A BAYESIAN MODEL OF INTRADAY SPECIALIST PRICING

by

Ananth Madhavan
Seymour Smidt

2-91

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

A Bayesian Model of Intraday Specialist Pricing

Ananth Madhavan
and
Seymour Smidt *

January, 1991

*Wharton School, University of Pennsylvania and Cornell University, respectively. Research support from Q-Group and the Research Foundation of the University of Pennsylvania and the Geewax-Terker Research Fund (Madhavan) is gratefully acknowledged. We have benefitted a great deal from the helpful comments of Yakov Amihud, Margaret Forster, and participants at seminars in Cornell University and the American Finance Association Meetings. We also thank the NYSE specialist (who wishes to remain anonymous) for providing us with the inventory data used in this study. Song-Qi Shen and Annu Ratta provided expert programming assistance. Any errors are entirely our own.

Abstract

We develop and test a model of intraday price formation based on an explicit description of a representative market maker whose beliefs evolve according to Bayes' rule. We derive an estimating equation where the weight the market maker places on the order flow as an information signal can be recovered from the parameter estimates. This weight is a natural measure of information asymmetry since is the ratio of the quality of private information to the quality of public information. The model is interesting for other reasons as well. First, the model encompasses several other models of intraday price formation. Second, the error term arises endogenously and possesses a natural economic interpretation. Third, the model permits us to partially distinguish the price effects of information asymmetry and inventory control by market makers. Fourth, the model provides a method to assess the implicit costs of trading. We show that there are substantial non-linearities in pricing that may reflect the way in which large blocks are traded in the upstairs market. We estimate the model with a new data set obtained from a NYSE specialist. The data set comprises almost 75,000 records for most of the year 1987 and is of independent interest given the paucity of inventory data. The results provide strong support of information asymmetries, as perceived by the market.

1 Introduction

This paper examines the determinants of intraday security price movements. Prices change when new information reaches the market. Prices also respond to trading activity. The growing literature on market microstructure has identified three major influences of trading on prices. First, transaction costs give rise to 'bid-ask bounce' as orders randomly arrive. Second, inventory carrying costs give market makers incentives to alter prices as trading causes their inventory positions to diverge from desired levels. This prediction also arises if market makers are risk averse. Papers by Amihud and Mendelson (1980), Zabel (1981), Ho and Stoll (1983), and O'Hara and Oldfield (1986)) formally model the effect of market maker inventory control on prices. Third, traders with private information about the value of the security induce market makers to revise prices in the direction of order flow, since it provides a noisy signal as to the information of informed traders. Papers by Kyle (1985), Glosten and Milgrom (1985), Easley and O'Hara (1987), Admati and Pfleiderer (1989), among others) have addressed this issue. Both the inventory and information theoretic models predict that prices move in the direction of order flow, but the two theories have not been integrated in a single empirical model.

We develop a model of price formation that incorporates these trade related factors, as well as the effect of unanticipated news shocks. The model is tested with a new data set that provides a record of a New York Stock Exchange (NYSE) specialist's intraday inventory position for almost a year.[1] In the model, we explicitly describe the process by which market makers learn from order flow. In doing so, we provide a natural summary measure of the information structure of the market.

The idea underlying our metric for information asymmetry is simple. Consider the learning process of a representative market maker who uses Bayesian rules to update his beliefs.

[1]On the NYSE, each listed stock is assigned to one market maker, the specialist. Stoll (1985) provides a detailed description and analysis of the specialist's activities.

is evidence that the price impact of large trades is non-linear, a finding attributable to the institutional features of the market. We discuss two possibilities: constraints imposed on the specialist in the form of price continuity-depth requirements and the pricing of large block trades in the 'upstairs' market.

The model is tested with a new data set drawn from a NYSE specialist's trading records, together with data from the Institute for the Study of Securities Markets (ISSM). The data set is of independent interest since it contains almost 75,000 specialist trading records for most of the year 1987, including the October crash. With these data, we can construct one of the largest and most detailed time-series of inventory and signed transaction volume currently available.

The paper proceeds as follows. In section 2, we develop a model of price formation that incorporates transaction costs, inventory effects, and the effect of public and private information. We show that the Bayesian weight can be inferred from the estimated parameters of the model. In section 3, we describe our data and the procedures used to verify its accuracy. Section 4 describes the model estimation technique and the results. In section 5, we modify the model to incorporate certain institutional features of the market and then apply the model to compute the implicit costs of trading. Finally, section 6 summarizes the paper.

2 A Bayesian Model of Intraday Price Formation

In this section, we develop a theoretical model based on the recent literature in market microstructure. Consider a multi-period model with two assets, a riskless bond (the numeraire) and a stock with a stochastic liquidation value. The risky security is traded at times $t = 1, 2, \ldots, T$, and its full information price at time T is denoted by $\tilde{v}$. The full information value is composed of a series of increments or 'dividends' and we write $\tilde{v} = d_0 + \sum_{i=1}^{T} \tilde{d}_i$, where d_0 is a positive constant. The increments are independently and identically distributed with zero mean. The increment d_t is realized immediately after trading in period t. In our model, prices that straddle the value of the asset. In the face of stochastic demand and supply the market maker's inventory will follow a random walk. Consequently, if the market maker's capital is finite, eventual bankruptcy is certain, as shown by Garman (1976). Consequently, the market maker adopts a non-stationary pricing policy that depends on the current *level* of inventory. Formally, bid and ask prices are the controls while inventory is the state variable. If market makers are long (short) relative to their desired inventory level they try to attract buy (sell) orders by lowering (increasing) the price. This gives rise to a negative relationship between price and the market maker's current inventory position, as derived by Amihud and Mendelson (1980).

The price p_t consists of the expected value of the security conditional upon all information available to the market maker at time t plus the 'contamination' by microstructure elements such as transaction costs and inventory effects. In the prototypical inventory control model, price is linearly related to the market maker's current share inventory:

$$p_t = \mu_t - \gamma(I_t - I_d) + \psi D_t. \tag{1}$$

In equation (1), μ_t is the expectation of $\tilde{v}_t$ conditional upon the market maker's information at time t, I_d is the long-run desired inventory level (assumed constant), I_t is the market maker's current share inventory, D_t is an indicator variable where $D_t = +1$ for a buy order and -1 for a sell order, and $\psi \geq 0$ and $\gamma \geq 0$ are constants. Note that in the absence of any market imperfections, $\gamma = \psi = 0$, and the model reduces to $p_t = \mu_t$.[4] The basic property of conditional expectations implies that prices, in this case, follow a random walk.

