

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection Lee Kong Chian School Of  
Business

Lee Kong Chian School of Business

---

10-2020

### Algorithmic trading and market quality: International evidence

Ekkehart BOEHMER

*Singapore Management University, EBOEHMER@smu.edu.sg*

Kingsley FONG

*University of New South Wales*

Juan Julie WU

*University of Nebraska - Lincoln*

Follow this and additional works at: [https://ink.library.smu.edu.sg/lkcsb\\_research](https://ink.library.smu.edu.sg/lkcsb_research)



Part of the [Corporate Finance Commons](#), and the [Finance and Financial Management Commons](#)

---

#### Citation

BOEHMER, Ekkehart; FONG, Kingsley; and WU, Juan Julie. Algorithmic trading and market quality: International evidence. (2020). *Journal of Financial and Quantitative Analysis*. 56, (8), 2659-2688.  
Available at: [https://ink.library.smu.edu.sg/lkcsb\\_research/6685](https://ink.library.smu.edu.sg/lkcsb_research/6685)

This Journal Article is brought to you for free and open access by the Lee Kong Chian School of Business at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection Lee Kong Chian School Of Business by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylds@smu.edu.sg](mailto:cherylds@smu.edu.sg).

**ALGORITHMIC TRADING AND MARKET QUALITY: INTERNATIONAL  
EVIDENCE**

Ekkehart Boehmer  
Kingsley Fong  
Juan (Julie) Wu\*

---

\* Boehmer is from Lee Kong Chian School of Business, Singapore Management University, Singapore (ekkehart.boehmer@gmail.com). Fong is from the University of New South Wales Business School, Sydney, Australia (k.fong@unsw.edu.au). Wu is from the University of Nebraska-Lincoln (juliewu@unl.edu).

# ALGORITHMIC TRADING AND MARKET QUALITY: INTERNATIONAL EVIDENCE

## Abstract

We study the effect of algorithmic trading (AT) on market quality between 2001 and 2011 in 42 equity markets around the world. We use exchange co-location service that increases AT as an exogenous instrument to draw causal inference of AT on market quality. On average, AT improves liquidity and informational efficiency but increases short-term volatility. Importantly, AT also lowers execution shortfalls for buy-side institutional investors. Our results are surprisingly consistent across markets and thus across a wide range of AT environments. We further document that the beneficial effect of AT is stronger in large stocks than in small stocks.

JEL classification: G12, G14

Keywords: Algorithmic trading, market quality, high frequency trading, buy-side institution execution costs

## I. Introduction

By most accounts, high frequency trading (HFT) represents the majority of the trading volume in today's equity markets.<sup>1</sup> Besides the sheer trading volume, HFT is important because their strategies are often not transparent, nor are the underlying strategies well understood. These reasons elicit 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 regularly scrutinize algorithmic and high-frequency order submission strategies and their consequences. Despite the intensity of this debate and a large theoretical and empirical literature in this area, many questions remain unanswered (see, e.g., Chordia et al. (2013), Jones (2013), O'Hara (2015), and Menkveld (2016) for reviews).

In this paper, we take a basic but comprehensive approach that contributes new insights and novel broad-sample evidence to this debate. Following Hendershott, Jones, and Menkveld (2011), we construct proxies for algorithmic trading (AT) intensity, a precondition for HFT, from order-related message traffic, and estimate an instrumental variable model that allows us to make causal inferences. We use eleven years of intraday data on security-level quotes and trades in 42 equity markets around the world, on average covering more than 21,507 firms per year. This new and comprehensive sample allows us to exploit the effect of AT intensity on market quality across stocks, markets, and trading infrastructures.<sup>2</sup>

---

<sup>1</sup> HFT, also called low latency trading (LLT), refers to the activity of algorithms that emit orders or order cancellations, reacting within milli- or nano-seconds to market updates or new information.

<sup>2</sup> Strictly speaking, HFT is a subset of algorithmic trading. HFT likely accounts for most algorithmic message traffic and volume. Because our analysis of AT has implications for HFT, we often use the terms "algorithmic trading" and "high frequency trading" interchangeably.

In addition to examining the relationship between algorithmic trading and standard market quality metrics of liquidity, informational efficiency, and short-term volatility, we investigate how AT affects buy-side institutional investors. Market quality metrics are indicators of specific aspects of overall market quality. With HFT accounting for a large fraction of trading, market wide statistics do not necessarily reflect the trading outcome for other market participants. For instance, institutional investors trade large quantities of shares and execute their orders over an extended period of time. This trading behavior increases the chances that these participants become targets of predatory trading or “back running” strategies that some high frequency traders use, hence leading to greater trading costs (Brunnermeier and Pedersen (2005), Yang and Zhu (2019)). Therefore, the net effect of higher AT intensity on institutional investors’ execution costs is an open empirical question. As an overall assessment of the impact of AT intensity on institutional investors, we estimate and analyze a commonly used trading cost measure that can be applied to large institutional orders, execution shortfall, using trade records of institutional investors around the world.

To help draw causal inferences of AT intensity on market quality, we use the formal introduction of co-location services as an exogenous instrument for AT intensity, an instrument with the same interpretation across markets that differ in trading protocols and market structure. Co-location allows fast traders to minimize data turnaround time by physically locating their computer hardware close to the exchange’s hardware. The event itself, the actual introduction of co-location services, marks the exchange’s commitment to low-latency infrastructure. Other things equal, this commitment makes the exchange more attractive to AT and should thus lead to an increase in AT intensity. Conversely, the introduction of co-location services has little direct impact on market quality (Menkveld and Zoican (2017)). Therefore, we develop instruments based on co-location events and use them to help assess the effect of AT on market quality.

We find that AT affects various dimensions of market quality. First, higher AT intensity improves liquidity. Specifically, we find that higher AT leads to smaller quoted and effective spreads. Second, more AT results in higher price efficiency. Stocks with higher AT intensity have smaller absolute value of autocorrelations of intraday returns, an indicator of prices being more random. Third, higher AT intensity increases short-term volatility: daily price range, intraday return variances and daily return volatility all increase with AT intensity.

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, producing presumably desirable price volatility. It is thus conceivable that the greater efficiency induced by 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 to 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 a market-making strategy. In this case, high volatility could be desirable because it attracts additional liquidity that would otherwise be absent. 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 leads to higher 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 by withdrawing their supply of liquidity or that the traders attracted by high volatility take liquidity rather than provide liquidity. Either way, AT activity that takes place during high-volatility episodes and lower liquidity appears undesirable to the market.

Although we document higher volatility induced by greater AT, we are agnostic about the exact causes of such high volatility. More intense AT-induced volatility is consistent with the

model by Rosu (2019), who shows that volatility increases when more high frequency traders enter the market or when their information becomes more precise. One might be concerned, however, that the AT effect on volatility we document could be driven by the release of firm fundamental news. To mitigate this concern, we control for firm news announcements in our analysis. We find that a stock's volatility is indeed high at news announcements, but AT's effect on volatility extends well beyond the news event itself.

The multifaceted effects of AT on traditional measures of market quality motivate us to look at buy-side institutional transaction costs. For buy-side institutions, smaller spreads may help lower their trading costs, but higher information efficiency could increase their costs of trading with informed traders. Furthermore, higher volatility could increase their trading costs because large orders take longer time to execute and prices may be pushed away while their orders are being executed. Thus, ex ante it is not clear how execution shortfalls, which capture cumulative price impact of a large order over the entire execution period rather than the half-spread for a single market order, are affected by AT. We find that the net effect of AT is beneficial for buy-side institutional investors: higher AT intensity reduces execution shortfalls.

While the average effect of AT on market quality has both benefits and costs, we observe substantial variation in AT intensity. Thus, understanding the cross-sectional differences of the benefits and costs of greater AT intensity is important. We examine how AT's effect on market quality varies with market capitalization. For example, providing liquidity in large-cap stocks requires less effort and their trading is active enough to allow for significant algorithmic activity. Implementing high-frequency market-making strategies in particular is likely easier in large-cap stocks than in small-cap stocks. To analyze these issues, we divide stocks into terciles based on market capitalization within each equity market and allow the effect of AT to vary across sizes.

We observe that the benefits of high AT intensity are more pronounced among large and medium stocks than in small stocks. When AT increases, small stocks experience less liquidity

improvement and little price efficiency enhancement. The main costs of high AT intensity, in terms of elevated volatility, are higher in small stocks than in large stocks. These findings suggest some segmentation of liquidity provision across large and small stocks, indicating that potential regulations on HFT should not impose one-size-fits-all rules.

