

The information content of the limit order book: evidence from NYSE specialist trading decisions[☆]

Lawrence E. Harris^a, Venkatesh Panchapagesan^{b,*}

^a*Fred V. Keenan Chair in Finance, Marshall School of Business at USC, USA*

^b*Olin School of Business, Washington University, St. Louis, MO 63130, USA*

Available online 15 September 2004

Abstract

Specialists compete with limit order traders to provide liquidity at the New York Stock Exchange. Since specialists see all system limit orders, they enjoy a unique advantage in this competition. We examine whether the limit order book is informative about future price changes and whether specialists use this information when trading. We use order quantities as well as option values to capture the information content of the limit order book. Our analyses consider three actions specialists can take when a market order arrives: stop the order, immediately fill the order at the quoted price, or immediately fill the order at an improved price. Using SuperDOT limit orders in the TORQ database, we find that the limit order book is informative about future price movements. We also find that specialists use this information in ways that favor them (and sometimes the floor community) over the limit order traders. The results are more evident for active stocks where the competition between specialists and limit order traders is more intense. Moreover, we find strong evidence that specialists in

[☆]We thank Jeff Bacidore, Jeff Benton, Mark Carhart, Amy Edwards, Andrew Ellul, Larry Glosten, Marc Lipson, Ananth Madhavan, Gideon Saar, Robert Seijas, George Sofianos, Avaniidhar Subrahmanyam (the editor) and three anonymous referees, as well as participants at the 2000 American Finance Association Meeting, 1999 NASDAQ-Notre Dame Conference on Market Microstructure, 1999 NBER Meeting on Market Microstructure, 1999 European Finance Association Meeting, the University of Missouri at Columbia Finance Seminar and the NYSE-JB seminar for their comments and discussion. The views expressed herein are those of the authors and do not necessarily reflect the views of the Commission, or the staff of the Commission.

*Corresponding author.

E-mail addresses: lharris@usc.edu (L.E. Harris), panchapagesan@olin.wustl.edu (V. Panchapagesan).

lower-priced stocks are less likely to initiate such actions because of the large relative tick size.
© 2004 Elsevier B.V. All rights reserved.

Keywords: Specialist trading; Limit order book; Market transparency; Limit order option values; Minimum price increments; Tick sizes

Specialists at the New York Stock Exchange (NYSE) are broker-dealers. They trade as brokers for their clients and as dealers for their own accounts. The dual nature of their operations attracted regulatory interest as early as the 1930s (see [Abolafia, 1996](#)). In particular, concerns have been raised about the extent to which specialists may profit from opportunities and information they obtain when acting as brokers for other traders.

To mitigate potential conflicts of interest, the NYSE has adopted several rules. For example, the public priority rule prohibits Exchange members, including specialists, from trading on the floor ahead of public traders at the same price. Although these rules regulate what specialists can and cannot do, specialists still enjoy an informational advantage over other traders. They see all SuperDot orders whereas most other traders only see the market quotation.

Whether specialists can trade profitably on this order book information given the constraints Exchange rules impose upon them is an empirical question. Using the TORQ database, we examine for the first time whether aggregate order information is indeed informative as academic and regulatory studies suggest, and whether specialists use this information to their benefit. We find evidence for both points.

Specialists deposit public orders that are not filled immediately into their limit order books. These orders consist primarily of limit orders with limit prices that are away from the market. Orders stay in the book until they are filled, until they expire, or until they are cancelled.

Until June 1991, only specialists (and Exchange officials) could see the contents of the order books. Since then, Exchange rules require that specialists share general information about their books with other floor traders on an informal basis when asked¹. In practice, the books remain largely hidden from most traders in the market. This practice is in contrast with some automated markets such as the Paris Bourse and the Toronto Stock Exchange that display their books.

Traders who know the contents of the book may exploit this information at least two ways. First, an asymmetry in the order book may reflect trader sentiment or the presence of well informed traders. In either event, order book asymmetry may indicate the likely direction of future price changes. Traders who can see the book therefore might want to trade in front of its heavy side. Second, even if asymmetries

¹Recently, the NYSE initiated a pilot program in which stocks were flagged when there were orders for at least 20,000 shares within 15 cents of the best bid or ask. Beginning in early 2002, the exchange also started selling information (\$ 5,000 per month for data feeds) on its limit order book through its new system called the NYSE OpenBook™. The system provides aggregate order book information at different prices and is updated every 10 seconds. The specialist still retains an advantage in that he gets to see the individual orders that make up the book.

in the book are not correlated with future price changes, asymmetries may still allow traders to profit from the quote-matching strategy discussed in Amihud and Mendelson (1990) and Harris (1990). Traders who use this strategy attempt to extract option values from standing limit orders. They do so by trading ahead of the heavy side of the book. If prices subsequently move in their favor, they profit to the full extent of the price rise. If prices move against them, they may limit their losses by trading with the heavy side of the book. Quote-matchers may profit from this strategy if they can submit market orders faster than other traders can cancel their limit orders. Whether the limit order book conveys information about future price changes, or merely information about the trading options inherent in limit orders, traders who know the book may trade more successfully than those who do not.

Since NYSE specialists know more about their order books than all other traders, they have an informational advantage over traders who compete with them to offer liquidity. Whether this advantage produces significant trading profits is uncertain. Recent SEC investigations suggest that specialists do indeed profit at the expense of investors because of their unique positional advantage on the floor of the Exchange (see Solomon and Craig, 2003). The NYSE maintains that specialists provide liquidity only when none is available. The low specialist participation rate supports this view: most trades at the NYSE are between public traders with little specialist intervention.² Not surprisingly, the participation rate is lowest for actively traded stocks with narrow spreads. The public order precedence rule prevents specialists from competing effectively against public traders in these stocks. The constraints on specialist trades may make it difficult for them to profit from their informational advantages.

Empirical analyses of specialist trading profits have been limited because the relevant data are not readily available. The few empirical studies that examine specialist profits suggest that specialists can predict short-term price movements to some extent (see Sofianos, 1995). These studies, however, do not explicitly identify how they make such predictions.

Using the NYSE trades, orders, reports and quotes (TORQ) database, we provide evidence that specialists profit from knowing their limit order books. We find that the limit order book is informative about future price movements, and that specialists use this information when deciding how to respond to market order arrivals. We examine three possible responses: specialists can stop the order,³ they

²The participation rate is the ratio of the dollar value of all specialist trades to the total dollar trading volume. Up until the June 1997 decrease in the minimum price variation to one-sixteenth, the rate had been steady near 17–18 percent. (see Madhavan and Sofianos, 1998) Following the change to a sixteenth tick, the participation rate rose to around 26–30 percent. (see NYSE Fact Book 1999.) The recent change to a penny tick further increased the participation rate (see Battalio and Jennings, 2002).

³A specialist stops an incoming market order by guaranteeing execution at a price not worse than the current quote. While an order is stopped, it must be matched with any incoming market orders on the other side. If no such orders arrive, the specialist must execute the order for his own account within 30 minutes. The specialist may also elect to execute the order any time before then. Ready (1999) analyzes various aspects of specialist order-stopping behavior. During our sample period, specialists could not stop an order in a one-tick market. This rule was subsequently relaxed under certain conditions.

can execute the order at the best quoted price, or they can fill the order at an improved price. Since Exchange rules and market conditions often limit these choices, we focus mainly on situations in which specialists have discretion in their choice of response. In addition, we exclude specialist actions that could be based on information not necessarily from the limit order book. Interestingly, specialist trading decisions depend on limit order option values but less on limit order sizes.

We find that specialists are more likely to stop, trade immediately, and improve prices when they can trade in front of an order book that is heavy on the same side. Specialists are more likely to buy for their own account when the book is heavy on the buy side and sell when it is heavy on the sell side. This evidence is stronger for more active stocks than for less active stocks. Our evidence also supports the suggestion in [Harris \(1990\)](#) that relative tick size (tick size as a fraction of price) affects quote-matching strategies.

These results are particularly relevant to investors who seek liquidity in the US markets following the recent switch to decimal pricing. The decrease in tick size from one-sixteenth dollar to one penny greatly decreased the economic importance of the public precedence rule. Not surprisingly, specialist participation rates increased substantially. Many public traders complain about “penny-jumping” by specialists and other traders (see [Ip, 2001](#)). Partly to address such complaints, the NYSE has started selling aggregate order book information on a real-time basis to the public. Our results suggest that the availability of these data may only partly mitigate the informational advantage enjoyed by the specialist since the data do not provide information about option characteristics such as duration and price position of individual limit orders. Although our study is based on data collected almost 10 years before the decimalization, the results clearly demonstrate the importance of the issue.

Though we show that specialists exploit information from their limit order books, we do not address the broader question of the optimal degree of order transparency. To do so would require that we examine the costs of affirmative obligations that the NYSE imposes upon specialists. Specialists assume these costly obligations in part because they can profit from exploiting information in their order books. Greater transparency would decrease the specialists’ informational advantage over other traders and thereby make their trading less profitable.⁴ Any analysis of the optimal degree of order transparency at the NYSE therefore must also consider the costs and benefits of the specialists’ affirmative obligations, which is beyond the scope of this study.

We proceed as follows. Section 1 surveys the existing literature. We motivate our key hypotheses in Section 2. Section 3 describes the data as well as the algorithms we

⁴Wider exposure of the limit order book to unregulated traders also would decrease the numbers and sizes of orders that traders are willing to leave in the book. This would further decrease specialist trading profits and likely increase transactions costs for large traders. Allowing traders to specify that their orders not be fully displayed in the order book could offset these effects to some extent. The electronic trading systems used at the Paris Bourse, at Globex and in several markets allow large traders to specify such display instructions.

use to identify specialist trades and to construct the limit order book. We present and interpret our results in Section 4 and conclude in Section 5.

1. Extant literature

This study is closely linked to theoretical models of competition in markets with dealers and limit order books. Rock (1999) models the competition between limit order traders and risk-averse specialists arising out of risk-related inventory issues. In his paper, the informational advantage of the specialist comes from market orders and not from limit orders, as proposed in this paper. Seppi (1997) examines this competition in the presence of informative limit orders and additional constraints such as public priority rules. In his model, however, the limit order book is common knowledge. It therefore abstracts from the transparency issues that are important here.

This study is also related to Madhavan and Panchapagesan (2000). They examine the role of specialists in price discovery at the NYSE open. Specialists facilitate price discovery using their superior information about the pre-opening order flow. As with our results for the continuous market, they find evidence of specialists using order book information in their trades in the opening auction.

Benveniste et al. (1992) and Chan and Weinstein (1992) consider how “professional relationships” among floor brokers and specialists provide specialists with information that helps them trade profitably. Our empirical paper focuses exclusively on information in the booked order flow whereas these theoretical papers consider the importance of information about orders held by floor brokers.

Our work complements analyses of specialist trading and of limit order books. Using the same database, Hasbrouck and Sofianos (1993) find specialist trades to be profitable, especially in the short-run. Sofianos (1995) finds that specialist dealing profits are small while their total trading profits are large, especially in the more active stocks. These results suggest that specialists time their trades well, a phenomenon that we attribute partly to their knowledge of the limit order book. Madhavan and Sofianos (1998) model specialist trading as a function of several variables, but do not explicitly consider information from the limit order book. However, consistent with our results, a more recent work by Boehmer et al. (2003) shows that specialists’ participation rate decreased after the NYSE opened up its limit order book to off-Exchange traders in real time. Ready (1999) finds evidence that suggests that specialists selectively stop orders to sample future order flow. Moreover, he finds that the market orders that the specialist does not allow to cross with the limit order book are those that are most profitable ex-post, a finding that is consistent with the results of this paper. We extend Ready’s work to establish the link between stopping and order book information. Finding results similar to ours, Edwards and Harris (2001) show that specialists step ahead of the limit orders more often since the tick size was reduced to a sixteenth of a dollar.

Huang and Stoll (1994) find that asymmetries in quoted depth predict future price changes, especially over short time intervals. We extend their work by decomposing quoted depth into depth provided by limit order traders and depth provided by the crowd including the specialist. This decomposition allows us to determine whether specialists complement or compete with SuperDot limit order traders when offering liquidity. Using data from the Australian Stock Exchange, an open electronic limit order book, Cao et al. (2003) provide empirical evidence that the orders behind the quotes are informative and helpful in predicting future short-term returns. However, their analysis is based on quantity asymmetry while we consider a richer characterization of the order book information using option values.

Though Copeland and Galai (1983) recognized the option-like characteristics of limit orders, only recently have researchers used option valuation techniques to study order book asymmetries. Such techniques are difficult to apply to limit orders because “time to maturity” is endogenous with random characteristics. Lo et al. (2002) present a model to compute time to expiration for limit orders using standard survival analysis methods. Jarnecic and McInish (1997) compute the option value of the limit order book at the Australian Stock Exchange using exogenously chosen expiration times such as 30 minutes, one hour and so on. We differ from the above two studies by estimating expected time to removal for a limit order using regression models.

