Algorithmic Trading and the Market for Liquidity [1]

Terrence Hendershott
Haas School of Business
University of California, Berkeley
hender@haas.berkeley.edu
510-643-0619

Ryan Riordan
Department of Economics and Business Engineering
Karlsruhe Institute of Technology
ryan.riordan@kit.edu
(+49) 721-608-48385

[1]We thank Hank Bessembinder (the editor), an anonymous referee, Bruno Biais, Doug Foster, and seminar and conference participants at the University of Texas at Austin, University of Sydney 4th Annual Microstructure Conference, New York University Courant Institute of Mathematical Sciences 2nd Annual Algorithmic Trading Conference, 2009 Workshop on Information Systems and Economics, 2009 German Finance Association, and IDEI-R Conference on Investment Banking and Financial Markets for helpful comments. Hendershott gratefully acknowledges support from the Net Institute, the Ewing Marion Kauffman Foundation, and the Lester Center for Entrepreneurship and Innovation at the Haas School at UC Berkeley. Riordan gratefully acknowledges support from the Stuttgart Exchange. Data was provided by the Deutsche Boerse and by the Securities Industry Research Centre of Asia-Pacific (SIRCA) on behalf of Reuters.

Algorithmic Trading and the Market for Liquidity

Abstract

We examine the role of algorithmic traders (AT) in liquidity supply and demand in the 30 DAX stocks on the Deutsche Boerse in January 2008. AT represent 52% of market order volume and 64% of nonmarketable limit order volume. AT more actively monitor market liquidity than human traders. AT consume liquidity when it is cheap, i.e., when the bid-ask quotes are narrow, and supply liquidity when it is expensive. When spreads are narrow AT are less likely to submit new orders, less likely to cancel their orders, and more likely to initiate trades. AT react more quickly to events and even more so when spreads are wide.

I. Introduction

Frictions related to investors participation and monitoring of financial markets are important for trading and asset price dynamics (Duffie (2010)). Imperfect monitoring prevents investors from immediately contacting all counterparties. This prolongs search and causes investors to offer greater price concessions to trade quickly, reducing liquidity. Uncertainty in the search process increases liquidity risk. Both the level and uncertainty of liquidity depress prices and lead to misallocations of capital. Technological progress in the form of algorithmic trading (AT; AT denotes algorithmic traders as well) reduces monitoring frictions, which can improve efficiency in the market for liquidity and facilitate gains from trade.[1]

We examine the intersection of AT and investor monitoring for DAX stocks (the 30 largest market capitalization stocks) traded on the Deutsche Boerse (DB) with data identifying whether or not the order was generated with an algorithm. Directly identifying AT is not possible in most markets. We study how technology that lowers monitoring costs affects the market for liquidity supply and demand by characterizing the role of investors with lower costs.

Lower monitoring costs for AT should lead to more frequent activity as AT can react quickly to liquidity supply and demand dynamics resulting from trading and order submissions. For example, the breaking up of large orders into smaller trades that execute when liquidity is high facilitates the search process without revealing the full trading interest. More generally, the ability to monitor and react to events means that trading can be continuously and dynamically optimized, leading AT to consume liquidity when it is expensive and supply liquidity when it is cheap. AT's rapid response to new events should be fastest when liquidity is low.

Algorithms are used to trade in both agency and proprietary contexts (Hasbrouck and Saar (2011)). Institutional investors utilize AT to trade large quantities gradually over time, thereby minimizing market impact and trading costs. Proprietary algorithms, often used for intermediation,

[1]Technology has revolutionized the financial market structure and trading: investors use computers to automate their trading processes and virtually all markets are now electronic limit order books (Jain (2005)). The speed and quality of access to such markets encourages the use of AT.

are usually referred to as high-frequency traders (HFT). HFT use algorithms to quickly process information contained in order flow to identify when a security's price deviates from the efficient price and trade against the deviations. Studying AT facilitates our overall understanding of the importance of technological advances in financial markets and how these affect the frictions in participation and monitoring faced by investors and traders.[2]

AT are identified in our data because of an unusual pricing scheme. Most markets offer volume discounts to attract the most active traders. During our sample period the German competition authority did not allow for generic volume discounts, rather required that discounts have a cost sensitive component. The DB successfully asserted that algorithm-generated trading is lower cost and highly sensitive to fee reductions and, therefore, could receive quantity discounts.[3] The fee rebate program also subsidized the investment in costly AT technology encouraging more investors to automate, boosting trading volume and liquidity at the DB. The DB provided data on AT orders, generated by members in the fee rebate program, in the DAX stocks for the first three weeks of January 2008.

In our sample AT initiate 52% of trading volume via marketable limit orders. AT initiate smaller trades with AT initiating 68% of volume for trades of less than 500 shares and 23% of volume for trades of greater than 10,000 shares. AT cluster their trades together and initiate trade quickly when bid-ask spreads are small. AT are more sensitive to human trading activity than humans are to AT activity. By splitting large orders into smaller slices AT reduce their own market impact but also the volatility of liquidity in general.

AT submit 64% of nonmarketable limit order volume. AT cancel orders more frequently leading to AT supplying liquidity in only 50% of trading volume. When spreads are narrow AT are less likely to submit new orders and less likely to cancel their orders. The net effect of their order

[2]Examining HFT and lower frequency traders separately, which is not possible with our data, can provide insights into AT's application to particular investment and trading strategies.

[3]In December of 2006, the DB introduced its fee rebate program for automated traders. The DB modified the fee rebate program on November 2, 2009, to a volume discount program. This effectively ends the AT specific fee rebate at the DB.

submissions and cancellations leads AT to be at the best bid and offer more often than humans with the difference being more pronounced when spreads are wider. AT cluster their orders together and AT are more sensitive to human order submission activity than humans are to AT order activity.

These results on trading and order submissions are consistent with AT closely monitoring market liquidity supply and demand. The dependency of AT trades and orders on the size of the spread, in terms of activity and speed, shows that their strategy is not random, but rather part of an efficient demand and supply strategy. Better monitoring allows traders to quickly react to changes in market conditions, leading AT to supply liquidity more when spreads are wide and demand liquidity more when spreads are narrow. This reduces uncertainty in liquidity provision thereby reducing liquidity risk.

To extend our examination of AT's monitoring we move beyond the public information contained in the limit order book by studying recent price changes in the futures index market. Futures markets generally lead the underlying stocks, so continuous monitoring of futures price changes is required to prevent limit order from becoming stale. We estimate probit models of AT liquidity demanding and supplying trades and order cancellations controlling for market condition variables incorporating the state of the limit order book, past volatility, and trading volume. We find evidence that AT liquidity demanding trades take advantage of stale limit orders as AT buys are more likely after recent positive index future returns and AT sells are more likely after recent negative index future returns. We find that AT are more likely to initiate trades when liquidity is high in terms of narrow bid-ask spreads. AT liquidity demanding trades are negatively related to volatility and volume in the prior 15 minutes. For liquidity supplying trades AT are more likely to trade when liquidity is low. AT liquidity supplying trades are positively related to prior volatility and negatively related to prior volume. An important component in the supply of liquidity is the ability to continuously monitor market conditions and cancel limit orders in the book to avoid being adversely selected. We find that AT are more likely to cancel orders that are in the opposite direction of recent index future returns, making AT better able to avoid being adversely selected based on public information.

Section II relates our work to existing literature. Section III describes the algorithmic trading on the Deutsche Boerse. Section IV describes our data. Section V analyzes when and how AT demands liquidity. Section VI examines order submissions strategies. Section VII uses multivariate probit analyses to study when AT supply and demand liquidity in transactions. Section VIII concludes.

II. Related Literature

Important implications of investors' monitoring/attention and participation decision for the movement of capital and asset prices are extensively reviewed in Duffie (2010). Electronic limit order markets represent a market for immediacy where limit order submitters offer terms of trade to potential market orders. Limit order submitters must either continuously monitor the market for changing conditions or face being taken advantage of by later arriving traders. Parlour and Seppi (2008) provide a general survey on limit order markets and the importance of the monitoring friction.

Foucault, Roëll, and Sandas (2003) study the equilibrium level of effort that liquidity suppliers should expend in monitoring the market. AT lowers the cost of this kind of monitoring and the adjustment of limit orders in response to market conditions.[4] The monitoring of the state of liquidity in the market and taking it when cheap and making it when expensive is consistent with AT playing an important role in the make/take liquidity cycle modeled by Foucault, Kadan, and Kandel (2008). If AT lowers the costs of monitoring, then frictions may be reduced and Rosu (2009)'s modeling of limit orders as being constantly adjusted is a reasonable simplification for theoretical modeling.

Due to the difficulty in identifying AT, initial research directly addressing AT used data from brokers who sell AT products to institutional clients. Engle, Russell, and Ferstenberg (2007)

[4]Biais, Hombert, and Weill (2010) theoretically examine the relation between AT, market monitoring, and liquidity dynamics under limited cognition. See Biais, Foucault, and Moinas (2011) and Pagnotta and Philippon (2011) for models where investors compete on their trading algorithm's speed. Monitoring also has important cross market competition implications as in Foucault and Menkveld (2008) and others.

use execution data from Morgan Stanley algorithms to study the tradeoffs between algorithm aggressiveness and the mean and dispersion of execution cost. Domowitz and Yegerman (2005) study execution costs of ITG buy-side clients, comparing results from different algorithm providers.

