ELSEVIER

Journal of Financial Markets 3 (2000) 83–111

Journal of
FINANCIAL
MARKETS

www.elsevier.nl/locate/econbase

Inferring investor behavior: Evidence from TORQ data☆

Charles M.C. Lee[a], Balkrishna Radhakrishna[b,*]

[a]Cornell University, USA
[b]Carlson School of Management, University of Minnesota, 3-271, 321, 19th Avenue S., Minneapolis, MN 55455, USA

Abstract

We use a unique dataset (TORQ) to calibrate several techniques commonly used to infer investor behavior from transactions data. Specifically, we evaluate the Lee–Ready (1991: Journal of Finance 46, 733–746) algorithm for distinguishing trade direction, and we examine the use of trade size as a proxy for trader identity. We find that, due to complexities in the NYSE auction process, up to 40% of reported trades cannot be unambiguously classified as either buyer- or seller-initiated. However, for those trades that can be classified, the Lee–Ready algorithm is 93% accurate. In addition, we construct a firm-specific trade size proxy that is highly effective in separating the trading activities of individual and institutional investors. These findings should be useful to researchers interested in inferring trader sophistication and buy/sell direction from transactions data.

JEL classification: G15

Keywords: Intraday transaction analysis; Trade direction; Investor behavior; Individual investors; Institutional investors

☆We thank the participants of the University of Minnesota accounting workshop, Bruce Lehman (Editor) and an anonymous referee for helpful comments.

* Corresponding author. Tel.: (612) 624-0302; fax: (612) 626-1335.

E-mail address: bradha@tc.umn.edu (B. Radhakrishna)

PII: S1386-4181(00)00002-1

1. Introduction

Traditional capital market research treats the market aggregation process as a 'black box' (see Fig. 1). News and information are impounded in some mysterious fashion and prices adjust to reflect the aggregate effect of this information. Recent developments in market microstructure research are beginning to unravel this black box. Armed with micro-level transactions data, many capital market researchers have begun to focus on the information structure and trading mechanisms that underpin the price formation process. The study of observed trades and quotes is leading to many new insights about market participants: how they respond to information, submit orders, engage in trades, affect market makers, and determine prices.

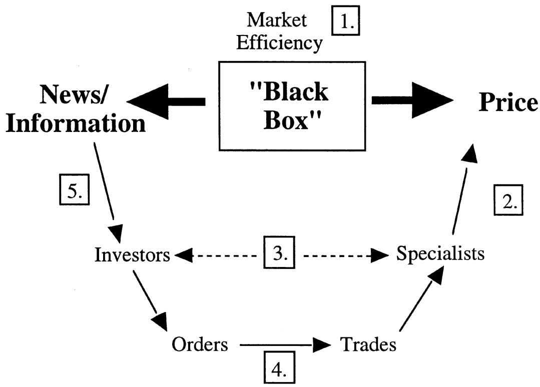

1. **Market Efficiency**
   - Speed of adjustment to news (firm-specific and market-wide).
   - Effect of spreads and other transaction costs on pricing anomalies.

2. **Price Formulation and Specialist Behavior**
   - What factors cause specialists to change quoted prices.

3. **Information Asymmetry**
   - Models in which traders are differentially informed. Market-makers lose to the informed and profit from noise/liquidity traders.
   - Empirical estimation of spread components. Tests of existing models.

4. **Market Regulation and Design**
   - Market transparency; Intermarket information transfer; Measuring market quality; Payments for order flows; Circuit breakers and trading halts; Front running; Minimum tick sizes; Alternative auction mechanisms.

5. **Information Content and Investor Behavior**
   - Response of investors to news signals.
   - Characterizing investor behavior using transactions data.

Fig. 1. Major areas of current research in market microstructure.

To progress further, this line of inquiry must overcome some important methodological challenges. Most of these challenges revolve around the distinction between *orders* and *trades*. Inquiries into investor behavior logically focus on the flow of the submitted orders. Yet commercially available transactional databases typically provide information, not on orders, but on trades and quotes.[1] Consequently, researchers interested in investor behavior patterns must infer this information through empirical proxies, such as trade size (e.g., Cready, 1988; Cready and Mynatt, 1991; Lee, 1992), and inferred buy–sell directions (e.g., Lee and Ready, 1991; Madhavan and Smidt, 1991, 1993; Huang and Stoll, 1994; Madhavan et al., 1997; Madhavan and Cheng, 1997; Easley et al., 1996, 1997; Gardner and Subrahmanyam, 1997; Lee, 1993). Unfortunately, without information about the original orders, it is impossible to verify the accuracy of these inference techniques.

In this study, we use the TORQ database made available by the New York Stock Exchange (NYSE) to examine the usefulness of several commonly used techniques for inferring investor behavior. Our investigation centers on three issues: (1) the mapping between orders and trades (that is, whether orders are 'batched' or 'split up' during execution); (2) the buy/sell direction of observed trades (that is, whether a trade is initiated by a buyer or a seller); and (3) the usefulness of trade-size as means for separating trades initiated by individuals and institutional investors.

The first issue deals with the basic reliability of trades as proxies for investor orders. If submitted orders are frequently batched or split up during execution, then the size and direction of reported trades cannot be used to make reliable inferences about the identity or intent of the order initiator. The second issue relates to the question of whether researchers can reliably infer the buy/sell intent of the trade initiator. In an earlier study, Lee and Ready (1991) suggest that, in continuous auction markets such as the NYSE, it may be possible to classify trades as to whether they are buyer- or seller-initiated. Lee and Ready describe an algorithm for making these inferences, which has become widely used in the microstructure literature. We provide direct evidence on the reliability of this algorithm.

Finally, we examine the usefulness of a trade-size proxy in separating trades initiated by individual investors from those initiated by institutional investors. The distinction between these two main investment groups is reflected in the institutional structure of the brokerage industry. Several prior studies suggest individual and institutional investors may differ in their level of sophistication in

[1] For example, the Institute for the Study of Securities Markets (ISSM) database or the Trades and Quotes (TAQ) database from the New York Stock Exchange (NYSE). Order/trade differentiation is particularly problematic for researchers working with data from the NYSE. Data from markets with order driven open limit books (e.g., Toronto, Paris and Tokyo) do not suffer from these problems to the same degree.

response to information.[2] A number of academic studies in finance also suggest that the differential trading behavior of these two groups may help explain well-known pricing anomalies.[3] Some studies have used trade size to distinguish between individual and institutional activities (Bhattacharya, 1998; Cready and Mynatt, 1991; Cready, 1988; Radhakrishna, 1998; Lee, 1992). However, the extant literature contains no evidence on the usefulness and limitations of this inference technique.

The TORQ dataset offers a unique opportunity to test these inference techniques because, in addition to information about trades and quotes, TORQ also contains information about individual orders. Specifically, for electronically routed (SuperDot) orders at the NYSE, TORQ provides details on the parties to each side of a trade, the time when their orders were submitted, and how these orders were handled prior to execution. Certain facts about the trader's identity (e.g., individuals versus institutional traders) and order characteristics (e.g., whether it was a seller- or buyer-initiated order) can be directly observed in this data.

The TORQ dataset itself is limited in scope – the dataset covers only 144 firms over a three-month period during 1990–91. These firms represent a size-stratified random sample of NYSE stocks, which should be representative of the larger population of all NYSE firms. However, the time-period covered is brief and detailed information on orders is only available for those processed by SuperDOT.[4] Our strategy is to use this limited sample as a laboratory, to test the power and accuracy of fairly general algorithms that can then be applied to bigger datasets such as the ISSM and TAQ. Consequently, our findings should help empirical researchers using either the ISSM or the TAQ dataset to calibrate the accuracy of their inferences.

Our main findings are as follows:

(1) The mapping between orders and trades is more complex than previous studies suggest. Few (6%) of the total market orders are *split up* in execution. However, a much larger proportion (24%) of total market orders is *batched* in execution. Batching is particularly prevalent in larger trades. For trades of more than 1900 shares, we find that most (56%) have multiple participants on the active-side.
(2) A surprisingly large proportion of reported trades (approximately 40%) are 'non-directional' – that is, they cannot be unambiguously classified as either

[2] For example, Bhattacharya (1998), Hand (1990), Lee (1992), Radhakrishna (1998) and Walther (1997).

[3] See, for example, Ritter (1988), Lakonishok and Maberly (1990), and Lee et al. (1991), Lease et al. (1974), Yunker and Krehbiel (1988), and Shiller and Pound (1989) present survey evidence on the differential decision styles of institutional and individual investors.

