Does Algorithmic Trading Improve Liquidity?

TERRENCE HENDERSHOTT, CHARLES M. JONES, AND ALBERT J. MENKVELD*

ABSTRACT

Algorithmic trading has sharply increased over the past decade. Does it improve market quality, and should it be encouraged? We provide the first analysis of this question. The NYSE automated quote dissemination in 2003, and we use this change in market structure that increases algorithmic trading as an exogenous instrument to measure the causal effect of algorithmic trading on liquidity. For large stocks in particular, algorithmic trading narrows spreads, reduces adverse selection, and reduces trade-related price discovery. The findings indicate that algorithmic trading improves liquidity and enhances the informativeness of quotes.

*Hendershott is at Haas School of Business, University of California Berkeley. Jones is at Columbia Business School. Menkveld is at VU University Amsterdam. We thank Mark van Achter, Hank Bessembinder, Bruno Biais, Alex Boulatov, Thierry Foucault, Maureen O'Hara, Sébastien Pouget, Patrik Sandas, Kumar Venkataraman, the Nasdaq Economic Advisory Board, and seminar participants at the University of Amsterdam, Babson College, Bank of Canada, CFTC, HEC Paris, IDEI Toulouse, Southern Methodist University, University of Miami, the 2007 MTS Conference, NYSE, the 2008 NYU-Courant algorithmic trading conference, University of Utah, the 2008 Western Finance Association, and Yale University. We thank the NYSE for providing system order data. Hendershott gratefully acknowledges support from the National Science Foundation, the Net Institute, the Ewing Marion Kauffman Foundation, and the Lester Center for Entrepreneurship and Innovation at the Haas School at UC Berkeley; Menkveld gratefully acknowledges the College van Bestuur of VU University Amsterdam for a VU talent grant.

Technological change has revolutionized the way financial assets are traded. Every step of the trading process, from order entry to trading venue to back office, is now highly automated, dramatically reducing the costs incurred by intermediaries. By reducing the frictions and costs of trading, technology has the potential to enable more efficient risk sharing, facilitate hedging, improve liquidity, and make prices more efficient. This could ultimately reduce the cost of capital for firms.

Algorithmic trading (AT) is a dramatic example of this far-reaching technological change. Many market participants now employ AT, commonly defined as the use of computer algorithms to automatically make certain trading decisions, submit orders, and manage those orders after submission. From a starting point near zero in the mid-1990's, AT is thought to be responsible for as much as 73% of trading volume in the U.S in 2009.[1]

There are many different algorithms, used by many different types of market participants. Some hedge funds and broker-dealers supply liquidity using algorithms, competing with designated market-makers and other liquidity suppliers. For assets that trade on multiple venues, liquidity demanders often use smart order routers to determine where to send an order (e.g., Foucault and Menkveld (2008)). Statistical arbitrage funds use computers to quickly process large amounts of information contained in the order flow and price moves in various securities, trading at high frequency based on patterns in the data. Last but not least, algorithms are used by institutional investors to trade large quantities of stock gradually over time.

Before algorithmic trading took hold, a pension fund manager who wanted to buy 30,000 shares of IBM might hire a broker-dealer to search for a counterparty to execute the entire quantity at once in a block trade. Alternatively, that institutional investor might have hired a New York Stock Exchange (NYSE) floor broker to go stand at the IBM post and quietly "work" the order, using his judgment and discretion to buy a little bit here and there over the course of the trading day to keep from driving the IBM share price up too far. As trading became more electronic, it became easier and cheaper to replicate that floor trader with a computer program doing algorithmic trading (see Hendershott and Moulton (2009) for evidence on the decline in NYSE floor broker activity).

Now virtually every large broker-dealer offers a suite of algorithms to its institutional customers to help them execute orders in a single stock, in pairs of stocks, or in baskets of stocks. Algorithms typically determine the timing, price, quantity, and routing of orders, dynamically monitoring market conditions across different securities and trading venues, reducing market impact by optimally and sometimes randomly breaking large orders into smaller pieces, and closely tracking benchmarks such as the volume-weighted average price (VWAP) over the execution interval. As they pursue a desired position, these algorithms often use a mix of active and passive strategies, employing both limit orders and marketable orders. Thus, at times they function as liquidity demanders, and at times they supply liquidity.

Some observers use the term algorithmic trading to refer only to the gradual accumulation or disposition of shares by institutions (e.g., Domowitz and Yegerman (2005)). We have a broader view of algorithmic trading, including in our definition all participants who use algorithms to submit and cancel orders. We note that algorithms are also used by quantitative fund managers and others to determine portfolio holdings and formulate trading strategies, but we focus on the execution aspect of algorithms, because our data reflect counts of actual orders submitted and cancelled.

The rise of AT has obvious direct impacts. For example, the intense activity generated by algorithms threatens to overwhelm exchanges and market data providers,[2] forcing significant upgrades to their infrastructures. But researchers, regulators, and policymakers should be keenly interested in the broader implications of this sea change in trading. Overall, does AT have salutary effects on market quality, and should it be encouraged? We provide the first empirical analysis of this question.

As AT has grown rapidly since the mid nineties, liquidity in world equity markets has also dramatically improved. Based on these two coincident trends, it is tempting to conclude that algorithmic trading is at least partially responsible. But it is not at all obvious *a priori* that AT and liquidity should be positively related. If algorithms are cheaper and/or better at supplying liquidity, then AT may result in more competition in liquidity provision, thereby lowering the cost of immediacy. However, the effects could go the other way if algorithms are used mainly to demand liquidity. Limit order submitters grant a trading option to others, and if algorithms make liquidity demanders better able to identify and pick off an in-the-money trading option, then the cost of providing the trading option increases, and spreads must widen to compensate. In fact, AT could actually lead to an unproductive arms race, where liquidity suppliers and liquidity demanders both invest in better algorithms to try to take advantage of the other side, with measured liquidity the unintended victim.

In this paper, we investigate the empirical relationship between algorithmic trading and liquidity. We use a normalized measure of NYSE electronic message traffic as a proxy for algorithmic trading. This message traffic includes electronic order submissions, cancellations, and trade reports. Because we normalize by trading volume, variation in our AT measure is mainly driven by variation in limit order submissions and cancellations. This means that our measure is mainly picking up variation in algorithmic liquidity supply. This liquidity supply is likely coming both from proprietary traders making markets algorithmically and from buy-side institutions that are submitting limit orders as part of "slice and dice" algorithms.

We first examine the growth of AT and the improvements in liquidity over a five-year period. As AT grows, liquidity improves. While AT and liquidity move in the same direction, it is certainly possible that the relationship is not causal. To establish causality we study an important exogenous event that increases the amount of algorithmic trading in some stocks but not others. We use the start of autoquoting on the NYSE as an instrument for algorithmic trading. Previously, specialists were responsible for manually disseminating the inside quote. This was replaced in early 2003 by a new automated quote whenever there was a change to the NYSE limit order book. This market structure provides quicker feedback to algorithms and results in more electronic message traffic. The change was also phased in for different stocks at different times, and we take advantage of this non-synchronicity to cleanly identify causal effects.

We find that algorithmic trading does in fact improve liquidity for large-cap stocks.

Quoted and effective spreads narrow under autoquote. The narrower spreads are a result of a sharp decline in adverse selection, or equivalently a decrease in the amount of price discovery associated with trades. Algorithmic trading increases the amount of price discovery that occurs without trading, implying that quotes become more informative. There are no significant effects for smaller-cap stocks, but our instrument is weaker there, so the problem may be a lack of statistical power.

Surprisingly, we find that algorithmic trading increases realized spreads and other measures of liquidity supplier revenues. This is surprising because we initially expected that if AT improved liquidity, the mechanism would be competition between liquidity providers. However, the evidence clearly indicates that liquidity suppliers are capturing some of the surplus for themselves. The most natural explanation is that, at least during the introduction of autoquote, algorithms had market power. Over a longer time period liquidity supplier revenues decline, suggesting that any market power was temporary, perhaps because new algorithms require considerable investment and time-to-build.

The paper proceeds as follows. Section I discusses related literature. Section II describes our measures of liquidity and algorithmic trading and discusses the need for an instrumental variables approach. Section III provides a summary of the NYSE's staggered introduction of autoquote in 2003. Section IV examines the impact of AT on liquidity. Section V explores the sources of the liquidity improvement. Section VI studies AT's relation to price discovery via trading and quote updating. Section VII discusses and interprets the results, and Section VIII concludes.

I. Related literature

There are only a few papers that address algorithmic trading directly. For example, Engle, Russell, and Ferstenberg (2007) use execution data from Morgan Stanley algorithms to study the effects on trading costs of changing algorithm aggressiveness. Domowitz and Yegerman (2005) study execution costs of ITG buy-side clients, comparing results from different algorithm providers. Chaboud et al. (2009) study AT in the foreign exchange market and focus on its relation to volatility, while Hendershott and Riordan (2009) measure the contributions of AT to price discovery on the Deutsche Boerse.

Several strands of literature touch related topics. Most models take the traditional view that one set of traders provides liquidity via quotes or limit orders and another set of traders initiates a trade to take that liquidity – for either informational or liquidity/hedging reasons. Many assume that liquidity suppliers are perfectly competitive, e.g., Glosten (1994). Glosten (1989) models a monopolistic liquidity supplier, while Biais, Martimort, and Rochet (2000) model competing liquidity suppliers and find that their rents decline as the number increases. Our initial expectation is that AT facilitates the entry of additional liquidity suppliers, increasing competition.

The development and adoption of AT also involves strategic considerations. While algorithms have low marginal costs, there may be substantial development costs, and it may be costly to optimize the algorithms' parameters for each security. The need to recover these costs should lead to the adoption of algorithmic trading at times and in securities where the returns to adoption are highest (see Reinganum (1989) for a review of innovation and technology adoption).

As discussed briefly in the introduction, liquidity supply involves posting firm commitments to trade. These standing orders provide free trading options to other traders. Using standard option pricing techniques, Copeland and Galai (1983) value the cost of the option granted by liquidity suppliers. Foucault, Roëll, and Sandas (2003) study the equilibrium level of effort that liquidity suppliers should expend in monitoring the market to reduce this risk. Black (1995) proposes to minimize picking-off risk with a new limit order type that is indexed to the overall market. Algorithms can efficiently implement this kind of monitoring and adjustment of limit orders.[3] If AT reduces the cost of the free trading option implicit in limit orders, then measures of adverse selection depend on AT. If some users of AT are better at avoiding being picked off, they can impose adverse selection costs on other liquidity suppliers as in Rock (1990) and even drive out other liquidity suppliers.

AT may also be used by traders who are trying to passively accumulate or liquidate a large position.[4] There are optimal dynamic execution strategies for such traders. For example, Bertsimas and Lo (1998) find that, in the presence of temporary price impacts and a trade completion deadline, orders are optimally broken into pieces so as to minimize cost.[5] Many brokers build models with such considerations into the AT products that they sell to their clients.

II. Data

We start by characterizing the time-series evolution of algorithmic trading and liquidity for a sample of NYSE stocks over the five years from February 2001 through December 2005.

We limit ourselves to the post-decimalization regime because the change to a one penny minimum tick was a structural break that substantially altered the entire trading landscape, including liquidity metrics and order submission strategies. We end in 2005 because there are substantial NYSE market structure changes shortly thereafter.

We start with a sample of all NYSE common stocks that can be matched in both the Trades and Quotes (*TAQ*) and *CRSP* databases. To maintain a balanced panel, we retain the stocks that are present throughout the whole sample period. Stocks with an average share price of less than $5 are removed from the sample, as are stocks with an average share price of more than $1,000. The resulting sample consists of monthly observations for 943 common stocks. The balanced panel eliminates compositional changes in the sample over time. It could induce some survivorship effects if disappearing stocks are less liquid. This could overstate time-series improvements in liquidity, although the same liquidity patterns are present without a survivorship requirement (see Comerton-Forde et al. (2010)).

Stocks are sorted into quintiles based on market capitalization. Quintile 1 refers to large-cap stocks and quintile 5 corresponds to small-cap stocks. All variables used in the analysis are 99.9% winsorized: values smaller than the 0.05% quantile are set equal to that quantile, and values larger than the 99.95% quantile are set equal to that quantile.

A. Proxies for algorithmic trading

We cannot directly observe whether a particular order is generated by a computer algorithm. For cost and speed reasons, most algorithms do not rely on human intermediaries but instead generate orders that are sent electronically to a trading venue. Thus, we use the rate of electronic message traffic as a proxy for the amount of algorithmic trading taking place.[6] This proxy is commonly used by market participants, including consultants Aite Group and Tabb Group, as well as exchanges and other market venues.[7]

For example, in discussing market venue capacity limits following an episode of heavy trading volume in February 2007, a *Securities Industry News* report quotes Nasdaq SVP of transaction services Brian Hyndman, who noted that exchanges have dealt with massive increases in message traffic over the past five to six years, coinciding with algorithmic growth.

> "It used to be one-to-one," Hyndman said. "Then you'd see a customer send ten orders that would result in only one execution. That's because the black box would cancel a lot of the orders. We've seen that rise from 20- to 30- to 50-to-one. The amount of orders in the marketplace increased exponentially."[8]

In the case of the NYSE, electronic message traffic includes order submissions, cancellations, and trade reports that are handled by the NYSE's SuperDOT system and captured in the NYSE's System Order Data (*SOD*) database. The electronic message traffic measure for the NYSE excludes all specialist quoting, as well as all orders that are sent manually to the floor and are handled by a floor broker.

[insert Figure 1]

As suggested by the quote above, an important issue is whether and how to normalize the message traffic numbers. The top half of Figure 1 shows the evolution of message traffic over time. We focus on the largest-cap quintile of stocks, as they constitute the vast bulk of stock market capitalization and trading activity. Immediately after decimalization at the start of 2001, the average large-cap stock sees about 35 messages per minute during the trading day. There are a few bumps along the way, but by the end of 2005, there are an average of about 250 messages per minute (more than 4 messages per second) for these same large-cap stocks. We could, of course, simply use the raw message traffic numbers, but there has been an increase in trading volume over the same interval, and without normalization a raw message traffic measure may just be capturing the increase in trading rather than the change in the nature of trading. Therefore, for each stock each month we calculate our our algorithmic trading proxy, $algo\_trad_{it}$, as the number of electronic messages per \$100 of trading volume.[9] The normalized measure still rises rapidly over the five-year sample, while measures of market liquidity such as proportional spreads have declined sharply but appear to asymptote near the end of the sample (see, for example, the average quoted spreads in the top half of Figure 2), which occurs as more and more stocks are quoted with the minimum spread of \$0.01.

The time-series evolution of $algo\_trad_{it}$ is displayed in the bottom half of Figure 1. For the largest-cap quintile, there is about \$7,000 of trading volume per electronic message at the beginning of the sample in 2001, decreasing dramatically to about \$1,100 of trading volume per electronic message by the end of 2005. Over time, smaller-cap stocks display similar time-series patterns.

It is worth noting that our algorithmic trading proxies may also capture changes in trading strategies. For example, messages and $algo\_trad_{it}$ will increase if the same market participants use algorithms but modify their trading or execution strategies so that those algorithms submit and cancel orders more often. Similarly, the measure will increase if existing algorithms are modified to "slice and dice" large orders into smaller pieces. This is useful, as we want to capture increases in the intensity of order submissions and cancellations by existing algorithmic traders, as well as the increase in the fraction of market participants employing algorithms in trading.

B. Liquidity Measures

We measure liquidity using quoted half-spreads, effective half-spreads, 5-minute and 30-minute realized spreads, and 5-minute and 30-minute price impacts, all of which are measured as share-weighted averages and expressed in basis points as a proportion of the prevailing midpoint. The effective spread is the difference between the midpoint of the bid and ask quotes and the actual transaction price. The wider the effective spread, the less liquid is the stock. For the NYSE, effective spreads are more meaningful than quoted spreads because specialists and floor brokers are sometimes willing to trade at prices within the quoted bid and ask prices. For the t^{th} trade in stock j, the proportional effective half-spread, $espread_{jt}$, is defined as:

$$espread_{jt} = q_{jt}(p_{jt} - m_{jt})/m_{jt}, \tag{1}$$

where q_{jt} is an indicator variable that equals $+1$ for buyer-initiated trades and -1 for seller-initiated trades, p_{jt} is the trade price, and m_{jt} is the quote midpoint prevailing at the time of the 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). For each stock each day, we use all NYSE trades and quotes to calculate quoted and effective spreads for each reported transaction and calculate a share-weighted average across all trades that day. For each month we calculate the simple average across days. We also measure share-weighted quoted depth at the time of each transaction in thousands of dollars.

[insert Figure 2]

Figure 2 shows quite clearly that our measures of liquidity are generally improving over this time period. Figure 1 shows that algorithmic trading increases almost monotonically. The spread measures are not nearly as monotonic, with illiquidity spikes in both 2001 and 2002 that correspond to sharp stock market declines and increased volatility over the same time period (see Figure IA-5 in the Internet Appendix). Nevertheless, one is tempted to conclude that these two trends are related. The analysis to come investigates exactly this relationship using formal econometric tools.

If spreads narrow when algorithmic trading increases, it is natural to decompose the spread along the lines of Glosten (1987) to determine whether the narrower spread means less revenue for liquidity providers, smaller gross losses due to informed liquidity demanders, or both. We estimate revenue to liquidity providers using the 5-minute realized spread, which assumes the liquidity provider is able to close her position at the quote midpoint 5 minutes after the trade. The proportional realized spread for the t^{th} transaction in stock j is defined as:

$$rspread_{jt} = q_{jt}(p_{jt} - m_{j,t+5\text{min}})/m_{jt}, \tag{2}$$

where p_{jt} is the trade price, q_{jt} is the buy-sell indicator (+1 for buys, −1 for sells), m_{jt} is the midpoint prevailing at the time of the t^{th} trade, and $m_{j,t+5\text{min}}$ is the quote midpoint five minutes after the t^{th} trade. The 30-minute realized spread is calculated analogously using the quote midpoint 30 minutes after the trade.

We measure gross losses to liquidity demanders due to adverse selection using the 5-minute price impact of a trade, $adv\_selection_{jt}$, defined using the same variables as:

$$adv\_selection_{jt} = q_{jt}(m_{j,t+5\text{min}} - m_{jt})/m_{jt}. \tag{3}$$

The 30-minute price impact is calculated analogously. Note that there is an arithmetic identity relating the realized spread, the adverse selection (price impact), and the effective spread $espread_{jt}$:

$$espread_{jt} = rspread_{jt} + adv\_selection_{jt}. \tag{4}$$

Figure 3 graphs the decomposition of the two spread components. Both realized spreads, $rspread_{it}$), and price impacts, $adv\_selection_{it}$, decline from 2001 to 2005. Most of the narrowed spread is due to a decline in adverse selection losses to liquidity demanders. Depending on the size quintile being studied, 75% to 90% of the narrowed spread is due to a smaller price impact.

[insert Figure 3]

So far, the graphical evidence shows time-series associations between algorithmic trading and liquidity. The natural way to formally test this association is by regressing the various liquidity measures, L_{it}, on algorithmic trading A_{it} and variables controlling for market conditions X_{it}:

$$L_{it} = \alpha_i + \beta A_{it} + \delta' X_{it} + \varepsilon_{it}. \tag{5}$$

The problem is that algorithmic trading is an endogenous choice made by traders. A trader's decision to adopt AT could depend on many factors, including liquidity. For example, the evidence in Goldstein and Kavajecz (2004) indicates that humans are used more often when markets are illiquid and volatile. Econometrically, this means that the slope coefficient β from estimating equation (5) via OLS is not an unbiased estimate of the causal effect of a change in algorithmic trading on liquidity. Unless we have a structural model, the only way to identify the causal effect is to find an instrumental variable that affects algorithmic trading but is uncorrelated with ε_{it}. Standard econometrics texts, e.g., Greene (2007, §12), show that under these conditions, the resulting IV estimator consistently estimates the causal effect, in this case the effect of an exogenous change in algorithmic trading on liquidity. We discuss such an instrument in the next section.

III. Autoquote

In this section we provide an overview of our instrument, which is a change in NYSE market structure that causes an exogenous increase in algorithmic trading.

As a result of the reduction of the minimum tick to a penny in early 2001 as part of decimalization, the depth at the inside quote shrank dramatically. In response the NYSE proposed that a "liquidity quote" for each stock be displayed along with the best bid and offer. The NYSE liquidity quote was designed to provide a firm bid and offer for substantial size, typically at least 15,000 shares, accessible immediately.[10]

At the time of the liquidity quote proposal, specialists were responsible for manually disseminating the inside quote.[11] Clerks at the specialist posts on the floor of the exchange were typing rapidly and continuously from open to close and still were barely keeping up with order matching, trade reporting, and quote updating. In order to ease this capacity constraint and free up specialists and clerks to manage a liquidity quote, the exchange proposed to "autoquote" the inside quote, disseminating a new quote automatically whenever there was a relevant change to the limit order book. This would happen when a better-priced order arrived, when an order at the inside was canceled, when the inside quote was traded with in whole or in part, or when the size of the inside quote changed.

