THE JOURNAL OF FINANCE • VOL. XLVIII, NO. 5 • DECEMBER 1993

An Analysis of Changes in Specialist Inventories and Quotations

ANANTH MADHAVAN and SEYMOUR SMIDT*

ABSTRACT

We develop a dynamic model of market making incorporating inventory and information effects. The market maker is both a dealer and an investor, quoting prices that induce mean reversion in inventory toward targets determined by portfolio considerations. We test the model with inventory data from a New York Stock Exchange specialist. Specialist inventories exhibit slow mean reversion, with a half-life of over 49 days, suggesting weak inventory effects. However, after controlling for shifts in desired inventories, the half-life falls to 7.3 days. Further, quote revisions are negatively related to specialist trades and are positively related to the information conveyed by order imbalances.

MARKET MAKERS PLAY CRUCIAL roles in forming prices and in providing liquidity in securities markets. This paper analyzes, both theoretically and empirically, the trading behavior of a market maker (or specialist) on the New York Stock Exchange (NYSE).

Views about specialists' roles have changed dramatically over time, perhaps reflecting real changes in their functions or their environment, or a realization that they perform several complex functions. The specialist was initially described as an auctioneer or as "the broker's broker," responsible for maintaining the limit order book and enforcing price and time priority rules for order execution. The specialist is also a "supplier of immediacy," who provides liquidity by acting as a dealer. Early analyses (e.g., Demsetz (1968)) presume that specialists would passively provide liquidity to accommodate transitory order imbalances, thereby stabilizing prices. To perform their dealership function, however, specialists must bear unwanted inventories. Theoretical models show that when market makers face inventory carrying

*Wharton School, University of Pennsylvania and Johnson School, Cornell University, respectively. We thank Marshall Blume, Fernando Diz, Prajit Dutta, Thomas Finucane, Bruce Grundy, Lawrence Harris, Joel Hasbrouck, Richard Highfield, Eric Hughson, Shmuel Kandel, Donald Keim, Craig MacKinlay, Krishna Ramaswamy, Matthew Richardson, Mark Roomans, Oded Sarig, Laura Starks, René Stulz (editor), and an anonymous referee for their helpful comments. We have also benefitted from the comments of seminar participants at CalTech, Princeton, Syracuse, the Cornell-Rochester conference, the 1992 USC/UCLA/NYSE Conference on Market Microstructure, the 1992 Western Finance Association Meetings, and the 1992 European Finance Association Meetings. Research support from the Geewax-Terker Research Fund (Madhavan) is gratefully acknowledged. Any errors are entirely our own.

costs or are risk averse they will actively control their inventory by setting prices to induce movements towards desired inventory levels.[1]

Another class of models emphasizes the importance of asymmetric information in analyzing market maker behavior.[2] In these models, the perceived presence of informed traders with private information regarding fundamental asset values affects price dynamics and the size of the bid-ask spread. The inventory control and asymmetric information theories are not mutually exclusive; they yield similar predictions about asset returns and volume, but have not yet been integrated into a formal model with optimizing agents. This paper contributes to the theoretical literature by providing such a model.

Empirical analyses of market maker behavior are mainly limited to indirect implications of these theories because the available public databases do not distinguish market maker transactions from those of the other market participants. Very few studies are based on the actual trading records of market makers in any financial markets.[3] Studies of market makers' daily inventory positions (U.S. Securities and Exchange Commission (1971), Stoll (1976), Ho and Stoll (1983), and Ho and Macris (1984)) find evidence of inventory effects, but do not distinguish these effects from the effects associated with asymmetric information. More recent studies that recognize both effects (see, e.g., Hasbrouck (1988), Stoll (1989), Madhavan and Smidt (1991), and Foster and Vishwanathan (1993); of these studies, only Madhavan and Smidt use specialist inventory data) find only weak evidence of short-run inventory effects but strong information effects. These results argue for examining inventory behavior over longer horizons.

Though those studies shed light on the relative importance of the information and inventory effects, several crucial questions concerning market maker trading behavior remain unresolved: If market makers pursue inventory control policies, why do previous empirical studies find inventory effects to be so weak? Are market makers at an informational disadvantage relative to other traders, or do they possess valuable information about market conditions? Do market makers take speculative positions, and if so, how do these positions affect short- and long-run return dynamics? This paper contributes to the empirical literature by providing evidence on these issues.

We develop an intertemporal model of market maker trades and quotes. The model differs from previous theoretical models in four important respects:

1. The model incorporates formally the effects of both asymmetric information and inventory control, where the dynamic strategic behavior of a

[1] See, e.g., Amihud and Mendelson (1980) and Ho and Stoll (1983).

[2] See, e.g., Kyle (1985) and Glosten and Milgrom (1985).

[3] See Working (1977a, 1977b), U.S. Securities and Exchange Commission (1971), Stoll (1976), Ho and Macris (1984), Silber (1984), Madhavan and Smidt (1991), Neuberger (1992), and Hasbrouck and Sofianos (1993).

market maker and a representative informed trader are explicitly characterized.
2. The market maker acts as both a dealer who provides liquidity on demand and an active investor, trading for his own account. Previous models assume that market makers act as pure dealers, ignoring the possibility that their trades may also reflect investment and speculative motives.
3. The market maker's long-run optimal inventory is endogenously chosen, and may be periodically revised due to changes in the risk profile of the underlying stock.
4. The market maker's trades also reflect short-run speculative demands based on information on the impending order flow.

The model provides new insights into the role of the specialist. We show that a specialist acts both as a dealer and as an active investor managing his portfolio exposure. As an active investor, the specialist seeks to maintain a long-term position in the stock consistent with his portfolio objectives, while profiting in the short-term from information about impending order imbalances obtained through his central position on the trading floor. As a dealer, the specialist's optimal quotes induce mean reversion in inventory towards the long-term target inventory level. We show theoretically that the target inventory may shift over time in response to changes in the risk profile of the stock.

We find very slow mean reversion in inventories assuming the specialist's desired stock holdings are constant. It takes over 49 trading days, on average, for an imbalance in inventory to be reduced by 50 percent. The slow adjustment process is consistent with the weak inventory effects documented by Madhavan and Smidt (1991) using intraday data. However, our theoretical analysis suggests that periodic shifts in desired inventories may bias our estimates of mean reversion. We develop an intervention model to correct for unobserved shifts in the specialist's desired inventory levels, and find strong evidence of mean reversion in inventories to these time-varying targets. Even so, mean reversion takes place over far longer horizons than was previously believed; it takes on average 7.3 trading days for an imbalance in inventory to be reduced by 50 percent.

The model predicts that quote revisions are inversely related to specialist trades and positively related to the information content of order flow. We find strong support for this hypothesis. Interestingly, it is the nonblock portion of order flow that appears to have information content. We also find evidence suggesting that specialist quotes anticipate future order imbalances, a result consistent with the specialist receiving information about the impending order flow.

The paper proceeds as follows. In Section I, we develop a dynamic model where agents act strategically. In Section II, we describe the data and our estimation technique. Section III provides the results of our analyses of specialist inventories, while Section IV examines the relation between specialist trades and stock prices. Finally, Section V summarizes the paper.

I. A Framework for Analysis

A. The Trading Environment

Consider a market where a risky asset trades in a series of auctions on days $t = 1, 2, \ldots, \tau$. Let v_t denote the value of the risky security on day t. In addition, a riskless (numeraire) asset also trades. On the final trading date, τ, the risky security pays a liquidating dividend, v_τ. The security's fundamental value follows a random walk until date τ, i.e., $v_t = v_{t-1} + \eta_t$ for $t \leq \tau$, where η_t is a normally distributed error term with mean zero and variance σ_η^2. The asset's underlying value, v_t, is not public information, and is viewed by uninformed agents as a random variable, denoted by $\tilde{v}_t$. The dividend date, τ, is a random variable, and we assume that at the beginning of day t, before any trading occurs, there is a constant probability $(1 - \rho)$ that day t will be the liquidation date, i.e., $\Pr[t = \tau] = (1 - \rho)$.[4]

At the start of each auction, traders submit their demands to floor brokers specifying what quantities are to be traded at what prices. Equivalently, traders may submit price-contingent orders to the specialist through an electronic trading system. The exact form of order entry is not relevant, as long as traders can submit price-contingent orders. If, with probability $(1 - \rho)$, day t is the liquidation date, a trading halt is called, all orders are canceled, and the price immediately is revised to the announced value of the asset, v_t. Otherwise, with probability ρ, the event does not occur on day t, and trading occurs at the price set by a risk-neutral market maker or specialist, who absorbs any excess demand into inventory after observing the aggregate excess demand schedule. The specialist bears inventory carrying costs so that he will not hold infinite positions. Note that the specialist need not participate in every transaction; some orders will cross at the posted price.

Let $z_t(p_t)$ denote the excess demand at price p_t with the convention that $z_t > 0$ represents positive excess demand and $z_t < 0$ positive excess supply. Excess demand originates from uninformed liquidity-motivated traders and a representative risk-neutral informed trader.[5] The assumption of a single informed trader is made purely for expositional ease; the extension to multiple informed traders is straightforward and does not alter our results. Let Q_t denote the order quantity of the informed trader, and let X_t denote the aggregate trade of the liquidity-motivated traders, where we assume that the unconditional mean of X_t is zero. Thus, $z_t = Q_t + X_t$.

[4]This distributional assumption allows us to express the agents' dynamic maximization problem as one with an infinite horizon since it implies that the expected time to the event is a constant, independent of the past history.

[5]The assumption of risk neutrality considerably simplifies the derivation of the dynamic order placement strategy.

The informed trader observes the current realization of v_t, which is the payoff that would occur if t is the liquidation date, before submitting her order for auction t. We can extend the model to the case where the informed trader obtains a noisy estimate of v_t. The informed trader takes into account the effect of the current trade on future quoted prices and acts strategically given rational conjectures about the specialist's quote-setting behavior. In turn, the specialist's optimal behavior depends on conjectures about the dynamic trading strategy of the informed trader. In equilibrium, the conjectures of both agents regarding their dynamic strategies are consistent with each other.

The informed trader's objective is to maximize expected final period wealth. The informed trader's optimal order quantity is assumed to be proportional to the deviation between the asset's current fundamental value and the price set by the specialist. Formally,

$$Q_t(v_t, p_t) = \delta(v_t - p_t), \tag{1}$$

where $\delta > 0$ is a constant which represents the slope of the trader's demand function. In the Appendix, we examine conditions under which equation (1) represents the informed trader's optimal dynamic strategy, given rational expectations about the specialist's behavior. Intuitively, the informed trader recognizes that the auction price is an increasing function of the order quantity; at the margin, the informed trader balances the profit from a marginal increase in the quantity traded against the price impact on the entire amount of the trade. This tradeoff ensures that the endogenously determined demand coefficient, δ, is finite.

The specialist observes the aggregate excess demand function before setting the price quote. Using equation (1), the excess demand function on day t can be expressed in the form

$$z_t(p_t) = Q_t(v_t, p_t) + X_t = D_t - \delta p_t \tag{2}$$

where $D_t \equiv \delta v_t + X_t$ is the intercept of the demand schedule. Irrespective of whether orders are submitted verbally (as in a call auction) or electronically (through an automated order routing system), the specialist observes the excess demand schedule before setting the price. The ability to condition on the aggregate orders submitted before setting the price provides the specialist with a noisy signal about the unobserved asset value. Let $\mu_t = E[\tilde{v}_t | \Phi_t^s]$ represent the market maker's conditional expectation about v_t given his information set, Φ_t^s. The specialist's information set consists of the history of trading and information obtained from observing the excess demand schedule.

