How Does Algorithmic Trading Improve Market Quality?

Matthew R. Lyle, James P. Naughton, Brian M. Weller *

Kellogg School of Management

August 19, 2015

Abstract

We use a comprehensive panel of NYSE limit order book data to investigate the channel by which algorithmic trading (AT) improves market quality. We find that enhanced market maker monitoring explains the majority of improvements in liquidity and quoting efficiency during the 2000s. Market maker monitoring subsumes the ratio of order cancellations to total volume (a broad measure of AT) in accounting for improvements in market quality. Moreover, the residual variation in AT not associated with our AT market making proxy is typically associated with higher spreads, suggesting that different categories of algorithmic traders have distinct effects on market function. To distinguish decreased monitoring costs from potential confounds, we develop a stylized model of constrained market maker attention and empirically verify unique predictions concerning market maker behaviors around idiosyncratic versus multi-asset price jumps and small versus large stock price jumps. Our results provide a novel explanation for why spreads have not continued to fall since 2007 despite sustained increases in AT.

JEL: G12, G14, G27

Keywords: High-Frequency Market Making, Algorithmic Trading, Adverse Selection, Monitoring Costs

*We appreciate helpful suggestions and comments from Snehal Banerjee, Jonathan Brogaard, Terry Hendershott, Bob Korajczyk, Pamela Moulton, Ryan Riordan, Christian Westheide (discussant) and workshop participants at Northwestern University and the 2015 Conference on The Industrial Organization of Securities and Derivatives Markets. Michal Ziembiński contributed outstanding research assistance. We are grateful for the funding of this research by the Kellogg School of Management and the Lawrence Revsine Research Fellowship.

I. Introduction

Improvements in market quality enhance resource allocation, risk transfer, and price determination. Technology and market structure innovations of the 2000s contributed to dramatic improvements in one measure of market quality, the bid-ask spread. All else equal, lower spreads result in lower fees being paid to financial intermediaries, which in turn contribute to more effective financial market function. The economic savings associated with the decline in spreads during the 2000s is substantial. The median daily total spread costs across all securities in the trades and quotes database (TAQ) were \$127 million on a dollar volume of \$101 billion in January 2001, compared with only \$41 million on a much larger volume of \$168 billion in January 2007.[1]

Prior market structure research has focused on the role of algorithmic trading (AT) in improving market quality. AT is defined as the use of computer algorithms to automatically make certain trading decisions, submit orders, and manage those orders after submission (Hendershott, Jones, and Menkveld (2011), Jones (2013)).[2] AT has grown from a negligible size in the mid-1990s to a dominant component of equity market structure, with implications for all aspects of equity market performance. The majority of empirical evidence suggests that AT reduces spreads and improves price efficiency (e.g., Jones (2013)), though some studies using more recent data indicate that the beneficial effect of AT may not be robust across time periods (e.g., Brogaard, Hendershott, and Riordan (2014), Menkveld and Zoican (2015)). In contrast, the majority of theoretical studies suggest that AT can increase adverse selection costs, and hence lead to a deterioration in market quality (e.g., Biais, Foucault, and Moinas (2015), Foucault, Hombert, and Rosu (2015), Han, Khapko, and Kyle (2014)).

We contribute to the line of research that examines the relation between AT and mar-

[1] We exclude outliers potentially attributable to data errors—specifically, we drop all observations with effective spreads exceeding three mean absolute deviations from the stock-day median.

[2] Some studies focus on high-frequency trading (HFT), which is a subset of AT Hendershott et al. (2011). HFT often rely on high-speed computers, co-location services and low latency trading. A formal definition of HFT is provided in Securities Exchange Act Release No. 34-61358, 75 FR 3594, 3606 (January 21, 2010).

ket quality by identifying improved market-maker monitoring as a key channel through which trading technology improves market quality.[3] Few papers in the AT literature distinguish among the impacts of diverse AT participants. Our paper joins Hagströmer and Norden (2013), Biais and Foucault (2014), and others in addressing this shortcoming and demonstrating its importance to understanding the role of AT in financial markets. Our results suggest that AT liquidity suppliers (e.g., market making strategy) improve market quality, whereas AT liquidity demands (e.g., early-information/speed-exploitation strategy) have harmed market quality. Of equal significance, we explain why subsequent increases in AT after 2007 are not associated with further reductions in trading costs, thus rationalizing the seemingly inconsistent empirical evidence stemming from different time periods.

Isolating the mechanism by which AT improves market quality is important because large-scale automation and investment in trading technology have continued apace well beyond 2007 (see, e.g., Pagnotta and Philippon (2015)), the year in which average bid-ask spreads reached their minimum in U.S. stocks. In addition, the potential mechanisms differ sharply in their welfare implications. For example, increased competition among market makers redistributes rents from intermediaries to traders, whereas increased quoting efficiency decreases costs for less-informed traders at the expense of sophisticated market participants.

We conjecture that technological improvements enhance the monitoring ability of market makers, thus enabling them to update quotes more efficiently and avoid being picked off on stale quotes. Market makers are "picked off" when the quoted price does not reflect information available to other market participants, thereby resulting in losses to the intermediary when trade occurs. Foucault, Röell, and Sandås (2003) and others show theoretically that picking-off risk is quite costly to market makers, and particularly so toward the end of the 1990s. This conjecture is also consistent with Biais, Declerck, and Moinas (2015), who find that monitoring is more effective than speed at mitigating

[3] We use the term "market maker" to stand in for the collection of marginal liquidity suppliers.

picking off risk.

We measure picking-off risk using detailed NYSE limit-order book data for a comprehensive panel of S&P 1500 stocks over the 2002–2007 period. In our time series regressions, we use the returns to a strategy where trades are based on the price-weighted order book imbalance to proxy for picking-off risk. The trading strategy buys stocks in the highest quintile of order book imbalance and sells stocks in the lowest quintile of order book imbalance and holds this portfolio for five minutes. We find that the returns to this strategy are very large, but attenuate over time: the resulting five-minute Sharpe ratios decrease from 0.98 in 2002 to 0.23 in 2007, with similar magnitudes and decay rates at longer holding horizons. We use these returns to proxy for market maker attentiveness because they translate directly into the cost of potential market maker losses.[4] In addition to our time-series analyses, we also conduct tests using panels of stock-level data to achieve better identification against potential alternatives. In these analyses, we use the firm's interquartile range of daily log order book imbalance (IQR) to proxy for market maker attentiveness as it is not possible to calculate portfolio returns when the unit of analysis is the firm.

Both sets of tests establish two main results relating market maker monitoring to improvements in liquidity and quoting efficiency. First, we find that our measure of market maker monitoring better explains decreases in effective spreads relative to the ratio of order cancellations to total volume, which is a commonly used measure of the overall level of AT. In our empirical tests, we find that market maker monitoring is strongly associated with spreads using both levels and changes specifications for both time-series average and panel regressions. The overall level of AT is negatively associated with spreads in some specifications, consistent with prior work, but this association is sensitive to the inclusion of our proxy for market maker monitoring. Moreover, AT is

[4]Inventory risk may also lead to order book imbalance (e.g., Comerton-Forde et al. (2010)). For example, if market makers have significant long positions, they may offer to sell more shares or offer to sell at more aggressive prices, which would give rise to an unbalanced order book even if they are very attentive to the market. However, this type of activity would only serve to lower the returns we document, because the strategy we employ goes with rather than against the weight of the order book.

generally *positively* associated with spreads in a differences specification that nets out common secular trends to isolate weekly variation in liquidity provision. Our results indicate that it is improvements in market maker monitoring, and not simply increases in AT, which has led to the reduction in spreads. In fact, our results suggest that increases in AT increase spreads, holding constant improvements in market maker monitoring.

Second, we find that our measure of market maker monitoring better explains improvements in quoting efficiency relative to the ratio of order cancellations to total volume. We examine quoting efficiency using the adverse selection component of the effective spread (Glosten (1987)). This measure captures quoting efficiency because, all else equal, smaller permanent price impacts imply improved steady-state price efficiency of quotes. We find that time series and panel variation in market maker monitoring have a powerful effect for explaining the sharp decline in market maker adverse selection costs. In fact, the association between the adverse selection component of spreads and market maker monitoring is almost identical to the association between effective spreads and market maker monitoring. This result is shown graphically in Figure I. Market maker monitoring, which we proxy for using the mean signal return, tracks the overall decline in liquidity much better than the ratio of order cancellations to total volume. This finding accords with our conjectured channel because adverse selection encapsulates picking-off risk, and as such, it should respond similarly to effective spreads if equilibrium bid-ask spreads decline via improvements in market maker monitoring.

The credibility of our conclusions is based, in part, on the accuracy of our measure of market making monitoring. Therefore, we use two additional sets of analyses to mitigate the concern that our measure is capturing another construct. First, we develop a stylized model that assumes a capacity constraint on market maker attention, and develop from this model several cross-sectional predictions that we use to provide additional support for our conclusions. In the model, market makers allocate attention so as to minimize picking off losses on a portfolio of intermediated stocks. Consistent with our model, we find that order book information is incorporated into prices more rapidly in recent periods (due to technological advances) and for larger stocks (due to the greater expected losses from picking off). In addition, we find that our measure of market maker monitoring varies across the spectrum of jump types in such a way that market makers monitor the least for idiosyncratic price jumps, more for multi-asset jumps, and the most for non-idiosyncratic jumps affecting many assets. The differential response to non-idiosyncratic versus idiosyncratic jumps and small- versus large-stock jumps are consistent with optimal allocation of attention, but it is difficult to rationalize with alternative interpretations of our measure of market maker monitoring.

Second, we conduct two instrumental variables analyses to address potential confounds arising from simultaneous causality and latent drivers of our measure of market maker monitoring and adverse selection. Our first IV approach mitigates the concern that variation in adverse selection mechanically drives the market maker monitoring measure for reasons other than market maker attentiveness. Under this approach, we instrument IQR (our proxy for market maker monitoring) using a jump indicator. Intuitively, this approach isolates the component of lagged IQR associated with features arising only from exposure to jump risk. Our second IV approach mitigates the concern that IQR and market quality might be simultaneously caused by adverse selection or jointly driven by a contemporaneous causal relation. Under this approach, we instrument for IQR using its one-week lagged value. Because algorithmic market makers can enter and exit the market at will, future adverse selection linked to current market conditions should be entirely avoidable a full week ahead. The conclusions from both IV specifications mirror those in our main analyses. Improvements in market maker monitoring drive innovations in market quality, and AT leads to a deterioration in market quality holding market maker monitoring constant.

Our results suggest that improved market-maker monitoring plays a key role in understanding the decline in trading costs and improvements in price efficiency during the 2000s. Our findings extend Hendershott et al. (2011) and similar studies that examine the effect of AT on market quality in two ways. First, our results suggest that there are different capital market consequences based on the type of algorithmic trader considered. We find that market quality improves for algorithmic liquidity supplying strategies, but worsens for algorithmic liquidity consuming strategies. This result suggests that improvements in market maker monitoring, and not simply increases in AT, have led to improvements in market quality. Second, our use of a long time series enables us to show that AT technology improves market quality only up to a point. Once market makers are sufficiently attentive across the spectrum of intermediated securities, quoting efficiency stagnates and spreads do not decline further. This result provides a potential explanation as to why spreads have not continued to fall since 2007 despite sustained increases in AT, even controlling for the volatility associated with macroeconomic phenomena.

In addition, our approach of isolating the mechanism by which trading technology improves market function and testing the effect of this mechanism over a long time series provides a potential explanation for conflicting empirical results in the AT literature across different time periods. In particular, our results suggest that studies that utilize data during the early part of our sample period are likely to find a negative association between AT and measures of market quality (e.g., Hendershott et al. (2011)). However, these results may reverse in later periods as the continued increase in AT is not accompanied by a corresponding improvement in market maker monitoring (e.g., Brogaard et al. (2014), Menkveld and Zoican (2015)).

Our study also has direct policy relevance for interpreting the role of algorithmic trading in capital markets. Our results call into question whether continued investment in high-frequency trading technology produces tangible benefits for financial markets. We find that average picking-off costs are indistinguishable from zero by late 2006, which suggests that most gains from improved market-maker monitoring have long since been achieved.

Our paper proceeds as follows. Section II discusses related literature. Section III describes our NYSE limit order book data. Section IV establishes our main results relating market maker monitoring to declines in effective spreads and adverse selection.

Section V develops a model of market maker attentiveness and provides additional panel and instrumental variables analyses to provide additional resolution on our hypothesized mechanism. Section VI concludes.

II. Related Literature

There are a number of different AT strategies employed by a diverse set of market participants. For example, some market participants design algorithms to compete with designated market-makers and other liquidity suppliers (e.g., Jovanovic and Menkveld (2011)). Others design algorithms to take advantage of differential prices for the same assets across exchanges, or to process large amounts of statistical data to generate profitable trade opportunities (e.g., Foucault and Menkveld (2008)). Because the data available to researchers typically do not identify the specific organization undertaking each trade, it is not possible to directly observe these different market participants or their underlying trade strategies. As a result, prior studies have generally used measures that capture the aggregate presence of these diverse market participants to conclude that AT reduces liquidity and improves market quality.

For example, Hendershott et al. (2011) study the implementation of autoquote at the New York Stock Exchange in 2003. Autoquote facilitated AT by allowing computer algorithms to submit and cancel orders and to see those orders quickly reflected in the NYSE's disseminated quote. Hendershott et al. (2011) find that effective spreads narrow, adverse selection is reduced, and more price discovery takes place through quotes after the implementation of autoquote. They attribute these changes to the increase in AT following autoquote, and thus conclude that AT improves market quality. Similarly, Hasbrouck and Saar (2013) find that low-latency trading, a characteristic of HFT, reduces short term volatility and enhances market liquidity.