Our study contributes to the growing literature on HFT in the following ways. First, we provide new large-sample, cross-country evidence of AT's effect on market quality. Existing literature offers mixed evidence on the effect that HFT has on market quality. Many studies find that HFT increases liquidity and price discovery, but others raise concerns about the quality or usefulness of HFT-provided liquidity, noting that such liquidity is often short-lived and negatively affects information acquisition and volatility.<sup>3</sup> This is partly due to the fact that most studies either employ a small sample of U.S. stocks traded by HFT firms or focus on a sample from a particular market outside the U.S.<sup>4</sup> Our analysis of a large international sample, particularly covering a period when algorithmic trading becomes more mature around the globe, helps draw broader inference on the effect of AT on market quality.

Second, our examination and results of the relationship between algorithmic trading and volatility shed light on an overlooked issue that has potential policy implications. In contrast to prior findings that market-making HFT reduces volatility based on 30 stocks listed on Nasdaq-OMX Stockholm (Hagströmer and Nordén (2013)), we find that across a large number of equity markets, more intense AT activity induces higher volatility. Our finding contributes to the

---

<sup>3</sup> For example, among others, Brogaard et al. (2014) find improved liquidity by HFT. 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. Weller (2018) finds that AT can reduce price informativeness due to less information acquisition.

<sup>4</sup> For example, among other studies that look at markets outside the U.S., Hendershott and Riordan (2013) examine 30 German stocks in a stock index on the Deutsche Boerse in January 2008; Menkveld (2013) examines one HFT firm on Euronext and Chi-X; Brogaard et al. (2015) examine NASDAQ OMX Stockholm; Anand and Venkataraman (2016) look at Toronto Stock Exchange; van Kervel and Menkveld (2019) examine Swedish market for large cap stocks.



theoretical debate on the relation between high frequency market-making and volatility (Baruch and Glosten (2016), Yueshen (2017), and Rosu (2019)).

Third, our study goes beyond examining AT's effect on traditional trading cost measures such as bid-ask spreads. Our analysis of buy-side institutions improves our understanding of how AT affects this important group of market participants. We show that these institutions incur lower execution costs when AT intensity is higher. Our international evidence complements recent studies that examine execution costs on a specific market (e.g., van Kervel and Menkveld (2019) on Swedish index stocks; Korajczyk and Murphy (2019) on the Canadian market; and Tong (2015) on Nasdaq).

Fourth, we establish important differences across stock sizes in the effect of AT on market quality under different trading environments. This is important, because given the increasing focus on how to regulate HFT, our results suggest segmentation of liquidity provision and imply that optimal regulation may need to impose different rules on different categories of stocks.

Our paper is organized as follows. In Section II, we discuss our data sources and define the key variables used in our analysis. We discuss our empirical design in Section III and present our main results in Section IV. Section V provides additional analysis and discussion. Section VI concludes.

## **II. Data and Variables**

### **A. Data sources**

We combine several data sources for our analysis. The Thomson Reuters Tick History (TRTH) database contains intraday trades and quotes for many equity markets around the world. Intraday trades and quotes for U.S. stocks listed on NYSE and NASDAQ are from the Trades and Quotes (TAQ) database. We use Thomson Reuters Datastream and Center for Research and Security Prices (CRSP) to obtain stock-level and market-level data. In addition, data on buy-side

institutional transaction costs come from the Ancerno database compiled by Ancerno Ltd. (formerly the Abel/Noser Corporation). Finally, we use RavenPack to obtain data on a firm's news release.

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 in our sample period. 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.<sup>5</sup> This procedure, together with TAQ, produces stocks trading on 42 equity exchanges in 37 countries spanning from 2001 to 2011.<sup>6</sup>

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 Reuters' Speedguide and the Handbook of World Stock, Derivative and Commodity Exchanges.

We also employ data from Ancerno in our analysis. 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 data using date, CUSIP, ISIN, and ticker. To accommodate investors who

---

<sup>5</sup> 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.

<sup>6</sup> 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.

split orders across brokers, we follow Anand et al. (2012) by aggregating trades into daily orders by client, stock, date, and trade direction.

Ravenpack, a leading news provider that covers corporate news releases, is used for some of our analysis. RavenPack generates company relevance scores ranging from 0 to 100 to measure the informational content of a news story. The company relevance scores allow researchers to extract news articles that are related to a specific company in the database. A relevance score equal to 100 indicates that a company is quoted as the main subject of a news release. Following Dai, Parwada, and Zhang (2015), we record the presence of news articles that are related to a specific company with relevance scores equal to 100 on a particular day.

## **B. Variable construction**

Our main objective is to examine the effect of algorithmic trading on market quality. Following previous literature, we construct variables that describe several dimensions of market quality, focusing on liquidity, informational efficiency, and volatility. We also construct measures for buy-side institutional trading costs. We describe these variables in this subsection, along with our proxies for AT.

### *Liquidity measures*

We compute several standard measures of liquidity. For each stock, we have the best quoted spread throughout the trading day. For a given time interval  $s$ , the relative quoted spread,  $RQS$ , 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. We standardize the quoted bid-ask spread by the quote midpoint. When aggregating over a trading day, we use time-weighted averages of  $RQS$ . Wider  $RQS$  suggests less liquidity.

To take into account possible price improvement arising from hidden liquidity, we compute the relative effective spread,  $RES$ . 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). We then standardize the measure by the quote midpoint at the time of the trade.  $RES$  measures the total price impact of a trade.

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

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

where  $M_{k+5}$  is the midpoint five-minutes after the  $k^{th}$  trade.  $RRS$  can be interpreted as the reward for providing liquidity.<sup>7</sup>

The permanent component,  $RPI$ , is defined as

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

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

---

<sup>7</sup> We follow the convention in the literature to use a 5-minute post-trade horizon (e.g., Boehmer, et al (2007), O'Hara and Ye (2011), Battalio, et al (2016)). This appears reasonable during our sample period 2001-2011. In the U.S. market after 2011, Conrad and Wahal (2019) show that a shorter post-trade horizon can better capture liquidity provision by high frequency market makers.

### *Informational efficiency*

We compute intraday measures of relative 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 and less efficiency. We thus use the absolute value of quote midpoint return autocorrelations. The larger this value is, the less efficient the prices are. We estimate this measure for each stock-day,  $|AR10|$  ( $|AR30|$ ), based on 10- (30-) minute return intervals (see Chordia et al. (2005)).

### *Volatility*

Our primary measure of short-term 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. As higher-frequency measures, we also use the log of intra-day return variances computed from 10-minute and 30-minute mid-quote returns,  $Ln(Ret10\_Var)$  and  $Ln(Ret30\_Var)$ . Additionally, we compute measures of realized volatility using the absolute value of daily returns.

### *Execution shortfall*

In addition to the above traditional measures, we use execution shortfall for the buy-side firms that report to Ancerno (e.g. Anand et al. (2012)). 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:

$$SHORTFALL = D_k (XP - RP) / RP \quad (5)$$

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,  $XP$  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.

### *Proxy for AT*

The spread of AT over the past two decades has spurred many studies to develop measures of AT and examine their consequences. Some studies are able to observe the actual identifier for a group of HF traders. For example, within the U.S. market, Brogaard et al (2014), Brogaard et al (2015) and Carrion (2013) use a 2008-2009 Nasdaq sample that summarizes the aggregate order flow generated by 26 HFT firms that capture about three quarters of trading volume in the sample stocks. A few other studies have access to actual HFT data for a specific foreign market.<sup>8</sup> Most of the remaining studies (e.g., Hendershott, Jones, and Menkveld (2011), Hasbrouck and Saar (2013)) use some variations of message-to-trade ratios as proxies for AT/HFT. 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 algorithmic traders than for non-algorithmic 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 and, in principle, allow the strongest inferences.

We follow Hendershott, Jones, and Menkveld (2011) and construct a proxy for AT that reflects fast order submissions and cancellations that are associated with algo-trading strategies. Specifically, the proxy for AT is the negative of trading volume in USD100 divided by the number of messages. 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.

---

<sup>8</sup> For example, outside the U.S., Hendershott and Riordan (2013) examine 30 Deutscher Aktien Index stocks on the Deutsche Boerse in January 2008. Menkveld (2013) examines one HFT firm on Euronext and Chi-X. Brogaard et al. (2015) find NASDAQ OMX Stockholm introduction of optional colocation improve effective spread and depth for slow traders and the market overall. Anand and Venkataraman (2016) find that synchronous withdrawal of liquidity provision by HFT firms under difficult market conditions contributes to fragility of liquidity supply on Toronto Stock Exchange.

