

Automated Liquidity Provision*

Austin Gerig[†]

Division of Economic and Risk Analysis
U.S. Securities and Exchange Commission
geriga@sec.gov

David Michayluk

Finance Discipline Group
University of Technology, Sydney
david.michayluk@uts.edu.au

Traditional market makers are losing their importance as automated systems have largely assumed the role of liquidity provision in markets. We update the model of Glosten and Milgrom (1985) to analyze this new world: we add multiple securities and introduce an automated market maker who prices order flow for all securities contemporaneously. This automated participant transacts the majority of orders, sets prices that are more efficient, increases informed and decreases uninformed traders' transaction costs, and has no effect on volatility. The model's predictions match very well with recent empirical findings and are difficult to replicate with alternative models.

Keywords: algorithmic trading; automated trading; high-frequency trading; market making; specialist; statistical arbitrage.

JEL Classification: G14, G19.

*An earlier version of this manuscript was entitled *Automated Liquidity Provision and the Demise of Traditional Market Making*. We thank Nadima El-Hassan, Daniel Gray, Terrence Hendershott, Bruce Lehmann, Danny Lo, Karyn Neuhauser, Lubomir Petrusek, Talis Putnins, Hui-Ju Tsai, Jian-Xin Wang, the participants of the *Market Design and Structure* workshop at the Santa Fe Institute, seminar participants at UTS, UNSW, ANU, Oxford, SEC, and participants at the 2011 FMA conference and 2014 Frontiers of Finance conference for their helpful comments and suggestions. This work was supported by the U.S. National Science Foundation Grant No. HSD-0624351 and the European Commission FP7 FET-Open Project FOC-II (no. 255987).

[†]The Securities and Exchange Commission, as a matter of policy, disclaims responsibility for any private publication or statement by any of its employees. The views expressed herein are those of the author and do not necessarily reflect the views of the Commission or of the author's colleagues on the staff of the Commission.

I Introduction

Traditionally, financial markets have appointed specialists or market makers to keep orderly markets and continually supply liquidity in traded securities. For many years, these individuals played a key role in determining prices and participated in a large fraction of trading. Today, this is no longer the case. Over the past decade, the task of liquidity provision has largely shifted from traditional market makers to proprietary automated systems that trade at high-frequency and across different exchanges and securities.¹

Why have automated systems replaced traditional market makers? What are the effects of automation: are prices more efficient, do transaction costs increase or decrease, who benefits and who is harmed?

In this paper, we present a straightforward model that helps answer these questions. In the model, we assume that automation alleviates a market friction due to the limited attention and processing power of traditional market makers. Specifically, we assume that a traditional market maker sets prices using only information about order flow in the security they trade, but that an automated liquidity provider uses

¹See Mackenzie (2013) for an account of the mechanizing of liquidity provision on the Chicago Mercantile Exchange in the late 1990's. The automated strategies on the CME eventually spread to equity markets and constitute a large share of what is now called "high-frequency trading." The TABB Group reports that high-frequency trading increased from 21% of US equity market share in 2005 to 61% in 2009. See "US Equity High Frequency Trading: Strategies, Sizing and Market Structure," September 2, 2009, available at <http://www.tabbgroup.com>. At the NYSE, the specialists' fraction of volume declined from approximately 16% to 2.5% over the period January 1999 to May 2007 (see Figs. 1 and 3 in Hendershott and Moulton (2011)). The drop in specialist activity is even more dramatic when considering the large shift of NYSE listed volume onto competing electronic trading platforms during the same period (see Fig. 20 in Angel, Harris, and Spatt (2011)). The catalyst for these changes was largely Regulation NMS (2005), which, among other things, eliminated the trade through protection for manual quotes at exchanges.

information about all order flow in the market. This assumption captures the main advantage that machines have over their human counterparts: they can quickly and accurately process large amounts of relevant information when setting prices.²

The model we present is an extension of Glosten and Milgrom (1985). Instead of using a single security, we include multiple related securities, each with their own traditional market maker, and we introduce an *automated market maker*. The automated market maker is the only individual in the model who trades in all of the securities, and who uses information from all order flow when setting prices. When adding the automated market maker, we find that:

- (1) Traditional market makers are largely priced out of the market and the automated market maker transacts the majority of order flow.
- (2) Liquidity traders have lower transaction costs and informed investors have higher transaction costs.
- (3) Prices are more efficient.
- (4) Short-term volatility is unaffected.

If the demand of investors is elastic, then the following results also hold:

- (5) Expected volumes increase.
- (6) Overall average transaction costs are reduced.

As we discuss in the literature review below, the assumptions of the model are justified by, and the predictions of the model match very well with, recent empirical findings.

²Many of the results of the model would also hold if the superior information of automated liquidity providers was something else relevant to prices other than market-wide order flow. For example, automated systems could be better at accurately determining prices based on information about the state of the aggregate order book or the likelihood that an observed imbalance in buying and selling continues, etc. We focus on cross-security information because the relevance of this information to prices is uncontroversial and easily modeled. Furthermore, high-frequency trading activity strongly correlates with the substitutability of a stock (see Gerig (2013) and Jovanovic and Menkveld (2012)).

The intuition behind the model is straightforward. By paying attention to all order flow, the automated market maker can form a better estimate of the price than traditional market makers and will use this to her advantage. She prices orders more accurately – offering worse prices to the toxic order flow of informed investors and better prices to the benign order flow of liquidity traders. In equilibrium, traditional market makers withdraw from the market because otherwise, they would transact against a disproportionate share of toxic order flow.

As an example, consider two securities XYZ and ZYX that are highly correlated with one another. If the automated market maker observes a buy order in XYZ and a sell order in ZYX, she can be somewhat confident that one or both of these orders are uninformed.³ She therefore offers both orders a better price and receives all of their order flow.

On the other hand, if the automated market maker observes buy orders in both XYZ and ZYX (or sell orders in both), she can be somewhat confident that both are informed, and that the price of each security should be higher (lower). She therefore draws her price away from the orders and lets the traditional market makers transact with them. On average, the traditional market makers receive the toxic order flow and the automated market maker receives the benign order flow. To compensate, traditional market makers must set wider quotes and therefore are largely squeezed out of the market.

³If both traders were informed, why would one be buying and the other one selling? Because XYZ and ZYX are similar, informed trading in these securities is more likely to be observed as contemporaneous buying or selling in both rather than buying in one and selling in the other.

Notice that in the above example, liquidity traders are offered better prices and informed investors are offered worse prices by the automated market maker. Therefore, the automated market maker lowers transaction costs for the uninformed (liquidity traders) and increases transaction costs for the informed (informed investors).

All of these results are driven by the simple fact that the automated market maker can reduce her adverse selection costs, i.e., her losses due to toxic order flow, by observing a larger amount of relevant information. Obviously, computers can collect, process, and act on more information than humans. Therefore, we should not be surprised to find that (1) liquidity provision is now automated, and that (2) this process reduces adverse selection in the market, with the relevant knock-on effects derived here. Indeed, as we detail in the literature review below, this appears to be the case.

