

The Limit Order Effect*

Juhani Linnainmaa[†]

The Anderson School at UCLA

November 2005

*I am truly indebted to Mark Grinblatt for his guidance and encouragement. I also thank Tony Bernardo, Michael Brennan, John Campbell, Eric Hughson (AFA discussant), Matti Keloharju, Hanno Lustig, Tuomo Vuolteenaho, and seminar participants at the American Finance Association 2004 Meetings, Harvard University, and Helsinki School of Economics for their comments. An earlier version of this paper was circulated under the title “Who Makes the Limit Order Book? Implications for Contrarian Strategies, Attention-Grabbing Hypothesis, and the Disposition Effect.” I gratefully acknowledge the financial support from the Allstate Dissertation Fellowship and the Foundation of Emil Aaltonen. Henri Bergström from the Finnish Central Securities Depository and Pekka Peiponen, Timo Kaski, and Jani Ahola from the Helsinki Exchanges provided data used in this study. Special thanks are due to Matti Keloharju for his help with the FCSD registry. All errors are mine.

[†]**Correspondence Information:** Juhani Linnainmaa, The Anderson School at UCLA, Suite C4.01, 110 Westwood Plaza, CA 90095-1481, tel: (310) 909-4666, fax: (310) 206-5455, <http://personal.anderson.ucla.edu/juhani.linnainmaa>, email: julinnai@anderson.ucla.edu.

The Limit Order Effect

Abstract

The limit order effect is the appearance that limit order traders react quickly and incorrectly to new information. This paper combines investor trading records with limit order data to examine the importance of this effect. We show that institutions earn large trading profits by triggering households' stale limit orders and that individuals' passivity significantly affects inferences about their behavior. Individuals are net buyers on days when prices fall because institutions unload shares to households with market orders—and vice versa on days when prices rise. An analysis of earnings announcements shows that institutions react to announcements, triggering individuals' limit orders: the orders executed during the first two minutes lose an average of -2.5% on the same day. Investors' use of limit orders may help to understand many findings in the extant literature, such as individuals' seemingly coordinated tendency to trade against short-term returns.

1 Introduction

Numerous studies find that individual investors' behavior is very sensitive to news and short-term returns. For example, Odean (1998) and Grinblatt and Keloharju (2000) find that individuals sell when prices rise and buy when prices fall. Hirshleifer, Myers, Myers, and Teoh (2003) show that individuals trade against earnings surprises, selling companies that release better-than-expected earnings and vice versa. Barber and Odean (2002) and Seasholes and Wu (2005) report that individuals trade stocks that grab their attention—i.e., stocks that release news or experience significant price movements. Moreover, many studies find that individuals trade systematically in the wrong direction.¹ Collectively, the evidence suggests that individuals monitor the market closely but systematically misinterpret new information, losing to other investors. This description of individual investors is puzzling. First, we would expect individuals to be trend-followers, not contrarians: uninformed investors' beliefs are, by definition, more sensitive to news, inducing them to trade more in the direction of the news (Brennan and Cao 1997).² Second, it seems inexplicable how uninformed investors could *systematically* lose money in their trades (Fama 1970).

This paper shows that individual investors' use of limit orders can explain *why* individuals seem to misinterpret new information and systematically trade in the “wrong direction”. The following example demonstrates how limit orders affect inferences about investor behavior (see Figure 1).³ Suppose there is no disagreement about a stock's fair value (V_0) so that investors trade only for liquidity reasons. Those willing to wait place limit orders, hoping that an impatient investor arrives and trades against them.⁴ Now, suppose that the company

¹For example, Odean (1999, pp. 1296) concludes: “*What is more certain is that these [individual] investors do have useful information which they are somehow misinterpreting.*”

²Suppose that there are two investors who initially have the same level of beliefs about the stock's intrinsic value but one (the uninformed) is more uncertain about the value than the other (the informed). Now, if the company releases negative news, the uninformed investor revises her beliefs downwards more. Hence, after the announcement, the uninformed investor must be selling shares to the informed investor until their beliefs converge. See Brennan and Cao (1997) and Brennan, Cao, Strong, and Xu (2005) for a detailed analysis.

³We develop a limit order model in Appendix A to formalize this example. We use the model to show that the limit order characteristics discussed here—for example, that more limit orders execute when there is a shock to the fundamentals—are *equilibrium* characteristics and not ad hoc intuition.

⁴Even when there may be privately informed investors in the market, an investor trading for liquidity reasons is likely to prefer a limit order to a market order. See, e.g., Glosten (1994) and Handa and Schwartz (1996) for reasons why uninformed traders should use limit orders. Kyle (1985)-type of models with multiple informed agents, such as Holden and Subrahmanyam (1992), address the market order side. They predict that any private information is rapidly revealed through informed agents' aggressive competition. Bloomfield, O'Hara, and Saar (2005) is one of the dissidents in the limit versus market order choice literature. The paper finds that informed traders use more *limit* orders than liquidity traders in a laboratory experiment. However, even

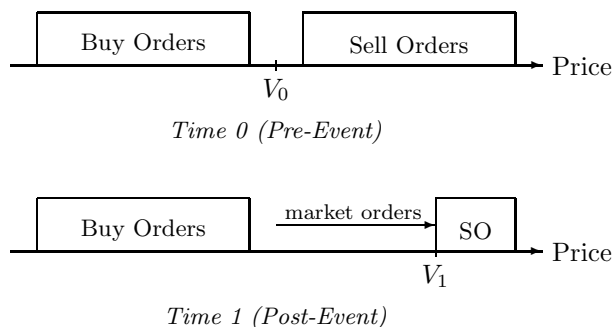


Figure 1: **An Example of Limit Orders Triggered by News**

unexpectedly revises its growth estimates upwards. Those investors who observe this news in real-time react by submitting market buy orders, triggering sell limit orders that are not withdrawn in time. All the sell limit orders in the book up to the new valuation V_1 execute and lose money.⁵ Thus, the arrival of news creates an appearance that the liquidity traders are reacting to the news and losing money because of their incorrect decisions. Moreover, more limit orders execute during this brief event compared to a normal “no-news period” where investors trade only for liquidity reasons. Hence, the limit order traders are more active and tend to trade in the wrong direction when something happens in the market.

This paper uses unique data set that combines investor trading records with limit order data to examine whether this mechanism—**the limit order effect**—significantly affects inferences about investor behavior. Limit orders are potentially important. The investor class differences in, e.g., trading strategies and the disposition effect may partly arise because individuals are passive investors, trading mostly for liquidity reasons and being willing to wait for execution. If so, individuals’ poor timing may simply mirror institutions’ great timing: institutions observe and react to news, triggering individuals’ limit orders.

We begin by demonstrating two strong patterns in individuals and institutions’ trading behavior: (1) individuals use significantly more limit orders and (2) institutions earn large trading gains from households’ limit orders. Individuals’ use of limit orders may be suboptimal: individuals often place their orders far away from the quoted spread and these orders, when

if investors with *private information* prefer limit orders, investors competing to gain from a release of public information must employ market orders.

⁵This is the adverse selection risk, first described by Bagehot (1971) and formalized by Copeland and Galai (1983), Glosten and Milgrom (1985), and others. Note also that those investors who had submitted buy limit orders now need to pay more than before for their shares.

triggered, generate large losses. Limit orders are a mechanism for a significant transfer of wealth from households to institutions: individuals' limit orders lose 74.7 million euros (or 169 euros per individual investor) in our sample while their market orders lose only 17.7 million euros (or 40 euros per investor).

We then study how limit orders affect inferences about investor behavior. First, we examine whether limit orders explain individual investors' attention-grabbing behavior—i.e., individuals appear to trade stocks that grab their attention. For example, Barber and Odean (2002) and Seasholes and Wu (2005) suggest that the difficulty of searching through all potential stocks may generate such behavior. In contrast to this cognitive explanation, we find that most of the effect arises from institutions' use of market orders against individuals' limit orders. Individuals' market order buy-sell imbalance—their “active reaction”—responds only modestly to what is happening in the market.

Second, we use intraday data on earnings announcements to show that institutions react to news with a flood of market orders, triggering households' limit orders. Households' limit orders suffer significant losses: the average same-day return is -2.5% for the orders executed during the first two minutes after the announcement. This result may explain why Hirshleifer, Myers, Myers, and Teoh (2003) find that “*individuals... tend to make contrarian trades in opposition to the direction of earnings surprises.*” The finding that limit orders lose money is important also because there is no mechanical trading strategy to exploit this phenomenon: the gains go to investors who analyze news quickly and react before the others do. The limit order effect may be a simple yet powerful explanation for many findings about individual investors' behavior.

The rest of the paper is organized as follows. Section 2 discusses the data set. Section 3 examines differences in institutions' and households' order strategies and trading gains. Section 4 studies individuals' attention-grabbing behavior and Section 5 analyzes limit orders around earnings announcements. Section 6 relates our results to the existing literature. Section 7 concludes.

2 Data

This section discusses the rules of the Helsinki Exchanges and the Finnish market during the sample period from September 1998 to October 2001. We also discuss the contents of our data set, categorize order types, and explain how we match trading records with the limit order data.

2.1 Helsinki Exchanges

Trading on the Helsinki Exchanges (HEX) is divided into sessions. Each trading day starts at 10:10 am with an opening call. Orders that are not executed at the opening remain on the book and form the basis for the continuous trading session. This trading session takes place between 10:30 am and 5:30 pm in a fully automated limit order book, the automated trading and information system (HETI). After-hours trading (5:30 – 5:45 pm) takes place after the continuous trading session and again the next morning (9:30 – 10:00 am) before the next opening call. (Two changes to the trading schedule were made during the sample period. On August 31, 2000, the regular trading session was extended to 6:00 pm and the after-hours session was moved to match this change. On April 10, 2001, an evening session that extended trading hours to 9:00 pm was introduced.)