Let I_t denote the specialist's inventory before trading takes place on day t, so that $I_t = I_{t-1} - z_{t-1}$. The specialist's wealth at the beginning of period t consists of the (stochastic) value of his opening inventory, $\tilde{v}_t I_t$, plus the (stochastic) value of the specialist's holdings in other assets, denoted by $\tilde{Y}_t$. In turn, the value of the specialist's outside assets at the start of period t

consists of the opening capital, K_t, plus the (stochastic) income from other assets, denoted by $\tilde{y}_t$, i.e., $\tilde{Y}_t = K_t + \tilde{y}_t$. Thus, wealth is

$$\tilde{W}_t = \tilde{v}_t I_t + K_t + \tilde{y}_t. \tag{3}$$

At the start of period t, the specialist bears inventory carrying costs, denoted by c_t. These costs are assumed to be proportional to the variance of the specialist's wealth at the start of period t (or wealth after trading in period $t-1$):

$$c_t = \omega\sigma^2\left[\tilde{W}_t\right], \tag{4}$$

where $\omega > 0$ is a constant representing the inventory cost per unit variance. Risk may impose costs on a specialist firm in the form of greater capital requirements to avoid cash constraints. Equivalently, we could interpret these costs as arising from risk aversion on the part of the specialist. By convention, carrying costs, c_t, are accrued at the start of period t (including date τ), prior to any trading activity.

Then, the specialist's capital evolves as follows:

$$K_{t+1} = K_t - c_t + p_t z_t + y_t, \tag{5}$$

i.e., it consists of the previous period's capital, less inventory carrying costs, plus any trading revenues, plus the income from other financial assets. We assume that $\tilde{y}_t$ is distributed identically with mean normalized to zero; this income is realized in period t irrespective of any trading activity. The asset's return may be correlated with the specialist's outside income, and we denote the covariance between these random variables by $\mathrm{Cov}[\tilde{y}_t, \tilde{v}_t] \equiv \sigma_{vy}$.

The specialist is risk neutral and is assumed to maximize the expected final period wealth, denoted by $\tilde{W}_\tau$, subject to the transitional relations governing the evolution of the state variables summarizing the current environment. The maximization decision is modeled in the Appendix as a stochastic dynamic programming maximization problem whose solution yields a policy function giving the optimal price quote and the optimal specialist trade.

B. Inventory Dynamics

In the Appendix, we show that the specialist's optimal dynamic trades in equilibrium are given by

$$I_{t+1} - I_t = \beta(I_t - I^d) - \frac{1+\beta}{2}x_t, \tag{6}$$

where $-1 < \beta < 0$ and I^d are constants, and x_t is an unbiased estimate of X_t given the information available to the specialist.

Equation (6) illustrates the dual nature of specialist trading behavior. The specialist's dealership role is reflected in the first term on the right-hand side of equation (6), which implies that inventories exhibit mean reversion towards I^d, which is interpreted as the specialist's long-term desired or target inventory. In our model, $I^d = -\sigma_{vy}\sigma_v^{-2}$, i.e., it is interpreted as a hedge ratio, and reflects the specialist's role as an active investor. Intuitively, the specialist seeks to maintain a long-term inventory position to reduce the risk induced by the correlation between the stock's return and shocks to the specialist's other income. The target inventory is unlikely to be constant over long intervals. In the Appendix, we demonstrate that periodic changes in the risk characteristics of the stock, possibly the result of changes in the composition of the underlying firm's assets or liabilities, will produce periodic endogenous shifts in target inventories. The form of the optimal policy function in this more complicated case is, however, a straightforward generalization of equation (6).

From equation (6), it is clear that the speed of inventory adjustment is related to β; lower values of β imply more rapid adjustments to the desired inventory level. From the solution to the specialist's optimization problem in the Appendix, the speed of adjustment is a decreasing function of the slope of the excess demand function, δ, the specialist's inventory costs, ω, and the conditional asset variance, σ_v^2. The greater the demand responsiveness, the disutility of carrying unwanted inventory, or the uncertainty about asset value, the greater the incentives to control inventories and the stronger the mean reversion. In the special case where ω is near zero, β is near zero and inventory changes approximate a random walk. Intuitively, if there is no disutility from carrying excess inventory, the market maker simply sets price equal to the expected value of the security and inventory fluctuates in response to excess demand.

The coefficient of x_t in equation (6) illustrates the short-horizon investment strategy of the specialist. This term represents the expected change in the specialist's inventory position that results from information about impending order imbalances. As the coefficient of x_t lies between $-\frac{1}{2}$ and 0, the specialist is willing to accommodate less than half the anticipated noise shock. Lower values of β imply the specialist is less willing to stabilize temporary order flow shocks.

C. Information and Learning

The dynamic behavior of inventories is mirrored in the behavior of prices since the market maker uses price as a control to move towards desired inventory levels. In the Appendix, we show that

$$p_t = \mu_t - \zeta_1(I_t - I^d) + \zeta_2 x_t, \tag{7}$$

where $\zeta_1 = -\beta/\delta > 0$ represents the inventory effect and $\zeta_2 = (1-\beta)/2\delta > 0$ represents the effect of order imbalances on price. Equation (7) shows that the price set by the specialist is the expected value of the security less a term proportional to the deviation between actual and target inventory, plus a term proportional to the anticipated order imbalance. The specialist lowers price when inventories are high relative to the target to induce traders to buy the security, thereby reducing future inventory carrying costs. Similarly, the specialist raises prices when faced with a positive anticipated imbalance since this increases trading profits.

To render equation (7) into a model suitable for empirical estimation, we must specify how the specialist's beliefs are formed. The specialist's inference problem consists of updating prior beliefs given the signals conveyed by order flow. However, as the unobservable state changes through time, the specialist's posterior beliefs do not "converge" to the fundamental asset value, but instead form a distribution centered about the conditional expected value of the asset. Under rational expectations, the conditional expectation at any point in time is an unbiased forecast of asset value.

We show in the Appendix that the market maker's inference problem can be modeled within a state-space framework using the Kalman filter algorithm. The solution yields the following model for quote revisions:

$$p_t - p_{t-1} = \gamma_1 s_t - \zeta_1(I_t - I_{t-1}) + \zeta_2(x_t - x_{t-1}). \tag{8}$$

In equation (8), s_t is the difference between the realized excess demand, z_t, and the predicted informed demand, and γ_1 is a positive constant. Thus, the revision in price is a linear function of the unanticipated component of order flow, the change in inventories, and the change in anticipated order imbalances. In equation (8), $\gamma_1 = \Omega/(1-\Omega)\delta > 0$, where $\Omega \in (0, 1)$ is the weight placed on the signal content of order flow. This constant represents the information effect on prices. If order flow is informative, Ω is near unity, and the information parameter, γ_1, is large. In summary, the specialist's trading behavior can be analyzed in two ways: by examining whether inventories exhibit mean reversion as suggested by equation (6), or by assessing the impact of specialist trades on prices, as suggested by equation (8). While both approaches represent the same equilibrium model, they differ from an econometric perspective, and our empirical analyses use both approaches.

II. Data and Empirical Methods

A. Data Sources

The data used in this paper are drawn from two sources: (1) a file covering all transactions of a New York Stock Exchange (NYSE) specialist firm in its 16 assigned stocks, from February 1, 1987 to December 31, 1987, and (2) a file of the bid, ask, and transaction prices and volumes of the specialist's stocks in all domestic markets, obtained from the Institute for the Study of Securities Markets (ISSM). Together these files enable us to compile a complete record of specialist and nonspecialist trading activity in the 16 stocks over the sample period.

The specialist data set consists of trading records, which are analogous to invoices, and settlement records, which are analogous to canceled checks.

Since the settlement records are based on corrected trades and represent actual cash flows, they are the most accurate representation of the specialist's trading activity, and we use these records in our analyses. The specialist settlement records typically contain multiple information on fields of interest to the specialist. For example, trade size is signed, but there is also a separate buy-sell indicator. These redundancies permit extensive cross-checks of the accuracy of the records. In total, there are almost 75,000 specialist transactions in the sample period.

The data on nonspecialist transactions and quotes used in this study are obtained from the ISSM transaction files. The file contains transaction and quote data for NYSE and American Stock Exchange stocks, as well as transactions and quotes from other markets in the National Market System (NMS). Although the specialist data sign volume, the ISSM files do not indicate whether a transaction was buyer or seller initiated. These transactions are classified as buyer or seller initiated using a method developed by Lee and Ready (1991) that compares the transaction price with the prevailing bid and ask quotations. Using the signed volume, we compute the net order imbalance on a given day as the buyer-initiated share volume minus the seller-initiated share volume on that day.

The two data files provide a complete time series of inventories, quotes, and nonspecialist trading in the 16 stocks in the sample period. We form a time series of the specialist's inventory at the beginning of each trading day for the 16 stocks traded between February 1 and December 31, 1987 which have an average of at least four transactions per day on the NYSE. These data, used in our inventory analyses, cover the 232 trading days in the sample period. From the ISSM data, we form a time series of opening bid and ask quotes and daily order imbalances in the NMS. The ISSM file is missing data on 33 days (mostly in August) so our measure of order imbalance, used in our quote analyses, covers 199 trading days. The only exception is stock 4, which is allocated to the specialist unit for part of the sample period. For this stock alone, there are 99 inventory and order imbalance observations.

Table I reports descriptive data for the 16 stocks in our sample, ranked by the average number of transactions per day in all markets, with stock 1 being the least active. The statistics tabulated are averages of daily observations for each stock; the standard errors associated with these estimates are very small and are not reported here. The closing inventories reported are the absolute daily inventories; changes in closing inventories are the absolute difference between the inventory levels on a given day and the levels on the previous trading day.[6] Because of price changes, the change in inventory value can exceed the value of the specialist's trading on that day.

The specialist's closing inventories are substantial, averaging \$2 million per stock, but there is only a weak tendency for these closing inventory levels to be larger in more active stocks. Over all 16 stocks, the smallest average

[6] In the case of stock splits, data measured in shares are adjusted to the presplit levels.

Table I

Descriptive Statistics on Specialist and Market Transactions for 16 NYSE Stocks

The figures represent averages of daily values for the period February to December 1987 for the 16 stocks, ranked by transaction frequency from lowest to highest, in the NMS. Data on inventories and specialist trades are obtained from the settlement records of the specialist. Data on prices and volumes in the NMS are obtained from the ISSM.

