

*Do retail traders benefit from improvements in liquidity?**

Katya Malinova[†] Andreas Park[‡] Ryan Riordan[§]

*November 18, 2013
(preliminary)*

Abstract

Using intraday trading data from the Toronto Stock Exchange for 2006-2012, we study whether recent improvements in liquidity benefitted retail traders. The answer is not obvious: in our sample, retail traders trade 45% of their volume with limit orders, and our findings on retail traders' costs and benefits to limit orders are mixed. Retail traders' per dollar average intraday returns from limit orders increase over time and are positive in 2010-2012, even though the realized spread that retail traders earn on their limit orders is negative and it declines over time. Retail traders' aggregate dollar intraday returns, on the other hand, are persistently negative, and these losses are substantially higher than retail traders' losses on market orders. The discrepancy between the average and the aggregate profits stems from a long thick left tail of the distribution of retail traders' limit order profits. We also find that retail traders lose on their market orders, that these losses are closely related to the bid-ask spread, and that they decline over time.

*Financial support from the SSHRC is gratefully acknowledged. The Toronto Stock Exchange (TSX) and Alpha Trading kindly provided us with databases. The views expressed here are those of the authors and do not necessarily represent the views of the TMX Group. TSX Inc. holds copyright to its data, all rights reserved. It is not to be reproduced or redistributed. TSX Inc. disclaims all representations and warranties with respect to this information, and shall not be liable to any person for any use of this information. Portions of this paper were formerly circulated as part of a paper titled "Shifting Sands: High Frequency, Retail, and Institutional Trading Profits over Time."

[†]katya.malinova@utoronto.ca; University of Toronto

[‡]andreas.park@utoronto.ca (corresponding author); University of Toronto

[§]ryan.riordan@uoit.ca; University of Ontario Institute of Technology

A central feature of most equity markets is that sophisticated, professional investors and traders operate in the same environment as unsophisticated retail investors. The interactions present many challenges for investment dealers, market operators and regulators as they try to balance a fair and equitable treatment of traders with different abilities and different trading needs. Moreover, over the past decade, the ability gap between unsophisticated and sophisticated has, arguably, risen further as professional traders increasingly use computerized algorithmic trading.

As the computerization of trading became more prevalent, most measures of liquidity have shown significant improvements (see Hendershott, Jones, and Menkveld (2011)). In particular, the bid-ask spreads, the most standard measure for transaction costs, have almost halved over just a few years.

In this context, we ask a number of simple but important questions. First, did the market-wide improvement in bid-ask spreads benefit the group of unsophisticated retail traders? Second, what factors contribute to the unsophisticated traders gains or losses? Third, do retail traders change their behavior in response to these factors?

We address these questions using intraday trading data from the Toronto Stock Exchange from July 2006 to May 2012. The long time span of the dataset ensures that the market participants have enough time to adjust to changing market conditions, and it allows us to discuss equilibrium outcomes.

A common belief is that unsophisticated retail traders do most of their trading with market orders. In our data, however, retail traders trade, on average, only about 55% of their volume with market orders; they trade the remainder with limit orders. Although a decrease in the bid-ask spread should benefit traders that use market orders, *ceteris paribus*, this decrease harms the profitability of limit orders. Furthermore, aggressive competition from algorithmic traders to post good quotes may limit the unsophisticated traders' ability to get their limit orders executed. To understand the impact of the

improved liquidity on retail traders, we thus need to evaluate the costs and benefits to both limit and market orders.

Focusing on market orders, we observe that unsophisticated traders' intraday cumulative losses on market orders are closely related to the bid-ask spread, consistently with the predictions of information-based theoretical models (see, e.g., Glosten and Milgrom (1985) or Kyle (1985)). These losses decline, for instance, as algorithmic trading increases, where we measure the extent of algorithmic trading by the negative of the dollar volume per submitted order, as in Hendershott, Jones, and Menkveld (2011).

Since retail traders make losses on market orders yet use them persistently over the six-year time horizon that we study, these traders must also make losses when using limit orders. To test this prediction, we compute the intraday return to limit orders by comparing price paid by retail traders on the limit orders relative to the end-of-the day price. We document that in the early part of our sample retail traders indeed earned a negative average intraday return on their limit orders. In the later part of our sample, however, retail traders appear to earn a positive average intraday return on their limit orders! This outcome is surprising. The positive ex-post return on limit orders can be consistent with the negative return on market orders if traders incur high costs when their limit orders do not execute. Since retail traders are uninformed, however, their costs to non-execution should, intuitively, be limited to the transaction costs for market orders.

To understand why retail traders continue to use market orders, despite the seemingly much better performance of their limit orders, we next compute the aggregate dollar return per day, instead of the per dollar per stock return discussed above. Our results for market orders are unaffected by the change of the metric. Yet, our conclusions on the profitability of retail traders' limit orders are reversed. Even though retail traders earn positive average intraday returns per dollar traded, they lose a substantial amount of money whenever they face negative returns, and their aggregate return to limit orders

over the day, in dollar terms, is negative! In fact, on aggregate, in dollar terms, retail traders lose about twice as much on limit orders as they lose on their market orders. We further observe that the retail traders' aggregate losses for limit orders account for a large fraction of the market-wide losses on limit orders, on average about 42%, even though these traders account for only 13% of limit order trading volume.

The discrepancy between the per-stock-per-day and the aggregate measures stems from the distribution of intraday returns. For market orders, the distribution is concentrated around zero and symmetric, indicating that, loosely, unsophisticated traders act as random noise when trading with market orders. For limit orders, however, the distribution of returns has a positive average but a long and thick left tail.

Limit orders earn the bid-ask spread, and losses on limit orders stem from traders being adversely selected, for instance because new information arrives after the limit order has been submitted. While limit order submitters face adverse selection risk, this risk is arguably highest for unsophisticated traders that monitor the market less intensely. In other words, retail traders are arguably more likely than the average limit order submitter to get their orders filled when the market move strongly against them. Indeed, we find that the retail traders' limit orders trade against market orders that have higher than average price impact (the signed price change subsequent to the trade). The retail traders' limit orders thus have higher than average downside risk and the intraday returns on retail traders' limit orders have thick left tails.

Analyzing the relationship between algorithmic trading and retail traders' returns to limit orders, we find that the average per dollar intraday returns for limit orders for retail traders increase with algorithmic trading and that they become positive from 2010 onwards.¹ The improvement in the average intraday return on retail trader limit orders

¹By 2010, algorithmic trading in Canada was widespread, with several major U.S. high-frequency traders being present in Canada (this information is part of a public record) and markets showed a significant degree of fragmentation.

is consistent with Malinova, Park, and Riordan (2013) who show, for a much shorter time horizon, that retail traders benefit from the activities of high-message algorithmic traders.

In addition to studying retail traders' profits and losses, using the closing price as the benchmark, we also study these profits and losses relative to the price five minutes after the trades. The (signed) difference between the price paid by a trader on her market buy order and the midpoint of the bid-ask spread five minutes after the trade is commonly referred to as the realized (half-)spread and it proxies for the revenues received by the limit order submitter. The realized spread paid on market orders can thus proxy for market order profits and the realized spread received on limit orders proxies for profits on limit orders. We find that the realized spread paid by retail traders for their market orders is positive, that it declines over time and with algorithmic trading, consistently with our results on market order profits that use the end-of-the-day price as the benchmark.

The realized spread received by retail traders on their limit orders is always negative, it too declines over time and with algorithmic trading. The sum of the realized spreads paid on market orders and received on limit orders proxies for the difference in profits on limit and market orders; as discussed above, we expect this difference to be positive, since limit order profits are realized with the probability below 1. We observe that this difference is indeed positive in the early part of our sample, yet it becomes negative in 2009 and it remains negative thereafter. Using realized spreads as a proxy for limit order revenues, we would thus conclude that in the later part of our sample retail traders were better off trading with market orders than with limit orders — this observation contradicts our earlier conclusion that was based on the per dollar profits to limit orders, using the end-of-the-day price as the benchmark. Our results indicate that the five-minute realized spreads may not reflect the limit order revenues, particularly, in the later part of our sample when algorithmic trading became very prevalent.

