

Does Algorithmic Trading Improve Liquidity?¹

Terrence Hendershott
Haas School of Business
University of California, Berkeley

Charles M. Jones
Graduate School of Business
Columbia University

Albert J. Menkveld
Vrije Universiteit Amsterdam
Tinbergen Institute

First draft: August 20, 2007
This draft: September 4, 2007

¹We thank the NYSE for providing system order data. Hendershott gratefully acknowledges support from the National Science Foundation.

Does Algorithmic Trading Improve Liquidity?

Abstract

Algorithmic trading has sharply increased over the past decade. Equity market liquidity has improved as well. Are the two trends related? For a recent five-year panel of New York Stock Exchange (NYSE) stocks, we use a normalized measure of electronic message traffic as a proxy for algorithmic liquidity supply and trace the associations between liquidity and message traffic. Based on within-stock variation, we find that algorithmic trading and liquidity are positively related. To sort out causality, we use the start of autoquoting on the NYSE as an exogenous instrument for algorithmic trading. Previously, specialists were responsible for manually disseminating the inside quote. As stocks were phased in gradually during early 2003, the manual quote was replaced by a new automated quote whenever there was a change to the NYSE limit order book. This market structure change provides quicker feedback to traders and algorithms and results in more message traffic. For large-cap stocks in particular, quoted and effective spreads narrow under autoquote and adverse selection declines, indicating that algorithmic trading does causally improve liquidity.

Technological change has revolutionized the way financial assets are traded. Back office improvements can support vastly increased trading volume. Retail investors place orders via computer rather than speaking to a broker on the phone. Trading floors have largely been replaced by electronic trading platforms (Jain (2005)).

The nature of order execution has changed dramatically as well, as many market participants now employ algorithmic trading (AT), commonly defined as the use of computer algorithms to manage the trading process. From a starting point near zero about ten years ago, AT is now thought to be responsible for $\frac{1}{3}$ of trading volume in the U.S and is expected to account for perhaps half of trading volume by 2010 (Economist (2007a)). The intense activity generated by algorithms threatens to overwhelm exchanges and market data providers (Economist (2007b)), forcing significant upgrades to their infrastructures (Economist (2007a)).

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 (2007) 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, and quantity of orders, dynamically monitoring market conditions across different securities and trading venues, reducing market impact by optimally (and possibly randomly) breaking large orders into smaller pieces, and closely tracking benchmarks such as the volume-weighted average price (VWAP) over the execution interval.

Many observers think of algorithms from the standpoint of this institutional buy-side investor.¹ But there are other important users of algorithms. 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 a marketable order. All of these are also forms of algorithmic trading.²

¹See, for example, Domowitz and Yegerman (2005).

²Algorithms can also be used to formulate trading decisions and strategies as well as implement them.

As AT has grown rapidly over the past ten years or so, 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 attempt to gauge empirically the 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 for the most part 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 using “slice and dice” algorithms.

AT’s effect on liquidity is assessed using two empirical approaches. First, panel regressions are used to establish that time-series increases in algorithmic trading are associated with more liquid markets. To disentangle causality, we use the start of autoquoting on the NYSE as an exogenous 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 change provides quicker feedback to traders and 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 the effects.

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

There are clearly feedback effects between the portfolio strategy side and the execution side. For example, algorithmic execution could be the difference in making a trading-intensive algorithmic portfolio strategy feasible. Our data reflect counts of actual orders submitted and canceled, so we focus on the execution side of algorithms.

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. There are no significant effects for smaller-cap stocks, but our instrument is weaker here, 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 expected that if AT improved liquidity, the mechanism would be competition between liquidity providers. But the evidence clearly indicates that liquidity suppliers are capturing some of the surplus for themselves. To help make sense of this counter-intuitive result, we introduce the generalized Roll model developed in Hasbrouck (2007), modified to allow algorithmic liquidity supply. The model matches up with all our empirical findings. In particular, it shows that liquidity supplier revenues depend on the degree of competition between the marginal liquidity suppliers. To put it starkly, in a world without algorithms, liquidity supplier revenues depend on the degree of competition between liquidity supplier humans. In a world with algorithms, liquidity supplier revenues depend on the degree of competition between algorithms. Our results suggest that, at least immediately following the start of autoquote, there was less competition between the best algorithms, perhaps because new algorithms require considerable investment and time-to-build.

The paper proceeds as follows. Section 1 discusses related literature. Section 2 describes our data and analyzes algorithmic trading and its impact from 2001 through 2005. Section 3 examines algorithmic trading and liquidity surrounding the NYSE's staggered introduction of autoquote in 2003. Section 4 discusses and interprets the results, and Section 5 concludes.

1 Related literature

There are very few papers that address algorithmic trading directly, but there are several strands of academic literature that 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 motivations. 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) allow for an arbitrary number of symmetric competing liquidity suppliers and

find that liquidity suppliers' rents decline as the number increases. Our initial expectation is that AT would lead to entry of additional liquidity suppliers, and reduce their overall rents.

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. The arrival of public information renders existing orders stale and can put the free trading option into the money. Foucault, Roëll, and Sandas (2003) study the equilibrium level of effort that liquidity suppliers should expend in monitoring the market to avoid this picking off risk. Improvements in technology allow algorithms to inexpensively perform such monitoring.³ In this way AT may be a way to implement Black (1995) idea of limit orders indexed to a market index. If AT reduces the cost of the free trading option implicit in limit orders, then measures of adverse selection should decrease with AT. If some users of AT are better at avoiding being picked off, they can impose adverse selection costs on other liquidity suppliers (similar to the mechanism in Rock (1990)) and even drive other liquidity suppliers out.

AT may also be used by traders who trying to passively accumulate or liquidate a large position (see Keim and Madhavan (1995) for evidence of large orders being broken up). How aggressively to execute trades involves complex dynamic optimization problems to determine order size, order frequency, and order type. Bertsimas and Lo (1998) study optimal execution strategies in the presence of temporary price impacts. Conditional on the constraint of completing the entire transaction by a fixed date, orders are broken into pieces so as to minimize cost. Almgren and Chriss (2000) extend this by considering the risk that arises from breaking up order 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. Many brokers build models with such considerations into their AT products that they sell to their clients.

For each component of the larger transaction a trader (or algorithm) must choose

³See Rosu (2006) for a model where limit orders are constantly adjusted.

the type and aggressiveness of the order. Cohen et al. (1981) and Harris (1998) focus on the simplest choice: market order versus limit order. If a trader chooses a non-marketable limit order, the aggressiveness of the order is determined by its limit price (Griffiths et al. (2000) and Ranaldo (2004)). How aggressively a limit order should be priced depends on how price affects the time to execution (Lo, MacKinlay, and Zhang (2002)). If limit orders do not execute, traders may cancel them and resubmit them with more aggressive prices. Hasbrouck and Saar (2007) find that on the INET trading platform (now Nasdaq's trading mechanism) a large number of limit orders are cancelled within two seconds.⁴

2 Data and full sample analysis

To study the relationship between algorithmic trading and liquidity, we start with a monthly panel of NYSE common stocks. 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. The sample extends for almost five years, beginning in February 2001 and ending in December 2005. We start with a sample of all NYSE common stocks that can be matched across the NYSE's Trade and Quotes (TAQ) and the CRSP databases and retain the stocks that are present throughout the whole sample period (in order to maintain a balanced panel). 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 comprises 943 common stocks.

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 (that is, 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).

⁴Boehmer, Saar, and Yu (2005) show that NYSE's introduction of OpenBook, which disseminates limit order book information beyond the best quote, is associated with an increase in the cancellation rate of limit orders.

2.1 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, the rate of electronic message traffic may be a useful proxy for the amount of algorithmic trading taking place. 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 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]

The main 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 these 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 a marked 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, we normalize by calculating for each stock each month the number of electronic messages per \$1,000 of trading volume. This 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. Since we are essentially regressing spreads on algorithmic trading measures, we want both measures to have this same general shape over time. Thus, our preferred measure is just the negative reciprocal of the messages per dollar traded. Specifically, $algo_trad_{it}$ is calculated as the negative of trading volume (in thousands of dollars) divided by the number of electronic messages. The time-series evolution of this measure 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. Cross-sectionally, there is a positive monotonic relationship between market cap and trading volume per message. In general, we focus more on time-series behavior rather than the cross-sectional 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.

2.2 Summary statistics

[insert Table 1]

Table 1 contains means by quintile and within-stock standard deviations for all of the variables used in the analysis. We measure liquidity using quoted half-spreads, effective half-spreads, 5-minute realized spreads, and 5-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 an estimate of the true value of the security (the midpoint of the bid and ask) and the actual transaction price. For the t^{th} trade in stock j , the proportional effective half-spread ($espread_{jt}$) is defined as:

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

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 quoted depth at the time of each transaction and report share-weighted averages measured in thousands of dollars. For example, the mean effective half-spread of 3.67 basis points for stocks in the largest-cap quintile corresponds to a half-spread of 1.68 cents (or a whole spread of 3.36 cents) on a stock with the mean share price of \$45.90.

[insert Figure 2]

The figures show quite clearly that, over time, algorithmic trading is gradually increasing while liquidity is gradually improving. Figure 1 shows that algorithmic trading increases almost monotonically. The spread measures in Figure 2 are not nearly as monotonic, with illiquidity spikes in both 2001 and 2002 that correspond to sharp stock market declines. Nevertheless, one is tempted to conclude that these two trends are related. The analysis to come investigates exactly this relationship using formal econometric tools rather than casual armchair empiricism.

2.3 Correlations

[insert Table 2]

We begin with Table 2, which uses our monthly panel to provide a set of univariate correlations between spreads, algorithmic trading variables, volume, volatility, and share price variables. It is interesting to note that the cross-sectional (between) correlation between spreads and $algo_trad_{it}$ is positive. This matches the cross-quintile evidence in Figure 1 and Figure 2. This is clearly driven by the cross-section of volume because the correlation with raw message traffic is strongly negative. Perhaps there is a fixed component to message traffic, in that a certain amount of message traffic is required for price discovery regardless of how much trading occurs, giving rise to the positive cross-sectional correlation between spreads and $algo_trad_{it}$. While there is undoubtedly value to further analyzing the causes and effects of cross-sectional variability in algorithmic trading, we focus hereafter on within-stock correlations, because we want to understand the impact of the change in algorithmic trading over time.