Note that the specialist's structural advantages were otherwise unaffected by autoquote. A specialist could still disseminate a manual quote at any time in order to reflect his own trading interest or that of floor traders. Specialists continued to execute most trades manually, and they could still participate in those trades subject to the unchanged NYSE rules. NYSE market share remains unchanged at about 80% around the adoption of autoquote.

[insert Figure 4]

Autoquote was an important innovation for algorithmic traders, because an automated quote update could provide more immediate feedback about the potential terms of trade. This speedup of a few seconds would provide critical new information to algorithms, but would be unlikely to directly affect the trading behavior of slower reacting humans. Autoquote allowed algorithmic liquidity suppliers to, say, quickly notice an abnormally wide inside quote and provide liquidity accordingly via a limit order. Algorithmic liquidity demanders could quickly access this quote via a conventional market or marketable limit order or by using the NYSE's automated execution facility for limit orders of 1,099 shares or less. In the next section, we show that autoquote is positively correlated with our algorithmic trading measure, which is one of the requirements for autoquote to be a valid instrument.

The NYSE began to phase in the autoquote software on January 29, 2003, starting with six active, large-cap stocks. During the next two months, over 200 additional stocks were phased in at various dates, and all remaining NYSE stocks were phased in on May 27, 2003.[12] Figure 4 provides some additional details on the phase-in process. The rollout order was determined in late 2002. Early stocks tended to be active large-cap stocks, because the NYSE felt that these stocks would benefit most from the liquidity quote. Beyond that criterion, conversations with those involved at the NYSE indicate that early phase-in stocks were chosen mainly because the specialist assigned to that stock was receptive to new technology.

The phase-in is particularly important to our empirical design. It allows us to take out all market-wide changes in liquidity, and identify the causal effect of algorithmic trading by comparing autoquoted stocks to non-autoquoted stocks using a difference-in-differences methodology. The IV methodology discussed below incorporates data before and after each NYSE stock's autoquote adoption so the estimated effect of algorithmic trading on liquidity incorporates every stock's autoquote transition, whenever it occurs. Thus, even if the phase-in order is determined by other unknown criteria, our empirical methodology remains valid in most cases. For example, there is no bias if the phase-in is determined by the specialist's receptiveness to new technology, and this is correlated with the amount of algorithmic trading in his stocks. There are a small number of problematic phase-in scenarios, however, and we discuss these next.

For the staggered introduction of autoquote to serve as a valid instrument, it must satisfy the exclusion restriction. Specifically, a stock's move to autoquote must not be correlated with the error term in that firm's liquidity equation (equation (5)). This does not mean that the autoquote rollout must be assigned randomly. The liquidity equation includes a firm fixed effect, calendar dummies, and a set of control variables. The instrument remains valid even if the rollout schedule is related to these particular explanatory variables. For instance, if the stocks chosen for early phase-in tend to have high mean liquidity, this would be picked up by the firm fixed effect and the exclusion restriction would still hold. In fact, due to the explanatory variables, the exclusion restriction is violated only if the autoquote phase-in schedule is somehow related to contemporaneous changes in firm-specific, idiosyncratic liquidity that are not due to changes in AT.

Thus, it is quite helpful that the rollout schedule for autoquote was fixed months in advance, as it seems highly unlikely that the phase-in schedule could be correlated with idiosyncratic liquidity months into the future. The only way this might happen is if there are sufficiently persistent but temporary shocks to idiosyncratic liquidity. For example, if temporarily illiquid stocks are chosen for early phase-in, these stocks might still be illiquid when autoquote begins, and their liquidity would improve post-autoquote as they revert to mean liquidity, thereby overstating the causal effect.[13] To investigate this, we study the dynamics of liquidity using an AR(1) model of effective spreads for each firm in the sample. Table IA-3 of the Internet Appendix shows that the average AR(1) coefficient is 0.18, corresponding to a half-life of less than a day.[14] We also do not find any statistical support for the conjecture that stocks that migrate experience unusual liquidity just ahead of the migration. More precisely, the predicted effective spread based on all information up until the day before the introduction, including liquidity covariates, is not significantly different from its unconditional mean. All of this supports the exogeneity of our instrument.

Lastly, the exclusion restriction requires autoquote to affect liquidity only via algorithmic trading. We have argued that autoquote's time scale is only relevant for algorithms and autoquote does not directly affect liquidity via nonalgorithmic trading.[15] However, we cannot test for this using the available data. Thus, it is important to emphasize that our conclusions on causality rely on the intuitively appealing but ultimately untestable assumption that autoquote affects liquidity only via its effect on algorithmic trading.

IV. AT's Impact on Liquidity

To study the effects of autoquote, we build a daily panel of NYSE common stocks. The sample begins on December 2, 2002, which is approximately two months before the autoquote phase-in begins, and it extends through July 31, 2003, about two months after the last batch of NYSE stocks moves to the autoquote regime. We use standard price filters: stocks with an average share price of less than $5 or more than $1,000 are removed. To make our various autoquote analyses comparable, we use the same sample of stocks throughout this section. The Hasbrouck (1991a, 1991b) decomposition (discussed below in section VI) has the most severe data requirements, so we retain all stocks that have at least 21 trades per day for each day in the eight-month sample period. This leaves 1,082 stocks in the sample. The shorter time period for the autoquote sample allows for a larger balanced panel compared to the five-year balanced panel used to create Figures 1-3.

[insert Table I]

Stocks are then sorted into quintiles based on market capitalization. Quintile 1 refers to large-cap stocks and quintile 5 corresponds to small-cap stocks. Table I contains means by quintile and standard deviations for all of the variables used in the analysis. We measure liquidity using quoted half-spreads, effective half-spreads, 5-minute and 30-minute realized spreads, and 5-minute and 30-minute price impacts, all of which are measured as share-weighted averages and expressed in basis points as a proportion of the prevailing midpoint. All variables used in the analysis are 99.9% winsorized: values smaller than the 0.05% quantile are set equal to that quantile, and values larger than the 99.95% quantile are set equal to that quantile.

[insert Table II]

Autoquote clearly leads to greater use of algorithms. Figure 1 shows that message traffic increases by about 50% in the most active quintile of stocks as autoquote is phased in; it is certainly hard to imagine that autoquote would change the behavior of humans by anything close to this magnitude. But nowhere in the paper do we rely on this time-series increase in AT. Instead, we always include stock fixed effects and time fixed effects (day dummies), so that we identify the effect of the market structure change via its staggered introduction. The presence of these two-way fixed effects means we are always comparing the changes experienced by autoquoted stocks to the changes in not-yet-autoquoted control stocks.

We begin by estimating the following first-stage regression:

$$M_{it} = \alpha_i + \gamma_t + \beta Q_{it} + \varepsilon_{it} \tag{6}$$

where M_{it} is the relevant dependent variable, e.g., the number of electronic messages per minute, Q_{it} is the autoquote dummy set to zero before the autoquote introduction and one afterward, α_i is a stock fixed effect, and γ_t is a day dummy. There are also separate regressions for each size quintile.

Table II reports the slope coefficients for this specification. When the dependent variable M_{it} is the number of electronic messages per minute for stock i on day t, we find a significant positive relationship. The coefficient of 2.135 on autoquote implies that autoquote increases message traffic by an average of two messages per minute. In December 2002, the month before autoquote begins its rollout, our sample stocks average 36 messages per minute, so autoquote causes a 6% increase in message traffic on average. Associations are stronger for large-cap stocks, consistent with the conventional wisdom that algorithmic trading was more prevalent at the time for active, liquid stocks.

Table II also shows that there is a significant positive relationship between the autoquote dummy and our preferred measure of algorithmic trading $algo\_trad_{it}$, which is the negative of dollar volume in hundreds per electronic message. Thus, it is clear that autoquote leads to more algorithmic trading in all but the smallest quintiles. [16] There is no consistent relationship between autoquote and any other variable, such as turnover, volatility, and share price.

Our principal goal is to understand the effects of algorithmic liquidity supply on market quality, and so we use the autoquote dummy as an instrument for algorithmic trading in a panel regression framework. Our main instrumental variables specification is a daily panel of 1,082 NYSE stocks over the eight-month sample period spanning the staggered implementation of autoquote. The dependent variable is one of five liquidity measures: quoted half-spreads, effective half-spreads, realized spreads, or price impacts, all of which are share-volume weighted and measured in basis points, or the quoted depth in thousands of dollars. We have fixed effects for each stock as well as time dummies, and we include share turnover, volatility based on the daily price range (high minus low, see Parkinson (1980)), the inverse of share price, and the log of market cap as control variables. Results are virtually identical if we exclude these control variables. Based on anecdotal information that algorithmic trading was relatively more important for active large-cap stocks during this time period, we estimate this specification separately for each market-cap quintile.

The estimated equation is:

$$L_{it} = \alpha_i + \gamma_t + \beta A_{it} + \delta' X_{it} + \varepsilon_{it} \tag{7}$$

where L_{it} is a spread measure for stock i on day t, A_{it} is the algorithmic trading measure

$algo\_trad_{it}$, and X_{it} is a vector of control variables, including share turnover, volatility, the inverse of share price, and log market cap. We always include fixed effects and time dummies. The set of instruments consists of all explanatory variables, except that we replace $algo\_trad_{it}$ with $auto\_quote_{it}$. Inference is based on standard errors that are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (see Arellano and Bond (1991)). Section 2 of the Internet Appendix shows that the IV regression is unaffected by the use of a proxy for algorithmic trading, as long as the noise in the proxy is uncorrelated with the autoquote instrument.

[insert Table III]

The results are reported in Panel A of Table III and the most reliable effects are in larger stocks. For large-cap stocks (quintiles 1 and 2), the autoquote instrument shows that an increase in algorithmic liquidity supply narrows both the quoted and effective spread. To interpret the estimated coefficient on the algorithmic trading variable, recall that the algorithmic trading measure $algo\_trad_{it}$ is the negative of dollar volume per electronic message, measured in hundreds of dollars, while the spread is measured in basis points. Thus, the IV estimate of -0.53 on the algorithmic trading variable for quintile 1 means that a unit increase in algorithmic trading, e.g., from the sample mean of $1,844 to $1,744 of volume per message, implies that quoted spreads narrow by 0.53 basis points.[17] The average within-stock standard deviation for $algo\_trad_{it}$ is 4.54 or $454, so a one standard deviation change in our algorithmic trading measure is associated with a 4.54*0.53=2.41 basis point change in proportional spreads. This represents nearly a 50 percent decline from the mean quoted spread of 5.19 basis points for quintile 1.

In the spirit of an event study, we also estimate an analogous non-IV panel regression with the autoquote dummy directly on the right-hand side. We do not report these results, but quoted and effective spreads are reliably narrower for the three largest quintiles. For quintile 1, quoted spreads are 0.50 basis points smaller (t = -9.18) after the autoquote introduction, and effective spreads are 0.17 basis points smaller (t = -4.33). Effective spreads narrow even more for quintiles 2 and 3, 0.21 and 0.23 basis points, respectively.

The IV estimate on $algo\_trad_{it}$ is statistically indistinguishable from zero for quintiles 3 through 5. This could be a statistical power issue. Figure 4 shows that most small-cap stocks were phased-in at the very end, reducing the non-synchronicity needed for econometric identification. Perhaps as a result, the autoquote instrument is only weakly correlated with algorithmic trading in these quintiles. Alternatively, it could be that algorithms are less commonly used in these smaller stocks, in which case the introduction of autoquote might have little or no effect on these stocks' market quality.

Quoted depth also declines with autoquote. One might worry that the narrower quoted spread simply reflects the smaller quoted quantity, casting doubt on whether liquidity actually improves after autoquote is introduced. Here, a calibration exercise is useful. The results for quintile 1 indicate that a one-unit increase in the algorithmic trading variable—a $100 decrease in trading volume per message—reduces the quoted spread by 10%, as the average quoted spread from Table I is 5.19 basis points. The same change reduces the quoted depth by about 5%, based on an average quoted depth of $71,220. A small liquidity demander is unaffected by the depth reduction and is unambiguously better off with the narrower spread. A liquidity demander who trades the average quoted depth of $71,220 is probably better off as well. She pays 10% less on 95% of her order, and as long as she pays less than 290% of the original spread on the remaining 5% of her order, she is better off overall. Based on the $40.01 average share price for this quintile, the average 5.19 basis point quoted spread translates to 2.1 cents. For these stocks, it seems extremely unlikely that the last 5% of her trade executes at a spread of more than 6.1 cents. Most likely this last 5% would execute only one cent wider. This makes it quite clear that the depth reduction is small relative to the narrowing of the spread, at least for these trade sizes.

To further explore the decline in depth, in Panel B of Table III we check to see whether autoquote is associated with narrower effective spreads after controlling for trade size. Specifically, we sort trades into bins based on trade size, calculate effective spreads for these trades alone, and estimate the IV specification for each trade size bin. For the larger quintiles, AT significantly narrows the effective spread for all trade sizes below 5,000 shares. Point estimates go in the same direction for the largest trade sizes but are not reliably different from zero. In any case, we only have transaction-level data, not order-level data, so it is not possible to conduct the analysis on large orders, given that large orders are typically broken up into many smaller pieces and executed over time. If more AT means large orders are broken up more, then even if the effective spread narrows for individual trades, the aggregate price impact of the entire order could increase.

V. Sources of Liquidity Improvement

As discussed earlier in the paper, narrower effective spreads imply either less revenue per trade for liquidity providers, smaller gross losses due to informed liquidity demanders, or both. In Table III Panel C, we decompose effective spreads into a realized spread component and an adverse selection or price impact component in order to understand the source or sources of the improvement in liquidity under autoquote. These spread components are calculated using a 5-minute and a 30-minute horizon, and the IV regressions are repeated using each component of the spread.

The results are somewhat surprising. For large and medium-cap stocks, quintiles 1 through 3, the realized spread actually increases significantly after autoquote, indicating that liquidity providers are earning greater net revenues. These greater revenues are offset by a larger decline in price impacts, implying that liquidity providers are losing far less to liquidity demanders after the introduction of autoquote. As before, nothing is significant for the two smallest-cap quintiles. Panel C also shows that the results are the same at both 5-minute and 30-minute horizons.

We describe these results as surprising because they do not match our priors going into the analysis. We thought that if autoquote improved liquidity, it would be because algorithmic liquidity suppliers were low-cost providers who suddenly became better able to compete with the specialist and the floor under autoquote, and thereby improving overall liquidity by reducing aggregate liquidity provider revenues. Instead, it appears that liquidity providers in aggregate were able to capture some of the surplus created by autoquote.

However, this liquidity supplier market power reflected in larger realized spreads appears to have been fleeting. Figure 3 shows that while liquidity provider revenues increased during the autoquote introduction in the first half of 2003, realized spreads decline during the second half of 2003. Existing algorithms could have found themselves with a distinct competitive advantage in response to the increased information flow, given that new algorithms take considerable time and expense to build and test. Thus, one can interpret the data as a temporary increase in algorithms' market power under autoquote that disappears as new entrants developed competing algorithms.

[insert Table IV]

Which liquidity providers benefit? We do not have any trade-by-trade data on the identity of our liquidity providers, but we do know specialist participation rates for each stock each day, so we can see whether autoquote changed the specialist's liquidity provision market share. We conduct an IV regression with the specialist participation rate on the left-hand side, and the results in Table IV confirm that, at least for the large-cap quintile of stocks, specialists appear to participate less under autoquote, suggesting that it is other liquidity providers who capture the surplus created by autoquote.

Table IV also puts a number of other non-spread variables on the LHS of the IV specification. The most interesting is trade size. At least for the two largest quintiles, the autoquote instrument confirms most observers' strong suspicions that the increase in algorithmic trading is one of the causes of smaller average trade sizes in recent years.[18]

The decomposition of the effective spread used above has the advantage of being simple, but it also has distinct disadvantages. In particular, it chooses an arbitrary time point in the future, five minutes and 30 minutes in this case, and implicitly ignores other trades that might have happened in that five or 30 minute time period. Lin, Sanger, and Booth (1995) develop a spread decomposition model that is estimated trade by trade and accounts for order flow persistence (the empirical fact, first noted by Hasbrouck and Ho (1987), that buyer-initiated trades tend to follow buyer-initiated trades).[19] Results in Section 3 and Table IA-6 of the Internet Appendix show that our spread decomposition results continue to hold using their approach.

VI. AT and Price Discovery: Trades versus Quotes

In this section, we investigate whether algorithmic trading changes the nature of price discovery. We use the framework of Hasbrouck (1991a, 1991b), who introduces a VAR-based model that makes almost no structural assumptions about the nature of information or order flow, but instead infers the nature of information and trading from the observed sequence of prices and orders. In this framework, all stock price moves end up assigned to one of two categories: they are either associated or unassociated with a recent trade. Though the model does not make any structural assumptions about the nature of information, we usually refer to price moves as private information-based if they are associated with a recent trade. Price moves that are orthogonal to recent trade arrivals are sometimes considered based on "public information," as in Jones, Kaul, and Lipson (1994) and Barclay and Hendershott (2003)).

To separate price moves into trade-related and trade-unrelated components, we construct a VAR with two equations: the first describes the trade-by-trade evolution of the quote midpoint, while the second equation describes the persistence of order flow. Continuing our earlier notation, define q_{jt} to be the buy-sell indicator for trade t in stock j (+1 for buys, -1 for sells), and define r_{jt} to be the log return based on the quote midpoint of stock j from trade $t-1$ to trade t. The VAR picks up order flow dependence out to 10 lags:

$$r_t = \sum_{i=1}^{10} \alpha_i r_{t-i} + \sum_{i=0}^{10} \beta_i q_{t-i} + \varepsilon_{rt}, \tag{8}$$

$$q_t = \sum_{i=1}^{10} \gamma_i r_{t-i} + \sum_{i=1}^{10} \phi_i q_{t-i} + \varepsilon_{qt}, \tag{9}$$

where the stock subscripts j are suppressed from here on. The VAR is inverted to get the vector moving average (VMA) representation:

$$y_t = \begin{bmatrix} r_t \\ q_t \end{bmatrix} = \theta(L)\varepsilon_t = \begin{bmatrix} a(L) & b(L) \\ c(L) & d(L) \end{bmatrix} \begin{bmatrix} \varepsilon_{rt} \\ \varepsilon_{qt} \end{bmatrix}, \tag{10}$$

where $a(L), b(L), c(L)$, and $d(L)$ are lag polynomial operators. The permanent effect on price of an innovation is given by $a(L)\varepsilon_{rt} + b(L)\varepsilon_{qt}$, and because we include contemporaneous q_t in the return equation, $cov(\varepsilon_{rt}, \varepsilon_{qt}) = 0$ and the variance of this random-walk component can be written as:

$$\sigma_w^2 = (\sum_{i=0}^{\infty} a_i)^2 \sigma_r^2 + (\sum_{i=0}^{\infty} b_i)^2 \sigma_q^2. \tag{11}$$

The second term captures the component of price discovery that is related to recent trades, and the first term captures price changes that are orthogonal to trading. The VAR is estimated in transaction time; we follow Hasbrouck (1991b) and multiply the resulting variances by the number of trades each day to get a decomposition of the daily variance of log changes in the efficient price.

The information content of individual trades can also be calculated using the VAR. As shown in Hasbrouck (1991a) the cumulative response of the quoted price to a one-time unit shock in the order flow equation is a measure of adverse selection that accounts for the persistence in order flow as well as possible positive or negative feedback trading. The cumulative impulse response is easiest to see in the VMA representation and is equal to $(\sum_{i=0}^{\infty} b_i)$. As discussed in Hasbrouck (1991a, 1991b), the VAR approach is robust to price discreteness, lagged adjustment to information, and lagged adjustment to trades.

[insert Figure 5]

The VAR, the cumulative impulse response, and the trade-related and non-trade-related standard deviations are estimated for each stock each day. As in the prior graphs, we calculate monthly averages for each quintile and graph these in Figure 5. The most striking feature of the graphs is the decline in the cumulative impulse response and the trade-related standard deviation during the first half of 2003, while the non-trade-related standard deviations do not change much as autoquote is introduced. This suggests that under autoquote much more information is being incorporated into prices without trade, consistent with Boulatov and George (2007), who show that mean-squared error in the quote midpoint is smaller when informed traders place limit orders, compared to a world where liquidity is only provided by perfectly competitive, uninformed, risk-neutral market makers.