Our work is closely related to the literature on trader and investor performance and

profitability. Hasbrouck and Sofianos (1993) study the profitability of market makers and break down their profits into liquidity and trading related. Barber and Odean (2000) show that active retail traders' portfolios underperform the market. Barber and Odean (2002) show that as investors switch to online brokerages, and trade more, their performance falls. Using a Taiwanese investor-level dataset, Barber, Lee, Liu, and Odean (2009) find that retail traders lose on their aggressive trades. Several studies, e.g. Kaniel, Saar, and Titman (2006), Hvidkjaer (2008), or Kelley and Tetlock (2013), find that retail traders net buying has predictive power for future returns. Foucault, Sraer, and Thesmar (2011) show that trades by retail investors contribute to the idiosyncratic volatility of stocks. Hau (2002) finds that traders close to corporate headquarters have higher trading profits than those located farther away, Dvořák (2005) finds that domestic traders are more profitable than foreign traders.

I. Theoretical Predictions

A. *Retail Traders' Profits*

Most models of information-based intraday trading assume the presence of traders who trade for non-informational reasons such as liquidity needs; these reasons are often not modeled explicitly and the traders are referred to as noise traders (see, e.g., Glosten and Milgrom (1985) and Kyle (1985)). In these models, traders who use market orders incur the bid-ask spread transaction costs because quotes reflect that possibility that a market order submitter is better-informed than the market maker. The uninformed traders who use market orders then make a loss because they must pay the transaction costs.

The informativeness of a market order is commonly measured by the order's price impact, which is defined as (twice) the signed price movement subsequent to the trade. If retail traders were always uninformed, then the price impact of their orders should

be zero; we use a weaker version and conjecture that retail traders are less-well-informed than the average trader so that the information content of retail market orders is therefore smaller than the market-wide average.

We predict that the retail traders make intraday losses on their market orders, and we use two measures to quantify these losses. First, we compute the realized spread, which is the difference between the bid-ask spread cost incurred by retail traders on their market orders and the price impact of their orders; we predict that retail traders pay a positive realized spread. Second, we compute the intraday return from retail traders' market orders, by comparing the price paid by these traders to the end-of-the-day closing price as a price benchmark.

In a market that is organized as a limit order book, traders may either post limit orders, specifying the terms at which they are willing to trade, or trade immediately against a previously posted limit order by submitting a market order. In equilibrium, the average trader should be indifferent between the two order types. If retail traders lose money when trading with market orders, they would only use these orders if they also lose money when trading with limit orders (or, if non-execution of their limit orders compels them to incur costs in excess of the market order costs).

Empirical Prediction 1 *Retail traders are uninformed. Therefore,*

- 1. the price impact of retail traders' is smaller than the market average price impact;*
- 2. their intraday returns from market orders are negative, and they pay a positive realized spread;*
- 3. their intraday returns from limit orders are negative, and they receive a negative realized spread on their limit orders.*

B. Comparison of Retail Traders' Profits on Market and Limit Orders

Importantly, when traders have the choice between submitting a market and a limit order, traders must weigh the better price offered by the limit order relative to the price of the

market order against the uncertain execution of a limit order (e.g., as in Parlour (1998)): market orders provide with certainty of execution whereas limit orders only trade when in the future there is a market order to match the limit order. In models with information, a limit order submitter must further account for the possible arrival of new information after the submission of their limit orders (e.g., Foucault (1999), Kaniel and Liu (2008), Rosu (2013), or Brolley and Malinova (2013)). The limit order is more likely to execute if the new information is unfavorable to the limit order submitter, and limit order submitters thus must account for these adverse selection costs.

We predict that retail traders are uninformed and that they make losses; they must thus have private motives for trading such as liquidity needs. We derive our empirical predictions, using a simple framework that models these needs by assigning a private value to trading (an alternative approach would be to assume a cost of non-execution). Since the execution of a limit order is uncertain, the expected profit from the market order must be smaller than the expected profit from the limit order.

Formally, denoting profits from market and limit orders as π^m, π^ℓ , if a trader has private value Y , the stock's fundamental value is V , $\text{pr}^{\text{execution}}$ is the probability that the limit order executes, and p^ℓ and p^m are the prices for market and limit orders, then a trader is indifferent between a market buy and a limit buy order if

$$\pi^m = \text{pr}^{\text{execution}} \pi^\ell \Leftrightarrow Y + \mathbb{E}[V] - p^m = \text{pr}^{\text{execution}} \times (Y + \mathbb{E}[V \mid \text{execution}] - p^\ell). \quad (1)$$

With the probability of the limit order execution $\text{pr}^{\text{execution}} \in (0, 1)$, the profit, conditional on the execution of the limit order must exceed the profit from the market order:

$$\mathbb{E}[V] - p^m < \mathbb{E}[V \mid \text{execution}] - p^\ell. \quad (2)$$

The difficulty lies in the choice of the benchmark for the value of the stock against which

to compare the cost of the trade. We consider two benchmarks.

First, we use the stock's closing price as the common benchmark for all trades. Second, we use the midpoint of the bid-ask spread five minutes after the trade, which is a common benchmark in the literature. In this case, we further decompose the market and limit order prices into the prevailing midpoint m_t at the time of the trade and the bid-ask half-spread s_t . The price for the buy market order at time t is the ask price at time t , and it can be written as $p^m = \text{ask}_t = m_t + s_t$. Using the midpoint 5 minutes after the trade as the benchmark for the fundamental value, we obtain (for the buy market order):

$$\mathbb{E}[V] - p^m = m_{t+5} - m_t - s_t = \frac{1}{2}\text{price impact}_t - \frac{1}{2}\text{effective spread}_t = -\frac{1}{2}\text{realized spread}_t, \quad (3)$$

where the factor $1/2$ is due to the common normalization of the price impact as twice the signed change of the midpoint after the trade. The price for the buy limit order that executes at time t is the bid price at time t , and it can be written as $p^l = \text{bid}_t = m_t - s_t$. For a buy limit order to execute at time t , the transaction at time t must be initiated by a sell market order. The (half-) price impact is thus: $\frac{1}{2}\text{price impact}_t = -(m_{t+5} - m_t)$. With this in mind, conditional on the buy limit order execution, we obtain:

$$\mathbb{E}[V \mid \text{execution at } t] - p^l = m_{t+5} - m_t + s_t = \frac{1}{2}\text{realized spread}_t \quad (4)$$

Rewriting the inequality (2) in terms of the realized spreads, we thus predict:

$$-\text{realized spread paid on market orders} < \text{realized spread earned on limit orders}. \quad (5)$$

Empirical Prediction 2 *For retail traders, the profits from limit orders must exceed profits to market orders. Thus:*

1. *the intraday return earned from trading with limit orders exceeds the return earned on market orders;*

2. *the sum of the realized spread received on limit orders and the realized spread paid on market orders is positive.*

C. Factors that Influence Retail Traders' Profits

Algorithmic Trading. We have discussed above that the submitter of a limit order faces a form of the winner's curse: she may trade precisely when she does not want to because prices move against her. Hoffman (2013) develops a model with fast and slow traders, where fast traders are able to avoid being adversely selected by slow traders because, upon arrival of the unfavorable new information, they are able to modify their limit orders before the (slow) arriving trader submits a market order. Slow traders, on the other hand, are always adversely selected if the new information is unfavorable.

Based on Hoffman (2013), *ceteris paribus*, an increase in algorithmic trading should not affect the adverse selection costs for the slow traders such as retail. The impact of algorithmic trading on the market-wide average price impact depends, intuitively, on its impact on the amount of liquidity provision by fast traders. If, as algorithmic trading increases, fast traders are responsible for a larger fraction of limit order trading, then the average price impact of market orders should decline (since the limit orders of fast traders face lower price impacts than the limit orders of slow traders). If, on the other hand, the fraction of executed limit orders that stem from the fast liquidity providers declines as the algorithmic trading increases, then the effect is the opposite.

If fast traders are indeed able to modify their limit orders before they are adversely selected, they will require lower compensation and the bid-ask spread will decline (see, e.g., Copeland and Galai (1983), Foucault (1999), or Bernales and Daoud (2013)). Assuming that the retail traders are close to uninformed and that the price impact of their market orders is close to zero, the decline in the effective spread must benefit those retail traders that use market orders.

Since retail traders who submit limit orders have a choice of switching to market orders, for them to continue to use both market and limit order in equilibrium, either the probability of execution of their limit order or the retail traders' profits from limit orders must increase.