Several recent studies use comprehensive data on AT. Chaboud, Chiquoine, Hjalmarsson, and Vega (2009) study the development of AT in the foreign exchange market on the electronic broking system (EBS) in three currency pairs euro-dollar, dollar-yen, and euro-yen. They find little relation between AT and volatility. Chaboud, Chiquoine, Hjalmarsson, and Vega (2009) find that AT seem to follow correlated strategies, which is consistent with our results on AT clustering together in time. Hendershott, Jones, and Menkveld (2011) use a proxy for AT, message traffic, which is the sum of order submissions, order cancellations, and trades. Unfortunately, such a proxy makes it difficult to directly examine when and how AT behave and their role in liquidity supply and demand. Hendershott, Jones, and Menkveld (2011) use an instrumental variable to show that AT improves liquidity and makes quotes more informative. Our results on AT liquidity supply and demand show the channels by which AT could lead to more liquid markets.

Algorithms are used by traders who are trying to passively accumulate or liquidate a large position. Bertsimas and Lo (1998) find that the optimal dynamic execution strategies for such traders involves optimally breaking orders into pieces so as to minimize cost.[5] While such execution strategies pre-dated the wide-spread adoption of AT (cf. Keim and Madhavan (1995)), brokers now automate the process with AT products.

For each component of the larger transaction, a trader (or algorithm) must choose the type and aggressiveness of the order. Cohen, Maier, Schwartz, and Whitcomb (1981) and Harris (1998) focus on the simplest static 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, Smith, Turnbull, and White (2000) and Ranaldo (2004)). Lo and Zhang (2002) find that execution times are very sensitive to the choice of limit price. If limit orders do not execute, traders can

[5] Almgren and Chriss (2000) extend this by considering the risk that arises from breaking up orders and slowly executing them.

cancel them and resubmit them with more aggressive prices. A short time between submission and cancellation suggests the presence of AT. Hasbrouck and Saar (2009) find that a large number of limit orders were canceled within two seconds on the INET trading platform (which is now Nasdaq's trading mechanism).

A number of papers analyze the high-frequency trading subset of AT. Biais and Woolley (2011) provide background and survey research on HFT and AT. Brogaard (2010) examines a number of topics in HFT. Brogaard (2011) studies the activity of HFT and finds that HFT are more likely to demand liquidity when spreads are wide and supply liquidity when spreads are narrow, consistent with our findings for AT. Hendershott and Riordan (2011) study the role of overall, aggressive, and passive HFT trading in the permanent and transitory parts of price discovery. Kirilenko, Kyle, Samadi, and Tuzun (2011) analyze HFT in the E-mini S&P 500 futures market during the May 6, 2010 flash crash. Jovanovic and Menkveld (2011) model HFT as middlemen in limit order markets and study their welfare effects. Menkveld (2011) shows how one HFT firm enabled a new market to gain market share.

III. Deutsche Boerse's Automated Trading Program

The Deutsche Boerse's order-driven electronic limit order book system is called Xetra (see Hau (2001) for details).[6] Orders are matched using price-time-display priority. Quantities available at the 10 best bid and ask prices and the number of participants at each level are disseminated continuously. See the Appendix for further details on Xetra.

During our sample period Xetra had a 97% market share of German equities trading. With such a dominant position the competition authorities (Bundeskartellamt) required approval of all fee changes prior to implementation. The criteria used to evaluate fee changes were: (i) all participants are treated equally; (ii) changes must have a cost-related justification; and (iii) fee changes are transparent and accessible to all participants. Criterion (i) and (iii) ensure a level playing field for

[6] Iceberg orders are allowed as on the Paris Bourse (cf. Venkataraman (2001)).

all members and are comparable to regulation in the rest of Europe and North America. The second criteria is the most important for this paper. AT were viewed as satisfying the cost justification for the change, so the DB could offer lower trading fees for AT.[7]

In December of 2007 the DB introduced its Automated Trading Program (ATP) to increase the volume of automated trading on Xetra. By offering fee rebates the DB was implicitly subsidizing investment in AT technologies. To qualify for the ATP an electronic system must determine the price, quantity, and submission time for orders. In addition, the Deutsche Boerse ATP agreement required that: (i) the electronic system must generate buy and sell orders independently using a specific program and data; (ii) the generated orders must be channeled directly into the Xetra system; and (iii) the exchange fees or the fees charged by the ATP member to its clients must be directly considered by the electronic system when determining the order parameters.

Before being admitted to the ATP, participants were required to submit an high-level overview of the electronic trading strategies they plan to employ. The level of disclosure required was intended to be low enough to not require ATP participants to reveal important details of their trading strategies. Following admission to the ATP, the orders generated by each participant were audited monthly for plausibility. If the order patterns generated did not match those suggested by the participant's strategy plan submitted or were considered likely to have been generated manually, the participant was terminated from the ATP and possibly suspended from trading on Xetra. The ATP agreement and the auditing process ensure that most, if not all, of the orders submitted by an ATP participant are electronically generated and that most, if not all, electronically generated orders are included in our data.[8]

[7]The logic was that electronic order generation by algorithms could be less costly for an exchange, making lower fees justifiable. This is debatable. At the end of 2009 the ATP program was ended and lower AT fees were replaced by the volume discounts used in most other markets.

[8]Conversations with the DB revealed that a small portion of AT orders may not be included in the data set. The suspicion on the part of the DB is due to the uncommonly high number of orders (message traffic) to executions of certain participants which is typical of AT. However, these participants make up less than 1% of trades in total and are, therefore, unlikely to affect our results.

The DB charges fees for executed trades and not for submitted orders. The rebate for ATP participants could be significant and increased with the total trading volume per month. The first Euro volume rebate level began at 250 million in Euro volume and rebated 7.5% of fees. Rebates rose to a maximum of 60% for euro monthly volume above 30 billion. Table 1 provides an overview of the rebate per volume level.

[Insert Table 1 Here]

For an ATP participant with 1.9 billion euro volume, the percentage rebate was:

$$(250 * 0\% + 250 * 7.5\% + 500 * 15.0\% + 900 * 22.5\%)/1,900 = 15.6\% \tag{1}$$

In the example above, an ATP participant received a rebate of 15.6% of fees. This translates into roughly 14,000 euros in trading cost savings on 91,200 in total fees plus an additional 5,323 euro savings on 61,500 in total in clearing and settlement costs. This rebate (14,000 + 5,323) translates into a 0.1 basis point saving on the 1.9 billion in turnover. For high-frequency trading firms, whose turnover is much higher than the amount of capital invested, these savings are significant.

The fee rebate for ATP participants was the only difference in how orders were treated. AT orders were displayed equivalently in the publicly disseminated Xetra limit order book. The Xetra matching engine did not distinguish between AT and human orders. Therefore, there were no drawbacks for an AT firm to become an ATP participant. Thus, we expect all AT took advantage of the lower fees by becoming ATP participants. From this point on we equate ATP participants with algorithmic traders and use AT for both. We will refer to non-ATP trades and orders as human or human-generated.

IV. Data and Descriptive Statistics

The DB provided data contain all AT orders submitted in DAX stocks, the leading German stock market index composed of the 30 largest and most liquid stocks, between January 1st and January

18th, 2008, a total of 13 trading days. This is combined with Reuters DataScope Tick History data provided by the Securities Industry Research Centre of Asia-Pacific (SIRCA). The SIRCA data contains two separate datasets, one for transactions and another for order book updates.

We generate a dataset similar to Biais, Hillion, and Spatt (1995). Using the orderbook snapshots we recreate the order causing the observed book changes. To identify trades, which, when observing order book updates, are similar to cancels at the best, we match these orders with the public transaction dataset. We match the generated events (insert, cancel, trade) with the DB provided AT order dataset. As in Biais, Hillion, and Spatt (1995) we truncate the orderbook at the first 5 levels to focus on orders closest to the best bid and ask. The resulting data contains all orders submitted by AT and humans at the first 5 levels on the bid and ask sides of the book.

Table 2 describes the 30 stocks in the DAX index. Market capitalization is as of December 31st, 2007, in billions of Euros.[9] The smallest firm (TUI AG) is large at 4.81 billion Euros but is more than 20 times smaller than the largest stock in the sample, Siemens AG. The standard deviation of daily returns is calculated for each stock during the sample period. All other variables are calculated daily during the sample period for each stock (30 stocks for 13 trading days for a total of 390 observations). Means and standard deviations along with the minimum and maximum values are reported across the 390 stock-day observations.

[Insert Table 2 Here]

DAX stocks are quite liquid. The average trading volume is 250 million euros per day with 5,344 trades per day on average. This implies that our data set contains roughly 2 million transactions (5,344*390). Quoted half-spreads are calculated when trades occur. The average quoted half-spread of 2.98 basis points is comparable to large and liquid stocks in other markets. The effective spread is the absolute value of the difference between the transaction price and the mid quote price (the

[9] Firms' market capitalization is gathered from the Deutsche Boerse website and cross-checked against data posted directly on each company's website.

average of the bid and ask quotes). Average effective spreads are only slightly larger than quoted spreads, evidence that market participants seldom submit marketable orders for depth at greater than the best bid or ask.