The within-stock correlation between quoted spreads and algorithmic trading is negative and significant. This is a contemporaneous correlation, and we do not have anything

yet to say about causality. But it appears that, stock by stock in the panel, algorithmic trading is high when spreads are narrow. We also find that algorithmic trading is negatively correlated with volatility, where volatility is measured as the standard deviation of daily midpoint returns for a given month.

2.4 Panel regressions

To confirm that these univariate correlations are robust, we specify regressions for our monthly panel that are of the form:

$$L_{it} = \alpha_i + \gamma_t + \beta A_{it} + \delta X_{it} + \epsilon_{it}, \quad (2)$$

where L_{it} is a liquidity measure (quoted spread, effective spread, or quoted depth) for stock i in month t , A_{it} is the algorithmic trading measure $algo_trad_{it}$, and X_{it} is a vector of control variables, including trading volume, return volatility, inverse price, and log market cap. Firm fixed effects are always included (α_i); results are always reported with and without calendar fixed effects (γ_t). We estimate separate regressions for each of the market-cap quintiles, and standard errors are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (see Arellano and Bond (1991)).

[insert Table 3]

The results are in Table 3 Panel A, and they are qualitatively consistent across the size quintiles. The sign of the algorithmic trading coefficient depends on whether time dummies are included. When there are no calendar fixed effects, the regression is identified using only within-stock variation. Here the results match the univariate within-stock correlations: the coefficient on $algo_trad_{it}$ is negative and significant, indicating that an increase in algorithmic trading is associated with narrower quoted or effective spreads. When time dummies are added, changes to common market-wide liquidity factors are removed, and the regression is identified using only idiosyncratic changes in liquidity. Interestingly, the coefficient on $algo_trad_{it}$ changes sign in this case. This does not cast doubt on our other results; it implies only that algorithmic trading and *idiosyncratic* liquidity are negatively related. However, it does suggest that the positive time-series association within a given stock between liquidity and algorithmic trading is driven by changes in market-wide common liquidity factors.

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 to liquidity demanders, or both. We estimate revenue to liquidity providers using the 5-minute realized spread. The proportional realized spread for the t^{th} transaction in stock j is defined as:

$$rspread_{jt} = q_{jt}(p_{jt} - m_{j,t+5\text{min}})/m_{jt}, \quad (3)$$

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 price at which the liquidity provider is assumed able to close her position).

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:

$$adv_selection_{jt} = q_{jt}(m_{j,t+5\text{min}} - m_{jt})/m_{jt}. \quad (4)$$

Note that there is an arithmetic identity relating the realized spread, the adverse selection (price impact), and the effective spread $espread_{jt}$:

$$espread_{jt} = rspread_{jt} + adv_selection_{jt}. \quad (5)$$

For these spread components, we estimate panel regressions of the same form as before, and the results are in Panel B of Table 3. Again we report results with and without calendar fixed effects, but we focus on the results without time dummies. Both realized spreads ($espread_{jt}$) and price impacts ($adv_selection_{jt}$) are negatively associated with algorithmic trading. More algorithmic trading is associated with narrower effective spreads, and these narrower spreads imply lower revenue per trade for the liquidity provider as well as smaller gross losses to liquidity demanders. But the relative contributions of the two components are very different. 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. We discuss this in considerable detail below, but we suspect that while manually submitted limit orders are eventually picked off

by informed traders, algorithms are better able to avoid some of these informed traders by a so-called “cancel and replace” of the stale limit order.

So far, all we have identified are time-series associations between algorithmic trading and liquidity. We cannot yet say anything about the direction of causality. It is certainly easier to tell a story that goes in the standard direction. For example, algorithmic trading could be providing competition in liquidity provision, thereby improving liquidity. But algorithmic trading is an endogenous choice variable that depends on liquidity, among other parameters. Liquidity is also endogenous and depends on a variety of factors, including the technological and other costs incurred by liquidity providers. Sorting out causality requires an exogenous instrument, and we turn to that in the next section.

3 Autoquote

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 October 2002, 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 more information about expressed trading interest at prices outside of the best bid and offer (it was also designed to recapture some of the block trading business that the NYSE had lost to upstairs markets and to algorithms). A liquidity quote was to be a firm bid and offer for substantial size, typically at least 15,000 shares, accessible immediately via a new type of market order called Institutional Xpress.⁵

At the time of the liquidity quote proposal, specialists were responsible for manually disseminating the inside quote.⁶ 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 allow the inside quote to be disseminated automatically. Under NYSE Rule 60, display book software would “autoquote” the NYSE’s highest bid or lowest offer whenever

⁵For more details, the NYSE proposal is contained in Securities Exchange Act Release No. 47091 (December 23, 2002), 68 FR 133.

⁶There was one main exception. If a new or cancelled SuperDOT limit order would change the inside quote, system software automatically disseminated an updated quote after 30 seconds if the specialist had not already done so.

there was a change to the limit order book via SuperDOT. This would happen when a better-priced order arrived, when the inside quote was traded with in whole or in part, or when the size of the inside quote changed.

[insert Figure 5]

In early 2003, the liquidity quote proposal became enmeshed in a dispute over property rights between the exchange and data vendors such as Bloomberg. The SEC eventually issued a stay delaying the implementation of the liquidity quote, and liquidity quote did not become operational until June 13, 2003. Meanwhile, the NYSE began to phase in the liquidity quote and autoquote software on January 29, 2003, starting with 6 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. Figure 5 provides some additional details on the phase-in process. Because liquidity quotes were not yet accessible or widely disseminated during this phase-in period, the only change to market structure from January to May 2003 was the non-synchronous addition of autoquote, making this an ideal exogenous event for study.

In fact, even when liquidity quotes became firm in June 2003, they had little impact on trading. Only a handful of large orders ever used the Institutional Xpress order type, mainly because the spread was typically quite large relative to what could be negotiated in the upstairs market. Most of the information in the liquidity quote was already available to traders and algorithms via NYSE's Openbook product, which provided periodic snapshots of the NYSE limit order book (see Boehmer, Saar, and Yu (2005) for an analysis of the introduction of OpenBook). Ultimately, liquidity quotes were deemed an unsuccessful innovation and were abandoned in July 2005. But the autoquote feature stayed in place.

For algorithmic traders, autoquote was an important innovation, because it provided much more immediate feedback about the potential terms of trade. 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 Direct+ facility, which provided automatic execution for limit orders of 1,099 shares or less against the exchange's disseminated quote.

3.1 Autoquote sample

To study the effects of autoquote, we build a daily panel for 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. For consistency, we start with the same share price filters as before: stocks with an average share price of less than \$5 or more than \$1000 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 3.3.2) 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.

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. All variables used in the analysis are 99.9% winsorized (that is, 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).

3.2 Autoquote results

[insert Table 4]

Autoquote clearly leads to greater use of algorithms. Message traffic increases by about 50% in the most active quintile of stocks as autoquote is phased in (see Figure 5), and the daily within-stock correlation between message traffic and the autoquote dummy reported in Table 4 is 0.08. Correlations are higher for large-cap stocks, consistent with the conventional wisdom that algorithmic trading was more effective at the time for active, liquid stocks. Table 4 also shows that there is also a significant positive correlation 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. This plus the exogenous phase-in makes the introduction of autoquote an ideal instrument for assessing the impact of algorithmic liquidity suppliers.

Within-stock correlations in Table 4 also show that after the introduction of autoquote turnover is higher, volatility is lower, and share prices are higher. However, we have no intention of ascribing these results to autoquote. These results likely reflect the fact

that the market rose during the early part of 2003 for unrelated reasons, and they highlight the importance of the staggered introduction of autoquote for cleanly identifying the effect of the market structure change. By including time dummies in the panel specification, we can use non-autoquoted stocks as controls, comparing phased-in autoquoted stocks to not-yet-autoquoted stocks.

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: the quoted (half) spread, the effective half-spread, the 5-minute realized spread, or the 5-minute price impact of a given trade, 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 (measured as the standard deviation of daily midquote returns in percent), 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.

[insert Table 6]

The results are reported in Panel A of Table 6, 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 -0.52 on the algorithmic trading variable for quintile 1 means that an increase in algorithmic trading from the whole-sample mean of \$2,634 of volume per message implies that quoted spreads narrow by 0.52 basis points. Over the whole five-year sample interval, the average within-stock standard deviation for $algo_trad_{it}$ is 11.2 or \$1,120, so a unit standard deviation change in our algorithmic trading measure is associated with a 5.82 basis point change in proportional spreads. 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 5 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.

3.3 Decomposition of the spread improvement

As discussed earlier in the paper, narrower effective spreads imply either less revenue per trade for liquidity providers, smaller gross losses to liquidity demanders, or both. In Table 6 Panel B, 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. IV regressions are repeated using each component of the spread.

The results are somewhat surprising. For quintiles 1 through 3 (large and medium-cap stocks), 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 autoquote. As before, nothing is significant for the two smallest-cap quintiles.

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.

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 6 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 6 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.

3.3.1 Lin-Sanger-Booth decomposition

The decomposition of the effective spread introduced 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 in this case) and implicitly ignores other trades that might have happened in that five-minute 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).⁷ Let

$$\delta = Prob[q_{t+1} = 1|q_t = 1] = Prob[q_{t+1} = -1|q_t = -1] \quad (6)$$

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:

$$m_{t+1} - m_t = \lambda_t q_t. \quad (7)$$

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:

$$\begin{aligned} E_t(p_{t+1}) &= \delta(m_t + q_t(\lambda_t + s_t)) + (1 - \delta)(m_t + q_t(\lambda_t - s_t)) \\ &= m_t + q_t(\lambda_t + (2\delta - 1)s_t). \end{aligned} \quad (8)$$

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

⁷See Barclay and Hendershott (2004) for discussion of how the Lin, Sanger, and Booth spread decomposition relates to other spread decomposition models.

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 6]

For each of the market-cap quintiles, the three panels of Figure 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 7]

As discussed above, we are fortunate and do not need to hang our hats on these time-series declines. 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 A of Table 7 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.