Empirical Prediction 3 *As algorithmic trading increases,*

- 1. the price impact incurred by retail traders' limit orders is not affected;*
- 2. the market-wide average price impact may increase or decline, depending on the fraction of executed limit orders that stem from the fast liquidity providers;*
- 3. the market-wide effective spread and the effective spread paid by retail traders for their market orders decline;*
- 4. retail traders' intraday return from market orders increase, and the realized spread paid on their market orders declines;*
- 5. either the probability of executions of retail limit orders or retail traders' profits from limit orders must increase.*

Price Changes. Intuitively, the amount of information on a given day is positively related to the absolute price change on that day as well as to the fraction of a one-sided, directional order flow. Although the retail traders are uninformed, they would incur higher bid-ask spread costs on their market orders if the average price impact is higher, and thus incur higher losses on their market orders. Retail traders that use limit orders will incur higher adverse selection costs. The impact of the absolute price changes on the retail traders' profits to limit orders will depend on whether they are able to adjust their quotes to account for the higher adverse selection costs. Furthermore, even if the retail traders are able to correctly adjust the pricing of their limit orders, the tradeoff between market and limit orders implies that they may be willing to accept higher losses on their limit orders when they incur higher losses on their market orders. The market vs. limit order tradeoff depends on the probability of execution for limit orders, and we do not have directional predictions.

Empirical Prediction 4 *As markets move strongly, either in the form of directional order flow or price movements,*

- 1. Retail traders’ profits from market orders decline;*
- 2. the price impact incurred by retail traders on their limit orders increases.*

II. Data and Sample Selection

A. Data

Our analysis is based on a proprietary dataset, provided to us by the TMX Group; we use additional proprietary methods to identify retail traders.² Index constituent status is obtained from the monthly TSX e-Review publications. Data on the U.S. volatility index VIX is from the CBOE database in WRDS.

The TSX data is the output of the central trading engine, and it includes all messages from the (automated) message protocol between the brokers and the exchange. Messages include all orders, cancellations and modifications, all trade reports, and all details on dealer (upstairs) crosses. The data specifies the active (liquidity demanding) and passive (liquidity supplying) party in a trade, thus identifying each trade as buyer-initiated or seller-initiated. The “prevailing quote” identifies the best bid and ask quotes and is updated each time there is a change in the best quotes.

Unique Identifiers. Our data has unique identifiers for the party that submitted an order to the exchange. A unique identifier links orders with a trading desk in charge of the order at a brokerage. Our unique identifiers are similar to those used by the Investment Industry Regulatory Organization of Canada (IIROC), and, according to them, are “the most granular means of identifying trading entities.” For IIROC’s data, a unique identifier may identify a single trader, a direct-market access (DMA) client, or a business flow (for

²Legal disclaimer: TSX Inc. holds copyright to its data, all rights reserved. It is not to be reproduced or redistributed. TSX Inc. disclaims all representations and warranties with respect to this information, and shall not be liable to any person for any use of this information.

example, orders originating from an online discount brokerage system). A DMA client may have multiple unique identifiers if they trade through multiple TSX participating organizations (i.e., multiple brokers), or for business or administrative purposes.³ To the best of our knowledge, most brokerages funnel specific types of order flow through separate unique identifiers, and they do not mix, for instance, retail and institutional order flow.

Our sample contains a total of 14,182 unique identifiers, but only around 2,000 unique identifiers are active in the market per month. It is our understanding that a unique identifier may be associated with a particular individual or with a particular algorithm, and that identifiers may be replaced when traders or algorithms change. Figure 1 plots the number of unique identifiers in our sample across time.

Securities. Our main analysis focusses on the 48 equities that are continuously in the S&P/TSX60 index (Canada’s Blue Chip index) for our sample period.

Dates. We study the period from July 1, 2006 to June 1, 2012. Since we classify traders based on a 20-day rolling window of past behaviour, we omit the month of July in our formal regression analysis. For each security, we exclude the entire day if trading in that security was halted. We are interested in “normal” activity, and most of our trader classification is by week. We thus exclude weeks that saw either extraordinarily low or high activity. Specifically, we exclude the weeks that had 3 or fewer days of both U.S. and Canadian trading: the week of the 4th of July (the first of July is a national holiday in Canada) in 2006, 2007, and 2008; the weeks of US Thanksgiving; the last week of all years, and the first week of the year in 2007, 2008, and 2009. We further exclude the week of the May 6th 2010 Flash Crash and the five weeks following the Lehman Brothers bankruptcy on September 15, 2008, we exclude the week of December 15, 2008 because the TSX experienced a technical glitch on December 17, 2008 and was closed for the entire

³See the description of User IDs and account types in IIROC (2012). In IIROC’s data, a User ID may use different account types, e.g. proprietary and client. In our data, different account types for the same user would correspond to separate unique identifiers.

day; trading activity on the following day (December 18, 2008) was also extremely low. Finally, our data files for January 2008 were corrupted and we have no data for the month.

III. Classification of Unsophisticated (Retail) Traders

Our usage of the term “retail trader” is to be understood as being synonymous to “unsophisticated”, and we use the terms interchangeably. In the past, researchers associated retail trades with small order sizes. In today’s market, large institutional orders, referred to as “parent” orders, are commonly split into smaller “child” orders before being submitted to the exchange, thus retail orders may exceed the average order size. Since our data only allows us to see the orders sent to the exchange, an observed order size is a poor proxy for the level of trader sophistication.

We use a two-pronged approach to identifying unsophisticated traders. For the later years of our sample, we can identify some unique identifiers as retail, based on proprietary methods.⁴ This set is small and prior to 2010, only few of these identifiers appear in our data.

As a second criteria we use the time that a trader’s passive order remained in the limit order book as a proxy for the trader’s monitoring activities and the level of sophistication. Specifically, we classify a unique identifier as managing non-sophisticated (“retail”) order flow if in the past 20 trading days plus the current day, the trader traded with an order that stayed in the order book overnight; such orders are also referred to as “good-till-cancelled” (GTC) or “good-till-filled” (GTF) orders, While these unique identifiers do not necessarily represent retail clients, trading with stale orders indicates low market monitoring, which we believe to be correlated with traders’ levels of sophistication.⁵

⁴Based on this method we capture some, but not necessarily all of unique identifiers that represent a retail flow and that trade on the TSX.

⁵We acknowledge some that some sophisticated trader may occasionally trade with GTC/GTF orders because they post so-called “stub” orders at extreme prices.

The practice of trading with GTC/GTF orders is arguably endogenous to market conditions, and we indeed observe a decline over time in the number of traders that use them. However, the total number of traders that either use “stale” orders or that we identify using proprietary methods remained relatively stable over time, with an average of 42 traders per day per stock; see Figure 1 for the time series of unsophisticated traders.

IV. Trading Costs and Benefits

A. Realized Spread

A common measure for the benefits of liquidity provision is the realized spread, defined as:

$$rspread_{it} = 2q_{it}(p_{it} - m_{i,t+5 \text{ min}})/m_{it}, \quad (6)$$

where p_{it} is the transaction price, where $m_{i,t+5 \text{ min}}$ is the midpoint of the quoted bid-ask spread that prevails 5 minutes after the trade, and q_{it} is an indicator variable, which equals 1 if the trade is buyer-initiated and -1 if the trade is seller-initiated. Our data includes identifiers for the active side (the market order that initiated the trade) and for the passive (the limit order) side of each transaction, precisely signing the trades as buyer- or seller-initiated. The data also contains the prevailing (Canadian) National best quotes at the time of each transaction.

The realized spread is commonly computed in relation to the price impact and the effective spread, where the price impact is the signed change in the midpoint of the bid-ask spread from the time of the trade to five minutes later:

$$price \ impact_{it} = q_{it}(m_{t+5 \text{ min},i} - m_{it})/m_{it}, \quad (7)$$

and where the effective spread

$$espread_{it} = q_{it}(p_{it} - m_{it})/m_{it}, \quad (8)$$

where m_{it} is the midpoint of the quoted spread prevailing at the time of the trade. The sum of the price impact and the realized spread is the effective spread.

The price impact is commonly interpreted to capture the adverse selection component of a trade, the realized spread is then the rent that pertains to the liquidity provider. For traders who use market orders, the realized spread thus reflects the fee that they have to pay to liquidity providers in excess of the compensation for the adverse selection component of the market orders. We compute the realized spread separately for market orders and limit orders for retail traders (for all traders together, the two measures would coincide, but this need not be the case for individual trader groups).

The five minute benchmark, which assumes that the price five minute subsequent to the trade fully reflects the information content of the market order that initiated that trade. We use the five minute benchmark rather than a shorter horizon one because we aim to capture the adverse selection against traders who trade to build long-term positions. The five-minute realized spread is likely not a valid metric to assess the benefits from liquidity provision for traders who hold the position for very short time horizons, such as high-frequency market makers, because these traders may manage their inventories in such a way so that they wouldn't hold the position even until the five minute benchmark.