While these time-series effects appear large, again we prefer to identify the effect using the staggered autoquote instrument. The IV panel regression is estimated first with the daily trade-related standard deviation as the dependent variable. We then repeat using the non-trade-related standard deviation on the left-hand side. The panel regressions continue to include stock fixed effects, calendar dummies, and the same set of control variables.

[insert Table V]

The results can be found in Table V, and at least for the two largest quintiles they confirm the results from the time-series graphs. When a large-cap stock adopts autoquote and experiences an exogenous increase in algorithmic trading, the impulse response measure of adverse selection decreases. In addition, there is much less trade-correlated price discovery, and much more price discovery that is uncorrelated with trading. Consistent with other methodologies, we do not find consistently reliable effects for the smaller-cap quintiles.

For the largest quintiles, algorithmic trading has an economically important effect on the nature of price discovery. During the autoquote sample period, the within standard deviation in our algorithmic trading variable is 4.54, so a one standard deviation increase in algorithmic trading during this sample period leads to an estimated change in trade-correlated price discovery equal to a $4.54 * 0.22 = 1.00$ percentage point reduction in the daily standard deviation of trade-correlated returns for the largest-cap stocks. Figure 5 shows that this is the same order of magnitude as the actual level of trade-correlated standard deviations, so this is indeed a substantial change in how prices are updated to reflect new information over time.

VII. Discussion and interpretation

Why is the nature of price discovery changing? It seems likely that algorithms respond quickly to order flow and price information, updating their limit orders to prevent them from becoming stale and being picked off. We suspect that two kinds of information are of first-order importance, though it is hard to be sure, given the opacity of algorithm creators and providers. First, we think algorithms can easily take into account common factor price information and adjust trading and quoting accordingly. For example, if there is an upward shock to the S&P futures price, an algorithmic liquidity supplier in IBM that currently represents the inside offer may decide to cancel its existing sell order before it is picked off by an index arbitrageur or another trade, replacing the sell order with a higher-priced ask. Shocks to other stocks in the same industry could cause similar reactions from algorithms. Second, some algorithms are designed to sniff out other algorithms or otherwise identify order flow and other information patterns in the data. For example, if an algorithm identifies a sequence of buys in the data and concludes that more buys are coming, an algorithmic liquidity supplier might adjust its ask price upward. Information in newswires can even be parsed electronically in order to adjust trading algorithms.[20]

To help understand the counterintuitive realized spread result, it is important to consider how cost structures differ for human and algorithmic trading. If monitoring costs are sufficiently high that humans do not always monitor the market, limit orders submitted by humans will not always reflect all public information and may become stale. While algorithms have large fixed development costs, one of their main advantages is that there is virtually no marginal cost in monitoring public information and adjusting their orders or quotes. Thus, an increase in algorithmic activity causes more changes in the efficient price to be revealed through a quote update rather than via trade. To develop this intuition further, Section 4 of the Internet Appendix develops a very simple generalized Roll model that is a slight variation on one developed in Hasbrouck (2007).

If the algorithmic traders have more market power than the non-algorithmic traders, an increase in algorithmic trading leads to larger imputed revenue to liquidity suppliers (larger realized spread). Is this market power argument plausible? As autoquote was implemented in 2003, the extant algorithms might have found themselves with a distinct competitive advantage in trading in response to the increased information flow, given that new algorithms take considerable time and expense to build and test. While we view the differential market power argument as possible, we do not have any additional evidence to support it. Over the longer run, liquidity supplier revenues decline (see Figure 3), suggesting that any market power was temporary.

Ideally, we would also directly analyze AT behavior to better understand the specific AT order submission and trading strategies that lead to these results. Unfortunately, no U.S. stock exchange has yet been able and willing to identify AT. However, Hendershott and Riordan (2009) can accurately identify AT on the Deutsche Boerse. They find that algorithmic traders generally place more efficient quotes, and algorithms supply more liquidity when spreads are wide. While they describe equilibrium behavior and cannot measure the causal effect of AT on market quality, Hendershott and Riordan (2009) do provide a natural mechanism by which algorithmic traders would improve liquidity and quote efficiency.

VIII. Conclusions

The declining costs of technology have led to its widespread adoption throughout financial industries. The resulting technological change has revolutionized financial markets and the way financial assets are traded. Many institutions now trade via algorithms, and we study whether algorithmic trading at the NYSE improves liquidity. In the five years following decimalization, algorithmic trading has increased, and markets have become more liquid. To establish causality we use the staggered introduction of autoquoting as an instrumental variable for algorithmic trading. We demonstrate that increased algorithmic trading lowers adverse selection and decreases the amount of price discovery that is correlated with trading. Our results suggest that algorithmic trading lowers the costs of trading and increases the informativeness of quotes. Surprisingly, the revenues to liquidity suppliers also increase with algorithmic trading, though this effect appears to be temporary.

We have not studied it here, but it seems likely that algorithmic trading can also improve linkages between markets, generating positive spillover effects in these other markets. For example, when computer-driven trading is made easier, stock index futures and underlying share prices are likely to track each other more closely. Similarly, liquidity and price efficiency in equity options probably improves as the underlying share price becomes more informative.

A couple of caveats are in order, however. Our overall sample period covers a period of generally rising stock prices, and stock markets are fairly quiescent during the 2003 introduction of autoquote. While we do control for share price levels and volatility in our empirical work, it remains an open question whether algorithmic trading and algorithmic liquidity supply are equally beneficial in more turbulent or declining markets. Like Nasdaq market makers refusing to answer their phones during the 1987 stock market crash, algorithmic liquidity suppliers may simply turn off their machines when markets spike downward. With access to the right data, 2007 and 2008 stock markets could prove to be a useful laboratory for such an investigation.

A second caveat relates to trading by large institutions. Some market participants complain that the decline in depth has hampered the ability to trade large amounts without substantial costs. While the rise of algorithmic trading has contributed to the decline in depth, we are optimistic that other technological innovations can offset some of these effects. For instance, some "dark pools" such as LiquidNet and Pipeline represent a modern version of an upstairs market, allowing traders with large orders to electronically search for counterparties without revealing their trading interest (see, e.g., Bessembinder and Venkataraman (2004)).

Finally, our results have important implications for both regulators and designers of trading platforms. For example, the U.S. Securities and Exchange Commission's Regulation NMS (SEC (2005)) is designed to increase competition among liquidity suppliers. Our results highlight the importance of algorithmic liquidity suppliers and the benefits of ensuring vigorous competition between them. Of course, markets need not leave this problem to the regulator. Trading venues can attract these algorithms by lowering development and implementation costs. For example, exchanges and other trading platforms can calculate useful information and metrics to be fed into algorithms, distributing them at low cost. A market can also allow algorithmic traders to co-locate their servers in the market's data center. Finally, offering additional order types, such as pegged orders, can lessen the infrastructure pressures that algorithms impose.

Notes

[1]See "SEC runs eye over high-speed trading," *Financial Times*, July 29, 2009. The 73% is an estimate for high-frequency trading, which, as discussed below, is a subset of algorithmic trading.

[2]See "Dodgy Tickers-Stock Exchanges," *Economist*, March 10, 2007.

[3]Rosu (2009) develops a model that implicitly recognizes these technological advances and simply assumes limit orders can be constantly adjusted. Consistent with AT, Hasbrouck and Saar (2009) find that by 2004 a large number of limit orders are cancelled within two seconds on the INET trading platform.

[4]Keim and Madhavan (1995) and Chan and Lakonishok (1995) study institutional orders that are broken up.

[5] Almgren and Chriss (2000) extend this by considering the risk that arises from breaking up orders and slowly executing them. Obizhaeva and Wang (2005) optimize assuming that liquidity does not replenish immediately after it is taken but only gradually over time. For each component of the larger transaction, a trader or algorithm must choose the type and aggressiveness of the order. Cohen et al. (1981) and Harris (1998) focus on the simplest static choice: market order versus limit order. However, a limit price must be chosen, and the problem is dynamic; Goettler, Parlour, and Rajan (2009) model both aspects.

[6]See Biais and Weill (2009) for theoretical evidence on how algorithmic trading relates to message traffic.

[7]See, for example, Jonathan Keehner, "Massive surge in quotes, electronic messages may paralyse US market," `http://www.livemint.com/2007/06/14005055/Massive-surge-in-quotes-elect.html`, June 14, 2007.

[8]See Shane Kite, "Reacting to market break, NYSE and Nasdaq act on capacity," *Securities Industry News*, March 12, 2007.

[9]Our results are virtually the same when we normalize by the number of trades or use raw message traffic numbers (see Table IA-4 in the Internet Appendix). The results are also the same when we use the number of cancellations rather than the number of messages to construct the algorithmic trading measure.

[10]For more details, the NYSE proposal is contained in Securities Exchange Act Release No. 47091 (December 23, 2002), 68 FR 133.

[11]One exception: NYSE software would automatically disseminate an updated quote after 30 seconds if the specialist had not already done so.

[12] Liquidity quotes were delayed due to a property rights dispute with data vendors, so they did not become operational until June 2003, after autoquote was fully phased-in. Liquidity quotes were almost never used and were formally abandoned in July 2005.

[13]Late phase-in stocks will not offset this effect. Even if late phase-in stocks are temporarily liquid when chosen, this temporary effect has more time to die out by the time autoquote is implemented for them.

[14]Also, because the average daily AR(1) coefficient is quite small, there is little scope for the bias that can arise in dynamic panel data models with strong persistence. See, e.g., Arellano (2003, §6.2).

[15]For example, autoquote could simply make the observed quotes less stale. We investigate this possibility in Section 1 of the Internet Appendix and find that this mechanical explanation is unlikely to account for our results.

[16]In the IV regressions in Tables III-V we report F statistics that reject the null that the instruments do not enter the first stage regression. Bound, Jaeger, and Baker (1995, p.446) mention that "F statistics close to 1 should be cause for concern." Our F statistics range from 5.88 to 7.32, and we are thus not afflicted with a weak instruments problem.

[17]Table IA-4 in the Internet Appendix contains additional analysis showing that the message traffic component of $algo\_trad_{it}$ drives the decline in spreads.

[18]Turnover has negative and significant coefficients in Q2 and Q5 in Table IV. To ensure that possible endogeneity of turnover is not affecting the $algo\_trad$ coefficients in Table III, Table IA-5 in the Internet Appendix repeats Table III omitting turnover as a control variable. Including turnover does not change the results.

[19]See Barclay and Hendershott (2004) for a discussion of Lin, Sanger, and Booth (1995) vs. other spread decomposition models.

[20]See, "Ahead of the Tape-Algorithmic Trading," *Economist*, June 23, 2007.

References

Almgren, Robert, and Neil Chriss, 2000, Optimal execution of portfolio transactions, *Journal of Risk* 3, 5–39.

Arellano, Manuel, 2003, *Panel Data Econometrics* (Oxford University Press, New York).

Arellano, Manuel, and Stephen 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.

Barclay, Michael J., and Terrence Hendershott, 2003, Price discovery and after trading hours, *Review of Financial Studies* 16, 1041–1073.

Barclay, Michael J., and Terrence Hendershott, 2004, Liquidity externalities and adverse selection: evidence from trading after hours, *Journal of Finance* 59, 681–710.

Bertsimas, Dimitris, and Andrew W. Lo, 1998, Optimal control of execution costs, *Journal of Financial Markets* 1, 1–50.

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

Bessembinder, Hendrik, and Kumar Venkataraman, 2004, Does an electronic stock exchange need an upstairs market?, *Journal of Financial Economics* 73, 3–36.

Biais, Bruno, David Martimort, and Jean-Charles Rochet, 2000, Competing mechanisms in a common value environment, *Econometrica* 68, 799–837.

Biais, Bruno, and Pierre-Olivier Weill, 2009, Liquidity shocks and order book dynamics, Unpublished manuscript, Toulouse University, IDEI.

Black, Fischer, 1995, Equilibrium exchanges, *Financial Analysts Journal* 51, 23–29.

Boulatov, Alex, and Thomas J. George, 2007, Securities trading when liquidity providers are informed, Unpublished manuscript, University of Houston.

Bound, John, David A. Jaeger, and Regina M. Baker, 1995, Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak, *Journal of the American Statistical Association* 90, 443–450.

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

Chan, Louis K.C., and Josef Lakonishok, 1995, The behavior of stock prices around institutional trades, *Journal of Finance* 50, 1147–1174.

Cohen, Kalman, Steven Maier, Robert Schwartz, and David Whitcomb, 1981, Transaction costs, order placement strategy and existence of the bid-ask spread, *Journal of Political Economy* 89, 287–305.

Comerton-Forde, Carole, Terrence Hendershott, Charles M. Jones, Mark S. Seasholes, and Pamela C. Moulton, 2010, Time variation in liquidity: the role of market maker inventories and revenues, *Journal of Finance* 65, 295–331.

Copeland, Thomas E., and Dan Galai, 1983, Information effects on the bid-ask spread, *Journal of Finance* 38, 1457–1469.

Domowitz, Ian, and Henry Yegerman, 2005, The cost of algorithmic trading: a first look at comparative performance, in Brian R. Bruce, ed.: *Algorithmic Trading: Precision, Control, Execution* (Institutional Investor).

Engle, Robert F., Jeffrey R. Russell, and Robert Ferstenberg, 2007, Measuring and modeling execution cost and risk, Unpublished manuscript, New York University.

Foucault, Thierry, and Albert J. Menkveld, 2008, Competition for order flow and smart order routing systems, *Journal of Finance* 63, 119–158.

Foucault, Thierry, Ailsa Roëll, and Patrik Sandas, 2003, Market making with costly monitoring: an analysis of the SOES controversy, *Review of Financial Studies* 16, 345–384.

Glosten, Lawrence R., 1987, Components of the bid ask spread and the statistical properties of transaction prices, *Journal of Finance* 42, 1293–1307.

Glosten, Lawrence R., 1989, Insider trading, liquidity, and the role of the monopolist specialist, *Journal of Business* 62, 211–235.

Glosten, Lawrence R., 1994, Is the electronic limit order book inevitable?, *Journal of Finance* 49, 1127–1161.

Goettler, Ronald L., Christine A. Parlour, and Uday Rajan, 2009, Informed traders in limit order markets, *Journal of Financial Economics* 93, 67–87.

Goldstein, Michael A., and Kenneth A. Kavajecz, 2004, Trading strategies during circuit breakers and extreme market movements, *Journal of Financial Markets* 7, 301–333.

Greene, William H., 2007, *Econometric Analysis* (Prentice Hall, London).

Harris, Lawrence, 1998, Optimal dynamic order submission strategies in some stylized trading problems, *Financial Markets, Institutions, and Instruments* 7, 1–76.

Hasbrouck, Joel, 1991a, Measuring the information content of stock trades, *Journal of Finance* 46, 179–207.

Hasbrouck, Joel, 1991b, The summary informativeness of stock trades: an econometric analysis, *Review of Financial Studies* 4, 571–595.

Hasbrouck, Joel, 2007, *Empirical Market Microstructure* (Oxford University Press, New York).

Hasbrouck, Joel, and Thomas Ho, 1987, Order arrival, quote behavior and the return generating process, *Journal of Finance* 42, 1035–1048.

Hasbrouck, Joel, and Gideon Saar, 2009, Technology and liquidity provision: the blurring of traditional definitions, *Journal of Financial Markets* 12, 143–172.

Hendershott, Terrence, and Pamela C. Moulton, 2009, Speed and stock market quality: the NYSE's hybrid, Unpublished manuscript, University of California, Berkeley.

Hendershott, Terrence, and Ryan Riordan, 2009, Algorithmic trading and information, Unpublished manuscript, University of California, Berkeley.

Jones, Charles M., Gautam Kaul, and Marc L. Lipson, 1994, Information, trading, and volatility, *Journal of Financial Economics* 36, 127–154.

Keim, Donald B., and Ananth Madhavan, 1995, Anatomy of the trading process: empirical evidence on the behavior of institutional traders, *Journal of Financial Economics* 37, 371–398.

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

Lin, Ji-Chai, Gary C. Sanger, and G. Geoffrey Booth, 1995, Trade size and components of the bid-ask spread, *Review of Financial Studies* 8, 1153–1183.

Obizhaeva, Anna, and Jiang Wang, 2005, Optimal trading strategy and supply/demand dynamics, Unpublished manuscript, MIT.

Parkinson, Michael, 1980, The extreme value method for estimating the variance of the rate of return, *Journal of Business* 51, 61–65.

Reinganum, Jennifer F., 1989, The timing of innovation: research, development, and diffusion, in Richard Schmalensee and Robert Willig, ed.: *Handbook of Industrial Organization* (Elsevier Publishing).

Rock, Kevin, 1990, The specialist's order book and price anomalies, Unpublished manuscript, Harvard University.

Rosu, Ioanid, 2009, A dynamic model of the limit order book, *Review of Financial Studies* 22, 4601–4641.

Securities and Exchange Commission Release No. 34-51808, 2005.

Table I: Summary Statistics

This table presents summary statistics on daily data for the period December 2002 through July 2003. This period covers the phase-in of autoquote, used as an instrument in the instrumental variable analysis. The dataset combines *TAQ*, *CRSP*, and the NYSE System Order Data (*SOD*) database. The balanced panel consists of 1,082 stocks sorted into quintiles based on market capitalization, where quintile 1 contains largest-cap stocks. All variables are 99.9% winsorized.

variable	description (units)	source	mean Q1	mean Q2	mean Q3	mean Q4	mean Q5	st. dev. within[a]
$qspread_{it}$	share-volume-weighted quoted half spread (bps)	*TAQ*	5.19	6.82	9.17	11.68	19.89	4.84
$qdepth_{it}$	share-volume-weighted depth ($1,000)	*TAQ*	71.22	41.85	31.43	24.12	15.76	23.42
$espread_{it}$	share-volume-weighted effective half spread (bps)	*TAQ*	3.63	4.79	6.56	8.46	14.50	3.73
$rspread_{it}$	share-volume-weighted realized half spread, 5min (bps)	*TAQ*	1.21	1.44	1.88	1.97	4.34	4.71
$adv\_selection_{it}$	share-volume-weighted adverse selection component half spread, 5min, "effective-realized" (bps)	*TAQ*	2.42	3.35	4.69	6.50	10.16	5.12
$messages_{it}$	#electronic messages per minute i.e. proxy for algorithmic activity (/minute)	*SOD*	119.30	53.90	29.81	19.33	10.44	15.55
$algo\_trad_{it}$	dollar volume per electronic message times (-1) to proxy for algorithmic trading ($100)	*TAQ/SOD*	-18.44	-10.99	-8.05	-6.39	-4.61	4.54
$dollar\_volume_{it}$	average daily volume ($million)	*TAQ*	94.71	24.09	10.12	5.32	2.17	22.72
$trades_{it}$	#trades per minute (/minute)	*TAQ*	5.72	2.92	1.78	1.24	0.72	0.72
$share\_turnover_{it}$	(annualized) share turnover	*TAQ/CRSP*	1.11	1.52	1.48	1.45	1.30	1.16
$volatility_{it}$	standard deviation open-to-close returns based on daily price range, i.e. high minus low, Parkinson (1980), (%)	*CRSP*	1.47	1.56	1.63	1.74	2.06	0.85
$price_{it}$	daily closing price ($)	*CRSP*	40.01	32.05	25.86	23.93	16.41	3.46
$market\_cap_{it}$	shares outstanding times price ($billion)	*CRSP*	28.99	4.09	1.71	0.90	0.41	1.96
$trade\_size_{it}$	trade size ($1,000)	*TAQ*	37.56	19.41	13.06	9.73	6.61	8.03
$specialist\_particip_{it}$	specialist participation rate (%)	*SOD*	13.07	12.97	13.08	13.73	15.84	3.92
#observations: 1082*167 (stock*day)								

[a]: Based on day t's deviation relative to the time mean, i.e., $x^*_{i,t} = x_{i,t} - \overline{x}_i$.

Table II: Autoquote Impact on Messages, Algorithmic Trading Proxy, and Covariates

This table shows the impact of autoquote on other variables, and the second column can be interpreted as the first-stage instrumental variables (IV) regression when $algotrade_{it}$ is the dependent variable. The analysis is based on daily observations from December 2002 through July 2003, which covers the phase-in of autoquote. We regress each of the variables used in the IV analysis on the autoquote dummy ($auto\_quote_{it}$) using the following specification:

$$M_{it} = \alpha_i + \gamma_t + \beta Q_{it} + \varepsilon_{it}$$

where M_{it} is the relevant dependent variable, e.g., the number of electronic messages per minute, Q_{it} is the autoquote dummy set to zero before the autoquote introduction and one afterward, α_i is a stock fixed effect, and γ_t is a day dummy. There are also separate regressions for each size quintile, and statistical significance is based on standard errors that are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (see Arellano and Bond (1991)). Table I provides other variable definitions.