We measure depth in two ways. The first is the standard measure of the depth at the inside quotes: the average depth in euros at the best bid price and the best ask price. As with spreads, depth is measured at the time of transactions. More depth allows traders to execute larger trades without impacting the price, corresponding to higher liquidity. However, if the width of the spread varies over time, then comparisons of depth at the inside do not clearly correspond to levels of liquidity, e.g., 50,000 euros at an inside spread of 10 basis points need not represent more liquidity than 5,000 euros at an inside spread of 5 basis points if in the latter case there is sufficient additional depth between 5 and 10 basis points. To account for time variation in the spread we calculate a second depth measure using the limit order book. For each stock we aggregate the depth at bid and ask prices that have a distance of less than three times that stock's average quoted half-spread from the quote midpoint at the time of transaction. We refer to this measure of depth that does not depend on the spread at the time of the transaction as depth3. A similar measure is used in Foucault and Menkveld (2008) to capture depth away from the best prices.

V. Trading

Trades represent liquidity demand and are arguably the most important events in limit order markets. Trades allow investors to manage risk and adjust their portfolio throughout the trading day and they are not subject to later cancellation as with non-marketable limit orders. Large liquidity demanding orders placed during periods of low liquidity can have substantial price impact and disrupt market liquidity and stability for long periods of time. Breaking the same order up into smaller pieces and submitting these conditional on market conditions can reduce the negative impact of the overall order. Therefore, AT's better monitoring should lead their trades to be more sensitive to market conditions than human trades.

To measure AT liquidity demand we create marketable order (trade) and limit order variables for AT and human, labeled AT and HUM, respectively. The AT variable takes the value 1 when a trade or order is from an AT, and is 0 otherwise. The HUM variable takes the value 1 when a trade or order is from a human and 0 otherwise. Panel A of Table 3 reports the fraction of euro trading volume for AT initiated trades by trade size and overall. For simplicity and comparability we use the U.S. SEC Rule 605 trade-size categories based on the number of shares traded. Panel B of Table 3 reports the fraction of trades initiated by AT by trade size and overall. Overall AT initiates 52% of euro volume and more than 60% of all trades. AT initiation declines as trade size increases. AT is greater than 68% and 57% in the two smallest trade-size categories (0-499 shares and 500-999 shares) and decreases to 23% in the largest trade-size category (10,000+ shares). AT's decline with trade size is consistent with several possibilities: AT being used to breakup large orders into smaller trades as suggested by Bertsimas and Lo (1998) and high-frequency traders using tight risk-management strategies as in Menkveld (2011). AT's use of smaller liquidity demanding trades may help to reduce the volatility of liquidity.

[Insert Table 3 Here]

Panels A and B of Table 4 provide the same statistics as Table 3 for non-marketable limit order submissions. The AT share of limit orders is substantially higher than its share of trades, 64% versus 52%. This difference declines in trade size. For nonmarketable limited orders that eventually execute, AT represent only 50% of volume. AT submit more orders than execute either by submitting uncompetitive orders away from best prices or by canceling and replacing orders close to the best prices.

[Insert Table 4 Here]

Because non-marketable limit order submissions are reversible and because limit orders vary by the aggressiveness of their prices we first focus on transactions before turning to order submission strategies more generally. To better understand how monitoring affects liquidity demand conditional on past trading we perform a series of analyses similar to those found in Biais, Hillion, and Spatt (1995) for AT and human trades. Examining AT and human trades separately doubles the number of variables, requiring some adaptations. First, we report the results of two separate and related analyses in Table 5. The first column of Table 5, labeled unconditional, provides what fraction of trades sequences, i.e., AT followed by AT, AT followed by human, etc., we expect if AT and human trades are randomly ordered. The other columns in Panel A are essentially a contingency table documenting the probability of observing a trade of a specific type after observing a previous trade of a given type. Rows sum up to 100% and can be interpreted as probability vectors.

[Insert Table 5 Here]

The first row in Table 5 shows that if AT and human trades were randomly ordered 37% of the transactions would be AT followed by AT while in the data this occurs 40.7% of the time. This shows that AT trades are more likely to follow AT trades than we would expect unconditionally. In addition, AT trades are more likely to be repeated on the same side of the market. The same is true for human trades. This suggests that human and AT liquidity demanding trading strategies execute at different times while having related characteristics. Table 5 also shows AT to be relatively more sensitive to human order flow than humans are to AT order flow. The conditional probability of AT following a human trade is 51.4% as compared to 66.9% following AT. Human trades follow AT with a conditional probability of 33.1% as compared to 48.6% following a human trade.

Table 6 extends the analysis of AT and humans trade sequences to include trade-size categories. As in Biais, Hillion, and Spatt (1995) we highlight in bold the three largest values in a column to illustrate the interdependence of trade sequences. The results are similar to the diagonal results reported in Biais, Hillion, and Spatt (1995) and predicted theoretically in Parlour (1998). The diagonal finding implies that trades of the same type—AT or human trades in the same trade-size category—follow other similar trades. This leads to a diagonal effect where the highest probabilities lie on the diagonal. The largest probability by far is for small AT trades: the $AT^1_{t-1}AT^1_t$ probability of 48.7% is much higher than the unconditional probability of 31.6%. This suggests that: (i) AT repeatedly use small trades to hide their information; (ii) AT limit their transitory price impact; or (iii) that different AT are following related strategies. All of these strategies are consistent with AT closely monitoring market conditions.

[Insert Table 6 Here]

VI. Order Submissions

AT's greater sensitivity to past trading activity is consistent with better monitoring and lower frictions in the trading process. AT monitoring should enable their trading and entire submission strategy, including liquidity provision, to incorporate the most current information on other traders' orders in the limit order book. This should lead their trades, order submissions, and order cancellations to be more sensitive to past orders and the current state of the limit order book, e.g., the bid-ask spread, than human submission strategies.

Table IV in Biais, Hillion, and Spatt (1995) examines orders and trades conditional on the prior order or trade. Tables 5 and 6 study this using our data for AT and human trades. Table 7 incorporates order submissions. To make the table size manageable we narrow the scope of trade sizes and orders relative to Biais, Hillion, and Spatt (1995) by using one trade size and not including limit order submissions away from the best prices.

[Insert Table 7 Here]

Table 7 shows that the diagonal effect for trades also holds for orders as actions of similar type are more likely to be repeated. As in Tables 5 and 6 AT reacts more to human orders than humans respond to AT orders. This can be seen visually in the pattern of bold numbers which are more prevalent in the lower left quadrant than the upper right quadrant. The conditional probabilities can also be calculated across various human orders and AT orders. Similar to the calculations above for Table 5, using all order types the relative difference in conditional probabilities for AT following AT versus human is smaller than the difference in conditional probabilities for humans following human versus AT.

Tables 5-7 provide evidence on the clustering and interdependence of AT and human trades and orders. To study how the monitoring of market conditions captured in the limit order book affect AT and human order submissions Table 8 examines order frequencies conditional on the bid-ask spread. As in Biais, Hillion, and Spatt (1995) spread-size categories are calculated for each stock separately. Large spreads are when spreads are in their widest quartile for that stock whereas small spreads are the lowest quartile.

[Insert Table 8 Here]

Panel A of Table 8 provides the order frequencies for AT and humans for each spread category. Within each spread category the order frequencies across AT and humans sum to 100. Panel B calculates the order frequency differences between small spreads and large spreads. We calculate the order frequencies for each stock each day. Statistical inference is conducted across the 390 stock-day observations controlling for contemporaneous cross-sectional correlation and within stock correlation by double clustering standard errors on stock and day as suggested by Petersen (2009) and Thompson (2011).

Overall AT order activity is greater in all spread categories and the AT-human difference increases in the spread with AT orders representing 66.4% of orders when spreads are small and 75.3% when spreads are large. AT orders worse than the best bid and ask prices are not sensitive to the spread whereas AT limits orders at or within the best prices become more frequent as spreads increase: AT orders at or inside the best prices make up 16.3% of orders when spreads are small and 22.7% of orders when spreads are wide.[10] AT trade initiations decline as spreads widen, falling from 10.3% of orders when spreads are narrow to 3.6% of orders when spreads are large. Because order frequencies sum to 100 the human frequencies generally decline as spreads widen.

These results are consistent with AT monitoring market conditions to optimize their liquidity supply and demand activities. There are several possible underlying strategies which would account for the increase in new AT liquidity supply and the decline in AT liquidity demand as spreads widen. The same algorithm could switch from supply to demand as spreads narrow. Alternatively, liquidity supply and liquidity demand algorithms could be entirely separate,[11] but both types of algorithms are sensitive to market conditions related to the competition for liquidity supply and demand. For example, a liquidity demanding algorithm could increase its trade initiations when spreads are tighter and a liquidity supplying algorithm could increase its limit order submissions when spreads widen. Without data identifying specific algorithms we cannot establish which of these strategies drives the results, but we expect it is likely that all are present. These represent healthy competitive responses to market conditions and likely lead to less volatility in liquidity supply and demand.

