellations are made in less than 50 milliseconds. The daily variation that is measured by the standard deviation is smaller in 2012 and 2013 than in 2008, suggesting that the rise of HFT cancellation in more recent years is persistent. The median percentage shows a similar trend as the average from 2008-2013.

We sort stocks into three groups based on market capitalization: large, mid, and small. Panels B-D present the summary statistics for each group respectively. The share of HFT cancellations in each group shows a similar time trend over the period from 2008-2013. However, the proportion is quite different among market capitalization groups. In general, HFT cancellations occur the most frequently in large cap stocks. This suggests that HFT liquidity supplying activities are the most active in large cap stocks. Compared to small cap stocks, mid caps have a larger proportion of HFT cancellations before 2009 but a lower share afterward. For example, in 2011 and 2012 the share of HFT cancellations in small cap stocks is 8% higher than that in mid caps and reaches to the level in large caps. It indicates that HFT liquidity suppliers engage more in small cap stocks in recent years. In addition, large cap stocks often have the smallest daily variation in the share of HFT cancellations, suggesting that HFT liquidity provision is relatively persistent in large caps.

We then analyze HFT aggressive liquidity taking activities by computing the proportion of executed orders with less than 50 millisecond life. The summary statistics are presented in Table 2.4.

[INSERT Table 2.4 HERE]

Panel A reports the proportion of HFT aggressive executions with orders pooled across all stocks on each day. The average share of HFT executions has risen steadily from 17% in 2008 to more than 23% in 2013. The daily variation has fallen in more recent years. The standard deviation in 2013 is less than half of that in 2008. The median HFT share is close to the average and has a similar time trend over the period from 2008-2013.

Panels B-D contain the summary statistics for large, mid, and small cap stocks respectively. The proportion of HFT executions in each group shows a consistently increasing trend from 2008-2013. In every year during the sample period, the percentage of HFT executions is the highest in large cap stocks and the lowest in small cap stocks. However, the differences among the three groups have reduced along time. In 2013 the difference in the average HFT share between large and small cap stocks is less than 1%.

2.4.2 Impact of speed on adverse selection

The academic literature is divided on the effects of speed on adverse selection. We follow Hendershott, Jones, and Menkveld (2011) and measure adverse selection by the signed change in the midpoint of the best bid and ask prices 5 minutes after the time of the initial trade. For trade t in stock i, the cost of adverse selection $AS_{i,t}$ is defined as

$$AS_{i,t} = \frac{x_{i,t}(m_{i,t+5\,\text{min}} - m_{i,t})}{m_{i,t}}, \tag{2.7}$$

where $x_{i,t}$ is an indicator variable for the trade direction and equals 1 for buyer-initiated trades and -1 for seller-initiated trades. The trade direction is explicitly provided in the ITCH data. $m_{i,t}$ is the prevailing midpoint of quotes at the time of the trade, i.e. $m_{i,t} = (p^b_{i,t}+p^a_{i,t})/2$, where $p^b_{i,t}$ and $p^a_{i,t}$ are the best bid and the best ask price respectively. $m_{i,t+5\,\text{min}}$ is the quote midpoint 5 minutes after the trade.

We also use the price impact from the Hasbrouck (1991) model as an alternative measure of adverse selection. We compute the permanent price impact of executions based on the cumulative impulse response of the quoted price to a unit shock in trade. Following Hasbrouck (1991) we use the vector autoregressive (VAR) model that provides a framework to infer the informational content of trades from observed quote revisions. For stock i, let $r_{i,t}$ be the log return in the quote midpoint between trades, i.e. $r_{i,t} = \log m_{i,t} - \log m_{i,t-1}$. The time index t is based on trade events, not clock time. The VAR model with 10-lag order flow dependence is given as

$$r_{i,t} = c_{r,0} + \sum_{j=1}^{10} \alpha_{r,j} r_{i,t-j} + \sum_{j=0}^{10} \beta_{r,j} x_{i,t-j} + \varepsilon_{r,i,t}, \quad (2.8)$$

$$x_{i,t} = c_{x,0} + \sum_{j=1}^{10} \alpha_{x,j} r_{i,t-j} + \sum_{j=1}^{10} \beta_{x,j} x_{i,t-j} + \varepsilon_{x,i,t}. \quad (2.9)$$

We use Hasbrouck's identifying assumption that the current trade has an impact on the current quote, but not vice versa. After estimating the model by seemingly unrelated regressions (SUR), we invert the VAR to obtain the vector moving average (VMA) representation:

$$\begin{bmatrix} r_{i,t} \\ x_{i,t} \end{bmatrix} = \begin{bmatrix} A(L) & B(L) \\ C(L) & D(L) \end{bmatrix} \begin{bmatrix} \varepsilon_{r,i,t} \\ \varepsilon_{x,i,t} \end{bmatrix} \quad (2.10)$$

where $A(L), \ldots, D(L)$ are lagged operators. The VMA representation provides an easy way to present the cumulative impulse response of price to a unit shock in trade. We compute the cumulative impulse response up to 50 transactions after the shock, which is equal to $\sum_{j=0}^{50} B_j$. The model is estimated for each stock on each day.

We report in Table 2.5 the average 5-minute adverse selection cost and price impact in each year from 2008-2013. Both measures are expressed in basis points (bps).

[INSERT Table 2.5 HERE]

Each of the adverse selection measures exhibits a downtrend over the period from 2008-2013, though the decline is not monotonic. Panel A presents the measures across all sample stocks. The average adverse selection cost is 7.5 bps in 2008. It rises to 8.2 bps in 2009, falls significantly to 5.5 bps in 2010, and then stays relatively stable afterwards. The price impact has a similar trend during the sample period.

Panels B-D contain the results for each of three market capitalization groups respectively. The adverse selection cost and price impact are the lowest in large cap stocks and the highest in small cap stocks. The time trend in each group is similar to the full sample. In particular, both measures drop significantly in 2010 and stay relatively flat afterwards.

We then analyze the impact of speed on adverse selection using econometric models. Given that both liquidity suppliers and demanders may use high speed technology to trade, it is their relative speed that matters. For stock i on day d, we define the relative speed of liquidity suppliers to demanders as

$$\text{rel\_spd}_{i,d} = \tau^E_{i,d} / \tau^C_{i,d}. \tag{2.11}$$

where $\tau^E_{i,d}$ and $\tau^C_{i,d}$ are the 5th percentile life of executed and cancelled orders respectively.

We regress each adverse selection measure on the relative speed variable and control for other effects. Since adverse selection varies across stocks and along time, we run panel data regression with stock and year fixed effects as follows.

$$Y_{i,d} = \beta_0 + \beta_1 \text{rel\_spd}_{i,d} + \text{controls}_{i,d} + \sum_{k=1}^{107} \phi_k \text{stock}_{i,d} + \sum_{j=1}^{5} \kappa_j \text{year}_{i,d} + \varepsilon_{i,d}, \tag{2.12}$$

where the dependent variable $Y_{i,d}$ is the average adverse selection cost or the price impact for stock i on day d. The control variables include volatility, log of trading volume, price, and log of market capitalization. $\text{stock}_{i,d}$ is a set of dummy variables that are used to capture unobserved stock fixed effects. We also include year dummy variables $\text{year}_{i,d}$ to control for time effects. The regression results are reported in Table 2.6.

[INSERT Table 2.6 HERE]

We find that both the 5-minute adverse selection cost and the price impact are negatively associated with the relative speed. This suggests that adverse selection costs decline when liquidity suppliers have superior speed advantage. The results confirm the first prediction *H1* from our theoretical model. The coefficients of relative speed in both regressions are statistically significant at the 1% level. Quantitatively, an 10% increase in the speed of liquidity suppliers to demanders is associated with a 1.48 bps decrease in adverse selection costs, and a 0.80 bps decline in price impact. Based on the fact that the average estimate across all stocks over the sample period is 6.22 bps for adverse selection costs and 5.78 bps for price impact, the results reflect that adverse selection costs decrease by 24% and price impact by 14%.

2.4.3 Impact of speed on spreads

Spreads are important indicators of market liquidity. Our theoretical model predicts that fast liquidity suppliers charge narrow spreads because of low adverse selection risk. Therefore, spreads would narrow when the relative speed of liquidity suppliers to demanders increases. We test this hypothesis using two spread measures: quoted half-spreads and effective half-spreads. Both measures are normalized by the prevailing quote midpoint, and expressed in basis points.

The quoted half-spread is half of the difference between the best bid and the best ask price. The effective half-spread is the signed difference between the actual trade price and the midpoint of the best bid and ask when the trade takes place. For trade t in stock i, we calculate the normalized effective spread $ES_{i,t}$ as

$$ES_{i,t} = \frac{x_{i,t}(p_{i,t} - m_{i,t})}{m_{i,t}}, \tag{2.13}$$

where $x_{i,t}$ is the buy/sell indicator variable that equals 1 for buys and -1 for sells, $p_{i,t}$ is the trade price, and $m_{i,t}$ is the prevailing midpoint of the best bid and ask quotes.

We report in Table 2.7 average quoted half-spreads and effective half-spreads in each year from 2008-2013.

[INSERT Table 2.7 HERE]

The results show that both spread measures have similar time trend with adverse selection. Panel A presents the measures across all sample stocks. The average quoted half-spread is 8.6 bps in 2008. It rises to 9.4 bps in 2009, falls significantly to 6.1 bps in 2010, and then stays relatively stable afterwards. The effective half-spreads exhibit similar trend during the period from 2008-2013.

Panels B-D contain the results for each of three market capitalization groups respectively. Both quoted and effective spreads are the lowest in large cap stocks and the highest in small cap stocks. The time trend in each group is similar to the full sample.

We then analyze the impact of speed on spreads by running panel regression with stock and year fixed effects as specified in (2.12). The dependent variable $Y_{i,d}$ is respectively the average quoted and effective half-spreads for stock i on day d. We report the regression results in Table 2.8.

[INSERT Table 2.8 HERE]

The results validate the second hypothesis *H2* that spreads narrow when liquidity suppliers have superior speed advantage. We find that both quoted and effective spreads are negatively associated with the relative speed. The coefficients of relative speed in both regressions are statistically significant at the 0.1% level. Quantitatively, an 10% increase in the speed of liquidity suppliers to demanders is associated with a 1.03 bps decline in quoted half-spreads, and a 0.77 bps decline in effective half-spreads. Based on the fact that the average estimate across all stocks over the sample period is 7.11 bps for quoted half-spreads and 5.62 bps for effective half-spreads, the results suggest that both quoted and effective half-spreads decrease by 14%.

2.5 Hidden Orders

Our theoretical model predicts that limit orders of slow liquidity suppliers are less likely to be executed in the presence of fast liquidity suppliers. Therefore, slow liquidity suppliers face the choice of exiting the market or using some other strategy to compete with fast liquidity suppliers. Hidden orders have lower adverse selection risk and thus provide a tool for slow liquidity suppliers to improve their competitive spreads. Fast liquidity suppliers are less likely to use hidden orders because they are able to manage the adverse selection risk by revising limit orders quickly and hidden orders reduce their trading volume. In this section, we first provide summary statistics of hidden order executions on NASDAQ based on the ITCH data. We then compare adverse selection costs and spreads between hidden and visible orders. Finally, we study whether HFT or non-HFT firms use hidden orders more frequently.

2.5.1 Summary statistics

Our analysis of hidden liquidity is based on the ITCH data that provides trade messages for executions involving hidden orders. Each of these messages tells us the price of the hidden order and whether it is a buy or sell order. We calculate the proportion of hidden executions as the number of executions involving hidden orders divided by the total number of executions during regular trading hours on each day. Table 2.9 presents summary statistics of hidden executions on NASDAQ from 2008-2013.

[INSERT Table 2.9 HERE]

Panel A summarizes the share of hidden executions across all stocks in each year over the sample period. The average share of hidden executions is 17% in 2008, and falls below 11% during the period from 2009-2011. It then gradually rises to 15% in 2013. The daily variation in the proportion of hidden executions is relatively stable from 2008-2013. Our results are consistent with the finding by Yao (2012) who analyzes all NASDAQ-listed common stocks from 2010-2011 and reports that 11% of trades involve hidden orders executions.

We report summary statistics for large, mid, and small cap stocks respectively in Panels B-D. The average percentage of hidden executions shows a similar time trend in each group. Hidden executions occur more frequently in small cap than large cap stocks from 2008-2010; however, their difference becomes smaller from 2011-2013. In 2013 large and small cap stocks have the same share of hidden executions.

2.5.2 Adverse selection of hidden orders

We test the hypothesis that hidden orders bear lower adverse selection risk than visible orders. Consistent with Section 2.4.2, we use two measures for adverse selection: 5-minute adverse selection costs and price impact from the VAR model. We use (2.7) to compute adverse selection costs of visible and hidden executions respectively. For price impact, we extend the VAR model in (2.8) and (2.9) to a three-equation model that separate visible and hidden executions.

$$r_{i,t} = c_{r,0} + \sum_{j=1}^{10} \alpha_{r,j} r_{i,t-j} + \sum_{j=0}^{10} \beta_{r,j} x^H_{i,t-j} + \sum_{j=0}^{10} \gamma_{r,j} x^V_{i,t-j} + \varepsilon_{r,i,t}, \tag{2.14}$$

$$x^H_{i,t} = c_{h,0} + \sum_{j=1}^{10} \alpha_{h,j} r_{i,t-j} + \sum_{j=1}^{10} \beta_{h,j} x^H_{i,t-j} + \sum_{j=1}^{10} \gamma_{h,j} x^V_{i,t-j} + \varepsilon_{h,i,t}, \tag{2.15}$$

$$x^V_{i,t} = c_{v,0} + \sum_{j=1}^{10} \alpha_{v,j} r_{i,t-j} + \sum_{j=1}^{10} \beta_{v,j} x^H_{i,t-j} + \sum_{j=1}^{10} \gamma_{v,j} x^V_{i,t-j} + \varepsilon_{v,i,t}, \tag{2.16}$$

where $r_{i,t}$ is the percentage change in the midpoint of the best bid and ask prices, $x^H_{i,t}$ is the buy/sell indicator for hidden executions, and $x^V_{i,t}$ for visible executions. The indicator variables equal 1 for buys and -1 for sells. If a hidden execution occurs in the absence of a visible one, $x^V_{i,t}$ is set to 0, and vice versa if a visible execution occurs in the absence of a hidden one, $x^H_{i,t}$ is set to 0. For both hidden and visible executions, the current trade has an impact on the current quote, but not vice versa. We then invert the VAR to obtain the VMA representation:

$$\begin{bmatrix} r_{i,t} \\ x^H_{i,t} \\ x^V_{i,t} \end{bmatrix} = \begin{bmatrix} A(L) & B(L) & C(L) \\ D(L) & E(L) & F(L) \\ G(L) & H(L) & I(L) \end{bmatrix} \begin{bmatrix} \varepsilon_{r,i,t} \\ \varepsilon_{h,i,t} \\ \varepsilon_{v,i,t} \end{bmatrix} \tag{2.17}$$

where $A(L)$, ..., $I(L)$ are lagged operators. We compute the price impact of hidden and visible executions respectively based on the cumulative impulse response of the quoted price to a unit shock in trade. We calculate the cumulative impulse response up to 50 transactions after the shock. It is equal to $\sum_{j=0}^{50} B_j$ for hidden executions and $\sum_{j=0}^{50} C_j$ for visible executions.

Over our sample period from 2008-2013, the average adverse selection cost across all stocks is 1.35 bps for hidden and 7.19 bps for visible executions. The average price impact is 4.70 bps for hidden and 5.92 bps for visible executions. To capture whether hidden and visible executions differ significantly in adverse selection, we estimate the following regression model.

$$Y_{i,d} = \beta_0 + \beta_1 \text{hidden}_{i,d} + \text{controls}_{i,d} + \sum_{k=1}^{107} \phi_k \text{stock}_{i,d} + \sum_{j=1}^{5} \kappa_j \text{year}_{i,d} + \varepsilon_{i,d}, \tag{2.18}$$

where $Y_{i,d}$ is the average adverse selection cost or the price impact for stock i on day d, $\text{hidden}_{i,d}$ is a dummy variable that equals 1 for hidden executions and 0 for visible executions, and the control variables include volatility, log of trading volume, price, and log of market capitalization. Similar to the regression model in the speed section, we also include stock and year fixed effects as adverse selection varies across stocks and along time. $\text{stock}_{i,d}$ is a set of stock dummy variables and $\text{year}_{i,d}$ is a set of year dummy variables. We report the regression results in Table 2.10.

[INSERT Table 2.10 HERE]

The estimated coefficient of the *hidden* variable captures the differential adverse selection between hidden and visible executions. We find that hidden orders bear statistically significantly lower adverse selection risk than visible ones. The 5-minute adverse selection cost of hidden executions is 5.84 bps lower than visible executions, and the price impact of hidden executions is 1.22 bps weaker. The results validate the theoretical prediction that hidden orders have lower adverse selection risk than visible orders.

2.5.3 Spreads of hidden orders

Given that hidden orders are less likely to be adversely selected, liquidity suppliers may use hidden orders to improve their quotes and charge a narrower spread. We test this hypothesis using two spread measures: quoted half-spreads and effective half-spreads. Both measures are normalized by the prevailing quote midpoint, and expressed in basis points. We compute these two metrics for hidden and visible executions respectively.

Over the sample period the average quoted half-spread across all stocks is 6.30 bps for hidden and 6.87 bps for visible executions. The average effective half-spread is 1.92 bps for hidden and 6.37 bps for visible executions. To determine the significance of the spread difference between hidden and visible orders, we estimate the regression model that is specified in (2.18). The results are presented in Table 2.11.

[INSERT Table 2.11 HERE]

We find that hidden orders have significantly narrower spreads than visible ones. The quoted half-spread of hidden orders is 0.58 bps narrower than visible orders, and the effective half-spread of hidden orders is 4.45 bps narrower. The results suggest that liquidity suppliers use hidden orders to improve their quote prices.

2.5.4 Who uses hidden orders more frequently?

We find in our model that slow liquidity suppliers are worse off in the presence of fast ones because their limit orders have lower execution probability. Hidden orders provide a way for slow liquidity suppliers to reduce the adverse selection risk so that their competitive quotes have narrow spreads. Fast liquidity suppliers are less likely to use hidden orders because it reduces their profits from trading with uninformed liquidity demanders. Based on this reasoning, we conjecture that non-HFT firms submit hidden orders more frequently than HFT firms. We test this hypothesis using the NASDAQ HFT data set and the ITCH data. The HFT data set tells us whether traders are HFT or non-HFT firms at both sides of a trade. However, it does not include information on whether a trade involves a hidden or a visible order. By matching ITCH with the HFT data trade by trade, we obtain the HFT status at both sides in every hidden execution.

We analyze whether HFT or non-HFT firms supply more hidden liquidity. Among all executions HFT firms provide liquidity in HH and NH trades, while non-HFT firms in HN and NN trades. Therefore, the share of liquidity supplied by non-HFT firms $\omega^{\text{non-HFT}}$ is calculated as the number of HN and NN trades divided by the total number of trades during regular trading hours on each day, i.e.

$$\omega^{\text{non-HFT}} = \frac{\text{HN} + \text{NN}}{\text{HH} + \text{HN} + \text{NH} + \text{NN}}. \tag{2.19}$$

Accordingly, the HFT share is $\omega^{\text{HFT}} = 1 - \omega^{\text{non-HFT}}$. We compare the share of liquidity provided by HFT and non-HFT firms in hidden and visible executions respectively and present the results in Table 2.12.

[INSERT Table 2.12 HERE]

The results show that 70% of liquidity in hidden executions is supplied by non-HFT firms across all stocks. In contrast, HFT firms supply more than 50% liquidity in visible executions. The difference between non-HFT and HFT shares is much more significant in hidden executions than in visible executions. The results suggest that non-HFT liquidity suppliers use hidden orders to reduce their adverse selection risk. In contrast, HFT liquidity suppliers who have speed advantage may manage the adverse selection risk by updating their limit orders quickly, and therefore use hidden orders less frequently.

2.6 Conclusion

We study in this paper two strategies that may reduce the adverse selection risk: investment in speed technology and use of hidden orders. Liquidity suppliers who use high speed technology are able to update their orders quickly and therefore face lower adverse selection risk. Our theoretical model reveals that adverse selection costs fall and spreads narrow when liquidity suppliers become faster. However, slow liquidity suppliers are worse off in the presence of fast ones because their limit orders have wider competitive spreads and therefore lower execution probabilities. The model predicts that hidden orders provide a tool for slow liquidity suppliers to reduce the adverse selection risk and compete with fast liquidity suppliers.

Our empirical results on speed and hidden order usage on NASDAQ validate the predictions from the model. We find that both adverse selection costs and price impact decline and spreads narrow, when liquidity suppliers have superior speed advantage. We also find that non-HFT firms account for 70% of liquidity provision in hidden executions, while HFT firms supply more liquidity in visible executions. Hidden orders have significantly lower adverse selection risk and narrower spreads than visible orders. Our findings suggest that speed and hidden orders have equivalent roles for liquidity suppliers to reduce the adverse selection risk.

Table 2.1: Equilibria for Fast and Slow Liquidity Suppliers

This table summarizes the equilibria where both fast and slow liquidity suppliers have zero expected profits. Each type of liquidity suppliers may choose between visible and hidden orders. $s^*_{F,V}$ and $s^*_{S,V}$ are visible competitive half spreads for fast and slow liquidity suppliers respectively, and $s^*_{F,H}$ and $s^*_{S,H}$ are hidden competitive half spreads. σ measures the volatility of the fundamental value for a risky security. γ represents the probability of news arrival that affects the fundamental value. α is the probability that liquidity suppliers make correct inference on whether a liquidity demander is a noise trader or a informed trader. λ_F (λ_S) is the probability that fast (slow) liquidity suppliers are able to cancel their limit orders on the arrival of informed traders. θ (η) is the probability that informed (noise) traders still choose to trade in the absence of visible limit orders in the order book.

Fast Liquidity Suppliers	Slow Liquidity Suppliers: Visible Order	Slow Liquidity Suppliers: Hidden Order
Visible Order	$s^*_{F,V} < s^*_{S,V}$, where $s^*_{F,V} = \frac{\gamma(1-\alpha\lambda_F)}{(1-\gamma)\alpha+\gamma(1-\alpha\lambda_F)}\sigma$ $s^*_{S,V} = \frac{\gamma\lambda_F(1-\alpha\lambda_S)}{(1-\gamma)(1-\alpha)+\gamma\lambda_F(1-\alpha\lambda_S)}\sigma$	$s^*_{F,V} \leqslant s^*_{S,H}$, where $s^*_{F,V} = \frac{\gamma(1-\alpha\lambda_F)}{(1-\gamma)\alpha+\gamma(1-\alpha\lambda_F)}\sigma$ $s^*_{S,H} = \frac{\gamma\lambda_F\theta(1-\alpha\lambda_S)}{(1-\gamma)(1-\alpha)\eta+\gamma\lambda_F\theta(1-\alpha\lambda_S)}\sigma$ if $\theta \geqslant \frac{(1-\alpha)(1-\alpha\lambda_F)}{\alpha\lambda_F(1-\alpha\lambda_S)}\eta$ $s^*_{S,H} < s^*_{F,V}$, where $s^*_{F,V} = \frac{\gamma(1-\alpha\lambda_F)}{(1-\gamma)\alpha+\gamma(1-\alpha\lambda_F)}\sigma$ $s^*_{S,H} = \frac{\gamma(1-\alpha\lambda_F(1-\theta))(1-\alpha\lambda_S)}{(1-\gamma)\alpha(1-\alpha\lambda_F(1-\eta))+\gamma(1-\alpha\lambda_F(1-\theta))(1-\alpha\lambda_S)}\sigma$ if $\theta < \frac{(\lambda_S-\lambda_F(1-\eta))(1-\alpha\lambda_F)}{\lambda_F(1-\alpha\lambda_S)}$
Hidden Order	$s^*_{F,H} < s^*_{S,V}$, where $s^*_{F,H} = \frac{\gamma(1-\alpha\lambda_S(1-\theta))(1-\alpha\lambda_F)}{(1-\gamma)\alpha(1-\alpha\lambda_S(1-\eta))+\gamma(1-\alpha\lambda_S(1-\theta))(1-\alpha\lambda_F)}\sigma$ $s^*_{S,V} = \frac{\gamma(1-\alpha\lambda_S)}{(1-\gamma)\alpha+\gamma(1-\alpha\lambda_S)}\sigma$	$s^*_{F,H} < s^*_{S,H}$, where $s^*_{F,H} = \frac{\gamma\theta(1-\alpha\lambda_F)}{(1-\gamma)\alpha\eta+\gamma\theta(1-\alpha\lambda_F)}\sigma$ $s^*_{S,H} = \frac{\gamma\theta(1-\alpha\lambda_S)}{(1-\gamma)\alpha\eta+\gamma\theta(1-\alpha\lambda_S)}\sigma$

Table 2.2: Stocks in NASDAQ HFT Database

This table lists the 120 ticker symbols in NASDAQ HFT database. Stocks are stratified by listing exchanges and market capitalization. 60 stocks are listed on NASDAQ and the other 60 on NYSE. Stocks are sorted into three market capitalization groups: large, mid, and small. Each group contains 40 stocks. Stocks in the large-cap group are selected from the largest market capitalization stocks. The mid-cap group is composed of stocks around the 1000th largest stocks in the Russell 3000, and the small-cap group consists of stocks around the 2000th largest stocks in the Russell 3000.