3.3.2 Hasbrouck decomposition

While the Lin-Sanger-Booth model begins to consider persistence in order flow, it implicitly limits the form of that persistence to an AR(1) process. Hasbrouck (1991a, 1991b) 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” (examples of this interpretation include 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:

$$r_t = \sum_{i=1}^{10} \alpha_i r_{t-i} + \sum_{i=0}^{10} \beta_i q_{t-i} + \epsilon_{rt}, \quad (9)$$

$$q_t = \sum_{i=1}^{10} \gamma_i r_{t-i} + \sum_{i=1}^{10} \phi_i q_{t-i} + \epsilon_{qt}, \quad (10)$$

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

$$y_t = \begin{bmatrix} r_t \\ q_t \end{bmatrix} = \theta(L)\epsilon_t = \begin{bmatrix} a(L) & b(L) \\ c(L) & d(L) \end{bmatrix} \begin{bmatrix} \epsilon_{rt} \\ \epsilon_{qt} \end{bmatrix}, \quad (11)$$

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

$$\sigma_w^2 = \left(\sum_{i=0}^{\infty} a_i \right)^2 \sigma_r^2 + \left(\sum_{i=0}^{\infty} b_i \right)^2 \sigma_q^2, \quad (12)$$

where the second term captures the component of price discovery that is related to trade, and the first term captures price changes that are unrelated to trading (sometimes referred to as public information). As discussed in Hasbrouck (1991a, 1991b), this method is robust to price discreteness, lagged adjustment to information, and lagged adjustment to trades.

[insert Figure 7]

The VAR and the trade-related and non-trade-related standard deviations are estimated for each stock each day. We calculate monthly averages for each quintile during the autoquote sample period and graph these in Figure 7. The most striking feature of the graph is the decline in the trade-related standard deviation, while the non-trade-related standard deviations do not change much as autoquote is introduced. This indicates that under autoquote much more information is being incorporated into prices without trade.

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.

The results can be found in Panel B of Table 7, and at least for the two largest quintiles they confirm that the time-series graphs are not spurious. When a large-cap stock adopts autoquote and experiences an exogenous increase in algorithmic trading, there is much less trade-correlated price discovery, and much more price discovery that is uncorrelated with trading. We discuss this further below, but it seems likely that algorithms are responding quickly to order flow information and price moves of this and other stocks, and thereby updating quotes to prevent them from becoming stale and being picked off. Consistent with other methodologies, we do not find consistently reliable effects for the three smallest-cap quintiles.

Based on our estimates, 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 $4.54 * -0.22$, or almost exactly a 1 basis point reduction in the standard deviation of trade-correlated returns. Figure 7 shows that this is the same order of magnitude as the actual level of trade-correlated standard deviations measured from trade to trade, so this is indeed a substantial change in how prices are updated to reflect new information over time.

4 Discussion and interpretation

To help make sense of our counter-intuitive results (particularly the realized spread or temporary component results), we turn next to a very simple 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.

4.1 A generalized Roll model


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

$$m_t = m_{t-1} + w_t, \quad (13)$$

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 buy the information w_1 at cost c . If they buy the information, they can submit a new limit order.
- 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, $-\theta$.

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 $2\epsilon > \theta$ is also required so that only one of the two arriving liquidity demanders transacts at $t = 2$.




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 $t = 1$, since 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.

4.1.1 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.

What kind of information can algorithms efficiently observe? There are probably many answers, but it is hard to tell, given the general opacity practiced by algorithm providers and users. Nevertheless, we suspect that two kinds of information are of first-order importance. First, we think algorithms can easily take into account common factor price information and adjust trading and quoting accordingly. For example, if there is 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, replace 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 (see Economist (2007a)).

4.2 Alternative explanations

Up to now, we have focused 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 quotes less stale and has no other effects? First, it is worth noting that this is counterfactual, since autoquote increases message traffic and our algorithmic trading proxy. But more timely quotes alone can also generate many of the same results. Consider the effective spread, for example. Suppose for concreteness that transactions are equally likely on either side of the efficient price m_t :

$$p_t = m_t + q_t s_t. \tag{14}$$

If quotes are stale, the observed midpoint m_t^0 differs from the efficient price by a stochastic amount δ_t . If this error is uncorrelated and the support of δ_t extends beyond the interval $[-s_t, +s_t]$, then it is fairly easy to show that the calculated effective spread is biased upward:

$$Es_t^0 = E|p_t - m_t^0| > s_t, \quad (15)$$

and anything that reduces the measurement error for the observed quote midpoint will tend to reduce the measured effective spread.

Similarly, if quotes did not always reflect the arrival of new limit orders prior to autoquote, the incremental quotes post-autoquote result in more non-trade related price changes. If the overall variance of the efficient price does not change with the introduction of autoquote, more of the variance in a Hasbrouck (1991a, 1991b) decomposition is attributed to public information, and less is attributed to private information, simply because we observe more quote changes.

In contrast, quoted spreads are unlikely to be affected in the same way by a mechanical increase in the frequency of quotes disseminated. Nevertheless, we find that autoquote affects quoted spreads at least as much as it affects effective spreads. This suggests that the mechanical increase in quote disseminations does not account for our findings.

5 Conclusions

The declining costs of technology have led to its wide-spread 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. Using panel regressions over the five years following decimalization, we establish that time-series increases in algorithmic trading are associated with more liquid markets. To establish causality we use the staggered introduction of a structural change at the NYSE (autoquoting) as an exogenous instrument for algorithmic trading. We demonstrate that increased algorithmic trading lowers adverse selection and decreases the amount of price discovery that is correlated with trading. These results suggest that algorithmic trading lowers the costs of trading and increases the informativeness of quotes and prices. Surprisingly, the revenues to liquidity suppliers also increase with algorithmic trading. This is consistent with algorithmic liquidity suppliers

having market power as they introduce their algorithms.

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 increase competition among 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.

References

- Almgren, R., and N. Chriss. 2000. "Optimal Execution of Portfolio Transactions." *Journal of Risk* 3:5–39.
- Arellano, M., and S.R. Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economic Studies* 58:277–297.
- Barclay, M.J., and T. Hendershott. 2003. "Price Discovery and After Trading Hours." *Review of Financial Studies* 16(4):1041–1073.
- . 2004. "Liquidity Externalities and Adverse Selection: Evidence from Trading After Hours." *Journal of Finance* 59:681–710.
- Bertsimas, D., and A.W. Lo. 1998. "Optimal Control of Execution Costs." *Journal of Financial Markets* 1:1–50.
- Bessembinder, H. 2003. "Issues in Assessing Trade Execution Costs." *Journal of Financial Markets* 1:1–50.
- Biais, B., D. Martimort, and J.C. Rochet. 2000. "Competing Mechanisms in a Common Value Environment." *Econometrica* 68(4):799–837.
- Black, F. 1995. "Equilibrium Exchanges." *Financial Analysts Journal* 51:23–29.
- Boehmer, E., G. Saar, and L. Yu. 2005. "Lifting the Veil: An Analysis of Pre-Trade Transparency at the NYSE." *Journal of Finance* 60:783–815.
- Cohen, K., S. Maier, R. Schwartz, and D. Whitcomb. 1981. "Transaction Costs, Order Placement Strategy and Existence of the Bid-Ask Spread." *Journal of Political Economy* 89:287–305.
- Copeland, T.E., and D. Galai. 1983. "Information Effects on the Bid-ask Spread." *Journal of Finance* 38:1457–1469.
- Domowitz, I., and H. Yegerman. 2005. "The Cost of Algorithmic Trading: A First Look at Comparative Performance." Edited by Brian R. Bruce, *Algorithmic Trading: Precision, Control, Execution*. Institutional Investor.
- Economist. 2007a. "Ahead of the Tape—Algorithmic Trading." *Economist*. June 23, 2007.
- . 2007b. "Dodgy Tickers—Stock Exchanges." *Economist*. March 10, 2007.
- Foucault, T., A. Roëll, and P. Sandas. 2003. "Market Making with Costly Monitoring: An Analysis of the SOES Controversy." 16:345–384.

- Glosten, L. 1994. "Is the Electronic Limit Order Book Inevitable?" *Journal of Finance* 49:1127–1161.
- Glosten, L.R. 1987. "Estimating the Components of the Bid-Ask Spread." *Journal of Finance* 42:1293–1307.
- . 1989. "Insider Trading, Liquidity, and the Role of the Monopolist Specialist." *Journal of Business* 62:211–235.
- Griffiths, M., B. Smith, D.A. Turnbull, and R. White. 2000. "The Costs and Determinants of Order Aggressiveness." *Journal of Financial Economics* 56:65–88.
- Harris, L. 1998. "Optimal Dynamic Order Submission Strategies in Some Stylized Trading Problems." *Financial Markets, Institutions, and Instruments* 7(2):1–76.
- Hasbrouck, J. 1991a. "Measuring the Information Content of Stock Trades." *Journal of Finance* 46:179–207.
- . 1991b. "The Summary Informativeness of Stock Trades: An Econometric Analysis." *Review of Financial Studies* 4:571–595.
- . 2007. *Empirical Market Microstructure*. New York: Oxford University Press.
- Hasbrouck, J., and T. Ho. 1987. "Order Arrival, Quote Behavior and the Return Generating Process." *Journal of Finance* 42:1035–1048.
- Hasbrouck, J., and G. Saar. 2007. "Technology and Liquidity Provision: The Blurring of Traditional Definitions." Technical Report, Working Papers, New York University.
- Hendershott, T., and P. Moulton. 2007. "The Shrinking New York Stock Exchange Floor and the Hybrid Market." Technical Report, UC Berkeley Working Paper.
- Jain, P.K. 2005. "Financial Market Design and the Equity Premium: Electronic versus Floor Trading." *Journal of Finance* 60(6):2955–2985.
- Jones, C.M., G. Kaul, and M.L. Lipson. 1994. "Information, Trading, and Volatility." *Journal of Financial Economics* 36:127–154.
- Keim, D., and A. Madhavan. 1995. "Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Traders." *Journal of Financial Economics* 37:371–398.
- Lee, C., and M. Ready. 1991. "Inferring Trade Direction from Intraday Data." *Journal of Finance* 46:733–746.
- Lin, J.C., G. Sanger, and G.G. Booth. 1995. "Trade Size and Components of the Bid-Ask Spread." *Review of Financial Studies* 8:1153–1183.

- Lo, W.A., A.C. MacKinlay, and J. Zhang. 2002. "Econometric Models of Limit-Order Executions." *Journal of Financial Economics* 65:31–71.
- Obizhaeva, A., and J. Wang. 2005. "Optimal Trading Strategy and Supply/Demand Dynamics."
- Ranaldo, A. 2004. "Order Aggressiveness in Limit Order Book Markets." *Journal of Financial Markets* 7(1):53–74.
- Reinganum, J. 1989. "The timing of innovation: Research, development, and diffusion." Technical Report, appeared in *Handbook of Industrial Organization*, page 849-908, edited by R. Schmalensee and R. Willig, Elsevier Publishing, The Netherlands.
- Rock, K. 1990. "The Specialist's Order Book and Price Anomalies." Technical Report, Harvard University Working Paper.
- Rosu, I. 2006. "A Dynamic Model of the Limit Order Book." Technical Report, University of Chicago.
- SEC. 2005. Regulation NMS, <http://www.sec.gov/rules/final/34-51808.pdf>.