The HEX trading system displays the five best price levels of the book to the market participants on both sides. The public can view this book in a market-by-price form while financial institutions receive market-by-order feed.⁶ Simple rules govern trading on the limit order book. There are no designated market makers or specialists; the market is completely order-driven. An investor trades by submitting limit orders. The minimum tick size is €0.01. An investor who wants immediate execution must place the order at the best price level on the opposite side of the book. An investor who wants to buy or sell more shares than what is currently outstanding at the best price level must “walk up or down the book” by submitting separate orders for each price level. If a limit order executes against a smaller order, the unfilled portion stays on the book as a new order. Time and price priority between limit orders is enforced. For example, if an investor submits a buy order at a price level that already has

⁶A market-by-price book displays the five levels on both sides of the market but only indicates the total number of shares outstanding at each price level. A market-by-order book shows each order separately and also shows which broker/dealer submitted each order.

other buy orders outstanding, all the old orders must execute before the new order.

The total market value of the 158 companies in the Helsinki Exchanges was €383 billion in the middle of the sample period (May 2000). We report several sample statistics for the 30 most actively trade stocks for future reference:

- A total of 14.2 million trades took place in these stocks. The most active stock is Nokia with 2.7 million trades.
- These stocks have an average realized log-spread of 0.44%. (Nokia’s average spread is the lowest, 0.13%.)
- Three out of ten trades originate from households. Households’ participation—as measured by the proportion of household trades—ranges from a low of 15% to a high of 72% in these 30 stocks.

2.2 Investor Trading Records and Limit Order Data

We use the following data sets in this study:

1. *The complete trading records and holdings information of all Finnish investors.* The Finnish Central Securities Depository registry (FCSD) provided us these data for the period from January 1995 to November 2002. Each trade record includes a date-stamp, a stock identifier, and the price, volume, and direction of the trade. Each record also identifies the investor type—a domestic institution, a domestic household, or a foreigner—and gives other demographic information. We classify all investors as either individuals or institutions (including foreigners) for this study. Grinblatt and Keloharju (2000) give the full details of this data set.
2. *The limit order data for all HEX stocks.* These data are the supervisory files from the HEX from September 18, 1998 to October 23, 2001. Each entry is a single order entered into the trading system, containing a unique order identifier, date- and time-stamps, a session code, a code for the brokerage firm submitting the order, a trade type indicator (i.e., upstairs/downstairs/odd-lot), and the price, volume, and direction of the order. All entries also contain a set of codes for tracking the life of an order—an order can expire, be partly or completely filled, or modified. We use these data to reconstruct the




Figure 2: **An Example of Classifying a Buy Order**

limit order book for each second of every trading day for all the stocks. Data before July 10, 2001 is missing the time-stamp that identifies when an unfilled order is withdrawn.⁷

2.2.1 Order Classification

We use the following categorization to classify every order entered into the trading system:

- *Market order.* An order placed at the best price level on the opposite side of the book to get immediate execution.⁸
- *Inside the Spread Limit Order.* An order placed inside the current bid-ask spread.
- *At the Spread Limit Order.* An order placed at the best bid (for a buy order) or at the best ask (for a sell order).
- *Outside the Spread Limit Order.* An order placed outside the current bid-ask spread.
- *Pre-Open Limit Order.* A limit order entered into the system before the continuous trading session begins (the limit order book is not viewable to any market participants at this time). *Stale limit orders* is a subset of pre-open orders: these are unfilled orders carried over from the previous trading day.⁹

Figure 2 shows an example of how a buy order is classified when best bid is currently \$58 and the best ask is \$63. For example, an order with a price between \$58.01 and \$62.99 is classified as an inside-the-spread order.

⁷This code is important for the analysis of the composition of the limit order book. Hence, we only use the latter part of the data in Section 5's composition analysis.

⁸All orders are technically limit orders; however, many brokers use the term *market order* for these active orders. We adopt the same terminology because there is no risk of confusion.

⁹A pre-open limit order with the same price, volume, and broker as an order that expired at the end of the previous trading day is defined as a stale limit order. We use this subcategory only in Table 1.

2.2.2 Matching the Data Sets

We match the investor trading records against the limit order data using executed trades to obtain information on, e.g., what type of orders different investors use and when trades take place. Each trade record in the limit order data contains all the same information as the investor trading records *except* the investor identity. We use common elements to link the data sets.

Matching Executed Trades

There is no ambiguity in matching two types of trades: trades with unique price-volume combinations and non-unique trades that must originate from the same investor.¹⁰ We call these trades *uniquely matched trades* and we use only these trades in most of our analysis. There is no one-to-one link between the data sets for the remaining trades. We match these trades in two steps, using brokerage firm identities to improve the match. In the first step, we use the uniquely matched trades to determine each investor’s broker. In the second step, we match the non-unique trades so that, when possible, an investor’s trade ends up originating from the preferred broker. The matching of the non-unique trades can generate errors. However, a matching error where investors A and B switch places only matters if one is an institution and the other is a household. Moreover, these errors make all our tests more conservative: because this process adds noise, the true differences in trading behavior are *even stronger* than what we find in our tests.

Predicted Identities for Unfilled Orders

We use data on all unique, limit order-initiated trades to identify whether an unfilled limit order originates from a household or from an institution. Our method is straightforward: we use information from executed limit orders to back out what type of orders most likely originate from households. We proceed as follows. First, we estimate a logistic regression with a household indicator variable as the explanatory variable. We include an intercept and the

¹⁰We say that a trade has a *unique* price-volume combination if, for example, there is only one trade (in one stock-day) with a price of €82 and a volume of 1,200 shares. A trade is non-unique if, say, three trades have the same price-volume combination. In this example, “all must originate from the same investor” would mean that a single investor in the investor data set is the buyer or the seller in all the three trades.

following explanatory variables on the right-hand side of the model:¹¹

- Brokerage firm identities (25 dummies for the largest brokers)
- The first three powers of the log-trade size
- Stock dummies for the 15 most actively traded stocks
- The cross-products of log-trade size variables and broker identities
- Dummies for limit order type

Second, we use this calibrated model to obtain the predicted household-probability for every limit order and then round these probabilities into *predicted identities*. Each limit order in the data now contains an indicator set to one if the order most likely originates from a household. We use these predicted identities to study *changes* in the composition of the limit order book around earnings announcements. Hence, as long as the identification errors do not systematically vary around earnings announcements, any errors make our estimates conservative.

3 Institutions versus Households: Order Strategies and Trading Gains

This section analyzes differences in institutional and household order strategies and trading performance.¹²

¹¹Linnainmaa (2003) uses the same data set as we do and focuses on how accurately trade characteristics reveal whether a trade is from a household or from an institution. For example, the broker identities are important: the fraction of households-as-customers ranges from a high of 83.9% to a low of 0.4%. We evaluate the fit of our model by classifying all in-sample trades as originating from either an individual or an institution. We find that the logistic regression correctly classifies 87% of the trades. We also estimate the regression separately for several “random 10% of observations” sub-samples and find similar coefficients across these sub-samples.

¹²The extant literature suggests that institutions are more informed than individuals. We take this as our hypothesis. For example, Grinblatt and Keloharju (2000) find this using non-overlapping data from the Finnish market. Barber, Lee, Liu, and Odean (2005) find that institutions gain from households in a study of the households versus institutions trade in Taiwan.

3.1 Order Placement Strategies

3.1.1 Methodology

We analyze order strategy differences in two ways. First, we compute frequencies of different trade types for all executed trades. Second, we compute the initial distance of an (executed) limit order from the opposite side of the book as

$$\text{limit distance}_{i,s,t} = \begin{cases} \ln(a_i) - \ln(p_{i,s,t}) & \text{if a buy order} \\ \ln(p_{i,s,t}) - \ln(b_i) & \text{if a sell order} \end{cases} \quad (1)$$

where $p_{i,s,t}$ = trade price of executed limit order

b_i, a_i = the best bid / ask price at the time the limit order is entered.

(For example, if the best ask is currently \$30, a buy order placed at \$29 has an initial distance of 3.4%.) We break this distance measure into intervals and compute what fraction of executed limit orders in each interval originates from households. We compute this fraction separately for each stock-day and then compute averages across all stock-day observations.

We use data on all trades in the 30 most active stocks throughout Section 3. We restrict the sample to the uniquely matched trades and also exclude pre-negotiated block trades (these take place outside the limit order book).

3.1.2 Results

Table 1 shows that most (56.1%) of households' trades originate from limit orders, whereas this proportion is 48.8% for institutions. Moreover, there are significant differences in *what kind* of limit orders institutions and individuals use. Even when institutions use limit orders, they place their orders close to the opposite side of the book—for example, more than half of institutions' limit orders improve the existing spread. Households, on the other hand, often place their limit orders outside the spread.¹³

Figure 3 shows that households place their limit orders farther outside the spread: whereas

¹³The frequencies in Table 1 do not show what fraction of *all* limit orders originate from households because every market order, but not every limit order, executes. The fact that households place their limit orders deeper into the book suggests that this true proportion is higher than what it is for executed trades. Section 5 uses the calibrated logistic model to classify all the unfilled orders in the limit order book as originating from either a households or an institution. The average proportion of limit orders (across announcements) from households is 73.7% for the 30-minute period before announcements.

Table 1: Institutions’ and Households’ Use of Market and Limit Orders

This table shows how institutions and individuals use market and limit orders. The sample consists of all uniquely matched trades (see text) in the 30 most actively traded stocks on the Helsinki Exchanges between September 18, 1998 to October 23, 2001. This table reports the frequencies of different order types for executed trades. *Market Order* is an order that is placed on the opposite side of the book so that it executes immediately. The limit order types are defined as follows: (i) *inside* is an order placed inside the current bid-ask spread; (ii) *at* is an order placed inside the bid-ask spread; (iii) *outside* is an order placed outside the current bid-ask spread; (iv) *pre-open* is an order entered into the system before the trading day begins; (v) *stale* is an order carried over from the previous trading day. Stale limit orders are not (double-)counted as *pre-open* orders. Note that (1) “Inside” + ... + “Stale” = “Limit Order” and (2) “Limit Order” + “Market Order” = 100%. N is the number of trades.