Large-cap		Mid-cap		Small-cap	
AA	EBAY	AMED	FCN	ABD	IMGN
AAPL	ESRX	ARCC	FL	AINV	IPAR
ADBE	GE	AYI	FMER	ANGO	KNOL
AGN	GENZ	BARE	FULT	APOG	KTII
AMAT	GILD	BRE	GAS	AZZ	MAKO
AMGN	GLW	BXS	ISIL	BAS	MDCO
AMZN	GOOG	CBT	JKHY	BW	MFB
AXP	GPS	CCO	LANC	BZ	MIG
BHI	HON	CETV	LECO	CBEY	MOD
BIIB	HPQ	CHTT	LPNT	CBZ	MRTN
BRCM	INTC	CKH	LSTR	CDR	MXWL
CB	ISRG	CNQR	MANT	CPSI	NC
CELG	KMB	COO	MELI	CRVL	NXTM
CMCSA	KR	CPWR	NSR	CTRN	PBH
COST	MMM	CR	NUS	DCOM	PPD
CSCO	MOS	CRI	PNY	DK	RIGL
CTSH	PFE	CSE	PTP	EBF	ROCK
DELL	PG	CSL	ROC	FFIC	ROG
DIS	PNC	ERIE	SF	FPO	RVI
DOW	SWN	EWBC	SFG	FRED	SJW

Table 2.3: Percentage of HFT Cancellations

This table reports mean, standard deviation (SD), the 10th (P10) percentile, median, and the 90th percentile (P90) of the proportion of HFT cancellations. We define an HFT cancellation as any limit order that is cancelled in less than 50 milliseconds based on the NASDAQ TotalView-ITCH data. The percentage is calculated as the number of HFT cancellations divided by the total number of order cancellations during regular trading hours on each day. The full sample consists of 108 stocks over the period from 2008-2013. Stocks are sorted into three market capitalization groups: large, mid, and small. Panel B-D present summary statistics for each group.

Year	2008	2009	2010	2011	2012	2013
Panel A: Full sample						
Mean	19.05%	16.23%	17.50%	22.58%	29.80%	29.46%
SD	3.46%	2.00%	2.00%	3.71%	2.38%	3.40%
P10	14.49%	13.59%	15.48%	18.38%	28.38%	25.87%
Median	19.47%	15.90%	16.95%	21.64%	29.72%	28.64%
P90	22.53%	18.12%	19.42%	28.18%	31.66%	34.83%
Panel B: Large-cap stocks						
Mean	19.60%	17.00%	17.64%	23.58%	31.11%	31.58%
SD	3.40%	2.23%	2.35%	4.09%	2.25%	3.30%
P10	15.32%	14.15%	15.10%	19.07%	29.42%	27.55%
Median	19.89%	16.70%	17.07%	22.65%	31.06%	31.09%
P90	22.91%	19.01%	20.44%	29.65%	33.45%	36.60%
Panel C: Mid-cap stocks						
Mean	15.81%	12.85%	16.01%	15.91%	21.48%	19.40%
SD	5.52%	2.47%	2.40%	2.15%	6.16%	3.29%
P10	9.25%	9.80%	13.25%	13.24%	16.72%	14.75%
Median	14.91%	12.72%	15.55%	15.25%	20.45%	19.60%
P90	23.50%	15.77%	18.10%	19.10%	29.04%	22.81%
Panel D: Small-cap stocks						
Mean	14.07%	11.36%	18.44%	23.90%	29.69%	20.13%
SD	4.41%	3.16%	4.91%	8.42%	6.32%	3.46%
P10	10.21%	8.44%	13.79%	15.70%	22.80%	16.08%
Median	12.48%	10.23%	16.90%	23.41%	28.32%	19.96%
P90	19.62%	15.41%	22.80%	36.16%	37.75%	24.41%

Table 2.4: Percentage of HFT Executions

This table reports mean, standard deviation (SD), the 10th (P10) percentile, median, and the 90th percentile (P90) of the proportion of HFT executions. We consider an execution as an HFT activity if a limit order is executed in less than 50 milliseconds based on the NASDAQ TotalView-ITCH data. The percentage is calculated as the number of HFT executions divided by the total number of executions during regular trading hours on each day. The full sample consists of 108 stocks over the period from 2008-2013. Stocks are sorted into three market capitalization groups: large, mid, and small. Panels B-D present summary statistics for each group.

Year	2008	2009	2010	2011	2012	2013
Panel A: Full sample						
Mean	16.92%	18.09%	21.95%	23.36%	22.49%	23.63%
SD	1.76%	0.91%	1.69%	1.88%	0.89%	0.84%
P10	14.29%	17.27%	20.11%	21.39%	21.23%	22.72%
Median	17.42%	17.93%	22.09%	22.88%	22.63%	23.64%
P90	18.60%	19.08%	24.11%	26.18%	23.35%	24.51%
Panel B: Large-cap stocks						
Mean	17.22%	18.33%	22.62%	23.93%	22.84%	23.71%
SD	1.85%	0.98%	1.77%	1.94%	0.89%	0.94%
P10	14.48%	17.54%	20.82%	21.98%	21.66%	22.82%
Median	17.82%	18.15%	22.71%	23.31%	23.03%	23.64%
P90	19.07%	19.51%	25.06%	26.65%	23.64%	24.81%
Panel C: Mid-cap stocks						
Mean	14.57%	17.12%	17.19%	19.50%	20.28%	23.16%
SD	1.53%	1.53%	1.08%	2.00%	1.57%	1.22%
P10	12.89%	15.28%	15.82%	17.47%	18.59%	21.64%
Median	13.97%	17.07%	17.07%	18.90%	20.20%	23.06%
P90	16.48%	19.05%	18.58%	22.69%	22.19%	24.55%
Panel D: Small-cap stocks						
Mean	11.44%	10.93%	14.06%	17.35%	18.17%	22.74%
SD	1.21%	1.12%	1.53%	3.06%	1.45%	2.01%
P10	9.62%	10.01%	11.88%	13.40%	16.74%	20.58%
Median	11.50%	10.70%	14.04%	17.68%	18.38%	22.59%
P90	12.56%	12.42%	15.72%	21.29%	19.88%	25.48%

Table 2.5: Adverse Selection Measures

This table reports the average adverse selection cost and price impact for a sample of 108 stocks over the period from 2008-2013. The adverse selection cost is defined as the signed change in the midpoint of best bid and ask prices 5 minutes after the time of the trade, and calculated using (2.7). The price impact is computed based on the Hasbrouck (1991) vector autoregressive (VAR) model with 10-lag order flow dependence as specified in (2.8) and (2.9). The cumulative price impact captures the permanent price change in response to a unit shock in trade. Both measures are reported in basis points. The stocks are sorted into three market capitalization groups: large, mid, and small. Panels B-D present the results for each group respectively.

Year	2008	2009	2010	2011	2012	2013
Panel A: Full sample						
Adverse selection cost	7.518	8.173	5.533	5.600	5.432	5.078
Price impact	7.191	7.927	4.869	5.246	4.754	4.557
Panel B: Large-cap stocks						
Adverse selection cost	2.494	2.868	2.384	2.544	2.321	2.185
Price impact	2.507	2.478	1.954	2.127	1.997	2.012
Panel C: Mid-cap stocks						
Adverse selection cost	7.116	6.781	4.803	5.252	5.218	4.707
Price impact	7.203	6.788	4.478	5.010	4.773	4.700
Panel D: Small-cap stocks						
Adverse selection cost	14.330	16.554	10.391	9.870	9.608	9.172
Price impact	13.369	16.468	9.461	10.064	8.886	8.170

Table 2.6: Impact of Speed on Adverse Selection

This table presents estimates from regressions of adverse selection measures on the relative speed of liquidity suppliers to demanders. The sample consists of 108 stocks over the period from 2008-2013. We run panel data regressions with the stock and year fixed effect: $Y_{i,d} = \beta_0 + \beta_1 \text{rel\_spd}_{i,d} + \text{controls}_{i,d} + \sum_{k=1}^{107} \phi_k \text{stock}_{i,d} + \sum_{j=1}^{5} \kappa_j \text{year}_{i,d} + \varepsilon_{i,d}$. The dependent variable $Y_{i,d}$ is the average 5-minute adverse selection cost for stock i on day d in the first column, and the price impact based on the Hasbrouck (1991) model in the second column. The relative speed of liquidity suppliers to demanders, $\text{rel\_spd}_{i,d}$, is defined as the ratio of the 5th percentile life of executed orders to that of cancelled orders for stock i on day d. The control variables include volatility, log of trading volume (logvol), price, and log of market capitalization (logcap). $\text{stock}_{i,d}$ is a set of dummy variables that are used to capture unobserved stock fixed effects, and $\text{year}_{i,d}$ is a set of year dummy variables that are used to control for time effects. Numbers in parentheses are t-statistic.

	Adverse selection cost	Price impact
rel_spd	-0.148	-0.080
	(-11.68)	(-2.57)
volatility	52.266	51.112
	(42.29)	(70.28)
logvol	0.541	0.450
	(5.63)	(9.34)
price	0.011	0.007
	(5.46)	(7.54)
logcap	-2.397	-1.796
	(-12.67)	(-19.56)
constant	31.855	24.063
	(10.42)	(15.81)
R^2	0.203	0.424
Stock Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes

Table 2.7: Spread Measures

This table reports the average of quoted and effective half-spreads respectively for a sample of 108 stocks over the period from 2008-2013. The quoted half-spread is calculated as the half difference between the best bid and ask prices. The effective half-spread is the signed difference between the actual trade price and the midpoint of the best bid and ask prices when the trade takes place, as specified in (2.13). Both spread measures are normalized by the prevailing midpoint price and reported in basis points. The stocks are sorted into three market capitalization groups: large, mid, and small. Panels B-D present the results for each group.

Year	2008	2009	2010	2011	2012	2013
Panel A: Full sample						
Quoted half-spread	8.614	9.431	6.131	5.978	6.143	6.386
Effective half-spread	6.460	7.291	4.916	4.961	5.052	5.052
Panel B: Large-cap stocks						
Quoted half-spread	2.912	2.871	2.003	2.041	1.979	1.992
Effective half-spread	2.194	2.372	1.769	1.793	1.722	1.670
Panel C: Mid-cap stocks						
Quoted half-spread	8.492	7.823	4.640	4.996	5.155	5.471
Effective half-spread	6.132	5.950	3.842	4.253	4.343	4.448
Panel D: Small-cap stocks						
Quoted half-spread	15.936	19.653	13.154	12.135	12.594	13.037
Effective half-spread	12.229	15.123	10.192	9.815	10.111	10.050

Table 2.8: Impact of Speed on Spreads

This table presents estimates from regressions of spread measures on the relative speed of liquidity suppliers to demanders. The sample consists of 108 stocks over the period from 2008-2013. We run panel data regressions with the stock and year fixed effect: $Y_{i,d} = \beta_0 + \beta_1 \text{rel\_spd}_{i,d} + \text{controls}_{i,d} + \sum_{k=1}^{107} \phi_k \text{stock}_{i,d} + \sum_{j=1}^{5} \kappa_j \text{year}_{i,d} + \varepsilon_{i,d}$. The dependent variable $Y_{i,d}$ is the average of quoted half-spread for stock i on day d in the first column, and the average of effective half-spread in the second column. The relative speed of liquidity suppliers to demanders, $\text{rel\_spd}_{i,d}$, is defined as the ratio of the 5th percentile life of executed orders to that of cancelled orders for stock i on day d. The control variables include volatility, log of trading volume (logvol), price, and log of market capitalization (logcap). $\text{stock}_{i,d}$ is a set of dummy variables that are used to capture unobserved stock fixed effects, and $\text{year}_{i,d}$ is a set of year dummy variables that are used to control for time effects. Numbers in parentheses are t-statistic.

	Quoted half-spread	Effective half-spread
rel_spd	-0.103	-0.077
	(-3.97)	(-3.99)
volatility	61.620	48.728
	(99.45)	(105.62)
logvol	0.132	0.023
	(3.13)	(0.73)
price	0.008	0.006
	(9.38)	(9.44)
logcap	-1.837	-1.402
	(-22.31)	(-22.87)
constant	28.886	23.143
	(21.32)	(22.95)
R^2	0.579	0.594
Stock Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes

Table 2.9: Proportion of Hidden Executions on NASDAQ

This table reports mean, standard deviation (SD), the 10th (P10) percentile, median, and the 90th percentile (P90) of the proportion of hidden executions, based on the NASDAQ TotalView-ITCH data. The percentage is calculated as the number of executions matched with hidden orders divided by the total number of executions during regular trading hours on each day. The full sample consists of 108 stocks over the period from 2008-2013. Stocks are sorted into three market capitalization groups: large, mid, and small. Panels B-D present summary statistics for each group.

Year	2008	2009	2010	2011	2012	2013
Panel A: Full sample						
Mean	16.55%	10.53%	10.01%	10.70%	12.21%	14.72%
SD	2.32%	2.05%	2.00%	0.81%	1.74%	0.85%
P10	13.63%	8.62%	7.47%	9.85%	9.74%	13.86%
Median	16.37%	9.85%	9.61%	10.63%	12.26%	14.76%
P90	19.59%	13.73%	12.49%	11.52%	14.18%	15.37%
Panel B: Large-cap stocks						
Mean	16.44%	10.11%	9.85%	10.86%	12.33%	14.78%
SD	2.43%	2.05%	2.09%	0.86%	1.79%	0.88%
P10	13.37%	8.31%	7.28%	9.82%	9.71%	13.93%
Median	16.20%	9.34%	9.39%	10.77%	12.49%	14.74%
P90	19.72%	13.35%	12.64%	11.85%	14.24%	15.71%
Panel C: Mid-cap stocks						
Mean	17.69%	13.75%	10.76%	9.00%	10.70%	14.15%
SD	2.20%	2.85%	2.47%	1.31%	2.05%	1.60%
P10	14.86%	10.95%	8.66%	7.71%	8.50%	12.15%
Median	18.25%	12.93%	10.20%	9.00%	10.21%	14.39%
P90	20.11%	17.40%	12.02%	10.09%	13.48%	16.33%
Panel D: Small-cap stocks						
Mean	18.01%	15.34%	12.43%	10.89%	12.75%	14.78%
SD	1.40%	2.18%	1.83%	2.08%	1.90%	2.48%
P10	16.32%	12.64%	9.64%	9.04%	10.43%	11.70%
Median	18.21%	15.21%	12.90%	10.54%	12.63%	14.47%
P90	19.36%	17.93%	14.33%	13.42%	14.51%	17.96%

Table 2.10: Adverse Selection of Hidden Orders

This table presents the regression results for difference in adverse selection between hidden and visible executions. The sample consists of 108 stocks over the period from 2008-2013. We run the following regression with the stock and year fixed effect: $Y_{i,d} = \beta_0 + \beta_1 \text{hidden}_{i,d} + \text{controls}_{i,d} + \sum_{k=1}^{107} \phi_k \text{stock}_{i,d} + \sum_{j=1}^{5} \kappa_j \text{year}_{i,d} + \varepsilon_{i,d}$. The dependent variable $Y_{i,d}$ is the average 5-minute adverse selection cost for stock i on day d in the first column, and the price impact based on the Hasbrouck (1991) model in the second column. The variable $\text{hidden}_{i,d}$ is a dummy that equals 1 for hidden executions and 0 for visible executions. The control variables include volatility, log of trading volume (logvol), price, and log of market capitalization (logcap). $\text{stock}_{i,d}$ is a set of dummy variables that are used to capture unobserved stock fixed effects, and $\text{year}_{i,d}$ is a set of year dummy variables that are used to control for time effects. Numbers in parentheses are t-statistic.

	Adverse selection cost	Price impact
hidden	-5.841	-1.223
	(-51.95)	(-7.03)
volatility	38.657	54.239
	(29.94)	(24.83)
logvol	0.227	1.226
	(2.15)	(7.22)
price	0.006	0.013
	(2.68)	(3.58)
logcap	-1.457	-1.705
	(-6.90)	(-5.19)
constant	23.588	12.318
	(6.97)	(2.31)
R^2	0.228	0.111
Stock Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes

Table 2.11: Spreads of Hidden Orders

This table presents the regression results for difference in spreads between hidden and visible executions. The sample consists of 108 stocks over the period from 2008-2013. We run the following regression with the stock and year fixed effect: $Y_{i,d} = \beta_0 + \beta_1 \text{hidden}_{i,d} + \text{controls}_{i,d} + \sum_{k=1}^{107} \phi_k \text{stock}_{i,d} + \sum_{j=1}^{5} \kappa_j \text{year}_{i,d} + \varepsilon_{i,d}$. The dependent variable $Y_{i,d}$ is the average of quoted half-spread for stock i on day d in the first column, and the average of effective half-spread in the second column. The variable $\text{hidden}_{i,d}$ is a dummy that equals 1 for hidden executions and 0 for visible executions. The control variables include volatility, log of trading volume (logvol), price, and log of market capitalization (logcap). $\text{stock}_{i,d}$ is a set of dummy variables that are used to capture unobserved stock fixed effects, and $\text{year}_{i,d}$ is a set of year dummy variables that are used to control for time effects. Numbers in parentheses are t-statistic.

	Quoted half-spread	Effective half-spread
hidden	-0.575	-4.446
	(-12.37)	(-81.37)
volatility	64.870	46.673
	(121.50)	(74.38)
logvol	-0.137	-0.145
	(-3.14)	(-2.83)
price	0.007	0.0003
	(7.58)	(0.23)
logcap	-1.668	-0.319
	(-19.10)	(-3.11)
constant	28.988	10.078
	(20.72)	(6.13)
R^2	0.758	0.576
Stock Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes

Table 2.12: Proportion of Liquidity Supplied by HFT and Non-HFT Firms

This table presents the shares of HFT and non-HFT participation in the liquidity supplying side for hidden and visible executions respectively. The HFT (non-HFT) share is calculated as the number of trades where an HFT (non-HFT) firm participates in the liquidity supplying side divided by the total number of trades. The first column reports the HFT and non-HFT shares in hidden liquidity, and the second column reports their shares in visible liquidity. Only trades during regular trading hours are used. T-statistics test the null that non-HFT firms provide an equal proportion of liquidity with HFT firms. The full sample consists of 120 stocks during the period from 2008-2009 and the week of February 22-26, 2010. Stocks are sorted into three market capitalization groups: large, mid, and small. Panels B-D report the results for each group.

	Hidden	Visible
Panel A: full sample		
Non-HFT	69.84%	49.34%
HFT	30.16%	50.66%
Non-HFT−HFT	39.68%	-1.32%
t-stat	45.41	3.85
Panel B: Large-cap stocks		
Non-HFT	67.92%	46.41%
HFT	32.08%	53.59%
Non-HFT−HFT	35.84%	-7.18%
t-stat	37.90	21.42
Panel C: Mid-cap stocks		
Non-HFT	80.97%	69.72%
HFT	19.03%	30.28%
Non-HFT−HFT	61.95%	39.45%
t-stat	123.75	67.94
Panel D: Small-cap stocks		
Non-HFT	85.92%	83.07%
HFT	14.08%	16.93%
Non-HFT−HFT	71.84%	66.13%
t-stat	195.38	131.97

Figure 2.1: Event Tree of Trading Process

This figure presents the structure of trading process. γ is the probability of news arrival that affects the fundamental value. α is the probability that liquidity suppliers make correct inference on whether a liquidity demander is a noise trader (NT) or a informed trader (IT). λ is the probability that liquidity suppliers are able to cancel their limit orders on the arrival of informed traders. R (UR) denotes that liquidity suppliers make correct (incorrect) inference and recognize (un-recognize) the identity of liquidity demanders. C represents cancellation of limit orders. E^B (E^S) represents execution of limit orders with respect to buy (sell) market or marketable orders.

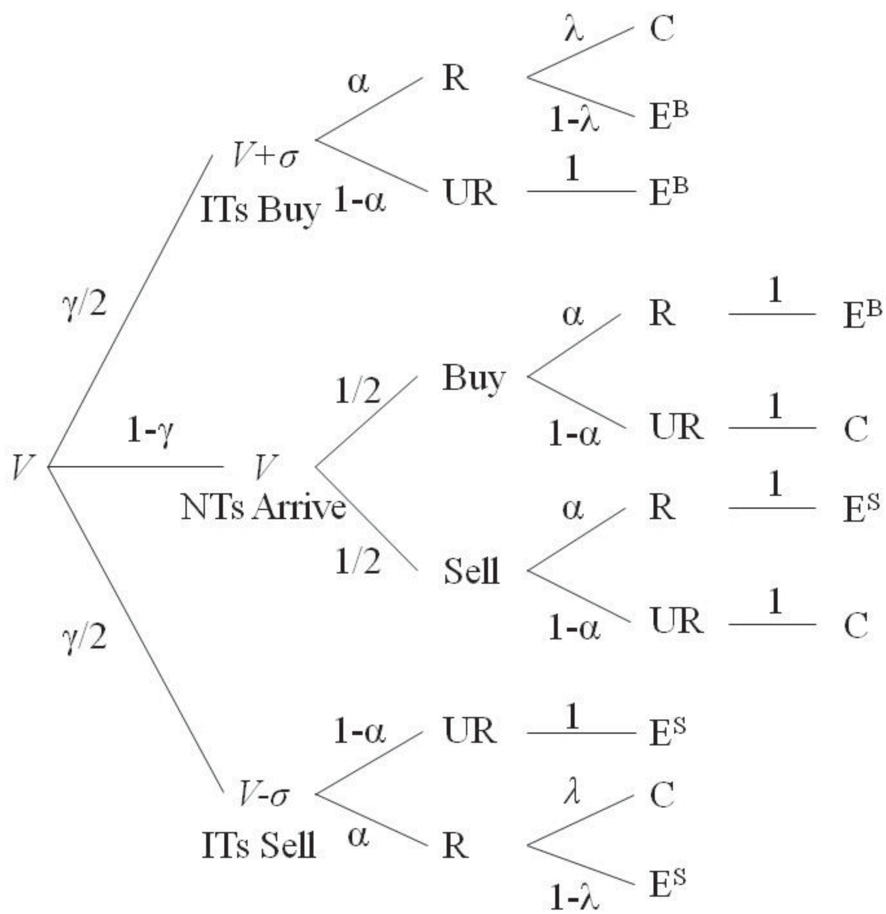

Figure 2.2: Cancelled and Executed Order Lives

This figure demonstrates the improvement in trading speed on NASDAQ for a sample of 108 stocks from 2008-2013. The blue line represents the 5th percentile life of cancelled orders that is used to measure the latency of liquidity suppliers. The red line shows the 5th percentile life of executed orders that is used as a proxy for the latency of liquidity demanders.