[4] While SuperDOT handled a majority of the volume at the NYSE during our sample period, these orders are smaller and tend to contain a greater proportion of orders from individuals.

buyer- or seller-initiated. Much of the ambiguity arises from 'stopped' trades and market 'crosses'.[5] However, in those instances where trades can be unambiguously classified, we find the Lee–Ready algorithm has a 93% overall agreement rate with TORQ classifications. The accuracy is highest for trades executed at the spread (98%), and is lower for midspread classifications using the tick test (76%) or the zero-tick test (60%).[6]

(3) Despite the prevalence of order batching, we find that trade size is still highly effective in separating institutional and individual investor activities. We show that the likelihood of misclassifications can be greatly reduced by using a buffer zone of medium-sized trades. Moreover, we show it is possible to further enhance the accuracy of the individual/institutional classification by conditioning trade size cutoff values on firm size. We construct such a proxy and provide error probability benchmarks for several trade size cutoffs for small, medium and large NYSE firms.

In sum, we show that batched orders, stopped orders, and market crosses all add noise to the inference process. However, despite these problems, our results suggest that substantial information about the original orders can still be inferred from trades and quotes data. Specifically, the active-side of each trade, as identified by the Lee–Ready (1991) algorithm, is generally a good proxy for the frequency, size, and direction of incoming market orders. In addition, our results suggest that it is possible to construct a trade size proxy that is highly effective in separating institutional and individual activities.

Finally, we wish to highlight the potential effect of two important changes in market dynamics since our sample period. First, trading volume has increased tremendously in the 1990s. This is likely to lead to more frequent batching of orders and an increase in the proportion of orders that are not clearly buyer or seller initiated. Second, U.S. equity markets now trade in smaller price increments (the NYSE adopted trading in 'teenies' in 1997). This development may also lead to changes in trading strategy and behavior that impact the validity of the assumed linkages. While it is difficult to estimate the net effect of these developments on our results, researchers interested in using our findings with more recent data should be mindful of these potential problems.

The remainder of the paper is organized as follows. Section 2 provides a brief introduction to the TORQ dataset and highlights its unique features.

[5] On receiving a market order, the NYSE specialist may 'stop' the order, thereby guaranteeing execution at the stop price (the prevailing quote) while attempting to execute the order at a better price. The subsequent execution of a stopped market order is difficult to classify unambiguously as either a buy or a sell. Market 'crosses' refer to market orders matched against another market order (or a limit order that is executable upon receipt). Market crosses are also difficult to classify unambiguously since both sides of the trade can be deemed to have initiated the transaction.

[6] Our results on the efficacy of the LR algorithm are largely consistent with Odders-White (1999) and Lightfoot et al. (1999). We discuss the differences between these studies in more detail later.

Section 3 investigates the mapping between orders and trades. Section 4 examines the accuracy of the Lee-Ready algorithm for inferring trade direction. Section 5 evaluates the usefulness of a trade size proxy for distinguishing between institutional and individual activities, and Section 6 concludes.

2. The TORQ dataset

The TORQ dataset was prepared under the supervision of Professor Joel Hasbrouck during his tenure as Visiting Economist to the NYSE. This dataset contains trades, quotes, order processing, and audit trail data for a sample of 144 NYSE stocks for the three months from November 1990 through January 1991. The firms were randomly selected from a size-stratified population of NYSE firms in the summer of 1990. Specifically, 15 firms were selected from each of ten deciles formed on the basis of market capitalization. Six of the original 150 securities were delisted before final data collection began on November 1, 1990, leaving the final sample of 144 firms.[7]

The database consists of four principal files – the consolidated trade file (CT), the consolidated quote file (CQ), the system order database file (SOD), and the consolidated audit file (CD). Fig. 2 provides an illustration of how the four files are related. The CT and CQ files contain details of each trade (price, shares traded, time, originating exchange and special trading condition codes) and each quote revision (bid and ask prices, bid and ask depths, time, originating exchange and special trading condition codes). Together, these two files provide all the information normally disseminated to market participants, in real time, through the Consolidated Tape System (CTS). Historical archives of this data are available to researchers through sources such as the ISSM or TAQ. Most recent empirical microstructure studies have been based on this data.

The incremental value of the TORQ database derives from the CD and SOD files. The CD file is an expanded version of the CT file, in that it provides information on the number and type of parties to a trade. These two files are linked by a trade sequence number.[8] For each trade, the CD file provides information on orders that make up the buy and sell side. Specifically, for trades executed at the NYSE, the CD file provides the number of shares traded, the order type (market, limit or other special types), and the trader type (individual, institutional, member, or programmed trade). For off-Board executions, the CD

---

[7] For more details on the TORQ database and on trading procedures at the NYSE, see Hasbrouck (1992), Hasbrouck and Sosebee (1992), and Hasbrouck et al. (1993).

[8] Some trades appeared in the CD file at the end of the day that were not in the CT file. NYSE officials indicate these are most likely post-transaction settling up adjustments. We focus our analysis on NYSE trades that appear in the CT file, since the CT file contains information made available in real time to all investors through the Consolidated Tape.

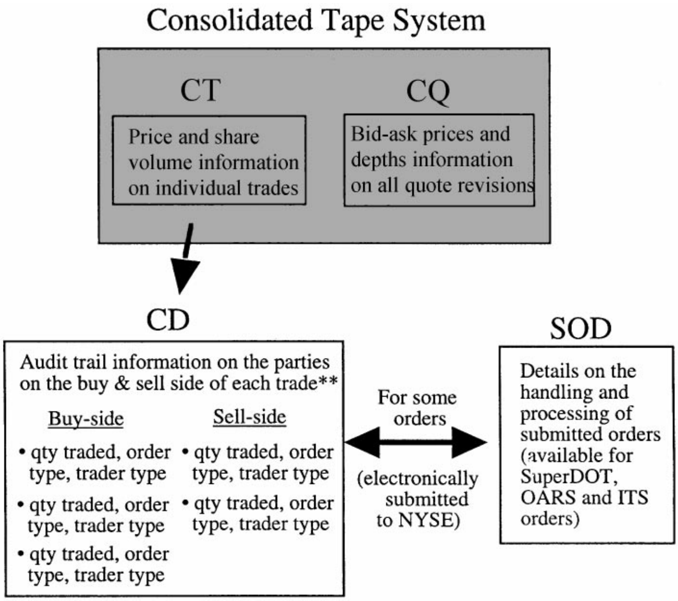

The TORQ database consists of four principal files. The Consolidate Trades (CT) and Consolidated Quotes (CQ) files contain broadly disseminated trade and quote data. The Audit file (CD) contains details of parties to each trade and is linked to the CT file by a trade sequence number. The System Order Database file (SOD) contains details on order submission and handling for orders that are electronically handled at the New York Stock Exchange (NYSE). Information in the SOD comes from three sources: SuperDOT, which is the main routing system at the NYSE; OARS, which handles orders at market openings; and ITS, which handles intermarket trades.

** All CT trades have one or more CD record. However, detailed order and trader information is available only for NYSE executions.

Trade and quote information is available in the broadly disseminated transactional datasets (e.g., ISSM or TAQ)

Fig. 2. An overview of the TORQ dataset.

file reports the shares traded, but not the order or trader information. Most of our analyses are based on NYSE executed trades, since no information is available on the different investor categories on off-Board trades.

The SOD file documents the arrival and processing of orders received from one of three systems: SuperDot (the main electronic order processing system for the NYSE), OARS (the Opening Automated Report System, used at market openings) and ITS (the Intermarket Trading System, used to transfer orders between market centers). Non-electronic orders, such as those submitted to floor brokers, are not in the SOD file. The SOD file is useful in providing specific details on order placement, routing, partial executions and stopped orders.

However, the link between the CD and SOD files is more difficult to establish, because there is no unique identifier between these two files. In this study, we use the SOD file only to obtain the sample of orders executed, to test the Lee–Ready (1991) algorithm.[9]

3. The mapping between orders and trades

The mapping between orders and trades is of elemental importance in inferring investor behavior. If orders are routinely batched or split up during execution, then observed trades would be a poor proxy for the incoming order flow. In particular, we will not be able to rely on observed trades to infer information about the *frequency* of order arrivals, their *size*, or their buy-sell *direction*.

In this section, we evaluate the mapping between orders and trades in two ways. First, we begin with a population of submitted orders, and trace them through to their execution. The main purpose of these tests is to evaluate the extent to which orders are *split up* in execution. Second, we begin with a population of trades, and trace the active-side of each trade back to the original submitted order(s). These tests evaluate the extent to which the active-side of reported trades reflect the *batching* of multiple orders.