Order Type	Households			Institutions		
	Buy	Sell	Both	Buy	Sell	Both
Limit Order	55.7%	56.4%	56.1%	47.0%	50.6%	48.8%
Inside	19.3%	21.3%	20.2%	24.7%	27.0%	25.9%
At	7.2%	7.0%	7.1%	10.6%	11.6%	11.1%
Outside	16.4%	16.0%	16.2%	7.9%	7.8%	7.9%
Pre-Open	10.2%	9.4%	9.8%	3.4%	3.8%	3.6%
Stale	2.6%	2.7%	2.6%	0.3%	0.3%	0.3%
Market Order	44.3%	43.6%	43.9%	53.0%	49.4%	51.2%
N	689,846	643,947	1,333,793	2,299,997	2,400,234	4,700,231

approximately 25% of limit orders placed close to the opposite side of the book originate from households, this fraction increases to over 70% for orders placed more than 5% away from the opposite side.¹⁴ This figure strikingly shows that institutions and individuals use limit orders very differently: households have discretion in their trading whereas institutions’ limit orders are mostly substitutes to market orders.

3.2 Trading Gains

3.2.1 Methodology

We now study whether institutions make money by trading against households. We focus on two predictions. First, if institutions have the informational advantage, limit orders triggered

¹⁴In unreported work, we find that both individuals and institutions place their sell limit orders significantly closer to the opposite side of the book. This result may be due to differences in impatience: an investor is often forced to sell for liquidity reasons but is rarely constrained to buy. This finding may be important in understanding why many studies (e.g., Saar 2001) find asymmetries between purchases and sales.)

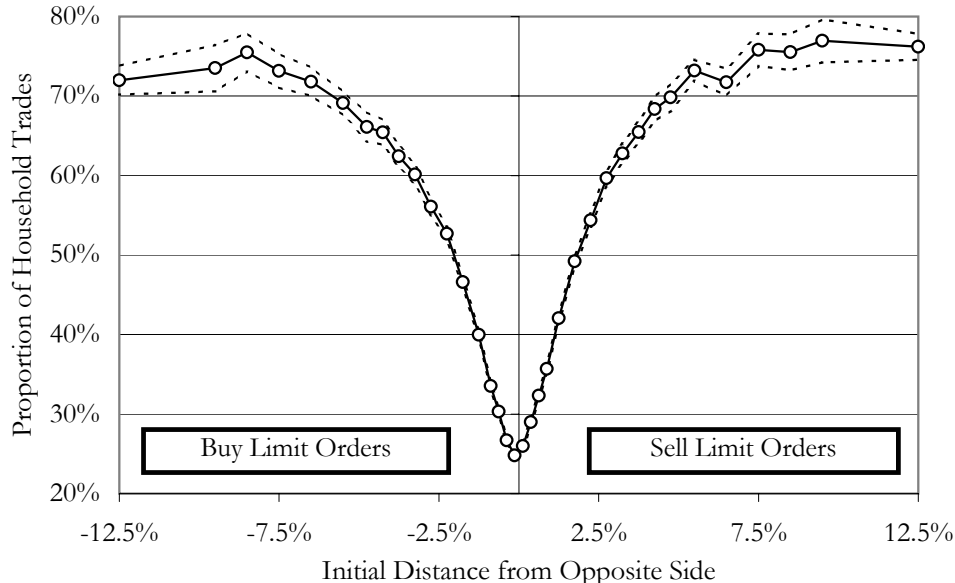


Figure 3: **The Proportion of Executed Household Limit Orders.** This figure shows the proportion of limit order-initiated trades originating from households, conditional on the initial distance from the opposite side of the book. An order’s log-distance from the opposite side of the limit order book is recorded at the time it is entered into the book. The fraction of household orders is computed separately for purchases (the left side of the figure) and sales (the right side). Each circle in the figure is a midpoint of an interval; the interval length increases from 0.25% (for orders close to the opposite side of the book) to 1% (for orders far from the opposite side). The farthest points (−12.5% and 12.5%) denote limit orders entered more than 10% away from the opposite side of the book. The proportion of household limit orders is computed for each stock/day/distance category. This figure shows the means and 95% confidence intervals across stock-days for each distance category. The sample consists of all uniquely matched trades (see text) in the 30 most actively traded stocks on the Helsinki Exchanges between September 18, 1998 to October 23, 2001.

by institutions generate trading losses while limit orders triggered by households generate gains. Second, limit orders placed farther outside the spread should have lower returns: more of these orders execute when there is a large information event. We study trading performance by computing the trading gain for each executed limit order:

$$r_{i,s,t} = \begin{cases} \ln(c_{s,t+k}) - \ln(p_{i,s,t}) & \text{if a buy order} \\ \ln(p_{i,s,t}) - \ln(c_{s,t+k}) & \text{if a sell order} \end{cases} \quad (2)$$

where $p_{i,s,t}$ = the trade price

$c_{s,t+k}$ = the closing price in stock s on date $t + k$.

Trading gains are ideal for analyzing informational differences: because our data are for the entire market, trading gains sum to zero for each stock-day, $\sum_i r_{i,s,t} \equiv 0$, by definition. This

identity allows us to track who loses and who gains from trade. We analyze trading gains in event-time, taking a single executed order as the unit of observation. This is the relevant measure for an investor submitting an order. We compute the average same-day and one-week trading gains for all executed limit orders conditional on (1) who places the limit order, (2) what is the limit order type, and (3) who triggers the limit order.

3.2.2 Results

Table 2 shows three distinct patterns: (1) limit orders placed farther outside the spread perform worse, (2) households perform worse than institutions with both limit and market orders, and (3) the performance of limit orders worsens when the horizon increases. We discuss each of these patterns in turn.

First, outside-the-spread limit orders perform worse than orders at or inside the spread: an executed inside-the-spread order has an average same-day gain of 0.08% whereas an outside-the-spread order loses 0.29%. (The pre-open orders lose even more.) Institutions' market orders cause these losses. For example, when an institution triggers an inside-the-spread order, the institution's market order returns 0.03%. However, when the institution triggers an outside-the-spread order, the market order gain is 0.18%. Note also that institutions trigger more of the outside-the-spread orders: the proportion of household limit orders triggered by institutions climbs from 57% ("inside the spread") to 73% ("outside the spread").

Second, households lose money on both limit and market orders. A household loses an average of 0.36% when an institution triggers the limit order but *gain* 0.16% when another household demands immediacy. Institutions' limit orders, on the other hand, have overall gains close to zero: 0.07% for the same-day and -0.02% for the one-week horizon. This dramatic difference shows up also on the market order side: institutions' market orders earn 0.09% on the same day while households' market orders lose 0.32%.¹⁵

Third, the performance of limit orders worsens with the horizon. This indicates that

¹⁵Our finding that institutions' limit orders lose when picked up by another institution but gain when picked up by a household mirrors the theoretical literature on the adverse selection component of the bid-ask spread. For example, Copeland and Galai (1983, pp. 1458) state the informed/uninformed trade-off as "*Because traders with special information have the option of not trading with the dealer, he will never gain from them. He can only lose. On the other hand the dealer gains in his transactions with liquidity-motivated traders.*" Glosten and Milgrom (1985, pp. 72) emphasize that "*... the specialist must recoup the losses suffered in trades with the well informed by gains in trades with liquidity traders.*"

Table 2: Trading Gains in the Trades between Institutions and Households

This table reports trading gains for institutions' and households' limit orders. The sample consists of all uniquely matched trades (see text) in the 30 most actively traded stocks on the Helsinki Exchanges between September 18, 1998 to October 23, 2001. The trading gain for each executed limit order is computed as

$$r_{i,s,t} = \begin{cases} \ln c_{s,t+k} - \ln p_{i,s,t} & \text{if a buy order} \\ \ln p_{i,s,t} - \ln c_{s,t+k} & \text{if a sell order} \end{cases}$$

where $p_{i,s,t}$ is the trade price and $c_{s,t+k}$ is the closing price in stock s on date $t+k$. Panel A reports the same-day trading gains ($k = 0$) and Panel B the one-week trading gains ($k = 5$). The trading gains are computed conditional on (1) who places the limit order (an institution or individual; this is reported in the first column), (2) what is the limit order type (*inside the spread limit order, . . . , pre-open limit order*; this is reported in the second column), and (3) who triggers the limit order (an institution or individual; this is reported at the top of the table). Column “→ Inst.” reports the proportion of limit orders that are picked up by institutions. The sample sizes are in Table 1.

Panel A: Same-Day Trading Gains

<i>Limit Order from</i>		→ Inst.	<i>Limit Order Triggered by</i>					
Order Type	Institution		Household		Total			
			Mean	s.e.	Mean	s.e.	Mean	s.e.
Institution	In	78%	0.01%	0.00%	0.45%	0.01%	0.11%	0.00%
	At	84%	0.00%	0.00%	0.45%	0.01%	0.07%	0.00%
	Out	84%	-0.08%	0.01%	0.22%	0.01%	-0.04%	0.00%
	Pre	67%	-0.29%	0.02%	0.25%	0.03%	-0.11%	0.02%
	All	80%	-0.02%	0.00%	0.41%	0.01%	0.07%	0.00%
Household	In	57%	-0.26%	0.01%	0.25%	0.01%	-0.04%	0.01%
	At	66%	-0.32%	0.01%	0.23%	0.02%	-0.13%	0.01%
	Out	73%	-0.37%	0.01%	0.04%	0.02%	-0.26%	0.01%
	Pre	61%	-0.61%	0.02%	0.00%	0.02%	-0.37%	0.01%
	All	64%	-0.36%	0.01%	0.16%	0.01%	-0.17%	0.00%
Both	In	74%	-0.03%	0.00%	0.39%	0.01%	0.08%	0.00%
	At	81%	-0.04%	0.00%	0.39%	0.01%	0.04%	0.00%
	Out	80%	-0.18%	0.00%	0.13%	0.01%	-0.12%	0.00%
	Pre	63%	-0.50%	0.01%	0.08%	0.02%	-0.29%	0.01%
	All	76%	-0.09%	0.00%	0.32%	0.00%	0.01%	0.00%