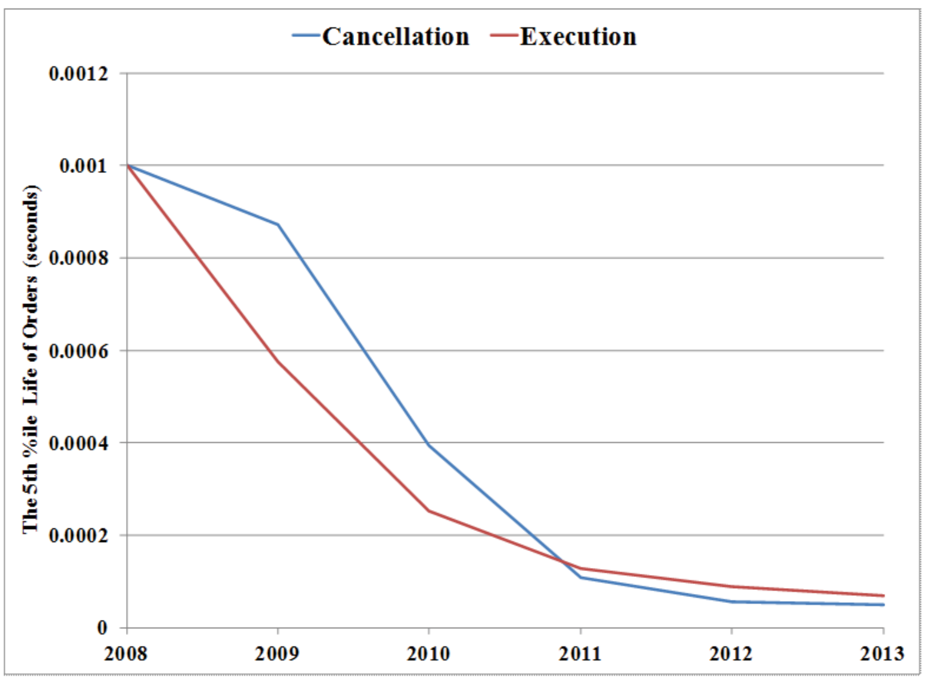

Chapter 3

Market Quality Breakdowns in Equities

Joint with Bruce Mizrach

3.1 Introduction

The collapse and sudden rebound of market indices and nearly 2,000 equity prices during the May 6, 2010 "Flash Crash" was a singular event in the history of the equity markets. The Securities and Exchange Commission (SEC) and Commodity Futures Trading Commission (CFTC) both launched full-scale investigations into the causes of the collapse, and a growing body of academic literature has also examined causal mechanisms behind the crash. Although the regulators have implemented a set of rules after the Flash Crash, including single stock circuit breakers and a ban on stub quotes, trading glitches are still happening. These serious errors in 2012, such as BATS Global Markets' initial public offering (IPO) failure on March 23, Facebook's IPO miscue on May 18, and Knight Capital's erroneous order flood on August 1, are merely the latest in a series of breakdowns. These events were not confined to a single stock exchange and the effect of each of these incidents was not transitory: BATS withdrew their IPO, NASDAQ faces billions of dollar in lawsuits from market making firms, and Knight nearly went bankrupt. The run of technology snafus have raised the concern that they could rattle investor's confidence and result in reduced liquidity of the equity market.

The new circuit breakers now trip at a 10% movement intra-daily. Many of these movements will be due to information released about individual stocks. To filter these out, we isolate stocks where prices recover to at least 2.5% below the 09:35 price. Using this definition, we find that the "Flash Crash" is not an isolated event. For example, on April 4, 2000, more than 1,500 stocks fall more than 10% intra-day before recovering most of their losses. These breakdowns occur by more than 30,000 events in 2000.

We ask a simple and straightforward question: how frequently do these breakdowns in market quality occur. We analyze every change in the listing exchanges' best bid and offer for 1993-2011. In total, we examine more than 30 million files of intra-day bid and offer quotes. What we find is that market quality breakdowns have been endemic to the equity markets. The daily average breakdown frequency is 0.64% throughout our sample period, an average of 44 stocks per day.

There is an uptrend in breakdown frequency from 1993-2000. This trend reverses from 2000-2006 but then begins to rise again in late 2007, in the early stages of the financial crisis. Breakdowns continue to rise through 2008, a particularly volatile period for the market. In 2009 though and continuing through 2010, the breakdown frequency is declining. Despite the Flash Crash, 2010 has the fewest breakdowns of any year since 2007. The breakdown frequency is 0.39% in 2011, half the rate of 1998 when humans provided the majority of quotes. Breakdowns in 2010-2011 occur less than once per year in a typical stock.

The academic literature has suggested a number of possible explanations for market quality breakdowns: (1) regulatory changes; (2) fragmentation; (3) excessive correlation; (4) exchange traded funds (ETFs); and (5) high frequency trading (HFT). We develop an explanatory model, with controls for volume and volatility, to assess the marginal effects of these potential causes of breakdowns. Also, by looking at a longer historical time frame, we hope to identify which explanations are robust.

Changes in the regulatory environment have dramatically effected quote and trade behavior. In 1996, the SEC adopted the display rule which placed electronic trading networks on an even playing field with dealers. Many others have suggested that the SEC's regulations governing the equity national market system has led to market quality deterioration.

The biggest change was the adoption of Reg. NMS in April 2005. The new regulations were extended in stages and were fully in place by October 15, 2007. One of our most striking findings is that market quality breakdowns are 41.78% less frequent after Reg. NMS. This implies that approximately $4,000$ fewer stocks are breaking down each year compared to the prior period.

Bennett and Wei (2006) claim that order flow consolidation improves market quality. Golub, Keane, and Poon (2012) attribute mini-Flash crash episodes in the period 2006-11 to the use of inter-market sweep orders in fragmented markets. Madhavan (2012) suggests that fragmentation from equity market structure changes has made markets more fragile and may have contributed to the Flash Crash. Jiang, McInish, and Upson (2014) take the contrary view, noting that order routing away from the primary exchange may result in better executions. O'Hara and Ye (2011) find that the volume share in off-exchange venues does not impact market quality. We use two measures of fragmentation in our analysis. The first is the Herfindahl index of each equity market's contribution to the national best bid and offer and the second is the market share of off-exchange volume. We do not find that either measure of fragmentation helps explain the frequency of breakdowns for the market as a whole.

Even though market structure differences have been reduced, exchanges still matter. Controlling for market capitalization, price, as well as volume and volatility, New York Stock Exchange (NYSE) stocks break down 20.03% less frequently than NASDAQ stocks, 43.91% less frequently than American Stock Exchange (AMEX) listings, and 69.04% less frequently than Archipelago (ARCA) listings.

Since we find that market structure changes appear not to be a primary factor in market breakdowns, we then consider the effects of rising security correlations. Acharya and Schaefer (2006) have noted that individual stocks become more highly correlated during financial crises. There has also been an uptrend in market correlation in recent years. The average correlation among the Fama-French industry portfolios rises from 37.16% in 1993 to 76.32% in 2011.

We construct a theoretical model with correlated liquidity shocks based on Sandås (2001). This model helps us to unify a number of factors in the literature which appear to operate through cross-equity correlation, including both ETFs and HFT.

A market maker in stock A responds to a liquidity shock in stock B, and the limit order books thins to a larger degree when the shocks are more highly correlated. We confirm the model empirically, finding that correlation does spike during market quality breakdowns, raising the frequency of breakdowns by almost 25.62%.

Ben-David, Franzoni, and Moussawi (2014) note that ETFs exacerbate the volatility of the underlying stocks through a propagation mechanism of liquidity shocks. We find that ETFs break down 90.33% more frequently than non-ETFs. ETF trading activity unidirectionally Granger causes the market correlation, revealing that ETFs are a source of stronger individual stock correlation and not vice versa.

The impact of HFT remains a widely debated issue. Brogaard (2011) analyzes a high-quality data set that identifies the trade and quote activity of high frequency trading firms. He shows that high frequency traders have become a dominant fraction of market activity, with approximately 70% of dollar volume in 2009. They also engage in highly correlated trading strategies. Several academic papers have suggested that HFT firms generally enhance market quality. Hasbrouck and Saar (2013) find that low-latency activity improves liquidity and dampens short-term volatility. Brogaard, Hendershott, and Riordan (2014) argue that HFTs increase the efficiency of prices through their marketable orders.

Other papers suggest that HFT activity might be more harmful. Gao and Mizrach (2013) find that, when markets experience stressful events, HFT firms tend to scale back their liquidity provision. For example, during the Federal Reserve large scale auction purchases of Treasuries, HFT firms were 8% less likely to be providing the inside bid or offer. Zhang (2010) observes that HFT is positively correlated with stock price volatility and hinders the ability of the market prices to reflect fundamental information. Sornette and Von der

Becke (2011), in a report prepared for the U.K. Office of Science, argue that HFT has led to crashes and can be expected to do so more and more in the future. We find that HFT trading activity has a significantly positive impact, raising the breakdown frequency by 18.33%.

We also analyze breakdown frequency using a predictive model. Two lagged breakdown probabilities are statistically significant. Along with volatility at the market open, these factors improve upon a constant forecast by nearly 50%.

We examine the robustness of our results. Rapid increases in offer prices, which we call "breakups", are also positively related to correlation shocks. Restricting the sample to stocks with a market cap of over $10 billion, we confirm the explanatory model for aggregate breakdown frequency. Large caps breakdown only 1/6 as often as other stocks, and only 118 large caps have broken down in the period 2009-11.

We also consider alternative microstructure definitions of breakdowns. We look at the national best bid and offer (NBBO) rather than just the primary exchange. These results are very similar to our original specification. Secondary markets reduce breakdowns by 31.96% in 2008-11. We also show that our model fits breakdowns using trade prices rather than quotes. Finally, when we look at the worst bid or offer in the market place, all of our models fit poorly.

Section 4.4 introduces our definition of market quality breakdowns and compares it to the Flash Crash. Section 3.3 measures the unconditional daily average breakdown frequency. We develop a baseline model for the aggregate breakdown frequency in Section 3.4 and test the impact of changes in market structure including Reg. NMS, fragmentation and exchange effects. We then construct, in Section 4.3, a theoretical model to study the effects of cross-security correlation on the limit order book. The empirical results in Section 3.6 confirm our model of correlated liquidity shocks. We analyze the impact of ETFs in Section 3.7 and HFT activity in Section 3.8. Section 3.9 builds a predictive model. We conduct robustness checks in Section 4.6 before concluding.

3.2 Data and Definitions

Our empirical analysis relies on quotes rather than trades. This, of course, increases the computational burden, but we feel breakdowns in market quality impede trading, and that the consolidated quotes provide the best real time portrait of the market. Our focus is on the best bid and offer from the listing exchange, but we examine alternative definitions[1] in our robustness section.

We analyze stocks that are in both the Center for Research in Security Prices (CRSP) and the New York Stock Exchange Trade and Quote Database (TAQ). Stocks from all three major exchanges are available from April 6, 1993 forward. Our sample ends on December 30, 2011.

We look at movements in the time frame 09:35-15:55. We do this because opening and closing procedures vary across exchanges and may not be comparable.

A stock is identified as having a market quality breakdown if the best bid prices fall 10% below the 09:35 price. 10% is a natural metric because that is where circuit breakers are now placed. In addition, the tick must be repeated at least once in a subsequent calendar second. This avoids fleeting quotes or errors.

We want to try to filter out news driven price declines. We do this by looking at stocks that rebound to within 2.5% of the 09:35 price at 15:55.

We have a symmetric definition for break ups, using the best offer price.

---

[1]The high frequency data provider Nanex has been analyzing mini-flash crashes for several years now. Nanex (2011) used the following criteria: "to qualify as a down-draft candidate, the stock had to tick down at least 10 times before ticking up – all within 1.5 seconds and the price change had to exceed 0.8%." They have a symmetric definition for up-drafts. We attempted to replicate their results, but we eventually used an alternative framework. Many Nanex breakdowns occur away from the primary exchange. It also does not work particularly well during the flash crash. In 2010, they identify 1,041 down-draft events, fewer than we find on the single day of May 6, 2010. Nanex does confirm our main message. Up and down drafts have been trending down since 2008.

3.2.1 Market quality metrics for the Flash Crash

On May 6, 2010, major U.S. stock indices, stock index products, and individual stocks experienced a sudden price drop of more than 5% followed by a rapid recovery within minutes. The unusual and severe event, commonly known as the Flash Crash, occurred in both futures and spot markets. The price of E-Mini S&P 500 futures fell in excess of 5% between 14:41 and 14:46. The Dow Jones Industrial Average (DJIA) plunged 998.5 points, the largest intraday point decline in history. Many individual stocks reached lows that exceeded 10%, and some were even traded down to a penny, e.g. Accenture.

The Flash Crash raises questions about the quality of U.S. financial markets. The CFTC and the SEC[2] jointly explored the market events of May 6, 2010 and identified the evaporation of liquidity in both the E-Mini and individual stocks. By analyzing the aggregate order books they found that reductions in liquidity may lead some stocks to trade at severe prices.

When we apply our filter to the Flash Crash day, we get very similar conclusions to the academic and policy literature. Our sample consists of 6,527 securities for May 6, 2010. Among them, 1,857 stocks experienced market quality breakdowns on the listing exchange.[3] The breakdown frequency differs on each of the four primary exchanges, as shown in Table 3.1.

[INSERT Table 3.1 HERE]

ARCA is affected more than any other exchange with more than 60% of stocks crashed, while AMEX has the lowest frequency of 12.73%. The breakdowns on NYSE and Nasdaq are close to the average level of the market.

---

[2] Commodity Futures Trading Commission and Securities and Exchange Commission (2010).

[3] It is shown in the CFTC-SEC report that approximately 14% of stocks traded at lows that are more than 10% away from the 14:40 prices. Given the fact that we analyze 10% price decline below the 09:35 price and our study relies on quotes instead of trades, it makes sense that our filter gives a higher frequency of 28.45%.

We then analyze in more detail the distribution of percentage decline in the best bid prices on the Flash Crash day. We want to compare our results to the CFTC-SEC finding, so we use the same stock filter here, i.e. a share price of more than $3.00 and a market capitalization of at least $10 million. The results are illustrated in Figure 3.1 Panel A.

[INSERT Figure 3.1 HERE]

The distribution displays a similar pattern to the finding in the CFTC-SEC report.[4] In particular, 227 stocks have the lowest best bid that are almost 100% below the 09:35 price on the listing exchange. The number is a little greater than that given in the CFTC-SEC report since we analyze quotes rather than trades.

Figure 3.1 Panel B presents a scatter plot of the time and percentage decline of the best bid for all stocks during the period from 14:00-15:00 on May 6, 2010. Each point on the graph represents a stock. The result is consistent with the finding by the CFTC and the SEC.[5] A few stocks began to crash shortly after 14:00 and the number of stocks increased steadily over the one hour interval. Many of the lows in the best bid occurred after 14:45, as represented by a dense area between −20% and 0% and a thick line around −100%.

Our results also suggest the same conclusions about ETFs as discussed in the CFTC-SEC report. ETFs were affected the most among all types of securities on May 6, 2010. Based on our filter, 559 out of 893 ETFs experienced market quality breakdowns. The number accounts for nearly one-third of the crashed stocks on the day. We present in Figure 3.2 the distribution of ETF lows measured by the best bid from the listing exchange. The spike of the left-most column in the figure indicates that a large portion of ETFs had almost 100% quality deterioration.

[INSERT Figure 3.2 HERE]

[4]Commodity Futures Trading Commission and Securities and Exchange Commission (2010), Figure 8, p.18.

[5]Commodity Futures Trading Commission and Securities and Exchange Commission (2010), Figure 10, p.24.

The timing of ETF lows by our filter, as shown in Figure 3.2 Panel B, is consistent with the finding in the CFTC-SEC report[6] as well. The number of ETF crashes started to rise after 14:40. Beginning about 14:45 a great number of ETFs experienced 100% price drops, which is represented by a dense line around −100%.

Since our filter works for the Flash Crash, we then apply it to the full TAQ sample of 1993-2011. The natural question is how often do events like this occur.

3.3 Unconditional Breakdown Frequency

We analyze the unconditional probability of a random stock experiencing a market quality breakdown. The breakdown frequency is calculated by the number of broken stocks divided by the total number of traded stocks. We report results for all exchanges and all types of stocks in Figure 3.3 Panel A.

[INSERT Figure 3.3 HERE]

The daily average breakdown frequency is 0.64% throughout our sample period. There is an uptrend in breakdown frequency from 1993-2000 followed by a downtrend from 2000-2006. Breakdowns begin to rise again in late 2007 with the onset of the financial crisis and peak in 2008 during the near collapse of the financial system. As the market stabilizes in the second half of 2009, the breakdown frequency returns to the level in the late 1990s. By 2011, the breakdown frequency has fallen to 0.39%, half the rate in 1998. In 2010 and 2011, a typical stock will break down less than once per year.

On average, 44 stocks per day experience breakdowns. Figure 3.3 Panel B shows the number of breakdown events by year from 1993-2011. It presents a similar trend with the breakdown frequency in Panel A. The breakdowns reached their peak in 2000 with more than 30,000 events, and there is a significant rise in 2008 as well.

---

[6] Commodity Futures Trading Commission and Securities and Exchange Commission (2010), Figure 16, p.39.

Even including the Flash Crash, there are fewer market quality breakdowns in 2010 than in 1998. Excluding the Flash Crash day, the breakdown frequency in 2010 is the fourth lowest in our sample.

When there are a large number of breakdown events in a year, it could be the case that some particular stocks break down more frequently or that stocks are essentially equally likely to break down. To distinguish between the two cases, we measure the distribution of breakdown incidence by the Gini coefficient in Figure 3.4.

[INSERT Figure 3.4 HERE]

The Gini coefficient of 0.52 implies that stocks are not equally likely to break down during our sample period. We show below that non-NYSE stocks are more likely to break down, and large capitalization stocks are less likely. Even though breakdown frequencies vary substantially year-by-year, the distribution across stocks is relatively stable during the period from 1993-2011.

3.4 Market Structure

We plan to explore the various theories in the literature by first developing a baseline model for the frequency of breakdowns. We will then extend this baseline model to see the time series impact of changes in market structure, both regulatory and competitive.

3.4.1 Baseline model

We now model the frequency of market quality events conditional on volatility and aggregate volume. We calculate the breakdown frequency on day t, π_t, by dividing the number of broken stocks by the total number of traded stocks. We measure market volatility using the opening value of the Chicago Board Options Exchange volatility index (VIX), $\sigma_t^{VIXopen}$. The daily aggregate volume, v_t, is the sum of trading activity on each exchange in its own listings. In the model, we use a dummy variable, $\widetilde{v}_t$, to represent volume spikes.

$$\widetilde{v}_t = I_{\alpha=0.05}\left(\frac{v_t - \sum_{j=1}^{20} v_{t-j}/20}{\sigma_t^v}\right), \tag{3.1}$$

where $I(\cdot)$ is an indicator function and σ_t^v is the standard deviation of the volume over the proceeding 20 days. In other words, $\widetilde{v}_t$ is set as 1 if the volume becomes significantly higher than the average of proceeding 20 days at the 5% level, and is 0 otherwise.

Given the fact that the breakdown frequency in most times is close to zero and not normally distributed, we use a generalized linear model with the assumption of $\pi_t \sim \Gamma(k, \theta)$, where k and θ are respectively the shape and scale parameter of the gamma distribution. The baseline model is written as

$$\log(\mathrm{E}[\pi_t]) = \alpha + \beta_1 \sigma_t^{VIXopen} + \beta_2 \widetilde{v}_t. \tag{3.2}$$

The model is estimated by quasi-maximum likelihood method using robust standard errors, and the results are shown in Table 3.2.

[INSERT Table 3.2 HERE]

All the estimated coefficients are statistically significant. The market volatility affects positively the aggregate breakdown frequency, which is consistent with intuition. The aggregate volume is positively associated with the breakdown frequency as well.

We measure goodness-of-fit using McFadden's measure, R_M^2, which is defined as

$$R_M^2 = 1 - \frac{\log L(M_f)}{\log L(M_i)} \tag{3.3}$$

where $\log L(M_f)$ is the log-likelihood of the full model and $\log L(M_i)$ is the log-likelihood of the model with just an intercept. Since the log-likelihood is non-positive, R_M^2 ranges from 0 to 1 and has a higher value for the model with better fit. Our baseline model shows a 42.22% improvement over the intercept-only model to explain the breakdown frequency.

Next we examine whether breakdowns have increased since Reg. NMS.

3.4.2 Reg. NMS

On April 6, 2005, the Securities and Exchange Commission (SEC), in a 3-2 vote, adopted Regulation National Market System (Reg. NMS). The SEC rules were adapting the national market system concept to the modern electronic marketplace. There are four major provisions: (1) Rule 610, which provides equal access to markets; (2) Rule 611, which prohibits trade-throughs of displayed and accessible quotations; (3) Rule 612, which prohibits subpenny quotations except in limited circumstances; (4) Rule 600, 601 and 603, which set up rules for market data.

We model whether breakdowns increased after the rules were fully adopted on October 15, 2007[7], by including a dummy variable d_t^{NMS} into the baseline model. The result in Table 3.2 shows that breakdowns become significantly less frequent after Reg. NMS, despite the Flash Crash. Quantitatively, breakdowns have fallen $-41.78\% = e^{-0.5410} - 1$ since the passage of Reg. NMS. With approximately $7,000$ U.S. equity listings, this implies that 18 fewer stocks each day are experiencing breakdowns or approximately $4,500$ fewer market quality breakdowns each year.

3.4.3 Market fragmentation

The academic literature is divided on the effects of fragmentation. We follow Madhavan (2012) and use the Herfindahl index as a measure of fragmentation.[8]

We first compile the national best bid and offer (NBBO) across all exchanges using the consolidated quotes from the TAQ database. The Herfindahl index is then computed as

---

[7]Reg. NMS was implememted in steps from 2005 to 2007. Our result is robust to the choice of a break point. Even if we break at the beginning of Reg. NMS in 2005, breakdowns have fallen significantly.

[8]The only difference from Madhavan (2012) is the way that we count the number of times for an exchange with the national best bid or offer. For example, if two exchanges have the best national quote at the same time, he counts one for each, while we assign one half to each to reflect the fact of their competition for orders.

the sum of squared frequencies of the best national bid or offer that each exchange posts. Mathematically, the Herfindahl index for stock i on day t is expressed as

$$H_{i,t} = \sum_{j=1}^{M} (f_{i,t}^{j})^2, \tag{3.4}$$

where the frequency $f_{i,t}^{j}$ is calculated as the proportion of times exchange j is the national best bid or offer on day t, and M is the total number of exchanges where stock i has quotation activity. If multiple venues are the best bid or offer at the same time, we give an equal weight to each of these venues since they are competing to attract order flows. The market fragmentation on day t is measured by the average of Herfindahl index values across all stocks,

$$H_t = N_t^{-1} \sum_{i=1}^{N_t} H_{i,t}, \tag{3.5}$$

where N_t is the total number of stocks on day t. It is worth noting that the Herfindahl index is smaller when the market is more fragmented.