B. Intraday Returns

If prices include all information at any point in time, then any price movement subsequent to a trade is the result of new information (or noise). By holding the security, an investor then earns a return on his/her investment. On the other hand, if, for instance, an informed order is split into many small orders and the total information content of the order is only revealed over time, then anyone trading against the split order will lose. This concern is particularly prevalent for limit orders because these get filled in particular if the market moves against them (i.e., limit buy orders execute particularly when the price drops due

to market sell orders). Furthermore, uninformed traders must thus take into account that they may trade at the wrong time, before prices reflect all the available information.

We use two measures for intraday returns. First, we compute a trader’s profit from buying and selling a security and we value the end-of-day portfolio holdings at the closing price; we refer to this measure as the $\$return$; formally

$$\$ return_{it} = (sell\$vol_{it} - buy\$vol_{it}) + (buy vol_{it} - sell vol_{it}) \times cprice_{it} \quad (9)$$

where $sell\$vol_{it}$ and $buy\$vol_{it}$ are the total sell and buy dollar-volumes for trader-group i , $buy vol_{it}$ and $sell vol_{it}$ are the share-volumes. Second, we scale this profit measure by the daily dollar volume to obtain the $\%return$:

$$\% return_{it} = \$ return_{it} / \$vol_{it}, \quad (10)$$

where $\$vol_{it} = sell\$vol_{it} + buy\$vol_{it}$ is the overall dollar volume. The profit from intraday trading is $(sell\$vol_{it} - buy\$vol_{it})$; a positive value means that the trader group “bought low and sold high.” The term $(buy vol_{it} - sell vol_{it})$ is the end-of-day net position (assuming a zero inventory position at the beginning of each day), which we evaluate at the closing price, $cprice_{it}$. We compute these returns separately for market and limit orders.

In the preceding section, we described the realized spread which is a standard measure that captures price movements subsequent to a trade. The intraday return capture price movements using the closing price as a common benchmark and we thus implicitly assume that the closing price reflects the total information that was generated during a trading day, including the information revealed by split orders.

We acknowledge that our analysis is based on TSX trading only — traders may well trade on other Canadian venues or in the U.S. as part of a cross-venue or cross-country arbitrage strategy, and thus their actual intraday returns may be different from the ones that we record. However, by regulation, Canadian retail trades must be posted to visible

exchanges (they cannot be systematically internalized). Moreover, for our profit measure to exhibit a systematic bias, the entire group of interest must persistently choose, for instance, to buy all the securities on the TSX and to sell them on a different exchange.

V. Testing Predictions on Retail Trader Profits

A. Testing Empirical Prediction 1: “Retail traders earn non-positive intraday returns”

We perform a t-test on the sign of %returns, \$ returns and realized spreads paid and received. We employ standard errors that are double-clustered by firm and date to control for cross-sectional and time-series correlation.

Panel A in Table III displays our results from these t-tests. We observe, as predicted in Empirical Prediction 1, that retail traders lose on their market orders, both in terms of %returns and \$returns, and that they pay positive realized spreads. Furthermore, they earn negative realized spreads on limit orders and they lose in terms of \$ intraday returns on limit orders. The %returns are indistinguishable from 0.

In Panel B, we further test whether there are changes across time by splitting our (almost) 6-year sample into 12 half-years. The % returns and \$ returns display the same patterns as the full sample, and their size are of similar magnitude across the different half-years. However, the % returns for limit orders increase over time (but they do not become statistically significantly positive). Realized spreads paid and received decline over time but they maintain the signs that are predicted by Empirical Prediction 1.

Finally, Empirical Prediction 1 also states that market orders by retail traders have lower price impacts than the average market order. Indeed, the last column in Table III indicates that the difference between the marketwide price impact and the price impact by retail traders is positive.

B. Testing Empirical Prediction 2: “Limit order profits exceed market order profits”

Similarly to the preceding subsection, we test Empirical Prediction 2 by performing a t-test on the difference in %returns and \$returns for market and limit orders. Table IV shows that, as predicted, limit orders have higher % returns than market orders. However, the result for \$returns is the reverse.

This puzzling discrepancy can be explained when considering the distribution of losses. Figure 4 plots the densities for all observations in our sample of \$returns for market and limit orders. As can be seen, the mode and the average of the limit order \$returns is larger than for the market order \$ returns. However, limit order \$returns have a long left tail. As Panel B in Figure 2 highlights, these tail losses are persistent (the line is non-increasing) for limit orders. Panel B in Table IV outlines that across time, %return and \$returns maintain the same pattern as for the full sample.

Finally, Empirical Prediction 2 also asserts that the sum of the realized spreads for market and limit orders must be positive. The last column in Table IV shows that we must reject this notion: the sum of the realized spreads appears to be negative. Inspecting the evolution of the sum across time we observe that for the early part of our sample, the sum of spreads was positive but that, starting in 2009, the sum became negative. Figure 3 illustrates this finding graphically.

The discrepancy in the results hints at two concerns with regard to the measurement of trading benefits. First, limit orders face a substantial tail risk that average or normalized measures do not capture. Second, the realized spread indicates that retail traders persistently use limit orders even though they should not. They are thus either persistently irrational or the measure is inadequate. Since our %return yields the predicted relation (opposite to that implied by the realized spreads) realized spreads may not capture gains and losses from limit orders well.

VI. Testing for Contributing Factors

Methodology. In our regression analysis we seek to understand the factors that influence the returns to market and limit orders for retail traders and all traders. We employ a standard OLS regression, using stock fixed effects and stock and date double clustered errors to account for autocorrelation and heteroscedasticity. To ensure that our results are not driven by outliers, we winsorize all variables at the 1% level.⁶ Specifically, we estimate the following type of equation

$$\text{dependent variable}_{it} = \alpha_{(i)} + \beta X_{it} + \epsilon_{it}, \quad (11)$$

where $\alpha_{(i)}$ are the stock fixed effects and X_{it} is a vector of covariates of interest.⁷

Covariates. We are interested in three variables that may capture the extent of the winner’s curse when trading with limit orders. First, we compute the imbalance of trading, measured as the absolute difference in the buyer- and seller-initiated dollar volume relative to the total dollar volume

$$|\text{imbalance}_{it}| = |\text{buy \$-volume}_{it} - \text{sell \$-volume}_{it}| / (\text{buy \$-volume}_{it} + \text{sell \$-volume}_{it}).$$

This measure would capture the extent of buying and selling pressure.

Second, we compute the absolute size of the closing price return, measured as

$$|\text{return}_{it}| = |\text{close}_{it} - \text{close}_{it-1}| / \text{close}_{it-1}.$$

⁶Results without winsorization are qualitatively similar.

⁷To further understand the evolution of effects across time, we further tested a specification in which we split the covariates by quarters; since our data is from August 2006 to June 2012, we have a total of 24 quarters. The estimated effects for the covariates commonly showed only little variation and we thus omit the tables.

This measure reflects the extent of price movements for the trading day.

Third, we are interested in the relation of our variables to algorithmic trading, which, following Hendershott, Jones, and Menkveld (2011), we measure by the negative of the dollar volume per exchange message

$$\text{algo trading}_{it} = -\$ \text{ volume}_{it} / (100 \times \#\text{messages}_{it}),$$

where $\#\text{messages}_{it}$ is the sum of the number of orders submissions, cancellations and modifications and trades. Algorithmic trading is marked by the frequent posting and cancellation of quotes and messages. Hendershott, Jones, and Menkveld (2011)'s measure thus captures how many messages are used to trade \$100 of volume.

Table II lists the per stock per day correlations among our covariates. We have considered using the U.S. volatility index VIX, as is standard in many microstructure studies, but it is highly correlated with returns; other measures of intraday volatility exhibited similar correlations. To ensure that our coefficients are meaningful we thus omit a volatility measure.

Summary Statistics Table I displays the summary statistics for trading variables for our sample. There is a substantial decline in different measures of the bid-ask spread over time, and there is a decline in the price impact.

Results on Algorithmic Trading and Price Changes. In Tables V, VI, and VII, we test Empirical Predictions 3 and 4 on the relationship of algorithmic trading and absolute price changes with retail traders' profits to market and limit orders and with retail traders' adverse selection costs for limit orders.

Adverse Selection. Consistent with Empirical Predictions 3 and 4, we find that the price impact that retail traders face on their limit orders does not change with algorithmic

trading and that it increases in the absolute value of the price change. The market-wide price impact, which we did not have a directional prediction on, also does not change with algorithmic trading, and this price impact increases with the absolute value of the price change.

Market Orders. Consistently with our predictions, we find that as algorithmic trading increases, the market-wide effective spread declines, that the realized spread paid by retail traders declines, and that the retail traders' intraday returns to market orders increase. Table VII documents an (economically small) increase in the effective spread as the absolute price change increases. Tables V and VI show that this increase translates into a decline in intraday returns to retail traders' market orders and in the realized spread paid by retail traders' on their market orders.

