INTERNATIONAL EVIDENCE ON ALGORITHMIC TRADING

Ekkehart Boehmer
Kingsley Fong
Julie Wu*

September, 2015

Abstract

We study the effect of algorithmic trading (AT) intensity on equity market liquidity, short-term volatility, and informational efficiency between 2001 and 2011 in 42 equity markets around the world. On average, AT improves liquidity and informational efficiency but increases volatility. We can attribute the AT-related increase in volatility neither to more "good" volatility that would arise from faster price discovery nor to algorithmic traders' inclination to enter the market when volatility is high. On the contrary, these volatility-seeking traders are associated with declines in market quality. Our results are surprisingly consistent across markets and thus across a wide range of AT practices. But results vary in the cross-section of stocks. In contrast to the average effect, greater AT intensity reduces liquidity and worsens the volatility increase in the smallest tercile of stocks. Finally, AT becomes less beneficial when market making is difficult.

* Boehmer is from Lee Kong Chian School of Business, Singapore Management University, Singapore (ekkehart.boehmer@gmail.com). Fong is from the Australian School of Business, UNSW, Sydney, NSW 2052 (k.fong@unsw.edu.au). Wu is from the Terry College of Business, University of Georgia, Athens, GA 30605 (juliewu@uga.edu). We thank Joel Hasbrouck, Terry Hendershott, Elvira Sojli, and seminar participants at the 2012 Institute Louis Bachelier Market Microstructure Conference, the 2012 ATC/ECB Workshop on High Frequency Trading, the Bank of England, the 2013 AFA meeting, 2015 Workshop on the Mathematics of High Frequency Financial Markets, Bocconi University, EDHEC Business School, EDHEC Risk Institute London, ESSEC, HKU, HKUST, HSE Moscow, Humboldt Universitaet Berlin, Imperial College, the Monetary Authority of Singapore, National University of Singapore, NTU, SAIF at Shanghai Jiao Tong University, SMU (Dallas), SMU (Singapore), Tinbergen Institute Amsterdam, Universitaet St. Gallen, University of Houston, TCU, and Warwick Business School for their helpful comments.

Abstract

We study the effect of algorithmic trading (AT) intensity on equity market liquidity, short-term volatility, and informational efficiency between 2001 and 2011 in 42 equity markets around the world. On average, AT improves liquidity and informational efficiency but increases volatility. We can attribute the AT-related increase in volatility neither to more "good" volatility that would arise from faster price discovery nor to algorithmic traders' inclination to enter the market when volatility is high. On the contrary, these volatility-seeking traders are associated with declines in market quality. Our results are surprisingly consistent across markets and thus across a wide range of AT practices. But results vary in the cross-section of stocks. In contrast to the average effect, greater AT intensity reduces liquidity and worsens the volatility increase in the smallest tercile of stocks. Finally, AT becomes less beneficial when market making is difficult.

1. Introduction

By most accounts, high frequency trading (HFT) represents most of the trading volume in today's markets. HFT refers to the activity of algorithms that emit orders or order cancellations, reacting within milliseconds to market updates or new information. Primarily because of their overall importance in terms of trading volume, but also because HFT strategies are neither transparent nor well understood, there is substantial public policy interest in the effects that HFT has on other market participants, trading strategies, and the quality of markets. Security-market regulators around the world actively debate whether and, if so, how HFT should be regulated, and place increasing scrutiny on algorithmic and high-frequency order submission strategies and their consequences. Despite this debate and a recent flurry of theoretical and empirical work in this area, many questions remain unanswered.

In this paper, we take a basic but comprehensive approach that contributes new evidence to this debate. We follow Hendershott, Jones, and Menkveld (2011) and construct proxies for algorithmic trading (AT), a precondition for HFT, from the intensity of order-related message traffic. We use eleven years of intraday data on security-level quotes and trades in 42 markets around the world, on average covering more than 21,507 firms per year. This new and comprehensive sample allows us to exploit variation in algorithmic trading intensity in the cross-section of stocks and in the cross-section of markets.[1]

We have several objectives. First, we describe the relationship between algorithmic trading and market quality, measured in terms of liquidity, informational efficiency, and short-term volatility with a large international sample. While some studies of HFT have looked outside the U.S. (e.g., Hendershott and Riordan, 2013; Menkveld, 2013), they are based on relatively small samples. Even the most comprehensive study thus far, Hendershott, Jones, and Menkveld (2011), does not use data beyond 2006 and its main analysis is based on a 2003 change in trading protocol. As these dates arguably precede the

[1] Strictly speaking, high-frequency and low-latency trading refer to a subset of algorithmic trading involving reacting to market changes within milliseconds. HFT/LLT likely accounts for most algorithmic message traffic and volume. Because our analysis of AT has implications for HFT/LLT, we often use the terms "algorithmic trading," "high frequency trading," and "low latency trading" interchangeably.

steep growth of HFT during the previous decade exploring the subsequent relationship between algorithmic trading and market quality is important.

Second, we exploit the presence of several separate cross-sections of firms. We investigate whether features that are known to affect order submission strategies—such as market cap, share price, and idiosyncratic volatility—impact the effects of AT on market quality.

Third, existing evidence suggests that HFT provides liquidity to other traders and that fast traders act as informal market makers (e.g. Brogaard et al, 2014;). However, in contrast to exchange-regulated market makers, informal market makers are not subject to affirmative obligations, such as requirements for continuous liquidity provision on both sides of the market. Therefore, it is likely that the liquidity provided by informal market makers varies over time, especially when market making becomes difficult. Anand and Venkataraman (2015) show that voluntary market makers are less likely to provide liquidity than exchange-designated ones under unfavorable market conditions. Kirilenko et al. (2014) provide evidence that high frequency traders ceased to supply liquidity during the "Flash Crash" of 2010. In this paper, we complement these studies by providing evidence on changes in the liquidity provision by algorithmic traders.

We find that greater AT intensity is, on average, associated with more liquidity, faster price discovery, and greater volatility. These results control for share price, trading volume, market capitalization, and volatility (where appropriate), and they are remarkably consistent across different markets. They are robust to using other econometric models for estimation and to different measures of market quality.

Importantly, we causally link AT to market quality by using the formal introduction of co-location events as an instrument for the intensity of AT. Co-location allows fast traders to minimize data turnaround time by physically locating their computer hardware next to the exchange's hardware. These events, which are essential for facilitating AT, represent exogenous shocks (to AT) that do not directly affect market quality. We develop instruments based on co-location events and use them to assess the effect of exogenous changes in AT. We find results consistent with our panel results. Given these findings, we argue that AT causally affects market quality—more AT improves liquidity and efficiency, but it increases volatility.

Our result that more intensive AT leads to greater volatility is consistent with the model by Rosu (2015), who find that volatility increases when more high frequency traders enter the market or when their information becomes more precise. Volatility is important to traders and issuers of equity securities and can have adverse effects on market quality through several channels. Limit orders provide liquidity to the market and represent options to trade for other market participants. Greater volatility makes this option more expensive and thus makes liquidity provision more costly. Using a different model, Ait-Sahalia and Saglam (2013) arrive at a similar result: higher volatility leads fast traders to reduce their liquidity provision. Because issuers prefer more liquid markets, the potential costs of lower liquidity could even extend to equity issuers (Stulz, Vagias, and van Dijk, 2014). Greater volatility also increases price uncertainty for traders, making trading more costly for risk-averse market participants (Hasbrouck, 2015).

We see at least two reasons why elevated volatility could be desirable. First, the more efficient markets are, the faster prices change in response to new information, and the higher is price volatility. It is thus conceivable that the greater efficiency that is associated with more AT also produces higher desirable volatility. To address this issue empirically, we hold constant each stock's level of informational efficiency, and still find that AT increases volatility. Therefore, it is unlikely that the AT-induced change in volatility is due solely the "good" volatility associated with faster price discovery.

Second, volatility could increase when AT is more intense because some fast traders may prefer high-volatility environments. It is conceivable that these fast traders enter the market when volatility is high as a by-product of market-making strategies. In this case, high volatility could be desirable because it attracts additional liquidity that would be absent otherwise. We analyze this possibility empirically by relating AT's effect on volatility to its contemporaneous effect on liquidity. We find that on days when AT is associated with high volatility, AT also induces lower liquidity. This link between volatility and liquidity suggests that either the same traders who generate higher volatility also cause lower liquidity or that the traders attracted by high volatility take liquidity rather than provide liquidity. Either way, the AT activity that takes place during high-volatility episodes does not improve liquidity and is therefore not desirable.

We next assess the cross-sectional differences in AT's effect on market quality. While the average effect of AT on market quality is positive, we observe substantial variation in AT intensity. Thus understanding the cross-sectional determinants of the benefits and costs of greater AT intensity is important. Specifically, stocks that are larger in terms of market capitalization, have higher share prices, or have low volatility are also typically easier to trade. Providing liquidity in these stocks is easier, and the trading intensity is high enough to allow for significant algorithmic activity. Implementing high-frequency market-making strategies in particular is likely easier in these stocks than in small, low-priced, or high-volatility stocks. To analyze these issues, we divide stocks into terciles based on market cap, price, and volatility within each market day, and allow the effects of AT to differ across these characteristics.

We find that much of the benefits of high AT intensity accrue to large, high-priced, and low-volatility stock. When AT increases, liquidity actually declines for the smallest terciles of stocks, remaining unchanged for high-volatility and low-price stocks. The main costs associated with AT, in terms of elevated volatility when AT intensity is high, are significantly higher in stocks that are small, low-priced, or high volatility. Establishing these cross-sectional differences in the effect of AT is important, because these differences imply that optimal regulation may need to impose different requirement on different categories of stocks.

Our third objective is to look more closely at the market-making strategies associated with algorithmic traders. We can neither observe actual strategies nor identify specific traders that might employ them. We only observe the stock-level aggregate effect of AT and its variation over time in a particular market. We exploit this advantage of our data by designing a firm-level time-series proxy for days when market making strategies are likely to be more costly to execute. Then we examine whether the effects of AT are different on those difficult market-making days compared to normal days. We find that when market making is more costly, AT provides less liquidity and is more informed, thereby enhancing efficiency. Moreover, volatility increases more on those days than on easy market-making days. Our evidence is consistent with the public policy perspective that competition among traders may endogenously create sufficient liquidity. If, however, regulators see the need to guarantee the availability of liquidity at all times, they may have to specify a binding set of affirmative obligations, perhaps in exchange for a monetary reward.

Overall, our results show that while algorithmic trading often improves liquidity, this beneficial effect is smaller when market making is either difficult or applied to low-priced or high-volatility stocks. AT's effect on market quality reverses for small cap stocks, where more AT is associated with a decrease in liquidity. Across our tests, AT usually improves efficiency. The main costs associated with AT appear to be elevated levels of volatility. Although this effect prevails even for large market cap, high price, or low volatility stocks, it is more pronounced in smaller, low-price, or high-volatility stocks.

Our paper is organized as follows. We review the theoretical and empirical literature in Section 2. In Section 3, we discuss our data and define the key variables we use. We discuss our empirical design in Section 4 and present our results in Section 5. The final section concludes.

2. Literature on algorithmic and high frequency trading

Algorithmic trading (AT) is not a recent phenomenon, but its intensity, volume, and especially the speed at which it is conducted has experienced precipitous growth over the past decade. The availability of an efficient platform for implementing AT strategies is a precondition for high frequency trading (HFT), which one can view as the subset of AT that can respond to changes in news flow and market conditions within milliseconds. In this paper, while recognizing the difference between AT and HFT, we will use the two terms interchangeably to refer to algorithmic traders with access to fast trading technology. Several theoretical and empirical models analyze HFT's effects on market quality measures, including execution costs, volatility, and informational efficiency. Although it is a young literature, analyses of HFT and algorithmic trading reveal an interesting dichotomy. While theoretical models mostly predict negative (or mixed) consequences of having fast traders in the market, the average effects estimated in empirical results tend to be positive.

A. Theory

Cartea and Penalva (2012) design a model with liquidity traders, market makers, and HFT. They find that HFT increases not only overall trading volume but also volatility and the price impact of liquidity traders. Market-makers come out even—they lose market share (and thus revenues) for liquidity provision to the HF traders but are compensated with higher rewards for their remaining liquidity supply. The cost for the higher rewards to market making, and for the greater revenues to HF traders, are all borne by the liquidity traders. Ait-Sahalia and Sagdam (2013) find that greater volatility reduces liquidity supplied by fast traders. In Jarrow and Protter's (2012) model, HF traders also observe order-flow information faster than other traders. They show that when demand curves are downward sloping, HF traders' activity affects price and creates a temporary mispricing that HF traders can profitably exploit. In this case, the detrimental effect lies in less efficient pricing in addition to a potentially undesirable transfer from slow to fast traders.[2]

Several recent papers address the welfare issues associated with HFT. Pagnotta and Philippon (2012) examine how exchanges determine their investment in fast-trading technology and how traders make their order submission decisions. Allowing market structure and speed to arise endogenously, they show that outcomes are generally inefficient relative to the efficient outcome where all venues break even. Depending on the market structure, in equilibrium participation is too low and, in some cases, trading speed is suboptimal.

