The Specialist's Discretion: Stopped Orders and Price Improvement

Mark J. Ready
University of Wisconsin-Madison

When a market order arrives, the NYSE specialist can offer a price one tick better than the limit orders on the book and trade for his own account. Alternatively, the specialist can "stop" the market order, which means he guarantees execution at the current quote but provides the possibility of price improvement. My model shows that specialists can use stops to sample the future order flow before making a commitment to trade. I present empirical evidence that both stops and immediate price improvement impose adverse selection costs on limit order traders.

Supporters of the NYSE claim that the role of the specialist is to facilitate trade; he is there to jump in when there is too little liquidity to prevent unwarranted swings in the price. In addition, supporters claim that the specialist helps to bring buyers and sellers together in cases where neither is comfortable with unilaterally exposing the size of their desired trade. Finally, supporters insist that the specialist's own trading activity gives better execution prices to incoming orders because the specialist must improve the best available limit order price in order to trade for his own account. Detractors claim that, if anything, the specialist gets in the way of willing buyers and sellers, and that the specialist's own trading activity worsens the adverse selection cost borne by the limit order traders. Supporters of the NYSE counter that this type of activity is controlled by the exchange's own regulations, including periodic specialist performance evaluations. The purpose of this article is to shed some light on this debate.

In this article, I model the interactions between market participants, including the NYSE specialist, "liquidity"[1] traders who trade for portfolio reasons, and profit-motivated outsiders who trade either because they have superior information, or because the design of the exchange allows them to profitably compete with the specialist. In my model, trading is driven by the exogenous demands of the informed and liquidity traders. In this framework, trading is inherently a zero-sum game, so it is not possible to make

I thank Jim Angel, Craig Holden, Charles Lee, and Robert Seijas for helpful comments and suggestions. I gratefully acknowledge financial support from the Puelicher Center for Banking Education, the Institute for Quantitative Research in Finance and the Wisconsin Alumni Research Foundation. Address correspondence to Mark J. Ready, University of Wisconsin-Madison, 975 University Ave., Madison, WI 53706, or email: mjready@facstaff.wisc.edu.

[1] The term "liquidity" follows Glosten and Milgrom (1985) and emphasizes the traders' motivation. Kyle (1985) uses "noise" and Easley and O'Hara (1992b) use "uninformed" to emphasize the traders' lack of information.

The Review of Financial Studies Winter 1999 Vol. 12, No. 5, pp. 1075–1112

definitive welfare comparisons between different market structures. To do so would require an understanding of how each trader's demands change in response to changes in the market mechanism. In spite of this limitation, my model provides important insights into the competition between the NYSE specialist and the limit order traders, and it allows a comparison between the NYSE and an electronic limit book from the standpoint of the trading costs and/or trading profits for each type of market participant.

One of the main differences between the structure of the NYSE and that of an open electronic limit book is the ability of the NYSE specialists (and floor traders) to set their pricing schedules after seeing some of the order flow. Posted limit orders are precommitted to trade, and they are the sole suppliers of liquidity in an open electronic limit book. In contrast, the NYSE specialist has the option to trade for his own account by offering a price that is one tick[2] better than the best limit order. In addition, the specialist can "stop" the incoming market order, which means he guarantees execution at the current posted quote but provides the market order with a chance for price improvement. Once he has stopped a market order, the specialist can observe the subsequent order flow before making the decision of whether to trade for his own account.[3]

Liquidity traders who use market orders on the NYSE appear to benefit from stops and immediate price improvement, at least as compared to certain execution at the posted quote. In order to compare the NYSE with an electronic limit book, however, it is also necessary to understand how the specialist's activity affects where the limit order traders choose to place their orders. The theoretical model that I use to make this comparison is an extension of Easley and O'Hara (1992b). The model contains a risk-neutral specialist, liquidity traders whose demands are random and exogenously determined, and informed traders who use market orders to trade if the market price differs from their valuation based on a private signal. The model allows the specialist to stop incoming market orders. There is price discreteness, so if the specialist wishes to trade for his own account he must offer price improvement of one tick compared to the best limit order.

The equilibrium position of the limit orders depends on their motivation, and I consider two types of limit order traders. Those of the first type simply supply liquidity in an attempt to make profits, whereas those of the second type are committed to trade but are using limit orders instead of market orders in an attempt to lower their expected execution cost. The second type of limit order trader faces the risk that the limit order won't

[2] Prior to June 1997, the minimum price tick was \$1/8 for most NYSE stocks. In June 1997, the NYSE began trading and quoting in \$1/16 increments ("teenies") as a first step toward trading in decimal increments.

[3] Note that the specialist's learning in this article is from passive observation of the order flow. This is in contrast to the active learning modeled by Leach and Madhavan (1993). In their article, the specialist adjusts the quotes and observes the resulting change in the order flow.

execute, forcing him to follow up with a market order at an inferior price. I first calculate the equilibrium assuming limit order traders are all of the first type (called a "profit-motivated" equilibrium), and then I calculate the equilibrium assuming both types are present (a "combined" equilibrium).

My analysis of the model focuses on the parameter values for which the limit order traders optimally place their orders one or two ticks apart under the current rules of the NYSE. These cases should be the most interesting ones from the standpoint of a market designer or regulator, because spreads wider than two ticks are relatively rare,[4] and a one tick spread is the minimum possible. Regardless of the motivation assumed for the limit order traders, the model shows that market order traders tend to benefit from the specialist's ability to stop orders, even after allowing for the fact that the limit order prices would sometimes be more aggressive if stops were forbidden. Of interest, the liquidity traders receive a much larger benefit from the stops than the informed traders. This is because an informed trader's order is more likely to be followed by more orders of the same type, so a stopped order from an informed trader is unlikely to receive price improvement.

I test the empirical implications of the model using the TORQ database, which is a 3-month sample of trades, quotes, and SuperDOT orders for 144 NYSE firms.[5] Nearly 30% of the market orders in this sample were stopped, so while the practice is common, it is by no means automatic. As predicted by my model, I find that orders that are kept from hitting the quotes, by either a stop or immediate price improvement, are those that would have been the most profitable for the limit order traders. Moreover, the magnitudes of these adverse selection effects are in line with the model's predictions. As predicted by my model, I also find that stops are *most* common in \$1/4 spread markets when the last trade was at the midpoint of the current spread.

The remainder of this article is organized as follows. In Section 1, I develop the model of the specialist's optimal use of stops and immediate price improvement, examine the costs and benefits of stops for each type of market participant, and develop the empirical implications. Section 2 describes the data used for this study and presents some summary statistics. I present the empirical tests of the model's predictions in Section 3, and Section 4 shows that the results from Section 3 can be used to improve the estimation of expected execution prices. Section 5 concludes. Most derivations are included in the appendix. Results that are referenced but not shown are available on request.

[4] This is certainly true for the November 1990–January 1991 period examined in this article, and for the January 1995 data examined by Ross, Shapiro, and Smith (1996). With the advent of a \$1/16 minimum tick in June 1997, it remains to be seen how frequently spreads are larger than two ticks.

[5] Harris and Hasbrouck (1996) provide summary statistics on stopped market orders for this sample, but do not investigate the determinants of the specialist's decision.

1. A Model of the Specialist's Strategic Behavior

1.1 Background and Related Literature

The adverse selection problem facing limit orders is modeled by Glosten (1994), Dupont (1995), Seppi (1995), and Rock (1996). It stems from the fact that limit orders will only execute against orders that are large enough to clear all orders ahead of them on the book and will always execute against all orders that are larger. In my model, there is always enough depth at the first limit price to handle the incoming order. Accordingly, whereas the primary focus of these earlier models is the specialist's decision to clean up the remainder of an order at a particular price, the main focus of my model is the specialist's decision to provide an execution price that is better than all of the limit orders.

Limit order competition with the specialist is examined empirically by Handa and Schwartz (1996) and Harris and Hasbrouck (1996). Handa and Schwartz (1996) use simulated limit order trades to show that providing liquidity using limit orders could be profitable, whereas Harris and Hasbrouck (1996) use actual limit order submissions and show that such a strategy would not be profitable. As the authors of both articles point out, one possible explanation for their conflicting results is that Handa and Schwartz's methodology implicitly assumes that the presence of the limit orders will not affect the prices of executed trades. The evidence presented in this article suggests that the presence of limit orders will affect the specialist's use of stops and price improvement, and that the result of this behavior will be lower profit for the limit orders.

1.2 Model Assumptions

There are four types of market participants: the specialist, limit order traders, liquidity traders, and informed traders. Of these, the simplest are the liquidity traders who randomly submit market buy or sell orders for reasons that are exogenous to the model. These orders convey no information about the value of the security. As in Easley and O'Hara (1992b), there is uncertainty about whether an information event has occurred. If there has been an information event, informed traders randomly arrive at the market and buy (sell) if the ask (bid) price is below (above) their private valuation of the security.

I assume that all orders are for the same number of shares, which simplifies the exposition and the numerical solutions of the model. The cost of this simplification is that the model ignores the informed traders' incentive to use larger order sizes to capitalize on their information. Indeed, Knez and Ready (1996) show that larger orders are much less likely to receive price improvement. I address this issue empirically by investigating the impact of order size on the specialist's decision to stop an order. Extension of the model to allow for multiple order sizes, perhaps along the lines of Easley and O'Hara (1992a), is a possible direction for future research.

Limit order traders submit orders to buy and sell the security before any trading occurs. Since the limit order traders move first, the model can handle a variety of motives for their trading. It could be that they trade only for expected profit, or it could be that they are committed to trade and are using a limit order (followed by a market order if necessary) as a way to reduce their expected cost of trading. In the latter case, their trades could result in expected losses, but these losses should be lower than those incurred by the liquidity traders. I discuss this issue in more detail in the next section.

I assume that the limit order traders are unable to determine where their orders sit in the queue. Thus, when they solve the model to determine where their orders should be placed, they consider only the average expected profit (or loss) for all of the orders placed at that price, and there are always enough limit orders at that price to handle incoming market orders. The reality of the NYSE is that the limit book is neither completely opaque nor completely transparent. Although specialists will often provide information about the book to floor brokers who request it, there are both explicit and implicit restrictions placed on the behavior of the broker who is given this information. In addition, the status of the limit order book is not readily available to traders away from the floor. In 1993, after the data used in this article were collected, the NYSE improved the SuperDOT system to allow limit order traders at the quote to find out how many shares are ahead of their order on the book by issuing a "how stand" request. It would be interesting to use more recent data to see whether limit order traders use this feature and cancel their order when they learn of traders ahead of them.

The model has $T + 1$ time periods, indexed by $t \in \{0, 1, 2, \dots, T\}$. At the start of period 0 the limit order traders submit orders to buy and sell, and then the specialist posts the quotes. All prices are restricted to a discrete grid of eighths $P \equiv \{.125n\colon n \in N\}$, where N is the set of natural numbers. The limit orders cannot be withdrawn until after the trade (if any) is cleared at time $t = T$. This assumption captures the real problem that it is costly to continuously monitor limit orders, and is consistent with the empirical observation that few limit orders are withdrawn. In the database used for this article, I find that only 21% of day limit orders are withdrawn before they have executed, whereas 60% have at least partial execution and 19% expire at the end of the day with no executions.

The specialist's objective is to maximize expected profits, and he sets bid and ask prices based on the limit orders and his own willingness to trade. He must give price priority to the limit orders, and can stop an incoming market order. In order to focus solely on the competition with the limit orders, I assume the specialist has no obligations imposed on him by the exchange. Accordingly, the specialist will always set his ask and bid quotes equal to the limit buy and sell orders, and can always earn zero expected profit by never trading. I also assume that the specialist stops an order whenever he is indifferent between stopping and letting it trade with the limit book. This last assumption can be motivated on one of two grounds: the specialist may simply be more concerned with the interests of the market order traders, or he may rationally use this strategy to discourage aggressive limit orders.

As in Easley and O'Hara (1992b), I assume that the liquidity traders and informed traders use only market orders.[6] If their decisions were explicitly modeled, the exclusive use of market orders by both types of traders might be optimal if they were sufficiently risk averse. Moreover, informed traders would tend to favor the certain execution available with market orders to avoid the risk that their information might become public before their order executes.

The specialist and the limit order traders share the same initial beliefs concerning the two possible future values of the security (V_L and V_H), the probabilities of these values n_0 and $(1 - n_0)$, the probability that an information event had occurred e_0 (in which case the informed traders know whether the value is V_L or V_H), and the probabilities of arrivals of buyers and sellers. I assume that market spreads are tight enough so that it always makes sense for the informed traders to trade. Without price discreteness, this is always weakly true because limit order competition would drive the bid price to at least V_L and the ask price to at most V_H.