	Closing Inventories		Change in Specialist's Inventories		Value of Transactions		Number of Transactions	
Stock Number	$1,000s	Round Lots	$1,000s	Round Lots	Specialist $1,000s	Market $1,000s	Specialist	Market
1	114	63	15	8.3	25	186	4	5
2	1,999	1,847	44	16.1	24	1,190	4	7
3	2,156	814	106	46.4	195	4,551	8	15
4	345	235	58	38.4	102	7,945	9	16
5	180	110	72	44.4	138	9,109	9	19
6	2,230	2,165	97	88.3	122	4,513	8	23
7	6,223	1,532	217	48.3	358	13,637	14	22
8	912	179	270	53.0	687	9,716	20	30
9	4,327	2,237	106	49.3	226	3,805	17	32
10	1,627	646	327	53.0	467	11,866	9	30
11	343	234	76	49.3	143	2,737	12	31
12	4,202	695	835	141.9	1,879	26,176	24	38
13	409	215	67	35.9	179	1,558	18	35
14	2,593	1,196	150	54.2	307	47,702	16	36
15	4,190	1,490	572	177.3	1,876	19,172	54	92
16	824	508	385	209.1	1,088	34,758	49	110
Totals	32,674	13,862			7,818	198,622	275	541

closing inventory was $114,000 for the least active stock while the largest average inventory was $6.2 million for stock 7. The average value of transactions per stock per day was $12.4 million.

Cross-sectionally, the average change in closing inventories in a stock is more closely correlated with trading activity in that stock than is the average level of closing inventories. The average change in closing inventories per stock ranges from $15,000 in stock 1 to $835,000 in stock 12. The average ratio of inventory change to value of trading is 2.6 percent. Changes in inventories understate the specialist's participation because he is typically buying and selling on the same day. We measure participation by comparing the specialist's trading (in either dollars or transactions) with the total trading in the stock. The participation ratio would be 100 percent if the specialist were the only party on the other side of every trade. The average dollar participation ratio is 5.1 percent but the average transaction participation ratio is 52.7 percent. The difference occurs because the specialist has only limited participation in large-block trades, which are a small fraction of all trades, but account for approximately 50 percent of the trading value.

B. Empirical Methods

We use Hansen's (1982) Generalized Method of Moments (GMM) to estimate the model. GMM is particularly appropriate here because the procedure is based on weak assumptions for the stochastic process generating the data, and provides a multivariate test suitable for cross-sectional data. To illustrate the GMM approach, note that equation (6) can be written as $u_t = I_t - I_{t-1} - \beta(I_{t-1} - I^d)$, where $u_t = -(1+\beta)x_{t-1}/2$ represents the error term from the viewpoint of the econometrician. If the model is correctly specified, u_t has a conditional mean of zero given information at time $t-1$, so that $E[u_t \mid Z_{t-1}] = 0$, where Z_{t-1} is an r-dimensional vector of instrumental variables. This condition implies that $E[u_t \otimes Z_{t-1}] = 0$. The GMM procedure consists of replacing this expectation with its sample analog, denoted by $g_T(\theta)$, where T is the number of observations and $\theta = (\beta, I^d)$ is the vector of unknown parameters of the model, and choosing parameter values for θ that minimize the quadratic criterion function

$$J_T(\theta) = g_T(\theta)' W_T(\theta) g_T(\theta), \tag{9}$$

where $W_T(\theta)$ is the weighting matrix. The estimator with the minimum variance-covariance matrix obtains from choosing the weighting matrix to be the inverse of the covariance matrix of the orthogonality conditions. Hansen (1982) proves that the minimized objective function, $Tg_T(\hat{\theta})' W_T(\hat{\theta}) g_T(\hat{\theta})$, is asymptotically χ^2 distributed with degrees of freedom equal to the number of overidentifying restrictions. Hansen proposes this as a test statistic for the goodness of fit of the model. A high χ^2 statistic is indicative of model misspecification since it implies the disturbance term is correlated with the instrumental variables. The asymptotic covariance matrix for the GMM estimator provides standard errors used to test separately the significance of the parameters.

Newey and West (1987) show that the weighting matrix can be adjusted to account for serial correlation and conditional heteroskedasticity. Multivariate hypothesis tests are straightforward to perform. In general, consider a null hypothesis, H_0: $\alpha(\theta) = 0$, where $\alpha(\theta)$ is a vector of model restrictions of dimension k, and let $J_T(\theta^*)$ denote the minimized objective function under the parameter restrictions implied by the null hypothesis. Similarly, let $J_T(\hat{\theta})$ denote the unrestricted minimum value of the objective function. Then, the test statistic

$$R = T\left[J_T(\theta^*) - J_T(\hat{\theta})\right] \tag{10}$$

is asymptotically distributed χ^2 with k degrees of freedom. This test statistic is used to assess the goodness of fit of the model.

III. Analysis of Inventories

A. Mean Reversion in Inventories

In this section, we test the model's hypotheses for specialist inventories; we examine the price effects associated with these inventory movements in Section IV. Table II presents the GMM estimates of equation (6), with t-statistics based on Newey-West autocorrelation-heteroskedasticity-consistent standard errors.[7] For all 16 stocks, the estimate of β is negative, and is significantly negative for half the stocks. The average value of β is about -0.05. These estimates are not consistent with inventory effects taking place within the day since this would imply that $\beta = -1$, i.e., that inventory is stationary around the target level, I^d. The specialist's desired inventory is positive for all 16 stocks and is statistically significant for 12 stocks. The estimates of I^d vary widely, ranging from 248 shares to 790,941 shares. Table II also reports for the minimized value of the GMM criterion function for the individual stocks. This statistic, as noted above, provides a test of the model's overidentifying restrictions and is asymptotically distributed as a χ^2 with three degrees of freedom. The χ^2 statistics indicate that the overidentifying restrictions are not rejected for any of the 16 stocks in the sample.

The results of Newey and West (1987) provide a multivariate test of the null hypothesis H_0: $\beta = 0$. A multivariate test is appropriate for these data because of the possibility of cross-stock inventory effects. Following the procedure described above, the statistic R has a value of 106.07 under the joint restriction. This statistic has a χ^2 distribution with 15 degrees of freedom.[8] The corresponding p-value is below 0.001, so that the null hypothesis of nonstationarity is strongly rejected. Intuitively, if dealer inventories were nonstationary a market maker with finite capital would eventually become bankrupt with certainty, as noted by Garman (1976).[9]

B. Cross-Stock Effects

The possibility that the specialist uses his inventory positions in some stocks to hedge his positions in other stocks suggests that we examine cross-stock correlations in specialist inventories. However, excluding October 19, 1987, when inventories rose sharply for most stocks, the cross-stock correlations in the estimated residuals from equation (6) are quite weak. Similar conclusions are reached by Hasbrouck and Sofianos (1993). Another

[7] The instruments used are a constant, current inventory, and lagged inventory over the previous three days. The standard errors reflect a correction for conditional heteroskedasticity and for serial correlation up to three lags.

[8] Stock 4 is excluded from the multivariate tests because it alone has 99 observations as opposed to 199 for the other stocks.

[9] We also estimate the model using Zellner's method of Seemingly Unrelated Regression Equations (SURE). The results from the joint generalized least squares estimation are not reported here and are essentially the same as reported in Table II. However, the system-wide R^2 for the SURE estimates is only 0.03, suggesting that mean reversion explains only a small fraction of the variation in specialist trades.

Table II

Mean Reversion in Specialist Inventories Using Daily Data from February to December 1987

GMM estimates of the model

$$I_t - I_{t-1} = \beta(I_{t-1} - I^d) + u_t,$$

where I_t represents the opening inventory on day t and u_t is the error term. The figures in parentheses are t-statistics based on Newey-West standard errors. The χ^2 statistic (p-value in parentheses) represents a test of the overidentifying restrictions for each stock. A multivariate test of the hypothesis that $\beta = 0$ for all stocks is rejected with a p-value below 0.001.

Stock	$\hat{\beta}$	$\hat{I}^d \times 10^{-3}$	χ_3^2
1	−0.017	7.77	4.53
	(−1.20)	(1.68)	(0.21)
2	−0.008	198.43	1.96
	(−0.46)	(3.62)	(0.58)
3	−0.020	72.51	1.41
	(−2.01)	(3.32)	(0.70)
4	−0.030	0.25	4.83
	(−1.82)	(0.01)	(0.18)
5	−0.115	10.81	2.48
	(−3.46)	(3.44)	(0.47)
6	−0.035	211.91	1.89
	(−1.78)	(10.03)	(0.60)
7	−0.006	166.74	3.87
	(−0.81)	(2.07)	(0.28)
8	−0.076	14.53	4.75
	(−2.96)	(2.38)	(0.19)
9	−0.019	225.98	1.97
	(−1.90)	(7.56)	(0.58)
10	−0.012	91.52	1.61
	(−0.69)	(0.66)	(0.66)
11	−0.069	20.45	2.49
	(−2.35)	(2.88)	(0.48)
12	−0.065	58.23	1.24
	(−2.58)	(2.65)	(0.74)
13	−0.047	24.75	4.13
	(−2.33)	(3.43)	(0.25)
14	−0.036	97.20	3.45
	(−2.50)	(5.11)	(0.33)
15	−0.002	790.94	2.87
	(−0.23)	(0.24)	(0.41)
16	−0.125	39.30	1.80
	(−3.21)	(2.85)	(0.62)

way to assess these effects is to estimate equation (6) using the dollar value of inventory held by this firm in all stocks. If there were significant cross-stock effects, the specialist's overall position should exhibit faster mean reversion than the individual stocks. This is not the case. The estimate of β is -0.032, and the corresponding t-ratio, using autocorrelation-heteroskedasticity–consistent standard errors based on the Newey-West (1987) procedure, is -1.355. The implied desired inventory level is \$33.23 million, with a t-ratio of 5.841. We cannot reject the null hypothesis of a unit root for dollar inventory.

In our model, steady state holdings are determined by long-term portfolio considerations, so it may appear surprising that the cross-stock effects are negligible. One explanation for this finding is that the value of the specialist firm's stock portfolio may not be large relative to the total assets of the firm's owners. Consequently, desired inventories may be determined primarily by the correlation between individual stock returns and the (unobservable) returns on the nonstock portion of the specialist's portfolio. Further, the specialist's trades in the individual stock may be strongly affected by idiosyncratic factors, so that portfolio considerations appear small. The absence of significant cross-stock inventory effects is consistent with the comments of observers familiar with the exchange, who note that specialists do not use their individual stock positions to offset market risk. Similarly, the use of derivatives or futures to hedge risk appears rare.

C. Shifts in Desired Inventory

The GMM estimation indicates that specialist inventories are jointly stationary, but the adjustment process appears to be very slow. We compute the adjusted coefficient of determination corresponding to the estimates reported in Table II by regressing the estimated residuals on the set of instruments. The adjusted R^2 (which is not reported here) is only 0.02 on average, indicating that lagged inventories by themselves account for only a small portion of the day-to-day variation in specialist trades. A possible explanation for these findings is that specialist trades may reflect periodic revisions in the target inventory level, I^d, so the mean of the process is not constant.

An examination of the time series behavior of individual stock inventory levels appears to confirm these suspicions. For many of the stocks there are substantial and persistent deviations of inventory from the mean over the entire sample period. Tsay (1986) has shown theoretically that failure to identify and correct for model misspecifications, such as periodic shifts in the mean of the process, can severely bias the estimate of the parameter β. Further, the estimation equation itself may take a different form if the specialist recognizes the possibility that desired inventory may change in the future. To obtain consistent estimates of the inventory effect, we must take into account, both theoretically and empirically, the effect of potential shifts in the target inventories. We address these issues in this section, starting with the theory and then moving on to estimation issues.