Consistent with the volume variable defined by (3.1) in our baseline model, we use a dummy variable, $\widetilde{H}_t$, to represent spikes of market fragmentation. The dummy is set as 1 if the Herfindahl index H_t is significantly lower than the average of proceeding 20 days at the 5% level, and is 0 otherwise. Since the market quality breakdown is defined using the bid price, the Herfindahl index in the breakdown model is based on the best bid. When we add the Herfindahl measure to our baseline regression, its coefficient is not statistically significant, as shown in Table 3.2. Therefore, we conclude that fragmentation is not associated with the breakdown frequency for the market as a whole.[9]

Since March 5, 2007, the SEC has required that all off-exchange trades must report to a

[9] We also explored the impact of fragmentation for individual stocks on the Flash Crash day, as Madhavan (2012) did. We include as control variables of opening price, market capitalization, volatility, and volume. The estimated coefficient for the Herfindahl index is negative (−0.6542) and statistically significant at the 1% level. Our result is consistent with Madhavan's conclusion about fragmentation for the Flash Crash. However, we do not find this measure helps explain the breakdown frequency at the aggregate level for a longer historical period.

trade reporting facility (TRF). O'Hara and Ye (2011) have suggested using the share of off-exchange volume as an alternative measure of fragmentation. We use a similar metric in our analysis as well. The TAQ database provides aggregate information of trades reported to the Financial Industry Regulatory Authority (FINRA).[10] We re-estimated the baseline model for the period from March 5, 2007 to December 30, 2011, and then explored the impact of this alternative measure of fragmentation. We find in Table 3.2, similar to O'Hara and Ye (2011), that the share of TRF volume is not a statistically significant contributor to market quality breakdowns.

3.4.4 Do exchanges still matter?

We now ask whether exchanges influence the frequency of breakdowns. We investigate it by modeling the number of breakdown occurrences of individual stocks, $n_{i,t}$. We analyze breakdowns at the monthly frequency, because the number of daily breakdowns is generally very small.

Since the time of the NASDAQ price fixing case (see e.g. Christie and Schultz (1994)), there has been an ongoing dialog of market quality across exchanges. The conclusion of the early literature was that the NYSE, despite having a monopoly market making specialist, typically had higher market quality. The debate continues to this day, especially involving the role and importance of market makers, e.g. Menkveld and Wang (2011). Recently some exchanges have proposed to offer market makers financial incentives to provide more liquidity in illiquid stocks. In December 2011 BATS filed and later was approved for the Competitive Liquidity Provider program that was designed to encourage market makers to post tight quoting spreads[11]. The NASDAQ filed a revised plan with the SEC in December

[10]The FINRA trades include those from the Nasdaq TRF, the NYSE TRF and the Alternative Display Facility (ADF). The TRF data sample in O'Hara and Ye (2011) also includes trades reported to National Stock Exchange (NSX) TRF. However, based on their results, it accounts for only 2.46% of consolidated volume, compared to the total share of 24.75% in the sources captured by the FINRA.

[11]See e.g. "BATS Gets SEC Approval for Liquidity Provider Program," *Traders Magazine*, February 6, 2012.

2012 that would pay market makers in thinly traded ETFs[12].

Figure 3.5 presents the breakdown frequency on each of the four primary exchanges from 1993-2011.

[INSERT Figure 3.5 HERE]

We see that breakdowns occur less frequently on the NYSE than any other exchange. Even in 2008, the frequency is less than 0.9%, and the frequency in 2011 is very close to the levels in early 1990s. The breakdown frequency on the NASDAQ shows a similar trend to the result for all exchanges, but the magnitude is higher. Interestingly, the ARCA experiences a breakdown frequency as high as 2.6% in 2007 when the frequency is relatively low on other exchanges. After that, there is a substantial improvement of market quality for ARCA and it is the second best exchange in 2011.

The unconditional probabilities are not by themselves indications of exchange related effects. Stocks differ across exchanges, and we must control for these in our market quality inferences. To test for marginal effects from exchange structure, we include the covariates from the baseline model and add the log opening price of the stock, $p_{i,t}^{open}$, and its log market capitalization, $\kappa_{i,t}$. Since the dependent variable is the number of market quality events for stock i in month t, we use Poisson regression with the assumption of $n_{i,t} \sim \text{Pois}(\lambda)$,

$$\log(\mathrm{E}[n_{i,t}]) = \alpha + \beta_1 p_{i,t}^{open} + \beta_2 \kappa_{i,t} + \beta_3 \sigma_{i,t} + \beta_4 \widetilde{v}_{i,t}. \tag{3.6}$$

The model is estimated by the quasi-maximum likelihood method using robust standard errors. The results in Table 3.3 show that all the estimated coefficients are statistically significant. At the individual stock level, the opening price and market capitalization are negatively associated with the number of breakdowns, while volatility and volume[13] are

[12] See e.g. "Nasdaq Seeks Approval for Revised 'Paid-for-Market-Making' Plan," *Traders Magazine Online News*, December 10, 2012

[13] Volume here is the trading activity on the primary exchange, as calculated from the TAQ data.

positively related to breakdowns.

[INSERT Table 3.3 HERE]

We then add three dummy variables, d^{NYSE} for the NYSE, d^{NASD} for the NASDAQ, and d^{ARCA} for the ARCA to (3.6), using AMEX as the omitted listing exchange. The results in Table 3.3 indicate that the exchange listing significantly affects the number of breakdowns for individual stocks. Even though the NYSE has lost market share in its own issues, NYSE listed stocks break down approximately 20.03% less frequently than NASDAQ stocks, 43.91% less frequently than AMEX listings, and 69.04% less frequently than ARCA listings.

We now turn from issues of market and regulatory structure to examine whether cross-security correlation might be explaining market quality breakdowns.

3.5 The Theoretical Model

The model presented in this section follows Sandås (2001). We extend Sandås' model to include two risky equities and their corresponding limit order books. Our model contributes to the literature by introducing the correlation between securities and analyzing the cross-equity impact on limit order books. The theoretical results discussed here provide a framework for the subsequent empirical analysis.

3.5.1 Model setup

We consider two risky equities, A and B, in the model. Equity i, $i = A$ or B, has a fundamental value X_t^i in period t, which incorporates all information available up to period t. The fundamental value for security i in the next period is

$$X_{t+1}^i = X_t^i + \mu^i + \varepsilon_{t+1}^i, \tag{3.7}$$

where μ^i represents the expected change in the fundamental value and ε^i_{t+1} is a random innovation in period $t+1$.

There are two types of agents, market makers and traders. Market makers provide liquidity by placing limit orders on one or both assets. They are risk neutral and profit maximizing. Traders are risk averse and their trades may be due to exogenous reasons, e.g. margin calls, rather than their best estimate of the fundamental value. Therefore, they want to trade quickly at the current price.

There are three stages in each period t. In the first stage, market makers submit new limit orders on one or both equities. They repeat the process until no market maker finds it optimal to place an additional order. Then a trader arrives and submits market orders on either security or both. The market order quantity on each equity may rely on the correlation between the two. Finally, market makers update their expectation about the fundamental values of the assets given the size of incoming trades and the process starts over.

For each limit order book we use a discrete pricing grid as follows. The bid prices in the book of equity i are denoted by $\left\{p^i_1, p^i_2, \ldots, p^i_k\right\}$, where p^i_1 is the best bid price. Let $\left\{Q^i_1, Q^i_2, \ldots, Q^i_k\right\}$ denote the order quantities associated with each price. The variables for the offer side may be denoted analogously, but we focus on the bid side only for market quality breakdowns. The market order quantity for equity i is denoted by m^i. It is positive for buy orders, and negative for sell orders.

3.5.2 Traders

Suppose that a trader buys or sells with equal probability. Following Sandås (2001) we assume that the market order quantities are exogenous and are exponentially distributed for simplicity. We focus on sell orders only when modeling market quality breakdowns. To incorporate the correlation between the two equities into the model, we use the bivariate exponential distribution[14] with the following joint density function,

$$f\left(m^A, m^B\right) = \frac{1}{\lambda^A \lambda^B} \mathrm{e}^{\frac{m^A}{\lambda^A}+\frac{m^B}{\lambda^B}} \left[1 + 4\rho\left(1 - 2\mathrm{e}^{\frac{m^A}{\lambda^A}}\right)\left(1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}}\right)\right],\ m^A \leq 0 \text{ and } m^B \leq 0. \tag{3.8}$$

It is not difficult to show that the marginal distributions of m^A and m^B are exponential with mean λ^A and λ^B respectively, and the correlation between m^A and m^B is ρ, where $-1 \leq \rho \leq 1$. We then mainly concentrate on the decision problem of the market makers.

3.5.3 Market Makers

Market makers observe trades on both equities and then update their best estimates of the fundamental values based on the market order quantities. Since equity A and B are symmetric in our model, only security A's fundamental value in the next period is given below,

$$\mathrm{E}\left[X_{t+1}^A | X_t^A, m^A, m^B\right] = X_t^A + \mu^A + h\left(m^A, m^B\right), \tag{3.9}$$

where $h\left(m^A, m^B\right)$ is a non-decreasing price impact function. It captures the market order impact of both equities on the fundamental value of security$A$. We assume that the price impact function is linear with respect to the market order quantity of each of the equities, i.e.

$$h\left(m^A, m^B\right) = \alpha m^A + \beta m^B, \tag{3.10}$$

where α and β represent respectively the marginal price impact of market orders on security A and B. Since buy (sell) orders typically contain positive (negative) information about the fundamental value of the asset, both α and β are expected to be positive.

To simplify the analysis, we assume that market makers face a quantity-invariant order processing cost c. Then the profit of a limit order at the best bid price level p_1^A of equity A

[14] We use one of the Gumbel (1960) versions of the bivariate exponential distribution because it clearly identifies the effect of the cross-asset correlation.

is given by

$$\pi_1^A = p_1^A - c - \mathrm{E}\left[X_{t+1}^A | X_t^A, m^A, m^B\right]. \tag{3.11}$$

Under this setup, we can calculate the expected profit of a limit order placed on equity A's last unit q^A given a market order on equity B.

$$\begin{aligned}
\mathrm{E}\left[\pi_1^A | m^B\right] &= \mathrm{E}_{-m^A \geq q^A}\left[p_1^A - c - \mathrm{E}\left[X_{t+1}^A | X_t^A, m^A, m^B\right]\right] \qquad (3.12)\\
&= \int_{-\infty}^{-q^A} \left(p_1^A - c - X_t^A - \alpha m^A - \beta m^B\right) \frac{1}{\lambda^A \lambda^B} \mathrm{e}^{\frac{m^A}{\lambda^A} + \frac{m^B}{\lambda^B}} \\
&\qquad \left[1 + 4\rho\left(1 - 2\mathrm{e}^{\frac{m^A}{\lambda^A}}\right)\left(1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}}\right)\right] dm^A \\
&= \frac{1}{\lambda^B} \mathrm{e}^{-\frac{q^A}{\lambda^A} + \frac{m^B}{\lambda^B}} [\left(p_1^A - c - X_t^A + \alpha\left(q^A + \lambda^A/2\right) - \beta m^B\right) \\
&\quad \left(1 + 4\rho\left(1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}}\right)\left(1 - \mathrm{e}^{-\frac{q^A}{\lambda^A}}\right)\right) + \frac{\alpha\lambda^A}{2}\left(1 + 4\rho\left(1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}}\right)\right)].
\end{aligned}$$

A zero-profit condition in (4.6) will characterize equilibrium.

3.5.4 Equilibrium

The model is in equilibrium if no market maker can profit by submitting an additional limit order at any price level. Therefore, the quantity placed at any price level must satisfy that the last unit breaks even, i.e. the expected profit of the marginal limit order at the end must be zero. From (4.6) we can obtain the quantity Q_1^A submitted at the best bid price level p_1^A by solving the equation.

$$\begin{aligned}
&\left(p_1^A - c - X_t^A + \alpha\left(Q_1^A + \lambda^A/2\right) - \beta m^B\right)\left(1 + 4\rho\left(1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}}\right)\left(1 - \mathrm{e}^{-\frac{Q_1^A}{\lambda^A}}\right)\right) \\
&\qquad + \frac{\alpha\lambda^A}{2}\left(1 + 4\rho\left(1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}}\right)\right) = 0.
\end{aligned} \tag{3.13}$$

We are more interested in the cross-equity effect of market orders on the limit order book. To analyze it we take the derivative of Q_1^A with respect to m^B,

$$\frac{\partial Q_1^A}{\partial m^B} = \frac{\beta + 4\rho C}{\alpha + 4\rho D}, \tag{3.14}$$

where

$$\begin{aligned} C &= \beta \left(1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}}\right)\left(1 - \mathrm{e}^{-\frac{Q_1^A}{\lambda^A}}\right) + \alpha \frac{\lambda^A}{\lambda^B} \mathrm{e}^{\frac{m^B}{\lambda^B}} \\ &\quad + \frac{2}{\lambda^B} \mathrm{e}^{\frac{m^B}{\lambda^B}} \left(1 - \mathrm{e}^{-\frac{Q_1^A}{\lambda^A}}\right)\left(p_1^A - c - X_t^A + \alpha\left(Q_1^A + \lambda^A/2\right) - \beta m^B\right) \\ D &= \alpha \left(1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}}\right)\left(1 - \mathrm{e}^{-\frac{Q_1^A}{\lambda^A}}\right) \\ &\quad + \frac{1}{\lambda^A} \mathrm{e}^{-\frac{Q_1^A}{\lambda^A}} \left(1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}}\right)\left(p_1^A - c - X_t^A + \alpha\left(Q_1^A + \lambda^A/2\right) - \beta m^B\right). \end{aligned}$$

We see that both $C > 0$ and $D > 0$ if $1 - 2\mathrm{e}^{\frac{m^B}{\lambda^B}} > 0$ and $p_1^A - c - X_t^A + \alpha\left(Q_1^A + \lambda^A/2\right) - \beta m^B > 0$. Therefore, when $m^B < \min\left\{-\lambda^B \log 2, \beta^{-1}\left(p_1^A - c - X_t^A + \alpha\left(Q_1^A + \lambda^A/2\right)\right)\right\}$, we have $\partial Q_1^A/\partial m^B > 0$. Since $m^B \leq 0$, it suggests that an increase in the quantity of a market sell order on equity B will also reduce the depth at the best bid price level of equity A.

In addition, (4.9) indicates the impact of the correlation ρ on the order book depth. When ρ increases, $\partial Q_1^A/\partial m^B$ becomes greater, so the cross-equity effect of a market sell order is even stronger.

The condition of zero expected profit also satisfies for the price levels deeper in the limit order book. We obtain the same conclusions for the cross-equity effect of market orders and for the impact of the correlation as well. When the limit order book thins to the extreme price level, a market quality breakdown occurs. If the trades are based on exogenous reasons rather than the expected fundamental value, the market would learn later nothing fundamental has happened and therefore the price movements are reversed.

3.6 Market Correlation

The critical parameter in the model is the cross-equity correlation ρ. We first try to measure how the correlation across stocks has changed during our sample period.

Since 2008, there have been instruments[15] that directly measure the market's implied correlation. These instruments are not available for our entire sample, so we construct our own measure using daily returns of 30 industry portfolios from Fama and French's website. We calculate the 20-day rolling correlation pairwise and use the median as measure of the market correlation.

We plot the market correlation over the sample period in Figure 3.6.

[INSERT Figure 3.6 HERE]

We include the median correlation in our explanatory model

$$\log(\mathrm{E}[\pi_t]) = \alpha + \beta_1 \sigma_t^{VIXopen} + \beta_2 \widetilde{v}_t + \beta_3 d_t^{NMS} + \beta_4 \widetilde{\rho}_t. \tag{3.15}$$

and present the results in Table 3.2. We construct our correlation variable, $\widetilde{\rho}_t$, as a dummy variable that represents spikes in market correlation, just as we did with volume in (3.1).

The results suggest that it does explain a higher frequency of market quality events. Spikes in market correlation raise the breakdown frequency by 25.62%.

3.7 Exchange Traded Funds

The growing volume of trading in ETFs has changed the character of the market. ETFs broke down more frequently during the Flash Crash. We confirm, after controlling for individual equity and exchange effects, that ETFs suffer substantially more market quality breakdowns.

[15]The Chicago Board Options Exchange began to disseminate an implied correlation measure for the S&P 500 options basket and its components. There are also instruments which use this correlation as an underlying.

It is an open question whether ETFs make other stocks unstable. We find they do contribute to market-wide equity breakdowns through their effects on market correlation.

We investigate whether ETFs as an equity class break down more often by including a dummy variable for ETFs, d_i^{ETF} into the individual stock model (3.6). We control for market capitalization, opening price, individual stock volume and volatility. We also want to ensure that our ETF results are not simply proxying for the effects of exchanges, so we include dummies variables for the NYSE, Nasdaq, and AMEX. As shown in Table 3.3, ETFs exhibit significantly higher likelihood of breakdowns, after controls, than non-ETFs. ETFs break down 90.33% more frequently. If the market consisted exclusively of ETFs, there would be greater than $9,000$ more breakdowns per year.

We also explore the relationship between market correlation and the trading activity of ETFs. We use the Granger causality test to study whether an increase in trading volume of ETFs causes a change in market correlation, or the other way around. In order to do the test, we first fit the correlation and aggregate ETF volume into a vector autoregressive model,

$$\begin{aligned} \widetilde{\rho}_t &= a_{\rho,0} + \sum_{i=1}^{M} a_{\rho,i}\widetilde{\rho}_{t-i} + \sum_{i=1}^{M} b_{\rho,i}\widetilde{v}_{t-i}^{ETF} + \varepsilon_{\rho,t}, \\ \widetilde{v}_t^{ETF} &= a_{v,0} + \sum_{i=1}^{M} a_{v,i}\widetilde{\rho}_{t-i} + \sum_{i=1}^{M} b_{v,i}\widetilde{v}_{t-i}^{ETF} + \varepsilon_{v,t}, \end{aligned} \tag{3.16}$$

where $\widetilde{\rho}_t$ and $\widetilde{v}_t^{ETF}$ are dummy variables that represent spikes of market correlation and aggregate ETF volume respectively. The number of lags $M = 4$ is determined by the Akaike Information Criterion. Based on this specification, we obtain the Granger causality test results which are reported in Table 3.4 Panel A.

[INSERT Table 3.4 HERE]

There is a statistically significant causation of aggregate ETF volume for market correlation, but the reverse is not true. Because ETFs cause correlation spikes, they increase the probability that other stocks will break down the next day. A lagged ETF volume spike raises the probability of a correlation spike by as much as 73.60%.

3.8 High Frequency Trading

To examine whether HFT activity helps explain the market quality breakdowns, we use the data set analyzed by Brogaard, Hendershott, and Riordan (2014) and Carrion (2013). The HFT data set includes all trades on the NASDAQ exchange for 120 stocks on each trading date in 2008 and 2009, as well as one week in February 2010. We focus only on 2008 and 2009 in our analysis. The data tells whether an HFT firm is a liquidity taker or a provider in each trade. We measure HFT activity as the share of volume executed by HFT firms in a trading day, either at the liquidity seeking side or at the passive side. We analyze the HFT effect using a dummy variable $\widetilde{HFT}_t$ that measures spikes in the HFT share of trading volume, constructed in the same way as volume and correlation. We assume that HFT firms have similar trading activity in other stocks that we do not observe in this sample, and extrapolate the measure to the broader equity market.[16]

The impact of HFT activity on market quality breakdowns may operate through the correlation channel. Our theoretical model predicts that a high correlation between market orders will result in a larger cross-equity effect, thus contributing to a higher breakdown frequency. We investigate this possibility by the Granger causality test in the framework of the following vector autoregressive model,

$$\begin{aligned} \widetilde{\rho}_t &= a_{\rho,0} + \sum_{i=1}^{M} a_{\rho,i}\widetilde{\rho}_{t-i} + \sum_{i=1}^{M} b_{\rho,i}\widetilde{HFT}_{t-i} + \varepsilon_{\rho,t}, \\ \widetilde{HFT}_t &= a_{h,0} + \sum_{i=1}^{M} a_{h,i}\widetilde{\rho}_{t-i} + \sum_{i=1}^{M} b_{h,i}\widetilde{HFT}_{t-i} + \varepsilon_{h,t}, \end{aligned} \tag{3.17}$$

where $\widetilde{\rho}_t$ and $\widetilde{HFT}_t$ are dummy variables representing spikes of market correlation and HFT

[16] We test the assumption by estimating the model using only the 120 stocks in the data set for which we observe volume directly. The effects of HFT are similar to the extrapolated sample.

activity respectively. The number of lags $M = 4$ is determined by the Akaike Information Criterion (AIC). The results in Table 3.4 Panel B suggest that HFT activity Granger causes market correlation significantly, but the reverse is only marginally significant. A lagged HFT volume spike raises the probability of a correlation spike by as much as 297.2%.

We re-estimate the explanatory model for the period from 2008 to 2009, controlling for correlation as our theoretical model suggests,

$$\log(\mathrm{E}[\pi_t]) = \alpha + \beta_1 \sigma_t^{VIXopen} + \beta_2 \widetilde{v}_t + \beta_3 \widetilde{HFT}_t + \beta_4 \widetilde{\rho}_t. \tag{3.18}$$

The results in Table 3.5 indicate that HFT activity has a significantly positive impact on market quality breakdowns, even when we include the effects of market correlation.

[INSERT Table 3.5 HERE]

The marginal effect of correlation spikes is 31.31% in 2008-2009, and spikes in HFT activity raise the breakdown frequency an additional 18.33%. If both variables trigger in a trading day, the breakdown frequency rises nearly 50%.

3.9 Prediction

We examine whether breakdowns are predictable. We take the only lagged variable from the explanatory model (3.15), the 09:30 opening value of the VIX, and then add the two prior days' breakdown probabilities π_{t-j}, $j = 1, 2$,

$$\log(\mathrm{E}[\pi_t]) = \alpha + \beta_1 \sigma_t^{VIXopen} + \sum_{j=1}^{2} \eta_j \pi_{t-j}. \tag{3.19}$$

The results in Table 3.6 demonstrate that the breakdown frequency is positively autocorrelated.

[INSERT Table 3.6 HERE]

Both lags are statistically significant.[17] The breakdown frequency rises by 25.26% if the frequency last day doubles the average. Following two days of breakdowns at twice the average rate, the breakdown frequency becomes 52.35% higher. The R_M^2 in the predictive model of breakdowns is 49.36%, slightly higher than 46.24% that we found in the explanatory model which included contemporaneous variables.

3.10 Robustness Checks

We recognize that researchers might use alternative definitions of a market quality breakdown. They might want to use 15% rather than 10% or have the day close flat rather than down 2.5%, etc. While not reported, our results are quite robust to perturbations in these values.

Order book breakdowns can also occur on the offer side of the book. This section studies these "breakups."

Menkveld and Wang (2011) note that liquidity risks are particularly acute for small cap stocks. We want to make sure that our results are not being driven solely by a set of volatile, low-liquidity securities. We repeat our explanatory models for a purely large market capitalization sample.

We also consider here alternative microstructure definitions of our 10% decline by looking first at the NBBO. We also look at a less familiar but related concept, the worst bid or offer (WBO). Our results are quite strong and similar for the NBBO, but the WBO appears to be a challenge for any model.

In addition, we analyze market quality events using trade data rather than quotes. The results confirm our explanatory model.

---

[17]The AIC suggests including up to 8 lags, but the lags beyond the second are not statistically significant.

3.10.1 Breakups

Market quality deterioration also results in rapid increases in offer prices. These are surprisingly as frequent as the breakdowns and tend to follow the same pattern over time in Figure 3.7.

[INSERT Figure 3.7 HERE]

The average breakup frequency is 0.63% in our sample from 1993-2011.

We effectively split the sample on October 15, 2007 using a dummy variable d^{NMS} for the full implementation of Reg. NMS. Breakups are more frequent than breakdowns after 2002, and we do not find a statistically significant decrease after Reg. NMS in Table 3.7.

[INSERT Table 3.7 HERE]

The effect of market correlation spikes is qualitatively similar to the results we found in Table 3.2. The results are not as strong as they are for breakdowns, with correlation spikes raising the breakup probability by 8.915%. The estimated effect is statistically significant at the 6% level.

3.10.2 Large caps

The literature often focuses on large cap stocks when studying equities. We want to explore whether our conclusions would change if we include only large caps in our sample. We construct a sub-sample that contains stocks with market capitalization of more than $10 billion. On average, there are 246 stocks per day in this category.