Kavejecz (1999) and Chung et al. (1999) examine the factors that determine specialist quotations. In contrast, this paper primarily examines the factors that determine specialist trades. Since quotation and trading decisions are closely related, the three studies are complementary. Quotation and trading decisions, however, differ in significant ways. Perhaps most importantly, quotations become public information when the specialist makes them. Specialist trades, however, are not public information and generally cannot be inferred from publicly available transactions data. Since any trader can trade ahead of the specialist by expressing interest at the same price, specialists may be reluctant to reveal private information in their quotes. Specialists trading decisions therefore may reveal more about their private information than their quotation decisions.

Finally, our study provides empirical validation to the claim in Harris (1990, 1998) that an economically meaningful tick size is necessary to protect limit order traders from quote-matchers. Quote-matchers trade only to exploit other traders’ interest in a security. A large relative tick makes it expensive for them to trade ahead of large orders in the book. A small relative tick encourages quote-matching, and hence lowers incentives to offer liquidity. Goldstein and Kavejecz (2000) provide evidence supporting this hypothesis in their study of the 1997 switch in tick size from an eighth to a sixteenth at the NYSE. Perhaps most notably, they show that the rate at which specialists stop orders increased 50 percent. The main complaint that buy-side traders make about the recent switch to a penny tick is that their orders too often are being “pennied”. Recent evidence on specialist participation following decimalization seems partly to vindicate traders’ complaints (see Battalio and Jennings, 2002). Our results bear directly on this issue.

2. Theoretical discussion

2.1. *Information in the limit order book*

Harris (1990) discusses two types of traders who may want to use limit orders: pre-committed traders and value-motivated traders. Pre-committed traders use limit orders to reduce trading costs, while value-motivated traders use limit orders to trade only at prices that are acceptable given their value estimates.

Both trader types often place aggressive orders. Pre-committed traders place limit orders close to the market to increase the probability that they will trade. If they do not trade, they must chase after the market to fulfill their commitment. Value-motivated traders place aggressive limit orders (and sometimes even market orders) when they want to capitalize quickly on profitable trading opportunities that they perceive.

Whether pre-committed traders or value-motivated traders generate orders, the aggregate order imbalance in the limit order book conveys some information about future price changes. For example, an imbalance due to pre-committed trading activity could trigger price changes if traders convert their limit orders to market orders upon their failure to execute. Similarly, value-motivated traders affect prices as information in their estimates is impounded into prices. In both cases, prices tend to increase following buy side imbalances (when the buy side of the book is heavier than the sell side), and decrease following sell side imbalances. The knowledge of such imbalances is therefore quite valuable to any trader who can act on it. At the NYSE, the specialists know more about when order imbalances arise than any other traders. Although recent moves by the NYSE to publish order book imbalances may partly negate this advantage, the fact still remains that specialists know more about their order books than do other traders. Information such as order size and the time of arrival affect order option values and hence can provide an informational advantage to specialists even when the Exchange publishes aggregate order imbalance information. We now turn to describe the actions specialists may take to exploit this information.

2.2. *Specialist discretionary actions*

Though Exchange rules often constrain specialist trading, several actions remain within their discretion. Upon the arrival of a market order, a specialist can stop the order from executing at the prevailing quote by guaranteeing execution later at a better price. He may then decide to fill the order for his account any time before another order arrives that can fill it. Alternatively, he can execute the incoming market order at the market quote, which may or may not reflect his own trading interest. Finally, he also may choose to execute the market order immediately for his account at an improved price. We examine each alternative action in detail below.

The public precedence rules of the NYSE often constrain these choices. As Exchange members, specialists can never trade ahead of public orders at the same

price. When public traders offer the best quoted prices, specialists therefore must improve price if they wish to trade immediately.

2.2.1. *The decision to stop*

Several reasons may explain why specialists stop market orders. First, stopping orders allows the specialist to seek price improvement for market orders. The specialist may then update the market quote to reflect the stopped market order.⁵ For example, a specialist who has stopped a market sell order at $10\frac{1}{8}$ when the quotes were $10-10\frac{1}{4}$ might now quote the stock at $10-10\frac{1}{8}$ with the offer size reflecting the size of the stopped market sell order. The stopped order gets price improvement when it trades with an incoming market buy order or when the specialist trades with it. Since the Exchange evaluates specialist performance based on how often they improve prices, this stop mechanism allows specialists to improve prices occasionally without trading themselves.

Specialists may also stop market orders when they are uncertain about where to place their quotes. Imbalanced market orders flows suggest that specialists have set their quoted prices too low or too high. When the quotes reflect the specialists' commitments to supply liquidity, rather than their booked limit orders, such imbalances may adversely affect their inventories. By stopping market orders, specialists gain time to analyze subsequent market conditions before deciding to trade for their accounts. For example, a specialist may stop a market sell order instead of immediately trading with it to observe whether related markets rise or fall. If values appear to be rising, the specialist will fill the order for his account and then move his quotes up to limit his exposure to market buy orders. If values appear to be falling, the specialist will hope that a marketable buy order will arrive with which he can match the stopped market sell order. Ready (1999) discusses these options in greater detail.

2.2.2. *The decision to trade*

After deciding whether to stop the market order or not, the specialist then must decide whether to trade with it or not, assuming that he has that option. Whether he has an option to trade with the order depends on the orders in his book and in the crowd, and in the case of stopped orders, on the subsequent order flow.

When a specialist stops an order, he may trade with it only until a marketable order on the other side arrives. At that time, the new order must be matched with the stopped order. If no such order arrives, the Exchange requires that the specialist fill the stopped order after a set time interval. During our sample period, this interval was 30 minutes.

If the specialist does not stop the market order, he must match it with one or more limit orders from the book or crowd. He also can fill it himself, but only at a better price than the best price on his limit order book. If no limit order is in the book or

⁵The NYSE Rule 116.30, enacted after our sample period, now requires specialists to update their quotes after stopping a market order to reflect the stopped price.

crowd at the specialist's quoted price, he must fill the incoming market order at the quoted price or better, up to his quoted size.

Whether the specialist wants to trade depends on his assessment of whether the trade will prove to be profitable, on the effect the trade will have on the quantitative and qualitative measures that the Exchange uses to evaluate his performance, and on the risks imposed upon him by his current inventory position.

The expected profitability of the trade depends on the price at which he can trade and upon any information that the specialist can extract from the order flow and from his limit order book. We expect that the specialist will more likely trade if he can buy when the limit order book is heavy on the buy side or sell when it is heavy on the sell side. He will be less likely to trade when he must substantially improve price to do so. Since the minimum price increment determines the cost of obtaining precedence through price priority, it should be an important determinant of whether specialists trade when they must improve prices to do so.

Exchange rules require that specialists trade to provide liquidity when no other traders are willing to trade. Under these circumstances, specialists will buy when their limit order books are heavy on the sell side and sell otherwise. This obligation complicates our analyses. We cannot simply examine whether specialists more often trade in front of the heavy side of their books than in front of the light side since Exchange rules often force them to the light side. Our analyses therefore only examine situations in which specialists were not required to trade.

Finally, specialists—like all dealers—do not like to have large inventory positions. Therefore they are more likely to trade if they can buy when they are short of their desired inventories or sell when they are long.

2.2.3. The decision to improve prices

Having decided to trade, the specialist then must decide whether to improve prices or not. When the quotes reflect public orders, this decision is trivial: the specialist can trade only at improved prices. Otherwise, the specialist may choose to trade at better than quoted prices to improve his measured performance as a supplier of liquidity.

The ability to selectively improve prices after the arrival of a market order gives specialists a distinct advantage in their competition with limit order traders. We expect that specialists will improve prices more often when the book is highly asymmetric.

2.3. Empirical design

We undertake three major analyses in this study. The first analysis determines whether limit order book asymmetry is indeed informative about future price changes. We expect this result as specialists often are reluctant to display the full contents of the order book. The second analysis examines whether specialist trading decisions depend on limit order book asymmetry. The final analysis examines whether relative tick size helps explain cross-sectional differences in specialist trading decisions.

To separate discretionary actions from non-discretionary specialist actions, we employ two criteria to classify every market order that reaches the specialist. These criteria identify whether Exchange regulations allow the specialist to stop the order and whether Exchange regulations allow the specialist to participate without improving the price. This two-way classification produces distinct four cells within which we perform separate analyses of the three decisions specialists may make when a SuperDot market order arrives: to stop the order, to trade with the order and/or to improve execution price for the order. Since Exchange regulations do not allow specialists any discretion in certain circumstances, our design examines only a subset of the 12 analyses that this design would otherwise entail. Appendix A describes these scenarios.

If the limit order book is indeed informative about subsequent price changes, then specialist actions should reflect that information when specialists are not otherwise constrained by Exchange regulations. In particular, we expect that specialists are more likely to stop a market order, trade immediately, and improve prices when they can trade in front of an order book that is heavy on the same side. We also expect that the specialist will more often trade in front of the book when the relative tick size is small.

3. Data

We use the NYSE TORQ database for this study. The TORQ database includes all trades, quotes and system orders for 144 randomly selected securities for the period November 1990 to January 1991.⁶ System orders include all orders placed through the NYSE's automated trading system, SuperDOT. Because non-system orders are not in the database, we restrict our attention to only SuperDOT orders. The database also includes audit trail information on the identity of the members behind each trade. We use this information to identify specialist trades. Our cross-sectional analyses use market capitalization and pre-sample daily prices taken from the CRSP database.

3.1. The superDOT market order sample

To construct our market order sample, we start with all orders coded as market orders in the TORQ database. To these we add all marketable limit orders. We identify marketable limit orders by matching all limit orders with the prevailing quotes at the time they arrived.⁷ The limit prices of such orders cross the quotes and

⁶Hasbrouck (1992) provides a detailed description of the TORQ database. Though the database is a decade old, it still remains popular with researchers for lack of an alternative. The NYSE has not since provided order data to the public.

⁷Studies that match trades to quotes often lag the quotes slightly because the quote revision process is often faster than the trade reporting process (see Lee and Ready, 1991). Such lags are not necessary in this study because we match arriving orders to prevailing quotes. Since we study only system orders, chances of mismatch are low.

thereby make them immediately marketable. We do not differentiate between marketable limit orders and market orders.

We use four filters to narrow our sample. First, we include only standard buy and sell orders. We do not examine short-sell orders or tick-sensitive orders since special rules govern them. Second, we restrict our attention only to post-opening orders. Pre-opening orders are filled in a batch market that is quite different from the subsequent continuous market. Third, we constrain the bid quote at the time of the arrival to greater than one dollar so that the tick size is one-eighth dollar for all orders.⁸ Finally, we include only orders with matched contra-party audit information.

This last filter requires further explanation. The audit trail created by the Exchange includes information about each trade from sources that include member firms and clearing agents. Subsequent matching of buyers and sellers is accurate in most cases but not all. It is not uncommon to find trades with unmatched or partially matched contra-parties. The latter case arises when audit trail information exists only for part of a trade. Since contra-party identification is crucial in flagging specialist participation, we include only market orders with full contra-party information in our sample.

3.2. *Limit order book reconstruction*

For each market order, we must characterize the information contained in the limit order book at the time it arrived. To do so, we reconstruct the limit order book for each stock as in Kavejcz (1999).

The procedure involves several steps. First, from order cancellations and from trades that had no corresponding submissions in the sample period, we infer all limit orders that were submitted before the sample period. To these orders we add all limit orders submitted before the time at which we want to reconstruct the book. From this set of orders, we then remove all orders that filled, expired or were cancelled before the time of interest. The book of standing limit orders consists of the remaining orders. We construct this book at the time of each market order arrival.

The resulting limit order books include all system orders except those good-till-cancelled orders that were placed before the start date of the sample and which were not executed or cancelled during the sample period. In most cases, such orders would have been placed so far from the market that they would not have been very informative.⁹

3.3. *Valuing the information in the limit order book*

To examine whether specialists rely upon their order books, we must empirically characterize order book asymmetries. For this purpose, we reduce the information in

⁸Tick sizes at the NYSE decreased to one-sixteenth dollar for stocks trading above one dollar in June 1997 and again to one cent in January 2001.

⁹See Kavejcz (1999) and Chung et al. (1999) for a detailed analysis of the quoting behavior of limit order traders.

the order book to a single quantity that characterizes its asymmetry. Although we can imagine many such measures, the best measure for our purpose is one that efficiently summarizes information in the order book about trading options and future price changes. We believe that such a measure should embody the following qualitative properties:

1. Larger orders are more informative and have greater option value than smaller orders.
2. Size offered close to the market is more informative and has greater option value than size offered far from the market.
3. Orders that specialists expect to remain in the book for a long time provide more option value than those that they expect soon will be filled or cancelled.
4. All other things equal, buy and sell orders are equally informative.
5. All other things equal, order option values are greater when the market is volatile than when it is stable.

