

Trading Fast and Slow: Colocation and Market Quality*

Jonathan Brogaard

Björn Hagströmer

Lars Nordén

Ryan Riordan

First Draft: August 2013

Current Draft: November 2013

Abstract: Using user-level data from NASDAQ OMX Stockholm, we investigate how different network connectivity speeds influence market participant dynamics. We find that collocated traders have an informational advantage relative to non-collocated participants. We use an exchange system upgrade that allows collocated traders to upgrade to an even faster connection to identify a shock to the speed hierarchy. Participants that upgrade reduce their adverse selection costs and improve their inventory management ability, allowing them to increase their market share in liquidity provision. Non-collocated traders incur higher adverse selection costs after the event. Overall, however, the introduction of speed differentiation improves both bid-ask spreads and market depth. Our results suggest that the liquidity improvements are related to the fastest traders' increased market share and their enhanced inventory management abilities.

*We thank NASDAQ OMX for providing the data, and Petter Dahlström, Richard Gaudy, Frank Hatheway, and Alessandro Romani for numerous discussions about the data and the market structure. Björn Hagströmer and Lars Nordén are grateful to the Jan Wallander and Tom Hedelius foundation and the Tore Browaldh foundation for research support.

Contact: Jonathan Brogaard, Foster School of Business, University of Washington, (Email) brogaard@uw.edu, (Tel) 206-685-7822; Björn Hagströmer, Stockholm University School of Business, (Email) bjh@fek.su.se, (Tel) 0046 8 163030; Lars Nordén, Stockholm University School of Business, (Email) ln@fek.su.se, (Tel) 0046 8 6747139; and Ryan Riordan, University of Ontario Institute of Technology, (Email) ryan.riordan@uoit.ca (Tel) 905-721-8668, ext. 5397.

I. Introduction

Proximity to markets matters. Those closest to the market are the first to observe market movements and to receive news. In the past such proximity was purchased through acquiring a seat on the exchange. Nowadays, proximity involves colocating a server at an exchange and subscribing to direct data connections. Many exchanges provide degrees of proximity in both physical location as well as speed access. This paper asks the questions of how latency heterogeneity affects trading dynamics among traders and how it influences market quality.

In modern equity markets, exchanges offer proximity services at multiple speeds, allowing them to price discriminate between customers. This is done in the competition between exchanges (as discussed by Pagnotta and Philippon, 2012), but also at individual trading venues. At NASDAQ OMX Stockholm (NOMX-St) the first colocation service was introduced in 2010, and in March 2011, a faster connection was offered as an add-on to the existing service. In September 2012, an even faster colocation service was introduced, enabling trading firms to choose between four (three colocated and one non-colocated) speeds. We use a proprietary data set allowing us to disentangle the trading activities of traders who subscribe to different colocation services. This unique data set enables us to study the industrial organization of investors, and how heterogeneous network connectivity influences market participants and ultimately market quality.

Several theoretical models of trader speed have been proposed. A common theme in such models is that there are two speeds, fast and slow, and that the speed advantage of fast traders amounts to a short-term informational advantage. In Foucault, Hombert and Rosu (2013), fast traders consist solely of liquidity taking news traders, while market makers are all slow. Similarly, in Biais, Foucault, and Moinas (2013) fast traders use market orders to trade on information that slow traders have yet to process. Martinez and Rosu's (2013) model fast traders

(referred to as high-frequency traders, HFTs) that observe information and trade actively on it, making markets more efficient, but generating large trading volumes and contributing to volatility. In Cartea and Penalva (2012), the fast traders use their speed advantage to profit from trading ahead of other traders, engaging in short-term intermediation.

Thus, these four models share the assumption that fast traders are primarily active traders, who use market orders to pick off the limit orders of slower traders. This is in contrast to a competing set of models, where fast traders are primarily seen as liquidity providers. Jovanovic and Menkveld (2012) consider a market where market makers are fast in the sense that they are never picked off based on public news as they are always the first to update their quotes based on hard (machine interpretable) news. Along the same lines, the fast traders modeled by Hoffman (2013) can never be picked off by slow traders. In the model by Aït-Sahalia and Saglam (2013), the liquidity-providing HFT has the ability to (imperfectly) predict incoming market orders by slower traders, and revising outstanding quotes in response to the prediction. Finally, Menkveld and Zoican (2013) consider both types of fast traders, referring to active fast traders as bandit HFTs, and passive fast traders as market-making HFTs.

Our first contribution to the literature is that we provide an empirical account of what fast and slow traders actually do. We classify colocated traders as fast and non-colocated traders as slow. This classification objectively captures fast traders regardless of how they use their speed. Typically fast traders are assumed to be HFTs and are either identified on the discretion of the exchange (such as the NASDAQ high-frequency trading [HFT] flag used in several papers, e.g., Brogaard, Hendershott, and Riordan, 2013) or based on trading behavior as in for example Kirilenko et al. (2011). The colocation indicator used in this paper, we argue, yields a trader speed classification that is free from any subjective view of the strategy used by fast traders.

Our results lend support to the notion that the speed advantage amounts to an informational advantage. We find strong evidence that fast traders trade actively ahead of news (imposing adverse selection costs on their counterparties), and that fast traders are good at avoiding adverse selection when supplying liquidity. Consistent with previous literature, we also find that fast traders have high order-to-trade ratios and make up substantial trading volume, in our case 43% of the total trading volume. In contrast with most theoretical models' assumptions, we find that fast traders use a mixture of market and limit orders, acting as liquidity suppliers in 47% of their trades. Thus, any assumption classifying fast traders as either active or passive is likely to be misleading. We find this empirical result unsurprising, as both active and passive trading strategies are likely to benefit from speed, and economies of scope should lead fast traders to do both.

Next, as a second contribution, we provide an analysis of how multiple segments of fast traders (subscribing to different colocation services) differ and how they interact with each other. This analysis takes advantage of the colocation upgrade in September 2012, which allows traders to reduce the roundtrip latency in order entry by more than 20% compared to the next best level of service. We show that when colocated traders are able to choose whether to upgrade or not; those that upgrade significantly increase their market share in liquidity provision. We also find that the upgraded traders reduce their adverse selection costs and, perhaps as a result, that they are able to relax their inventory control. An increase in the inventory management capabilities of certain traders should lead the fastest to supply more liquidity at given levels of inventory. The increase in risk-bearing capacity will lead fast market makers to submit tighter bid-ask spreads for longer periods of time before inventory constraints cause them to widen their spreads or to use spread-widening market orders to manage their inventory risk. We find evidence that the upgrade leads fast market makers to take on more inventory, and to mean revert their inventories less often.

For active trading, our results are less clear. We find that the fast traders who do not upgrade increase the adverse selection imposed on other traders, and more so than the upgraded traders do. The non-upgraded fast traders also increase their market share in active trading more than the upgraded traders do. On the other hand, in a multinomial logit regression framework, we find that the upgrade allows fast traders to improve their ability to trade actively on hard information, proxied by index futures returns news (following Hendershott and Riordan, 2012). We think that these conflicting results reflect segmentation in news trading between the upgraded and the non-upgraded colocated traders. The colocation upgrade may be useful for trading on order book news that travels fast within the exchange. News on the macroeconomy or firm fundamentals, which in general are larger information shocks than order book news, travels slower and requires other algorithms. The fastest colocation service may yield little edge in trading on such news, allowing the non-upgraded fast traders to dominate that business segment. This reasoning is analogous to Martinez and Rosu (2013), where HFTs that fail to compete on speed change their business model to less latency-sensitive trading strategies.

Finally, as a third contribution to the literature on trading speed, we investigate how the colocation upgrade affects market quality. Overall, we find that market liquidity improves in both the tightness and depth dimensions, and that short-term volatility is unaffected by the upgrade. Our result on liquidity is consistent with empirical results from several previous studies on the introduction of colocation services. Boehmer, Fong, and Wu (2012) study the introduction of colocation services at 39 exchanges, and find that such services leads to improved liquidity but also increasing volatility. Frino, Mollica, and Webb (2013) find that the introduction of colocation services at the Australian Securities Exchange (ASX) leads to decreasing bid-ask spreads and improved market depth. Riordan and Storckenmaier (2012) do not study colocation, but find that a trading system upgrade at Deutsche Börse in 2007 leads to

improved liquidity. Hendershott, Jones, and Menkveld (2011) find that the automatization of quotes dissemination at the NYSE, which they use as an instrument for algorithmic trading, is associated with liquidity improvements. Our results are in contrast to Menkveld and Zoican (2013), who analyze a trading system upgrade at NASDAQ OMX Nordic exchanges in 2010 (INET, including the first introduction of colocation services), and Gai, Yao, and Ye (2013), who study two latency-improving events at the NASDAQ in 2010. Menkveld and Zoican (2013) find that adverse selection costs increase, whereas Gai, Yao and Ye (2013) conclude that bid-ask spreads are unaffected and that market depth decreases.

We find that the improvement in market quality may be due to improved inventory capabilities among the fastest traders. Our evidence shows that the traders who upgrade to the fastest colocation type are able to maintain their liquidity supply even when they approach their inventory constraints, leading to improvements in overall market quality. The relation between trading speed, inventory management and market quality is not considered in any of the current theoretical models, and we see this as an interesting topic for future research.

Our results are limited in two ways. NOMX-St trades roughly 60% of volume in our sample stocks, meaning that we only observe 60% of trading. Whether or not these results are good characterizations of all trading is unknown. Also, the stocks we study are the largest and most liquid in Sweden yet still relatively small in terms of market capitalization.

The paper is organized as follows. In Section II, we describe our data. Section III presents the static results of market activity by type of participant. Section IV exploits the colocation upgrade to study how speed differentiation affects market dynamics and latency competition. Section V explores why speed enhances market quality, and Section VI concludes.

II. Institutional detail, data and descriptive statistics

Exchanges have several incentives to cater to the fastest traders. Fast traders' ability to trade on short-lived informational advantages generates substantial order flows that are profitable for the exchange. Fast traders generally lead price discovery, making the exchange where they trade attractive to liquidity traders. Finally, fast traders are keen to maintain their speed advantage, allowing the exchanges to differentiate their pricing between fast and slow traders. Charging higher prices for higher speeds also allows exchanges to extract some of the rents that accrue to informed traders. For these reasons, competitive exchanges offer proximity services, allowing trading firms to place their trading servers as close to the matching engine as possible.

NOMX-St runs a fully electronic limit order book market for equities and derivatives. NOMX-St adopted its current trading system INET in February 2010, and at the same time introduced its first type of colocation service. Exchange members who pay for the colocation service get the infrastructure for running a trading server from within the exchange. The package includes everything from the actual connection to the information flows and matching engine, to server cages, electricity, maintenance, and safety installations. Effectively, colocated firms are able to cut their latency in the access to news about order flow and in their order submission to the matching engine.

In March 2011, members who subscribed to the original colocation service were offered an even faster connection, called *Premium Colocation*. The Premium Colocation is an add-on to the original colocation that according to the exchange cuts the latency of both order entry and order book information retrieval. In September 2012, another optional upgrade was provided by NOMX-St, referred to as the *10G Premium Colocation*. According to the exchange's marketing materials, the 10G Premium Colocation service can cut the time it takes from order submission

to order confirmation by more than 20%, relative to the Premium Colocation service. Whereas the Premium Colocation can be used for both the cash and the derivatives markets, the 10G is for equity trading only. Thus, for the equity markets at NOMX-St, there are three different segments of proximity services. In this paper, we utilize the different colocation segments to proxy for trader speed.

NOMX-St operates continuous trading from 9 am to 5.30 pm every weekday, except on Swedish bank holidays. If the day before a Swedish bank holiday is a weekday, trading on that day closes at 1 pm. Opening and closing prices are determined in call auctions. Limit orders entered into INET specify a trading interest in terms of a quantity and a limit price, and may be cancelled or modified at any given time. Limit orders may also be entered with various conditions; see Hagströmer and Nordén (2013) for further details. Execution during the continuous trading session takes place when two limit orders cross, following an order of priority by price, member, visibility, and time.¹

There are roughly one hundred member firms at NOMX-St. Exchange members may sell trading services to clients either as traditional brokers, through direct market access, or through sponsored access. Direct market access gives customers access to the market through the infrastructure of the member firm. In the case of sponsored access, the customer uses its own infrastructure but trades under the member identity of the sponsor. Subscriptions to proximity services may be acquired by either exchange members or sponsored access clients.

(a) Data

We access all messages that are entered into INET, including limit order submissions, cancellations, modifications, and executions. We refer to limit orders that are priced in a way

¹ The "member" priority rule implies that a member with a limit order posted at the best price has execution priority if a market order originating from the same member is posted on the other side of the book.

that they execute immediately as market orders; and to other orders as non-marketable limit orders. Note that both these order types have their origin in an order submission, for which we observe the time of entry, the quantity, the limit price, trader identity information, visibility conditions, and time in force. Through an order sequence number we are able to track an order and see when it is cancelled, modified, or executed. Each such event is associated with one or many messages in our data set, and may be for less than the full order volume. By aggregation of all limit orders that are active at a given point in time, we are able to reconstruct the state of the order book at any given instant of the trading day. For trades there is one execution message for each of the two limit orders that cross. The execution messages convey an execution matching identifier allowing us to see which messages belong together, for example to see the account numbers that trade with each other.

For each message, we access the broker ID, whether the order is proprietary or on behalf of a client, and through which computer port the order is entered. We also access a list of colocation accounts as of Dec. 3, 2012, showing which computer ports are associated with which colocation subscription. A key feature of our paper is that we can use that list to proxy for whether a colocated trading firm enters a specific message, and which type of colocation service that is used. Finally, for each message there is a field indicating automated order entry, as NOMX-St offers a fee discount for such orders. We use this field to identify algorithmic orders.

We use the colocation information to categorize traders into different speed segments. This classification is done at the trading account level rather than the member firm level, implying that sponsored access clients are considered to be separate trading entities comparable to member firms. Furthermore, in order to achieve clean definitions of colocation segments, trading accounts of different types belonging to the same member firm are treated as separate trading entities. A limitation of our trader classification is that the colocation status of trading firms can change over time, whereas we know the status at one date only. For example, we can

see who subscribes to the fastest colocation service on Dec. 3, 2012, but we do not know whether those firms signed up for the service at the date of introduction or at a later date.

As a complement to the full message history proprietary data set, we access public data on stock market capitalization from the NOMX-St website. Furthermore, we use data on prices of OMXS 30 index futures, traded at NASDAQ OMX Derivatives Exchange in Stockholm. The futures price data is retrieved from Thomson Reuters' Tick History database, maintained by the Securities Research Centre of Asia-Pacific (SIRCA).