	$messa-$ ges_{it}	$algo\_$ $trad_{it}$	$share\_$ $turnover_{it}$	$vola$- $tility_{it}$	$1/price_{it}$	ln_mar- ket_cap_{it}
Slope coefficient from regression of column variable on $auto\_quote_{it}$						
all	2.135**	0.291**	-0.016**	0.001	0.000**	-0.003**
Q1 (largest-cap)	6.286**	0.414**	0.016*	-0.003	0.000**	-0.005**
Q2	0.880**	0.396**	-0.029*	0.007	-0.000**	0.003**
Q3	0.944**	0.292**	0.002	-0.001	0.000	-0.004**
Q4	0.223**	0.029	-0.006	-0.003	-0.000	0.002
Q5 (smallest-cap)	-0.031	0.219**	-0.080**	0.003	0.002**	-0.013**

*/**: Significant at a 95%/99% level.

Table III: Effect of Algorithmic Trading on Spread

This table regresses various measures of the (half) spread on our algorithmic trading proxy. It is based on daily observations from December 2002 through July 2003 which covers the phase-in of autoquote. The nonsynchronous autoquote introduction instruments for the endogenous $algo\_trad_{it}$ to identify causality from algorithmic trading to liquidity. The specification is:

$$L_{it} = \alpha_i + \gamma_t + \beta A_{it} + \delta X_{it} + \varepsilon_{it}$$

where L_{it} is a spread measure for stock i on day t, A_{it} is the algorithmic trading measure $algo\_trad_{it}$, and X_{it} is a vector of control variables, including share turnover, volatility, 1/price, and log market cap. Fixed effects and time dummies are always included. The set of instruments consists of all explanatory variables, except that $algo\_trad_{it}$ is replaced with $auto\_quote_{it}$. There are separate regressions for each size quintile, and t-values in parentheses are based on standard errors that are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (see Arellano and Bond (1991)).

	Coefficient on $algo\_trad_{it}$					Coefficients on control variables[a]			
	Q1	Q2	Q3	Q4	Q5	$share\_turnover_{it}$	$vola-tility_{it}$	$1/price_{it}$	$ln\_mkt\_cap_{it}$
Panel A: Quoted spread, quoted depth, and effective spread									
$qspread_{it}$	-0.53**	-0.42**	-0.43	-0.21	9.92	-2.81**	0.90**	108.40**	-3.61**
	(-3.23)	(-2.21)	(-1.44)	(-0.06)	(1.22)	(-2.98)	(9.71)	(7.42)	(-2.28)
$qdepth_{it}$	-3.49**	-1.43	-1.99	15.60	0.61	-5.22	-1.64*	-3.44	12.01
	(-2.51)	(-1.16)	(-1.07)	(0.39)	(0.19)	(-0.64)	(-1.86)	(-0.02)	(0.82)
$espread_{it}$	-0.18**	-0.32**	-0.35	-1.67	4.65	-1.01**	0.69**	72.77**	-1.30
	(-2.67)	(-2.23)	(-1.56)	(-0.42)	(1.16)	(-2.30)	(9.39)	(10.80)	(-1.46)
Panel B: Effective spread by trade size category[b]									
$espread1^b_{it}$	-0.12**	-0.14**	-0.17	-1.83	4.99	-0.83**	0.28**	50.45**	-1.10
	(-3.06)	(-2.02)	(-1.09)	(-0.44)	(1.20)	(-2.70)	(3.91)	(12.43)	(-1.61)
$espread2^b_{it}$	-0.22**	-0.30**	-0.41	-4.21	4.21	-1.62**	0.43**	53.84**	-2.24
	(-3.25)	(-2.61)	(-1.62)	(-0.45)	(1.16)	(-2.68)	(2.80)	(8.15)	(-1.58)
$espread3^b_{it}$	-0.25**	-0.26	-0.66	-3.27	6.89	-1.85**	0.64**	61.48**	-1.96
	(-2.76)	(-1.42)	(-1.57)	(-0.36)	(1.13)	(-2.67)	(4.03)	(7.55)	(-1.29)
$espread4^b_{it}$	-0.11	-0.29	-0.24	643.14	-3.27	28.99	-10.35	224.78	93.04
	(-1.31)	(-0.90)	(-0.62)	(0.00)	(-0.38)	(0.00)	(-0.00)	(0.01)	(0.00)
$espread5^b_{it}$	-0.04	-0.32	-0.36	1.45	10.48	-0.21	1.05**	73.01**	0.12
	(-0.45)	(-1.04)	(-0.86)	(0.26)	(0.25)	(-0.25)	(4.01)	(8.06)	(0.05)

<continued on next page>

<continued from previous page>

	Coefficient on $algo\_trad_{it}$					Coefficients on control variables[a]			
	Q1	Q2	Q3	Q4	Q5	$share\_turnover_{it}$	$vola-tility_{it}$	$1/price_{it}$	$ln\_mkt\_cap_{it}$
Panel C: Spread decompositions based on 5-min and 30-min price impact									
$rspread_{it}$	0.35**	0.76**	1.03**	14.25	15.88	3.13*	-1.06**	45.81**	5.05
	(3.52)	(3.97)	(2.06)	(0.46)	(1.36)	(1.92)	(-2.15)	(4.14)	(1.18)
$adv\_selection_{it}$	-0.53**	-1.07**	-1.39**	-15.51	-11.21	-4.12**	1.76**	26.65*	-6.30
	(-3.56)	(-4.08)	(-2.06)	(-0.47)	(-1.33)	(-2.23)	(3.28)	(1.84)	(-1.34)
$rspread\_30m_{it}$	0.33**	0.47*	0.91	11.11	12.63	2.69**	-2.33**	52.24**	2.83
	(2.82)	(1.94)	(1.61)	(0.47)	(1.31)	(1.98)	(-5.99)	(4.21)	(0.83)
$adv\_selection\_30m_{it}$	-0.51**	-0.81**	-1.27*	-12.60	-8.28	-3.66**	3.02**	20.21	-4.10
	(-3.43)	(-2.76)	(-1.80)	(-0.47)	(-1.25)	(-2.33)	(6.91)	(1.35)	(-1.05)

#observations: 1082*167 (stock*day)

F test statistic of hypothesis that instruments do not enter first stage regression: 7.32 ($F(5, 179587)$),
p-value: 0.0000

*/**: Significant at a 95%/99% level.

[a]: Coefficients for the control variables and time dummies are quintile-specific. For brevity, only (across the quintiles) market-cap-weighted coefficients are reported for the control variables.

[b]: The suffix indicates the effective spread for a particular trade size category, i.e.
"1" if 100 shares $\leq$ trade size $\leq$ 499 shares;
"2" if 500 shares $\leq$ trade size $\leq$ 1999 shares;
"3" if 2000 shares $\leq$ trade size $\leq$ 4999 shares;
"4" if 5000 shares $\leq$ trade size $\leq$ 9999 shares;
"5" if 9999 shares $<$ trade size.

Table IV: Effect of Algorithmic Trading on Nonspread Variables

This table regresses nonspread variables on our algorithmic trading proxy. It is based on daily observations from December 2002 through July 2003 which covers the phase-in of autoquote. The nonsynchronous autoquote introduction instruments for the endogenous $algo\_trad_{it}$ to identify causality from algorithmic trading to these nonspread variables. The specification is:

$$M_{it} = \alpha_i + \gamma_t + \beta A_{it} + \varepsilon_{it}$$

where M_{it} is a nonspread variable for stock i on day t, and A_{it} is the algorithmic trading measure. Fixed effects and time dummies are always included. There are separate regressions for each size quintile, and t-values in parentheses are based on standard errors that are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (see Arellano and Bond (1991)).

	Coefficient on $algo\_trad_{it}$				
	Q1	Q2	Q3	Q4	Q5
$share\_turnover_{it}$	0.04	-0.07*	0.01	-0.20	-0.36**
	(1.02)	(-1.77)	(0.07)	(-0.26)	(-2.89)
$trades_{it}$	0.58**	-0.01	-0.01	-0.51	-0.15**
	(2.60)	(-0.23)	(-0.15)	(-0.33)	(-2.60)
$trade\_size_{it}$	-2.04**	-0.80**	-0.33	2.27	-0.22
	(-4.64)	(-3.23)	(-0.69)	(0.20)	(-0.60)
$specialist\_particip_{it}$	-0.59**	-0.23	-0.92	-13.24	-1.89**
	(-2.22)	(-1.24)	(-1.43)	(-0.29)	(-2.02)

#observations: 1082*167 (stock*day)

F test statistic of hypothesis that instruments do not enter first stage regression: 5.88 ($F(5, 179607)$), p-value: 0.0000

*/**: Significant at a 95%/99% level.

Table V: Effect of Algorithmic Trading on Permanent Price Impact and Efficient Price Variance Composition

This table regresses the permanent price response to a trade and the two components of efficient price variance on our algorithmic trading proxy. The daily sample extends from December 2002 through July 2003 which covers the phase-in of autoquote. A Hasbrouck VAR model on midquote returns and signed trades is estimated in order to identify the long-term price impact of a trade ($impulse\_response\_hasbr_{it}$) and the trade-related ($stdev\_tradecorr\_comp_{it}$) and trade-unrelated ($stdev\_nontradecorr\_comp_{it}$) components of the daily percentage variance of changes in the efficient price (see Hasbrouck (1991a, 1991b) for details). For the regressions, the nonsynchronous introduction of autoquote is used as an instrument for the endogenous $algo\_trad_{it}$ to identify causality from algorithmic trading to these variables. The specification is:

$$M_{it} = \alpha_i + \gamma_t + \beta A_{it} + \delta X_{it} + \varepsilon_{it}$$

where M_{it} is the dependent variable for stock i on day t, A_{it} is the algorithmic trading measure, and X_{it} is a vector of control variables, including share turnover, volatility, 1/price, and log market cap. Fixed effects and time dummies are always included; control variables are excluded from the Hasbrouck component regressions. There are separate regressions for each size quintile, and t-values in parentheses are based on standard errors that are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (see Arellano and Bond (1991)).

	Coefficient on $algo\_trad_{it}$					Coefficients on control variables[a]			
	Q1	Q2	Q3	Q4	Q5	$share\_turnover_{it}$	$vola-tility_{it}$	$1/price_{it}$	$ln\_mkt\_cap_{it}$
Panel A: Effect of algorithmic trading on the long-term price impact of a trade (bps)									
$impulse\_response\_hasbr_{it}$	-0.54**	-1.21**	-1.44**	-18.92	-11.93	-4.90**	1.11*	16.66	-7.83
	(-3.50)	(-4.05)	(-2.08)	(-0.46)	(-1.36)	(-2.20)	(1.66)	(1.15)	(-1.35)
Panel B: Effect of algorithmic trading on trade- and nontrade-correlated component of daily efficient price variance (%)									
$stdev\_tradecorr\_comp_{it}$	-0.22**	-0.26**	-0.30*	-3.40	-0.57**				
	(-2.62)	(-3.08)	(-1.69)	(-0.29)	(-2.73)				
$stdev\_nontradecorr\_comp_{it}$	0.12**	0.13**	0.13	1.04	0.13				
	(2.48)	(2.36)	(1.47)	(0.28)	(1.12)				

#observations: 1082*167 (stock*day)

F test statistic of hypothesis that instruments do not enter first stage regression: Panel A: 7.32 ($F(5, 179587)$), p-value: 0.0000; Panel B: 5.88 ($F(5, 179607)$), p-value: 0.0000

*/**: Significant at a 95%/99% level.

[a]: Coefficients for the control variables and time dummies are quintile-specific. For brevity, only (across the quintiles) market-cap-weighted coefficients are reported for the control variables.

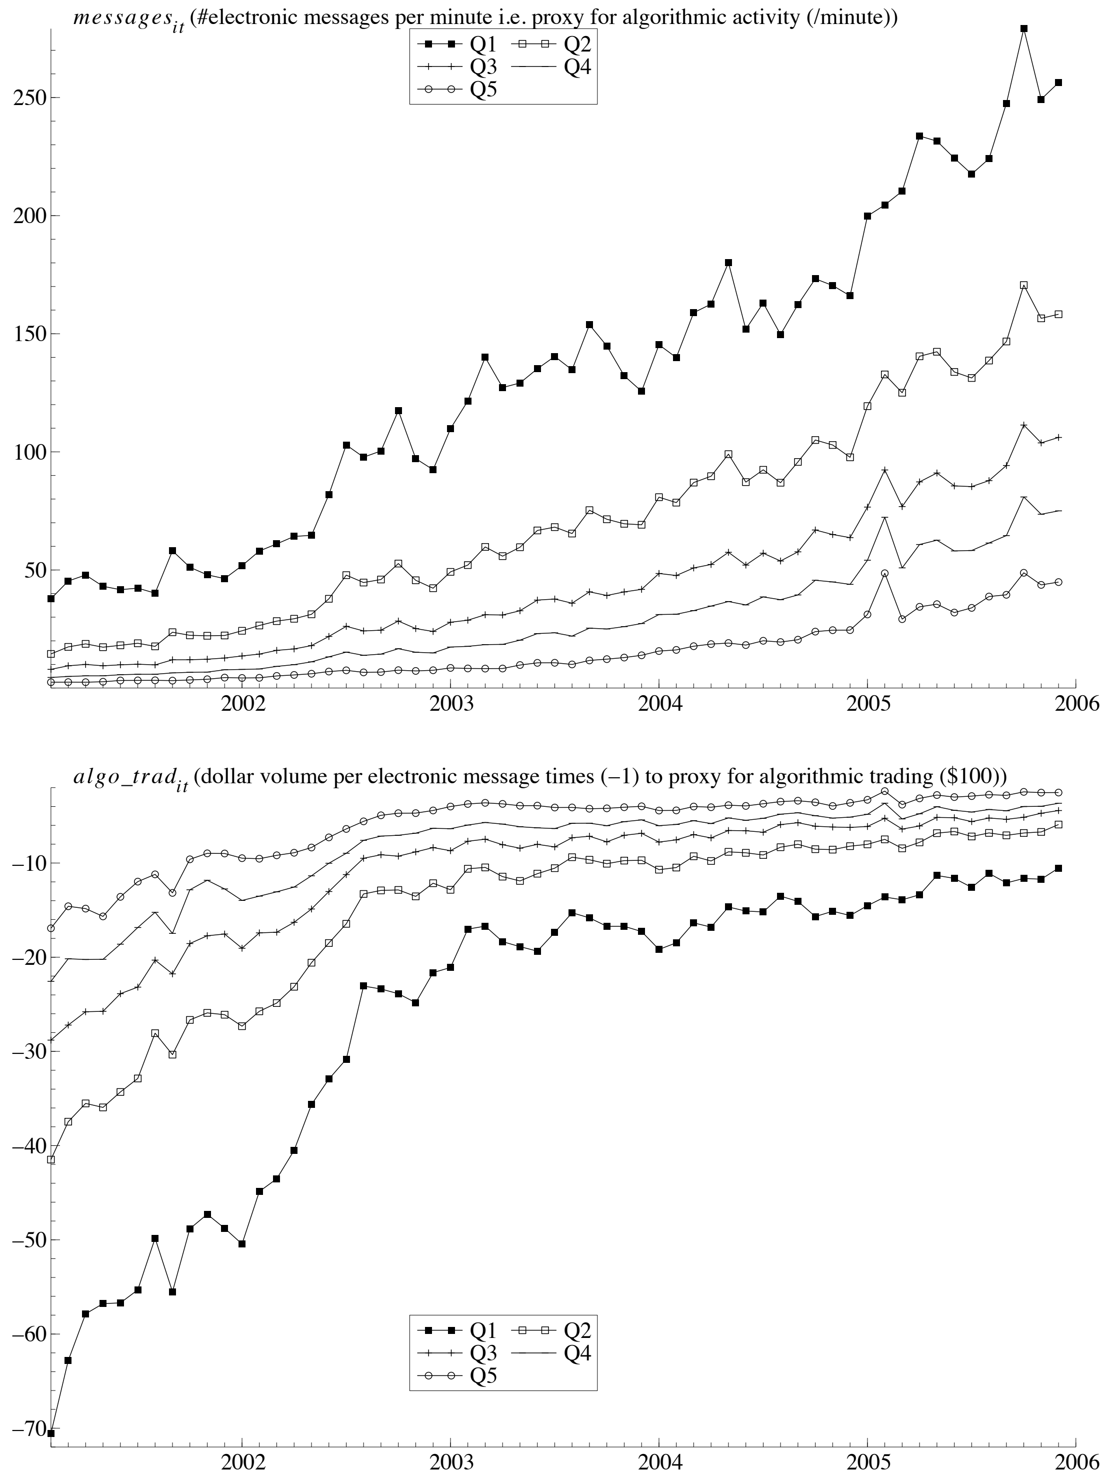

Figure 1: For each market-cap quintile, where Q1 is the largest-cap quintile, these graphs depict (i) the number of (electronic) messages per minute and (ii) our proxy for algorithmic trading, which is defined as the negative of trading volume (in hundreds of dollars) divided by the number of messages.

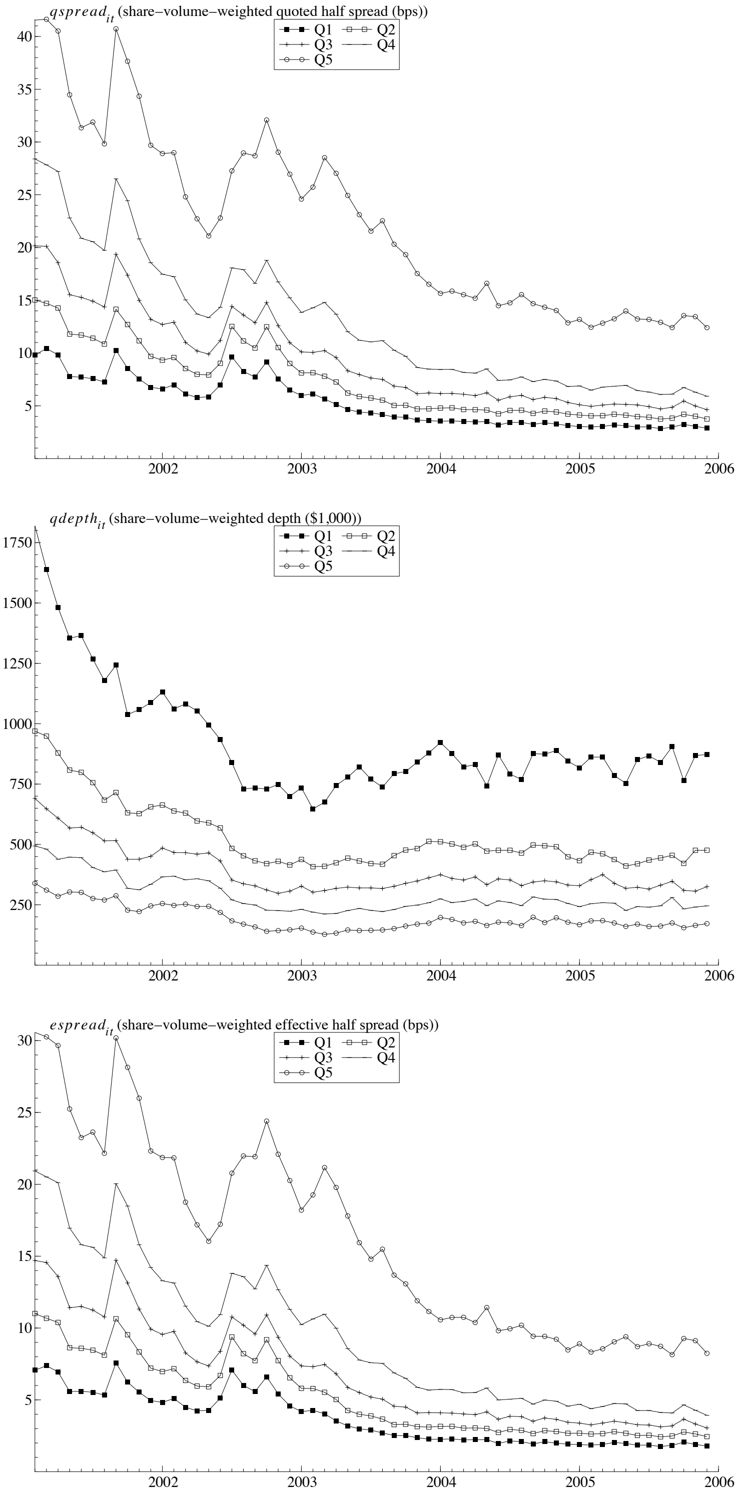

Figure 2: These graphs depict (i) quoted half spread, (ii) quoted depth, and (iii) effective spread. All spread measures are share-volume weighted averages within-firm, and then averaged across firms within each market-cap quintile, where Q1 is the largest-cap quintile.