Table 1: Summary Statistics (943*59 stock*month)

This table presents summary statistics on our dataset, which combines TAQ, CRSP, and the NYSE System Order Data database. We create a balanced panel that 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. wi- thin ^a
<i>qsread_{it}</i>	share-volume-weighted quoted half spread (bps)	TAQ	5.31	7.33	9.47	12.92	22.44	8.40
<i>qdepth_{it}</i>	share-volume-weighted depth (\$1,000)	TAQ	92.37	52.93	38.62	28.69	19.43	21.88
<i>espread_{it}</i>	share-volume-weighted effective half spread (bps)	TAQ	3.67	5.19	6.79	9.40	16.16	6.42
<i>rsread_{it}</i>	share-volume-weighted realized half spread, 5min (bps)	TAQ	0.96	1.24	1.56	2.19	4.95	2.82
<i>adv_selection_{it}</i>	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
<i>messages_{it}</i>	#electronic messages per minute i.e. proxy for algorithmic activity (/minute)	NYSE	131.99	71.70	43.46	28.86	15.84	43.79
<i>algo_trad_{it}</i>	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
<i>dollar_volume_{it}</i>	average daily volume (\$mio)	TAQ	112.13	31.70	13.85	7.03	2.82	23.18
<i>trades_{it}</i>	#trades per minute (/minute)	TAQ	5.84	3.19	2.02	1.43	0.80	1.28
<i>share_turnover_{it}</i>	(annualized) share turnover	TAQ/CRSP	1.02	1.48	1.46	1.44	1.22	0.69
<i>volatility_{it}</i>	volatility daily midquote returns (%)	CRSP	1.75	1.95	1.96	2.16	2.54	1.01
<i>price_{it}</i>	daily closing price (\$)	CRSP	45.90	38.60	33.09	27.98	20.62	9.53
<i>market_cap_{it}</i>	shares outstanding times price (\$bln)	CRSP	36.75	5.48	2.30	1.17	0.53	5.09
<i>trade_size_{it}</i>	trade size (\$1,000)	TAQ	46.52	24.95	16.97	12.25	8.32	11.52
<i>specialist_particip_{it}</i>	specialist participation rate (%)	NYSE	12.42	12.19	12.28	13.16	15.15	4.42
#observations: 943*59 (stock*month)								

^a: Based on the deviations from time means i.e. $x_{i,t}^* = x_{i,t} - \bar{x}_i$.

Table 2: Overall, Between, and Within Correlations

This table presents the overall, between, and within correlations for some variables in our sample, which contains monthly observations from February 2001 through December 2005. For variable definitions, we refer to Table 1.

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

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

^b: Based on the deviations from time means i.e. $x_{i,t}^* = x_{i,t} - \bar{x}_i$.

*: Significant at a 95% level.

Table 3: Effect of AT on Spread in the Full Sample

This table regresses various measures of the (half) spread on our algorithmic trading proxy. We use our full sample, which consists of monthly observations from February 2001 through December 2005. We estimate

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

where L_{it} is a spread measure for stock i in month 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. We always include fixed effects and report results with and without time dummies. We regress by quintile and report t -values 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				time dum- mies	R ²
		Q1	Q2	Q3	Q4	Q5	$turnover_shares_{it}$	$vola-tility_{it}$	$1/price_{it}$	$ln_mkt_cap_{it}$		
<i>Panel A: Quoted spread, quoted depth, and effective spread</i>												
$qspread_{it}$	(1)	-0.06**	-0.17**	-0.29**	-0.46**	-0.86**	-1.83**	1.78**	57.54**	-2.66**	No	0.71
		(-19.37)	(-19.98)	(-18.53)	(-16.78)	(-17.89)	(-17.11)	(35.44)	(11.98)	(-15.85)		
	(2)	0.03**	0.04**	0.05*	0.12**	0.22**	-0.17	1.06**	74.68**	-0.17	Yes	0.80
		(6.40)	(3.14)	(1.87)	(3.41)	(2.39)	(-1.60)	(18.70)	(12.53)	(-0.80)		
$qdepth_{it}$	(1)	-1.37**	-1.45**	-1.43**	-1.45**	-1.32**	11.74**	-8.44**	76.04*	8.46**	No	0.49
		(-22.44)	(-28.46)	(-26.66)	(-24.37)	(-20.32)	(7.25)	(-13.48)	(1.81)	(3.06)		
	(2)	-1.19**	-1.14**	-1.09**	-1.01**	-1.03**	13.53**	-7.70**	55.81*	3.03	Yes	0.54
		(-10.40)	(-11.96)	(-10.54)	(-13.00)	(-6.72)	(7.73)	(-11.46)	(1.66)	(1.00)		
$espread_{it}$	(1)	-0.05**	-0.14**	-0.23**	-0.36**	-0.67**	-1.44**	1.37**	41.43**	-1.95**	No	0.70
		(-18.99)	(-20.97)	(-19.44)	(-16.58)	(-17.79)	(-16.22)	(36.89)	(14.57)	(-13.76)		
	(2)	0.03**	0.03**	0.05**	0.09**	0.20**	-0.10	0.77**	55.65**	0.14	Yes	0.80
		(6.46)	(3.07)	(2.27)	(3.43)	(2.88)	(-1.12)	(17.39)	(33.13)	(0.98)		
<i>Panel B: Spread decompositions</i>												
$rsread_{it}$	(1)	-0.01**	-0.02**	-0.02**	-0.04**	-0.10**	-0.47**	-0.01	32.14**	-0.23**	No	0.39
		(-6.70)	(-5.96)	(-4.53)	(-5.35)	(-6.34)	(-8.60)	(-0.32)	(7.45)	(-2.08)		
	(2)	0.01**	0.02**	0.04**	0.06**	0.12**	-0.13**	-0.07**	35.90**	0.46**	Yes	0.46
		(4.04)	(3.58)	(4.09)	(4.64)	(4.20)	(-2.75)	(-2.34)	(9.05)	(3.66)		
$adv_selection_{it}$	(1)	-0.04**	-0.12**	-0.21**	-0.32**	-0.57**	-0.96**	1.38**	9.36**	-1.71**	No	0.67
		(-16.99)	(-23.96)	(-22.94)	(-18.02)	(-18.75)	(-12.05)	(38.59)	(2.34)	(-11.65)		
	(2)	0.01**	0.00	0.01	0.03	0.08	0.03	0.84**	19.86**	-0.31*	Yes	0.78
		(3.41)	(0.70)	(0.48)	(1.60)	(1.55)	(0.33)	(17.13)	(4.61)	(-1.79)		

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

^a: We use quintile-specific coefficients for the control variables and time dummies. For brevity, we only report the (across the quintiles) market-cap-weighted coefficient for the control variables and its t -statistic.

Table 4: Overall, Between, and Within Correlations Autoquote Analysis

This table presents the overall, between, and within correlations for the variables used in the autoquote analysis. It is based on daily observations in the period when autoquote was phased in, i.e. December 2, 2003, through July 31, 2003. For variable definitions, we refer to Table 1. We exploit the exogenous autoquote dummy (0 before the autoquote introduction, 1 after) to instrument for $algo_trad_{it}$ in order to identify causality from $algo_trad_{it}$ to our liquidity measures. In the IV estimation, we exclude identification off of a time trend (by adding time dummies) and thus solely rely on the nonsynchronous introduction of autoquote (see Figure 5). Before we report the IV estimation results in subsequent tables, this table reports correlations between the instrument ($auto_quote_{it}$) and the endogenous variable ($algo_trad_{it}$) after removing the time trend.

		<i>messa-</i> <i>ges_{it}</i>	<i>algo-</i> <i>trad_{it}</i>	<i>share-</i> <i>turnover_{it}</i>	<i>vola-</i> <i>tility_{it}</i>	<i>1/price_{it}</i>	<i>ln-mar-</i> <i>ket_cap_{it}</i>
<i>Panel A: Overall, between, and within correlation after removing the time trend</i>							
<i>auto_quote_{it}</i>	$\rho(\text{overall})$	0.15*	-0.05*	0.02*	0.03*	0.02*	0.10*
	$\rho(\text{between})$	0.23*	-0.16*	0.06	0.09*	0.04	0.18*
	$\rho(\text{within})$	0.08*	0.03*	-0.01*	0.00	0.01*	-0.01*
<i>Panel B: Within correlation by quintile after removing the time trend</i>							
<i>auto_quote_{it}</i>	Q1 $\rho(\text{within})$	0.15*	0.03*	0.01*	-0.00	0.03*	-0.03*
<i>auto_quote_{it}</i>	Q2 $\rho(\text{within})$	0.03*	0.04*	-0.01*	0.00	-0.02*	0.01*
<i>auto_quote_{it}</i>	Q3 $\rho(\text{within})$	0.05*	0.03*	0.00	-0.00	0.01	-0.02*
<i>auto_quote_{it}</i>	Q4 $\rho(\text{within})$	0.01*	0.00	-0.00	-0.00	-0.01	0.01
<i>auto_quote_{it}</i>	Q5 $\rho(\text{within})$	-0.00	0.03*	-0.02*	0.00	0.05*	-0.04*

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

^b: Based on the deviations from time means i.e. $x_{i,t}^* = x_{i,t} - \bar{x}_i$.

*: Significant at a 95% level.