Limit Orders. Our empirical predictions for limit order profits depend on the probability of execution for limit orders. Table VII shows that the probability of execution for limit orders declines with algorithmic trading. Since retail traders' profits for market orders increase with algorithmic trading, for retail traders to continue to use market orders in equilibrium, their profits for limit orders must increase, too. Panel B in Table V supports this prediction: measuring limit order profits as per dollar intraday returns, we find that these increase with algorithmic trading. Panel A in Table VI indicates, however, that the realized spread received on retail traders' limit orders declines with algorithmic trading. We believe that in combination, our results suggest that the realized spread need not reflect limit order profits.

Turning to the impact of the absolute price changes, we find, in Table VII, that the probability of limit order execution for retail traders increases as the absolute price change increases. Tables V and VI illustrate that profits to limit orders decline with the absolute value of the intraday price change for all of our measures of limit order profits. Furthermore, for all these measures, the decline in retail traders' limit order profits that

is associated with an increase in the absolute value of the price change is substantially higher than the corresponding decline in market order profits.

VII. Discussion and Conclusion

Over last decade, equity markets have undergone substantial changes. Most work to date indicates that these changes have been for the better and, in particular, that liquidity has improved following these market changes. In this context, we ask a simple question: did the market-wide improvements in liquidity benefit the less-sophisticated traders?

For instance, while a decline in bid-ask spreads (a standard measure of liquidity) makes trading with market orders cheaper, it may make trading with limit orders more expensive. Do retail traders respond optimally to the changes in market structure? Using a standard measure for limit order revenues, the realized spread, we find that, taken at face value, changes in this measure over time could appear to indicate that retail traders behave irrationally. At the same time, computing the intraday returns to limit orders shows that retail traders need not act irrationally. Finally, in trying to identify changes in trading costs and benefits, we also observe that unsophisticated traders face substantial tail risk when using limit orders, a risk that averages may not pick up.

Appendix: Trading on the TSX and Notable Events

The Canadian market is similar to the U.S. market, in terms of regulations and market participants. Trading on the TSX is organized in an electronic limit order book, according to price-time priority.⁸ The Canadian market is closely connected to the United States, and most U.S. events have an impact on Canadian markets.

⁸Differently to U.S. markets, trades below 100 shares (odd-lot trades) are executed outside of the limit order book, by a registered trader. Canadian markets further employ so-called broker-preferencing, so orders that originate from the same brokerage may be matched with violations of time-priority.

Our data spans July 1st, 2006 to June 1st, 2012. During this time financial markets overall and Canadian markets, in particular, have experienced several major shocks and changes. We discuss changes that are most relevant to our analysis below.

2006-2007: Introduction of Maker-Taker Fees. In 2006, almost all Canadian trading in TSX-listed securities occurred on the TSX. High-frequency traders were few, with the most prominent known participant being Infinium Capital Corporation, which became a participating organization on the TSX on March 11, 2005.⁹ On July 01, 2006, the TSX introduced maker-taker pricing, with rebates for liquidity provision, for all securities (a pilot started in 2005). According to the S.E.C., liquidity rebates facilitated the development of high frequency liquidity provision, and we thus chose to start our analysis in July 2006.¹⁰

2008-2009: Financial Crisis, Market Fragmentation, and Arrival of HFT. The year 2008 was pivotal for Canadian markets: with the 2008 financial crisis, the growth of alternative trading systems, and the arrival of major U.S. high frequency trading firms.

Pure Trading opened a cross-printing facility in October 2006 and launched continuous trading in September 2007; Omega ATS launched in December 2007; Chi-X Canada started in February 2008; Alpha Trading launched in November 2008. Continuous dark trading plays only a small role in Canada: the main dark pool, MATCH Now, launched in July 2007, and it has had an average market share of under 3%. The graph shows that the TSX market share declined from 93% of dollar volume traded in January 2009 to 65% in January 2010, and that the TSX market share has varied between 60 and 70% since January 2010. The two major competitors of the TSX are Alpha and Chi-X. Alpha's market share increased from 1.7% in January 2009 to 22.5% in January 2010, and has varied between 15 and 22% since then; Chi-X Canada's market share went from 2.7% in January 2009 to 8% in January 2010, and has been between 7 and 11% since then.

⁹See the TMX Group Notice 2005-010.

¹⁰See the S.E.C. 2010 Concept Release on Market Structure.

The arrival of alternative trading systems coincided with several major changes on the TSX. First, in 2007 the TSX introduced a new, faster trading engine. Second, in 2008, the TSX offered co-location services, allowing traders to place their servers physically close to the main TSX trading engine.¹¹ Third, on October 29, 2008, the TSX started the “Electronic Liquidity Provider” (ELP) program that offered “fee incentives to experienced high-velocity traders that use proprietary capital and passive electronic strategies to aggressively tighten spreads on the TSX Central Limit Order Book.”

The above developments have arguably facilitated the arrival of major U.S. HFT firms. From public sources, we know that Getco LLC has participated in Canadian markets since 2008, and that Tradebot Systems expanded trading to Canada in October 2008.¹² Dave Cummings, Chairman of Tradebot Systems Inc. commented in particular, that “[t]he ELP program [...] was a major factor in [Tradebot’s] decision to trade in the Canadian market.”¹³

2010-2012: Flash Crash, Debt Crisis, and Per-Message-Fees. The years 2010-2012 were much calmer, with some notable exceptions. First, the May 6, 2010 “Flash Crash” drew public attention to the presence of high-frequency trading. The flash crash was visible in Canadian data, albeit the price fluctuations were of smaller magnitude. Second, and more generally, May 2010 was the beginning of the European sovereign debt crisis, with high volatility levels. Volatility was also high from August to October 2011, fuelled by concerns over the U.S. debt ceiling and over the future of the European Monetary Union. Finally, on April 01, 2012 IIROC introduced a policy change that had a pronounced effect on high frequency trading. Namely, as of April 01, 2012, brokers incurred a fee per

¹¹The TSX QuantumTM engine migration was completed in May 2008; see TSX Group notice 2008-021. The TMX Group 2008 annual report describes the introduction of co-location services.

¹²See, e.g., Getco LLC regulatory comment to the CSA, available at http://www.osc.gov.on.ca/documents/en/Securities-Category2-Comments/com_20110708_23-103_kinge.pdf for Getco LLC arrival and “Tradebot Systems often accounts for 10 percent of U.S. stock market trading”, Kansas City Business Journal, May 24, 2009, for Tradebot Systems Inc. arrival.

¹³<http://www.newswire.ca/en/story/269659/tmx-group-targets-liquidity-with-reduced-equity-trading-fees>, Canada Newswire, October 29, 2008.

submitted message (such as a new order, trade, or a cancellation of an order). This fee is meant to recover the costs of surveillance and its exact magnitude varies with the total number of messages submitted by all participants.¹⁴

Fragmentation levels did not change much during this period, although, new venues did come in. In June 2011, Alpha launched its IntraSpreadTM dark pool, and almost concurrently, in July 2011, the TMX Group launched the alternative trading system TMX Select.¹⁵

References

- Barber, Brad M., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean, 2009, Just how much do individual investors lose by trading?, *Review of Financial Studies* 22, 609–632.
- Barber, Brad M., and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *The Journal of Finance* 55, 773–806.
- , 2002, Online investors: Do the slow die first?, *Review of Financial Studies* 15, 455–488.
- Bernales, Alejandro, and Joseph Daoud, 2013, Algorithmic and high frequency trading in dynamic limit order markets, *SSRN eLibrary*.
- Brolley, Michael, and Katya Malinova, 2013, Informed trading and maker-taker fees in a low-latency limit order market, Working paper University of Toronto.
- Copeland, T.E., and D. Galai, 1983, Information effects on the bid-ask spread, *Journal of Finance* 38, 1457–1469.
- Dvořák, Tomáš, 2005, Do domestic investors have an information advantage? Evidence from indonesia, *The Journal of Finance* 60, 817–839.
- Foucault, Thierry, 1999, Order flow composition and trading costs in a dynamic limit order market1, *Journal of Financial Markets* 2, 99–134.

¹⁴According to a research report by CIBC (2013), this fee is on the order of \$0.00022 per message. Malinova, Park, and Riordan (2012) study the impact of this change on market quality.

¹⁵According to IIROC's published market shares statistics, TMX Select's 2012 market share of dollar volume was around 1.4%. IntraSpread is not listed separately in IIROC's market shares because IntraSpread is technically part of Alpha's main market. According to Alpha's April 2012 newsletter, the market share of IntraSpreadTM was about 3.5-4%.