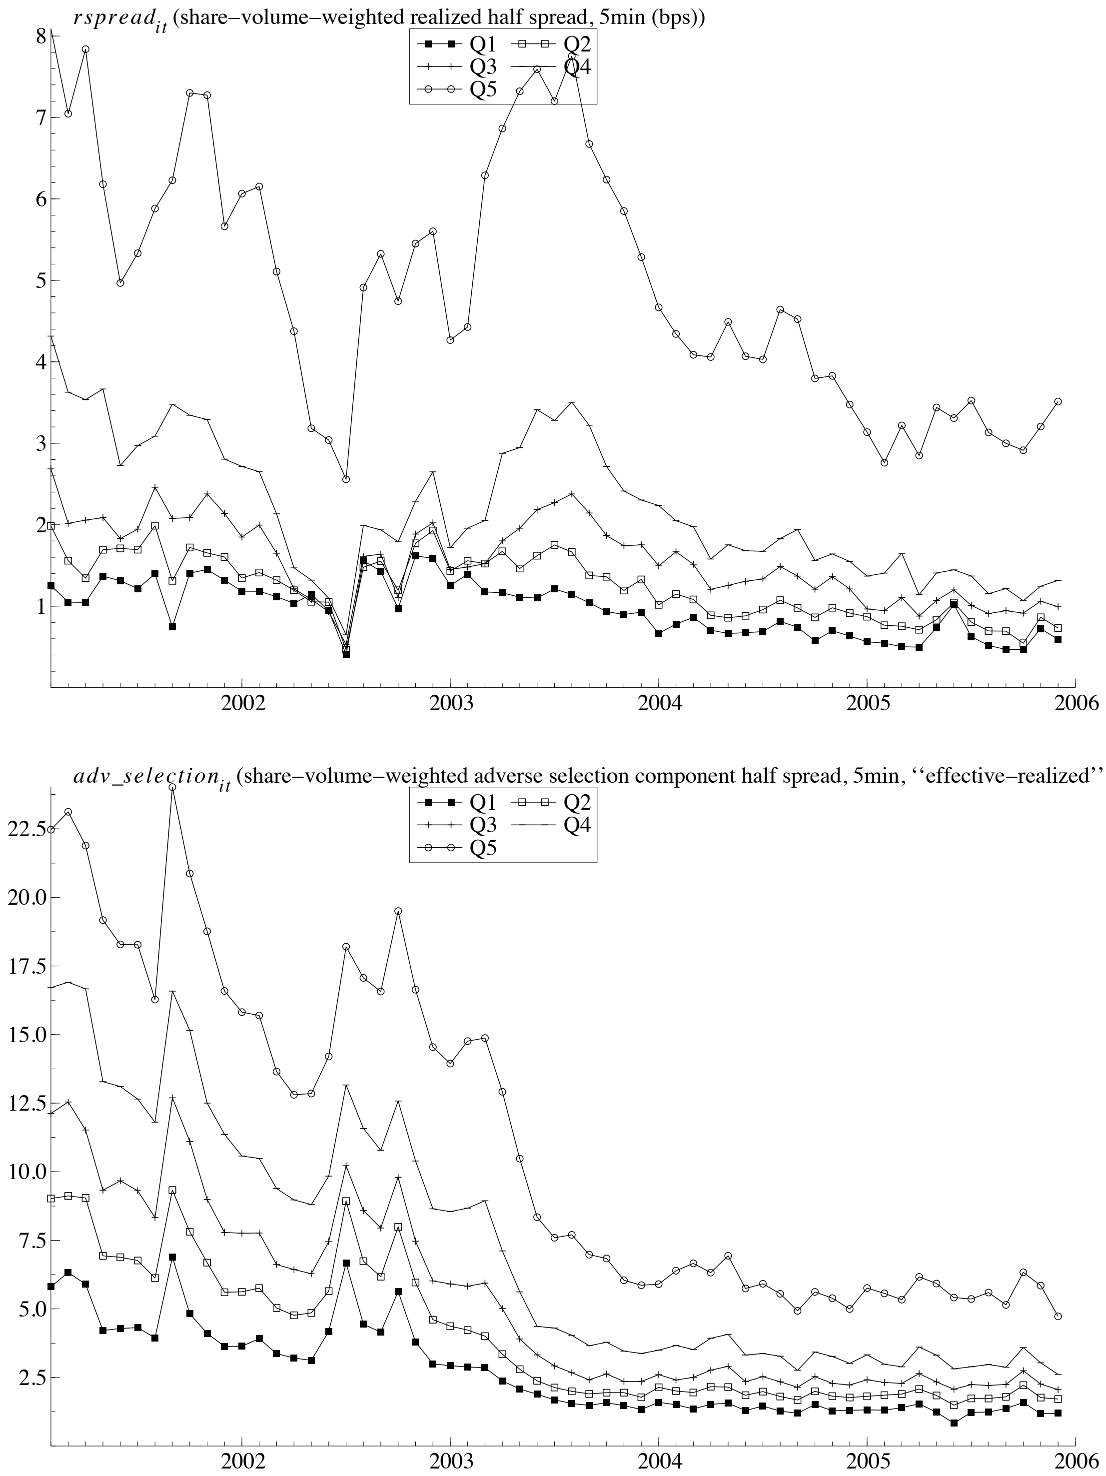

Figure 3: These graphs depict the two components of the effective spread: (i) realized spread and (ii) the adverse selection component, also known as the (permanent) price impact. Both are based on the quote midpoint 5 minutes after the trade. Results are graphed by market-cap quintile, where Q1 is the largest-cap quintile.

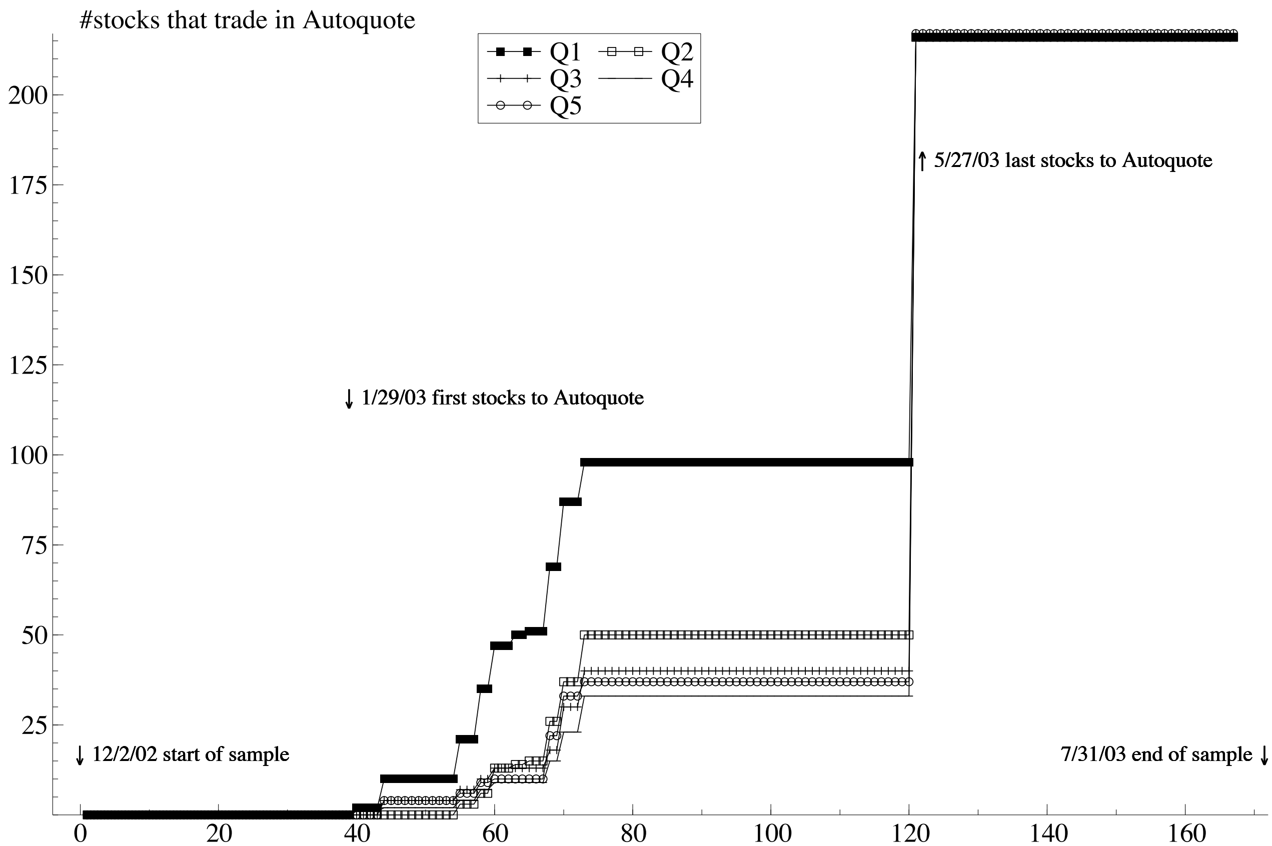

Figure 4: This graph depicts the staggered introduction of autoquote on the NYSE. It graphs the number of stocks in each market-cap quintile that are autoquoted at a given time. Quintile 1 contains largest-cap stocks.

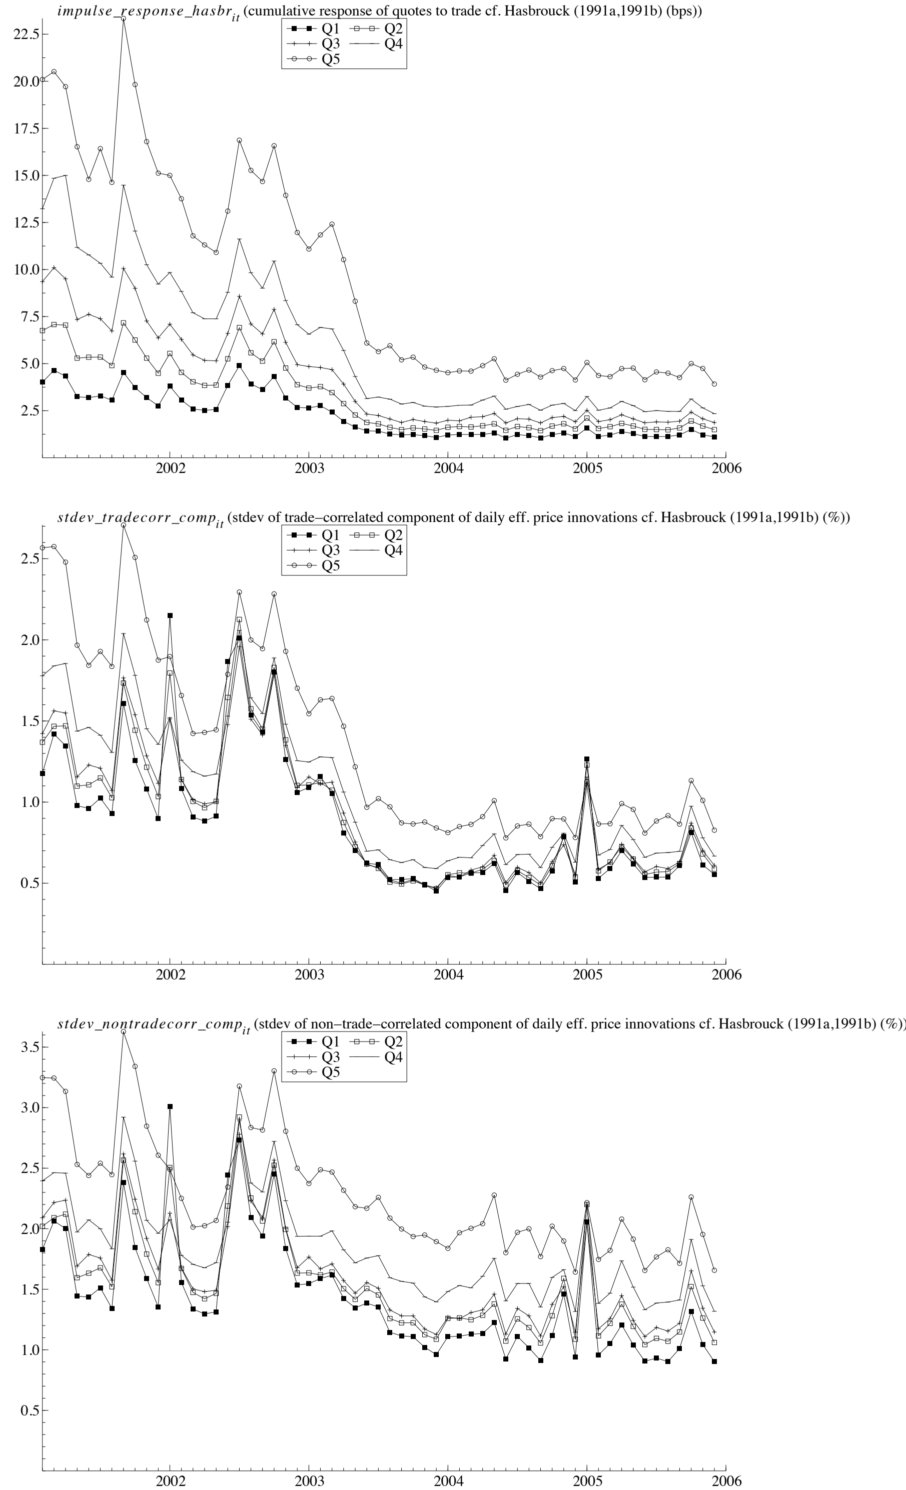

Figure 5: These graphs illustrate the estimation results of the Hasbrouck (1991a,1991b) VAR model for midquote returns and signed trades. The top graph illustrates the time series pattern of the long-term price impact of the midquote to a unit impulse in the signed trade variable. The bottom two graphs illustrate the decomposition of the daily percentage variance of changes in the efficient price into a trade-related ($stdev\_tradecorr\_comp_{it}$) and trade-unrelated ($stdev\_nontradecorr\_comp_{it}$) component (see Hasbrouck (1991a, 1991b) for details on the methodology). The graph depicts the autoquote sample period which runs from December 2002 through July 2003. Results are graphed by market-cap quintile, where Q1 is the largest-cap quintile.

Internet Appendix for "Does Algorithmic Trading Improve Liquidity?"*

This internet appendix contains the following supplementary content:

- Section I considers mechanical explanations for the autoquote results, including stale quotes and slow quote replenishment.
- Section II shows that IV estimates are consistent even if the instrument is a noisy proxy.
- Section III discusses how algorithmic trading affects the various components of the bid-ask spread based on the spread decomposition of Lin, Sanger, and Booth (1995).
- Section IV proposes a simple generalized Roll model as a framework for interpreting the empirical results.
- Table IA-I provides summary statistics (similar to Table 1 in the main text) for the five year sample (monthly from February 2001 through December 2005).

*Citation format: Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld, 2010, Internet Appendix to "Does Algorithmic Trading Improve Liquidity?" Journal of Finance [vol #], [pages], http://www.afajof.org/IA/[year].asp. Please note: Wiley-Blackwell is not responsible for the content or functionality of any supporting information supplied by the authors. Any queries (other than missing material) should be directed to the authors of the article.

- Table IA-II provides univariate correlations for the five year sample between spreads, algorithmic trading, volume, volatility, and share price.

- Table IA-III investigates the exogeneity of the timing of the autoquote introduction.

- Table IA-IV reports IV regression results using the numerator and the denominator of the algorithmic trading proxy separately as regressors.

- Table IA-V reports the IV regression results for spreads with share turnover, a potentially endogenous variable, removed from the set of covariates.

- Table IA-VI provides results for the spread decomposition proposed by Lin, Sanger, and Booth (1995).

- Figures IA-1 through IA-4 replicate figures in the main document (Figures 1,2,3, and 5, respectively), except that these figures include 95% confidence intervals.

- Figure IA-5 graphs the evolution of the non-spread variables (trade size, number of trades, volume, and volatility) over the five year sample period.

- Figure IA-6 graphs the three components of the Lin, Sanger, and Booth (1995) spread decomposition over the five year sample period.

I. Stale Quotes and Slow Quote Replenishment

In the main text, we focus on the algorithmic trading channel, but it is important to consider whether a more mechanical explanation might account for our autoquote results. What might we expect if autoquote simply makes the observed quotes less stale and has no other effects?

We start by examining what occurs when the inside quote updates are driven by the submission of better quotes or cancellations of the orders at the inside quote. Let a_t and b_t be the ask and bid prices at time t, and assume this quote is disseminated by the specialist. Limit orders arrive or are cancelled, and at a later time t', $a_{t'}$ and $b_{t'}$ are the best ask and bid prices. Assume that $a_{t'}$ and $b_{t'}$ are disseminated only after the adoption of autoquote; otherwise, the econometrician identifies a_t and b_t as the ask and bid in effect at time t'.

To simplify the exposition, assume that the ask side of the book changes ($a_t \neq a_{t'}$) while the bid side of the book remains unchanged ($b_t = b_{t'}$). Symmetric arguments apply for changes to the bid side of the book alone, and the results also hold when both the bid and the ask change between t and t'.

There are two possibilities for the change in the inside ask. If the time t inside ask is cancelled, then $a_{t'} > a_t$. If instead a new sell order arrives at time t' that would improve the inside quote, then $a_{t'} < a_t$. Overall, if cancels are more common than improvements, then prior to the adoption of autoquote the disseminated quoted spread is artificially narrow, and autoquote should be associated with a widening of quoted spreads. However, we find the reverse. Autoquote is associated with a narrowing of the quoted spread, so we focus hereafter on the arrival of new orders at time t' that improve the existing time t quote. Prior to autoquote, we continue to observe the old, wider quote (a_t, b_t) at time t'. Under autoquote, the new, narrower quote $(a_{t'}, b_t)$ is disseminated at time t'.

Let $m_{t'} = 1/2(a_{t'} + b_{t'})$ be the midquote at time t'. Under autoquote, we see the true state of the order book, and if a trade at time t' occurs at price $p_{t'}$ (at either the bid price $b_{t'}$ or the ask price $a_{t'}$), assume that the effective half-spread $s_{t'} = q_{t'}(p_{t'}$ - $m_{t'})$ is correctly measured. In contrast, before the adoption of autoquote the observed midquote at time t' is $m_t = 1/2(a_t + b_t)$, which is stale. Because we focus on the arrival of a sell order that improves the ask, $m_{t'} < m_t$, which means that in the absence of autoquote the observed quote midpoint is biased upwards. Define the measured effective spread pre-autoquote as $s_{t',pre} = q_{t'}(p_{t'}$ - $m_t)$.

Thus, the change in the measured effective spread under autoquote is the difference $s_{t'}$ - $s_{t',pre}$ $= q_{t'}$ $(m_t$ - $m_{t'}) = q_{t'}$ $(a_t$ - $a_{t'})/2$. The term in parentheses is positive, since the arriving sell order improves the quote by lowering the ask price, so the effective spread declines under autoquote if and only if $\mathrm{E}(q_{t'}) < 0$. But this cannot be the case as long as the demand for immediacy is downward sloping in the price of immediacy. To say it another way, a better ask price should on average draw in a marketable buy order, which implies $\mathrm{E}(q_{t'}) > 0$. Thus, if autoquote is simply displaying quotes that were previously undisseminated, the result should be a widening of the effective spread under autoquote.

Note that there is an implicit assumption in the above analysis that without autoquote, the difference between the true midquote $m_{t'}$ and the disseminated midquote m_t does not affect $q_{t'}$, the sign of the trade. The trade sign can indeed be affected if the new ask price $a_{t'}$ is below the disseminated midquote m_t. In this case both the true ask and bid prices are below the disseminated midquote, and with the right choice of parameter values effective spreads could be mechanically narrower under autoquote. However, this scenario seems unlikely to dominate. First, it is quite likely that the specialist would disseminate an updated quote if an incoming limit order crosses the midquote in this way, as the new quoted spread would be less than half as wide as the old quoted spread. Second, if this scenario were empirically important, the resulting trade-signing errors would bias downward the pre-autoquote estimates of the adverse selection component of the spread, because future price changes would be less correlated with trade signs. In this scenario, we would expect to see an *increase* in adverse selection with the elimination of stale quotes under autoquote. This is the opposite of our findings in Tables 3 and 5 in the main text.