[2] A wealth transfer similar to those in these HFT studies arises in an earlier model by Brunnermeier and Pederson (2005), who allow traders to follow order anticipation strategies ("predatory trading" in their model). This strategy requires the ability to predict order flow in real time at high frequency, and it is easily implemented as a trading algorithm. Order anticipators attempt to predict large uninformed orders and then trade ahead of these orders, in the same direction. This trading activity increases the costs for the large liquidity trader, who will end up trading at relatively inferior prices, perhaps even with the order anticipator. Brunnermeier and Pederson show that order anticipation trading leads to price overshooting and withdraws liquidity from the market when liquidity is most needed (by the large trader). As a result, a wealth transfer occurs from the large liquidity trader to the order anticipator. Moreover, Brunnermeier and Pederson show that the low-liquidity event can trigger systemic liquidity shocks for other traders and markets, thereby multiplying the negative consequences that the order anticipator imposes on the market.

Hoffman (2014) extends Foucault's (1999) limit order market and allows algorithmic (fast) and human (slow) traders to compete. Traders can endogenously choose to invest in fast-trading technology. Being fast means that traders can react to news first, thereby reducing the risk that their limit orders will be picked off after adverse price moves. In this model, the welfare effect of introducing algorithmic traders depends on the level of market efficiency. While investments in fast technology improve welfare when efficiency is sufficiently high, they reduce welfare when efficiency is too low.

Biais, Foucault, and Moinas (2014) show that HF traders can generate gains not only from trade but also from adverse selection, due to their faster access to information. However, a social planner would only consider gains from trade; as a result, HF traders overinvest in technology, a strategy leading to socially undesirable outcomes. Overall, existing theoretical models agree that HFT has undesirable consequences for liquidity traders, informational efficiency, and volatility, and these effects may well result in lower social welfare. Finally, Jovanovic and Menkveld (2015) presents a model where middlemen intermediate between fast limit order and slow market order traders. Depending on parameter values, their entry may increase or decrease trading volume, and also has a mixed effect on welfare.

B. Empirical studies

The recent spread of HFT has spurred a number of empirical studies that examine its consequences. Using detailed account-level data, Kirilenko et al. (2014) look at E-mini S&P 500 futures trading around the flash crash of May 6, 2010. They identify actual HFT using account identifiers and assess these traders' role in the market decline and subsequent recovery. They argue that although the HF traders were passive and did not cause the downturn, they did not provide the liquidity for accelerating recovery either. Baron, Brogaard, and Kirilenko (2014) use similar data to show that revenues in HFT in S&P 500 e-minis during 2010-2012 are concentrated in a small number of HFT firms via liquidity taking and higher speed. To date, however, no account-level data are available for equities.

Hendershott and Riordan (2013) and Boehmer and Shankar (2014) use order-level identifiers for orders that originate from algo traders. Likewise, Menkveld's (2013) sample relies on broker identities to infer the trades by a single HFT in the European market. Although these samples allow inferences about algos and HFT, respectively, they are limited to relatively narrow samples.[3]

Hasbrouck and Saar (2013) infer HFT from the speed with which traders react to market events. Most remaining studies, including this paper, use some variations of message-to-trade ratios as proxies for AT. Messages refer to trades, order arrivals, or order cancellations. Because many algo-trading strategies involve frequent cancel-and-replace order traffic, the proportion of traffic that leads to a trade is typically much smaller for algo traders than for non-algo traders. Message-to-trade ratios are well accepted as AT proxies in the trading industry. They allow researchers to use the full panel of stock-days for which standard intraday trade and quote data are available.

Some studies exploit broader categories, using an identifier for a group of HF traders (e.g. Brogaard et al 2014; Brogaard et al, 2015; Carrion, 2013). These studies use a 2008-2009 Nasdaq sample that summarizes the aggregate order flow generated by 26 HFT firms. These firms capture about three quarters of trading volume in the sample stocks. The main advantage is that actual HFT can be observed for a random sample of 120 stocks. Potential drawbacks include the selection of HFT firms, which have been picked by the exchange that provided the data and, presumably, have been willing to have their order flows disclosed to academics and, implicitly, regulators. Because high frequency strategies are typically considered sensitive both from a legal and a competitive perspective, this selection process could conceivably result in orders that are more benign and of lower competitive value than a random sample of HFT orders. Other potential issues also complicate drawing inferences from this dataset. First, the sample of 26 HFT firms does not include any of the large proprietary trading desks that presumably are responsible for a sizeable portion of HFT. Second, the 26 sample firms appear to be large trading firms that specialize in HFT, and these firms often operate in multiple countries and on multiple exchanges. Yet we do not know what fraction of their order flow the sample firms submit to other markets, for stocks either included or not included in the Nasdaq sample. Overall, while these data are

[3] For example, Hendershott and Riordan (2013) examine 30 Deutscher Aktien Index stocks on the Deutsche Boerse in January of 2008. Boehmer and Shankar (2014) look at one event in 2010 on National Stock Exchange of India. Menkveld (2013) examines 1 HF firm on Euronext and Chi-X.

appealing because they reveal certain HFT activity, they also have significant shortcomings that complicate drawing inferences.

In summary, the broadest data sets that rely on proxies for identifying AT, a strategy that in principle would allow the strongest inferences, make the least clear distinction between HF algorithmic, and slow trading. At the other extreme, data sets that identify actual HF activity tend to be either small, limited to specific securities or periods, or not necessarily representative for other reasons. Moreover, some of these available data sets are subject to endogeneity concerns, because identifying whether causality goes from market quality to HFT activity or from HFT activity to market quality is generally difficult.

Given these basic data concerns, most studies present a positive picture of AT/HFT on liquidity and price efficiency. Hendershott, Jones, and Menkveld (2011) are among the first to document this relationship. They show that algorithmic trading leads to better liquidity and faster price discovery. They use the 2003 introduction of autoquote at the NYSE as an instrument for establishing causality from algorithmic trading to market quality improvements. Using HFT activity inferred from millisecond-level responses, Hasbrouck and Saar (2013) find improvements in volatility, spreads, and depth when these fast traders are active. Brogaard, Hendershott and Riordan (2014) document that HFT plays an important role in price discovery. Additionally, for a much smaller Deutsche Boerse sample, Hendershott and Riordan (2013) find that algorithmic trading makes prices more informative. For retail traders in the Canadian market, Malinova, Park, and Riordan (2013) show that a decline in HFT reduces liquidity and profits.

On the negative side, Kirilenko et al. (2014) argue that HFT worsened (but did not cause) the May 6, 2010 flash crash. Dichev, Huang, and Zhou (2014) find that trading per se generates excess volatility, suggesting that HFT can lead to undesirable levels of volatility. Hasbrouck and Saar (2009) document the “fleeting” nature of many limit orders in electronic markets, questioning the traditional view that limit orders provide liquidity to the market. This finding raises questions about the quality or usefulness of HFT-provided liquidity that is often short-lived, with availability periods sometimes measured in milliseconds.

Consistent with this concern, Egginton, Van Ness, and Van Ness (2014) show that periods of extremely active quoting behavior are associated with degraded liquidity and elevated volatility. Importantly, they show that such episodes are surprisingly frequent. Yet despite good economic reasons for such quote-bunching to occur as a benign by-product of HF liquidity provision, as Hasbrouck and Saar (2013) argue, that this quote-bunching arises as a consequence of intentional "quote stuffing" is also possible. This practice involves submitting a large volume of messages to disguise trading strategies. Gai, Yao, and Ye (2013) show that quote stuffing has negative effects on trading, arguing that no offsetting social benefits exist.[4] McInish and Upson (2012) examine trading around quote changes. Comparing fast and slow responses, they find that fast traders strategically leave stale orders on the book and that slow traders often interact with these orders at prices that are inferior to those available elsewhere.

Such "structural strategies" (see SEC 2010 for a discussion) exploit wealth transfers among traders and may not have off-setting market-quality implications. Interestingly, Hirschey (2013) finds that the profits of HF traders are most easily explained by their ability to predict other traders' order flow, rather than by arbitrage or market-making activities that provide benefits to markets. Finally, Chaboud et al. (2014) look at HFT in the foreign exchange market and document that the correlation among algorithmic "machine" orders is much higher than that among "human" orders. Similarly, Anand and Venkataraman (2015) find that synchronous withdrawal of liquidity provision by HF firms under difficult market conditions contributes to fragility of liquidity supply on Toronto Stock Exchange. These findings raise questions about the contribution of AT to the transmission of systemic risk.

Overall, we make two observations. First, the generally positive picture of AT emerging from the empirical evidence appears inconsistent with the generally negative expectation arising from theoretical work in this area. Second, the empirical evidence is not in agreement either. While many studies find that algorithmic and HF traders increase liquidity and price discovery, others raise concerns about the quality of liquidity, AT's effect on volatility, and potential wealth transfers from slow to fast traders. We believe

[4] However, recent study by Conrad et al (2015) finds higher quoting activity is associated with better price efficiency.

that these observations demand additional analysis of the broader issues related to algorithmic and HF trading. In this paper, we examine how algorithmic trading is related to market quality and contributes new large-sample, cross-country evidence to the literature.

3. Data

Our main data source is the Thomson Reuters Tick History (TRTH) database, which contains intraday trades and quotes data for many markets around the world. We combine these data with U.S. intraday data from Trades And Quotes (TAQ) database, and merge the result with firm-level data in Datastream and Center for Research and Security Prices (CRSP) database. Our initial sample includes all domestic common stocks covered in the resulting database. Data on buy-side transaction costs come from the Ancerno database complied by Ancerno Ltd. (formerly the Abel/Noser Corporation).

The TRTH database (supplied by the Securities Industry Research Centre of Asia-Pacific, SIRCA) provides access to the data feeds from various stock and derivatives exchanges that are time-stamped to the millisecond and transmitted through Reuters' terminals. TRTH organizes data by the Reuters Instrument Code (RIC). Each RIC is associated with a list of characteristics, such as asset class (e.g., equity), market, currency denomination, the date of the first and the last record, and the ISIN and SEDOL where applicable. The database contains more than 5 million equity and derivatives instruments around the world. A company may have multiple RICs representing common shares, preferred shares, different share classes, or securities in special trading status. To both create a comprehensive sample of RICs for each market and avoid double counting, we focus on one common stock per company, traded in the home country and in local currency. As TRTH has limited coverage of these screening variables, we construct our sample by first merging TRTH with Datastream by identifying matches between RIC and Datastream firm identifiers.

Datastream identifies securities by DSCODE, which uniquely identifies a security-trading venue combination. Each DSCODE is associated with a comprehensive list of static securities information. We retain only the DSCODE in the local market, traded in the local currency and identified as “major security" and "primary quote." These screening criteria lead to one DSCODE per domestic company, each having a unique ISIN. We are interested in the primary trading location, which coincides with the listing exchange in all markets except Germany. For Germany, we use XETRA (the primary trading location) rather than Frankfurt (the primary listing location), because XETRA handles roughly 90% of volume for most stocks. We merge the two data sources as follows: For each exchange, we obtain the ISIN and the history of high, low, and last trade price for each RIC from the TRTH database. We find the corresponding trading venue on Datastream and identify the unadjusted daily price, market capitalization, and the adjustment factor (dilution) for each screened DSCODE. Then we match RIC to DSCODE using ISIN. There may be more than one RIC per DSCODE if a company changes the trading symbol. We validate the match by comparing the Datastream price history to the TRTH price history after adjusting for currency-reporting differences. [5] This procedure produces stocks trading on 42 equity exchanges in 37 countries.[6]

The TRTH data have qualifiers that contain market-specific codes denoting the first trade of the day, auction trades, and irregular trades (such as off-market trades or option exercises). We remove irregular trades before computing intraday variables.

Trading hours differ across exchanges and over time. We determine each exchange's historical trading hour regime by examining the trade frequency across all stocks on the exchange at 5-minute intervals. We identify the opening and closing times of regular trading from spikes and drops in trading activity across all stocks at each exchange. We cross-check our approach against the trading hour regime and the trading mechanism entries listed in Reuter's Speedguide and the Handbook of World Stock, Derivative and Commodity Exchanges.

[5] TRTH prices are historical prices in the original currency. Datastream unadjusted prices are historical prices in the current currency unit, e.g., French stocks prior to 1999 were traded in French franc. We convert Datastream prices to Euro equivalents.
[6] We drop Ireland, where data is available for fewer than 30 stocks prior to 2008. China has three exchanges covered in Datastream (Hong Kong, Shenzhen, and Shanghai); India (Mumbai and National exchanges), Japan (Tokyo and Osaka), and the U.S. (NYSE and Nasdaq) have two; and all other countries have one exchange included in the sample.

Ancerno provides transaction costs analysis for its institutional buy-side clients. Each Ancerno data record includes an anonymized client code, a broker code, the CUSIP and ISIN for the stock, the date of execution, the execution price, and the number of shares executed, as well as whether the execution is a buy or sell. Multiple trades are often associated with a client on a particular stock day. We match the Ancerno data to CRSP and Datastream stock and market data using the date, CUSIP, ISIN, and ticker. To accommodate investors who split orders across brokers, we follow Anand et al. (2012) by aggregating trades into daily orders by client, stock, date, and trade direction.

A. Variables

Our objective is to make inferences about the relationship between algorithmic trading and market quality. We use variables that describe several dimensions of market quality, focusing on liquidity, volatility, and informational efficiency. We describe these variables in this section, along with our proxies for algorithmic trading activity.