The order frequencies in Table 8 demonstrate the sensitivity of AT orders to market conditions. However, calculating frequencies does not capture how the rate at which events occur depends on market conditions. To better understand the impact of monitoring on time in order submission dynamics we continue to follow Biais, Hillion, and Spatt (1995) in Table 9 to report the average time between orders conditional on the same spread categories as in Table 8. For brevity we do not report details for each spread-size categories, but report the overall time between events and the small-large spread differences. The small spread times between events are equal to the average times plus one half of the small-large differences.

[Insert Table 9 Here]

[10] It is interesting to note that AT is less likely to cancel orders at the best price when spreads are narrow. This could represent the value of time priority when spreads are lower.

[11] See Menkveld (2011) for an example of a high-frequency algorithm that almost exclusively supplies liquidity.

Panel A of Table 9 gives results for the times between two trades, between the spread narrowing and a trade, and between the spread widening and an order narrowing the spread. On average 6.78 second pass between when a trade occurs and the next AT trade. Human trades are less frequent and consequently there is an average of 9.27 seconds from the time of a trade until a human trade. Both AT and humans trade more quickly when spreads are narrower. The human time decreases by 2.34 seconds from large to small spreads versus a decline of 4.36 seconds for AT. The difference-in-difference of 2.02 seconds between AT large-spread minus AT small-spread and human large-spread minus human small-spread is statistically significant and economically large. Panel A of Table 9 also shows that AT trade more quickly than humans following the spread narrowing. These results are consistent with the frequency results in Table 8 suggesting that AT more actively monitor market conditions when demanding liquidity.

The make/take liquidity cycle examined in Foucault, Kadan, and Kandel (2008) alternates between two phases. First, an order from a liquidity supplier narrows the spread by offering a better price. Second, a liquidity demander monitors the market and reacts to the narrow spread by initiating a trade. The trade causes the spread to widen and the cycle repeats. The above discussion of the time from a narrowing order to a trade is consistent with AT having lower monitoring costs in the second phase of the Foucault, Kadan, and Kandel (2008) cycle. To examine the initial phase of the make/take cycle, where a liquidity supplier monitors the market for a wide quote and offers a better price, Panel A of Table 9 also provides the time between the spread widening due to a trade or cancellation and an order narrowing the spread. AT react faster to a widening spread with the difference increasing in the width of the spread. This is consistent with AT attempting to capture the liquidity supply profits in the Foucault, Kadan, and Kandel (2008) make/take liquidity cycle.

Monitoring the limit order book and the resulting liquidity cycles are manifestations of search frictions for investors seeking gains from trade. Lower monitoring costs implicitly lowers search costs. Lower search costs typically result in greater competition among traders due to lower bargaining frictions. In limit order books these frictions are a form of market power and are a component of the bid-ask spread. Therefore, lower monitoring by AT can lead to better liquidity as found by Hendershott, Jones, and Menkveld (2011). However, the benefits of AT are not necessarily equally distributed between AT and humans.

Table 8 documents AT and human activity conditional on spreads. We see that AT are more likely to submit a new bid/ask at the inside or better overall and when spreads are wide. Because AT are also more likely to cancel their orders Table 8 is not fully informative about the overall impact of AT activity (inserts and cancels) on the best quotes over the entire trading day. In Panel B of Table 9 we report the number of seconds AT and humans spend alone at the best bid/ask across the trading day. Panel B of Table 9 shows that AT are on average at the inside for almost 1 hour more per day than humans. The small-large spread difference examines whether or not AT are more likely to be present at the inside when spreads are wide or narrow. As before for each stock we identify times when spreads are wider and narrower that stock. We then calculate the amount of time AT and humans are on the inside during the large- and small-spread times. Table 9 shows that AT are at the inside more often during both large- and small-spread periods, but the AT-human difference is significantly higher during the large-spread periods. This shows that AT are more likely to offer to supply liquidity when it is expensive.

For AT to be on the inside more often yet only provide liquidity for 50% of volume AT orders must be smaller or times when humans are alone at the inside are more likely to have transactions. One natural explanation for trades occurring more often when humans are alone at the inside quote is that the human quote is stale and is picked off. Our next analysis of liquidity supply and demand will move beyond the unconditional and single-dimension conditioning thus far and attempt to examine potentially stale limit orders.

VII. Multivariate Liquidity Supply and Demand Analysis

Up to this point we have studied how the increased ability of AT to monitor the limit order book impacts liquidity demand and supply dynamics. To broaden our examination of AT's differential reaction to public information, we move beyond the limit order book itself by studying recent price changes in the futures index market. To attempt to measure public information about when unexecuted limit orders may be stale we use the fact that index futures price changes typically lead price changes in the underlying stocks, e.g., Kawaller, Koch, and Koch (1987). We measure returns on DAX futures in the 30 seconds prior to each trade.[12] A sell limit order could be thought of as stale if the previous DAX return is positive as the systematic component of the stock price will have increased since limit order placement. If the limit order executes before incorporating changes in the index value, it could be stale.

If futures prices lead the underlying, then sell limit orders after positive futures returns and buy limit orders after negative futures returns are more likely to be stale. To capture these we interact lagged 30-second futures returns with a $BuySell$ indicator variable set to +1 if the trade is buyer initiated and -1 if the trade is seller initiated. To limit the number of estimates while providing information on the potential staleness of the DAX returns we also decompose the 30-second return into the return over the prior second, two to ten seconds earlier, and 11 to 30 second earlier. The lagged futures returns are calculated as $Rtn_{t-x,t-X} = ln(midpoint_{t-x}/midpoint_{t-X})$, where x = 0, 2, and 11 and X = 1, 10, 30. We interact positive and negative futures returns separately with trade direction as follows $Rtn^{+}_{t-1,t-30} * BuySell$ and $Rtn^{-}_{t-1,t-30} * BuySell$.

Because these reflect market-wide factors that may be correlated with the state of the limit order book in each stock, we also account for contemporaneous and lagged liquidity measures and market conditions. Following Barclay, Hendershott, and McCormick (2003) we use the liquidity variables summarized in Table 2 along with past return volatility and trading volume. Lagged volatility is the absolute value of the stock return over the 15 minutes prior to the transaction. Lagged volume is the euro trading volume in the 15 minutes prior to the transaction.

[12]To analyze the lead-lag relationship we calculate the cross autocorrelation of the front month DAX future and DAX index prints at 5-second frequencies. For the future we take the prevailing midpoint on Eurex and for the index Deutsche Boerse uses the last transaction price for each index constituent. The cross autocorrelations of the lagged future (in five seconds intervals) and the contemporaneous index are 0.21, 0.08, 0.04, 0.02, 0.02, and 0.01, all of which are significant at the 1% level. The cross autocorrelations of the lagged index (in five seconds intervals) and the contemporaneous future are 0.06, 0.03, 0.01, 0.01, 0.00, and 0.00, the first four lags are significant at the 1% level.

Table 10 shows the univariate correlations between dummy variables for AT initiated trades, AT_{Init}, trades where an AT non-marketable limit order executes, AT_{Pass}, AT order cancellations, AT_{Cancel}, the futures return variables, and the market condition variables. Consistent with Tables 3 and 4 larger trades are less likely to be initiated by AT. AT_{Pass} is positively correlated with trade size. This reflects the fact that AT_{Pass} captures trades that are entirely supplied by AT and trades that are supplied by both AT and humans because they are large. Consistent with Tables 8 and 9 narrower spreads are positively correlated with AT initiated trades and negatively correlated with passive AT trades and AT cancellations.

[Insert Table 10 Here]

Table 11 reports coefficients estimates from probit regressions for AT initiated trades, passive trades, and AT cancellations along with their corresponding linear probability slopes and p-values. To control for stock effects and time of day effects, we include, but do not report, firm dummy variables (30) and time of day dummy variables (17, one for each half-hour period). The only significant time of day effects are that AT becomes less likely at the end of the trading day, primarily in the last half hour of continuous trading.

[Insert Table 11 Here]

The probit results show that AT_{Init} is more likely when spreads are narrow and when trading volume over the prior 15 minutes is low. As in Table 3 larger trades are less likely to be initiated by AT. Volatility over the prior 15 minutes is negatively related to AT_{Init}. Depth at the inside (depth) and depth measured independently of the inside spread (depth3) are negatively related to AT_{Init}. The negative relations between AT initiation and spreads and between AT initiation and lagged volatility provide no evidence to support a hypothesis that AT exacerbates volatility.

The probit results show that AT_{Pass} is more likely when spreads are wide. Volatility over the prior 15 minutes is somewhat positively related to AT_{Pass}. Depth at the inside (depth) and depth measured independently of the inside spread (depth3) are negatively related to AT_{Pass}. The positive relation between AT liquidity supply and spreads and lagged volatility, and the negative relation to depth could lead to AT liquidity supply reducing volatility and smoothing liquidity. AT_{Pass} having a negative coefficient in the AT_{Init} regression shows that AT is less likely to supply liquidity when AT is demanding liquidity even after controlling for market conditions.

The coefficients on lagged futures return variables are consistent with AT_{Init} picking off stale human limit orders. Following past positive futures returns AT_{Init} buy markets orders are more likely. Conversely, when past futures returns are negative AT_{Init} sell orders are more likely. If futures prices lead the underlying stock prices then these initiated AT trades impose adverse selection costs on the non-marketable limit orders they execute against. AT_{Pass} has some relation to the lagged futures return variables suggesting that AT non-marketable limit orders may also be adversely selected.