(b) Sample selection

We limit our investigation to the thirty large-cap constituents of the leading Swedish equity index, OMXS 30.² The restriction to large-cap stocks is due to the fact that fast traders concentrate their activity in such stocks. A limitation of our data is that we access information concerning NOMX-St only, not the other trading venues where the same stocks are traded. At the time of our sample, NOMX-St hosts about 60% of the trading volume in the sample stocks. The main competitors are BATS Chi-X (around 30%), Turquoise (around 4%), and Burgundy (around 4%).³

We choose the time frame of our study to be able to study the characteristics of traders in different speed segments and to see the impact of the colocation upgrade event in September 2012. We consider two trading periods, one before and one after the colocation upgrade event. The *Before* period is Aug. 20 – Sep. 14, i.e., the four weeks just before the event. The *After* period runs from Oct. 1 – Oct. 26 and is chosen to see the effect of the colocation event. The reason that we skip two weeks directly following the event is that we have no information on how soon after the 10G Premium Colocation was made available trading firms subscribed and

² All thirty issues except one are Swedish stocks that have their primary listing at NOMX-St. The exception is NOKI SEK, which is a Swedish depository receipt issued by Nokia Oyj, a Finnish cell-phone manufacturer. Nokia Oyj has its primary listing at NASDAQ OMX Helsinki.

³ Market share data is taken from http://www.batstrading.co.uk/market_data/venue/index/OMXS/

started to use the new technology. Even if they subscribed to it immediately, there may be a testing period where they adapt their trading programs to the new infrastructure. Our event period is not entirely free from other changes in the market microstructure. On Sep. 27, 2012, NOMX-St introduced the possibility to modify limit orders prices. Before that possibility existed, traders would cancel their previous order and submit a new order at a new price level. This should be kept in mind when evaluating our event findings.

(c) Descriptive statistics

Table 1 shows summary statistics for the thirty issues at NOMX-St considered in our sample, based on data from the After period. Column 1 provides the ticker symbols of these 30 issues. Market capitalization is stated in millions of Swedish Krona (MSEK) and is calculated per the closing prices of Oct. 31, 2012. On that date, 1 MSEK was worth about 114,000 EUR or 151,000 USD. Note that market capitalization statistics concern the value of the issue in question, which is not always equal to the firm value. For example, *NOKI SEK* represents the depositary receipts of *Nokia Oyj* on NOMX-St, which is a much smaller issue than the Nokia stocks traded at NASDAQ OMX Helsinki. Note also that one firm may have several issues, as seen for *ATCO A* and *ATCO B*, both issued by *Atlas Copco*, a mining equipment manufacturer. As seen in Column 2 of Table 1, market capitalization varies a lot among the 30 issues with the highest being 327,921 MSEK (HM B) and the lowest being 2,056 MSEK (NOKI SEK).

INSERT TABLE 1 ABOUT HERE

Columns 3 and 4 report daily trading volume and daily turnover, respectively. Daily trading volumes and daily turnover statistics provided include the continuous trading and the opening and closing call auctions at NOMX-St, but exclude trading at other venues. Daily turnover is the fraction of market capitalization that is traded on average on a daily basis. The daily trading volume varies between 44 MSEK (SECU B) to 604 MSEK (VOLV B). The daily

turnover is typically around 0.25% of the market capitalization, however NOKI SEK is a large outlier having 6.63% of its shares traded daily. This is likely due to NOKI SEK being a small portion of the Nokia Oyj's total market capitalization.

We provide four measures of market quality, all measured on a second-by-second basis during continuous trading, excluding the first and last minute of each trading day. Each reported statistic corresponds to the average across all seconds in the studied period. Column 5 reports the *Bid-ask spread*, half the difference between the best buy and best sell prices available in the limit order book, divided by the spread midpoint, and expressed in basis points. There is sizeable variation among firms' bid-ask spreads, ranging from 2.3 basis points (TLSN) to 6.1 (NOKI SEK).

Column 6 reports the *Depth at BBO*, the volume available at the best bid-offer prices (those that constitute our bid-ask spread measure). Typically the Depth at BBO is around 1 MSEK. Column 7 shows the *Depth at 0.5%*, the volume required to move the price in either direction by at least 0.5%. The difference between Depth at BBO and Depth at 0.5% is large. For instance HM B has 1.16 MSEK Depth at BBO, but almost 16 MSEK Depth at 0.5%. This suggests that much of the liquidity rests away from the BBO. Both depth measures are reported as the average across the buy and the sell side. Finally, Column 8 reports *Volatility*, the average of the squared one-second basis point returns. Volatility is normally around 0.50 squared basis points. However, NOKI SEK is an outlier with Volatility of 3.45.

As can be expected from a set of large-cap stocks at a technologically advanced exchange, the measures presented in Table 1 indicate that the market is highly liquid both in terms of tightness (the price of crossing the spread) and depth (the price impact of large trading volumes). As expected, there is a tendency for larger stocks to be more heavily traded, more liquid and less volatile.

III. Cross-Sectional Results: Fast and Slow Traders

In this section we compare the trading activity and performance of fast and slow traders. As discussed above, a common modeling approach in the theoretical literature on HFT is to distinguish trader types by their relative ability to detect information and react to it faster than their competitors. We contribute to that literature by distinguishing fast and slow traders by their colocation status and reporting empirical measures of their activities.

Our categorization is distinct from empirical papers on HFT in that it does not impose any priors on fast traders. HFTs are generally thought of as the market participants closest to the trading mechanism with high costs of carrying inventories and sophisticated proprietary trading strategies. In previous empirical studies with data that allows for identification of HFTs, such traders are classified either on the discretion of the exchange (as in Brogaard, Hendershott, and Riordan, 2013; and Hagströmer and Nordén, 2013), or based on their trading behavior (as in Kirilenko et al., 2011; and Baron, Brogaard, and Kirilenko, 2013). Both classification methods have the drawback that they impose certain properties on HFTs.

An alternative approach to categorize traders by their trading speed is to measure latency from the empirical data. The drawback with empirical measures of trading speed, however, is that they are inherently strategy-specific. For example, the strategic runs measure by Hasbrouck and Saar (2013) captures the speed of quote revisions in response to order flow events. Another example is Scholtus and van Dijk (2012), who investigate the trading performance of technical analysis strategies in relation to trading speed. Scholtus, Frijns, and van Dijk (2012), in turn, focus on the importance of reaction times when trading on macroeconomic news events. To measure latency at the trader level without specifying a trading strategy of interest, off-exchange data is likely required. Based on this discussion, we regard colocation status as an appealing proxy of trader speed.

(a) *Trader characteristics*

Table 2 reports summary statistics regarding the different types of traders based on colocation status. All messages entered during the continuous trading of OMXS 30 index stocks during Oct. 1 – Oct. 26, 2012, are considered. Trader grouping with respect to colocation type is based on the status on Dec. 3, 2012.⁴ Column 1 shows the distribution of accounts across the colocation status groups.

INSERT TABLE 2 ABOUT HERE

We document a number of volume statistics, all reported as the trader group fraction of the total activity. Column 2 shows the *Trading volume*, the percent of SEK volume by group. Column 3 reports the fraction of executions by group (*Trades*). Column 4 is *Submissions*, the number of limit orders submitted to NOMX-St. Column 5 reports the fraction of limit orders cancelled by each group (*Cancellations*). Column 6 shows *Modifications*, the percent of limit order modifications. Columns 7 and 8 are *Active trades* and *Passive trades*. These capture the percent of executions where the trader participates through a market order or a non-marketable limit order, respectively.

In addition to volume measures we document three measures showing the distribution of trading and quoting activity within each group. First, Column 9 reports the order-to-trade ratio (q/t), which is the sum of all submissions, cancellations, and modifications, divided by the number of executions. Columns 10 and 11 show *Liq. supply*, which is the fraction of all Trading volume and Trades, respectively, within the trader group where the trader participates through a non-marketable limit order. All statistics are based on sums across stocks and trading days.

⁴ To check the validity of the colocation status list for the period studied (Oct. 1 – Oct. 26, 2012), we repeat the analysis for a date range surrounding the colocation status list (Nov. 15 – Dec. 14, 2012). We obtain qualitatively similar results, available from the authors on request.

The main breakdown of interest is between colocated trading accounts, which we refer to as *Colo* or fast traders; and non-colocated trading accounts, which we refer to as *NonColo* or slow traders. We report the aforementioned statistics for these two groups in Panel A of Table 2. There are 112 accounts that are NonColo, and 58 that are Colo. Even so, Colo traders make up 41.3% of trading volume. At the same time they make up the vast majority of submissions, cancellations, and modifications, over 80% of each. Colo traders initiate 45.7% of all trades but appear on the passive side of the trade in only 41.8% of all trades. Colo traders' aggregate q/t ratio is 40.8, while NonColo accounts' is only 6.0. To summarize, fast traders appear to be behaving in a distinct manner from the slow traders. This is expected as the fast traders have made a choice to emphasize speed in their trading process by investing in colocation.

Panel B focuses exclusively on the fast traders. It separates the results observed in Row 2 into the three subgroups based on the type of colocation service that they subscribe to. There are 17 accounts who subscribe to the original colocation service (referred to as *BasicColo*), 22 who have the Premium colocation (*PremiumColo*), and 19 who subscribe to the fastest connection available (*10GColo*). The fastest segment plays the leading role among the fast traders; traders in that segment represent 21.3% of the total trading volume, while the PremiumColo traders represent 16.5%, and BasicColo traders are behind 3.5% of the total trading volume (Column 1).

The Submissions, Cancellations, and Modifications are similarly distributed with the 10GColo accounts making up the majority of these actions, and BasicColo traders doing relatively little of it. Columns 6 and 7 provide insights into how speed is used. The Active trades are equally distributed between the PremiumColo and the 10GColo traders, while the Passive trades are mostly done by the 10GColo traders. The order-to-trade ratio appears to be increasing with speed. It almost triples going from BasicColo (9.3) to PremiumColo (27.7), and then doubles when going to 10GColo (55.5). Finally, it is notable from the liquidity supply ratios that traders of all colocation segments trade with a mix of market orders and non-marketable limit

orders. This is in contrast with theoretical models assuming that the speed advantage is used exclusively for either aggressive trading, or for passive trading.

Panel C disaggregates Colo and NonColo trading volume into client and proprietary order flows. Most of the Colo trading is done on the accounts' own behalf (40.6%) while only 0.7% is done for clients of colocated trading accounts. This shows that services that provide clients with “sophisticated” order entry systems are using technology that is slower than that of the fastest traders. Overall, most of the fast trading properties are driven by the proprietary trading activity. The activity of slow traders is a mixture of client and proprietary flows. The NonColo client trading activity displays liquidity supply ratios close to 50%, indicating that clients use mixtures of market orders and non-marketable limit orders. Proprietary trading is nowadays intimately associated with HFT. The fact that 24.6% of the trading volume in our data set comes from NonColo proprietary traders indicate that not all proprietary flows are necessarily from fast traders.

Panel D disaggregates the Colo and NonColo trading activities into *algorithmic* order flows and *other* flows. This distinction is related to a discount in trading fees for automated order flows. Perhaps the most surprising result is that not all fast trading is done using the algorithmic order type. We observe that 4.4% of the total trading goes through Colo accounts but does not use the algorithmic order entry that is eligible for a fee discount. We do not know whether these flows are entered manually or if there are other reasons for not using the discounted fees. Our results show that the Colo trading that is not using the discounted order entry has properties that are more similar to NonColo trading (either algorithmic or other) than to algorithmic Colo trading. For example, the algorithmic Colo trading has a much higher q/t ratio (44.6) than the other Colo trading activity (9.3).

Overall, the results in Table 2 show that there are large differences between fast and slow trading. In line with common conceptions about HFT, colocated flow is dominated by algorithmic activity for proprietary purposes. Furthermore, most of the colocation activity is done using the fastest technology, 10G Premium Colocation. The results may either indicate that there is a strong appetite for trading speed, or that fast traders are able to capture a large market share. Our next point of interest is what incentives are behind the investments in speed.

While colocation is a feature associated with HFT, it is neither necessary nor sufficient. As the literature has suggested different definitions of HFT, we repeat our trader characteristic breakdown based on two alternative definitions. Firstly, we apply the definition of HFT used in Hagströmer and Nordén (2013). That definition stipulates that HFTs are algorithmic traders that do proprietary trading only, i.e., that never trade for clients. Secondly, we use a definition based on trading activity similar to Kirilenko et. al. (2011). Here, a trader is defined as a HFT firm if, for 50% of the days they are active, they: (a) have end of day inventory less than 5% of their total volume; (b) have maximum net inventory at any time of the day less than 15% of total volume; (c) have trading volume in the top quartile of all proprietary traders that are active on that day. Furthermore, they are required to be active in more than 50% of all stock-days. In the Appendix, Table A1 reports the trader characteristics according to the alternative HFT definitions. Panel A holds results for the definition following Hagströmer and Nordén (2013), and Panel B the definition following Kirilenko et. al. (2011).

The most striking result in Table A1 is that the use of colocation status as an indicator of trading speed differs substantially from the competing HFT definitions. We find that not all HFTs, in either definition, are colocated, though a majority is. Moreover, there are numerous Non-HFTs (according to both definitions) that use colocation accounts. The HFT definition used in Panel B yields only nine HFT accounts, even though we use less restrictive criteria than in Kirilenko et al. (2011). The Panel A definition, in contrast, yields 64 HFT accounts, which is

more than the number of colocated accounts.⁵ We conclude that our definition of fast traders overlaps with the HFT definitions, but that the definitions are far from the same, at least when implemented in the current market setting.

Looking at the trader properties presented in Table A1, we see that most of the HFT trading and quoting volumes, according to Panel A, go through colocated accounts. Furthermore, q/t ratios are higher in colocated accounts, both for HFTs and Non-HFTs.

(b) Trader Speed and Adverse Selection

According to the theoretical literature, a speed advantage allows traders to react to public information faster than their competitors. For liquidity providers, this ability enables them to revise outstanding quotes as they become stale (Aït-Sahalia and Saglam, 2013; Hoffman, 2013; Jovanovic and Menkveld, 2012; Menkveld and Zoican, 2013). A relatively fast liquidity provider thus incurs less adverse selection costs, allowing for more aggressively priced quotes, and potentially a larger market share. For liquidity demanders, the speed advantage amounts to a better ability to trade against stale quotes, i.e., to trade on news before liquidity providers have time to revise their quotes (Cartea and Penalva, 2012; Foucault, Hombert and Roşu, 2013; Menkveld and Zoican, 2013).

We investigate differences between colocated and non-colocated traders' ability to trade aggressively on news (active side), and to avoid being picked off by revising stale quotes (passive side). We focus on the Colo and the NonColo trader groups and consider a number of measures of the bid-ask spread that reflect adverse selection in different ways. All spreads are expressed as fractions of the spread midpoint, which is recorded immediately before the trade. As active side

⁵ Hagströmer and Nordén (2013) have 29 HFT members in their paper based on the same market as the current paper. We record 64 HFT accounts, which is possible because: (i) one member may have several accounts; and (ii) as opposed to their paper, we consider sponsored access clients as separate trading entities.

traders pay the spread, low spreads indicate good performance. For passive side traders, who earn the spread, the opposite holds.