Table 2: (cont'd)

Panel B: *One-Week Trading Gains*

<i>Limit Order from</i>	Order Type	→ Inst.	<i>Limit Order Triggered by</i>					
			Institution		Household		Total	
			Mean	s.e.	Mean	s.e.	Mean	s.e.
Institution	In	78%	-0.01%	0.01%	0.25%	0.02%	0.05%	0.01%
	At	84%	-0.07%	0.02%	0.14%	0.04%	-0.04%	0.01%
	Out	84%	-0.14%	0.02%	-0.25%	0.05%	-0.16%	0.02%
	Pre	67%	-0.53%	0.06%	-0.59%	0.10%	-0.55%	0.05%
	All	80%	-0.06%	0.01%	0.12%	0.02%	-0.02%	0.01%
Household	In	57%	-0.19%	0.03%	0.09%	0.04%	-0.07%	0.02%
	At	66%	-0.45%	0.05%	-0.26%	0.08%	-0.39%	0.04%
	Out	73%	-0.57%	0.03%	-0.78%	0.06%	-0.63%	0.03%
	Pre	61%	-0.99%	0.05%	-0.60%	0.07%	-0.84%	0.04%
	All	64%	-0.47%	0.02%	-0.27%	0.03%	-0.40%	0.02%
Both	In	74%	-0.04%	0.01%	0.20%	0.02%	0.02%	0.01%
	At	81%	-0.12%	0.01%	0.02%	0.04%	-0.09%	0.01%
	Out	80%	-0.28%	0.02%	-0.52%	0.04%	-0.33%	0.01%
	Pre	63%	-0.83%	0.04%	-0.59%	0.06%	-0.74%	0.03%
	All	76%	-0.14%	0.01%	-0.03%	0.02%	-0.12%	0.01%

market order traders tend to be on the “correct side” of the market. For example, the average gain to an institution’s market order increases from 0.09% to 0.14% (from one-day to one-week) while a household’s performance increases from -0.32% to -0.03%. Because this result holds for both households and institutions, it does not simply show that institutions are on the correct side of the market—rather, it shows that *limit orders* are on the wrong side. (The wealth transfer from households to institutions in Table 2 is not trivial. Households’ market orders lose 14.8 million euros on the same day and their limit orders lose 22.6 million euros. These losses are gains to institutions: institutions earn 117.1 million from their market orders while their limit orders lose 79.8 million.¹⁶

¹⁶The results are similar for other horizons. For example, measured from one-week trading gains and using the whole sample (i.e., including also the non-uniquely matched trades), the losses to households are 17.7 (market) and 74.7 (limit) million euros. A total of 442,456 individuals traded during the sample period, so the *per trader* losses are 40.0 and 168.8 euros, respectively. In the same sample, institutions’ market orders gain 353.4 million and limit orders lose 261.0. Note that, by definition, households’ and institutions’ gains and losses sum up to zero.

3.2.3 Robustness to Alternative Specifications

The trading gain results are robust to different trading horizons and to analyzing purchases and sales separately. For example, the two-week horizon results are very similar: (1) institutions outperform households for all order types, (2) both households' and institutions' limit orders do better if they are triggered by households, and (3) limit orders placed farther outside the spread perform worse. The results are also very similar for stocks other than the 30 most actively traded. The question about separating purchases and sales is more intricate. For example, if institutions just happened to be on the right side of the market during the sample period, they might appear to have better timing abilities in our tests. However, this is not the case. Institutions display superior performance *relative* to individuals in both purchases and sales.¹⁷

4 An Example of the the Limit Order Effect: Individuals' Attention-Grabbing Behavior

We now demonstrate that limit orders bias inferences about investor behavior. We study individuals' attention-grabbing behavior: i.e., their tendency to trade stocks that grab their attention (Barber and Odean 2002). This phenomenon is susceptible to the limit order effect: investors who first observe the news gain by cutting through the limit order book with market orders. It seems that limit order investors react to news although they are only the passive party.

4.1 Methodology

We test the role of limit orders as follows. First, we compute the dividend-adjusted close-to-close returns for each of the 30 most actively traded stocks. Second, we assign these stocks each day into return-sorted quintiles, with Q1 representing the six stocks with the lowest

¹⁷We also use a calendar-time methodology to compute trading gains. First, we assign stocks into quintiles based on their absolute daily price movements. Second, we compute the average trading gains for each stock/day/quintile/order type/investor type. We then compute the trading gain difference between market orders and limit orders for each stock/day/quintile/investor type and compare averages. This methodology controls for differences in the probability of execution—i.e., it accounts for the fact that more limit orders execute (and lose money) when there is a shock to fundamentals—and shows how limit orders perform relative to market orders. We find, not surprisingly, that limit orders perform well when the stock price change is small but suffer large losses when prices move more.

returns. Third, for each quintile-day, we compute the buy-sell imbalance for institutions’ and households’ market and limit orders. We compute these imbalances as

$$\text{Buy-Sell Imbalance} = \frac{\#\text{Purchases} - \#\text{Sales}}{\#\text{Purchases} + \#\text{Sales}}. \quad (3)$$

This generates $2*2*5 = 20$ time series of buy-sell imbalances. Finally, we compute time-series averages and standard errors across trading days for each investor type/order type/quintile.

4.2 Results

The order imbalance results in Figure 4 show that limit orders contaminate inferences about individuals’ attention-grabbing behavior. Households’ limit order imbalances respond significantly to the return sort. The imbalance decreases from a high of 34% (Q1) to a low of -29% (Q5).¹⁸ The results also modestly support the hypothesis that individuals’ behavior is affected by limited attention: households actively buy shares that go either up or down with market orders. However, these imbalances are small compared to the limit order imbalances. Hence, most of individual investors’ attention-grabbing behavior in our sample comes from the triggering of individuals’ (passive) limit orders. In contrast, *institutions* exhibit true “attention-grabbing” behavior with their market orders but this is probably not because of their limited attention but because of reaction to news. Institutions buy shares with market orders on days stock prices go up and vice versa; their market order imbalance increases from a low of -27% (in Q1) to a high of 24% (in Q5). Institutions’ limit order imbalance, on the other hand, varies only modestly across the quintiles.

The attention-grabbing results are robust to alternative specifications. For example, although our return sort (the same-day return) yields the strongest results, the differences between limit and market order imbalances are similar for volume, news, and yesterday’s return sorts. This is not surprising: because volume, news, and returns are all interrelated¹⁹ and individual stock volatility is highly autocorrelated, all of these sorts are linked to the mechanistic nature of limit orders.

¹⁸The results are nearly identical in a sample limited to households-institutions trades. Thus, our results show that institutions actively sell to and buy from households. Cohen, Gompers, and Vuolteenaho (2002) find a similar result.

¹⁹He and Wang (1995), for example, is a rational model with this feature.

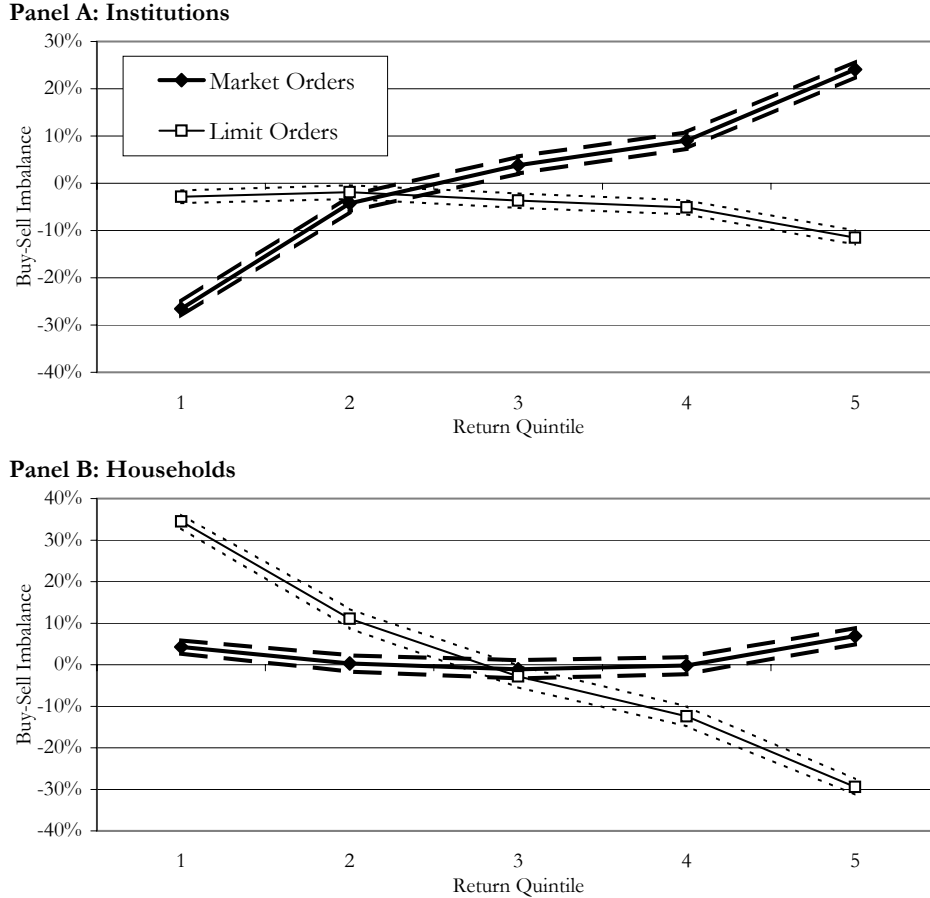


Figure 4: **Limit Orders and Attention-Grabbing Behavior.** This figure shows buy-sell imbalances for institutions' and households' limit and market orders for same-day return quintiles. The sample consists of all uniquely matched trades (see text) in the 30 most actively traded stocks on the Helsinki Exchanges between September 18, 1998 to October 23, 2001. These stocks are assigned into return quintiles each day based on the return from the previous day's close to the same day close. The buy-sell imbalance is computed separately for each day/quintile/order type/investor type as the difference between the number of purchases and the number of sales, divided by the total number of trades. This figure plots the means and standard errors from these time-series.