In deriving equation (6), we assume that the moments of the distribution over wealth are constant. This assumption may be reasonable for short time intervals. Over a longer horizon, however, the correlation between the specialist's outside income and equity returns may exhibit periodic changes, resulting in shifts in target inventories. From an economic perspective, such changes may arise from new projects being undertaken by the corporation whose stock is being traded. Alternatively, changes in the composition of the specialist's other assets may also alter the covariance term, and hence I^d.

The model can be extended to allow for shifts in I^d. Formally, denote by σ_t the covariance σ_{vy} at time t. We assume that this moment can change in period $t + 1$ with probability w, where $w \in (0, 1)$. In the event of a change in the joint distribution, the covariance in the next period has a mean equal to σ_t. Thus, we can write $\sigma_{t+1} = \sigma_t + \pi_t$, where π_t is a random variable with mean zero. The specialist knows the covariance at the time the quote is determined, but recognizes the possibility that this moment may shift in the future.

In the Appendix, we derive the following model with these additional assumptions about the covariance structure:

$$I_{t+1} - I_t = \beta(I_t - I_t^d) + u_t \tag{11}$$

$$I_t^d = -\frac{\sigma_0}{\sigma_v^2} - \sum_{i=1}^{s} \frac{\pi_{t_i}}{\sigma_v^2} S_{it}. \tag{12}$$

In equation (11), I_t^d is the target inventory at time t and u_t is a term proportional to the imbalance shock, x_t. Equation (12) describes the evolution of the target inventory, I_t^d; s is the total number of changes in the covariance structure (i.e., the number of dates on which $\pi_t \neq 0$), t_i is the time of the ith shift, and $S_{it} = 1$ if $t \geq t_i$ and 0 otherwise. The model is a generalization of equation (6) to allow mean reversion towards time-varying targets. The additional uncertainty induced by possible changes in the joint distribution of v and y reduces the specialist's expected utility of wealth, but does not alter the form of the strategy since the specialist is unable to hedge against this type of risk.

It would be easy to estimate equations (11) and (12) using dummy variables if the days on which shifts in steady state holdings occurred were known. However, these events arise in our model because of changes in the correlation between the specialist's nonstock wealth and equity returns, so that neither the dates nor the magnitude of these shifts are observable. We use intervention analysis to address this problem. Intervention analysis focuses on systematic patterns in residuals to identify and correct for events, such as periodic shifts in desired stockholdings, which have intervened into changing the underlying stochastic process.[10] The type of intervention illustrated in equation (12) is called a Level Shift (LS). A level shift is a "step"—it has a permanent effect on the mean of the series after a certain point in time.

Unlike most time series analyses, which focus on the overall patterns of serial correlation, intervention models require an examination of the residuals associated with specific observations. Intuitively, the estimated residuals should not display systematic patterns if the model is correctly specified. Various kinds of misspecifications will be reflected in systematic patterns in

[10] For a description of these techniques see, e.g., Fox (1972), Box and Tiao (1975), Tsay (1986), and Chang, Tiao, and Chen (1988).

the estimated residuals. For example, suppose in equation (12), that $s = 1$ and $t_1 = 98$, $\sigma_0 = 0$ and $-\sigma_1/\sigma_v^2 = 18{,}000$, implying that from day 99 onward, desired inventory increased by 18,000 shares. In this case, if we estimate equation (6) assuming a constant mean, the estimated residuals from day 99 onward would tend to be positive, while the estimated residuals on days 1 through 98 would tend to be negative. The timing of the shift in the mean can be identified by analyzing the time series of the model's residuals.

Other kinds of interventions can also be accommodated in a straightforward manner. For example, we can explicitly correct for heteroskedasticity in order imbalance shocks by assuming the error term u_t has two components: an error term, ϵ_t, with a constant variance, and an abnormal shock with a large standard deviation which occurs on p different days. Let e_j be the abnormal order flow on day t_j when the jth shock occurs, so that

$$u_t = \epsilon_t + \sum_{j=1}^{p} \epsilon_j P_{jt}, \tag{13}$$

where $P_{jt} = 1$ if $t = t_j$ and 0 otherwise. Equation (13) illustrates an intervention known as an Additive Outlier (AO). Unlike an LS-type intervention, an additive outlier is a "pulse" that affects the series for only a single period. From an economic viewpoint, allowing the error term to exhibit heteroskedasticity of this type is natural because inventories may exhibit occasional shocks arising from block positioning or large transitory order imbalances, such as those associated with the October 1987 crash.

Tsay (1986) proposes an iterative procedure for detecting and correcting for interventions:

1. The types of errors allowed and the criteria for identification are specified. The criterion for identification may be in terms of a significance level for a statistical test or a prespecified upper bound on the number of interventions.
2. Using the estimated model residuals, the largest outliers are identified and classified using a likelihood ratio test for the null hypothesis that there is no intervention on that date. For example, consider an AO-type intervention. To perform the likelihood ratio test for the ith observation, a consistent estimate of the residual error variance and an estimate of the magnitude of the intervention on that date are required. It can be shown that the best estimate of the size of an AO is a linear weighted combination of the residual for observation i and for future time periods. Similarly, the estimate of the error variance also depends on all the residuals.
3. If there are no significant outliers (where significance is determined by the choice of the critical value for the likelihood ratio test), then the procedure terminates. Otherwise, the data are corrected for the estimated magnitude and type of the interventions detected from the previous step (to provide consistent standard errors) and the model is reesti-

mated to identify the next largest outliers, until the remaining outliers are below the specified significance level or the maximum number of outliers is reached.

To identify the AO- and LS-type interventions in each stock's inventory series, we follow this iterative procedure, setting a maximum of five outliers per stock and choosing a relatively high significance level for the critical value in the detection step. Conservative values are chosen for two reasons: First, we wish to estimate a parsimonious model, and second, the inventory process is subject to a few, relatively large shocks to inventory in the sample period.

Of the 63 interventions identified for the individual stocks, 44 are of the LS-type, the remainder being AO-types. Interventions occur at a relatively constant rate over the sample period, except for the month of October 1987, where there are 20 AO- and LS-type interventions. For the dollar inventory series, there are three interventions, only one of which is of the LS-type. All three interventions occur within a week of the crash. This suggests that the shifts in desired inventory levels for individual stocks are small relative to the total dollar value of inventory or that the shifts in individual stocks offset each other.

Based on the results of the intervention analysis, we reestimate equation (6) for each stock, adding indicator variables based on the identified timings of the interventions. Table III presents the GMM estimates of the nonlinear model

$$I_{t+1} - I_t = \beta\left(I_s - \alpha_0 - \sum_{j=1}^{5} \alpha_j D_{jt}\right) + \epsilon_t, \tag{14}$$

where the α_j are constants, D_{jt} are indicator variables, and ϵ_t is the error term. The indicator variables are defined as follows: If outlier j is of the AO-type and occurs on day t_j, then $D_{jt} = 1$ for $t = t_j$ and 0 otherwise; similarly, if outlier j is of the LS-type and occurs on day t_j, then $D_{jt} = 1$ for $t \geq t_j$ and 0 otherwise. The system is just identified.[11] An estimate of the initial desired inventory is given by $\hat{\alpha}_0$, while $\hat{\alpha}_i$ is interpreted as the change in desired inventory on day t_i if the intervention is the LS-type, or the order flow shock on day t_i if the intervention is of the AO-type. Fewer than five indicator variables are included in the regressions for some stocks if fewer than five outliers are identified. No interventions are estimated for stock 4.

Table III should be compared with Table II. The mean reversion coefficient β in equation (14) is negative in all 16 cases, and is significantly negative at the 5 percent level in 13 cases. The average estimate of β in Table III is -0.134 as opposed to an average of -0.05 for equation (6). The Newey-West standard errors of β are also much lower. The great majority of the indicator

[11] The instruments for each stock are a constant, opening inventory, and the intervention dummies for that stock. With the inclusion of the intervention dummy variables, we found no need to adjust for serial correlation. The standard errors do correct for conditional heteroskedasticity.

Table III

Mean Reversion in Inventories Corrected for Interventions

GMM coefficient estimates (with Newey-West t-ratios in parentheses) of the model

$$I_t - I_{t-1} = \beta\left(I_{t-1} - \alpha_0 - \sum_{j=1}^{5} \alpha_j D_{jt}\right) + \epsilon_t,$$

where I_t is the opening inventory (scaled by 10^{-5}), D_{jt} is an indicator variable for intervention j on day t_j, and ϵ_t is the error term. The system is exactly identified. A multivariate test of the hypothesis that $\beta = 0$ is rejected with a p-value below 0.001.

Stock	$\hat{\beta}$	$\hat{\alpha}_0$	$\hat{\alpha}_1$	$\hat{\alpha}_2$	$\hat{\alpha}_3$	$\hat{\alpha}_4$	$\hat{\alpha}_5$
1	−0.06	−0.03	0.12	−0.05	0.29		
	(−2.47)	(−0.98)	(3.43)	(−1.23)	(1.57)		
2	−0.24	1.67	0.24[a]	0.26	0.24	0.17	−0.56
	(−2.16)	(26.11)	(1.55)	(13.18)	(10.48)	(5.73)	(−6.06)
3	−0.18	0.99	0.22	−0.48	−0.62	2.30	−2.30
	(−2.73)	(15.53)	(2.72)	(−8.16)	(−0.74)	(2.11)	(−3.78)
4	−0.03	0.00					
	(−1.82)	(0.01)					
5	−0.13	0.10	1.90[a]	−1.77[a]			
	(−3.75)	(3.75)	(3.54)	(−3.76)			
6	−0.15	1.82	1.78[a]	4.40	−4.85[a]	−4.04	0.43
	(−2.57)	(27.85)	(2.37)	(2.53)	(−2.35)	(−2.37)	(3.19)
7	−0.07	0.03	0.66	1.17	0.30	−0.75	3.89[a]
	(−1.84)	(0.15)	(1.63)	(3.29)	(1.52)	(−4.29)	(1.71)
8	−0.10	0.14	0.32				
	(−2.95)	(2.99)	(1.25)				
9	−0.16	2.53	−0.51	−0.08	0.87	1.94[a]	0.24
	(−3.27)	(13.86)	(−2.44)	(−0.61)	(1.28)	(2.74)	(0.35)
10	−0.13	3.07	−2.43	−5.75[a]	−0.60	3.21[a]	4.12
	(−1.44)	(7.90)	(−5.42)	(−1.35)	(−2.41)	(1.87)	(2.53)
11	−0.12	0.07	1.61[a]	0.18	0.32	−3.92[a]	
	(−3.93)	(1.13)	(4.17)	(2.06)	(3.12)	(−3.77)	
12	−0.11	−0.35	1.20	8.53[a]	8.33[a]	−4.72[a]	−0.66
	(−2.99)	(−2.02)	(4.60)	(3.31)	(3.27)	(−2.49)	(−2.11)
13	−0.06	0.23	2.65[a]	0.04	−3.20[a]		
	(−2.95)	(2.90)	(2.78)	(0.33)	(−3.21)		
14	−0.05	1.02	0.96	−0.96	−6.08	3.12	2.85
	(−2.75)	(5.51)	(0.57)	(−0.57)	(−2.69)	(1.91)	(2.24)
15	−0.31	−0.31	0.33	0.85	2.81	1.92	2.99
	(−3.97)	(−2.19)	(2.14)	(4.85)	(5.03)	(3.17)	(2.79)
16	−0.17	0.26	6.29[a]	−10.22[a]	−4.43[a]	1.11	−0.56
	(−4.28)	(2.76)	(4.50)	(−4.06)	(−4.43)	(2.00)	(−0.92)

[a] Denotes an AO intervention; the remaining are LS interventions.

variables have statistically significant coefficients. Since the model is just identified, the value of the χ^2 statistic is exactly zero. Accordingly, we compute the adjusted R^2 corresponding to these estimates; the average $\overline{R}^2$ for this set of regressions is 0.18, which is considerably higher than the average adjusted R^2 for the regressions in Table II. Again, a multivariate χ^2

test of the null hypothesis H_0: $\beta = 0$ is rejected ($\chi^2 = 337.8$) with a p-value below 0.001.