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). It thus 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 intensity 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 traded on the same exchange. 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.<sup>9</sup>

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 impose serious problems on our analysis. For example, the HFT strategies mentioned in the SEC 2010 concept release involve most activity *at* the BBO, rather than *behind*

---

<sup>9</sup> The AT proxy also has potential drawbacks. First, as suggested by Hendershott, Jones, and Menkveld (2011), it may have low power for small stocks due to less algorithmic trading. For small stocks the AT proxy could potentially capture market uncertainty, when orders are submitted and cancelled yet few of those orders result in transactions. To mitigate this concern, we explicitly control for market cap in our regression analysis. Second, the AT proxy may not be very good during the release of news announcements. Fundamental volatility is high during news announcements, and the event, rather than algorithmic trading, could cause a high number of electronic messages to transaction volume ratio and high price volatility. To help address this concern, we control for news announcements in additional analysis. Our results suggest that the AT proxy is a reasonable proxy for algorithmic trading even in the presence of high fundamental volatility during news releases.



it. Therefore, the AT activity in our BBO trade data set is highly correlated with AT activity in an order trade dataset.<sup>10</sup>

### **III. Sample and Methodology**

#### **A. Descriptive statistics**

To be included in our analysis, a stock needs to have data for more than 21 trading days during the sample period to minimize the concern of thin trading. 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 more than half, or from 473 to 560.

Our key 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 average number of messages per stock day in 2001 and in 2011 along with the growth over this eleven-year period. We make several important observations. First, message traffic grows over time, often many folds, with the exception of Greece. Overall, message traffic grows by more than ten times across markets, from 289 messages per stock day in the beginning of the sample period to 4,042 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

---

<sup>10</sup> We formally address the correlation between AT measures based on order level data, and AT based on trades and best quotes. We repeat Hendershott, Jones, and Menkveld (2011)'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 order level data and the TAQ data overlap is very similar to the period presented in Hendershott, Jones, and Menkveld (2011). Our exercise using only NYSE activity yields qualitatively identical results for their 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.

of the decade when exchanges started to offer co-location services. Figure 1 also reveals that most of the message growth comes from quote messages rather than trade messages, further motivating our AT proxy for the unobservable AT intensity.<sup>11</sup>

Figure 2 presents the monthly time-series of liquidity, efficiency, and volatility, respectively. For each market day, we first compute an equally weighted average across firms, and then calculate the average within each market month. In the figure, we plot the monthly time series of averages across markets. The relative quoted and effective spreads,  $RQS$  and  $RES$ , in Panel A show similar patterns. For example,  $RES$  began at 250 bps in the beginning of 2001 and declined to 150 bps by the end of 2007. Afterwards, it peaked at the end of 2008, when the financial crises around the world started to unfold.  $RQS$  declined from 500 bps in 2001 to 246 bps in 2007, and then rose to 660 bps during the financial crisis before declining to 400 bps 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 decreased until mid-2007 and then increased again. We also observe that  $RRS$  exceeds  $RPI$  in every year by about 50%.

Panel B plots two efficiency measures,  $|AR10|$  and  $|AR30|$ . Both measures show a slight decrease over the sample period. The intraday volatility measures in Panel C decline slightly over the first half of the sample period, with a large spike towards the end of 2008 and a smaller increase in 2011.

---

<sup>11</sup> Conrad, Wahal and Xiang (2015) make similar observations that HFT mainly act via quotes than trades for the U.S. and Japan markets.

Panel D graphs the monthly measures of execution shortfall. We see an increase in execution shortfall from below 20 bps in the pre-crisis period to more than 35 bps during the financial crisis period and fall back in the post-crisis period.

## **B. Methodology**

To establish a causal relation between AT and market quality, we use an instrumental variable approach. 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. “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. Similarly, Baron et al. (2019) find that HFTs that improve their latency rank due to co-location upgrades (at NASDAQ OMX Stockholm) deliver better trading performance that comes through short-lived information channels as well as risk management channels.<sup>12</sup>

To introduce a co-location 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

---

<sup>12</sup> 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. We believe that co-location allows the cleanest measurement. Moreover, the first co-location also tends to be prominently and consistently reported by the media.

co-location-related latency. Co-location introductions mark events that are fairly homogenous across exchanges, in that the event specifically provides infrastructure for fast traders and signals an exchange’s commitment to accommodating such traders. As a potential caveat, the precision of reporting on co-location event dates could differ across countries. Yet, 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.

During our sample period, 22 markets have adopted co-location services (see the list of co-location dates in Appendix 1).<sup>13</sup> To allow a reasonable amount of time to observe any potential changes in the effects of AT due to co-location, we select a two-year window centered on the first co-location implementation. Specifically, the co-location dummy takes a value of zero in the 12 months leading to the adoption of co-location and switches to one during the 12 months after co-location. To account for the clustering of co-location events across markets, robust standard errors are clustered by date.

In the first stage, we regress our AT proxy on the 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 the following equation:

$$MQ_{it} = \alpha_{it} + \beta AT^*_{it} + \delta X_{it} + \varepsilon_{it}, \quad (6)$$

where  $AT^*$  is the predicted value from the first-stage regression, and  $X$  is a vector of control variables. This vector includes variables such as share turnover, inverse price, the log value of market capitalization, the lagged dependent variable, and intraday price range (a proxy for volatility, omitted from the volatility regressions which add  $RES$  and  $|AR30|$ ). These control variables have been shown to be related to market quality (e.g., Hendershott et al. (2011)).

---

<sup>13</sup> For two markets (Germany and Taiwan) we couldn’t obtain the specific month during which co-location was implemented, so we use the year-end as the event month. NYSE’s co-location was first offered through its acquisition of TransactTools in January 2007 and its most recent co-location data service started in August 2010.

Following Boehmer and Kelley (2009), we lag all control variables by one day to ensure that explanatory variables are predetermined.<sup>14</sup> To make comparisons meaningful across firms and markets, all continuous variables are standardized within each firm during the co-location event window.

In the first stage regression, co-location dummy has a coefficient of 0.108 ( $t$ -stat = 13.1). This suggests that the adoption of co-location services significantly contributes to AT growth.

#### **IV. Empirical results**

In this section, we first analyze how AT affects various spread measures, price efficiency, short-term volatility, and institutional buy-side execution costs. We then examine if there exists any cross-sectional variations in the AT effect on these dimensions of market quality.

##### **A. Main results**

Table 3 presents the regression estimation results based on Equation (6). Panel A of Table 3 reports the effect of AT on liquidity after controlling for other variables.<sup>15</sup> Each row represents a measure of liquidity. We find that AT significantly lowers the quoted and effective spreads. A one-standard deviation increase in AT improves effective spreads ( $RES$ ) by 0.7 standard deviations. When we decompose  $RES$  into its adverse selection and transitory components, we find that AT's effect on the permanent portion ( $RPI$ ) is larger than the effect on the transient portion, the realized spread ( $RRS$ ).  $RRS$  is a measure of the premium earned by liquidity suppliers. These results suggest that, on average, greater AT intensity leads to improved liquidity by reducing the information content of trades rather than reducing the rewards earned by liquidity suppliers.

---

<sup>14</sup> Using contemporaneous values produces similar results.

<sup>15</sup> Coefficients on control variables are not tabulated.

The positive effect of AT on liquidity we document across the international markets is in line with several existing studies. Hendershott, Jones, and Menkveld (2011) are among the first to show that algorithmic trading leads to better liquidity. They use the 2003 introduction of autoquote at the NYSE as an instrument to establish causality from algorithmic trading to market quality improvements. Using HFT activity inferred from millisecond-level responses, Hasbrouck and Saar (2013) find improvements in liquidity when these fast traders are more active. Malinova, Park, and Riordan (2013) show that a decline in HFT reduces liquidity for retail traders in the Canadian market.

Panel B presents the effect of AT on informational efficiency. The significantly negative coefficients on AT suggest that more intense AT leads to smaller absolute value of intraday return autocorrelations, indicating that prices are closer to a random walk. These results indicate that across markets, more AT consistently leads to an improvement in informational efficiency. Our finding complements Brogaard, Hendershott and Riordan (2014) who show that HFT facilitates price discovery among a sample of Nasdaq stocks. Conrad, Wahal and Xiang (2015) also find higher quoting activity is associated with better price efficiency.