Since limit orders provide valuable trading options (see [Copeland and Galai, 1983](#)), we believe that a measure patterned on an option pricing model is most appropriate. We use two different option pricing models—the simple Black-Scholes model as well as a barrier option¹⁰ pricing model—to characterize the values of the puts and calls that buy and sell limit orders respectively offer to the market. Both these models have their limitations in our current context: The Black-Scholes formula will not accurately represent limit order option values because its derivation is based on arbitrage and price continuity assumptions that are not reasonable in microstructure applications and because limit order traders have an option to cancel their orders before execution. On the other hand, applying the barrier option pricing formula includes making assumptions on the barrier price for limit orders. For the barrier option to have value when the underlying value is equal to the limit price, the barrier price must be different from, but close to, the limit price. Although both these models will not accurately estimate actual limit order option values, any estimation errors will likely equally affect both buy and sell orders and therefore have little effect on our measure of asymmetry. Our empirical results confirm this and, therefore, we present our results only for asymmetries using the simple Black-Scholes formulation. We can provide the results under the barrier option formulation upon request.

We construct our limit order book asymmetry measure by separately summing the “option values” of all buy and sell limit orders in the book. The difference between these two numbers is our measure of limit order book asymmetry. Because this

¹⁰Barrier options come into existence (or cease to exist) when the underlying security’s price hits a pre-determined level or barrier before expiration. Barrier options that come into existence upon reaching the barrier are called “in” options while options that are knocked out are called “out” options. Options provided by limit orders have value only until they get filled (when the stock value reaches the limit price) or cancelled. A natural way to interpret this would be that the limit order resembles a standard barrier “out” option with the barrier price being a price close to its limit price and with time to expiration being its expected time to cancellation.

measure is based on an option pricing formula, it has all the desired qualitative properties enumerated above.¹¹

To compute option values using the barrier option method, we specify the barrier price as the price that is one tick away from the limit price—one tick higher for limit sell orders and one tick lower for limit buy orders. A limit buy order would, therefore, provide the market with a free down-and-out put option while a limit sell order would provide a free up-and-out call option.¹²

To apply the option pricing formulae, we must specify values for the various inputs. The strike price is, of course, the limit price of the order. For the underlying security value, we take the quotation midpoint.

We use the annualized 3-month Treasury bill rate as of October 31, 1990 for the risk-free interest rate. The assumed interest rate has essentially no impact on the results since the short time intervals involved ensure that it has little effect on the option values. Moreover, any bias in option values that might be caused by using the wrong interest rate will have equal and opposite effects on the buy and sell orders and therefore should have no systematic effect on our results.

We estimate volatility for each stock by its annualized daily return variance measured over the twelve months preceding the sample. Although the volatility assumption substantially affects the level of the “option value” estimates, the effects are symmetric on both sides of the market. We could estimate more precise “option values” using a time-varying volatility model, but we do not expect that the additional information would significantly change the results, which depend on differences in values between the two sides of the book.

The most important assumption we make concerns the time to maturity. Unlike options contracts, limit orders do not expire at any set time. Instead, they stand until they are filled, are cancelled, or are no longer valid because they have expired. Determining the expected time to removal is therefore critical to the option value provided by a limit order. To compute limit order option values at the Australian Stock Exchange, [Jarnecic and McNish \(1997\)](#) assume arbitrary maturities (5, 30 min and so on) for all limit orders. In practice, however, the times when limit orders are removed depend on their prices, sizes, validity instructions (day or good-till-canceled), and current market conditions. Since specialists undoubtedly are aware of the conditions that determine time-in-the-book, we must model these conditions to better understand how specialists use information in their limit order books.

¹¹We also examined other ways of characterizing the order book asymmetry such as the difference in modal depth and the simple difference in the total quantities bid and offered. Modal depth represents the cumulative quantity offered or bid at prices that would produce the modal spread (or a multiple of it) for the security in the sample. The measure we used was $Q^O - Q^B$, where O and B represent offer and bid prices such that $O - B = n \times \text{modal spread}$ and $n = (1, 2, \dots)$. Since these measures do not include information about order duration and about total distance from the market, our results were not as strong as those presented below. In the interest of conserving space and because the results are qualitatively similar, we do not report these results.

¹²[Haug \(1998\)](#) describes the valuation of standard barrier options.

For each stock, we use the following simple linear model to estimate the expected time to removal:

$$\begin{aligned} E(\text{LogDuration}) = & \beta_0 + \beta_1 \text{LogAge} + \beta_2 \text{OrderSize} + \beta_3 \text{Queue} + \beta_4 \text{PriPos} \\ & + \beta_5 \text{LogTimeToClose} + \beta_6 \text{MkOrdAsym} + \beta_8 \text{SubToday} \\ & + \beta_8 \text{Queue} \times \text{TimeToClose} + \beta_9 \text{PriPos} \times \text{TimeToClose} \end{aligned}$$

LogDuration represents the log of the time (in minutes) remaining until the limit order is removed from the book. To apply the Black-Scholes pricing formula, we use all modes of removal—including a fill or a cancellation (either voluntarily or automatically at the end of the day)—to determine the expected duration for a limit order. On the other hand, we use only the time until cancellation to determine duration of the barrier option provided by a limit order, as time to fill is indirectly captured in the option's value itself. We log the dependent variable to control residual heteroskedasticity. *LogAge* represents the log of the time since the limit order was entered into the system. We expect recently entered limit orders to be cancelled more frequently than older orders. *OrdSize* is the limit order size. We expect that traders cancel larger orders more frequently than smaller orders because larger orders have greater option value. To facilitate cross-stock comparisons, we scale this variable by dividing it by average trade size for the stock. *Queue* is the cumulated sizes of all booked orders with higher price priority and time precedence. We expect that orders that stand behind large size are less likely to execute and more likely to be cancelled. We also scale this variable by dividing it by average trade size for the stock. *PriPos* is the absolute difference between the order price and the inside quotation on the same side of the market. We expect that traders are more likely to cancel orders that stand far from the market. *LogTimeToClose* is log minutes until the daily close of trading. Since day orders expire at the close of trading, time-to-close must be highly informative. We use this variable only when estimating maturities for day orders. *MkOrdAsym* characterizes asymmetry in recent market order arrivals. If market orders are serially correlated over short intervals, we expect that traders will more likely cancel orders that supply liquidity to the less dominant side than orders on the other side. We compute *MkOrdAsym* as the summed difference between the market sell and buy orders that arrived in the 15 preceding minutes. We scale the variable by dividing by average 15 min market order volume and we multiply the variable by -1 for limit sell order observations. *SubToday* is a dummy that indicates whether the order was submitted today. We use this variable only when estimating maturities for good-till-cancelled orders. We expect that traders cancel many good-till-cancelled orders on the same day they submit them.

The two interaction terms model cross-effects among the variables that are most relevant in determining the expected duration of limit orders.

We estimate separate models for day orders and for good-till-cancelled orders since we expect that their parameter estimates will differ considerably. In particular, we expect that *TimeToClose* will be much more important for the day orders than for the good-till-cancel orders. We also estimate these two models separately for each stock to account for cross-sectional variation among the stocks.

We composed the estimation samples as follows: At fifteen minute intervals from 9:45 AM through 4:00 PM on each trading day and for each stock, we reconstruct the limit order book as described above. We then identify and include in the estimation sample all standing limit orders that subsequently were removed from the book during the sample.¹³

We then used our estimates to predict the remaining duration of each limit order in the book at the arrival time of each market order. Using this information and the above assumptions, we compute the value of the Black-Scholes formula as well the value of the barrier option for all standing orders in the book at the arrival time of each market order.

From these we then compute our measure of order book asymmetry,

$$OrdBookAsym = \left(\sum_{i=1}^{N_b} V_i^b S_i - \sum_{j=1}^{N_s} V_j^s S_j \right),$$

where V is the Black-Scholes option value (or the barrier option value) per share, S is the order size, and the superscripts b and s represent buy and sell orders respectively. *OrdBookAsym* characterizes asymmetry in the limit order book by weighting limit orders by their “option values”. Note that we do not scale our measure as most of our analyses are performed for each stock separately. We use this measure to predict price future changes and to explain specialist trading decisions.

3.4. Identifying specialist trades

The TORQ database does not explicitly identify specialist trades. Instead, we infer them using an algorithm originally developed by Edwards (1999) and later refined by Panchapagesan (2000). The idea behind the algorithm is straightforward. The TORQ data include detailed information taken from the NYSE equity consolidated audit trade file (CAUD) about the identities of most traders that participated in each trade. Certain trader identities, including those of the specialists, are omitted from TORQ, however. The omission of a trader identity code therefore may indicate that a specialist participated in the trade. Although trader identities may be omitted for other traders as well, certain other tests allow us to infer confidently most specialist

¹³Since the sample does not include orders that were removed from the book after the end of our sample period, we underestimate time-to-removal for long-standing good-till-cancelled orders placed far from the market. Since these orders are not very common and since orders standing far from the market usually do not have significant option values, we do not expect that this selection bias will significantly affect the results. We use an alternative assumption as well where GTC orders that remained at the end of sample period were assumed to be cancelled arbitrarily after two trading days but we find no change in our results. Since many orders appear in the sample at many different sampling times, the regression error terms undoubtedly have complex properties characteristic of pooled cross-sectional–time series models. We ignore these issues because our interest is in predicting order durations rather than in drawing statistical inferences from the model coefficients.

trades.¹⁴ A complete description of our algorithm and measures of its efficiency appear in Panchapagesan (2000).

We use the algorithm to identify all specialist trades with SuperDOT market orders in the TORQ data. We also use it to measure specialist inventories by summing all signed specialist trades. Since we do not observe opening inventories, we cannot estimate the actual inventory levels. Our measure, however, is sufficient to estimate the changes in inventory levels that are relevant to our analyses.

Though the TORQ data is a decade old, it provides us with the identification of specialist trades that newer data (including the NYSE OpenBook™)—available in the public domain either for free or for a price—cannot provide. Since limit order book information is more valuable in a penny environment where liquidity at the inside has diminished substantially, we believe that using more recent data would only strengthen our inferences.

4. Empirical results

4.1. Descriptive statistics

Table 1 summarizes specialist trades with SuperDOT market and marketable limit orders as identified by our specialist trade identification algorithm. In the TORQ sample, SuperDot traders submitted more marketable buy orders than marketable sell orders, whether measured by numbers of orders or by total share volume. Specialist participation rates were higher for buy orders than for sell orders so that specialists sold more than they bought when trading with SuperDOT marketable orders. These results are consistent with their role as liquidity suppliers. The specialists during this period either bought more from floor traders than they sold to them or they divested inventory. These results are robust across stocks classified by average number of trades and across orders classified by size.¹⁵

As expected, specialists participate more in less actively traded stocks. This result is consistent with many previous studies that identify the importance of dealers in illiquid stocks.

Our simple linear model of limit order durations explains 43 percent (72 percent when we consider only the duration until cancellation), on average across stocks, of the variation in the log remaining time that day limit orders spend in the book (Table 2). Most of the estimated coefficients, especially our median estimates, have the sign we expected. Cross-sectional Wilcoxon signed-rank tests indicate it very unlikely that the frequency of positive coefficients is one-half.¹⁶

¹⁴The audit information is incomplete for trades reported by other exchanges. This algorithm will therefore not capture trades of NYSE specialists routed through other exchanges. However, such trades are likely to constitute a low percentage of their total trades.

¹⁵Other measures of trading activity such as dollar volume produce similar results.

¹⁶As discussed above, standard measures of statistical significance are mis-specified in the OLS estimation of this pooled cross-sectional–time series model. Notwithstanding these concerns, we take comfort in noting that the estimated coefficients with the expected sign were mostly “statistically significant” but few of the other estimated coefficients were “significant”.

Table 1
Specialist trades with superDOT market and marketable limit orders

This table presents summary statistics of specialist trades with SuperDOT market orders and marketable limit orders for all 144 stocks in the TORQ database between the period November 1, 1990 and January 31, 1991. Only standard non-opening buy and sell orders with matched contra-party information in the audit file are included. Specialist trades are identified using the algorithm presented in Panchapagesan (2000). We use the average daily number of trades to classify stocks into three trading groups. Mean specialist participation for each group is the average of the individual stock participation rates weighted by the number of market orders for each stock. SE represents the cross-sectional standard error. Panel B classifies specialist participation rates by market order size.