In order to simplify the model, I assume that the specialist can only stop orders arriving at time $t < T$, so at $t = T$ all orders must be cleared against his own account or against the limit orders. This provides a clean ending for the problem to make the recursive calculation of the optimal strategy unambiguous. Since period T contains this artificial restriction on the specialist's behavior, I focus on the predictions of the model for periods 1 through $T-1$. I also assume that the specialist trades only against arriving market orders, which simplifies the model but may be ignoring some interesting strategies that the specialist could follow. In particular, if the limit orders can't be withdrawn and several orders arrive in the same direction, this might convey enough information so that the specialist would like to trade against the limit orders himself.

I assume that at most one market order arrives in each time period. The first possible order arrival is at the end of period 0, and the last is at the end of period $T - 1$. Figure 1 shows the sequence of events, including limit order placement, market order arrivals, and specialist decisions. If no information event has occurred, then the probability of a (liquidity trader) buyer arriving is given by b_N and the probability of a seller arriving is given by s_N. Accordingly, if no information event has occurred the probability of no arrival at each time period is $z_N = (1 - b_N - s_N)$. If an information

[6] Chakravarty and Holden (1995) provide conditions under which it will be optimal for an informed trader to simultaneously submit a market order and a limit order in a single-period model.

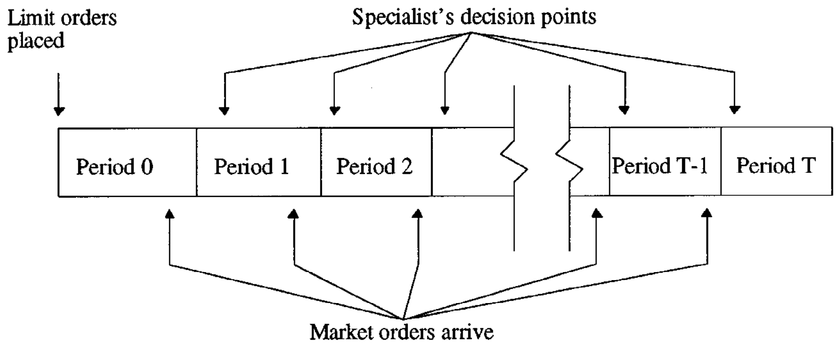

Figure 1
The sequence of events
Limit orders are placed at the start of time zero. Market orders arrive at the end of periods $0, 1, 2, \ldots, T-1$. Specialist decisions are made at the start of periods $1, 2, \ldots, T$.

event has occurred, then the corresponding probabilities are b_L, s_L, and z_L if $V = V_L$, and they are b_H, s_H, and z_H if $V = V_H$.

Although differences in buy and sell arrival rates are potentially interesting, preliminary analysis of solutions to the model showed that the most important parameters are the relative intensities of informed and liquidity trading. Accordingly, in all of the numerical solutions that follow I restrict the parameters by assuming that liquidity buyers and sellers are equally likely ($b_N = s_N$), that informed trader volume is simply added to that of the liquidity traders ($s_H = s_N$ and $b_L = b_N$), and that the intensity of informed buying after positive events equals the intensity of informed selling after negative events ($b_H = s_L$). Given these assumptions, the arrival rates can be fully characterized by specifying the probability of no trade following no event and the probability of no trade following an information event.

1.3 The Specialist's Problem

Let A denote the ask price and B the bid price as defined by the positions of the limit sell and buy orders. Given these order positions, the specialist chooses an optimal strategy for handling market buy and sell orders in each period. If a market order arrives in the first period, he can allow it to execute with the corresponding limit order (at the quoted bid or ask), he can take the other side of the order at a price \$1/8 better than the quote, or he can stop the order.

If no order arrives in the first period he does nothing. The decisions in any subsequent period will be predicated on whether an order arrives, as well as what has happened over the prior periods. If there are stopped orders at the start of the period, then an arriving order in the other direction automatically crosses with one of the stopped orders. The specialist then decides whether to execute any uncrossed order and/or any remaining stops, and if executed whether to trade for his own account or to let the limit book trade.

The specialist's dynamic programming problem can be described as follows [see Bertsekas (1976)]:

Random order arrival:
(at end of period t)

$$w_t, t = 0, 1, \ldots, T-1 \text{where } \begin{cases} w_t = -1 \text{ denotes a sell order} \\ w_t = 0 \text{ denotes no arrival} \\ w_t = 1 \text{ denotes a buy order} \end{cases}$$

$$\text{State vector: } x_t = \begin{pmatrix} x_t^1 = \text{stopped orders from period } t-1 \text{ plus (or net} \\ \qquad \text{of) the order arriving at the end of } t-1^7 \\ x_t^2 = \text{vector of order history through the end of} \\ \qquad t-1 \end{pmatrix}$$

$$\text{Control vector: } u_t(x_t) = \begin{pmatrix} u_t^1 = \text{number of shares bought in period } t \\ u_t^2 = \text{number of shares sold in period } t \\ u_t^3 = \text{number of shares stopped in period } t \\ \qquad \text{(negative for sell orders)} \\ u_t^4 = \text{number of shares crossed with the} \\ \qquad \text{limit orders in period } t \end{pmatrix}$$

Admissible control laws: $\pi = \{u_1(x_1), u_2(x_2), \ldots u_T(x_T)\}$ where each $u_t(x_t)$ is constrained as follows:

$$\begin{aligned} &\text{if } x_t^1 \geq 0 \text{ then } u_t^1 = 0 \text{ and } 0 \leq u_t^2 \leq x_t^1. \\ &\text{if } x_t^1 < 0 \text{ then } u_t^2 = 0 \text{ and } 0 \leq u_t^1 \leq -x_t^1. \\ &u_t^2 - u_t^1 + u_t^3 + u_t^4 = x_t^1. \\ &\text{if } t = T \text{ then } u_t^3 = 0. \end{aligned}$$

$$\text{System equations: } x_0 = \begin{pmatrix} 0 \\ \varnothing \end{pmatrix}, \quad x_{t+1} = \begin{pmatrix} u_t^3 + w_t \\ (x_t^2, w_t)^8 \end{pmatrix}$$

---

[7] A positive (negative) value denotes stopped buy (sell) orders.

[8] (x_t^2, w_t) denotes the row vector $(w_0, w_1, \ldots, w_t)$. For all t, there are 3^t possible histories, so a scalar value can be used to represent the history by letting $x_0^2 = 1$ and $x_{t+1}^2 = 3x_t^2 + w_t - 1$.

Objective[9]: choose π to maximize

$$J_\pi = E_0\left[\sum_{t=1}^{T} u_t^2(x_t)\left\{(A-.125)-E_t[V \mid x_t^2]\right\} + \sum_{t=1}^{T} u_t^1(x_t)\left\{E_t[V \mid x_t^2]-(B+.125)\right\}\right]. \tag{1}$$

To complete the definition of the problem, it is necessary to calculate the probability measure $P_t(\cdot|x_t)$ for the random order arrivals and the conditional expectation of firm value used in the above objective function. As shown in the appendix, these can be calculated using Bayes' rule.

1.4 The Role of Stops in the Specialist's Optimal Policy

The optimal policy can be solved recursively for a particular position (ask A and bid B) of the nearest limit orders. By Bellman's principle of optimality, the optimal policy defines an optimal value function $J_t(x_t)$ with the following property:

$$\begin{aligned} J_t(x_t) = \max_{u_t}\Big[&u_t^2(x_t)\{(A-.125)-E_t[V \mid x_t^2]\} \\ &+ u_t^1(x_t)\{E_t[V \mid x_t^2]-(B+.125)\} \\ &+ E_t[J_{t+1}(x_{t+1}) \mid x_t, u_t]\Big]. \end{aligned} \tag{2}$$

Equation (2) is used to recursively determine optimal policy for a given combination of parameters.

One can also use Equation (2) to illustrate the conditions under which stops will be a part of the specialist's optimal policy. Consider the situation depicted in Figure 2, that is, an ask price of \$30.125, a bid price of \$29.875, no orders stopped from period $t-1$, and a sell order arrived at the end of period $t-1$. This means that $x_t^1 = -1$ and the specialist's choices for u_t are

$$u_t = \begin{pmatrix}1\\0\\0\\0\end{pmatrix} \quad \text{or } u_t = \begin{pmatrix}0\\0\\-1\\0\end{pmatrix} \quad \text{or } u_t = \begin{pmatrix}0\\0\\0\\-1\end{pmatrix}$$

(specialist buy) (stop) (cross with limit orders)

[9] As shown in Appendix A, Section A.1, this value function is equivalent to the expectation of the total value of the specialist's position as of period T (net of purchase costs plus sales revenues).

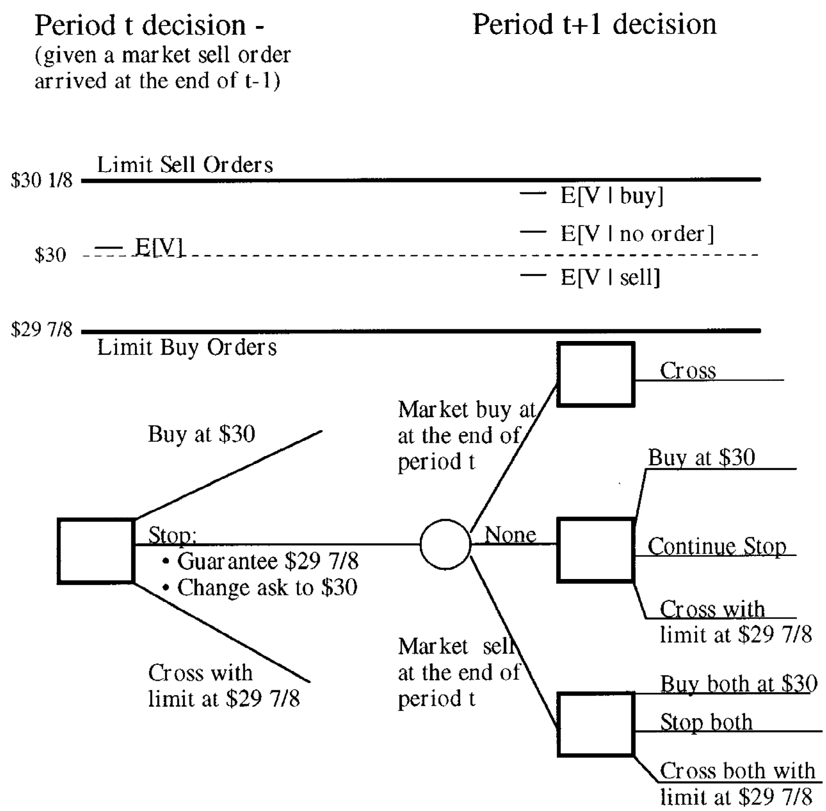

Figure 2
The specialist's decision tree
The diagram illustrates the specialist's potential choices in periods t and $t+1$ assuming there are no stopped orders from period $t-1$ and that a market sell order arrived at the end of period $t-1$.

The corresponding right-hand side values for Equation (2) are

$$\begin{aligned} &\left\{E_t\left[V \mid x_t^2\right] - \$30.00\right\} \\ &\quad + E_t\left[J_{t+1}\begin{pmatrix} w_t \\ (x_t^2 + w_t)\end{pmatrix} \mid x_t\right] && \text{(specialist buy)} \\ &E_t\left[J_{t+1}\begin{pmatrix} -1 + w_t \\ (x_t^2 + w_t)\end{pmatrix} \mid x_t\right] && \text{(stop)} \\ &E_t\left[J_{t+1}\begin{pmatrix} w_t \\ (x_t^2 + w_t)\end{pmatrix} \mid x_t\right] && \text{(cross with limit orders)} \end{aligned}$$

Since they both result in the same value for J_{t+1}, the choice between the specialist buying for his own account or crossing the order with a limit buy simply depends on whether the current expected value of the security is above the midpoint. In Figure 2, the time t expected value is drawn to be above the midpoint, so in this case the specialist will prefer buying for his own account to letting the incoming market sell order cross with a limit buy.