- , David Sraer, and David J. Thesmar, 2011, Individual investors and volatility, *The Journal of Finance* 66, 1369–1406.
- Glosten, L., and P. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogenously informed traders, *Journal of Financial Economics* 14, 71–100.
- Hasbrouck, Joel, and George Sofianos, 1993, The trades of market makers: An empirical analysis of NYSE specialists, *The Journal of Finance* 48, 1565–1593.
- Hau, Harald, 2002, Location matters: An examination of trading profits, *The Journal of Finance* 56, 1959–1983.
- Hendershott, T., C. M. Jones, and A. J. Menkveld, 2011, Does algorithmic trading improve liquidity?, *Journal of Finance* 66, 1–33.
- Hoffman, Peter, 2013, A dynamic limit order market with fast and slow traders, working paper 1526 European Central Bank.
- Hvidkjaer, Soeren, 2008, Small trades and the cross-section of stock returns, *Review of Financial Studies* 21, 1123–1151.
- IIROC, 2012, The HOT study, Discussion paper, Investment Industry Regulatory Organization of Canada.
- Kaniel, R., and H. Liu, 2008, Individual investor trading and stock returns, *Journal of Business* 1, 1867–1913.
- Kaniel, Ron, Gideon Saar, and Sheridan Titman, 2006, So what orders so informed traders use?, *Journal of Finance* 79, 273–310.
- Kelley, Eric K., and Paul C. Tetlock, 2013, How wise are crowds? insights from retail orders and stock returns, *The Journal of Finance* 68, 1229–1265.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.
- Malinova, Katya, Andreas Park, and Ryan Riordan, 2012, Do retail traders suffer from high frequency traders?, Working paper University of Toronto <http://ssrn.com/abstract=2183806>.
- , 2013, Do retail traders suffer from high frequency traders?, *SSRN eLibrary*.
- Parlour, C., 1998, Price dynamics in limit order markets, *Review of Financial Studies* 11, 789–816.

Rosu, I., 2013, Liquidity and information in order driven markets, *SSRN eLibrary*.

Securities and Exchange Commission, 2010, Concept release on market structure, Release No. 34-61358, File No. S7-02-10, Discussion paper, Securities and Exchange Commission <http://www.sec.gov/rules/concept/2010/34-61358.pdf>.

Ye, Mao, Chen Yao, and Jiading Gai, 2012, The externalities of high frequency trading, Discussion paper, University of Illinois Champaign-Urbana SSRN: <http://ssrn.com/abstract=2066839>.

Table I
Summary Statistics

This table presents summary statistics by year. Each number is per day per security for the 48 stocks in our sample. The sample is from July 1st 2006 to June 1st 2012, but omits a number of weeks as described in the Section II.. The daily VWAP is the volume weighted average trade price; |price change| is the absolute price change from close-to-close, scaled by yesterday's close; |imbalance| is the absolute difference between the buyer and seller initiated \$volume, scaled by the daily volume; algo trading is Hendershott, Jones, and Menkveld (2011)'s measure of algorithmic trading, defined as the negative of the \$-volume per message (trades and order submissions, cancellations and modifications). MO abbreviates market order, LO abbreviates limit order. All volume measures involve the double-counting of volume. The minimum price increment/tick size in our sample is 1 cent.

	Units	2006	2007	2008	2009	2010	2011	2012
<i>Panel A: Marketwide statistics</i>								
daily VWAP	dollars	51.4	54.3	48.5	39.0	44.4	42.9	39.2
price change	percent	1.3	1.3	2.5	2.1	1.1	1.4	1.2
effective spread	cent	3.5	3.7	3.4	2.3	1.6	1.7	1.3
effective spread	bps	7.6	7.4	8.2	6.7	4.5	4.7	4.6
realized spread	bps	-1.5	-0.7	-1.4	-2.2	-2.2	-2.2	-2.3
price impact	bps	9.1	8.0	9.6	8.9	6.7	6.9	6.9
dollar-volume	dollar (million)	83.6	102.8	136.1	104.7	88.9	84.3	68.6
share-volume	million	2.0	2.4	3.4	3.4	2.4	2.3	2.2
transactions	thousands	5.0	7.0	13.3	13.6	11.3	11.6	10.5
imbalance	percent	14.3	13.2	11.7	10.9	11.0	11.1	11.0
total messages	thousands	74.8	124.3	336.4	496.9	465.2	364.8	214.8
algo trading	-\$100 per msg)	-19.4	-15.4	-8.1	-3.9	-3.0	-2.9	-3.5
%retail of \$volume	percent	14.5	14.2	12.0	14.3	12.9	12.2	10.5
% retail of LO volume	percent	12.3	12.6	10.8	13.6	12.0	11.9	10.4
<i>Panel B: Retail trader statistics</i>								
LO vol/(LO vol+MO vol)	percent	41.5	43.2	42.9	46.5	45.1	48.8	49.3
LO vol traded/(LO vol submitted)	percent	31.9	38.0	39.6	31.6	33.0	35.4	34.8
%return MO	bps	-2.6	-3.3	-4.6	-2.4	-1.4	-1.7	-1.3
%return LO	bps	-1.6	-0.5	-1.3	-0.8	1.5	0.5	1.7
\$return MO	cent	-0.1	-0.2	-0.4	-0.2	-0.1	-0.1	0.0
\$return LO	cent	-0.3	-0.2	-0.6	-0.5	-0.3	-0.3	-0.2
effective spread MO	bps	9.8	9.8	11.9	8.8	5.2	5.7	5.2
effective spread LO	bps	8.3	7.8	8.7	7.3	4.6	4.8	4.8
realized spread MO	bps	5.0	5.5	6.4	3.5	1.4	1.9	1.1
realized spread LO	bps	-3.5	-3.0	-6.0	-5.6	-5.0	-5.4	-5.0
price impact MO	bps	4.9	4.3	5.5	5.3	3.9	3.9	4.1
price impact LO	bps	11.8	10.9	14.8	13.0	9.6	10.2	9.8

Table II
Correlation Table for Covariates

	algo trading	imbalance	price change
algo trading	1.00		
imbalance	-0.16***	1.00	
price change	0.04***	-0.0**	1.00

Table III
T-Tests for Empirical Prediction 1

The table tests Empirical Prediction 1 which asserts that retail traders are uninformed. We test the hypothesis based on (a) whether %returns and \$returns are positive (split by market and limit orders; columns 1,2, 4, and 5), (b) realized spreads paid are positive (column 3), (c) realized spreads received are negative (column 6), and (d) whether the market-wide price impact exceed the price impact from retail market-orders (column 7). Standard errors are double-clustered by firm and time. * indicates significance of non-zero correlation at the 10% level, ** at the 5% level, and *** at the 1% level. T-statistics are in brackets.

	market orders			limit orders			price impact marketwise – retail
	%return	\$return	rsread paid	%return	\$return	rsread received	
<i>Panel A: Full sample</i>							
mean effect	-2.46*** [-9.43]	-0.17*** [-5.97]	3.45*** [22.86]	-0.00 [-0.01]	-0.33*** [-5.42]	-4.89*** [-20.06]	3.41*** [20.49]
<i>Panel B: Sample split by half-year</i>							
2006H2	-2.57*** [-4.83]	-0.11*** [-2.77]	4.96*** [11.48]	-1.58* [-1.69]	-0.25*** [-3.83]	-3.47*** [-7.36]	4.20*** [14.68]
2007H1	-3.52*** [-6.93]	-0.22*** [-4.50]	4.91*** [14.87]	-1.17 [-1.48]	-0.30*** [-3.30]	-2.43*** [-6.12]	3.60*** [15.97]
2007H2	-3.07*** [-4.87]	-0.19*** [-2.90]	6.08*** [19.45]	0.30 [0.25]	-0.14 [-1.15]	-3.67*** [-8.42]	3.84*** [16.53]
2008H1	-3.90*** [-6.64]	-0.32*** [-3.87]	5.50*** [21.48]	-0.36 [-0.33]	-0.41** [-2.36]	-4.55*** [-12.31]	3.52*** [17.75]
2008H2	-5.34*** [-3.63]	-0.52*** [-3.44]	7.36*** [18.60]	-2.28 [-0.98]	-0.82*** [-3.29]	-7.66*** [-12.23]	4.74*** [14.99]
2009H1	-2.81*** [-3.23]	-0.31*** [-3.38]	4.37*** [16.08]	-1.22 [-0.58]	-0.60*** [-2.84]	-6.59*** [-11.12]	4.06*** [16.15]
2009H2	-2.01*** [-4.31]	-0.19*** [-3.91]	2.61*** [12.55]	-0.32 [-0.26]	-0.45*** [-3.30]	-4.60*** [-13.39]	3.06*** [17.11]
2010H1	-1.07** [-2.24]	-0.09** [-2.37]	1.45*** [9.06]	1.53 [1.48]	-0.22** [-2.09]	-5.05*** [-15.46]	2.75*** [17.20]
2010H2	-1.69*** [-5.27]	-0.12*** [-3.69]	1.31*** [7.97]	1.44 [1.53]	-0.28*** [-3.09]	-4.89*** [-17.17]	2.95*** [17.76]
2011H1	-1.59*** [-3.76]	-0.07** [-2.57]	1.57*** [7.43]	0.06 [0.07]	-0.28*** [-3.03]	-4.71*** [-14.98]	2.95*** [15.33]
2011H2	-1.85** [-2.47]	-0.06 [-1.23]	2.26*** [8.43]	0.89 [0.60]	-0.18* [-1.72]	-6.17*** [-13.69]	3.05*** [12.92]
2012H1	-1.33*** [-3.38]	-0.04** [-2.19]	1.12*** [7.14]	1.66 [1.60]	-0.14** [-2.24]	-5.01*** [-13.71]	2.82*** [11.63]
Observations	58,040	58,040	58,040	58,037	58,040	58,037	58,040