Panel C presents the effect of AT on short-term volatility. More intense AT leads to higher volatility, and the results are uniform across volatility proxies whether we look at intraday realized volatility or daily realized volatility. Larger AT-induced short-term volatility is also documented in several others studies in different contexts. For example, Kirilenko et al. (2014) argue that HFT worsened (but did not cause) the May 6, 2010 flash crash. Dichev, Huang and Zhou (2014) show that trading per se generates excess volatility, arguing that HFT can lead to undesirable levels of volatility.

Volatility is important to traders 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.<sup>16</sup> We will return to the volatility discussion later.

In Panel D, we focus on the impact of AT intensity on execution shortfall, an overall assessment of the execution costs incurred by buy-side institutional investors. As alluded to earlier, the majority of trading activity involves AT, so that market wide metrics of trading may not properly represent the outcome of buy-side institutional traders, because they typically trade large quantities of shares and their orders are executed over a longer trading horizon. High frequency traders can choose to provide liquidity, conduct order-anticipation strategies or back-running strategies on large institutional orders (van Kervel and Menkveld (2019), Yang and Zhu (2019)). While narrower spreads can lower trading costs for an average trader who submits a market order, higher information efficiency may increase the costs of trading with informed traders. Moreover, higher volatility could lead to larger trading cost variability and increase execution risk. Thus, the net effect of higher AT intensity on institutional investors' execution cost is best answered by the data.

Panel D reports a significant and negative coefficient on AT. This suggests that more AT helps reduce buy-side execution shortfall. Thus, the net effect of smaller spreads, better information efficiency and higher volatility induced by higher AT intensity appears beneficial to this group of market participants across many markets. Our result based on a broad international sample helps resolve different findings reported in several recent studies that examine institutional

---

<sup>16</sup> Consistent with this concern, Egginton, Van Ness, and Van Ness (2016) 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 (2014) show that quote stuffing has negative effects on trading.

investors' execution costs on a specific market (e.g., van Kervel and Menkveld (2019) on Swedish index stocks; Korajczyk and Murphy (2019) on the Canadian market, and Tong (2015) on Nasdaq).

## **B. Heterogeneity of AT impact**

Both anecdotal information as well as academic evidence show that AT plays a more important role in larger stocks than in smaller stocks. We address such heterogeneity by estimating AT's effect on market quality for different market-cap groups. We group stocks into small, medium, and large categories based on their market capitalization within each market. Then we interact the instrumented AT with the three size dummies in the second stage to capture the potentially different AT impact across market-capitalizations. Specifically, we expand Equation (6) as follow:

$$MQ_{it} = \alpha_{it} + \beta_1 AT_{it}^* \times S_i + \beta_2 AT_{it}^* \times M_i + \beta_3 AT_{it}^* \times L_i + \delta X_{it} + \varepsilon_{it}, \quad (7)$$

where  $S$ ,  $M$ ,  $L$  represent dummies for small, medium, and large stocks, respectively. Table 4 presents the results.<sup>17</sup>

Panel A focuses on liquidity measures. Results show that AT improves liquidity in all size categories, but its effect is stronger in large stocks. Take *RES* for example, the negative coefficient on  $AT^* \times L$  is almost twice as large as the negative coefficient on  $AT^* \times S$ , suggesting a larger reduction in effective spreads in large-cap stocks compared to small-cap stocks. Panel B reports AT's effect on price efficiency across three sizes. Again, AT's effect is mainly present in medium and large stocks. AT has minimal improvement in price efficiency in small stocks. Panel C presents the volatility results. Although we see higher volatility across all stocks, smaller stocks experience a larger increase in volatility (except for the daily price range, which is more sensitive for large firms).

---

<sup>17</sup> Regressions also include the size dummies but are not tabulated.



Lastly, Panel D shows that the beneficial effect of AT on execution shortfalls is mostly concentrated in large and medium stocks. Note that the negative (but insignificant) coefficient on the interaction between AT and small stocks indicates that AT does not lead to lower execution cost when buy-side institutional investors trade small stocks.

Overall, results in this subsection suggest an important heterogeneity in AT's effect on market quality. AT leads to better liquidity, faster price efficiency, moderately higher volatility and lower execution costs in large firms. Small firms, however, experience less liquidity improvement and much higher volatility. These results indicate that segmentation in liquidity provision across stocks of different sizes goes beyond the U.S. market (Hendershott, Jones, and Menkveld (2011)).

### **C. AT and volatility**

We show earlier that AT leads to higher volatility. This finding raises the question regarding the nature of this volatility. Given that AT leads to better informational efficiency, it is conceivable that the elevated volatility reflects faster price adjustments when new information arrives. Under this scenario, the higher volatility reflects new information, not noise, and could therefore be desirable. Another possibility is that narrower spreads, which also result from 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 generate 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 trade-to-trade volatility. We take several approaches to address these possibilities.

First, we directly control for price efficiency and liquidity in the volatility regression. Specifically, efficiency proxy  $|AR30|$  and liquidity proxy  $RES$  are already included in the regression model whenever we estimate AT on volatility in Equation (6). By controlling for

efficiency and round-trip transaction costs, we hold constant price efficiency and liquidity which likely are the main sources of “good” volatility. Since the reported results already control for informational efficiency and liquidity associated with “good” volatility, it is difficult to attribute the elevated volatility that we observe to either faster incorporation of new information or to tighter spreads.

Second, we examine whether higher AT-induced volatility is associated with improved liquidity. Our results suggest that more AT leads 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 a simple test that sheds additional light on the relation among AT, volatility, and liquidity. We employ a two-step procedure. In the first step, we estimate cross-sectional regression within each market day, regressing liquidity, price efficiency, and volatility on AT and control 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. In the second step, we compute Spearman rank correlations between liquidity and volatility effects and between efficiency and volatility effects.<sup>18</sup>

Panel A of Table 5 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 wider spread, rather than narrower spread. Or, conversely, if high volatility indeed attracts algorithmic traders, these traders demand—rather than supply—liquidity.

---

<sup>18</sup> Pearson correlations produce the same inferences.

Therefore, at least in our sample, the costs of high volatility are not contemporaneously offset by greater liquidity, as suggested by Castura et al. (2010).

In contrast to the relationship between liquidity and volatility, Panel B shows that 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 already 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.

Third, we address a potential concern that fundamental volatility is high during news announcements, and the news event, rather than algorithmic trading, causes the high volatility. To examine this possibility, we explicitly control for news announcements. Specifically, we obtain news data from RavenPack, a professional news provider whose products have been used in some recent studies (e.g., Dai, Parwada, and Zhang (2015)), and augment our main second-stage regression into the following specification:

$$MQ_{it} = \alpha_{it} + \beta_1 AT^*_{it} + \beta_2 AT^*_{it} \times News_{it} + \beta_3 News_{it} + \delta X_{it} + \varepsilon_{it}, \quad (8)$$

where *News* is a dummy variable equal to one if there is news announcement for a firm on a day and zero otherwise.

Table 6 reports the regression results on volatility controlling for news. As expected, the significant coefficient on *News* suggests that volatility is heightened when news comes to the market. More importantly, the coefficient on instrumented AT is still significantly positive, suggesting that AT leads to higher volatility even in the absence of news. The interaction term of  $AT^* \times News$  indicates that AT's effect on volatility is even stronger on news days.

Overall, these three approaches complement one another in addressing the nature of the volatility. Although we are agnostic about the exact causes of AT-induced volatility, the overall evidence suggests that AT-induced volatility is difficult to attribute to more “good” volatility that

would arise from faster price discovery, to algorithmic traders' inclination to enter the market when volatility is high, or to trading on fundamental news. Although we see beneficial effects of AT in improving liquidity and efficiency, AT also appears to elevate less desirable volatility.

## **V. Additional analysis**

### **A. Market-level IV estimation**

To complement our firm-level analysis, we also conduct market-level estimation with the adoption of co-location as an instrumental variable for AT. Specifically, each day, we aggregate all variables within each market by forming value-weighted averages across firms. This produces one time-series for each market. To maintain a balanced panel with all 42 markets, this analysis uses data from 2005 to 2011. We estimate a two-way panel across markets and days using co-location as the instrument for AT. In the first stage, we regress AT on the co-location dummy that equals one once that market officially starts co-location service. For markets that have not adopted co-location services by the end of our sample period, its co-location dummy has a value of zero. As expected, co-location significantly contributes to the intensity of AT, with a highly significant coefficient of 0.217 ( $t\text{-stat} = 27.9$ ). In the second stage, we use the predicted value of AT as a regressor. While such market-level IV estimation allows us to conduct a difference in differences analysis, it has lower statistical power.