Liquidity measures

We compute several standard measures of liquidity and execution costs. For each stock, we have the best quoted spread throughout the trading day. For a given time interval s, the relative quoted spread, standardized by the quote midpoint, is defined as

$$RQS_s = (Ask_s - Bid_s) / ((Ask_s + Bid_s)/2), \quad (1)$$

where Ask_s is the best ask quote and Bid_s is the best bid quote in that time interval. When aggregating over a trading day, we use time-weighted averages of RQS. The wider the spread, the less liquid is the stock.

To take into account possible price improvement potentially arising from hidden liquidity, we compute the relative effective spread, standardized by the quote midpoint at the time of the trade. The RES on the k^{th} trade is defined as

$$RES_k = 2D_k\,(P_k - M_k)/\,M_k, \quad (2)$$

where D_k is an indicator variable that equals +1 if the k^{th} trade is a buy and -1 if the k^{th} trade is a sell, P_k is the price of the k^{th} trade, and M_k is the prevailing midpoint at the time of the k^{th} trade. We follow the standard trade signing approach of Lee and Ready (1991) and use contemporaneous quotes to sign trades and calculate effective spreads (see Bessembinder (2003), for example). *RES* measures the total price impact of a trade.

We further decompose this price impact into a permanent (information-related) price impact, *RPI*, and a transitory component, the relative realized spread, *RRS*. We follow standard practice and base both components on the quote midpoint that prevails five minutes after the trade. *RRS* on the k^{th} trade is defined as

$$RRS_k = 2D_k\,(P_k - M_{k+5})/\,M_k, \tag{3}$$

where $M_{(k+5)}$ is the midpoint five-minutes after the k^{th} trade. RRS can be interpreted as the reward for providing liquidity. The permanent component, *RPI*, is defined as

$$RPI_k = (RES_k - RRS_k) = 2D_k\,(M_{k+5} - M_k)/\,M_k, \tag{4}$$

and measures the change in quote midpoints that is attributable to the information content of the trade. We first compute trade-weighted averages of *RES*, *RRS*, and *RPI* for each stock-day and then equally weighted averages across stocks.

In addition to these intraday liquidity measures, we compute the Amihud (2002) illiquidity ratio, a lower-frequency measure of liquidity, estimated as the absolute value of daily return divided by the contemporaneous dollar trading volume. A larger Amihud ratio indicates that a given volume moves prices by a larger magnitude, thus implying lower liquidity.

As a robustness test, we use execution shortfalls for the buy-side firms that report to Ancerno. In contrast to the trade-and-quote-based measures, the shortfall represents actual execution expenses for an institution's order flow. We define daily execution shortfall, SHORTFALL, as:

$$SHORTFALL = D_k * (XP - RP) / RP, \tag{5}$$

where *XP* is the volume weighted average price across component trades of a daily order and *RP* is the reference price, defined as the opening price on the day of the order.

Volatility

Our primary measure of volatility is the intraday range between the highest and lowest prices of a day standardized by the daily closing price. This measure is useful because it reflects intraday fluctuations in share prices that may trigger or result from algorithmic trading. In addition, we compute four different measures of realized volatility. For lower-frequency measures, we employ the absolute value of daily returns, |Ret|, and daily return squared, Ret^2. We compute analogous measures for daily market-adjusted returns, |MktadjRet|, using market-cap weighted index returns (dilution and dividend adjusted) on all stocks from Datastream as a benchmark. As higher-frequency measures, we use the log of intra-day return variances computed from 10-minute and 30-minute mid-quote returns, Ln(Ret10_Var) and Ln(Ret30_Var).

Informational efficiency

We compute intraday measures of informational efficiency following Boehmer and Kelley (2009). For most of our analysis we rely on intraday measures of quote midpoint autocorrelation. If prices are efficient and follow a random walk, these measures should be close to zero at all horizons . Deviations from zero in either direction indicate partial predictability. We thus use the absolute value of quote midpoint return autocorrelations. We estimate this measure for each stock-day, |AR30|, based on 30-minute return intervals (see Chordia, Roll, and Subrahmanyam, 2005). Results are qualitatively similar for 10, 20, and 60-minute return intervals.

Proxy for AT

Algorithmic activity is generally associated with fast order submissions and cancellations (see Hasbrouck and Saar, 2013). The proxy for AT used by Hendershott, Jones, and Menkveld (2011) reflects this concept. We follow their approach and use AT, the negative of trading volume in USD100 divided by the number of messages, as a proxy for algorithmic trading activity. It represents the negative of the dollar volume associated with each message (defined as either a trade or a quote update). An increase in this measure reflects an increase in algorithmic activity.

Our AT measure is well suited for international and inter-firm comparisons, because it provides a continuous scale of *relative* AT intensity for each market (rather than an on-off switch, or an absolute measure that does not recognize differences across markets). Doing so allows us to use the same measure across a variety of market structures that differ substantially in the degree to which AT is prevalent. Perhaps more importantly, using a relative measure of AT allows the nature of “fast” or “low-latency” trading to differ across markets. For example, some markets impose hurdles to fast quoting. AT will remain more intense in some stocks and some episodes than in others. Moreover, because our proxy represents a relative measure of AT, we can use it to compare the effect of AT across countries even when comparing a market with latency measured in nanoseconds to one where it is measured in seconds. In either market, HF traders gain by being faster than other traders and our relative, continuous measure of AT captures this contrast well.

Our measure of AT differs in an important way from the one used by Hendershott, Jones, and Menkveld (2011), who have access to order-level messages. For our worldwide sample, we have access only to a subset of these messages, observing each exchange’s best quotes and trades, rather than all order-related messages. Conceptually, using just trades and changes in the best quotes should not be a problem. For example, the HFT strategies mentioned in the SEC 2010 concept release involve most activity *at* the BBO, rather than *behind* it. Therefore, the AT activity in our BBO trade data set is highly correlated with AT activity in an order trade dataset.[7]

[7] We formally address the correlation between AT measures based on order level data, and AT based on trades and best quotes. We repeat HJM’s time-series and panel results for the U.S. using order-level data and compare the results to the ones we obtain with our data and our version of the AT measure. The time series in which our orderlevel data and the TAQ data overlap is very similar to the period presented in HJM. Our exercise using only NYSE activity yields qualitatively identical results for HJM’s order-level count and our count of trades and inside quote changes. This result is not surprising because the correlation of these two series, for the average stock, exceeds 0.9. Therefore, we have little reason to expect our AT proxy to deliver substantially different results than the Hendershott, Jones, and Menkveld (2011) proxy.

B. Descriptive statistics

For inclusion in our final analysis, we impose additional data requirements. We exclude stocks that have data for fewer than 21 trading days during the sample period. We then winsorize all variables each day at 0.5% and at 99.5% within each market. To illustrate the breadth of our sample, Table 1 lists the number of stocks for each market. For the average year, our sample includes about 21,507 firms, and we have substantial variation across markets. Over the sample period, the number of listed firms increases, on average, by 50%, or from 473 to 560 for the average market.

Our key analysis variable is the distribution of message traffic (the number of all quote changes and all trades), the main component of our algorithmic trading proxy. For each market, Table 2 lists the median number of messages per day in 2001 and in 2011 along with the change and percentage change over this eleven-year period. We make several important observations. First, message traffic grows over time, often steeply, in all but three markets. Only Greece, Singapore, and Sweden experience declines in average message traffic per stock day. Overall, message traffic grows by 412% across markets, from 30 messages per stock day in 2001 to 121 in 2011. This development is consistent with AT playing an increasingly important role around the world. Figure 1 shows that this growth accelerates exponentially during the second half of the decade. We also observe that most of the message growth comes from quote messages rather than trade messages, further motivating an order-to-trade ratio as a proxy for the unobservable AT intensity.

Although our analysis is cross-sectional, we begin by describing the time-series of our key dependent variables. In Figure 2 we present the average monthly time-series for our measures of liquidity, efficiency, and volatility, respectively. For each market day, we first compute an equally weighted mean across firms, and then average within each market month. In the figure, we plot the monthly time series of averages across markets. The quoted and effective spreads, RQS and RES, in Panel A show similar patterns. For example, RES begins at 250 bp in the beginning of 2001 and declines to 150 bp by the end of 2007. Afterwards, it peaks at the end of 2008, when the financial crises around the world start to unfold. RQS declines from 500 bp in 2001 to 246 bp in 2007, and then increases to 660 bp during the financial crisis before declining to 400 bp in 2011. The difference between RQS and RES primarily reflects the absence of trades during high-spread periods or the presence of traders who execute against non-displayed liquidity inside the quotes. When we decompose RES into its transient (RRS) and permanent (RPI) components, we again find very similar patterns. Both components decrease until mid-2007 and then increase again. We also observe that RRS exceeds RPI in every year by about 50%.

Panel B shows |AR|, the absolute value of quote midpoint autocorrelation of returns measured over 10 and 30 minute intervals. Both measures are almost flat, with a slight decline over the sample period. The volatility measures in Panel C also decline slightly over the sample period, with a large spike towards the end of 2008 and a smaller increase in 2011. Both efficiency and volatility trends provide a clear contrast to the liquidity measures, which have a pronounced "U" shape over time.

4. Methodology

A. Country-specific analysis

We first identify, for each of the 42 markets, the relation between AT intensity and market quality, summarized by measures of liquidity, volatility, and informational efficiency. We document this relationship in panel regressions that control for firm and day fixed effects. These fixed effects prevent us from interpreting systematic patterns in market quality across firms or secular patterns over time as the result of AT.

Separately from our main analysis, we use an instrumental variable (IV) approach to handle potential endogeneity issues. In addition, our cross-country design at least partly allays the concern about endogeneity in the market quality/AT system, a possible concern about single-country studies. The reason is that with a cross-country study the hurdle for the reverse-causality argument is higher: that causality runs the same way in all countries is unlikely, ceteris paribus. Although this argument still does not imply causality, one can learn a lot from taking a detailed look at the actual (not instrumented) relationships, especially when they are consistent and significant across countries.

For our main cross-sectional analysis, we employ the following panel regression within each market:

$$MQ_{it} = \alpha_i + \gamma_t + \beta AT_{i,t-1} + \delta X_{i,t-1} + \varepsilon_{it,} \qquad (7)$$

where the α_i are firm fixed effects, the γ_t are day fixed effects, AT is our proxy for algorithmic trading, and X is a vector of control variables. This vector includes share turnover, inverse price, the log of market value of equity, the lagged dependent variable, and the daily price range standardized by the daily closing price (a proxy for volatility, omitted from the volatility regressions). To ensure that all explanatory variables are predetermined, they are lagged by one period. To make coefficients comparable across countries, we standardize all continuous variables each day in the cross-section. For inference within countries, we use standard errors that are robust to cross-sectional and time-series heteroskedasticity and within-group autocorrelation (Arellano and Bond, 1991). Cross-market inference, our focus in this paper, is based on an equal-weighted means of the 42 market-specific coefficients and simple cross-sectional t-statistics also based on these 42 observations. This approach is conservative in terms of standard errors, because all inference is based only on the 42 market-specific observations.[8]

Because the relation between AT and market quality could differ across firm attributes, we differentiate observations according to cross-sectional firm characteristics including market cap, volatility, and share price. Unless stated otherwise, we determine daily, separately for each market, the lowest and highest tercile of firms based on the most recent 20 trading days. We assign “LOW” and “HIGH” dummies, respectively, to firms in these terciles. We augment our regression model (7) with the two interactions between AT and each dummy. The interaction coefficients capture the potentially different effect of AT on market quality in the LOW or HIGH terciles relative to the middle tercile. The

[8] Our approach is conservative especially in relation to a three-way panel across markets, stocks, and time, pooling all observations. The three-way estimates agree in sign with the ones provided in this paper. However, because of the much larger number of observations, these estimates have substantially lower standard errors. In another robustness check, we use a market-specific Fama-MacBeth model. For each day, we estimate a cross-sectional regression analogous to equation (7) within each market, but without the firm and time effects. For tests within markets, we compute the time-series average of each coefficient and use Newey-West standard errors for inferences. Across-country inference uses cross-sectional t-statistics as in the main analysis. This approach produces qualitatively identical results (but economically larger in magnitude) that are not tabulated. Again, the method tabulated and discussed in the paper is the more conservative approach.

total effect of AT for LOW firms is given by the sum of the coefficient on AT and the coefficient on AT*LOW. We interpret results for the HIGH dummy analogously.

B. Instrumenting AT

We next establish that the relation between AT and market quality is not spurious. We seek an instrument that satisfies the exclusion restriction, i.e., is not causally related to any of our market quality variables. In addition, the instrument should be closely related to AT intensity. As our sample represents a multitude of trading protocols and market structures, finding an instrument that has the same interpretation across markets is important. We rely on the event of "co-location" in each country.[9] "Co-location" refers to locating a trader's computer hardware physically close to a trading center's hardware. Doing so allows the trader's order submission algorithm to interact with the trading center with minimal latency. Brogaard et al. (2015) show that co-location (at NASDAQ OMX Stockholm) allows fast traders to reduce their cost of liquidity provision and thus trade more profitably. To introduce such a program, some markets announce a program or pricing scheme, while others announce that a specific trading firm is now co-located (and typically invite successors).

From these announcements, we identify the first implementation date (rather than use the first announcement date itself) to capture the change in trading that is prompted by the lower co-location-related latency. Co-location introductions mark an event that is fairly homogenous across exchanges, in that the event specifically provides infrastructure for fast traders and signals an exchange's commitment to accommodate such traders. Obtaining event dates from news searches could lead to varying precision across countries. However, to the extent that the resulting errors are random, they should not affect the consistency of the IV estimator, because such random errors would be captured by the regression error. A complete list of all co-location dates is reported in the Appendix.