Table IV
T-Tests for Empirical Prediction 2

The table tests empirical prediction 2 which asserts that for retail traders, limit orders should be more attractive than market orders. We test the hypothesis based on whether (a) the difference of %returns and \$returns for limit and market orders (columns 1 and 2) is positive and (b) whether the sum of realized spreads paid and received is positive (column 3). Standard errors are double-clustered by firm and time. * indicates significance of non-zero correlation at the 10% level, ** at the 5% level, and *** at the 1% level. T-statistics are in brackets.

	difference returns LO – MO		sum of rspread paid and received
	%return	\$return	
<i>Panel A: Full sample</i>			
mean effect full sample	2.46*** [5.49]	-0.15*** [-3.57]	-1.44*** [-6.83]
<i>Panel B: Sample split by half-year</i>			
2006H2	0.99 [1.00]	-0.14** [-2.41]	1.49** [2.18]
2007H1	2.35*** [3.15]	-0.08 [-1.14]	2.48*** [4.55]
2007H2	3.37*** [2.97]	0.05 [0.46]	2.42*** [4.75]
2008H1	3.54*** [3.58]	-0.1 [-0.66]	0.94** [2.35]
2008H2	3.06 [1.29]	-0.31 [-1.36]	-0.3 [-0.46]
2009H1	1.59 [0.84]	-0.29* [-1.65]	-2.22*** [-3.32]
2009H2	1.69 [1.49]	-0.26** [-2.26]	-1.99*** [-5.17]
2010H1	2.60*** [2.70]	-0.13 [-1.56]	-3.60*** [-9.76]
2010H2	3.13*** [3.69]	-0.16** [-2.35]	-3.57*** [-11.82]
2011H1	1.65* [1.83]	-0.21** [-2.50]	-3.14*** [-8.31]
2011H2	2.74* [1.74]	-0.12 [-1.20]	-3.91*** [-8.13]
2012H1	2.99*** [2.84]	-0.1 [-1.61]	-3.89*** [-9.97]
Observations	58037	58040	58037

Table V
Regressions for Determinants of Intraday Return

This table test Empirical Prediction 3 on the determinants of intraday returns. We consider three variables: the Hendershott, Jones, and Menkveld (2011) measure of algorithmic trading, defined as the negative of the \$-volume per message (trades and order submissions, cancellations and modifications); the absolute value of the order imbalance, scaled by the daily volume; and the absolute value of the price change from close to close, scaled by the preceding days' closing price. All variables are per stock per day for our sample of 48 securities. Returns are defined in term of \$s and in terms of % of the aggregate daily \$volume. The estimated equation is

$$\text{dependent variable}_{it} = \alpha_{(i)} + \beta X_{it} + \epsilon_{it}.$$

$\alpha_{(i)}$ are the stock fixed effects and β are the coefficient of interest for the covariate vector X_{it} . * indicates significance of non-zero correlation at the 10% level, ** at the 5% level, and *** at the 1% level. Standard errors are in parentheses; they are double-clustered by firm and time.

<i>Panel A: Intraday returns from market orders</i>								
	retail traders				all traders			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
algo trading	0.03 (0.02)			0.04** (0.02)	-0.04** (0.01)			0.02 (0.01)
imbalance		0.06*** (0.02)		0.07*** (0.02)		0.38*** (0.03)		0.37*** (0.03)
price change			-0.01*** (0.00)	-0.01*** (0.00)			0.01*** (0.00)	0.01*** (0.00)
Observations	58,040	57,997	57,992	57,949	58,040	57,997	57,992	57,949
Adjusted R-squared	0.002	0.003	0.014	0.014	0.003	0.070	0.047	0.111
<i>Panel B: Intraday returns from limit orders</i>								
	retail traders				all traders			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
algo trading	0.18*** (0.04)			0.12*** (0.04)	0.04** (0.01)			-0.02 (0.01)
imbalance		-0.56*** (0.06)		-0.49*** (0.06)		-0.38*** (0.03)		-0.37*** (0.03)
price change			-0.09*** (0.01)	-0.09*** (0.01)			-0.01*** (0.00)	-0.01*** (0.00)
Observations	58,037	57,994	57,989	57,946	58,040	57,997	57,992	57,949
Adjusted R-squared	0.004	0.016	0.160	0.172	0.003	0.070	0.047	0.111

Panel C: Intraday raw payoff from market orders

	retail traders				all traders			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
algo trading	0.00			0.00	-0.03**			-0.00
	(0.00)			(0.00)	(0.01)			(0.01)
imbalance		0.00**		0.00***		0.19***		0.19***
		(0.00)		(0.00)		(0.03)		(0.03)
price change			-0.00***	-0.00***			0.01***	0.01***
			(0.00)	(0.00)			(0.00)	(0.00)
Observations	58,040	57,997	57,992	57,949	58,040	57,997	57,992	57,949
Adjusted R-squared	0.004	0.004	0.026	0.026	0.002	0.024	0.028	0.049

Panel D: Intraday raw payoff from limit orders

	retail traders				all traders			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
algo trading	0.01**			0.00	0.03**			0.00
	(0.00)			(0.00)	(0.01)			(0.01)
imbalance		-0.05***		-0.05***		-0.19***		-0.19***
		(0.01)		(0.01)		(0.03)		(0.03)
price change			-0.01***	-0.01***			-0.01***	-0.01***
			(0.00)	(0.00)			(0.00)	(0.00)
Observations	58,040	57,997	57,992	57,949	58,040	57,997	57,992	57,949
Adjusted R-squared	0.005	0.014	0.137	0.145	0.002	0.024	0.028	0.049

Table VI
Regressions for Determinants of Realized Spreads

This table test Empirical Prediction 3 on the determinants of realized returns (used as a proxy for trading costs and benefits). The table used the same covariates as Table V. All variables are per stock per day for our sample of 48 securities. The realized spread is computed as the signed difference of the trade price and the midpoint five minutes subsequent to the trade, scaled by the midpoint at the time of the trade. We compute it for retail trades with market and limit orders and for the entire market. The estimated equation is

$$\text{dependent variable}_{it} = \alpha_{(i)} + \beta X_{it} + \epsilon_{it}.$$

$\alpha_{(i)}$ are the stock fixed effects and β are the coefficient of interest for the covariate vector X_{it} . * indicates significance of non-zero correlation at the 10% level, ** at the 5% level, and *** at the 1% level. Standard errors are in parentheses; they are double-clustered by firm and time.