3.1. Are orders split up in execution?

The first issue we address is the frequency with which orders (specifically market orders) are split up in execution. Table 1 presents an overview of the total orders in the TORQ dataset reported as having been executed. A total of 687,980 executed orders are in the SOD file, representing all executed SuperDot orders for the 144 sample firms.[10] Of these, 258,058 (37%) are limit or other special orders and 429,922 (63%) are market orders. We focus on market orders because these orders typically trigger trades.

Panel A of Table 1 shows that the vast majority of market orders receive single executions – a total of 403,443 market orders (94%) are filled in single executions; only 26,479 (6%) are filled in multiple executions. Panel B provides a more detailed breakdown of the market orders that received multiple execution. Most of these orders are ‘stopped’ by the specialist or are for more than

[9] Specifically, we use SOD records to obtain the time an order is received. This time stamp is then used to identify the quote prices at the time of order arrival.

[10] A total of 875,133 orders are in the SOD file, but only 687,980 orders were actually executed. The orders we analyze include all ITS commitments to trade and orders received as part of the daily opening batch trade through the OARS system.

Table 1
The Mapping of Orders to Trades.

This table presents evidence on how often market orders are split up in execution. Part A is an overview of the executed orders in the TORQ database. Numbers in italics represent the percentage breakdown relative to the total reported in the previous hierarchical level. Part B presents a more detailed analysis of market orders executions, conditional on the size of the original order. The last column of Part B reports the market orders that are only partially executed by the end of the day

Part A: Overview of total executed orders

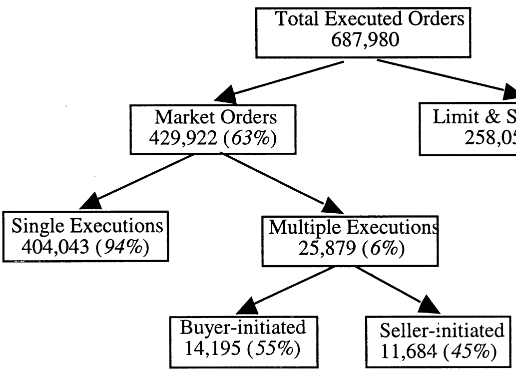

Part B: Breakdown of market orders by size

Size of original market order (in shares)	Number of executions required to fill the order				
	Total orders	One execution	Two executions	Three or more	Partially executed
100	117,267	117,008	216	38	5
	100%	*100%*	*0%*	*0%*	*0%*
200–400	155,646	150,465	4803	327	51
	100%	*97%*	*3%*	*0%*	*0%*
500–900	69,499	64,754	3886	824	35
	100%	*93%*	*6%*	*1%*	*0%*
1000–1900	49,970	43,955	4507	1450	58
	100%	*88%*	*9%*	*3%*	*0%*
> 1900	37,540	27,861	5768	3829	82
	100%	*74%*	*15%*	*10%*	*0%*
All	429,922	404,043	19,180	6468	231
	100%	*94%*	*4%*	*2%*	*0%*

1000 shares. In approximately one-third (8162) of the cases involving multiple executions, the executions were at different prices. This evidence suggests that the reason the original order was broken up was because of insufficient liquidity at the initial price.[11] In summary, Table 1 shows that market orders are rarely split up in execution. When they do occur, these split-ups often reflect stopped orders or orders partially filled at different prices.

3.2. Do trades reflect batched orders?

Thus far, the evidence shows most orders map into single trades. We are also interested in the flip side: whether single trades reflect batched orders. Table 2 presents such an analysis. The unit of observation in this table is trades, not orders. Panel A shows the total NYSE and off-Board trading volume in number of transactions and share volume. As reported in other studies (e.g., Lee, 1993), off-Board executions typically involve smaller-sized trades.

Panel B further breaks down NYSE trades by type and size category. We focus on NYSE executions, because these are the trades for which we have detailed order information.[12] This panel shows the number of orders making up the 'active-side' of each NYSE trades. In constructing this panel, we use the Lee–Ready (1991) algorithm (see Appendix) to identify the active-side of each trade.[13] The active-side is emphasized because it reflects the initiator of the trade, and is the focus of most prior studies on price formation and investor behavior. In contrast, the passive-side facilitates the transaction, but does not instigate it.

Table 2 shows that 76% of the reported trades have a single-participant on the active-side and 24% have multiple-participants. It is clear from this table that a significant amount of batching takes place. In particular, 56% of all large trades – trades in excess of 1900 shares each – have multiple- participants on the active-side. This finding raises questions about the reliability of large trades as a proxy for institutional investor activities. We examine this issue in more detail later.

[11] The ISSM manual notes that sellers are responsible for reporting trades. This has led to concerns that buys may be split up in reporting when they are satisfied by two or more parties on the sell side. However, we find no relation between trade direction and the likelihood of a split execution.

[12] We have no reason to believe batching is more frequent in off-Board executions. Indeed, our sense is that, because of it receives a greater share of the public order flow, batching occurs more frequently at the NYSE than at other market centers.

[13] We use this algorithm because it is a comprehensive technique and can be applied to virtually every trade. We evaluate this technique in the next section.

Table 2
The Mapping of Trades to Orders.

This table presents evidence on how often the active-side of each trade reflects batched orders. Part A is an overview of the total trades in the TORQ database, excluding 1667 out-of-sequence trades. The Lee–Ready (1991) algorithm is used to infer the active-side of each trade. The number of participants corresponds to the reported number of separate orders on the active-side of the trade. Numbers in italics represent the percentage breakdown relative to the total reported in the previous hierarchical level. Part B presents a more detailed analysis of the number of participants on the active-side of each trade, conditional on trade size. All non-opening trades at the NYSE are included, except 330 trades not classifiable by the Lee-Ready algorithm[a]

Part A: Overview of total TORQ trades

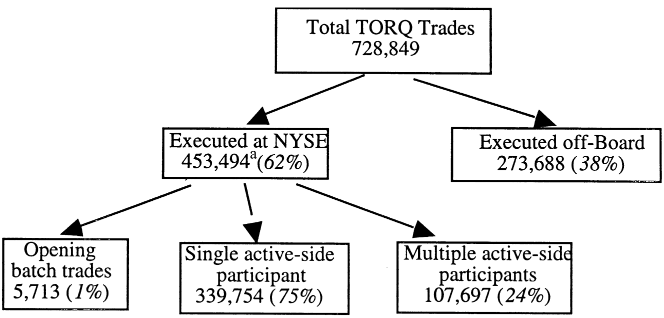

Part B: Breakdown of non-opening trades by size

Trade type or size (shares)	Number of participants on the active-side of the trade					
	Total Trades	One	Two	Three	Four	Five or more
100	77,701	77,701	0	0	0	0
	100%	*100%*	*0%*	*0%*	*0%*	*0%*
200–400	110,027	97,521	11,167	1208	131	0
	100%	*89%*	*10%*	*1%*	*0%*	*0%*
500–900	73,516	59,963	9384	2979	1028	466
	100%	*75%*	*13%*	*4%*	*1%*	*1%*
1000–1900	73,516	54,963	11,084	4133	1758	1578
	100%	*75%*	*15%*	*6%*	*2%*	*2%*
> 1900	112,447	49,666	19,679	13,686	9413	20,003
	100%	*44%*	*18%*	*12%*	*8%*	*18%*
All	447,451	339,754	51,314	22,006	12,330	22,047
	100%	*76%*	*11%*	*5%*	*3%*	*5%*

[a] Excludes 1667 out of sequence trades.

4. Inferring trade direction

Lee and Ready (1991) [LR] examine the question of how to infer the buy-sell direction of an observed trade using intraday data. They propose an algorithm based on trade price relative to the prevailing quote or prior trade prices (see Appendix). This algorithm, or variations of it, has been used in many research contexts: for example, to examine investor reaction to information events (e.g., Lee, 1992; Radhakrishna, 1998), specialist behavior (e.g., Madhavan and Smidt, 1991, 1993), intraday price formation (e.g., Huang and Stoll, 1994; Madhavan et al., 1997), block trading (Holthausen et al., 1987; Madhavan and Cheng, 1997), information asymmetry (Easley et al., 1997), and intermarket price execution (Lee, 1993; Gardner and Subrahmanyam, 1997; Easley et al., 1996). Despite its popularity, the accuracy of the algorithm has not been checked against original order flows.