Table 5: Effect of AT on Spread: Nonsynchronous Autoquote Introduction as Instrumental Variable

This table regresses various measures of the (half) spread on our algorithmic trading proxy. It is based on daily observations in the period when autoquote was phased in, i.e. December 2, 2003, through July 31, 2003. We use the exogenous nonsynchronous autoquote introduction to instrument for the endogenous $algo_trad_{it}$ to identify causality from algorithmic trading to liquidity. We estimate

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

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. We always include fixed effects and time dummies. The set of instruments we use consists of all explanatory variables, except that we replace $algo_trad_{it}$ with $auto_quote_{it}$. We regress by quintile and report t -values 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				time dummies
	Q1	Q2	Q3	Q4	Q5	$share_turnover_{it}$	$vola_tility_{it}$	$1/price_{it}$	$ln_mkt_cap_{it}$	
<i>Panel A: quoted spread, quoted depth, and effective spread</i>										
$qspread_{it}$	-0.52** (-3.23)	-0.42** (-2.21)	-0.43 (-1.44)	-0.16 (-0.05)	9.92 (1.22)	-2.80** (-3.01)	0.90** (9.70)	108.30** (7.49)	-3.55** (-2.27)	Yes
$qdepth_{it}$	-3.47** (-2.50)	-1.43 (-1.16)	-1.99 (-1.07)	15.49 (0.39)	0.61 (0.19)	-5.16 (-0.64)	-1.64* (-1.87)	-3.90 (-0.03)	12.12 (0.83)	Yes
$espread_{it}$	-0.18** (-2.65)	-0.32** (-2.23)	-0.35 (-1.56)	-1.63 (-0.42)	4.65 (1.16)	-1.01** (-2.32)	0.69** (9.51)	72.72** (10.91)	-1.27 (-1.45)	Yes
<i>Panel B: spread decompositions</i>										
$rspread_{it}$	0.35** (3.53)	0.76** (3.97)	1.03** (2.06)	14.26 (0.46)	15.88 (1.36)	3.13* (1.92)	-1.06** (-2.15)	45.81** (4.14)	5.06 (1.18)	Yes
$adv_selection_{it}$	-0.53** (-3.57)	-1.07** (-4.08)	-1.39** (-2.06)	-15.48 (-0.47)	-11.21 (-1.33)	-4.12** (-2.24)	1.75** (3.29)	26.61* (1.84)	-6.27 (-1.34)	Yes

#observations: 1082*167 (stock*day)

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

^a: We use quintile-specific coefficients for the control variables and time dummies. For brevity, we only report the (across the quintiles) market-cap-weighted coefficient for the control variables and its t -statistic.

Table 6: Effect of AT on Nonspread Variables: Nonsynchronous Autoquote Introduction as Instrumental Variable

This table regresses nonspread variables on our algorithmic trading proxy. It is based on daily observations in the period when autoquote was phased in, i.e. December 2, 2003, through July 31, 2003. We use the exogenous nonsynchronous autoquote introduction to instrument for the endogenous $algo_trad_{it}$ to identify causality from algorithmic trading to these nonspread variables. We estimate

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

where M_{it} is a nonspread variable for stock i on day t , and A_{it} is the algorithmic trading measure. We always include fixed effects and time dummies. We regress by quintile and report t -values 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 $algo_trad_{it}$					time dum- mies
	Q1	Q2	Q3	Q4	Q5	
$share_turnover_{it}$	0.04 (1.03)	-0.07* (-1.77)	0.01 (0.07)	-0.20 (-0.26)	-0.36** (-2.89)	Yes
$trades_{it}$	0.58** (2.60)	-0.01 (-0.23)	-0.01 (-0.15)	-0.51 (-0.33)	-0.15** (-2.60)	Yes
$trade_size_{it}$	-2.04** (-4.63)	-0.80** (-3.23)	-0.33 (-0.69)	2.26 (0.20)	-0.22 (-0.60)	Yes
$specialist_particip_{it}$	-0.59** (-2.22)	-0.23 (-1.24)	-0.92 (-1.43)	-13.19 (-0.29)	-1.89** (-2.02)	Yes

#observations: 1082*167 (stock*day)

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

Table 7: Lin-Sanger-Booth and Hasbrouck Decompositions: Nonsynchronous Autoquote Introduction as Instrumental Variable

This table regresses the components of Lin-Sanger-Booth (LSB) and Hasbrouck decompositions on our algorithmic trading proxy. It is based on daily observations in the period when autoquote was phased in, i.e. December 2, 2003, through July 31, 2003. The LSB decomposition accounts for order persistence in decomposing the bid-ask spread. It 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 3.3.1 and Lin, Sanger, and Booth (1995) for details). As the LSB decomposition limits persistence to that of an AR(1) process, we also estimate a Hasbrouck VAR-based model to identify the size of the trade-related ($stdev_tradecorr_comp_{it}$) and trade-unrelated ($stdev_nontradecorr_comp_{it}$) components of *permanent* price changes in between transactions (see Section 3.3.2 and Hasbrouck (1991a, 1991b) for details). For the regressions, we use the exogenous nonsynchronous introduction of autoquote to instrument for the endogenous $algo_trad_{it}$ to identify causality from algorithmic trading to these LSB and Hasbrouck components. We estimate

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

where M_{it} is a LSB or Hasbrouck component 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. We always include fixed effects and time dummies, but leave out the control variables in the Hasbrouck component regressions. We regress by quintile and report t -values 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				time dummies
	Q1	Q2	Q3	Q4	Q5	$share_turnover_{it}$	$vola_tility_{it}$	$1/price_{it}$	$ln_mkt_cap_{it}$	
<i>Panel A: Effect of AT on effective half spread components cf. Lin, Sanger, and Booth (1995)</i>										
$LSB95_fixed_{it}$	0.26** (3.63)	0.59** (4.16)	0.69** (2.26)	9.91 (0.46)	8.97 (1.36)	2.35** (2.07)	-0.28 (-0.80)	26.23** (3.81)	3.85 (1.29)	Yes
$LSB95_adv_sel_{it}$	-0.26** (-3.46)	-0.61** (-3.80)	-0.84** (-2.14)	-12.19 (-0.46)	-7.72 (-1.32)	-2.58* (-1.85)	0.57 (1.32)	15.70** (1.99)	-4.26 (-1.15)	Yes
$LSB95_order_persist_{it}$	-0.18** (-3.06)	-0.30** (-3.10)	-0.21 (-1.60)	0.66 (0.28)	3.30 (1.21)	-0.82** (-2.33)	0.41** (8.66)	30.68** (6.24)	-0.90 (-1.43)	Yes
<i>Panel B: Effect of AT on trade- and nontrade-correlated component of efficient price innovation cf. Hasbrouck (1991a,1991b)</i>										
$stdev_tradecorr_comp_{it}$	-0.22** (-2.62)	-0.26** (-3.08)	-0.30* (-1.69)	-3.39 (-0.30)	-0.57** (-2.73)					Yes
$stdev_nontradecorr_comp_{it}$	0.13** (2.48)	0.13** (2.36)	0.13 (1.47)	1.03 (0.28)	0.13 (1.12)					Yes

#observations: 1082*167 (stock*day)

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

^a: We use quintile-specific coefficients for the control variables and time dummies. For brevity, we only report the (across the quintiles) market-cap-weighted coefficient for the control variables and its t -statistic.