<i>Panel A: Realized spread paid by retail traders for market orders</i>				
	(1)	(2)	(3)	(4)
algo trading	-0.09*** (0.02)			-0.09*** (0.02)
imbalance		0.01** (0.01)		0.00 (0.00)
price change			0.00*** (0.00)	0.00*** (0.00)
Observations	58,040	57,997	57,992	57,949
Adjusted R-squared	0.027	0.018	0.027	0.037
<i>Panel B: Realized spread received by retail traders for limit orders</i>				
	(1)	(2)	(3)	(4)
algo trading	-0.06*** (0.01)			-0.06*** (0.01)
imbalance		0.03*** (0.01)		0.03*** (0.01)
price change			-0.01*** (0.00)	-0.01*** (0.00)
Observations	58,037	57,994	57,989	57,946
Adjusted R-squared	0.020	0.019	0.063	0.065
<i>Panel C: Realized spread marketwide</i>				
	(1)	(2)	(3)	(4)
algo trading	-0.05*** (0.01)			-0.05*** (0.01)
imbalance		0.00 (0.01)		-0.00 (0.00)
price change			-0.00*** (0.00)	-0.00*** (0.00)
Observations	58,040	57,997	57,992	57,949
Adjusted R-squared	0.043	0.034	0.049	0.058

Table VII
Regressions for Determinants of the Price Impact, Effective Spread and Execution Probability

This table test Empirical Prediction 3 on the determinants of the 5-minute price impact (market-wide and for limit order trades by retail) (we use the price impact as a proxy for asymmetric information), the effective spread and the probability of execution. The table used the same covariates as Table V. All variables are per stock per day for our sample of 48 securities. The price spread is computed as the signed difference of the midpoint of the bid-ask spread at the time of the trade and the midpoint five minutes subsequent to the trade, scaled by the midpoint at the time of the trade. The execution probability is measured by the ratio of the volume traded with limit orders to the total volume of submitted orders, by retail traders, The estimated equation is

$$\text{dependent variable}_{it} = \alpha_{(i)} + \beta X_{it} + \epsilon_{it}.$$

$\alpha_{(i)}$ are the stock fixed effects and β are the coefficient of interest for the covariate vector X_{it} . * indicates significance of non-zero correlation at the 10% level, ** at the 5% level, and *** at the 1% level. Standard errors are in parentheses; they are double-clustered by firm and time.

	Price impact – marketwide				Price impact – LO retail			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
algo trading	-0.01 (0.02)			-0.01 (0.02)	-0.01 (0.02)			-0.02 (0.02)
imbalance		0.01 (0.01)		0.00 (0.00)		-0.02* (0.01)		-0.03*** (0.01)
price change			0.01*** (0.00)	0.01*** (0.00)			0.02*** (0.00)	0.02*** (0.00)
Observations	58,040	57,997	57,992	57,949	58,037	57,994	57,989	57,946
Adjusted R-squared	0.376	0.375	0.415	0.415	0.111	0.111	0.173	0.174
	effective spreads – marketwide				% limit order volume that trades			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
algo trading	-0.06*** (0.01)			-0.06*** (0.01)	-0.16*** (0.05)			-0.13** (0.05)
imbalance		0.01*** (0.00)		0.00 (0.00)		0.19*** (0.02)		0.16*** (0.02)
price change			0.00*** (0.00)	0.00*** (0.00)			0.01*** (0.00)	0.01*** (0.00)
Observations	58,040	57,997	57,992	57,949	56,134	56,092	56,088	56,046
Adjusted R-squared	0.579	0.568	0.593	0.605	0.059	0.063	0.067	0.077

Figure 1
Trader Classification

The figure displays the average number of all traders and of unsophisticated retail traders that we observe in our sample, per stock per week.

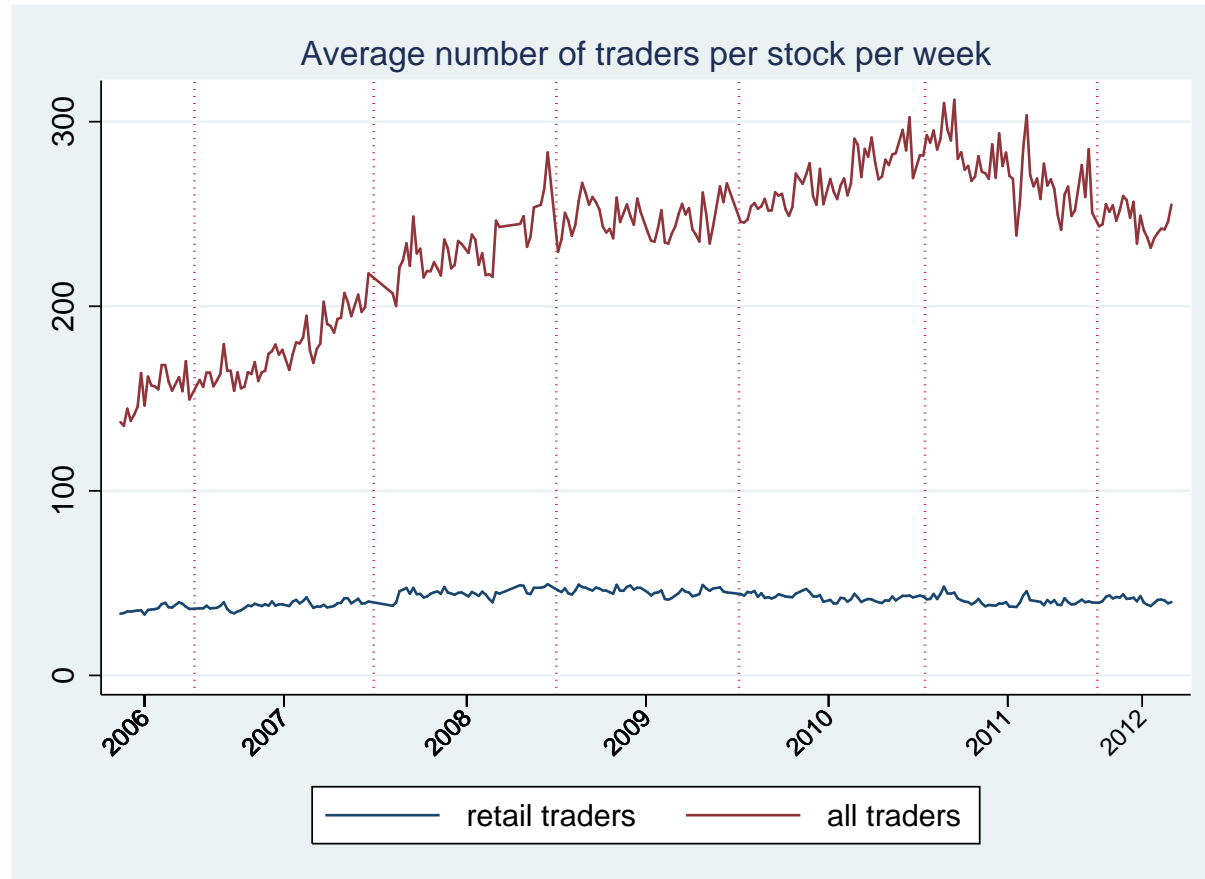

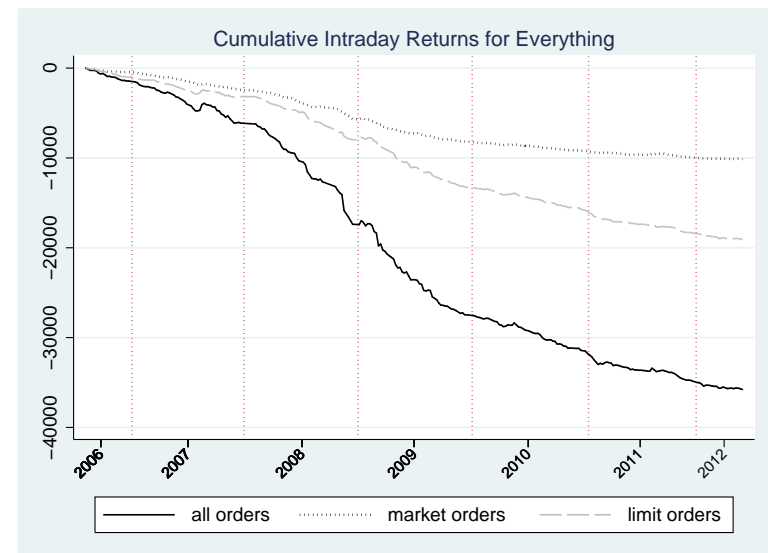


Figure 2
Cumulative Intraday Returns for Retail Traders

The left panel plots the cumulative intraday returns for all traders split by market and limit orders, the right panel plots cumulative returns for retail traders. We sum over all 48 stocks in our sample and all retail traders per week; the scale is in 10,000s of dollars, i.e. by 2012, cumulative losses by retail traders amounted to \$400million.



Panel A: % returns



Panel B: \$ returns

Figure 3
Realized Spreads Paid and Received by Unsophisticated Traders




Figure 4
Density of %returns to market and limit orders for retail traders— per day per stock

The figure plots the density of the (winsorized) per day per stock intraday %returns across the entire panel. Returns are defined in (10) (in terms of % of the daily \$volume); we the measure is scaled to basis points.