The algorithm relies on the quote and trade prices immediately prior to a given trade. If quote prices prevailing at the time of the trade are available, LR recommends classifying trade direction with reference to these quotes: trades at, or closer to, the quoted ask (bid) price are deemed buys (sells). For trades at the midpoint of the quoted spread, LR recommends using either a 'tick' or 'zero-tick' test. Specifically, 'uptick' ('downtick') trades – trades transacted at a higher (Lower) price than the previous trade – are deemed buys (sells). A trade at the same price as the previous trade (a zero-tick trade) is classified with reference to the last price change – 'zero-uptick' trades are deemed buys, 'zero-downtick' trades are deemed sells.

LR explains the intuition behind the tick tests. They note that most inside-the-spread executions occur because some liquidity suppliers are unwilling to reveal their proclivity to trade by submitting a formal limit order. The tick test provides a way to infer whether these hidden liquidity suppliers (what LR refer to as 'standing orders') are buyers or sellers. However, the accuracy of the tick tests depends on the order size of these liquidity suppliers. If these suppliers of hidden liquidity have relatively small orders which are filled quickly, then the algorithm can be quite noisy.

With the TORQ dataset, it is possible to trace some trades directly back to the originating order(s) and examine whether the initiator of the trade was a buyer or a seller. Specifically, the information in the audit (CD) file in many cases allows us to identify the active-side of a trade. For these trades, we can observe directly whether it was initiated by a buy or a sell order. In other cases, the supplemental information in the SOD file allows us to determine the trade direction. For example, by tracing limit orders back to the SOD file, we can assess whether they were executable (i.e., within the quoted spread) on arrival.

It is important to understand the limitations of the TORQ approach. Many trades in the CD file cannot be unambiguously classified as to trade direction. For example, trades that involve market orders on both sides, or limit orders on both sides, are difficult to classify. Similarly, it is difficult to determine the direction of batched trades that contain a mixture of market and limit orders on the active side. Finally, stopped orders can introduce trade direction ambiguities. Around 25% of the market orders are stopped, and price improvement among this group occurs frequently.[14] In such instances, it is difficult to judge whether the original market order is passive (i.e., it has helped facilitate a subsequent trade) or active (it triggered the subsequent trade).

Table 3 presents evidence on the magnitude of these problems. To construct this table, we first identify the active and passive side of each trade using the LR algorithm. Over 99% of all market orders can be classified in this manner. For each trade on which order information in available, we then examine the proportion of market orders (or executable limit orders hereafter called 'ELOs') on the active and passive side.[15] This table reports the proportion of total trades that fall into each of three categories: (1) trades in which a market order appears on the active side only, (2) trades in which a market order appears on both sides, and (3) trades in which a market order appears on the passive side only. Panel A provides an overview of trades classifiable by LR for which TORQ order information is available. Panel B reports results for single participant trades and Panel C reports results for multiple participant trades.[16]

Row 2 (Panels B and C) of Table 3 shows that for approximately 40% of the reported executions, market orders or ELOs can be found on both sides of the trade. Some of these trades reflect executions of stopped market orders, others represent market crosses, or the pairing of an ELO against a market order. Given the ambiguity inherent in these situations, it does not seem appropriate to classify such executions as either buys or sells.

Comparing row 1 to row 3, we see that the LR algorithm generally captures the more aggressive side of each trade. For example, from the last column of Panel B, we see that when only one side of the trade has a market order, the LR algorithm correctly identifies the aggressive side 87% of the time.[17] The performance of the LR algorithm is best for single participant trades executed at the quoted bid or ask price (over 90% accurate); it is worse for multiple participant trades inside the quoted spread (around 66% accurate).

In Table 4, we provide a further test of the algorithm by eliminating all trades which are ambiguous as to buy/sell direction based on TORQ data. Our strategy is to include for testing only trades with a clearly-active order on one

---

[14] Hasbrouck et al., 1993 reports 62% of the stopped orders were executed at a price better than the stopped price.

[15] A limit order to buy (sell) is considered an executable limit order (ELO) if the price specified in the order is higher (lower) than the best available bid (offer) price at the time the order is received.

[16] A multiple participant trade is a trade in which the active-side consists of more than one order.

[17] To compute this estimate, divide 8% by (8% + 54%).

Table 3
The Distribution of TORQ Market Orders on the Active and Passive Side of Trades as Classified by the Lee–Ready (1991) Algorithm.

This table compares the buy–sell direction classification of the Lee–Ready (1991) (LR) algorithm to buy–sell information obtained from original orders in the TORQ database. To construct this table, we first use the LR algorithm to identify the active and passive side of each trade. We then examine the proportion of TORQ market orders on the active side and the passive side. Any limit orders that are executable on arrive (ELOs) are also deemed to be market orders. For this purpose, a limit order to buy (sell) is considered an ELO if the price specified in the order is higher (lower) than the best available bid (offer) price at the time the order is received. All trades for which order data is available in the TORQ dataset are included in the analysis. Panel A presents an overview – trades that are classifiable by Lee–Ready, for which TORQ order information is available. Panel B presents results when the active side consists of a single participant. Panel B presents results when the activeside has more than one participant. Table values represent the total number of trades in each category. Percentages are in parentheses

Panel A: Distribution of TORQ market orders – on overview

NYSE trades classifiable by Lee–Ready algorithm (> 99% of all trades)	447,656
Order information not available on TORQ	152,105
Trades analyzed	295,551
Of which, trades with single participants (Panel B)	194,787
And, trades with multiple participants (Panel C)	100,764
Trades analyzed	295,551

Panel B: Single participant trades

	Trades at bid or ask	Midspread Tick test	Midspread zero tick Test	Inside-the spread, but not midspread	Total trades
Market order on active-side only	93,737 *(62%)*	1826 *(27%)*	7118 *(25%)*	1673 *(19%)*	104,354 *(54%)*
Market order on both sides	50,206 *(33%)*	3865 *(56%)*	14,588 *(52%)*	5451 *(61%)*	74,110 *(38%)*
Market order on passive-side only	6701 *(5%)*	1187 *(17%)*	6598 *(23%)*	1837 *(20%)*	16,323 *(8%)*
Total Trades	150,644 *(77%)*	6878 *(3%)*	28,304 *(15%)*	8961 *(5%)*	194,787 *(100%)*

Panel C: Multiple participant trades

	Trades at bid or ask	Midspread tick test	Midspread Zero-tick test	Inside-the spread, but not midspread	Total trades
Market order on active-side only	36,027 *(48%)*	1320 *(33%)*	4987 *(29%)*	1026 *(26%)*	43,360 *(43%)*
Market order on both sides	30,981 *(41%)*	1844 *(45%)*	7668 *(45%)*	1870 *(47%)*	42,363 *(42%)*
Market order on passive-side only	8698 *(11%)*	891 *(22%)*	4367 *(26%)*	1085 *(27%)*	15,041 *(15%)*
Total trades	75,706 *(75%)*	4055 *(4%)*	17,022 *(17%)*	3981 *(4%)*	100,764 *(100%)*

side, and a clearly passive order on the other. For this purpose, we define the following as clearly active or passive:

Clearly active	Clearly passive
Non-stopped market orders	Non-executable limit orders
Executable limit orders	Stopped market-orders
ITS commitments	ITS executions

This approach eliminates all batched trades, trades in which both sides have active (or passive) orders, and trades in which neither side has an active (or passive) order. Of the 447, 781 NYSE trades in the data set, only 93,408 (21%) single participant trades and 36,292 (8%) multiple participant trades qualify for this analysis.

Table 4 (Panel C) shows that, of these 93,408 single participant trades, 80,460 (86%) were executed at the quoted bid or ask price, 11,030 (12%) were executed at the middle of the bid and ask price (midspread), and 1918 (2%) were inside-the-spread but not at the midspread. 93% of these TORQ trades agree with the Lee–Ready classification. As might be expected, the level of agreement is highest for trades transacted at the bid or ask price (98%). Inside the spread, the tick test agrees with the TORQ classification 76% of the time. However, zero-tick tests are less accurate and are only marginally better than a random assignment (60% agreement). The performance of the Lee–Ready algorithm is just as good with multiple participant trades. 88% of TORQ trades agree with the LR classification. The best performance is, once again, for trades at the quote with 92% accuracy. Similarly, the zero-tick test for mid spread trades shows the least accuracy at 61%.