The *Quoted spread* is half the difference between the best bid and ask prices just before the trade; the *Effective spread* is the difference between the trade price and the spread midpoint; the *Realized spread* is the difference between the trade price and the spread midpoint at a certain horizon after the trade (horizons considered include 5 sec, 10 sec, 30 sec, 1 min, and 5 min). Midpoints are recorded every second, implying that a 1-second realized spread would be the difference between the trade price and the spread midpoint prevailing at the turn of the next second. Spreads are measured for all trades in the continuous trading of OMXS 30 index stocks during Oct. 1 – Oct. 26, 2012, i.e., the period after the colocation upgrade event.

To capture whether fast and slow traders differ in their ability to deal with adverse selection, we formulate the following ordinary least squares regression model:

$$y_{it} = \alpha + \beta Colo_{it} + \sum \theta_a Stock_{it}^a + \epsilon_{it}, \quad (1)$$

where y is the spread measure of interest for stock i on day t . $Colo$ takes the value one if the trade is by a colocated trader, and zero otherwise. $Stock_{it}^a$ is a set of dummy variables designed to capture stock fixed effects in the coefficient vector θ , with $a = 2, \dots, 30$. For each dependent variable the regression is performed twice. Once based on the trader-type that is taking liquidity, and again based on the trader-type that is supplying liquidity.

The regression results are reported in Table 3. Each row reports two regressions, with the dependent variable listed at the beginning of the row. The left half of the table reports the regression results based on active trading, and the right half holds regression results for passive trading.

INSERT TABLE 3 ABOUT HERE

Because there are only two types of traders, Colo and NonColo, and the variable *Colo* captures the differential adverse selection of the Colo trader, then the intercept, α , captures the effect of the NonColo traders. We adjust the intercept to reflect the average across all stocks, rather than just the benchmark stock, by adding the average stock fixed effect to the estimated intercept (see the Table 3 caption for definition). We report the NonColo spreads in the first column, and the Colo spreads, which is the coefficient on *Colo* (β , Column 3) plus the α in Column 1, in the second column. Column 4 shows the t-statistic for the difference between α and β (calculated using the Newey and West, 1987; 1994, heteroskedasticity and autocorrelation consistent, HAC, covariance matrix).

We focus on the results on Active trading first. Table 3 Row 1 displays the quoted spread. There is an economically small, but statistically significant difference between fast and slow traders. The quoted spreads are slightly wider (0.43 basis points) when slow traders demand liquidity than when fast traders do so. The second row reports the effective spread results. Difference between quoted and effective spreads emerge when (i) a market order is trading through the BBO, partially executing against volume deeper in the book, and when (ii) a market order hits hidden liquidity within the BBO. The former makes the effective spread wider than the quoted spread, whereas the latter makes it tighter. For active trading, we find that NonColo traders incur slightly wider effective than quoted spreads, indicating that they occasionally trade through the BBO. Colo traders record tighter effective spreads than quoted spreads. Finally, the next five rows repeat the analysis using the realized spread at different horizons.

The negative coefficients in Column 2 indicate that when fast traders demand liquidity, the price they pay is on average better than the spread midpoint shortly after the trade. That is, fast traders appear to have the ability to trade ahead of price changes, adversely selecting the liquidity providers. Following an active buy trade (sell trade) from a fast trader, the midpoint

price is increasing (decreasing) to 0.47 basis points above (below) the trade price already after 5 seconds, and the price keeps increasing (decreasing) at least one minute after the trade.

We think this result demonstrates that fast traders have the ability to profit on aggressive news trading, which is in line with theoretical models (Cartea and Penalva, 2012; Foucault, Hombert and Roşu, 2013; Menkveld and Zoican, 2013). Slow traders do not have this ability, as their trade price for buy trades (sell trades) is on average more than 0.4 basis points higher (lower) than the midpoint price on all considered horizons (see column one). The difference between fast and slow traders in active trading realized spreads is highly significant. It is also consistent with Brogaard, Hendershott and Riordan (2013), who show that high frequency traders' aggressive orders predict permanent price movements, whereas their passive orders do not.

The right-hand side of Table 3 repeats the analysis but for the trader-type on the passive side, the liquidity provider in the trades. Whereas a positive value in the previous results was taken as a cost, here a positive value indicates revenue. In terms of quoted and effective spreads, Colo traders charge higher spreads than NonColo traders do. Thus, fast traders, relative to slow traders, provide liquidity when it is expensive and demand it when it is cheap. An important aspect in liquidity provision is to avoid being picked off, to revise quotes quickly to reflect the latest news. The realized spreads presented for passive trading show that the spread earned by Colo traders remains above 0.13 basis points five minutes after the trade, whereas it is -0.65 basis points for NonColo traders. The difference is statistically significant at all horizons considered and demonstrates that fast traders are better than slow traders at revising their stale quotes. This result is in line with theoretical models put forward by Hoffman (2013), Jovanovic and Menkveld (2012), and Menkveld and Zoican (2013).

As Table 2 shows, algorithmic trades differ from manual trades, and proprietary trading has different characteristics than client trading. To make sure that the differences observed in Table 3, Panel A, are driven by colocation status and not by trade type, we repeat the exercise focusing exclusively on trades that are algorithmic and proprietary (Panel B). We consider the same liquidity measures as in Panel A for both aggressive and passive trades. Most results remain qualitatively the same. Active realized spreads are positive (negative) for NonColo (Colo) accounts. Passive realized spreads are negative (positive) for NonColo (Colo) accounts. Noteworthy, the coefficients on the realized spreads for Passive trades are uniformly higher. Proprietary algorithmic NonColo lose more, and proprietary algorithmic Colo earn more, than the general NonColo or Colo trading account.

Of all the ordering of coefficients, only the active quoted spread changes direction. Overall, NonColos have higher active quoted spreads than Colos, but when comparing the proprietary algorithmic NonColos and Colos, the Colos have a higher active quoted spread. The fact that this is the only measure that switches sign suggests that proprietary algorithmic NonColos focus on trading based on quoted spread, but that does not imply better executions than the colocated firms. In this case, the effective and realized spreads show that fast traders execute their trades at a lower cost.

Overall, the results presented in this section constitute novel evidence on the nature, behavior, and performance of fast traders. Our results are in line with common theoretical assumptions saying that fast traders use their speed advantage to pick off stale quotes and to avoid being picked off when providing liquidity.

In contrast to assumptions made in the theoretical literature, we find that fast traders use mixtures of strategies, both supplying and demanding liquidity. Given that investments in speed are costly, such diversification of strategies is in line with economies of scope. We also

find that there is substantial trading activity in several different segments of colocation services. The fact that not all colocated traders subscribe to the fastest technology demonstrates that traders may pursue latency-sensitive trading strategies even though some traders possess faster trading services. That indicates that the common limitation to two speed segments (fast and slow) may be a strong assumption. In the next section we study specifically how the introduction of a new trading technology influences the competition among fast traders.

IV. **The effects of a colocation upgrade**

In this section we investigate the effects of the introduction of 10G colocation at NOMX-St. The 10G colocation service is an add-on service to the Basic and 1G colocation packages, enabling traders to cut their latency by 20%.

An important insight of our findings in the previous section is that latency is a more diverse concept than the two categories of Colo and NonColo traders. Among the Colo traders, there are differences in reaction time due to differences in technology and programming skills. Furthermore, different strategies require different levels of analysis with different processing times. In this section we treat colocated traders who belong to the Basic or 1G colocation schemes as one group of traders, referred to as *SlowColo*. These traders supposedly have much smaller latency than non-colocated traders, but they decide that paying for the 10G upgrade is not worthwhile for their business. The colocated traders who do pay for the upgrade, to remain among the fastest traders in the market, are referred to as *10GColo*. To see how the colocation upgrade affects the competition among traders, we study the behavior of *SlowColo* and *10GColo*, as well as NonColo traders, before and after the date of the upgrade.

What predictions can be made regarding the effect of the colocation upgrade event? The colocation upgrade constitutes an opportunity to achieve a nominal latency improvement, which

any trader can pursue at a fixed cost.⁶ If we assume that agents are rational, only traders whose benefits of having cutting-edge technology exceed the fixed cost will pursue the upgrade. If the market is in equilibrium before the event, it is reasonable to expect that the latency shock to 10GColo traders improves their performance relative other traders. We analyze trading performance in terms of market share of active and passive trading, adverse selection costs in active and passive trading, as well as inventory management. Before turning to the empirical analysis, we use the current theoretical literature in combination with our previous results to form expectations on performance effects.

Expectation 1 (E1): 10GColo traders increase their market share. As discussed by Biais, Foucault and Moinas (2013), there are two reasons to believe that traders who pursue the colocation upgrade increase their trading opportunities. Firstly, as the cutting-edge level of technology improves, short-lived arbitrage opportunities that were previously not pursuable may become available. The 10GColo traders are the only market participants who can capture such opportunities. Secondly, the latency shock of the event opens up a latency gap between the 10GColo and the SlowColo traders, implying that the competition for trading opportunities is affected. Both these effects imply that the market share of 10GColo traders, all else equal, increases. According to Martinez and Rosu (2013), news traders who are no longer competitive in terms of speed would reorient their business away from latency sensitive news trading, to trade on other dimensions than the fastest traders. The increasing trading opportunities for fast traders are perhaps most intuitive for active strategies (e.g., trading on news), but they extend to passive trading too. For example, according to Cartea and Penalva (2012) faster traders are better able to tailor their intermediation of trades, leading to larger trading volumes. Furthermore, Hoffmann (2013) shows that the ability of fast traders to quickly revise their limit orders in response to news makes slower traders more exposed to adverse selection. The

⁶ The cost is the same for all traders, but traders who have not previously paid for colocation services would also need to pay for the Basic colocation package to pursue the upgrade.

consequence is that slower traders post limit orders with lower execution probability, leading to less trading volume.

Expectation 2 (E2): 10GColo traders impose more adverse selection costs on other traders and avoid being picked-off. The most central aspect in the theoretical literature on trading speed is the ability to react quickly to news. In active trading, lower latency allows traders to trade aggressively on public news by picking off stale quotes (Menkveld and Zoican, 2013; Foucault, Hombert, and Roşu, 2013). In passive trading, lower latency allows traders to revise their quotes faster, which lowers their adverse selection costs related to public news (Jovanovic and Menkveld, 2012; Hoffman, 2013). We expect the colocation upgrade to lead to a redistribution of adverse selection, primarily between the SlowColo and the 10GColo traders. Before the event, traders in these groups possess similar technology but are heterogeneous in terms of strategies. Due to the heterogeneity of strategies, the upgrade is worthwhile for some traders but not for others. The event leads to a latency gap between the two groups that did not exist, or was smaller, before the event. We expect the latency gap to improve the trading performance of both active and passive trading of the 10GColo traders.

Importantly, the expected effects of adverse selection costs and picking-off ability concern public information. Adverse selection costs are traditionally viewed as uniform across liquidity providers, as they all access the same public information but are exposed to privately informed liquidity demanders. When latency differs across traders, faster traders are able to revise or cancel their quotes in instantaneous response to news releases, whereas slower traders are not. Faster liquidity providers can thus avoid being adversely selected due to public news, but their speed advantage does not allow them to avoid being adversely selected by privately informed traders. We do not expect the event to affect adverse selection costs due to private information. Furthermore, we do not expect the event to affect adverse selection on neither active nor passive trading among NonColo traders, because they are already an order of

magnitude slower than colocated traders. In the view of the NonColo traders, the news release and the market reaction are simultaneous events, whereas for colocated traders they are sequential.

Expectation 3 (E3): 10GColo traders improve their inventory management. Inventory control is an important aspect in the theory of market-making. Still, even though passive trading is central to HFTs (Hagströmer and Nordén, 2013; Menkveld, 2013), inventory effects are not covered in the current theoretical literature on trading speed. We postulate that a latency advantage allows 10GColo traders to increase their inventory positions. As their control of the trading process is improved, their ability to turnaround quickly is also improved, reducing the inventory risk. Based on this reasoning, we expect them to have less mean-reversion in inventory, and to take larger inventory positions.

(a) Volume Effects

We now turn to the empirical investigation of the expectations outlined above. To investigate E1, we set up a regression model where the dependent variable (y_{it}) is either *active trading volume* or *passive trading volume* (both expressed in MSEK). To capture the change in trading activity following the colocation upgrade, we consider a panel with observations on each stock and each trading day. The regression model is defined as

$$y_{it} = \beta + \alpha Event_t + \sum \theta_a Stock_{it}^a + \epsilon_{it}, \quad (2)$$

where all regressors are dummy variables and ϵ_{it} denotes the unexplained variation of stock i ($i = 1, \dots, 30$) and trading day t (with t indexing all trading days in the Before and After periods). $Event_t$ is the event dummy, taking the value one for trading days in the After period and zero otherwise. $Stock_{it}^a$ is a set of dummy variables designed to capture stock fixed effects in the coefficient vector θ , with $a = 2, \dots, 30$. The estimated average trading activity per day and stock is

calculated for the Before period as the mean of the vector $(\beta, \beta + \theta)$, and for the After period as the mean of $(\alpha + \beta, \alpha + \beta + \theta)$.

As we are interested in how trading activity is redistributed following the colocation upgrade, we also specify a regression model on the market shares of each trader group. Thus, we add a dimension to the input data by disaggregating the stock-day trading activity into the activities of each trader group. We express these trading activity measures as market shares, considering again the active and the passive trading volumes separately. The regression model takes the form

$$y_{itg} = \beta_0 + \beta_1 \text{SlowColo}_{itg} + \beta_2 \text{10GColo}_{itg} + \alpha_0 \text{Event}_t \text{NonColo}_{itg} + \alpha_1 \text{Event}_t \text{SlowColo}_{itg} + \alpha_2 \text{Event}_t \text{10GColo}_{itg} + \epsilon_{itg} \quad (3)$$

where all regressors are dummy variables and ϵ_{itg} denotes the unexplained variation of stock i , trading day t , and trader group g ($g = \text{NonColo}, \text{SlowColo}, \text{10GColo}$). Event_t is the event dummy, taking the value one for trading days in the After period. NonColo_{itg} , SlowColo_{itg} , and 10GColo_{itg} are dummy variables indicating the trader group. Note that the trader classification is done in retrospect, such that those traders who subscribe to the 10G colocation belong to the 10GColo group already before the event. We do not consider stock fixed effects here, because the market shares of each stock-day sum to unity. The Before period statistics are as follows: β_0 for NonColo; $\beta_0 + \beta_1$ for SlowColo; and $\beta_0 + \beta_2$ for 10GColo. The After period statistics are retrieved as $\beta_0 + \alpha_0$ for NonColo; $\beta_0 + \beta_1 + \alpha_1$ for SlowColo; and $\beta_0 + \beta_2 + \alpha_2$ for 10GColo.

The results of the regressions specified in Equations (2) and (3) are given in Table 4. We report Before and After estimates along with t-tests of the difference between the two periods. The t-statistics for each test is based on the Newey and West (1987, 1994) HAC covariance matrix.