These empirical findings contrast with the conclusions of several theoretical models. Biais et al. (2015) develop a model where the intermediation provided by HFT helps traders find counterparties, leading to gains from trade. However, they also show that

HFT can trade on new information more quickly, which generates adverse selection costs. Similarly, Foucault et al. (2015) develop a model where an HFT processes interim information before other traders. The HFT's activities sharpen adverse selection experienced by market makers, thereby worsening liquidity. Han et al. (2014) reach similar conclusions in a model of fleeting quotes by HFT market makers.

Our approach differs from prior work because we develop a proxy for market maker monitoring that allows us to separate changes in market maker monitoring from the overall change in AT. By isolating a key mechanism, we leave open the possibility that theoretical results relating AT to worsened market quality still obtain empirically but are swamped by the effects of monitoring innovations. Indeed, as we will outline in more detail later in the paper, we find in several empirical specifications that controlling for market monitoring effaces liquidity gains that would otherwise be associated with the aggregate level of algorithmic trading.

III. Data

We use New York Stock Exchange (NYSE) Openbook Data. Unlike standard trade and quote (TAQ) data, which provides information about inside quotes only (i.e., the best bid and offer on an exchange), the data we use constitutes a history of the *entire* limit order book for all NYSE traded securities. Specifically, we use NYSE aggregated order book history files, which provide full order book data every second and are available from 2002 to present.[5] The aggregate history files are unique in providing resolution on the pre-2008 period during which spreads decline most significantly and during which the NYSE is the dominant exchange for NYSE-listed securities.[6] Additional details on these data are available at http://www.nyxdata.com/openbook/.

Use of limit-order book data uniquely reveals potential sources of picking-off risk experienced by market makers. For example, the standard trades and quotes (TAQ)

[5]The NYSE Ultra history files provide millisecond resolution but only become available in 2008.
[6]The NYSE is considered the dominant exchange for NYSE listed stocks up to 2006–2007. (see, e.g., Holden and Jacobsen (2014).)

database only provides quotes at the best bid and offer, thereby masking large trader demands away from the quote midpoint that are informative for pricing. We limit our sample to companies that were in the S&P 1500 index for each year of the sample period, resulting in 773 unique firms. Summary statistics for our sample are provided in Table I.

Our non-order book data comes from the TAQ monthly files. We calculate the bid-ask spread following Holden and Jacobsen (2014) and we calculate effective spreads and their adverse selection component using the same approach as Hendershott et al. (2011). For the NYSE, effective spreads are more meaningful than quoted spreads because specialists and floor brokers are sometimes willing to trade at prices within the quoted bid and ask prices. For each stock and five-minute window, we use all NYSE trades and quotes to calculate quoted and effective spreads for each reported transaction and compute a volume-weighted average.

IV. Main Results

We present our results in three parts. First, we describe how we generate our estimate of market maker monitoring using a long-short trading strategy based on order book imbalance. Second, we show that this measure of market maker monitoring is a strong predictor of changes in effective spreads using both time-series and panel regressions. Finally, we show that it is also a strong predictor of changes in adverse selection using both time-series and panel regressions.

A. Estimating Market Maker Monitoring using Picking-off Risk

Order book imbalance indicates that the price is not centered with respect to traders' displayed information. Imbalances suggest that traders value the asset significantly more or less, on average, than the mid price, which is unsustainable in a market with resale. Therefore, an attentive market maker should utilize order book information and update the price to reflect value-weighted demands for buying and selling the asset. If the market maker does not update in response to either a demand or supply imbalance, then the market maker is more likely to be counter party to a particular trade. As a result, traders who incorporate information from the limit order book can pick off inattentive market makers.

We generate a proxy for the potential losses to the market maker due to picking off from a trading strategy that uses asymmetry in the limit-order book to make trades. We begin by calculating order book imbalance every 5 minutes from 10:00 am to 3:00 pm using the following formula:

$$OrderImbalance = \log\left(\frac{depth_{1\%}^{buy}}{depth_{1\%}^{sell}}\right), \tag{1}$$

where $depth_{1\%}$ is defined as dollar-weighted depth within one percent of the reference price,

$$depth_{1\%} = \sum_{i=1} price_i \cdot depth_i \cdot \mathbf{1}_{\frac{|price_i - \text{refprc}|}{\text{refprc}} \in [0,1\%]}, \tag{2}$$

and we take logs to obtain a more symmetric measure.[7]

Next, we split stocks into five portfolios sorted on order imbalance. Because our measure of order book imbalance is signed, portfolio 1 (5) includes stocks where there are significantly more sell (buy) orders than buy (sell) orders. We calculate the average return in each portfolio and the average return on a zero-cost portfolio that buys portfolio 5 (i.e., the portfolio where the securities have significantly more limit buy orders) and sells portfolio 1 (i.e., the portfolio where the securities have significantly more limit sell orders) for each 5-minute interval. We measure returns over the subsequent 5-, 15-, 30- and 60- minute interval. We refer to the average return as the mean signal return. We refer to the mean signal return scaled by the daily standard deviation from this strategy as the Sharpe Ratio. The annual values are shown in Table II.

The average returns gross of transactions costs are large. In 2002, the average strategy

[7]This calculation only includes buy and sell orders within 1% of the reference price, which is the last executed price on the exchange. The choice of a 1% cutoff is based on visual inspection of limit-order books over our sample period and the desire to avoid incorporating non-credible orders deeper in the book. To ensure that our results are not sensitive to the choice of cutoff, we repeat our analysis using cutoffs of 0.25%, 0.5%, and 0.75%, and our conclusions are unchanged.

return is 4.7 basis points per 5 minute interval–this return represents revenues that market makers give up to traders by not monitoring more attentively. Our calculations do not reflect trading costs, which of course would attenuate trader returns to this strategy. Table II shows that the average returns and Sharpe Ratio decay over time for each look ahead horizon. Average strategy returns decay to near-zero levels by 2005 (0.7 basis points) and are flat through 2007 (0.6 basis points).

B. The Association between Market Maker Monitoring and Effective Spreads

Our first set of tests examines how market maker monitoring impacts effective spreads. We use two empirical approaches. The first approach uses the mean signal return to explain time-series variation in average spreads, where the mean signal return is the average return to the 5-minute portfolio described in Section A. The second approach uses the firm's interquartile range of daily log order book imbalance (IQR) to proxy for market maker attentiveness in a panel regression. The mean signal return cannot be used for panel tests because it is based on a comparison between diversified portfolios. While the IQR is not as precise a measure of market maker monitoring, it does allow us to identify the relative contributions of improved monitoring and AT to the decline in effective spreads in a manner that facilitates comparison with prior work (e.g., Hendershott et al. (2011), Brogaard et al. (2014)).

B.1. Time-Series Evidence

We examine the time-series relation between picking off risk and effective spreads using the following specification:

$$Spread_t = \beta_0 + \beta_1 MeanRet_t + \beta_2 VIX_t + \beta_3 AT_t + \epsilon_t. \tag{3}$$

We calculate effective spreads using the same approach as Hendershott et al. (2011). The determination of the mean signal return ($MeanRet$) is described in Section A. We measure AT using the ratio of order cancellations to total volume (Hendershott et al. (2011)). Specifically, for each second order book snapshot, we look at the number of orders at each price point in the book and subtract from this all orders at the same price points in the prior second (the last snapshot). The total reduction in orders in the book must come from cancellations or from executions (C+E). In order to determine the cancellations, we subtract from C+E the total NYSE executions recorded in the TAQ files (E). Our measure of algorithmic trading is then $AT = C/E$, which is the total reduction in liquidity in the book that was not caused by NYSE order executions divided by total volume.[8] We use VIX as an additional explanatory variable because aggregate market risk is highly correlated with spreads over time (e.g., Chung and Chuwonganant (2014)). All variables are measured on a daily basis.

The results of Equation (3) are provided in Table III. Panel A uses a levels specification, and Panel B uses a changes specification. In each panel, Column (1) uses the mean signal return, Column (2) uses AT, and Column (3) combines both explanatory variables. The purpose of presenting the results in this way is to provide insights into the separate and joint effect of improved monitoring and AT on effective spreads.

The coefficient on the mean signal return in Panel A is positive and highly significant in Column (1). This indicates that lower values of picking-off risk are associated with lower spreads. The coefficient on AT in Panel A is negative and highly significant in Column (2). The negative sign indicates that an increase in AT is associated with lower spreads, consistent with prior research. The results in Column (3) show that the inclusion of both the mean signal return and AT heavily influence the coefficients. The coefficient on the mean signal return remains positive and significant, but the magnitude drops by approximately 30 percent when AT is included. The effect of including both regressors on the AT coefficient is even stronger. While that coefficient also remains statistically significant, the magnitude of the coefficient drops by almost 60 percent.

[8] This measurement is subject to multiple sources of error: odd lot trades, hidden orders, and very short-lived quotes are not captured in our level book snapshots. We anticipate that these biases are relatively small early in the sample, but potentially increase in importance as order shredding and sub-minute order cancellations become more prevalent.

In Panel B, we provide additional support for these conclusions using a changes specification across adjacent trading weeks. This approach identifies off of weekly variation in liquidity rather than low-frequency variation in liquidity or common secular trends.[9] Focusing on higher-frequency variation in liquidity makes unobserved confounds less likely to contaminate our results. At the same time, in light of the empirical results of Foucault et al. (2014), we anticipate that market maker attentiveness and picking-off costs should meaningfully (co)vary at this frequency.

The results of the changes specification in Panel B are striking. The coefficient on the mean signal return is positive and highly significant in both Column (1) and Column (3). Moreover, the inclusion of AT has virtually no effect on the coefficient on the mean signal return. This finding suggests that picking-off risk explains a significant proportion of weekly variation in spreads. By contrast, the coefficient on AT is insignificant, which suggests that there is no association between AT and effective spreads.

The coefficient on the mean signal return is also economically meaningful. Both the mean signal return and the effective spread are expressed in basis points. In Panel B, column (3), this implies that a 1 basis point increase in the mean signal return corresponds to a 0.1 basis point increase in the effective spread. Since the mean signal return has a standard deviation of 3.2 basis points, this implies that a one standard deviation increase in the mean signal return corresponds to an increase in the effective spread of approximately 0.4 basis points. In addition, over the 2002 through 2007 period, the mean signal return dropped by approximately 10.5 basis points. This implies that approximately 1.3 basis points (or about 15%) of the overall decline in the effective spread is attributable to the reduction in the mean signal return.

[9] In robustness tests, we eliminate Fridays and Mondays from our weekly averages, and we also examine daily rather than weekly changes. Neither of these robustness checks affects our inferences.

B.2. Panel Data Evidence

We examine the relation between picking off risk and effective spreads using firm-day data and the following specification:

$$Spread_{it} = \beta_0 + \beta_1 IQR_{it} + \beta_2 AT_{it} + \beta_3 Controls_{it} + \epsilon_{it}. \quad (4)$$

IQR is our proxy for market maker monitoring. This variable is equal to the interquartile range of the log order imbalance of stock i on date t. We measure AT using the ratio of order cancellations to total volume using the same procedure outlined in Section IV.B.1. The control variables we use are daily share turnover, log intraday price range (a dispersion proxy based on Parkinson (1980)), the inverse of the closing share price (1/Price), log market capitalization (Size), VIX, and PIN (Easley and O'hara (1987)). We include PIN to control for the possibility that there is a mechanical relation between market maker monitoring and spreads through their shared associated with adverse selection. We include fixed effects for each stock, and we also include year fixed effects in certain specifications to control for unobserved factors.

Table IV presents the results using effective spreads and a levels specification. IQR is positive and statistically significant in each of Columns (1), (3) and (4). This indicates that improvements in market maker attentiveness are associated with lower effective spreads. The results for AT are not nearly as consistent. Column (2) presents the results of a regression using AT and the control variables used in Hendershott et al. (2011). This specific result indicates that higher levels of AT are associated with lower spreads, consistent with their work.

However, when we combine IQR with AT in Column (3), the sign of the coefficient on AT becomes insignificant. This indicates that after controlling for improved monitoring by market markers, AT is not associated with the effective spread. The coefficient on AT is highly significant in Column (4) when we include VIX and year fixed effects. However, the sign of this coefficient suggests the opposite relation between the effective spread and

AT when compared with Column (2). That is, Column (4) suggests that higher levels of AT are associated with higher effective spreads. Overall, the results in Table IV provide strong support for the association between market maker monitoring and the decline in spreads and weak support for the association between AT and the decline in spreads.

We investigate the robustness of these conclusions using a changes specification, the results of which are provided in Table V. The coefficients on both IQR and AT are positive and statistically significant in each specification. This suggests that an improvement in market maker monitoring is associated with a decline in spreads, consistent with our hypothesis. However, the coefficients on AT indicate that increased AT increases effective spreads. Across each column, the coefficients are relatively stable, suggesting that the bias associated with including both variables is an issue in levels rather than changes regressions.

C. The Association between Market Maker Monitoring and Quoting Efficiency

We use the same empirical strategy as the prior section to examine the ability of improved market maker monitoring and increased AT to explain improved quoting efficiency, measured using the adverse selection component of spreads. As before, the first empirical approach uses time-series analyses with variables aggregated to the daily level and the second empirical approach uses firm-day panel analyses.