Taken together, the results in Tables 3 and 4 confirm the usefulness of the LR algorithm as a general technique for identifying the more aggressive side of a trade. However, these findings also show that, in a continuous auction, it is often futile to attempt to label either side of a trade as its initiator. The implication of this analysis will depend on the research context in which the LR algorithm is used. While the technique is clearly useful as a general indicator of the direction of incoming order flows, we would caution researchers against using it to identify small samples of individual trades as clear buys or sells. Moreover, if the research design calls for greater accuracy, we recommend omitting the trades classified using a zero-tick test.

Two recent studies provide corroborating evidence on the effectiveness of the LR algorithm. Odders-White (1999) also uses TORQ data, and reports that the algorithm correctly classifies 85% of the trades in her sample. The main difference between this study and the Odders-White study arises from the treatment of trades that involve market orders on both sides (these trades often

Table 4
Performance of the Lee–Ready (1991) algorithm.

This table examines the performance of the Lee–Ready algorithm. Panels A and B summarize results all NYSE trades in TORQ. Panels C and D examine cases where the buy–sell classification using TORQ data is unambiguous. Panel C (D) is restricted to single- (multiple-) participant trades with a clearly active order on one side, and a clearly passive order on the other. Clearly active orders include: non-stopped market orders, executable limit orders, and ITS commitments. Clearly passive orders include: non-executable limit orders, stopped market orders, and ITS executions. Each trade is classified into a separate category based on the trade price relative to the prices of the prevailing quote. Trades at the middle of the quoted spread (Midspread trades) are further classified based on either the tick test (when the price changed from the last trade) or the zero-tick test (when the price is the same as the last trade). Table values represent the total number of trades in each category. Percentages for each trade price category are in parentheses

Panel A: Trades that are classifiable in TORQ

Total NYSE Trades in TORQ		455,161
Less: Out of sequence and opening trades		7380
		447,781
Multiple participants in a trade	145,803	
Less: Ambiguous trades, i.e., both sides active or passive	109,511	36,292
Single Participant in a trade	301,978	
Less: Ambiguous trades, i.e., both sides active or passive	208,570	93,408
TOTAL TRADES ANALYZED		129,700

Panel B: Summary of total NYSE trades in TORQ

	At the bid or ask price	Trade price			All trades
		Midspread Tick test	Midspread Zero-tick test	Inside-the-spread, but not midspread	
Agree	106,552 *(33%)*	3185 *(13%)*	7309 *(9%)*	1912 *(9%)*	118,958 *(27%)*
Not classifiable	213,289 *(66%)*	21,074 *(84%)*	65,006 *(84%)*	18,587 *(87%)*	317,956 *(71%)*
Disagree	4217 *(1%)*	973 *(4%)*	4740 *(6%)*	812 *(4%)*	10,742 *(2%)*
Number of trades	324,058 *(72%)*	25,232 *(6%)*	77,055 *(17%)*	21,311 *(5%)*	447,656 *(100%)*

Table 4 (*continued*)

Panel C: Single participant trades with unambiguous classifications

	At the bid or ask price	Trade price			All trades
		Midspread tick test	Midspread Zero-tick test	Inside-the-spread, but not midspread	
Agree	78,695	2010	5071	1369	87,145
	(97.8%)	*(76.4%)*	*(60.4%)*	*(71.4%)*	*(93.3%)*
Disagree	1765	622	3327	549	6263
	(2.2%)	*(23.6%)*	*(39.6%)*	*(28.6%)*	*(6.7%)*
Number of trades	80,460	2632	8398	1918	93,408
	(86%)	*(3%)*	*(9%)*	*(2%)*	*(100%)*

Panel D: Multiple participant trades with unambiguous classifications

	At the bid or ask price	Trade price			All trades
		Midspread Tick test	Midspread Zero-tick test	Inside-the-spread, but not midspread	
Agree	27,857	1175	2238	543	31,813
	(92%)	*(77%)*	*(61.3%)*	*(67.4%)*	*(88%)*
Disagree	2452	351	1413	263	4479
	(8%)	*(23%)*	*(38.7%)*	*(32.4%)*	*(12%)*
Number of trades	30,309	1526	3651	806	36,292
	(84%)	*(4%)*	*(10%)*	*(2%)*	*(100%)*

arise when a specialist ‘stops’ a market order). The Odders-White study assigns a trade direction to these executions by treating the later arriving order as the ‘initiating side’ of the trade. For reasons discussed earlier, this assignment strikes us as being too arbitrary. We believe the concept of an ‘initiator’ is not meaningful when both sides of the trade are market orders. Accordingly, we treat these trades as unclassifiable. Our results are otherwise quite similar to those reported in the Odders-White study.

Another related study is Lightfoot et al. (1999). This study focuses on the effect of order preferencing on transaction costs, and does not address the accuracy of the LR algorithm directly. However, using audit trail data from October 28 to November 1, 1996, the authors provide some indirect evidence on the frequency of misclassifications in minimum variation markets (i.e., when trading is at a $\frac{1}{8}$ spread). Specifically, the authors show that approximately 25% of the time, the NYSE provides price improvement even when the spread is $\frac{1}{8}$th. In other words, when the spread is only $\frac{1}{8}$th, market buys are executed at the bid (or market sells are executed at the ask) approximately 25% of the time.

After taking into account the effect of matched market orders, these results are consistent with those reported in our study and the Odders-White study. In other words, allowing for batched trades and market order crosses, we also find approximately 20–25% of the market buys (sells) in minimum variation markets being executed at the bid (ask). Our analysis shows that these results are due primarily to 'stopped' orders at the NYSE, or executions in which both sides contain market orders.

5. Inferring trader type

5.1. Individual and institutional trading activities

Several prior studies used a trade size proxy to distinguish individual and institutional trades. For example, Lee (1992) assumed that small dollar-valued trades (less than or equal to $10,000) are initiated by individuals, Cready (1988) and Cready and Mynatt (1991) assumed trades in excess of 900 shares are initiated by institutional traders. In this section, we examine the validity of these assumptions and develop a general technique for separating the activities of these two groups of traders.

Our approach relies on the TORQ database's account type code to classify trader groups. Trades coming from individuals are coded with an 'I' account code. This code is associated with orders that brokerage firms have indicated as being initiated by individual investors. Brokers are encouraged to provide this information where appropriate, so that their clients might receive preferential routing through the Individual Investor Express Delivery Service (IIEDS).[18] Trades coded in TORQ as agency trades, program trades (other than those identified as initiated by individuals), and member trades (trades by exchange members on their own account) are classified as institutional.[19] We exclude all

[18] This reporting is not independently monitored, but there is little reason to suspect broker misrepresentation. The total volume of 'individual' trades in the TORQ data base is similar to the total volume of individual trades estimated by Securities Industry Association (SIA) for this time period. The SIA estimate is based on an independent data source (regulatory filings by institutional investors).

[19] In theory, individual investors may participate in program trading. However, in our sample, only 23 program trades (0.000% of total trades, and 0.003% of all program trades) were initiated by individuals.

batched trades that have a mixture of trader types and all trades involving uncoded orders.[20]

Table 5 reports the trading activity of institutions and individuals by firm size. A total of 310,292 NYSE trades are clearly identifiable by trader type. Column 2 shows that 43% of these trades (30% of total trades) are initiated by individuals, and 57% (39% of total trades) are initiated by institutions.[21] Columns 3–5 provide the same information for three subsets of firms, classified by market capitalization on October 31, 1990. Column 3 shows that most of the trades in large firms (60% of identified trades) are initiated by institutional investors. In contrast, Column 5 shows that most of the trades in small firms (67% of identified trades) are initiated by individual investors.

Table 5 also shows that the average dollar value of the trades initiated by individuals and institutions vary by firm size. Both groups tend to engage in higher dollar value trades in larger firms. For example, the mean (median) institutional trade in large firms is \$103,163 (\$29,150), while the mean (median) institutional trade in small firms is only \$10,306 (\$3062). The pattern is similar in individual trades. This evidence suggests that we may be able to better separate the two trader groups by conditioning the trade size cutoff on firm size. We will explore this possibility in Section 5.3.

5.2. The reliability of trade-size proxies

Two types of trade-size proxies have been used in prior studies: dollar-based cutoffs, and share-based cutoffs. Dollar-based cutoffs are conceptually superior because they incorporate differences in stock prices. All trade size proxies operate on the premise that individuals have tighter wealth-constraints, and therefore initiate smaller trades. Because share-based cutoffs ignore differences in stock prices, they should lead to noisier classifications.