[9] Other possible instruments include the introduction of direct market access for traders, DMA, or other updates to the trading protocol that imply a structural change in how traders implement AT / HFT strategies.

To compute the actual instrument, we choose two different strategies. First, we use a simple switch variable that is zero before co-location and one afterwards. Second, we use a continuous version that equals zero before, and the number of days since co-location afterwards. This latter variable is motivated by the observation that with more time after the first event, we would expect more widespread use of co-location and thus a stronger relation with AT.

The introduction of co-location happens at different dates across markets. Because co-location applies to all firms within a market simultaneously, we use a between-estimator at the market level. Specifically, we compute the market-value weighted averages for all variables within each market, standardize the resulting time series within each market, and then perform a two-stage generalized instrumental variable estimation. In the first stage, we regress our AT proxy on time-fixed effects and a co-location dummy (adding the remaining explanatory variables to the first stage leaves inferences unchanged but increases standard errors). In the second stage, we estimate

$$MQ_{ct} = \alpha_c + \gamma_t + \beta AT^*_{ct} + \delta X_{ct} + \varepsilon_{ct,} \qquad (8)$$

where the α_c are market fixed effects, the γ_t are day fixed effects, AT* is the vector of predicted values from the first-stage regression, and the other variables are market-specific weighted averages of the control variables described earlier. For inference we use standard errors that are robust to cross-sectional and time-series heteroskedasticity and within-group autocorrelation based on Arellano and Bond (1991). Besides formalizing our approach to address endogeneity, this market-level regression also serves as a robustness check for the main analysis, in which we treat individual markets separately.

5. Regression results

A. Within-market effects of AT

As described in section 4, we conduct a two-dimensional analysis within each market using daily two-way fixed-effects regressions. This approach allows more conservative global inferences based on market-level averages.

Panel A of Table 3 presents a summary of the liquidity-related coefficients that we estimate for each market.[10] For example, the mean coefficient of AT on RES is -0.0097, meaning that a one-standard deviation increase in AT implies almost a 0.01 standard deviation decrease in relative effective spreads. The associated t-statistic is -3.5, using the cross-sectional standard error across the 42 markets. AT is associated with better RES in 69% of the markets. The opposite (a significant positive coefficient) is true only in 26% of the markets. We find consistent results with the other liquidity measures, i.e., an increase in AT is associated with decreases in RQS and the Amihud measure. Finally, we document that more AT is associated with lower information content of trades (RPI) and the transitory price impact (RRS), a measure of the premium earned by liquidity suppliers. These results suggest that, on average, greater AT intensity is associated with improved liquidity but does not increase the information content of trades.[11]

In Panels B, C, and D, respectively, we assess how AT-liquidity relation varies with market cap, share price, or return volatility. For each measure, we contrast the effect for the lowest tercile ("LOW") with the effect of the largest tercile ("HIGH"). Within each market, we first determine the lowest and highest terciles from the moving average (standard deviation) of market cap and share price (returns) over the past 20 trading days and interact the LOW and HIGH dummies with AT, as described earlier. We report the mean coefficient of AT, which now represents the relation between AT and liquidity for the middle tercile of the sort variable, and coefficients for the two interactions. We also report the total effects for LOW and HIGH stocks.

In Panel B, we sort by market cap, with the LOW and HIGH dummies representing firm sizes within each market. A look at Amihud illiquidity reveals that the largest firms have a marginal coefficient of -0.006, significant at the 5% level, suggesting that the largest firms experience a reduction in Amihud

---

[10] Coefficients on firm fixed effects, daily fixed effects, and control variables are estimated as described in equation (7) but not tabulated.

[11] While not tabulated, the corresponding coefficient averages based on the Fama-MacBeth approach yield identical inferences for the coefficient signs. However, that the magnitude of the AT effect is much larger for the FMB models than for the panel estimation that we present in Table 3. For example, the mean effect of AT on RES is -0.035 standard deviations with FMB—more than three times larger than the estimate from the panel regressions. The differences arise from different treatment of time effects. FMB allows slope coefficients to vary across days, while the two-way panel nets out an aggregate time trend. We tabulate the panel regressions because they are econometrically more appropriate and, ex post, present the more conservative results.

illiquidity (i.e., an improvement in liquidity) that is 0.006 standard deviations greater than the reduction of 0.01 standard deviations experienced by firms in the middle terciles. The total effect of AT on large firms is the sum of these coefficients, -0.016, presented in the last row of Panel B. With a t-statistic of -9.6, the total effect is statistically significant at the 1% level. The results for the other liquidity measures indicate that the AT effect in the large terciles is not significantly different from that in the middle terciles, but that AT in the largest firms remains associated with a significant liquidity improvement. Results are quite different for small firms. The marginal effect of AT is positive and significant at the 5% level across all measures, i.e., AT is associated with wider spreads in small firms compared to middle-tercile firms. This finding suggests that in small firms, compared to other firms, more intense AT is associated with lower liquidity. Indeed, the total effect of AT is *positive* for small firms. Greater AT intensity in small firms is associated with higher transaction costs and, therefore, a decline in liquidity.

Panel C presents a similar cross-sectional analysis based on share price. The marginal effects (the coefficients on the interactive terms) tend to be significant for the LOW group but not for the HIGH group, indicating that the LOW tercile is different from the middle one. Although the differences are not as pronounced as those for market cap, AT is associated with better liquidity in the mid- and high-priced categories. AT in low-priced stocks is associated with lower liquidity.

Panel D looks at the relation between AT and liquidity, sorted by each stock’s past 20-day return volatility. Similar to the price sorts, we find better liquidity associated with AT is concentrated in low-volatility stocks. The marginal effect of AT is significantly more positive for high-volatility stocks, implying that when AT increases, they experience a significant lower liquidity benefit.

Table 4 presents the relationship between AT and informational efficiency. On average, we find significantly negative coefficients in Panel A. This finding suggests that more AT is associated with lower |AR|, implying an improvement in informational efficiency for both measures of autocorrelation. Panels B, C, and D show how the relation between AT and efficiency varies in the cross section. Generally, AT is associated with better efficiency in all terciles, independent of the sort variable. One exception is the tercile of small firms, which experience no efficiency improvement when AT increases.

Share price has no statistically significant effect on the AT-efficiency relationship. High volatility firms, however, have significantly better efficiency than low-volatility firms when AT becomes more intense.

Table 5 summarizes the AT coefficients for regressions that explain short-term volatility. As Panel A shows, more intense AT is associated with higher volatility, and the results are uniform across volatility proxies whether we look at intraday realized volatility, daily realized volatility, or the standardized intraday price range. Moreover, the coefficients are consistently positive across most markets. The percentage of positive coefficients ranges from 76% (variance of 10-minute returns) to 86% in most other models. The remaining panels in Table 5 reveal that greater AT intensity is associated with greater volatility in each firm tercile, whether sorted by size, price, or volatility. However, the relationship is significantly stronger for stocks that are small, are low-priced, or have high return volatility to begin with.

The question arises as to whether the positive coefficient of AT reflects an association of AT intensity with "good" volatility. Given that AT is associated with better informational efficiency, it is conceivable that the elevated volatility associated with more AT may reflect faster price adjustments to new information. In such a case, the higher volatility would likely reflect new information, not noise, and could therefore be desirable. Another possibility is that narrower spreads, which are also associated with greater AT, are related to smaller quoted sizes, so that subsequent trades result in trade prints that experience lower trade-by-trade execution costs but result in greater price fluctuations. Such a trade-off between liquidity and volatility could be desirable if the benefit of smaller spreads outweighs the potential costs of elevated volatility. To control for both possibilities, we add lagged |AR30|, our main efficiency measure, and lagged RES, a measure of liquidity and execution costs, to model (7) whenever we estimate AT on volatility. By controlling for efficiency and the cost of arbitrage (i.e., execution costs), we hold constant price efficiency which likely is the main source of "good" volatility, we find that it does not change our inference (coefficients inTable 5 are already based on these augmented regressions). Therefore, attributing the elevated volatility associated with more intense AT to faster reflection of news or to tighter spreads is difficult.

Is higher AT-induced volatility associated with improved liquidity? Our results suggest that more AT is related to better liquidity and greater efficiency but also to greater volatility. If a stock experiences these effects on the same trading day they could conceivably offset one another. For example, high-volatility periods could attract AT, which would then lead to a liquidity improvement. Although we cannot fully disentangle these effects and causal directions, we conduct—in addition to the instrumental variable approach that we describe in Section 5.C—a simple test that sheds additional light on the relationship among AT, volatility, and liquidity. We employ a two-step procedure. In the first step, we estimate a cross-sectional regression within each market day, using liquidity, efficiency, and volatility as dependent variables, and record the AT coefficients. Doing so produces a time series of daily AT coefficients for each market; one set each for liquidity, efficiency, and volatility. These regressions use the same controls as those in Tables 3-5, respectively. In the second step, we compute Spearman rank correlations between liquidity and volatility effects and then between efficiency and volatility effects . Pearson correlations produce identical inferences.

Panel A of Table 6 reports Spearman rank correlations between AT coefficients for liquidity and volatility. All but one of the correlations is positive, and most are significantly so. This means that on days when AT is associated with higher volatility, AT is also contemporaneously related to larger spread. Or, conversely, if high volatility indeed attracts algorithmic traders, these traders demand—rather than supply—liquidity. Therefore, at least in our sample, the costs of high volatility are not contemporaneously offset by greater liquidity, as suggested by Castura, Litzenberger, Gorelick, and Dwivedi (2010).

In contrast to the relationship between liquidity and volatility, Panel B shows that AT-induced high volatility and AT-induced greater efficiency are complements, i.e., days with high efficiency also tend to have high volatility. This result is intuitive, because greater efficiency implies faster incorporation of news into prices, resulting in greater realized volatility. Because our volatility regressions control for the level of efficiency, this observation does not affect our inferences from Panel A. In other words, the greater volatility-related efficiency happens on days when liquidity declines.

Taken together, our results indicate that while AT is associated with better informational efficiency for most stocks, liquidity improvements are limited to firms in the two largest and the two least volatile terciles. In contrast, the negative effects of AT on intraday and realized volatility are significant for all firms independent of the market cap, price, or volatility category. Moreover, the higher volatility cannot be attributed to more “good” volatility, and traders who induce higher volatility do not appear to induce greater liquidity.

B. Difficult market-making days

Market-making strategies constitute a prominent subset of algorithmic and HFT strategies (see Kirilenko et al., 2014; SEC, 2010). It remains unclear, however, whether and, if so, how their prevalence varies over time and across stocks. The algorithmic traders who supply liquidity are not subject to the same affirmative obligations that force “traditional” market makers to provide liquidity at all times. The absence of such obligations explains regulators’ concerns that the liquidity provided by these strategies is less stable over time than that provided by traditional market makers (SEC, 2010). Anand and Venkataraman (2015)find that non-traditional (i.e., high-frequency) market makers scale back liquidity provision while traditional market makers provide more liquidity when market conditions are not favorable. Their results suggest that the propensity for supplying liquidity without affirmative obligations varies over time, depending on the market trading conditions.

To complement this discussion, we examine how the AT’s associations with liquidity, efficiency, and volatility vary under different market conditions. Specifically, we test whether they are different on days when market making is more difficult or more costly and try to gauge the magnitude of this effect. To identify difficult market-making days, we rely on a simple proxy based on daily returns. Market making is easiest when prices do not change and when buyers are as likely to arrive as sellers are. Market making is more difficult when a trading day is one-sided. For example, if buyers are more aggressive than sellers on a particular day, prices are likely to increase, and market makers are likely to build up a short position and end the day with an unusually large loss (either unrealized on a short inventory or realized from covering short positions at the prevailing high prices). As these losses reduce capital, increasing losses make market making more difficult. If the positive return/positive order imbalance day was followed by another day with returns and imbalances in the same direction, market-making strategies would become even more difficult to operate. Market makers, facing losses for a second day in a row, would be even more reluctant to provide liquidity. We identify events where two days experience one-sided trading in the same direction and define the second day in this pair as a "difficult market making day."

Specifically, for each stock, we identify all days when the daily return has the same sign as the previous day's return. In addition, we require that the two-day cumulative return exceed the 20-day historical mean by at least one standard deviation. This criterion eliminates smaller return episodes that likely do not have much effect on liquidity supply. We create a dummy variable, HARD, that is one on the second day of these episodes, because we expect market making to be unusually difficult on that day. To estimate how AT affects market quality on difficult days, we proceed analogously to our previously discussed cross-sectional analyses and expand model (7) by interacting the HARD dummy with AT.

The results in Table 7 show the mean coefficients for liquidity, efficiency, and volatility effects in separate panels. Panel A shows that AT is associated with tighter displayed quotes (RQS) on difficult market-making days but that they do not lead to lower execution costs. Instead, AT is associated with significantly lower liquidity on difficult market-making days (RES and Amihud). The source of the greater execution costs on these days is greater information content, as represented by RPI, implying that AT on difficult market-making days is either more informed or provides a stronger inducement to other informed traders to trade on these days. We also find that the transient spread component, RRS, decreases on difficult market-making days. This finding implies a smaller reward for providing liquidity, consistent with liquidity providers withdrawing (or switching to strategies other than market making) on these days. The greater information content of trades is consistent with the negative interaction coefficients in Panel B. These coefficients imply that AT improves efficiency more on difficult days. If AT is associated with more information on difficult days, then it is reasonable that prices become more informative.