C.1. Time-Series Evidence

We examine the time-series relation between picking off risk and the adverse selection component of spreads using the following specification:

$$AdvSel_t = \beta_0 + \beta_1 MeanRet_t + \beta_2 VIX_t + \beta_3 AT_t + \epsilon_t. \tag{5}$$

We calculate the adverse selection component of spreads using the same approach as Hendershott et al. (2011). All other variables are defined the same as in Equation (3).

The results in Table VI are shown using a levels specification in Panel A and a changes specification in Panel B. As with Table III, Column (1) uses the mean signal return, Column (2) uses AT, and Column (3) combines both explanatory variables.

The coefficient on the mean signal return is highly significant in both Column (1) and Column (3). The sign of the coefficient indicates that improved market maker monitoring is associated with lower adverse selection, and hence improved quoting efficiency. The fact that the results in Table VI are somewhat stronger than Table III is reassuring, since prior research has shown that most of the variation in effective spreads is due to the adverse selection component.

The coefficient on AT is highly significant in Column (2). The negative sign indicates that an increase in AT is associated with lower spreads, consistent with prior research. The inclusion of the mean signal return variable in Column (3) once again has a dramatic impact on the coefficient on AT. The coefficient on AT is now insignificant. This is consistent with AT having no association with changes in spreads over our sample period.

The results of the changes specification in Panel B of Table VI are consistent for the mean signal return, but different for AT. The coefficient on the mean signal return is positive and highly significant in both Column (1) and Column (3). Moreover, the inclusion of AT has virtually no effect on the coefficient on the mean signal return. In contrast, the coefficient on AT is sensitive to the inclusion of the mean signal return. The economic magnitude of the coefficient drops by approximately 40 percent in Column (3) when compared with Column (1).

The coefficient on the mean signal return in Table VI is more economically powerful than the coefficient in Table III. Both the mean signal return and adverse selection variables are expressed in basis points in Table VI. In Panel B, column (3), this implies that a 1 basis point increase in the mean signal return corresponds to a 0.25 basis point increase in adverse selection. In addition, because the mean signal return dropped by approximately 10.5 basis points from 2002 through 2007, this implies that approximately 2.6 basis points (or about 32%) of the overall decline in the adverse selection component of spreads is attributable to the reduction in the mean signal return.

C.2. Panel Data Evidence

We examine the relation market maker monitoring and quoting efficiency, measured using the adverse selection component of spreads, using firm-day level data and the following specification:

$$Spread_{it} = \beta_0 + \beta_1 IQR + \beta_2 AT_{it} + \beta_3 Controls_{it} + \epsilon_{it}. \tag{6}$$

All variables are defined the same as in Equation (3). Panel A of Table IV presents the results using a levels specification, and Panel B presents the results using a changes specification. In Panel A, IQR is positive and statistically significant in each of Columns (1), (3) and (4). This indicates that improved market maker monitoring is associated with lower adverse selection.

Once again, the results for AT are not nearly as consistent. When we combine IQR and AT in Column (3), the magnitude of the coefficient on AT drops by approximately 60 percent relative to Column (2). Even more striking, the sign on AT becomes positive in Column (4) when we include VIX and year fixed effects. This indicates that after controlling for improved monitoring by market markers, higher levels of AT increase adverse selection.

This result is similar to those presented in Table IV. However, unlike Table IV, the switching of the sign of the coefficient on AT is not driven by the inclusion of IQR. Rather, it is the inclusion of VIX and year fixed effects that cause the coefficient on AT to become positive.

We investigate the robustness of our conclusions on the relation between market maker monitoring and adverse selection using a changes specification. The results in Table VIII show that the coefficient on IQR is positive and statistically significant. By contrast, our results suggest that increased AT participation has no association with weekly variation in adverse selection costs to market makers.

V. Additional Analyses

The results in Tables IV through VIII suggest that improved market maker monitoring has contributed significantly to the decline in spreads. In contrast, while we find that AT is negatively associated with spreads in a levels specification (consistent with prior work), we find that this association is sensitive to the inclusion of our proxy for market maker monitoring and certain control variables. In fact, several of our specifications suggest that AT may worsen market quality, consistent with the theoretical predictions of Foucault et al. (2015) and empirical evidence of Brogaard et al. (2014), but inconsistent with several prior empirical studies.

In this section, we provide further support for these conclusions by conducting a series of cross-sectional and instrumental variables regression analyses. Section A outlines a stylized model of market maker attentiveness that we use to motivate our cross-sectional analyses in section B. The model assumes that market makers face a capacity constraint in their ability to monitor their portfolio and that they allocate attention to minimize picking off losses. The concept of limited attention has a foundation in prior empirical studies. For example, Chakrabarty and Moulton (2012) show that when one stock in a specialist's portfolio has an earnings announcement, the liquidity of other stocks covered by the same specialist worsens. The third section outlines two separate instrumental variables regressions that are designed to address measurement and reverse-causality issues with our proxy for market maker monitoring.

A. Model of Market Maker Attentiveness

In this section we introduce a reduced-form model highlighting essential trade-offs of market makers in allocating attention among a set of intermediated securities. We purposefully abstract from several realistic additions to focus on trader arrival rates, information arrival rates, and non-idiosyncratic versus idiosyncratic shocks. Our guiding model for distinguishing tests of the market maker monitoring channel relates most closely to Van Nieuwerburgh and Veldkamp (2010), who show that investors concentrate information gathering for assets most likely to be held or for factors most likely to influence their returns. Although we fix the set of assets considered, our model can readily be extended to the case in which market makers specialize and intermediate only for a subset of available securities. To the best of our knowledge, this simple model is the first to endogenize information acquisition in market making.

A.1. Model Setup

A competitive market making sector optimally allocates attention across N securities in a discrete-time economy with time indexed by $T = 1, 2, 3, \ldots$. At the start of each period, the underlying value of each security updates as

$$\Delta p_{it} = \xi_{it} + \beta_{im}\xi_{mt} \tag{7}$$

where ξ_{it} and ξ_{mt} represent conditional jump realizations for idiosyncratic shocks and a market factor.[10] We assume without loss of generality that these realizations are large relative to potential equilibrium spreads.[11] Price jumps arrive with fixed probabilities λ_i and λ_m. These time intervals are assumed to be quite small (e.g., on the order of a second), such that moderate price jumps can, but do not always, occur, and co-jumps between idiosyncratic shocks and the market shock are of negligible probability.

Stocks are differentiated by the arrival rate of liquidity traders and arbitrageurs and by the exposure to the market jump β_{im}.[12] Arbitrageurs arrive with Poisson intensity η_i, and liquidity traders arrive at rate $k\eta_i$. Upon arrival, an arbitrageur compares the

[10] Our model is best understood as treating jumps associated with unscheduled news; however, we can allow attention parameters to vary over time to consider known announcement dates. This extension would rationalize the pre-announcement withdrawal of liquidity documented by **?** and **?**.

[11] jumps smaller than the spread are not "picked off" and hence are not costly to the market maker. Our formulation eliminates these jumps by suitably adjusting the jump arrival rate downward and truncating the jump CDFs.

[12] In a richer model, the idiosyncratic jump intensities and jump distributions may vary across stocks, but we do not exploit this variation in our empirical tests.

displayed price against the "intrinsic" value of the asset, and he trades a single unit if these values differ by more than the spread. At the end of each period, trades execute and the market maker posts new limit orders to facilitate liquidity provision.

Market makers scan for fundamental value jumps at an endogenous intensity of ψ_i for idiosyncratic jumps and ψ_m for non-idiosyncratic jumps. We motivate arrival rates with a Poisson process at the end of each period that governs the competition between the market maker and the arbitrageur. Specifically, for a scanning rate of ψ_i, when a non-idiosyncratic jump in the asset value occurs, the probability that the market maker incurs a loss by arriving after the arbitrageur is

$$\Pr\left(FT \text{ arrives before } MM\right) = \frac{\eta_i}{\eta_i + \psi_m} \tag{8}$$

The cost function is increasing in combined attentiveness. For simplicity, we assume a linear form:

$$C\left(\left\{\psi_i\right\}, \psi_m\right) = c \times \left(\sum_i \psi_i + \psi_m\right) \tag{9}$$

This formulation has a physical interpretation associated with allocating processor clock cycles, internet bandwidth, or employee focus under a scalable computing technology.[13] For example, over a given millisecond, a single computer can assess the underlying asset value once, on average, whereas three computers can check and assess three times for one security during the same span at three times the cost. Given a half-spread h, the market maker purchases computing power for each asset ψ_i^* to infer value from the order book for firm-specific shocks and ψ_m^* to infer potential market innovations.

[13]Equivalently, market makers face attention constraints, where the attention budget is given by

$$\sum_i \psi_i + \psi_m \leq C. \tag{10}$$

This formulation better describes allocation of fixed computing capacity, although there is a one-to-one correspondence of attention constraints C to shadow costs of capacity c.

A.2. Model Solution

When a market jump occurs and the market maker is not the first to respond, the expected loss is $\beta E\left[|\xi_{mt}|\right]$. Likewise, for idiosyncratic jumps, the respective probabilities and expected losses are $\frac{\eta_i}{\eta_i+\psi_i}$ and $E\left[|\xi_{it}|\right]$.

For a half-spread h_i, the expected gains per unit time are $k\eta_i h_i$ and the expected losses per unit time are

$$\frac{\eta_i}{\eta_i+\psi_m}\left(\beta_{im}\lambda_m E\left[\xi_{mt}\right]\right)+\frac{\eta_i}{\eta_i+\psi_i}\left(\lambda_i E\left[\xi_{it}\right]\right) \tag{11}$$

Each market maker takes the half-spreads $\{h_i\}_{i=1}^N$ as given. Expected profits are maximized with

$$\max\sum_i \eta_i\left((1+k)h_i-\frac{\beta_{im}\lambda_m E\left[|\xi_{mt}|\right]}{\eta_i+\psi_m}-\frac{\lambda_i E\left[|\xi_{it}|\right]}{\eta_i+\psi_i}\right)-c\left(\sum_i\psi_i+\psi_m\right) \tag{12}$$
$$\text{s.t. } \psi_i,\psi_m\geq 0$$

Taking first-order conditions with respect to the decision variables (holding spreads fixed) obtains the following system:

$$\begin{aligned}\psi_i^* &= \max\left(\sqrt{\frac{1}{c}\eta_i\lambda_i E\left[|\xi_{it}|\right]}-\eta_i,0\right) & (13)\\ 0 &= \sum_i\frac{\eta_i\beta_{im}}{(\eta_i+\psi_m^*)^2}-\frac{c}{\lambda_m E\left[|\xi_{mt}|\right]} & (14)\end{aligned}$$

where the second-order conditions for a maximum are satisfied for both expressions. To obtain meaningful results, we assume that stale quote risk for the market maker is significant enough to ensure that ψ_m^* is strictly positive:

$$\lambda_m E\left[|\xi_{mt}|\right]\sum_i\frac{\beta_{im}}{\eta_i}>c.$$

ψ_i^* and ψ_m^* have unique positive solutions under this assumption. Moreover, attentiveness coefficients strictly increase in (respective) expected jump losses, and both coefficients strictly decrease with attention costs.

Spreads are not uniquely determined because each market maker takes spreads as given and spreads do not enter into market makers' attention choice. We do not impose additional structure to pin down spreads and test associated hypotheses because cross-sectional patterns in spreads are well-known (e.g., Hendershott et al. (2011)). Our model nonetheless produces intuitive implications for a competitive market making sector in a single-asset setting: slower trader arrivals or greater jump risks translate into higher equilibrium spreads.[14]

B. Empirical Analysis of Model Predictions

This section discusses and tests three sets of predictions that derive directly from our model. First, we show that order imbalance strongly (weakly) predicts large idiosyncratic price movements in small (large) stocks, but non-idiosyncratic price jumps are anticipated by liquidity withdrawals by market makers in all securities. Second, we show that market makers incorporate information into prices more rapidly in recent periods and for larger stocks. Third, we show that market maker monitoring varies across the spectrum of jump types. The empirical results for each prediction are consistent with optimal allocation of attention, but they are difficult to rationalize with alternative interpretations of our measure of market maker monitoring.

B.1. Jump Events and Market Maker Monitoring

Our model predicts that both ψ_i^* and ψ_m^* (i.e., the attentiveness coefficients) strictly increase in expected jump losses. This implies that market makers will pay more attention

[14]To establish this claim, we set the number of assets to one and impose the zero-profit condition to determine h:

$$(1+k)h_i = \frac{c}{\eta_i}\left(\psi_i + \psi_m\right) + \frac{\beta_{im}\lambda_m E\left[|\xi_{mt}|\right]}{\eta_i + \psi_m} + \frac{\lambda_i E\left[|\xi_{it}|\right]}{\eta_i + \psi_i}.$$

Because optimal attention increases with the square root of η, η is everywhere in the denominator. Likewise, optimal attention increases with the square root of anticipated jump costs, and taking ratios, we find that these costs translate into higher spreads in a square-root fashion.

to market jumps than to any individual stock's jumps if arrival rates and expected jump sizes are comparable for market and idiosyncratic jumps. This occurs because market jumps affect all assets and hence generate a proliferation of picking-off opportunities if not aggressively monitored. Therefore, to the extent that order book imbalance is an effective proxy for market maker monitoring, then we should see differences in order book imbalance for market versus idiosyncratic jumps.

