AN ANALYSIS OF DAILY CHANGES IN SPECIALIST INVENTORIES AND QUOTATIONS

by

Ananth Madhavan
Seymour Smidt

22-92

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

An Analysis of Daily Changes in Specialist Inventories and Quotations

Ananth Madhavan
and
Seymour Smidt*

Original Version: September, 1991
Current Version: October, 1992

*Wharton School, University of Pennsylvania and Johnson School, Cornell University, respectively. We thank Marshall Blume, Fernando Diz, Prajit Dutta, Thomas Finucane, Lawrence Harris, Joel Hasbrouck, Richard Highfield, Eric Hughson, Shmuel Kandel, Donald Keim, Craig MacKinlay, Krishna Ramaswamy, Oded Sarig, and Laura Starks for their helpful comments. Special thanks are due to Matt Richardson for his help with the estimation. We have also benefitted from the comments of seminar participants at the California Institute of Technology, Cornell University, the University of Rochester, Syracuse University, 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.

An Analysis of Daily Changes in Specialist Inventories and Quotations

Abstract

This paper develops a model of market making that incorporates both inventory control and asymmetric information effects. We show that the specialist acts both as a market maker and as an active investor trading for his own account. As a market maker, the specialist quotes prices that induce mean reversion toward a desired level of inventory; as an active investor, he periodically adjusts the target inventory levels towards which inventories revert. We test the model using data obtained from a NYSE specialist. We find that specialist inventories exhibit mean reversion, but the adjustment process is slow, even controlling for shifts in target inventories. The model also predicts that quote revisions are negatively related to specialist trades and positively related to the information conveyed by order imbalances. We find strong evidence for this hypothesis; further, our results suggest that specialist quotes anticipate future order imbalances.

1 Introduction

Interest in the behavior market makers reflects their crucial roles in price formation and in the provision of liquidity in securities markets. This paper analyzes, both theoretically and empirically, the trading behavior of a 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 market maker. Early analyses of the specialist's market making function presumed that they would passively provide liquidity to accommodate transitory order imbalances, thereby stabilizing prices. To perform their market making function, however, specialists must bear unwanted inventories. Theoretical models show that when market makers face inventory carrying 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 until now have not been integrated into a formal model with optimizing agents.

Empirical analyses of market maker behavior are mainly limited to indirect implications of these theories because the available public databases do not distinguish market

[1] See, e.g., Amihud and Mendelson (1980) and Ho and Stoll (1983).
[2] See, e.g., Kyle (1985) and Glosten and Milgrom (1985).

maker transactions from those of other market participants.[3] Very few studies are based on the actual trading records of market makers in any financial markets.[4] 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 (e.g., Hasbrouck (1988), Stoll (1989), and Madhavan and Smidt (1991)) find only weak evidence of short-run inventory effects but strong information effects. These results argue for examining inventory behavior over longer horizons.

While these 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, as presumed by asymmetric information models, 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 examines these issues, both theoretically and empirically.

We develop an intertemporal model of specialist trades and quotes. The model differs from previous models in three important respects: (1) The model incorporates formally the effects of both asymmetric information as well as inventory control, where the behavior of all agents is explicitly derived from utility maximization. (2) We model the specialist as both a market maker who provides liquidity on demand and an active investor for his own account. Previous models assume that the specialist acts as a pure market maker, ignoring the possibility that specialist trades may also reflect investment and speculative motives. (3) We explicitly model the impact on quotes and trades when the specialist receives information about liquidity-based trading. Such information may arise from the

[3]See, e.g., Glosten and Harris (1988) and Hasbrouck (1988).

[4]See Working (1977a), Working (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 (1992).

central position of the specialist on trading floor, his privileged access to the limit order book, and the fact that he is the first to receive indications about order imbalances on computerized trading systems.[5]

The model provides new insights into the role of the specialist. We show that a specialist acts both as a market maker and as an active investor managing his portfolio exposure. As a market maker, the specialist quotes prices that induce mean reversion in inventory. 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.

The optimal quotations of the specialist induce mean reversion in inventory towards target inventory levels determined by relatively long-term considerations. We find very slow mean reversion in inventories if the specialist's desired stockholdings are assumed constant. It takes over 49 trading days, on average, for an imbalance in inventory to be reduced by 50%. However, the specialist is an active investor as well as a market maker, and shifts in desired inventories may bias our estimates of mean reversion. We use 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%.