Our argument above makes use of the observed decline in adverse selection post-autoquote. If this decline is an artifact of measurement error, our argument is weakened. In addition, the reduction in adverse selection associated with autoquote is quite striking. Thus, it is worth considering a mechanical explanation for the observed changes in adverse selection.[1]

Recall that in order to measure adverse selection, we use quotes 5 minutes or 30 minutes after the trade. In the VAR approach, we use the next 10 trades to calculate the permanent price impact of a unit shock to signed order flow. If it takes longer than this to replenish the quotes after a trade exhausts the depth at the inside, our estimates of adverse selection would be biased upward. AT replenishes quotes more rapidly, removing this upward bias, and making it appear that adverse selection is declining in AT. However, our 30-minute results are virtually identical to our 5-minute results, implying that there is little quote replenishment during that 25 minute interval. Thus, while we think changes in quote replenishment are unlikely to drive the adverse selection results, we cannot rule out the possibility.

To summarize, neither a mechanical increase in quote disseminations nor faster quote replenishment is likely to be the source of our results.

II. Instrumental variable regression with a noisy proxy for algorithmic trading

As discussed in the text, suppose we begin with a linear relationship between liquidity L_{it} and algorithmic trading A_{it}:

$$L_{it} = \alpha_i + \beta A_{it} + \delta' X_{it} + \varepsilon_{1it}, \tag{1}$$

where X_{it} is a vector of control variables. The usual full-rank conditions apply, and $E(X_{it}\varepsilon_{1it}) = 0$, but $cov(A_{it}, \varepsilon_{1it}) \neq 0$ because A_{it} also depends on L_{it}:

$$A_{it} = \omega_i + \theta L_{it} + \phi' X_{it}. \tag{2}$$

Furthermore, the observed proxy for algorithmic trading A_{it} measures algorithmic trading with error:

$$A^o_{it} = A_{it} + \varepsilon_{2it} \tag{3}$$

so that

$$A^o_{it} = \omega_i + \theta L_{it} + \phi' X_{it} + \varepsilon_{2it}. \tag{4}$$

Suppose there exists an instrument Z_{it} s.t. $cov(Z_{it}, A_{it}) \neq 0$, $cov(Z_{it}, \varepsilon_{1it}) = 0$, $cov(Z_{it}, \varepsilon_{1it}) = 0$, and $var(\varepsilon_Z) > 0$ where ε_Z is the residual of a regression of Z_{it} on X_{it}. We rewrite equation (1) as

$$L = W\xi + \varepsilon_1, \tag{5}$$

where we stack all equations indexed by it into vectors and matrices so that the subscripts disappear:

$W = [1 \quad A^o \quad X]$, $\xi' = [\alpha' \quad \beta \quad \delta']$, and $\tilde{Z} = [1 \quad Z \quad X]$ where 1 is a dummy matrix to match the stock-specific fixed effects. Now premultiply by $n^{-1}\tilde{Z}'$:

$$n^{-1}\tilde{Z}'L = n^{-1}\tilde{Z}'W\xi + n^{-1}\tilde{Z}'\varepsilon_1. \tag{6}$$

By assumption, plim $n^{-1}\tilde{Z}'\varepsilon_1 = 0$, so a consistent estimate is:

$$\hat{\xi} = (\tilde{Z}'W)^{-1}\tilde{Z}'L. \tag{7}$$

This is well-defined, since the $[Z \quad X]$ matrix is of full rank, and $cov(Z_{it}, A^o_{it}) \neq 0$ because we assumed that the instrument is correlated with the desired endogenous variable ($cov(Z_{it}, A_{it}) \neq 0$). So the consistency of the IV estimator is unaffected by using a noisy proxy for AT.

III. Lin-Sanger-Booth spread decomposition

The decomposition of the effective spread introduced in equations (2) and (3) in the main text has the advantage of being simple, but it also has distinct disadvantages. In particular, it chooses an arbitrary time point in the future (five minutes or 30 minutes in this case) and implicitly ignores other trades that might have happened in that time period. Lin, Sanger, and Booth (Lin, Sanger, and Booth (1995)) develop a spread decomposition model that is estimated trade by trade and accounts for order flow persistence (the empirical fact, first noted by Hasbrouck and Ho (1987), that buyer-initiated trades tend to follow buyer-initiated trades).[2] Let

$$\delta = Prob[q_{t+1} = 1 | q_t = 1] = Prob[q_{t+1} = -1 | q_t = -1] \tag{8}$$

be the probability of a continuation (a buy followed by a buy or a sell followed by a sell). Further suppose that the change in the market-maker's quote midpoint following a trade is given by:

$$m_{t+1} - m_t = \lambda_t q_t. \tag{9}$$

The dollar effective half-spread $s_t = q_t(p_t - m_t)$ and is assumed constant for simplicity. If there is persistence in order flow, the expected transaction price at time $t + 1$ does not equal m_{t+1} but instead is:

$$\begin{aligned} E_t(p_{t+1}) &= \delta(m_t + q_t(\lambda_t + s_t)) + (1 - \delta)(m_t + q_t(\lambda_t - s_t) \\ &= m_t + q_t(\lambda_t + (2\delta - 1)s_t). \end{aligned} \tag{10}$$

This expression shows how far prices are expected to permanently move against the market-maker. While the market-maker earns s_t initially, in expectation he then loses $\lambda_t + (2\delta - 1)s_t$ due to adverse selection and order persistence, respectively. Note that this reduces to Glosten (1987) if $\delta = 0.5$ so that order flow is independent over time. We can identify the adverse selection component λ by regressing midpoint changes on the buy-sell indicator, and we can identify the order persistence parameter with a first-order autoregression on q_t. The remaining portion of the effective spread is revenue for the market maker, referred to by LSB as the fixed component of the spread. Thus, spreads are decomposed into three separate components: a fixed component associated with temporary price changes, an adverse selection component, and a component due to order flow persistence. The fixed, temporary component continues to reflect the net revenues to liquidity suppliers after accounting for losses to (the now persistent) liquidity demanders. The adverse selection component captures the immediate gross losses to the current liquidity demander, while the order flow persistence component captures the expected gross losses to those demanding liquidity in the same direction in the near future. We estimate the model and calculate components of the effective spread for each sample stock each day.

[insert Figure IA-6]

For each of the market-cap quintiles, the three panels of Figure IA-6 show how the three LSB spread components evolve over the whole 2001 to 2005 sample period. There are no consistent trends in the fixed component: around the implementation of autoquote, there is an increase for the smallest quintile, but this increase does not extend to the other quintiles. In contrast, the adverse selection component falls sharply during the implementation of autoquote in the first half of 2003. This is true across all five quintiles, and the change appears to be permanent. Beginning in the second half of 2002 and continuing to the end of 2005, there is also a steady decline in the order persistence component of the spread. This suggests less persistence, which could indicate that over this period algorithms and human traders both become more adept at concealing their order flow patterns, perhaps by using mixed order submission strategies that sometimes demand liquidity and sometimes supply it.

[insert Table IA-VI]

The staggered introduction of autoquote allows us to take out all market-wide effects and focus on cross-sectional differences between the stocks that implement autoquote early vs. the stocks that implement autoquote later on. As we did for the simpler decomposition, we can put any one of the LSB spread components on the LHS of our IV specification to determine the sources of the liquidity improvement when there is more algorithmic trading. The results are in Panel B of

Table IA-VI and are quite consistent with the earlier decomposition. For the largest two quintiles, autoquote (and the resulting increases in algorithmic trading) are associated with an increase in the fixed component of the spread, and a decrease in the adverse selection component and the order persistence component. The drop in the adverse selection component is economically quite large. During the autoquote sample period, the within standard deviation in our algorithmic trading variable is 4.54, so a one standard deviation increase in algorithmic trading during this sample period leads to an estimated change in the adverse selection component equal to $4.54 * -0.26$, or about a 1.2 basis point narrowing of the adverse selection component. This is quite substantial, given that the adverse selection component for the biggest quintile is only about 2 basis points on average out of an overall 3.62 basis point effective half-spread. The coefficients on the other two components are of similar magnitude, indicating similar economic importance. As in the earlier decomposition, there are no significant effects for the smaller-cap quintiles.

IV. A generalized Roll model

To further explore our counter-intuitive results, particularly the increase in realized spreads caused by AT, here we develop a generalized Roll model that is a slight variation on one developed in Hasbrouck (2007). Though the model is quite simple, it provides a useful framework for thinking about algorithmic trading and delivers a number of empirical predictions, all of which match our empirical results.

A. The model without algorithmic trading

The "game" has two periods, each with an i.i.d. innovation in the efficient price:

$$m_t = m_{t-1} + w_t, \tag{11}$$

where $w_t \in \{\epsilon, -\epsilon\}$, each with probability 0.5. The game features three stages:

- At t = 0, risk-neutral humans can submit a bid and ask quote and, given full competition, the first one arriving bids her reservation price.

- At t = 1, humans can observe w_1 at cost c. If humans choose to buy this information, they can submit a new limit order.[3]

- At t = 2, two informed liquidity demanders arrive, one with a positive private value associated with a trade, $+\theta$, the other with a negative private value, -θ.

We assume that $2c > \theta$, i.e., the cost of "observing" information for humans is sufficiently high that they do not update their quotes. The technical assumption $\epsilon > \theta$ ensures that trade occurs at $t = 2$ iff the efficient price changes between $t = 0$ and $t = 2$, and that only one of the two arriving liquidity demanders transacts in that case.

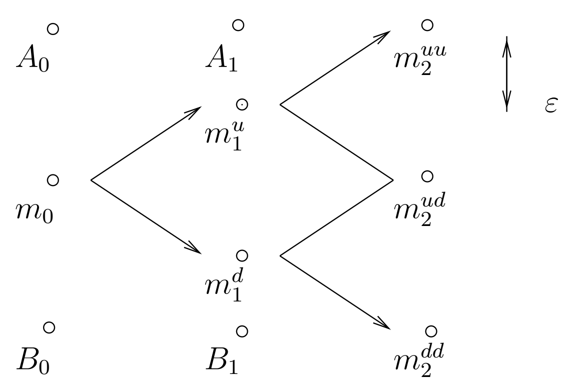

There are four equally likely paths through the binomial tree: uu, ud, du, and dd, where u represents a positive increment of ϵ to the fundamental value and d is a negative increment. In equilibrium, humans do not buy the w_1 information and update the quote at $t = 1$, because they have to quote so far away from the efficient price to make up for c that neither liquidity demander will transact at that quote as $2c > \theta$. Given that they do not acquire the w_1 information, humans protect themselves by setting the bid price equal to $m_0 - 2\epsilon$ and the ask price equal to $m_0 + 2\epsilon$. One of the liquidity demanders trades at $t = 2$ if the path is either uu or dd; the quote providers break even. If the path is ud or du, then there is no trade, because the liquidity demander's private value is too small relative to the spread.

Clearly, under these assumptions all price changes are associated with order flow, and there is no public information component.

B. The model with algorithmic trading

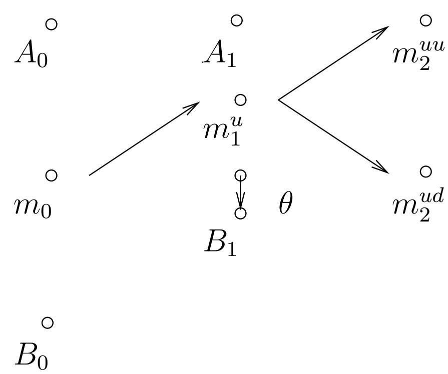

Now we introduce an algorithm that can buy the w_1 information at zero cost ($c = 0$). The results at $t = 0$ remain unchanged. At $t = 1$, the algorithm optimally issues a new quote. To illustrate the idea, suppose $w_1 > 0$. The algorithm knows that it is the only liquidity provider in possession of w_1, and so it puts in a new bid equal to $m_0 - \theta$. If $w_2 > 0$ as well, then a transaction takes place at the original ask of $m_0 + 2\epsilon$. If $w_2 < 0$, then a liquidity demander will hit the algorithm's bid. This bid is below the efficient price, so there will eventually be a reversal, and there is a temporary component in prices. Contrariwise, if $w_1 < 0$, the algorithm places a new ask at $m_0 + \theta$, which is traded with if it turns out that $w_2 > 0$.

In the presence of algorithmic trading, part of the change in the efficient price is revealed through a quote update without trade. Public information now accounts for a portion of price discovery, and imputed revenue to liquidity suppliers is now positive. Thus, the model can explain even the surprising empirical findings on realized spreads and trade-correlated price moves. The model also delivers narrower quoted spreads and more frequent trades, both of which are also observed in the data.

To deliver an increase in realized spread, it is important in the model that competition between algorithms be less vigorous than the competition between humans. This seems plausible in reality as well. As autoquote was implemented in 2003, the extant algorithms might have found themselves with a distinct competitive advantage in trading in response to the increased information flow, given that new algorithms take considerable time to build and test.

Notes

[1]We thank an anonymous referee for suggesting this alternative.

[2]See Barclay and Hendershott (2004) for discussion of how the Lin, Sanger, and Booth spread decomposition relates to other spread decomposition models.

[3]Periods here are on the order of seconds, and the information is best thought of as information contained in order flow and prices, rather than as a direct signal about future cash flows.

References

Arellano, Manuel, and Stephen 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.

Barclay, Michael J., and Terrence Hendershott, 2004, Liquidity externalities and adverse selection: evidence from trading after hours, *Journal of Finance* 59, 681–710.

Glosten, Lawrence R., 1987, Components of the bid ask spread and the statistical properties of transaction prices, *Journal of Finance* 42, 1293–1307.

Hasbrouck, Joel, 2007, *Empirical Market Microstructure* (Oxford University Press, New York).

Hasbrouck, Joel, and Thomas Ho, 1987, Order arrival, quote behavior and the return generating process, *Journal of Finance* 42, 1035–1048.

Lin, Ji-Chai, Gary C. Sanger, and G. Geoffrey Booth, 1995, Trade size and components of the bid-ask spread, *Review of Financial Studies* 8, 1153–1183.

Table IA-I: Summary Statistics for Five Year Sample

This table presents summary statistics for the five year dataset that merges TAQ, CRSP, and NYSE System Order Data (cf. Table 1 in the main text that is based on the autoquote daily sample). The balanced panel consists of monthly data on 943 stocks from February 2001 through December 2005. Stocks are sorted into quintiles based on market capitalization, where quintile 1 contains large-cap stocks. All variables are 99.9% winsorized.

variable	description (units)	source	mean Q1	mean Q2	mean Q3	mean Q4	mean Q5	st. dev. within[a]
$qspread_{it}$	share-volume-weighted quoted half spread (bps)	TAQ	5.31	7.33	9.47	12.92	22.44	8.40
$qdepth_{it}$	share-volume-weighted depth ($1,000)	TAQ	92.37	52.93	38.62	28.69	19.43	21.88
$espread_{it}$	share-volume-weighted effective half spread (bps)	TAQ	3.67	5.19	6.79	9.40	16.16	6.42
$rspread_{it}$	share-volume-weighted realized half spread, 5min (bps)	TAQ	0.96	1.24	1.56	2.19	4.95	2.82
$adv\_selection_{it}$	share-volume-weighted adverse selection component half spread, 5min, "effective-realized" (bps)	TAQ	2.71	3.96	5.23	7.22	11.21	5.02
$messages_{it}$	#electronic messages per minute i.e. proxy for algorithmic activity (/minute)	NYSE	131.99	71.70	43.46	28.86	15.84	43.79
$algo\_trad_{it}$	dollar volume per electronic message times (-1) to proxy for algorithmic trading ($100)	TAQ/NYSE	-26.34	-15.22	-10.88	-8.38	-5.95	11.20
$dollar\_volume_{it}$	average daily volume ($million)	TAQ	112.13	31.70	13.85	7.03	2.82	23.18
$trades_{it}$	#trades per minute (/minute)	TAQ	5.84	3.19	2.02	1.43	0.80	1.28
$share\_turnover_{it}$	(annualized) share turnover	TAQ/CRSP	1.02	1.48	1.46	1.44	1.22	0.69
$volatility_{it}$	standard deviation daily midquote returns (%)	CRSP	1.75	1.95	1.96	2.16	2.54	1.01
$price_{it}$	daily closing price ($)	CRSP	45.90	38.60	33.09	27.98	20.62	9.53
$market\_cap_{it}$	shares outstanding times price ($billion)	CRSP	36.75	5.48	2.30	1.17	0.53	5.09
$trade\_size_{it}$	trade size ($1,000)	TAQ	46.52	24.95	16.97	12.25	8.32	11.52
$specialist\_particip_{it}$	specialist participation rate (%)	NYSE	12.42	12.19	12.28	13.16	15.15	4.42
#observations: 943*59 (stock*month)								

[a]: Based on month t's deviation relative to the time mean, i.e., $x^*_{i,t} = x_{i,t} - \overline{x}_i$.

Table IA-II: Overall, Between, and Within Correlations for Five Year Sample

This table presents overall, between, and within correlations for some variables in the monthly sample that extends from February 2001 through December 2005. Table IA-I provides variable definitions.

		$messa-$ ges_{it}	$algo\_$ $trad_{it}$	$share\_$ $turnover_{it}$	$vola$- $tility_{it}$	$1/price_{it}$	ln_mar- ket_cap_{it}
$qspread_{it}$	ρ(overall)	-0.43*	0.10*	-0.14*	0.54*	0.74*	-0.57*
	ρ(between)^a	-0.51*	0.51*	-0.09*	0.65*	0.83*	-0.68*
	ρ(within)^b	-0.33*	-0.23*	-0.20*	0.48*	0.63*	-0.59*
$messages_{it}$	ρ(overall)		-0.08*	0.13*	-0.20*	-0.24*	0.72*
	ρ(between)^a		-0.87*	0.08*	-0.17*	-0.32*	0.90*
	ρ(within)^b		0.63*	0.19*	-0.24*	-0.13*	0.43*
$algo\_trad_{it}$	ρ(overall)			-0.12*	-0.12*	0.24*	-0.52*
	ρ(between)^a			-0.11*	0.19*	0.36*	-0.86*
	ρ(within)^b			-0.14*	-0.28*	0.12*	0.02*
$share\_turnover_{it}$	ρ(overall)				0.35*	-0.07*	-0.07*
	ρ(between)^a				0.44*	-0.03*	-0.13*
	ρ(within)^b				0.31*	-0.12*	0.15*
$volatility_{it}$	ρ(overall)					0.47*	-0.29*
	ρ(between)^a					0.72*	-0.41*
	ρ(within)^b					0.30*	-0.33*
$1/price_{it}$	ρ(overall)						-0.44*
	ρ(between)^a						-0.45*
	ρ(within)^b						-0.66*

^{a}: Based on the time means, i.e., $\overline{x}_i = \frac{1}{T}\sum_{t=1}^{T} x_{i,t}$.

^{b}: Based on month t's deviation relative to the time mean, i.e., $x^*_{i,t} = x_{i,t} - \overline{x}_i$.

*: Significant at a 95% level.

Table IA-III: Effective Spread Forecast minus Its Long-Term Mean on Autoquote Introduction Day

This table forecasts effective spread on the autoquote introduction day based on the pre-introduction period in order to analyze whether introductions coincide with temporarily wide spreads. One time-series regression estimates a univariate AR(1) model. The second specification also includes lagged values of variables that correlate with liquidity – share turnover, volatility, the inverse of price, and log market cap (cf. control variables in Table 3 of the main text):

$$L_{it} = \alpha_i + \hat{\gamma}_t + \beta_i L_{i,t-1} + \delta_i X_{i,t-1} + \varepsilon_{it}, \quad t \in [2, \ldots, \tau_i - 1]$$

where L_{it} is the effective half spread for stock i on day t, X_{it} is a vector of predictor variables (i.e. share turnover, volatility, 1/price, and log market cap), α_i is the stock-specific mean, $\hat{\gamma}$ is cross-sectional average for each day t, and τ_i is the autoquote introduction day for stock i. Regressions are estimated stock by stock. Panel A and B report the results for a univariate AR(1) model (i.e. setting δ_i to zero). Panel A reports the AR(1) parameter (β_i) estimates and their standard errors both by quintile and overall. Panel B reports the out-of-sample liquidity forecast on the autoquote introduction day. That is, the forecast is based on all days up until the last day before the introduction:

$$f_{\tau_i} = \hat{\beta}_i L_{\tau_i - 1} - \hat{\alpha}_i$$

where the hats indicate estimates based on the pre-introduction period. Panel C replicates Panel B, but includes the control variables in the estimation and in the forecast.