We plot the aggregate breakdown frequency in Figure 3.7. On average during 1993-2011, the large cap daily breakdown frequency is 0.11% versus 0.64% for all stocks. Large caps break down at a lower frequency in each year of the sample. Only 118 large caps have broken down in the period 2009-11. 77 of these occur during the Flash Crash.

Our explanatory model explains the breakdown frequency of large caps as well. The results are reported in Table 3.7. Adding the Reg. NMS dummy variable to the model reveals a significant reduction in breakdown frequency of 91.81% since October 15, 2007. A correlation spike raises large cap breakdowns by 121.9%.

3.10.3 NBBO

We plot the frequency of breakdowns based on the NBBO in Figure 3.7. [18] After 2007, the NBBO breakdown frequency averages 31.96% lower than the frequency based on the listing exchange.

Market quality breakdowns in the NBBO are less frequent in 2011 than they were in 1993 when exchanges dominated liquidity in their own listings. From 2008-11, secondary exchanges were able to provide a liquidity buffer in cases where the primary exchange was experiencing a breakdown. This supports the conclusion of Jiang, McInish, and Upson (2014) that competition could enhance market quality.

With our explanatory model, the change in definition confirms, in Table 3.7, our previous conclusions. NBBO breakdowns have decreased significantly since Reg. NMS by 60.61%. A correlation spike raises the NBBO breakdown frequency by 17.62%.

We feel that looking at the primary listing exchange for market quality effects is quite natural, but these results show that using the NBBO would result in nearly identical conclusions.

3.10.4 Trading events

We also apply our filter to trade prices as another robustness check. Specifically, we look at trade price movements from 09:35-15:55 in the listing exchange. A stock is identified to experience a breakdown if the trade prices fall 10% below the 09:35 price and rebound to

[18] Because the national best bid price may be higher than the listing exchange price at 09:35 or 15:55, the NBBO breakdown frequency can be larger than the same measure computed solely on the listing exchange.

within 2.5% of the 09:35 price by 15:55. In addition, the trade with the lowest price must be repeated at least once in a subsequent calendar second.

Figure 3.7 presents the frequency of trading breakdowns on the primary listing exchange. The trading events occur relatively more frequently in the 1990s. Consistent with the results based on the NBBO, breakdowns are less frequent in 2011 than they were in 1993.

The results using our explanatory model are shown in Table 3.7. Trading breakdowns have decreased significantly since Reg. NMS by 79.91%. A correlation spike raises the trading breakdown frequency by 31.19%.

3.10.5 WBO

There is a lot more exchange competition in the latter half of our sample, but that does not mean that every exchange plays an active role in liquidity provision for every stock. In particular, the presence of stub quotes of $0.01 is important on the non-primary exchanges. These come into play on many days, not just on the Flash Crash.

Breakdown frequencies grow virtually monotonically in Figure 3.8 between 1997 and 2006.

[INSERT Figure 3.8 HERE]

The frequencies are also nearly 100 times higher. In 2010, stocks on the primary listing exchange break down 0.36% of the time, whereas the WBO breaks down nearly 35% of the time.

Not surprisingly, our model offers no explanatory power for these events. In Table 3.7, volume and correlation offer no improvement to the likelihood of the breakdown frequency.

We don't feel that the WBO is a proper measure of overall market quality, but we still think it provides some perspective on quote activity away from the listing exchange or the NBBO.

3.11 Conclusion

Market quality, in our view, should be assessed using quotes as well as trade prices. We analyze the intra-daily consolidated bids and offers of every security in CRSP and TAQ during the period of 1993-2011. We examine stocks which fall more than 10% between 09:35 and 15:55 but recover within the day. These market quality breakdowns have a daily average frequency of 0.64%, approximately 44 stocks per day. Breakdowns in 2010-2011 average 0.38%, which is less than once per year in a typical stock.

Volume and volatility are still the prime causes of market quality breakdowns in our explanatory model, improving the likelihood by more than 40% over a model with just a constant term. Market quality has improved since the passage of Reg. NMS. A stock is 41.78% less likely to breakdown after mid-October 2007. We find no impact on breakdowns from a Herfindahl index of quote fragmentation or the market share of off-exchange volume. The NYSE, despite many changes and losses of market share, still has the highest market quality. ETFs break down 90.33% more frequently than non-ETFs.

We confirm our theoretical model of correlated liquidity shocks. Breakdowns are 25.62% more frequent when correlation between market sectors spikes. ETF and HFT trading volume Granger cause this correlation. Surges in HFT activity raises the breakdown frequency by 18.33%.

Breakdown effects are persistent up to two days. Lagged factors improve upon a constant forecast by up to 50%.

Stocks with market capitalization of more than $10 billion have an average daily breakdown frequency of only 0.11%. During 2009-11, only 118 large cap stocks break down.

In 2008-11, NBBO quote activity away from the listing markets lowers the average breakdown frequency by 31.96%. This suggests that exchange competition may be beneficial to market quality.

Table 3.1: Market Quality Breakdowns on May 6, 2010

This table gives the number of breakdowns in market quality on listing exchanges on May 6, 2010. Our sample is selected from all securities listed on the NYSE, NASDAQ, ARCA, and AMEX. We remove the securities that are not included in the Center for Research in Security Prices (CRSP) on May 6, 2010. Those include preferred, warrants, and units bundled with warrants. We apply our filter, as discussed in section 4.4, to identify market quality breakdowns for individual stocks.

	Listings	Breakdowns	Frequency
Total	6,527	1,857	28.45%
NYSE	2,382	691	29.01%
NASDAQ	2,821	582	20.63%
ARCA	829	521	62.85%
AMEX	495	63	12.73%

Table 3.2: Aggregate Breakdown Frequency Models 1993-2011

This table presents estimates from the aggregate breakdown frequency models and shows the impact of changes in market structure and correlation on breakdowns. The dependent variable π_t is the daily breakdown frequency from April 6, 1993 to December 30, 2011. Given the fact that it is close to zero and not normally distributed, we use the generalized linear model with the assumption that π_t follows the gamma distribution. The models are estimated by quasi-maximum likelihood method using robust standard errors. Column (1) shows estimates for the baseline model: $\log(\mathrm{E}[\pi_t]) = \alpha + \beta_1 \sigma_t^{VIXopen} + \beta_2 \widetilde{v}_t$. $\sigma_t^{VIXopen}$ is the opening value of the VIX. $\widetilde{v}_t$ represents volume spikes, and equals 1 if the volume becomes significantly higher than the average of proceeding 20 days at the 5% significance level and 0 otherwise. The volume is measured as the sum of trading activity on each exchange in its own listings. R_M^2 denotes McFadden's R-squared. Column (2) displays the effect of Reg. NMS after the rules were fully adopted on October 15, 2007, by including a dummy variable d_t^{NMS}. We measure fragmentation in two ways, the Herfindahl index and the share of consolidated volume executed in trade reporting facilities (TRFs), as discussed in section 4.3. Consistent with the volume variable, $\widetilde{H}_t$ and $\widetilde{TRF}_t$ represent spikes of market fragmentation, respectively. They equal 1 if the market is significantly more fragmented compared to the proceeding 20 days and 0 otherwise. The Herfindahl index is in column (3) and the TRF volume in column (4). The TRF data are available from March 5, 2007 to December 30, 2011. The R_M^2 for the baseline model in this sub-sample is 0.5168. Column (5) shows the impact of market correlation on the breakdown frequency. We calculate the 20-day rolling correlation pairwise by using daily returns of 30 Fama-French industry portfolios and then use the median as a measure of the market correlation. Consistent with the volume variable and fragmentation measures, $\widetilde{\rho}_t$ represent spikes of market correlation.

	(1) Baseline	(2) Reg. NMS	(3) Herfindahl	(4) TRF%[†]	(5) Correlation
$\sigma_t^{VIXopen}$	0.0749	0.0845	0.0845	0.0668	0.0841
	(47.66)	(49.97)	(50.19)	(27.35)	(50.02)
$\widetilde{v}_t$	0.5396	0.5932	0.5936	0.9207	0.5412
	(4.44)	(4.23)	(4.23)	(3.32)	(4.48)
d_t^{NMS}		-0.5410	-0.5416		-0.5463
		(-13.27)	(-13.37)		(-14.36)
$\widetilde{H}_t$			-0.0137		
			(-0.27)		
$\widetilde{TRF}_t$				0.1706	
				(1.61)	
$\widetilde{\rho}_t$					0.2281
					(3.21)
constant	-2.2814	-2.3848	-2.3843	-2.5000	-2.3947
	(-70.07)	(-71.74)	(-71.37)	(-36.58)	(-72.07)
R_M^2	0.4222	0.4587	0.4587	0.5185	0.4624

t-statistics in parentheses.

† denotes sample period from March 5, 2007 to December 30, 2011.

Table 3.3: Models for Individual Stocks 1993-2011

This table contains estimates from pooled panel regression models at the individual stock level. We analyze monthly data and the dependent variable $n_{i,t}$ is the number of breakdowns on stock i in month t. The sample period is from April 1993 to December 2011. We estimate the models by Poisson regression with robust standard errors using quasi-maximum likelihood method. Column (1) shows estimates for the baseline model: $\log(\mathrm{E}[n_{i,t}]) = \alpha + \beta_1 p_{i,t}^{open} + \beta_2 \kappa_{i,t} + \beta_3 \sigma_{i,t} + \beta_4 \widetilde{v}_{i,t}$, where i is the subscript representing each stock and t is the subscript for each month. $p_{i,t}^{open}$ is the log opening price of the stock, $\kappa_{i,t}$ is the log market capitalization, $\sigma_{i,t}$ is the monthly volatility, and $\widetilde{v}_{i,t}$ is volume spikes on the listing exchange. Column (2) presents the exchange effects on breakdowns by including three dummy variables, $d_{i,t}^{NYSE}$ for the NYSE, $d_{i,t}^{NASD}$ for the NASDAQ, and $d_{i,t}^{ARCA}$ for the ARCA. The AMEX is used as the basis in the model. Column (3) shows market quality breakdowns in ETFs by including a dummy variable d_i^{ETF}. The estimates for the model with both exchange listing and ETFs are reported in column (4). R_M^2 denotes McFadden's R-squared.

	(1) Baseline	(2) Exchange	(3) ETF	(4) All
$p_{i,t}^{open}$	-0.3466	-0.3567	-0.4105	-0.3803
	(-51.71)	(-53.91)	(-59.33)	(-54.30)
$\kappa_{i,t}$	-0.3193	-0.2942	-0.2782	-0.2792
	(-30.53)	(-26.79)	(-27.38)	(-25.25)
$\sigma_{i,t}$	0.2557	0.2546	0.2475	0.2514
	(9.97)	(9.76)	(9.49)	(9.60)
$\widetilde{v}_{i,t}$	0.2146	0.2264	0.2026	0.2195
	(38.35)	(28.28)	(27.11)	(27.37)
$d_{i,t}^{NYSE}$		-0.5783		-0.5305
		(-44.47)		(-39.43)
$d_{i,t}^{NASD}$		-0.3548		-0.3072
		(-34.71)		(-27.76)
$d_{i,t}^{ARCA}$		0.5942		0.0680
		(18.02)		(1.55)
d_i^{ETF}			1.0071	0.6436
			(38.02)	(15.50)
constant	-0.8051	-0.8949	-1.0329	-0.9826
	(-17.60)	(-19.06)	(-22.45)	(-20.33)
R_M^2	0.1974	0.2035	0.2016	0.2043

t-statistics in parentheses.

Table 3.4: Granger Causality Tests

This table demonstrates the relationship between market correlation, exchanged traded funds (ETFs) and high frequency trading (HFT) by Granger causality tests. Panel A shows the test results for ETFs during the sample period from 1993 to 2011. The tests are based on the vector autoregressive model in (4.18). Panel B presents the test results for HFT. The sample period is from 2008 to 2009 and the tests are based on the vector autoregressive model in (4.20).

Panel A: Exchange Traded Funds			
H_0: ETF volume does not Granger cause market correlation.			
F-stat	7.20	p-value	0.0000
H_0: market correlation does not Granger cause ETF volume.			
F-stat	1.36	p-value	0.2446

Panel B: High Frequency Trading			
H_0: HFT% does not Granger cause market correlation.			
F-stat	3.65	p-value	0.0058
H_0: market correlation does not Granger cause HFT%.			
F-stat	2.07	p-value	0.0833

Table 3.5: Aggregate Breakdown Frequency Models 2008-2009

This table presents estimates from the aggregate breakdown frequency models from 2008-2009, the period covered in the HFT data set. The dependent variable π_t is the daily breakdown frequency. We use the generalized linear model with gamma probability distribution and estimate models by quasi-maximum likelihood method using robust standard errors. Column (1) re-estimates the baseline model over the period from 2008-2009. Column (2) shows the impact of market correlation spikes $\widetilde{\rho}_t$ over the period from 2008-2009. Column (3) reports the estimates when both HFT activity and market correlation are included in the model, $\log(\mathrm{E}[\pi_t]) = \alpha + \beta_1 \sigma_t^{VIXopen} + \beta_2 \widetilde{v}_t + \beta_3 \widetilde{HFT}_t + \beta_4 \widetilde{\rho}_t$. The HFT activity is measured by the share of volume executed by HFT firms. Consistent with the volume and correlation variables, $\widetilde{HFT}_t$ is a dummy variable measuring spikes in HFT activity. R^2_M denotes McFadden's R-squared.

	(1) Baseline	(2) Correlation	(3) HFT
$\sigma_t^{VIXopen}$	0.0526	0.0528	0.0532
	(19.75)	(19.40)	(19.26)
$\widetilde{v}_t$	0.5700	0.4766	0.4408
	(3.87)	(4.22)	(3.80)
$\widetilde{HFT}_t$			0.1683
			(2.15)
$\widetilde{\rho}_t$		0.2324	0.2427
		(2.59)	(2.65)
constant	-1.7866	-1.8082	-1.8366
	(-19.14)	(-18.68)	(-18.24)
R^2_M	0.2316	0.2330	0.2340

t-statistics in parentheses.

Table 3.6: Predictive Models

This table contains estimates from a predictive models of breakdowns. We include the prior days' breakdown probabilities along with the opening value of the VIX, $\log(\mathrm{E}[\pi_t]) = \alpha + \beta_1 \sigma_t^{VIXopen} + \sum_{j=1}^{2} \eta_j \pi_{t-j}$. R^2_M denotes McFadden's R-squared.

	Breakdown
$\sigma_t^{VIXopen}$	0.0388
	(10.64)
π_{t-1}	0.3519
	(3.74)
π_{t-2}	0.3059
	(3.33)
constant	-1.9544
	(-48.12)
R^2_M	0.4936

t-statistics in parentheses.

Table 3.7: Robustness Checks for Aggregate Frequency Models

This table presents estimates for a variety of robustness checks for our explanatory model (3.15). Column (1) examines "breakups," rapid increases in offer prices from the listing exchange that are subsequently reversed. Columns (2)-(5) look at breakdown frequency. Column (2) is a sample of stocks with market capitalization of more than $10 billion. In column (3), the breakdown frequency is based on the NBBO rather than the listing exchange. Column (4) uses trades rather than quotes as the breakdown measure. In column (5), the frequency is based on the WBO on any exchange. We use the generalized linear model with gamma probability distribution, and the models are estimated by quasi-maximum likelihood method using robust standard errors. R_M^2 denotes McFadden's R-squared.

	(1) Breakups	(2) Large Caps	(3) NBBO	(4) Trades	(5) WBO
$\sigma_t^{VIXopen}$	0.0686	0.1325	0.0836	0.0804	0.0039
	(47.86)	(13.69)	(51.86)	(40.67)	(2.76)
$\widetilde{v}_t$	0.2094	1.9653	0.5138	0.6829	0.0623
	(4.70)	(4.32)	(5.30)	(4.30)	(1.19)
d_t^{NMS}	0.0100	-2.5028	-0.9316	-1.6050	0.6200
	(0.33)	(-10.05)	(-30.12)	(-36.46)	(28.36)
$\widetilde{\rho}_t$	0.0854	0.7971	0.1623	0.2715	0.0584
	(1.94)	(2.99)	(3.16)	(2.84)	(1.33)
constant	-2.1231	-5.3716	-2.3821	-2.2085	2.5283
	(-73.12)	(-27.01)	(-78.69)	(-58.91)	(72.74)
R_M^2	0.3721	0.6466	0.5442	0.6095	0.0111

t-statistics in parentheses.

Figure 3.1: Market Quality Metrics on May 6, 2010

This figure presents the market quality metrics on May 6, 2010 by our filter. For comparison with the CFTC-SEC results, we apply the same stock filter here, i.e. a share price of more than $3.00 and a market capitalization of at least $10 million. Panel A shows the distribution of the percentage decline in the best bid prices on the Flash Crash day. Panel B displays the time and percentage decline of the best bid from 14:00 to 15:00. Each point on the graph represents a stock.

Panel A: Distribution of % Decline in the Best Bid

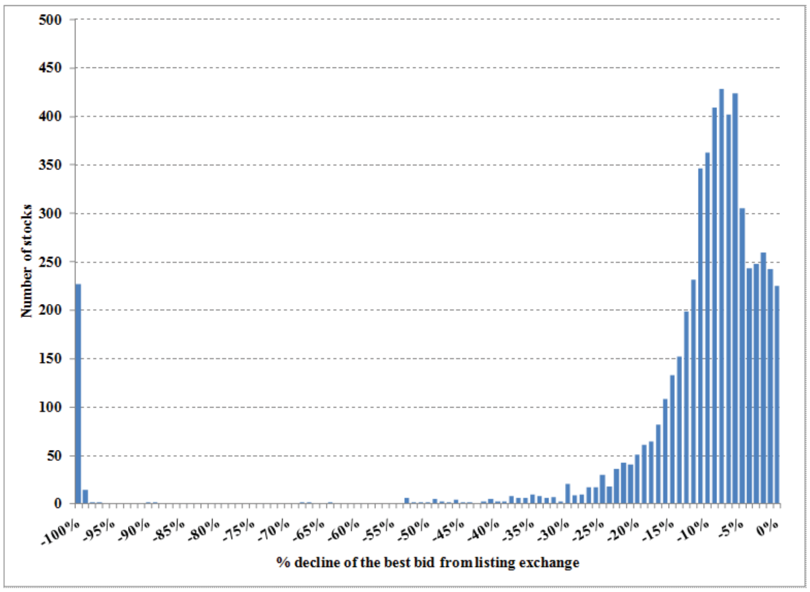

Panel B: Timing of Lows in Best Bid 14:00-15:00

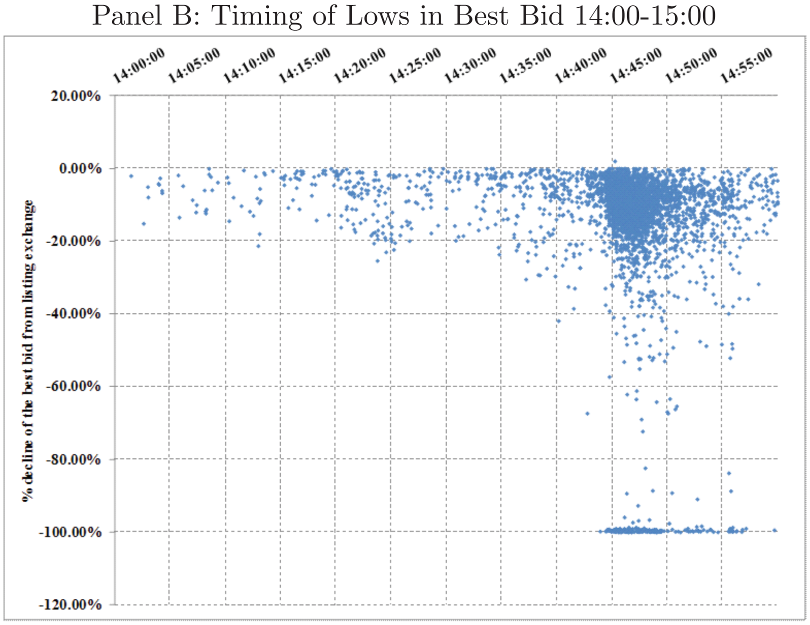

Figure 3.2: Market Quality of ETFs on May 6, 2010

This figure presents the market quality of ETFs on May 6, 2010 by our filter. Panel A shows the distribution of the percentage decline in the best bid prices for a sample of all ETFs. Panel B displays the time and percentage decline of the best bid from 14:00 to 15:00. Each point on the graph represents an ETF.

Panel A: Distribution of ETF % Decline in the Best Bid

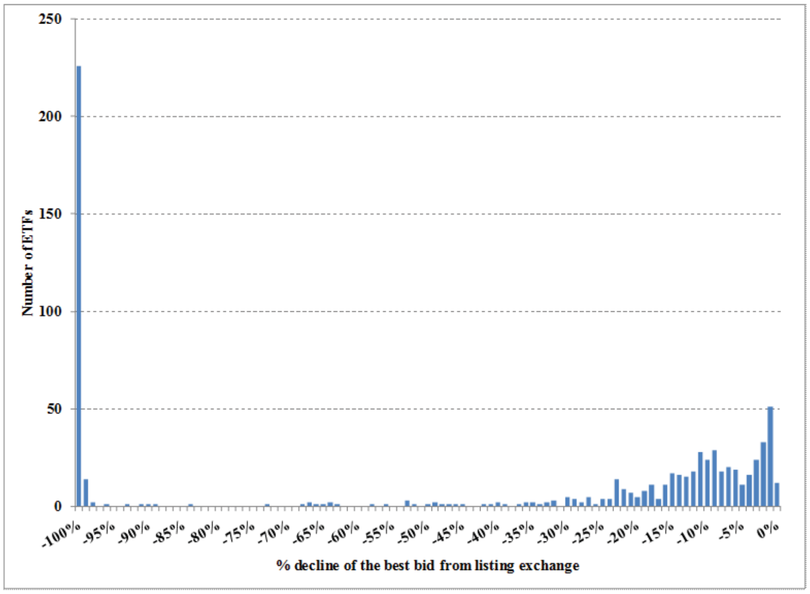

Panel B: Timing of ETF Lows in Best Bid 14:00-15:00

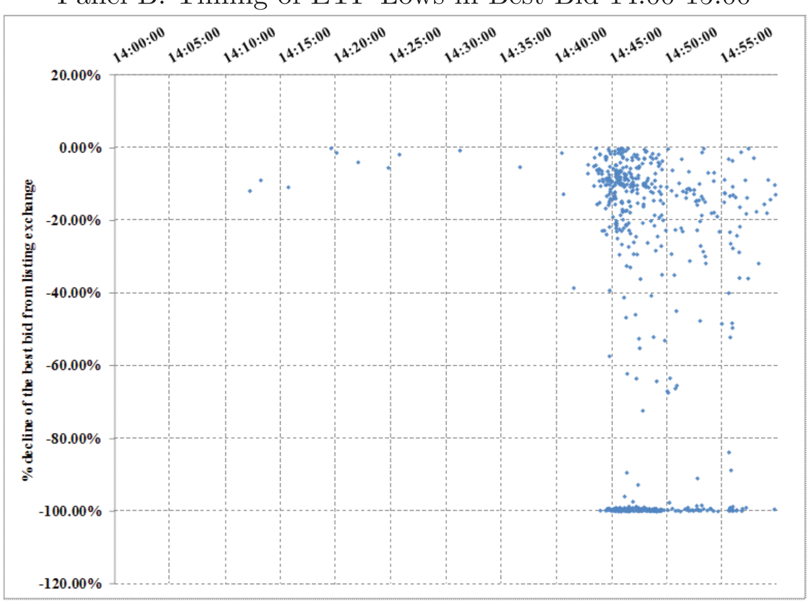

Figure 3.3: Breakdown Frequency 1993-2011

This figure presents market quality breakdowns from 1993 to 2011. Panel A shows the breakdown frequency, which is calculated by the number of breakdowns divided by the total number of securities. Panel B plots the number of breakdowns in each year.