It is easiest to see the trade-off between stopping and immediate execution in period t if one temporarily assumes that stopping in periods $t+1$ and forward is not part of the optimal policy (which is always the case for $t+1=T$, because stops are not allowed in the last period). With that assumption, the difference in value between a stop and an immediate buy in period t can be written as (see appendix)

$$\begin{aligned}
&\text{value(stop)} - \text{value(buy)}\\
&\quad\text{(no stops after period } t)\\
&\quad= -\Pr\left\{w_t = 1 \mid x_t^2\right\} \max\left[\$30.00 - E_{t+1}\left[V \mid x_t^2, w_t = 1\right], 0\right]\\
&\qquad + \Pr\left\{w_t = 0 \mid x_t^2\right\} \max\left[E_{t+1}\left[V \mid x_t^2, w_t = 0\right] - \$30.00, 0\right]\\
&\qquad + \Pr\left\{w_t = -1 \mid x_t^2\right\} \max\left[E_{t+1}\left[V \mid x_t^2, w_t = -1\right] - \$30.00, 0\right]\\
&\qquad - \left\{E_t\left[V \mid x_t^2\right] - \$30.00\right\}. \qquad (3)
\end{aligned}$$

In order to improve readability, the period $t+1$ expected values have been written as conditioned jointly on the order flow through the start of period t and the order arriving at the end of period t. It would be equivalent to condition on the order flow through the start of period $t+1$.

The first conclusion that can be drawn from Equation (3) is that stopping the incoming sell order will be suboptimal if the next period's expected value of the security $E_{t+1}[V \mid x_t^2, w_t]$ will be above the midpoint after all three possible values for w_t. In this case, the first term on the right-hand side of Equation (3) equals zero and the second and third terms are positive. As shown in the appendix, Equation (3) then simplifies to

$$-\Pr\left\{w_t = 1 \mid x_t^2\right\}\left\{E_{t+1}\left[V \mid x_t^2, w_t = 1\right] - \$30.00\right\},$$

which is negative. A stop in this situation merely runs the risk that a profitable trade will be lost if a buy order arrives at the end of period t and is crossed with the stopped sell.

A stop has potential value if the sampling of the future order flow allows the specialist to avoid undesirable trades. In Figure 2, the expected value of the security is below the midpoint in period $t+1$ if a sell order arrives at the end of period t, so a stop could be valuable in period t. Continuing with the assumption that there will be no stops from $t+1$ forward [Equation (3)], and assuming $E_{t+1}[V \mid x_t^2, 1] > 30$, $E_{t+1}[V \mid x_t^2, 0] > 30$, and $E_{t+1}[V \mid x_t^2, -1] < 30$, the appendix shows that stopping a sell is optimal at the start of period t if and only if

$$E_{t+1}\left[V \mid x_t^2, w_t = 0\right] - \$30.00 \geq \frac{E_t[V \mid x_t^2] - \$30.00}{\Pr\{w_t = 0 \mid x_t^2\}}. \qquad (4)$$

Equation (4) illustrates two more conditions that are required for a stop to be preferred over immediate execution. First, it implies that $E_{t+1}[V \mid x_t^2, w_t = 0] > E_t[V \mid x_t^2]$, which means that no order arrival is a positive signal for value. The expected value drifts back toward the ex ante value following no order arrival, so Equation (4) implies that stopping a sell can be optimal in period t only if the ex ante value is above the spread midpoint (in fact, it must be above $E_t[V \mid x_t^2]$). The second implication follows from the fact that $\Pr\{w_t = 0 \mid x_t^2\} < 1$. This means that Equation (4) will tend to hold only when the current value is near the midpoint, so the potential profit from an immediate trade, given by $E_t[V \mid x_t^2] - \$30.00$, is fairly small.

Figure 2 depicts a case where $E_t[V \mid x_t^2] - \$30.00 > 0$, so that immediate execution is preferred to crossing with the limit orders. Stops can also be optimal when $E_t[V \mid x_t^2] - \$30.00 < 0$, provided the expected security value drifts above the midpoint following no order, that is, $E_{t+1}[V \mid x_t^2, w_t = 0] - \$30.00 > 0$. This will tend to occur only when the period t expected value is not too far below the midpoint.

Before turning to the determination of the equilibrium position of the limit orders, it is important to address the impact of the assumption that the specialist stops a market order whenever he is indifferent between stopping and letting the order trade with the limit orders. In so doing, I first continue with the assumptions that the spread is $1/4, a sell arrived at the end of period $t-1$, and no stops will be used after period t. In this case, indifference between a stop in period t and another action is relatively rare. I then address the case of a $1/8 spread, and explain why indifference between stops and other actions is much more common.

It turns out that it doesn't really matter how one handles the specialist's indifference between stopping and trading for his own account, because this almost never happens. In the example depicted in Figure 2, the specialist will be indifferent only if Equation (4) holds with strict equality, which would be very unusual.[10] If, however, $E_t[V \mid x_t^2]$, $E_{t+1}[V \mid x_t^2, w_t = 0]$, and $E_{t+1}[V \mid x_t^2, w_t = -1]$ are all less than the spread midpoint and $E_{t+1}[V \mid x_t^2, w_t = 1]$ is greater than the spread midpoint, then the specialist is indifferent between stopping the market sell order and letting it trade with the limit orders. The reason for this indifference is that the specialist can tell in period t that he will never be interested in trading in period $t+1$. Accordingly, he can stop the order in period t and then let any remaining net order cross with the limit orders in period $t+1$. In the numerical simulations described below, 8% of the stops in $1/4 spread markets occur when the specialist is indifferent between stopping and letting the order cross with a limit order.

[10] To be more concise, if we consider repeated solutions of the model where in each case the ex ante expected value is drawn from any continuous distribution at time zero, then strict equality in Equation 4 has probability zero.

When the spread is $\$1/8$, it is much "more common" for the specialist to be indifferent between stopping an order and letting it trade with the limit orders. This is because the specialist must pay the ask A to purchase for his own account and will only receive the bid B when he sells. Continuing with the assumptions that a sell order is available from $t-1$ and that no stops will be used beyond period t, the specialist will be willing (indifferent) to stop the sell provided $E_t[V \mid x_t^2]$, $E_{t+1}[V \mid x_t^2, w_t = 0]$, and $E_{t+1}[V \mid x_t^2, w_t = -1]$ are all less than the ask A and $E_{t+1}[V \mid x_t^2, w_t = 1]$ is greater than the bid B. In the numerical simulations below, about 70% of the orders are stopped in $\$1/8$ spread markets, and over 99.9% of these stops occur when the specialist is indifferent. Accordingly, stops due to the specialist's strict preference would almost never occur in $\$1/8$ spread markets. Empirically, about 21% of the orders were stopped in $\$1/8$ spread markets, which is less than predicted by the model. Accordingly, it may be the case that specialists *sometimes* stop when they are indifferent. I return to this issue below after defining the equilibrium position of the limit orders and describing the numerical simulations.

1.5 Equilibrium Position of the Limit Orders

I consider two related definitions for the equilibrium position of the limit orders. My first definition of an equilibrium assumes that all of the limit order traders are motivated only by the potential profit from supplying liquidity. Thus, a *profit-motivated* equilibrium is a pair of ask A and bid B such that the following conditions are satisfied:

Condition 1. *The T limit orders on each side of the market make nonnegative average expected profit given the specialist acts optimally.*

Condition 2. *For any other ask A' and bid B' with $A' - B' < A - B$, Condition 1 does not hold.*

Some limit order traders may be committed to trade and may be using limit orders as a way to reduce their expected execution costs. If such traders are also present, then the expected profit condition needs to be augmented with a comparison of expected execution prices between a committed limit order trader and an uninformed liquidity trader who uses a market order. For this calculation, I impute a fill price for the limit orders that don't execute through period T by assuming any unexecuted limit orders trade at the ex post expected value, rounded up to the nearest eighth for buy orders and down to the nearest eighth for sells.[11] Realized execution prices are calculated for each possible complete order sequence x_T^2, and then the

[11] Consistent with the earlier assumption that limit orders don't know where they sit in the queue, execution prices are averaged across the T limit orders on each side of the market that comprise the initial limit book.

expectation is calculated using the ex ante probabilities of each x_T^2. I define a *combined* equilibrium as a pair of ask A and bid B such the following conditions are satisfied:[12]

Condition 3. *For the T limit orders on each side of the market, either the average expected execution price for committed limit orders placed at that price is no worse than the average expected execution price for the uniformed market orders, or the average expected profit for limit orders placed at that price (with no forced execution) is nonnegative.*

Condition 4. *For any other ask A' and bid B' with $A' - B' < A - B$, Condition 3 does not hold.*

For any parameter combination, the combined equilibrium spread will be at least as narrow as the profit-motivated equilibrium spread. This is because the profit-motivated traders are present in both cases. Although the committed limit order traders may seem to be more "aggressive" than the profit-motivated limit order traders, for some parameter combinations, combined equilibria are supported by profit-motivated traders on one or both sides of the market. In these cases, the imputed fill price for limit orders that don't execute through period T is bad enough and the probability of nonexecution is high enough so that the committed traders would be better off using market orders.

In an actual market, one might expect that the resulting spreads would tend to be narrower than the profit-motivated equilibrium, because it would be reasonable for some liquidity traders to consider limit orders as a way to improve their expected execution prices. On the other hand, actual spreads might tend to be wider than the combined equilibrium, because the above definition implicitly ignores the greater price volatility faced by the limit order traders as compared to the market order traders. These committed traders might be willing to accept a slightly worse than expected execution price if it came with much lower variance.

1.6 Numerical Simulations

I solve for the specialist's optimal policy for the following set of parameter combinations:

$$\begin{aligned}\Theta = \{\theta = (T, \gamma, e_0, n_0, z_N, z_H, \delta) &\text{ such that}\\ T &\in \{3, 4, 5\} \text{ (number of periods)}\\ \gamma = V_H - V_L &\in \{\$2, \$4, \$6, \$8, \$10\}\\ &\quad\text{(magnitude of the information event)}\\ e_0 &\in \{.02, .04, .06, .08, .10\} \text{ (probability of an information event)}\end{aligned}$$

[12] This equilibrium definition is in the same spirit as the "ecology" described by Handa and Schwartz (1996).

$n_0 \in \{.2, .3, .4, .5, .6, .7, .8\}$ (probability that the value is low)
$z_N \in \{.2, .4, .6, .8\}$ (probability of no trade following no event)
$z_H (= z_L) \in \{.03, .06, .09, .12, .15, .18\}$
(probability of no trade following an event)
$\delta = [n_0 V_L + (1 - n_0) V_H] - P_0 \in \{\pm\$.01, \pm\$.02 \ldots, \pm\$.06\}$,
for some $P_0 \in P$ (difference between the ex ante value and the nearest value on the price grid)}

The price parameter P_0 given above can be chosen arbitrarily, because the specialist's objective function [Equation (1)] will be invariant to translations of A, B, V_L, and V_H by any constant. Figure 2 is drawn so that $P_0 = \$30.00$, and for the sake of discussion I will assume $P_0 = \$30.00$ for the remainder of the article.

Hopefully I have chosen the ranges of the parameters to be broad enough to include what most readers would consider to be "reasonable" values. The number of periods is constrained by computational considerations, as well as by the belief that although limit orders may not be revised after every trade, they will begin to react after a few trades. The position of the ex ante value relative to the closest point on the price grid spans all possible values. The magnitude and probability of an information event are both bounded below by zero. The upper bounds are chosen based on qualitative observation of actual NYSE trading. Price moves of \$10 are extremely rare for NYSE stocks. In addition, it is hard to imagine that the probability of a substantial information event known only to a subset of traders is much greater than 10% at any point in time. Finally, the possible combinations of z_N and z_H (z_L) range from very informative order flow ($z_N = .8$ and $z_H = .03$), where most traders after an information event are informed, to very uninformative order flow ($z_N = .2$ and $z_H = .18$).