II Literature Review

As we detail now, the literature consistently reports that when liquidity providers are enabled so that they can quickly and frequently update their quotes in an automated way, prices are more efficient, liquidity is enhanced, spreads decrease, and adverse selection decreases – all of which are predicted by our model. These results are economically significant and robust; they occur in many different markets over different circumstances. Furthermore, when the quotes of liquidity providers are constrained, the opposite occurs: liquidity is inhibited, spreads increase, and adverse selection increases. Finally, retail investors, who are often considered the most

uninformed, are particularly hurt – again a prediction of our model.

A Empirical literature

In the 1970’s and 1980’s there were a number of papers that discussed and even forecast the eventual automation of financial markets (e.g., Demsetz (1968, pg. 38), Fama (1970, footnote 22), Garbade and Silber (1978), Amihud and Mendelson (1988)). Although many authors anticipated the shift from human to computer-mediated trade, relatively few predicted the complete mechanization of liquidity provision. Black (1971) and Hakansson, Beja, and Kale (1985) are notable exceptions. As Mackenzie (2013) states, “Originally, the universal assumption had been that automated trading would involve a human being inputting orders into a computer terminal - all the early efforts to automate exchanges of which I am aware assumed this ...” Mackenzie (2013) attributes the shift towards fully mechanized trade to the rise in automated scalping and spreading strategies on the Chicago Mercantile Exchange, both of which are forms of proprietary liquidity provision.⁴

In the mid 1990’s, many influential financial markets introduced or switched to automated exchange. The effects of the transition were immediate and dramatic: liquidity improved considerably, which reduced trading costs for investors and the cost of equity for listed firms.⁵ In a large and comprehensive study, Jain (2005)

⁴It is perhaps pertinent to note that one of the authors of the current paper spent a summer as an intern in Chicago automating a spreading strategy on electronic futures markets. The general features of that strategy serve as the basis for the model presented here.

⁵Domowitz and Steil (2002) report that trading costs were 32.5% lower for NASDAQ listed stocks and 28.2% lower for NYSE listed stocks between 1992 and 1996 when using the new electronic exchanges. Furthermore, they estimate that every 10% drop in trading costs produces an 8% increase in share turnover and an 1.5% decline in the cost of capital.

analyzes the shift from floor to automated trading in 120 countries and estimates the cost of equity for listed firms dropped 0.49% per month (based on the ICAPM). Furthermore, he reports that after automation, share turnover increased by almost 50% and spreads reduced by 39%. The price of listed stocks increased by 8.99% on average when automation was announced.

In the 2000's, liquidity continued to improve as many automated exchanges upgraded their technological infrastructure and US regulation became more friendly for automated trade (especially in NYSE listed securities). Angel, Harris, and Spatt (2011) discuss this period in depth and report that “virtually every dimension of US equity market quality” improved. Castura, Litzenberger, and Gorelick (2012) study NYSE listed stocks before, during, and after Regulation NMS was implemented (which removed protections for NYSE manual quotes). They find that transaction costs for NYSE stocks reduced significantly compared to their NASDAQ counterparts before and after the change and that price efficiency increased as well. The Vanguard Group, a large mutual fund company that primarily manages index funds has reported a drop in their transactions costs of 35% to 60% over the past 15 years. They estimate the reduction in costs translates into a 30% higher return over a 30 year investment horizon in an actively managed fund.⁶

Although it would be difficult to attribute all of these improvements to advances in liquidity provision, and specifically to the reduction in the constraints of liquidity providers posited here, recent empirical papers help make the case that the automa-

⁶See Gus U. Sauter's comment on the “SEC Concept Release on Equity Market Structure” dated April 21, 2010 and available at www.sec.gov.

tion of liquidity provision deserves at least a portion of the credit.

There are a number of papers that find a positive and significant relationship between algorithmic trading (which includes automated liquidity provision) and market efficiency and liquidity.⁷ Hendershott, Jones, and Menkveld (2011) develop a measure of algorithmic trading on the NYSE based on message traffic and document a connection between the increase in algorithmic trading on the NYSE from 2000 to 2006 and the increase in both liquidity and the accuracy of quotes during the same period.⁸ To establish causality they use the staggered introduction of automated quote dissemination on the NYSE in 2003 (called Autoquote) as an instrument variable for algorithmic trading. Hasbrouck and Saar (2013) study order-level NASDAQ data, stamped to the millisecond, for October 2007 and June 2008 and develop a measure of “low-latency trading” based on runs of order cancellations and resubmissions that occur so quickly they could only be due to automated trading. They find that low-latency trading decreases spreads, increases liquidity, and lowers short-term volatility.

An argument can be made that the metrics for automated trading used in Hendershott, Jones, and Menkveld (2011) and in Hasbrouck and Saar (2013) should especially correspond with the activities of automated liquidity providers. An auto-

⁷Algorithmic trading usually includes all types of automated trading, including the automated splitting and execution of large directional orders as well as high-frequency arbitrage and automated liquidity provision strategies.

⁸There are other papers that report similar results. Boehmer, Fong, and Wu (2012) study AT on 42 markets and find that it improves liquidity and informational efficiency, but also increases volatility. Chaboud, Chiquoine, Hjalmarsson, and Vega (2013) study algorithmic trading on the foreign exchange market and find that it reduces arbitrage opportunities and the autocorrelation of high-frequency returns. Also see Hendershott and Riordan (2013).

mated liquidity provider who responds to order flow across the entire market needs to quickly and frequently update their quotes and will generate a large amount of message traffic in the process whereas directional investors or predatory traders using automated strategies will not.⁹ The introduction of Autoquote on the NYSE which was used in Hendershott, Jones, and Menkveld (2011) as an instrument variable should have particularly benefited automated liquidity providers. Furthermore, as shown in Hendershott and Moulton (2011), when the NYSE introduced its Hybrid Market structure (which also made automated trading more feasible), bid-ask spreads and adverse selection actually *increased*. This seemingly contradictory result can be explained in our framework because Hybrid (unlike Autoquote) had very little effect on the ability of quotes to automatically update, but did dramatically increase the speed of aggressive marketable orders.

Further evidence of the benefits of automated liquidity provision come from studies that look at market latency. Riordan and Storckenmaier (2012) study an upgrade at the Deutsche Boerse that reduced system latency from 50ms to 10ms in 2007. They find that both quoted and effective spreads decrease due to a dramatic reduction in adverse selection costs. Quotes go from 43% of price discovery to 89% after the upgrade, and liquidity providers are able to make more profits overall due to the lower adverse selection costs.¹⁰ Brogaard et al. (2013a) study an exchange system upgrade on NASDAQ OMX Stockholm that allows co-located traders to upgrade to a faster

⁹There are on average over 10,000 transactions per second in US equities alone during the trading day (see “U.S. Consolidated Tape Data” available at www.utpplan.com). Because of this, an automated liquidity provider can easily be required to update their quotes at millisecond timescales.

¹⁰See also Frino, Mollica, and Webb (2013) for a similar study on Australian securities.

connection. They find that those firms that pay for the upgrade reduce their adverse selection costs, improve their inventory management ability, and as a consequence, provide more liquidity. Overall, the upgrade improves both bid-ask spreads and market depth. These results support the hypothesis that reduced latency, even at millisecond timescales, is extremely important for automated liquidity providers – a circumstance difficult to explain outside the confines of our model.¹¹

Also, as reported in Jovanovic and Menkveld (2012), when Dutch stocks started trading on the electronic exchange Chi-X in 2007, quoted spreads, effective spreads, and adverse selection reduced significantly (by 57%, 15%, and 23% respectively using a difference-in-differences analysis). Jovanovic and Menkveld (2012) attribute these improvements to characteristics of Chi-X that made automated liquidity provision easier: unlike the incumbent Euronext market, Chi-X did not charge for quote changes and had lower latency.