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 non-block portion of order flow that appears to have information content. We also find evidence that the specialist anticipates future

[5]Indeed, the 'tick test' was introduced in the 1930s in response to congressional concerns that the specialist's trading was detrimental to the interests of their brokerage customers. Thus, federal regulation of specialists in the United States starts from the premise that individual investors must be protected from exploitation by better informed specialists. Forster and George (1991) provide a model where market makers are at an informational disadvantage relative to traders with security-specific private information, but possess private information about the distribution of noise trading.

order imbalances. This finding provides support for early arguments that the specialist's unique position provides him with market information not available to most traders.

The paper proceeds as follows. In Section 2, we develop a theoretical model of specialist trading. In Section 3, we describe the data and in Section 4 we discuss our estimation technique. Section 5 provides the results of our analyses of specialist inventories while Section 6 examines the relation between specialist trades and stock prices. Finally, Section 7 summarizes the paper and offers some suggestions for further research.

2 A Framework for Analysis

2.1 The Trading Environment

We begin by developing an intertemporal model of specialist trades and quotes. As noted above, the model differs from previous theoretical work in three important respects: First, the model incorporates the two main theories of market maker behavior, namely the asymmetric information and inventory control theories, in a framework where all agents solve an explicit utility maximization problem. Second, the specialist acts both a market maker and an active investor. Third, the specialist's information about the distribution of liquidity-based trading is incorporated in the decision process.

Consider the market for a single risky asset that trades on days $t = 1, 2, \ldots, \tau$. Let v_t denote the value of the security on day t. On the final trading date, τ, the security pays a liquidating dividend. The dividend date is a random variable, and we assume that at the beginning of day t, before any trading occurs, there is a positive probability $(1-\rho)$ that day will be the liquidation date, i.e., $\Pr[t=\tau] = (1-\rho)$. We assume that the security's fundamental value follows a random walk, i.e., $v_t = v_{t-1} + \eta_t$, where η_t is normally distributed error term with mean zero and variance σ_η^2.

At the beginning of day t, the specialist quotes a price p_t for the security and investors then submit their orders given this price.[6] The specialist's quotes are determined by his

[6] It is not difficult to relax the assumption of a single price and incorporate a bid-ask spread without altering our basic conclusions. We show the specialist can make positive profits through his knowledge of liquidity demands, even in the absence of a spread.

conjectures about trader demands, which in turn maximize utility given the quotes set by the specialist. Equilibrium requires that these conjectures and actions be consistent.

The submitted orders may or may not get executed depending on whether day t is the liquidation date. If with probability $(1-\rho)$, day t is the dividend day, then submitted orders are canceled, and the security pays a liquidating dividend. Otherwise, with probability ρ, day t is not the liquidation date, and the submitted orders are executed at the quoted price by the specialist who absorbs any excess demand into his inventory. 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. Let I_t denote the specialist's inventory before trading takes place on day t, so that $I_t = I_{t-1} - z_{t-1}$. After trading is completed, the specialist updates his beliefs based on order flow and quotes a price for the next round of trade. Investors then submit their orders for execution, and the process continues until, with probability one, the liquidation date is reached.

Traders and the specialist maximize the expected utility of final period wealth, $\widetilde{W}_\tau$. Consistent with previous theoretical research, we assume that traders have mean-variance expected utility functions of the form:

$$U(\widetilde{W}_\tau) = E[\widetilde{W}_\tau|\Phi_t] - \omega\sigma^2[\widetilde{W}_\tau|\Phi_t], \quad (1)$$

where $\omega > 0$ is a parameter that measures the degree of risk aversion and $E[\cdot]$ and $\sigma^2[\cdot]$ denote the conditional expectation and variance operators with respect to the information set, Φ_t.