I solve for the specialist's optimal policy for each of the $3 \times 5 \times 5 \times 7 \times 4 \times 6 \times 12 = 151{,}200$ elements of Θ and up to three different pairs of ask A and bid B. I first assume that $A - B = \$1/8$ and that A and B bracket the ex ante value, and I solve for the specialist's optimal policy. For each equilibrium definition, if the \$1/8 spread is not an equilibrium, then I move on to consider the two possible pairs of A and B with $A - B = \$1/4$ that bracket the ex ante value. If both \$1/4 spreads are equilibria, then I choose the one where the ex ante value rounds to the midpoint. If neither is an equilibrium, then I simply classify the equilibrium for the parameter combination as "wider than 1/4." Let $\Gamma(E, S)$ denote $\{\theta\colon\ \theta \in \Theta$ and $S \in \{1/8, 1/4, > 1/4\}$ is an equilibrium of type $E \in$ {profit-motivated, combined}}. Table 1 classifies all of the elements of Θ according to whether a \$1/8 spread represents a profit-motivated and/or a combined equilibrium, and if not, whether the limit orders whould be placed \$1/4 apart or wider than \$1/4. As an example using the set notation, the value in the first column of the

Table 1
Profit-motivated and combined equilibria

Profit-motivated equilibrium spread	Combined equilibrium spread $1/8	$1/4	Wider than $1/4	Total
$1/8	26,452 (100/86.2)	0 (0.0/0.0)	0 (0.0/0.0)	26,452 (100/17.5)
$1/4	4,240 (15.6/13.8)	22,970 (84.4/89.5)	0 (0.0/0.0)	27,210 (100/18.0)
Wider than $1/4	6 (0.0/0.0)	2,684 (2.8/10.5)	94,848 (97.2/100.0)	97,538 (100/64.5)
Total	30,698 (20.3/100)	25,654 (17.0/100)	94,848 (62.7/100)	151,200 (100/100)

The table summarizes the results of equilibrium calculations for 151,200 combinations of the parameters. The rows of the table separate the parameter combinations based on equilibrium spread using the profit-motivated equilibrium definition. The columns of the table separate the parameter combinations based on the combined equilibrium definition. Numbers in parentheses give the number in the cell as a percent of the row and then as a percent of the column.

second row is the number of elements in $\{\theta\colon \theta \in \Gamma(\text{profit-motivated}, 1/4)$ and $\theta \in \Gamma(\text{combined}, 1/8)\}$. The table shows that the range of parameters used is broad enough to cover all of the possible situations.

Let $\Theta[\cdot]$ denote the restriction of Θ to a particular subset based on the condition given in the brackets. For example, $\Theta[T = 3]$ equals $\{\theta\colon \theta \in \Theta$ and $T = 3\}$. To check whether any of the parameter ranges are too large, I recalculate Table 1 for $\Theta[T = 3]$, $\Theta[T = 4]$, $\Theta[T = 5]$, $\Theta[\gamma = 2], \ldots, \Theta[z_H = .18]$. I do not consider the restrictions on the parameter δ, because I believe that in any reasonable model the difference between the ex ante value and the nearest point on the price grid will evolve randomly through time. I find that in each of these $3+5+5+7+4+6 = 30$ cases, all of the sets $\Gamma(E, S)[\cdot]$ are nonempty. These results (not shown) indicate that none of the parameter values are individually unreasonable and that the existence of the various equilibria are not unduly sensitive to particular parameters.

1.7 Do Stops and Immediate Price Improvement Help Liquidity Traders?

As mentioned above, the impact of stops and price improvement on the trading costs of the uniformed liquidity traders depends not only on the average price improvement received, but also on how these actions affect the equilibrium spread between the limit orders. Table 2 examines various market statistics in order to compare the NYSE market structure (with both stops and immediate price improvement) to a specialist market where stops are forbidden only in $1/8 spread markets, to a specialist market where stops are completely forbidden, and to a pure electronic limit book (no stops and no immediate improvement).

Table 2
Comparison of alternative market structures

Market structure	Quoted spread	Effective spread: Uninformed	Effective spread: Informed	Profit/market order: Limit orders	Profit/market order: Specialist
Panel A: $1/8 spread is a profit-motivated equilibrium (26,452 parameter combinations)					
NYSE	0.1250	0.0553	0.0665	0.0200	0.0001
No stops with $1/8 spread	0.1250	0.1231	0.1246	0.0526	0.0001
No stops at all	0.1250	0.1231	0.1246	0.0526	0.0001
Pure limit book	0.1250	0.1250	0.1250	0.0529	0.0000
Panel B: $1/4 spread is a profit-motivated equilibrium (27,210 parameter combinations)					
NYSE	0.2500	0.1478	0.1703	0.0342	0.0084
No stops with $1/8 spread	0.1909	0.1371	0.1533	0.0399	0.0030
No stops at all	0.1909	0.1596	0.1704	0.0491	0.0028
Pure limit book	0.1895	0.1895	0.1895	0.0645	0.0000
Panel C: $1/8 spread is a combined equilibrium (30,698 parameter combinations)					
NYSE	0.1250	0.0562	0.0676	0.0179	0.0001
No stops with $1/8 spread	0.1250	0.1222	0.1237	0.0497	0.0001
No stops at all	0.1250	0.1222	0.1237	0.0497	0.0001
Pure limit book	0.1250	0.1250	0.1250	0.0505	0.0000
Panel D: $1/4 spread is a combined equilibrium (25,654 parameter combinations)					
NYSE	0.2500	0.1469	0.1709	0.0277	0.0096
No stops with $1/8 spread	0.1276	0.1141	0.1209	0.0264	0.0012
No stops at all	0.1276	0.1150	0.1216	0.0267	0.0012
Pure limit book	0.1260	0.1260	0.1260	0.0306	0.0000

Equilibrium spreads are calculated for each of the 151,200 parameter combinations using both equilibrium definitions. Each panel is restricted to those combinations that result in the particular equilibrium spread size using the particular equilibrium definition, and assuming the NYSE structure with stops and price improvement. For each parameter combination in the panel, equilibrium spreads are recalculated assuming no stops are allowed when the spread is $1/8, then assuming that no stops are allowed at all, and finally assuming that no stops or immediate price improvement are allowed (a pure limit book). Table values are simple averages across all parameter combinations included in the panel. The effective spread columns reflect the expected average purchase price minus the expected average selling price. The profit/market order columns reflect the expected total profit divided by the expected number of market orders.

The four panels in the table reflect the four parameter sets: Γ(profit-motivated,$1/8$), Γ(profit-motivated,$1/4$), Γ(combined,$1/8$), and Γ(combined,$1/4$). Note that these parameters are chosen based on the equilibrium that would occur assuming full use of stops and price improvement, which is reflected in the first line of each panel. Without changing the parameter set, the remaining lines of each panel assume that restrictions are placed on the specialist's actions and that the spread between the limit orders adjusts accordingly. All values in the table are calculated for each parameter value (see the appendix for details) and then averaged across all of the parameter values for the panel.

Assume for the moment that situations where we observe a spread of S are equally likely to be the result of any one of the parameter combinations contained in $\Gamma(E, S)$ for one of the two equilibrium definitions E. With this assumption, comparisons of the rows in each panel of Table 2 gives an indication of the potential impact of a change in market structure on the trading costs of the uninformed traders. For example, if the limit order traders are purely profit-motivated (panels A and B), then the uninformed traders prefer the NYSE market structure to a pure limit book. Uninformed traders who would be transacting in stocks with a $1/4 spread under the NYSE market structure might prefer a rule change that forbids stops in $1/8 spread markets. This is because this rule change might encourage tighter limit orders, and they would still get the benefit of stops if the limit orders remain at $1/4. The cost of such a change is the lost price improvement from stocks where the spread is already $1/8 under the NYSE structure. As long as the situations resulting in a $1/8 spread are more common than those resulting in $1/4, the NYSE market structure will result in the lowest average effective spreads for uninformed traders. The same conclusion holds if the limit order spread is determined by a combined equilibrium (panels C and D of Table 2).

As mentioned above, the large amount of price improvement in $1/8 markets shown in Table 2 is a result of the assumption that the specialist uses stops when he is indifferent. If this assumption did not hold, then forbidding stops in $1/8 markets might be much less costly to the uninformed traders than would appear from panels A and C. The assumption, however, is also an important determinant of the impact of forbidding stops shown in panels B and D, because forbidding stops is less likely to cause the limit order spread to tighten. A recalculation of Table 2 under the assumption that stops are only used when they are strictly preferred (not shown) yields the same ultimate conclusion that the NYSE market structure would yield the lowest average effective spread for the uninformed traders.

I test whether the conclusions drawn from Table 2 are robust over different "regions" of the sets $\Gamma(E, S)$ by recalculating Table 2 for each of 30 subsets of the parameter space given by $\Theta[T = 3]$, $\Theta[T = 4]$, $\Theta[T = 5]$, $\Theta[\gamma = 2], \ldots, \Theta[z_H = .18]$. For each of these parameter restrictions, as long as one assumes that at least as many orders arrive in $1/8 spread markets as in $1/4 spread markets, the NYSE market structure results in lower average effective spreads for the uninformed traders than the other three market structures.

Table 2 shows that stops, and to a lesser degree immediate price improvement, have a differential impact on the trading costs of the uninformed and informed traders. The uninformed traders are much more likely to receive price improvement from stops than the informed traders, and the reason is simple. Suppose you are a buyer and your order has been stopped. The specialist is waiting to make sure your order is not followed by another buy before he offers price improvement. If you are uninformed, then the likelihood of a buy in the next period is no greater than the likelihood of a sell, and the likelihood of no trade is fairly high. On the other hand, if you are informed, then there is a much higher chance that your order will be followed by another buy.

Table 3
Theoretical limit order profit and potential profit

	$1/4 spread is a	
Order Handling	Profit-motivated equilibrium (27,210 parameter combinations)	Combined equilibrium (25,654 parameter combinations)
Immediately at bid or ask	$0.0606	$0.0508
Stopped, then at bid or ask	0.0607	0.0517
Stopped, then at midpoint with specialist	0.1398	0.1417
Stopped, then at midpoint due to cross	0.1415	0.1434
Immediately at midpoint with specialist	0.1496	0.1531

For each possible order sequence and each market order in the sequence, the profit for orders that trade against the limit orders is calculated by comparing the initial trade price (which equals the limit price) to the ex post expected value conditioned on the order sequence. Similarly, potential profit (the last three rows in the table) is calculated for market orders that aren't allowed to execute against the limit orders by comparing the limit price to the ex post expected value conditioned on the order sequence. Limit order profits (potential profits) are averaged across all the market orders, weighting by the ex ante probability of each order sequence. Orders that arrive when there is already a stopped order and those that arrive in the last period are excluded.

Values in the table are simple averages across all parameter combinations in each column.

1.8 Empirical Predictions of the Model

I focus on the predictions of the model for cases where the equilibrium spread is $1/4$ and I restrict my attention to periods where there was no open stop. Exchange rule 116 currently states that specialists are required to update their quotes when an order is stopped and the spread at the time of the stop was greater than the minimum tick size. For example, if the sell order depicted in Figure 2 was stopped, the rule would require that the ask quote be reduced to $30.00. This rule was not binding during the period covered by my sample, but the practice was followed approximately 10% of the time.

In order to examine the adverse selection cost that the specialist's behavior imposes on the limit orders, I calculate the limit order trader's ex post profit from each market order that is allowed to trade with the limit orders. I also examine the market orders that execute against the specialist or another market order, and I calculate the potential profit that the limit orders would have made had they been allowed to trade with these orders. Table 3 shows the pattern in these profits as a function of the five possible ways the order could have been handled. The first column shows the average across Γ(profit-motivated,$1/4$) and the second across Γ(combined,$1/4$). In averaging across parameter combinations, I exclude cases where the particular order handling never occurs. This has a relatively minor effect on the averages, because most parameter combinations exhibit all five possible cases.

Table 3 shows that orders that are allowed to execute against the limit orders tend to be the least profitable, including both those that are immediately allowed to execute and those that are first stopped. Orders that are immediately improved by the specialist are those that would have given the limit orders the highest profit had they been allowed to trade. I also recalculate Table 3 for the 30 restricted parameter sets $\Theta[\cdot]$, and the results are quite consistent across all of these restricted sets.

The discussion in Section 1.5 indicates that in \$1/4 spread markets, most stops should occur when the expected value of the security is close to the midpoint. This feature is not directly testable, because the expected value is unobservable. Accordingly, I focus on the related and testable prediction that nearly all stops should occur when the last trade was at the midpoint of the spread. The easiest way to understand this result is to consider the cases when the last trade was *not* at the midpoint. First, when a buy (sell) order arrives and the last trade price was at the ask (bid), the model predicts that the specialist will nearly always let the order execute immediately against the limit orders. Note that he could have traded the previous buy (sell) order at the midpoint of the spread for his own account, but evidently did not. Another order in the same direction should make him more worried that there has been an information event, and therefore even less anxious to trade for his own account. The second case is when a buy (sell) order arrives and the last trade price was at the bid (ask). Here the model predicts that the specialist will be much more likely to offer immediate price improvement than to stop the order. This is because a stop runs the risk that the profitable trading opportunity will be lost if another market order arrives and crosses with the stopped order.

By comparing the stop frequencies from the model to those observed in the data, I am implicitly assuming that the data represent the model being repeated many times in sequence. Unfortunately the model does not address how each set of trading periods joins with the next. In particular, there is no definition of the last trade price before the limit order set in period 0. I handle this problem by counting the stop frequencies in two ways. First, when using all of the orders, I assume the last trade price before time 0 was equal to the rounded ex ante value of the security. Second, I also report the results for only those orders that arrive after a trade has occurred in an earlier period.