The second-stage results are reported in Table 7. Overall, these market-level estimates deliver similar results to the firm-level analysis. Importantly, despite the lower power of this approach, most estimates remain significant. The liquidity measures in Panel A decline as AT increases, suggesting that more AT leads to narrower spreads. The efficiency measures in Panel B also decrease with higher AT, likewise implying greater efficiency. Panel C shows that AT increases volatility. In Panel D, we observe better execution quality for buy-side investors who experience smaller execution shortfalls when AT increases.

## **B. Market-specific analysis**

We also use another approach to help draw inferences across markets. Specifically, we first conduct country-specific analysis on the relation between AT intensity and market quality for each of the 42 markets, and then aggregate the results across markets. Cross-market inference is based on equal-weighted means of the 42 market-specific AT coefficients and simple cross-sectional t-statistics also based on these 42 observations. Although this approach does not directly address causality, we can still learn from the actual (not instrumented) relationships between AT and market quality, especially when the relationships are consistent across countries. We find that better liquidity, better price efficiency and higher volatility are observed when AT is high. These results are reported in Appendix 2.

## **VI. Conclusion**

Previous studies have produced mixed evidence on the effect of AT/ HFT on market quality. This is partly due to data limitations that constrain researchers to either a small sample of stocks or a specific market outside the U.S. We examine a large international sample over an extended period of time that helps draw broader inferences about the effect of AT on market quality.

Surprisingly consistent across the 42 markets in our sample, more intense algorithmic trading leads to improved liquidity, better price efficiency, and elevated volatility. The overall net effect of AT on buy-side institutions' execution shortfalls is positive, suggesting that more AT benefits this group of market participants. AT's effect on volatility cannot be attributed to more efficient prices that adjust faster to new information, to the activities of liquidity suppliers seeking out more volatile stocks, or to heightened volatility at news announcements.

Aside from AT's overall influence on market quality, its effects are not uniform across all stocks. AT's positive effects on liquidity and efficiency are more pronounced for large stocks, and AT also increases volatility more in smaller stocks.

Overall, our findings support prior results that attribute liquidity-enhancing and efficiency-enhancing effects to algorithmic and high frequency trading. We show that this finding is surprisingly pervasive across countries. We complement prior studies with evidence that AT's liquidity provision does not apply uniformly to all firms, suggesting that regulation should tread carefully and recognize that the effects of AT differ across firm characteristics. Equally importantly, we show that AT systematically increases volatility, thereby imposing costs on market participants.

These results suggest that measures of the true cost of algorithmic trading should incorporate realized volatility. Our evidence suggests that the positive effect of AT on volatility doesn't derive from reverse causality and doesn't rely on specific types of volatility that market participants would consider beneficial. We further show that the effect of AT on market quality, and hence the need to impose new regulation, differ dramatically across stocks and markets. This suggests that borrowing regulatory approaches from abroad is unlikely to produce a balanced playing field.

## References

Anand, A., and K. Venkataraman. "Market Conditions, Fragility and the Economics of Market Making." *Journal of Financial Economics*, 121 (2016), 327-349.

Anand, A.; P. Irvine; A. Puckett; and K. Venkataraman. "Persistence in Trading Cost: An Analysis of Institutional Equity Trades." *Review of Financial Studies*, 25 (2012), 557-598.

Baron, M.; J. Brogaard; B. Hagströmer; and A. Kirilenko. "Risk and Return in High-Frequency Trading." *Journal of Financial and Quantitative Analysis*, 54 (2019), 993-1024.

Baruch, S., and L. R. Glosten. "Strategic Foundation for the Tail Expectation in Limit Order Book Markets." Working paper, Columbia Business School (2016).

Battalio, R.; S.A. Corwin; and R., Jennings. "Can Brokers Have It All? On the Relation between Make - Take Fees and Limit Order Execution Quality." *Journal of Finance*, 71 (2016), 2193-2238.

Bessembinder, H. "Issues in Assessing Trade Execution Costs." *Journal of Financial Markets*, 6 (2003), 233–257.

Boehmer, E.; R. Jennings; and L. Wei. "Public Disclosure and Private Decisions: Equity Market Execution Quality and Order Routing." *Review of Financial Studies*, 20 (2006), 315-358.

Boehmer, E., and E. Kelley. "Institutional Investors and the Informational Efficiency of Prices." *Review of Financial Studies*, 22 (2009), 3563-3594.

Brogaard, J.; T. Hendershott; and R. Riordan. "High-Frequency Trading and Price Discovery." *Review of Financial Studies*, 27 (2014), 2267-2306.

Brogaard, J.; B. Hagstroemer; L. Norden; and R. Riordan. "Trading Fast and Slow: Colocation and Market Quality." *Review of Financial Studies*, 28 (2015), 3407-3443.

Brunnermeier, M., and L. H. Pedersen. "Predatory Trading." *Journal of Finance*, 60 (2005), 1825-1863.

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

Castura, J.; R. Litzberger; R. Gorelick; and Y. Dwivedi, “Market Efficiency and Microstructure Evolution in U.S. Equity Markets: A High-Frequency Perspective.” Working paper, RGM Advisors, LLC (2010).

Chordia, T.; A., Goyal; B. Lehmann; and G. Saar. “High-Frequency Trading.” Working paper, Emory University (2013).

Chordia, T.; R. Roll; and A. Subrahmanyam. “Evidence on the Seed of Convergence to Market Efficiency.” *Financial Journal of Economics*, 76 (2005), 271-292.

Conrad, J.; S. Wahal; and J. Xiang. “High Frequency Quoting, Trading, and the Efficiency of Prices.” *Journal of Financial Economics*, 116 (2015), 271-291.

Conrad, J., and S. Wahl. “The Term Structure of Liquidity Provision.” *Journal of Financial Economics*, forthcoming (2019).

Dai, L.; J.T. Parwada; and B., Zhang. “The Governance Effect of the Media's News Dissemination Role: Evidence from Insider Trading.” *Journal of Accounting Research*, 53(2015), 331-366

Dichev, I.; K. Huang; and D. Zhou. “The Dark Side of Trading.” *Journal of Accounting, Auditing & Finance*, 29 (2014), 492-518.

Egginton, J.; B. F. Van Ness; and R. A. Van Ness. “Quote Stuffing.” *Financial Management*, 45 (2016), 583-608.

Gai, J.; C. Yao; and M. Ye. “The Externalities of High Frequency Trading.” working paper, University of Illinois at Urbana-Champaign (2014).

Hagströmer, B., and L., Nordén. “The Diversity of High-Frequency Traders.” *Journal of Financial Markets*, 16(2013),741-770.



Hasbrouck, J., and G. Saar. "Technology and Liquidity Provision: The Blurring of Traditional Definitions." *Journal of Financial Markets*, 12 (2009), 143-172.

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

Hendershott, T.; C. Jones; and A. Menkveld. "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance*, 66 (2011), 1-33.

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

Jones, C.M. "What Do We Know about High-Frequency Trading?" Working paper, Columbia Business School (2013).

Kirilenko, A.; A. S. Kyle; M. Samadi; and T. Tuzun. "The Flash Crash: The Impact of High Frequency Trading on an Electronic Market." *Journal of Finance*, 72 (2017), 967–998.

Korajczyk, R., and D. Murphy. "High Frequency Market Making to Large Institutional Trades." *Review of Financial Studies*, 32 (2019), 1034–1067.

Lee, C. M.C., and M. J. Ready. "Inferring Trade Direction from Intraday Data." *Journal of Finance*, 46 (1991), 733–746.

Malinova, K.; A. Park; and R. Riordan. "Do Retail Traders Suffer from High Frequency Traders?" Working paper, University of Toronto (2013).