At time t, the specialist views the asset's current fundamental value as a random variable, denoted by $\tilde{v}_t$. Traders observe the current realization of v_t, which is the payoff that would occur if t is the liquidation date, and trade based on this information. We could, with no loss of generality, assume that traders had only a noisy estimate of v_t. The specialist's beliefs will be a function of the observed history of order flow, as we show below. Excess demand on day t originates from N traders, who enter the market with an existing endowment of the security. Let E_{it} denote the share endowment of trader

i in auction t, where we normalize the mean endowment over all traders to be zero. A negative endowment is interpreted as a short position. For simplicity, we assume that traders trade only once, when they have private information about fundamental values, and do not place orders at other times. We normalize the discount rate to zero to simplify the notation; this assumption has no effect on our qualitative results because there is, with probability one, a terminal date.[7]

We show in the appendix that the aggregate excess demand can be represented by the equation:

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

where $\delta = \frac{N(1-\rho)}{2\omega\sigma_\eta^2}$ measures the responsiveness of demand to price and $X_t = -\sum_i E_{it}$ represents liquidity trading. Note that liquidity trading arises endogenously within the model; it represents the portfolio hedging demands of traders which are the exact analogue of the specialist's motives for inventory control.

As discussed above, the specialist may have information about the composition of liquidity-motivated demands. Formally, we assume that $X_t = x_t + \xi_t$, where x_t is the portion of the shock observed by the specialist and ξ_t is the unanticipated shock. We assume that at time t, ξ_t and x_{t+1} are pure white noise and are temporally uncorrelated with all other random variables. Then, the specialist views excess demand as a random variable, denoted by $\tilde{z}_t(p)$, given by:

$$\tilde{z}_t(p_t) = \delta(\tilde{v}_t - p_t) + x_t + \tilde{\xi}_t, \tag{3}$$

where $\tilde{\xi}_t$ is the unknown order flow shock.

Like traders, the specialist maximizes the expected utility of his final period wealth, $\widetilde{W}_\tau$. At time t, the specialist maximizes:

$$E[u(\widetilde{W}_\tau)|\Phi_t^s] = E[\widetilde{W}_\tau|\Phi_t^s] - \omega_s\sigma^2[\widetilde{W}_\tau|\Phi_t^s], \tag{4}$$

where $\omega_s > 0$ measures the specialist's risk aversion and $E[\,\cdot\,]$ and $\sigma^2[\,\cdot\,]$ denote the conditional expectation and variance operators given the specialist's information set at

[7] It is straightforward to interpret the probability ρ as an effective discount factor for a problem where the discount rate is strictly positive.

time t, Φ_t^s. The specialist's information set, Φ_t^s, consists of information on the history of trading and noisy information signals about impending order imbalances. Note that the specialist's risk aversion parameter is not constrained to be the same as that of the traders, as specialists may be less risk averse than the average trader.

The specialist's wealth at the beginning of period t consists of the value of his opening inventory, $\tilde{v}_t I_t$, plus his opening capital, denoted by $\tilde{K}_t$. Thus:

$$\widetilde{W}_t = \tilde{v}_t I_t + \tilde{K}_t. \tag{5}$$

The capital at the start of period t represents the value of the specialist's holdings in other assets and is a random variable because the current income from other assets, denoted by $\tilde{y}_t$, is stochastic. The specialist's capital evolves according to the following transition rule:

$$\tilde{K}_{t+1} = K_t + p_t z_t + \tilde{y}_t, \tag{6}$$

i.e., it consists of the previous period's capital, trading revenues, $p_t z_t$, and income from other financial assets. We assume that $\tilde{y}_t$ is distributed identically with mean zero. The assumption of a zero mean is just a convenient normalization, and does not affect our results in any way. Note that this income is received in period t irrespective of trading activity, i.e., even in the final period, τ. 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 solution to the utility maximization problem (4) represents the indirect utility function, whose value depends on four state variables summarizing the current environment: the opening inventory, I_t, anticipated order imbalances, x_t, and the conditional expectations of the security, μ_t, and non- equity capital, K_t, given the specialist's prior beliefs. Formally, let $J(I_t, x_t, \mu_t, K_t) = \max_{\{p_t\}} E[u(\widetilde{W}_\tau)|\Phi_t^s]$ denote maximized value of equation (4). We assume that J is well-defined and that the specialist's utility maximization problem at time t can be written, using the Bellman principle of optimality, as:

$$J(I_t, x_t, \mu_t, K_t) = \max \sum_{j=t}^{\infty} \left(E[\widetilde{W}_j] - \omega_s \sigma^2[\widetilde{W}_j] \right) \Pr[\tau = j]. \tag{7}$$

2.2 Inventory Dynamics

In the appendix, we show that in equilibrium, the solution to the dynamic maximization problem (7) implies that the change in the inventories (i.e., specialist trades) are:

$$I_{t+1} - I_t = \beta(I_t - I_d) + \gamma x_t + \epsilon_t, \tag{8}$$

where $-1 < \beta < 0$ and $-\frac{1}{2} < \gamma < 0$ are constants, $I_d = -\sigma_{yv}\sigma_v^{-2}$, and $\epsilon_t = [\delta(E[\tilde{v}_t|\Phi_t^s] - v_t) - \xi_t]$ represents an error term from the viewpoint of the specialist. This error term represents the unanticipated component of order flow. Clearly, $E[\epsilon_t] = 0$ because the revision in beliefs is an innovation.