Not all changes in market structure have been beneficial for automated liquidity providers. As previously mentioned, the introduction of Hybrid on the NYSE increased the speed of market orders but not quotes, to the detriment of automated liquidity provision. Also, in 2012, the investment regulatory body in Canada introduced a fee for message traffic that was especially costly for traders who frequently update their quotes. As a result, quoted spreads, effective spreads, and adverse selection costs increased (see Malinova, Park, and Riordan (2013a) and Lepone and

¹¹Gai, Yao, and Ye (2013) find no evidence that markets benefit from reducing latency beyond millisecond levels. Within our model, trades from directional investors would need to arrive 1 million times per second for microsecond speeds to be of consequence. As mentioned in a previous footnote, evidence suggests they arrive 10,000 to 100,000 times per second in US equities.

Sacco (2013)). Malinova, Park, and Riordan (2013a) show that the reduction in frequent quote updates especially hurt retail investors, who made larger losses after the introduction of the fee. Assuming that retail investors are uninformed, this otherwise surprising result is predicted by our model.

B Who are the automated liquidity providers?

Unfortunately, the literature can be inconsistent or vague in the way it classifies automated trading. For example, the term *algorithmic trading* sometimes refers to all types of automated trading and sometimes only to the automated splitting and execution of large directional orders. The term *high-frequency trading* often refers to all non-directional low-latency automated trading strategies, which can be problematic when interpreting results. For example, aggressive low-latency arbitrage strategies that pick-off the quotes of automated liquidity providers or that front-run large directional trades can be harmful for liquidity (see Hirschey (2013)), yet are classified together with automated liquidity provision under the umbrella-term “high-frequency trading.”

For the purposes of this paper, we define automated liquidity provision as high-frequency trading that primarily transacts with the orders of directional investors when those orders offer a sufficient price concession. Although we expect automated liquidity providers to maintain a presence at the best quotes (and to frequently update these quotes) this does not necessarily mean they are the passive party in all of their transactions. For example, if a directional trader places a limit order and the market

moves against that order, we expect an automated liquidity provider to transact against the order as a consequence of updating their quote. This circumstance is just another form of liquidity provision as we have defined it.

In the literature, there is convincing evidence that automated liquidity providers exist and make up a significant portion of high-frequency trading. Furthermore, there is evidence that their strategies are heavily dependent on cross-market and cross-security information (as posited in the model).

Hagströmer and Nordén (2013) study order book data from NASDAQ-OMX Stockholm that allows them to identify the trading activity of high-frequency trading firms. They divide the firms into two subsets based on their market activity: those who act as automated liquidity providers (at the best quotes more than 20% of the time) and those who are opportunistic (all others). They find that approximately 2/3 of all high-frequency volume and more than 80% of high-frequency trading activity is due to the automated liquidity providers.

Javonic and Menkveld (2012) and Menkveld (2013) identify a single large high-frequency trader that acts as an automated liquidity provider in two competing Dutch markets (Chi-X and Euronext). The high-frequency trader participates in 8.1% of trades on Euronext and 64.4% of trades on Chi-X, carefully controls its inventory across the two markets, predominately trades with passive orders (80% of its trades), and is profitable (with an estimated Sharpe ratio of 7.6). Javonic and Menkveld (2012) find evidence that quotes for Dutch stocks on Chi-X (where the high-frequency trader dominates trading) are more responsive to changes in the Dutch index futures market

than Euronext quotes, and that the trading of the high-frequency trader is higher on days when prices are closely linked to the index. These results support our assumption that automated liquidity providers depend heavily on cross-security information when setting their prices. Further evidence is found in Gerig (2013). Using data that flags high-frequency trading activity on NASDAQ, Gerig (2013) finds that high-frequency trading activity synchronizes the price response of different securities and that the fraction of high-frequency trading in a particular security correlates strongly ($\rho = 0.80$) with how correlated that security's price movements are to other securities.

Finally, the effects of high-frequency trading on the trading costs of different types of investors has been studied in various papers with somewhat mixed results. According to our model, automated liquidity provision should increase transaction costs for informed investors and reduce transaction costs for uninformed investors. To the extent that high-frequency trading in a particular market corresponds to automated liquidity provision and to the extent that retail and institutional investors correspond to uninformed and informed traders in that market, our model predicts that retail investor transaction costs should decrease and institutional investor transaction costs should increase as a result of high-frequency trading.

Malinova, Park, and Riordan (2013b) study the changing profits of institutional traders, retail traders, and high-frequency traders from 2006 to 2012 on the Toronto Stock Exchange. They report that retail profits are positively associated with high-frequency trading activity and that institutional profits are overall unrelated to high-frequency trading activity, although for some years it is negatively related. Brogaard

et al. (2013b) study the effects of high-frequency trading activity on institutional investor execution costs on the London Stock Exchange from November 2007 to August 2010 and find no evidence of a relationship between the two. Tong (2013) combines high-frequency trading data from NASDAQ with institutional trading data from An-cerno and finds that high-frequency trading activity increases the transaction costs of institutional investors.

C Alternative explanations

Why has automated liquidity provision improved market quality? In the model, transaction costs are lower because the automated liquidity provider passes along her reduced adverse selection costs. However, transaction costs can also decrease if liquidity providers make less profits – perhaps because they experience increasing competition or lower fixed costs. Estimates of liquidity provider trading revenue, however, do not support the notion that automated liquidity providers extract any less money from the market than their previous human counterparts.¹²

Menkveld (2013) reports that an identified automated liquidity provider on Chi-X (mentioned previously) made total gross profits of €9,542 per day from trading 14 Dutch stocks. Brogaard, Hendershott, and Riordan (2013) and Carrion (2013) study a proprietary dataset of high-frequency trading activity on the NASDAQ exchange. Although these papers report slightly different numbers for high-frequency trading

¹²The profit per transaction and per share traded for automated liquidity providers is lower, but this circumstance supports our story of the limited attention and processing power of human liquidity providers rather than a story of increased competition.

revenue, they both find that high-frequency trading is profitable and particularly so when on the passive side of transactions (these passive trades are paid a rebate on NASDAQ). Carrion (2013) reports that high-frequency traders make trading profits on the NASDAQ exchange of \$2,623.84 per stock per trading day and Brogaard, Hendershott, and Riordan (2013) report profits of \$1990.10 per stock per trading day that increases to \$2,284.89 when including NASDAQ fees and rebates. Extrapolating from these numbers, they estimate high-frequency trading revenue of \$5 billion per year in US equities.¹³

Baron, Brogaard, and Kirilenko (2012) study transaction level data with user identifiers for the E-mini S&P 500 Futures contract and estimate that high-frequency traders earn \$23 million in trading revenue per day. Contrary to what is reported in US equities, a large fraction (a little less than half) of these profits are from high-frequency trading firms that are the aggressive party in at least 60% of their trades.¹⁴

The estimates of high-frequency trader profits are similar to, if not more than, those reported for human liquidity providers before automation took hold. Sofianos (1995) uses April 1993 NYSE specialist transaction data in 200 stocks to estimate their gross trading revenues. He finds that specialist trading revenue averaged \$670 per stock per day. Hansch, Naik, and Viswanathan (1999) analyze the trading profits of all dealers in all stocks on the London Stock Exchange in August 1994. They

¹³They multiply by approximately 3 to include volume on other exchanges, by 3000 to include US equity listings where high-frequency trading is active, and by 250 trading days per year.