Q1	Q2	Q3	Q4	Q5	all
Panel A: AR(1) coefficient estimates (β_i)					
0.186	0.199	0.153	0.181	0.180	0.180
(0.152)	(0.146)	(0.137)	(0.148)	(0.140)	(0.065)
Panel B: Forecast[a] minus long-term mean, AR(1)					
0.018	-0.002	-0.043	-0.038	-0.036	-0.020
(0.224)	(0.356)	(0.686)	(1.152)	(1.963)	(0.045)
Panel C: Forecast[a] minus long-term mean, AR(1) + controls					
0.032	0.037	-0.060	-0.000	-0.070	-0.012
(0.505)	(0.767)	(1.101)	(1.585)	(2.962)	(0.101)

*/**: Significant at a 95%/99% level.

[a]: The forecast is out-of-sample, i.e. the model estimate and the forecast are based on all days up until the last day before the autoquote introduction.

Table IA-IV: Effect of Algorithmic Trading on Spread: Results for Numerator and Denominator of Algorithmic Trading Proxy

This table separately regresses the effective half spread on the numerator and the denominator of the AT proxy. The regression is based on daily observations in the period from December 2002 through July 2003, covering the phase-in of autoquote. The nonsynchronous autoquote introduction instruments for the endogenous $algo\_trad_{it}$, its denominator, and its numerator, respectively. The specification is (cf. Table 3 in main text)

$$L_{it} = \alpha_i + \gamma_t + \beta A_{it} + \delta X_{it} + \varepsilon_{it}$$

where L_{it} is a spread measure for stock i on day t, A_{it} is the algorithmic trading measure $algo\_trad_{it}$, its denominator, or its numerator, and X_{it} is a vector of control variables, including share turnover, volatility, 1/price, and log market cap. Fixed effects and time dummies are always included. The set of instruments consists of all explanatory variables, except that A_{it} is replaced with $auto\_quote_{it}$. There are separate regressions for each size quintile, and t-values in parentheses are based on standard errors that are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (see Arellano and Bond (1991)).

	Coefficient on A_{it}					Coefficients on control variables[a]			
	Q1	Q2	Q3	Q4	Q5	$share\_turnover_{it}$	$vola-tility_{it}$	$1/price_{it}$	$ln\_mkt\_cap_{it}$
Panel A: $A_{it} = algo\_trad_{it} = -\frac{dollar\_volume_{it}}{messages_{it}}$ *(cf. Table 3 main text)*									
$espread_{it}$	-0.18**	-0.32**	-0.35	-1.67	4.65	-1.01**	0.69**	72.77**	-1.30
	(-2.67)	(-2.23)	(-1.56)	(-0.42)	(1.16)	(-2.30)	(9.39)	(10.80)	(-1.46)
Panel B: $A_{it} = messages_{it}$									
$espread_{it}$	-0.01**	-0.10**	-0.10**	-0.21	1.68	0.38**	0.76**	68.32**	0.95**
	(-3.70)	(-2.05)	(-2.24)	(-0.87)	(1.49)	(5.11)	(17.46)	(13.37)	(2.77)
Panel C: $A_{it} = dollar\_volume_{it}$									
$espread_{it}$	-0.05**	206.98	0.87	0.35	9.86	-350.29	-1.16	-535.11	-679.84
	(-2.15)	(0.00)	(1.42)	(0.77)	(0.80)	(-0.00)	(-0.00)	(-0.00)	(-0.00)
#observations: 1082*167 (stock*day)									

*/**: Significant at a 95%/99% level.

[a]: Coefficients for the control variables and time dummies are quintile-specific. For brevity, only (across the quintiles) market-cap-weighted coefficients are reported for the control variables.

Table IA-V: Effect of Algorithmic Trading on Spread: Turnover Removed as Covariate

This table regresses various measures of the (half) spread on our algorithmic trading proxy. It mirrors Table 3 in the main text where the only difference is that share turnover is removed due to an endogeneity concern. It is based on daily observations from December 2002 through July 2003 which covers the phase-in of autoquote. The nonsynchronous autoquote introduction instruments for the endogenous $algo\_trad_{it}$ to identify causality from algorithmic trading to liquidity. The specification is:

$$L_{it} = \alpha_i + \gamma_t + \beta A_{it} + \delta X_{it} + \varepsilon_{it}$$

where L_{it} is a spread measure for stock i on day t, A_{it} is the algorithmic trading measure $algo\_trad_{it}$, and X_{it} is a vector of control variables, including volatility, 1/price, and log market cap. Fixed effects and time dummies are always included. The set of instruments consists of all explanatory variables, except that $algo\_trad_{it}$ is replaced with $auto\_quote_{it}$. There are separate regressions for each size quintile, and t-values in parentheses are based on standard errors that are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (see Arellano and Bond (1991)).

	Coefficient on $algo\_trad_{it}$					Coefficients on control variables[a]		
	Q1	Q2	Q3	Q4	Q5	$vola\text{-}tility_{it}$	$1/price_{it}$	$ln\_mkt\_cap_{it}$
Panel A: Quoted spread, quoted depth, and effective spread								
$qspread_{it}$	-0.67**	-0.29**	-0.43	-0.11	4.52**	-0.93	157.19**	-1.34
	(-2.16)	(-2.19)	(-1.16)	(-0.05)	(2.04)	(-1.13)	(4.24)	(-0.65)
$qdepth_{it}$	-3.85**	-1.65**	-1.99	10.09	-0.61	-5.10	104.21	16.66
	(-2.11)	(-1.98)	(-1.06)	(0.32)	(-0.47)	(-1.04)	(0.51)	(1.38)
$espread_{it}$	-0.23**	-0.22**	-0.36	-1.13	2.18*	0.05	88.58**	-0.51
	(-1.99)	(-2.23)	(-1.28)	(-0.37)	(1.84)	(0.15)	(6.36)	(-0.57)
Panel B: Effective spread by trade size category[b]								
$espread1^c_{it}$	-0.16**	-0.08*	-0.17	-1.23	2.32**	-0.27	63.80**	-0.48
	(-2.07)	(-1.82)	(-0.90)	(-0.36)	(1.95)	(-1.28)	(6.69)	(-0.67)
$espread2^c_{it}$	-0.29**	-0.19**	-0.42	-2.87	2.09*	-0.60	78.50**	-0.96
	(-2.13)	(-2.41)	(-1.22)	(-0.39)	(1.85)	(-1.58)	(4.80)	(-0.67)
$espread3^c_{it}$	-0.33*	-0.15	-0.67	-1.83	3.22**	-0.55	91.14**	-0.43
	(-1.94)	(-1.19)	(-1.20)	(-0.34)	(2.06)	(-1.19)	(4.57)	(-0.32)
$espread4^c_{it}$	-0.16	-0.17	-0.24	-6.52	-0.02	0.01	78.32**	-1.10
	(-1.19)	(-0.76)	(-0.58)	(-0.22)	(-0.01)	(0.01)	(5.31)	(-0.23)
$espread5^c_{it}$	-0.05	-0.21	-0.37	-1.10	1.90	0.85**	77.01**	-0.36
	(-0.42)	(-0.94)	(-0.79)	(-0.17)	(0.82)	(2.37)	(5.58)	(-0.30)

<continued on next page>

<continued from previous page>

	Coefficient on $algo\_trad_{it}$					Coefficients on control variables[a]		
	Q1	Q2	Q3	Q4	Q5	$vola-tility_{it}$	$1/price_{it}$	$ln\_mkt\_cap_{it}$
Panel C: Spread decompositions based on 5-min and 30-min price impact								
$rspread_{it}$	0.45**	0.53**	1.04*	9.88	6.86**	0.80	5.96	2.36
	(2.33)	(3.59)	(1.68)	(0.42)	(2.50)	(1.21)	(0.26)	(0.62)
$adv\_selection_{it}$	-0.67**	-0.75**	-1.40*	-10.72	-4.69**	-0.74	82.21**	-2.83
	(-2.34)	(-3.75)	(-1.64)	(-0.42)	(-2.21)	(-0.84)	(2.39)	(-0.65)
$rspread\_30m_{it}$	0.42**	0.34**	0.92	7.72	5.46**	-0.71	16.05	0.51
	(2.12)	(2.01)	(1.45)	(0.42)	(2.15)	(-1.15)	(0.71)	(0.17)
$adv\_selection\_30m_{it}$	-0.64**	-0.57**	-1.28	-8.72	-3.42*	0.79	71.48**	-1.01
	(-2.33)	(-2.75)	(-1.53)	(-0.42)	(-1.78)	(0.97)	(2.23)	(-0.27)

#observations: 1082*167 (stock*day)

F test statistic of hypothesis that instruments do not enter first stage regression: 7.32 ($F(5, 179587)$), p-value: 0.0000

*/**: Significant at a 95%/99% level.

[a]: Coefficients for the control variables and time dummies are quintile-specific. For brevity, only (across the quintiles) market-cap-weighted coefficients are reported for the control variables.

[b]: The suffix indicates the effective spread for a particular trade size category, i.e.

"1" if 100 shares $\leq$ trade size $\leq$ 499 shares;
"2" if 500 shares $\leq$ trade size $\leq$ 1999 shares;
"3" if 2000 shares $\leq$ trade size $\leq$ 4999 shares;
"4" if 5000 shares $\leq$ trade size $\leq$ 9999 shares;
"5" if 9999 shares $<$ trade size.

Table IA-VI: Effect of Algorithmic Trading on Spread: Lin, Sanger, and Booth (1995) Spread Decomposition

This table regresses various components of the effective half spread on the algorithmic trading (AT) proxy. It uses the spread decomposition model of Lin, Sanger, and Booth (1995) (LSB) which accounts for order persistence. The LSB model identifies a fixed (transitory) component ($LSB95\_fixed_{it}$), an adverse selection component ($LSB95\_adv\_sel_{it}$), and a component due to order persistence ($LSB95\_order\_persist_{it}$) (see section 1 for details). The regressions are based on daily observations in the period from December 2002 through July 2003 which covers the phase-in of autoquote. The nonsynchronous autoquote introduction instruments for the endogenous $algo\_trad_{it}$ to identify causality from these explanatory variables to liquidity. The specification is (cf. Table 3 in main text)

$$L_{it} = \alpha_i + \gamma_t + \beta A_{it} + \delta X_{it} + \varepsilon_{it}$$

where L_{it} is a spread measure for stock i on day t, A_{it} is the algorithmic trading measure $algo\_trad_{it}$, and X_{it} is a vector of control variables, including share turnover, volatility, 1/price, and log market cap. Fixed effects and time dummies are always included. The set of instruments consists of all explanatory variables, except that $algo\_trad_{it}$ is replaced with $auto\_quote_{it}$. There are separate regressions for each size quintile, and t-values in parentheses are based on standard errors that are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (see Arellano and Bond (1991)).

	Coefficient on $algo\_trad_{it}$					Coefficients on control variables[a]			
	Q1	Q2	Q3	Q4	Q5	$share\_turnover_{it}$	$vola-tility_{it}$	$1/price_{it}$	$ln\_mkt\_cap_{it}$
$LSB95\_fixed_{it}$	0.26**	0.59**	0.69**	9.92	8.97	2.36**	-0.28	26.21**	3.86
	(3.62)	(4.16)	(2.26)	(0.46)	(1.36)	(2.06)	(-0.80)	(3.80)	(1.29)
$LSB95\_adv\_sel_{it}$	-0.26**	-0.61**	-0.84**	-12.21	-7.72	-2.58*	0.57	15.71**	-4.27
	(-3.45)	(-3.80)	(-2.14)	(-0.46)	(-1.32)	(-1.85)	(1.31)	(1.99)	(-1.15)
$LSB95\_order\_persist_{it}$	-0.18**	-0.30**	-0.21	0.64	3.30	-0.82**	0.41**	30.73**	-0.93
	(-3.06)	(-3.10)	(-1.60)	(0.27)	(1.21)	(-2.32)	(8.81)	(6.16)	(-1.47)
#observations: 1082*167 (stock*day)									

*/**: Significant at a 95%/99% level.

[a]: Coefficients for the control variables and time dummies are quintile-specific. For brevity, only (across the quintiles) market-cap-weighted coefficients are reported for the control variables.

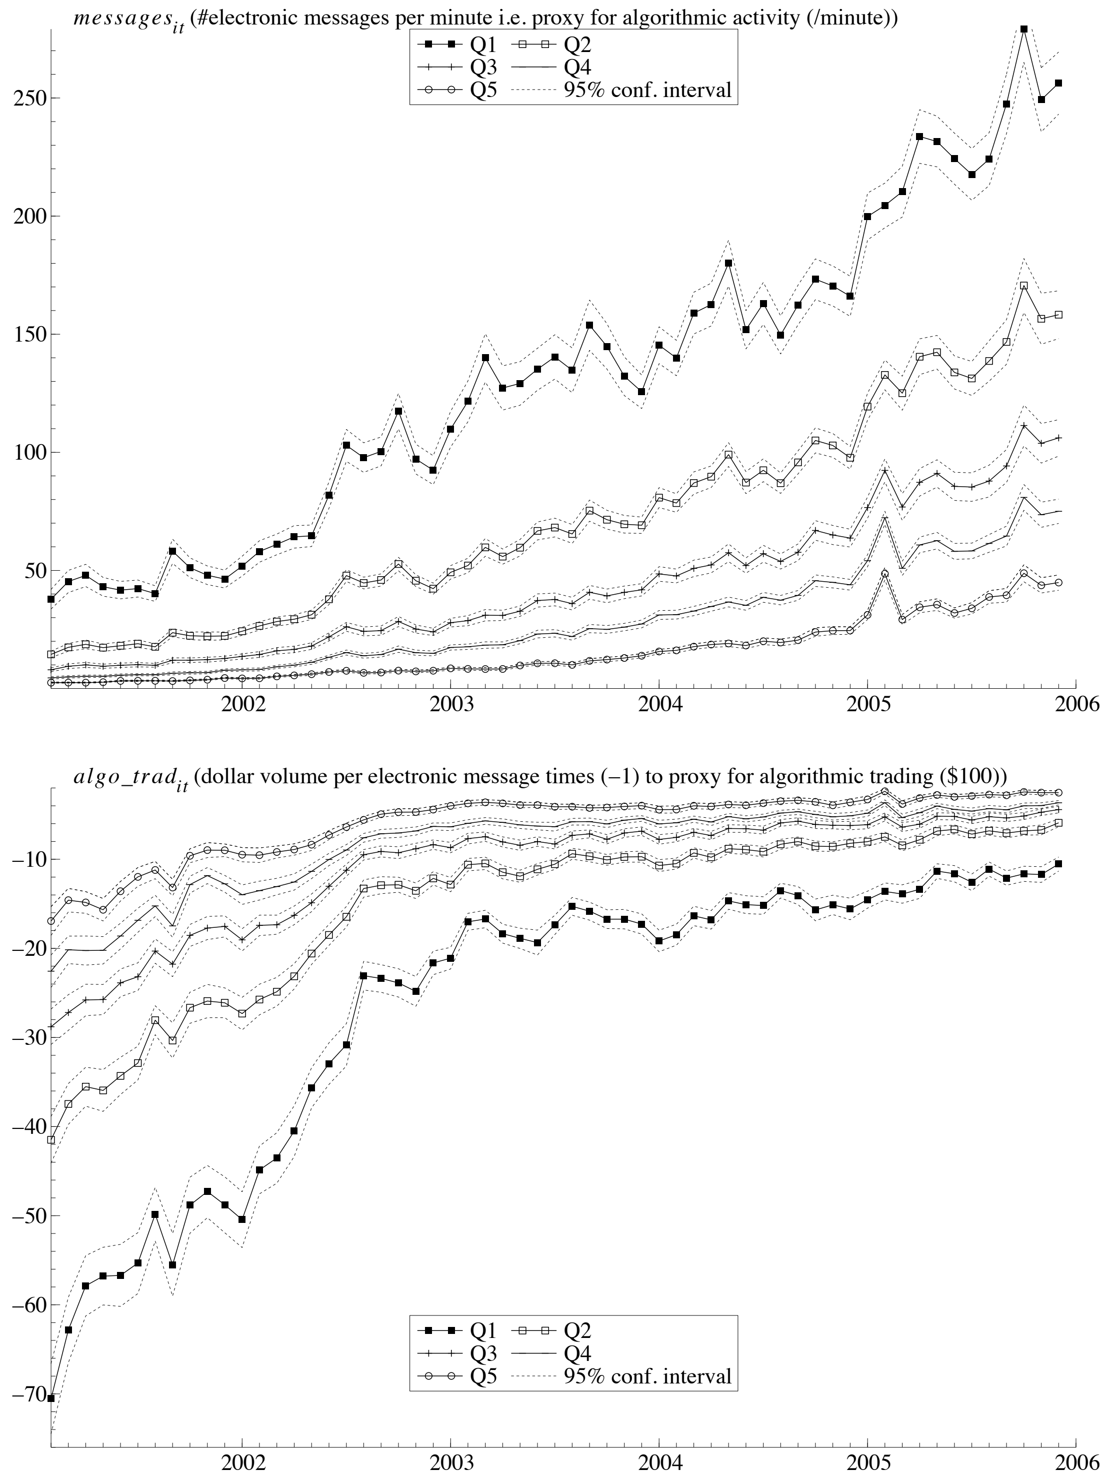

Figure IA-1: For each market-cap quintile, where Q1 is the large-cap quintile, these graphs depict averages for (i) the number of (electronic) messages per minute and (ii) our proxy for algorithmic trading, which is defined as the negative of trading volume (in hundreds of dollars) divided by the number of messages.

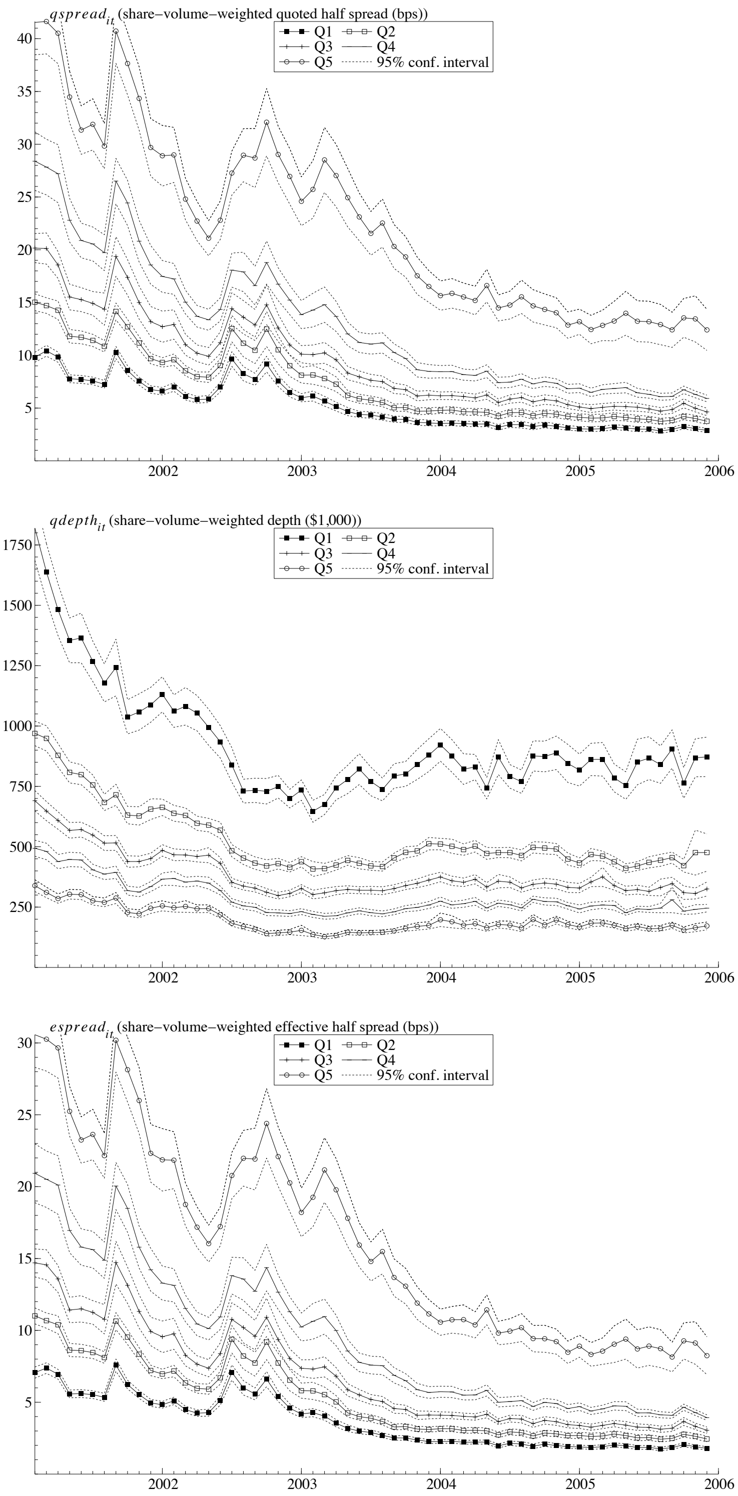

Figure IA-2: These graphs depict (i) quoted half spread, (ii) quoted depth, and (iii) effective spread. All spread measures are share-volume weighted averages within-firm, and then averaged across firms within each market-cap quintile, where Q1 is the large-cap quintile.