Panel A: Market Quality Breakdown Frequency

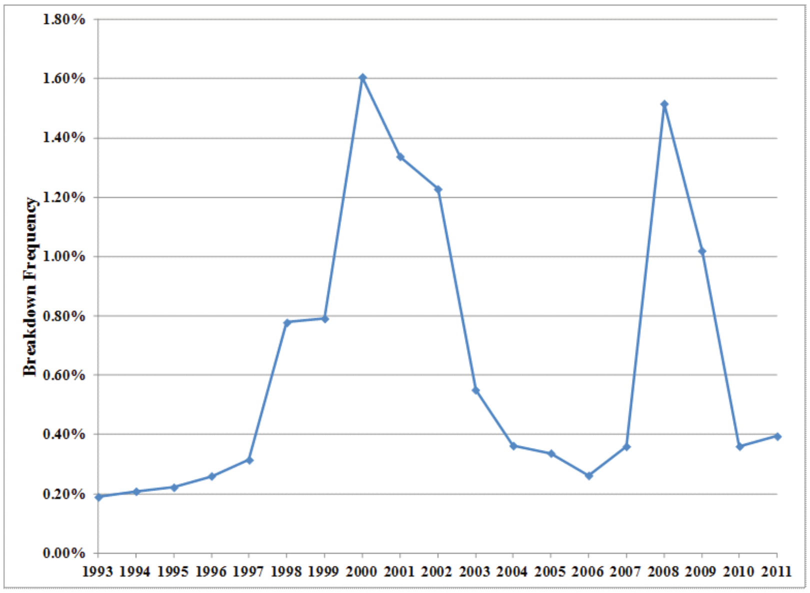

Panel B: Number of Market Quality Breakdowns

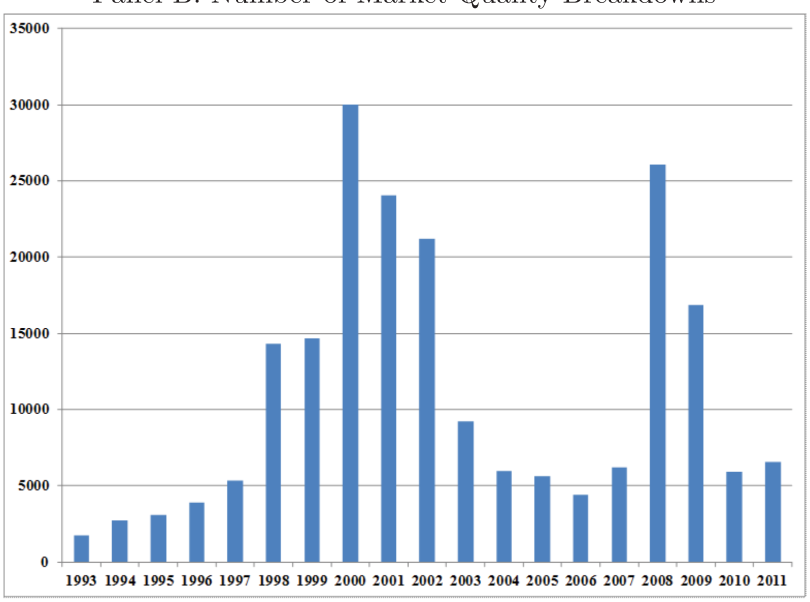

Figure 3.4: Breakdown Inequality by Gini Coefficient 1993-2011
This figure demonstrates the inequality of breakdown incidence among individual stocks by the Gini coefficient. For each year we include the securities that experience at least one breakdown in market quality.

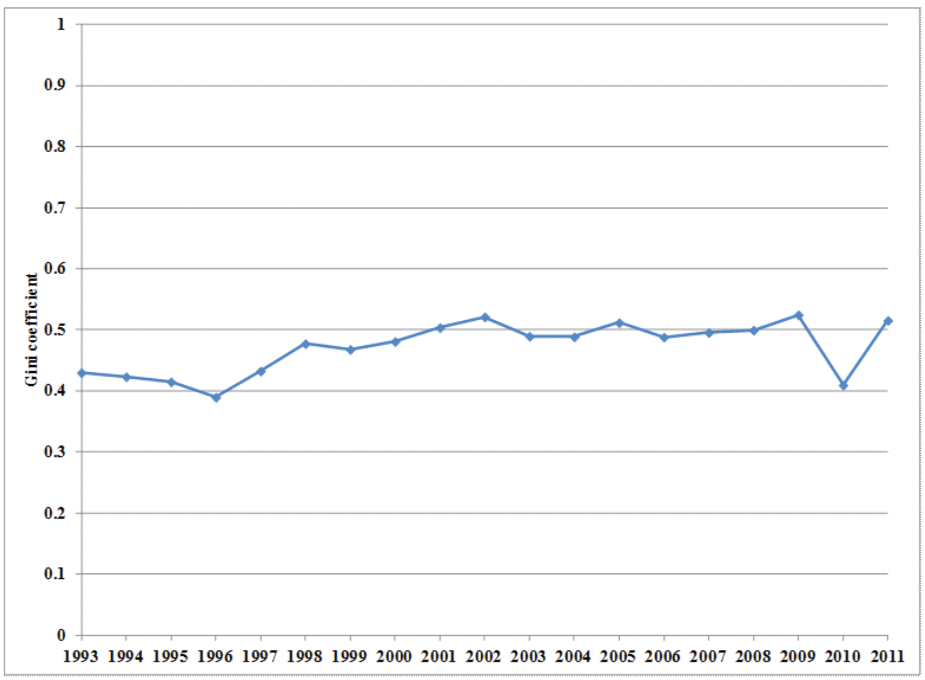

Figure 3.5: Market Quality Breakdown Frequency by Exchange 1993-2011
This figure presents the breakdown frequencies from 1993 to 2011 on each of the four listing exchanges, NYSE, NASDAQ, AMEX, and ARCA. ARCA becomes a listing exchange in 2006 after its merger with NYSE. The breakdown frequency on each exchange is calculated by the number of breakdowns divided by the total number of securities on that exchange.

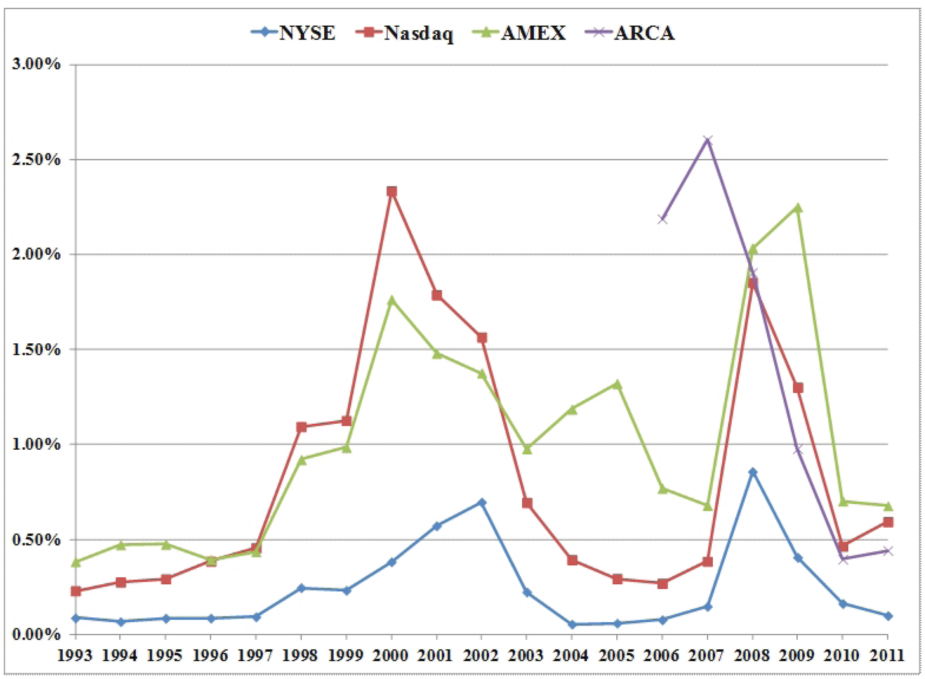

Figure 3.6: Rolling Correlation of Industry Portfolios 1993-2011

This figure illustrates the market correlation from 1993 to 2011. We calculate the 20-day rolling correlation pairwise using daily returns of 30 Fama-French industry portfolios, and then use the median as a measure of the market correlation.

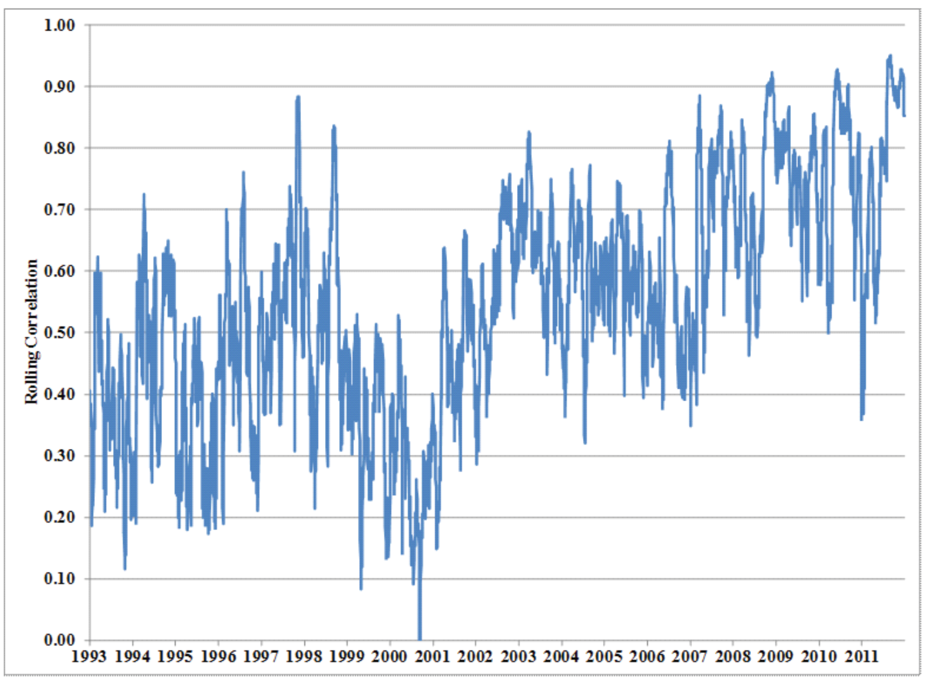

Figure 3.7: Robustness Checks for Breakdown Frequency 1993-2011
This figure presents event frequencies from 1993 to 2011 for alternative measures of market quality. We plot "breakups," rapid increases in offer prices from the listing exchange that are subsequently reversed. We also plot three breakdown frequencies: (1) a sample of stocks with market capitalization of more than $10 billion, "Large caps"; (2) the national best bid or offer (NBBO); and (3) trade prices ("Trading events").

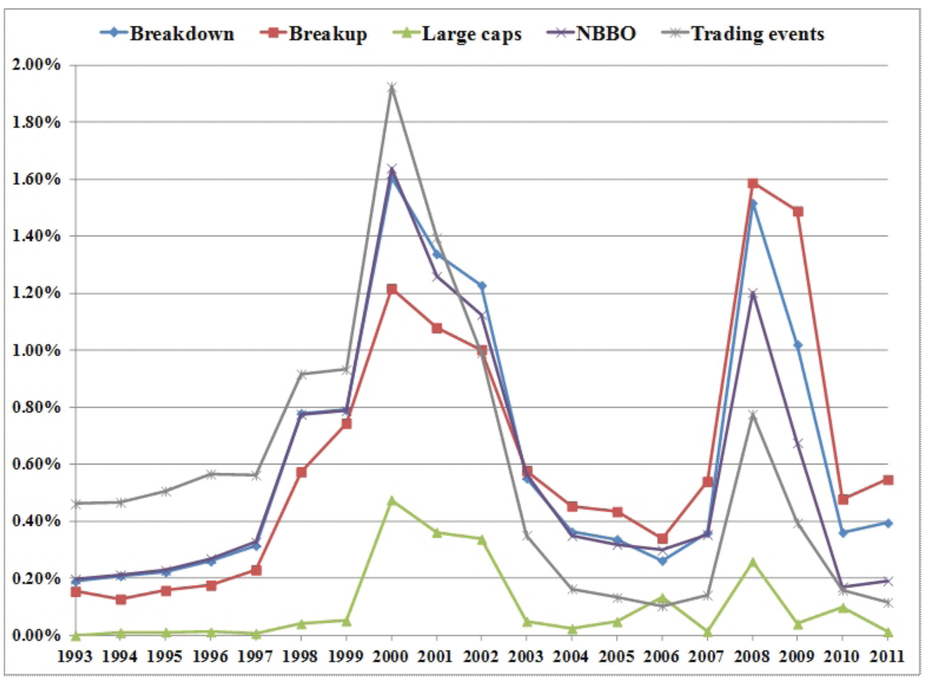

Figure 3.8: WBO Breakdown Frequency 1993-2011
This figure presents market quality breakdown frequencies in the worst bid or offer (WBO) from 1993 to 2011.

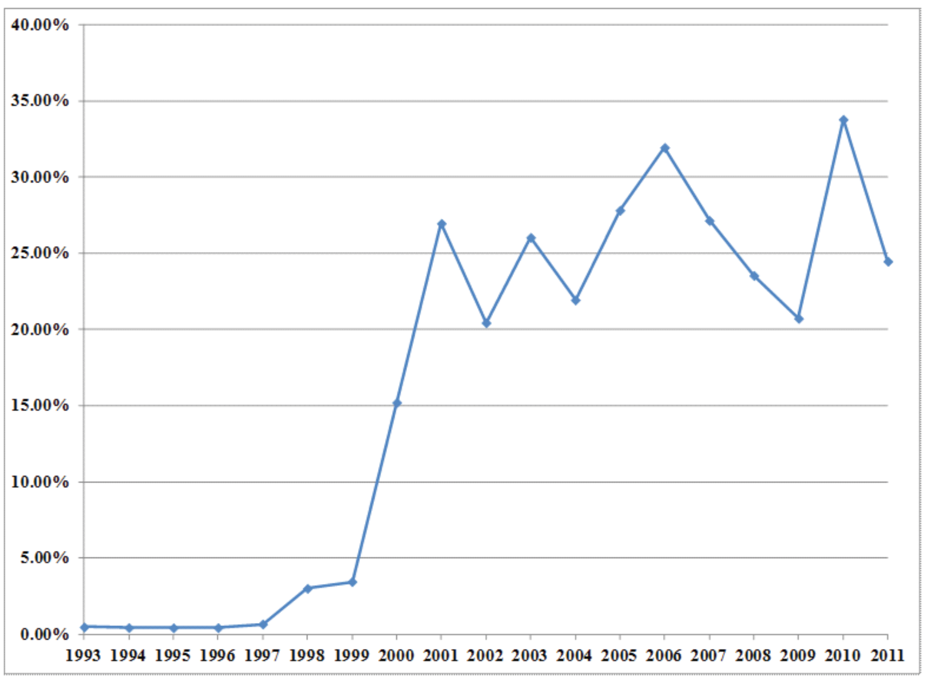

Chapter 4

High Frequency Trading in the Equity Markets During U.S. Treasury POMO

Joint with Bruce Mizrach

4.1 Introduction

High frequency trading (HFT) represents more than 70% of equity market trading volume, according to industry estimates from Sussman, Tabb, and Iati (2009). The impact of HFT on the equity markets has become a central question in the policy debate about market structure and in the academic literature on market microstructure. The policy debate has been motivated in part by the "Flash Crash" of May 6, 2010 when over 200 stocks traded down to a penny bids before the market quickly rebounded. The Commodity Futures Trading Commission and Securities and Exchange Commission (2010) task force report analyzed HFT activity from the 12 largest firms during the crash. Half significantly curtailed their trading activity during the crash including two firms that stopped trading for the rest of the day.

The academic literature has focused on the effect of HFT firms on liquidity provision. Brogaard (2011) analyzes a data set that explicitly identifies HFT activity. He concludes that HFT firms improve market quality on average. Hasbrouck and Saar (2013) suggest that HFT activity improves traditional market quality measures such as short-term volatility, spreads, and displayed depth in the limit order book. Brogaard, Hendershott, and Riordan (2014) argue that HFT increases price efficiency through their marketable orders. Menkveld

(2013) analyzes the arrival of the Chi-X high frequency platform in Europe and concludes that HFT firms act as market makers in the new market. Zhang (2010) finds more harmful effects. He observes that HFT is positively correlated with stock price volatility and hinders the ability of the market prices to reflect fundamental information. Cvitanic and Kirilenko (2010) provide a theoretical perspective and show that HFT activity effects volume and the distribution of transaction prices. Hirschey (2013) provides evidence that HFT firms anticipate the order flow from non-HFT investors and their aggressive trades are highly correlated with future returns.

While HFT firms are generally deemed to be passive liquidity providers, we find that they act as trade initiators in nearly 47% of trades in normal times. High frequency traders appear to have superior information. Whether they are at the active or passive side, the trades are more profitable when the counterpart is a non-HFT firm rather than a HFT firm.

The "Flash Crash" helps to clarify why reporting the average effect of HFT firms on the market may provide a misleading portrait of their contribution to market quality. Analyzing their impact when the market is under stress or reacting to news needs to be isolated from their contribution during less turbulent periods.

Brogaard (2011) does find that HFT firms reduce their liquidity provision during volatile times of the trading day. He claims that they increase their trading activity during periods of exogenous volatility. He looks at days surrounding firms' quarterly earnings announcements and the week of the Lehman Brothers failure. Both approaches show HFTs tend to increase their trading during times of exogenous volatility. We re-examine Brogaard's claim by looking at another period of potential market stress, the U.S. Treasury purchases made by the Fed beginning in late 2008 as part of its quantitative easing program.

The Federal Reserve's asset purchase program began in November 2008 with $600 billion of GSE debt and mortgage backed securities. In March 2009, the Federal Open Market Committee enlarged these programs and authorized purchases of $300 billion in long-term Treasuries.

The goal of the large-scale asset purchases was "an effort to drive down private borrowing rates, particularly at longer maturities." Using an event study, Gagnon, Raskin, Remache, and Sack (2011) conclude that 10-year Treasury bond yields fell 91 basis points and that 10-year agency debt yields declined 156 basis points. Joyce, Lasaosa, Stevens, and Tong (2011) find that a program of similar scale in the U.K. lowered gilt yields by 100 basis points.

Central bank asset purchases can also have impact on related asset markets. Neely (2010) and Joyce, Lasaosa, Stevens, and Tong (2011) have both emphasized the portfolio balance channel in which declining exposure to Treasury also raises other asset prices. Neely (2010) shows that announcements related to the U.S. asset purchase program also lowered 10-year government bond yields in Australia, Canada, Germany, Japan and the U.K. between 19 and 78 basis points. Krishnamurthy and Vissing-Jorgensen (2011) estimate that U.S. corporate bond yields fell between 43 and 130 basis points. Neely (2010) also finds evidence for reallocation into the U.S. stock market: the S&P 500 index rises a cumulative 3.42%.

Even though the size of Fed's overall program was largely known by to the market by March 2009, the specific securities they would buy and the bids they would accept at each auction were not. The participation levels and prices paid, just like any auction, reveal information to the markets.

We develop a theoretical model in which high frequency trading firms receive valuable private information before other traders. This information can lead HFT firms to temporarily abandon their role as liquidity providers and trade aggressively in the direction of the news.

Our model predicts that HFT firms should reduce their inside quoting activity. We find that during Treasury POMO auctions HFT firms reduce their inside bid participation by 8%.

The model also predicts that HFT firms should trade more aggressively against non-HFT participants in the direction of the news. We do find that HFT firms buy (sell) more frequently using aggressive trades with positive (negative) news. We also find that HFT firms are less likely to supply liquidity passively to non-HFT firms trading in the direction of news. These results are even stronger once we control for microstructure effects.

Consistent with Biais, Foucault, and Moinas (2013), we find that the release of private information raises market impact. A $1,000$ share order moves the market on average \$0.0318, but on POMO days this rises to \$0.0341.

High frequency traders are able to generate the most profit from private information because of their ability to trade quickly. HFT firms are consistently profitable trading during POMO events. They are profitable 88% of the time on aggressive trades and 100% of the time on passive trades. We estimate a daily average profit per stock of $\$1,300.35$ which rises to $\$1,895.37$ on POMO days. The profits per share from aggressive trading rise 300%. Extrapolating these results to the market as a whole, we estimate profits of more than \$105 million.

This chapter is organized as follows. Section 4.2 examines the POMO purchases by the Federal Reserve. We develop a theoretical model to explain HFT behavior when private information is released in Section 4.3. Section 4.4 describes the HFT data set. In Sections 4.5, we provide empirical support for the model by examining HFT quote and trade activities, market impact and profits during POMO. We perform robustness checks in Section 4.6, and Section 4.7 concludes.

4.2 Permanent Open Market Operations (POMO)

Economists renewed their interest in the zero interest rate lower bound during Japan's prolonged slump. Krugman (1998) revived discussion of the liquidity trap, and Ben Bernanke, as a Governor on the Federal Reserve Board, considered the effectiveness of a money-financed tax cut,[1] earning the nickname, "Helicopter Ben." Eggertsson and Woodford (2003)

---

[1]Remarks by Governor Ben S. Bernanke Before the National Economists Club, Washington, D.C. November 21, 2002.

have suggested that the zero-lower bound could be combated by expectations. Their prescription, a commitment to a long-term lowering of interest rates, has influenced the Fed's pledge to keep interest rates low for an "extended period."

The use of asset purchases to lower long-term interest rates was used in the 1960's as part of "Operation Twist," a U.S. effort to shorten the average maturity of the government debt. Roley (1982) found that the program lowered yields on both government and corporate debt. The Bank of Japan also dramatically expanded its balance sheet. Between 2001 and 2006, Shiratsuka (2010) reports, the central bank progressively added 35 trillion yen (roughly $350 billion) to its reserves through asset purchases.

The Federal Reserve implements its monetary policy targets through the purchase and sale of Treasury securities. The unconventional policies it pursued after the Lehman bankruptcy were different in both size, scope and duration. The Fed increased reserve bank credit from $893 billion on September 4, 2008 to $2,298 billion on March 25, 2010 during (what turned out to be) the first round of quantitative easing (QE1). The Federal Reserve not only purchased U.S. Treasuries as it normally would, it also bought GSE mortgage backed securities and debt. Because these assets were intended to remain on the balance sheet for an extended period, they were called "permanent" open market operations (POMO). The history and motivation of the Federal Reserve program is analyzed in Gagnon, Raskin, Remache, and Sack (2011).

We focus on U.S. Treasury purchases because Treasury securities play a unique role in the asset markets. We are motivated by the work of Lou, Yan, and Zhang (2013) who find that regularly scheduled Treasury auctions have price impacts on the Treasury, repo and equity markets. They link these effects to the limited risk bearing capacity of the primary dealers and fund flows influenced by temporary price distortions. Krishnamurthy and Vissing-Jorgensen (2010) highlight the unique role of Treasuries and note that changes in their supply effect the equilibrium price of liquidity and safety.

4.2.1 U.S. Treasury Market Purchases

U.S. Treasury security purchases began in March 2009. The Fed purchased $295.4 billion in 57 auctions in which dealers offered $1,137 billion for sale. Maturities ranged from 2 to 30 years for 160 different CUSIPs. This represented about 3% of the outstanding Treasury debt, and about 8% of the available Treasury supply. There are three purchases of Treasury Inflation Protected securities (TIPS) totaling $4.5 billion, but we did not include them in our analysis.