Equation (1) is a reduced form expression for a number of inventory control models. The linearity is not particularly stringent; linear decision rules are optimal in a number of formal inventory models including Zabel (1981) among others. The spread element ψ, is interpreted

[4]There are strong theoretical reasons to believe $\psi > 0$. Otherwise, if $\psi = 0$, the market maker would have negative expected profits, assuming buy (sell) orders are more likely when he is long (short) and sets prices below (above) the equilibrium price.

magnitude of the order flow which is a noisy signal of the private information of informed traders. Since both μ_t and D_t are functions of order quantity, q_t, equation (1) cannot be estimated directly without formally specifying the market maker's learning process.

We assume that just before time t, all agents observe the realization of a noisy public information signal concerning the value of the increment d_t at time t.[9] Since v_{t-1} is public information at time t, a signal about d_t can be expressed as a signal, denoted by $\tilde{y}_t$, of the form:

$$\tilde{y}_t = v_t + \tilde{\varepsilon}_t \ . \qquad (3)$$

In equation (3), $\tilde{\varepsilon}_t$ is an independently normally distributed error term with mean zero and variance σ_ε^2. The market maker's prior mean is the realization of $\tilde{y}_t$, denoted by y_t. In addition to public information, the trader at time t receives a private (noisy) information signal, $\tilde{w}_t$, about the value of v_t. The private signal has similar structure to the common signal, i.e., $\tilde{w}_t = v_t + \tilde{\omega}_t$, where $\tilde{\omega}_t$ is independently and identically normally distributed with zero mean and constant variance, denoted σ_ω^2. Let w_t be the realization of $\tilde{w}_t$. As the trader's prior distribution of $\tilde{v}_t$ at time t is normal and the private signal is drawn from a normal distribution, we can apply a fundamental theorem in statistical decision theory to compute the trader's posterior mean.[10] This mean is denoted m_t, and is given by:

$$m_t = \theta w_t + (1 - \theta) y_t \qquad (4)$$

where $\theta = \sigma_\varepsilon^2 / (\sigma_\varepsilon^2 + \sigma_\omega^2)$.

The trader's demand depends on the functional form of the price schedule quoted by the market maker. We assume the order quantity, q_t, can be expressed in the form:

$$q_t = \alpha(m_t - p_t) - x_t \qquad (5)$$

where α is a positive constant, m_t is the trader's expectation of $\tilde{v}_t$ at time t (given the

[9] Equivalently, all market participants share a common prior distribution over the value of the increment.
[10] See, for example, DeGroot (1970), page 169.

where $\pi \equiv (\zeta+\theta-1)/\theta$. From the definitions of ζ and θ it can be shown that $\pi \in (0,1)$. From equation (4), the market maker's minimum variance estimate of m_t given q_t is $(p_t + q_t/\alpha)$, so that equation (7) shows that the posterior mean is a weighted average of prior information and the information revealed by the trade. The parameter π is the weight (corrected for the statistical dependence of m_t on y_t) placed on prior beliefs. In a market with significant volume of liquidity-based trading, (i.e., σ_x is relatively large), accurate public information (i.e., σ_ϵ is relatively small) and imprecise private information (σ_ω is relatively large), π is high, indicating the market maker relies heavily upon prior beliefs. So, π is inversely related to the degree of information asymmetry in the market. We will derive an equation where π can be recovered from the estimates directly. This estimate will yield the market maker's perception of the degree of information asymmetry in the market.

Substituting equation (7) into equation (1), We obtain:

$$p_t = \pi y_t + (1-\pi)[p_t + \alpha^{-1} q_t] - \gamma(I_t - I_d) + \psi D_t \tag{9}$$

Since y_t, the market maker's prior mean at time t, is unobservable, we cannot estimate (9) directly. Our solution is to find a proxy for this unobservable variable. This proxy is based on the previous price after adjusting for 'contamination' by transaction costs and inventory effects. Accordingly, let us use equation (1) to write y_t in the form:

$$y_t = p_{t-1} + \gamma(I_{t-1} - I_d) - \psi D_{t-1} + \eta_t \tag{10}$$

where $\eta_t \equiv y_t - \mu_{t-1}$ is the difference between the prior mean at time t and the posterior mean at time $t-1$. The prior and posterior differ because of public information signals, specifically the revelation of the increment d_t. Thus, η_t represents the innovation in the market maker's conditional expectations of the security's value. This innovation cannot be predicted by an econometrician, and gives rise to the error term in our model.

2.1 Error Structure

From equation (6), we see that $\hat{w}(q_{t-1}) = v_{t-1} + \omega_{t-1} - (\alpha\theta)^{-1}x_{t-1}$. So, using the definition of y_t and the definition of μ_t given in (6), we see that:

$$\eta_t = (v_t - v_{t-1}) + \varepsilon_t - \zeta\varepsilon_{t-1} - (1-\zeta)[\omega_{t-1} - (\alpha\theta)^{-1}x_{t-1}]. \tag{12}$$

Define u_t by:

$$u_t \equiv (v_t - v_{t-1}) - (1-\zeta)[\omega_{t-1} - (\alpha\theta)^{-1}x_{t-1}] \tag{13}$$

Then, it follows that:

$$\eta_t = \varepsilon_t - \zeta\varepsilon_{t-1} + u_t. \tag{14}$$

Under our assumptions about the stochastic process $\{d_t\}$, it follows that $E[\tilde{v}_t|v_{t-1}] = v_{t-1}$. Our assumptions about signal structure and the martingale property of $\{\tilde{v}_t\}$ imply that $E[\tilde{u}_t] = 0$ and $E[\tilde{u}_t\tilde{u}_{t-1}] = 0$.[15] Taking expectations in (14), and using the martingale property, we obtain:

$$E[\tilde{\eta}_t] = 0 \tag{15}$$

and

$$E[\tilde{\eta}_t\tilde{\eta}_{t-1}] = -\zeta\sigma_\varepsilon^2 < 0 \tag{16}$$

From (14), the error term η follows a MA(1) process with parameter ζ, where ζ is inversely related to the degree of information asymmetry in the market. Using the definition of ζ, write $\zeta = \sigma_\varepsilon^{-2}/(\sigma_s^{-2} + \sigma_\varepsilon^{-2})$. So, $\zeta \in (0,1)$, with lower values of ζ corresponding to greater asymmetries between public and private information. If information asymmetries were less severe for actively traded stocks than for thinly traded stocks, we would expect ζ to rise with trading volume or market value. Unfortunately, since we cannot distinguish the model errors due to public information shocks from specification or measurement errors, we cannot use the estimates of ζ to draw conclusions about the information structure in the market. Accordingly, we restrict our attention to π.