5 An Example of the Limit Order Effect: Earnings Announcements

This section uses data on earnings announcements to show how institutions react to news and take advantage of households' stale limit orders. Earnings announcements are ideal for a study of the limit order effect because an announcement endows some market participants (i.e., those who monitor the market closely) a transient trading opportunity. An investor who reacts to the information before her competitors can gain from all limit orders between

the current and post-announcement values. An earnings announcement also renders all pre-announcement orders *stale*. In fact, because we focus on pre-scheduled announcements, the market might even endogenously shut down completely just before the announcement.²⁰

We use data on all pre-scheduled earnings announcements during the sample period to maximize the sample size, excluding announcements released outside the regular trading hours. The resulting sample consists of 586 announcements. Each observation contains the date and time of the announcement.²¹

5.1 Methodology

5.1.1 The Use of Limit Orders

We analyze institutions' and individuals' order strategies around earnings announcements in two ways. The first approach uses data on executed trades as follows. First, we categorize all trades into three categories: (1) market order-initiated trades, (2) stale limit order-initiated trades (i.e., trades from pre-announcement orders), and (3) all limit order-initiated trades. Second, we divide the one-hour window around the earnings announcement into two-minute intervals and compute the proportions of market and stale limit order-initiated trades for each announcement/interval/investor type. Finally, we compute average proportions for each interval across announcements, separately for households and institutions.

The second approach uses the predicted identities for unfilled limit orders to analyze households participation in the limit order book. We compute what proportion of (1) all limit orders and of (2) stale limit orders originate from households. We first compute this proportion for each interval/announcement. We then demean each observation around the time-series mean

²⁰There are several reasons why this may not happen in practise. First, it is possible that market participants with orders in the book have a precise signal about the announcement and try to gain from a liquidity-driven shock (e.g., they hope that some investors misinterpret the announcement). Second, it may be that market participants with orders in the book do not know that the announcement is to be released (e.g., because of high monitoring costs) or fail to withdraw their orders in time.

²¹The time-stamp is rounded downwards to the nearest minute. For example, if the time-stamp is 12:03pm, the exact time of the announcement must be $t \in [12:03:00, 12:03:59]$. An announcement is usually released both in English and in Finnish. We use the time-stamp from the announcement that arrives first. Note that we do not classify announcements as positive or negative surprises or exclude announcements that cause no price movements.

to account for heterogeneity in household participation rates:

$$\widehat{HH}_{t,a,\text{all}} = \frac{HH_{t,a,\text{all}} - \frac{1}{30} \sum_{\tau=-15}^{14} HH_{\tau,a,\text{all}}}{\frac{1}{30} \sum_{\tau=-15}^{14} HH_{\tau,a,\text{all}}} \quad (4)$$

where $HH_{t,a,\text{all}}$ is the proportion of household limit orders for interval t in announcement a . We demean the proportion of stale household limit orders in the same way. Finally, we compute the average demeaned proportions for each two-minute interval across announcements.

5.1.2 Trading Gains around Earnings Announcements

We compute average trading gains for all households’ and institutions’ executed orders for each two-minute interval/announcement. We also separately divide the time relative to earnings announcement to *before*, *during*, and *after* intervals: *before* contains all the same day trades executed before the announcement, *during* contains trades executed during the first five minutes after the announcement, and *after* contains all the same day trades executed after these five minutes.

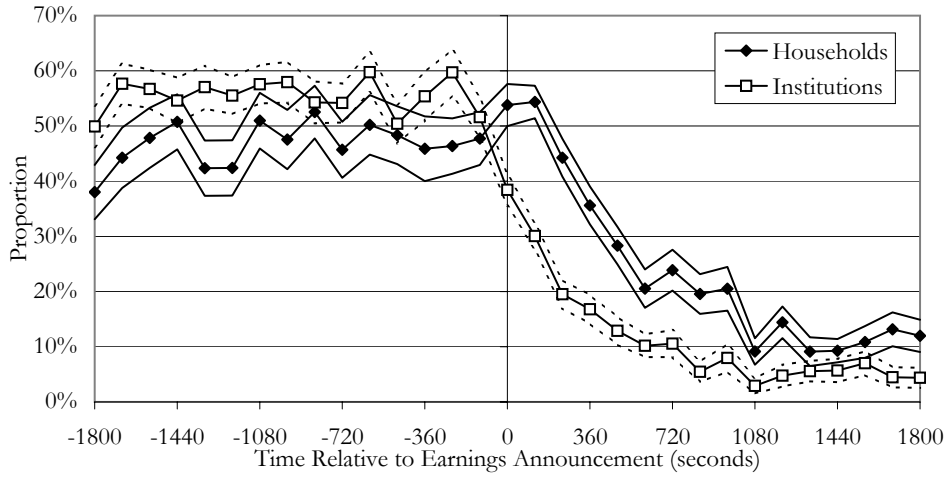
5.2 Results on the Use of Limit Orders

Figure 5 shows an abrupt change in the way institutions and households behave when the announcement is released. Institutions are net suppliers of liquidity before the announcement but submit significant amounts of market orders when the announcement arrives. For example, 57% of institutions’ trades are market order-initiated in the first two-minute interval after the release. The households’ market order ratio drops to 40%.²² These results may explain why Hirshleifer, Myers, Myers, and Teoh (2003) find that households trade *against* earnings surprises.

Figure 6 shows that institutions’ participation in the limit order book changes significantly over the announcement. First, institutions have less limit orders in the book before the

²²The results in Figure 5 are for “households excluding the largest online broker.” The results are slightly moderated but qualitatively the same and significant even if the trades from the online broker are included. For example, the “proportion of market orders” graph still drops to 42% after the announcement. This dampening effect is not surprising: individuals trading through online brokers are in better position, relative to those using traditional brokers, to respond to earnings announcements. We find in unreported work that institutions’ overall participation (i.e., we use unstandardized time series) in the limit order book is low. The average proportion of limit orders from households (across announcements) during the 30-minute period before the announcement is 73.7%.

Panel A: Proportion of Stale Limit Orders



Panel B: Proportion of Market Orders

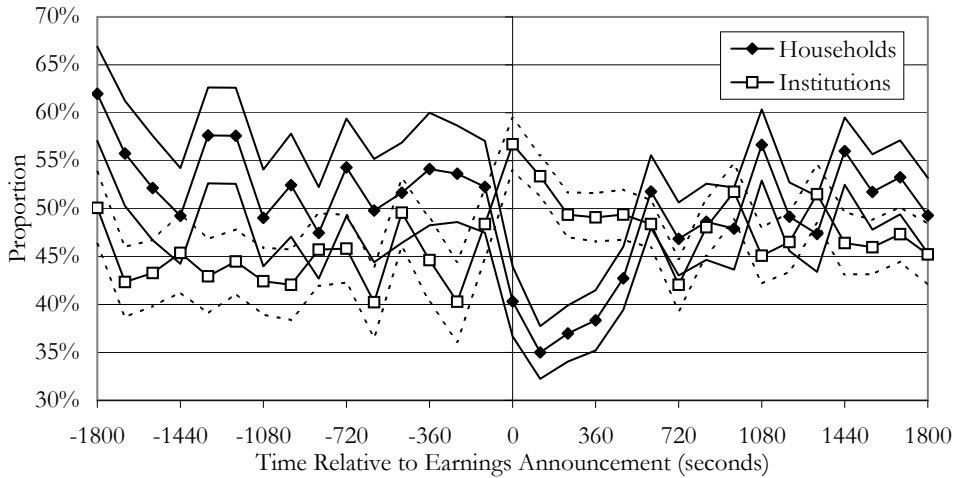


Figure 5: Market Order and Stale Limit Order-Initiated Trades around Earnings Announcements. This figure shows the proportions of stale limit order (Panel A) and market order-initiated trades (Panel B) for institutions and households around earnings announcements. A stale limit order is an order entered into the system before the announcement. The sample consists of all 586 pre-scheduled earnings announcements released during the regular trading hours in the Helsinki Exchanges between September 18, 1998 to October 23, 2001. Trades from the largest online broker are excluded (see text). We compute the number of market order and stale limit order-initiated trades separately for each two-minute period/announcement/investor type. This figure shows the average proportions and their 95% confidence intervals across the announcements.

announcement than after announcement. For example, households' contribution to all limit orders (including stale limit orders) changes by -9% from 30 minutes before to 30 minutes after the announcement. Second, institutions pull out their stale limit orders as the announcement arrives: the change in the fraction of stale limit orders from households is $+8\%$ over the announcement window.