¹⁴This result could be evidence that a significant portion of high-frequency traders in futures markets partake in activities that are harmful for liquidity, or alternatively might represent a difference in how liquidity is provided in futures markets, particularly in contracts that have large tick sizes.

find that “dealers earn on average, a “dealer profit” (before subtracting the costs of market making: salaries, technology, cost of capital, etc.) that is not statistically different from zero.” Ellis, Michaely, and O’Hara (2002) study dealer activity in 290 newly listed firms on NASDAQ between September 1996 and July 1997. They report total dealer profit of \$570,000 per stock over the first 140 trading days (approximately \$4,070 per day). Dealer profits were largest in the first 20 trading days after the IPO (\$13,100 per day) and then shifted to \$2,570 per stock per day afterward.

D Theoretical literature

There are many recent papers that model high-frequency trading, but nearly all of these papers focus on predatory types of high-frequency trading that cannot be classified as automated liquidity provision. Such papers include Cvitanić and Kirilenko (2010), Jarrow and Protter (2012), Cartea and Penalva (2012), Biais, Foucault, and Moinas (2013), Foucault, Hombert, and Roşu (2013), Martinez and Roşu (2013), Hoffmann (2013), and Budish, Cramton, and Shim (2013).

More closely related to our work, Jovanovic and Menkveld (2012) model high-frequency traders as “middlemen” in limit order markets. These middlemen observe hard information with certainty (Peterson (2004)), whereas an early seller who places a limit order does not and a late buyer who places a market order does with a certain probability. For certain parameter values, the middlemen reduce the asymmetry of information between the two counterparties and increase welfare. There are two main similarities between our model and theirs that underlie the liquidity enhance-

ments observed in both models. First, their middlemen (like our automated liquidity provider) transacts with investors who possess asymmetric information, and second, their middlemen (like our automated liquidity provider) can reduce this asymmetry by observing hard information.

More generally, informed trading in a multi-asset setting has been modeled in several papers as an extension of the one-asset equilibrium models of Kyle (1985, 1989), Easley et al. (1996), and Easley et al. (1998). In these papers, competitive market makers set prices using information from all order flow,¹⁵ or alternatively, using only information about order flow in a specific market or security.¹⁶ Here, we are interested in analyzing differences in market quality when liquidity providers do and do not observe all order flow. In this regard, the multi-asset Kyle (1985) model described in Baruch and Saar (2009) is closest to our model (although developed for an entirely different purpose). They show that the liquidity of an asset is enhanced when market makers can observe order flow in correlated assets, and the more correlated the assets, the stronger the effect. Fricke and Gerig (2013) show similar results using a multi-asset extension of Garbade and Silber (1979).

The cost of liquidity provision has been described in various ways in previous papers and can be subdivided into three main categories: order-handling costs, inventory costs, and adverse selection costs (see Biais, Glosten, and Spatt (2005) for an overview). In this paper, we focus only on the adverse selection component and do not

¹⁵Caballé and Krishnan (1994), Heidle (2004), Tookes (2008), Bernhardt and Taub (2008), and Pasquariello and Vega (2013)

¹⁶Boulatov, Hendershott, and Livdan (2013) and Chordia, Sarkar, and Subrahmanyam (2011)

consider order-handling or inventory costs. Similar results would hold if these other costs were considered. For example, Andrade, Chang, and Seasholes (2008) develop a multi-asset equilibrium model where risk-averse liquidity providers accommodate noninformational trading imbalances. They show that demand shocks in one security will lead to cross-stock price pressure due to the hedging desires of liquidity providers, even without considering information effects. Automated liquidity provision, therefore, could also enhance market quality by increasing hedging opportunities instead of, or in addition to, reducing adverse selection costs.

III The Model

As in Glosten and Milgrom (1985), we assume a pure dealership market where *informed investors* and *liquidity traders* submit unit-sized buy and sell orders for a security with a random end-of-period value. These orders are priced and cleared by a competitive, risk neutral *market maker*. Unlike the original Glosten and Milgrom model, we include multiple securities—each with a separate market maker and separate informed investors and liquidity traders—and introduce a new type of competitive, risk neutral liquidity provider called an *automated market maker*. The automated market maker is unique in that she is the only market participant who trades in multiple securities and understands how their end-of-period values are related; everyone else is unaware of this, unable to model it correctly, or does not have the sophistication to trade broadly and quickly on this information.

Each round of trading takes place in two steps. In the first step, one unit-sized

order is submitted for each of the $N > 1$ securities in the market. These orders arrive randomly and anonymously from the pool of liquidity traders and informed investors within each security. Liquidity traders are uninformed of the end-of-period value for the security they trade, are equally likely to buy or sell this security, and are always willing to accept the best price set by the liquidity providers.¹⁷ Informed investors know with certainty the end-of-period value for the security that they trade, and they submit a buy or sell order based on this information. If the end-of-period value is above the expected transaction price, they will place a buy order, and if it is below the expected transaction price, they will place a sell order. The probability that an order for security i is from an informed investor is public knowledge, is assumed to be larger than zero but less than one, and is denoted γ_i . A buy order for security i is denoted B_i and a sell order is denoted S_i .

In the second step, liquidity providers observe the set of orders placed, determine the prices at which they are willing to transact, and then transact with an order if their price is the most competitive. If the market maker for security i and the automated market maker set the same price for a transaction, then we assume each partakes in half of the trade. Notice that we assume prices are set by liquidity providers after all orders have been placed, whereas in the original Glosten and Milgrom (1985) model, quotes are set by the liquidity provider before transactions take place. In real markets, transactions occur randomly in time across securities, and quotes for any particular security will be dynamically updated as transactions occur in other related securities.

¹⁷This is the case of perfectly inelastic demand in the original Glosten and Milgrom (1985) model. We relax this assumption in Section IV.

For simplicity, we assume these updates are collapsed into one moment in time.

We assume that the true value of security i can take on one of two possible values at the end of the trading period, $\tilde{V}_i = V_i$, where $V_i \in \{V_i^+, V_i^-\}$ and,

$$V_i^+ = p_i + r_i, \tag{1}$$

$$V_i^- = p_i - r_i. \tag{2}$$

r_i is a constant that sets the scale of price changes for security i . For simplicity, we assume that each of these occur with probability $1/2$ so that the unconditional expected value of security i is p_i . Everyone in the model is aware of these possible outcomes and their probabilities.

We assume that the value of securities are related to one another such that the conditional probability of a security value being V_i^+ or V_i^- changes as information about the value of other securities is added,

$$P(V_i) \neq P(V_i|V_j) \neq P(V_i|V_j, V_k) \neq \dots \neq P(V_i|V_j, V_k, \dots, V_N). \tag{3}$$

These conditional probabilities are not known to the traders or the market makers, but are known to the automated market maker and used to price transactions.