[15]This follows from our assumption that $E[\tilde{x}_t\tilde{x}_{t'}] = E[\tilde{\omega}_t\tilde{\omega}_{t'}] = 0$ for $t \neq t'$.

Equation (18) represents Roll's (1984) model, with the additional assumption that D_t is serially independent and has mean zero. The first-order serial covariance of successive price changes given by equation (18) is $\text{Cov}(\Delta p_t, \Delta p_{t-1}) = E[\Delta p_t \Delta p_{t-1}] = -\psi^2$. Since the bid-ask spread, denoted by s, is simply 2ψ, an unbiased estimator of the spread given by:

$$s = 2\sqrt{-\text{Cov}(\Delta p_t, \Delta p_{t-1})}. \tag{19}$$

The advantage of Roll's model is that a bid-ask spread estimate can be obtained using only transaction prices.[16]

Observe that the absolute size of the coefficients of D_t and D_{t-1} are the same in (18). This restriction is implicitly imposed by Ho and Macris (1984) and Glosten and Harris (1988). By contrast, the key to estimating the information parameter π in our model is the absence of this restriction in equation (11). To understand this, consider a simple variant of Roll's model where some traders possess private information concerning the (unknown) value of the security. Using equation (1), we obtain:

$$p_t - p_{t-1} = \psi D_t - \psi D_{t-1} + \eta_t. \tag{20}$$

where $\eta_t \equiv \mu_t - \mu_{t-1}$ represents the innovation in the market maker's beliefs. The market maker's innovation η_t is positively correlated with the current trade because order flow partly originates from informed traders. Suppose, for example, that trade size is fixed at one round lot and that $\eta_t = \xi D_t$, where $\xi > 0$ is a constant. If the trader buys, the market maker's conditional value increases by ξ; if the trader sells, the conditional value falls by ξ. Equation (18) can be written as:

$$p_t - p_{t-1} = (\psi + \xi) D_t - \psi D_{t-1}. \tag{21}$$

We can use equation (21) to obtain estimates of ψ and ξ using only transaction prices. Under Roll's assumptions, we have $\sigma^2(\Delta p_t) = (\psi+\xi)^2 + \psi^2$ and $\text{Cov}(\Delta p_t, \Delta p_{t-1}) = -(\psi+\xi)\psi$, and

[16] However, Harris (1990) demonstrates that estimates from Roll's model are very noisy if the underlying value is subject to random innovations.

2.3.3 The Glosten-Harris (1988) Model

Glosten and Harris (1988) assume a linear price adjustment rule to capture the information effect and assume a fixed cost of executing a trade. Their model, expressed in our notation, resembles (11):

$$\Delta p_t = \lambda q_t + \psi D_t - \psi D_{t-1} + \eta_t \ . \tag{23}$$

Note that in (23), the coefficients of D_t and D_{t-1} are restricted to be the same. As discussed earlier, this restriction is important if there are information asymmetries. For this reason, the Glosten-Harris model is not equivalent to our model without an inventory effect. To see this, with $\gamma = 0$ equation (11) reduces to:

$$\Delta p_t = \kappa + \lambda q_t + \psi D_t - (\frac{\psi}{\pi}) D_{t-1} \tag{24}$$

which resembles the Glosten-Harris model except in the difference in the coefficients of the signed variables. This is also the restriction implicit in Roll's model, and the remarks made in section 2.3.1 apply here as well. In particular, only if both $\gamma = 0$ and $\pi = 1$ does (11) reduce to the Glosten-Harris model, but this corresponds to a case with no inventory effects and perfect public information. Glosten and Harris estimate (23) for a sample of NYSE stocks in the period 1981-1983. They then analyze the cross-sectional determinants of the estimated time-series parameters, concluding that a significant portion of the bid-ask spread may be attributable to information asymmetry.

2.3.4 Hasbrouck's (1990) Model

A recent paper by Hasbrouck (1990a) uses an elegant approach to assess the effect of microstructure effects on stock prices. Hasbrouck models the observed transaction price, p_t, as the sum of a random-walk component, denoted m_t and a stationary component, denoted by s_t:

$$p_t = m_t + s_t \ . \tag{25}$$

special mention. First, since the data are taken directly from the computerized records maintained by the specialist for operational purposes, those fields in the records that are important to the specialist are of very high quality. These include the price per share and quantity of shares purchased or sold. Second, as we explain below, the data structure allows us to perform several cross-checks and verify the accuracy of some of the fields in the trading record.[18] Finally, the nearly 75,000 transactions covered in the data set, comprising all transactions by the firm in its specialty stocks for nearly an entire calendar year, makes it the largest time-series presently available on specialist transactions.

The original data set consisted of two types of records, trading records and settlement records. The trading records are analogous to business invoices. They represent the specialist firms's contemporaneous record of transactions in which the firm believed that it bought or sold stock. Settlement records are analogous to canceled checks. They indicate that the other party to the transaction confirms that the transaction took place and that the specialist firm has paid for the stock (in the case of the purchase) or been paid (in the case of a sale). A trading or settlement record includes the following fields:

- CUSIP number of the stock.
- Trade time and date: separate fields identify the settlement date.
- Amount: The dollar amount paid or received by the specialist as a result of the transaction.[19]
- Price: The transaction price, recorded in eighths.
- Quantity: The signed volume of shares traded; a separate buy-sell indicator provides a check on the accuracy of the sign.