Panel C presents the AT coefficients for volatility on difficult days. The interaction terms have significantly positive coefficients for all measures except the 10- and 30-minute return variances (where the interaction coefficients are not significant). For the other five volatility variables, AT is associated with substantially higher volatility on difficult days than on regular trading days.

The magnitudes of the incremental effects (the coefficients on the interaction terms) tend to be large relative to the effect on regular days. For example, coefficient of RES is about 25% smaller on HARD days than on normal days. For the Amihud ratio, it is about 71% smaller, and the permanent price impact is about three times larger. Similarly, compared to normal trading days, on HARD days the marginal effect of AT on efficiency and volatility is quite large.

Taken together, these results show that the association between AT and market quality are different on days when market making is difficult: AT seems to provide less liquidity and brings more informed order flow that improves efficiency more, but it also increases execution costs and elevates volatility more. These effects are economically large, indicating that AT involves different strategies on difficult days than on other days. These differences imply significant changes in how AT affects market quality. These results are broadly consistent with the conclusions in Anand and Venkataraman (2015), who recommend some trading options (e.g. option to auto-participate) in a hybrid model where AT market makers compete with traditional market makers, especially under unfavorable conditions. Our analysis suggests that such mechanism can be beneficial if it encourages either type of market maker to provide liquidity on HARD days, when presumably it is most needed.

C. Instrumental variable estimation

Co-location facilitates AT without directly affecting market quality. We create a time-varying dummy variable within each market to indicate the availability, if ever, of co-location. We use this dummy as an instrument for AT in a between-markets panel as described in section 4.B. The second-stage results in Table 8 are uniformly consistent with the within-market analysis: AT improves market quality and efficiency, but elevates short-run volatility. Each of the liquidity measures in Panel A declines as AT increases, implying narrower spreads and smaller Amihud price impacts. The efficiency measures in

Panel B also decline, likewise implying greater efficiency. Panel C shows that AT increases each of the volatility measures.[12]

Overall, these estimates mirror the within-market estimates using the original AT variables. Importantly, despite the lower power of the IV approach, most estimates remain significant and suggest that while AT causally improves liquidity and efficiency, it worsens volatility. Moreover, the market-day panel that underlies the estimation in Table 8 is also an important robustness check on the aggregation of firm-day panels that we use in the main analysis. We obtain qualitatively identical results with either approach, thereby documenting the robustness of our estimates.[13]

6. Conclusions

Quite consistently across the 42 markets in our sample, more intense algorithmic trading (AT) is associated with improved liquidity, improved efficiency, and elevated volatility. AT's effect on volatility cannot be attributed to more efficient prices that adjust faster to new information or to the activities of liquidity suppliers seeking out more volatile stocks. We use co-location events, which represent exogenous shocks to AT, as instruments. This analysis suggests that these effects arise because AT causally affects market quality.

Aside from AT's influence on efficiency, its effects are not uniform across stocks or over time. AT has systematically negative effects on the liquidity of small or low-priced stocks, and AT also increases volatility more in those stocks. The effects of AT on market quality are not stable over time. In

[12] Finally, Panel D presents the IV estimates of the effect of AT on the Ancerno measures of actual institutional trading costs. We find a significant decline in the price impact as AT intensity increases. Additionally, in untabulated estimations, we repeat the liquidity regressions in Table 3, using Ancerno as an additional liquidity measure. We do not report these models because none of them produces a significant coefficient on the AT variable. Nonetheless, together with the IV estimation, this finding reinforces our conclusion that AT, on average, enhances liquidity for mid-cap and large-cap stocks. The Ancerno price impact results in Table 8 suggest that this inference also holds for the passive buy-side investors that constitute Ancerno's clientele.

[13] Under certain conditions, using a binary variable as an instrument could invalidate the IV estimation. To handle this problem, we create a continuous instrument that equals zero before colocation, and otherwise the number of days since colocation. This variable definition assumes that AT intensity increases as more time passes since colocation. Using this alternative definition leave all main results unchanged.

particular, AT provides smaller liquidity benefits on days when market making is more difficult. In addition to the main effects of AT on market quality, the nature of time variation in algorithmic traders' propensity to supply liquidity is an important consideration for optimal regulation of AT.

Overall, our results support prior results that attribute liquidity-enhancing and efficiency-enhancing effects to algorithmic and HF trading. We complement and qualify these results with evidence that AT's liquidity provision does not apply to all firms and that it actually declines on days when market making is unusually costly. Market-making strategies are an important subset of the strategies available to algorithmic and high frequency traders. Our results suggest that particularly for the smallest tercile of firms, and on days when market making is costly, algorithmic traders' strategies do not primarily appear to focus on market-making strategies. Equally importantly, we show that AT systematically increases volatility, thereby imposing costs on most market participants.

References

Ait-Sahalia, Y. and M. Saglam, 2013, High frequency traders: Taking advantage of speed, Working paper, Princeton University.

Amihud, Y., 2002, Illiquidity and stock returns: Cross-section and time-series effects, Journal of Financial Markets 5, 31-56.

Anand, A. and K. Venkataraman, 2015, Market conditions, fragility and the economics of market making Journal of Financial Economics, forthcoming.

Anand, A., P. Irvine, A. Puckett and K. Venkataraman, 2012, Persistence in trading cost: An analysis of institutional equity trades, Review of Financial Studies 25, 557-598.

Arellano, M. and S. R. Bond, 1991, Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations, Review of Economic Studies 58, 277–297.

Baron, M., J. Brogaard, and A. Kirilenko, 2014, Risk and return in high frequency trading, Working paper, Cornell University.

Biais, B., T. Foucault, and S. Moinas, 2014, Equilibrium fast trading, Working paper, IDEI Toulouse.

Bessembinder, H., 2003, Issues in assessing trade execution costs, Journal of Financial Markets 6, 233–257.

Boehmer, E., and E. Kelley, 2009, Institutional investors and the informational efficiency of prices, Review of Financial Studies 22, 3563-3594.

Boehmer, E. and R. L. Shankar, 2014, Low latency trading and the comovement of order flow, prices, and market conditions, Working paper, Singapore Management University.

Brogaard, J , A. Carrion, T. Moyaert, R. Riordan, a. Shkilko, and K Sokolov, 2015, High-frequency trading and extreme price movements, Working paper, University of Washington.

Brogaard, J., T. Hendershott, and R. Riordan, 2014, High-frequency trading and price discovery, Review of Financial Studies 27, 2267-2306.

Brogaard, J., B. Hagstroemer, L. Norden, and R. Riordan, 2015, Trading fast and slow: Colocation and market quality, Review of Financial Studies, forthcoming.

Brunnermeier, M. and L. H. Pedersen, 2005, Predatory trading, Journal of Finance 60, 1825-1863.

Carrion, A. 2013, Very fast money: High -f requency trading on the NASDAQ, Journal of Financial Markets16, 680–711

Cartea, A. and J. Penalva, 2012, Where is the value in high frequency trading? Quarterly Journal of Finance 2, 1250014.

Castura, J., R. Litzenberger, R. Gorelick, and Y. Dwivedi, 2010, Market efficiency and microstructure evolution in U.S. equity markets: A high-frequency perspective, Working paper, RGM Advisors, LLC.

Chaboud, A., B. Chiquoine, E. Hjalmarsson, and C. Vega, 2014, Rise of the machines: Algorithmic trading in the foreign exchange market, Journal of Finance 69, 2045-2084.

Chordia, T., R. Roll, and A. Subrahmanyam, 2005, Evidence on the seed of convergence to market efficiency, Journal of Financial Economics 76, 271-292.

Conrad, J., S. Wahal, and J. Xiang, 2015, High frequency quoting, trading, and the efficiency of prices, Journal of Financial Economics, forthcoming.

Dichev, I., K. Huang, and D. Zhou, 2014, The dark side of trading, Journal of Accounting, Auditing & Finance 29, 492-518.

Egginton, J., B. F. Van Ness, and R. A. Van Ness, 2014, Quote stuffing, Working paper, University of Mississippi.

Foucault, T. 1999, Order flow composition and trading costs in a dynamic limit order market, Journal of Financial Markets 2, 99-134.

Hasbrouck, J., 2015, High-frequency quoting: short-term volatility in bids and offers, Working paper, NYU.

Hasbrouck, J. and G. Saar, 2009, Technology and liquidity provision: The blurring of traditional definitions, Journal of Financial Markets 12, 143-172.

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

Hendershott, T., C. Jones, and A. Menkveld, 2011, Does algorithmic trading improve liquidity? Journal of Finance 66.

Hendershott, T. and R. Riordan, 2013, Algorithmic Trading and the Market for Liquidity, Journal of Financial and Quantitative Analysis 48, 1001-1024.

Hoffmann, P., 2014, A dynamic limit order market with fast and slow traders, Journal of Financial Economics 113, 156–169.

Hirschey, N.H., 2013, Do high-frequency traders anticipate buying and selling pressure? Working paper, LBS.

Jarrow, R. and P. Protter, 2012, A dysfunctional role of high frequency trading in electronic markets, International Journal of Theoretical and Applied Finance 15, 1250022 .

Jovanovic, B. and A. Menkveld, 2015, Middlemen in limit-order markets, Working paper, NYU.

Kirilenko, A., A. S. Kyle, M. Samadi, and T. Tuzun, 2014, The Flash Crash: The impact of high frequency trading on an electronic market, Working paper, University of Maryland.

Lee, C. M.C., and M. J. Ready, 1991, Inferring trade direction from intraday data, Journal of Finance 46, 733–746.

Malinova, K., A. Park, A. and R. Riordan, 2013, Do retail traders suffer from high frequency traders? Working paper, University of Toronto.

McInish, T. and J. Upson, 2012, Strategic liquidity supply in a market with fast and slow traders, Working paper, University of Memphis.

Menkveld, A., 2013, High frequency trading and the new-market makers, Journal of Financial Markets 16, 712–740.

Pagnotta, E. and T. Philippon, 2012, Competing on speed, Working paper, NYU.

Rosu, I. 2015, Fast and Slow Informed Trading, Working paper, HEC Paris.

SEC, 2010, Concept release on equity market structure, U.S. Securities and Exchange Commission Release No. 34-61358. http://www.sec.gov/rules/concept/2010/34-61358.pdf

Stulz, R., D. Vagias, and M. van Dijk, 2014, Do firms issue more equity when markets are more liquid? Working paper, Ohio State University.

Ye, M., C. Yao, Chen and J. Gai, 2013. The Externalities of high frequency trading, Working paper, University of Illinois at Urbana-Champaign.

Table 1. Number of stocks listed on sample markets
This table reports the average number of stocks listed on each exchange by year from 2001 to 2011.

Market	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	Change 2001-2011	% Change 2001-2011
Argentina	19	55	64	45	61	66	68	57	65	71	65	46	242%
Athens	287	285	287	284	276	263	250	226	219	186	168	-119	-41%
Australia	714	666	921	990	1105	1343	1415	1215	1337	1254	1220	506	71%
Brussels	121	105	124	117	132	132	139	128	139	130	131	10	8%
Copenhagen	154	143	164	169	166	183	207	201	202	193	177	23	15%
Dt Boerse Xetra	239	264	257	264	314	345	377	368	366	381	383	144	60%
Euronext Amsterdam	152	148	139	136	136	138	129	116	114	108	104	-48	-32%
Euronext Lisbon					42	41	41	43	42	40	35	-7	-17%
Euronext Paris	480	462	456	454	567	635	669	611	615	585	575	95	20%
Helsinki	152	149	143	148	147	148	147	141	138	134	129	-23	-15%
Hong Kong	389	419	493	520	549	580	629	701	765	832	802	413	106%
Istanbul	277	286	287	299	305	314	313	307	314	338	359	82	30%
Jakarta	261	253	277	274	267	280	295	185	304	327	347	86	33%
Johannesburg	270	220	230	224	245	263	299	266	257	256	261	-9	-3%
Korea/ Daehan	311	300	667	670	678	703	713	742	711	754	764	453	146%
Kuala Lumpur	692	728	782	824	857	870	855	811	824	829	813	121	17%
London	838	781	821	892	1006	1334	905	799	735	891	736	-102	-12%
Madrid	120	119	110	110	108	116	125	123	120	117	114	-6	-5%
Mexican	69	61	67	66	72	82	81	77	91	97	87	18	26%
Milan	271	280	265	264	280	289	310	281	274	272	264	-7	-3%
Mumbai	240	425	716	821	1077	1161	999	939	1348	1414	1242	1002	418%
NASDAQ	3606	3237	2948	2879	2822	2797	2743	2622	2466	2359	2237	-1369	-38%
NSE (India)	500	542	598	671	750	880	1034	1107	1188	1329	1308	808	162%
New Zealand		55	54	63	77	82	95	74	84	86	69	14	25%
NYSE	1524	1498	1475	1475	1459	1443	1392	1347	1328	1330	1311	-213	-14%
Osaka	183	195	232	242	277	275	263	250	253	252	246	63	34%
Oslo	174	174	164	172	198	206	236	226	218	213	202	28	16%
Philippines	140	112	135	143	153	184	187	162	195	194	205	65	46%
Santiago		81	69	84	84	90	91	92	100	89	100	19	23%
Sao Paulo						306	400	379	385	368	350	44	14%
Shanghai	580	669	742	803	709	716	734	754	821	843	833	253	44%
Shenzhen	459	466	487	516	459	508	596	693	780	1105	1314	855	186%
Singapore	315	326	368	424	472	514	553	525	552	572	545	230	73%
Stockholm	328	315	314	319	345	384	443	462	468	472	467	139	42%
Swiss Exchange	224	218	228	227	235	235	227	227	221	218	213	-11	-5%
Taiwan	518	592	636	667	661	670	676	704	735	755	765	247	48%
Tel-Aviv	310	283	319	342	417	469	528	495	484	480	453	143	46%
Thailand	314	335	363	407	446	466	467	480	480	486	487	173	55%
Tokyo	1957	1996	2072	2205	2310	2371	2360	2289	2284	2286	2261	304	16%
Toronto	614	619	658	712	780	872	930	936	913	915	942	328	53%
Warsaw	137	135	137	179	209	237	307	323	354	366	369	232	169%
Wiener Borse	45	61	53	50	54	61	71	73	68	56	58	13	29%
Column average	473	451	483	504	520	549	555	537	556	571	560	120	
Column total	17984	18058	19322	20151	21307	23052	23299	22557	23357	23983	23511	5043	

Table 2. Number of quote-change and trade messages per stock-day

This table reports median messages for each exchange. We count all intraday messages that represent trades or changes in the price or size of the best quotes for each stock on all 42 sample markets.