In general, the end of the trading period can occur at some distant point in the future, so that there can be multiple rounds of trading before the end-of-period value for each security is revealed. To simplify the analysis, and because our main results

hold for only one round of trading, we only analyze the case where the trading period ends after one round. When presenting results, we implicitly assume that many of these trading periods have taken place and we take averages over the outcomes.

Because liquidity providers are competitive and risk neutral, they set prices such that they expect to make zero profit on their trades (see Glosten and Milgrom (1985)). For example, in a market with only one security and given that the fraction of informed trading is γ_1 , both the market maker and the automated market maker would set the price of an order at the expected value of the security given the order,

$$E \left[\tilde{V}_1 \mid B_1 \right] = p_1 + \gamma_1 r_1, \quad (4)$$

$$E \left[\tilde{V}_1 \mid S_1 \right] = p_1 - \gamma_1 r_1. \quad (5)$$

Therefore, a buy order transacts at price $p_1 + \gamma_1 r_1$ and a sell order at price $p_1 - \gamma_1 r_1$ in this scenario. Because orders are cleared at these prices, there is no expected profit for the liquidity providers.

In the propositions below, we assume that everyone is aware of the setup of the model and has traded long enough to know what actions are optimal for them. After stating the propositions and giving a short explanation of each, we prove them for a two security market. Proofs for markets with more than two securities are straightforward extensions of these proofs.

PROPOSITION 1. *The market maker for security i transacts a minority of order flow and competes with the automated market maker only on the widest spreads.*

The market maker cannot compete with the automated market maker on pricing because he is not using information about the relationships between securities. If he sets prices to their expected value given his information, he will find himself transacting unprofitable orders and not transacting profitable ones. This is because the automated market maker knows which orders are profitable at these prices; she takes the profitable ones away from the market maker by pricing them more aggressively and leaves the unprofitable ones by pricing them less aggressively. Because the market maker is transacting unprofitable orders (which contain a larger fraction of informed trades), he must widen his spread in price between buy orders and sell orders. The equilibrium point for him is not reached until he is using the overall highest and lowest prices set by the automated market maker.

PROPOSITION 2. *With the addition of the automated market maker, the spread in average transaction price between buys and sells in security i remains the same, but is smaller for liquidity traders and larger for informed investors.*

The automated market maker is unable to reduce spreads because the unconditional probability of an order arriving from an informed versus a liquidity trader for each security remains unchanged. She, however, can use information about the orders placed in other securities to help distinguish if a particular order is likely from an informed investor vs. a liquidity trader. She uses this information to price these situations differently. Informed investors receive worse prices because they now compete with one another across securities; this means the average spread in price between a buy and a sell increases for them. Liquidity traders receive better prices because

their trades can be connected to other liquidity traders in different securities through the actions of the automated market maker. For these traders there is a decrease in the average spread in price between buys and sells.

PROPOSITION 3. *With the addition of the automated market maker, prices are more efficient, i.e., on average, the transaction price for security i is closer to its end-of-period value.*

The automated market maker uses a better information set than that used by the market maker and can therefore price order flow more precisely. This means the transaction price is, on average, closer to the security's end-of-period value. This effect is most pronounced in securities that have a small proportion of informed trading.

In the original Glosten and Milgrom (1985) paper, as the liquidity provider observes more order flow, transaction prices approach the security's fundamental value. This proposition documents the same effect, but now it occurs because the automated market maker can observe "order flow substitutes" for security i by observing the orders for the rest of the securities in the market.

PROPOSITION 4. *With the addition of the automated market maker, the volatility of price changes in security i is unaffected.*

Liquidity providers set prices such that they follow a martingale with respect to their information set. The automated liquidity provider has a more refined information set, but this has no effect on the overall variance of price changes – the martingale property ensures that price volatility is the same regardless of information set.

IV Two Security Market

In what follows, we prove the propositions for a two security market where

$$P(V_1^+|V_2^+) = \phi.$$

A Proposition 1

There are four possible order flow states in this market, (B_1, B_2) , (B_1, S_2) , (S_1, B_2) , (S_1, S_2) . Because the values of the two securities are related to each other, and because informed investors trade in both, certain order flow states will be observed more often than others. For example, when the values of security 1 and security 2 are positively correlated, then the states (B_1, B_2) and (S_1, S_2) are more likely to occur, and the orders in these states are more likely to be informed. The automated market maker is aware of this and sets prices accordingly. She sets the transaction price of orders for security 1, T_1 , to the expected value of security 1 conditioned on the particular order flow state. This is calculated as follows,

$$T_1(B_1, B_2) = E \left[\tilde{V}_1 \middle| B_1, B_2 \right] = p_1 + \frac{\gamma_1 + (2\phi - 1)\gamma_2}{1 + (2\phi - 1)\gamma_1\gamma_2} r_1, \quad (6)$$

$$T_1(B_1, S_2) = E \left[\tilde{V}_1 \middle| B_1, S_2 \right] = p_1 + \frac{\gamma_1 - (2\phi - 1)\gamma_2}{1 - (2\phi - 1)\gamma_1\gamma_2} r_1, \quad (7)$$

$$T_1(S_1, B_2) = E \left[\tilde{V}_1 \middle| S_1, B_2 \right] = p_1 - \frac{\gamma_1 - (2\phi - 1)\gamma_2}{1 - (2\phi - 1)\gamma_1\gamma_2} r_1, \quad (8)$$

$$T_1(S_1, S_2) = E \left[\tilde{V}_1 \middle| S_1, S_2 \right] = p_1 - \frac{\gamma_1 + (2\phi - 1)\gamma_2}{1 + (2\phi - 1)\gamma_1\gamma_2} r_1. \quad (9)$$




Figure 1: Diagram of the transaction prices for security 1 in the 4 possible order flow states, (B_1, B_2) , (B_1, S_2) , (S_1, B_2) , (S_1, S_2) , in a two security market. The values shown at earlier nodes are not transaction prices, but are the expected price at those nodes.

The transaction prices in Eqs. 6-9 are shown in diagram form in Fig.1. The market maker for security 1 is unaware of the importance of conditioning prices on the order in security 2 (or is simply unable to do so), and therefore does not use Eqs. 6-9 to price order flow. He *is* aware of the following,

$$E \left[\tilde{V}_1 \mid B_1 \right] = p_1 + \gamma_1 r_1, \quad (10)$$

$$E \left[\tilde{V}_1 \mid S_1 \right] = p_1 - \gamma_1 r_1. \quad (11)$$

The market maker would set prices at these values if the automated market maker were not present.