We identify jumps using the Lee and Mykland (2008) jump detection procedure. Lee and Mykland (2008) disentangles continuous changes in the stock price from jump variation by calculating the ratio of the return at each price observation to a measure of instantaneous volatility over a preceding period. This jump detection test captures all significant and rapid movements in stock prices, not just those that reflect permanent changes in price. Therefore, we screen each event ex post to ensure that there is a permanent price jump and not simply a transitory liquidity demand that a market maker might wait out. We do this because picking off is based on the assumption that information was not incorporated into prices by the marker maker and would not be attenuated by a short-term buy-and-hold strategy. This would only be the case if there is a permanent movement in price. We categorize a jump as positive (negative) if

$$\frac{\max[P_{t-12}, ..., P_{t-1}]}{\min[P_{t+1}, ..., P_{t+12}]} \overset{(<)}{>} 1 \tag{15}$$

We limit our sample of jumps to those occurring between 11:00am and 3:00pm. We eliminate data before 10:00am because prior literature has shown that there is significant activity in the order book during early morning trading, and we do not want our results to be contaminated by the differential trading patterns during this period. We are limited to jumps occurring between 11:00am and 3:00pm because we require 1 hour of order book data both before and after the jump occurs so that we can examine whether there are changes in the order book before and after the jump. We categorize jumps into two groups: non-idiosyncratic and idiosyncratic. Non-idiosyncratic jumps occur when either our market proxy (the SPDR S&P 500 ETF) jumps or when two firms' prices co-jump in the same interval. Idiosyncratic jumps occur when a firm's price jumps in isolation.

Table I provides an overview of the 10,843 jumps we identify using the Lee and Mykland (2008) methodology at the 1% significance level. The jumps are divided between positive and negative and between idiosyncratic and non-idiosyncratic.[15] The average number of non-idiosyncratic jumps by stock-year ranges from a high of 1.92 for negative jumps in 2002 to a low of 1.12 for negative jumps in 2006. The number of idiosyncratic jumps declines over the sample period for both positive and negative jumps. The average number of idiosyncratic jumps is slightly more than 2 per stock-year during 2002-2003, and approximately 1.6 during the remainder of our sample period. The average detected jump is associated with a price change of approximately 2% as measured over the two-hour window centered on the jump event.

Figures II and III illustrate patterns in spreads and order book imbalance over a two-hour window centered on price jumps. These figures separate jumps on the dimensions of jump sign (positive or negative) and jump breadth (non-idiosyncratic or idiosyncratic). We find that order book imbalances change prior to both idiosyncratic and non-idiosyncratic jumps. The trend in order book imbalance in the direction of the price jump begins approximately 15 minutes prior to the jump. This pattern is relatively consistent across each of the four quadrants. These plots thus suggest that information in the order book may be useful for predicting large price movements regardless of the jump's sign or scope.

By contrast, we only find support for changes in the traditional bid-ask spreads for non-idiosyncratic jumps. There is no movement in these spreads prior to the five-minute period immediately preceding idiosyncratic jumps. This difference in the spreads suggests that that the marker maker does not adjust spreads to reflect the information in limit-order book. These figures suggest that market makers incorporate non-idiosyncratic information from the order book, but do not consistently incorporate firm-specific infor-

[15]To simplify visual presentation, we include market jumps in the non-idiosyncratic jump category.

mation. These features are predicted by our model of constrained market maker attention.

The model also provides insights into the predictability horizon, which is the expected time after a underlying price jump until a new trade or quote revision occurs. Under the continuous-time interpretation of the "race" between market makers and arbitrageurs to update the order book through quotes or trades, the predictability horizons for idiosyncratic and non-idiosyncratic jumps are as follows:

$$H_{ii} = \frac{1}{\eta_i + \psi_i^*} = \frac{1}{\max\left(\sqrt{\frac{1}{c}\eta_i \lambda_i E\left[|\xi_{it}|\right]}, \eta_i\right)}, \tag{16}$$

$$H_{im} = \frac{1}{\eta_i + \psi_m^*}. \tag{17}$$

The first-order conditions for ψ_i^* and ψ_m^* indicate that the predictability horizon is shorter for non-idiosyncratic jumps than for idiosyncratic jumps under the intuitive condition that idiosyncratic jumps are not extremely large or common relative to market jumps.[16] Equivalently, the probability of liquidity withdrawal rather than order imbalance is much higher for non-idiosyncratic jumps. These equations also imply that the predictability horizon is shorter for high volume stocks and stocks with large jump sizes or fast information arrival rates.

B.2. Empirical Implications of Market Maker Monitoring

Our model has several empirical predictions related to monitoring in the cross section and time series. Figures III and V illustrate the first prediction, namely, market makers should better anticipate non-idiosyncratic jumps than idiosyncratic jumps in setting spreads, and market makers should better control order book imbalances for non-idiosyncratic jumps than idiosyncratic jumps. Consistent with the model, Figure III depicts elevated spreads preceding both positive and negative jump events for non-idiosyncratic news, but *no* detectable response for firm-specific news. Reassuringly, the spread jumps contemporaneously with large price changes for both event types. Like-

[16] We verify this feature in the order book data for individual stock and S&P 500 SPDR (SPY) jumps.

wise, Figure V shows that the predictability horizon differs for non-idiosyncratic and idiosyncratic jumps. Pre-jump order book imbalances are significantly detectable for at least *five minutes* longer for idiosyncratic jumps than for more widespread jumps. The difference in signal persistence suggests that market makers take longer to "check" order flow changes associated with more idiosyncratic events.

Second, because monitoring technology improves over time, market makers should incorporate information into prices more rapidly in more recent periods. In our model, improvements in monitoring technology reduce c. A lower value for c implies market makers are more attentive for all securities. We investigate how the pattern in order book imbalance has changed over time by looking at the information in Figures II and III by year. These graphs are provided in Figures IV (spreads) and V (order book imbalance). In each case, only three years of data are provided: 2002, 2004 and 2007. This was done entirely for presentation purposes. The pattern is preserved when we include all years within the figure.

Figure IV indicates that the reaction of spreads to non-idiosyncratic jumps is relatively consistent over the sample period. If anything, it appears the spreads are adjusting somewhat earlier in the later years, although this different is visually modest. The differences in the reaction of spreads to idiosyncratic jumps by year is more noticeable. There is a much stronger reaction in 2002 relative to 2004, and in 2004 relative to 2007. However, in each case, the reaction begins approximately five minutes prior to the jump event.

Figure III indicates that the movement in order book imbalance diminishes sequentially over our sample period. In 2002, the upward trend in the order book imbalance begins earlier than any other year, the change in the period immediately surrounding the jump is the highest, and the reversion after the jump is the weakest. These charts illustrate that the strength of the relation between order book imbalance and jump predictability declines over our sample period. By 2007, there is virtually no change in order book imbalance before or after a jump. This suggests that information in the limit-order book is incorporated into prices by the 2007.

Third, the model predicts that market makers pay more attention to high volume securities, all else equal—high volume securities have larger η_i and are thus more costly to not monitor attentively. This relation implies that permanent price changes for low-volume stocks are associated with larger average order book imbalances.[17] Empirically, we replace volume with market capitalization because volume is highly time-varying whereas market capitalization is persistent relative to the horizon of our tests. The panel monthly correlation between these metrics (in logs) is 74%, so little information is lost in the substitution.

We empirically validate this prediction by splitting the portfolio analysis from Table II between large and small firms. The results are provided in Table X. For exposition, we only show the returns associated with 5-minute look-ahead portfolios, but results are similar for all horizons considered. Because the market maker is more likely to monitor large stocks (i.e., those with significantly more dollars in trading volume), we expect that the average returns will be significantly less for large versus small firms. The results are consistent with this expectation. Moreover, Table X indicates that the overall reduction in the Sharpe Ratio is primarily due to small firms. The large firm portfolios generate modest average returns starting in 2003.

B.3. Market Maker Monitoring and the Type of Jump Event

Lastly, our model predicts that market maker monitoring will vary across the spectrum of jump types. We expect that market makers monitor the least for idiosyncratic price jumps, more for multi-asset jumps, and the most for non-idiosyncratic jumps affecting many assets. Formally, the coefficient on the number of co-jumping securities should be negative in a regression of order imbalance on jump indicators.

We test this prediction with separate regressions of order imbalance on the four types of jump return variables. The first variable, jump, captures idiosyncratic jumps. This

[17]The model also predicts that market makers pay more attention to securities with high average jump size or fast information arrival, but we do not test this implication.

variable is a binary indicator that takes the value of 1 when there is a jump return measured at the firm-day level. The remaining three variables all capture broader measures of market-level jump activity. Interval jumps counts the number of co-jumps of S&P 1500 stocks during the same 5-minute interval; day jumps counts the number of co-jumps during the same day; and market jump equals 1 if the SPY market proxy jumps during the same day. The first specification uses firm-level data, and the remaining three specifications use daily-level data.

The results in Table XI show that there is a positive association between order book imbalance and jumps, and an incremental negative association between order book imbalances and interval jumps, day jumps and market jumps. In each case, the coefficient has the predicted sign, and the statistical association is very strong. Our model predicts these results because it predicts that monitoring by market markers is more likely to capture non-idiosyncratic rather than idiosyncratic jump return events.

B.4. Discussion

Taken together, we find strong evidence for all predictions associated with our stylized model of optimal market maker allocation of attention. Market makers are most responsive to (and anticipatory of) the costliest jumps, both in the aggregate comparison of portfolio returns accruing to picking off in stocks of different sizes and in the panel for individual firms over time. Systematic jumps receive greater attention than non-systematic jumps, but importantly, this distinction holds even in the context of jumps affecting only a few assets against jumps affecting one: widely anticipated market news powerfully predicts market maker attention, but it does not drive our results. In addition, as technology improves, market makers update prices in response to ever smaller order imbalances, and the returns to a picking off portfolio fall to zero by the end of our sample, suggesting that the returns to improved monitoring have long since declined.

These cross-sectional results distinguish our model from competing explanations for the pronounced decline in bid-ask spreads since 2001. For example, improved market maker diversification facilitated by automation would be expected to decrease bid-ask spreads, but improved diversification would reduce the cost of idiosyncratic jumps. By contrast, we see market makers improving their pre-jump responsiveness to order book imbalances particularly strongly for single-stock events. Likewise, as we discuss further in the next section, improved market maker competition should translate into lower realized spreads, but realized spreads are roughly constant over the period. Instead, changes in effective spreads move roughly one-for-one with changes in the adverse selection component of spreads, and the picking-off returns associated with jumps decline especially strongly. In addition, to the best of our knowledge, no existing alternative explanation can rationalize the negative slope of pre-jump order book imbalance with respect to jump systematicness that our constrained-attention model successfully predicts.

C. Instrumental Variables Analyses

We conduct two instrumental variable analyses to isolate the mechanism driving the strong empirical link between adverse selection and IQR, our panel data measure of market maker attentiveness. In particular, we seek to guard against potential contamination arising from channels different from market maker monitoring. Our first approach extends the discussion of jumps as a key driver of picking-off risk suffered by inattentive market makers. Our second approach speaks primarily to concerns about simultaneous causality between adverse selection and our IQR measure.

C.1. Instrumental Variables Approach using Return Jumps

An imbalanced order book can reflect market maker de-risking, correlated liquidity demands, or informed order flow. To the extent that these latent variables oscillate throughout the day yet have persistent price impacts unforeseen by intermediaries, the coefficient on the order imbalance IQR may be biased upwards. Likewise, an alternative reading of IQR as jointly reflecting liquidity supply and demand complicates one-way causality from our lag-IV specification because lagged adverse selection may contribute both to lagged IQR and to current adverse selection.

Our first IV approach addresses the concern that variation in adverse selection mechanically drives the IQR measure for reasons other than market maker attentiveness. Based on the strength of the association between jumps, market maker monitoring, and adverse selection described in Section B, we repeat our analyses from column (4) of Table IV and Table VII using a jump indicator as an instrument for order imbalance. Intuitively, this IV isolates the component of lagged IQR only associated with exposure to jump risk. Focusing on jumps has the virtue of refining away potential contamination of intraday IQR, deriving for example, from specialist inventory management (as suggested by Comerton-Forde et al. (2010) and others). This specification thus provides additional assurance that our interpretation of the empirical results is consistent with picking-off risk and market maker attentiveness rather than a confounding variable.

The results in Tables XII and XIII confirm our earlier conclusions. The coefficient on order imbalance is positive and highly significant, indicating that improvements in monitoring by market markers are associated with a decline in spreads. In addition, the coefficient on AT is positive and highly significant, indicating that higher levels of AT actually increase spreads.

Focusing on the variation associated with jumps appears to markedly improve the estimated link between market maker monitoring and adverse selection, as the coefficients are significantly larger in the IV specification.[18] The strengthened economic relationship supports our conjecture that market maker attentiveness is particularly important for managing risk and mitigating picking-off costs around likely jump events. Unfortunately, we are at risk of identifying the effect of monitoring on adverse selection associated with jumps rather than in the typical non-jump state; the effect of IQR on the "treated" stock-dates (with jumps) may well differ from the average effect of IQR on across all dates. For this reason, we supplement our first IV approach with a second approach that better captures the typical effect of market-maker monitoring on adverse selection.

[18] The difference in coefficients is not caused by a weak instruments problem. The first-stage Cragg-Donald F statistic exceeds 1,200.

C.2. Instrumental Variables Approach using Lagged IQR

The second IV approach addresses the concern that IQR might be simultaneously caused by adverse selection or jointly driven by a contemporaneous causal relation. Order book imbalances are associated with near-term price changes, and the imbalance is a component of the mechanism by which prices adjust. Our estimates are contaminated if the *range* of order book imbalance is affected as part of the adjustment process.