Panel A: overall statistics												
	Buy orders						Sell orders					
	Number of orders			Dollar volume			Number of orders			Dollar volume		
	All orders (N)	Specialist participation (%)		All orders (\$000)	Specialist participation (%)		All orders (N)	Specialist participation (%)		All orders (\$000)	Specialist participation (%)	
		Mean	SE		Mean	SE		Mean	SE		Mean	SE
All stocks	204,538	22.6	0.9	5,110,972	18.4	0.9	172,840	20.3	0.9	4,090,448	17.1	0.8
Activity group												
Least	3,547	42.7	2.9	24,612	36.5	2.8	4,985	32.7	2.8	26,479	27.0	2.7
Moderate	20,130	38.1	2.1	167,574	33.8	2.1	18,140	31.4	2.4	157,205	28.3	2.2
Most	180,861	20.5	1.2	4,918,786	16.3	1.0	149,715	18.5	1.3	3,906,763	15.4	1.0
Panel B: specialist participation as a fraction of total dollar volume at different order sizes												
	Specialist participation (%)				Specialist participation (%)				Specialist participation (%)			
	≤1,000 Shares				1,001–10,000 Shares				≥10,000 Shares			
	Buy		Sell		Buy		Sell		Buy		Sell	
	Mean	SE	Mean	SE	Mean	SE	Mean	SE	Mean	SE	Mean	SE
All stocks	17.3	0.9	15.2	0.9	17.6	0.8	17.3	0.8	12.6	1.4	17.9	2.8
Activity group												
Least	39.3	2.8	13.3	1.1	26.9	3.2	21.8	2.4	14.8	21.0	0.0	0.0
Moderate	34.2	1.9	27.7	2.2	25.9	2.5	23.8	2.6	10.1	7.3	12.9	6.8
Most	15.0	1.1	13.3	1.1	16.5	1.0	16.2	1.1	12.6	1.4	18.7	3.2

Table 2

Regression estimates of the time remaining to removal of superDOT limit orders

This table presents summary statistics of the OLS estimates of a regression model for the log expected time (in minutes) remaining until removal of standing limit orders. Orders are removed from the order book through execution, cancellation or expiration. We also present estimates (provided in parentheses) for the log expected time to removal through a cancellation only. The model is estimated separately for day orders and good-till-canceled orders for each stock using TORQ data sampled at 15-minute time intervals starting at 9:45 A.M. between November 1, 1990 and January 31, 1991. The independent variables are: *LogAge*, the logarithm of time already spent in the system by the limit order (in minutes); *OrdSize*, the order size relative to average trade size for the stock (in shares); *Queue*, the number of shares that are ahead of the limit order in price and time priority relative to the average trade size for the stock; *PriPos*, the price position of the order relative to the market (bid for a buy order and ask for a sell order) expressed in ticks; *LogTimeToClose*, the logarithm of time (in minutes) remaining till market-close; *MkOrdAsym*, the ratio of the difference between market sell (buy) orders and market buy (sell) orders in the preceding 15 minutes over average market order volume for a limit buy (sell) order; and *SubToday*, a dummy variable that is one if the order was submitted today. In addition, we also include variables to capture the interaction of *LogTimeToClose* with *Queue* and *PriPos*. The cross-sectional mean and median estimates are computed over all stocks. The superscript, * and **, after the variable indicates whether the frequency of positive coefficients is one-half under a Wilcoxon signed-rank test for the models for the log expected time until removal and for the log expected time until cancellation respectively.

Variable	CS mean estimate	CS median estimate	Standard error of mean	Percentage of stocks with coefficients			
				Negative	Positive	Significantly negative	Significantly positive
A. Day orders ^a							
<i>Intercept</i>	−0.243 (−0.131)	−0.190 (−0.099)	0.041 (0.024)	66% (66)	34% (35)	48% (43)	12% (12)
<i>LogAge</i> ^{*,**}	0.174 (0.063)	0.168 (0.051)	0.008 (0.005)	1 (6)	99 (94)	1 (1)	92 (81)
<i>OrdSize</i> [*]	0.027 (0.001)	0.024 (0.001)	0.003 (0.003)	9 (47)	91 (53)	3 (22)	83 (31)
<i>Queue</i> [*]	−0.033 (0.001)	−0.031 (−0.004)	0.009 (0.007)	86 (65)	14 (35)	60 (22)	1 (7)
<i>PriPos</i> [*]	−0.100 (0.001)	−0.066 (−0.002)	0.015 (0.009)	87 (55)	13 (45)	55 (20)	2 (7)
<i>LogTimeToClose</i> ^{*,**}	0.773 (0.932)	0.780 (0.928)	0.008 (0.004)	0 (0)	100 (100)	0 (0)	100 (100)
<i>MkOrdAsym</i>	−0.004 (0.004)	−0.003 (0.003)	0.001 (0.004)	74 (23)	26 (77)	31 (1)	5 (40)

<i>Queue</i> × <i>LogTimeToClose</i> [*]	0.009 (−0.001)	0.008 (0.001)	0.003 (0.002)	12 (35)	88 (65)	2 (9)	68 (31)
<i>PriPos</i> × <i>LogTimeToClose</i> [*]	0.030 (0.001)	0.019 (0.001)	0.004 (0.003)	9 (40)	91 (60)	3 (7)	64 (24)
B. Good-till canceled orders ^b							
<i>Intercept</i> [*]	7.780 (7.501)	7.782 (7.370)	0.093 (0.092)	0% (0)	100% (100)	0% (0)	100% (100)
<i>LogAge</i>	0.039 (0.046)	0.045 (0.057)	0.012 (0.012)	32 (29)	68 (71)	26 (26)	65 (67)
<i>OrdSize</i>	−0.014 (−0.005)	−0.005 (−0.007)	0.038 (0.031)	55 (54)	45 (46)	50 (50)	41 (52)
<i>Queue</i>	−0.001 (−0.004)	−0.001 (−0.002)	0.004 (0.003)	53 (56)	47 (44)	51 (53)	45 (42)
<i>PriPos</i> [*]	0.044 (0.046)	0.014 (0.018)	0.009 (0.009)	15 (13)	85 (87)	13 (11)	84 (87)
<i>MkOrdAsym</i>	−0.001 (−0.000)	−0.001 (−0.000)	0.001 (0.001)	57 (50)	43 (50)	34 (27)	20 (26)
<i>SubToday</i> [*]	−0.630 (−0.332)	−0.625 (−0.255)	0.033 (0.033)	97 (84)	3 (16)	96 (76)	0 (9)

^a*N* = 139 stocks; Average R^2 = 0.43 (0.72).

^b*N* = 141 stocks; Average R^2 = 0.10 (0.17).

The results confirm that long-standing day orders remain in the book longer than recently submitted orders (*LogAge*) and that larger day orders remain in the book longer than smaller orders (*OrdSize*). The negative average coefficients for *Queue*—shares ahead—and *PriPos*—order price placement—must be interpreted together with the positive average coefficients for their cross-products with *LogTimeToClose*. These results together imply that *Queue* and *PriPos* have positive effects on remaining time when *TimeToClose* is greater than 39 and 28 minutes respectively.¹⁷ For smaller *TimeToClose*, traders probably cancel their limit orders and resubmit them when they are far back in the queue or when they are far from the market. *LogTimeToClose* has a very strong positive effect on remaining day order time, as expected. Finally, *MkOrdAsym*—the signed ratio of net signed market order size to total market order size in the last 15 min—has a slight negative effect on remaining time. As expected, a higher recent market order arrival rate on the opposite side of the limit order hastens the limit order's exit from the book. All our median estimates have similar signs whether we consider all forms of removal or only cancellations for determining limit order durations. However, the frequency of positive coefficients among firms differs from one-half for just two variables, *LogAge* and *LogTimeToClose*, when we consider only cancellations. In contrast, the frequency differs for seven variables when we estimate duration using all forms of removal.

For the good-till-cancel (GTC) orders, the regression model only explains 10 percent (17 percent when we consider only the duration until cancellation) of the variation in log order durations. Large GTC orders stay in the book for less time than do small orders, probably because institutional traders cancel their orders more often and because they prefer to use day orders. The coefficient for *PriPos* is positive as expected for GTC orders. Since many of these orders are not aggressively priced, they tend to stay long in the system. The negative *SubToday* coefficient estimates indicate that GTC orders are often cancelled on the same day they are submitted. Here too, our median coefficients are quite similar whether we consider all forms of removal or only cancellations for determining limit order durations. Although the poor explanatory power of the good-till-cancel model concerns us, we are comforted by the fact that only 20 percent of the TORQ limit orders are good-till-cancel orders (see Harris and Hasbrouck, 1996).

4.2. Limit order book information

Information that may be useful for predicting future price changes may appear in various places and forms in the limit order book. It may be at the quoted market. It may be behind the quoted market. It may be in the expected remaining durations of the orders in the book. It may be in the distribution of size by price within the

¹⁷The average estimated total effect of *Queue* is given by $-0.033 \text{ Queue} + 0.009 \text{ Queue} \times \text{LogTimeToClose}$. This expression is zero when *LogTimeToClose* is 3.67 (when *TimeToClose* is approximately 39 minutes), at which point the effect of *Queue* changes sign. Similarly, the sign change for *PriPos* also occurs at approximately 28 minutes.

book. Finally, it may be in the limit orders that were placed inside the quoted market but hidden to the public at large by the specialist.¹⁸

To determine whether the limit order book is informative, and to determine where the information lies, we created several measures of limit order book asymmetry. The first three measure raw asymmetry, which is simply the difference between the aggregate quantities that limit order traders are willing to buy and sell. These are

$BookAsym_{i,t}^{Away}$	Asymmetry in limit orders standing behind the quoted market
$BookAsym_{i,t}^{atQuote}$	Asymmetry in limit orders that are at the quoted market
$BookAsym_{i,t}^{insQuote}$	Asymmetry in limit orders that are inside the quoted market, and hence, hidden to the market at large

The sum of these three measures is the total raw limit order book asymmetry. The subscripts in each of the above variables represent security i and the 15-min time interval t that we use to estimate the asymmetry in the limit order book respectively. The informativeness of raw order book asymmetry is especially relevant today with the NYSE offering to display aggregate order book information to investors for a price.

We next created three corresponding variables that characterize option characteristics of the limit order book. As mentioned before, we use both the Black-Scholes formula and the barrier option formula to determine limit order option values. However, given that the two characterizations yield similar results, we present results only under the simpler Black-Scholes formulation.

For each stock, we estimated a linear regression of the limit order option values on raw limit order quantities.¹⁹ The estimated residuals from this model represent information about option characteristics of the limit order book (price placement and expected duration) that cannot be explained by aggregate order quantities. We estimate separately for limit buy orders and limit sell orders to correct for any differences in how the specialist may view large orders to sell from large orders to buy. Using this estimated model, we then computed an option characteristic asymmetry variable corresponding to each of the three raw variables listed above. The option characteristic asymmetry variable represents that part of the book asymmetry in option values that is not explained by asymmetry in quantities. The resulting variables are

$OptChar_{i,t}^{Away}$	Option characteristic asymmetry in limit orders standing behind the quoted market
$OptChar_{i,t}^{atQuote}$	Option characteristic asymmetry in limit orders that are at the quoted market

¹⁸During our sample period, the specialist was not obligated to display all standing limit orders in his market quote. Limit order handling rules subsequently introduced now require that the specialist display all size at the best prices. McInish and Wood (1995) document and discuss “hidden” limit orders.

¹⁹We estimate the model without an intercept to ensure that we do not identify option characteristics when the order book is empty. Also, we estimate the model separately for buys and sells in the limit order book.

$OptChar_{i,t}^{insQuote}$

Option characteristic asymmetry in limit orders that are inside the quoted market, and hence, hidden to the market at large

To determine whether order book information is useful for predicting future values, we estimated a regression of future returns for each stock using the above variables as explanatory variables. In addition, we included a lagged return to model the well-known return mean reversion in short horizon transaction price returns. We examined three time intervals (five minutes, one hour and one day) to determine the intervals, if any, over which the order book and quoted depth are informative. We observed our time-series at 15-min time intervals starting at 9:45 A.M. for all trading days between November 2, 1990, and January 31, 1991.²⁰

Though we present cross-sectional averages of the individual stock-wise regression coefficients, it is difficult to interpret them given the wide differences in trading activity and the depth of the limit order book across our sample stocks. Hence we focus on the results from individual stock wise regressions. The results (Table 3) show that order book asymmetry has substantial explanatory value in predicting future price movements, especially for active stocks. This result is consistent with Cao et al. (2003) who find limit order book asymmetry to explain future returns even in an open electronic market such as the Australian Stock Exchange.

The raw asymmetry measure coefficient estimate—the sum of the three $BookAsym_{i,t}$ estimates—is significantly positive for a little more than a fifth of all stocks examined at five-minute intervals. Positive coefficients indicate that a book heavy on the buy-side forecasts future price increases. Not surprisingly, the more aggressive orders—orders placed at or inside the quotes—convey the most information about future price changes. The percentage of stocks with positive estimated coefficients is higher for actively traded stocks than less active stocks. The result is most probably due to the more abundant data for the actively traded stocks.

The one-hour time interval results are stronger. About 42 percent of all stocks and 47 percent of the active stocks have a positive coefficient for the order book asymmetry variable. Here again, most of the information on the limit order book comes from orders placed at or inside the quotes.

The asymmetry measure is informative over the one-day horizon as well, but not as informative as over the short-term horizons. This result suggests that order book imbalances, expressed in raw quantities, convey more information about price movements over short intervals than over longer intervals. Interestingly, orders placed away from the quotes are more informative about long-run price changes than aggressively priced orders in the active stocks. As expected, evidence for return mean reversion is strong: the coefficient for lag returns is negative and significant in most stocks.

²⁰November 1, 1990, the first day in the TORQ database, does not appear in the sample due to the lagged dependent variable among the explanatory variables.