In what follows, we assume that ϕ is restricted to $(1/2, 1]$, such that the values

of security 1 and 2 are positively correlated. The results would be similar, although with some signs and descriptions reversed, if they were negatively correlated. The propositions hold in either case. Because γ_1 and γ_2 are restricted to $(0, 1)$, it is straightforward to show that,

$$E \left[\tilde{V}_1 \middle| B_1, B_2 \right] > E \left[\tilde{V}_1 \middle| B_1 \right] > \dots$$

$$E \left[\tilde{V}_1 \middle| B_1, S_2 \right] \overset{\leftarrow}{\rightarrow} E \left[\tilde{V}_1 \middle| S_1, B_2 \right] > E \left[\tilde{V}_1 \middle| S_1 \right] > E \left[\tilde{V}_1 \middle| S_1, S_2 \right], \quad (12)$$

where $\overset{\leftarrow}{\rightarrow}$ denotes that $E \left[\tilde{V}_1 \middle| B_1, S_2 \right]$ and $E \left[\tilde{V}_1 \middle| S_1, B_2 \right]$ may be switched in this ordering. Notice that the automated market maker sets the price for B_1 more aggressively when it is accompanied by S_2 and more timidly when accompanied by B_2 (with the opposite results for S_1). The automated market maker realizes that the opposite direction of orders makes it more likely both investors are uninformed and that the same direction makes it more likely they are informed, and she prices them accordingly. Because the market maker does not realize this, his role is marginalized. Suppose the market maker tried to set prices according to Eqs. 10-11. If he transacted a random sampling of orders at these prices, he would make zero expected profit. Unfortunately for him, the automated market maker ensures that his transactions are not a random sampling. She transacts orders that would have been profitable to him by pricing them more aggressively (this occurs for (B_1, S_2) and (S_1, B_2) , see Eq. 12), and she allows him to transact orders that are unprofitable because she prices them less aggressively (this occurs for (B_1, B_2) and (S_1, S_2) , see Eq. 12). Because the market

maker is now receiving a larger fraction of informed trading, he is forced to widen his spread until he expects to make zero profit. This occurs at the widest spread set by the automated market maker. This same argument applies to the market maker for security 2.

We have assumed that when the market maker and automated market maker set the same price for a transaction, they each partake in half of the trade. This means that the fraction of order flow transacted by the market maker is equal to $1/2$ of the proportion of order flow at the widest spread. This is,

$$P(B_1, B_2)/2 + P(S_1, S_2)/2 = [1 + (2\phi - 1)\gamma_1\gamma_2]/4. \quad (13)$$

Because γ_1 and γ_2 are restricted to $(0, 1)$ and ϕ is restricted to $(1/2, 1]$, the maximum value of this equation is less than $1/2$. Therefore, the market maker in security 1 trades a minority of order flow. The same argument applies to the market maker for security 2.

B Proposition 2

We can calculate the expected transaction price of informed investors when buying and selling security 1 as follows,

$$E[T_1|B_1, I_1] = E[T_1|B_1, B_2, I_1]P(B_2|B_1, I_1) + E[T_1|B_1, S_2, I_1]P(S_2|B_1, I_1), \quad (14)$$

$$E[T_1|S_1, I_1] = E[T_1|S_1, S_2, I_1]P(S_2|S_1, I_1) + E[T_1|S_1, B_2, I_1]P(B_2|S_1, I_1), \quad (15)$$

where the conditioning variable I_1 means the order for security 1 was from an informed investor. The expected spread in price between buying and selling security 1 for an informed investor, $\Delta_{1,I}$, is therefore,

$$\Delta_{1,I} = 2\gamma_1 r_1 \left\{ \frac{1}{\gamma_1} - \frac{(1/\gamma_1 - 1)[1 - (2\phi - 1)^2 \gamma_2^2]}{1 - (2\phi - 1)^2 \gamma_1^2 \gamma_2^2} \right\} \quad (16)$$

In a similar way, we calculate the spread between buying and selling security 1 for a liquidity trader, $\Delta_{1,L}$,

$$\Delta_{1,L} = 2\gamma_1 r_1 \left[\frac{1 - (2\phi - 1)^2 \gamma_2^2}{1 - (2\phi - 1)^2 \gamma_1^2 \gamma_2^2} \right] \quad (17)$$

$\Delta_{1,I}$ and $\Delta_{1,L}$ are to be compared to the unconditional spread, Δ_1 , which can be calculated as follows,

$$\Delta_1 = \gamma_1 \Delta_{1,I} + (1 - \gamma_1) \Delta_{1,L}, \quad (18)$$

$$= 2\gamma_1 r_1. \quad (19)$$

This is just the spread that would be observed without the automated market maker (subtract Eq. 11 from Eq. 10). The same argument applies for security 2. In Fig. 2, we diagram these results.

Because γ_1 and γ_2 are restricted to $(0, 1)$, and ϕ is restricted to $[0, 1/2)$ or $(1/2, 1]$,




Figure 2: Diagram of the average buy and sell transaction price for informed investors, liquidity traders, and unconditional on trader type. Notice that informed investors realize a larger spread than liquidity traders, and that the spread without the automated market maker is the expected spread over both of these trader types.

then it is straightforward to show that,

$$\Delta_{1,I} > \Delta_1 > \Delta_{1,L} \quad (20)$$

so that the spread is increased for informed investors and decreased for liquidity traders. The same argument applies for security 2.

C Proposition 3

To determine the inefficiency of transaction prices, we calculate the expected absolute difference between \tilde{V}_1 and the transaction price, T_1 ,

$$E \left[|\tilde{V}_1 - T_1| \right] = E \left[|\tilde{V}_1 - T_1| \middle| I_1 \right] P(I_1) + E \left[|\tilde{V}_1 - T_1| \middle| L_1 \right] P(L_1). \quad (21)$$

where I_1 and L_1 denote that the order in security 1 was placed by an informed investor and a liquidity trader respectively. Calculating this, we have,

$$E \left[|\tilde{V}_1 - T_1| \right] = (1 - \gamma_1) \left[1 + \frac{1 - (2\phi - 1)^2 \gamma_2^2}{1 - (2\phi - 1)^2 \gamma_1^2 \gamma_2^2} \gamma_1 \right] r_1. \quad (22)$$

Without the automated market maker, the expected absolute difference between the end-of-period value and the transaction price is,

$$E \left[|\tilde{V}_1 - T_1| \middle| \text{No Auto MM} \right] = (1 - \gamma_1)(1 + \gamma_1)r_1. \quad (23)$$

Because γ_1 and γ_2 are restricted to $(0, 1)$, and ϕ is restricted to $[0, 1/2)$ or $(1/2, 1]$, then it is straightforward to show,

$$E \left[|\tilde{V}_1 - T_1| \right] < E \left[|\tilde{V}_1 - T_1| \mid \text{No Auto MM} \right], \quad (24)$$

so that transaction prices in security 1 are closer to the end-of-period value for security 1 when the automated market maker is included. The same argument applies for security 2.

D Proposition 4

Consider the general case where a liquidity provider sets the transaction price of security 1 based on information Ω .

$$T_1(\Omega) = E[\tilde{V}_1 | \Omega] = p_1 + [P(V_1^+ | \Omega) - P(V_1^- | \Omega)]r_1. \quad (25)$$

The expected variance of the first price change, $T_1(\Omega) - p_1$, is,

$$E[(T_1(\Omega) - p_1)^2] = [(2P(V_1^+ | \Omega) - 1)r_1]^2. \quad (26)$$

The expected variance of the final price change, $\tilde{V}_1 - T_1(\Omega)$, is,

$$E[(\tilde{V}_1 - T_1(\Omega))^2] = P(V_1^+|\Omega)[2(1 - P(V_1^+|\Omega))r_1]^2 + \dots \quad (27)$$

$$[1 - P(V_1^+|\Omega)][2P(V_1^+|\Omega)r_1]^2, \quad (28)$$

$$= 4P(V_1^+|\Omega)[1 - P(V_1^+|\Omega)]r_1^2. \quad (29)$$

The average variance of a price change (averaging the variance for both the first and final price change) is therefore,

$$\text{Avg Variance} = \frac{1}{2}r_1^2. \quad (30)$$

Because volatility is just the square root of the average variance and because the average variance doesn't depend on Ω , volatility also doesn't depend on the information set, Ω , of the liquidity provider. The same argument holds for security 2.