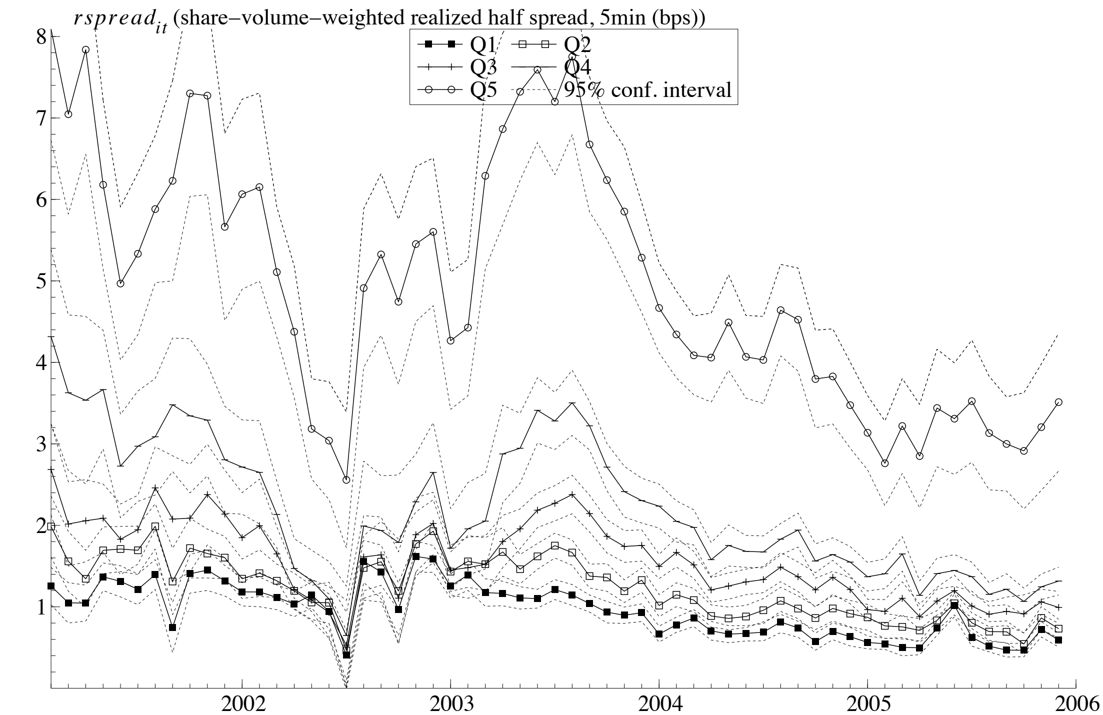

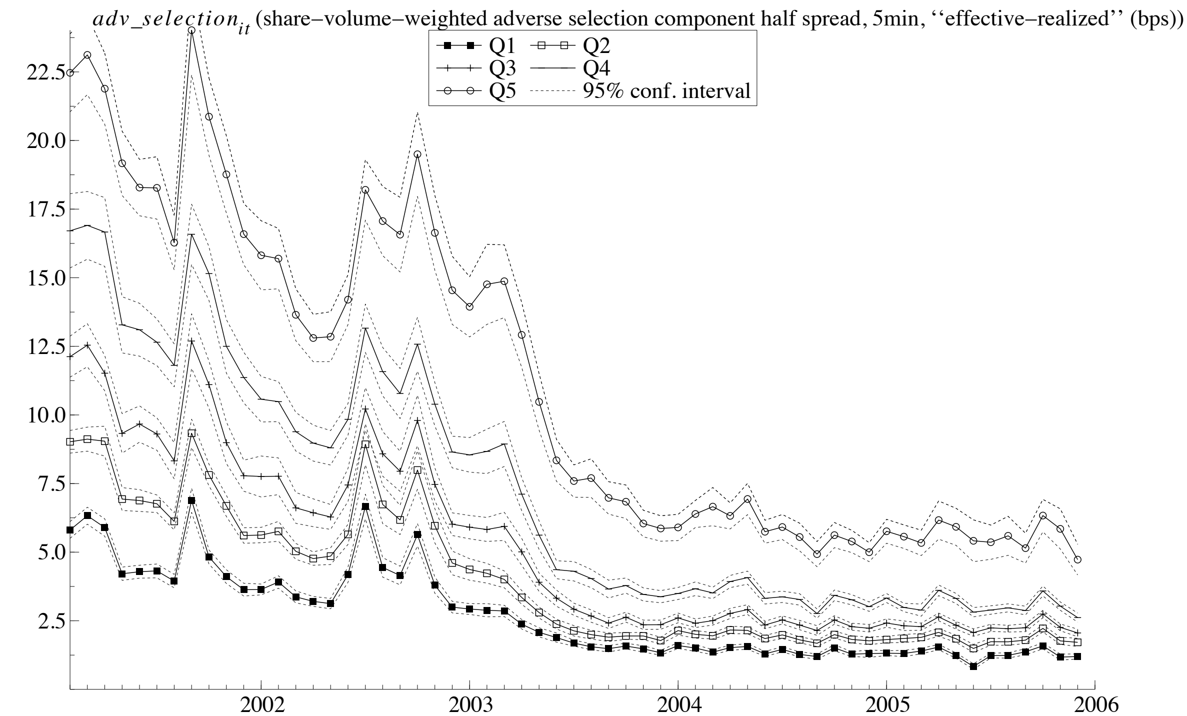

Figure IA-3: These graphs depict the two components of the effective spread: (i) realized spread and (ii) the adverse selection component, also known as the (permanent) price impact. Both are based on the quote midpoint 5 minutes after the trade. Results are graphed by market-cap quintile, where Q1 is the large-cap quintile.

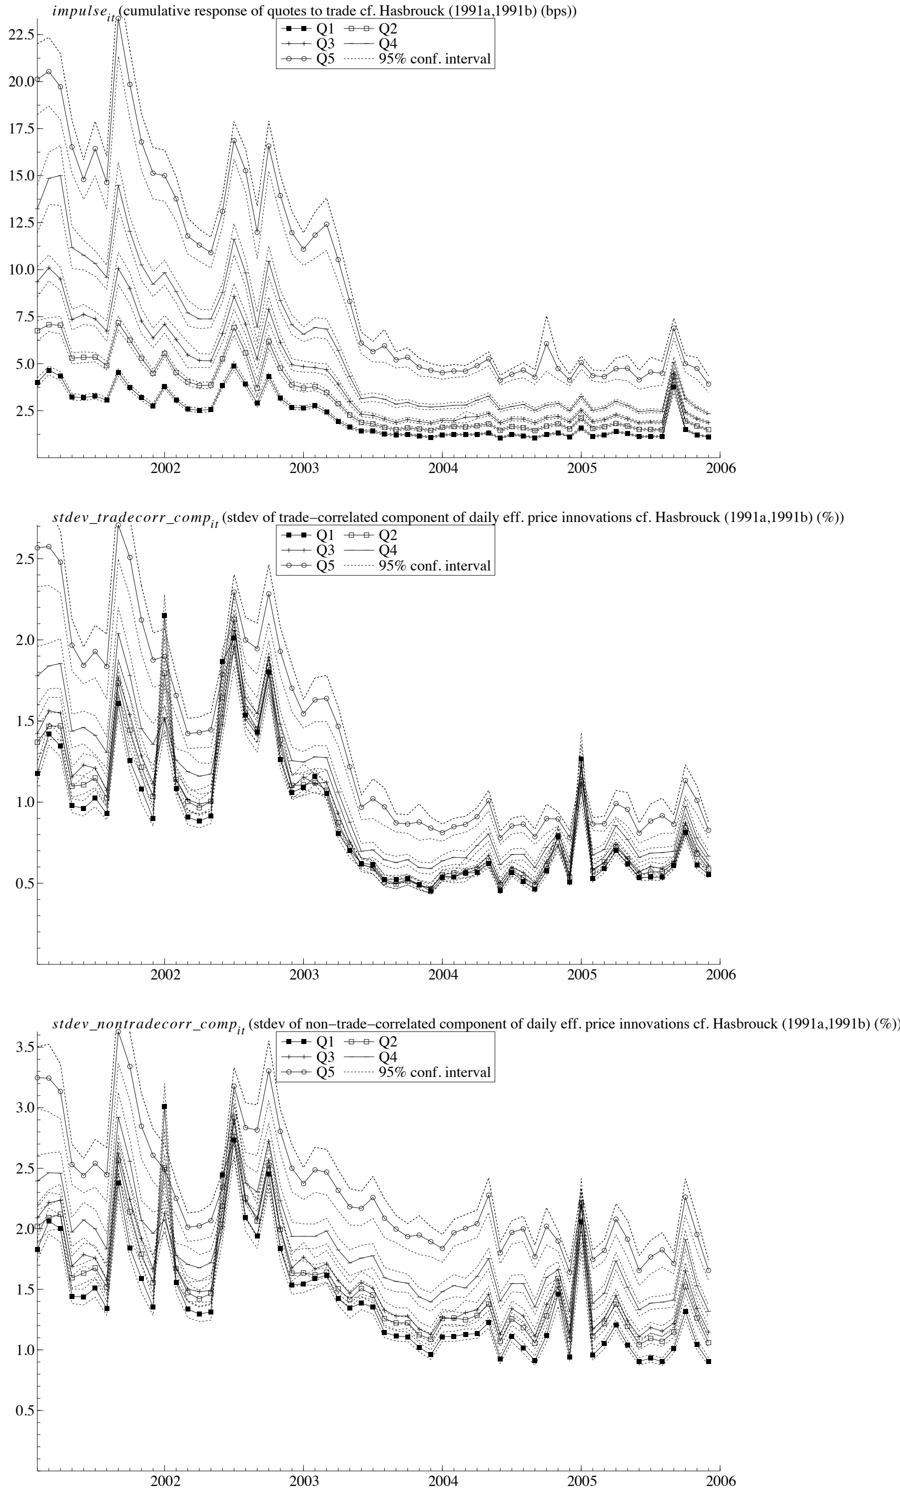

Figure IA-4: These graphs illustrate the estimation results of the Hasbrouck (1991a,1991b) VAR model for midquote returns and signed trades. The top graph illustrates the time series pattern of the long-term price impact of the midquote to a unit impulse in the signed trade variable. The bottom two graphs illustrate the decomposition of the daily percentage variance of changes in the efficient price into a trade-related ($stdev\_tradecorr\_comp_{it}$) and trade-unrelated ($stdev\_nontradecorr\_comp_{it}$) component (see section 6 in the main text and Hasbrouck (1991a, 1991b) for details). Results are reported by market-cap quintile, where Q1 is the large-cap quintile.

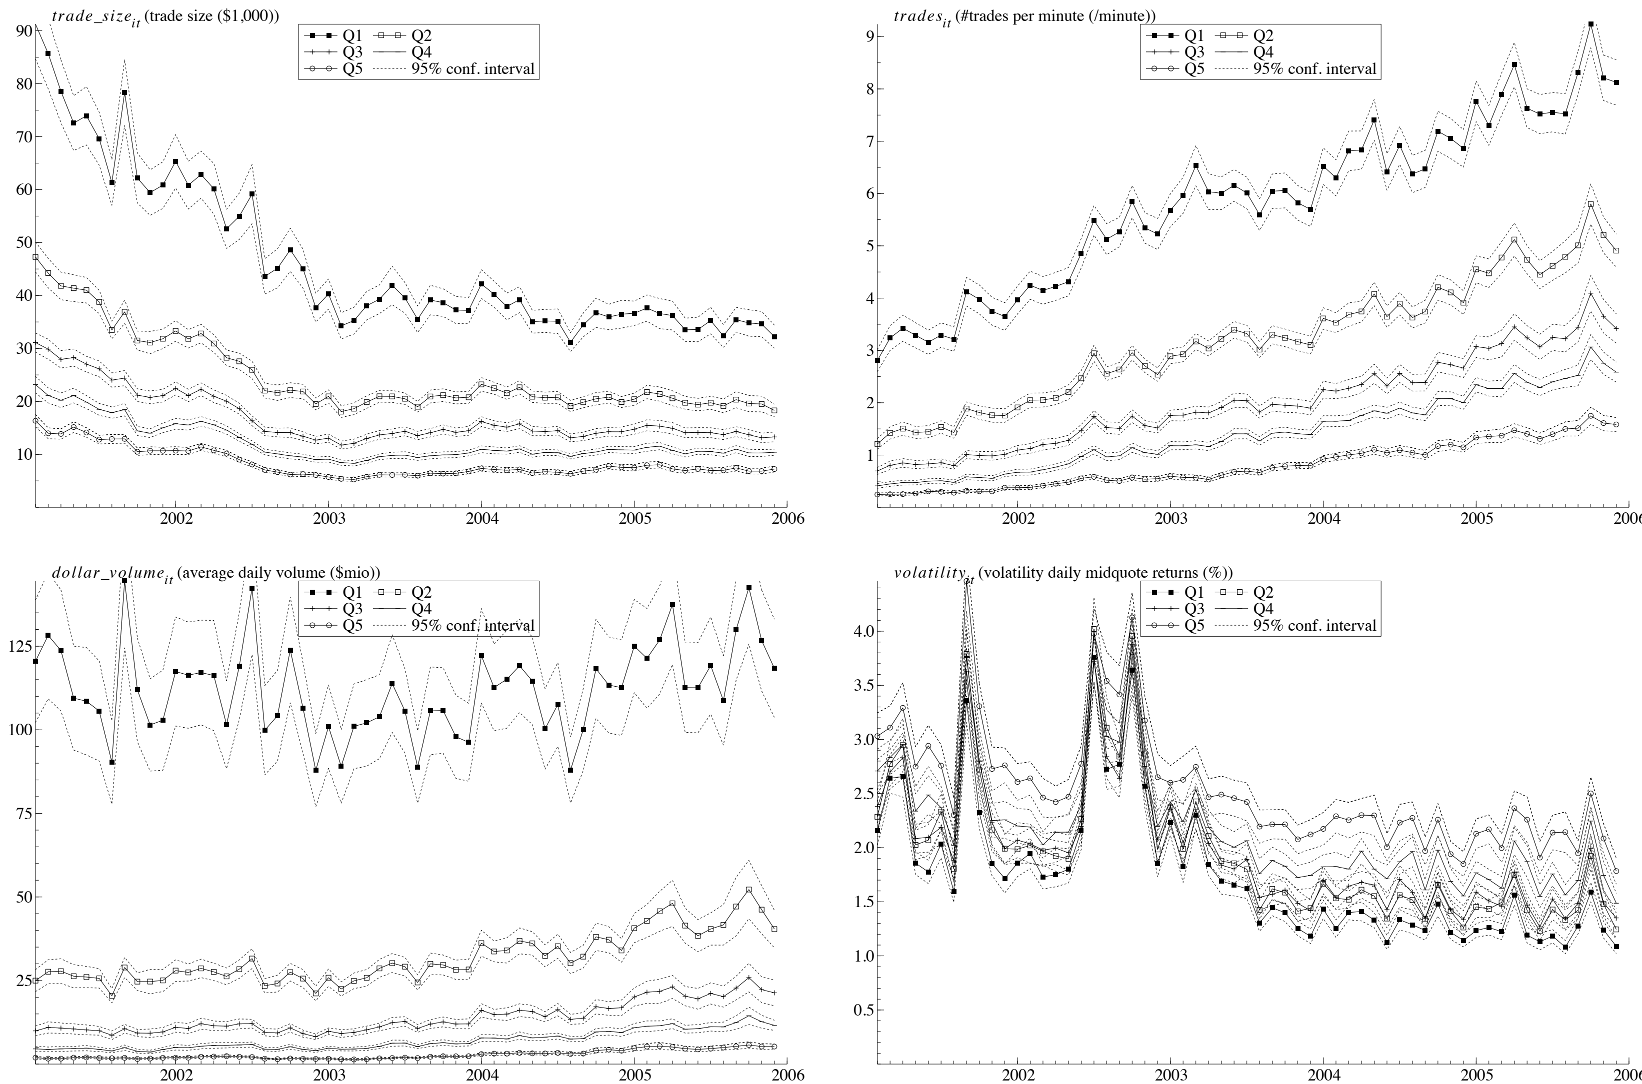

Figure IA-5: These graphs depict (i) trade size, (ii) the number of trades per minute, (iii) daily dollar volume, and (iv) daily midquote return volatility. Results are reported by market-cap quintile, where Q1 is the large-cap quintile.

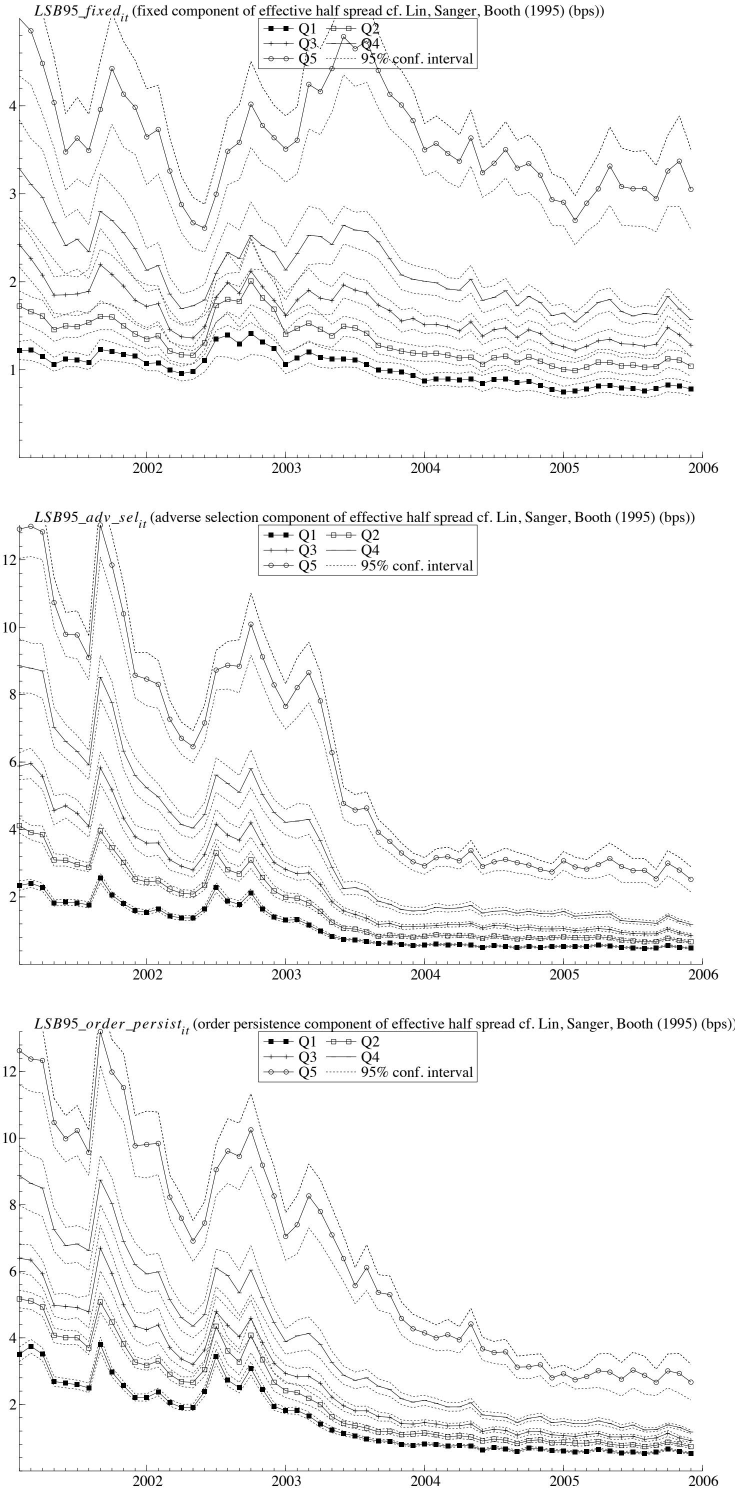

Figure IA-6: These graphs depict the three components of a Lin, Sanger, and Booth (1995) spread decomposition, which identifies a fixed (transitory) component ($LSB95\_fixed_{it}$), an adverse selection component ($LSB95\_adv\_sel_{it}$), and a component due to order persistence ($LSB95\_order\_persist_{it}$) (see section 1 for details). Results are reported by market-cap quintile, where Q1 is the large-cap quintile.