Market	Median number of messages per stock in 2001	Median number of messages per stock in 2011	Change 2001-2011	% Change 2001-2012
Euronext Amsterdam	41	2762	2721	6635%
Athens	179	25	-154	-86%
Australian	14	19	5	36%
Argentina	17	28	11	67%
Thailand	33	101	68	208%
Mumbai	13	53	40	308%
Brussels	11	156	145	1318%
Copenhagen	7	21	14	200%
Dt Boerse Xetra	54	745	691	1279%
Helsinki	23	94	71	307%
Hong Kong	52	199	147	282%
Istanbul	147	957	810	551%
Johannesburg	13	31	18	138%
Jakarta	17	119	102	600%
Kuala lumpur	31	35	4	13%
Korea/ Daehan	264	932	669	253%
London*	21	122	101	481%
Euronext Lisbon**	59	586	528	902%
Madrid	156	904	748	479%
Milan	108	529	422	392%
Mexican	14	448	435	3219%
NASDAQ	101	5071	4970	4921%
NSE (India)	51	702	651	1276%
NYSE	773	25496	24723	3197%
New Zealand ***	16	17	1	6%
Oslo	18	53	35	192%
Osaka	12	30	18	150%
Euronext Paris	27	83	56	207%
Philippines	6	28	22	367%
Swiss Exchange	20	113	93	465%
Sao Paulo ****	28	153	125	446%
Singapore	33	23	-11	-32%
Santiago ***	5	30	25	505%
Shanghai	504	5331	4827	958%
Stockholm	59	55	-4	-6%
Shenzhen	383	2159	1776	464%
Tokyo	85	453	368	431%
Tel-aviv	6	50	44	733%
Toronto	31	568	537	1731%
Taiwan	326	565	238	73%
Wiener Borse	14	276	262	1870%
Warsaw	11	31	20	182%
Median	30	121	102	412%

* Quote messaged computed based on only price changes because quote sizes are not available until recent y

** Data begin in 2005.

*** Data begin in 2002.

**** Data begin in 2006.

Table 3. The relationship between algorithmic trading and liquidity

Our data cover 42 markets from 2001-2011. We first estimate, for each market, a firm-day fixed effects panel regression. We show the mean coefficients across the 42 markets, the associated t-statistic using the cross-market standard deviation, and in Panel A additionally the percentage of positive market-specific coefficients. Liquididty measures include time-weighted quoted spread (RQS), trade-weighted relative effective spread (RES), permanent price impact (RPI), and temporary price impact (RRS), and Amihud. AT is the negative of dollar trading volume ($100) per message. For message counts, we include all inside quote changes and trade messages. Control variables include daily share turnover, intraday price range, 1/price, log market cap, and the first lag of the dependent variable, all measured at t-1. All continuous variables are standardized every day to have a mean of zero and a standard deviation of one within each exchange. In Panels B, C, and D we interact AT with two dummy variables, LOW and HIGH. In Panel B, the dummies indicate the smallest and largest market cap tercile based on a moving average from the past 20 trading days. In Panel C, the dummies indicate the smallest and largest share price tercile based on moving average from the past 20 trading days, and in Panel D the dummies indicate the smallest and largest volatility tercile based on the standard deviation of the 20 most recent daily returns.

	RQS	RES	RPI	RRS	Amihud
Panel A. Aggregate of market-specific, firm fixed panel regressions					
Mean coefficient on AT	-0.0093	-0.0097	-0.0009	-0.0156	-0.0110
Cross-sectional t-stat	-6.69	-3.52	-0.26	-5.37	-6.81
Percent positive	2%	31%	57%	7%	7%
Percent positive and significant	2%	26%	52%	7%	5%
Panel B. The effect of AT for the smallest and largest market cap terciles					
Mean coefficient on AT	-0.0119	-0.0109	-0.0019	-0.0194	-0.0104
Cross-sectional t-stat	-4.88	-2.26	-0.29	-3.58	-4.17
Mean coefficient on AT*LOW	0.0189	0.0467	0.0371	0.0415	0.0673
Cross-sectional t-stat	4.03	5.91	3.76	6.15	6.36
Mean coefficient on AT*HIGH	0.0016	-0.0035	-0.0031	-0.0002	-0.0060
Cross-sectional t-stat	0.90	-0.93	-0.63	-0.06	-1.96
Mean total effect for LOW	0.0070	0.0359	0.0351	0.0221	0.0569
Cross-sectional t-stat	1.28	3.32	2.43	2.32	4.78
Mean total effect for HIGH	-0.0103	-0.0144	-0.0051	-0.0196	-0.0163
Cross-sectional t-stat	-7.57	-5.86	-1.78	-7.35	-9.59
Panel C. The effect of AT for the smallest and largest price terciles					
Mean coefficient on AT	-0.0092	-0.0079	0.0013	-0.0158	-0.0105
Cross-sectional t-stat	-5.37	-2.17	0.31	-4.11	-4.66
Mean coefficient on AT*LOW	0.0063	0.0152	0.0174	0.0086	0.0261
Cross-sectional t-stat	1.92	2.96	2.64	1.71	4.20
Mean coefficient on AT*HIGH	-0.0010	-0.0050	-0.0049	-0.0020	-0.0047
Cross-sectional t-stat	-0.84	-2.26	-1.70	-0.86	-2.31
Mean total effect for LOW	-0.0029	0.0072	0.0187	-0.0071	0.0156
Cross-sectional t-stat	-0.78	1.15	1.97	-1.19	2.27
Mean total effect for HIGH	-0.0102	-0.0129	-0.0036	-0.0178	-0.0152
Cross-sectional t-stat	-6.98	-4.98	-1.17	-6.57	-9.20
Panel D. The effect of AT for the smallest and largest volatility terciles					
Mean coefficient on AT	-0.0121	-0.0144	-0.0044	-0.0195	-0.0126
Cross-sectional t-stat	-8.07	-5.54	-1.40	-6.46	-9.12
Mean coefficient on AT*LOW	0.0029	0.0033	-0.0004	0.0042	0.0033
Cross-sectional t-stat	4.69	3.24	-0.36	3.46	2.21
Mean coefficient on AT*HIGH	0.0097	0.0232	0.0240	0.0162	0.0101
Cross-sectional t-stat	5.78	7.19	7.45	5.21	2.52
Mean total effect for LOW	-0.0092	-0.0112	-0.0048	-0.0153	-0.0093
Cross-sectional t-stat	-6.08	-4.35	-1.48	-5.70	-5.42
Mean total effect for HIGH	-0.0024	0.0088	0.0196	-0.0033	-0.0025
Cross-sectional t-stat	-1.11	1.80	3.57	-0.73	-0.60

Table 4. The relationship between algorithmic trading and informational efficiency

Our data cover 42 markets from 2001-2011. We first estimate, for each market, a firm-day fixed effects panel regression. We show the mean coefficients across the 42 markets, the associated t-statistic using the cross-market standard deviation, and in Panel A additionally the percentage of positive market-specific coefficients. Efficiency measures are daily observations of the absolute value of intraday autocorrelations |AR##|, measured for quote-midpoint returns over 10 and 30 minute periods. AT is the negative of dollar trading volume ($100) per message. For message counts, we include all inside quote changes and trade messages. Control variables include daily share turnover, intraday price range, 1/price, log market cap, and the first lag of the dependent variable, all measured at t-1. All continuous variables are standardized every day to have a mean of zero and a standard deviation of one within each exchange. In Panels B, C, and D we interact AT with two dummy variables, LOW and HIGH. In Panel B, the dummies indicate the smallest and largest market cap tercile based on a moving average from the past 20 trading days. In Panel C, the dummies indicate the smallest and largest share price tercile based on moving average from the past 20 trading days, and in Panel D the dummies indicate the smallest and largest volatility tercile based on the standard deviation of the 20 most recent daily returns.

	|AR10|	|AR30|
Panel A. Aggregate of market-specific, firm-fixed effects panel regressions		
Mean coefficient on AT	-0.0126	-0.0042
Cross-sectional t-stat	-7.23	-4.01
Percent positive	14%	21%
Percent positive and significant	7%	14%
Panel B. The effect of AT for the smallest and largest market cap terciles		
Mean coefficient on AT	-0.0123	-0.0022
Cross-sectional t-stat	-3.48	-0.92
Mean coefficient on AT*LOW	0.0063	0.0057
Cross-sectional t-stat	1.76	2.21
Mean coefficient on AT*HIGH	-0.0009	-0.0022
Cross-sectional t-stat	-0.31	-1.17
Mean total effect for LOW	-0.0060	0.0035
Cross-sectional t-stat	-1.24	0.89
Mean total effect for HIGH	-0.0132	-0.0044
Cross-sectional t-stat	-7.74	-4.72
Panel C. The effect of AT for the smallest and largest price terciles		
Mean coefficient on AT	-0.0100	-0.0020
Cross-sectional t-stat	-3.86	-0.97
Mean coefficient on AT*LOW	0.0010	0.0026
Cross-sectional t-stat	0.34	1.01
Mean coefficient on AT*HIGH	-0.0034	-0.0029
Cross-sectional t-stat	-1.95	-1.67
Mean total effect for LOW	-0.0090	0.0006
Cross-sectional t-stat	-2.15	0.17
Mean total effect for HIGH	-0.0133	-0.0049
Cross-sectional t-stat	-8.21	-5.25
Panel D. The effect of AT for the smallest and largest volatility terciles		
Mean coefficient on AT	-0.0098	-0.0032
Cross-sectional t-stat	-5.20	-2.63
Mean coefficient on AT*LOW	-0.0010	-0.0003
Cross-sectional t-stat	-0.96	-0.24
Mean coefficient on AT*HIGH	-0.0102	-0.0038
Cross-sectional t-stat	-7.70	-4.84
Mean total effect for LOW	-0.0108	-0.0035
Cross-sectional t-stat	-5.49	-2.93
Mean total effect for HIGH	-0.0200	-0.0070
Cross-sectional t-stat	-9.23	-5.04

Table 5. The relationship between algorithmic trading and short-term volatility

Our data cover 42 markets from 2001-2011. We first estimate, for each market, a firm-day fixed effects panel regression. We show the mean coefficients across the 42 markets, the associated t-statistic using the cross-market standard deviation, and in Panel A additionally the percentage of positive market-specific coefficients. Volatility measures include |Ret|, |MktadjRet|, Ret^2, MktadjRet^2, the daily price range standardized by the daily closing price, and Ln(Ret##_Var), the log of the daily averages of the variances of 10-minute and 30-minute quote midpoint returns, respectively. AT is the negative of dollar trading volume ($100) per message. For message counts, we include all inside quote changes and trade messages. Control variables include daily share turnover, 1/price, log market cap, |AR30|, and the first lag of the dependent variable, all measured at t-1. All continuous variables are standardized every day to have a mean of zero and a standard deviation of one within each exchange. In Panels B, C, and D we interact AT with two dummy variables, LOW and HIGH. In Panel B, the dummies indicate the smallest and largest market cap tercile based on a moving average from the past 20 trading days. In Panel C, the dummies indicate the smallest and largest share price tercile based on moving average from the past 20 trading days, and in Panel D the dummies indicate the smallest and largest volatility tercile based on the standard deviation of the 20 most recent daily returns.