V An Example

As an example, consider a two security market with the following parameters: $\gamma_1 = \gamma_2 = .5$, $P(V_1^+|V_2^+) = \phi = 0.9$, $p_1 = 50$, and $r_1 = 1$. Fig. 3 shows the transaction prices set for security 1 by the automated market maker conditioned on different order flow states. The market maker will price orders using only the highest and lowest price in the figure, which is 50.75 for a buy and 49.25 for a sell. The proportion of order flow he transacts is equal to 1/2 the proportion of orders at the




Figure 3: Diagram of the transaction prices for security 1 in the 4 possible order flow states, (B_1, B_2) , (B_1, S_2) , (S_1, B_2) , (S_1, S_2) , in a two security market. The values shown at earlier nodes are not transaction prices, but are the expected price at those nodes. The parameters used are: $\gamma_1 = \gamma_2 = .5$, $\phi = 0.9$, $p_1 = 50$, and $r_1 = 1$.

widest spread. Using Eq. 13, this is 30% of the orders for security 1. This means the automated market maker transacts the other 70%, and that the automated market maker transacts the majority of order flow.

The average spread in the price between buy orders and sell orders is 1 overall, 0.875 for liquidity traders, and 1.125 for informed investors. These values are shown in Fig. 4. Notice that the spread increases for informed investors and decreases for liquidity traders with the addition of the automated market maker.

The average difference between the transaction price and the end-of-period value of security 1 is 0.71875 when the automated market maker is present, and is 0.75 when she is not. This means prices are more efficient with the addition of the automated market maker.




Figure 4: Diagram of the average buy and sell transaction price for informed investors, liquidity traders, and unconditional on trader type. Notice that informed investors realize a larger spread than uninformed investors. The parameters used are $\gamma_1 = \gamma_2 = .5$, $\phi = 0.9$, $p_1 = 50$, and $r_1 = 1$.

The expected variance of the first price change is 0.34375 when the automated market maker is present and 0.25 when she is not. For the second price change, the expected variances are 0.65625 and 0.75 respectively. The overall average variance for a price change is 0.5 regardless of the presence of the automated market maker.

VI Extension of Model

Whereas before we assumed that liquidity traders were always willing to buy or sell, we now extend the model to allow liquidity traders to refrain from trading if their expected transaction cost is too high. Specifically we assume that the fraction of liquidity traders who are willing to submit an order is a monotonically decreasing function of their expected transaction cost. This means the following two propositions hold:

PROPOSITION 5. *In the extended model, with the addition of the automated market maker, the probability of a transaction in security i increases, i.e., expected volumes increase.*

When the automated market maker is added to the market, the expected transaction cost of a liquidity trader is reduced (Proposition 2). In the extension of the model, we have assumed this reduced cost increases the probability that a liquidity trader will place an order. Trivially, this increases the overall probability of a transaction.

PROPOSITION 6. *In the extended model, with the addition of the automated mar-*

ket maker, the spread in average transaction price between buys and sells in security i is reduced.

From the previous explanation for Proposition 5, we know that the addition of the automated liquidity provider increases the probability of a liquidity trader placing an order. This decreases the fraction of orders that come from informed investors. When the fraction of orders from informed investors decreases, adverse selection is reduced and liquidity providers can set smaller spreads for buy and sell orders.

Below, we prove the new propositions for a two security market. Proofs for markets with more than two securities are straightforward extensions of the two security case.

A Proposition 5

We denote the fraction of informed investors in security 1 by δ_1 and the fraction of liquidity traders who are willing to submit an order by π_1 . The probability of a transaction in security 1, P_1 , is therefore,

$$P_1 = \delta_1 + (1 - \delta_1)\pi_1. \tag{31}$$

The expected transaction cost of a liquidity trader is 1/2 the expected spread, which is $\Delta_{1,L}/2$ (Eq. 17). Without the automated market maker, this would be $\Delta_1/2$ (Eq. 18). Because π_1 is a monotonically decreasing function of the liquidity traders' expected transaction costs, and because $\Delta_{1,L}/2 < \Delta_1/2$, then π_1 is larger when the automated market maker is added. Therefore, the probability of a transaction in security 1,

P_1 , is also larger when the automated market maker is added. The same argument applies for security 2.

B Proposition 6

The fraction of orders for security i that come from informed investors is,

$$\gamma_1 = 1 - (1 - \delta_1)\pi_1. \quad (32)$$

As shown in the proof of Proposition 5, π_1 is larger when the automated market maker is added. Therefore, from Eq. 26, γ_1 is smaller. The unconditional average spread in price between buying and selling in security 1 is given in Eq. 19, $\Delta_1 = 2\gamma_1 r_1$, and is therefore reduced when the automated market maker is added. The same argument applies for security 2.

VII Conclusion

The purpose of this paper is to explain the increasing dominance of automated liquidity provision and to understand its effects on the market. In our model, automated liquidity providers price securities better than traditional market makers because they observe and react to order flow across securities. Consequently, they transact the majority of orders in the market and cause prices to be more efficient.

The presence of automated liquidity providers has material effects on investors in the market. Informed investors make smaller profits and uninformed investors lose

less money. Informed investors make smaller profits because they must now compete with one another across securities. Uninformed investors lose less money because they are able to transact through the liquidity provider to other uninformed investors in related securities. If the uninformed increase their trading activity due to lower transaction costs, overall volumes increase and overall transaction costs are reduced. These results match nicely with recently observed changes in global financial markets, where automated liquidity provision now dominates.

REFERENCES

- Amihud, Yakov and Haim Mendelson, 1988, Liquidity, volatility, and exchange automation, *Journal of Accounting, Auditing, and Finance*, 3(4), 369–395.
- Andrade, Sandro C., Charles Chang, and Mark S. Seasholes, 2008, Trading imbalances, predictable reversals, and cross-stock price pressure, *Journal of Financial Economics*, 88, 406–423.
- Angel, James, Lawrence Harris, and Chester S. Spatt, Equity trading in the 21st century, *Quarterly Journal of Finance*, 1(1).
- Baron, Matthew, Jonathan Brogaard, and Andrei Kirilenko, 2012, The trading profits of high frequency traders, working paper.
- Bernhardt, Dan, and Bart Taub, 2008, Cross-asset speculation in stock markets, *The Journal of Finance*, 63, 2385–2427.
- Biais, Bruno, Larry Glosten and Chester Spatt, 2005, Market microstructure: A survey of microfoundations, empirical results, and policy implications, *Journal of Financial Markets* 8, 217–264.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2013, Equilibrium fast trading, working paper.
- Black, Fischer, 1971, Toward a fully automated stock exchange, part I, *Financial Analysts Journal*, 27(4), 28–35.