[18]This is especially important since the data cover the period of the crash in October, 1987.

[19]In the case of a stock purchase it is equal to the price times the quantity; in the case of a stock sale it is the price times the quantity less a small SEC fee paid by the seller. A separate code indicates whether the amount is a cash inflow or an outflow.

Next, we performed internal checks for consistency using the fact that some fields within a settlement record are redundant. For example, the amount, for purchase transactions, is simply the price times the quantity. Further, given signed quantities, the buy-sell field is unnecessary. Consequently, omitted variables rarely present a problem with these data.[21]

3.1.2 Corrections and Additions

Given the nature of the data, there were very few problems in determining the price, the quantity involved, and whether the specialist bought or sold. On the other hand there were considerable problems in determining the precise time at which the transaction took place. For some of our analyses we sought to match the transactions reported in the specialist's file with the transactions reported in the Institute for the ISSM file. For this purpose, an accurate trade time data would have been extremely helpful.

While a time code is included on the specialist's trading and settlement records, the time field format allows for time to be recorded only in terms of hours and minutes and the time field entry is missing for some round lot transactions. During periods of active trading there may be more than one transaction per minute. In the settlement process, the time field is used when there is a problem matching buyer and seller records; but when there is no problem matching buyer and seller, settlement can take place even if the times recorded on the buyer and seller records differ.

In addition, there were two problems in ordering the transaction records: some records are missing time stamps and sometimes two or more records corresponding to different transactions have the same time stamp. Specialist odd-lot purchases (sales) through the day in each stock are reported as a single trade with no price or time fields entries in the

[21]Special treatment was necessary for odd-lot transactions. NYSE specialists act as odd-lot dealers in their specialty stocks. Odd-lot transactions are not reported on the NYSE ticker or in the ISSM trade and quote file. The specialist transaction file contained daily summaries of the specialist's odd-lot transactions for each stock. These records are described in more detail in section 3.1.2.

the identity of the firm, we report only data for the 16 stocks for which a complete sequence of transactions was available from February 1 to December 31, 1987 and which had an average of at least 4 transactions per day on the NYSE.

3.2 ISSM data

The data on transactions and quotes used in this study was obtained from the ISSM "Trades and Quotes" transaction files. This database includes all information distributed by the Securities Industry Automation Corporation (SIAC) over their high speed line. The version of the tape we used was missing data for certain days because that data had not been supplied to ISSM by SIAC at the time we obtained the trades and quotes file.[23] These dates were omitted from our study, and for practical purposes our data set covers 199 trading days.

3.2.1 Classification of Active Trades

Although the specialist data allows us to sign volume, the ISSM files do not indicate whether a transaction was initiated by a buyer or seller. However, we can determine this from information contained in these files. The traditional method of making that determination is the tick test, which is based on the sign of the price change. Specifically the current trade price is compared with the most recent different price. Trades with positive (negative) price changes are assumed to be initiated by buyers (sellers).

We used an alternative methodology developed by Lee and Ready (1989) that classifies trades by comparing the trade price with the prevailing quote whenever one is available. Lee and Ready consider quotes that are eligible for inclusion in the National Market System and NASD Best Bid and Offer calculations, so called BBO quotes.[24]

[23]The 33 missing days are February 25, April 7, April 16, May 1, May 13, June 17, July 6–8, August 3–7, 10–14, 17–21, 24–28, November 23–25 and 27.

[24]For a description of the BBO criteria, see "Preliminary Documentation for NYSE and AMEX Trades and Quotes File," Institute for the Study of Securities Markets, October, 1989, pp.13-16. The ISSM quote records include associated condition codes which allow users to identify BBO eligible quotes.

all stocks, we had 74,360 specialist settlement records in our files. Of these, 61,961 had time stamps. Of the time stamped records, 5,721 occurred on days for which ISSM data was missing. Therefore, there were 56,240 settlement records that were eligible to be matched. Using the criteria described above, we found matches for 44,164 records, or 78.5% of those eligible to be matched.[25]

In running the model, we used a matched data set that contained all ISSM transactions with the record fields for each transaction augmented to include information about the specialist's inventory level at the time of the transaction (for all transactions) and the number of shares the specialist bought or sold if it was a matched transaction. If inventory level information had been computed using only matched transactions, failures to match would have induced errors in the estimated inventory levels on subsequent transactions. To avoid this problem, the inventory level information in the matched file was taken from the specialist transaction file.

We adjusted the data for stock splits and dividends. An additional correction was necessitated by NYSE opening procedures. The NYSE opens with an auction characterized by a set of multilateral transactions at a single price. Transactions during the day, however, take place in a dealer market with bilateral transactions over time. Amihud and Mendelson (1987) and Stoll and Whaley (1990) provide evidence that these differences in market protocols have a significant impact on the observed return distribution. Accordingly, we dropped the overnight price change from our study, so that all price changes reflect only changes during trading hours.

[25]Our high success rate in identifying and matching transactions relative to other researchers may be attributable to the fact that the problematic transactions, from an econometrician's point of view, are those that occur within the quoted bid-ask spread. Since the specialist is often involved in such transactions, signing these trades presented no difficulty.

cross-sectional regression of the medians of the matched specialist transaction values on the median total value of trading has a slope of 0.303 (t-ratio = 9.57) and a constant term that is insignificantly different from zero. Thus there is no evidence that this specialist participates more heavily in thinly traded stocks. Adjusting for unmatched transactions, we estimate that the specialist takes the other side of about 39 percent (0.303/0.785) of the value of the trades taking place in his stocks.