Table 4 shows the average frequencies of immediate executions, stops, and immediate improvements, for each of the three possible last trade prices. I define the "receiving" side of the quote as the ask (bid) when the incoming market order is a buy (sell). Panel A of Table 4 reflects the average across Γ(profit-motivated,1/4) and panel B reflects the average across Γ(combined,1/4). The table shows that when the specialist follows the optimal policy given by the model, the frequency of stops is quite high when the last trade price was at the midpoint of the spread, but otherwise is very low. Furthermore, the frequency of stops following midpoint trades is nearly the same when the first order to arrive is excluded. Recalculation of Table 4 using

Table 4
Theoretical frequencies of stops and immediate improvement

	All trades			Excluding first trades		
Last trade price	Immed. at quote	Stopped	Immed. improved	Immed. at quote	Stopped	Immed. improved
$1/4 spread is a profit-motivated equilibrium (27,210 parameter combinations)						
Other side	0.02%	1.46%	98.52%	0.02%	1.42%	98.56%
Midpoint	49.71	31.87	18.43	48.36	40.38	11.26
Receiving side	99.96	0.04	0.00	99.96	0.04	0.00
$1/4 spread is a combined equilibrium (25,654 parameter combinations)						
Other side	0.02%	2.22%	97.76%	0.02%	1.54%	98.44%
Midpoint	49.70	31.36	18.95	48.60	40.09	11.30
Receiving side	99.95	0.05	0.00	99.95	0.05	0.00

The model is used to calculate the specialist's initial order handling decisions assuming the limit orders are placed $1/4 apart. For each individual parameter combination and each category of last trade price, the frequency calculations reflect the ex ante probabilities of each possible sequence of orders. The results in the table are the equally weighted averages of these frequency calculations across all of the parameter combinations in each category. The "receiving" side is the ask (bid) when the incoming market order is a buy (sell).

the restricted parameter sets $\Theta[\cdot]$ yields very consistent results. In fact, for every individual parameter combination where stops are part of the optimal strategy, the stop frequency is highest when the last trade was at the midpoint.

2. Data

2.1 The TORQ Database

The data used in the study are from the TORQ database, available from the NYSE. The database covers a 3-month period from November 1990 through January 1991 for a sample of 144 firms. These firms were selected by partitioning all NYSE firms into size deciles and then randomly selecting 15 firms from each decile. Some data problems encountered after the initial selection was made caused six firms to be dropped from the final sample. In addition to intraday consolidated trade and quote information, the database contains information on the parties on each side of a trade and a record of all orders submitted to the floor by way of the NYSE's SuperDOT system.

Importantly, the SuperDOT order information includes the time that each order was submitted, the time when it was executed, and whether it was stopped. The primary drawback of this sample is that not all market orders arrive at the post via the SuperDOT system. Larger orders, and those that require "special handling" are more likely to avoid the SuperDOT system entirely and be hand carried to the specialist's post by a floor broker. Since this routing decision presumably depends on a broker's assessment of current market conditions, there are clearly selection biases at work.

I use a sample of SuperDOT market orders from TORQ to evaluate the determinants of the specialist's decision to stop market orders. In order to determine the corresponding quotes, it is important to have a good idea when the specialist first became aware of the order. Most market orders entered in the SuperDOT system are immediately routed to the specialist's post and appear on a monitor visible to both the specialist and his trading assistant. In contrast, some orders are first routed to a broker's booth located along the edge of the trading floor. Sometimes these orders are forwarded electronically to the specialist's post and sometimes they are hand carried to the post by a floor broker to be "worked." Although most of these orders will ultimately arrive at the post within 3 minutes of when they entered SuperDOT, there is no way to accurately determine the quote in effect when the order was first presented to the specialist. Fortunately, the TORQ database includes a field indicating whether the order was first routed to a broker's booth, and I exclude these market orders (0.3% of all SuperDOT market orders) from my sample.

2.2 Building the Limit Book

The TORQ database includes all limit orders submitted to the floor over SuperDOT, as well as any hand carried to the post that were subsequently electronically entered in the limit order book. The database also tracks executions and cancellations of these orders, so it is possible to build a picture of the limit book at each point in time [see also Kavajecz (1997)]. The biggest problem with this picture is that some limit orders may be left at the post and never "officially" entered in the limit book.[13] With the available data, there is no good way to automatically handle this problem; it must always be kept in mind when interpreting test results that are based on the apparent lack of a limit order at the current quote.

Another problem in building the limit book is that some orders are canceled without the proper record appearing in the database. A single such order with a "stale" price can result in a long sequence of market orders for which the position of the nearest limit order information is incorrect. I try to minimize this by automatically deleting any limit order when a trade execution occurs on the NYSE that is 1 minute after the limit order is entered on the book and the trade price is such that it would violate the price priority of the limit order (for example a trade at $1/8 below a limit buy order). This results in the automatic deletion of approximately 5% of the limit orders entered in the book. In part these instances represent reporting lags; about 20% of these orders had an execution or cancellation record within 60 seconds of the automatic deletion.[14]

---

[13] Sometimes specialists will be given "participation" orders. These orders request to be included as part of many trades throughout the day. They are often entered by traders whose execution is evaluated against the average trade price for the day. These are not enterable in the book, but are often at least partially reflected in the quote size.

[14] For "all-or-none" limit orders, theoretically there can be trades past the limit price if the quantity is small. However, only 2% of the automatic deletions were all-or-none orders.

2.3 Sample Summary Statistics

Table 5 summarizes the sample definitions and presents some statistics on the frequencies of stops and price improvement for various categories of the TORQ market orders. For a few orders in the sample (1.3%), a portion of the order was executed immediately, and the balance was stopped. These are included in the nonstopped category. Most importantly, the data show how common stops are in this sample. The specialists stopped about 20% of market orders submitted at $1/8 spread and about 40% at $1/4 and $3/8 spreads.[15] Although all of the analysis in this article combines buy and sell orders, in all tests separate analyses of buy and sell orders yielded qualitatively similar results.

The first line of Table 5 shows that 332,569 SuperDOT market orders were sent directly to the specialist's post after the open while there was a good quote posted by the specialist.[16] Section 3 empirically tests the model's predictions as presented in Tables 3 and 4. For these tests, the sample is restricted to those orders where the quoted spread is $1/4, both sides of the quote reflect the nearest known limit orders, there are no quotes from competing exchanges inside the NYSE quote, and none of the previous stopped SuperDOT market orders remains open. This sample, called the "model sample" in Table 5, consists of 19,053 market orders.

In order to try to ensure that timing and reporting problems are not driving the results, I repeat each of the analyses in Tables 6 and 7 of Section 3 using a subsample of market orders for which there was no change in the current quote during the 30-second period before the order arrived and there was no other SuperDOT traffic in the 60-second interval centered on the order's arrival. I also restrict the sample by deleting any order larger than the 95th percentile for each firm. This last sample, referred to as the "stable situations" sample, consists of 7,940 market orders. Although the various screens do eliminate some of the 144 TORQ firms, a total of 129 firms are represented in both the model sample and the stable situations subsample.

It is important to note that by restricting the sample to match the model, I am only focusing on cases where the specialist is using stops and price improvement in direct competition with the limit orders. There are other explanations for these actions that are tied to the structure and rules of the NYSE. As evidence of these other explanations, note that of the 24,971

[15] Ross, Shapiro, and Smith (1996) examine a sample of SuperDOT orders from January 1996, and they find that 12% of orders at $1/8 spread were stopped and 21% of orders at greater than $1/8 spread were stopped. Apparently, while the practice is still common, the frequency has declined. They also examine the relation between stops and price improvement and find results similar to those shown in Table 5. Evidence of the importance of stops in providing price improvement is also given Angel (1994) and Harris and Hasbrouck (1996).

[16] This sample excludes orders if there was missing trade volume or execution price data, or if the average execution price differed from the quote by $1 or more. In total, all of the screens applied (primarily the requirement that the order arrived between the open and close) eliminate 12% of the SuperDOT market orders.

Table 5
TORQ sample and subsamples

Sample criteria	Number	Percent stopped	Percent receiving price improvement: Stopped	Percent receiving price improvement: Not stopped
SuperDOT market orders, routed directly to post during continuous trading, with a good quote, and that pass the error screens	332,569	28.8%	55.1%	22.2%
Less orders at other than $1/4 NYSE quoted spread				
$1/8 spread	196,411	20.8	41.3	5.1
$3/8 spread	17,788	41.3	68.4	63.9
All other	2,004	37.4	69.6	49.0
Orders with $1/4 NYSE quoted spread	116,366	40.2	64.7	53.6
Less orders that don't meet the model's conditions				
"Hidden" limit order inside the NYSE quote	38,777	38.0	65.5	64.0
No limit order at the quote on the receiving side	24,971	36.7	56.0	37.5
No limit order at the quote on the other side	14,642	39.8	63.7	44.3
Receiving side ITS quote inside NYSE quote	5,166	48.1	82.2	84.7
Other side ITS quote inside NYSE quote	5,586	47.5	63.7	47.7
Prior stop still open	8,171	63.2	62.1	80.6
"Model" sample	19,053	35.1	71.8	48.9
Less orders with a change in the current quote up to 30 seconds before arrival, or other SuperDOT traffic within 30 seconds either before or after arrival, or order size greater than the 95th percentile for the firm	11,113	34.0	73.7	53.5
"Stable situations" subsample	7,940	36.6	69.4	42.1

The NYSE quoted spread is measured as of the time the order entered the SuperDOT system. Orders that were partially executed immediately, with the balance of the order stopped (1.3% of the sample) are included in "Not stopped." The "receiving" side is the ask (bid) when the market order is a buy (sell).

orders that arrived when there was no apparent limit order on the "receiving" side of the market, 56% were stopped. In these cases, the specialist was providing the inherent price guarantee against his own account, which would be inconsistent with the profit-maximizing behavior assumed in the model. Also, the table shows that of the 5,166 orders that arrived when there was a better Intermarket Trading System (ITS) quote on the receiving side, 82% were stopped. NYSE specialists often try to find an NYSE contra party to meet the better ITS quote, but even when they send the order directly across the ITS, they routinely stop the order at the NYSE quote. In these cases the stop provides an electronic record that some action has been taken, and the order is guaranteed at least the NYSE quote if the ITS quote is withdrawn. As discussed in Section 3, even the "model" sample shows evidence of other motives for stops and immediate price improvement.

In Section 4, I show that the results of Section 3 can be used to improve the estimation of expected execution price that was developed by Knez and Ready (1996). Market participants do not have limit book information, so

Table 6
Observed limit order profit and potential profit

Market order handling	Limit order profit (or potential profit) 5 min.	Limit order profit (or potential profit) One day	Relative size	Number of market orders
Model sample (19,053 market orders)				
Executed immediately at quote	0.080	0.044	1.73	6,319
Stopped, then executed at quote	0.078	0.001	1.45	1,885
Stopped and improved (crossed with)				
• Other than SuperDOT Market Order	0.133	0.114	1.47	3,264
• SuperDOT Market Order	0.119	0.088	0.85	1,542
Immediately Improved	0.139	0.150	1.28	6,043
Stable situations sample (7,940 market orders)				
Executed immediately at quote	0.098	0.066	1.11	2,914
Stopped, then executed at quote	0.095	0.015	1.04	889
Stopped and improved (crossed with)				
• Other than SuperDOT Market Order	0.133	0.126	1.00	1,419
• SuperDOT Market Order	0.121	0.099	0.81	600
Immediately Improved	0.134	0.150	0.89	2,118

The first two columns compare the quote at the time of the order arrival to the midpoints of quotes 5 minutes and 1 day later. These columns show the potential profit to the limit orders if the order had been allowed to execute at the quote, so for incoming buy orders they are (current ask − subsequent midpoint) and for incoming market sells they are (subsequent midpoint − current bid). For each firm, average order size is computed across all of the orders in the stable situations sample. This average is used to normalize each of the orders for that firm in both the model and the stable situations samples. The Relative size column reflects the averages of these normalized values across all of the orders in each category.

Statistical significance is judged at the 5% level in a two-tailed test using Monte Carlo simulations that control for potential differences across firms. For the model sample, all of the comparisons of values within each column are statistically significant *except* for the following: in the 5-minute profit column, the difference between the values 0.080 and 0.078; in the 1-day column, the difference between the values 0.044 and 0.088; in the 1-day column, the difference between the values 0.114 and 0.088; in the 1-day column, the difference between the values 0.114 and 0.150; and in the relative size column, the difference between the values 1.45 and 1.47.

this estimation uses only trade and quote data. Accordingly, the analysis in Section 4 is based on all of the 116,366 orders that arrived when the NYSE quoted spread was $1/4.