In the AT_{Cancel} regressions the coefficients on lagged positive and negative DAX futures confirm the conjecture that AT are able to cancel limit orders quickly before they become stale. AT being able to cancel orders before they become stale may also allow AT to offer tighter spreads throughout the trading day by reducing their adverse selection costs. AT are also more likely to cancel their orders when spreads are wide and less likely to cancel when depth at the inside and depth independent of the inside spread is high. Volatility and volume over the prior 15 minutes is negatively related to AT_{Cancel}.

Consistent with prior univariate results, the probit results suggest that AT helps smooth out liquidity over time and is consistent with AT having lower monitoring costs in the liquidity make/take cycle proposed by Foucault, Kadan, and Kandel (2008). There is some evidence of aggressive AT adversely selecting stale limit orders. However, we cannot determine whether or not the introduction of AT increased this adverse selection or if AT is simply used for spot/future arbitrage that was previously executed manually. This latter possibility is quite plausible as spot/future arbitrage is one of the easiest strategies to automate.

VIII. Conclusion

We study algorithmic traders (AT) use of technology which reduces their monitoring frictions and AT's role in liquidity supply and demand dynamics. We find that AT consume liquidity when it is cheap and provide liquidity when it is expensive, likely reducing the volatility of liquidity. AT closely monitor the market and respond more quickly to changes in market conditions. The results are consistent with technology facilitating AT to more closely resemble the Friedman (1953) stabilizing speculator in terms of market liquidity. Further examinations of particular types of AT, e.g., high-frequency trading, should provide insight into the potentially differing impact that types of AT strategies may have.

Our results have important implications for academics, regulators, and market operators. Theoretical models of limit order books should allow for a significant fraction of traders who closely monitor the market. These traders would constantly reprice their orders and prevent spreads from widening beyond a certain point; both of these features can help simplify theoretical models as they reduce the dimensionality of the state space (cf. Goettler, Parlour, and Rajan (2009)). Given the slow progress in the modelling of theoretical limit order markets, this may have significant value.

Monitoring costs and limited attention are frictions limiting trade. In models without information asymmetry better market monitoring increases trading and investors' gains from trade (for examples, see Foucault, Kadan, and Kandel (2008) and Biais, Hombert, and Weill (2010)). Our results support the intuitive notion that AT reduces trading frictions and the use of the reduced form assumption that AT increases the probability of finding a counterparty (as in Biais, Foucault, and Moinas (2011)).

While lower monitoring costs for investors can be beneficial for important aspects of liquidity supply and demand, heterogeneous monitoring costs could also impose information asymmetry on slower traders. Slower traders face adverse selection if faster traders have an information advantage due to access to better and more current information about market conditions. Our results on net buying and selling of AT initiated trades being correlated with the direction of past index futures returns is consistent with this. If AT increases the scope of adverse selection sufficiently, liquidity and welfare could decline. While there is little empirical evidence supporting such a negative impact of technology in the trading process, it is an important avenue for future research, particularly if the adverse selection costs fall disproportionally on certain types of investors.

The increase in AT has important implications for both regulators and designers of trading platforms. For example, the U.S. Securities and Exchange Commission's Regulation NMS (SEC (2005)) tries to promote competition among liquidity suppliers.[13] AT lowering the monitoring costs for liquidity suppliers must also ensure vigorous competition among them. Trading venues should compete for AT by lowering development and implementation costs by facilitating the production of useful information and metrics for AT. Markets allowing algorithmic traders to co-locate their servers in the markets' data center should attempt to place all market participants on equal footing. Finally, markets and brokers can offer additional order types with features designed to lower investors' monitoring costs, e.g., pegged orders. Incorporating AT features into the market mechanism itself can lessen infrastructure costs for investors and mitigate arms races in technology investment.

[13]AT lowering monitoring costs should also improve linkages and integration among markets, potentially reducing concerns about liquidity fragmenting across many trading platforms. The Deutsche Boerse's dominant market share during our sample period precludes us from studying this.

References

Almgren, R., and N. Chriss. "Optimal Execution of Portfolio Transactions." *Journal of Risk*, 3 (2000), 5–40.

Barclay, M., T. Hendershott, and D. McCormick. "Competition among Trading Venues: Information and Trading on Electronic Communications Networks." *Journal of Finance*, 58 (2003), 2637–2666.

Bertsimas, D., and A. Lo. "Optimal Control of Execution Costs." *Journal of Financial Markets*, 1 (1998), 1–50.

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

Biais, B., T. Foucault, and S. Moinas. "Equilibrium Algorithmic Trading." Manuscript, Toulouse University, IDEI (2011).

Biais, B., P. Hillion, and C. Spatt. "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse." *Journal of Finance*, 50 (1995), 1655–1690.

Biais, B., J. Hombert, and P.-O. Weill. "Trading and Liquidity with Limited Cognition." Manuscript, Toulouse University, IDEI (2010).

Biais, B., and P. Woolley. "High Frequency Trading." Manuscript, Toulouse University, IDEI (2011).

Brogaard, J. "High Frequency Trading and its Impact on Market Quality." Manuscript (2010).

Brogaard, J. "The Activity of High Frequency Traders." Manuscript (2011).

Chaboud, A., B. Chiquoine, E. Hjalmarsson, and C. Vega. "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market." Technical report, FRB International Finance Discussion Paper No. 980 (2009).

Cohen, K., S. Maier, R. Schwartz, and D. Whitcomb. "Transaction Costs, Order Placement Strategy and Existence of the Bid-Ask Spread." *Journal of Political Economy*, 89 (1981), 287–305.

Domowitz, I., and H. Yegerman. "The Cost of Algorithmic Trading: A First Look at Comparative Performance." *Edited by Brian R. Bruce, Algorithmic Trading: Precision, Control, Execution. Institutional Investor*, (2005).

Duffie, D. "Asset Price Dynamics with Slow-Moving Capital." *Journal of Finance*, 65 (2010), 1238 – 1268.

Engle, R., J. Russell, and R. Ferstenberg. "Measuring and Modeling Execution Cost and Risk." Manuscript, NYU Stern (2007).

Foucault, T., O. Kadan, and E. Kandel. "Liquidity Cycles and Make/Take Fees in Electronic Markets." Manuscript, Toulouse University, IDEI (2008).

Foucault, T., and A. Menkveld. "Competition for order flow and smart order routing systems." *Journal of Finance*, 63 (2008), 119–158.

Foucault, T., A. Roëll, and P. Sandas. "Market Making with Costly Monitoring: An Analysis of the SOES Controversy." *Review of Financial Studies*, 16 (2003), 345–384.

Friedman, M. "The Case for Flexible Exchange Rates." In M. Friedman, editor, "Essays in Positive Economics," University of Chicago Press, Chicago (1953).

Goettler, R., C. Parlour, and U. Rajan. "Informed traders and limit order markets." *Journal of Financial Economics*, 93 (2009), 67–87.

Griffiths, M., B. Smith, D. Turnbull, and R. White. "The Costs and Determinants of Order Aggressiveness." *Journal of Financial Economics*, 56 (2000), 65–88.

Harris, L. "Optimal Dynamic Order Submission Strategies in Some Stylized Trading Problems." *Financial Markets, Institutions, and Instruments*, 7 (1998), 1–76.

Hasbrouck, J., and G. Saar. "Technology and liquidity provision: The blurring of traditional definitions." *Journal of Financial Markets*, 12 (2009), 143–172.

Hasbrouck, J., and G. Saar. "Low Latency Trading." Manuscript (2011).

Hau, H. "Location Matters: An Examination of Trading Profits." *Journal of Finance*, 56 (2001), 1959–1983.

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

Hendershott, T., and R. Riordan. "High Frequency Trading and Price Discovery." Manuscript (2011).

Jain, P. "Financial Market Design and the Equity Premium: Electronic versus Floor Trading." *Journal of Finance*, 60 (2005), 2955–2985.

Jovanovic, B., and A. Menkveld. "Middlemen in Limit-Order Markets." Manuscript (2011).

Kawaller, I., P. Koch, and T. Koch. "The Temporal Price Relationship between S&P 500 Futures and the S&P 500 Index." *Journal of Finance*, 42 (1987), 1309 – 1329.

Keim, D., and A. Madhavan. "Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Traders." *Journal of Financial Economics*, 37 (1995), 371–398.

Kirilenko, A., A. S. Kyle, M. Samadi, and T. Tuzun. "The Flash Crash: The Impact of High Frequency Trading on an Electronic Market." Manuscript (2011).

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

Lo, M. A., A., and J. Zhang. "Econometric Models of Limit-Order Executions." *Journal of Financial Economics*, 65 (2002), 31–71.

Menkveld, A. "High Frequency Trading and The New-Market Makers." Manuscript (2011).

Pagnotta, E., and T. Philippon. "Competing on Speed." Manuscript, New York University (2011).

Parlour, C. "Price dynamics in limit order markets." *Review of Financial Studies*, 11 (1998), 789–816.