D'Amico and King (2010) provide details on the implementation of the Treasury purchases. On every other Wednesday, the Open Market Desk at the Federal Reserve Bank of New York would announce the range of the yield curve they were purchasing and the dates on which bids could be submitted. At 10:15 AM on each auction day, the Fed would publish a list of CUSIPs that it would consider purchasing. Most days, the bidding would commence at 10:30 AM. Shortly after bidding closed at 11:00 AM, the Fed used a confidential algorithm to determine which bids to accept.

Table 4.1 provides details on the first Treasury purchase on March 25, 2009. The Federal Reserve announced that it would consider purchasing securities with maturities between February 29, 2016 and February 15, 2019.

[INSERT Table 4.1 HERE]

It listed 18 CUSIPs in this maturity range, but excluded one security, the 5.125% note, maturing on June 15, 2016. On March 25, they accepted bids on 13 of the securities, buying $7.5 billion overall. This was 31% of the $21.9 billion submitted.

These auctions results show that while the total amount of assets to be purchased was known prior to the auction, the maturity composition of Treasury purchases did provide news to the market. We then try to model in the next section how HFT firms might alter their trading activity upon receipt of this news.

4.3 The Model

4.3.1 Model Setup

Consider a risky security, with the terminal value V, that changes from its initial value V_0 based on random innovations $\varepsilon \sim N(0, \sigma^2)$ and fundamental information,

$$V = V_0 + \varepsilon + \eta. \tag{4.1}$$

η is the expected change in the fundamental value due to the information arrival of POMO auctions. η is assumed to be independent from ε and can take three values: $\eta = +\delta > 0$ if the news is positive, $\eta = -\delta$ if the news is negative, and $\eta = 0$ if no information arrives.

There are three types of traders in our model: noise traders (NTs), limit order traders (LOTs), and high frequency traders (HFTs). NTs submit orders for liquidity reasons and use market orders that hit the bid or offer on the limit order book. We assume that a noise trader arrives exogenously with probability θ, and will submit a buy order with probability γ or a sell order with probability $1 - \gamma$.

LOTs provide liquidity by placing bid and offer quotes competitively. HFTs are profit maximizing. They trade either passively to earn the bid-ask spread by submitting limit orders or aggressively to realize a positioning profit using marketable limit orders. We assume in our model that HFTs trade faster than NTs and LOTs in the sense that they are more quickly informed of the value of η than noise traders and limit order traders.

To simplify the analysis we assume that all orders by each type of traders are for one unit of the security. Because the order flow of noise traders is exogenous, we only need to focus on two players: LOTs and HFTs. We then analyze their decision problems under different market conditions.

4.3.2 Limit Order Traders

LOTs do not observe the value of η, but they infer its value based on trading activity. Their conjecture about the probability distribution of η: $\eta = +\delta$ with probability α, $\eta = -\delta$ with probability β, and therefore $\eta = 0$ with probability $1 - \alpha - \beta$. The unconditional expectation of η is calculated as $\overline{\eta} = \delta\left(\alpha - \beta\right)$. LOTs post bid and offer quotes at B and A respectively, and are aware that HFTs have superior information about η. Following Glosten and Milgrom (1985), we consider buys and sells separately. Given that other traders buy, the expected profit of LOTs is

$$\mathrm{E}\left[\pi_{LOT}|\mathrm{Buy}\right] = A - \mathrm{E}\left[V|\mathrm{Buy}\right] = A - V_0 - \mathrm{E}\left[\eta|\mathrm{Buy}\right]. \tag{4.2}$$

To calculate the expectation of η given that the trade is a buy, we first compute the conditional probabilities of η.

$$\begin{aligned} \Pr\left(\eta = +\delta|\mathrm{Buy}\right) &= \frac{\Pr\left(\eta = +\delta, \mathrm{Buy}\right)}{\Pr\left(\mathrm{Buy}\right)} = \frac{\alpha(1+\theta\gamma)}{\alpha+\theta\gamma}, \\ \Pr\left(\eta = -\delta|\mathrm{Buy}\right) &= \frac{\Pr\left(\eta = -\delta, \mathrm{Buy}\right)}{\Pr\left(\mathrm{Buy}\right)} = \frac{\beta\theta\gamma}{\alpha+\theta\gamma}. \end{aligned} \tag{4.3}$$

We assume that competition among LOTs drives their expected profit to a positive amount c_{LOT}. Therefore, the best offer is set as

$$A = V_0 + \mathrm{E}\left[\eta|\mathrm{Buy}\right] + c_{LOT} = V_0 + \overline{\eta} + \frac{\delta\alpha\left(1-\alpha+\beta\right)}{\alpha+\theta\gamma} + c_{LOT}. \tag{4.4}$$

Similarly, we can obtain the best bid by LOTs. The conditional probabilities of η given that other traders sell is

$$\begin{aligned} \Pr\left(\eta = +\delta|\mathrm{Sell}\right) &= \frac{\Pr\left(\eta = +\delta, \mathrm{Sell}\right)}{\Pr\left(\mathrm{Sell}\right)} = \frac{\alpha\theta(1-\gamma)}{\beta+\theta(1-\gamma)}, \\ \Pr\left(\eta = -\delta|\mathrm{Sell}\right) &= \frac{\Pr\left(\eta = -\delta, \mathrm{Sell}\right)}{\Pr\left(\mathrm{Sell}\right)} = \frac{\beta\left(1+\theta\left(1-\gamma\right)\right)}{\beta+\theta(1-\gamma)}. \end{aligned} \tag{4.5}$$

If the expected profit is driven to c_{LOT} by competition, the best bid is set as

$$B = V_0 + \mathrm{E}\left[\eta|\text{Sell}\right] - c_{LOT} = V_0 + \overline{\eta} - \frac{\delta\beta\left(1+\alpha-\beta\right)}{\beta+\theta(1-\gamma)} - c_{LOT}. \tag{4.6}$$

The bid-ask spread is

$$A - B = \frac{\delta\left[2\alpha\beta + \theta\left(\alpha\left(1-\gamma\right)\left(1+\beta-\alpha\right)+\beta\gamma\left(1+\alpha-\beta\right)\right)\right]}{\left(\alpha+\theta\gamma\right)\left(\beta+\theta(1-\gamma)\right)} + 2c_{LOT}. \tag{4.7}$$

It is not hard to show that the spread is always positive. As seen in (4.7), the bid-ask spread can be decomposed into two components for LOTs. The first term captures the adverse selection risk and the second one covers the inventory cost.

In the symmetric case that LOTs' conjecture on positive or negative news has equal probability, i.e. $0 \le \alpha = \beta \le \frac{1}{2}$, the bid-ask spread reduces to

$$A - B = \frac{\delta\alpha\left(2\alpha+\theta\right)}{\left(\alpha+\theta\gamma\right)\left(\alpha+\theta(1-\gamma)\right)} + 2c_{LOT}. \tag{4.8}$$

4.3.3 High Frequency Traders

HFTs maximize their profit using either limit orders or marketable limit orders. They can expect to earn the bid-ask spread on passive trades by placing limit orders. By using marketable limit orders HFTs must pay the spread. A trader may want to do so because valuable limit orders can disappear quickly given the competition from other HFTs and the cancellation of limit orders. In this way they expect to gain trading profits. We assume that HFTs are informed of the value of η under news release and have the same conjecture as LOTs about the distribution of η when no information arrives. We then study the optimal order placement decision of HFTs based on different market conditions.

When there is no news expected on POMO auctions, HFTs' conjecture about the terminal security value is

$$\mathrm{E}\left[V|\text{non-POMO}\right] = V_0 + \overline{\eta}. \tag{4.9}$$

Since it lies between the best bid and offer quotes by LOTs, they could expect a loss if they submit marketable limit orders by crossing the spread. For example, the expected profit for a buy marketable limit order would be

$$\mathrm{E}\left[V|\text{non-POMO}\right] - A = -\frac{\delta\alpha\left(1-\alpha+\beta\right)}{\alpha+\theta\gamma} - c_{LOT} < 0. \tag{4.10}$$

Instead, HFTs are better off under no expected news if they provide liquidity by posting bid and offer quotes and earn the bid-ask spread on passive trades.

HFTs place their quotes at the same bid and offer prices as the limit order traders. They are not adversely selected by other traders, so they would earn a higher expected profit than LOTs at the bid and offer quotes specified in (4.6) and (4.4). At the bid side HFTs expect to have a profit of

$$c^{B}_{HFT} = \frac{\delta\beta\left(1+\alpha-\beta\right)}{\beta+\theta(1-\gamma)} + c_{LOT}, \tag{4.11}$$

and their expected profit at the offer side would be

$$c^{A}_{HFT} = \frac{\delta\alpha\left(1-\alpha+\beta\right)}{\alpha+\theta\gamma} + c_{LOT}. \tag{4.12}$$

When a positive information of POMO auctions is expected, HFTs' conjecture about the terminal security value is

$$\mathrm{E}\left[V|\text{positive}\right] = V_0 + \delta. \tag{4.13}$$

For a marketable limit order purchase, their expected profit is

$$\mathrm{E}\left[\pi_{HFT}|\text{positive}\right] = V_0 + \delta - A = \delta\left(1-\alpha+\beta\right) - c^{A}_{HFT}. \tag{4.14}$$

It is positive when $\delta > c^{A}_{HFT}/\left(1-\alpha+\beta\right)$. This suggests that HFTs would take the profitable opportunity to buy the security at the offer quote A by LOTs when they are informed of a good news with a relatively big rise of the equity value. Although they pay the spread in this way, HFTs can earn trading profits because of their superior information about the news.

It also indicates that in this situation HFTs would withdraw their liquidity provision at the inside offer and then post a higher offer quote at $V_0 + \delta + c^A_{HFT}$.

The analysis for HFTs' order strategy with a negative expected information is similar. Their expected profit for a sell marketable limit order is

$$\mathrm{E}\left[\pi_{HFT}|\text{negative}\right] = B - (V_0 - \delta) = \delta\left(1 + \alpha - \beta\right) - c^B_{HFT}, \tag{4.15}$$

which is greater than zero when $\delta > c^B_{HFT}/\left(1 + \alpha - \beta\right)$. It suggests that HFTs would sell the security to LOTs at the bid quote B when they expect a bid drop of the equity value. In this case they would choose to scale back from the inside bid and then post a lower bid at $V_0 - \delta - c^B_{HFT}$. The theoretical results on HFTs' order placement decision are confirmed by our empirical analysis of their quoting and trading activities during the period of Treasury POMO auctions.

4.4 HFT Database

We utilize the same data set provided by NASDAQ OMX as in Chapter 2. The data tracks 120 stocks,[2] and has information at different intervals and samples about quotes and trades from 26 HFT firms.

The trade information is most complete. It includes all trades on the NASDAQ exchange during regular market hours, apart from the opening and closing crosses, from January 2008 to December 2009, plus the week of February 22-26, 2010. We begin our analysis in December 2008 with the onset of POMO activity by the Federal Reserve. This sample covers the entire time first round of asset purchases by the Federal Reserve. The data set tells whether an HFT firm initiated or filled a trade. These 26 firms are involved in 76% of

[2] The symbols of stocks are listed in Table 2.2.

all the trading activity during the period January 2008 through February 2010.

There are detailed NASDAQ order book data snapshots sampled from the first week of each quarter from January 2008 to December 2009, and then February 22-26, 2010. We observe whether an HFT firm is providing liquidity at each tier of the order book. To supplement the HFT data for our market impact analysis, we make use of the ITCH data set. ITCH provides full order book level detail for the NASDAQ market, but it does not provide any HFT information. We only analyze inside quote activity in both data sets though.

Market participants[3] and regulators[4] have been concerned about the size and scope of HFT activity in recent years. Our data set documents a growing role for HFT activity. Figure 4.1 plots the monthly average percentage of HFT trades. HFT trading activity appears to trend up in 2008, back down in 2009, before stabilizing in early 2010.

[INSERT Figure 4.1 HERE]

Another measure of HFT liquidity is the extent to which HFT firms make up the inside quote. We also graph this frequency in Figure 4.1. Inside quote activity continues to uptrend in 2009, unlike the trade series.

We will control for these trends in our analysis of the POMO auctions.

4.5 Empirical Results

We analyze HFT quote and trade activities during U.S. Treasury POMO auctions. We find that (1) HFT firms pull back as market makers during periods of information release; (2) HFT firms use information to trade aggressively in the direction of the news; (3) HFT firms provide less passive liquidity on the opposite side of the news; (4) market impact rises

[3] See e.g. Christopher Matthews, “High Frequency Trading: Wall Street’s Doomsday Machine?”, *Time Magazine*, August 8, 2012.

[4] SEC Chairman Mary Jo White, in testimony before the Senate Banking Committee on March 13, 2013, noted “..high frequency trading, complex trading algorithms, dark pools, and intricate new order types raise many questions and concerns.”

during U.S. Treasury POMO auctions; (5) HFT profits rise during POMO events. The empirical evidence provides support to our theoretical model.

4.5.1 HFT Firms Pull Back from the Inside Quote

Our model implies that, around the release of news, market makers should become more cautious. Admati and Pfleiderer (1988) have emphasized that the risk of trading against valuable private information is higher, and market makers should widen spreads and reduce their depth.

To examine this empirically, we estimate how frequently HFT firms participate in the inside quote on the NASDAQ market. Our order book overlaps with U.S. Treasury POMO auctions on 5 trading days.

We calculate the percentage of ticks in which the HFT firms is at the inside bid or offer. Interpreting this raw percentage requires some care. First, we have to account for the trend in the data that we noted in Figure 4.1. We find that a quadratic trend fits the data well.

The data are also seasonal intra-daily. The vast majority of POMO auctions occur between 10:30 and 11:00 AM. This is a relatively quiet time during the day in which HFT participation tends to fall off. Therefore, we include a time dummy for the period from 10:30-11:00 AM in the model.

We also need to control for the typical microstructure factors that influence order aggressiveness. These include realized volatility, which we measure as a ten-tick moving average, trade volume, and the order imbalance of buyer less seller initiated trades. These variables are all lagged one period.

We estimate the model in the probit form with robust standard errors on a pooled cross-section of the 120 stocks, using maximum likelihood. We report the result for the bid and offer side respectively. Table 4.2 shows that the participation rate of HFT firms in the inside quote falls significantly during Treasury purchases, whether it is a positive or negative news. The result is consistent with our model.

[INSERT Table 4.2 HERE]

HFT firms are almost 8% less likely to quote on the inside bid and 5% less frequently on the inside offer during U.S. Treasury POMO auctions. This contrasts with Brogaard, Hendershott, and Riordan (2014) finding that HFT participation increases on high volatility days.

4.5.2 HFT Firms Trade More Aggressively in the Direction of News

Given the fact that the HFT firms tend to withdraw liquidity from the inside quotes during POMO auctions, the other question to ask is whether they demand more liquidity from other non-HFT market participants. The trade data set tells us whether traders are HFT or non-HFT firms at both sides of a trade. We treat the HFT firms as a group and focus particularly on HN and NH trades, where the first letter refers to the liquidity seeker and the second to the liquidity provider. We study the trading behavior of the HFT firms when they expect a positive and a negative news respectively.

We report, in Table 4.3, the average difference of the number of HFT buyer and seller initiated trades between 10:30-11:00 AM and the rest of the day, at a 1-minute frequency.

[INSERT Table 4.3 HERE]

We divide U.S. Treasury purchases into positive and negative news events based on the 1-hour equity return after the start of the auction. The event is treated as positive news if the average return across the 120 stocks is greater than zero, and a negative one otherwise. Among 57 Treasury purchases, there are 32 positive and 25 negative news events.

We compute the average difference on non-POMO days and on days with positive and negative news from U.S. Treasury auctions. We then test the differences in these net buy counts during event and non-event periods. We find a statistically significant reduction in buyer initiated trades on negative news days, with a reduction of 345 net buy trades during the POMO period. There is an increase of 482 net buy trades on positive news days, although this result is only significant at the 10% level.

POMO announcements days are volatile periods for the market, and this should lead to a less aggressive trading posture for HFT firms. To confirm and perhaps strengthen the results in Table 4.3, we need to then control for microstructure factors. We add lagged returns, realized volatility, volume and order imbalances as before, as well as a seasonal time dummy. We also include two dummy variables for positive and negative news.

The dependent variable is the 1-minute net differential between buyer and seller initiated trades. We then estimate a least squares model for HN net buyer counts trades in Table 4.4.

[INSERT Table 4.4 HERE]

We estimate a significantly positive effect of U.S. Treasury POMO events on HN trades, indicating the more aggressive stance of the HFT firms during the auctions. Once we control for microstructure factors, HFT firms increase their net buying by 600 trades on good news days and decrease their net buying by 891 trades when there is bad news.

4.5.3 HFT Firms Reduce Their Passive Liquidity Supply

We then do the same comparison in Table 4.3 for NH trades in which HFT firms are the passive liquidity suppliers. We find that non-HFT firms reduce their net buys by 273 trades on positive news days and increase their net buys by 229 trades on bad news days. This indicates that HFT firms have become more reluctant to supply passive liquidity to noise traders trading in the direction of the news. Neither of these results is significant at the 10% level though.

Introducing microstructure variable controls helps to isolate the effects predicted by our model. When we regress NH net buyers counts, the effect of the POMO auctions becomes much more strongly significant. Non-HFT firms decrease their net buying by 478 trades on good news days and increase their net buying by 490 trades with bad news.

We have now confirmed three of the primary predictions of the model. HFT firms become less active participants in the inside market on either the bid or offer. HFT firms increase their net buying activity on good news days and decrease on bad news days. Finally, we show that non-HFT firms are not able to trade as aggressively as HFT firms in the direction of the news because the HFT firms reduce their passive liquidity supply.

4.5.4 Market Impact of Trades by HFT Firms Becomes Higher

Another measure of liquidity is the market impact of trades. This is a dynamic indicator which incorporates the bid-ask spread, market depth, the persistence in order flow, and the resiliency of the order book.

Let $r_{i,t}$ be the change in the midpoint of the bid-ask spread, $(p^b_{i,t} + p^a_{i,t})/2 - (p^b_{i,t-1} + p^a_{i,t-1})/2$. $x_{i,t} \in \{-1, +1\}$ is an indicator variable which measures the trade direction. It is assigned as $+1(-1)$ if the transaction is a buy(sell). Let $V_{i,t}$ denote the size of the trade. We follow Hasbrouck (1991) using a vector autoregressive (VAR) model of their dynamic interaction. We also use Hasbrouck's identifying assumption that the current trade can effect the current quote, but not vice versa,

$$r_{i,t} = a_{r,0} + \sum_{j=1}^{10} a_{r,j} r_{i,t-j} + \sum_{j=0}^{10} b_{r,j} x_{i,t-j} V_{i,t-j} + \varepsilon_{r,t}, \quad (4.16)$$

$$x_{i,t} V_{i,t} = a_{x,0} + \sum_{j=1}^{10} a_{x,j} r_{i,t-j} + \sum_{j=1}^{10} b_{x,j} x_{i,t-j} V_{i,t-j} + \varepsilon_{x,t}. \quad (4.17)$$

We use 10 lags in the VAR. The estimates are not sensitive to this choice.

Market impact is a dynamic process

$$\partial r_{i,t+j} / \partial x_t V_t \quad (4.18)$$

which we will now compute during POMO and non-POMO intervals. We sum the aggregate effect

$$\overline{\Lambda} = \frac{1}{120} \sum_{i=1}^{120} \sum_{j=1}^{50} \partial r_{i,t+j} / \partial x_{i,t} V_{i,t} \tag{4.19}$$

arbitrarily after 50 trades, filtering out negative impacts.

The number of POMO days we can include is limited by the availability of our NASDAQ inside quote data. We have only 5 U.S. Treasury POMO days to contrast with 14 non-POMO days. To do reasonable comparisons, we expand the sample to 14 POMO days using NASDAQ ITCH data.

Our HFT data set classifies trades into four categories. The trade has an aggressor and a passive supplier. Either can be a HFT or not. We report the comparison of average market impact by HFT trades, $x_t \in \{x_t^{HH}, x_t^{HN}, x_t^{NH}\}$, across the 120 stocks in Table 4.5.

[INSERT Table 4.5 HERE]

We find, as our quote analysis indicated, that market impact from HFT is significantly higher during U.S. Treasury POMO auctions than the corresponding period on non-POMO days. A $1,000$ share order moves the market on average \$0.0318, but on POMO days this rises to \$0.0341. The rise in market impact of trades could make the trading costs of non-HFT firms even higher.

4.5.5 HFT Firms Make More Profits During POMO

Menkveld (2013) makes a useful division of trading profits for an HFT firm. On passive trades, designated NH in our sample, they can expect to earn the bid-ask spread. On aggressive trades, designated HN, they must pay the spread, hoping to realize a positioning profit.

Under some assumptions, we can estimate the profitability of the HFT firms as a group using our trade data. We assume that HFT firms try to end the day flat and assess their profits by valuing any position at the day's average price. By construction, we consider only HN and NH trades.

The HFT daily profits for stock i are estimated as,

$$\pi_i^{HFT} = \sum_{t=1}^{T} [D_{i,t}^S p_{i,t} q_{i,t} - D_{i,t}^B p_{i,t} q_{i,t}] + \frac{\sum_{t=1}^{T} p_{i,t} q_{i,t}}{\sum_{t=1}^{T} q_{i,t}} \sum_{t=1}^{T} [D_{i,t}^B q_{i,t} - D_{i,t}^S q_{i,t}], \tag{4.20}$$

where D^B and D^S are buy and sell indicators respectively, $p_{i,t}$ is the price of stock i at time t, and $q_{i,t}$ is the quantity. It closes out open positions at the end of the day using the daily volume weighted average price. The method for calculating profits is similar to Brogaard, Hendershott, and Riordan (2014) and Baron, Brogaard, and Kirilenko (2014).

Profits per stock for POMO U.S. Treasury days are compared to profits on non-POMO days in Table 4.6.

[INSERT Table 4.6 HERE]

On non-POMO days, we estimate profits of $\$1,300.35$ per stock for the entire sample of trading days between December 2008 and February 2010. This compares to Brogaard, Hendershott, and Riordan (2014) estimate of $\$2,484.28$ for the entire trading sample back to January 2008. We find that HFT firms increase their average daily profits by 46% on U.S. Treasury POMO days. The increase in profit of $595.02 per stock is marginally significant at 10% level.

To approximate returns from HFT activity, we also estimate in Table 4.7 the profits per share $\pi_{i,ps}^{HFT}$ from their aggressive and passive trades,

$$\pi_{i,ps}^{HFT} = \frac{\pi_i^{HFT}}{\sum_{t=1}^{T} q_{i,t}/2}. \tag{4.21}$$

[INSERT Table 4.7 HERE]

Given an average share price of around $30, the returns are quite modest. The trades, however, are very short term and rarely lose money. On U.S. Treasury POMO days, profits per share are positive 96.49% of the time. On the 2 days where the HFT firms lose money, they lose only 2/10 and 3.6/10 of one cent per share respectively, compared with the largest gain of nearly $0.04 per share on April 30, 2009.

In HN trades, HFT firms also rarely lose money on Treasury POMO days. They make profits 87.72% of the time. Crossing the spread on non-POMO days is much more risky. Profits are positive on only 60.63% of non-POMO trading days. The average profit per share when crossing the spread is typically small, only $0.0032 per share, but this rises by a statistically significant 300% during POMO auctions.

In NH trades where HFT firms are passive liquidity providers, profits per share are always positive on U.S. Treasury POMO days. Compared to their performance on non-POMO days, HFT firms increase the average profits per share by 38% on U.S. Treasury POMO days, and the effect is statistically significant.

On POMO days, HFT firms became more aggressive. While this should raise the profits on their passive activity, it should actually reduce their profits on aggressive trades unless their positioning profits are higher. This is evidence that HFT firms receive valuable private information during the POMO auctions because their profits per share rise despite the wider spreads.

Extrapolating the daily profit estimates to the broader market requires an estimate of the percentage of high frequency trading in the market captured by our sample. We present an estimate here in based on the 12.3% of total market capitalization represented by the firms in our sample. We sum daily profits across the 120 stocks, the 57 U.S. Treasury POMO days, and we assume similar activity in the sample we do not observe. We estimate profits of over $105 million during U.S. Treasury POMO auctions.