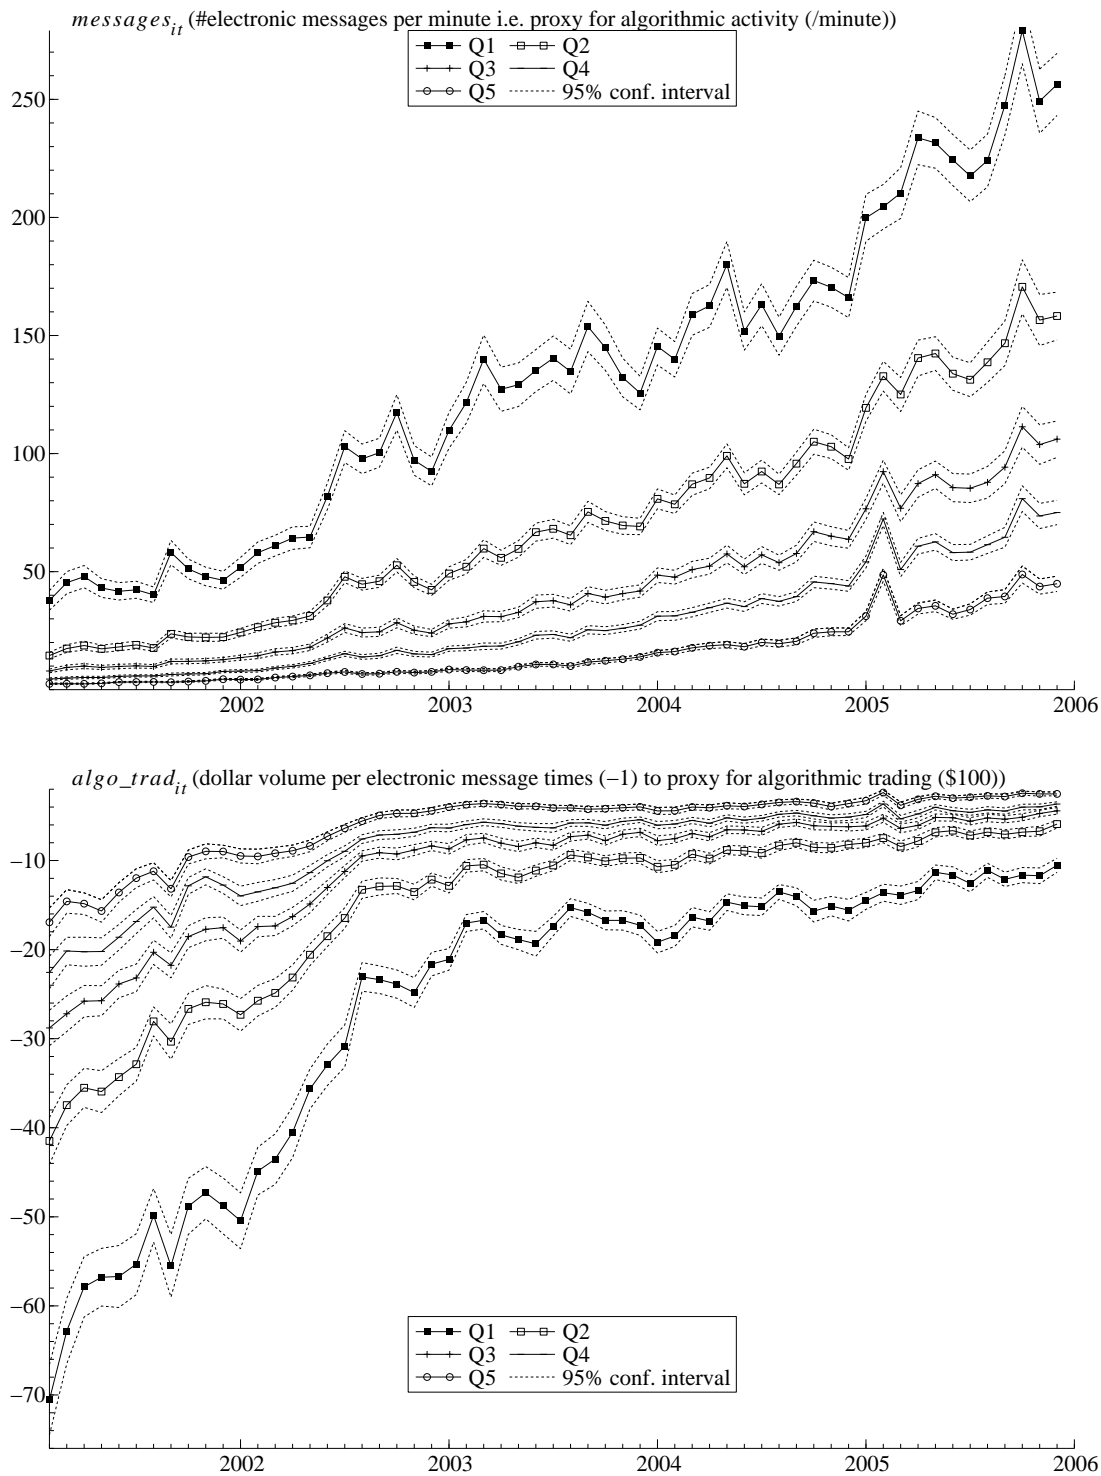


Figure 1: These graphs depict (i) the number of (electronic) messages per minute and (ii) our proxy for algorithmic trading, which is defined as the negative of trading volume (in thousands of dollars) divided by the number of messages. The graphs are done by market-cap quintile, where Q1 is the large-cap quintile.

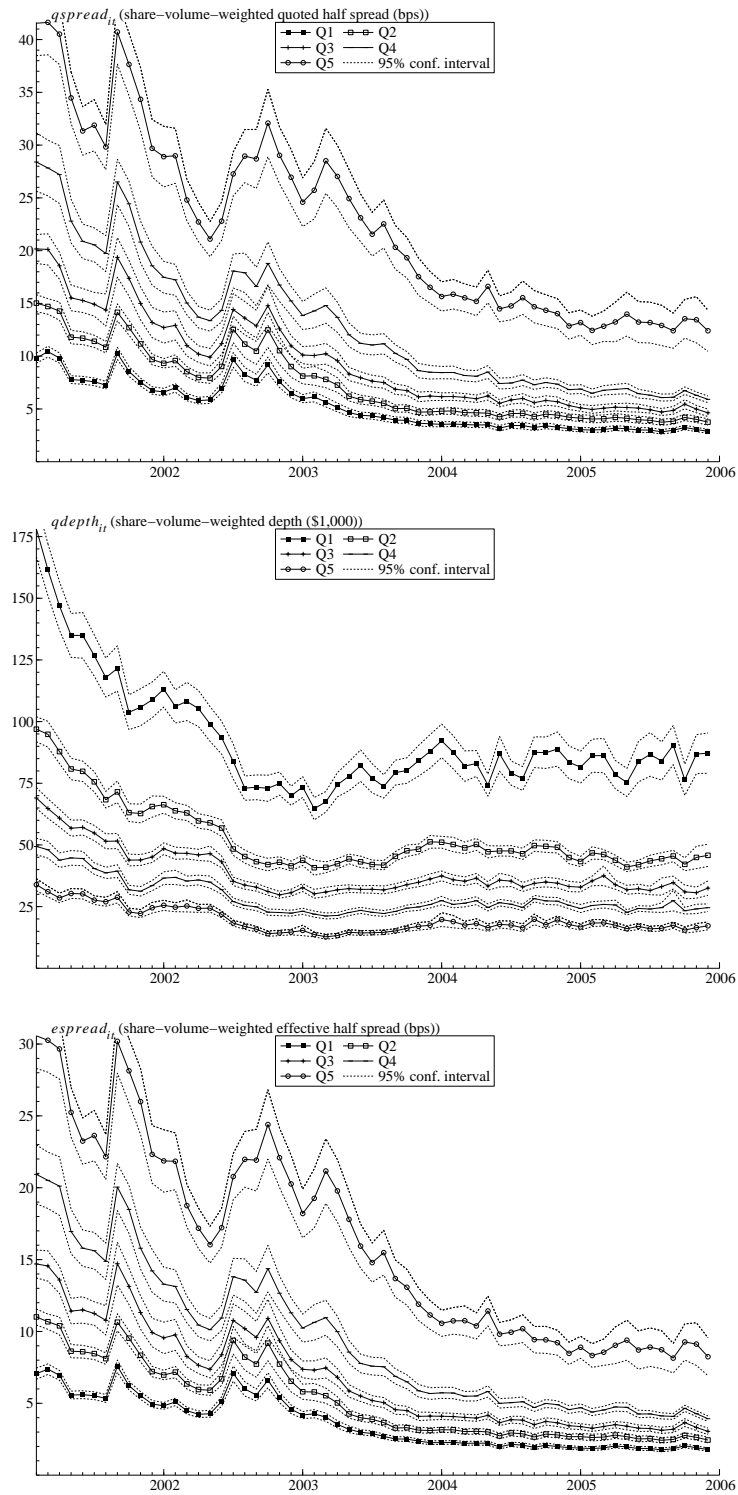


Figure 2: These graphs depict (i) quoted half spread, (ii) quoted depth, and (iii) effective spread. All spread measures are share-volume weighted averages. The graphs are done by market-cap quintile, where Q1 is the large-cap quintile.

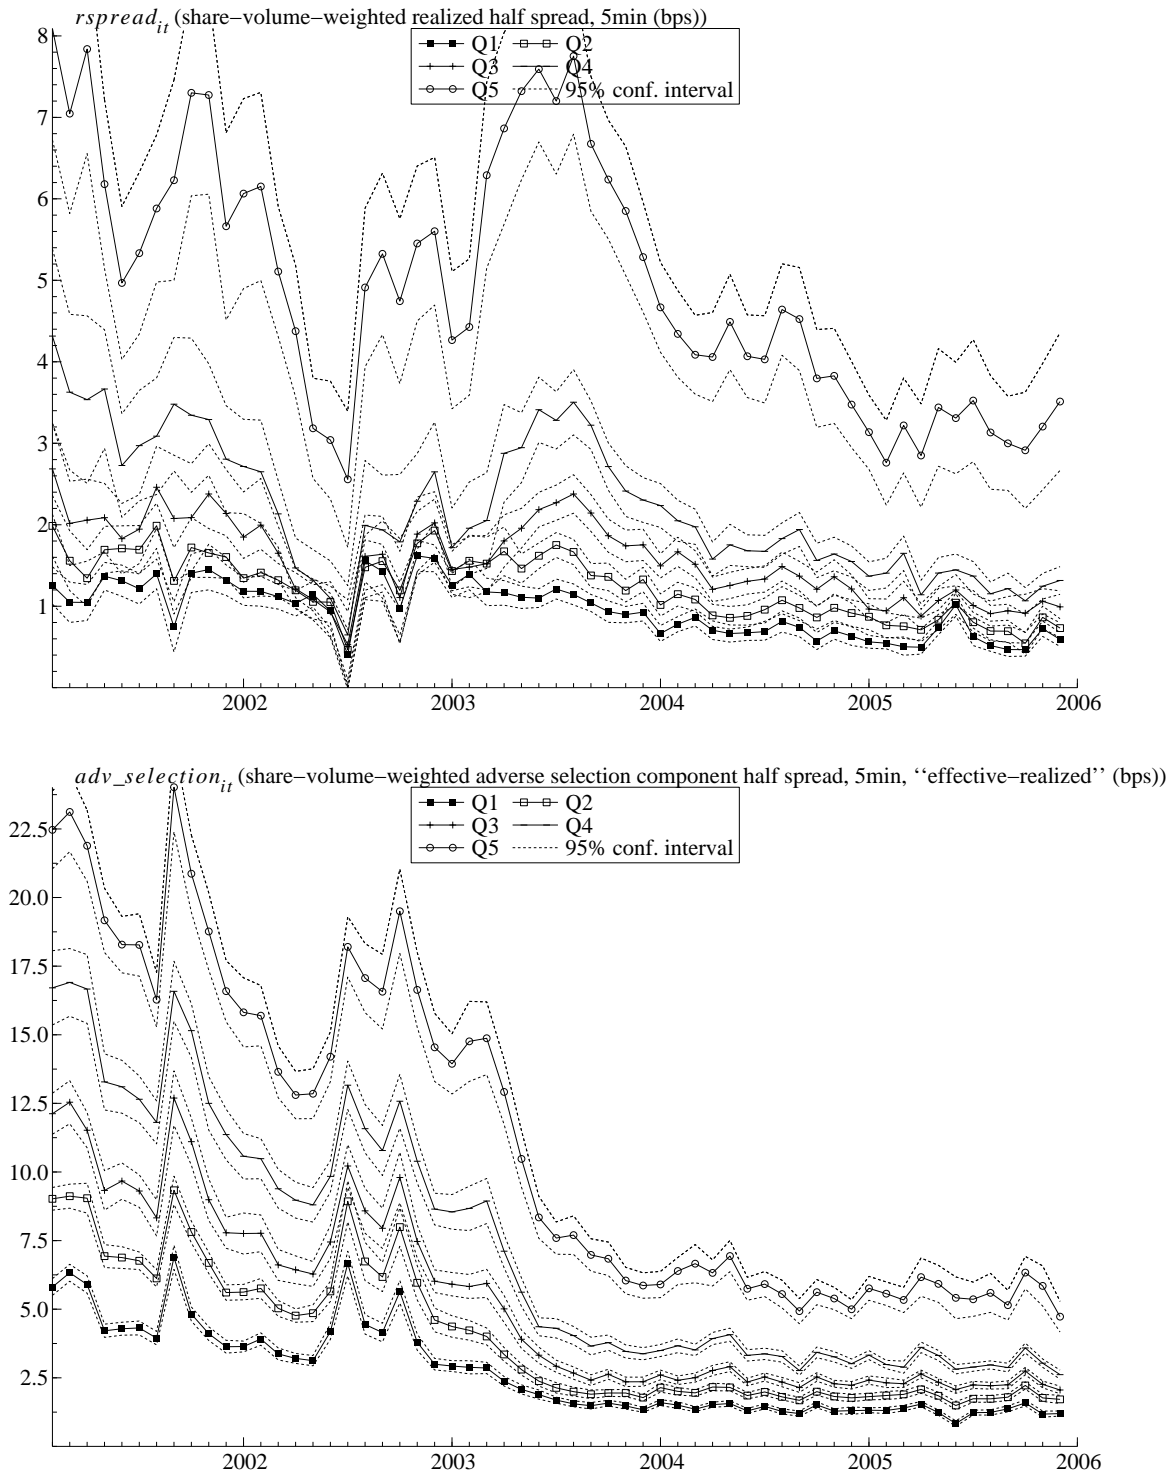


Figure 3: These graphs depict the two components of the effective spread: (i) realized spread and (ii) the adverse selection component, i.e. the (permanent) price impact. The spread decomposition is based on the 5 minute delayed spread midpoint. The graphs are done by market-cap quintile, where Q1 is the large-cap quintile.