3. Empirical Tests of the Model

3.1 Adverse Selection Costs

As shown in Table 3, the model predicts that stops and immediate price improvement will impose an adverse selection cost on the limit orders. To test this prediction, Table 6 separates the orders in the model and stable situation samples into categories based on how they were ultimately handled. For each category, the table reports the average profit or potential profit, calculated by comparing the limit price at the time each market order arrived to the midpoints of subsequent quotes, measured 5 minutes and one trading day later. The table also reports the average normalized size for the orders in each category, where the normalization uses the average order size for

Table 7
Observed frequencies of stops and price improvement

Last NYSE trade price relative to the NYSE quoted spread	Number of orders	No stop or price improvement	Stopped	Immediate price improvement
Model sample (19,053 market orders)				
Other side	1,875	20.0%	29.2%	50.8%
Midpoint	12,971	21.9	42.4	35.7
Receiving side	2,910	89.5	7.2	3.3
Other	1,297	38.9	33.4	27.8
Stable situations sample (7,940 market orders)				
Other side	998	20.2%	32.3%	47.5%
Midpoint	4,874	24.6	45.9	29.6
Receiving side	1,283	92.1	6.7	1.2
Other	785	42.5	33.6	23.8

The NYSE quoted spread is measured as of the time the order entered the SuperDOT system. Orders that were partially executed immediately, with the balance of the order stopped, are counted as not stopped. The "receiving" side is the ask (bid) when the market order is a buy (sell). The "other" category primarily reflects cases where there was no earlier trade price for the day.

Statistical significance is evaluated using Monte Carlo simulations that control for cross-firm differences. For the model sample, all differences in percentages between rows are significant at the 5% level in two-tailed tests with two exceptions. The 33.4% stop percentage for the other category is not significantly different from the 29.2% for the other side category, and the 27.8% immediate improvement percentage for the other category is not significantly different from the 35.7% for the midpoint category.

all of the firm's orders included in the stable situations sample. Statistical significance for the differences between the values in the rows of the table is calculated using Monte Carlo simulations that control for cross-firm differences (see appendix). For these tests, the null hypothesis is that the statistics may differ across firms, but for each firm the expected value is constant across the categories.

The results in Table 6 are roughly consistent with the predictions in Table 3. As expected, the specialist immediately allows a market order to trade against a limit order only in the cases that are least profitable for the limit orders. There are, however, two apparent inconsistencies between Tables 3 and 6. The model predicts that market orders that are stopped but then later allowed to trade at the quote should be slightly better than those that immediately execute against the limit orders, because in the former case the specialist at least considered trading for his own account at a price $\$1/8$ better than the limit order. Table 6 shows that these orders are the least profitable, particularly using the one-day measure.

The model also predicts that when stops are used, the potential profit should be higher when future price improvement results from crossing with another market order than it is when it results from a specialist trade. This result follows from the fact that an order arrival that is in the other direction from the stopped is generally a stronger signal for value than the lack of an order. I am unable to separately identify the specialist's trades, but I can tell from the audit data whether an order crossed with another SuperDOT market order. Accordingly, I separate stopped orders that were subsequently improved into two categories: those that crossed with SuperDOT market orders,[17] and all others. Specialist trades are in the "all other" category, but so are trades crossed with a market order that is hand-carried by a floor broker. Table 6 shows that orders in the "all other" category seem to have slightly higher potential profit than those that crossed with SuperDOT market orders. Possibly this indicates that market orders that bypass the SuperDOT system and are hand-carried to the post are relatively more informative.

Although the model does not address trade size, the "relative size" column in Table 6 indicates a clear association between order size and the specialist's handling of the order. I return to this issue in Section 3.3 below.

3.2 Conditioning on Last Trade Price

Table 7 shows the frequencies of various order handling decisions as a function of the position of the last NYSE trade price. Stops are most common when the last trade was at the midpoint of the spread, which is consistent with the model's prediction as shown in Table 4. When the last trade was at the midpoint of the spread, the observed price improvement frequency is somewhat larger than predicted. In some cases, this price improvement may represent crossing with stopped market orders that could not be detected in the database because they were not routed through the SuperDOT system. In other cases, the price improvement may be provided by a floor broker "working" an order at the post.

The model predicts that when the last trade was on the side of the spread that would receive the incoming order, the order will nearly always execute immediately at the quote. When the last trade was on the other side, the model predicts that the specialist will nearly always provide immediate price improvement. These extreme predictions are a result of the assumption that the order flow is the specialist's only source of information. The results in Table 7 for the receiving and opposite side categories are not as clear-cut as the model's predictions, which may indicate that the specialist is responding to other sources of information besides the order flow.

Hasbrouck, Sofianos, and Sosebee (1993) assert that specialists use stops in \$1/4 spread markets to avoid price changes of more than \$1/8. Table 7 shows that about 30% of the orders are stopped when the last trade was on the other side of the spread. In these cases, immediate execution at the quote would result in a \$1/4 price change, so the high stop frequency is consistent with the hypothesis that specialists sometimes use stops to manage their price continuity statistics. Note, however, that the immediate

[17] I use the audit file in TORQ to determine the contra party to the market order. Orders are said to have crossed with a SuperDOT market order (or orders), if such orders constitute at least 50% of the volume on the contra side.

price improvement predicted by the model would also avoid a $1/4 price change.

3.3 Order Handling as a Function of Trade Size

Past theoretical and empirical work suggests that larger orders contain more information.[18] If the model were extended to incorporate different trade sizes, it is likely that stops would still be used as a way to sample future order flow, but the specialist's order handling decision would clearly depend on the size of the incoming order. In this section, I empirically investigate this issue by using nonparametric regressions to estimate conditional order handling frequencies.

Rather than using order size per se, I use "excess depth" as defined by Knez and Ready (1996) as the conditioning variable. Excess depth is the difference between the NYSE quoted depth (on the side that would receive the order) and the size of the incoming order, divided by the average quoted depth for the firm. This measure has two advantages over order size. First, normalization by the average depth for the firm allows aggregation across the different firms in the sample. Second, the comparison of order size to quoted depth allows for changes in the state of the market. If for some reason the specialist is more concerned that there has been an information event, he may reflect this concern by lowering his quoted depth. See Lee, Mucklow, and Ready (1993) and Kavajecz (1997). In these cases he may also be more sensitive to a larger order. Finally, Knez and Ready (1996) show that excess depth does a good job in capturing the information from both depth and order size when estimating expected price improvement for an order.

Nonparametric regression is well suited for this application because the database is large and there is no prior theoretical guidance as to the form of the functional dependence of the handling frequencies on excess depth. I use a variable span local linear smoother known as *loess* [see Cleveland (1979,1994)] to estimate the unknown regression function m from the data generating process $Y_i = m(X_i) + \varepsilon_i$, where $E[\varepsilon_i \mid X] = 0$. For a particular value (gridpoint) of the independent variable x_0, the loess estimate is given by

$$\hat{m}(x_0) = \hat{\alpha} + \hat{\beta} x_0,$$

where $\hat{a}$ and $\hat{\beta}$ minimize the weighted sum of squares,

$$\sum_{i=1}^{N} W\left(\frac{|X_i - x_0|}{\Delta(x_0)}\right)(Y_i - \alpha - \beta X_i)^2.$$

[18] For example, see Hasbrouck (1991) for empirical evidence and Easley and O'Hara (1992a) for theoretical justification.

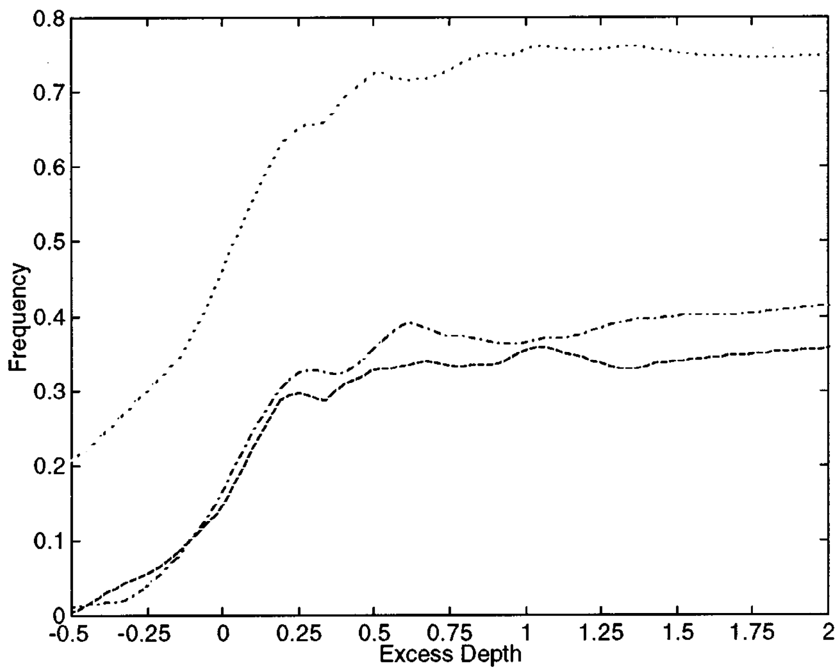

Figure 3
Order handling as a function of excess depth
Excess depth is the difference between the quoted depth and the size of the incoming order, divided by the average quoted depth for the firm. The figure is based on the model sample. The three curves show the estimated frequency of a particular action as a function of excess depth, estimated using a loess nonparametric regression. The dashed curve shows immediate price improvement, the dot and dash curve shows stops and the dotted curve shows the frequency of price improvement for only those orders that were stopped.

The weight function $W(u)$ is the *tricube kernel* given by $(1 - |u|^3)^3$ for $|u| < 1$ and 0 otherwise. $\Delta(x_0)$ determines which observations are included in the estimation according to a spanning parameter. For the estimates in both Figures 3 and 4, I use a spanning parameter of 0.25, so at each gridpoint x_0, $\Delta(x_0)$ is equal to the 25th percentile of the distances of the observations from the gridpoint given by $|X_i - x_0|$ for $i = 1, 2, \ldots, N$.

Figure 3 shows loess estimates of the frequency of immediate improvement (dashed curve) and the frequency of stops (dot and dash curve) as a function of excess depths for the model sample. The curves are plotted over a range of excess depths from -0.5 to 2.0 in order to focus on the interesting features of the estimate in the region that contains most of the orders. This region contains 86% of the orders with negative excess depth and 86% of the orders with positive excess depth. The dotted curve in Figure 3 shows the frequency of subsequent price improvement for only those orders in the

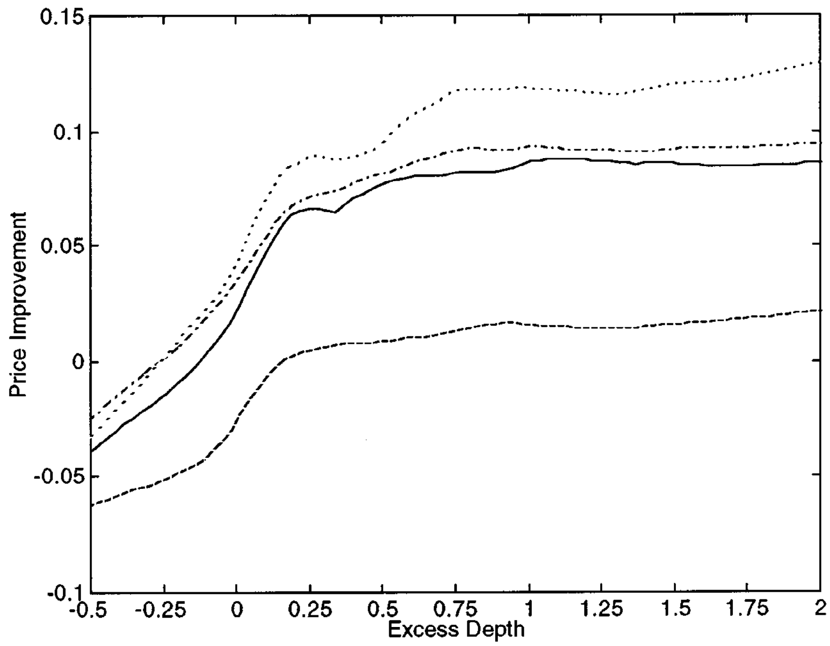

Figure 4
Price improvement as a function of excess depth
Excess depth is the difference between the quoted depth and the size of the incoming order, divided by the average quoted depth for the firm. Price improvement is trade price minus the bid for sell orders and ask minus trade price for buy orders. The curves in the figure show expected price improvement as a function of excess depth, estimated using a loess nonparametric regression. The solid curve is the relationship for all orders when the quoted spread was $1/4. The dashed curve shows the relationship for the subset of these orders where the last trade was on the receiving side of the quote, the dash and dot curve shows orders where the last trade was at the midpoint, and the dotted curve shows orders where the last trade was on the other side of the quote.

subsample that were initially stopped. All three curves increase dramatically as excess depth becomes positive, and then level off as excess depth exceeds about 0.5. These results are consistent with the hypothesis that larger orders contain more information, but there doesn't seem to be a further reduction in the information content as one moves from medium-size orders to the smallest orders.