Table 3
Information in limit order book asymmetry

This table presents OLS estimates of a model of future transaction price returns on past returns and limit order book asymmetry variables. The following model is estimated for each security i of the 139 TORQ stocks with adequate data: $R_{i,t+k} = \alpha_i + \beta_{0,i}R_{i,t-k} + \beta_{1,i}BookAsym_{i,t}^{Away} + \beta_{2,i}BookAsym_{i,t}^{atQuote} + \beta_{3,i}BookAsym_{i,t}^{insQuote} + \beta_{4,i}OptChar_{i,t}^{Away} + \beta_{5,i}OptChar_{i,t}^{atQuote} + \beta_{6,i}OptChar_{i,t}^{insQuote}$ where subscript t denotes 15-minute time intervals starting at 9:45 A.M. between November 2, 1990 and January 31, 1991, $R_{i,t+k}$ is the transaction price return in basis points over k periods starting at time t , $R_{i,t-k}$ is the transaction price return in basis points over k periods ending at time t . The variables $BookAsym_{i,t}^{Away}$, $BookAsym_{i,t}^{atQuote}$ and $BookAsym_{i,t}^{insQuote}$ represent the quantity asymmetry (buys–sells) in limit orders that stand away from, at and inside the quotes respectively, while the variables $OptChar_{i,t}^{Away}$, $OptChar_{i,t}^{atQuote}$ and $OptChar_{i,t}^{insQuote}$ represent the residual option value asymmetry not captured by quantity asymmetry in limit orders away from, at and inside the quotes respectively. To obtain $OptChar$ variables, we first estimate a simple regression model, without an intercept, of option values on simple quantities using all limit orders for each stock, separately for buys and sells. The $OptChar$ variables represent the difference between the actual asymmetry in option values and the predicted asymmetry from this model for each of three price position categories of limit orders (at, away and inside the quotes). We consider five-minute, one-hour, and one-trading day periods for k . All quantity asymmetry measures represent the difference in aggregate quantities on the buy and sell sides. Order option values are computed using the Black-Scholes pricing formula with a risk-free interest rate of 7.5%, historic stock return volatility, and maturity equal to the expected time to removal estimated from the regression model presented in Table 2. We also compute option values using the standard barrier option pricing formula with similar inputs except that we use the expected time to cancellation alone for order maturity. Figures in parentheses represent results based on the barrier option model. The percentage of stocks that have significant negative and positive coefficients (at 5% confidence level under a two-tailed test) is reported for each trading activity group. The column ‘Total Book’ presents results for the above regression (run separately) using aggregate quantity and residual option value asymmetry. We use average daily number of trades to measure trading activity. The cross-sectional coefficient averages and standard errors for both quantity and residual option characteristic asymmetry are scaled by a factor of 100 only for presentation.

$k = 5$ minutes									
	Lagged return β_0	Quantity order book asymmetry variables				Residual option characteristic asymmetry			
		Away β_1	At quote β_2	Inside β_3	Total book	Away β_4	At quote β_5	Inside β_6	Total book
Cross-sectional	−0.62	0.01	0.01	0.27*	0.01	−0.15	−0.15	−4.75	0.30
Average	(−0.62)	(0.00)	(0.00)	(0.37*)	(0.01)	(3.41)	(1.99)	(−47.57)	(0.03)
Cross-sectional	0.02	0.01	0.01	0.12	0.00	0.33	0.30	4.13	0.17
Standard error	(0.02)	(0.01)	(0.01)	(0.13)	(0.00)	(4.17)	(5.02)	(69.46)	(0.56)

Table 3 (Continued)

	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	N
All stocks	99 (99)	1 (1)	6 (6)	17 (20)	9 (4)	23 (21)	5 (3)	32 (38)	5 (5)	22 (23)	8 (7)	7 (7)	13 (9)	6 (4)	11 (13)	16 (10)	10 (2)	14 (14)	139 (138)
<i>Activity group</i>																			
Least	100 (100)	0 (0)	2 (7)	18 (19)	20 (9)	5 (9)	11 (9)	7 (16)	7 (9)	14 (19)	5 (7)	5 (12)	14 (9)	9 (2)	11 (16)	9 (9)	14 (5)	7 (5)	44 (43)
Moderate	100 (100)	0 (0)	10 (2)	13 (19)	6 (4)	23 (15)	2 (0)	29 (29)	4 (2)	15 (21)	13 (6)	10 (0)	13 (10)	4 (2)	13 (6)	21 (15)	15 (2)	15 (13)	48 (48)
Most	96 (96)	2 (2)	4 (9)	19 (21)	2 (0)	40 (38)	2 (0)	60 (66)	4 (4)	38 (30)	6 (9)	6 (11)	13 (6)	4 (9)	9 (17)	17 (6)	2 (0)	21 (23)	47 (47)
<i>k</i> = 1 hour																			
	Lagged return β_0		Quantity order book asymmetry variables								Residual option characteristic asymmetry								
			Away β_1		At quote β_2		Inside β_3		Total book		Away β_4		At quote β_5		Inside β_6		Total book		
Cross-sectional	−0.13*		0.04		0.09*		1.01*		0.04		−1.69		−1.34		−1.69		0.95		
Average	(−0.13*)		(0.09)		(0.10*)		(0.84*)		(0.04)		(48.66*)		(6.32)		(−136.14)		(1.38)		
Cross-sectional	0.01		0.03		0.03		0.23		0.02		3.00		1.60		5.79		0.64		
Standard error	(0.01)		(0.05)		(0.03)		(0.30)		(0.02)		(25.46)		(22.43)		(183.62)		(2.24)		
	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	N
All stocks	63 (65)	6 (6)	14 (17)	33 (36)	9 (9)	49 (42)	4 (3)	42 (43)	14 (9)	42 (43)	13 (13)	24 (23)	17 (10)	16 (17)	20 (13)	12 (12)	12 (10)	40 (35)	139 (138)
<i>Activity group</i>																			
Least	91 (95)	0 (0)	20 (26)	39 (37)	25 (19)	30 (30)	11 (7)	25 (23)	30 (19)	32 (35)	16 (16)	25 (26)	23 (23)	30 (21)	18 (19)	18 (26)	16 (21)	43 (33)	44 (43)
Moderate	73 (75)	2 (2)	10 (15)	31 (35)	2 (4)	63 (44)	2 (2)	52 (44)	8 (6)	48 (46)	19 (6)	13 (17)	17 (6)	13 (13)	31 (6)	10 (8)	17 (2)	33 (31)	48 (48)
Most	28 (28)	15 (15)	11 (11)	30 (34)	0 (4)	53 (51)	0 (0)	47 (60)	6 (4)	47 (49)	4 (17)	34 (28)	13 (2)	6 (17)	11 (15)	9 (2)	2 (9)	45 (40)	47 (47)

$k = 1$ day

	Lagged return β_0		Quantity order book asymmetry variables								Residual option characteristic asymmetry								N
			Away β_1		At quote β_2		Inside β_3		Total book		Away β_4		At quote β_5		Inside β_6		Total book		
Cross-sectional	−0.02		0.02		0.58*		2.26*		0.08		2.52		−7.86		−11.86		1.54		
Average	(−0.02)		(0.16)		(0.59*)		(2.45*)		(0.14*)		(227.94*)		(−45.55)		(−112.43)		(16.20)		
Cross-sectional	0.01		0.09		0.16		0.82		0.08		5.54		10.37		26.18		5.30		
Standard error	(0.01)		(0.10)		(0.16)		(0.89)		(0.01)		(83.83)		(70.45)		(457.82)		(9.35)		
	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	N
All stocks	37	27	27	48	9	42	8	26	27	55	23	51	22	20	16	14	20	58	139
	(35)	(28)	(26)	(50)	(9)	(43)	(8)	(28)	(23)	(52)	(22)	(46)	(20)	(19)	(12)	(12)	(18)	(51)	(138)
<i>Activity group</i>																			
Least	45	18	41	39	16	52	14	39	39	43	20	43	27	36	30	23	23	59	44
	(49)	(23)	(30)	(42)	(12)	(60)	(16)	(33)	(26)	(47)	(21)	(44)	(28)	(19)	(21)	(19)	(23)	(53)	(43)
Moderate	31	35	21	50	10	48	8	23	19	65	35	38	23	19	15	15	29	46	48
	(23)	(35)	(21)	(54)	(6)	(40)	(4)	(19)	(19)	(54)	(17)	(54)	(17)	(25)	(8)	(15)	(19)	(50)	(48)
Most	34	28	19	55	2	28	2	17	23	57	13	72	17	6	4	4	9	68	47
	(34)	(23)	(28)	(53)	(9)	(32)	(4)	(34)	(26)	(55)	(28)	(38)	(15)	(13)	(6)	(4)	(13)	(51)	(47)

An asterisk for averages indicates significance at the 5% level.

The residual option characteristic asymmetry variable coefficients, as determined by a simple Black-Scholes characterization, are often statistically significant in predicting longer time horizon returns. Like the order book asymmetry measures, the results are stronger for active stocks than for less active stocks. The option characteristic variable is positive and significant in roughly one-fifth of all active stocks for the five-minute returns (and less than a quarter of all stocks when barrier option pricing formulation is used). This number increases to 45 percent for the one-hour returns and to 68 percent for the one-day returns (35 percent and 51 percent under barrier pricing formulation respectively). Not surprisingly, option characteristics such as price position and duration are most useful for limit orders placed away from the quotes. We find similar results when we use the barrier option pricing formulation, suggesting that our option value asymmetry measure is robust to the model specification.

Note that though order size asymmetry may be partially revealed in the quoted size, only the specialist is privy to option values. This could mean that the specialist can choose to offer liquidity based on information from option values rather than from simple quantities of limit orders in his book. This suggests that recent attempts by the Exchange to display the aggregate quantities in the book do not dilute the informational advantage enjoyed by the specialist. We provide a more detailed analysis of specialist actions and the source of their informational advantage later in the paper.

These results suggest that the specialists' unique access to the information in the limit order book gives them an informational advantage in their proprietary trading. We now consider whether specialists use this information when making trading decisions. We examine two types of decisions: quote revisions and actual trading decisions.

4.3. Limit order imbalances and quotation revisions

Specialists frequently revise their quotes following market order arrivals. To determine whether these revisions depend on asymmetries in their limit order books, we analyze whether order book imbalances predict quote revisions. In particular, we estimate a linear regression model of future quotation midpoint returns using the same explanatory variables described above. We expect that order book imbalances have a positive effect on quote revisions. In particular, if the buy side is strong, we expect that quotes will be revised upward, and vice versa for a strong sell side.

The results (Table 4) are very similar to, though stronger than, those reported in Table 3. As before, our focus is more on the results from the individual stock wise regressions than on the cross-sectional averages. The order book asymmetry variable is positive and significant for all time horizons. The estimated coefficient is positive and significant for about a third of all active stocks in the five-minute data and for about three-fifths of such stocks in the one-hour data. For the one-day returns, the coefficient is positive and significant in about 57 percent of all active stocks. As in the transaction return analysis, the residual option characteristic variable often provides significant additional explanatory power over simple quantity asymmetry in the

Table 4
Quotation revisions and limit order book asymmetry

This table presents OLS estimates of a model of future quotation midpoint returns on past returns and limit order book asymmetry variables. The following model is estimated for each security i of the 139 TORQ stocks with adequate data: $R_{i,t+k} = \alpha_i + \beta_{0,i}R_{i,t-k} + \beta_{1,i}BookAsym_{i,t}^{Away} + \beta_{2,i}BookAsym_{i,t}^{atQuote} + \beta_{3,i}BookAsym_{i,t}^{insQuote} + \beta_{4,i}OptChar_{i,t}^{Away} + \beta_{5,i}OptChar_{i,t}^{atQuote} + \beta_{6,i}OptChar_{i,t}^{insQuote}$ where subscript t denotes 15-minute time intervals starting at 9:45 A.M. between November 2, 1990 and January 31, 1991, $R_{i,t+k}$ is the midpoint quotation return in basis points over k periods starting at time t , $R_{i,t-k}$ is the midpoint quotation return in basis points over k periods ending at time t . The variables $BookAsym_{i,t}^{Away}$, $BookAsym_{i,t}^{atQuote}$ and $BookAsym_{i,t}^{insQuote}$ represent the quantity asymmetry in limit orders that stand away from, at and inside the quotes respectively, while the variables $OptChar_{i,t}^{Away}$, $OptChar_{i,t}^{atQuote}$ and $OptChar_{i,t}^{insQuote}$ represent the residual option value asymmetry not captured by quantity asymmetry in limit orders away from, at and inside the quotes respectively. To obtain $OptChar$ variables, we first estimate a simple regression model, without an intercept, of option values on simple quantities using all limit orders for each stock, separately for buys and sells. The $OptChar$ variables represent the difference between the actual asymmetry in option values and the predicted asymmetry from this model for each of three price position categories of limit orders (at, away and inside the quotes). We consider five-minute, one-hour, and one-trading day periods for k . All quantity asymmetry measures represent the difference in aggregate quantities on the buy and sell side. Order option values are computed using the Black-Scholes pricing formula with a risk-free interest rate of 7.5%, historic stock return volatility, and maturity equal to the expected time to removal estimated from the regression model presented in Table 2. We also compute option values using the standard barrier option pricing formula with similar inputs except that we use the expected time to cancellation alone for order maturity. Figures in parentheses represent results based on the barrier option model. The percentage of stocks that have significant negative and positive coefficients (at 5% confidence level under a two-tailed test) is reported for each trading activity group. The column ‘Total Book’ presents results for the above regression (run separately) using aggregate quantity and residual option value asymmetry. We use average daily number of trades to measure trading activity. The cross-sectional coefficient averages and standard errors for both quantity and residual option characteristic asymmetry are scaled by a factor of 100 only for presentation.