These results suggest that level shifts in the desired inventory level explain the poor fit of the model and the apparently weak inventory effects documented in Table II and found by previous researchers. Similar conclusions are reached by Hasbrouck and Sofianos (1993) using a very different technique and database. They perform a spectral analysis of specialist profits for a sample of 144 stocks, and note that specialist positions "are driven by rapid inventory adjustments towards targets that are themselves time varying."

D. Inventory Effects and the Speed of Adjustment

The speed of adjustment of inventories is directly related to the mean reversion coefficient β, which represents the fraction of the deviation between actual and desired inventories that is eliminated each day. A useful measure of adjustment speed is the inventory half-life, denoted by h, defined as the expected number of days required to reduce a deviation between actual and desired inventories by 50 percent, where

$$h = -\frac{\ln(2)}{\ln(1 + \beta)}. \tag{15}$$

Table IV provides two estimates of the inventory half-life for the 16 stocks. The first estimate, denoted by h^0, is based on the estimates of the mean reversion parameter β in Table II, assuming a constant mean for I^d. The second estimate, denoted by h^1, is based on the estimates of the mean reversion model correcting for interventions in Table III. Without any correction for interventions, the inventory effect appears weak. It takes, on average, 49.7 trading days for an inventory imbalance to be eliminated by half, with the smallest half-life being just over 5 days and the largest 334 days. By contrast, the inventory effect correcting for interventions is much stronger. On average, the half-life is only 7.3 trading days, ranging from 1.9 to 22.4 days. This adjustment process, however, is still longer than many researchers had previously thought.[12]

IV. Specialist Trades and Stock Returns

Our approach to analyzing the specialist's behavior so far has focused on inventory dynamics. Equation (8) suggests an alternative approach based on an examination of the relation between price changes and specialist trades. From a theoretical viewpoint, both approaches are representations of the same optimal policy rule. However, there are econometric advantages to this approach over our previous approach based solely on specialist trades. First, by differencing prices and inventories, we resolve some of the more serious

[12] The half-life, h, is zero if inventory adjustments occur within a day. This occurs if and only if $\beta = -1$, and this hypothesis is rejected for all stocks.

Table IV

Estimates of Inventory Half-Life for Specialist Stocks Based on Coefficient Estimates of the Mean Reversion Parameter

The estimated half-life, denoted by h, is defined as

$$h = -\frac{\ln(2)}{\ln(1+\beta)},$$

where β is the mean reversion parameter. Two figures are reported, based on the GMM estimates: h^0 denotes the half-life using the estimates of β contained in Table II, and h^1 denotes the half-life using the estimates of β corrected for AO and LS interventions in Table III. The sample mean ($\bar{h}$) and standard deviation (σ) are reported in the last two rows.

Stock	h^0	h^1
1	39.36	12.33
2	83.31	2.54
3	34.28	3.48
4	22.37	22.37
5	5.65	4.89
6	19.15	4.21
7	101.09	9.05
8	8.77	6.76
9	34.68	3.95
10	55.01	5.14
11	9.64	5.42
12	10.29	6.15
13	14.10	12.17
14	18.45	13.37
15	333.79	1.86
16	5.17	3.64
$\bar{h}$	49.69	7.33
σ	78.16	5.20

problems created by periodic shifts in the desired inventory level, without requiring an intervention model. Second, a model of price changes allows us to examine directly the impact of specialist trades on stock prices.

Let s_t denote the shock to order imbalances. An analysis of the estimated imbalances (constructed from the intraday data by signing volume) reveals the presence of some extremely large, stock-specific, outliers corresponding to large-block trades. It is likely that these trades are prearranged in the so-called upstairs market and executed on the floor, and should be excluded from an analysis of specialist quote determination.

Accordingly, we construct a time series of all intradaily large-block trades. We define a block buy (or sell) for a particular stock as a buy order from the top (or bottom) 0.5 percent of the size distribution of intradaily buys (or sells) for that stock over the sample period. Thus, our definition of a block trade varies by stock and by trade type. This procedure appears more reasonable than simply defining a block as 10,000 shares, irrespective of the volume of trading in that stock.

We then construct a daily net imbalance variable, denoted by q_t, for each stock defined as the order imbalance from day $t-1$ to day t less the aggregate volume of block trades for that day. The shock to order imbalances is modeled as the estimated residual in the following regression:

$$q_t = \gamma_0 + \sum_{i=1}^{M} \gamma_i q_{t-i} + \sum_{j=1}^{N} \delta_j (p_{t-j} - p_{t-j-1}) + s_t. \tag{16}$$

In estimating these regressions (which are not reported here) we used $M = N = 3$. This procedure is motivated by Hasbrouck (1991), who estimates a vector autoregressive model using intraday data for NYSE stocks. From an empirical viewpoint, the two approaches differ primarily in that: (1) we define order imbalances excluding large-block trades that are likely to have been negotiated in the upstairs market, and (2) we use daily data while Hasbrouck uses intradaily data. The shock variable s_t is then defined as the residual in the regression, i.e., $s_t = q_t - E[q_t \mid \Phi_{t-1}]$. Thus, s_t reflects the innovation in order imbalances net of large-block trades measured from the opening on day $t-1$ to day t.

The regression model follows from equation (8):

$$p_t - p_{t-1} = \gamma_0 + \gamma_1 s_t + \gamma_2 (I_t - I_{t-1}) + \gamma_3 \mathrm{OCT19}_t + \epsilon_t, \tag{17}$$

where p_t is the opening midquote from the ISSM file, s_t is the innovation in order imbalance from the opening on the previous day to the opening on day t, I_t is the opening inventory on day t, and $\mathrm{OCT19}_t$ is a dummy variable that equals one on October 19, 1987 and 0 otherwise, and ϵ_t is an error term which captures the unobservable specialist order imbalance signals. Equation (8) implies that $\gamma_0 = 0$, $\gamma_1 > 0$, and $\gamma_2 < 0$. The crash dummy is included to assess whether the drop in stock prices on the day of the crash could be explained by order imbalances alone. Blume, MacKinlay, and Terker (1989) find a significant relation between order imbalances and price movements on the day of the crash, but it is not clear whether this is true on other days as well. If $\gamma_3 = 0$, the drop in prices on October 19 can be explained entirely by selling pressure on that day.

Table V presents the GMM estimates of equation (17).[13] The results provide strong support for the model. As hypothesized, the constant γ_0 is generally close to zero. The information effect, as measured by γ_1, is positive in 15 cases and is statistically significantly positive at the 5 percent level for 13 stocks. The inventory parameter, as measured by γ_2, is negative as predicted for 14 of the 16 stocks, and is statistically significantly negative for 11 stocks. The coefficient of the dummy variable is negative for 14 stocks and

[13] The additional instrumental variables are lagged inventory changes, lagged imbalance shocks, and the lagged quote revision. We also estimate the model using a limited information (k-class) instrumental variables estimator which is the least variance ratio estimator. Our results are not sensitive to the estimation method.

Table V

Estimates of the Model of Quote Revisions

GMM estimates (with t-statistics in parentheses) of the model

$$p_t - p_{t-1} = \gamma_0 + \gamma_1 s_t + \gamma_2(I_t - I_{t-1}) + \gamma_3 \mathrm{OCT19}_t + \epsilon_t,$$

where, on day t, p_t is the opening mid-quote, s_t is the unanticipated order imbalance (the estimated residual from regressing nonblock order imbalances on lagged order imbalances and lagged quote revisions), I_t is opening inventory, $\mathrm{OCT19}_t$ is a dummy variable that equals 1 on October 19, 1987 and 0 otherwise, and ϵ_t is the error term. The χ^2 statistic (p-value in parentheses) tests each stock's overidentifying restrictions; a multivariate test of the hypothesis that $\gamma_1 = \gamma_2 = 0$ for all stocks is rejected with a p-value below 0.001.

Stock	$\hat{\gamma}_0$	$\hat{\gamma}_1 \times 10^5$	$\hat{\gamma}_2 \times 10^5$	$\hat{\gamma}_3$	χ^2_3
1	−0.02	6.82	−6.11	−0.94	1.39
	(−1.05)	(3.22)	(−2.40)	(−2.98)	(0.70)
2	−0.04	1.50	−2.42	−1.33	0.05
	(−1.90)	(2.74)	(−3.56)	(−4.22)	(0.99)
3	−0.01	1.43	−0.64	−4.10	1.91
	(−0.50)	(3.09)	(−1.47)	(−7.58)	(0.58)
4	0.01	−0.65	−5.80	0.20	1.34
	(0.13)	(−1.20)	(−4.81)	(0.36)	(0.71)
5	−0.00	0.10	−1.80	−0.98	1.41
	(−0.03)	(0.98)	(−3.68)	(−2.27)	(0.70)
6	−0.03	0.22	−0.54	−0.26	0.11
	(−1.91)	(2.32)	(−4.35)	(−1.02)	(0.98)
7	−0.00	3.84	−4.84	−8.55	1.41
	(−0.10)	(4.40)	(−5.65)	(−8.76)	(0.70)
8	0.06	4.32	−7.33	−6.27	2.93
	(0.67)	(4.75)	(−4.90)	(−4.48)	(0.40)
9	0.09	2.68	0.82	−7.72	11.59
	(1.88)	(5.24)	(1.21)	(−10.60)	(0.00)
10	0.11	0.45	0.18	−3.67	1.79
	(1.90)	(2.49)	(0.64)	(−4.36)	(0.61)
11	−0.02	1.48	−0.63	0.08	0.48
	(−0.86)	(6.23)	(−2.07)	(0.22)	(0.92)
12	−0.01	1.09	−1.37	−3.89	3.04
	(−0.19)	(3.48)	(−3.84)	(−4.41)	(0.38)
13	−0.00	0.91	−0.62	−1.33	0.64
	(−0.57)	(3.91)	(−1.76)	(−5.46)	(0.88)
14	−0.00	0.45	−1.26	−3.49	2.46
	(−0.26)	(1.57)	(−3.09)	(−7.92)	(0.48)
15	−0.02	0.80	−0.20	−7.33	7.85
	(−0.35)	(2.97)	(−0.91)	(−7.52)	(0.04)
16	−0.02	0.28	−0.43	−2.04	1.71
	(−0.74)	(3.62)	(−4.09)	(−4.68)	(0.63)

is significant for 13 stocks, suggesting that selling pressure and inventory effects cannot entirely explain the drop in prices on October 19, 1987.