- Boehmer, Ekkehart, Kingsley Fong, and Julie Wu, 2012, International evidence on algorithmic trading, working paper.
- Boulatov, Alex, Terrence Hendershott, and Dmitry Livdan, 2013, Informed trading and portfolio returns, *The Review of Economic Studies*, 80(1), 35–72.
- Brogaard, Jonathan, Björn Hagströmer, Lars L. Nordén, and Ryan Riordan, 2013, Trading fast and slow: colocation and market quality, working paper.
- Brogaard, Jonathan, Terrence Hendershott, Stefan Hunt, Torben Latza, Lucas Pedace, and Carla Ysusi, 2013, High-frequency trading and the execution costs of institutional investors, working paper.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2013, High frequency trading and price discovery, working paper.
- Budish, Eric, Peter Cramton, and John Shim, 2013, The high-frequency trading arms race: Frequent batch auctions as a market design response, working paper.
- Caballé, Jordi, and Murugappa Krishnan, 1994, Imperfect competition in a multi-security market with risk neutrality, *Econometrica*, 62, 695–704.
- Carrion, Allen, 2013, Very fast money: High-frequency trading on the NASDAQ, *Journal of Financial Markets*, 16(4), 680–711.
- Cartea, Álvaro, and José Penalva, 2012, Where is the value in high frequency trading?, *The Quarterly Journal of Finance* 2(3).
- Chaboud, Alain, Erik Hjalmarsson, Clara Vega, and Ben Chiquoine, 2013, Rise of the machines: Algorithmic trading in the foreign exchange market, *The Journal of Finance*, forthcoming.
- Chordia, Tarun, Asani Sarkar, and Avanidhar Subrahmanyam, 2011, Liquidity dynamics and cross-autocorrelations, *Journal of Financial and Quantitative Analysis*, 46(3), 709.
- Cvitanic, Jaksza, and Andrei Kirilenko, 2010, High frequency traders and asset prices, working paper.
- Demsetz, Harold, 1968, The cost of transacting, *The Quarterly Journal of Economics*, 82(1), 33–53.
- Domowitz, Ian and Benn Steil, 2002, Securities trading, *Technological Innovation and Economic Performance*.
- Easley, David, Nicholas M. Kiefer, Maureen O’Hara, and Joseph B. Paperman, 1996, Liquidity, information, and infrequently traded stocks, *The Journal of Finance*, 51(3), 811–833.

- Easley, David, Maureen O'Hara, and P.S. Srinivas, 1998, Option volume and stock prices: Evidence on where informed traders trade, *The Journal of Finance*, 53(2), 431–465.
- Ellis, Katrina, Roni Michaely, and Maureen O'Hara, 2002, The making of a dealer market: From entry to equilibrium in the trading of Nasdaq stocks, *The Journal of Finance*, 57(5), 2289–2316.
- Fama, Eugene F., 1970, Efficient capital markets: A review of theory and empirical work, *The Journal of Finance*, 25(2), 383–417.
- Foucault, Thierry, Johan Hombert, and Ioanid Rosu, 2013, News trading and speed, working paper.
- Fricke, Daniel and Austin Gerig, 2013, Liquidity risk, speculative trade, and the optimal latency of financial markets, working paper.
- Frino, Alex, Vito Mollica, and Robert I. Webb, 2013, The impact of co-location of securities exchanges' and traders' computer servers on market liquidity, *Journal of Futures Markets*, forthcoming.
- Gai, Jaiding, Chen Yao, and Mao Ye, 2013, The externalities of high-frequency trading, working paper.
- Garbade, Kenneth D. and William L. Silber, 1978, Technology, communication and the performance of financial markets: 1840-1975, *Journal of Finance*, 33(3), 819–832.
- Garbade, Kenneth D. and William L. Silber, 1979, Structural organization of secondary markets: Clearing frequency, dealer activity and liquidity risk, *Journal of Finance*, 34(3), 577–593.
- Gerig, Austin, 2012, High-frequency trading synchronizes prices in financial markets, working paper.
- Glosten, Lawrence R. and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics*, 14, 71–100.
- Hagströmer, Björn and Lars Nordén, 2013, The diversity of high-frequency traders, *Journal of Financial Markets*, 16(4), 741–770.
- Hansch, Oliver, Narayan Y. Naik, and S. Viswanathan, 1999, Preferencing, internalization, best execution, and dealer profits, *The Journal of Finance*, 59(5), 1799–1828.
- Hakansson, Nils H., Avraham Beja, and Jivendra Kale, 1985, On the feasibility of automated market making by a programmed specialist, *The Journal of Finance*, 40(1), 1–20.

- Hasbrouck, Joel and Gideon Saar, 2013, Low-latency trading, *Journal of Financial Markets*, 16(4), 646–679.
- Heidle, Hans G., 2004, Inferring information from trading, working paper.
- Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld, 2011, Does algorithmic trading improve liquidity?, *Journal of Finance*, 66(1), 1–33.
- Hendershott, Terrence and Pamela C. Moulton, 2011, Automation, speed, and stock market quality: The NYSE’s Hybrid, *Journal of Financial Markets* 14(4), 568–604.
- Hendershott, Terrence and Ryan Riordan, 2013, Algorithmic trading and the market for liquidity, *Journal of Financial and Quantitative Finance*, forthcoming.
- Hirschey, Nicholas, 2013, Do high-frequency traders anticipate buying and selling pressure?, working paper.
- Jain, Pankaj K., 2005, Financial market design and the equity premium: Electronic versus floor trading, *The Journal of Finance*, 60(6), 2955–2985.
- Jarrow, Robert A. and Philip Protter, 2012, A dysfunctional role of high frequency trading in electronic markets, *International Journal of Theoretical and Applied Finance*, 15(3).
- Jovanovic, Boyan and Albert J. Menkveld, 2012, Middlemen in limit-order markets, working paper.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.
- Kyle, Albert S., 1989, Informed speculation with imperfect competition, *The Review of Economic Studies*, 56(3), 317–355.
- Lapone, Andrew and Alexander Sacco, 2013, The impact of message traffic regulatory restrictions on market quality: Evidence from Chi-X Canada, working paper.
- Litzenberger, Robert, Jeff Castura, and Richard Gorelick, 2012, The impacts of automation and high frequency trading on market quality, *Annual Review of Financial Economics*, 4(1), 59–98.
- Mackenzie, Donald, 2013, Mechanizing the Merc: The Chicago Mercantile Exchange and the Rise of High-Frequency Trading, working paper.
- Malinova, Katya, Andreas Park, and Ryan Riordan, 2013, Do retail traders suffer from high-frequency traders?, working paper.
- Malinova, Katya, Andreas Park, and Ryan Riordan, 2013, Shifting sands: high frequency, retail, and institutional trading profits over time, working paper.

Martinez, Victor H. and Ioanid Rosu, 2013, High frequency traders, news and volatility, working paper.

Menkveld, Albert J., 2013, High frequency trading and the new market makers, *Journal of Financial Markets*, 16(4), 712–740.

Pasquariello, Paolo, and Clara Vega, 2013, Strategic cross-trading in the U.S. stock market, working paper.

Petersen, Mitchell A., 2004, Information: hard and soft, working paper.

Riordan, Ryan and Andreas Storkenmaier, 2012, Latency, liquidity and price discovery, *Journal of Financial Markets*, 15(4), 416–437.

Sofianos, George, 1995, Specialist gross trading revenues at the New York Stock Exchange, working paper.

Tong, Lin, 2013, A blessing or a curse? The impact of high frequency trading on institutional investors, working paper.

Tooke, Heather E., 2008, Information, trading, and produce market interactions: Cross-sectional implications of informed trading, *The Journal of Finance*, 63(1), 379–413.