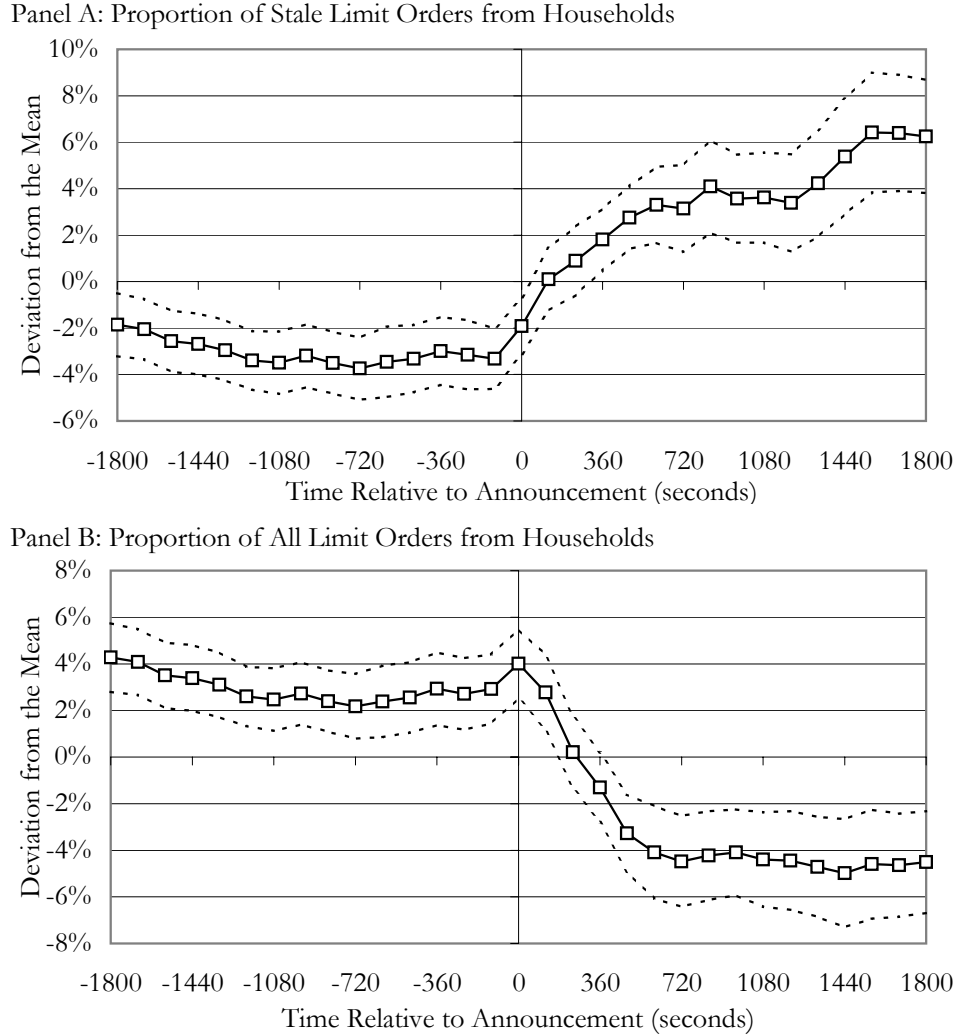


Figure 6: Individual Investors' Participation in the Limit Order Book around Earnings Announcements. Each order entered into the limit order book is identified as originating from either a household or an institution using a logistic model calibrated with executed trades (see text). This figure reports what proportion of (1) stale limit orders (Panel A) and of (2) all limit orders (Panel B) originate from households around earnings announcements. The proportion is first computed each two-minute interval/announcement. These announcement-specific time series are then demeaned by computing the percentage deviation from the overall mean (Eq. 4). This figure plots the average deviations and their 95% confidence intervals. The sample consists of all 586 pre-announced (i.e., no surprises) earnings announcements released during the regular trading hours in the Helsinki Exchanges between July 10, 2000 to October 23, 2001.

5.3 Results on Trading Gains

Figure 7 shows that households' stale limit orders triggered after an announcement perform very poorly. The same-day and one-week trading gains are significantly negative for orders executed during the first eight minutes after the announcement. For example, the cross-

Table 3: Returns on Market Order and Stale Limit Order-Initiated Trades around Earnings Announcements

This table reports trading gains for market orders and stale-limit orders that are executed around earnings announcements. The sample consists of all 586 pre-scheduled earnings announcements released during the regular trading hours in the Helsinki Exchanges between September 18, 1998 to October 23, 2001. A stale limit order is an order entered into the book before the release of an announcement. *Before* contains trades executed before the announcement, *during* contains trades executed during the first five minutes after the announcement, and *after* contains trades executed after these five minutes. The average trading gains are first computed for each interval/announcement with at least two trades. This table reports means and standard errors of these first-stage average trading gains. N is the number of announcements. *# of Trades* is the total number of trades.

	# of Trades	N	Trading Gain Horizon					
			Same-Day		One-Week		Two-Week	
			Mean	s.e.	Mean	s.e.	Mean	s.e.
<i>Households, Stale Limit Orders</i>								
Before	8,776	379	0.36%	0.23%	0.53%	0.37%	0.22%	0.53%
During	2,516	207	-1.58%	0.45%	-2.92%	0.83%	-4.06%	1.20%
After	3,103	368	-0.30%	0.19%	-0.92%	0.47%	-1.51%	0.77%
<i>Households, Market Orders</i>								
Before	9,004	376	-0.18%	0.25%	-0.28%	0.40%	0.32%	0.57%
During	1,627	200	0.72%	0.43%	1.61%	0.80%	2.06%	1.09%
After	28,910	465	-0.12%	0.08%	-0.52%	0.23%	0.12%	0.43%
<i>Institutions, Stale Limit Orders</i>								
Before	18,301	317	0.55%	0.24%	0.84%	0.41%	0.46%	0.49%
During	1,949	191	-0.82%	0.48%	-2.24%	0.80%	-3.21%	1.10%
After	1,738	277	-0.17%	0.23%	-1.16%	0.56%	-0.92%	0.81%
<i>Institutions, Market Orders</i>								
Before	18,188	321	-0.33%	0.23%	-0.58%	0.44%	-0.54%	0.55%
During	5,227	208	1.44%	0.44%	2.99%	0.84%	3.68%	1.22%
After	82,897	426	0.20%	0.07%	0.41%	0.24%	0.39%	0.34%

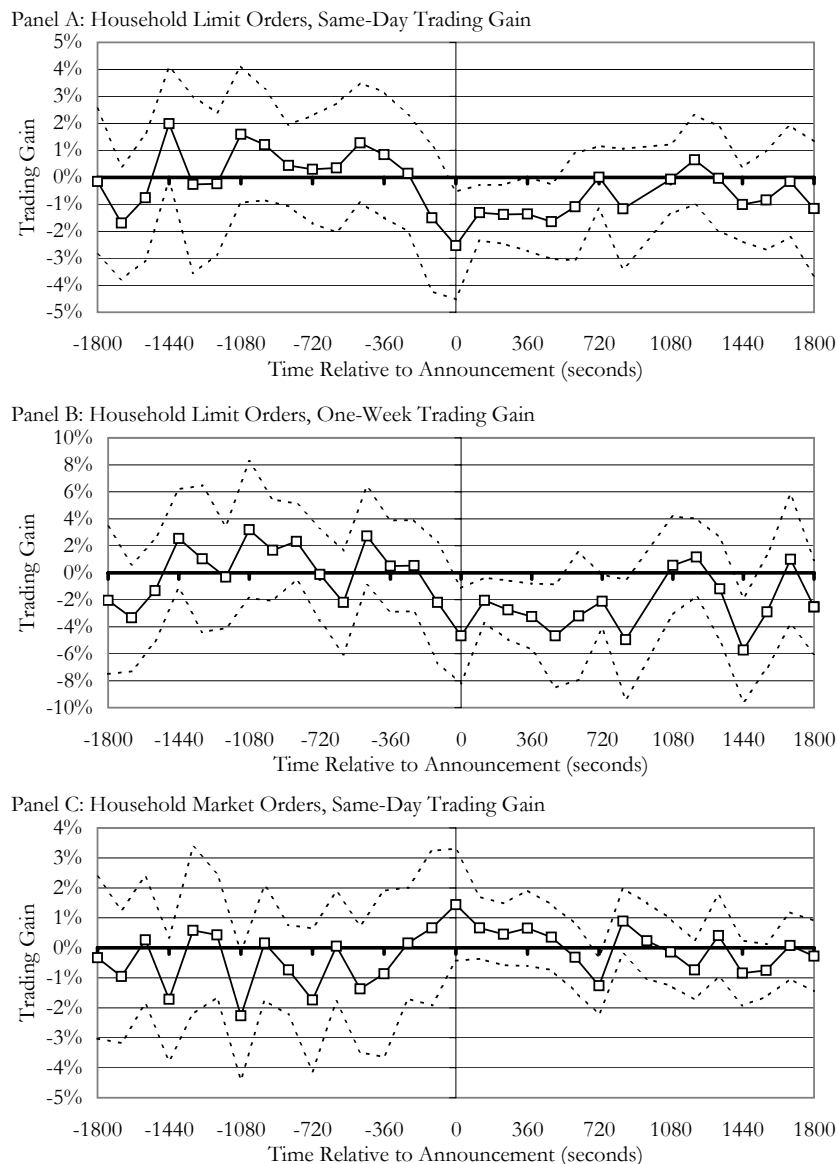


Figure 7: Trading Gains for Households' Market Orders and Stale Limit Orders around Earnings Announcements. This figure shows trading gains for households' market orders and stale limit orders earnings announcements. The sample consists of all 586 pre-scheduled earnings announcements released during the regular trading hours in the Helsinki Exchanges between September 18, 1998 to October 23, 2001. The average trading gain is first computed separately for each investor type/order type/announcement/interval with at least two trades. This figure shows the means and their 95% confidence intervals of these first-stage averages for each interval. Panels A and C report same-day trading gains for limit and market orders; Panel B reports the one-week trading gains for limit orders.

sectional average same-day return on orders executed during the $[0, 120s)$ interval is -2.5% (with a t -value of -2.5). Panel C shows that individuals lose *only* on their limit orders: their market orders have positive trading gains up to ten minutes after the announcement.

A comparison of Panels A and B shows that the performance of limit orders worsens with the horizon. For example, the average one-week return for the limit orders in the first interval is -4.7% . This increase may arise from the post-earnings announcement drift (PEAD). However, note that the standard errors also increase with the horizon—the *statistical* significance of the losses is almost unchanged. Hence, a strategy designed to exploit this effect would be very risky.²³ The mechanism that generates the limit order losses is important: it would not be possible back out and reverse individuals’ trading strategies to make money. Our results show that individuals’ poor timing only mirrors institutions’ great timing. Table 3 shows that institutions earn significant gains by taking advantage of individuals’ stale limit orders. For example, the average same-day announcement window trading gain (“during”-row in the table) to institutions’ market orders is 1.4% .