INSERT TABLE 4 ABOUT HERE

In Column 1 we report the nominal trading volumes before and after the exchange upgrade. As each trade has one active and one passive counterparty, the volumes on the active and passive sides are identical, but both are included as their distributions across trader groups differ. Our results show that the average stock-day trading volume increased by about 14%, but the change is not statistically significant.

Columns 2, 3, 4 show the market shares of each trader group before and after the event. For active trades the NonColo market share decreased from 54.9% to 50.1%. Both SlowColo and 10GColo increased their market share, by 3.4% and 1.4% respectively. All three changes are statistically significant at the 1% level. This is in line with expectation E1 in the sense that traders who become faster should increase their market share. The fact that SlowColo increased their market share even more, however, is not in accordance to our expectations.

For passive trades, NonColo traders dominate the market with 61.9%, but there is no difference between the two sample periods. There is a redistribution, however, in the market shares of SlowColo and 10GColo. Our results show that the traders who pursue the collocation upgrade capture 2.3% market share from the colocated traders who do not upgrade. This effect is statistically significant and completely in line with *E1*. Our interpretation is that the boost in trading speed allows 10GColo traders to improve their competitiveness in liquidity supply.

(b) Picking-Off and Adverse Selection Effects

We now turn to event effects on adverse selection. To investigate expectation E2, we run dummy variable regressions with adverse selection costs as dependent variables, using all trades in the continuous trading as our input data. For each trade we measure adverse selection costs

as the log midpoint change from just before the trade to 10 seconds or 5 minutes after the trade.⁷ This measure of adverse selection cost is in line with the specification in Menkveld and Zoican (2013). To see how the colocation upgrade influences adverse selection, we group trades by the combination of counterparties from the NonColo, SlowColo, and 10GColo groups. We set trades with NonColo traders on both the active and the passive side as our benchmark, and form dummy variables to represent each of the other eight combinations of trader groups. We organize the intercept and the trader combination dummies in a $(9 \times T)$ matrix denoted *TraderComb*, where T is the number of trades across all stocks and all trading days in the Before and After periods. To capture the impact of the colocation upgrade, we form another $(9 \times T)$ matrix denoted *EventTraderComb*. In this matrix, each column is a dummy variable for a trader combination, but it is zero for all trades happening in the Before period. Finally, we include 29 stock dummies to control for stock fixed effects, and nine hourly time-of-the-day dummies to control for intraday fixed effects. The benchmark case for the stock dummies is *ABB*, which is an average-sized firm relative our sample. For the time-of-the-day dummies the benchmark case is set to 1-2 pm. The model is specified as

$$y = \beta' \text{TraderComb} + \alpha' \text{EventTraderComb} + \theta' \text{Stock} + \rho' \text{Hour} + \epsilon, \quad (4)$$

where β and α are coefficient vectors of length nine reflecting before and after effects, θ and ρ are coefficient vectors of length 29 and 8 respectively that capture stock and intraday fixed effects, and ϵ is a residual vector of length T .

According to expectation *E2* a speed advantage may influence adverse selection costs from the perspective of both Active and Passive trading. An advantage with the regression specification outlined in Equation (4) is that we are able to analyze both the Active and Passive

⁷ 5 seconds, 30 seconds, and 1 minute adverse selection results are available from the authors upon request. The results of those regressions are qualitatively the same as for 10 seconds and 5 minutes. Trades where the midpoint change is unavailable due to trading halts or market closure are excluded.

trading in the same regression, as we capture all possible trader combinations. Table 5 presents the regression output. The table lists first the Active side participant and then the Passive side participant of each trader combination. The next two columns display results for the 10 second adverse selection, and the final two columns are for the five minute horizon.

The table output should be interpreted as follows. If we index the coefficients of β and α as $j = 1, \dots, 9$, the Before value for the benchmark trader combination, where NonColo traders are on both sides of the trade, is β_1 , and is equal to 2.187 for the 10 seconds horizon. That implies that during the Before period, such trades (in *ABB* between 1 pm and 2 pm) tend to have a price impact of 2.187 basis points. A positive coefficient indicates a gain (loss) to the active (passive) side of the trade. In the After column, the change from Before to After is given, in this case the estimate of $\alpha_1 = 0.061$, meaning that the After period adverse selection cost is $2.187 + 0.061 = 2.248$ basis points. For the other trader combinations ($j > 1$), the adverse selection of the *Before* period is $(\beta_1 + \beta_j)$; and for the After period it is $(\beta_1 + \beta_j + \alpha_j)$. Along with each coefficient estimate, we report a t-statistic that is based on the heteroskedasticity consistent standard errors (White, 1980).

INSERT TABLE 5 ABOUT HERE

As the results for the 10-second and 5-minute horizons are qualitatively the same we focus on the 10-second results. We find that before the event the highest price impact is recorded for trades where 10GColo traders are on the active side (1.536, 1.236, and 1.552 higher than the benchmark case when NonColo, SlowColo, and 10GColo trades, respectively, are on the passive side of the trade). Each of these differences to the benchmark case is statistically significant at the 1% confidence level. This indicates that the traders who upgraded are the traders who were fastest in the market already before the event, and are able to react quickly when public news is disclosed. Adverse selection when SlowColo traders are on the active side is

significantly higher than the benchmark case too, regardless of liquidity provider (1.019, 0.636, and 0.994 for NonColo, SlowColo, and 10GColo trades, respectively, on the passive side of the trade), but less so than for 10GColo traders.

After the event, SlowColo traders improve their active trading performance. Their adverse selection increases significantly regardless of counterparty, except for the 5-minute horizon where NonColo is on the passive side. Surprisingly, the active 10GColo firms' adverse selection decreases, in most cases significantly. Thus, our results are opposite to expectation *E2*.

From the perspective of the passive trader, the adverse selection costs should be as low as possible. Before the event, 10GColo traders have lower adverse selection than the SlowColo traders when supplying liquidity to NonColo traders. When a colocated trader is on the active side, in contrast, 10GColo traders incur higher adverse selection costs than SlowColo traders do. As this is a sign of poor per-trade performance in 10GColo market making, it may be a reason that they decided to invest in the colocation upgrade. After the event, the traders who upgraded are catching up with those who did not. To see this, note that the change recorded for Colo10G is lower than the change for SlowColo, regardless of who is on the active side of the trade and for both horizons. Thus, relative to SlowColo, the results on passive trading is in line with expectation *E2*.

Overall, the results presented in Table 5 are consistent with the results on market shares. In line with SlowColo traders' improved ability (adverse selection) in active trading, we saw that they increase their market share in Active trading. Similarly, as 10GColo improves its liquidity supply ability relative to SlowColo, 10GColo also gains passive trading market share at the expense of SlowColo. The results for SlowColo traders may be interpreted as a reorientation of their business from passive to active trading. That would be in line with the theoretical results

by Martinez and Rosu (2013), where HFTs who lose their edge in trading speed are predicted to change their orientation to less latency-sensitive trading strategies.

(c) Inventory Management Effects

Speed is used to manage inventory by quickly entering and exiting inventory positions while simultaneously avoiding being adversely selected. When public information arrives faster market-makers are able to cancel resting orders in the book and thereby avoid being “picked-off”, thereby lowering their inventory risks. Similarly when favorable liquidity conditions are presented faster market-makers are able to taken advantage of these, thereby lowering their liquidity costs. Taken together, speed allows faster traders to offer tighter bid-ask spreads due to the ability to manage inventory risk and lower their own liquidity costs. To investigate differential behavior in handling inventory among market participant types, we consider three measures of inventory management. We base all inventory metrics on net inventory, which is the sum of all shares traded up to a given point during the day, assuming that each trader starts the day flat, with buy quantities having a positive sign, and sell quantities having a negative sign. All metrics are measured on a trader-by-trader basis for proprietary trading only. *Inventory Cross Zero* is a count of how many times per day net inventory goes from positive to negative or vice versa. *Max Abs Inventory* is the natural logarithm of the maximum absolute net inventory recorded within an hour of the day. *Mean Abs Inventory* is the average time-weighted net inventory recorded in a day (that is, the daily average of the hourly Max Abs Inventory measure). Table 6 considers these three measures for NonColo, SlowColo, and 10GColo traders. The statistics presented in Panel A are equal-weighted averages across traders within each group, across constituent stocks of the OMXS 30 index, and across trading days during Oct 1 – 26. In the *SlowColo* column, “**” indicates that the given measure is significantly different from the corresponding *NonColo* measure, at the 1% level, according to a t-test. Similarly, in the

10GColo column, ‘**’ indicates that the given measure is significantly different from the corresponding *SlowColo* measure.

INSERT TABLE 6 ABOUT HERE

Table 6 shows that the subscribers of different types of technology handle their inventory differently. This result is particularly evident in the *Inventory Cross Zero* measure. NonColo accounts switch from positive to negative (or vice versa) inventory held about once a day. SlowColo accounts switch 3.34 times a day, while 10GColo firms do it 8.19 times. The mean inventory variable also varies among the trading types. We see that 10GColo accounts hold more than the double Mean Abs Inventory than NonColo trading accounts do (as the exponential scale 11.22 and 10.47, respectively, correspond to SEK 74,608 and SEK 35,242). For the hourly maximum inventory position we also observe that uniformly 10GColo accounts hold more inventory than SlowColo accounts, and the NonColo accounts hold the least amount.

In Table 7 we analyze how the inventory management differences between trader groups are affected by the colocation upgrade. For each of the three inventory measures we perform the following regression:

$$y_{itg} = \beta_0 + \beta_1 NonColo_{itg} + \beta_2 10GColo_{itg} + \alpha_0 Event_t SlowColo_{itg} + \alpha_1 Event_t NonColo_{itg} + \alpha_2 Event_t 10GColo_{itg} + \sum \theta_a Stock_{it}^a + \gamma Volatility_{it} + \epsilon_{itg}, \quad (5)$$

where ϵ_{itg} denotes the unexplained variation of stock i , trading day t , and trader group g , and all independent variables except $Volatility_{it}$ are dummies representing trader groups, the event, and stock fixed effects. $Volatility_{it}$ is the daily average of squared percentage log midpoint second-by-second changes, for each stock and each trading day. All dummy variables are defined as in previous regressions. In order to see the difference between SlowColo and 10GColo clearly, we set SlowColo as the benchmark case for the Before period. Columns 1, 2, and 3 use

the Inventory Cross Zero, the Mean Abs. Inventory, and the Max Abs. Inventory (hourly), respectively, as dependent variables. In Column 3 we include hourly dummies to control for intraday fixed effects. For all statistical tests, standard errors are estimated using the Newey and West (1987, 1994) HAC covariance matrix.

INSERT TABLE 7 ABOUT HERE

The results show that the differences between the trader groups existed already before the colocation upgrade. SlowColo has an Inventory CrossZero estimate of 3.31 in the Before period, and the corresponding 10GColo estimate is 11.75 higher than that. After the event, there is no significant change for SlowColo, but a significant decrease for 10GColo. The negative coefficient of -3.38 shows that the 10GColo accounts are crossing zero less often, which implies they are holding positions longer. The average inventory measure and the maximum inventory measure show that 10GColo accounts in the Before period take significantly larger inventory positions than SlowColo traders do. For both measures, the 10GColos increase after the event, suggesting that they are taking even larger positions. The results also indicate that SlowColo traders take larger positions. Finally, the NonColo interaction coefficient is only statistically significant (and positive) in the maximum inventory variable regression, but the coefficient is economically small compared to SlowColo and 10GColo. Overall, these results are in line with expectation *E3*, stating that the latency cut allows 10GColo traders to relax their inventory control.

(d) Trading Probability Effects

So far, we have analyzed trading activity, adverse selection costs, and inventory management in isolation. We now turn to a trade-by-trade analysis where we seek to explain the changes in trading activity (market shares) through the changes in adverse selection and inventory management. Specifically, we apply multinomial logit regressions to analyze market

shares of active and passive trading, conditional on hard news, inventory, and the state of the order book (liquidity and volatility). Our regression framework allows for a simultaneous investigation of expectation E_2 , that 10GColo traders impose more adverse selection costs on other traders, and avoid being picked off, and expectation E_3 , that 10GColo traders improve their inventory management, following the colocation upgrade.

We gather data on all trades during the continuous trading session, in all thirty stocks, from all trading days of the Before and the After periods. We use these data to run one regression for liquidity demand (*Active Trading*) and one regression for liquidity supply (*Passive Trading*). As above, Active Trading is where we focus on the counterparty submitting the market order leading to immediate execution, and Passive Trading is where we focus on the counterparty whose limit order is being hit by the market order. As a market order may be executed against several different limit orders, resulting in several sub-executions, we aggregate executions resulting from the same market order. The price of the aggregated trade is set to the value-weighted average price and the quantity is set to the sum of all sub-executions. For the passive side we do not aggregate in the same way, as different trader groups may supply liquidity to the same market order. Instead, for each market order, we aggregate all executions where the liquidity suppliers come from the same trader group. Having done this aggregation, we have 2,442,960 active trades and 3,035,063 passive trades.

The regression analysis pursued here is closely related to the probit regressions by Hendershott and Riordan (2012), who analyze differences between human and algorithmic traders. Important differences to their setup include that (i) we use three trader groups, which leads us to the multinomial logit rather than their two-group probit specification; (ii) we interact some of our variables with the Event dummy to investigate the influence of the colocation upgrade, in order to explicitly analyze our formulated expectations E_2 and E_3 ; and (iii) we add

inventory management to the set of explanatory variables, as well as stock dummies and hourly time-of-the day dummies.

The multinomial regression models are set up in the following way. Let y be a nominal dependent variable representing who initiates a trade (Active Trading) or who is supplying liquidity in a trade (Passive Trading), each with $J = 3$ categories of outcomes defined as $y = 1$ if NonColo, $y = 2$ if SlowColo, and $y = 3$ if 10GColo. Let $\Pr(y = m|\mathbf{x})$, $m = 1, 2, 3$ be the conditional probability of observing the outcome m given the explanatory variables \mathbf{x} . Assume that $\Pr(y = m|\mathbf{x})$ is a function of $\mathbf{x}\boldsymbol{\beta}_m$, where $\boldsymbol{\beta}_m = (\beta_{0,m} \cdots \beta_{k,m} \cdots \beta_{K,m})'$ includes an intercept term $\beta_{0,m}$ and coefficients $\beta_{k,m}$ for the effect of variable x_k on outcome m . The probabilities for trade i can be written as:

$$\Pr(y_i = m|\mathbf{x}_i) = \frac{1}{1 + \sum_{j=2}^J \exp(\mathbf{x}_i\boldsymbol{\beta}_j)} \text{ for } m = 1,$$

$$\Pr(y_i = m|\mathbf{x}_i) = \frac{\exp(\mathbf{x}_i\boldsymbol{\beta}_m)}{1 + \sum_{j=2}^J \exp(\mathbf{x}_i\boldsymbol{\beta}_j)} \text{ for } m > 1.$$

In line with Hendershott and Riordan (2012), we include four market quality variables as explanatory variables: the *Relative Spread* (the nominal spread divided by the spread midpoint, expressed in basis points), the *Depth* (the traded SEK amount required to have a price impact, averaged across the buy and the sell side and expressed in 0.1 MSEK), *Lagged Volatility* (the average squared one-second returns), and the *Lagged Volume* (expressed in 0.1 MSEK). The liquidity measures are based on the limit order book prevailing just before the trade, whereas the lagged volatility and volume measures are based on the ten seconds preceding the trade. We also use *Size*, the 0.1 MSEK value of the trade in question, as an explanatory variable.