4. Estimation of Expected Execution Price

Knez and Ready (1996) found an important relation between excess depth and price improvement, and they provided separate estimates for each spread size. They considered several conditioning variables beyond excess depth and quoted spread, including various statistics drawn from the 6.5-hour period preceding the order, but they found none that provided much improvement in the estimates.[19] The model in Section 1 and the results in Table 7 indicate that the position of last trade price is an important determinant of order handling, so it should also be an important determinant of price improvement. Accordingly, in this section I extend the results of Knez and Ready to include last trade price as a conditioning variable, and I show that this variable has a dramatic effect on the estimates.

Figure 4 shows the loess estimate of expected price improvement as a function of excess depth for the sample of 116,366 orders with a $1/4 quoted NYSE spread. As in Figure 3, the plot covers excess depths ranging from −0.5 to 2.0, which captures 87% of the market orders in the sample with positive excess depth and 87% of the orders with negative excess depth. The solid line shows the estimated curve for all orders where the quoted spread was $1/4, which is essentially the same as the curve shown in Figure 2 of Knez and Ready. As they report, the expected price improvement is relatively constant in the region of positive excess depths, but declines dramatically as excess depth falls below zero (as order size exceeds the quoted depth).

The broken curves in Figure 4 show the relationships between expected price improvement and excess depth for three subsamples separated based on the last trade price. The dashed curve shows orders where the last trade was on the side of the quote that would receive the order, the dash and dot curve shows orders where the last trade was at the midpoint, and the dotted curve shows orders where the last trade was on the other side of the spread. The curves all have the same characteristic shape as the curve for the overall sample, indicating that the relationship between price improvement and excess depth discovered by Knez and Ready persists after controlling for the size of the potential price move. The interesting result from Figure 4 is the dispersion between the three curves, which indicates that last trade price provides important additional information beyond excess depth in explaining expected price improvement.

5. Summary and Conclusions

In this article I develop a model of the NYSE specialists' use of stops and immediate price improvement in competition with limit order traders. Consistent with the predictions of the model, I find that stops and immediate price improvement impose adverse selection costs on limit order traders and I find that the stop frequency is highest when the last trade was at the midpoint of the current spread. I also show that the estimate of expected

[19] Petersen and Fialkowski (1994) provide an OLS analysis of the determinants of price improvement, but they do not consider last trade price as a conditioning variable.

price improvement can be sharpened by conditioning on the last trade price.

The model allows an analysis of the possible costs and benefits of allowing the specialist to stop orders, including the possibility that the adverse selection costs imposed by stops on limit order traders may cause them to place their orders farther away from the market, resulting in a wider quoted spread. Averaging across all of the parameter combinations where either \$1/8 or \$1/4 spread would be an equilibrium when stops are allowed, the model predicts that the elimination of stops would help the limit order traders at the expense of the specialist and the liquidity traders. Informed market order traders would be roughly unaffected. The reason that stops provide a greater benefit to the uninformed traders is that their orders are relatively less likely to be followed by another order in the same direction, so their orders are more likely to ultimately receive price improvement from a stop. Of course, these conclusions may be sensitive to the underlying model parameters, but they appear to be fairly robust.

The result that stops provide a relatively greater benefit to uninformed traders is an important contribution of the model. The analysis indicates that policymakers interested in attracting "liquidity traders" to the market might continue to allow the specialist to use stops, even if the specialist's choice of when to use them is based solely on maximizing his own trading profit. More importantly, there are other considerations that go beyond the factors considered in the model. For one thing, the specialist is given a positive obligation to be the "buyer or seller of last resort." Accordingly, it may be necessary to give him adequate ability to make trading profits during normal episodes, so that he can cover the losses that occur during volatile periods when the other suppliers of liquidity leave the market.

Appendix A: Model Calculations and Proofs

A.1 Alternative Forms of the Value Function (Section 1.3)

The expectation of the total value of the specialist's position as of period T (net of purchase costs plus sales revenues) can be expressed as

$$E_0\left[E_T[V \mid x_T^2]\sum_{t=1}^{T}\left[u_t^1 - u_t^2\right] - \sum_{t=1}^{T}u_t^1(B + .125) + \sum_{t=1}^{T}u_t^2(A - .125)\right].$$

This is equivalent to the following optimal value function (used in Section 1.3):

$$\begin{aligned} J_\pi \;=\; E_0\Bigg[&\sum_{t=1}^{T}u_t^2(x_t)\{(A - .125) - E_t[V \mid x_t^2]\} \\ &+ \sum_{t=1}^{T}u_t^1(x_t)\{E_t[V \mid x_t^2] - (B + .125)\}\Bigg].\end{aligned}$$

Proof.

$$E_0\left[E_T[V \mid x_T^2]\sum_{t=1}^{T}\left[u_t^1-u_t^2\right]-\sum_{t=1}^{T}u_t^1(B+.125)+\sum_{t=1}^{T}u_t^2(A-.125)\right]$$

Collect terms and note the dependence of u_t on x_t to give

$$E_0\left[\sum_{t=1}^{T}u_t^1(x_t)\{E_T[V \mid x_T^2]-(B+.125)\}+\sum_{t=1}^{T}u_t^2(x_t)\{(A-.125)-E_T[V \mid x_T^2]\}\right].$$

Move the unconditional expectation inside the sums and condition on x_t to give

$$\begin{aligned}
&= \sum_{t=1}^{T}E_0\left[E_t\left[u_t^1(x_t)\{E_T\left[V \mid x_T^2\right]-(B+.125)\} \mid x_t\right]\right]\\
&\quad+\sum_{t=1}^{T}E_0\left[E_t\left[u_t^2(x_t)\{(A-.125)-E_T\left[V \mid x_T^2\right]\} \mid x_t\right]\right]\\
&= \sum_{t=1}^{T}E_0\left[u_t^1(x_t)\{E_t[E_T[V \mid x_T^2] \mid x_t]-(B+.125)\}\right]\\
&\quad+\sum_{t=1}^{T}E_0\left[u_t^2(x_t)\{(A-.125)-E_t[E_T[V \mid x_T^2] \mid x_t]\}\right].
\end{aligned}$$

Note that x_t^2 is the part of x_t that is relevant for computing $E[V]$ and that x_T^2 contains the information in x_t^2. This means that

$$E_t[E_T[V \mid x_T^2] \mid x_t]=E_t[V \mid x_t^2]$$

and the final result follows from a simple substitution.

A.2 Calculation of $P_t(\cdot|x_t)$ for w_t and calculation of $E_t[V \mid x_t]$ (Section 1.3)

Calculating these objects requires calculating the updated probability assessments of whether an information event has occurred and whether the firm value is low or high. As in Easley and O'Hara (1992b), these probabilities are calculated in each period using Bayes' rule.

Following Easley and O'Hara, let $\Psi \in \{N, L, H\}$ be a random event where $\Psi = N$ means no information event has occurred, $\Psi = L$ means informed traders received a low signal, and $\Psi = H$ means informed traders received a high signal. At time $t = 0$, the specialist and the limit order traders assign probabilities $(1 - e_0)$, $e_0 n_0$, and $e_0(1 - n_0)$ to the three possible outcomes for Ψ. Thus the ex ante probabilities for each order sequence can be calculated by conditioning on Ψ. For example, the probability of the order sequence $x_2^2 = (1, -1)$ (a buy at the end of period 0 followed by a sell at the end of period 1) is given by

$$\begin{aligned}
\Pr\{x_2^2=(1,-1)\} &= \Pr\{\text{buy, sell} \mid \Psi=N\}\Pr\{\Psi=N\}\\
&\quad+\Pr\{\text{buy, sell} \mid \Psi=L\}\Pr\{\Psi=L\}
\end{aligned}$$

$$\begin{aligned} &+ \Pr\{\text{buy, sell} \mid \Psi = H\} \Pr\{\Psi = H\} \\ &= b_N s_N (1 - e_0) + b_L s_L e_0 n_0 + b_H s_H e_0 (1 - n_0). \end{aligned}$$

The conditional probabilities for the outcomes for Ψ are updated at $t = 1, 2, \ldots, T$ using the information provided by the order flow as follows:

$$\Pr\{\Psi = X \mid x_t^2\} = \frac{\Pr\{\Psi = X\} \Pr\{x_t^2 \mid \Psi = X\}}{\Pr\{x_t^2\}}.$$

Using these updated conditional probabilities, the updated probability of $V = V_L$ is given by $\Pr\{\Psi = L \mid x_t^2\} + \Pr\{\Psi = N \mid x_t^2\} n_0$. The updated probability measure $P_t(\cdot|x_t)$ for the random order arrivals is given by

$$\begin{aligned} \Pr\{w_t = -1 \mid x_t^2\} &= s_N \Pr\{\Psi = N \mid x_t^2\} + s_L \Pr\{\Psi = L \mid x_t^2\} + s_H \Pr\{\Psi = H \mid x_t^2\} \\ \Pr\{w_t = 0 \mid x_t^2\} &= z_N \Pr\{\Psi = N \mid x_t^2\} + z_L \Pr\{\Psi = L \mid x_t^2\} + z_H \Pr\{\Psi = H \mid x_t^2\} \\ \Pr\{w_t = 1 \mid x_t^2\} &= b_N \Pr\{\Psi = N \mid x_t^2\} + b_L \Pr\{\Psi = L \mid x_t^2\} + b_H \Pr\{\Psi = H \mid x_t^2\} \end{aligned}$$

A.3 Conditions for the Optimality of Stops (Section 1.4)

The difference in value between stopping the order and buying immediately can be written as follows

$$\begin{aligned} &\text{value(stop)} - \text{value(buy)} \\ &\quad = \Pr\{w_t = 1 \mid x_t^2\} \left[J_{t+1} \begin{pmatrix} 0 \\ (x_t^2, 1) \end{pmatrix} - J_{t+1} \begin{pmatrix} 1 \\ (x_t^2, 1) \end{pmatrix} \right] \\ &\qquad + \Pr\{w_t = 0 \mid x_t^2\} \left[J_{t+1} \begin{pmatrix} -1 \\ (x_t^2, 0) \end{pmatrix} - J_{t+1} \begin{pmatrix} 0 \\ (x_t^2, 0) \end{pmatrix} \right] \\ &\qquad + \Pr\{w_t = -1 \mid x_t^2\} \left[J_{t+1} \begin{pmatrix} -2 \\ (x_t^2, -1) \end{pmatrix} - J_{t+1} \begin{pmatrix} -1 \\ (x_t^2, -1) \end{pmatrix} \right] \\ &\qquad - \{E[V \mid x_t^2] - \$30.00\}. \end{aligned} \tag{A.1}$$

If stops are never optimal in period $t+1$ then the value of any unexecuted orders carried into period $t+1$ is just the immediate trading profit. Thus for any $x_{t+1}^1 > 0$ (some market buy orders carried over from prior periods).

$$J_{t+1} \begin{pmatrix} x_{t+1}^1 \\ x_{t+1}^2 \end{pmatrix} = J_{t+1} \begin{pmatrix} 0 \\ x_{t+1}^2 \end{pmatrix} + (x_{t+1}^1) \max[\$30.00 - E_{t+1}[V \mid x_{t+1}^2], 0] \tag{A.2}$$

and for any $x_{t+1}^1 < 0$ (market sell orders carried over)

$$J_{t+1} \begin{pmatrix} x_{t+1}^1 \\ x_{t+1}^2 \end{pmatrix} = J_{t+1} \begin{pmatrix} 0 \\ x_{t+1}^2 \end{pmatrix} + (-x_{t+1}^1) \max[E_{t+1}[V \mid x_{t+1}^2] - \$30.00, 0]. \tag{A.3}$$

Combining Equations A.1, A.2, and A.3 gives

$$\begin{aligned} &\text{value(stop)} - \text{value(buy)} \\ &\text{(no stops after period } t) \\ &\quad = -\Pr\{w_t = 1 \mid x_t^2\} \max[\$30.00 - E_{t+1}[V \mid x_t^2, w_t = 1], 0] \end{aligned}$$

$$
\begin{aligned}
&+ \Pr\{w_t = 0 \mid x_t^2\} \max[E_{t+1}[V \mid x_t^2, w_t = 0] - \$30.00, 0] \\
&+ \Pr\{w_t = -1 \mid x_t^2\} \max[E_{t+1}[V \mid x_t^2, w_t = -1] - \$30.00, 0] \\
&- \{E_t[V \mid x_t^2] - \$30.00\},
\end{aligned} \tag{3}
$$

which is Equation (3) in Section 1.4.