6 Relation to Earlier Literature

The limit order effect may help to understand several findings about investor behavior in the extant literature. This section such findings but is not an exhaustive list: we merely suggest that the future research must acknowledge the possibility that investors’ *passive* trading strategies may bias inferences drawn from executed trades.²⁴

Differences in Contrarian/Momentum Behavior

Odean (1998), Heath, Huddart, and Lang (1999), Nofsinger and Sias (1999), Grinblatt and Keloharju (2000), Barber and Odean (2002), and others find that less sophisticated investors have a higher propensity to sell after positive price movements and to buy after negative movements. Hirshleifer, Myers, Myers, and Teoh (2003) report that individuals trade against earnings announcements. This contrarian behavior is particularly strong for short-term (or even same-day) price movements. The limit order effect may contribute to these findings: un-

²³For example, an investor could use the data from the first minute after an announcement to observe the direction of the order-flow and then mimic this behavior. However, this strategy is linked to the PEAD because the first minute reaction is probably a good proxy for the direction of the earnings surprise. Hence, this strategy would only capture the well-known result (Ball and Brown 1968) that prices continue to drift after earnings announcements.

²⁴An earlier version of this paper documented that investor class differences in contrarian/momentum strategies and the disposition effect significantly dissipate when investors’ limit and market order-executed trades are separated. We omit these results for the sake of brevity.

informed investors' use of limit orders may cause them to be *passively contrarian*—as opposed to being consciously and actively contrarian.

Disposition Effect

Grinblatt and Keloharju (2001), Shapira and Venezia (2001), and Dhar and Zhu (2002) find that households display a stronger disposition effect than institutions. Limit orders may create such an appearance. For example, suppose an investor just bought ten stocks and immediately needs to sell one of them. She could place sell limit orders for each of these stocks 10% above the current market price. *If one of these sell limit orders executes*, (1) the stock that is sold will be the one with the highest capital gain while (2) the unsold stocks must have unrealized returns of $r_i < 10\%$.²⁵ Albeit an extreme example, the same mechanism *must* always amplify inferences about the disposition effect. The bias in the disposition effect is probably not trivial given that individuals' often place their limit orders far outside the spread.

Correlated Trading

Barber, Odean, and Zhu (2003) and Seasholes and Wu (2005) find that individuals' trading is highly coordinated and concentrated in few stocks each day. Barber and Odean (2002) suggest that this effect may be related to the attention-grabbing behavior. It is likely that the limit order effect contributes to these findings: institutions react to news, cutting through outstanding limit orders with market orders. A flood of market orders from a single smart trader can create an appearance that hundreds of individuals are trading the same stock in the “wrong” direction.

Limit Order Effect in the Literature

The studies mentioned above do not directly address the possibility that the results are amplified by (or due to) the limit order effect. However, several studies—e.g., Barber and Odean (2002), and Barber, Odean, and Zhu (2003)—acknowledge this possibility. For example, Lim

²⁵The fact that the return distributions of unsold stocks are truncated from above actually makes a stronger prediction. For example, suppose that all stock returns are drawn from an arbitrary distribution with mean μ and only one limit sell order executes. This means that the stock that is sold has a return of $r_i = 10\%$ while the *average* return for all other stocks is *strictly less* than μ ! If $\mu = 0$, this means that the investor sells a stock with capital gains while the investor's other stocks have, on average, capital losses.

(2004) refers to an earlier version of this paper and controls for daily price movements to argue that the effects of mental accounting on individuals' trading behavior are robust to the limit order effect. Similarly, Richards (2004, pp. 31) refers to the earlier version and concludes: *“Although data are not available for the other five markets studied here, the similarity between the [Korea Stock Exchange] data and the Finnish evidence. . . suggests that greater use of limit orders by households may be a fairly widespread phenomenon. It is therefore likely that order-submission effects are a substantial cause of the finding that domestic individual investors in Asian equity markets appear to be contrarian investors.”*

7 Conclusions

This paper analyzes how limit orders alter inferences about investor behavior. Because limit orders are mechanically contrarian and exposed to the adverse selection risk, limit orders

- are more likely to execute when there is an information event
- generate losses when there is an information event
- create an appearance that the investor placing the order is reacting to news

For example, a limit order investor appears to exhibit negative market timing when an informed investor triggers the order. A study that does not account for the fact that limit order investors exhibit “passive reaction” to news runs the risk of confounding cause and effect.

This paper examines the importance of the limit order effect. We find that institutions earn large gains by triggering households' limit orders. Households are the passive party in the market: institutions have and trade on information but individuals absorb the order-flow imbalances. This is market-clearing: if one side of the market reacts to news, the other must accommodate the order-flow imbalance. We show that limit orders significantly affect inferences about investor behavior. First, we find that limit orders are the main reason for why individuals appear to trade stocks with attention-grabbing events. Second, we use earnings announcements to demonstrate how institutions quickly react to news to take advantage of individuals' limit orders. This creates an appearance that individuals themselves are reacting to the news and losing money.²⁶ Because limit order books play significant roles in many

²⁶We have focused on a very narrow phenomenon (pre-scheduled earnings announcements). However, our results suggest that institutions profit from their close monitoring of the market also in other instances. In

markets, we suspect that limit orders significantly alter inferences about investor behavior in many trading record data sets.²⁷

Limit orders are also important for their welfare effects. We find that households' limit order strategies may be suboptimal. For example, individuals' tendency to place limit orders far outside the spread seems insensible: these orders usually execute because of an information shock, generating significant losses. The most unambiguous normative implication is that individuals could improve their welfare by reducing their use of stale limit orders. Any inefficiency in individuals' use of limit orders creates a wealth transfer to institutions and our results suggest this transfer may be very significant.

particular, all unexpected events—such as a release of an earnings warning or a natural disaster—are more profitable because limit order investors cannot withdraw their limit orders in anticipation of the news.

²⁷Limit orders are widely used also in the US markets. For example, the Securities Exchange Act Release (September 6, 1996) states that “*limit orders accounted for 50% of [the NYSE] customer trades of 100-500 shares and 66% of customer trades of 600-1000 shares.*” Many exchanges are completely order-driven and almost all supplement the market maker/specialist structure with a limit order book. The role of limit orders is unlikely to diminish in the future. First, the SEC has increasingly promoted individuals' ability to compete directly in the market place. Second, the arrival of online brokers has granted individuals a convenient and direct access to the trading systems.

Appendix

A A Simple Model of a Limit Order-Driven Market

This section formulates a tractable and stylized equilibrium model of a limit order-driven market and uses it to demonstrate the mechanical aspects of limit orders (the limit order effect). Our approach is motivated by the models of Glosten (1994) and Handa and Schwartz (1996).²⁸ The spread in our model arises endogenously from the risk of adverse selection—there are no order-processing or inventory effects. The motivation for the use of limit orders in our model is that the order may trigger because of a liquidity-shock.

A.1 Setup

We assume the following:

- There are three dates, $t = 0, 1, 2$. There is a single stock traded in an order-driven market with a current intrinsic value V_0 (known to everyone). The date 2 value is \tilde{V}_2 , unknown at date 0.
- The market consists of many uninformed and risk-neutral agents who want purchase or sell a single share for liquidity reasons.
- The investor can submit a limit order at date 0 at price L , or wait until date 2 and submit a market order. If the investor submits a market order at date 2, the execution price is $V_0 \pm s$, where s is the half-spread. If the investor's limit order does not execute, the investor always submits a date 2 market order.
- The investors maximize expected profits: $E[\tilde{V}_2 - P]$ for a buyer and $E[P - \tilde{V}_2]$ for a seller, where P is the execution price.
- At date 1, a large institutional investor enters the market. This investor trades on a private signal with probability π and is otherwise trading for liquidity reasons. The two

²⁸Recent studies on the use of limit orders include Biais, Hillion, and Spatt (1995), Harris and Hasbrouck (1996), Aitken, Berkman, and Mak (2001), Lo, MacKinlay, and Zhang (2002), Bae, Jang, and Park (2003), Ranaldo (2004).




Figure 8: **The Timeline of the Events in the Limit Order Model**

possible outcomes are:²⁹

- If there is a signal, the date 2 intrinsic value is drawn from a uniform distribution $U(V_0 - \delta_I, V_0 + \delta_I)$, where $\delta_I > 0$ is an exogenous parameter.
- If there is no signal, a liquidity-shock temporarily pushes the stock price to a level drawn from $U(V_0 - \delta_U, V_0 + \delta_U)$, where $\delta_U > 0$ is an exogenous parameter. The date 2 intrinsic value is the same as the date 0 value, $V_2 = V_0$.

Figure 8 illustrates the sequence of the model's events.

A.2 Optimal Behavior, the Choice of Limit Price, and Equilibrium

We analyze the problem by focusing on a single investor who needs to buy one share. The investor can do two things: wait until date 2 and submit a market order or submit a limit order at date 0. If the investor decides to wait until date 2, the expected loss is equal to the half-spread. Now, suppose that the investor places a limit order. There are four possible outcomes:

1. *There is a signal but the order does not execute.* The expected terminal value is $E_0[V_2 | V_2 > L] = \frac{V_0 + \delta_I + L}{2}$. The investor pays the spread s to buy the share with a market order at date 2. The probability of this event is $\pi \frac{V_0 + \delta_I - L}{2\delta_I}$.
2. *There is a signal and the order executes.* The expected terminal value is $E_0[V_2 | V_2 \leq L] = \frac{V_0 - \delta_I + L}{2}$ and the profit $\frac{V_0 - \delta_I - L}{2}$. The probability of this event is $\pi \frac{L - V_0 + \delta_I}{2\delta_I}$.
3. *There is no signal and the order does not execute.* The terminal value is V_0 and the investor pays the spread s to buy the share at date 2. The probability of this event is $(1 - \pi) \frac{V_0 + \delta_U - L}{2\delta_U}$.

²⁹Note that this model does not explicitly consider *how* market orders and limit orders interact; instead, we use “the outsider” as the source of both information and liquidity-driven shocks. We do this for tractability: our stylized model emphasizes the information/liquidity shock tradeoff faced by a limit order trader while shutting out unnecessary complications.

4. *There is no signal and the order executes.* The terminal value is V_0 and the profit $V_0 - L$. The probability of this event is $(1 - \pi) \frac{L - V_0 + \delta_U}{2\delta_U}$.