Since all trades are classified as either buyer initiated or seller initiated, we calculated a measure of trading imbalance as the absolute difference between the dollar value of the purchases and sales. Median trading imbalances are highly correlated with the median dollar value of trading. A natural question concerns the correlation of specialist inventory changes with order imbalances. The NYSE argues that specialists perform a valuable service of price stabilization by 'leaning against the wind,' i.e., standing ready to absorb transitory order imbalances. Critics counter that the specialist's actions may in fact exacerbate temporary price swings, at the expense of limit orders. The resolution of this issue is essentially an empirical question. To address this, we estimated a cross-sectional regression of the medians of the specialist's dollar imbalance (i.e., the absolute value of the difference between the dollar values of his purchases and sales) on the median New York dollar imbalance. The slope coefficient is 0.297 (t-ratio = 9.88) and the constant term is insignificantly different from zero. If the specialist were always on the opposite side of the New York imbalance, this could be interpreted as indicating that the specialist absorbed an average of 38 percent (i.e., 0.297/0.785) of the market imbalance, a figure almost identical to the specialist's share of total transaction value. This computation suggests the specialist's participation rate is relatively insensitive to the prevailing market conditions, i.e., there is no evidence that his trading activity is altered by market stress.

inventory level is weak. Similarly, the coefficient of the lagged inventory variable should be positive, and the results are analogous to the current inventory variable. These weak results may reflect multicollinearity between current and lagged inventory. While multicollinearity increases the estimated standard errors the coefficient estimates are unbiased and consistent. The estimates of γ suggest that inventory effects have a small impact on intradaily pricing. Further, the fact that λ is lower than expected, together with the strong evidence of information asymmetries as measured by π suggest that current inventory carrying costs not captured by γ have little economic significance.

Turning now to the sign variables, both variables are highly significant and of the correct sign. The estimated spread element ψ is simply the negative of the coefficient of D_{t-1}. Price discreteness implies that the minimum tick is \$0.125, so we would expect the spread element to be at least 0.0625. The mean value of the estimated coefficient ψ is 0.089, and is fairly tightly distributed around this value. The high significance values and closeness of the estimates (for 8 stocks, the estimates lie in the range 0.05–0.07) suggest that the spread element is very important. Our restrictions imply that the absolute value of the coefficient of D_{t-1} should be less than that of D_t to ensure that $\pi \in (0, 1)$. This condition is also satisfied in all 16 cases. The average coefficient estimate (weighted by transaction frequency) of the weight on prior information π is 0.76. Again, this finding provides strong evidence of information asymmetries. Finally, there is strong evidence that the error structure does indeed follow a MA(1) process. The moving average parameter, ζ, with the exception of the least traded stock, is of the correct sign. All but 3 are significant at the 5 percent level, with the significance levels rising with the trading activity. It can be seen from the table that ζ is, as hypothesized, positively related to trading activity.

To check the sensitivity of the parameter estimates to the estimation technique, we also estimated the reduced-form model using a high-order autoregressive error structure to approximate the MA(1) term. The coefficient estimates produced by this technique were works well. The significance levels confirm our earlier findings that there is considerable information asymmetry, as perceived by the specialist. Overall, the goodness-of-fit is good, even in October, and the R^2 is generally high.[30]

4.1 The Specialist's Participation Rate

Before turning to some applications of the model, we will briefly discuss the determinants of the specialist's trading activity. Let r_i denote the specialist's participation rate, i.e., the ratio of the specialist's trading volume to NYSE trading volume for the period February-December, for stock $i = 1, \ldots, 16$. In our sample, the participation rate ranges from 14 percent for stock 10 to 64 percent for stock 1. In this section, we examine empirically the influence of the specialist's perception of information asymmetry on his participation rate across securities. The relationship is unclear from a theoretical viewpoint. If the specialist, seeking to avoid losses to traders with private information, participates more actively in stocks where his information disadvantage is smaller there will be a positive relationship between π and r across securities. Alternatively, if the specialist provides a service of price stabilization by 'leaning into the wind' to absorb order imbalances, he will be more active in thinner markets and there will be a negative relationship between r and π.

Since r is constrained to lie in the interval $[0, 1]$, we used a logit model given by:

$$r_i = \frac{1}{1 + e^{-\beta_0 - \beta_1 \pi_i}} \tag{29}$$

where β_0 and β_1 are the coefficients and π_i is the value of π for stock i from the estimated time-series model. The estimated coefficients are:

β_0	β_1
1.299	−2.834
(2.34)	(−3.66)

[30]From table 2, half the stocks had R^2s above 0.50, the lowest being 0.24. The results for October are comparable. This fit is relatively good considering the dependent variable is the transaction-to-transaction price change.

5.1 Block Trades

The estimates of λ, although significant and correct in sign, are lower than expected.

The institutional structure of the NYSE provides one possible explanation. The NYSE evaluates specialists according to certain criteria that are primarily related to the degree of price stability they provide by absorbing temporary order imbalances. The New York Stock Exchange's Rule 104 requires that the specialist maintain a 'fair and orderly market' as part of his affirmative obligation to provide liquidity. This term has never been explicitly defined by the NYSE, but in practice, the NYSE provides specialists with price continuity-depth guidelines. These guidelines place limits on the transaction to transaction price change permissible for a given volume of trade. For example, for a stock trading between \$20 and \$29 $\frac{7}{8}$ with average daily share volume in the previous month (excluding trades of 25,000 or more) of 10,000–24,999, the maximum price change for 3,000 share volume is \$ $\frac{7}{8}$.[32] The NYSE reports that in 1988, 92.1% of all transactions of 1,000 shares or less traded with a price change of 0 or $\frac{1}{8}$ from the immediately preceding trade.

These guidelines are not strictly binding and there exist provisions to allow more rapid price movements in times of market stress. Nevertheless, the NYSE market surveillance unit evaluates specialists in real time according to these criteria, and specialists who fail to comply with the continuity-depth guidelines risk having their stocks reassigned to others or not being assigned more profitable stocks in the future. The effect of these formal and informal price continuity requirements is to limit the price impact for a given volume of trade. This implies lower values of λ than theoretically predicted for all order size ranges.

An alternative theory focuses on the trading mechanism in use for large blocks of stock. In 1988, 55% of all shares traded on the NYSE were traded in blocks of 10,000 shares or more, most of this accounted for by institutions.[33]

[32] See, e.g., Floor Official Manual, Market Surveillance, New York Stock Exchange, June 1989.