The test of the overidentifying restrictions (distributed as χ^2_3) rejects the model at the 5 percent significance level for stocks 9 and 15. The estimated coefficient of γ_2 is not significant for either of these stocks. In general, the low values of the χ^2 statistic suggest that the instruments have a low correlation with the error term. The model appears to be well specified, and the adjusted R^2 is above 0.25 for 10 of the 16 stocks. Using a multivariate test, we can reject ($\chi^2 = 298.3$) the null hypothesis that $\gamma_1 = \gamma_2 = 0$ for all stocks with a p-value below 0.001.

When the daily large-block volume is added to the regression equation, this variable has little or no significance suggesting that it is the nonblock order flow that conveys information to the specialist. This result is consistent with Madhavan and Smidt (1991) who find that the price impact associated with large blocks is bounded above as a function of trade size. This may reflect the leakage of information concerning the impending block transaction since negotiating a large-block trade in the upstairs market takes time. Keim and Madhavan (1992) find evidence that the information conveyed by block trades may already have been impounded in security prices before the transaction occurs. This finding is also consistent with Seppi (1990), who shows that because of reputational concerns there may exist a separating equilibrium where uninformed agents submit block orders and informed agents place a series of smaller trades.

To further investigate the nature of the information available to the specialist, we estimate equation (17) with the leading order imbalance (net of large-block trades), i.e., q_{t+1}, as an additional explanatory variable. In theory, the contemporaneous price change should, under rational expectations, already reflect any anticipated future imbalances. However, if the specialist does have information about impending liquidity-motivated demands, our model predicts that the current price change will be positively related to the signals about this imbalance. We find that the coefficient of leading order imbalances is positive in all 16 cases and is significantly positive in 10 of these cases, while our previous conclusions regarding the inventory and information effects are unaltered. This result is consistent with the specialist receiving information signals regarding future nonblock order imbalances, as suggested here.[14]

V. Conclusions

This paper develops and tests an intertemporal model of market maker behavior. The model differs from the previous theoretical literature because it incorporates formally both asymmetric information and inventory control effects. We derive the optimal intertemporal trading strategy of the market maker when faced with an informed trader whose dynamic strategy is endogenously derived. The model sheds new light on the role of specialists. As a dealer, the specialist quotes prices that induce mean reversion towards a

[14] We also found Granger-Sims causality running from quote revisions to inventory changes in 12 of the 16 stocks, but little evidence for causality in the opposite direction. This result is consistent with our model where specialist quotes are positively related to expected future order imbalances which are subsequently absorbed into inventory.

target inventory level; as an investor, the specialist chooses a desired long-term inventory based on portfolio considerations, and may periodically revise this target.

We test the model using data obtained from a NYSE specialist. Specialist inventories exhibit mean reversion, but the adjustment process is slow if the specialist's desired inventories are assumed constant; it takes over 49 trading days, on average, for an imbalance in inventory to be reduced by 50 percent. The theoretical model suggests the possibility of periodic shifts in desired inventory holdings. Such shifts may account for the apparent slowness of the inventory adjustment process and the apparently weak inventory effects found in the literature. We develop an econometric model that corrects for periodic, unobserved shifts in the specialist's desired stockholdings and find strong evidence of mean reversion in inventories to these time-varying targets. The average inventory half-life, correcting for such shifts, is 7.3 trading days.

Turning to the effect of specialist trades on stock prices, we find strong support for our hypothesis that quote revisions are negatively related to specialist trades and positively related to the information conveyed by unanticipated (nonblock) order imbalances. Interestingly, large-block trades appear to convey little information to the specialist, perhaps because they have been anticipated by the specialist through leakage in the upstairs market. Further, the specialist appears to possess market information unavailable to most traders; future order imbalances affect current price quotations.

The analysis suggests that the specialist plays a far more complex role in price formation than previously thought. The specialist is not only a dealer who adjusts quoted prices to control fluctuations in inventory, but is also an active investor who seeks to maintain a long-term position in the stock while profiting in the short term from information about impending order imbalances. As an active investor, the specialist may hold large positions, and may periodically adjust the size of these positions based on long-term portfolio considerations.

Appendix

The Derivation of the Specialist's Trading Strategy, Equation (6)

We derive first the optimal quote-setting behavior of the specialist, given conjectures about the dynamic trading behavior of the informed trader. We then derive the informed trader's trading strategy, given the quotation rule of the specialist and verify that each agent's conjectures are consistent in equilibrium.

The specialist's beliefs are endogenously determined and depend on prior beliefs and the signal content of the submitted orders. As the state v_t is not observed by the specialist, we can model the learning process in a state-space framework, using the Kalman filter algorithm. Observing the net excess demand curve $z_t(p) = D_t - \delta p$ is equivalent to observing $w_t \equiv \delta^{-1} D_t$. Under the conjectured strategy of the informed trader, $Q_t(p) = \delta(v_t - p)$, so that

$w_t = v_t + X_t/\delta$. Since X_t has a zero mean, w_t is an unbiased signal about the unobserved fundamental price. Let $\sigma_w^2 \equiv \delta^{-2}\sigma_x^2$ denote the variance of this signal and let $\Upsilon = \sigma_\eta^2/\sigma_w^2$ represent the signal-to-noise ratio. A higher value of Υ implies more informative order flow and provides a more precise signal about the asset's value.

The Kalman filter provides a recursive method to summarize the learning behavior of the specialist; it provides the minimum mean square error estimate of the unobserved state. As the specialist observes signals from order flow, he updates his prior beliefs regarding asset values, generating a posterior distribution. However, since the unobservable state changes through time, the posterior beliefs never "converge" to the actual state; rather, posterior beliefs converge to a steady state distribution whose time-varying mean is an unbiased estimate of the true value of the asset at that point in time.

Consistent with rational expectations, we assume that the specialist's prior distribution is the "steady state" distribution over asset values. Let μ_t denote the market maker's forecast of v_t given the market maker's information at time t (including information on the submitted demand schedule, $z_t(p)$) and let σ_f^2 denote the conditional variance of this forecast. It is convenient to express this conditional variance in the form: $\sigma_f^2 = \sigma_w^2\Omega$, where σ_w^2 is the variance of the signal and Ω is a constant. Combining the prediction and updating equations for the Kalman filter (see, e.g., Harvey (1989), chapter 3) yields

$$\mu_t = \Omega w_t + (1 - \Omega)\mu_{t-1}. \tag{A1}$$

where

$$\Omega = \frac{-\Upsilon + \sqrt{\Upsilon^2 + 4\Upsilon}}{2} \tag{A2}$$

is the steady state weight placed on the signal about the unobservable state. As the signal-to-noise ratio, Υ, increases, Ω increases and less weight is placed on prior beliefs. Ω can also be interpreted as the conditional variance of the specialist's forecast, relative to the signal variance. Clearly, $0 < \Omega < 1$, as the forecast, which also uses prior information, is more accurate than the signal alone. Note that the conditional mean μ_t does not depend on p_t, but only on the intercept D_t.

Having defined the market maker's beliefs, we are now in a position to formally state the optimization problem. The specialist maximizes expected wealth, subject to inventory carrying costs. The inventory costs in period t are proportional to the variance of the specialist's wealth immediately after trading on day $t - 1$. Using equation (3), the ex ante variance of the specialist's wealth can be expressed as

$$\sigma^2\left[\tilde{W}_t\right] = \sigma_v^2 I_t^2 + \sigma_y^2 + 2I_t\sigma_{vy}, \tag{A3}$$

where σ_y^2 is the variance of $\tilde{y}_t$, σ_{vy} is the covariance between $\tilde{v}_t$ and $\tilde{y}_t$, and σ_v^2 is the conditional variance of $\tilde{v}_t$ given the specialist's information set prior to period t.[15] It follows that the inventory costs can be expressed as

$$c_t = \omega\left[\phi_0 + \phi_1(I_t - I^d)^2\right], \tag{A4}$$

where

$$\phi_0 = \sigma_y^2 - \left(\frac{\sigma_{vy}}{\sigma_v^2}\right)^2 \tag{A5}$$

$$\phi_1 = \sigma_v^2 \tag{A6}$$

$$I^d = \frac{-\sigma_{vy}}{\sigma_v^2} \tag{A7}$$

are constants. The parameter I^d is interpreted as the specialist's desired or target inventory.

The expected terminal wealth at time t depends on the current environment, which is summarized by four state variables: Opening inventory, I_t, the submitted order schedule, $z_t(p)$, the specialist's conditional expectation of the asset's value, μ_t, and the capital, K_t. Since the information on the current excess demand schedule is summarized by the intercept, D_t, we can use this as a state variable. Let $J(I_t, D_t, \mu_t, K_t)$ denote maximum value of expected final period wealth given these state variables and the transition equations governing their evolution. By assumption,

$$J(I_t, D_t, \mu_t, K_t) = \max_{p_t} \sum_{j=t}^{\infty} E\left[\tilde{W}_j\right] \Pr[\tau = j]. \tag{A8}$$

It is easiest to express the dynamic programming problem corresponding to equation (A8) by redefining the state variable for the demand schedule. Define

$$x_t = D_t - \delta\mu_t, \tag{A9}$$

so that $z_t(p_t) = \delta(\mu_t - p_t) + x_t$. Using equation (1), $x_t = X_t + \delta(v_t - \mu_t)$. Since μ_t is an unbiased estimate of v_t, x_t is also an unbiased estimate of the noise shock, X_t. Then, applying the Bellman principle, the functional equation equivalent to the optimality equation associated with the specialist's maximization problem is

$$J(I_t, x_t, \mu_t, K_t) = \max_{\{p_t\}} E\Big\{(1-\rho)\Big[\mu_t I_t + K_t - \omega\Big(\phi_0 + \phi_1(I_t - I^d)^2\Big)\Big] + \rho J\Big(\tilde{I}_{t+1}, \tilde{x}_{t+1}, \tilde{\mu}_{t+1}, \tilde{K}_{t+1}\Big)\Big|\Phi_t^s\Big\}. \tag{A10}$$

Intuitively, the functional equation consists of two parts. The first term in brackets is the expected wealth if the event occurs in the current period, i.e., at time t (by convention, the orders submitted in period t are not executed if

[15] By convention, this variance is computed before information on current orders becomes available to the specialist. Since $v_t = v_{t-1} + \eta_t$, we see that $\sigma_v^2 = \sigma_f^2 + \sigma_\eta^2$, i.e., the sum of the forecast variance of the previous asset value plus the variance of the innovation.

$\tau = t$), weighted by the probability of this event, i.e., $(1 - \rho)$. This term is the immediate payoff if the stock pays a liquidating dividend. The final term in equation (A10) represents the continuation value, i.e., the expected terminal wealth if liquidation does not occur in the current period, multiplied by the probability there is trading in period t.

For now, we suppress subscripts to minimize the notational burden. Define by $\tilde{I}'$ the inventory position in the next period, and similarly define $\tilde{x}'$, $\tilde{\mu}'$, $\tilde{K}'$. The state variables, i.e., inventories, beliefs, order imbalance signals, and capital, evolve according to the following transitional relations:

$$E[\tilde{I}' \mid \Phi^s] = I - \delta(\mu - p) - x \tag{A11}$$

$$E[\tilde{x}' \mid \Phi^s] = 0 \tag{A12}$$

$$E[\tilde{\mu}' \mid \Phi^s] = \mu \tag{A13}$$

$$E[\tilde{K}' \mid \Phi^s] = K + p(\delta(\mu - p) + x) - c. \tag{A14}$$

Equation (A11) states that expected inventories are current inventories less expected demand, equation (A12) states that the expected future imbalance is zero, equation (A13) is the Law of Iterated Expectations and implies that the revision in beliefs is an innovation, and equation (A14) states that future capital is the capital after trading in the previous auction plus trading revenue less inventory costs.