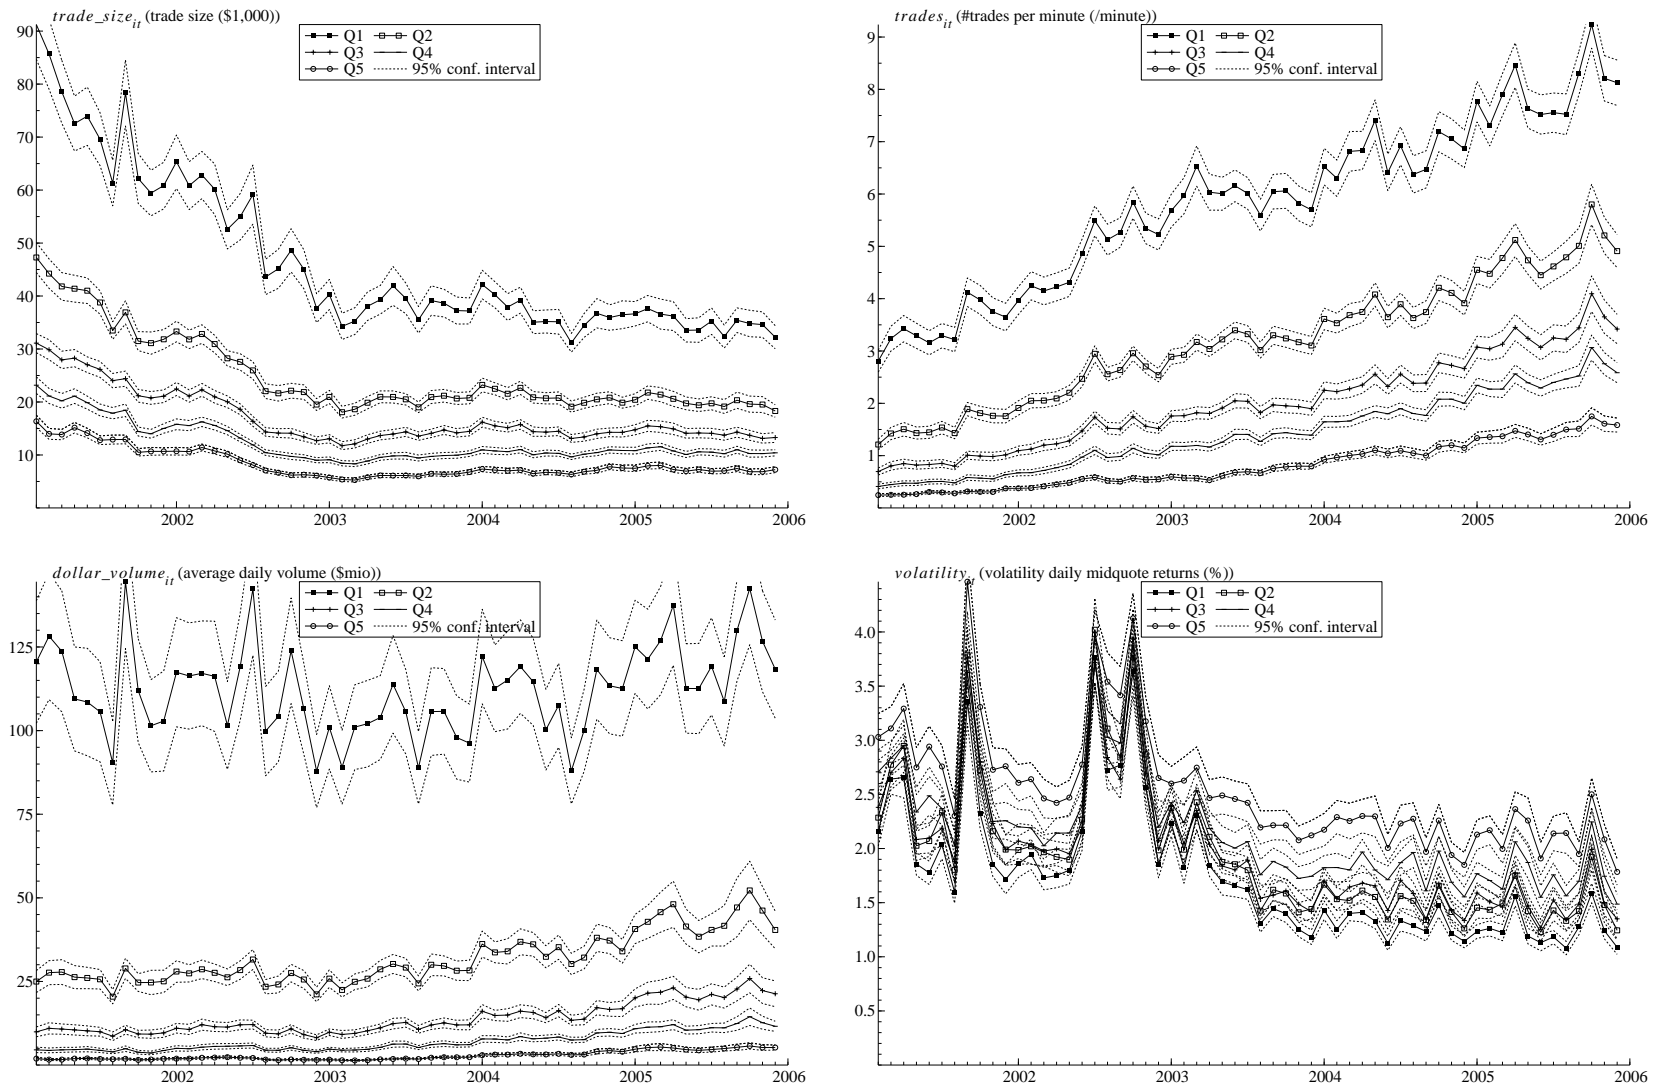


Figure 4: These graphs depict (i) trade size, (ii) the number of trades per minute, (iii) daily dollar volume, and (iv) daily midquote return volatility. The graphs are done by market-cap quintile, where Q1 is the large-cap quintile.

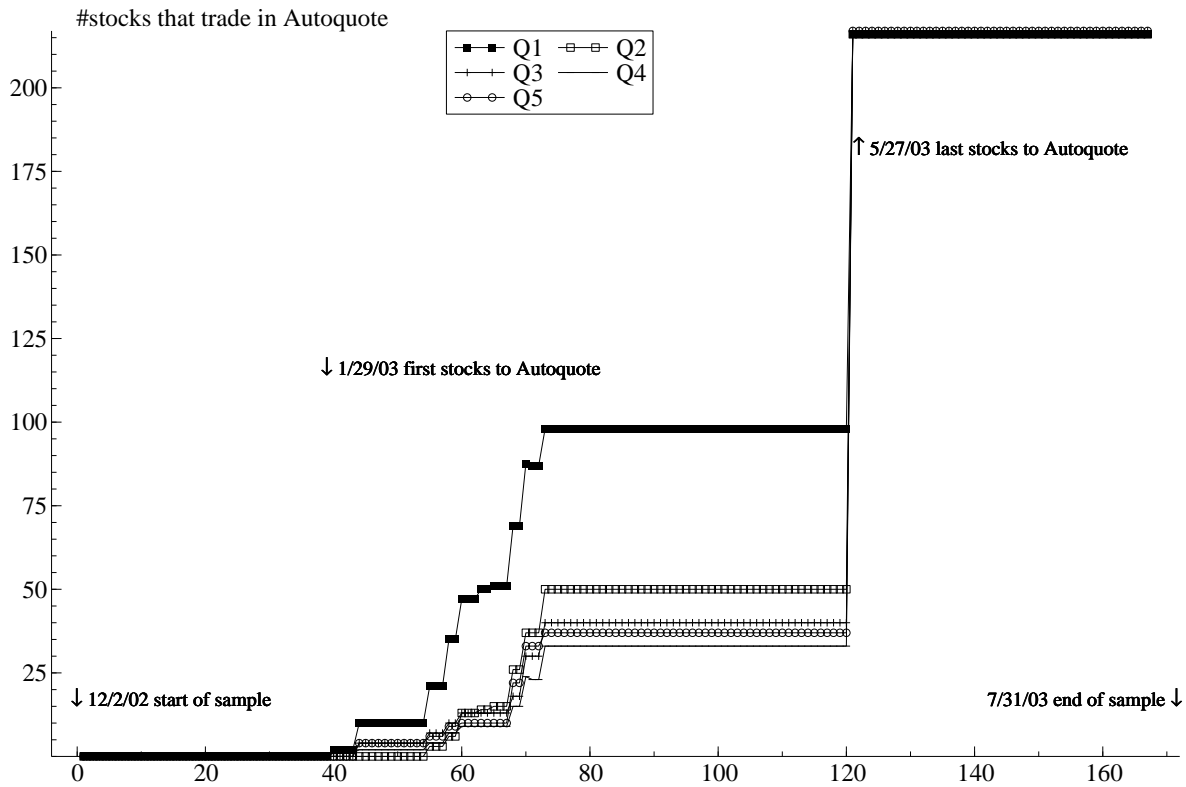


Figure 5: This graph depicts the staggered introduction of autoquote on the NYSE. It graphs the number of stocks in each market-cap quintile that are autoquoted at a given time. Quintile 1 contains large-cap stocks.