[33] Traditionally, a block trade is defined as a trade of 10,000 shares or more. This definition is not meaningful from an economic viewpoint since in active stocks, an order of this magnitude may be quite

block buys from block sells. The existence of a reputation based alternative to the dealer market implies that dealers will refuse to make markets for block trades because the fact that a trader wishes to trade a block anonymously means that he or she could not obtain a favorable execution price upstairs. Thus, the existence of a reputation based block trading mechanism worsens the lemons problem in the downstairs market, implying that specialists will place absolute limits on the size of trades choose to accommodate.

The institutional features of block trading imply the price change for a large trade predicted by our model overstates the true impact of the trade. Large trades do have a greater price impact than small trades, but the relationship is not proportional to volume; on a per share basis, the price impact of a block trade may be less than that of a single round lot. By contrast, the effect of price continuity restrictions is to reduce λ uniformly for all trade sizes. To correct for the effect of these factors on our coefficient estimates, we extended the econometric model to allow the price change to be a piecewise linear function of order quantity. This procedure allows us to partly distinguish the effect of these two institutional features of the market.

Formally, we estimated the following MA(1) model based on (27):

$$\Delta p_t = \beta_0 + \beta_1 q_t + \sum_{h=1} \delta_h \chi_h (q_t - q_h) + \beta_2 I_t + \beta_3 I_{t-1} + \beta_4 D_t + \beta_5 D_{t-1} + \eta_t \qquad (30)$$

$$\eta_t = \tilde{\varepsilon}_t - \zeta \tilde{\varepsilon}_{t-1} + \tilde{u}_t \qquad (31)$$

where $(\delta_1, \delta_2, \delta_3, \delta_4)$ are coefficients to be estimated and $\chi_1, \ldots, \chi_4$ are dummy variables defined by $\chi_h = 1$ if $q_t \in A_h$ and 0 otherwise. The sets A_h are defined as $A_1 = (-\infty, \bar{q}_1]$, $A_2 = (-\infty, \bar{q}_2]$, $A_3 = [\bar{q}_3, \infty)$ and $A_4 = [\bar{q}_4, \infty)$, where $\bar{q}_1 < \bar{q}_2 < 0 < \bar{q}_3 < \bar{q}_4$ are stock specific definitions of 'large' trades, based on the observed daily trading volume.[36] Table 4

[36]The procedure is as follows: for each stock, we computed the distribution of order sizes after excluding the smallest 50% of transaction sizes. The cutoffs are the 1st, 5th, 95th and 99th percentiles of this distribution. By eliminating the smallest transactions, the natural breaks for large blocks are more readily identified. This procedure also lets the definition of a block trade vary by security, and since the distribution of order size is right-skewed, our approach also corrects for differences between block purchases and sales.

positive for 14 of the 16 stocks. This suggests that large block buys do have a price impact, but since the significance levels are weak, this result should be viewed with caution. Along these lines, note that the value of δ_2 is generally greater than the value of δ_3, implying that a block buy has a greater price impact than the corresponding block sale. This occurs even though the cutoffs in shares for block buys are smaller than the cutoffs in shares for block sells.

For most stocks, a large block trade has almost the same price impact as a mid-sized trade, and the impact is much lower on a per share basis. Consider, for example, the price impact of a 5,000 share sell order in stock 15. From table 5, the price impact of this trade is -\$0.0273, i.e., a drop of 2.73 cents. For a 10,000 share sell order, the price impact is -\$0.0369 while for a 100,000 share sell, the impact is -\$0.0828. While the absolute price impact is increasing in trade size, it is decreasing on a per share basis. Another application of the model is to assess the implicit costs of trading in the dealer market. We address this issue below, where we define a measure of the implicit spread that takes into account the price impact of the trade.[38]

5.2 Implied Bid-Ask Spreads

We define by $s(q)$ the *effective bid-ask spread* for an order of size q. Formally, the effective spread is the difference between the price if the order were to buy and the price if the order were to sell, i.e., $s(q) = p(|q|) - p(-|q|)$. For $q \in (\bar{q}_1, \bar{q}_2)$, equation (11) yields:

$$s(q) = 2(\psi + \lambda|q|) \,. \tag{32}$$

This definition implies the bid-ask spread is not a constant, but varies by the size of the order.[39]

[38] We also estimated the piecewise linear model using a break points that provide a narrower definition of a block, and obtained similar results. Presumably, stronger results could be obtained if a better procedure for approximating the break points could be found. This is a topic for further research.

[39] Glosten and Harris (1988) make this point in their concluding remarks where they state: "Since an important part of liquidity is the ability to make large trades without affecting price, price- liquidity studies

but may understate the implicit transaction costs of large investors created by their market impact and conversely overstate the costs of trading for small orders.

6 Conclusion

In this paper, we developed and estimated a model of intraday price formation. The model incorporates the effects of transaction costs, specialist inventory control, and asymmetric information, and encompasses a number of recent microstructure models. We test the model using data obtained from a NYSE specialist firm. These data, not previously available to empirical researchers, are of extremely high quality. We combined these data with data obtained from the Institute for the Study of Securities Markets (ISSM) to create a complete time-series of transactions and associated specialist inventories.

The estimated parameters of the model provide a natural measure of the degree of information asymmetry present in the market, the weight on the information content of trade placed by a representative market maker who uses Bayes' rule to update his beliefs. The estimated model provides strong support for the existence of information asymmetries, as perceived by the market maker. This perception suggests that transitory order imbalances can cause relatively large intraday price movements because they are viewed as potentially originating from informed traders, providing a possible explanation for higher price volatility during trading hours. The results suggest that perceived information asymmetries were significantly higher after the crash because either the ratio of public to private information decreased or the amount of noise trading fell. In addition to significant information effects, there is strong evidence of a fixed spread element in intradaily price changes. However, the inventory effect does not appear to be strong, possibly because of multicollinearity problems. Overall, the model performs very well in explaining intraday price changes.

We show that the specialist participation rate is directly related to the degree of information asymmetry, possibly because there is less competition to the specialist in markets where

References

Admati, Anat R. and Paul Pfleiderer, 1989, Divide and conquer: a theory of intraday and day-of-the-week mean effects, *Review of Financial Studies*, 2, 189–223.

Amihud, Yakov and Haim Mendelson, 1980, Dealership market: market making with inventory, *Journal of Financial Economics*, 8, 31–53.