However, dollar-based cutoffs must contend with a different problem. Specifically, they are sensitive to price changes, so that small price movements may cause different round-lot trades to be classified as individual-initiated trades. For example, if a stock is trading at \$20, a \$10,000 cutoff would classify all trades of 100–500 shares as small trades. If the stock moves to 20 1/8, suddenly only 100–400 share trades will be classified as 'small'. Because individual investors are not likely to immediately adjust the size of their trades due to such small price changes, we will observe an immediate drop in the individual investor activities that is artificially induced.

---

[20] If all participants in a batched trade are either individuals or institutions, the trade is included in the analysis. The uncoded trades typically reflect orders placed through floor brokers or specialists' principal trades. For a detailed description of specific TORQ codes see Radhakrishna (1998).

[21] The active-side is determined for this purpose using the Lee–Ready (1991) algorithm.

Table 5
Individual and institutional activities by firm size.

This table provides descriptive information on the trading activities of individuals and institutional investors by firm size. The sample consists of all trades clearly identifiable as having been initiated by either an institutional investor or an individual investor. The total of 310,292 trades represents 69% of all NYSE trades in the database, and excludes daily opening trades, out-of-sequence trades, uncoded trades, as well as batch trades involving both individual and institutional investors. The Lee–Ready (1991) algorithm is used to identify the active-side (or initiator) of each trade. Table values represent the number of trades and the average trade size in each firm size category. The relative size of the firms are determined by their market capitalization on October 31, 1990

	All firms	Large firms	Medium firms	Small firms
Institutional trades				
Total trades	176,086	144,994	25,733	5359
(as % trades analyzed)	*(57%)*	*(60%)*	*(49%)*	*(33%)*
Dollar value per trade mean	$89,862	$103,163	$31,406	$10,306
(median)	($23,250)	($29,150)	($8,375)	($3,062)
No. of shares per trade mean	2378	2390	2321	2334
(median)	*(700)*	*(700)*	*(500)*	*(900)*
Individual trades				
Total trades	134,206	95,886	27,302	11,018
(as % trades analyzed)	*(43%)*	*(40%)*	*(51%)*	*(67%)*
Dollar value per trade mean	$15,779	$19,116	$8,884	$3,824
(median)	*($6000)*	*($7125)*	*($4600)*	*($1550)*
No. of shares per trade mean	579	482	781	917
(median)	*(200)*	*(200)*	*(400)*	*(500)*
Total trades analyzed	310,292	240,880	53,035	16,377
	(100%)	*(100%)*	*(100%)*	*(100%)*

To mitigate this price sensitivity problem, Lee (1992) proposed a firm-specific dollar-based trade-size proxy computed as follows:

1. Obtain a closing price for the firm during the study period.
2. Compare that closing price to $10,000 and determine the largest number of round lot shares that is less than or equal to $10,000. Trades transacted during the year for this firm at this number of shares or less will be deemed small trades.

We test a similar procedure based on December 31, 1990 prices and various dollar-based cutoffs. For example, using a $10,000 cutoff, if the December 31, 1990 price of a share is \$30, then all trades for this firm involving 300 shares or less will be deemed individual trades. Similarly, for a \$40 stock, all trades of 200 shares or less will be deemed individual trades.

Table 6 examines the efficacy of this dollar-based technique. Panel A reports the percentage of total institutional (individual) trades captured above (below) a given trade-size cutoff. This panel shows that 66% of all trades by individuals are below \$10,000 in value, and 84% of individual trades are below \$20,000 in value. Similarly, 53% of institutional trades are above \$20,000 in value, and 67% are above \$10,000. The evidence shows that while it's possible to separate these two groups using a single cutoff value, the result is a noisy proxy for the true dichotomy. For example, a \$10,000 value cutoff captures approximately two-third of the individual trades in the small trade stratum, and two-thirds of the institutional trades in the large trade stratum. Moving to an alternative cutoff improves the classification accuracy of one trade size stratum at the expense of the other.

Panel B reports the Type I and Type II error probabilities associated with each cutoff value. For this panel, we define a Type I error as the probability that an individual trade is erroneously classified in the large trade stratum. Similarly, we define a Type II error as the probability that an institutional trade is erroneously classified in the small trade stratum. For example, \$10,000 cutoff results in a 14% probability of Type I error and a 19% probability of Type II error. As expected, moving to larger cutoff values reduces Type I error, but increases Type II error.

Table 6 shows that error probabilities increase or decrease monotonically in trade value cutoffs. This observation suggests that we can enhance classification accuracy by introducing a 'buffer zone' of medium-sized trades. For example, we can classify trades of \$5000 or less as individual trades, and trades of \$50,000 or more as institutional. Panel B shows this strategy results in only a 2% probability Type I error and a 10% probability of Type II error. Further, Panel A shows that the \$5000 cutoff is able to capture 39% of all individual trades and the \$50,000 cutoff captures 35% of all institutional trades. We expect that, in many research applications, this three-tiered approach will result in a cleaner separation between individual and institutional trades.

Table 7 reports a similar analysis using share-based rather than dollar-based cutoffs. Similar to Table 6, we find that a share-based trade size proxy is able to separate the activities of individuals and institutions. However, as expected, the classifications are noisier than with dollar-based cutoffs, and the error probabilities are generally higher. For example, a small trade proxy using a 200 share cutoff captures 51% of all individual trades, but has a 15% probability of Type II error. Similarly, a large trade proxy using a 1,900 share cutoff captures only 28% of all institutional trades, but has a 3% probability of Type I error. Comparing Tables 6 and 7, we see that the dollar-based approach is generally more effective than the share-based approach.

Table 6
The Ability of Dollar-based Trade-Size Cutoffs to Discriminate Between Individual and Institutional Activities.

This table examines the ability of various dollar-based cutoffs in discriminating between the trading activities of individual and institutional investors. Panel A reports the proportion of total institutional (individual) trading activity captured above (below) each cutoff value. Panel B reports the error probabilities associated with each cutoff value. For this panel, a Type I error is the probability that an individual trade is erroneously included with the large trades, and a Type II error is the probability that an institutional trade is erroneously included with the small trades

Panel A: The percentage of total institutional (individual) trades captured above (below) a given trade-size cutoff

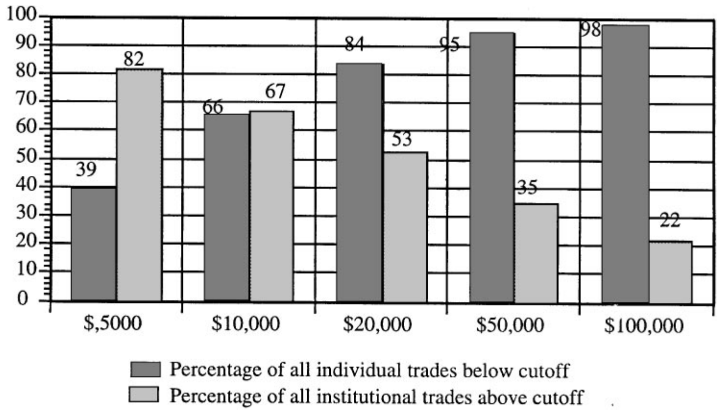

Panel B: The probability of Type I and Type II errors at each cutoff value

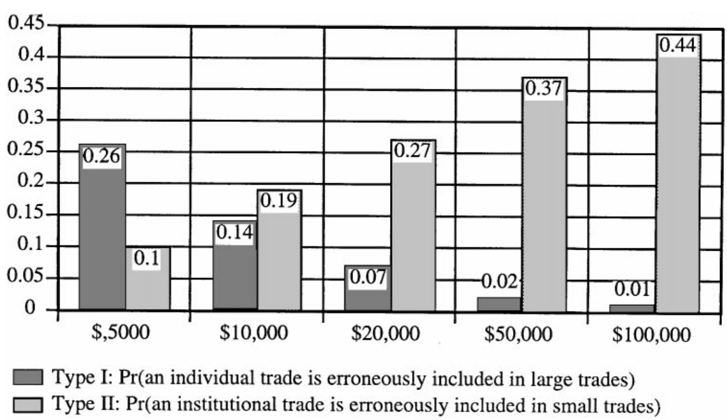

Table 7
The ability of share-based trade-size cutoffs to discriminate between individual and institutional activities.