4.6 Robustness Checks

4.6.1 Time Window

We analyze the HFT firm behavior in the half-hour after U.S. Treasury POMO as a robustness check. In terms of inside quote frequency by the HFT firms, we use a similar model described in Section 5.1 but replace the variable U.S. Treasury Purchase with a dummy variable indicating the half-hour after purchases. The results for the HFT inside bid and offer participations are reported in Table 4.8. The effect is not statistically significant for either bids or offers during the half-hour after U.S. Treasury POMO auctions.

[INSERT Table 4.8 HERE]

The market impact of HFT trades in the half-hour after U.S. Treasury POMO is not significantly different from the same period of non-POMO days either. The result is shown in the second set of columns in Table 4.5.

4.6.2 FOMC Days

We also contrast our results with the behavior of HFT firms on the eight Federal Open Market Committee (FOMC) dates in our sample listed in the third set of columns in Table 4.5.

We compare the market impact of HFT trades during the period from 13:45-14:15, the half-hour before the Fed announces its policy intentions. We felt this period was analogous to our half-hour before the release of POMO Treasury purchases. We used trades from the HFT database and quotes from the NASDAQ ITCH feed. This limits our analysis to the 60 NASDAQ stocks in the sample. For the 60 stocks, the market impact on FOMC and non-FOMC days is little changed. From this, we conclude that the POMO auctions were more important events for the market.

4.7 Conclusion

Federal Reserve operations provide an unusual lens on the primary dealers, many of whom are also HFT firms. We document that the POMO auctions provide valuable private information to the HFT participants. We find evidence of larger permanent price impact and higher profits and profitability, particularly for aggressive trades. HFT firms perform a dual role as market makers, which the academic literature has extensively studied. During the POMO auctions though, our model predicts that they may shift their focus from being liquidity providers to trading aggressively. We find that HFT firms reduce their presence at the inside quote and less frequently provide liquidity to non-HFT firms. Studying HFT activity in event windows like POMO may give us a better indication of how HFT firms will perform in stressful market conditions.

Figure 4.1: Trade and Inside Quote Activity of HFT Firms
This figure presents the share of HFT firm participation in trades and inside quotes respectively from January 2008 to December 2009 as well as one week in February 2010.

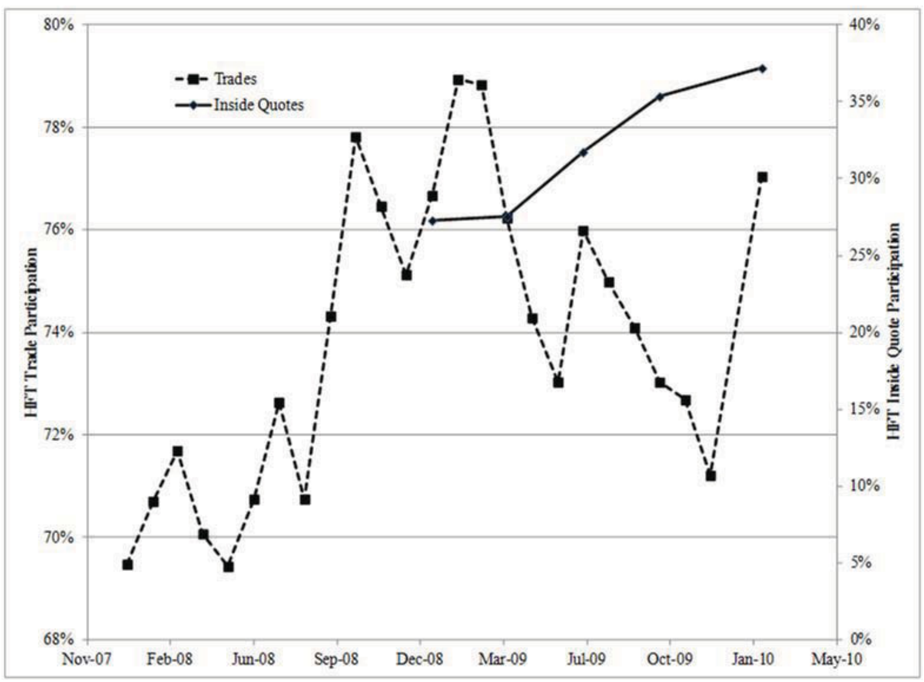

Table 4.1: U.S. Treasury Purchase Detail for March 25, 2009
This is the first of 214 Treasury purchases between March 2009 and June 2011. Details can be found on the New York Federal Reserve web site, http://www.newyorkfed.org/markets /pomo/display/index.cfm

Release Time:		10:30
Close Time:		11:00
Settlement Date:		March 26, 2009
Maturity/Call Date Range:		02/29/2016 - 02/15/2019
Total Par Amt $	Submitted	21,937,000,000
	Accepted	7,500,000,000

CUSIP	**Coupon**	**Maturity**	**Par Amt Accepted ($)**
912828KS8	2.6250	2/29/2016	2,836,000,000
912810DW5	7.2500	5/15/2016	115,000,000
912828FQ8	4.8750	8/15/2016	1,031,000,000
912828FY1	4.6250	11/15/2016	739,000,000
912810DX3	7.5000	11/15/2016	147,000,000
912828GH7	4.6250	2/15/2017	35,000,000
912828GS3	4.5000	5/15/2017	950,000,000
912810DY1	8.7500	5/15/2017	238,000,000
912828HA1	4.7500	8/15/2017	702,000,000
912810DZ8	8.8750	8/15/2017	159,000,000
912828HH6	4.2500	11/15/2017	0
912828HR4	3.5000	2/15/2018	0
912828HZ6	3.8750	5/15/2018	0
912810EA2	9.1250	5/15/2018	23,000,000
912828JH4	4.0000	8/15/2018	0
912828JR2	3.7500	11/15/2018	0
912810EB0	9.0000	11/15/2018	193,000,000
912828KD1	2.7500	2/15/2019	0
912810EC8	8.8750	2/15/2019	332,000,000
Exclusions:			
912828FF2	5.125	6/15/2016	0

Table 4.2: HFT Inside Quote Frequency During U.S. Treasury POMO

The table reports estimates of a model for the inside quote participation of the 26 HFT trading firms. We control for the growth in HFT activity using a linear and quadratic trend. We also include standard regressors for market making aggressiveness, past returns, volatility, volume, and order imbalance. We also include a time dummy for the quiet period from 10:30-11:00. Finally, we measure the effect of U.S. Treasury (UST) POMO auctions using two dummy variables, one for positive news and the other for negative. The classification of positive and negative news events is based on the 1-hour equity return after the start of the auction.

Variable	**Bid**	**Offer**
Trend	0.0022	0.0017
	(62.38)	(50.34)
$\text{Trend}^2/1000$	-0.0012	-0.0008
	(-16.00)	(-10.93)
Returns_{t-1}	0.0206	-0.0368
	(1.92)	(-3.55)
Realized Vol_{t-1}	-0.0146	-0.0121
	(87.94)	(-77.25)
$\text{Volume}_{t-1}/1000$	0.0745	0.0401
	(221.70)	(158.12)
Order $\text{Imbalance}_{t-1}/1000$	1.6756	-1.1129
	(13.36)	(-10.29)
S10:30-11:00	-0.0266	-0.0335
	(-4.62)	(-5.92)
Positive UST News	-0.0763	-0.0450
	(-4.00)	(-2.43)
Negative UST News	-0.0772	-0.0458
	(-6.87)	(-4.16)
Constant	-0.4383	-0.3251
	(-107.41)	(-82.64)
$\overline{R}^2$	0.1450	0.1155

t-statistics in parentheses.

Table 4.3: Unconditional HFT Net Buy Counts

This table presents the average HFT net buy counts per minute in aggressive HN and passive NH trades respectively during U.S. Treasury POMO auction periods and non-POMO times. The HFT net buys are computed as the difference between the number of HFT buyer and seller initiated trades. U.S. Treasury purchases are classified into positive and negative news events based on the 1-hour equity return after the start of the auction.

	Positive UST News		**Negative UST News**		**non-POMO**	
	HN	**NH**	**HN**	**NH**	**HN**	**NH**
Avg.	16.08	-9.11	-11.51	7.63	2.34	-1.71
SD	46.46	36.70	36.76	33.66	35.57	33.91
H_0: $\bar{c}_{POMO} = \bar{c}_{Non}$						
t-stat	1.94	-1.48	-2.10	1.68		

Table 4.4: HFT Net Buy Counts During U.S. Treasury POMO

The table reports estimates of models for HFT net trade counts in aggressive HN and passive NH trades respectively. We include standard regressors for trading aggressiveness, past returns, volatility, volume, and order imbalance. We also include a time dummy for the quiet period from 10:30-11:00. Finally, we measure the effects of positive and negative U.S. Treasury auctions using two dummy variables respectively.

Variable	**HN**	**NH**
Returns_{t-1}	-80.3836	-54.0185
	(-30.40)	(-18.46)
Realized Vol_{t-1}	-0.0332	-0.0526
	(-2.90)	(-3.04)
$\text{Volume}_{t-1}/1000$	-0.0020	0.0021
	(-2.21)	(1.86)
Order $\text{Imbalance}_{t-1}/1000$	0.0553	-0.0674
	(33.03)	(-32.03)
S10:30-11:00	0.0471	-0.0317
	(3.61)	(-2.04)
Positive UST News	0.1667	-0.1328
	(5.40)	(-3.65)
Negative UST News	-0.2476	0.1360
	(-7.59)	(3.50)
Constant	-0.0235	0.0075
	(-4.64)	(1.31)
$\overline{R}^2$	0.0024	0.0031

t-statistics in parentheses.

Table 4.5: HFT Market Impact

This table reports the average market impact calculated by 4.19. We use the eight FOMC announcements in 2009: January 29, March 18, April 29, June 24, August 12, September 23, November 4, and December 16. For the FOMC results we use the 60 NASDAQ stocks in the sample with quotes from the ITCH feed.

	10:30-11:00		11:00-11:30		13:45-14:15	
	POMO	**non-POMO**	**POMO**	**non-POMO**	**FOMC**	**non-FOMC**
Avg. (10^{-5})	3.4138	3.1761	3.0856	3.0129	2.8937	2.8927
SD (10^{-5})	0.3007	0.3037	0.4289	0.2251	0.2317	0.3630
H_0: $\overline{\Lambda}_{POMO} = \overline{\Lambda}_{Non}$						
t-stat	2.54		1.10		0.71	

Table 4.6: HFT Daily Profits Per Stock

This table presents estimated average profits per stock by 4.20 during U.S. Treasury POMO and non-POMO days.

	UST POMO	**non-POMO**
Mean	\$1,895.37	\$1,300.35
SD	2,736.70	3,642.63
H_0: $\overline{\pi}^{HFT}_{POMO} = \overline{\pi}^{HFT}_{Non}$		
t-stat	1.65	

Table 4.7: HFT Daily Profits Per Share

This table presents estimated average profits per share by 4.21 during U.S. Treasury POMO and non-POMO days.

	Total		**HN**		**NH**	
	UST	**non-POMO**	**UST**	**non-POMO**	**UST**	**non-POMO**
Mean	\$0.0178	\$0.0129	\$0.0099	\$0.0032	\$0.0341	\$0.0298
SD	0.0085	0.0105	0.0132	0.0158	0.0149	0.0177
Min	-\$0.0036	-\$0.0254	-\$0.0235	-\$0.0594	\$0.0064	-\$0.0084
Max	\$0.0385	\$0.0422	\$0.0550	\$0.0484	\$0.0897	\$0.1002
% Days>0	96.49%	88.13%	87.72%	60.63%	100.00%	96.88%
H_0: $\overline{\pi}^{HFT}_{ps,POMO} = \overline{\pi}^{HFT}_{ps,Non}$						
t-stat	3.86		3.42		2.11	

Table 4.8: HFT Inside Quote Frequency After U.S. Treasury POMO

The table reports estimates of a model for the inside quote participation of the 26 HFT firms in the half-hour after U.S. Treasury POMO auction periods.

Variable	**Bid**	**Offer**
Trend	0.0013	0.0011
	(32.88)	(28.59)
$\text{Trend}^2/1000$	-0.0007	-0.0006
	(-8.26)	(-7.06)
Returns_{t-1}	0.0348	-0.0561
	(2.95)	(-4.83)
Realized Vol_{t-1}	-0.0088	-0.0078
	(48.18)	(-44.28)
$\text{Volume}_{t-1}/1000$	0.0394	0.0323
	(116.47)	(109.12)
Order $\text{Imbalance}_{t-1}/1000$	1.3817	-1.1281
	(10.45)	(-8.93)
Inside Quote_{t-1}	1.6012	1.6150
	(465.57)	(473.96)
S11:00-11:30	0.0006	-0.0035
	(0.08)	(-0.52)
After UST Purchases	-0.0235	0.0076
	(-1.67)	(0.54)
Constant	-1.0536	-1.0210
	(-220.77)	(-217.52)
$\overline{R}^2$	0.4211	0.4118

t-statistics in parentheses.

Bibliography

ACHARYA, V., AND S. SCHAEFER (2006): "Liquidity Risk and Correlation Risk: Implications for Risk Management," Working Paper, New York University.

ADMATI, A., AND P. PFLEIDERER (1988): "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies*, 1, 3–40.

AITKEN, M. J., H. BERKMAN, AND D. MAK (2001): "The Use of Undisclosed Limit Orders on the Australian Stock Exchange," *Journal of Banking and Finance*, 25, 1589–1603.

ANAND, A., AND D. G. WEAVER (2004): "Can Order Exposure be Mandated?," *Journal of Financial Markets*, 7, 405–426.

BARON, M., J. BROGAARD, AND A. KIRILENKO (2014): "Risk and Return in High Frequency Trading," Working Paper, Princetion University.

BEN-DAVID, I., F. FRANZONI, AND R. MOUSSAWI (2014): "Do ETFs Increase Volatility?," Working Paper, NBER, http://nber.org/papers/w20071.

BENNETT, P., AND L. WEI (2006): "Market Structure, Fragmentation, and Market Quality," *Journal of Financial Markets*, 9, 49–78.

BESSEMBINDER, H., M. PANAYIDES, AND K. VENKATARAMAN (2009): "Hidden Liquidity: An Analysis of Order Exposure Strategies in Electronic Stock Markets," *Journal of Financial Economics*, 94, 361–383.

BIAIS, B., AND T. FOUCAULT (2014): "HFT and Market Quality," *Bankers, Markets & Investors*, 128, 5–19.

BIAIS, B., T. FOUCAULT, AND S. MOINAS (2013): "Equilibrium High-Frequency Trading," Working Paper, HEC Paris, http://ssrn.com/abstract=1834344.

BRECKENFELDER, J. (2013): "Competition between High-Frequency Traders, and Market Quality," Working Paper, Stockholm School of Economics, http://ssrn.com/abstract=2264858.

BROGAARD, J. (2011): "The Activity of High Frequency Traders," Working Paper, University of Washington.

BROGAARD, J., B. HAGSTRÖMER, L. L. NORDEN, AND R. RIORDAN (2014): "Trading Fast and Slow: Colocation and Market Quality," Working Paper, University of Washington, http://ssrn.com/abstract=2317797.

BROGAARD, J., T. HENDERSHOTT, AND R. RIORDAN (2013): "High Frequency Trading and the 2008 Short-Sale Ban," Working Paper, University of Washington.

——— (2014): "High Frequency Trading and Price Discovery," *Review of Financial Studies*, 27, 2267–2306.

Buti, S., and B. Rindi (2013): "Undisclosed Orders and Optimal Submission Strategies in a Limit Order Market," *Journal of Financial Economics*, 109, 797–812.

Carrion, A. (2013): "Very Fast Money: High-Frequency Trading on the NASDAQ," *Journal of Financial Markets*, 16, 680–711.

Cebiroğlu, G., N. Hautsch, and U. Horst (2014): "Does Hidden Liquidity Harm Price Efficiency? Equilibrium Exposure under Latent Demand," Working Paper, University of Vienna, http://ssrn.com/abstract=2305171.

Christie, W. G., and P. H. Schultz (1994): "Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes?," *Journal of Finance*, 29, 1813–1840.

Commodity Futures Trading Commission and Securities and Exchange Commission (2010): "Preliminary Findings Regarding the Market Events of May 6, 2010," Washington, D.C.

Copeland, T. E., and D. Galai (1983): "Information Effects on the Bid-Ask Spreads," *Journal of Finance*, 38, 1457–1469.

Cvitanic, J., and A. Kirilenko (2010): "High Frequency Traders and Asset Prices," Working Paper, California Institute of Technology, http://ssrn.com/abstract=1569067.

D'Amico, S., and T. B. King (2010): "Flow and Stock Effects of Large-Scale Treasury Purchases," Finance and Economics Discussion Series Paper No. 2010-52, Board of Governors of the Federal Reserve System.

De Winne, R., and C. D'Hondt (2007): "Hide-and-Seek in the Market: Placing and Detecting Hidden Orders," *Review of Finance*, 11, 663–692.

Eggertsson, G., and M. Woodford (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 1, 139–211.

Ellis, N. S. (2010): "The NYSE Response to Specialist Misconduct: An Example of the Failure of Self-Regulation," *Berkeley Business Law Journal*, 7, 102–149.

Foley, S., K. Malinova, and A. Park (2013): "Dark Trading on Public Exchanges," Working Paper, University of Toronto, http://ssrn.com/abstract=2182839.

Foucault, T., J. Hombert, and I. Roşu (2013): "News Trading and Speed," Working Paper, HEC Paris, http://ssrn.com/abstract=2188822.

Frey, S., and P. Sandås (2012): "Effective Search for Liquidity Using Iceberg Orders," Working Paper, University of Virginia.

Gagnon, J., M. Raskin, J. Remache, and B. Sack (2011): "The Financial Market

Effects of the Federal Reserve's Large-Scale Asset Purchases," *International Journal of Central Banking*, 7, 3–44.

Gai, J., C. Yao, and M. Ye (2014): "The Externalities of High Frequency Trading," Working Paper, University of Illinois, http://ssrn.com/abstract=2066839.

Gao, C., and B. Mizrach (2013): "High Frequency Trading in the Equity Markets During U.S. Treasury POMO," Working Paper, Rutgers University, http://ssrn.com/abstract=2294038.

Glosten, L. R., and P. R. Milgrom (1985): "Bid, Ask, and Transaction Prices in A Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71–100.

Golub, A., J. Keane, and S.-H. Poon (2012): "High Frequency Trading and Mini Flash Crashes," Working Paper, University of Manchester, http://ssrn.com/abstract=2182097.

Gumbel, E. J. (1960): "Bivariate Exponential Distributions," *Journal of the American Statistical Association*, 55, 698–707.

Hagströmer, B., and L. Norden (2013): "The Diversity of High Frequency Traders," *Journal of Financial Markets*, 16, 741–770.

Harris, L. E. (1997): "Order Exposure and Parasitic Traders," Working Paper, University of Southern California.

Hasbrouck, J. (1991): "Measuring the Information Content of Stock Trades," *Journal of Finance*, 46, 179–207.

——— (2013): "High Frequency Quoting: Short Term Volatility in Bids and Offers," Working Paper, NYU Stern, http://ssrn.com/abstract=2237499.

Hasbrouck, J., and G. Saar (2013): "Low-Latency Trading," *Journal of Financial Markets*, 16, 646–679.

Hendershott, T., C. M. Jones, and A. J. Menkveld (2011): "Does Algorithmic Trading Improve Liquidity," *Journal of Finance*, 66, 1–33.

Hirschey, N. H. (2013): "Do High-Frequency Traders Anticipate Buying and Selling Pressure?," Working Paper, London Business School, http://ssrn.com/abstract=2238516.

Jiang, C. X., T. H. McInish, and J. Upson (2014): "Market Fragmentation and Information Quality: The Role of TRF Trades," Working Paper, University of Memphis, http://ssrn.com/abstract=1960115.

Jovanovic, B., and A. J. Menkveld (2012): "Middlemen in Limit Order Markets," Working Paper, NYU and VU University Amsterdam, http://ssrn.com/abstract=1624329.

Joyce, M., A. Lasaosa, I. Stevens, and M. Tong (2011): "The Financial Market

Impact of Quantitative Easing in the United Kingdom," *International Journal of Central Banking*, 7, 113–161.

Krishnamurthy, A., and A. Vissing-Jorgensen (2010): "The Aggregate Demand for Treasury Debt," Working Paper, Northwestern University.

——— (2011): "The Effects of Quantitative Easing on Interest Rates," Working Paper, Northwestern University.

Krugman, P. R. (1998): "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 2, 137–205.

Lou, D., H. Yan, and J. Zhang (2013): "Anticipated and Repeated Shocks in Liquid Markets," *Review of Financial Studies*, 26, 1891–1912.

Madhavan, A. (2012): "Exchange-Traded Funds, Market Structure and the Flash Crash," Working Paper, http://ssrn.com/abstract=1932925.

Malinova, K., A. Park, and R. Riordan (2013): "Do Retail Traders Suffer from High Frequency Traders?," Working Paper, University of Toronto, http://ssrn.com/abstract=2183806.

Martinez, V. H., and I. Roşu (2013): "High Frequency Traders, News and Volatility," Working Paper, HEC Paris.

Menkveld, A. J. (2013): "High Frequency Trading and the New-Market Makers," *Journal of Financial Markets*, 16, 712–740.

Menkveld, A. J., and T. Wang (2011): "How Do Designated Market Makers Create Value for Small-Caps?," Working Paper, VU University Amsterdam, http://ssrn.com/abstract=890526.

Moinas, S. (2010): "Hidden Limit Orders and Liquidity in Order Driven Markets," Working Paper, Toulouse School of Economics.

Nanex (2011): "Flash Equity Failures in 2006, 2007, 2008, 2009, 2010, and 2011," http://www.nanex.net/FlashCrashEquities/FlashCrashAnalysis_Equities.html.

Neely, C. (2010): "The Large-Scale Asset Purchases Had Large International Effects," Working Paper, Federal Reserve Bank of St. Louis.

O'Hara, M., and M. Ye (2011): "Is Market Fragmentation Harming Market Quality?," *Journal of Financial Economics*, 100, 459–474.

Pardo, A., and R. Pascual (2006): "On the Hidden Side of Liquidity," Working Paper, Universidad de las Islas Baleares.

Roley, V. V. (1982): "The Effect of Federal Debt-Management Policy on Corporate Bond and Equity Yields," *Quarterly Journal of Economics*, 97, 645–668.

Sandås, P. (2001): "Adverse Selection and Competitive Market Making: Empirical Evidence from a Limit Order Market," *Review of Financial Studies*, 14, 705–734.

SEC (2010): "Concept Release on Equity Market Structure," Washington, D.C., http://www.sec.gov/rules /concept/2010/34-61358.pdf.

Shiratsuka, S. (2010): "Size and Composition of the Central Bank Balance Sheet: Revisiting Japan's Experience of the Quantitative Easing Policy," Working Paper, Federal Reserve Bank of Dallas Globalization and Monetary Policy Institute.

Sornette, D., and S. Von der Becke (2011): "Crashes and High Frequency Trading: An Evaluation of Risks Posed by High-speed Algorithmic Trading," Working Paper, Swiss Finance Institute, http://ssrn.com/abstract=1976249.

Sussman, A., L. Tabb, and R. Iati (2009): "US Equity High Frequency Trading: Strategies, Sizing and Market Structure," Tabb Group Report.

Tuttle, L. (2006): "Hidden Orders, Trading Costs and Information," Working Paper, American University of Sharjah.

Yao, C. (2012): "Hidden Agendas: A Study of the Informativeness of Concealed Orders," Working Paper, University of Warwick.

Zhang, F. (2010): "The Effect of High Frequency Trading on Stock Volatility and Price Discovery," Working Paper, Yale University, http://ssrn.com/abstract=1691679.