Amihud, Yakov and Haim Mendelson, 1987, Trading mechanisms and stock returns: an emprical investigation, *Journal of Finance*, 42, 533–553.

DeGroot, Morris H., 1970, *Optimal statistical decisions* (McGraw-Hill Book Company).

Easley, David and Maureen O'Hara, 1987, Price, quantity, and information in securities markets,, *Journal of Financial Economics*, 19, 69–90.

Garman, Mark, 1976, Market microstructure, *Journal of Financial Economics*, 3, 257–275.

George, Thomas and Francis Longstaff, 1990, The components of the bid-ask spread: direct evidence from the index options markets,, Working paper, Ohio State University.

Glosten, Lawrence, 1987, Components of the bid-ask spread and the statistical properties of transaction prices, *Journal of Finance*, 42, 1293–1307.

Glosten, Lawrence and Lawrence Harris, 1988, Estimating the components of the bid-ask spread, *Journal of Financial Economics*, 21, 123–142.

Glosten, Lawrence and Paul Milgrom, 1985, Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents, *Journal of Financial Economics*, 14, 71–100.

Harris, Lawrence, 1990, Statistical properties of the Roll serial covariance bid-ask spread estimator, *Journal of Finance*, 45, 579–590.

Hasbrouck, Joel, 1988, Order arrival, quote behavior, and the return generating process, *Journal of Financial Economics*, 22, 229–252.

Hasbrouck, Joel, 1990a, Assessing the quality of a market, Working paper, New York University.

Hasbrouck, Joel, 1990b, Measuring the information content of stock trades, Forthcoming, *Journal of Finance.*

Hasbrouck, Joel, 1990c, Stock trades and informational asymmetries: an econometric analysis, Working paper, New York University.

Hausman, Jerry, Andrew Lo, and A. Craig MacKinlay, 1990, An ordered probit analysis of transaction stock prices, Working paper, MIT.

Ho, Thomas and R. Macris, 1984, Dealer bid-ask quotes and transaction prices: an empirical study of some AMEX options, *Journal of Finance*, 39, 23–45.

Ho, Thomas and Hans Stoll, 1983, The dynamics of dealer markets under competition, *Journal of Finance*, 38, 1053–1074.

Holthausen, Robert, Richard Leftwich, and David Mayers, 1987, The effect of large block transactions on security prices: a cross-sectional analysis, *Journal of Financial Economics*, 19, 237–267.

Appendix

Derivation of the Demand Function:

We show here that the demand function (5) used to construct the price functional (11) can be derived from a mean-variance utility maximization problem. Let $p(q)$ denote the quoted price of the risky security as a function of order size. Using equation (11), the transaction price can be expressed in the form:

$$p(q) = \xi_0 + \psi \text{sign}(q) + \lambda q \tag{A.1}$$

where ξ_0 represents the lagged price plus the inventory effects which is regarded as a constant by the trader. Suppose the trader has a mean-variance utility function of the form:

$$u(\widetilde{W}) = E[\widetilde{W}] - \left(\frac{\rho}{2}\right) \text{Var}[\widetilde{W}] \tag{A.2}$$

where $\rho > 0$ is the coefficient of absolute risk aversion and $\widetilde{W}$ is the trader's risky wealth. Wealth is given by:

$$\widetilde{W} = \tilde{v}(q + X) - p(q)q + C \tag{A.3}$$

where X represents the trader's initial endowment of the risky asset and C the initial holdings of cash. Using (A.2), investor's quantity q solves:

$$\max_q \left\{ m(q + X) - p(q)q + C - \frac{\rho\sigma^2}{2}(q + X)^2 \right\} \tag{A.4}$$

where m and σ^2 are (respectively) the expectation and variance of the asset's value given the trader's information. If there is an interior solution where the trader chooses to trade (the presence of a fixed cost element implies that for some traders the optimal action is not to trade), the first order conditions yield:

$$q = \frac{m - \xi_0 - \psi \text{sign}(q) - \rho\sigma^2 X}{\rho\sigma^2 + 2\lambda} \tag{A.5}$$

Note that the optimal demand is a function of the slope and intercept of the price functional. Since $\text{sign}(q) = \text{sign}[m - \xi_0 - \psi \text{sign}(q) - \rho\sigma^2 X]$, it follows that $q = 0$ and there is no trade if:

$$|m - \xi_0 - \rho\sigma^2 X| < \psi. \tag{A.6}$$

If there is an interior solution, the transaction price is $p(q) = \xi_0 + \psi \text{sign}(q) + \lambda q$ and substituting this into equation (A.5) and rearranging, we obtain:

$$q = \frac{m - p - \rho\sigma^2 X}{\rho\sigma^2 + \lambda}. \tag{A.7}$$

Define $x \equiv \alpha\rho\sigma^2 X$ where $\alpha \equiv 1/(\rho\sigma^2 + \lambda)$. Then, we can write (A.7) as: $q = \alpha(m - p) - x$, as conjectured in (5). This demand function in turn supports the price equation (11) used to derive it initially.

TABLE 2

Reduced-Form Coefficient Estimates February–December, 1987

This table reports the estimated coefficients (with t-values in parentheses) for the reduced-form model for the period February–December, 1987, using Box-Jenkins methods. The model to be estimated is:

$$\Delta p_t = \beta_0 + \beta_1 q_t + \beta_2 I_t + \beta_3 I_{t-1} + \beta_4 D_t + \beta_5 D_{t-1} + \tilde{\varepsilon}_t - \zeta \tilde{\varepsilon}_{t-1} + \tilde{u}_t.$$

Inventory and volume are scaled by 10^{-6}, and the coefficients β_1, β_2, and β_3 should be interpreted accordingly.