Assuming the functional equation J is differentiable and the maximizing value of p is interior, the first-order condition for the dynamic programming problem is found by differentiating equation (A10) with respect to p. This yields

$$\delta E\left[J_1\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}'\right)\right] + E\left[J_4\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}'\right)(\delta(\mu - p) + x - \delta p)\right] = 0. \tag{A15}$$

The envelope conditions are found by differentiating the value function (A10) with respect to each of the four state variables. These yield

$$\begin{aligned} J_1(I, x, \mu, K) = (1 - \rho)\mu - \left((1 - \rho) + \rho E\left[J_4\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{W}'\right)\right]\right) \\ \cdot\left(2\omega\phi_1(I - I^d)\right) + \rho E\left[J_1\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{W}'\right)\right], \end{aligned} \tag{A16}$$

$$J_2(I, x, \mu, K) = -\rho\left(E\left[J_1\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}'\right)\right] - pE\left[J_4\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}'\right)\right]\right), \tag{A17}$$

$$\begin{aligned} J_3(I, x, \mu, K) = (1 - \rho)I - \delta\rho E\left[J_1\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}'\right)\right] \\ + \rho\left(E\left[J_3\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}'\right)\right] + p\delta E\left[J_4\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}'\right)\right]\right), \end{aligned} \tag{A18}$$

$$J_4(I, x, \mu, K) = (1 - \rho) + \rho E\left[J_4\left(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}'\right)\right]. \tag{A19}$$

Equations (A15) to (A19) suggest that the value function takes the form

$$J(I, x, \mu, K) = A_0 + \mu I + K + A_1(I - I^d)^2 + A_2 x(I - I^d) + A_3 x^2, \quad \text{(A20)}$$

where A_i $(i = 0, \ldots, 3)$ are constants. Differentiating J, we obtain $J_1 = \mu + 2A_1(I + I^d) + A_2 x$. Since $E[\mu'] = \mu$ and $E[x'] = 0$, $E[J_1] = E[\mu' + 2A_1(I' - I^d) + x'] = \mu + 2A_1(I' - I^d)$. Substituting this expression into equation (A15) and using the fact that $E[J_4] = 1$, we obtain

$$p = \mu + A_1(I' - I^d) + \frac{x}{2\delta}. \quad \text{(A21)}$$

Now, $I' = I - \delta(\mu - p) - x$, and substituting equation (A21) into this expression yields equation (6):

$$I' = I + \beta(I - I^d) - \frac{1 + \beta}{2}x, \quad \text{(A22)}$$

where $\beta = \delta A_1/(1 - \delta A_1)$ or equivalently $A_1 = \beta/[(1 + \beta)\delta]$. Substituting (A22) and the conjectured form of J_1 and $E[J_1]$ into equation (A17) and equating the coefficients of $(I - I^d)$ and x respectively, we obtain $A_2 = -\rho A_1(1 + \beta) = -\rho\beta/\delta$ and $A_3 = \rho(1 + \beta)/4\delta$. Similarly, setting $J_1 = \mu + 2A_1(I - I^d) + A_2 x$ in equation (A16), equating coefficients yields $A_1 = -\omega\phi_1/(1 - \rho(1 + \beta))$. Substituting this term into previous expression for A_1, we obtain

$$\beta + \frac{\delta\omega\phi_1}{1 - \rho(1 + \beta) + \delta\omega\phi_1} = 0. \quad \text{(A23)}$$

For an economically meaningful solution we require $A_1 < 0$ or equivalently that $\beta \in (-1, 0)$. To show this, note that the left-hand side of equation (A23) is strictly positive when $\beta = 0$ and is strictly negative when $\beta = -1$. From the Intermediate Value Theorem, it follows that there exists $\beta \in (-1, 0)$ satisfying equation (A23). Further, as $\delta \to 0$, $\beta \to 0$, and as $\delta \to \infty$, $\beta \to -1$.

We now turn to the derivation of the model for quote revisions. Substituting equation (A22) into (A21) yields equation (7). Taking first differences yields

$$p_t - p_{t-1} = \mu_t - \mu_{t-1} - \zeta_1(I_t - I_{t-1}) + \zeta_2(x_t - x_{t-1}). \quad \text{(A24)}$$

Using equation (A1) and the definition of w_t, we can write the revision in beliefs as $\mu_t - \mu_{t-1} = \Omega\delta^{-1}[z_t - \delta(\mu_t - p_t) + \delta(\mu_t - \mu_{t-1})]$. Since the expected informed flow is $\delta(\mu_t - p_t)$, this implies that $\mu_t - \mu_{t-1} = [\Omega/(1 - \Omega)\delta]s_t$, where s_t is defined as the difference between the realized order flow z_t and the predicted informed order flow. Substituting this expression into equation (A24) and rearranging yields equation (8).

Derivation of the Informed Trader's Strategy, Equation (1)

Our objective is to examine conditions under which equation (1) represents the optimal dynamic order placement strategy of an informed trader who maximizes expected terminal wealth. The informed trader observes the value of the security before each auction, and can trade at each date. The trader has rational conjectures about the specialist's strategy. In turn, the behavior of the specialist is consistent with the trading strategy of the informed trader in equilibrium.

Before trading at time t, there is a probability $(1 - \rho)$ of an information event, in which case there is no trading, and the trader's expected wealth is $v_t B_t + C_t$, where B_t is the trader's endowment of the risky asset and C_t is the cash position. With probability ρ, however, the event does not occur at time t and the informed trader's utility is the continuation value given an updated inventory of shares and cash of $B_t + Q_t$ and $C_t - p_t Q_t$, respectively. The trader then receives information about v_{t+1}, and the process repeats itself until the event occurs.

It is easiest to derive the trading strategy of the informed trader by showing that any deviation from the trading rule $\delta(v_t - p_t)$ is suboptimal. Accordingly, let $Q_t = \delta(v_t - p_t) + \Delta_t$, where the increment Δ_t may be a function of the vector of the variables v_t, p_t, B_t, and C_t. We will demonstrate that the optimal increment is zero, i.e., $\Delta_t = 0$. The trader observes the equilibrium price in each auction. This price is an increasing function of the trader's order quantity, through its effect on the specialist's inventories and beliefs. Let P_t denote the price set by the specialist corresponding to the order quantity $\delta(v_t - p_t) + \Delta_t$. Substituting $w_t = (Q_t + X_t + \delta p_t)/\delta$ into equation (A1), and using equation (7), the trader views the price P_t as an increasing function of the increment, Δ_t: $P_t(\Delta_t) = p_t + \lambda\Delta_t$, where $\lambda \equiv \Omega/\delta$ and p_t is the base price, i.e., the price prevailing if $\Delta_t = 0$.

Let $V(v_t, p_t, B_t, C_t)$ denote the maximum expected wealth given the current state represented by the asset's current value, the specialist's base price quote, the trader's share endowment, and the trader's cash position. The functional equation for the trader is analogous to the functional equation associated with the specialist's maximization problem, except that the trader chooses the order quantity while the specialist selects a quote. Applying the Bellman principle, we obtain

$$\begin{aligned} V(v_t, p_t, B_t, C_t) = \max_{\Delta_t} E[&(1-\rho)(v_t B_t + C_t) \\ &+ \rho V(v_{t+1}, p_{t+1}, B_{t+1}, C_{t+1})], \end{aligned} \tag{A25}$$

subject to the transitional equations $E[v_{t+1}] = v_t$, $B_{t+1} = B_t + Q_t$, and $C_{t+1} = C_t - P_t Q_t$, where $Q_t = \delta(v_t - p_t) + \Delta_t$ and $P_t = p_t + \lambda\Delta_t$.

In equation (A25), the trader's expected terminal wealth consists of the expected wealth if liquidation occurs immediately, multiplied by the probability of this event, plus the continuation value, i.e., the expected terminal wealth given that the event does not occur and the trader's orders are executed times the probability that no event occurs. To solve the dynamic programming problem, we need a transitional relation for the notational base price. The base price in the next period depends on the trader's current quantity through information and inventory effects. From equation (7), the future base price (which results when $\Delta_{t+1} = 0$) is given by

$$p_{t+1} = \mu_{t+1}(\Delta_t) - \zeta_1(I_{t+1} - I^d) + \zeta_2 x_{t+1}, \tag{A26}$$

where we write $\mu_{t+1}(\Delta_t)$ to emphasize the fact that the conditional expectation in period $t + 1$ depends on the current increment Δ_t. Using the Law of Iterated Expectations, $E[\mu_{t+1}(\Delta_t)] = \mu_t(\Delta_t)$. From equation (A1), $\mu_t(\Delta_t) = \mu_t + \lambda\Delta_t$, i.e., the conditional mean in period t is μ_t, the mean that would obtain if the increment $\Delta_t = 0$, plus a term proportional to the unanticipated "excess" trade, $\lambda\Delta_t$. Using equation (7), we can write μ_t, the conditional mean that would prevail if the increment were zero, as $\mu_t = p_t + \zeta_1(I_t - I^d) - \zeta_2 x_t$. Substituting this into the expression just derived for $E[\mu_{t+1}(\Delta_t)]$, and using this in equation (A26), we obtain $p_{t+1} = p_t + \lambda\Delta_t - \zeta_1(I_{t+1} - I_t) + \zeta_2(x_{t+1} - x_t)$. Taking expectations yields $E[p_{t+1}] = p_t + \lambda\Delta_t + \zeta_1 Q_t$, where we assume that since the informed trader does not have information on noise trades, the expected values of current and future noise shocks (i.e., x) is zero.[16] This expression provides the desired transitional relation for the base price. It shows that a derivation Δ_t affects the future base price in two ways. First, the incremental trade biases the market maker's conditional expectation since inferences are predicated on the assumption that $\Delta_t = 0$. Second, the deviation affects inventory, and hence future prices. The informed trader need not attempt to infer the specialist's inventory or current beliefs since the price quote summarizes this information, i.e., is a sufficient statistic.

It is important to note that the transitional relations depend on the conjectured demand coefficient δ both directly and indirectly through the parameters ζ_1 and λ. When we solve for the equilibrium strategy, by setting the incremental trade to zero, we will obtain restrictions on δ. We then verify that there is indeed a positive δ coefficient that jointly satisfies the specialist's conjectures and the trader's optimality conditions.

Omitting time subscripts and using primes to denote variables in the next period, the first-order condition is obtained by differentiating equation (A25) and using the transitional relations for the state variables (v, p, B, C). This yields

$$\begin{aligned}(\zeta_1 + \lambda)E[V_2(v', p', B', C')] + E[V_3(v', p', B', C')] \\ -(p + 2\lambda\Delta + \lambda\delta(v - p))E[V_4(v', p', B', C')] = 0.\end{aligned} \tag{A27}$$

[16] The analysis can be extended to the case where the trader has private information about noise trading. In this case, the trader's demand function takes the form $Q_t = \delta_1(v_t - p_t) - \delta_2\hat{x}_t$, where δ_1 and δ_2 are positive constants and $\hat{x}_t$ is the trader's estimate of X_t. It is straightforward to show that the reduced form expressions for price and inventory dynamics are unaffected.