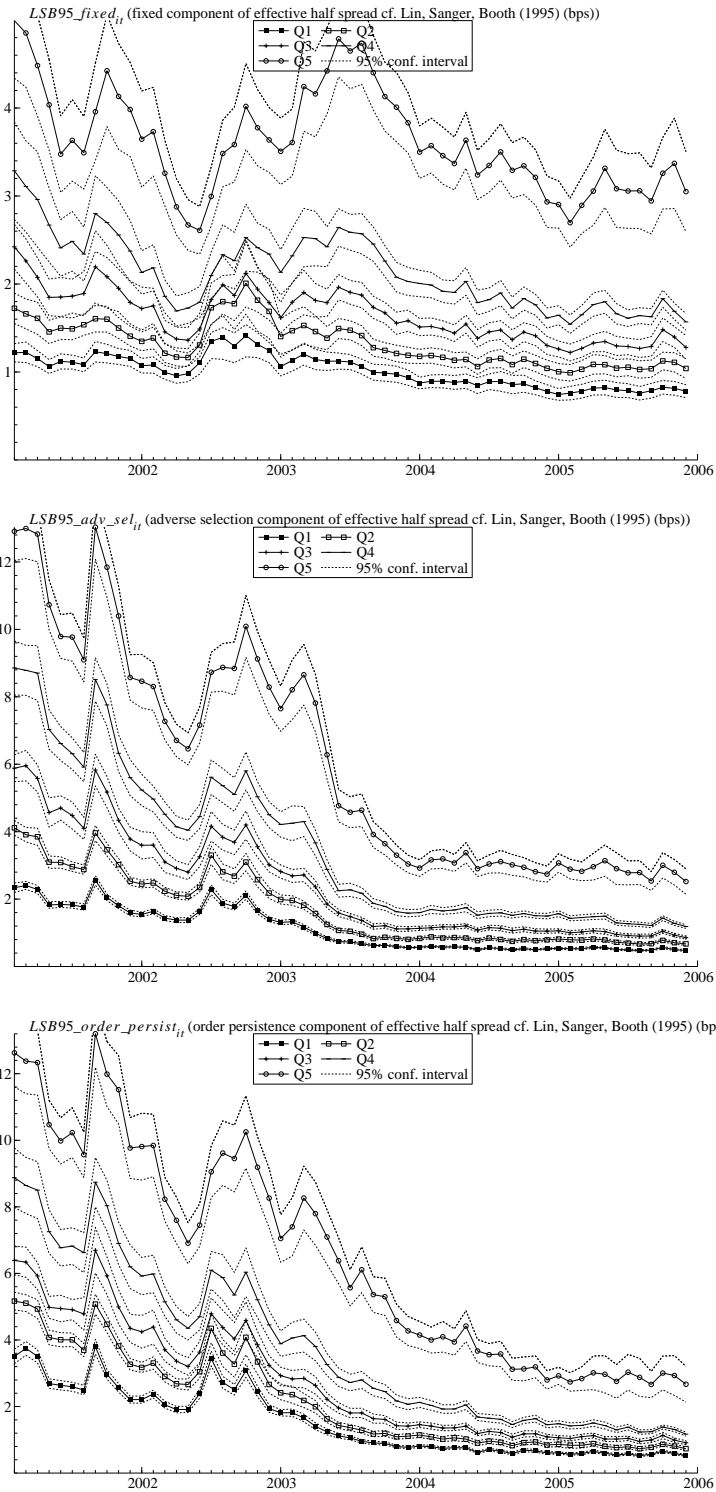


Figure 6: These graphs depict the three components of a Lin-Sanger-Booth spread decomposition, which accounts for order persistence. It identifies a fixed (transitory) component ($LSB95_fixed_{it}$), an adverse selection component ($LSB95_adv_sel_{it}$), and a component due to order persistence ($LSB95_order_persist_{it}$) (see Section 3.3.1 and Lin, Sanger, and Booth (1995) for details). The graphs are done by market-cap quintile, where Q1 is the large-cap quintile.

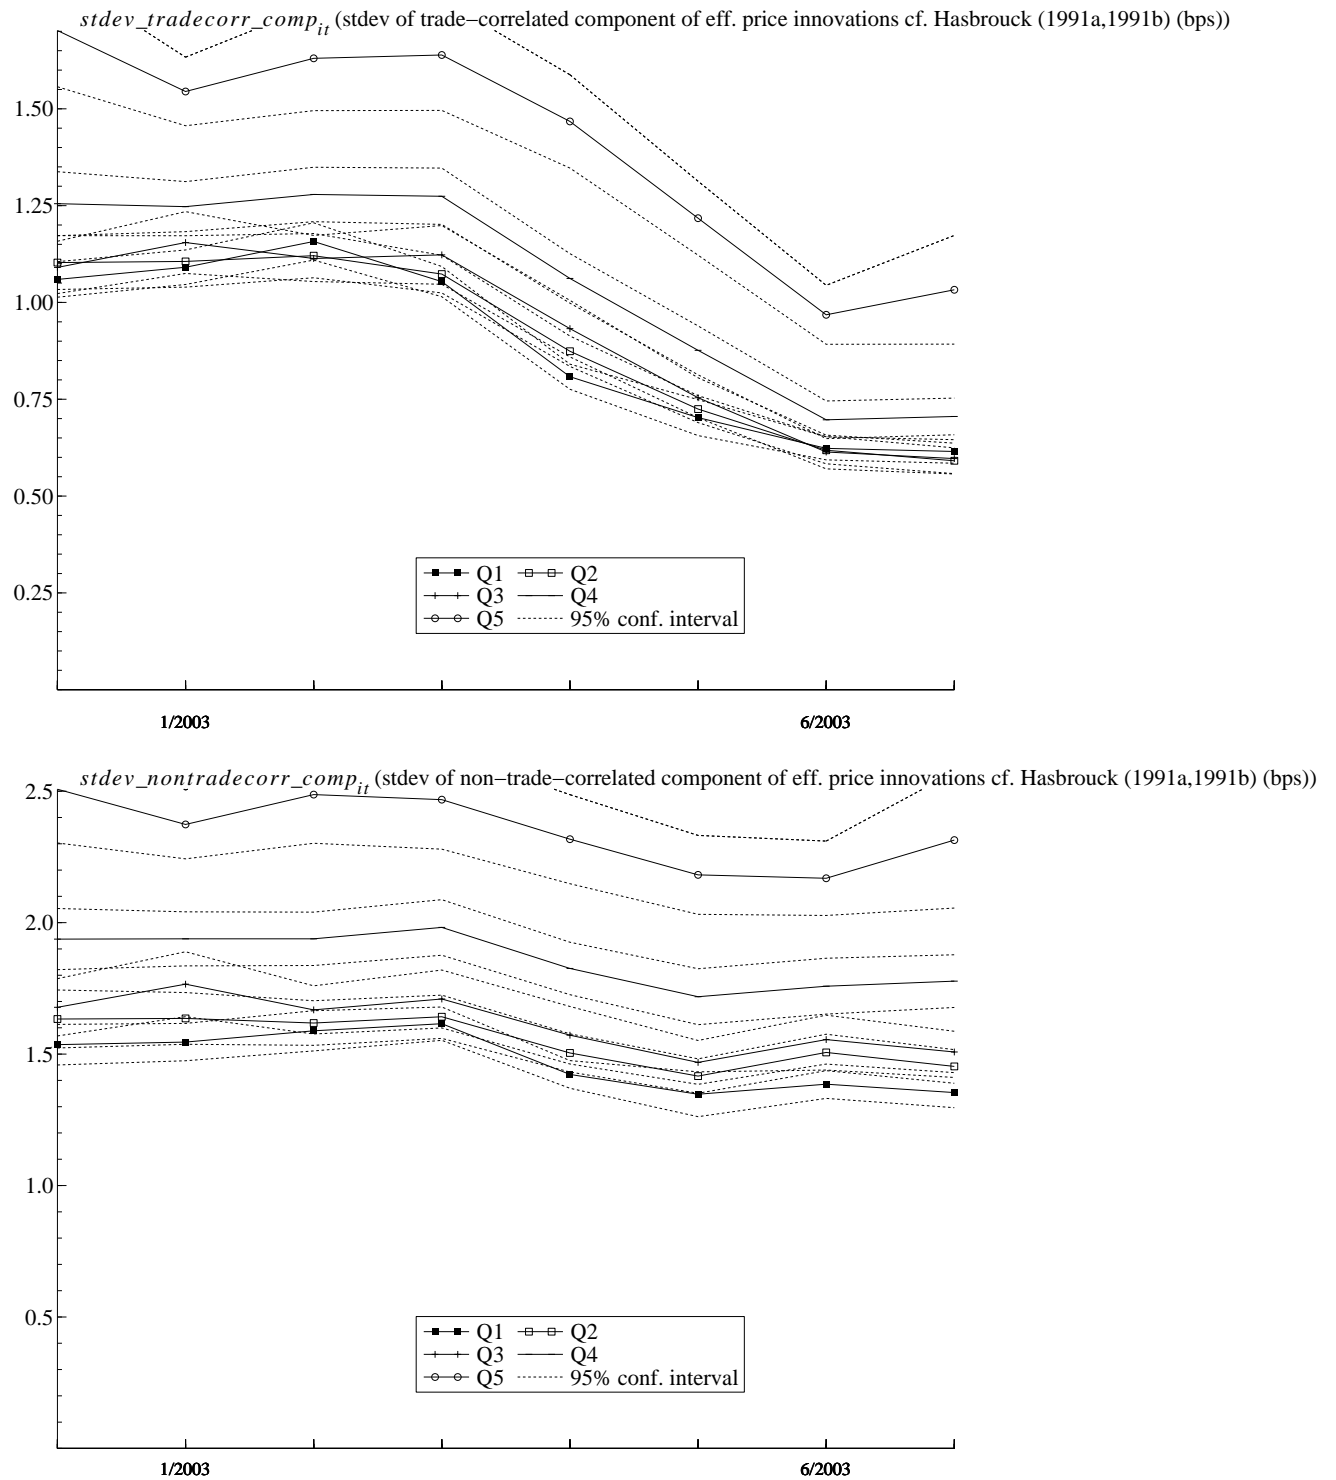


Figure 7: These graphs depict a VAR-based Hasbrouck decomposition of the *permanent* price change in between transactions into a trade-related ($stdev_tradecorr_comp_{it}$) and trade-unrelated ($stdev_nontradecorr_comp_{it}$) component (see Section 3.3.2 and Hasbrouck (1991a, 1991b) for details). The graph depicts the autoquote sample period which runs from December 2002 through July 2003. The graphs are done by market-cap quintile, where Q1 is the large-cap quintile.