To address this and similar concerns, we instrument for the IQR using its one-week lagged value. We justify the use of the lagged IQR measure on the following grounds. The relevance condition is satisfied because computing technology that enables algorithmic market makers to monitor stock i at date t minus one week is very strongly correlated with the technology available for monitoring in stock i at date t. Indeed, the first stage F-statistics are extremely large (e.g., $> 6,600$ in the first specification). Rather than laying out a general argument for the exclusion restriction, we argue exclusion against the alternative that lagged IQR might be simultaneously causal with adverse selection. Because algorithmic traders are highly skilled in withdrawing from markets in inopportune volatility environments, perceived future adverse selection should be entirely avoidable at the frequency of a week ahead. Thus, forward expectations should not drive lagged IQR if market makers drive the IQR measure.

Table XIV reports results from the instrumental variables regression of the adverse selection component of spreads on $\widehat{IQR_{it}}$, where $\widehat{IQR_{it}}$ is estimated using $IQR_{i,t-1}$ (and other covariates). The first specification uses no contemporaneous information in a within specification for the effect of market maker attentiveness on adverse selection. We achieve similar results as in Table IV in both economic and statistical significance. A one unit change in the lagged IQR is associated with a 5.78 basis point increase in average adverse selection. In isolation, AT instrumented using lagged AT also survives, as before, with comparable economic and statistical significance. Specifications (3) and (4) challenge our measure by running a horse race of our instrumented variable against *contemporaneous*

AT: the lagged IQR measure survives, but AT again switches sign to be associated with decreased market quality. Specification (5) implements a “fair” horse race using the following three-equation approach in which both potentially endogenous variables are instrumented using lagged values:

$$\begin{aligned}
\widehat{IQR_{it}} &= \gamma_0 + \gamma_1 IQR_{i,t-1} + \gamma_3 Controls_{it} + \epsilon_{it} \\
\widehat{AT_{it}} &= \delta_0 + \delta_2 AT_{i,t-1} + \gamma_3 Controls_{it} + \epsilon_{it} \\
Spread_{it} &= \beta_0 + \beta_1 IQR_{it} + \beta_2 AT_{it} + \beta_3 Controls_{it} + \epsilon_{it}.
\end{aligned} \tag{18}$$

Accounting for two potentially endogenous variables reduces coefficient magnitudes slightly, but otherwise has minimal effect on our results.

VI. Conclusion

Recent work has challenged the nascent empirical consensus that algorithmic trading robustly reduces bid-ask spreads. We resolve the tension in the algorithmic trading and market quality debate by deconstructing AT to identify the dominant mechanism by which technological improvements have improved market function. Using both time series and panel tests, we identify market maker monitoring as the key channel through which technological advances have affected the historic decline in bid-ask spreads throughout the 2000s. We buttress our initial regressions by generating and testing unique predictions of an intuitive model of attention-constrained market makers. We find that our novel implications about the market maker responses to jump risk are borne out across the spectrum of large and small firms, idiosyncratic and systematic jumps, and over time. In so doing, we discover why different sample periods deliver conflicting results for AT’s contribution to market quality, namely because the gains to improved monitoring are exhausted as of the late 2000s.

Our findings have direct policy relevance for interpreting the role of algorithmic trading in capital markets. Our results suggest that enhanced market maker monitoring, and not simply increases in AT activity, have led to improvements in market quality on the order of tens of billions of dollars a year.[19] This implies that policies facilitating AT liquidity provision rather than liquidity taking are more likely to generate improvements in market quality, at least up to a point. Once market makers are sufficiently attentive, quoting efficiency and spreads do not further improve. We thus explain why spreads[20] have stagnated for nearly a decade and call into question whether continued investment in trading technology will produce tangible benefits for financial markets.

[19] For a back-of-the-envelope calculation, consider that the IQR measure explains 70% of within-firm variation in adverse selection costs over the sample period. Bid-ask spreads have declined on the order of 10 basis points for the median stock from 2001–2007, and daily equity volume is on the order of hundreds of billions. As a conservative estimate for 2007, we obtain $1/2 \times 70\% \times 0.1\% \times \$160B \times 252 \approx \$14.1B$.

[20] With or without aggregate volatility adjustment.

References

Biais, B., F. Declerck, and S. Moinas (2015, June). Who supplies liquidity, how and when? Working paper.

Biais, B. and T. Foucault (2014). Hft and market quality. *Bankers, Markets & Investors* (128), 5–19.

Biais, B., T. Foucault, and S. Moinas (2015). Equilibrium fast trading. *Journal of Financial Economics 116* (2), 292 – 313.

Brogaard, J., T. Hendershott, and R. Riordan (2014). High frequency trading and the 2008 short sale ban. Working paper.

Chakrabarty, B. and P. C. Moulton (2012). Earnings announcements and attention constraints: The role of market design. *Journal of Accounting and Economics 53* (3), 612–634.

Chung, K. H. and C. Chuwonganant (2014). Uncertainty, market structure, and liquidity. *Journal of Financial Economics 113* (3), 476–499.

Comerton-Forde, C., T. Hendershott, C. M. Jones, P. C. Moulton, and M. S. Seasholes (2010). Time variation in liquidity: The role of market-maker inventories and revenues. *The Journal of Finance 65* (1), 295–331.

Easley, D. and M. O'hara (1987). Price, trade size, and information in securities markets. *Journal of Financial economics 19* (1), 69–90.

Foucault, T., J. Hombert, and I. Rosu (2015). News trading and speed. Working paper.

Foucault, T., R. Kozhan, and W. W. Tham (2014). Toxic arbitrage. *HEC Paris Research Paper No. FIN-2014-1040*.

Foucault, T. and A. J. Menkveld (2008). Competition for order flow and smart order routing systems. *The Journal of Finance 63* (1), 119–158.

Foucault, T., A. Röell, and P. Sandås (2003). Market making with costly monitoring: An analysis of the soes controversy. *Review of Financial Studies 16*(2), 345–384.

Glosten, L. R. (1987). Components of the bid-ask spread and the statistical properties of transaction prices. *The Journal of Finance 42*(5), 1293–1307.

Hagströmer, B. and L. Norden (2013). The diversity of high frequency traders. *Journal of Financial Markets 16*(4), 741–770.

Han, J., M. Khapko, and A. S. Kyle (2014). Liquidity with high-frequency market making. *Swedish House of Finance Research Paper* (14-06).

Hasbrouck, J. and G. Saar (2013). Low-latency trading. *Journal of Financial Markets 16*(4), 646 – 679. High-Frequency Trading.

Hendershott, T., C. M. Jones, and A. J. Menkveld (2011, February). Does algorithmic trading improve liquidity? *Journal of Finance 66*(1), 1–33.

Holden, C. W. and S. Jacobsen (2014). Liquidity measurement problems in fast, competitive markets: expensive and cheap solutions. *The Journal of Finance 69*(4), 1747–1785.

Jones, C. M. (2013). What do we know about high-frequency trading? Working paper.

Jovanovic, B. and A. J. Menkveld (2011). Middlemen in limit-order markets. *Western Finance Association (WFA).*

Lee, S. S. and P. A. Mykland (2008). Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies 21*(6), 2535–2563.

Menkveld, A. J. and M. A. Zoican (2015). Need for speed? exchange latency and liquidity. Working paper.

Pagnotta, E. and T. Philippon (2015). Competing on speed. Working paper.

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *The Journal of Business 53*(1), pp. 61–65.

Van Nieuwerburgh, S. and L. Veldkamp (2010). Information acquisition and under-diversification. *The Review of Economic Studies 77*(2), 779–805.

A. Figures

Figure I: Spreads and Marker Maker Monitoring over Time

This figure presents rolling 21 day averages of effective spreads and the adverse selection component of spreads (following Hendershott et al. (2011)), our proxy for market maker monitoring (the mean signal return measured using a 5 minute holding period, developed in Section IV.A) and an aggregate measure of algorithmic trading (the ratio of order cancellations to total volume, developed in Section IV.B.1). The cancellation-to-trade ratio is plotted in inverse form to facilitate comparison.

(a) Effective Spreads and Adverse Selection Component of Spreads Over Time

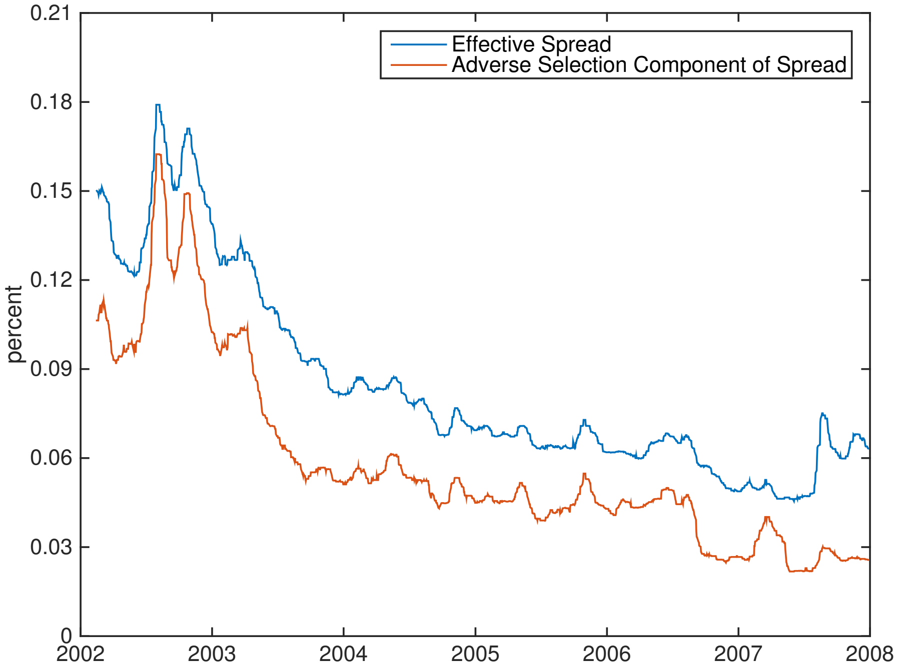

(b) Mean Spreads and Strategy Returns over Time

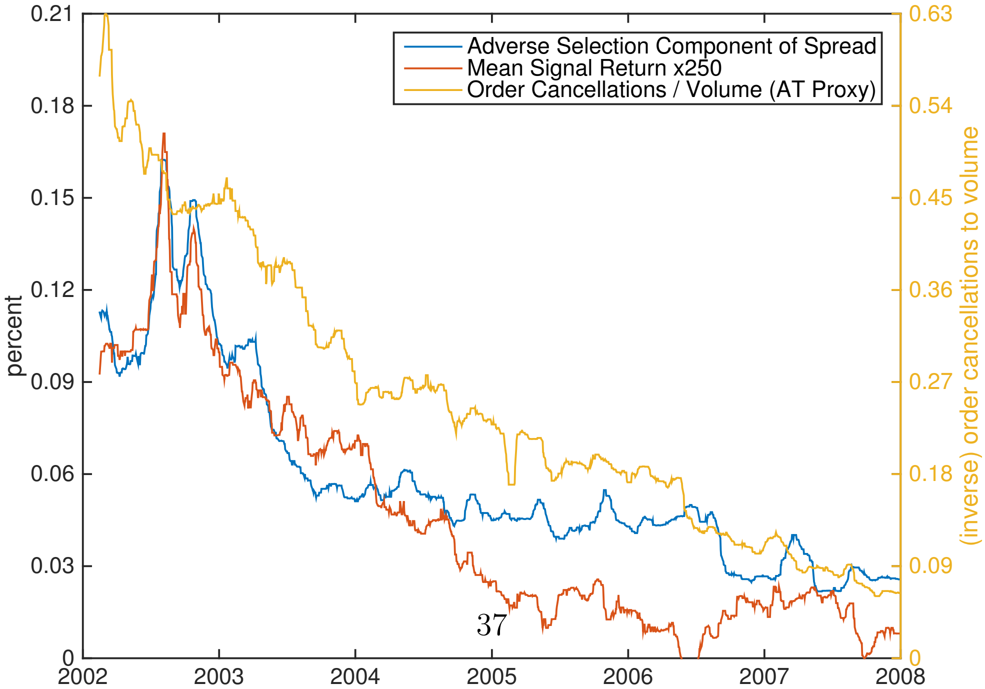

Figure II: Bid-ask Spreads around Price jumps

This plot presents average bid-ask spreads relative to a one-hour pre-jump reference bid-ask spread, $log(S_t/S_{T-60})$, to take into account stock and time specific effects. Each plot is composed of cross-sectional averages of 5-minute average spreads across stocks in event time. The top row presents spreads associated with co-jumps among multiple stocks, and the bottom row presents spreads associated with idiosyncratic jumps. Left and right columns correspond with negative and positive jumps.