The associated envelope conditions are

$$\begin{aligned} V_1(v, p, B, C) = {} & (1-\rho)B + \rho E[V_1(v', p', B', C')] \\ & + \rho\zeta_1\delta E[V_2(v', p', B', C')] + \rho\delta E[V_3(v', p', B', C')] \\ & - \rho\delta(p + \lambda\Delta)E[V_4(v', p', B', C')], \end{aligned} \tag{A28}$$

$$\begin{aligned} V_2(v, p, B, C) = {} & \rho(1-\delta\zeta_1)E[V_2(v', p', B', C')] - \rho\delta E[V_3(v', p', B', C')] \\ & + \rho(\delta(p+\lambda\Delta) - \delta(v-p) - \Delta)E[V_4(v', p', B', C')], \end{aligned} \tag{A29}$$

$$V_3(v, p, B, C) = (1-\rho)v + \rho E[V_3(v', p', B', C')], \tag{A30}$$

$$V_4(v, p, B, C) = 1 - \rho + \rho E[V_4(v', p', B', C')]. \tag{A31}$$

Equations (A27) to (A31) suggest that $V(v, p, B, C) = vB + C + A(v-p)^2$, where A is an undetermined coefficient, so that $V_1 = B + 2A(v-p)$, $V_2 = -2A(v-p)$, $V_3 = v$, and $V_4 = 1$. Using the transitional relation for the base price, $E[p'] = p + \lambda\Delta + \zeta_1(\delta(v-p) + \Delta)$, so that $E[(v'-p')] = (v-p)(1 - \delta\zeta_1) - (\zeta_1 + \lambda)\Delta$. Substituting this into equation (A27), and using the functional form for V, and setting the optimal increment $\Delta = 0$, we obtain the following restriction:

$$A = \frac{1-\delta\lambda}{2(1-\delta\zeta_1)(\zeta_1+\lambda)}. \tag{A32}$$

In other words, for $\Delta = 0$, the coefficients must satisfy equation (A32). Similarly, using the expressions derived for $E[v'-p']$ and B' in equation (A28) and setting $\Delta = 0$, we obtain

$$A = \frac{\rho\delta}{1-\rho(1-\delta\zeta_1)^2}. \tag{A33}$$

Subtracting equation (A33) from (A32), and noting that $\delta\lambda = \Omega$ (the weight placed on the signal content of order flow) and $\delta\zeta_1 = -\beta > 0$ (the mean reversion coefficient), we obtain

$$\frac{1-\Omega}{2\rho(1+\beta)} - \frac{\Omega-\beta}{1-\rho(1+\beta)^2} = 0. \tag{A34}$$

In equation (A34), β is a function of δ. Further, Ω depends on Υ, which in turn depends on δ. To verify that an equilibrium exists, we must show that there exists a $\delta > 0$ solving equation (A34).

Let $f(\delta)$ denote the left-hand side of equation (A34). For $\delta \in (0, \infty)$, f is a continuous function of δ. As $\delta \to 0$, it is straightforward to show that $\beta \to 0$ and $\Omega \to 0$, and $f(\delta)$ tends to $1/2\rho > 0$. Further, as $\delta \to \infty$, $\beta \to -1$ and $\Omega \to 1$. Using L'Hôpital's rule, it can be shown that as $\delta \to \infty$, $f(\delta)$ becomes negative. Applying the Intermediate Value Theorem, there exists $\delta^* \in (0, \infty)$ such that $f(\delta^*) = 0$, so that there exists a set of conjectures for the specialist and the trader regarding their optimal dynamic strategies which are mutually consistent.

The Derivation of the Intervention Model, Equations (11) and (12)

The specialist's target inventories are stationary in the base case model where the moments of the distribution over wealth are assumed constant. While this assumption may be reasonable over short intervals, it may not be so plausible over a longer horizon. In particular, even though the asset's fundamental volatility, σ_v^2, may remain constant, the correlation between the specialist's outside income and equity returns may change periodically. From an economic perspective, such changes may arise from periodic changes in the projects undertaken by the corporation whose stock is being traded. Alternatively, a change in the specialist's wealth invested in other assets may also change I^d.

Here, we extend the model to allow for endogenous shifts in target inventory levels arising from changes in the stochastic process governing the specialist's wealth. Formally, let σ_t represent the covariance σ_{vy} at time t. The specialist knows the covariance before choosing his quoted price at time t, but this moment may change in period $t+1$ with some probability w, where $w \in (0, 1)$. In the event of a change in the joint distribution, the covariance in the next period has a mean equal to σ_t. Thus, we can express the covariance in the next period as $\sigma_{t+1} = \sigma_t + \pi_t$, where π_t is random variable with mean zero, i.e., $E[\sigma_{t+1} \mid \Phi_t^s] = \sigma_t$. Let $I_t^d = -\sigma_t/\sigma_v^2$ denote the target inventory at time t.

The additional uncertainty added by stochastic population moments means that we have to rederive the functional equation J, with σ_t (or equivalently I_t^d) as an additional state variable. It is notationally less cumbersome to express the functional equation using I_t^d as a state variable instead of σ_t. Let J^* denote the functional equation for this more complicated problem. Then, the analog to equation (A10) is

$$J^*(I_t, x_t, \mu_t, K_t, I_t^d) = \max_{\{p_t\}} \Big\{(1-\rho)\Big[\mu_t I_t + K_t - \omega\Big(\phi_0 + \phi_1 (I_t - I_t^d)^2\Big)\Big] + \rho E\Big[J^*\Big(\tilde{I}_{t+1}, \tilde{x}_{t+1}, \tilde{\mu}_{t+1}, \tilde{K}_{t+1}, \tilde{I}_{t+1}^d\Big)\Big]\Big|\Phi_t^s\Big\}. \quad \text{(A35)}$$

Note that the expectation $E[J^*]$ is taken over all future random variables, including σ_{t+1}.

Clearly, the transition equations (A11) to (A14) are unaffected. The law of motion for the new state variable is simply derived. Since $E[\sigma_{t+1} \mid \sigma_t] = \sigma_t$, it follows that $E[I_{t+1}^d \mid I_t^d] = I_t^d$. Repeating our previous arguments, the value function conditional upon I^d is

$$J^*(I, x, \mu, K; I^d) = A_0 + \mu I + K + A_1 (I - I^d)^2 + A_2 x I + A_3 x^2 + A_4 x, \quad \text{(A36)}$$

where A_i ($i = 0, \ldots, 4$) are constants, as before. Recall that I_{t+1}^d is a random variable with mean I_t^d. Taking expectations with respect to π_t implies that $E[J^*] = E[J] + A_1\sigma^2[-\pi_t/\sigma_v^2]$. As $A_1 < 0$ and π_t has a constant variance, the additional uncertainty induced by possible changes in the joint distribution of y and v reduces the value function by a constant amount. The uncertainty over the future covariance does not affect the first-order and envelope conditions of the specialist's optimization problem because this form of risk cannot be hedged. Following our previous arguments, the optimal trade is $I_{t+1} - I_t = \beta(I_t - I_t^d) + \gamma x_t$, yielding equation (11). To derive equation (12), suppose that there are s days in the sample period on which σ_1 changes, i.e., when $\pi_t \neq 0$. Let t_i denote the period when the ith nonzero realization of π occurs, and let $S_{it} = 1$ if $t \geq t_i$ and 0 otherwise. Then, equation (12) follows directly from the fact that $\sigma_{t+1} = \sigma_t + \pi_t$.

Note that as long as other traders do not receive private information signals about impending shifts in the covariance structure, the derivation of the demand function, equation (1), is unaffected because the expected change in desired inventory is zero (i.e., $E[I_{t+1}^d] = I_t^d$) implying that the effect of shifts on future prices is also an innovation. This completes the derivation of the intervention model.

REFERENCES

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

Blume, M., A. C. MacKinlay, and B. Terker, 1989, Order imbalances and stock price movements on October 19 and 20, 1987, *Journal of Finance* 44, 827–848.

Box, G. E. P., and G. C. Tiao, 1975, Intervention analysis with application to economic and environmental problems, *Journal of the American Statistical Association* 70, 70–79.

Chang, I., G. Tiao, and C. Chen, 1988, Estimation of time series parameters in the presence of outliers, *Technometrics* 30, 193–204.

Demsetz, H., 1968, The cost of transacting, *Quarterly Journal of Economics* 82, 33–53.

Foster, D. F., and S. Viswanathan, 1993, Variations in trading volume, return volatility, and trading costs: Evidence on recent price formation models, *Journal of Finance* 48, 187–211.

Fox, A. J., 1972, Outliers in time series, *Journal of the Royal Statistical Society, Series B* 34, 350–363.

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

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

Hansen, L., 1982, Large sample properties of generalized method of moment estimators, *Econometrica* 50, 1029–1084.

Harvey, A., 1989, *Forecasting, Structural Time Series Models, and the Kalman Filter* (Cambridge University Press, Cambridge).

Hasbrouck, J., 1988, Trades, quotes, inventories and information, *Journal of Financial Economics* 22, 229–252.

———, 1991, Measuring the information content of stock trades, *Journal of Finance* 46, 178–208.

——— and G. Sofianos, 1993, The trades of market-makers: An analysis of NYSE specialists, *Journal of Finance* 48, 1565–1593.

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

Keim, D., and A. Madhavan, 1992, The upstairs market for large-block transactions: Analysis and measurement of price effects, Working paper, Wharton School, University of Pennsylvania.

Kyle, A., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

Lee, C., and M. Ready, 1991, Inferring trade direction from intradaily data, *Journal of Finance* 46, 733–746.

Madhavan, A., and S. Smidt, 1991, A Bayesian model of intraday specialist pricing, *Journal of Financial Economics* 30, 99–134.

Neuberger, A., 1992, An empirical examination of market maker profits on the London Stock Exchange, Working paper, London Business School.

Newey, W., and K. West, 1987, A simple, positive semi-definite, heteroskedastic and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Seppi, D., 1990, Equilibrium block trading and asymmetric information, *Journal of Finance* 45, 73–94.

Silber, W. L., 1984, Marketmaker behavior in an auction market: An analysis of scalpers in futures markets, *Journal of Finance* 39, 937–953.

Stoll, H., 1976, Dealer inventory behavior: An empirical investigation of NASDAQ stocks, *Journal of Financial and Quantitative Analysis* 11, 359–380.

———, 1989, Inferring the components of the bid-ask spread: Theory and empirical tests, *Journal of Finance* 44, 115–134.

Tsay, R. S., 1986, Time series model specification in the presence of outliers, *Journal of the American Statistical Association* 81, 132–141.

U.S. Securities and Exchange Commission, 1971, *Institutional Investor Study Report*, 92nd Congress, 1st Session (U.S. Government Printing Office, Washington, D.C.).

Working, H., 1977a, Price effects of scalping and day trading, in A. E. Peck, ed.: *Selected Writings of Holbrook Working* (Chicago Board of Trade, Chicago).

———, 1977b, Tests of a theory concerning floor trading on commodity exchanges, in A. E. Peck, ed.: *Selected Writings of Holbrook Working* (Chicago Board of Trade, Chicago).