	|Ret|	|MktadjRet|	Ret^2	MktadjRet^2	PriceRange	Ln(Ret10_Var)	Ln(Ret30_Var)
Panel A. Aggregate of market-specific, firm-fixed effects panel regressions							
Mean coefficient on AT	0.0270	0.0163	0.0182	0.0135	0.0401	0.0216	0.0295
Cross-sectional t-stat	7.65	5.00	6.67	5.21	8.52	4.25	5.17
Percent positive	86%	83%	86%	86%	83%	76%	81%
Percent positive and significant	81%	79%	81%	76%	83%	71%	79%
Panel B. The effect of AT for the smallest and largest market cap terciles							
Mean coefficient on AT	0.0323	0.0163	0.0201	0.0117	0.0519	0.0320	0.0395
Cross-sectional t-stat	6.49	2.90	4.86	2.52	8.17	4.17	4.60
Mean coefficient on AT*LOW	0.0174	0.0192	0.0131	0.0116	0.0236	0.0232	0.0229
Cross-sectional t-stat	2.38	2.76	1.97	1.76	2.76	2.84	2.50
Mean coefficient on AT*HIGH	-0.0069	-0.0017	-0.0032	0.0008	-0.0154	-0.0121	-0.0115
Cross-sectional t-stat	-1.53	-0.39	-0.93	0.21	-3.02	-1.89	-1.56
Mean total effect for LOW	0.0497	0.0354	0.0332	0.0233	0.0755	0.0552	0.0625
Cross-sectional t-stat	6.14	3.79	4.87	2.94	7.54	5.04	5.53
Mean total effect for HIGH	0.0254	0.0145	0.0170	0.0125	0.0365	0.0199	0.0280
Cross-sectional t-stat	7.05	4.68	6.28	5.05	7.94	3.98	4.92
Panel C. The effect of AT for the smallest and largest price terciles							
Mean coefficient on AT	0.0280	0.0155	0.0194	0.0137	0.0414	0.0215	0.0284
Cross-sectional t-stat	6.90	3.72	6.61	4.61	7.80	3.47	4.25
Mean coefficient on AT*LOW	0.0175	0.0150	0.0098	0.0080	0.0219	0.0214	0.0226
Cross-sectional t-stat	4.64	4.15	3.52	2.86	4.16	4.10	4.07
Mean coefficient on AT*HIGH	-0.0040	-0.0012	-0.0038	-0.0021	-0.0050	-0.0011	0.0000
Cross-sectional t-stat	-1.79	-0.48	-2.48	-1.29	-2.15	-0.29	0.01
Mean total effect for LOW	0.0454	0.0304	0.0292	0.0217	0.0633	0.0429	0.0510
Cross-sectional t-stat	7.34	4.86	6.16	4.62	7.88	4.77	5.37
Mean total effect for HIGH	0.0240	0.0143	0.0156	0.0116	0.0364	0.0204	0.0284
Cross-sectional t-stat	7.42	4.97	6.33	4.92	8.48	4.16	5.10
Panel D. The effect of AT for the smallest and largest volatility terciles							
Mean coefficient on AT	0.0173	0.0086	0.0123	0.0084	0.0287	0.0086	0.0138
Cross-sectional t-stat	4.71	2.45	4.30	3.08	6.16	2.05	2.84
Mean coefficient on AT*LOW	0.0051	0.0004	0.0003	-0.0021	0.0080	0.0185	0.0203
Cross-sectional t-stat	2.34	0.21	0.22	-1.66	2.93	6.05	5.72
Mean coefficient on AT*HIGH	0.0373	0.0376	0.0295	0.0312	0.0405	0.0189	0.0269
Cross-sectional t-stat	8.56	9.60	9.07	10.52	8.11	5.65	7.00
Mean total effect for LOW	0.0224	0.0090	0.0127	0.0062	0.0367	0.0271	0.0341
Cross-sectional t-stat	5.33	2.45	4.37	2.53	6.61	4.30	4.76
Mean total effect for HIGH	0.0546	0.0462	0.0418	0.0396	0.0692	0.0275	0.0407
Cross-sectional t-stat	10.37	8.88	10.11	9.34	10.43	4.93	6.30

Table 6. Correlation between AT coefficients in liquidity and volatility regressions

Our data cover 42 markets from 2001-2011. We first estimate, for each market, a daily time series of cross-sectional coefficients. We regress liquidity, efficiency, and volatility measures on AT and controls. Liquidity measures include time-weighted quoted spread (RQS), trade-weighted relative effective spread (RES), permanent price impact (RPI), temporary price impact (RRS), and Amihud. Efficiency measures are daily observations of the absolute value of intraday autocorrelations |AR##|, measured for quote-midpoint returns over 10 and 30 minute periods. Volatility measures include |Ret|, |MktadjRet|, Ret^2, MktadjRet^2, the daily price range standardized by the daily closing price, and Ln(Ret##_Var), the log of the daily averages of the variances of 10-minute and 30-minute quote midpoint returns, respectively. AT is the negative of dollar trading volume ($100) per message. For message counts, we include all inside quote changes and trade messages. Control variables include daily share turnover, 1/price, log market cap, and the first lag of the dependent variable, all measured at t-1. Regressions where the dependent variable is volatility do not include price range, but add RES and |AR30|. |AR30| is the absolute value of intraday autocorrelations measured for quote-midpoint returns over 30-minute periods. All continuous variables are standardized every day to have a mean of zero and a standard deviation of one within each exchange. This table reports the mean spearman rank correlation between AT coefficients from the regressions of volatility and liquidity (Panel A) and between the AT coefficients from the regressions of volatility and efficiency (Panel B). *, **, *** indicate significance at 10%, 5% and 1%, respectively.

	|ret|	|mktadjRet|	Ret^2	MktadjRet^2	PriceRange	Ln(Ret10_Var)	Ln(Ret30_Var)
Panel A. the cross-sectional correlation between AT's effect on volatility and on liquidity							
RQS	0.01	0.01 *	0.01 *	0.02 **	0.04 ***	0.10 ***	0.08 ***
RES	0.05 ***	0.06 ***	0.05 ***	0.05 ***	0.08 ***	0.14 ***	0.11 ***
Amihud	0.06 ***	0.03 ***	0.05 ***	0.04 ***	-0.02	0.02	0.02
Panel B. The cross-sectional correlation between AT's effects on volatility and efficiency							
|AR10|	-0.06 ***	-0.04 ***	-0.04 ***	-0.03 ***	-0.10 ***	0.04 **	-0.09 ***
|AR30|	-0.02 ***	-0.02 ***	-0.02 ***	-0.01 ***	-0.05 ***	0.01	0.01

Table 7. The effect of algorithmic trading when market making is difficult

Our data cover 42 markets from 2001-2011. We first estimate, for each market, a firm-day fixed effects panel regression. We show the mean coefficients across the 42 markets, and the associated t-statistic using the cross-market standard deviation. Market quality measures include time-weighted quoted spread (RQS), trade-weighted relative effective spread (RES), permanent price impact (RPI), temporary price impact (RRS), and Amihud. Efficiency measures include daily observations of the absolute value of intraday autocorrelations |AR|, measured for quote-midpoint returns over 10 and 30 minute periods. Volatility measures include |Ret|, |MktadjRet|, Ret^2, MktadjRet^2, the daily intraday price range standardized by the daily closing price, and Ln(Ret##_Var), the log of the daily averages of the variances of 10-minute and 30-minute quote midpoint returns, respectively. AT is the negative of dollar trading volume ($100) per message. For message counts, we include all inside quote and trade messages. Control variables include daily share turnover, 1/price, intraday price range, log market cap, and the first lag of the dependent variable, all measured at t-1. Regressions where the dependent variable is volatility do not include price range, but add RES and |AR30|. |AR30| is the absolute value of intraday autocorrelations measured for quote-midpoint returns over 30-minute periods. All continuous variables are standardized every day to have a mean of zero and a standard deviation of one within each exchange. We create a variable HARD that indicates days on which market making is difficult for a particular stock. Using daily returns, the HARD dummy equals one if a daily return has the same sign as the return on the previous day, and the absolute value of the 2-day return is at least one standard deviation larger than the 20-day trailing average of returns for that stock. We interact AT with the HARD dummy in the regression.

Panel A. Difficult market-making days and the effect of AT on liquidity

	RQS	RES	RPI	RRS	Amihud
Mean coefficient on AT	-0.0085	-0.0104	-0.0048	-0.0135	-0.0140
Cross-sectional t-stat	-5.85	-3.78	-1.42	-4.72	-6.92
Mean coefficient on AT*HARD	-0.0033	0.0026	0.0141	-0.0077	0.0100
Cross-sectional t-stat	-3.92	3.22	7.57	-4.87	5.37
Mean total effect on HARD days	-0.0118	-0.0079	0.0093	-0.0212	-0.0039
Cross-sectional t-stat	-8.68	-2.74	2.28	-6.55	-2.27

Panel B. Difficult market-making days and the effect of AT on informational efficiency

	|AR10|	|AR30|
Mean coefficient on AT	-0.0115	-0.0035
Cross-sectional t-stat	-6.38	-3.21
Mean coefficient on AT*HARD	-0.0041	-0.0026
Cross-sectional t-stat	-4.23	-3.42
Mean total effect on HARD days	-0.0156	-0.0061
Cross-sectional t-stat	-8.78	-5.51

Panel C. Difficult market-making days and the effect of AT on volatility

	|ret|	|mktadjRet|	Ret^2	MktadjRet^2	PriceRange	Ln(Ret10_ Var)	Ln(Ret30_ Var)
Mean coefficient on AT	0.0208	0.0083	0.0120	0.0063	0.0368	0.0217	0.0293
Cross-sectional t-stat	5.84	2.68	4.78	2.89	7.79	4.11	4.99
Mean coefficient on AT*HARD	0.0224	0.0292	0.0229	0.0263	0.0112	-0.0007	0.0005
Cross-sectional t-stat	4.54	6.84	5.82	6.85	4.09	-0.33	0.23
Mean total effect on HARD days	0.0432	0.0374	0.0349	0.0326	0.0481	0.0210	0.0298
Cross-sectional t-stat	8.04	7.31	7.70	7.12	9.17	4.24	5.20

Table 8. Instrumental variable estimation of the effect of algorithmic trading

For each day, we aggregate all variables within each market by forming market-value-weighted averages across firms. We estimate a two-way panel across markets and days using a generalized instrumental variable approach. AT, the negative of dollar trading volume (in $100) per message, is a proxy for algorithmic trading. As an instrument for algorithmic trading we use colocation, a market-specific dummy that switches on once that market officially or at least publicly, for the first time, allows or facilitates colocation for trading firms. Market quality measures include time-weighted quoted spread (RQS), trade-weighted relative effective spread (RES), permanent price impact (RPI), temporary price impact (RRS), Amihud, |AR10|, and |AR30|, and volatility measures include |Ret|, |MktadjRet|, Ret^2, MktadjRet^2, the intraday price range standardized by the daily closing price, and Ln(Ret##_Var), the log of the average variances of 10-minute and 30-minute quote midpoint returns, respectively. Execution short falls include average shortfall (mshortO), shared weighted average shortfall (msvwO), and dollar weighted average shortfall (mdvwO), benchmarked on opening price. Control variables include daily share turnover, intraday price range, 1/price, log market cap, and the first lag of the dependent variable, all measured at t-1. Price range is not included in regressions where the dependent variable is volatility. The sample period is from 2005 to 2011 to maintain a balanced panel where all markets are present in the data.

Dependent variable	AT coefficient	t
Panel A. Liquidity		
RQS	-0.0231	-4.06
RES	-0.0454	-7.12
RPI	-0.0104	-1.30
RRS	-0.0972	-11.07
Amihud	-0.0027	-0.32
Panel B. Efficiency		
Dependent		
|AR10|	-0.0409	-4.00
|AR30|	0.0102	0.99
Panel C. Volatility (controlling for lag RES and lag|AR|		
PriceRange	0.0595	9.99
ln(Ret10_Var)	0.0760	15.62
ln(Ret30_Var)	0.0931	16.79
|Ret|	0.0659	9.16
|MktAdjRet|	0.0364	5.11
Ret^2	0.0456	5.93
MktAdjRet ^2	0.0231	2.89
Panel D. Ancerno institutional price impact		
mshortO	-0.013	-1.10
msvwO	-0.024	-1.95
mdvwO	-0.024	-1.95

Figure 1. Messages and AT

We count all intraday messages that represent trades or changes in the price or size of the best quotes for each stock. Panel A reports the time series of messages. Panel B reports the time series of AT computed as the negative of dollar trading volume ($100) per message. Then we compute equally weighted means for each market month, and then compute the mean across markets.