Parlour, C., and D. Seppi. "Limit Order Markets: A Survey." *Handbook of Financial Intermediation and Banking, edited by A.W.A. Boot and A.V. Thakor*, (2008).

Petersen, M. "Estimating standard errors in finance panel data sets: Comparing approaches." *Review of Financial Studies*, 22 (2009), 435.

Ranaldo, A. "Order Aggressiveness in Limit Order Book Markets." *Journal of Financial Markets*, 7 (2004), 53–74.

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

Thompson, S. "Simple Formulas for Standard Errors that Cluster by Both Firm and Time." *Journal of Financial Economics*, 99 (2011), 1–10.

Venkataraman, K. "Automated Versus Floor Trading: An Analysis of Execution Costs on the Paris and New York Exchanges." *Journal of Finance*, 56 (2001), 1445–1485.

Appendix - Xetra and AT Matching Details

I. Xetra

The Xetra trading system is the electronic trading system operated by the Deutsche Boerse and handles more than 97% of German equities trading by euro volume in DAX stocks (2007 Deutsche Boerse Factbook). The DB is a publicly traded company that also operates the Eurex derivatives trading platform and the Clearstream European clearing and settlement system. DB admits participants that want to trade on Xetra based on regulations set and monitored by German and European financial regulators. After being admitted participants can only connect electronically to Xetra, floor trading is operated separately with no interaction between the two trading segments.

Xetra is implemented as an electronic limit order book with trading split into phases as follows:

- Opening call auction with a random ending that opens trading at 9:00
- A continuous trading period
- A two-minute intra-day call auction at 1:00 with a random ending
- A second continuous trading period
- A closing call auction beginning at 5:30 with a random ending after 5:35

We focus our analysis on trade occurring during the two continuous trading periods. Liquidity in DAX stocks is provided by public limit orders displayed in the order book of each stock. Orders execute automatically when an incoming market, or marketable limit order crosses with an outstanding non-marketable limit order. Order execution preference is determined using price-time-display priorities. Three types of orders are permitted, limit, market and iceberg orders. Iceberg orders are orders that display only a portion of the total size of an order. Iceberg orders sacrifice time priority on the non-displayed portion. Pre-trade transparency includes the 10 best bids and ask prices and quantities but not the ID of the submitting participant (as on the Paris Bourse (Venkataraman

(2001)). Trade price and size are disseminated immediately to all participants. The tick size for most stocks is 1 euro cent with the exception of two stocks that trade in tenths of a cent.[14]

II. Matching

To create the final dataset of trades and orders we use three separate data sources: AT order data from DB, public orderbook data from SIRCA, and public transactions data from SIRCA. Because SICRA timestamps reflect routing delays between the DB and Thompson-Reuters the SIRCA datsets are subject to time lags relative to the AT system order. The timestamp or SIRCA orderbook data is lagged by up to 250 milliseconds. The SIRCA transactions dataset is lagged by up to 500 milliseconds. The matching process for orders and trades is described in more detail below.

A Order Matching

We generate orders from successive orderbook updates similar to Biais, Hillion, and Spatt (1995). We match AT orders with the SIRCA public orderbook generated orders for the 5 best levels (bid and ask). To match the AT order data to the public data we use the following criteria:

- Symbol
- Price
- Size
- Side (bid or ask)
- Order type (insert or delete)
- Time stamp (micosecond)

[14]Both stocks, Deutsche Telekom AG and Infineon AG have trade prices below 15 euros. Stocks with prices lower than 15 euros have a tick size of one tenth of a cent.

Adjustments are made for a lag between the AT and SIRCA orderbook datasets. The publicly available data is time-stamped to the microsecond but, due to transmission and additional system processing, it lags the system order data. We allow for a time window of up to 250 milliseconds in the public data when looking for a match of the remaining criteria. The 250 ms maximum lag was determined by manual inspection of a large number of AT orders and SIRCA order books. We match the AT order with the next public order that matches the above criteria. If we do not find a match we delete the AT record. Approximately 5% of AT orders cannot be matched in the public data, many of these because they are outside the 5 best bid and ask prices in the SIRCA orderbook.

B Trade Matching

For trades we match two separate types of data, the DB supplied AT order data and the public transactions record. AT trades are matched with trades in the SIRCA public data record. To match AT trades to the SIRCA data we use the following criteria:

- Symbol
- Price
- Size
- Trade direction
- Time stamp (microsecond)

We identify the trade direction in the SIRCA public data using the Lee and Ready trade direction algorithm Lee and Ready (1991) with the Bessembinder (2003) modifications to determine the trade direction in the public data. Liquidity demanding (AT_{init}) trades match trade size and price in the public data. AT liquidity supplying trades (AT_{pass}) may be smaller than the total trade size as the marketable order may execute against multiple limit orders. We identify AT_{pass} using the same criteria as for AT_{init} and modify the size criteria to be less than or equal to the size reported in the public data.

Adjustments are made for a lag in the time stamp between the AT and SIRCA transaction datasets. As with orders, the publicly available trade data is time-stamped to the microsecond but, due to transmission and additional system processing, it lags the system order data. We allow for a time window of 500 milliseconds in the public transactions data when looking for a match on the remaining criteria. If we do not find a match we delete the AT trade. Roughly 97% of all AT trades are matched in the public data.

Table 1: **ATP-Rebate Program**. Fee rebate schedule for ATP participants by volume levels.

Cumulative Monthly ATP-Volume (in Mil. Euros)	**ATP-Rebate (per Volume level)**
$0 < 250$	0.0%
$250 < 500$	7.5%
$500 < 1000$	15.0%
$1000 < 2000$	22.5%
$2000 < 3750$	30.0%
$3750 < 7500$	37.5%
$7500 < 15000$	45.0%
$15000 < 30000$	52.5%
> 30000	60.0%

Table 2: **Summary Statistics**. This table presents descriptive statistics for the 30 constituents of the DAX index between January 1, 2008 and January 18, 2008. The data set combines Deutsche Boerse Automated Trading Program System Order data and SIRCA trade, quote, and order data. Market Capitalization data is gathered from the Deutsche Boerse website and cross-checked against data posted directly on the company's website and is the closing market capitalization on December 31, 2007. Other variables are averaged per stock and day (390 observations) and the mean, std. dev., maximum and minimum of these stock-day averages are reported.

Variable	Mean	Std. Dev.	Min	Max
Mkt. Cap. (Euro Billion)	32.85	26.03	4.81	99.45
Price (Euros)	67.85	42.28	6.45	155.15
Std. Dev of Daily Return (%)	3.12	1.40	1.47	9.29
Daily Trading Volume (Euro Million)	250	217	23	1,509
Daily Number of Trades per Day	5,344	3,003	1,292	19,252
Trade Size (Euro)	40,893	15,808	14,944	121,710
Quoted Spread (bps)	2.98	3.01	1.24	9.86
Effective Spread (bps)	3.49	3.05	1.33	10.05
Depth (Euro 10 Million)	0.0177	0.0207	0.0044	0.1522
Depth3 (Euro 10 Million)	0.1012	0.1545	0.0198	1.0689

Table 3: **AT Trade Participation by Size Category**. This table reports participation by AT and humans in 5 size categories. Panel A reports volume-weighted trade participation. Panel B reports transaction weighted trade participation.

Panel A	Trades		
Size Categories	AT	HUM	All
0 - 499	68%	32%	21%
500 -999	57%	43%	21%
1,000 - 4,999	42%	58%	43%
5,000 - 9,999	30%	70%	7%
10,000 +	23%	77%	8%
All	52%	48%	100%

Panel B	Trade		
Size Categories	AT	HUM	All
0 - 499	61%	39%	62%
500 -999	62%	38%	18%
1,000 - 4,999	53%	47%	18%
5,000 - 9,999	39%	61%	1%
10,000 +	31%	69%	1%
All	59%	41%	100%

Table 4: **AT Order Participation by Size Category**. This table reports participation by AT and humans in 5 size categories. Panel A reports order size-weighted participation for non-marketable limit orders. Panel B reports order weighted participation for non-marketable limit orders.

Panel A	Non-marketable Limit Orders		
Size Categories	AT	HUM	All
0 - 499	78%	22%	32%
500 -999	74%	26%	24%
1,000 - 4,999	55%	45%	35%
5,000 - 9,999	30%	70%	5%
10,000 +	20%	80%	4%
All	64%	36%	100%

Panel B	Non-marketable Limit Orders		
Size Categories	AT	HUM	All
0 - 499	77%	23%	62%
500 -999	74%	26%	20%
1,000 - 4,999	60%	40%	16%
5,000 - 9,999	31%	69%	1%
10,000 +	24%	76%	1%
All	73%	27%	100%

Table 5: **Trade Frequency Conditional on Previous Trade**. This table reports the conditional frequency of observing AT and human trades after observing trades of other participants. In column and row headings t indexes trades. AT represents AT trades and Hum represents human trades.