To proxy for hard news we follow Hendershott and Riordan (2012) and use lagged returns from the index futures market. The idea here is that fast traders are able to watch the

news at the index level, where market-wide information shows up first, and profit from trading on the news at the cash market. Traders who increase their active trading on the buy-side (sell-side) following positive (negative) futures market returns are believed to pursue such strategies. Traders who take the other side of such trades are likely to be adversely selected, potentially because they do not react fast enough to revise their stale quotes. We use index futures returns from OMXS 30 index futures, the index constituted by our thirty sample stocks. We calculate returns over the last 10 whole seconds preceding each trade.⁸ To capture the direction of the news-trading, we separate futures returns into positive and negative returns. *Positive Returns* is zero when there are negative or zero futures returns and equal to the returns otherwise, multiplied by a *BuySell* indicator variable that is +1 if the trader in question is on the buying side and -1 otherwise. *Negative Returns* is constructed in the same way but for negative returns. This specification implies that positive regression coefficients imply increasing trading activity in the direction of the futures market news, for both positive and negative news. Thus, successful traders would have positive coefficients for their Active Trading, and negative coefficients for their Passive Trading.

To account for inventory management, we measure the net inventory accumulated over the trading day, recording the inventory status just before each trade. Traders that tend to buy stocks when they have a positive inventory, and to sell when they have a short position, are seen as position-builders. The opposite behavior implies a mean-reverting inventory management. To analyze such behavior, we construct the inventory variables in the same fashion as the futures market news variables. *Positive Inventory* is equal to the accumulated inventory (number of shares) for trades where it is positive and zero otherwise, multiplied by the prevailing bid-ask spread midpoint and the *BuySell* indicator. *Negative Inventory* is constructed in the same way but for trades where the inventory is negative. We do not aggregate

⁸ That is, if the trade happens 9:30:15.345 we use the futures return for 9:30:05.000 – 9:30.14.999.

inventory across all trading firms in the respective trader groups, as that measure would likely tend to be zero on average. Instead, inventory is recorded for the trading firm associated with each trade. The interpretation for the Active Trading regression output is that positive regression coefficients indicate a tendency of position-building behavior, for both positive and negative inventory, whereas negative coefficients indicate mean-reverting inventory management. For Passive Trading regressions the opposite holds, meaning that positive coefficients indicate mean-reversion and that negative coefficients indicate position-building.

As news trading and inventory management are two features that we, based on our previous results, believe are related to trader speed, we interact these variables with the Event dummy, which is zero for trades during the Before period and one for trades during the After period. The coefficients on the interaction terms capture the difference between the two periods. According to *E2*, we expect 10GColo traders to increase their market share of Active Trading, and decrease their market share of Passive Trading, following large positive and negative futures returns following the colocation upgrade. Moreover, from *E3*, we anticipate 10Colo traders to be less prone to inventory management, i.e., show less mean-reverting inventory behavior, after the colocation upgrade.

To be able to interpret the regression results economically, we compute the implied event probabilities for the case when all explanatory variables are equal to their sample means, and all the dummy variables are equal to zero. This calculation provides us with base case probabilities (market shares) for each trader group. For Active Trading, the base case market shares are 28.5% for NonColo traders, 36.5% for SlowColo traders, and 35.0% for 10GColo traders. For Passive Trading, the base case market shares are 33.8% for NonColo traders, 28.7% for SlowColo traders, and 37.4% for 10GColo traders. Our specification of the Event and the fixed effect dummies implies that these market shares are for trading in the ABB stock, between 1 pm and 2 pm on trading days during the Before period. Deviations from the base case are

presented in the regression output, given in Table 8. The regression coefficients are difficult to interpret in economic terms, but we still provide these along with the p-values of corresponding t-tests.⁹ To be able to interpret the economic magnitude of the results, we compute the marginal effect of each explanatory variable. The marginal effect is defined as the difference in probability, relative the base case, incurred by a one standard deviation ceteris paribus increase in the variable. Coefficients are reported for NonColo and 10GColo traders, and should be interpreted as relative those of SlowColo traders.

INSERT TABLE 8 ABOUT HERE

Our first point of interest in the Table 8 multinomial logit regression output is the Active Trading following news at the index futures market. The interpretation of the marginal effect coefficients here are that the market shares of Active Trading for NonColo and 10GColo traders, relative SlowColo traders, change by -1.4% and 1.5%, respectively, if the Positive Return increases by one standard deviation. The effect is even stronger following Negative Returns, -2.3% and 2.3%, respectively. These results are consistent with the idea that faster traders are better equipped to monitor the futures market, react to public information and pick off stale limit orders. In line with expectation *E2*, after the colocation upgrade event, where the 10GColo traders get even faster, the news trading pattern is amplified. As seen from the marginal effects of return variables interacted with the Event dummy, the market share of 10GColo accounts increases by another 0.8% following Positive Returns, and by another 0.5% following Negative Returns. At the same time, SlowColo traders also improve their market share relative NonColo traders.

The results for hard news are consistent with expectations on the Passive Trading side too. Here, the market share of 10GColo traders is 2.1% lower following Positive Returns, and

⁹ For brevity the coefficients on stock dummies and time-of-the-day dummies are not reported.

2.4% lower following Negative Returns. This result implies that, relative SlowColo traders, 10GColo traders are able to revise their quotes in response to news such that they are not being picked off. Our results show that SlowColo traders enjoy the same ability relative NonColo traders. The event effect here is small. The gap between 10GColo and SlowColo traders is not significantly different from zero. The gap to NonColo traders is affected, but in economic terms the event effect is negligible. Thus, our overall result with respect to news trading and the colocation upgrade, in association with expectation $E2$, is that the 10GColo traders maintain their edge in revising stale quotes but do not improve it, whereas they improve their ability to trade on the news by taking liquidity.

Our results for inventory management show that the faster a trader is, the more mean-reverting is its inventory. The marginal effects of a positive one standard deviation change in Positive Inventory (Negative Inventory) induce the 10GColo traders to lower their Active Trading market share by 1.5% (1.0%). On the Passive Trading side, the effect is even stronger. 10GColo traders decrease their Passive Trading market share by 3.5% (3.6%) when their long (short) inventory position is increased. Traders that post two-way quotes may lean against the wind by applying asymmetric pricing schemes on the two sides of the limit order book. On all four counts, NonColo traders react in the opposite direction to inventory.

After the colocation upgrade event, the inventory constraints of 10GColo traders relative SlowColo traders appear to be relaxed. Our findings show negative marginal effects for 10GColo traders when the inventory variables are interacted with the Event dummy. For Negative Inventory, the effect is substantial, 2.6% on the demand side and 4.2% on the supply side. An interpretation of this is that as the 10GColo traders are getting even faster, they grow more confident in their ability to manage their inventory. If so, they do not need to be as lean in inventory, because they know that they can turn around their position quicker. The findings

here are in line with expectation E_3 , and our results on inventory management in the previous subsection.

Among the market quality variables, the two liquidity measures show the strongest influence on the distribution of market shares across trader groups. 10GColo traders' market share of Active Trading is 4.0% lower when the Relative Spread is increased by one standard deviation. Instead, 10GColo traders increase their share of Passive Trading in that setting. When Depth is increased by one standard deviation, 10GColo traders decrease their share of Passive Trading by 2.3%. They decrease their share of Active Trading too, but much less so, 0.7%. Overall, these results are in line with the notion that 10GColo traders consume liquidity when it is cheap, and supply liquidity when it is expensive, leading to a stabilizing role in the limit order book liquidity. The opposite may be said about NonColo traders, who demand much more and supply much less liquidity when the Relative Spread is wide.

The effects of Lagged Volatility and Lagged Volume tend to be statistically significant, but the economic impact of these variables on market shares is in general small. It is notable that 10GColo traders increase both their demand and supply of liquidity following market volatility. We also note that the size of the trade in question is rather evenly distributed across the trader groups submitting the market order. On the other side of the trade, however, 10GColo traders have the ability to avoid supplying liquidity to large trades (-5.34%). This may be a result of anticipatory behavior, where the 10GColo trader predicts large incoming orders and move out of their way to avoid being adversely selected.

The results show that the technology upgrade changes participants' behavior. SlowColo traders take liquidity when it is inexpensive, and provide it when it is expensive. They also better predict prices, as seen in their realized spread sizes. SlowColo, and especially 10GColo, accounts are very active making up almost half of trading volume even though there are only half as many

firms. After the technology change all firms (except SlowColo) experienced a decrease in adverse selection costs. Inventory management also seemed to improve for both the 10GColo and the SlowColo accounts. These results are confirmed in the multinomial logit regression analysis. In addition, we see that past stock returns become even more important to 10GColo accounts after the technology upgrade.

V. Overall Market Quality

The analysis so far has been on the market dynamics between the different types of market participants and we take these findings as given moving forward. While we think the industrial organization of the market is important, there is theoretical justification that trading speed differentiation may impact market quality (e.g., Biais, Foucault and Moinas, 2013; Jovanovic and Menkveld, 2012). As such we direct our analysis towards how the overall market changes following the colocation upgrade event.

The effect of the upgrade on market quality is related to how traders incorporate the speed changes into their trading strategies. If the upgrade is used predominantly by informed liquidity-takers, spreads may widen due to increase in adverse selection. If the upgrade is used primarily by liquidity-suppliers to avoid being adversely selected and to take advantage of favorable market conditions when unwinding inventories, spreads may decrease as a result. In all likelihood the speed upgrade is being used by both demanders and suppliers and our results provide evidence on which of these two groups is having the larger impact on market quality.

We perform an event study around the introduction of the new colocation service. The regression specification is

$$y_{it} = \beta + \alpha Event_t + \sum \theta_a Stock_{it}^a + \epsilon_{it}, \quad (6)$$

where y_{it} is one of five market quality variables, $Event_t$ takes the value one after the event, and $Stock_{it}^a$ represent 29 stock dummy variables ($a = 2, \dots, 30$), implying that θ is a vector of stock fixed effects. Each market quality metric is observed for a stock i on a trading day t . The five different dependent variables are: *Bid-ask spread*, the spread divided by its midpoint, averaged across seconds and expressed in basis points; *Depth at BBO* and *Depth at 0.5%*, the MSEK trade volume required to change the price at all and by 0.5%, respectively, averaged across seconds and the buy and sell sides of the book; *Absolute return* is the average logged one-second midpoint changes, expressed in basis points; *Volatility*, the average squared logged one-second midpoint changes, multiplied by 10,000.

The results of our event regressions are given in Table 9. For each market quality variable, the Before period estimate given in Column 1 is the intercept adjusted for stock fixed effects (see the Table 9 caption for definition). The After period estimate is the adjusted intercept plus the event coefficient α . Columns 3-5 report the event change coefficient, along with a t-test showing whether the change is significantly different from zero. The t-statistics are calculated using the Newey and West (1987, 1994) HAC covariance matrix.

INSERT TABLE 9 ABOUT HERE

The results show consistently that the colocation upgrade event is followed by improvements in market quality, in particular in terms of liquidity. This is observed in a decrease in spreads (0.15 basis points), and an increase in Depth at the BBO (0.06 MSEK) and within 0.5% of the BBO (1.02 MSEK). All these differences are statistically significant at the 1% level. The two measures of volatility, Absolute return and Volatility, also show improvement, but the changes are not statistically significant.

Our results on liquidity are consistent with previous empirical evidence. Boehmer, Fong, and Wu (2012), who use a cross-sectional data set of 39 exchanges, find strong evidence of

improved liquidity and market efficiency following the introduction of colocation services. Furthermore, Frino, Mollica, and Webb (2013) study colocation services at the ASX, and find that the lower latency is associated with improved liquidity. Studies on technological enhancements of trading systems, that in general facilitate fast trading, have found mixed evidence with respect to liquidity. Hendershott, Jones, and Menkveld (2011), who study the automatization of quotes dissemination at the NYSE in 2003, and Riordan and Storckenmaier (2012), who investigate a trading system upgrade at the Deutsche Börse in 2007, find that liquidity is improved. Hendershott and Moulton (2011; studying the introduction of the Hybrid Market at NYSE in 2006), Gai, Yao, and Ye (2013; two latency-improving events at the NASDAQ in 2010), and Menkveld and Zoican (2013; the introduction of INET at the NASDAQ OMX Nordic exchanges in 2010), on the other hand, find negative effects on liquidity. For volatility, Boehmer, Fong, and Wu (2012) present evidence that the colocation upgrades amplify volatility, which is not in line with our results.

Our evidence of improved liquidity point to that the speed improvement has stronger effects on liquidity providing fast traders, than on liquidity demanders. If so, that is consistent with our previous results showing that 10GColos after the event gain passive trading market share from SlowColos. Whereas previous empirical studies on colocation and market quality are limited to viewing the colocation upgrade as a market-wide event, we are able to single out the trader groups who benefit from the event. Given the indications that liquidity suppliers is the main channel of market quality effects related to the event, we proceed by studying the quoting behavior of market-makers.

The inventory management abilities of market makers are an important factor in overall market-quality. In order to provide liquidity market-makers take on unwanted inventory positions for short periods of time hoping to re-sell these in the future at more favorable prices. The risks associated with holding inventory positions determine how much liquidity a market-

maker is willing to supply. Increasing the speed with which market-makers can manage their inventory should reduce the risk associated with any given level of inventory, thereby increasing the amount of liquidity they are willing to supply conditional on their inventory. The speed change we study should shed light on how speed affects inventory management and how inventory management changes are translated into liquidity and market quality.

We investigate how the quoting behavior conditional on current inventory is changing with the event. For each trading day and each stock, before and after the colocation upgrade event, we check every five minutes during continuous trading whether a trading account has orders posted at the current spread prices. As we want to focus on how market makers adapt their quoting behavior, we exclude stock-days where an account has no trading activity, and stock-days where an account records two-way quotes on less than 5% of the snapshots. The latter restriction is designed to exclude accounts who come to the market to acquire one-way positions. The exact timing of snapshots is randomized to avoid turn-of-the-minute or turn-of-the-second effects as documented by Hasbrouck and Saar (2013). Our procedure of measuring presence at the BBO is closely related to market maker presence metric introduced by Hagströmer and Nordén (2013). On a stock-day-account basis, we calculate *Bid or Ask Presence* as the fraction of snapshots where they have orders posted at either the best bid or the best ask prices. We also calculate and *Bid and Ask Presence*, as the fraction of snapshots where the account has two-way quotes posted at BBO. These two metrics proxy for general market-making activity.