$k = 5$ minutes									
	Lagged return β_0	Quantity order book asymmetry variables				Residual option characteristic asymmetry			
		Away β_1	At quote β_2	Inside β_3	Total book	Away β_4	At quote β_5	Inside β_6	Total book
Cross-sectional	−0.28*	0.00	0.50*	0.56	0.01*	−0.18	−0.16	3.18	0.33
Average	(−0.27*)	(0.01)	(0.05*)	(0.46)	(0.01*)	(5.16)	(7.22*)	(19.04)	(1.25*)
Cross-sectional	0.02	0.00	0.01	0.08	0.00	0.37	0.25	2.14	0.21
Standard error	(0.02)	(0.01)	(0.01)	(0.09)	(0.00)	(4.12)	(2.97)	(38.32)	(0.40)

Table 4 (Continued)

	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	N
All stocks	83 (79)	3 (2)	6 (4)	17 (17)	1 (1)	51 (50)	1 (4)	68 (60)	1 (1)	32 (30)	7 (4)	6 (10)	14 (4)	3 (5)	28 (18)	19 (17)	5 (1)	29 (35)	139 (138)
<i>Activity group</i>																			
Least	75 (67)	5 (2)	5 (2)	16 (21)	0 (0)	34 (30)	5 (14)	50 (28)	5 (2)	30 (28)	2 (2)	9 (16)	9 (5)	2 (5)	39 (12)	9 (16)	2 (0)	27 (23)	44 (43)
Moderate	94 (90)	0 (0)	8 (6)	21 (21)	0 (0)	63 (58)	0 (0)	69 (69)	0 (2)	35 (29)	17 (8)	2 (10)	17 (4)	4 (4)	23 (15)	23 (23)	8 (4)	29 (48)	48 (48)
Most	79 (79)	4 (4)	6 (2)	13 (11)	2 (2)	55 (60)	0 (0)	83 (81)	0 (0)	30 (32)	2 (2)	9 (4)	15 (4)	2 (6)	23 (28)	23 (13)	4 (0)	30 (32)	47 (47)
<i>k</i> = 1 hour																			
	Lagged return β_0		Quantity order book asymmetry variables								Residual option characteristic asymmetry								
			Away β_1		At quote β_2		Inside β_3		Total book		Away β_4		At quote β_5		Inside β_6		Total book		
Cross-sectional	−0.06*		0.03*		0.24*		1.47*		0.05*		2.94		0.38		−11.47		2.34*		
Average	(−0.05*)		(0.04*)		(0.19*)		(0.95*)		(0.04*)		(43.64*)		(39.32*)		(60.78)		(8.91*)		
Cross-sectional	0.01		0.01		0.03		0.28		0.01		1.63		0.70		10.83		0.62		
Standard error	(0.01)		(0.01)		(0.03)		(0.24)		(0.01)		(15.57)		(11.52)		(195.72)		(1.84)		
	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	N
All stocks	51 (53)	10 (10)	12 (10)	38 (33)	1 (3)	70 (61)	2 (4)	58 (46)	4 (7)	52 (51)	10 (14)	29 (30)	22 (5)	9 (20)	24 (14)	16 (11)	6 (5)	50 (54)	139 (138)
<i>Activity group</i>																			
Least	64 (63)	5 (7)	20 (12)	30 (26)	0 (2)	73 (63)	5 (12)	55 (28)	9 (12)	41 (44)	7 (12)	39 (42)	20 (16)	18 (21)	39 (16)	16 (21)	2 (2)	52 (65)	44 (43)
Moderate	58 (65)	13 (6)	6 (8)	38 (35)	2 (2)	79 (69)	2 (0)	63 (50)	0 (4)	54 (50)	17 (10)	13 (27)	27 (0)	4 (19)	23 (8)	21 (8)	17 (6)	44 (54)	48 (48)
Most	32 (32)	13 (17)	11 (11)	47 (38)	0 (4)	57 (51)	0 (0)	55 (60)	4 (6)	60 (57)	6 (19)	36 (23)	17 (0)	4 (19)	11 (17)	11 (4)	0 (6)	55 (43)	47 (47)

$k = 1$ day

	Lagged return β_0		Quantity order book asymmetry variables								Residual option characteristic asymmetry								N
			Away β_1		At quote β_2		Inside β_3		Total book		Away β_4		At quote β_5		Inside β_6		Total book		
Cross-sectional	−0.03		0.06		0.72*		1.79*		0.12		9.49*		−2.27		2.80		5.76		
Average	(−0.03)		(0.13)		(0.65*)		(1.65*)		(0.15*)		(198.87*)		(−17.18)		(74.43)		(21.65*)		
Cross-sectional	0.02		0.09		0.18		0.58		0.08		3.57		6.40		13.57		3.64		
Standard error	(0.02)		(0.08)		(0.16)		(0.79)		(0.05)		(73.43)		(65.71)		(438.87)		(9.66)		
	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	N
All stocks	37	31	24	50	7	52	8	22	24	53	23	54	24	21	14	14	21	60	139
	(38)	(28)	(25)	(50)	(4)	(49)	(8)	(27)	(24)	(54)	(20)	(47)	(21)	(20)	(11)	(11)	(15)	(55)	(138)
<i>Activity group</i>																			
Least	43	23	34	36	9	73	11	25	30	43	20	45	27	34	25	18	27	61	44
	(47)	(21)	(26)	(42)	(5)	(70)	(14)	(28)	(23)	(49)	(16)	(49)	(28)	(23)	(19)	(14)	(14)	(63)	(43)
Moderate	33	40	21	50	10	54	8	23	21	58	35	42	29	23	15	17	27	48	48
	(31)	(38)	(21)	(52)	(4)	(46)	(6)	(19)	(21)	(56)	(15)	(54)	(21)	(25)	(8)	(15)	(19)	(50)	(48)
Most	36	30	19	62	2	30	4	17	21	57	13	74	15	6	4	6	9	70	47
	(36)	(23)	(28)	(55)	(4)	(32)	(4)	(34)	(28)	(55)	(28)	(38)	(15)	(11)	(6)	(4)	(13)	(53)	(47)

An asterisk for averages indicates significance at the 5% level.

book. It explains as much or more of future returns than simple quantities, especially in active stocks. In the five-minute returns, the coefficient of option characteristic asymmetry variable is positive and significant in about a third of all active stocks. This number increases to 55 percent in the one-hour returns and to 70 percent in the one-day returns. These results strongly suggest that specialists condition their quotation behavior on limit order book asymmetry. Here too, we find our results to be robust to the option model formulation that we use to compute limit order option values. In the next section, we examine whether they use the order book asymmetry to base their trading decisions as well.

4.4. Limit order imbalances and specialist trading decisions

We use logit analyses to determine whether specialist trading decisions made when an incoming SuperDot market order (or a marketable limit order²¹) arrives depend on imbalances in their limit order books.²² As described above, we consider the following four options:

1. Let the order trade with limit orders in the order book.
2. Stop the order.
3. Participate in the trade at the quoted price.
4. Participate in the trade and improve price.

The sequence of these decision options provides a natural ordering of the aggressiveness of specialist actions. A specialist who does not want to trade will let market orders go to his book. If he is ambivalent about trading, he will stop orders. If he wants to trade, he will do so if possible at the quoted price. If he is eager to trade, he may improve the price. This natural ordering allows us to use ordered logit to jointly model trading decisions.²³

Exchange regulations often limit the specialist's options. We therefore classify our analyses by the following two binary conditions:

²¹Peterson and Sirri (2002) document that market orders and marketable limit orders often are used by different sets of traders and under different market conditions. As such, they provide the specialists with another source of information beyond what is available in the limit order book. Though it is quite possible that these two sources of information are correlated, making them less distinct, we would like to retain our focus on the informativeness of the limit order book to the specialist. For this purpose, we feel that, conditional upon their arrival, both market and marketable limit orders provide specialists with similar trading opportunities to exploit their information about the book. Our results using only market orders confirm our conjecture.

²²Since we are interested in the information content of the limit order book, we do not consider actions that the specialist takes following his stopping of a market order. Stopping allows the specialist to postpone his decision to trade while he waits to see new information from the incoming order flow or from correlated asset values. We therefore cannot confidently attribute an action that the specialist may take on a stopped order to information in the limit order book.

²³For estimation of ordered logit models, see Maddala (1983).

- A. Could the specialist stop the incoming market order? Before March 1991 (and hence, relevant to our study), specialists could stop market orders only when the quoted spread was greater than the minimum tick for that security.
- B. Could the specialist participate in the incoming market order without improving the quoted price? If booked limit orders are at the best quoted price, and if their aggregate size is larger than the incoming market order, the public order precedence rule requires that the specialist improve price to trade the order.²⁴

These two marginal conditions imply four pairs of joint conditions. We separately analyze specialist decisions for each of these pairs.

The two conditions sometimes preclude some specialist options. For example, if the specialist cannot participate without improving price (i.e., Condition B does not hold) then the specialist cannot exercise Option 3 (participate at the quoted price). Likewise, whenever Condition A does not hold, the specialist cannot exercise Option 2.

Some of the stocks had few observations for one or more pairs of joint conditions. We therefore were unable to estimate the logit model for every stock-condition pair in our sample. Also, in some cases, for a given pair of conditions, the specialist in a given stock almost always took the same action. When there is not enough variation in the dependent variable, the logit model is not identified. To mitigate such problems, we estimate the logit model only when the specialist did not exercise the same decision option more than $N-50$ times where N is the number of orders for which the two conditions applied. We chose $N-50$ arbitrarily though the results do not change much for other cut-offs.

Table 5 summarizes the four ordered logit models and provides some statistics that characterize the frequencies of market orders that fell within each paired classification, and the frequency of specialist actions within each paired classification. There was roughly an even distribution of market orders across the four conditions that characterize our four models. Within each model, we find the specialist to use more than one decision option frequently. Only when the spread was equal to the minimum tick, and when the aggregate size of limit orders at the quote was larger than the incoming market order, did the specialists choose overwhelmingly not to participate, leaving us with only five stocks with adequate data to estimate the logit model.

We use the following explanatory variables in the logit analyses:

<i>Relative size</i>	The logarithm of the ratio of incoming market order size to the quoted size on the relevant side of the market. (Bid size for sell orders and offer size for buy orders.) We expect that
----------------------	--

²⁴We consider quantities of both limit orders at the quotes as well as of limit orders inside the quotes to determine the applicability of this condition. When there are (hidden) limit orders inside the quotes, an additional condition is imposed upon the specialist in that the size of the price improvement needed for him to trade is bound by a minimum amount (the difference between the best quote and the inside limit order's price). Including this additional condition improves the sophistication of our analysis without changing the qualitative nature of our results.

	specialists will offer less liquidity to larger orders than smaller orders. However, affirmative obligations of the specialists may often force them into trading portions of large orders, even when they do not want to, to prevent the order from filling at multiple prices
<i>Excess spread</i>	The excess of the spread over the minimum tick, expressed in number of ticks. We expect that specialists will trade more aggressively (stop orders, participate and improve prices) when the spread is larger than smaller. This variable does not appear in Models 2 and 4 because both models are conditioned on the spread being equal to one tick.
<i>Signed inventory</i>	The specialist's inventory (in shares) preceding the trade multiplied by -1 if the incoming market order is a sell order. ²⁵ We expect that specialists will trade more aggressively if trading would restore their inventories.
<i>Signed order book asymmetry</i>	The summed difference between buy and sell order quantities in the order book multiplied by -1 if the incoming market order is a buy order. We decompose this variable further into asymmetry in limit orders away from, at, or inside the quotes. We expect that specialists will trade more aggressively when they can step in front of the heavy side of the book.
<i>Signed residual option characteristic asymmetry</i>	The summed difference between buy and sell order "option characteristics" in the order book multiplied by -1 if the incoming market order is a buy order. The option characteristic variable—described above—captures information contained in option values but not in simple quantities. As with the order book asymmetry, we decompose this variable into option characteristic asymmetry in limit orders away from, at, or inside the quotes. We expect that specialists will trade more aggressively when they can step in front of the heavy side of the book as reflected in their option values.

Table 6 summarizes the ordered logit model estimation results by tabulating the numbers of stocks with positive and negative estimated coefficients for all stocks and for stocks classified by trade frequency. Results are reported for the four models separately. Model 1 describes specialist actions when Exchange regulations allowed the specialist to stop the order and participate without improving the price. This was true for 28 percent of all market orders in the sample.