Ordering	Uncond.	Freq.	Buy_{t-1} Buy_t	$Sell_{t-1}$ $Sell_t$	Buy_{t-1} $Sell_t$	Buy_{t-1} Buy_t
$AT_{t-1}AT_t$	37.0	40.7	13.7	10.9	7.8	8.2
$AT_{t-1}Hum_t$	23.8	20.1	5.5	5.0	5.4	4.1
$Hum_{t-1}AT_t$	23.8	20.1	6.4	5.6	3.8	4.1
$Hum_{t-1}Hum_t$	15.3	19.0	5.4	5.3	3.7	4.4
Total		100.0	31.1	27.0	20.9	20.9

Table 6: **Trade Frequency Conditional on Previous Trade**. This table reports conditional frequencies based on the previous trade's size and participant. The three highest values per column are highlighted in bold. In column and row headings t indexes trades. AT represents AT trades and Hum represents human trades. Superscripts represent trade sizes with lower numbers corresponding to smaller trade sizes. Each row adds to 100.

t-1	AT_t^5	AT_t^4	AT_t^3	AT_t^2	AT_t^1	Hum_t^5	Hum_t^4	Hum_t^3	Hum_t^2	Hum_t^1
AT_{t-1}^5	**8.3**	**9.4**	**18.1**	16.7	7.8	**8.1**	**6.0**	6.7	7.9	10.5
AT_{t-1}^4	**3.8**	**7.8**	**15.9**	23.3	11.7	**4.5**	4.6	7.3	9.9	10.9
AT_{t-1}^3	1.3	2.7	12.0	**28.9**	20.6	2.2	2.7	6.2	11.1	12.0
AT_{t-1}^2	0.2	0.7	4.6	**27.1**	**33.8**	0.6	1.1	4.0	11.8	15.7
AT_{t-1}^1	0.0	0.1	1.7	16.6	**48.7**	0.2	0.5	2.2	9.9	19.9
Hum_{t-1}^5	**5.4**	**6.5**	**13.7**	17.5	8.3	**10.3**	**7.2**	**8.7**	10.3	12.1
Hum_{t-1}^4	1.8	3.4	10.4	22.5	14.4	4.2	**6.4**	**9.8**	13.5	13.6
Hum_{t-1}^3	0.5	1.4	6.7	**23.5**	21.2	1.7	2.8	**10.2**	**16.4**	**15.6**
Hum_{t-1}^2	0.2	0.5	3.4	19.2	28.4	0.7	1.3	4.9	**19.8**	**21.5**
Hum_{t-1}^1	0.1	0.3	2.2	14.9	**33.1**	0.6	0.9	3.4	**13.9**	**30.5**
Uncond.	0.4	0.7	3.4	17.1	31.6	1.0	1.0	3.9	15.1	26.2

Table 7: **Order conditional on past order**. This table reports the conditional frequencies of 8 order types, based on the previous order type and participant. The three highest values per column are highlighted in bold. *AT* represents AT orders and *Hum* represents human orders. Each row adds to 100.

	Algorithmic Trading								Human Trading							
t-1	Buy	Order At Ask	Cancel At Ask	Order New Bid	Sell	Order At Bid	Cancel At Bid	Order New Ask	Buy	Order At Ask	Cancel At Ask	Order New Bid	Sell	Order At Bid	Cancel At Bid	Order New Ask
AT Buy	**24.9**	7.5	3.9	6.2	5.5	8.5	6.7	5.9	5.6	3.9	3.3	2.7	3.0	2.8	5.0	**4.5**
AT Order At Ask	5.8	9.1	**6.1**	**9.6**	4.4	**20.1**	16.5	4.8	2.8	3.2	2.0	2.6	2.9	4.2	4.1	1.7
AT Cancel At Ask	4.3	**18.3**	**15.1**	5.7	4.2	10.6	11.8	**8.6**	2.7	3.1	2.1	1.6	2.3	2.8	4.8	2.2
AT Order New Bid	5.0	8.5	8.4	5.3	4.2	**26.9**	17.4	7.1	2.1	2.1	1.5	1.4	2.4	3.1	2.5	2.2
AT Sell	**11.0**	7.6	3.4	5.0	**24.5**	6.5	6.4	6.4	3.1	2.8	1.7	**4.2**	**5.5**	3.1	5.8	2.9
AT Order At Bid	4.7	10.7	5.8	**11.2**	3.7	**17.3**	**17.8**	**8.5**	2.3	2.8	1.6	2.5	2.6	3.1	3.6	1.9
AT Cancel At Bid	3.6	**12.2**	5.5	**12.7**	3.1	14.7	**19.1**	**12.0**	1.8	2.4	1.2	2.3	2.2	2.1	3.2	2.0
AT Order New Ask	5.0	**23.7**	1.4	6.6	4.2	8.5	**24.7**	5.8	2.1	3.3	0.3	2.0	2.1	2.0	6.6	1.6
Hum Buy	**8.2**	5.5	2.8	3.8	3.6	5.8	4.6	4.0	**28.5**	3.1	**8.8**	1.9	3.2	2.5	8.6	**5.2**
Hum Order At Ask	6.5	11.9	5.7	3.1	4.2	7.7	8.0	6.7	3.6	**12.4**	**9.5**	1.8	3.4	4.9	6.8	4.0
Hum Cancel At Ask	6.6	5.8	4.2	3.7	3.0	7.0	4.8	2.5	**14.0**	**10.0**	**12.6**	2.4	3.0	5.2	**12.7**	2.5
Hum Order New Bid	5.5	7.1	**12.4**	3.5	**8.4**	13.8	6.8	4.8	2.7	2.9	3.5	2.2	5.5	**6.4**	10.7	3.9
Hum Sell	4.5	5.2	2.0	2.9	**7.1**	4.9	5.1	3.6	**11.1**	2.6	2.1	2.5	**26.6**	2.6	**15.3**	2.1
Hum Order At Bid	5.3	8.0	5.5	6.4	5.5	11.8	8.0	3.5	3.3	**5.6**	4.7	**3.7**	3.9	**10.5**	**12.0**	2.2
Hum Cancel At Bid	3.8	6.7	2.9	7.7	4.3	6.3	7.9	5.5	4.7	4.6	2.8	**12.5**	**9.8**	**5.5**	10.6	**4.2**
Hum New Bid	6.9	9.3	0.6	3.5	3.5	5.6	16.1	3.1	2.9	5.1	1.1	2.3	2.0	2.1	33.9	2.0
Uncond.	6.8	12.1	5.8	6.5	5.7	11.8	12.2	6.9	4.4	4.6	2.8	2.8	4.3	4.0	6.5	2.9

Table 8: **Order Frequency Conditional on Spread**. This table reports the frequency of 6 order types conditional on contemporaneous spread broken down by AT and human participants. Panel A reports order frequencies conditional on spread size (small, small medium, medium, large). We calculate conditional spread sizes by taking time series quartiles for each stock and comparing these to the contemporaneous spread. If the spread is lower than or equal to the 25th percentile we classify it as small. If the spread is greater than or equal to the 75th percentile we classify it as large. Rows report frequencies of orders. The AT-Human column reports the difference between AT and human frequencies in that row. The T-Stat column reports the t-statistics for the AT-Human column accounting for both within stock time-series and contemporaneous cross-sectional correlation. In Panel B we report the difference between small and large spread order frequencies for AT and humans. Asterisks * or ** indicate significance at the 5% and 1% levels, respectively.

Panel A Small Spreads	AT	Human	AT-Human	T-Stat
New Bid/Ask inside	6.6	5.2	1.3	4.42**
New Bid/Ask at	9.7	4.8	4.9	11.17**
New Bid/Ask outside	16.9	6.2	10.7	12.55**
Cancel At	5.6	5.2	0.4	1.39
Cancel Away	17.3	6.2	11.1	13.39**
Trade Initiation	10.3	6.0	4.3	10.79**
Small Medium Spreads	**AT**	**Human**	**AT-Human**	**T-Stat**
New Bid/Ask inside	6.4	3.2	3.2	9.42**
New Bid/Ask at	11.7	4.9	6.8	11.76**
New Bid/Ask outside	17.6	6.8	10.8	12.08**
Cancel At	7.9	5.0	2.9	6.11**
Cancel Away	18.3	6.6	11.7	13.62**
Trade Initiation	6.9	4.7	2.2	12.42**

Continued...

Continued from Panel A Table 8.

Medium Spreads	AT	Human	AT-Human	T-Stat
New Bid/Ask inside	7.0	2.1	4.9	12.06**
New Bid/Ask at	13.6	4.5	9.2	15.72**
New Bid/Ask outside	17.6	6.4	11.2	14.32**
Cancel At	10.4	4.6	5.8	11.14**
Cancel Away	18.8	6.1	12.7	15.23**
Trade Initiation	5.1	4.0	1.1	7.14**
Large Spreads	AT	Human	AT-Human	T-Stat
New Bid/Ask inside	7.8	1.4	6.4	11.15**
New Bid/Ask at	14.8	3.7	11.1	11.54**
New Bid/Ask outside	16.7	5.9	10.9	9.79**
Cancel At	13.6	4.7	8.9	8.44**
Cancel Away	18.7	5.5	13.2	10.47**
Trade Initiation	3.6	3.4	0.2	1.64

Panel B Small-Large Spread Differences	AT	Human	AT-Human	T-Stat
New Bid/Ask inside	-1.3	3.8	-5.0	-12.19**
New Bid/Ask at	-5.1	1.1	-6.2	-9.43**
New Bid/Ask outside	0.1	0.3	-0.2	-0.47
Cancel At	-8.0	0.4	-8.5	-10.21**
Cancel Away	-1.3	0.7	-2.1	-1.97
Trade Initiation	6.7	2.6	4.1	11.42**

Table 9: **Time Interval Between Events and at Best Bid/Ask**. This table reports the time in seconds between events in Panel A and how long AT or humans spend at the best bid/ask. Panel A reports the time in seconds between events. The columns report time between events for AT, humans, and AT - humans and t-statistics. The rows report average time between events and the difference for small - large spread times. Small spreads are defined as spreads that are equal to or below the 25th percentile, large are defined as spreads that are greater than the 75th percentile. Percentiles are calculated as the time series mean for each stock. Panel B reports reports the number of seconds AT are at the best bid and ask quotes minus the number of seconds human traders are at the best quotes. The remainder of time both AT and humans are both at the best quotes. Small/Small Medium spreads are spreads that are below the time series mean. Medium/Large spreads are above their time series mean. T-statistics account for both within stock time-series and contemporaneous cross-sectional correlation. Asterisks * or ** indicate significance at the 5% and 1% levels, respectively.