According to, e.g., Aït-Sahalia and Saglam (2013), liquidity providers may adapt their quoting behavior when they are close to their inventory constraint. A well-known strategy to account for inventory is to post orders asymmetrically around the current price, in order to adjust the execution probabilities (known as leaning against the wind). We investigate that effect by studying Bid Presence and Ask Presence conditional on that the account has a high inventory

on either the long or short side. Inventory is measured contemporaneous to the randomized snapshots in terms of number of shares held, assuming that each account starts the day flat and trades only at NOMX-St. For each stock and each account, we record the top decile of long positions as well as the top decile for short positions. To account for time-series effects unrelated to the technology upgrade, we calibrate the decile limits for observations before the technology upgrade and again afterwards.

We refer to *Ask Presence* and *Bid Presence* as the fraction of snapshots where the account has an order posted at the ask and the bid price, respectively. We use those to form four metrics of how liquidity providers change their quoting behavior when they are close to their inventory constraints: (i) the Ask Presence when the account has a large long position minus the Ask Presence when the same account has a large short position; (ii) the Bid Presence when the account has a large short position minus the Bid Presence when the same account has a large long position; (iii) the Ask Presence minus the Bid Presence when the account has a large long position; and (iv) the Bid Presence minus the Ask Presence when the account has a large short positions. For all these measures, we expect positive differences if the trader is leaning against the wind, and zero differences if quoting activity is independent of inventory positions.

We investigate how quoting activity differ across trader groups and before and after the colocation upgrade event by estimation of the regression model

$$y_{itg} = \beta_0 + \beta_1 NonColo_{itg} + \beta_2 10GColo_{itg} + \alpha_0 Event_t SlowColo_{itg} + \alpha_1 Event_t NonColo_{itg} + \alpha_2 Event_t 10GColo_{itg} + \sum \theta_a Stock_{it}^a + \epsilon_{itg}, \quad (7)$$

where the dependent variable y_{itg} can be any of the six quoting activity metrics, observed for stock i , trading day t , and trader group g . The event dummy, trader type dummies, stock fixed effects, and residuals are defined as in Equation (5). The results are presented in Table 10, where each column corresponds to a regression of one of the six dependent variables.

INSERT TABLE 10 ABOUT HERE

The main variable of interest is the Event*10GColo, which in Table 10 is reported in the bottom half of the table (After) in the 10GColo row. The benchmark coefficient is that in the Before part of the table in the 10GColo row. Note that the reported coefficients are the sum of the relevant coefficients, so one can directly compare the before and after coefficients. The two first columns, reporting the general activity at the best bid and/or ask price, show that 10G Colo traders behave similarly before and after the upgrade in their frequency of providing competitive liquidity. Even though the coefficient on After for 10GColo in Column two is statistically significantly different from zero, its change from Before is economically small. The coefficient falls from 0.36 to 0.34, whereas the SlowColo presence decreases substantially from 0.37 to 0.26.

The remaining columns test how the fastest traders react to liquidity provision when they near their inventory bounds. Column 3 looks at the difference in Ask quoting when inventory is very high compared to when it is very low (negative). A larger coefficient shows that a market participant varies his behavior more based on his inventory position. Column 4 is similar except that it considers Bid quoting behavior. For both the Bid and the Ask analysis, the 10G Colos' coefficient drops significantly from 0.21 (0.22) to 0.11 (0.07) for the Ask (Bid) presence. With the additional speed, the market participant changes his behavior less based on inventory positions. This is consistent with more stable and liquid markets due to the lower latency. SlowColo traders decrease too, from 0.16 (0.09) to 0.05 (0.07) for the Ask (Bid) presence.

An alternative approach to address the same question of quoting behavior based on inventory positions is performed in Columns 5 and 6. In Column 5 we look at the difference in fraction of snapshots a market participant is at the best ask minus time at the best bid, conditional on have a large long position. Column 6 reports the corresponding (reverse)

difference for large short positions. These measures capture participants' aversion to continue taking on inventory even though they are near their inventory constraints. The results are consistent with inventory constraints being loosened. 10GColo traders decrease their sensitivity to their inventory boundaries. For long (short) positions the coefficient decreases significantly from 0.21 (0.23) to 0.11 (0.07). The quoting asymmetry decreases significantly for SlowColo too, from 0.10 (0.14) to 0.04 (0.08) for long (short) positions. The asymmetry in quoting policy is thus decreasing substantially more for the fastest traders.

The previous analyses show that for the same inventory levels, the fastest traders are less likely to reduce their liquidity provision after the colocation upgrade event. Holding the demand for liquidity constant this should, and appears to, lead to lower bid-ask spreads as market makers are able to produce more liquidity with the same inputs. The increased ability and willingness to hold inventory, particularly amongst the fastest traders, can also lead to reductions in transitory volatility as larger changes in liquidity demand can be absorbed without similarly large price concessions. This increase in inventory management capabilities may also lead to a decrease in price pressures (Hendershott and Menkveld, 2013) being applied to mean-reverting inventories, also leading to reductions in transitory volatility.

VI. Conclusion

Technology has revolutionized trading in financial markets. Most evidence suggests that the increased use of technology has led to improvements in liquidity and price efficiency. Little evidence exists on how traders with different levels of automation and speed compete with each other to demand and supply liquidity in a dynamic setting. We provide some of the first evidence on these interactions.

This paper analyzes how latency heterogeneity affects trading dynamics among market participants and how it influences market quality. In our empirical analyses, we access a

proprietary data set allowing us to disentangle the trading activities of different types of traders in detail. The unique data set enables us to analyze the industrial organization of trader speed, and how heterogeneous network connectivity influences market participants, and ultimately market quality.

The paper contributes to previous research in three major ways. First, we provide empirical evidence of how fast and slow traders behave. We find that fast traders trade actively ahead of news, thereby imposing adverse selection costs on their counterparties, and that they are good at avoiding being adversely selected when supplying liquidity. These results are consistent with the notion that the speed advantage amounts to an informational advantage. Second, we exploit an optional colocation upgrade in September 2012, which allows traders to reduce the roundtrip latency in order entry by more than 20% compared to the next best level of service, and find that faster traders impose higher adverse selection costs on other traders by improving their ability to trade actively on hard information. However, faster traders also increase their market share in liquidity provision and relax their inventory bands.

Third, we investigate how the colocation upgrade affects market quality. Overall, we find that market liquidity improves in terms of both tighter bid-ask spreads, and deeper limit order books, while short-term stock return volatility is unaffected by the upgrade. Relating the changes in market quality to the changes in trading activity and trading performance of the upgraded, faster, traders, we find that an increased market-making activity by the faster traders leads to liquidity improvements. Moreover, faster traders' improved inventory management significantly improves market liquidity.

An important concern about faster and colocated traders has also been the fear that they impose adverse selection costs on slower traders causing deteriorations in liquidity. We find that faster traders impose higher adverse selection costs on other traders overall but that their

contribution in terms of liquidity supply outweigh these costs. That fast traders are getting faster appears to improve liquidity for all market participants.

References

- Aït-Sahalia, Yacine, and Mehmet Saglam. 2013. High frequency traders: Taking advantage of speed. *Working paper*.
- Baron, Matthew, Jonathan Brogaard, and Andrei Kirilenko. 2012. The trading profits of high frequency traders. *Working paper*.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas. 2013. Equilibrium fast trading. *Working paper*.
- Boehmer, Ekkehart, Kingsley Fong, and Julie Wu. 2012. International evidence on algorithmic trading. *Working paper*.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan. 2013. High frequency trading and price discovery. *Working paper*.
- Cartea, Alvaro, and Jose Penalva. 2012. Where is the value in high frequency trading. *Quarterly Journal of Finance*, 2(3).
- Foucault, Thierry, Johan Hombert, and Ioanid Roşu. 2013. News trading and speed. *Working paper*.
- Frino, Alex, Vito Mollica, and Robert I. Webb. 2013. The impact of co-location of securities exchanges' and traders' computer servers on market liquidity. *Journal of Futures Market*. Forthcoming.
- Gai, Jiading, Chen Yao, and Mao Ye. 2013. The externalities of high frequency trading. *Working paper*.
- Hagströmer, Björn, and Lars Nordén. 2013. The diversity of high-frequency traders. *Journal of Financial Markets* 16: 741-770.
- Hasbrouck, Joel and Gideon Saar. 2013. Low-latency trading. *Journal of Financial Markets* 16: 646-679.
- Hendershott, Terrence, and Albert Menkveld. 2013. Price Pressures. *Journal of Financial Economics*, Forthcoming.
- Hendershott, Terrence, and Pamela Moulton. 2011. Automation, speed, and stock market quality: The NYSE's hybrid. *Journal of Financial Markets* 14: 568-604.

- Hendershott, Terrence, and Ryan Riordan. 2012. Algorithmic trading and the market for liquidity. *Journal of Financial and Quantitative Analysis*. Forthcoming.
- Hendershott, Terrence, Charles Jones, and Albert Menkveld. 2011. Does algorithmic trading improve liquidity? *Journal of Finance* 66(1): 1-33.
- Hoffmann, Peter. 2013. A dynamic limit order market with fast and slow traders. *Working paper*.
- Jovanovic, Boyan, and Albert Menkveld. 2012. Middlemen in limit-order markets. *Working paper*.
- Kirilenko, Andrei A., Albert Kyle, Mehrdad Samadi, and Tugkan Tuzun. 2011. The flash crash: The impact of high frequency trading on an electronic market. *Working paper*.
- Martinez, Victor Hugo, and Ioanid Rosu. 2013. High frequency traders, news and volatility. *Working paper*.
- Menkveld, Albert. 2013. High frequency trading and the new-market makers. *Journal of Financial Markets* 16: 712-740.
- Menkveld, Albert and Marius Zoican. 2013. Need for speed? Low latency trading and adverse selection. *Working paper*.
- Newey, Whitney and Kenneth West. 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55: 703–708.
- Newey, Whitney and Kenneth West. 1994. Automatic lag selection in covariance matrix estimation. *Review of Economic Studies*, 61: 631–653.
- Pagnotta, Emiliano and Thomas Philippon. 2012. Competing on speed. *Working paper*.
- Riordan, Ryan, and Andreas Storkenmaier. 2012. Latency, liquidity and price discovery. *Journal of Financial Markets* 15(4): 416-437.
- Scholtus, Martin and Dick van Dijk. 2012. High-frequency technical trading: The importance of speed. *Working paper*.
- Scholtus, Martin, Dick van Dijk, and Bart Frijns. 2012. Speed, algorithmic trading, and market quality around macroeconomic news announcements *Working paper*.
- White, Halbert. 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, 817-838.

Table 1: Stock Characteristics

The table lists properties of the OMXS30 constituents as of Oct. 31, 2012. All constituents except NOKI SEK are stocks that have NOMX-St as their primary trading venue. NOKI SEK is a Swedish Depositary Receipt issued by the Finnish firm Nokia Oyj. *Market capitalization* is based on the closing price of Oct. 31, 2012 (expressed in million Swedish Krona, *MSEK*). All other statistics are calculated as averages across trading days in the period Oct. 1 – Oct. 26, 2012. *Daily trading volume* is reported in MSEK and number of shares. *Bid-ask spread* is the spread divided by its midpoint, averaged across seconds and expressed in basis points. *Depth at BBO* and *Depth at 0.5%* is the MSEK trade volume required to change the price at all and by 0.5%, respectively, averaged across seconds and the buy and sell sides of the book. *Volatility* is the average squared one-second basis points returns.

Exchange code	Market cap. (MSEK)	Daily trading vol. (MSEK)	Daily turnover (%)	Bid-ask spread (bps)	Depth at BBO (MSEK)	Depth at 0.5% (MSEK)	Volatility (sq. bps returns)
HM B	327,921	437	0.13	2.8	1.16	15.77	0.27
NDA SEK	244,010	332	0.14	4.6	1.59	20.27	0.53
TLSN	189,095	279	0.15	2.3	0.38	10.27	0.25
ERIC B	177,424	415	0.23	4.4	1.61	13.51	0.56
SHB A	141,042	255	0.18	3.1	0.82	9.99	0.36
ATCO A	136,905	449	0.33	4.2	1.07	13.61	0.74
VOLV B	131,064	604	0.46	3.3	1.10	15.83	0.67
SEB A	119,351	260	0.22	5.1	1.13	8.36	0.63
SWED A	116,991	305	0.26	4.7	1.51	14.50	0.53
SAND	115,404	324	0.28	3.7	0.67	11.75	0.79
SCA B	78,848	261	0.33	4.6	1.23	9.25	0.49
ASSA B	77,757	159	0.20	3.6	0.50	6.43	0.40
INVE B	66,637	128	0.19	4.5	0.87	9.35	0.29
ABB	62,698	172	0.27	4.9	1.57	16.89	0.42
SKF B	61,699	356	0.58	4.1	0.91	10.22	0.77
ATCO B	56,699	138	0.24	5.8	0.79	9.60	0.77
AZN	51,375	197	0.38	2.8	0.87	14.98	0.25
ELUX B	51,030	249	0.49	4.2	0.66	8.83	0.68
SCV B	50,520	135	0.27	5.3	0.64	6.31	0.85
LUPE	50,484	130	0.26	4.6	0.58	7.30	0.53
ALFA	48,321	142	0.29	5.2	0.64	7.13	0.65
TEL2 B	46,909	145	0.31	4.9	1.06	7.49	0.41
SWMA	46,556	157	0.34	3.2	0.46	5.76	0.31
GETI B	45,366	94	0.21	4.2	0.40	4.64	0.36
SKA B	41,475	97	0.23	5.6	0.85	6.34	0.44
BOL	31,700	203	0.64	5.3	0.82	7.53	1.01
SECU B	16,783	44	0.26	5.2	0.28	3.34	0.30
MTG B	12,302	131	1.06	5.3	0.16	2.32	1.09
SSAB A	11,427	158	1.38	3.4	0.15	4.13	0.69
NOKI SEK	2,056	136	6.63	6.1	0.26	4.50	3.45

Table 2: Trader Characteristics

The table reports trading and quoting activity by colocation status. All messages entered during the continuous trading of OMXS30 stocks during Oct. 1 - Oct. 26, 2012, are considered. Trader grouping with respect to colocation type is based on the status on Dec. 3, 2012. Panel A compares colocated (*Colo*) with non-colocated (*NonColo*) traders. Panel B compares different segments of Colocated traders, as defined by different colocation types. Panel C disaggregates NonColo and Colo trading volume into client and proprietary order flows. Panel D does the same with respect algorithmic and other order entry. All statistics are based on sums across stocks and trading days. *Number of accounts* is the number of trading entities considered in each trader group. The following statistics are reported as the share of total volume: *Trading volume* (the SEK volume), *Trades* (the number of executions), *Submissions*, *Cancellations*, and *Modifications* (the number of limit orders posted, cancelled, and modified), *Active trades* and *Passive trades* (the number of executions where the trader participates through a market order and limit order, respectively). The following three ratios are calculated within each group: *q/t* is the sum of all submissions, cancellations, and modifications, divided by the number of executions; *Liq. supply* is the fraction of *Trading volume (Trades)* where the trader participates through a limit order.