²⁵We construct this variable using all specialist trades (including non-SuperDot trades) but we do not know beginning-of-sample specialist inventories. Fortunately, the model intercept captures the effect of the omitted opening inventory.

Table 5
Ordered logit models of specialist discretionary actions

This table presents the conditions and ordered dependent variables of the four logit models of specialist discretionary actions following the arrival of a market orders. Condition A is true if the spread was greater than the minimum tick so that the specialist could stop the order. Condition B is true if the specialist was at least partly behind the quotes so that he could participate in the incoming market order without improving the price. The estimates of these models appear in Table 6. A given logit model was estimated for a given stock if the specialist did not exercise the same decision option more than $N-50$ times where N was number of orders for which the two conditions applied.

Model	Conditions		C/S average frequency of market orders that met these conditions	Number of estimated models	C/S average number of market orders used in estimation	C/S minimum number of market orders used in estimation	Ordered decision option	C/S average frequency of specialist decisions within this cell
	A Spread > tick	B Specialist at the quote						
1	True	True	28%	80	608	80	0—Did not participate 1—Participated by stopping 2—Participated without improving price 3—Participated by improving price	27% 30 31 12
2	False	True	19	56	1,284	116	0—Did not participate 1—Participated without improving price 2—Participated by improving price	67 32 1
3	True	False	22	69	876	83	0—Did not participate 1—Participated by stopping 2—Participated by improving price	49 36 15
4	False	False	31	5	10,918	4,148	0—Did not participate 1—Participated by improving price	98 2

Table 6

Results of ordered logit models of specialist discretionary actions

This table presents results of four ordered logistic regression models of specialist discretionary actions taken upon the arrival of SuperDOT market and marketable limit orders. The ordering of the dependent variable measures the aggressiveness of the specialist in his actions with higher values indicating greater aggressiveness. Only market and marketable limit orders that satisfy the following conditions are included: (1) The order must be a simple buy or sell order, (2) The order must not be an opening order, (3) The order must have matched contra-party information in the TORQ audit file, and (4) The order must be entered when the bid price was greater than \$1. The independent variables are: (1) *Relative size*, the logarithm of the ratio of order size of the incoming market order over the relevant quote depth (bid depth for a sell order and ask depth for a buy order), (2) *Excess spread*, the excess of spread over the minimum tick expressed in number of ticks, (3) *Signed inventory*, a constructed variable to represent specialist's inventory (in shares) preceding the trade multiplied by -1 if the order is a sell order, (4) *Signed order book asymmetry*, the difference between aggregate quantities of limit buy orders and limit sell orders multiplied by -1 if the order is a buy order, and (5) *Signed residual option characteristic asymmetry*, the residual option characteristic asymmetry variable discussed in Table 3, multiplied by -1 if the order is a buy order. Order option values are computed using the Black-Scholes pricing formula with a risk-free interest rate of 7.5%, historic stock return volatility, and maturity equal to the expected time to removal estimated from the regression model presented in Table 2. We also compute option values using the standard barrier option pricing formula with similar inputs except that we use the expected time to cancellation alone for order maturity. Figures in parentheses represent results based on the barrier option model. Both the *Signed order book asymmetry* and *Signed residual option characteristic asymmetry* variables are further decomposed based on the price position of limit orders (away from, at or inside the quotes) with respect to the prevailing market quote. The spread variable is not included in Models 2 and 4 as these models presuppose that the spread is equal to one tick. Likewise, there are no asymmetry variables for limit orders inside the quotes for these two models. Each model is estimated for each stock individually, but results are reported as percentage of stocks that have significant negative or positive coefficients at the 5 percent level under a two-tailed *t*-test. Results are further reported by trading activity groups, formed using the average daily number of trades. The column 'Total Book' presents results for the above regression (run separately) using aggregate signed quantity and residual option characteristic asymmetry.

Model 1. Exchange regulations *allowed* the specialist to stop the order (the spread was greater than one tick) and participate without improving price (the specialist was at least partially behind the quotes). The dependent variable values are: 0—did not participate, 1—stopped order, 2—participated without improving price, and 3—participated by improving price.

	Relative size		Excess spread		Signed inventory		Signed order book asymmetry								Signed residual option characteristic asymmetry								N
							Away		At quote		Inside		Total book		Away		At quote		Inside		Total book		
	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	
All stocks	0 (1)	71 (64)	6 (1)	44 (33)	11 (13)	15 (9)	15 (16)	5 (4)	4 (3)	12 (20)	4 (3)	15 (16)	5 (11)	6 (5)	9 (9)	4 (11)	1 (6)	15 (4)	7 (9)	10 (4)	6 (4)	20 (11)	80 (80)
Activity group																							
Least	0 (14)	29 (29)	0 (0)	71 (43)	29 (0)	29 (0)	0 (0)	14 (0)	0 (0)	0 (0)	0 (0)	0 (14)	0 (0)	13 (0)	14 (0)	0 (0)	0 (0)	29 (0)	0 (0)	14 (0)	13 (0)	13 (0)	7 (7)

Moderate	0	58	3	52	7	19	19	3	3	3	7	7	3	0	10	3	3	7	7	13	6	16	31
	(0)	(52)	(3)	(45)	(3)	(10)	(21)	(3)	(3)	(10)	(0)	(3)	(16)	(6)	(10)	(7)	(0)	(7)	(3)	(7)	(6)	(6)	(29)
Most	0	86	9	34	11	9	14	5	5	21	2	23	7	9	7	5	0	18	9	7	5	25	44
	(0)	(77)	(0)	(23)	(20)	(9)	(16)	(5)	(2)	(30)	(5)	(25)	(9)	(5)	(9)	(16)	(11)	(2)	(14)	(2)	(2)	(16)	(44)

Model 2. Exchange regulations did not allow the specialist to stop the order (the spread was equal to one tick) but permitted the specialists to participate without improving price (the specialist was at least partially behind the quotes). The dependent variable values are: 0—did not participate, 1—participated without improving price, and 2—participated by improving price.

	Relative size		Signed inventory		Signed order book asymmetry						Signed residual option characteristics asymmetry								N
					Away		At quote		Total book		Away		At quote		Total book				
	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos			
All stocks	0	93	14	16	18	14	4	30	(14)	14	14	18	5	18	9	25	56		
	(0)	(77)	(5)	(11)	(16)	(11)	(4)	(32)	(14)	(11)	(11)	(18)	(7)	(5)	(7)	(21)	(56)		
Activity group																			
Least	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)		
Moderate	0	91	9	18	9	9	18	18	9	0	0	0	9	9	9	9	11		
	(0)	(73)	(9)	(9)	(9)	(0)	(9)	(18)	(0)	(0)	(0)	(0)	(9)	(18)	(0)	(9)	(11)		
Most	0	93	16	16	20	16	0	33	16	18	18	22	4	20	9	29	45		
	(0)	(78)	(4)	(11)	(18)	(13)	(2)	(36)	(18)	(13)	(13)	(22)	(7)	(2)	(9)	(24)	(45)		

Model 3. Exchange regulations allowed the specialist to stop the order (spread was greater than one tick) but did not permit the specialist to participate without improving price (the limit orders were behind the quotes). The dependent variable values are 0—did not participate, 1—stopped the order, and 2—participated by improving price.

	Relative size		Excess spread		Signed inventory		Signed order book asymmetry								Signed residual option characteristic asymmetry								
							Away		At quote		Inside		Total book		Away		At quote		Inside		Total book		
	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	N		
All stocks	33 (35)	4 (7)	3 (4)	57 (54)	10 (14)	12 (9)	25 (28)	12 (9)	0 (1)	32 (39)	15 (13)	3 (6)	13 (22)	19 (17)	10 (19)	13 (9)	4 (10)	23 (10)	1 (1)	12 (4)	1 (6)	28 (13)	69 (69)

Table 6 (Continued)

	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	N		
Activity group																					
Least	0	0	0	0	0	100	0	100	0	0	0	0	0	0	0	100	0	0	0	100	1
	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(100)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(1)
Moderate	46	0	0	77	9	9	27	5	0	5	5	0	9	9	9	0	0	23	5	0	22
	(50)	(0)	(0)	(73)	(5)	(9)	(32)	(5)	(0)	(23)	(14)	(0)	(27)	(9)	(18)	(9)	(23)	(5)	(0)	(5)	(22)
Most	28	7	4	48	11	11	24	13	0	46	20	4	15	24	11	17	7	22	0	17	46
	(28)	(11)	(7)	(46)	(20)	(9)	(26)	(11)	(2)	(48)	(13)	(9)	(17)	(22)	(20)	(9)	(4)	(13)	(2)	(4)	(46)

Model 4. Exchange regulations did not allow the specialist to stop the order (the spread was equal to one tick) and did not permit the specialist to participate without improving price (the limit orders were behind the quotes). The dependent variable values are 0—did not participate, 1—participated by improving price.

	Relative size		Signed inventory		Signed order book asymmetry						Signed residual option characteristics asymmetry								N
					Away		At quote		Total book		Away		At quote		Total book				
	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos			
	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos	Neg	Pos			
All stocks	20	20	0	0	40	0	0	100	40	0	60	20	0	0	20	20	5		
	(20)	(20)	(0)	(20)	(40)	(20)	(0)	(100)	(20)	(20)	(40)	(20)	(20)	(0)	(40)	(40)	(5)		
Activity group																			
Least	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)		
Moderate	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)		
Most	20	20	0	0	40	0	0	100	40	0	60	20	0	0	20	20	5		
	(20)	(20)	(0)	(20)	(40)	(20)	(0)	(100)	(20)	(20)	(40)	(20)	(20)	(0)	(40)	(40)	(5)		

Our results indicate that specialists do condition their actions based on the informativeness of the limit order book. Interestingly it is the option values that mattered more than the simple quantity asymmetry of the limit order book. For about a fifth of all stocks, the specialist aggressively offered liquidity ahead of a heavy limit order book as characterized by their option values. The numbers were distinctly higher for active stocks than for less active stocks. There are at least two reasons for this. The first reason has to do with the smaller number of observations used for estimation for the less active stocks. The limit order book, for such stocks, is more likely to be thin or comprised of orders placed far away from the market (where the option has little or no value). The second reason is that limit order imbalances in less active stocks are less informative as well, as shown in Table 3.

In the most active stocks, we find that the specialists aggressively sought to offer liquidity to market orders in a quarter of all stocks in front of a heavy limit order book. This suggests that it is not the lack of liquidity in the book that induces specialists to offer liquidity.

Further decomposition suggests that the option values of the limit orders placed at the quotes are most informative. Order book asymmetry, expressed in quantities, is informative, but not as much as is asymmetry in option values. Interestingly, quantity asymmetry among aggressively priced limit orders seems to matter to specialist trading decisions more in the active stocks. This suggests that display of limit orders close to the market, such as the recently introduced NYSE OpenBook™, may diminish some, but not all, of the trading opportunities of the specialist.²⁶

While we find specialists to be more aggressive in their actions when the spread is wide, we do not find inventory considerations to be important, much like other researchers in the past (see, for example, Madhavan and Panchapagesan, 2000). We find the specialist to be more aggressive in offering liquidity to larger orders than smaller orders suggesting that affirmative obligations may play an important role in specialist trading decisions.

We get similar, though slightly weaker, results when we use the barrier option pricing formulation to characterize limit order option values. But we do find clear evidence that specialists are likely to stop, or trade ahead of, a heavy limit order book as characterized by their option values irrespective of how we compute them.

Model 2 describes specialist actions where Exchange regulations did not allow the specialist to stop the order but permitted him to participate without improving the price. These conditions were satisfied in about 19 percent of all market orders. Here too, we find residual option characteristic asymmetry to be more important in explaining specialist actions than simple quantity asymmetry in the limit order book. Option characteristic asymmetry is a significant explanatory variable in predicting specialist trading decisions in about a quarter of all stocks (a fifth of all stocks when we use the barrier option pricing formulation) while simple quantity asymmetry is

²⁶Boehmer et al. (2003) present evidence that NYSE specialist participation rate declined following the introduction of the OpenBook.

significant in only half of those stocks. We again find stronger results in active stocks than in inactive stocks. Moreover, specialists seem to condition their trading decisions more on orders at the quotes than on orders far from the market. These results suggest that option values are indeed important, and that dissemination of simple quantity imbalances by the specialist to other traders may not fully dilute the informational advantage.

As in Model 1, the relative size variable, which measures the order size relative to the existing quote size, is again positive for most stocks. The signed inventory variable is positive and negative in an equal number of stocks, indicating that inventory is not an important factor, at least in the short run, to specialists in their decision to offer liquidity.

Models 3 and 4 describe specialist actions when they can participate only by improving the price. Public precedence rules prohibit members, including specialists, from trading ahead of public traders at the same price. Model 3 describes actions when the specialist is permitted by Exchange regulations to stop a market order, while Model 4 studies actions when he is not allowed to stop a market order.

The residual option characteristic variable is important in determining specialist actions when the specialist could stop the market order (Model 3) but not when he was disallowed from stopping (Model 4). As mentioned earlier, our results for Model 4 are less meaningful as there were only five stocks where the specialist chose more than one option enough times to reasonably estimate the model. As in earlier models, the simple quantity asymmetry variable is less important in explaining specialist actions. Among all orders in the book, the orders at the market are the most informative. As before, our results are stronger for the active stocks than for the inactive stocks.