Panel A - Time between	AT	Human	AT-Human Diff.	T-Stat
Trades	6.78	9.27	-2.48	-18.97**
Small - Large Spreads	-4.36	-2.34	-2.02	-23.42**
T-Stat	-38.73**	-15.53**		
Spread Narrowing and Trade	4.33	4.63	-0.30	-1.24
Small - Large Spreads	5.52	1.35	4.17	11.23**
T-Stat	14.23**	2.06*		
Spread Widening and Narrowing Order	3.22	4.11	-0.89	-2.38*
Small - Large Spreads	3.36	5.36	-2.00	-6.84**
T-Stat	9.87**	11.21**		

Panel B - Time at	AT	Human	AT - Hum	T-Stat
Best Bid/Ask	4,482	1,118	3360	7.00**
Small - Large Spreads	-2,149	-704	-1445	-7.21**
T-Stat	-12.22**	-2.01*		

Table 10: **Correlation of Order Flow and Liquidity Measures**. This table reports the correlation of AT_{Init} and AT_{Pass} trading and AT_{Cancel} orders with liquidity variables and DAX future returns. AT_{Init} takes the value of one if the trade is initiated by an AT. AT_{Pass} takes the value of one if AT supplies at least one share of a trade. AT_{Cancel} takes the value of one if AT cancel an order at the best and takes the value of zero if a human cancels an order at the best. $Rtn^{+/-}_{t-1,t-30} * BS$ is the return on the DAX future between t-1 and t-30 seconds * a buy/sell indicator for positive and negative DAX returns, respectively. Depth is the depth at best. Depth3 is the depth at three times the average quoted at the bid and ask side. Depth and Depth3 are reported in 10 million euros. Lagged volatility is the absolute value of the stock return in the 15-minutes prior to the trade. Lagged volume is the sum of the volume in the 15-minutes prior to the trade.

	AT_{Init}	AT_{Pass}	AT_{Cancel}	Spread	$Rtn^{+}_{t-1,t-30}$ $*BS$	$Rtn^{-}_{t-1,t-30}$ * BS	Size	Depth	Depth3	Lagged Volatility	Lagged Volume
AT_{Init}	1.00										
AT_{Pass}	-0.02	1.00									
AT_{Cancel}	–	–	1.00								
Quoted Spread	-0.08	0.10	0.03	1.00							
$Rtn^{+}_{t-1,t-30}$ * BS	0.02	0.00	0.01	0.01	1.00						
$Rtn^{-}_{t-1,t-30}$ * BS	-0.01	0.00	-0.01	-0.02	0.12	1.00					
Size	-0.09	0.00	-0.13	0.16	0.01	-0.01	1.00				
Depth	-0.05	-0.08	-0.10	-0.01	0.00	0.00	0.22	1.00			
Depth3	-0.05	-0.04	-0.10	-0.01	-0.01	0.01	0.26	0.63	1.00		
Lagged Volatility	0.01	0.01	-0.06	-0.05	0.02	0.04	0.01	0.01	0.04	1.00	
Lagged Volume	-0.05	-0.04	-0.01	-0.12	0.04	-0.05	0.04	0.14	0.14	-0.17	1.00

Table 11: **AT Probit Regression** In the first two columns the dependent variable (AT_{Init}) is equal to one if the trade is initiated by an AT and zero otherwise. In the third and fourth columns the dependent variable (AT_{Pass}) is equal to one if at least one share in the trade is supplied by an AT and zero otherwise. In the last two columns the dependent variable (AT_{Cancel}) takes the value of one if AT cancel an order at the best and takes the value of zero if a human cancels an order at the best. Size is the euro volume of a trade divided by 100,000. Depth is the depth at the best bid and ask. Depth3 is the depth at three times the average quoted spread on the bid and ask side. Depth and Depth3 are reported 10 million euros. $Rtn^{+/-}_{t,t-x} * BuySell$ is the return on the DAX future for t and $t-x$ seconds * a buy/sell indicator variable for positive and negative DAX returns, respecitively. Lagged volatility is the absolute value of the stock return in the 15-minutes prior to the trade. Lagged volume is the sum of the volume in the 15-minutes prior to the trade. Firm fixed effects and time of day dummies for each half-hour of the trading day are not reported. P-values are calculated using standard errors that account for both time-series and cross-sectional correlation. Asterisks * or ** indicate significance at the 5% and 1% levels, respectively.

	AT_{Init}		AT_{Pass}		AT_{Cancel}	
Variable	Model A	Model A1	Model B	Model B1	Model C	Model C1
AT_{Pass}	-0.02	-0.02	-	-	-	-
– Probability Slope	-0.02	-0.02	-	-	-	-
– P-value	(0.00**)	(0.00**)	-	-	-	-
Quoted Spread	-0.016	-0.016	0.035	0.035	0.04	0.04
– Probability Slope	-0.01	-0.01	0.01	0.01	0.01	0.01
– P-value	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)
Size	-0.11	-0.12	0.02	0.00	-0.57	-0.57
– Probability Slope	-0.09	-0.09	0.02	0.01	-0.21	-0.21
– P-value	(0.15)	(0.05*)	(0.00**)	(0.00**)	(0.00**)	(0.00**)
Depth	-0.10	-	-0.73	-	-1.91	-
– Probability Slope	-0.06	-	-0.28	-	-0.80	-
– P-value	(0.00**)	-	(0.00**)	-	(0.00**)	-
Depth3	-	0.01	-	-0.02	-	-0.58
– Probability Slope	-	-0.01	-	0.00	-	-0.26
– P-value	-	(0.00**)	-	(0.00**)	-	(0.00**)
$Rtn^{+}_{t,t-1} * BuySell$	20.70	20.75	17.96	17.94	25.02	24.98
– Probability Slope	8.56	8.58	7.00	6.98	9.49	9.48
– P-value	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)
$Rtn^{+}_{t-2,t-10} * BuySell$	34.33	34.21	2.81	1.58	42.55	42.18
– Probability Slope	14.40	14.22	0.84	0.35	15.86	15.71
– P-value	(0.00**)	(0.00**)	(0.51)	(0.55)	(0.00**)	(0.00**)
$Rtn^{+}_{t-11,t-30} * BuySell$	50.50	50.45	3.80	2.84	49.72	49.52
– Probability Slope	20.06	19.95	1.38	1.01	18.36	18.28
– P-value	(0.00**)	(0.00**)	(0.55)	(0.66)	(0.00**)	(0.00**)

Continued...

Continued from Table 11.

Variable	AT_{Init} Model A	AT_{Init} Model A1	AT_{Pass} Model B	AT_{Pass} Model B1	AT_{Cancel} Model C	AT_{Cancel} Model C1
$Rtn^{-}_{t,t-1} * BuySell$	-21.18	-21.21	-13.52	-13.35	-29.19	-29.11
– Probability Slope	-8.79	-8.80	-5.22	-5.14	-11.02	-11.00
– P-value	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)
$Rtn^{-}_{t-2,t-10} * BuySell$	-38.71	-38.77	-7.09	-6.43	-68.52	-67.92
– Probability Slope	-15.87	-15.77	-2.59	-2.33	-25.55	-25.31
– P-value	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)
$Rtn^{-}_{t-11,t-30} * BuySell$	-39.69	-39.63	-5.55	-4.60	-58.12	-57.70
– Probability Slope	-16.11	-15.97	-2.02	-1.65	-21.59	-21.41
– P-value	(0.00**)	(0.00**)	(0.47)	(0.54)	(0.00**)	(0.00**)
Lagged Volatility	-0.01	-0.01	0.02	0.02	-0.01	-0.01
– Probability Slope	-0.01	-0.01	0.01	0.01	-0.01	-0.01
– P-value	(0.03*)	(0.05)	(0.01*)	(0.01*)	(0.00**)	(0.00**)
Lagged Volume	-0.08	-0.09	-0.37	-0.41	-0.05	-0.05
– Probability Slope	-0.11	-0.11	-0.12	-0.14	-0.02	-0.02
– P-value	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)	(0.00**)
Observations	2,084,347	2,084,347	2,084,347	2,084,347	3,208,761	3,208,761