	Number of accounts	Share of total volume							Ratios within group		
		Trading volume	Trades	Submissions	Cancellations	Modifications	Active trades	Passive trades	q/t	Liq. supply (volume)	Liq. supply (trades)
<i>Panel A: NonColo and Colo trading</i>											
NonColo	112	58.7%	56.3%	17.3%	14.7%	2.0%	54.3%	58.2%	6.0	54.3%	51.7%
Colo	58	41.3%	43.7%	82.7%	85.3%	98.0%	45.7%	41.8%	40.8	43.8%	47.8%
<i>Panel B: Different segments of colocated (Colo) trading</i>											
BasicColo	17	3.5%	3.2%	1.6%	1.1%	6.3%	1.9%	4.5%	9.3	67.9%	69.7%
PremiumColo	22	16.5%	17.7%	24.7%	22.3%	0.0%	20.4%	15.0%	27.7	37.6%	42.4%
10GColo	19	21.3%	22.8%	56.4%	61.9%	91.7%	23.3%	22.3%	55.5	44.7%	48.9%
<i>Panel C: NonColo and Colo trading broken down by capacity</i>											
NonColo client		34.1%	33.0%	9.7%	8.6%	2.0%	34.3%	31.8%	5.8	51.0%	48.1%
Colo client		0.7%	0.7%	0.1%	0.1%	0.0%	1.1%	0.3%	3.8	11.8%	19.9%
NonColo proprietary		24.6%	23.2%	7.6%	6.1%	0.0%	20.1%	26.4%	6.1	58.9%	56.8%
Colo proprietary		40.6%	43.1%	82.5%	85.2%	98.0%	44.6%	41.5%	41.4	44.4%	48.2%
<i>Panel D: NonColo and Colo trading broken down by order entry method</i>											
NonColo algorithmic		9.2%	9.4%	7.9%	7.4%	2.0%	10.0%	8.8%	17.1	60.9%	46.9%
Colo algorithmic		36.9%	39.1%	80.4%	83.6%	91.7%	41.8%	36.3%	44.6	42.4%	46.5%
NonColo other		49.5%	46.8%	9.4%	7.3%	0.0%	44.3%	49.3%	3.7	53.1%	52.7%
Colo other		4.4%	4.7%	2.3%	1.7%	6.3%	3.9%	5.5%	9.3	55.6%	58.9%

Table 3: Quoted, Effective and Realized Spreads

The table reports quoted, effective, and realized spreads by colocation status. Spreads are measured for all trades in the continuous trading of OMXS30 stocks during Oct. 1 - Oct. 26, 2012. Trader grouping with respect to colocation type is based on trader colocation status on Dec. 3, 2012. The trader that submits a market order is on the active side and pays the spread, and the counterparty is on the passive side and earns the spread. All spreads are expressed as fractions of the spread midpoint, which is recorded immediately before the trade. The *Quoted spread* is half the difference between the best bid and ask prices just before the trade; the *Effective spread* is the difference between the trade price and the spread midpoint; the *Realized spread* is the difference between the trade price and spread midpoint at some horizon after the trade (horizons considered include 5 sec, 10 sec, 30 sec, 1 min, and 5 min). Midpoints are recorded every second, implying that the 1-second realized spread is the difference between trade price and the spread midpoint prevailing at the turn of the next second. The spread difference between non-colocated and colocated traders (*NonColo* and *Colo*) is determined in a dummy regression analysis of stock-day observations, including stock dummies to control for fixed effects. The regression specification is $y_{it} = \alpha + \beta \text{Colo}_{it} + \sum \theta_i \text{Stock}_{it} + \epsilon_{it}$, where ϵ_{it} are residuals. The reported intercept is adjusted for stock fixed effects as follows: $\alpha^* = \alpha + \overline{(0, \theta)}$. The *t*-statistics are based on the Newey and West (1987, 1994) HAC covariance matrix. Panel A shows results for all eligible trades, whereas Panel B shows results for the subset of trades that are algorithmic and proprietary.

<i>Panel A: NonColo v. Colo (all trades)</i>								
	Active side spreads				Passive side spreads			
	NonColo (α^*)	Colo ($\alpha^* + \beta$)	Diff (β)	t stat. for β	NonColo (α^*)	Colo ($\alpha^* + \beta$)	Diff (β)	t stat. for β
Quoted spread	3.90	3.47	-0.43	-12.2	3.59	3.83	0.23	7.2
Effective spread	4.18	3.41	-0.77	-17.4	3.60	4.11	0.51	11.9
Realized spread, 5 sec	1.01	-0.47	-1.48	-20.8	-0.03	0.63	0.66	13.8
10 sec	0.85	-0.72	-1.57	-20.6	-0.30	0.50	0.80	13.8
30 sec	0.60	-1.07	-1.67	-21.3	-0.68	0.30	0.98	13.2
1 min	0.44	-1.25	-1.70	-20.7	-0.87	0.15	1.02	13.0
5 min	0.50	-1.02	-1.52	-15.1	-0.65	0.13	0.79	7.1

<i>Panel B: NonColo v. Colo (algorithmic trades for proprietary purpose)</i>								
	Active side spreads				Passive side spreads			
	NonColo (α^*)	Colo ($\alpha^* + \beta$)	Diff (β)	t stat. for β	NonColo (α^*)	Colo ($\alpha^* + \beta$)	Diff (β)	t stat. for β
Quoted spread	3.40	3.44	0.05	2.6	3.62	3.90	0.28	6.4
Effective spread	3.40	3.31	-0.09	-5.1	3.79	4.23	0.44	6.9
Realized spread, 5 sec	1.09	-0.56	-1.65	-15.1	-0.71	0.78	1.49	24.2
10 sec	0.86	-0.81	-1.67	-15.8	-0.93	0.67	1.60	21.4
30 sec	0.68	-1.18	-1.86	-11.6	-1.36	0.52	1.88	22.2
1 min	0.47	-1.35	-1.82	-9.4	-1.56	0.38	1.94	22.4
5 min	0.36	-1.09	-1.45	-4.7	-1.07	0.30	1.37	9.0

Table 4: Trading Volume and Market Share

The table reports nominal volumes and market shares by colocation status, before and after the colocation upgrade. Each statistic is measured in terms of the number of active and passive trades and the number of actively and passively traded stocks. A counterparty who submits a market order is considered to be on the active side, and a counterparty whose limit order is hit is considered to be on the passive side. All trades in the continuous trading of OMXS30 stocks are considered. Nominal volumes is observed for each stock and each trading day in the Before and After periods, and assessed in regressions specified as $y_{it} = \beta + \alpha Event_t + \sum \theta_a Stock_{it}^a + \epsilon_{it}$, where ϵ_{it} are residuals and the independent variables are dummies representing the event and stock fixed effects. The intercept is adjusted for stock fixed effects as follows: $\beta^* = \beta + \overline{(0, \theta)}$, and reported as the Before estimate. The After estimate is $\beta^* + \alpha$, and the difference with associated tests is based on α . For relative volumes, we consider three trader groups based on trader colocation status on Dec. 3, 2012: non-located (*NonColo*), collocated who did not upgrade (*SlowColo*), and collocated who did upgrade (*10GColo*). Observations for each stock, trading day, and trader group, are considered in regressions of the form $y_{itg} = \beta_0 + \beta_1 SlowColo_{itg} + \beta_2 10GColo_{itg} + \alpha_0 Event_t NonColo_{itg} + \alpha_1 Event_t SlowColo_{itg} + \alpha_2 Event_t 10GColo_{itg} + \epsilon_{itg}$, where ϵ_{itg} are residuals and all independent variables are dummies representing trader groups and the event. Reported statistics for the *Before* period are as follows: β_0 for NonColo; $\beta_0 + \beta_1$ for SlowColo; and $\beta_0 + \beta_2$ for 10GColo. Reported statistics for the *After* period are retrieved as $\beta_0 + \alpha_0$ for NonColo; $\beta_0 + \beta_1 + \alpha_1$ for SlowColo; and $\beta_0 + \beta_2 + \alpha_2$ for 10GColo. Statistical tests for differences between the Before and After periods are based on α_0 , α_1 , and α_2 . All t-tests are based on the HAC covariance matrix (Newey and West, 1987; 1994).

		Nominal volume	Market share		
			NonColo	SlowColo	10GColo
Active trading volume	Before	207,796,700	0.549	0.211	0.240
	After	234,892,800	0.501	0.245	0.254
	Diff.	27,096,070	-0.048	0.034	0.014
	t	1.10	-5.06	3.51	2.53
	p	0.27	0.00	0.00	0.01
Passive trading volume	Before	207,796,700	0.619	0.201	0.181
	After	234,892,800	0.618	0.178	0.203
	Diff.	27,096,070	0.000	-0.023	0.023
	t	1.10	-0.06	-5.56	6.15
	p	0.27	0.95	0.00	0.00

Table 5: Adverse Selection

The table reports adverse selection costs by collocation status of the active side as well as the collocation status of its counterparties, before and after the collocation upgrade. The adverse selection cost of a trade is calculated as the logged bid-ask spread midpoint change from just before the trade to some horizon after the trade (horizons considered include 10 sec and 5 min). All trades in the continuous trading of OMXS30 stocks are considered. Trader grouping with respect to collocation type is based on trader collocation status on Dec. 3, 2012. Three groups are considered: non-located (*NonColo*), collocated who did not upgrade (*SlowColo*), and collocated who did upgrade (*10GColo*). The *Before* period is Aug. 20 – Sept. 14, 2012, and the *After* period is Oct. 1 – Oct. 26, 2012. Differences between pairs of trading partners and between periods are reported based on a regression of adverse selection costs on a set of dummy variables: $y = \beta' \text{TraderComb} + \alpha' \text{EventTraderComb} + \theta' \text{Stock} + \rho' \text{Hour} + \epsilon$, where ϵ are residuals and all independent variables are dummies representing trader group combinations, time-of-the-day and stock fixed effects. *TraderComb* is a matrix where the first column is an intercept used as benchmark and corresponding to trades where NonColo traders trade with each other. The other eight columns are dummies representing other trader combinations. *EventTraderComb* is a matrix where each column is zero for trades in the Before period, and dummies for trader combinations in the After period. Benchmark stock for the fixed effect dummies is *ABB*, and benchmark time-of-the-day is 1-2 pm. Estimates of the hourly and stock dummies coefficient vectors, ρ and θ are not reported. T-statistics are given within parentheses. The *t*-statistics (given within parentheses) are estimated using heteroskedasticity consistent White (1980) standard errors.

		10 sec		5 min	
		Before (β)	After (α)	Before (β)	After (α)
Benchmark:	NonColo vs. NonColo	2.187 (133.73)	0.061 (6.56)	2.891 (54.63)	-0.110 (-2.99)
Active side	Passive side				
NonColo	SlowColo	0.346 (26.97)	0.093 (4.64)	-0.218 (-4.52)	0.385 (5.09)
	10Gcolo	0.236 (19.04)	-0.037 (-1.96)	-0.206 (-4.49)	0.197 (2.68)
SlowColo	NonColo	1.019 (86.30)	0.302 (16.98)	0.831 (18.78)	0.045 (0.66)
	SlowColo	0.636 (23.91)	0.838 (18.00)	-0.029 (-0.29)	1.236 (8.06)
	10Gcolo	0.994 (41.52)	0.356 (9.67)	0.270 (3.25)	0.627 (4.91)
10GColo	NonColo	1.536 (124.84)	-0.299 (-16.86)	1.666 (36.56)	-0.479 (-7.01)
	SlowColo	1.236 (43.24)	0.036 (0.91)	1.524 (14.32)	-0.485 (-3.47)
	10GColo	1.552 (51.27)	-0.348 (-8.62)	1.635 (15.69)	-0.684 (-5.07)

Table 6: Inventory Management

The table presents metrics of inventory management by colocation status. Trader grouping with respect to colocation type is based on trader colocation status on Dec. 3, 2012. Three groups are considered: non-colocated (*NonColo*), colocated who did not upgrade (*SlowColo*), and colocated who did upgrade (*10GColo*). We base all inventory metrics on net inventory, which is the sum of all shares traded up to a given point in the day, assuming that each trader starts the day flat, that buy quantities have a positive sign, and that sell quantities have a negative sign. All metrics are measured on a trader-by-trader basis for proprietary trading only. *Inventory Cross Zero* is a count of how many times per day net inventory goes from positive to negative or vice versa. *Max Abs Inventory* is the maximum log net inventory recorded on an hourly basis. *Mean Abs Inventory* is the daily mean across the hourly Max Abs Inventory observations. The statistics are equal-weighted averages across traders within each group, across constituent stocks of the OMXS 30 index, and across trading days in the time period Oct. 1 – Oct. 26, 2012. In the *SlowColo* column, ‘**’ indicates that the given measure is significantly different from the corresponding *NonColo* measure, at the 1% level, according to a t-test. Similarly, in the *10GColo* column, ‘**’ indicates that the given measure is significantly different from the corresponding *SlowColo* measure, at the 1% level.

Inventory measures	NonColo	SlowColo	10GColo
Inventory Cross Zero	1.04	3.34**	8.19**
Mean Abs Inventory (log SEK)	10.47	10.83**	11.22**
Max Abs Inventory (log SEK)			
9am – 10am	8.96	10.11**	10.61**
10am – 11am	9.85	10.50**	10.91**
11am – 12pm	10.22	10.57**	11.03**
12pm – 1pm	10.35	10.61**	11.06**
1pm – 2pm	10.47	10.69**	11.13**
2pm – 3 pm	10.74	10.92**	11.39**
3pm – 4pm	11.10	11.12	11.56**
4pm – 5pm	11.36	11.44**	11.73**
5 pm – 5.30pm	11.15	11.48**	11.56**

Table 7: Change in Inventory Management Behavior

The table presents metrics of inventory management before and after the colocation upgrade. Trader grouping with respect to colocation type is based on trader colocation status on Dec. 3, 2012. Three groups are considered: non-located (*NonColo*), collocated who did not upgrade (*SlowColo*), and collocated who did upgrade (*10GColo*). We base all inventory metrics on net inventory, which is the sum of all shares traded up to a given point in the day, assuming that each trader starts the day flat, that buy quantities have a positive sign, and that sell quantities have a negative sign. All metrics are measured on a trader-by-trader basis for proprietary trading only. *Inventory Cross Zero* is a count of how many times per day net inventory goes from positive to negative or vice versa. *Max Abs Inventory* is the maximum log net inventory recorded on an hourly basis. *Mean Abs Inventory* is the daily mean across the hourly Max Abs Inventory observations. Statistics are based on regressions of stock-day observations, which in turn are averages across traders within each group. The regressions are specified as $y_{it} = \beta_0 + \beta_1 NonColo_{it} + \beta_2 10GColo_{it} + \alpha_0 Event_{it} SlowColo_{it} + \alpha_1 Event_{it} NonColo_{it} + \alpha_2 Event_{it} 10GColo_{it} + \sum \theta_i Stock_{it} + \gamma Volatility_{it} + \epsilon_{it}$, where ϵ_{it} are residuals and all independent variables except $Volatility_{it}$ are dummies representing trader groups, the event, and stock fixed effects. $Volatility_{it}$ is the daily average of squared percentage log midpoint second-by-second changes, for each stock and each trading day. The *Before* period is Aug. 20 – Sept. 14, 2012, and the *After* period is Oct. 1 – Oct. 26, 2012. Column (1), (2), and (3) use the Inventory Cross Zero, the Mean Abs. Inventory, and the Max Abs. Inventory, respectively, as dependent variables. In specification (3) we include hourly dummies to control for intraday fixed effects. Standard errors are estimated using the Newey and West (1987, 1994) HAC covariance matrix.