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

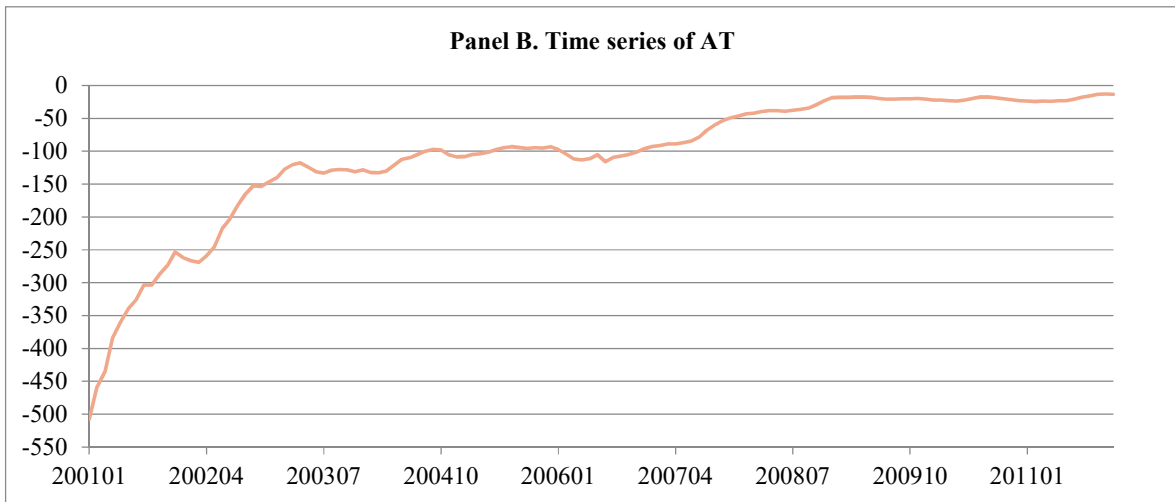

Menkveld, A. "The Economics of High-Frequency Trading: Taking Stock." *Annual Review of Financial Economics*, 8 (2016), 1-24.

Menkveld, A., and M. Zoican. "Need for Speed? Exchange Latency and Liquidity." *Review of Financial Studies*, 30 (2017), 1188–1228.

- O'Hara, M. "High Frequency Market Microstructure." *Journal of Financial Economics*, 116 (2015), 257-270.
- Rosu, I. "Fast and Slow Informed Trading." *Journal of Financial Markets*, 43 (2019), 1-30.
- SEC. "Concept Release on Equity Market Structure." U.S. Securities and Exchange Commission Release No. 34-61358. (2010), <http://www.sec.gov/rules/concept/2010/34-61358.pdf>
- Tong, L. "A Blessing or a Curse? The Impact of High Frequency Trading on Institutional Investors." Working Paper, Fordham University (2015).
- Van Kervel, V., and A. Menkveld. "High-Frequency Trading around Large Institutional Orders." *Journal of Finance*, 74 (2019), 1091-1137.
- Weller, B. M. "Does Algorithmic Trading Deter Information Acquisition?" *Review of Financial Studies*, 31 (2018), 2184–2226.
- Yueshen, B. Z. "Uncertain Market Making." Working paper, INSEAD (2017).
- Yang, L., and H. Zhu. "Back-Running: Seeking and Hiding Fundamental Information in Order Flows," *Review of Financial Studies*, 2019, forthcoming.

**FIGURE 1**  
**Messages and AT**

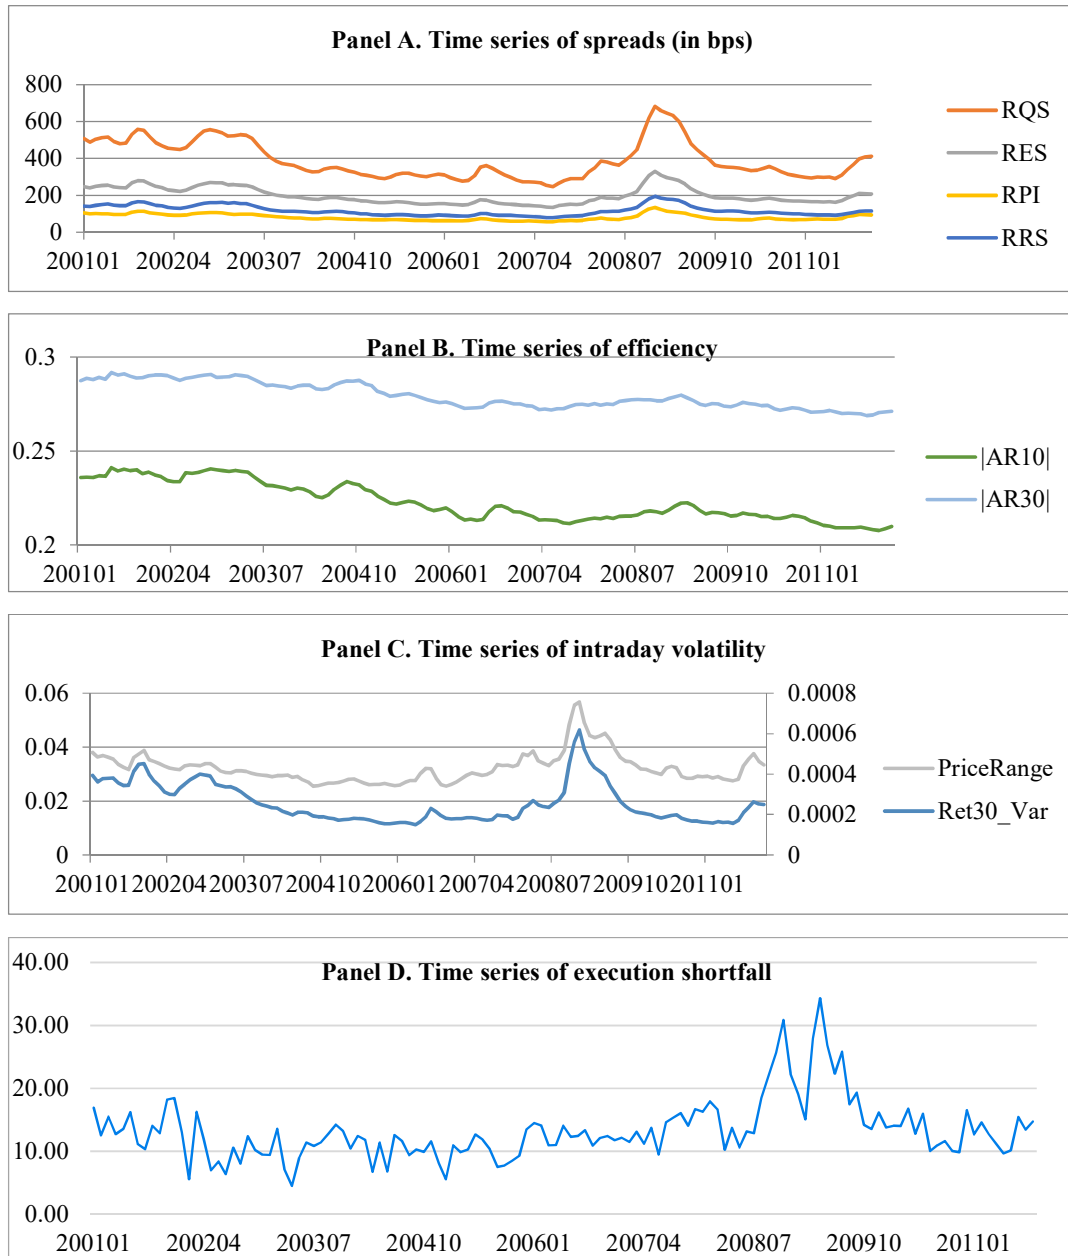
This figure presents the time series of aggregate messages and AT measure. Our sample covers the period 2001-2011. We count all intraday messages that represent trades or changes in the price or size of the best quotes for each stock. AT is the negative of dollar trading volume (\$100) per message. Then we compute equally weighted averages each market month, and report the mean across 42 markets.



**FIGURE 2**

**Market quality measures over time**

This figure graphs the time series of market quality over the sample period. 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 average across 42 markets. Our sample covers the period 2001-2011. Panel A is liquidity. RQS are time weighted relative quoted spreads, RES are relative effective spreads, RRS are 5-minute relative realized spreads, and RPI are 5-minute permanent price impacts. Panel B is efficiency.  $|AR10|$  ( $|AR30|$ ) is the absolute value of the daily average 10-minute (30-minute) quote-midpoint return autocorrelations. Panel C is volatility. Measures include the daily intraday price range standardized by the daily closing price, the variances of 30-minute quote midpoint returns ( $Ret30\_Var$ ). Panel D plots execution shortfall.



**TABLE 1**

**Number of stocks listed on markets in our sample**

This table lists the year-end number of stocks listed on each stock exchange. Our sample covers the period 2001-2011. Growth is calculated as the number of firms in 2011 divided by the number of firms in the beginning of the sample and then minus 1.