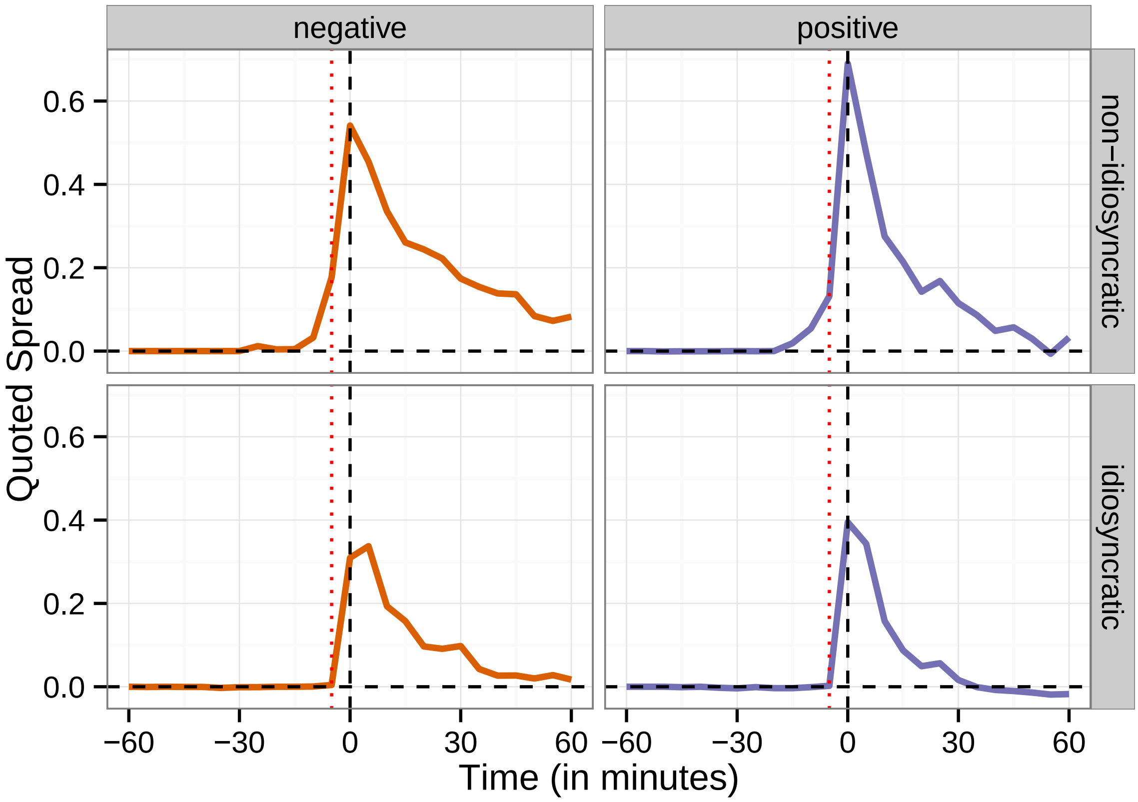

Figure III: Order Imbalances around Price Jumps

This plot presents average order book imbalances relative to a one-hour pre-jump reference imbalance, $log(S_t/S_{T-60})$, to take into account stock and time specific effects. Each plot is composed of cross-sectional averages of 5-minute average imbalances across stocks in event time. The top row presents imbalances associated with co-jumps among multiple stocks, and the bottom row presents imbalances associated with idiosyncratic jumps. Left and right columns correspond with negative and positive jumps.

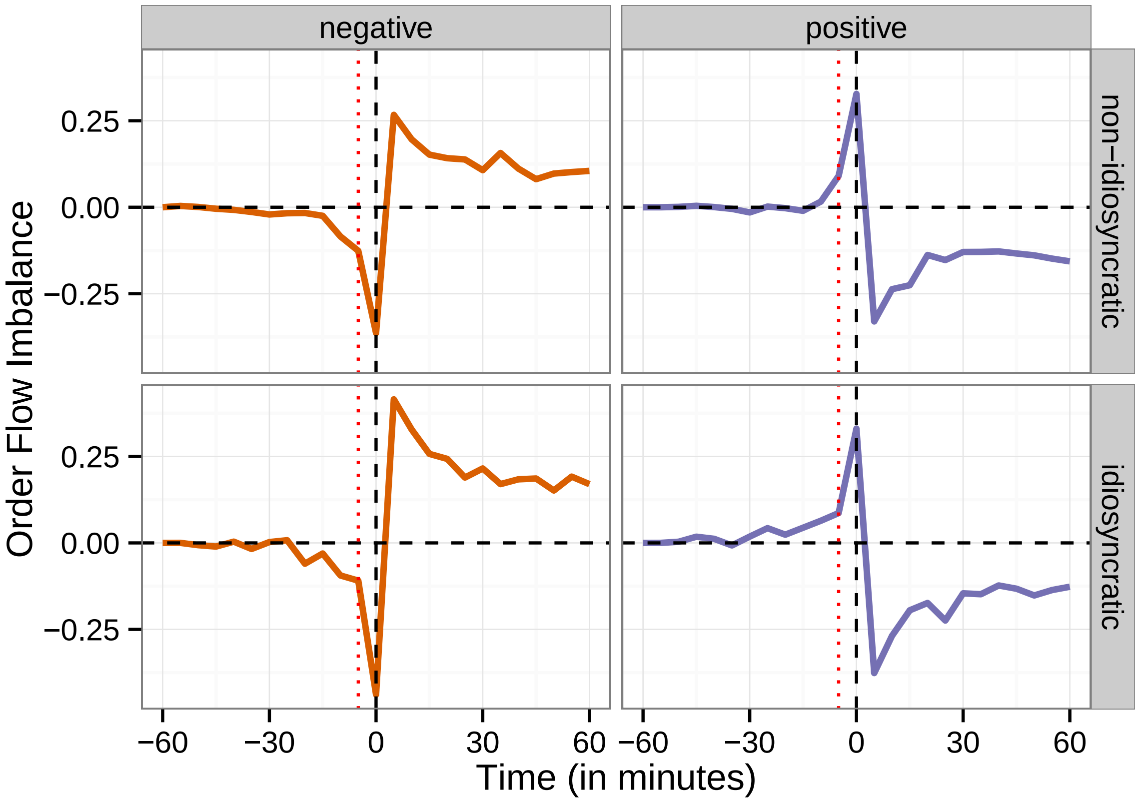

Figure IV: Bid-ask Spreads around Price Jumps by Year

This plot presents average bid-ask spreads relative to a one-hour pre-jump reference bid-ask spread, $log(S_t/S_{T-60})$, to take into account stock and time specific effects. Each plot is composed of cross-sectional averages of 5-minute average spreads across stocks in event time by calendar year. The top row presents spreads associated with co-jumps among multiple stocks, and the bottom row presents spreads associated with idiosyncratic jumps. Left and right columns correspond with negative and positive jumps.

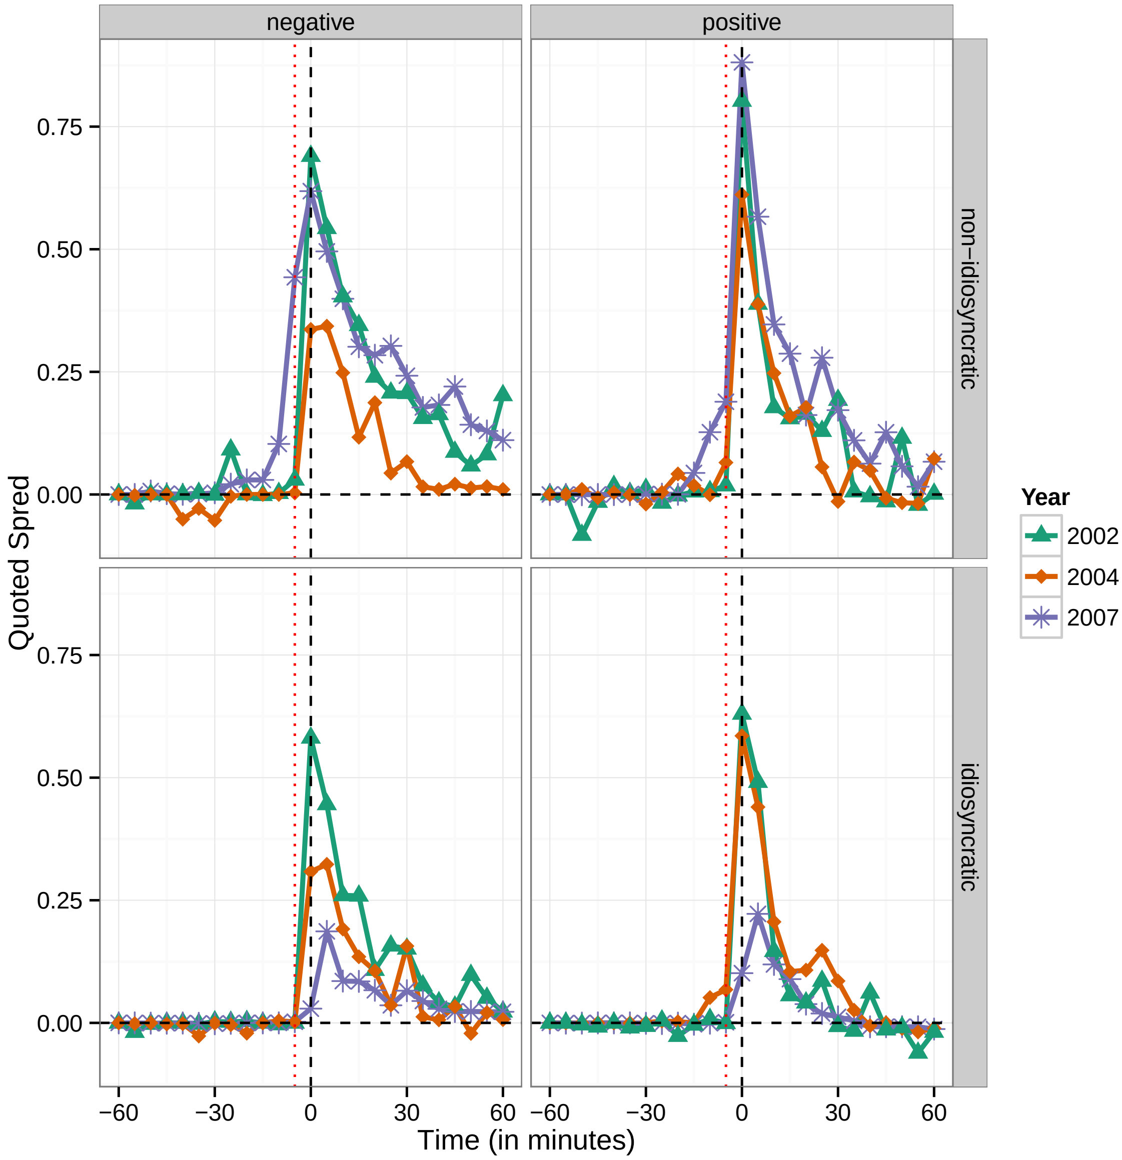

Figure V: Order Imbalances around Price Jumps by Year

This plot presents average order book imbalances relative to a one-hour pre-jump reference imbalance, $log(S_t/S_{T-60})$, to take into account stock and time specific effects. Each plot is composed of cross-sectional averages of 5-minute average imbalances across stocks in event time by calendar year. The top row presents imbalances associated with co-jumps among multiple stocks, and the bottom row presents imbalances associated with idiosyncratic jumps. Right and left columns correspond with negative and positive jumps.

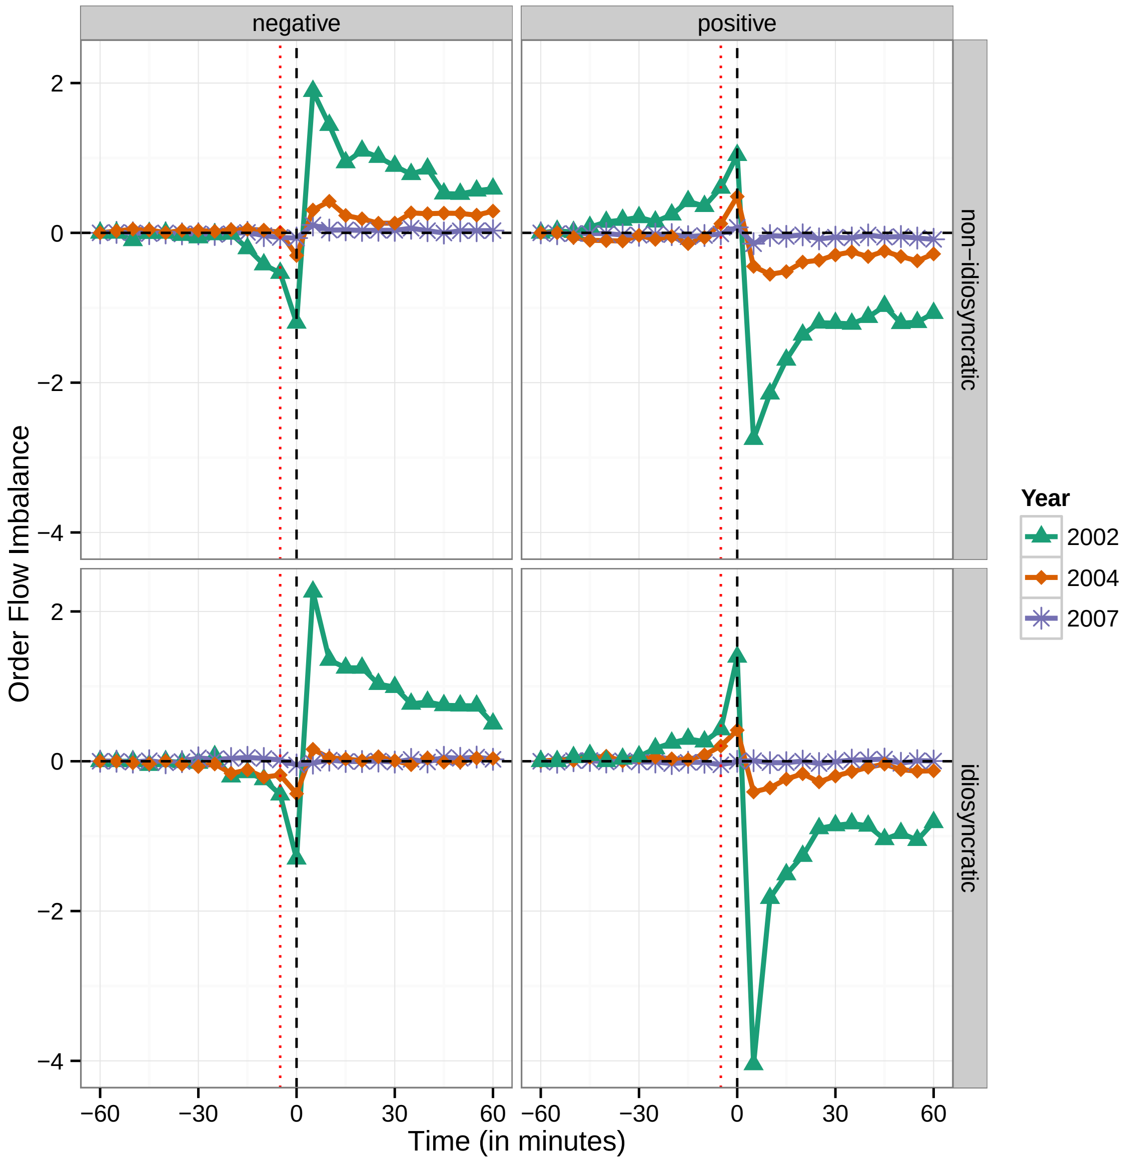

B. Tables

Table I: Descriptive Statistics

This table reports summary information for the variables used in our panel specifications. Effective spreads and Adverse Selection are both calculated using the same approach as Hendershott et al. (2011) IQR is the interquartile range of the log order book imbalance, where the order book imbalance is the ratio of buy to sell orders within 1 percent of the reference price. AT is ratio of order cancellations to total volume. Volatility is measured using the log intraday price range (Parkinson (1980)). 1/Price is based on the closing price. Size is the log market capitalization. VIX is the CBOE VIX. PIN is calculated following Easley and O'hara (1987).