The investor's optimal limit price maximizes the expected profit

$$L^* = \arg \max_L \left\{ -\pi \frac{(L - V_0 + \delta_I)^2}{4\delta_I} + \pi \frac{L - (V_0 - \delta_I)}{2\delta_I} (-s) \right. \\ \left. + (1 - \pi) \frac{V_0 + \delta_U - L}{2\delta_U} (-s) + (1 - \pi) \frac{L - V_0 + \delta_U}{2\delta_U} (V_0 - L) \right\} \quad (\text{A1})$$

$$= \frac{\pi\delta_U(V_0 + s - \delta_I) + 2(1 - \pi)\delta_I \left(V_0 - \frac{\delta_U - s}{2} \right)}{\pi\delta_U + 2(1 - \pi)\delta_I}. \quad (\text{A2})$$

The agent prefers a limit order to a market order when the expected return at L^* is higher than $-s$ (the expected loss from a market order):

$$-s \leq \frac{-\pi \frac{(V_0 + \delta_I - L^*)^2}{4\delta_I} + (1 - \pi) \frac{L^* + \delta_U - V_0}{2\delta_U} (V_0 - L^*)}{1 - \pi \frac{L^* + \delta_I - V_0}{2\delta_I} - (1 - \pi) \frac{V_0 + \delta_U - L^*}{2\delta_U}} \\ = \frac{-\pi(V_0 + \delta_I - L^*)^2\delta_U + 2(1 - \pi)(L^* - V_0 + \delta_U)\delta_I(V_0 - L^*)}{2\pi(L^* - V_0 + \delta_I)\delta_U + 2(1 - \pi)(L^* - V_0 + \delta_U)\delta_I}. \quad (\text{A3})$$

We now use two aspects of equilibrium in an order-driven market:

- The market is order-driven so the spread must arise from investors limit orders (i.e., $V_0 - L^* = s$).
- All investors in the model are homogeneous—except that some want to buy and others want to sell—so that at equilibrium, investors must be indifferent between submitting limit and market orders (i.e., Eq. A3 holds with equality).

The equilibrium spread from substituting L^* from Eq. A2 into Eq. A3 is

$$s^* = \frac{-(1 - \pi)\delta_I^2\delta_U + \sqrt{\pi(1 - \pi)\delta_I^2\delta_U(\delta_I - \delta_U)(\pi\delta_U + 2(1 - \pi)\delta_I)((1 - \pi)\delta_I + (1 + \pi)\delta_U)}}{(\pi\delta_U + (1 - \pi)\delta_I)^2}. \quad (\text{A4})$$

A necessary condition for the existence of equilibrium ($s^* > 0$) is that $\delta_I > \delta_U$.³⁰ We suppose that parameters $\{\pi, \delta_I, \delta_U\}$ constitute such equilibrium.

A.3 Implications

1. *More limit orders execute when there is an information event.* Suppose that there is a buy limit order at price $L_b > V_0 - \delta_U$ and a sell limit order at price $L_s < V_0 + \delta_U$ so that both a signal and a liquidity shock can trigger the orders. (Whether L_b and L_s are equilibrium values or arbitrary is inconsequential for our argument.) The probability of a limit order executing when there is no signal is $\frac{2\delta_U - (L_s - L_b)}{2\delta_U}$ and $\frac{2\delta_I - (L_s - L_b)}{2\delta_I}$ when there is a signal. Hence, more orders execute when there is a signal if $\delta_I > \delta_U$, which is a necessary condition for equilibrium.
2. *Limit orders placed farther from the fair value are more likely to trigger when there is an information event.* For example, a buy limit order placed in the interval $[V_0 - \delta_I, V_0 - \delta_U)$ in our model executes *only* when there is an information event. The ratio of information versus liquidity-driven execution probabilities increases in the distance from the spread.
3. *Investors using limit orders trade in the “wrong” direction in response to news.* Limit order traders trade in the wrong direction when there is an information event because only orders overshot by the intrinsic value execute. Limit orders triggered due to an information event on average lose $-\frac{1}{2}(\delta_I - s^*)$ in equilibrium.
4. *Investors (on average) gain less than the half-spread from using limit orders.* This is the risk of adverse selection. The first and fourth terms in Eq. A1 are the two opposing effects: (1) if the limit order executes because of a liquidity shock, it earns the half-spread but (2) the losses to the informed investor offset some of these gains. In our risk-neutral model, these losses offset all the gains (because Eq. A3 holds with equality).

³⁰The sufficient condition can be written as

$$\pi(\delta_I - \delta_U)(\pi\delta_U + 2(1 - \pi)\delta_I)((1 - \pi)\delta_I + (1 + \pi)\delta_U) > (1 - \pi)\delta_I^2\delta_U \quad (\text{A5})$$

which is a cubic equation in δ_I , δ_U , and π . Note that this condition is satisfied when $\delta_I > \delta_U$ and π tends close to 1; i.e., when the risk of adverse selection is sufficiently high.

References

- Aitken, M., H. Berkman, and D. Mak (2001). The use of undisclosed limit orders on the Australian Stock Exchange. *Journal of Banking and Finance* 25, 1589–1603.
- Bae, K.-H., H. Jang, and K. S. Park (2003). Traders’ choice between limit and market orders: Evidence from nyse stocks. *Journal of Financial Markets* 6, 517–538.
- Bagehot, W. (1971). The only game in town. *Financial Analysts Journal* 27, 12–14, 22.
- Ball, R. and P. Brown (1968). An empirical evaluation of accounting income numbers. *Journal of Accounting Research* 6, 159–178.
- Barber, B. M., Y.-T. Lee, Y.-J. Liu, and T. Odean (2005). Who loses from trade? Evidence from Taiwan. University of California, Berkeley, working paper.
- Barber, B. M. and T. Odean (2002). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. University of California, Berkeley, working paper.
- Barber, B. M., T. Odean, and N. Zhu (2003). Systematic noise. University of California, Berkeley, working paper.
- Biais, B., P. Hillion, and C. Spatt (1995). An empirical analysis of the limit order book and the order flow in the Paris Bourse. *Journal of Finance* 50, 1655–1689.
- Bloomfield, R., M. O’Hara, and G. Saar (2005). The “make or take” decision in an electronic market: Evidence on the evolution of liquidity. *Journal of Financial Economics* 75, 165–199.
- Brennan, M. J. and H. H. Cao (1997). International portfolio investment flows. *Journal of Finance* 52, 1851–1880.
- Brennan, M. J., H. H. Cao, N. Strong, and X. Xu (2005). The dynamics of international equity market expectations. *Journal of Financial Economics* 77, 257–288.
- Cohen, R. B., P. A. Gompers, and T. Vuolteenaho (2002). Who underreacts to cash-flow news? Evidence from trading between individuals and institutions. *Journal of Financial Economics* 66, 409–462.

- Copeland, T. E. and D. Galai (1983). Information effects and the bid-ask spread. *Journal of Finance* 38, 1457–1469.
- Dhar, R. and N. Zhu (2002). Up close and personal: An individual level analysis of the disposition effect. Yale School of Management, working paper.
- Fama, E. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance* 25, 383–417.
- Glosten, L. (1994). Is the electronic open limit order book inevitable? *Journal of Finance* 49, 1127–1161.
- Glosten, L. and P. Milgrom (1985). Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 13, 71–100.
- Grinblatt, M. and M. Keloharju (2000). The investment behavior and performance of various investor-types: A study of Finland’s unique data set. *Journal of Financial Economics* 55, 43–67.
- Grinblatt, M. and M. Keloharju (2001). What makes investors trade? *Journal of Finance* 56, 589–616.
- Handa, P. and R. A. Schwartz (1996). Limit order trading. *Journal of Finance* 51, 1835–1861.
- Harris, L. and J. Hasbrouck (1996). Market vs. limit orders: The superdot evidence on order submission strategy. *Journal of Financial and Quantitative Analysis* 31, 213–231.
- He, H. and J. Wang (1995). Differential information and dynamic behavior of stock trading volume. *Review of Financial Studies* 1995, 919–972.
- Heath, C., S. Huddart, and M. Lang (1999). Psychological factors and stock option exercise. *Quarterly Journal of Economics* 114, 601–627.
- Hirshleifer, D., J. N. Myers, L. A. Myers, and S. H. Teoh (2003). Do individual investors drive post-earnings announcement drift? Ohio State University, working paper.
- Holden, C. W. and A. Subrahmanyam (1992). Long-lived private information and imperfect competition. *Journal of Finance* 47, 247–270.
- Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica* 53, 1315–1336.

- Lim, S. S. (2004). Do investors integrate losses and segregate gains? Mental accounting and investor trading decisions. *Journal of Business*, forthcoming.
- Linnainmaa, J. (2003). Does it matter who trades? Investor sophistication, broker identity, and permanent price impacts. working paper.
- Lo, A. W., A. C. MacKinlay, and J. Zhang (2002). Econometric models of limit-order executions. *Journal of Financial Economics* 65, 31–71.
- Nofsinger, J. R. and R. W. Sias (1999). Herding and feedback trading by institutional and individual investors. *Journal of Finance* 54, 2263–2295.
- Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775–1798.
- Odean, T. (1999). Do investors trade too much? *American Economic Review* 89, 1279–1298.
- Ranaldo, A. (2004). Order aggressiveness in limit order book markets. *Journal of Financial Markets* 7, 53–74.
- Richards, A. (2004). Big fish in small ponds: The trading behaviour and price impact of foreign investors in Asian emerging equity markets. *Journal of Financial and Quantitative Analysis*, forthcoming.
- Saar, G. (2001). Price impact asymmetry of block trades: An institutional trading explanation. *Review of Financial Studies* 14, 1153–1181.
- Seasholes, M. S. and G. Wu (2005). How costly are behavioral biases? University of California, Berkeley, working paper.
- Shapira, Z. and I. Venezia (2001). Patterns of behavior of professionally managed and independent investors. *Journal of Banking and Finance* 25, 1573–1587.