Market	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	Growth
Argentina	19	55	64	45	61	66	68	57	65	71	65	2.42
Athens	287	285	287	284	276	263	250	226	219	186	168	-0.41
Australia	714	666	921	990	1105	1343	1415	1215	1337	1254	1220	0.71
Brussels	121	105	124	117	132	132	139	128	139	130	131	0.08
Copenhagen	154	143	164	169	166	183	207	201	202	193	177	0.15
Xetra	239	264	257	264	314	345	377	368	366	381	383	0.60
Euronext												
Amsterdam	152	148	139	136	136	138	129	116	114	108	104	-0.32
Euronext Lisbon					42	41	41	43	42	40	35	0.83
Euronext Paris	480	462	456	454	567	635	669	611	615	585	575	0.20
Helsinki	152	149	143	148	147	148	147	141	138	134	129	-0.15
Hong Kong	389	419	493	520	549	580	629	701	765	832	802	1.06
Istanbul	277	286	287	299	305	314	313	307	314	338	359	0.30
Jakarta	261	253	277	274	267	280	295	185	304	327	347	0.33
Johannesburg	270	220	230	224	245	263	299	266	257	256	261	-0.03
Korea	311	300	667	670	678	703	713	742	711	754	764	1.46
Kuala Lumpur	692	728	782	824	857	870	855	811	824	829	813	0.17
London	838	781	821	892	1006	1334	905	799	735	891	736	-0.12
Madrid	120	119	110	110	108	116	125	123	120	117	114	-0.05
Mexican	69	61	67	66	72	82	81	77	91	97	87	0.26
Milan	271	280	265	264	280	289	310	281	274	272	264	-0.03
Mumbai	240	425	716	821	1077	1161	999	939	1348	1414	1242	4.18
NASDAQ	3606	3237	2948	2879	2822	2797	2743	2622	2466	2359	2237	-0.38
NSE (India)	500	542	598	671	750	880	1034	1107	1188	1329	1308	1.62
New Zealand		55	54	63	77	82	95	74	84	86	69	1.25
NYSE	1524	1498	1475	1475	1459	1443	1392	1347	1328	1330	1311	-0.14
Osaka	183	195	232	242	277	275	263	250	253	252	246	0.34
Oslo	174	174	164	172	198	206	236	226	218	213	202	0.16
Philippines	140	112	135	143	153	184	187	162	195	194	205	0.46
Santiago		81	69	84	84	90	91	92	100	89	100	1.23
Sao Paulo						306	400	379	385	368	350	1.14
Shanghai	580	669	742	803	709	716	734	754	821	843	833	0.44
Shenzhen	459	466	487	516	459	508	596	693	780	1105	1314	1.86
Singapore	315	326	368	424	472	514	553	525	552	572	545	0.73
Stockholm	328	315	314	319	345	384	443	462	468	472	467	0.42
Swiss	224	218	228	227	235	235	227	227	221	218	213	-0.05
Taiwan	518	592	636	667	661	670	676	704	735	755	765	0.48
Tel-Aviv	310	283	319	342	417	469	528	495	484	480	453	0.46
Thailand	314	335	363	407	446	466	467	480	480	486	487	0.55
Tokyo	1957	1996	2072	2205	2310	2371	2360	2289	2284	2286	2261	0.16
Toronto	614	619	658	712	780	872	930	936	913	915	942	0.53
Warsaw	137	135	137	179	209	237	307	323	354	366	369	1.69
Wiener Borse	45	61	53	50	54	61	71	73	68	56	58	0.29

**TABLE 2**

**Number of messages per stock-day by market**

This table reports the summary statistics of messages at the beginning and end of our sample period. Messages represent trades or changes in the price or size of the best quotes in each stock. Reported numbers are the median of the daily mean of messages for each stock within each market. Growth is calculated as messages in 2011 divided by messages in 2001 and then minus 1. \* Quote messages computed based on only price changes because quote sizes are not available until recent years. \*\* Data begin in 2005. \*\*\* Data begin in 2002. \*\*\*\* Data begin in 2006.

Market	Messages 2001	Messages 2011	Growth
Euronext Amsterdam	480	18718	38.00
Athens	324	245	-0.24
Australian	95	1403	13.82
Argentina	89	90	0.01
Thailand	219	551	1.52
Mumbai	120	132	0.10
Brussels	110	4375	38.63
Copenhagen	49	1039	20.33
Xetra	229	3597	14.68
Helsinki	233	2565	10.01
Hong Kong	247	1013	3.11
Istanbul	246	970	2.94
Johannesburg	74	2232	29.32
Jakarta	115	689	4.98
Kuala Lumpur	85	200	1.35
Korea	523	3547	5.79
London*	259	2171	7.38
Euronext Lisbon**	281	3167	10.25
Madrid	1014	8649	7.53
Milan	553	5621	9.16
Mexican	113	2600	21.91
NASDAQ	1006	15363	14.27
NSE (India)	357	2768	6.75
NYSE	1207	45998	37.10
New Zealand ***	33	66	0.97
Oslo	119	2002	15.85
Osaka	73	468	5.37
Euronext Paris	566	9266	15.37
Philippines	35	159	3.56
Swiss	99	688	5.92
Sao Paulo ****	621	4599	6.40
Singapore	170	279	0.64
Santiago ***	19	134	6.04
Shanghai	489	7710	14.78
Stockholm	242	1791	6.40
Shenzhen	431	2522	4.85
Tokyo	280	2026	6.24
Tel-aviv	55	549	8.90
Toronto	221	7659	33.70
Taiwan	487	656	0.35
Wiener Borse	60	1273	20.27
Warsaw	100	200	1.00

**TABLE 3**

**The effect of AT on market quality**

This table reports AT's effect on market quality using panel regression, with co-location as instrument variable for AT. In the 1st stage, AT is regressed on co-location dummy. In the 2nd stage, each measure of market quality is regressed on predicted AT ( $AT^*$ ) and controls. Market quality includes liquidity, efficiency and volatility. Panel A presents the effect of AT on liquidity measures including time-weighted quoted spread (RQS), trade-weighted relative effective spread (RES), permanent price impact (RPI), and temporary price impact (RRS). Panel B presents the effect of AT on efficiency measures including the absolute value of intraday autocorrelations  $|AR_{\#\#}|$ , measured for quote-midpoint returns over 10-minute and 30-minute periods. Panel C presents the effect of AT on volatility measures including daily price range standardized by the closing price,  $\text{Ln}(\text{Ret}_{\#\#\_Var})$ , log value of the daily variances of 10-minute and 30-minute quote midpoint returns, and  $|\text{Ret}|$ , the absolute value of return. Panel D presents the effect of AT on execution shortfall calculated following Anand et al. (2012). AT is the negative of dollar trading volume (\$100) per message where messages include all inside quote changes and trade messages. Control variables include daily share turnover, price range, inverse 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  $|AR_{30}|$ . The estimation is performed for the 12 months before and 12 months after the co-location in the 22 markets that adopted this service. All continuous variables are standardized to have zero mean and unit variance within the estimation window for each firm. Standard errors are clustered by date. The table reports coefficients on predicted AT. \*, \*\*, \*\*\* indicate significance at 10%, 5% and 1%, respectively.

Panel A. Liquidity		
	Coef.	t-stat
RQS	-0.220**	-2.35
RES	-0.705***	-13.93
RPI	-0.341***	-10.33
RRS	0.091***	2.98
Panel B. Efficiency		
	Coef.	t-stat
$ AR_{10} $	-0.320***	-3.45
$ AR_{30} $	-0.227***	-2.45
Panel C. Volatility		
	Coef.	t-stat
Price Range	1.246***	4.73
$\text{LnRet}_{10\_Var}$	1.564***	6.01
$\text{LnRet}_{30\_Var}$	1.438***	5.61
$ \text{Ret} $	0.385**	2.41
Panel D. Execution shortfall		
	Coef.	t-stat
Shortfall	-0.116**	-2.44

**TABLE 4**

**The effect of AT on market quality by size**

This table reports how AT's effect on market quality varies with size, using a panel regression with co-location as instrument variable for AT. In the 1st stage, we regress AT on co-location dummy, and in the 2nd stage, each measure of market quality is regressed on predicted AT ( $AT^*$ ), its interaction with three size dummies S, M, and L, and controls. Market quality includes liquidity, efficiency and volatility. Panel A presents the effect of AT on liquidity measures including time-weighted quoted spread (RQS), trade-weighted relative effective spread (RES), permanent price impact (RPI), and temporary price impact (RRS). Panel B presents the effect of AT on efficiency measures including the absolute value of intraday autocorrelations  $|AR\#\#|$ , measured for quote-midpoint returns over 10-minute and 30-minute periods. Panel C presents the effect of AT on volatility measures including daily price range standardized by the closing price,  $\text{Ln}(\text{Ret}\#\#\_Var)$ , log value of the daily variances of 10-minute and 30-minute quote midpoint returns, and  $|Ret|$ , the absolute value of return. Panel D presents the effect of AT on execution shortfall calculated following Anand et al. (2012). AT is the negative of dollar trading volume (\$100) per message where messages include all inside quote changes and trade messages. Control variables include daily share turnover, price range, inverse 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  $|AR30|$ . The estimation is performed for the 12 months before and 12 months after the co-location in the 22 markets that already adopted co-location. All continuous variables are standardized to have zero mean and unit variance within the estimation window for each firm. Standard errors are clustered by date. \*, \*\*, \*\*\* indicate significance at 10%, 5% and 1%, respectively.