This table examines the ability of various share-based cutoffs to discriminate between the trading activities of individual and institutional investors. Panel A reports the proportion of total institutional (individual) trading activity captured above (below) each cutoff value. Panel B reports the error probabilities associated with each cutoff value. For this panel, a Type I error is the probability that an individual trade is erroneously included with the large trades, and a Type II error is the probability that an institutional trade is erroneously included with the small trades

Panel A: The percentage of total institutional (individual) trades captured above (below) a given trade-size cutoff

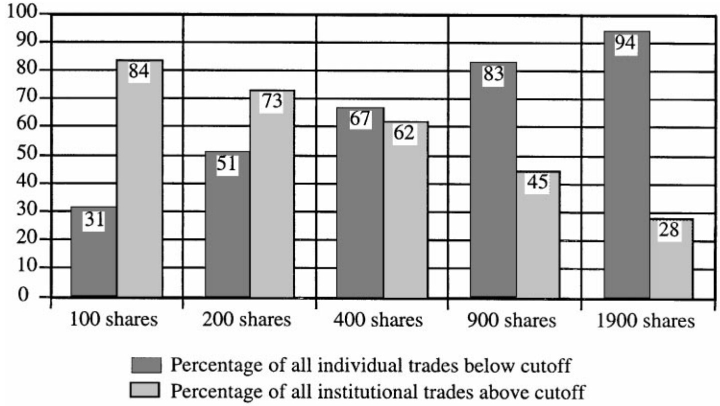

Panel B: The probability of Type I and Type II errors at each cutoff value

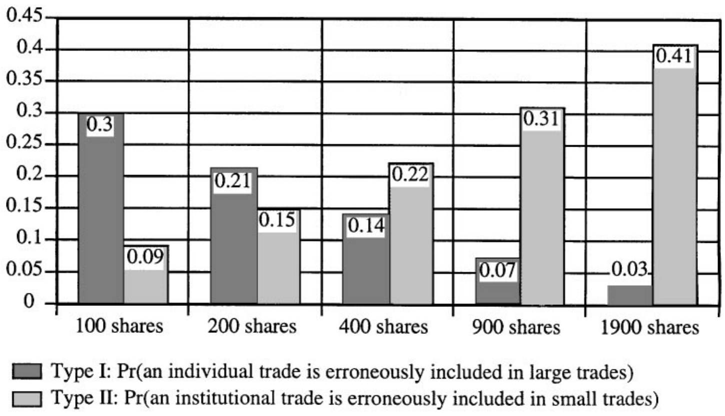

5.3. Sensitivity to firm size

Earlier results reported in Table 5 show both individuals and institutions scale down their average trade size when trading in smaller firms. In this subsection, we evaluate the effect of this behavior on the efficacy of various cutoff values. In Table 8, we report the classification accuracy of various dollar-based cutoff values conditional on firm size. This table shows that significant gains in classification accuracy can be made by conditioning the cutoffs on firm size.

For example, a \$5000 small trade proxy captures 31% of individual trades in large firms, with an associated 8% probability of Type II error. However, the same \$5000 cutoff will capture 81% of all individual trades in small firms, with a 20% probability of Type II error. These numbers suggests that while a \$5000 cutoff may be appropriate for large firms, it is probably too high for small firms because in small firms, trades of less than \$5000 are too often initiated by institutions. Indeed, a \$2500 cutoff seems to do a better job for the small firms in our sample, capturing 63% of all individual-initiated trades with only a 15% probability of Type II error.

The argument is analogous if a researcher is interested in isolating institutional trades using a large trade size proxy. For example, Panel B shows that if a researcher has a tolerance level of 2% for Type I errors (the probability an individual trade is included erroneously in the large trade stratum), the appropriate large trade proxy for large firms is \$100,000. However, given the same tolerance level, the appropriate large trade proxy for medium-sized firms is \$50,000, and for small firms it is only \$20,000. Table 8 shows that by using these cutoff, a researcher can minimize the likelihood of a misclassification while still capturing a substantial portion of the total trades in each trader group.

Table 9 reports the results of a similar analysis using share-based cutoffs. Compared to Table 8, these results show that share-based cutoffs generally have lower discriminatory power than dollar-based cutoffs. However, Table 9 shows that, particularly for small firms, it is possible to minimize the probability of a Type II error by using a very small share cutoff. For example, Panel B shows that trades of 100 shares are very unlikely to be initiated by institutions even in small firms (the corresponding probability of Type II error is only 5%). At the same time, Panel A shows that using trades of 100 shares, a researcher can still capture 18% of all trades initiated by individual investors in small firms. Therefore, if the research results are particularly sensitive to Type II errors, a small share-based cutoff may be the most effective approach.

Another approach is to discriminate between individual and institutional trades by using a constant percentage of the largest (smallest) trades in each firm. We conducted a test in which the largest (smallest) 20% of the trades in each firm were used to proxy for institution (individual) activities. The results (not reported) show that the constant-percentage technique leads to noisier

Table 8
The effect of firm size on the ability of dollar-based trade-size proxies to discriminate between individual and institutional trades.

This table examines the ability of various dollar-based cutoffs to discriminate between the trading activities of individual and institutional investors, conditional on firm size. Large, medium and small firms are those ranked in the upper, medium, and lower third, respectively, based on market capitalization as of October 31, 1990. Panel A reports the proportion of total institutional (individual) trading activity captured above (below) each cutoff value. Panel B reports the error probabilities associated with each cutoff value. For this panel, a Type I error is the probability that an individual trade is erroneously included with the large trades, and a Type II error is the probability that an institutional trade is erroneously included with the small trades

Panel A: The percentage of total institution (individual) trades captured above (below) a given trade-size cutoff

Trade size cutoffs	Percentage of all institution-initiated trades captured above cutoff value			Percentage of all individual-initiated trades captured below cutoff value		
	Large firms	Medium firms	Small firms	Large firms	Medium firms	Small firms
$2,500	97%	75%	55%	5%	30%	63%
$5,000	87%	60%	38%	31%	54%	81%
$10,000	73%	45%	22%	61%	77%	92%
$20,000	59%	30%	16%	81%	92%	97%
$50,000	39%	15%	4%	93%	98%	100%
$100,000	25%	7%	1%	97%	99%	100%

Panel B: The probability of Type I and Type II errors at each cutoff value

Trade size cutoffs	Pr(Type I error) = Pr.(An individual trade is erroneously included with large trades)			Pr(Type II error) = Pr. (An institution trade is erroneously included with small trades)		
	Large firms	Medium firms	Small firms	Large firms	Medium firms	Small firms
$2,500	0.38	0.36	0.25	0.02	0.12	0.15
$5,000	0.28	0.24	0.13	0.08	0.19	0.20
$10,000	0.16	0.12	0.05	0.16	0.27	0.26
$20,000	0.08	0.04	0.02	0.25	0.34	0.29
$50,000	0.03	0.01	0.00	0.37	0.41	0.31
$100,000	0.01	0.00	0.00	0.45	0.45	0.32

inferences than either the share-based or dollar-based method. The small trade proxy is particularly noisy. For example, a size stratum consisting of the smallest 20% of reported trades will capture close to 90% of all trades actually initiated by individuals. However, the probability that institutional trades are

Table 9
The effect of firm size on the ability of share-based trade-size proxies to discriminate between individual and institutional trades. This table examines the ability of various share-based cutoffs to discriminate between the trading activities of individual and institutional investors, conditional on firm size. Large, medium and small firms are those ranked in the upper, medium, and lower third, respectively, based on market capitalization as of October 31, 1990. Panel A reports the proportion of total institutional (individual) trading activity captured above (below) each cutoff value. Panel B reports the error probabilities associated with each cutoff value. For this panel, a Type I error is the probability that an individual trade is erroneously included with the large trades, and a Type II error is the probability that an institutional trade is erroneously included with the small trades

Panel A: The percentage of total institutional (individual) trades captured above (below) a given trade-size cutoff

Trade size cutoffs	Percentage of all institution-initiated trades captured above cutoff value			Percentage of all individual-initiated trades captured below cutoff value		
	Large firms (%)	Medium firms (%)	Small firms (%)	Large firms (%)	Medium firms (%)	Small firms (%)
100 shrs	85	78	86	36	20	18
400 shrs	63	57	66	74	52	47
900 shrs	46	42	49	88	74	67
1900 shrs	28	26	28	96	91	88

Panel B: The probability of Type I and Type II errors at each cutoff value

Trade size cutoffs	Pr(Type I error) = Pr.(An individual trade is erroneously included with large trades)			Pr(Type II error) = Pr. (An institution trade is erroneously included with small trades)		
	Large firms	Medium firms	Small firms	Large firms	Medium firms	Small firms
100 shrs	0.26	0.41	0.55	0.09	0.11	0.05
400 shrs	0.10	0.24	0.36	0.23	0.21	0.11
900 shrs	0.05	0.14	0.22	0.33	0.28	0.17
1900 shrs	0.02	0.04	0.08	0.43	0.36	0.24

erroneously included in small trades (i.e., the probability of Type II error) is also very high, at close to 30%.