Section 1.4 next considers the situation where the $t+1$ expected value is above the midpoint, regardless of the direction of the order arriving at the end of period t. In this case, $E_{t+1}[V \mid x_t^2, w_t = 1] > 30.00$, $E_{t+1}[V \mid x_t^2, w_t = 0] > 30.00$, and $E_{t+1}[V \mid x_t^2, w_t = -1] > 30.00$ mean that the first term in Equation (3) disappears and the max operator in not needed in the second and third. Combining this with the fact that

$$
\begin{aligned}
\{E_t[V \mid x_t^2]\} \;=\; &\Pr\{w_t = 1 \mid x_t^2\} E_{t+1}[V \mid x_t^2, w_t = 1] \\
&+ \Pr\{w_t = 0 \mid x_t^2\} E_{t+1}[V \mid x_t^2, w_t = 0] \\
&+ \Pr\{w_t = -1 \mid x_t^2\} E_{t+1}[V \mid x_t^2, w_t = -1]
\end{aligned}
$$

yields

$$
\begin{array}{c}\text{value(stop)} - \text{value(buy)} \\ \text{(no stops after period } t) \end{array} = -\Pr\{w_t = 1 \mid x_t^2\}\{E_{t+1}[V \mid x_t^2, w_t = 1] - \$30.00\}.
$$

Finally, Section 1.4 considers the case depicted in Figure 2, wherein $E_{t+1}[V \mid x_t^2, w_t = 1] > 30.00$, $E_{t+1}[V \mid x_t^2, w_t = 0] > 30.00$, and $E_{t+1}[V \mid x_t^2, w_t = -1] < 30.00$. In this case the first and third terms in Equation (3) disappear and the max operator disappears from the second term, leaving

$$
\begin{aligned}
\begin{array}{c}\text{value(stop)} - \text{value(buy)} \\ \text{(no stops after period } t) \end{array} &= \Pr\{w_t = 0 \mid x_t^2\}[E_{t+1}[V \mid x_t^2, w_t = 0] - \$30.00] \\
&\quad - \{E_t[V \mid x_t^2] - \$30.00\}.
\end{aligned}
$$

Stopping a sell is optimal at the start of period t if and only if the above is ≥ 0, or equivalently,

$$
E_{t+1}[V \mid x_t^2, w_t = 0] - \$30,000 \geq \frac{E_t[V \mid x_t^2] - \$30.00}{\Pr\{w_t = 0 \mid x_t^2\}}. \tag{4}
$$

A.4 Effective Spread Calculations for Table 2 (see Section 1.7)

Define "expanded" order arrivals W_t, where $W_t = 0$ means an informed seller, $W_t = 1$ an uninformed seller, $W_t = 2$ no trade, $W_t = 3$ an uninformed buyer, and $W_t = 4$ an informed buyer. Let X_t denote the "expanded" order history, so $X_0 = \varnothing$ and $X_{t+1} = (X_t, W_t)$. The ex ante probability for each such sequence can be calculated by conditioning on Ψ (see Section A.3). The specialist's optimal policy is a function of the sequence of buys, sells, and no trades. Using this information, let $R_t(X_T)$ be the ultimate execution price for the market order that arrived in period t (unless there was no arrival, then $R_t(X_T) = 0$) as part of the expanded sequence X_T. Also, let O_t be an indicator function with $O_t(X_T, k) = 1$ if the order in period t of the expanded sequence X_T has $W_t = k$ (0 otherwise). Let P_k denote the long-run average trade price for order type k. Finally, let H denote the set of all possible "expanded" order histories through period

T (H has 5^T elements). Then

$$P_k = \frac{\sum_{X_T \in H} \left[\sum_{t=1}^{T} O_t(X_T, k) R_t(X_T) \right] \Pr\{X_T\}}{\sum_{X_T \in H} \left[\sum_{t=1}^{T} O_t(X_T, k) \right] \Pr\{X_T\}} \qquad \text{for } k = 0, 1, 2, 4.$$

For a particular parameter combination, the uninformed traders' effective spread is $P_3 - P_1$, and the informed traders' effective spread is $P_4 - P_0$.

A.5 Profit/Market Order Calculations for Table 2 (see Section 1.7)

Let c_t be an indicator function with $c_t(x_T^2, m) = 1$ if a market order arrived in period t as a part of the order sequence x_T^2 and crossed with liquidity supplier of type m, where $m = 1$ means a limit buy order, $m = 2$ a limit sell order, $m = 3$ the specialist as a buyer, $m = 4$ the specialist as a seller, and $m = 5$ a stopped market order from a prior period. Profits are computed by comparing execution prices to the ex post expected value given the realized order sequence. Effectively this computation implies that there is no future microstructure-related profit or loss ascribed to "unwinding" the position. Let h denote the set of all possible order sequences through period T (h has 3^T elements). For a particular parameter combination, the long-run average limit order profit per market order is

$$\frac{\sum_{x_T^2 \in h} \left[\sum_{t=1}^{T} c_t(x_T^2, 1)[E_T[V \mid x_T^2] - B] + c_t(x_T^2, 2)[A - E_T[V \mid x_T^2]] \right] \Pr\{x_T^2\}}{\sum_{x_T^2 \in h} \left[\sum_{t=1}^{T} \sum_{m=1}^{5} c_t(x_T^2, m) \right] \Pr\{x_{\}}^2}.$$

The long-run specialist profit per market order is

$$\frac{\sum_{x_T^2 \in h} \left[\sum_{t=1}^{T} c_t(x_T^2, 3)[E_T[V|x_T^2] - B - .125] + c_t(x_T^2, 4)[A - .125 - E_T[V|x_T^2]] \right] \Pr\{x_T^2\}}{\sum_{x_T^2 \in h} \left[\sum_{t=1}^{T} \sum_{m=1}^{5} c_t(x_T^2, m) \right] \Pr\{x_T^2\}}.$$

A.6 Profit/Potential Profit Calculations for Table 3 (see Section 1.8)

Table 3 is restricted to parameter combinations that result in \$1/4 spreads, that is, the sets Γ(profit-motivated,1/4) and Γ(combined,1/4). Let $d_t(x_T^2, m) = 1$ if a market order *buy* arrived in period t as a part of the order sequence x_T^2 *when there is not a stopped order from a previous period* (so the quoted spread is 1/4) and is ultimately handled as follows: $m = 1$ means immediately crossed with a limit order, $m = 2$ stopped and then crossed with a limit order, $m = 3$ stopped and then crossed with the specialist, $m = 4$ stopped and then crossed with an incoming market order, and $m = 5$ immediately improved by the specialist. Sell orders are defined analogously except that the m values run from 6 to 10. Profits and potential profits are computed by comparing the posted quote to the ex post expected value given the realized order sequence. As in Section A.6, this computation implies that there is no future microstructure-related profit or loss ascribed to "unwinding" the position. For a particular parameter combination, the long-run average limit order profit/potential profit for order handling category m is

$$\frac{\sum_{x_T^2 \in h} \left[\sum_{t=1}^{T} d_t(x_T^2, m)[A - E_T[V \mid x_T^2]] + d_t(x_T^2, m+5)[E_T[V \mid x_T^2] - B] \right] \Pr\{x_T^2\}}{\sum_{x_T^2 \in h} \left[\sum_{t=1}^{T} d_t(x_T^2, m) + d_t(x_T^2, m+5) \right] \Pr\{x_T^2\}}.$$

Appendix B: Monte Carlo Simulations for Tables 6 and 7

Both Tables 6 and 7 separate the market orders for each firm $f = 1, 2, \ldots, 129$ in the sample into $j = 1, 2, \ldots, J$ distinct categories and then calculate various statistics averaged across all of the orders in the category. Let

$$c_j^f = \text{the set of market orders for firm } f \text{ in category } j,$$

$$c_j = \bigcup_{f=1}^{129} c_j^f, \text{ and}$$

$$c^f = \bigcup_{j=1}^{J} c_j^f.$$

Define a counting function $N(c)$ that gives the number of market orders in the set c, and let $S(c, N)$ be a set-valued operator that returns a set with N elements randomly drawn from the set c with replacement. Let v be a statistic associated with each market order in the sample and let $\bar{v}(c)$ be the average of v over the set c. For example, v could be the 5-minute ex post profit (as in Table 6) or an indicator variable based on whether the order was stopped (as in Table 7). For each category j, the tables report $\bar{v}(c_j)$.

In the Monte Carlo test, the null hypothesis is that the distribution of v is constant across each c^f, but it may vary across firms. The tests individually determine whether the difference between two categories, $\bar{v}(c_j) - \bar{v}(c_k)$ for $j \cdot k$, is statistically different from zero. The simulated difference $\bar{v}(\cup_{f=1}^{129} S(c^f, N(c_j^f))) - \bar{v}(\cup_{f=1}^{129} S(c^f, N(c_k^f)))$ is generated 2,000 times to create the reference distribution. If the difference in the reported statistics is greater than the 50th or less than the 1,950th largest element of the reference distribution, the difference is judged to be significant at the 5% level in a two-tailed test.

References

Angel, J., 1994, "Who Gets Price Improvement on the NYSE?" Working Paper FINC-1377-12-694, Georgetown University.

Bertsekas, D., 1976, *Dynamic Programming and Stochastic Control*, Academic Press, New York, chap. 2.

Chakravarty, S., and C. Holden, 1995, "An Integrated Model of Market and Limit Orders," *Journal if Financial Intermediation*, 4, 213–241.

Cleveland, W., 1979, "Robust Locally Weighted Regression and Smoothing Scatterplots," *Journal of the American Statistical Association*, 74, 829–839.

Cleveland, W., 1994, *Visualizing Data*, Chapman and Hall, London.

Dupont, D., 1995, "Market Making, Prices and Quantity limits," working paper, Federal Reserve Board.

Easley, D., and M. O'Hara, 1992a, "Adverse Selection and Large Trade Volume: The Implications for Market Efficiency," *Journal of Financial and Quantitative Analysis*, 27, 185–208.

Easley, D., and M. O'Hara, 1992b, "Time and the Process of Security Price Adjustment," *Journal of Finance*, 47, 577–606.

Glosten, L., 1994, "Is the electronic open limit order book inevitable?" *Journal of Finance*, 49, 1127–1162.

Glosten, L., and P. Milgrom, 1985, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71–100.

Handa, P., and R. Schwartz, 1996, "Limit order Trading," *Journal of Finance*, 51, 1835–1862.

Harris, L., and J. Hasbrouck, 1996, "Market vs. Limit Orders: The SuperDOT Evidence on Order Submission Strategy," *Journal of Financial and Quantitative Analysis*, 31, 213–232.

Hasbrouck, J., 1991, "Measuring the Information Content of Stock Trades," *Journal of Finance*, 22, 229–252.

Hasbrouck, J., G. Sofianos, and D. Sosebee, 1993, "New York Stock Exchange Systems and Trading Procedures," Working Paper 93-01, New York Stock Exchange.

Kavajecz, K., 1997, "A Specialist's Quoted Depth and the Limit Order Book," working paper, University of Pennsylvania.

Knez, P., and M. Ready, 1996, "Estimating the Profits from Trading Strategies," *Review of Financial Studies*, 9, 1121–1164.

Kyle, A., 1985, "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1336.

Leach, C., and A. Madhavan, 1993, "Price Experimentation and Security Market Structure," *Review of Financial Studies*, 6, 375–404.

Lee, C., B. Mucklow, and M. Ready, 1993, "Spreads, Depths, and the Impact of Earnings Information: An Intraday Analysis," *Review of Financial Studies*, 6, 345–374.

Petersen, M., and D. Fialkowski, 1994, "Posted Verses Effective Spreads: Good Prices or Bad Quotes?" *Journal of Financial Economics*, 35, 269–292.

Rock, K., 1996, "The Specialist's Order Book and Price Anomalies," forthcoming in *Review of Financial Studies*.

Ross, K., J. Shapiro, and K. Smith, 1996, "Price Improvement of SuperDOT Market Orders on the NYSE," Working Paper 96-01, New York Stock Exchange.

Seppi, D., 1995, "Liquidity Provision with Limit Orders and a Strategic Specialist," working paper, Carnegie Mellon University.