Panel A. Liquidity

		Coef.	t-stat
RQS	$AT^* \times S$	-0.212***	-3.19
	$AT^* \times M$	-0.147	-1.61
	$AT^* \times L$	-0.288**	-2.36
RES	$AT^* \times S$	-0.479***	-10.02
	$AT^* \times M$	-0.532***	-10.25
	$AT^* \times L$	-0.998***	-15.47
RPI	$AT^* \times S$	-0.390***	-11.60
	$AT^* \times M$	-0.419***	-12.20
	$AT^* \times L$	-0.244***	-5.99
RRS	$AT^* \times S$	0.080**	2.30
	$AT^* \times M$	0.249***	6.75
	$AT^* \times L$	-0.038	-0.97

Panel B. Efficiency

		Coef.	t-stat
$ AR10 $	$AT^* \times S$	-0.069	-0.65
	$AT^* \times M$	-0.497***	-4.79
	$AT^* \times L$	-0.322***	-2.76
$ AR30 $	$AT^* \times S$	0.054	0.52
	$AT^* \times M$	-0.324***	-3.22
	$AT^* \times L$	-0.301**	-2.37

Panel C. Volatility

		Coef.	t-stat
--	--	-------	--------



Price Range	$AT^* \times S$	1.150***	5.03
	$AT^* \times M$	1.059***	4.35
	$AT^* \times L$	1.512***	4.84
LnRet10_Var	$AT^* \times S$	1.879***	8.35
	$AT^* \times M$	1.320***	5.34
	$AT^* \times L$	1.499***	4.78
LnRet30_Var	$AT^* \times S$	1.642***	7.33
	$AT^* \times M$	1.287***	5.33
	$AT^* \times L$	1.330***	4.30
Ret	$AT^* \times S$	0.692***	4.22
	$AT^* \times M$	0.150	0.88
	$AT^* \times L$	0.159	0.89

---

Panel D. Execution shortfall

		Coef.	t-stat
Shortfall	$AT^* \times S$	-0.023	-0.13
	$AT^* \times M$	-0.132**	-2.12
	$AT^* \times L$	-0.126***	-3.04

---

**TABLE 5**

**Correlation between AT coefficients in liquidity and volatility regressions**

We first estimate cross-sectional regressions each day with each market to obtain AT 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). Efficiency measures are daily observations of the absolute value of intraday autocorrelations  $|AR_{\#\#}|$ , measured for quote-midpoint returns over 10-minute and 30-minute periods. Volatility measures include daily price range standardized by the closing price,  $\text{Ln}(\text{Ret}_{\#\#\_Var})$ , log value of the daily variances of 10-minute and 30-minute quote midpoint returns, and  $|\text{Ret}|$ , the absolute value of return. AT is the negative of dollar trading volume (\$100) per message where messages include all inside quote changes and trade messages. Control variables include daily share turnover, price range, inverse 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  $|AR_{30}|$ . All continuous variables are standardized every day to have zero mean and unit variance 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.

Panel A. correlation of AT coefficients from the regressions of volatility and liquidity

	Price Range		$\text{Ln}(\text{Ret}_{10\_Var})$		$\text{Ln}(\text{Ret}_{30\_Var})$		$ \text{Ret} $	
RQS	0.04	***	0.10	***	0.08	***	0.01	*
RES	0.08	***	0.14	***	0.11	***	0.06	***

Panel B. correlation of AT coefficients from the regressions of volatility and efficiency

	Price Range		$\text{Ln}(\text{Ret}_{10\_Var})$		$\text{Ln}(\text{Ret}_{30\_Var})$		$ \text{Ret} $	
$ AR_{10} $	-0.10	***	0.04	**	-0.09	***	-0.04	***
$ AR_{30} $	-0.05	***	0.01		0.01		-0.02	***

**TABLE 6**

**The effect of AT on volatility on news days**

This table reports how AT's effect on volatility varies with news, using a panel regression with co-location as instrument variable for AT. In the 1st stage, we regress AT on co-location dummy, and in the 2nd stage, each measure of volatility is regressed on predicted  $AT(AT^*)$ , its interaction with news dummy, news, and controls. Volatility measures include daily price range standardized by the closing price,  $\text{Ln}(\text{Ret}_{\#\#\_Var})$ , log value of the daily variances of 10-minute and 30-minute quote midpoint returns, and  $|\text{Ret}|$ , the absolute value of return. AT is the negative of dollar trading volume (\$100) per message where messages include all inside quote changes and trade messages. Control variables include daily share turnover, inverse price, log market cap, RES and  $|\text{AR30}|$ , and the first lag of the dependent variable, all measured at t-1. The estimation is performed for the 12 months before and 12 months after the co-location in the 22 markets that already adopted co-location. All continuous variables are standardized to have zero mean and unit variance within the estimation window for each firm. Standard errors are clustered by date. \*, \*\*, \*\*\* indicate significance at 10%, 5% and 1%, respectively.

		Coef.	t-stat
Price Range	$AT^*$	1.171***	5.10
	$AT^* \times \text{News}$	1.524**	2.29
	<i>News</i>	0.079***	4.31
LnRet10_Var	$AT^*$	1.392***	5.91
	$AT^* \times \text{News}$	2.292***	4.10
	<i>News</i>	0.053***	3.56
LnRet30_Var	$AT^*$	1.262***	5.44
	$AT^* \times \text{News}$	1.954***	3.87
	<i>News</i>	0.044***	3.23
Ret	$AT^*$	0.291*	1.84
	$AT^* \times \text{News}$	-0.354	-1.07
	<i>News</i>	0.120***	13.53

**TABLE 7**

**The effect of AT on market quality at the aggregate level**

This table reports the effect of AT on market quality at the market level. 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 an instrumental variable approach. We use co-location as an instrument for algorithmic trading, a market-specific dummy that switches on once that market officially or at least publicly starts co-location service. Panel A examines liquidity measures which include time-weighted quoted spread (RQS), trade-weighted relative effective spread (RES), permanent price impact (RPI), and temporary price impact (RRS). Panel B examines efficiency measures which are daily observations of the absolute value of intraday autocorrelations  $|AR_{\#}|$ , measured for quote-midpoint returns over 10-minute and 30-minute periods. Panel C examines volatility measures which include daily price range standardized by the closing price,  $\ln(Ret_{\#\_Var})$ , log value of the daily variances of 10-minute and 30-minute quote midpoint returns, and  $|Ret|$ , the absolute value of return. Panel D examines execution shortfall calculated following Anand et al. (2012). AT is the negative of dollar trading volume (\$100) per message where messages include all inside quote changes and trade messages. Control variables include daily share turnover, intraday price range, inverse 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  $|AR30|$ . The sample period is from 2005 to 2011 to maintain a balanced panel where all 42 markets are present in the data. \*, \*\*, \*\*\* indicate significance at 10%, 5% and 1%, respectively.

Panel A. Liquidity

	Coef.	t-stat
RQS	-0.023***	-4.06
RES	-0.045***	-7.12
RPI	-0.010	-1.30
RRS	-0.097***	-11.07

Panel B. Efficiency

	Coef.	t-stat
$ AR10 $	-0.041***	-4.00
$ AR30 $	0.010	0.99

Panel C. Volatility

	Coef.	t-stat
Price Range	0.060***	9.99
$\ln(Ret10\_Var)$	0.076***	15.62
$\ln(Ret30\_Var)$	0.093***	16.79
$ Ret $	0.036***	5.11

Panel D. Execution shortfall

	Coef.	t-stat
Shortfall	-0.024**	-1.95