Inventory management before and after the colocation upgrade			
	(1)	(2)	(3)
	Inventory CrossZero	AbsMeanInv (daily)	AbsMaxInv. (hourly)
<i>Intercept</i>	3.31 (13.34)	10.46 (135.96)	10.44 (272.62)
<i>NonColo</i>	-7.01 (-19.28)	0.00 (0.05)	0.00 (0.09)
<i>10GColo</i>	11.75 (34.73)	0.90 (25.96)	0.90 (47.77)
<i>SlowColo × Event</i>	-0.02 (-0.20)	0.28 (6.82)	0.28 (14.50)
<i>NonColo × Event</i>	0.08 (1.18)	0.05 (1.43)	0.05 (2.55)
<i>10GColo × Event</i>	-3.38 (-37.91)	0.12 (2.91)	0.12 (6.23)
<i>Volatility</i>	0.83 (3.73)	0.03 (1.21)	-0.01 (-1.41)
Number of obs.	40287	40254	362286
R^2	0.136	0.048	0.061
$Adj. R^2$	0.135	0.048	0.061
Stock fixed effects	Yes	Yes	Yes
Hourly fixed effects	No	No	Yes

Table 8: Multinomial Logit Regression on Trading Probabilities

The table presents regression results for the probability of trades from each trader group. Let y be a nominal dependent variable representing who initiates a trade (Active Trading) or who is supplying liquidity in a trade (Passive Trading), each with $J = 3$ categories of outcomes defined as $y = 1$ if NonColo, $y = 2$ if SlowColo, and $y = 3$ if 10GColo. Let $\Pr(y = m|\mathbf{x})$, $m = 1, 2, 3$ be the conditional probability of observing the outcome m given the explanatory variables \mathbf{x} . Assume that $\Pr(y = m|\mathbf{x})$ is a function of $\mathbf{x}\boldsymbol{\beta}_m$, where $\boldsymbol{\beta}_m = (\beta_{0,m} \cdots \beta_{k,m} \cdots \beta_{K,m})'$ includes an intercept term $\beta_{0,m}$ and coefficients $\beta_{k,m}$ for the effect of variable x_k on outcome m . The probabilities for trade i can be written as:

$$\Pr(y_i = m|\mathbf{x}_i) = \frac{1}{1 + \sum_{j=2}^J \exp(\mathbf{x}_i \boldsymbol{\beta}_j)} \text{ for } m = 1$$

$$\Pr(y_i = m|\mathbf{x}_i) = \frac{\exp(\mathbf{x}_i \boldsymbol{\beta}_m)}{1 + \sum_{j=2}^J \exp(\mathbf{x}_i \boldsymbol{\beta}_j)} \text{ for } m > 1$$

The variables in \mathbf{x} are Relative Spread (the difference between the best ask quote and the best bid quote, divided by the midpoint quote, $\times 10,000$), Size (SEK trading volume, divided by 100,000), Depth (the average number of shares available at the best bid quote and the best ask quote, multiplied with the midpoint quote, divided by 100,000), Positive Futures Return, Negative Futures Return, Lagged Volatility (realized volatility measured as average squared second-by-second returns over the previous ten seconds, $\times 1,000$), Lagged Volume (SEK trading volume over the previous ten seconds, $\times 100,000$). The positive and negative futures returns are defined as $Rtn_{t,t-10}^+ \times BuySell$ and $Rtn_{t,t-10}^- \times BuySell$, where $BuySell$ is a trade indicator variable equal to +1 if the observed trade is buyer-initiated, and equal to -1 if the trade is seller-initiated, and $Rtn_{t,t-10} = \ln(midpoint_t) - \ln(midpoint_{t-10})$ is defined as the log change in futures midpoint quotes during a 10-second window preceding the time t stock trade. Fixed (stock) effects and time of day dummy variables for each hour of the trading day are included in the regressions, but are not reported. Statistical significance of given estimates are indicated with ‘**’ for the 95% and ‘*’ for the 90% confidence level. Marginal Effect denotes the change in probability caused by a one standard deviation increase in each variable relative the average level.

Table 8 cont.

Variable		Active trading		Passive trading	
		NonColo	10GColo	NonColo	10GColo
Positive Return	Coefficient	-3.35**	3.02**	8.66**	-1.55**
	Marginal Effect	-0.014	0.015	0.031	-0.021
Negative Return	Coefficient	-5.24**	4.35**	8.93**	-1.83**
	Marginal Effect	-0.023	0.023	0.035	-0.024
Positive Return (After)	Coefficient	-6.15**	-0.48*	-0.71**	0.05
	Marginal Effect	-0.018	0.008	-0.003	0.002
Negative Return (After)	Coefficient	-6.51**	-1.33**	0.99**	0.33
	Marginal Effect	-0.018	0.005	0.003	-0.001
Positive Inventory	Coefficient	2.95**	0.32**	-3.79**	0.31**
	Marginal Effect	0.039	-0.015	-0.056	0.035
Negative Inventory	Coefficient	2.70**	0.65**	-3.93**	0.11*
	Marginal Effect	0.036	-0.010	-0.061	0.036
Positive Inventory (After)	Coefficient	-0.63**	0.10	1.09**	0.27**
	Marginal Effect	-0.010	0.006	0.013	-0.004
Negative Inventory (After)	Coefficient	-0.88**	1.18**	1.74**	-1.61**
	Marginal Effect	-0.023	0.026	0.039	-0.042
Relative Spread	Coefficient	0.13**	0.01**	-0.09**	-0.02**
	Marginal Effect	0.095	-0.040	-0.056	0.019
Depth	Coefficient	0.20**	0.04**	0.26**	-0.02**
	Marginal Effect	0.022	-0.007	0.037	-0.023
Lagged Volatility	Coefficient	-0.03**	0.00	0.01**	0.01**
	Marginal Effect	-0.021	0.010	0.001	0.003
Lagged Volume	Coefficient	0.05**	0.05**	-0.03**	-0.00
	Marginal Effect	0.004	0.005	-0.004	0.002
Size	Coefficient	-0.04**	-0.53**	1.14**	-0.58**
	Marginal Effect	0.009	-0.022	0.065	-0.053

Table 9: Market Quality

The table presents market quality metrics before and after the colocation upgrade. Each metric is calculated for each stock in the OMXS30 index, and for each trading day. *Bid-ask spread* is the half-spread divided by its midpoint, averaged across seconds and expressed in basis points. *Depth at BBO* and *Depth at 0.5%* is the MSEK trade volume required to change the price at all and by 0.5%, respectively, averaged across seconds and the buy and sell sides of the book. *Absolute return* is the average logged one-second midpoint changes, expressed in basis points. *Volatility* is the average squared logged one-second basis points returns. The difference between the *Before* (Aug. 20 – Sept. 14, 2012) and *After* (Oct. 1 – Oct. 26, 2012) periods is determined in a dummy regression analysis of stock-day observations, including stock dummies to control for fixed effects. The regression specification is $y_{it} = \beta + \alpha Event_{it} + \sum \theta_a Stock_{it}^a + \epsilon_{it}$, where ϵ_{it} are residuals and $Event_{it}$ is a dummy variable for the colocation upgrade event. The reported intercept is adjusted for stock fixed effects as follows: $\beta^* = \beta + \overline{(0, \theta)}$. Standard errors used for statistical tests are estimated using the Newey and West (1987, 1994) HAC covariance matrix.

	Before (β^*)	After ($\beta^* + \alpha$)	Diff. (α)	t stat. for α	p value for α
Bid-ask spread (bps)	4.52	4.37	-0.15	-3.14	0.00
Depth at BBO (MSEK)	0.76	0.82	0.06	4.27	0.00
Depth at 0.5% (MSEK)	8.52	9.54	1.02	7.95	0.00
Absolute return (bps)	0.12	0.11	-0.01	-1.38	0.17
Volatility (sq. bps returns)	0.72	0.65	-0.07	-1.35	0.18

Table 10: Why Market Quality Improves

The table presents metrics of quoting activity at the best bid and ask prices by different trader groups, before and after the colocation event, and conditional on the inventory of the trader. For each trading day, each stock, and each trading account, we check every five minutes whether an order is posted at the best bid and ask prices and what the current inventory level is. The order presence is used to calculate average unconditional presence on a stock-day-account basis, as well as conditional on the current inventory level. We exclude stock-days on which a market participant has no trading activity or on which a market participant is posting two-way quotes less than 5% of the time. We regress: $y_{itg} = \beta_0 + \beta_1 NonColo_{itg} + \beta_2 10GColo_{itg} + \alpha_0 Event_t SlowColo_{itg} + \alpha_1 Event_t NonColo_{itg} + \alpha_2 Event_t 10GColo_{itg} + \sum \theta_a Stock_{it}^a + \epsilon_{itg}$, where the event dummy, trader type dummies, and stock fixed effects. The dependent variable takes on one of six values, one in each column. (i) Bid or Ask Presence is the fraction of snapshots during which an account has a quote at either the best bid or the best ask price. (ii) Bid and Ask Presence is the fraction of snapshots during which an account has a quote at both the best bid and the best ask price. (iii) Ask Presence Long-Short measures the fraction of snapshots an account has shares available at the best ask price conditional on having a large positive inventory minus the fraction when having a large negative inventory. (iv) Bid Presence Short-Long measures the fraction of snapshots an account has shares available at the best ask price conditional on having a large positive inventory minus the fraction when having a large negative inventory. (v) Long bin Ask-Bid measures the fraction of snapshots an account has share available at the best ask conditional on having a large positive inventory minus the fraction of time available at the best bid conditional on having a large positive inventory. (vi) Short bin Bid-Ask measures the fraction of snapshots an account has share available at the best ask conditional on having a large positive inventory minus the fraction of time available at the best bid conditional on having a large positive inventory. Standard errors used for statistical tests are estimated using the Newey and West (1987, 1994) HAC covariance matrix.

		Bid or Ask Presence	Bid and Ask Presence	Ask Pres. Long- Short	Bid Pres. Short- Long	Long pos. Ask-Bid	Short pos. Bid-Ask
Before	SlowColo (β_0)	0.70	0.37	0.16	0.09	0.10	0.14
	t(β_0)	(116.42)	(61.48)	(1.90)	(2.33)	(1.97)	(2.27)
	10GColo ($\beta_0 + \beta_2$)	0.79	0.36	0.21	0.22	0.21	0.23
	t(β_2)	(12.42)	(-1.99)	(1.42)	(3.62)	(3.05)	(2.20)
Before	NonColo ($\beta_0 + \beta_1$)	0.60	0.13	-0.12	-0.03	-0.11	-0.04
	t(β_1)	(-11.90)	(-39.72)	(-5.80)	(-3.14)	(-5.62)	(-4.18)
	After	SlowColo ($\beta_0 + \alpha_0$)	0.64	0.26	0.05	0.07	0.04
	t(α_0)	(-6.83)	(-9.18)	(-3.13)	(-0.83)	(-2.66)	(-2.17)
After	10GColo ($\beta_0 + \beta_2 + \alpha_2$)	0.79	0.34	0.11	0.08	0.11	0.07
	t(α_2)	(-0.56)	(-3.41)	(-2.59)	(-4.45)	(-2.70)	(-3.52)
	NonColo ($\beta_0 + \beta_1 + \alpha_1$)	0.60	0.14	-0.09	-0.08	-0.10	-0.07
	t(α_1)	(0.87)	(2.98)	(0.67)	(-1.30)	(0.20)	(-0.61)
	Observations	3141	3141	703	703	703	703
	AdjR2	0.38	0.52	0.14	0.12	0.16	0.12

Appendix Table A1: Trader Characteristics

The table reports trading and quoting activity by colocation status. All messages entered during the continuous trading of OMXS30 stocks during Oct.1 - Oct. 26, 2012, are considered. Trader grouping with respect to colocation type is based on the status on Dec. 3, 2012. Panel A compares colocated (*Colo*) with non-colocated (*NonColo*) traders. Panel B compares different segments of Colocated traders, as defined by different colocation types. All statistics are based on sums across stocks and trading days. *Number of accounts* is the number of trading entities considered in each trader group. The following statistics are reported as the share of total volume: *Trading volume* (the SEK volume), *Trades* (the number of executions), *Submissions*, *Cancellations*, and *Modifications* (the number of limit orders posted, cancelled, and modified), *Active trades* and *Passive trades* (the number of executions where the trader participates through a market order and limit order, respectively). The following three ratios are calculated within each group: *q/t* is the sum of all submissions, cancellations, and modifications, divided by the number of executions; *Liq. supply* is the fraction of *Trading volume (Trades)* where the trader participates through a limit order. Panel A uses the Hagströmer and Nordén (2013) HFT definition. Panel B uses the Kirilenko et al. (2011) and Baron et al (2012) HFT definition.

	Number of accounts	Share of total volume							Ratios within group		
		Trading volume	Trades	Submissions	Cancellations	Modifications	Active trades	Passive trades	q/t	Liq. supply (volume)	Liq. supply (trades)
Panel A: Hagströmer and Nordén (2013) HFT Definition											
NonColo & HFT	40	7.3%	5.6%	3.9%	4.0%	2.0%	4.4%	6.8%	14.4	69.6%	60.8%
Colo & HFT	24	25.8%	28.3%	45.4%	40.4%	0.0%	27.2%	29.4%	30.8	47.7%	51.9%
NonColo & NonHFT	72	51.3%	52.0%	13.4%	10.7%	0.0%	52.3%	51.7%	5.0	52.1%	49.7%
Colo & NonHFT	29	15.7%	14.1%	37.3%	44.9%	98.0%	16.1%	12.1%	60.7	37.7%	42.7%
Panel B: Kirilenko et al. (2011) HFT Definition											
HFT*	9	37.7%	39.0%	43.3%	38.1%	0.3%	35.4%	42.6%	21.4	53.0%	54.6%
NonColo & NonHFT	110	50.9%	49.8%	15.3%	12.8%	2.0%	51.6%	48.1%	6.0	50.8%	48.3%
Colo & NonHFT	46	11.4%	11.2%	41.5%	49.1%	97.8%	13.0%	9.3%	87.6	36.5%	41.7%

*Due to the small number of firms in this HFT category, we are unable to disclose their distribution across NonColo and Colo accounts. This is to comply with the NASDAQ OMX policies on participant confidentiality.