	P5	P25	P50	P75	P95	Mean	Std. Dev.
Effective Spread	0.02298	0.04128	0.06691	0.11903	0.37926	0.11806	2.58834
Adverse Selection	0.00680	0.02019	0.03757	0.07171	0.19745	0.06236	0.36770
IQR	0.17572	0.25999	0.36103	0.65843	1.87950	0.59589	0.57594
AT	0.71028	2.11075	2.65992	3.13366	3.77357	2.55246	0.86134
Share Turnover	0.89659	1.82150	2.95391	4.88736	11.41439	4.56189	9.08337
Volatility	0.06690	0.09080	0.10904	0.13239	0.18345	0.11305	0.03981
1/Price	0.01288	0.02306	0.03648	0.06333	0.18727	0.06046	0.09160
Size	12.3123	13.7497	14.7302	15.8448	17.5119	14.8114	1.5778
VIX	16.6200	20.7100	24.4000	30.8100	49.6800	27.6541	11.0192
PIN	0.01500	0.06200	0.08700	0.11700	0.18700	0.09422	0.05550

Table II: Average Return to Portfolio Formed using Order Book Imbalance

This table reports the mean returns for portfolios sorted by quintile of order book imbalance, where the order book imbalance is the ratio of buy to sell orders within 1 percent of the reference price. Portfolios are formed every 5 minutes during the period beginning at 10:00am and ending at 3:00pm. All returns are in basis points. The average returns are calculated as portfolio 5 minus portfolio 1 returns, and all average returns are significant at the 1 percent level. The Sharpe Ratio is equal to the average T-minute portfolio return divided by the standard deviation of the portfolio average returns.

(a) Panel A: Mean Return in Basis Points using 5-Minute Window

Year	1	2	3	4	5	5-1 Return	Sharpe Ratio
2002	-2.4	-0.94	0.05	1.07	2.28	4.69	0.98
2003	-1.53	-0.49	0.13	0.69	1.61	3.14	1.02
2004	-0.89	-0.1	0.14	0.37	0.85	1.74	0.63
2005	-0.43	-0.04	0.03	0.06	0.31	0.74	0.3
2006	-0.23	0.03	0.06	0.06	0.16	0.4	0.15
2007	-0.41	-0.16	-0.08	-0.02	0.2	0.61	0.23

(b) Panel A: Mean Return in Basis Points using 15-Minute Window

Year	1	2	3	4	5	5-1 Return	Sharpe Ratio
2002	-3.59	-1.28	0.13	1.55	3.26	6.85	0.85
2003	-2	-0.42	0.37	1.1	2.37	4.37	0.88
2004	-0.81	0.22	0.31	0.4	0.86	1.67	0.37
2005	-0.38	0.04	0.01	-0.07	0.09	0.47	0.11
2006	-0.16	0.12	0.1	0.06	0.11	0.27	0.06
2007	-0.69	-0.33	-0.22	-0.15	0.19	0.88	0.19

Table II: Average Return to Portfolio Formed using Order Book Imbalance (Continued)

(c) Panel C: Mean Return in Basis Points using 30-Minute Window

Year	1	2	3	4	5	5-1 Return	Sharpe Ratio
2002	-4.12	-1.3	0.19	1.7	3.64	7.76	0.72
2003	-1.76	0.03	0.8	1.48	2.88	4.64	0.7
2004	-0.41	0.57	0.57	0.49	0.78	1.19	0.2
2005	-0.29	0.09	-0.03	-0.17	-0.13	0.16	0.03
2006	-0.15	0.12	0.14	0.12	0.11	0.26	0.04
2007	-0.89	-0.48	-0.32	-0.22	0.2	1.09	0.17

(d) Panel D: Mean Return in Basis Points using 60-Minute Window

Year	1	2	3	4	5	5-1 Return	Sharpe Ratio
2002	-3.59	-1.28	0.13	1.55	3.26	6.85	0.85
2003	-2	-0.42	0.37	1.1	2.37	4.37	0.88
2004	-0.81	0.22	0.31	0.4	0.86	1.67	0.37
2005	-0.38	0.04	0.01	-0.07	0.09	0.47	0.11
2006	-0.16	0.12	0.1	0.06	0.11	0.27	0.06
2007	-0.69	-0.33	-0.22	-0.15	0.19	0.88	0.19

Table III: Impact of Picking-off Risk on Effective Spread

This table reports the results of an OLS regression of average daily effective spreads on the daily mean signal returns for the period beginning on February 1, 2002 and ending on December 31, 2007. The mean signal return is measured using a 5 minute holding period. VIX is the CBOE volatility index and AT is the ratio of order cancellations to total volume. All variables are aggregated to the daily level. t-statistics using Newey-West standard errors with 21 lags are reported in parentheses. Levels of significance using a two-tailed test are indicated by *, **, and ***, representing 10 percent, 5 percent, and 1 percent, respectively.

(a) Panel A: Levels Specification

	(1) Effective Spread	(2) Effective Spread	(3) Effective Spread
Mean signal return	0.424***		0.282***
	(13.51)		(8.21)
VIX	0.231***	0.261***	0.234***
	(15.72)	(12.39)	(14.73)
AT		-2.631***	-1.134***
		(-11.13)	(-5.14)
Constant	1.637***	8.025***	4.353***
	(6.53)	(15.01)	(7.60)
# Observations	1,343	1,343	1,343
R^2	0.870	0.850	0.883

(b) Panel B: Changes Specification

	(1) ΔEffective Spread	(2) ΔEffective Spread	(3) ΔEffective Spread
ΔMean signal return	0.112***		0.118***
	(4.82)		(4.80)
ΔVIX	0.031	0.041**	0.032
	(1.39)	(1.98)	(1.37)
ΔAT		0.206	0.486
		(0.64)	(1.46)
Constant	-0.005	-0.006	-0.006
	(-1.06)	(-1.17)	(-1.16)
# Observations	1,342	1,342	1,342
R^2	0.063	0.008	0.068

Table IV: Firm-Level Panel Regression of Picking-off Risk on Effective Spreads using Levels

This table reports the results of an OLS regression of daily effective spreads on market maker monitoring, which we proxy for using IQR, for the period beginning on February 1, 2002 and ending on December 31, 2007. IQR is the interquartile range of the log order book imbalance, where the order book imbalance is the ratio of buy to sell orders within 1 percent of the reference price. AT is ratio of order cancellations to total volume. The control variables are: share turnover, volatility (log intraday price range, a dispersion proxy based on Parkinson (1980)), 1/price (based on the closing price), size (log market capitalization), the CBOE VIX and PIN (Easley and O'hara (1987)). The constant is absorbed and suppressed due to the inclusion of firm fixed effects. ***, **, and * indicate statistical significance at the 1 percent, 5 percent, and 10 percent level, respectively, using two-tailed tests and standard errors clustered at the firm- and year-level.

	(1) Effective Spread	(2) Effective Spread	(3) Effective Spread	(4) Effective Spread
IQR	0.03839***		0.02237***	0.01319***
	(42.585)		(42.088)	(31.435)
AT		-0.01092***	-0.00011	0.01173***
		(-13.129)	(-0.172)	(16.423)
Share Turnover		-0.00026**	-0.00022**	0.00000
		(-2.380)	(-2.396)	(0.098)
Volatility		0.78664***	0.62330***	0.36157***
		(26.444)	(26.261)	(18.421)
1/Price		0.63078***	0.68253***	0.91755***
		(13.46)	(16.943)	(24.476)
Size		-0.03092***	-0.02276***	-0.00870***
		(-15.254)	(-12.980)	(-4.698)
VIX				0.00072***
				(12.089)
PIN				0.02322***
				(3.754)
# Observations	1,064,994	1,064,994	1,064,994	1,064,994
Firm Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	No	No	No	Yes
R^2	0.7002	0.7727	0.7987	0.8237

Table V: Firm-Level Panel Regression of Picking-off Risk on Effective Spreads using Weekly Changes

This table reports the results of an OLS regression of changes in weekly effective spreads on changes in market maker monitoring, which we proxy for using IQR, for the period beginning on February 1, 2002 and ending on December 31, 2007. IQR is the interquartile range of the log order book imbalance, where the order book imbalance is the ratio of buy to sell orders within 1 percent of the reference price. AT is ratio of order cancellations to total volume. The control variables are: share turnover, volatility (log intraday price range, a dispersion proxy based on Parkinson (1980)), 1/price (based on the closing price), size (log market capitalization), the CBOE VIX and PIN (Easley and O'hara (1987)). The constant is absorbed and suppressed due to the inclusion of firm fixed effects. ***, **, and * indicate statistical significance at the 1 percent, 5 percent, and 10 percent level, respectively, using two-tailed tests and standard errors clustered at the firm- and year-level.

	(1) ΔEffective Spread	(2) ΔEffective Spread	(3) ΔEffective Spread	(4) ΔEffective Spread
ΔIQR	0.00699***		0.00803***	0.00791***
	(10.197)		(16.215)	(17.513)
ΔAT		0.02035***	0.02184***	0.02125***
		(7.592)	(8.368)	(7.640)
ΔShare Turnover		0.00059***	0.00054***	0.00054***
		(2.926)	(2.762)	(2.774)
ΔVolatility		0.07194	0.06365	0.06087
		(1.104)	(1.035)	(1.264)
Δ1/Price		1.06428***	1.04334***	1.08232***
		(10.366)	(10.188)	(10.865)
ΔMarket Cap (Log)		-0.03175***	-0.03184***	-0.01506*
		(-2.880)	(-3.004)	(-1.886)
ΔVIX				0.00102**
				(2.416)
ΔPIN				-0.01471
				(-0.801)
# Observations	183,538	183,538	183,538	183,538
Firm Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	No	No	No	Yes
R^2	0.0210	0.0508	0.0564	0.0576

Table VI: Impact of Picking-off Risk on Quoting Efficiency

This table reports the results of an OLS regression of the daily average adverse selection component of spreads on the mean signal return for the period beginning on February 1, 2002 and ending on December 31, 2007. The mean signal return is measured using a 5 minute holding period. VIX is the CBOE volatility index and AT is the ratio of order cancellations to total volume. All variables are aggregated to the daily level. t-statistics using Newey-West standard errors with 21 lags are reported in parentheses. Levels of significance using a two-tailed test are indicated by *, **, and ***, representing 10 percent, 5 percent, and 1 percent, respectively.

(a) Panel A: Levels Specification

	(1) Adverse Selection	(2) Adverse Selection	(3) Adverse Selection
Mean signal return	0.429***		0.397***
	(22.19)		(13.27)
VIX	0.172***	0.210***	0.172***
	(16.43)	(10.06)	(16.16)
AT		-2.360***	-0.252
		(-11.66)	(-1.38)
Constant	-0.421***	5.356***	0.182
	(-2.13)	(10.55)	(0.37)
# Observations	1,343	1,343	1,343
R^2	0.908	0.821	0.908

(b) Panel B: Changes Specification

	(1) ΔAdverse Selection	(2) ΔAdverse Selection	(3) ΔAdverse Selection
ΔMean signal return	0.261***		0.251***
	(8.32)		(7.85)
ΔVIX	0.049***	0.068***	0.048***
	(2.59)	(3.93)	(2.65)
ΔAT		-1.439***	-0.840***
		(-4.72)	(-3.27)
Constant	-0.002	-0.002	-0.001
	(-0.49)	(-0.37)	(-0.27)
# Observations	1,342	1,342	1,342
R^2	0.236	0.050	0.248

Table VII: Firm-Level Panel Regression of Picking-off Risk on Adverse Selection Component of Spreads using Levels

This table reports the results of an OLS regression of the daily average adverse selection component of spreads on market maker monitoring, which we proxy for using IQR, for the period beginning on February 1, 2002 and ending on December 31, 2007. IQR is the interquartile range of the log order book imbalance, where the order book imbalance is the ratio of buy to sell orders within 1 percent of the reference price. AT is ratio of order cancellations to total volume. The control variables are: share turnover, volatility (log intraday price range, a dispersion proxy based on Parkinson (1980)), 1/price (based on the closing price), size (log market capitalization), the CBOE VIX and PIN (Easley and O'hara (1987)). The constant is absorbed and suppressed due to the inclusion of firm fixed effects. ***, **, and * indicate statistical significance at the 1 percent, 5 percent, and 10 percent level, respectively, using two-tailed tests and standard errors clustered at the firm- and year-level.