In summary, our results provide support for the use of trade size as a proxy for individual and institutional trading. We show that misclassifications can be reduced by using a buffer zone of medium-sized trades, and that classification accuracy can be further enhanced by conditioning the cutoff values on firm size. We find that dollar-based cutoff values generally have greater discriminatory power than cutoffs based on number of shares. The appropriate cutoffs depend, of course, on the research context and the sensitivity of the results to Type I and

Type II errors. However, Tables 8 and 9 offer simple benchmarks for calibrating these error probabilities.

6. Conclusion

In this paper, we have examined the efficacy of several inference techniques commonly used to discern investor behavior from observed trade and quote data. Our investigation suggests that the frequency, size and direction of observed trades provide a reasonable basis for evaluating the incoming flow of market orders. Specifically, we find that orders are rarely split up in execution but batching is much more common. We show that many trades cannot be unambiguously classified as buys or sells. However, for those trades that can be so classified, the Lee–Ready algorithm performs quite well. Finally, we find that a trade size proxy has some ability to discern between individual and institutional trading activities, and we make several recommendations for improving inferences using this technique.

Our main purpose is not to test specific hypotheses. Rather, we aim to contribute toward future empirical market microstructure research by providing new evidence on the accuracy of some established inference methods. It is our hope and expectation that the evidence presented here will help future scholars working with trade and quote data to develop more powerful research designs. In particular, we expect the error probabilities associated with alternative methods of discriminating between individual and institutional trading to be useful for future researchers interested in the distinguishing the trading behavior of these two constituencies.

Finally, we believe it is important to highlight the potential effect of two changes in market dynamics since our sample period. First, trading volume has increased tremendously in the 1990s. This is likely to lead to more frequent batching of orders and an increase in the proportion of orders that are not clearly buyer or seller initiated. The exact extent of the noise introduced by increased trading is difficult to quantify without more recent order flow data.

Second, U.S. equity markets now trade in smaller price increments (the NYSE adopted trading in 'teenies' in 1997). This development may also lead to changes in trading strategy and behavior that impact the validity of the assumed linkages. In general, a reduction in the economic significance of the liquidity premium will tend to blur the line between suppliers and demanders of immediacy. As the price of immediacy declines, it becomes less meaningful to speak in terms of 'buyer' or 'seller' initiated trades. While it is difficult to estimate the net effect of these developments on our results, researchers interested in using our findings with more recent data should be mindful of these potential problems.

7. For further reading

The following reference is also of interest to the reader: Madhavan and Sofianos, 1998.

Appendix A. The Lee–Ready (1991) Algorithm

Lee and Ready (1991) proposed the following algorithm for inferring trade directions. Only NYSE issued quotes which are BBO-eligible were used (a quote is BBO-eligible if it qualifies for the National Association of Security Dealers' Best-Bid-Or-Offer calculation):

1. *Current Quote Match* – If the trade price is at the bid or ask, and the current quote was not revised within the last 5 s, then the direction of the trade is determined by the current quote (i.e., a buy of it is at the ask and a sell if it is at the bid).
2. *Delayed Quote Match* – If the current quote is less than 5 s old, it is ignored and the trade price is compared to the bid and ask prices of the previous quote.
3. *Outside the Spread* – If the trade price, when compared to the quote in either 1 or 2, is greater than the ask (less than the bid), then the transaction is deemed a buy (sell) – this is extremely rare.
4. *Tick Test* – If the trade is at the midpoint of the spread, or if a BBO-eligible quote is not available, the tick test is used to determine trade direction. In other words, if the last price change was positive (negative), then this trade is deemed a buy (sell). All out-of-sequence trades are ignored in updating price changes.
5. *Proximity to Bid/Ask* – If a trade is between the spread but not at the midpoint, then the trade is classified according to its proximity to the bid or ask price. Trades at prices above the midpoint are classified as buys and trades at prices below the midpoint are classified as sells.
6. *Indeterminable* – This classification is assigned to a trade when none of the above conditions apply. Specifically, it applies to the first trade of the year for each firm and any trade which is reported out-of-sequence.

References

Bhattacharya, N., 1998. Investors' trade size and informedness around earnings announcements: an intraday analysis. Working paper, University of Georgia.

Cready, W.M., 1988. Information value and investor wealth: the case of earnings announcements. Journal of Accounting Research 26, 1–27.

Cready, W.M., Mynatt, P.G., 1991. The information content of annual reports: a price and trading response analysis. Accounting Review 66, 291–312.

Gardner, A., Subrahmanyam, A., 1997. Multi-market trading and the informativeness of stock trades: an empirical intraday analysis. Journal of Economics and Business 49, 515–531.

Easley, D., Keifer, N., O'Hara, M., 1996. Cream-skimming or profit-sharing? The curious role of purchased order flow. Journal of Finance 51, 811–833.

Easley, D., Keifer, N., O'Hara, M., 1997. One day in the life of a very common stock. Review of Financial Studies 10, 805–835.

Hand, J.R.M., 1990. A test of the extended functional fixation hypothesis. Accounting Review 65, 740–763.

Hasbrouck, J., 1992. Using the TORQ database. NYSE and NYU working paper.

Hasbrouck, J., Sosebee, D., 1992. Orders, trades, reports and quotes at the New York stock exchange. NYSE working paper #92–01.

Hasbrouck, J., Sofianos, G., Sosebee, D., 1993. New York stock exchange systems and trading procedures. NYSE working paper #93–01.

Holthausen, R.W., Leftwich, R.W., Mayers, D., 1987. The effect of large block transactions on security prices: a cross-sectional analysis. Journal of Financial Economics 19, 237–267.

Huang, R.D., Stoll, H.R., 1994. Market microstructure and stock return predictions. Review of Financial Studies 7, 179–213.

Lakonishok, J., Maberly, E., 1990. The weekend effect: trading patterns of individual and institutional investors. Journal of Finance 45, 231–243.

Lease, R., Lewellen, W., Schlarbaum, G., 1974. The individual investor: attributes and attitudes. Journal of Finance 29, 313–333.

Lee, C.M.C., 1992. Earnings news and small traders: an intraday analysis. Journal of Accounting and Economics 15, 265–302.

Lee, C.M.C., 1993. Market integration and price execution for NYSE-listed securities. Journal of Finance 48, 1009–1038.

Lee, C.M.C., Ready, M., 1991. Inferring trade directions from intraday data. Journal of Finance 46, 733–746.

Lee, C.M.C., Shleifer, A., Thaler, R., 1991. Investor sentiment and the closed-end fund puzzle. Journal of Finance 46, 75–109.

Lightfoot, L.E., Martin, P.G., Peterson, M.A., Sirri, E.R., 1999. Order preferencing and market quality on United States equity exchanges. Working paper, Securities and Exchange Commission, Washington, DC.

Madhavan, A., Cheng, M., 1997. In search of liquidity: block trades in the upstairs and downstairs markets. Review of Financial Studies 10, 175–203.

Madhavan, A., Sofianos, G., 1998. An empirical analysis of NYSE specialist trading. Journal of Financial Economics 48, 189–210.

Madhavan, A., Richardson, M., Roomans, M., 1997. Why do security prices fluctuate? A Transaction-level analysis of NYSE stocks. Review of Financial Studies 10, 1035–1064.

Madhavan, A., Smidt, S., 1991. A bayesian model of specialist pricing. Journal of Financial Economics 31, 99–134.

Madhavan, A., Smidt, S., 1993. An analysis of daily changes in specialist incentives and quotations. Journal of Finance 48, 1595–1628.

Odders-White, E., 1999. On the occurrence and consequences of inaccurate trade classifications. Working paper, University of Wisconsin – Madison.

Radhakrishna, B., 1998. Investor behavior around earnings announcements. Working paper. University of Minnesota.

Ritter, J.R., 1988. The buying and selling behavior of individual investors at the turn of the year. Journal of Finance 43, 701–717.

Shiller, R.J., Pound, J., 1989. Survey evidence on diffusion of interest and information among investors. Journal of Behavior and Organization 12, 47–66.

Walther, B.R., 1997. Investor sophistication and market earnings expectation. Journal of Accounting Research 35, 157–179.

Yunker, J.A., Krehbiel, T.L., 1988. Investment analysis by the individual investor. Quarterly Review of Economics and Finance 28, 90–101.