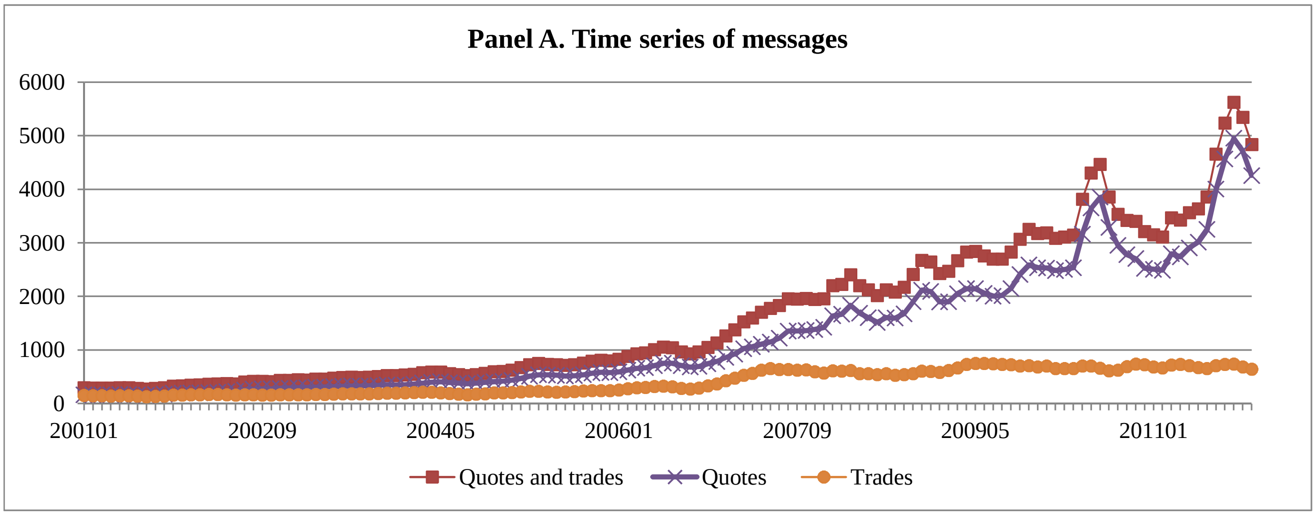

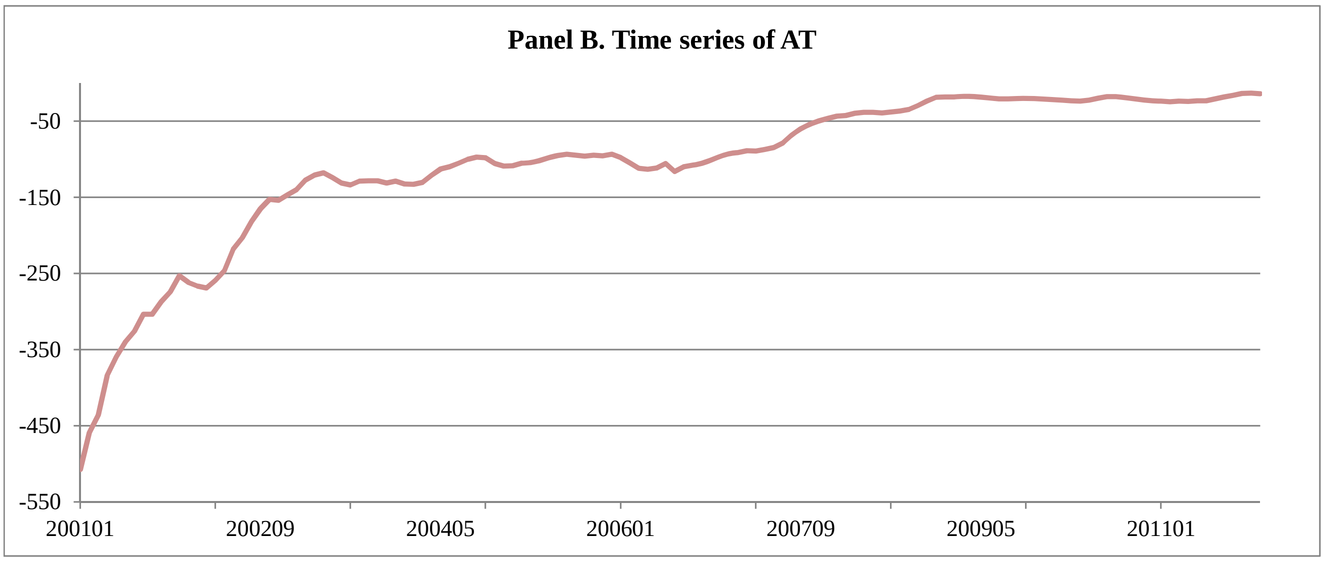

Figure 2. Market quality measures over time

Panel A reports the time series of spreads. RQS are time weighted relative quoted spreads (RQS), RES are relative effective spreads, RRS are 5-minute relative realized spreads, and RPI are 5-minute permanent price impacts. Panel B reports the time series of price efficiency. |AR10| (AR|30|) is the absolute value of the daily average 10-minute (30-minute) quote-midpoint return autocorrelations. We omit overnight returns and periods without price changes. Panel C reports the time series of volatility measures including |Ret|, Ret^2, the daily intraday price range standardized by the daily closing price, and the variances of 30-minute quote midpoint returns (Ret30_Var), respectively. All measures are computed intraday for each stock. Then we compute the mean for each day, and then the mean for each market month. The figures report the mean across markets.

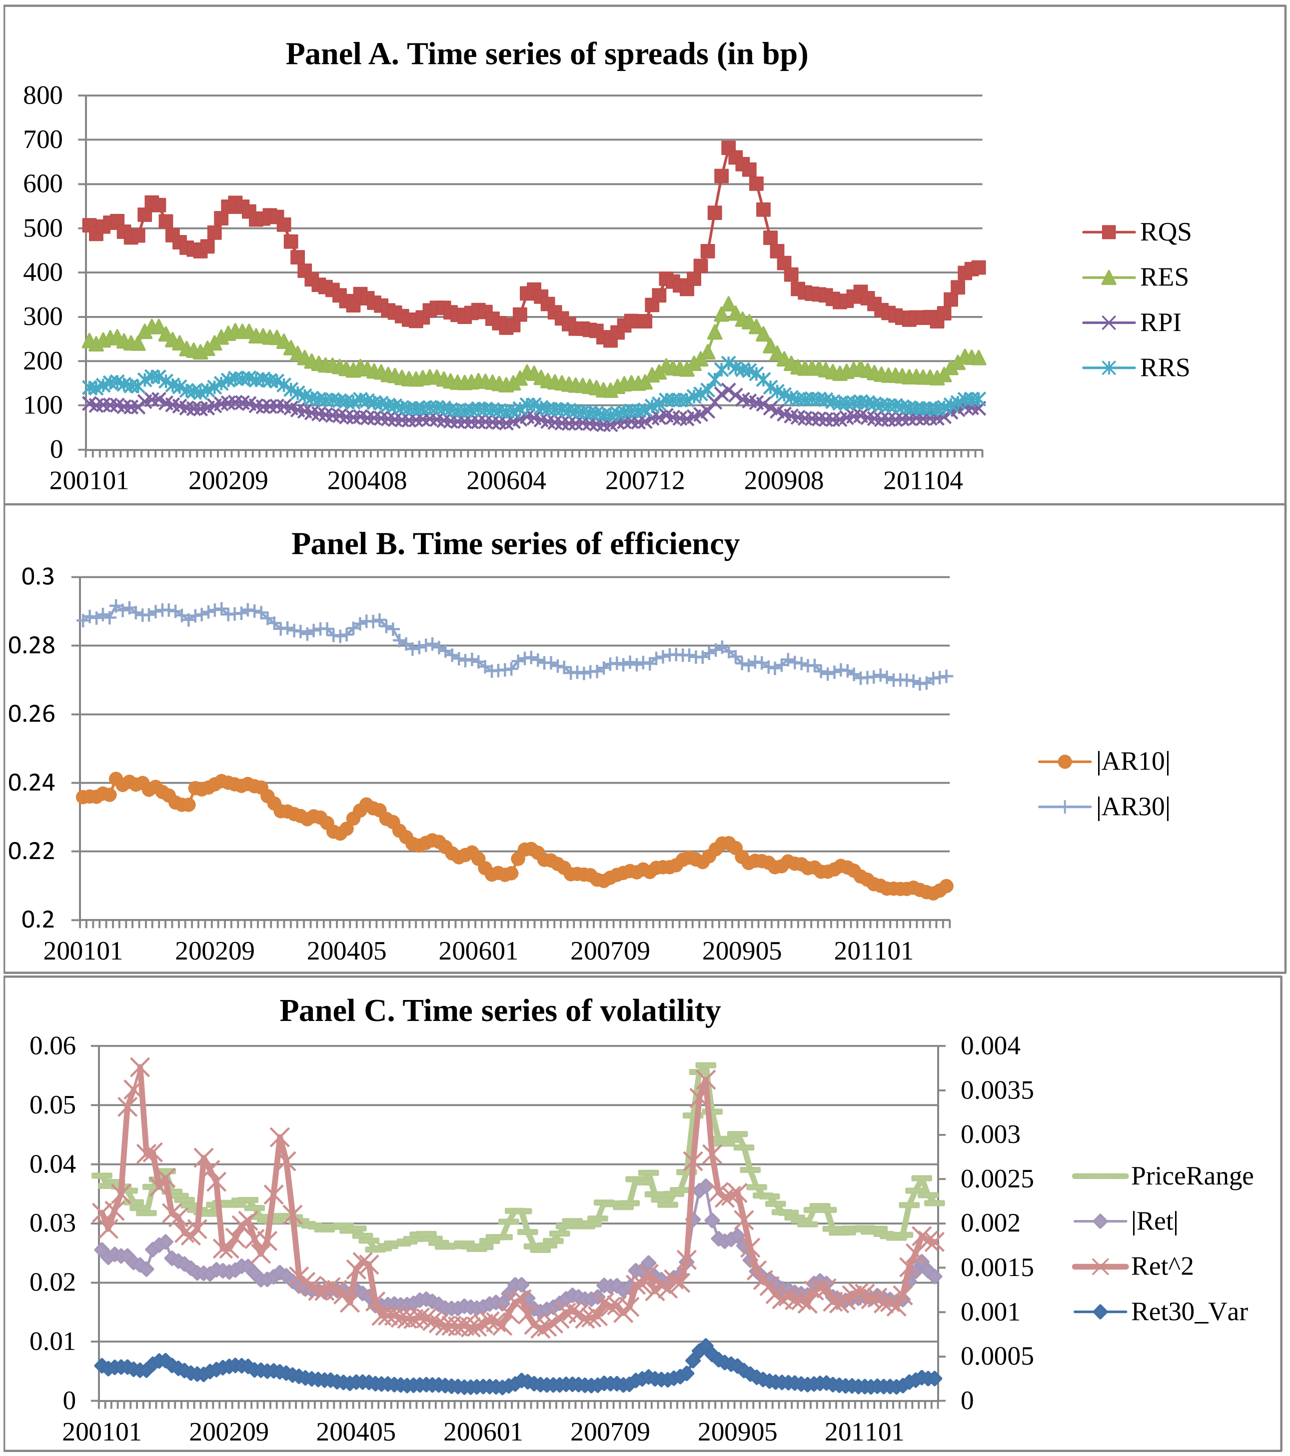

Appendix: Dates of the first co-location implementation across markets

Country	MKT	Earliest colocation date	Link (as of September 2013)
Australia	AX	200811	http://www.asxgroup.com.au/media/PDFs/mr030708_co-location_hosting.pdf
Belgium	BR	200804	https://europeanequities.nyx.com/sites/europeanequities.nyx.com/files/327777.pdf
Brazil	SA	20090629	http://ir.bmfbovespa.com.br/enu/1190/NMCoLocation.pdf
Canada	TO	200811	http://www.investorpoint.com/stock/X%253ACA-TMX+Group+Limited/news/6495414
Denmark	CO	20080625	http://www.ft.com/intl/cms/s/0/b2cde4f0-42ce-11dd-81d0-0000779fd2ac.html#axzz2RRv89HWs
Finland	HE	20080625	http://www.ft.com/intl/cms/s/0/b2cde4f0-42ce-11dd-81d0-0000779fd2ac.html#axzz2RRv89HWs
France	PA	200804	https://europeanequities.nyx.com/sites/europeanequities.nyx.com/files/327777.pdf
Germany	DE	2006Q4	http://deutsche-boerse.com/dbg/dispatch/en/listcontent/gdb_navigation/press/10_Latest_Press_Releases/Content_Files/13_press/August_2006/pm_news_Proximity_090806.htm?newstitle=deutscheboersesystemsandixeuro&location=press
India	BO	20101115	http://www.business-standard.com/article/markets/bse-s-co-location-facility-begins-to-get-brokers-111012000030_1.html
India	NS	200908	http://www.nseindia.com/circulars/circular.htm
Italy	MI	200909	http://www.borsaitaliana.it/borsaitaliana/ufficio-stampa/comunicati-stampa/2008/081110avviotradelect.en_pdf.htm
Japan	OS	200811	http://www.ose.or.jp/e/news/14734
Japan	T	200905	http://www.tse.or.jp/english/rules/equities/arrowhead/info.html
Netherlands	AS	200804	https://europeanequities.nyx.com/sites/europeanequities.nyx.com/files/327777.pdf
Portugal	LS	200804	https://europeanequities.nyx.com/sites/europeanequities.nyx.com/files/327777.pdf
Singapore	SI	201104	http://www.sgx.com/wps/wcm/connect/sgx_en/home/higlights/news_releases/news+update+sgxs+reach+trading+engine+goes+live
Sweden	ST	20080625	http://www.ft.com/intl/cms/s/0/b2cde4f0-42ce-11dd-81d0-0000779fd2ac.html#axzz2RRv89HWs
Switzerland	S	20080624	http://www.six-swiss-exchange.com/swx_messages/online/swx_message_200901161725_en.pdf
Taiwan	TW	2010Q4	http://asiaetrading.com/taiwan-stock-exchange-to-launch-co-location-services/
UK	L	200809	http://www.londonstockexchange.com/about-the-exchange/media-relations/press-releases/2008/launchofexchangehostingcreatessub-millisecondaccesstoitsmarkets.htm
USA	NASDAQ	200504	http://www.wallstreetandtech.com/electronic-trading/data-latency-playing-an-ever-increasing/199702208
USA	NYSE	201008	http://www.ft.com/intl/cms/s/0/2d62bcfa-ad26-11de-9caf-00144feabdc0.html#axzz2RRv89HWs