	(1) Adverse Selection	(2) Adverse Selection	(3) Adverse Selection	(4) Adverse Selection
IQR	0.03824***		0.02219***	0.01288***
	(38.303)		(34.659)	(34.330)
AT		-0.01752***	-0.00679***	0.00303***
		(-21.188)	(-10.961)	(4.704)
Share Turnover		-0.00019**	-0.00015**	0.00005**
		(-2.228)	(-2.216)	(2.525)
Volatility		0.86089***	0.69890***	0.34779***
		(20.745)	(20.733)	(15.421)
1/Price		-0.05922	-0.00795	0.21703***
		(-1.450)	(-0.231)	(6.622)
Size		-0.03165***	-0.02356***	-0.01227***
		(-18.280)	(-16.584)	(-8.693)
VIX				0.00121***
				(11.746)
PIN				0.00065
				(0.111)
# Observations	1,064,994	1,064,994	1,064,994	1,064,994
Firm Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	No	No	No	Yes
R^2	0.1436	0.1552	0.1665	0.1803

Table VIII: Firm-Level Panel Regression of Picking-off Risk on Adverse Selection Component of Spreads using Weekly Changes

This table reports the results of an OLS regression of weekly changes in the adverse selection component of spreads on changes in market maker monitoring, which we proxy for using IQR, for the period beginning on February 1, 2002 and ending on December 31, 2007. IQR is the interquartile range of the log order book imbalance, where the order book imbalance is the ratio of buy to sell orders within 1 percent of the reference price. AT is ratio of order cancellations to total volume. The control variables are: share turnover, volatility (log intraday price range, a dispersion proxy based on Parkinson (1980)), 1/price (based on the closing price), size (log market capitalization), the CBOE VIX and PIN (Easley and O'hara (1987)). The constant is absorbed and suppressed due to the inclusion of firm fixed effects. ***, **, and * indicate statistical significance at the 1 percent, 5 percent, and 10 percent level, respectively, using two-tailed tests and standard errors clustered at the firm- and year-level.

	(1) ΔAdverse Selection	(2) ΔAdverse Selection	(3) ΔAdverse Selection	(4) ΔAdverse Selection
ΔIQR	0.00746***		0.01095***	0.01053***
	(7.781)		(3.865)	(3.954)
ΔAT		-0.00556	-0.00357	-0.00539
		(-0.362)	(-0.240)	(-0.351)
ΔShare Turnover		-0.00159	-0.00167	-0.00167
		(-0.760)	(-0.787)	(-0.789)
ΔVolatility		0.52701	0.51467	0.44098
		(1.51)	(1.503)	(1.438)
Δ1/Price		0.02203	-0.01494	0.07117
		(0.063)	(-0.042)	(0.217)
ΔMarket Cap (Log)		0.00513	0.00475	0.04603
		(0.15)	(0.14)	(1.145)
ΔVIX				0.00227**
				(2.224)
ΔPIN				-0.05750
				(-1.160)
# Observations	183,538	183,538	183,538	183,538
Firm Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	No	No	No	Yes
R^2	0.0046	0.0054	0.0056	0.0057

Table IX: Summary Information on Detected Jump Events

This table reports the number of jumps by year, separated by direction (i.e., negative or positive) and type (i.e., idiosyncratic and non-idiosyncratic). We use the Lee and Mykland (2008) procedure to identify jumps.

(a) Panel A: Non-Idiosyncratic Jumps

Jump Direction	Year	Jump Firms	Avg. per Firm	Jump Return
negative	2002	414	1.92	-3.78
negative	2003	221	1.32	-2.22
negative	2004	238	1.33	-2.18
negative	2005	214	1.29	-2.01
negative	2006	121	1.12	-1.58
negative	2007	457	1.54	-2.78
positive	2002	433	1.58	2.26
positive	2003	214	1.26	2.01
positive	2004	163	1.19	1.76
positive	2005	198	1.17	1.63
positive	2006	173	1.32	2.09
positive	2007	533	1.42	2.64

(b) Panel B: Idiosyncratic Jumps

Jump Direction	Year	Jump Firms	Avg. per Firm	Jump Return
negative	2002	498	2.56	-3.7
negative	2003	405	2.04	-2.38
negative	2004	333	1.65	-2.22
negative	2005	287	1.7	-2.56
negative	2006	210	1.57	-2.7
negative	2007	209	1.67	-2.76
positive	2002	436	2.04	2.56
positive	2003	417	2	2.22
positive	2004	351	1.78	2.09
positive	2005	299	1.72	2.11
positive	2006	242	1.64	2.33
positive	2007	270	1.79	3.05

Table X: Average Returns Separated by Firm Size

This table reports the mean returns for portfolios based on level of order book imbalance separated by big and small firms based on market capitalization. All returns are in basis points and are calculated using a 5 minute window. The average returns are calculated as portfolio 5 minus portfolio 1 returns, and all average returns are significant at the 1 percent level. The Sharpe Ratio is equal to the average 5-minute portfolio return divided by its respective standard deviation.

Year	Firms	1	2	3	4	5	5-1 Return	Sharpe Ratio
2002	Big	-1.63	-0.61	0.03	0.64	1.66	3.28	0.50
2002	Small	-4.18	-1.99	-0.09	1.88	4.25	8.42	0.52
2003	Big	-0.72	-0.25	0.06	0.34	0.79	1.51	0.37
2003	Small	-2.69	-1.04	0.16	1.62	3.42	6.10	0.55
2004	Big	-0.32	-0.08	0.10	0.18	0.37	0.69	0.20
2004	Small	-1.74	-0.72	0.15	1.01	1.98	3.72	0.45
2005	Big	-0.20	-0.06	0.02	0.00	0.11	0.31	0.09
2005	Small	-1.16	-0.37	0.05	0.55	1.04	2.20	0.31
2006	Big	-0.12	-0.01	0.02	0.08	0.12	0.25	0.07
2006	Small	-0.69	-0.20	0.08	0.35	0.75	1.44	0.23
2007	Big	-0.29	-0.17	-0.07	0.04	0.19	0.49	0.12
2007	Small	-0.66	-0.34	-0.08	0.07	0.34	1.00	0.13

Table XI: Association between Order Book Imbalance and Jumps

This table reports the results of an OLS regression of average IQR on a set of return jump indicators. Jump equals 1 if a jump occurs in stock i on date t and 0 otherwise. Interval jumps count the number of co-jumps of S&P 1500 stocks during the same 5-minute interval. Day jump counts the number of co-jumps during the same day. Market jump equals 1 if the SPY market proxy jumps during the same day. ***, **, and * indicate statistical significance at the 1 percent, 5 percent, and 10 percent level, respectively, using two-tailed tests and standard errors clustered at the firm- and year-level.

	(1) Ord. Imb.	(2) Ord. Imb.	(3) Ord. Imb.	(4) Ord. Imb.	(5) Ord. Imb.	(6) Ord. Imb.
Jump	0.13899***				0.32955***	26.78686***
	(4.833)				(4.719)	(3.832)
Interval Jump		-0.06592***				
		(-7.273)				
Day Jump			-0.04652*			
			(-1.868)			
Market Jump				-0.32946***		
				(-10.911)		
Market Cap					-0.33611***	
					(-25.309)	
Jump*Market Cap					-0.01552**	
					(-2.566)	
Year						-0.12029***
						(-73.065)
Jump*Year						-0.01332***
						(-3.817)
# Observations	1065908	11033	11033	11033	1053683	1065908
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	No	No	No	No	No	Yes
R^2	0.0981	0.3315	0.2593	0.2958	0.2616	0.5315

Table XII: Instrumental Variables Regression of Spreads on Picking-off Risk using Levels

This table reports the results of a levels IV regression where we instrument order imbalance using a jump return indicator, which equals 1 if a jump occurs in stock i on date t and 0 otherwise. The order book imbalance is the ratio of buy to sell orders within 1 percent of the reference price. The AT proxy is ratio of order cancellations to total volume. The control variables are: share Turnover, log intraday price range (a dispersion proxy based on Parkinson (1980)), 1/price (based on the closing price), log market capitalization and the CBOE VIX. ***, **, and * indicate statistical significance at the 1 percent, 5 percent, and 10 percent level, respectively, using two-tailed tests and standard errors clustered at the firm- and year-level.

	(1) Effective Spread	(2) Adverse Selection
IQR	0.21568***	0.30400***
	(10.691)	(14.857)
AT	0.03545***	0.03776***
	(13.747)	(14.236)
Share Turnover	-0.00026**	-0.00034**
	(-2.257)	(-2.289)
Volatility	0.22938***	0.15581***
	(8.598)	(4.954)
1/Price	0.75126***	-0.02834
	(16.744)	(-0.583)
Size	-0.00106	-0.00112
	(-0.553)	(-0.572)
VIX	-0.00013	-0.00005
	(-1.140)	(-0.371)
PIN	0.00109	-0.03191***
	(0.138)	(-3.220)
# Observations	1,063,778	1,063,778
Firm Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes

Table XIII: Instrumental Variables Regression of Spreads on Picking-off Risk using Weekly Changes

This table reports the results of a changes IV regression where we instrument order imbalance using changes in a jump return indicator, which equals 1 if a jump occurs in stock i on date t and 0 otherwise. The order book imbalance is the ratio of buy to sell orders within 1 percent of the reference price. The AT proxy is ratio of order cancellations to total volume. The control variables are: share Turnover, log intraday price range (a dispersion proxy based on Parkinson (1980)), 1/price (based on the closing price), log market capitalization and the CBOE VIX. ***, **, and * indicate statistical significance at the 1 percent, 5 percent, and 10 percent level, respectively, using two-tailed tests and standard errors clustered at the firm- and year-level.

	(1) ΔEffective Spread	(2) ΔAdverse Selection
ΔIQR	0.21288***	0.47802***
	(5.603)	(4.135)
ΔAT	0.04253***	0.04429***
	(8.669)	(3.629)
ΔShare Turnover	0.00000	-0.00292
	(0.002)	(-1.223)
ΔVolatility	-0.02273	0.31962
	(-0.560)	(1.179)
Δ1/Price	0.88617***	-0.38044
	(6.659)	(-0.887)
ΔSize	-0.02918***	0.00904
	(-3.227)	(0.241)
ΔVIX	0.00008	0.00006
	(0.228)	(0.072)
ΔPIN	0.02133	0.01669
	(0.773)	(0.227)
# Observations	183303	183303
Firm Fixed Effects	Yes	Yes
Year Fixed Effects	No	No

Table XIV: Firm-Level Panel Regression of Picking-off Risk on Effective Spreads using Levels – IV using Lags

This table reports the results of an OLS regression of daily effective spreads on market maker monitoring, which we proxy for using IQR, for the period beginning on February 1, 2002 and ending on December 31, 2007. IQR is the one-week lagged interquartile range of the log order book imbalance, where the order book imbalance is the ratio of buy to sell orders within 1 percent of the reference price. AT is ratio of order cancellations to total volume. The control variables are: share turnover, log intraday price range (a dispersion proxy based on Parkinson (1980)), 1/price (based on the closing price), log market capitalization, the CBOE VIX and PIN (Easley and O'hara (1987)). The constant is absorbed and suppressed due to the inclusion of firm fixed effects. ***, **, and * indicate statistical significance at the 1 percent, 5 percent, and 10 percent level, respectively, using two-tailed tests and standard errors clustered at the firm- and year-level. The last specification clusters by firm ID only because the associated maximum likelihood routine accommodates only single clustering.

	(1) Adverse Selection	(2) Adverse Selection	(3) Adverse Selection	(4) Adverse Selection	(5) Adverse Selection
IQR (Lagged)	0.05779***		0.04891***	0.03295***	0.0129***
	(45.428)		(34.205)	(20.149)	(44.14)
AT		-0.02421***	0.00610***	0.00761***	0.0030***
		(-24.818)	(8.115)	(10.887)	(5.98)
Share Turnover		-0.00024**	-0.00010**	0.00001	0.0001**
		(-2.346)	(-2.155)	(0.422)	(2.76)
Volatility		0.85171***	0.50365***	0.32826***	0.3474***
		(21.567)	(19.068)	(15.668)	(21.19)
1/Price		-0.01113	0.05794*	0.18190***	0.2216***
		(-0.287)	(1.953)	(5.894)	(7.01)
Size		-0.02550***	-0.01371***	-0.01081***	-0.0122***
		(-14.847)	(-11.220)	(-8.434)	(-8.82)
VIX				0.00100***	0.0012***
				(10.155)	(32.39)
PIN				-0.00462	0.0058
				(-0.813)	(0.11)
# Observations	1,064,982	1,064,995	1,064,982	1,062,854	1,062,856
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	No	No	No	Yes	Yes
R^2	0.0477	0.0770	0.0726	0.0970	N/A