Stock	β_0	β_1	β_2	β_3	β_4	β_5	ζ	Number of Observations	R^2
1	0.000 (0.07)	32.873 (4.43)	7.830 (1.26)	−8.075 (−1.30)	0.132 (26.58)	−0.062 (−14.38)	−0.026 (−0.70)	770	0.654
2	−0.016 (−0.74)	0.275 (0.29)	−0.019 (−0.03)	0.108 (0.20)	0.103 (32.66)	−0.063 (−20.80)	0.046 (1.58)	1,179	0.525
3	0.004 (0.75)	2.650 (5.44)	−2.464 (−2.17)	2.408 (2.12)	0.116 (46.32)	−0.065 (−26.37)	0.095 (4.92)	2,782	0.487
4	−0.006 (−2.10)	0.039 (0.05)	2.592 (0.95)	−2.559 (−0.93)	0.120 (33.41)	−0.063 (−17.77)	0.165 (6.34)	1,484	0.463
5	−0.006 (−3.51)	1.171 (5.31)	0.864 (0.98)	−0.636 (−0.72)	0.089 (59.40)	−0.059 (−39.42)	0.130 (7.51)	3,461	0.568
6	−0.007 (−2.55)	1.211 (12.37)	−0.127 (−0.60)	0.162 (0.77)	0.072 (79.17)	−0.059 (−66.39)	0.200 (14.68)	5,260	0.627
7	0.005 (0.92)	4.152 (5.23)	−1.504 (−1.27)	1.420 (1.20)	0.229 (63.64)	−0.144 (−40.92)	0.244 (18.38)	5,537	0.469
8	−0.010 (−3.31)	3.394 (5.13)	−1.125 (−0.54)	1.355 (0.65)	0.261 (76.83)	−0.172 (−51.20)	0.319 (28.29)	7,137	0.482
9	−0.010 (−2.15)	3.278 (5.14)	0.690 (0.55)	−0.670 (−0.53)	0.126 (75.81)	−0.089 (−54.91)	0.299 (27.16)	7,616	0.493
10	0.004 (4.73)	0.852 (4.60)	−0.513 (−1.08)	0.512 (1.08)	0.070 (42.46)	−0.058 (−36.77)	0.515 (52.72)	7,755	0.241
11	0.001 (0.96)	0.710 (3.99)	0.658 (1.15)	−0.665 (−1.16)	0.083 (96.05)	−0.068 (−79.63)	0.289 (26.73)	7,925	0.601
12	0.004 (2.40)	1.360 (7.49)	−0.714 (−1.44)	0.711 (1.44)	0.139 (89.28)	−0.093 (−60.25)	0.211 (20.55)	9,190	0.513
13	−0.003 (−3.60)	2.867 (6.01)	−0.864 (−0.91)	0.874 (0.92)	0.099 (116.41)	−0.085 (−104.65)	0.465 (52.05)	9,983	0.660
14	−0.001 (−0.42)	0.205 (2.57)	−1.640 (−2.45)	1.627 (2.43)	0.120 (124.25)	−0.103 (−107.64)	0.410 (47.82)	11,441	0.641
15	−0.001 (−0.60)	0.898 (4.10)	−1.008 (−2.11)	1.014 (2.12)	0.145 (91.41)	−0.102 (−65.06)	0.487 (78.72)	20,039	0.319
16	−0.001 (−4.71)	0.402 (4.74)	−1.031 (−3.29)	1.051 (3.35)	0.085 (42.89)	−0.076 (−38.54)	0.877 (303.92)	27,553	0.441

TABLE 4

Break Points

This table reports the break points (in shares) for the piecewise linear regression presented in tables 5 and 6. The break points are selected from the 1st, 5th, 95th and 99th percentiles of the sample distribution for order size after excluding the smallest 50% of transactions.

Stock	$\bar{q}_1$	$\bar{q}_2$	$\bar{q}_3$	$\bar{q}_4$
1	−2,931	−1,400	1,295	2,459
2	−16,802	−5,000	2,500	10,000
3	−17,480	−4,860	5,000	19,236
4	−13,920	−3,200	4,890	20,000
5	−25,000	−7,260	6,500	30,000
6	−28,320	−8,800	6,500	17,260
7	−16,102	−4,695	3,000	9,406
8	−20,615	−7,500	7,150	17,940
9	−8,000	−2,000	2,000	5,931
10	−30,000	−13,000	16,000	42,988
11	−12,182	−3,200	3,800	15,052
12	−42,886	−9,000	8,000	25,000
13	−6,880	−2,400	2,000	4,860
14	−23,860	−3,700	3,600	20,000
15	−25,000	−6,000	5,000	19,126
16	−37,124	−6,300	5,955	29,655

Table 6

Effective Bid-Ask Spreads

This table reports the effective bid-ask spread estimates for the period February-December, 1987, and for October, 1987. Roll's effective spread is:

$$s_r = 2\sqrt{-\text{Cov}(\Delta p_t, \Delta p_{t-1})} \,.$$

The bid-ask spread is defined as:

$$s(q) = 2(\hat{\psi} + \hat{\lambda}|q|)$$

where q is the trade size and $\hat{\lambda}$ and $\hat{\psi}$ are the estimated coefficients of the reduced-form model. Here s_1 is the spread for a 100 share trade and $\bar{s}$ is the spread corresponding to average trade size. The average (weighted by number of transactions) is reported at the bottom of each column. Note: The serial price covariance is positive for stock 1 so Roll's formula inapplicable.

Stock	February-December, 1987			October, 1987		
	s_r	s_1	$\bar{s}$	s_r	s_1	$\bar{s}$
1	0.112	0.130	0.150	–	0.101	0.153
2	0.495	0.126	0.127	0.239	0.176	0.249
3	0.131	0.130	0.136	0.173	0.149	0.189
4	0.150	0.126	0.126	0.252	0.155	0.163
5	0.128	0.118	0.122	0.167	0.157	0.158
6	0.120	0.119	0.124	0.165	0.146	0.148
7	0.152	0.289	0.299	0.315	0.473	0.528
8	0.216	0.344	0.359	0.351	0.525	0.599
9	0.152	0.179	0.184	0.261	0.265	0.297
10	0.210	0.115	0.121	0.475	0.226	0.237
11	0.140	0.136	0.137	0.224	0.194	0.196
12	0.205	0.187	0.192	0.304	0.287	0.299
13	0.189	0.171	0.174	0.297	0.236	0.237
14	0.177	0.206	0.207	0.255	0.290	0.290
15	0.306	0.204	0.207	0.619	0.307	0.315
16	0.142	0.153	0.154	0.270	0.209	0.212
	0.191	0.179	0.183	0.381	0.284	0.300