To sum up, our results suggest that specialists use information from the limit order book to the detriment of limit order traders. In support of the option characterization of limit orders, we find that specialists use information derived from limit order option values rather than order sizes in their trading decisions. The results appear robust to the different option models that we use to characterize limit order option values. On the other hand, specialists face affirmative obligations that may erode some of this advantage. Since we do not focus on their obligations, we refrain from making an overall evaluation of the specialist system. Besides limit order book information, relative order size and spread also seem to influence specialists' trading decisions. Inventory does not seem to be an important factor, a fact that is consistent with many previous studies.

These regressions do not highlight the cross-sectional differences in specialist actions that are equally important. Harris (1990) discusses the importance of a security's relative tick size (tick size over price) to liquidity providers of that security. A higher relative tick size makes it difficult for traders, including the specialist, to capture option values inherent in limit orders through quote-matching strategies. Recent reports of "penny-jumping" by specialists and other traders indicate that a reduction in tick size may make it easier for specialists to profit at the expense of limit order traders. In the next section, we examine whether relative tick affects specialist discretionary actions.

4.5. *Cross-sectional analysis of specialist discretionary actions*

To determine whether relative tick size affects specialist trading decisions, we examine whether specialists tend to step ahead of the book more often in stocks with low relative tick than in stocks with high relative tick. Upon the arrival of a market order, a specialist can step ahead of the book two ways: he can stop the market order, or he can trade with it immediately. To examine only the discretionary aspect of specialist actions, we consider only market orders that arrive when Exchange regulations prohibit the specialist from trading without improving the price. As mentioned before, price priority and public precedence rules ensure that the specialist cannot trade ahead of a public order at the same price. We focus therefore on market orders that arrive when limit orders, and not the specialist, are behind the quotes.

The dependent variable in our cross-sectional regression is the percentage of market orders that the specialist stopped or traded with by improving the price for each stock. This measure best captures the propensity of the specialist to step ahead of his limit order book. Our explanatory variables include market capitalization, daily return standard deviation, and relative tick. Market capitalization is measured on October 31, 1990, the last day of the pre-sample period, and is reported in millions of dollars. Consistent with relative participation rates in small and large stocks documented in this study and others, we expect specialists to be more aggressive in offering liquidity in smaller stocks than in larger stocks. We use daily return standard deviation of log close-to-open returns to characterize the underlying volatility of the stock. We believe volatility to be important in explaining some of the variations in specialist actions across stocks. The relative tick is the ratio of tick size over the average price of the security in the sample period.

If relative tick size matters in the competition between limit orders and floor traders, then it should be evident in the cross-sectional regressions. Specialists are more likely to act against limit orders when the relative tick is small. This suggests a negative coefficient for relative tick in the regressions.

Table 7 presents results of the cross-sectional regression of the specialist propensity to step ahead of the limit order book. The coefficient for relative tick is negative and significant, indicating that the tick size affects the competition between the specialist and limit order traders to offer liquidity. Our results strongly support the argument raised in Harris (1990) that small tick sizes hurt limit order traders at the expense of faster, and more informed traders, such as the specialist. While other studies have documented a reduction in the use of limit orders following decimalization, we provide the first evidence of how some traders could exploit limit orders in the presence of a small tick size.

The coefficients of both market capitalization and daily return standard deviation are of the right sign but are not statistically significant. Our model R^2 is 26 percent, indicating that our independent variables are important in explaining differences in specialists' propensity to step ahead of their limit order books across stocks.

Table 7

Results of cross-sectional regression of specialist discretionary actions

This table presents results of a cross-sectional OLS regression model of specialist discretionary actions on incoming SuperDOT market and marketable limit orders in the TORQ database. Only market and marketable limit orders that satisfy the following conditions are included: (1) The order must be a simple buy or sell order; (2) The order must not be an opening order; (3) The order must have matched contra-party information in the TORQ audit file; and (4) The order must be entered when the bid price was greater than \$1. The dependent variable is the propensity of the specialist to step in front of a limit order book for each stock. We measure this propensity as the percentage of eligible market orders that the specialist stopped or traded by improving price, thereby preventing such orders from trading with the book. Our sampling universe includes only those market orders that satisfy the above criteria and are placed when limit orders are behind the quotes. Public precedence rules constrain the specialist from trading ahead of the limit orders without improving the price. The explanatory variables are: market capitalization (in millions of dollars), Standard deviation of daily log returns and Relative tick (tick over the stock's price as a percentage). Standard errors of the estimates are given in parentheses.

Dependent variable: percentage of market orders that the specialist stopped or traded with when Exchange regulations *do not permit* the specialist to participate without improving price (when limit orders are behind the quotes).

Intercept	Market capitalization (\$ M)	Daily return standard deviation	Relative tick	R ²	N
0.265* (0.017)	−0.002 (0.001)	1.724 (0.936)	−0.048* (0.010)	0.26	134

An asterisk represents significance at 5% level.

5. Conclusions

Specialists compete with limit orders for the provision of liquidity at the New York Stock Exchange. Though the Exchange uses a variety of mechanisms to regulate this competition, specialists have a unique advantage. They know the expressed trading interest of all traders, while a limit order trader only knows about his own order. If the aggregate order information is valuable, then specialists can profit from it through selective participation in trades, thereby exacerbating the adverse selection problem for limit orders. In this paper, we examine for the first time whether the limit order book is informative, and whether specialists use this information in their discretionary actions. By discretionary actions, we mean all actions that the specialist is not required to take as part of his obligations to the Exchange.

We characterize information in the book through asymmetry between buys and sells. We use asymmetry in simple quantities as well as asymmetry in the option values provided by limit orders. We use a simple Black-Scholes formulation to characterize limit order option values though using a more sophisticated barrier option pricing formulation does not change the qualitative nature of our results. We find that the limit order book is informative about future price changes. Both asymmetry in quantities and in option values help explain future returns. We find strong evidence that the specialist uses information from option values more than from quantities in his decision to stop a market order, or trade with it. This suggests

that individual order properties such as duration, price relative to the market, and order size have information content of which only the specialist is likely to be aware. These results are particularly relevant now as the NYSE opens its limit order book by displaying only aggregate quantities to the public. Our results are stronger for active stocks where the competition between the specialist and the limit orders is more intense. We also find strong evidence for such actions to be more likely in low-priced stocks than high-priced stocks because of the binding tick size as suggested by Harris (1990).

Our option characterization of limit orders undoubtedly only crudely captures limit order option properties. The fact that we obtained significant results suggests that better results could be obtained if better empirical measures of their option properties were available.

Though we find that the specialist uses information from the limit order book, we do not address the broader issue of whether such information should be made transparent to all traders. It may be the case that the specialists are given this advantage, as compensation for meeting their Exchange obligations such as providing liquidity when there is none available. We leave this issue for future research, as it requires an analysis of the cost of such obligations that was beyond the scope of this study.

Appendix A. Conditions upon which specialists make various trading decisions

This table presents the various specialist trading options on every market or marketable limit order subject to Exchange regulations. We characterize Exchange regulations in two ways: does it allow the specialist to stop the order, and does it allow the specialist to participate without improving the price. This two-way classification produces four distinct cells within which we analyze specialist trading options. We study four possible specialist options: (1) let the order trade with limit orders in the order book, (2) stop the order, (3) participate at the quoted price, and (4) participate and improve price. The sequence of options reflects the aggressiveness of specialist actions. Given the nature of the constraints imposed by the Exchange, some of these options may not be feasible to the specialist.

Conditions		Feasible specialist trading options
A	B	
Did Exchange regulations allow the specialist to stop the order? (Is the spread greater than 1 tick?)	Did Exchange regulations permit the specialist to participate in the order without improving price? (Is the specialist partially behind the quotes?)	

True	True	1—Let the order trade with the limit order book 2—Stop the order 3—Participate at the quoted price 4—Participate and improve price
False	True	1—Let the order trade with the limit order book 3—Participate at the quoted price 4—Participate and improve price
True	False	1—Let the order trade with the limit order book 2—Stop the order 4—Participate and improve price
False	False	1—Let the order trade with the limit order book 4—Participate and improve price

References

- Abolafia, M., 1996. *Making Markets: Opportunism and Restraint on Wall Street*. Harvard University Press, Cambridge.
- Amihud, Y., Mendelson, H., 1990. Option market integration: an evaluation. Working paper, New York University.
- Battalio, R., Jennings, R., 2002. Getting pennied: the effect of decimalization on traders' willingness to lean on the limit order book at the New York Stock Exchange. Working paper, University of Notre Dame.
- Benveniste, L., Marcus, A., Wilhelm, W., 1992. What's special about the specialist? *Journal of Financial Economics* 32, 61–86.
- Boehmer, E., Saar, G., Yu, L., 2003. Lifting the veil: an analysis of pre-trade transparency at the NYSE, working paper, New York Stock Exchange, Inc.
- Cao, C., Hansch, O., Wang, X., 2003. The information content of an open limit order book. Working paper, Penn State University.
- Chan, Y., Weinstein, M., 1992. Reputation, bid-ask spread and market structure. Working paper, University of Southern California.
- Chung, K., Van Ness, B., Van Ness, R., 1999. Limit orders and the bid–ask spread. *Journal of Financial Economics* 53, 255–287.
- Copeland, T., Galai, D., 1983. Information effects on the bid–ask spread. *Journal of Finance* 38, 1457–1469.
- Edwards, A., 1999. NYSE specialists competing with limit orders: a source of price improvement. Working paper, Securities Exchange Commission.
- Edwards, A., Harris, J., 2001. Stepping ahead of the book. Working paper, University of Notre Dame.

- Goldstein, M., Kavejecz, K., 2000. Eighths, sixteenths and market depth: changes in tick size and liquidity provision on the NYSE. *Journal of Financial Economics* 56, 125–149.
- Harris, L., 1990. Liquidity, trading rules and electronic trading systems. New York University Salomon Center Monograph Series in Finance, Monograph 1990-4.
- Harris, L., 1998. Does a large minimum price variation encourage order exposure? Working paper, University of Southern California.
- Harris, L., Hasbrouck, J., 1996. Market vs. limit orders: the SuperDot evidence on order submission strategy. *Journal of Financial and Quantitative Analysis* 31, 219–232.
- Hasbrouck, J., 1992. Using the TORQ database. Working paper, New York Stock Exchange, Inc.
- Hasbrouck, J., Sofianos, G., 1993. The trades of market makers: an empirical analysis of NYSE specialists. *Journal of Finance* 48, 1565–1594.
- Haug, E., 1998. *The Complete Guide to Option Pricing Formulas*. Irwin, Homewood, IL, McGraw-Hill, New York.
- Huang, R., Stoll, H., 1994. Market microstructure and stock return predictions. *Review of Financial Studies* 7, 179–213.
- Ip, G., 2001. Big board specialists make money despite being an anachronism. *Wall Street Journal*.
- Jarnecic, E., McNish, T., 1997. An empirical investigation of the option value of the limit order book on the Australian Stock Exchange. Working paper, University of Sydney.
- Kavejecz, K., 1999. A specialist's quoted depth and the limit order book. *Journal of Finance* 54, 747–771.
- Lee, C., Ready, M., 1991. Inferring trade direction from intraday data. *Journal of Finance* 46, 733–746.
- Lo, A., MacKinlay, C., Zhang, J., 2002. Econometric models of limit-order executions. *Journal of Financial Economics* 65, 31–71.
- Maddala, G., 1983. *Limited Dependent and Qualitative Variables in Econometrics*. Cambridge University Press, Cambridge.
- Madhavan, A., Panchagesan, V., 2000. Price discovery in auction markets: a look inside the black box. *Review of Financial Studies* 13, 627–658.
- Madhavan, A., Sofianos, G., 1998. An empirical analysis of NYSE specialist trading. *Journal of Financial Economics* 48, 189–210.
- McNish, T., Wood, R., 1995. Hidden limit orders on the NYSE. *Journal of Portfolio Management* 19–26.
- New York Stock Exchange, 1999. *New York Stock Exchange Fact Book*. New York Stock Exchange, New York.
- Panchapagesan, V., 2000. Identifying specialist trades in the TORQ data—a simple algorithm. Working paper, Washington University in St. Louis.
- Peterson, M., Sirri, E., 2002. Order submission strategy and the curious case of marketable limit orders. *Journal of Financial and Quantitative Analysis* 37, 221–241.
- Ready, M., 1999. The specialist's discretion: stopped orders and price improvement. *Review of Financial Studies* 12, 1075–1112.
- Seppi, D., 1997. Liquidity provision with limit orders and a strategic specialist. *Review of Financial Studies* 10, 103–150.
- Sofianos, G., 1995. Specialist gross trading revenues at the New York Stock Exchange. Working paper, New York Stock Exchange, Inc.
- Solomon, D., Craig, S., 2003. SEC blasts big board oversight of 'specialist' trading firms. *Wall Street Journal*.