Equation (8) illustrates the dual nature of specialist trading behavior. The specialist's market-making role is reflected in the first term on the right-hand side of equation (8), which implies that inventories exhibit mean reversion. The specialist's role as an active investor is reflected in I_d, the desired or target inventory level towards which inventories revert. In our model, I_d is the covariance between the stock's return and shocks to the specialist's outside income, relative to the conditional variance of the stock's return.[8] The theory suggests that desired inventories may shift with periodic adjustments in the composition of the non-equity portion of the specialist's portfolio or in the returns distribution. These adjustments may occur prior to the event date, τ, or may coincide with such events.

Equation (8) also illustrates the short-horizon investment strategy of the specialist, as reflected in the term γx_t. This term represents the expected change in the specialist's inventory position that results from his information about impending order imbalances. As γ, the coefficient of x_t, lies between $-\frac{1}{2}$ and 0, the specialist is willing to accommodate less than half the anticipated order imbalance. Nevertheless, to an outside observer, the

[8]The desired holdings may also reflect short-sale restrictions and minimum capital requirements.

specialist appears to be stabilizing prices by absorbing transitory demand shocks into inventory.

2.3 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. Substituting equation (3) (using the fact that $z_t = I_t - I_{t+1}$) into (8) yields:

$$p_t = \mu_t - \zeta_1(I_t - I_d) + \zeta_2 x_t, \tag{9}$$

where $\zeta_1 = \delta^{-1}|\beta| > 0$ represents the inventory effect and $\zeta_2 = \delta^{-1}(1+\gamma) = (1-\beta)/2 > 0$ represents the effect of order imbalances on price. The equation shows that the deviation between the price set by the specialist and the expected value of the security is inversely related to inventory and positively related to the anticipated order imbalance. Intuitively, the specialist lowers price when inventories are high 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.

Taking first differences in equation (9), we obtain:

$$p_t - p_{t-1} = (\mu_t - \mu_{t-1}) - \zeta_1(I_t - I_{t-1}) + \zeta_2(x_t - x_{t-1}). \tag{10}$$

Equation (10) shows that returns from day $t-1$ to t can be decomposed into three components: (1) the change in the specialist's beliefs, (2) a term proportional to the change in inventory, and (3) a term proportional to the change in predicted order imbalances. The latter term is unobservable and constitutes the error in the regression model. However, the first term is related to the current order flow. To estimate equation (10) we must specify how the specialist's beliefs evolve over time.

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:

$$p_t - p_{t-1} = \Omega\delta^{-1}(z_{t-1} - E[\tilde{z}_{t-1}|\Phi^s_{t-1}]) - \zeta_1(I_t - I_{t-1}) + \zeta_2(x_t - x_{t-1}), \qquad (11)$$

where $0 < \Omega < 1$ is interpreted as the weight placed on the noisy signal from order flow about the unobservable state, which is proportional to the variance of market maker's forecast of the asset's value. The first term on the right-hand side of equation (11) is the deviation between the actual and predicted excess demand. Equation (11) implies that 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. This representation suggests that specialist trading behavior can be examined not only through the degree of mean reversion in inventory but also through the impact of specialist trades on prices. Both approaches are equivalent theoretically (The inventory effect ζ_1 is proportional to the speed of mean reversion $|\beta|$.), but differ from an econometric perspective. In our empirical analyses, we use both approaches.

Our analysis of inventory dynamics also has implications for the autocorrelation pattern in prices. Using equation (10), we can compute $E[\tilde{p}_{t+1} - p_t|\Phi^s_t]$, i.e., the expected quote revision given the specialist's information just before trade in period t:

$$E[p_{t+1} - p_t|\Phi^s_t] = \delta^{-1}\beta^2(I_t - I_d) - (\zeta_1\gamma + \zeta_2)x_t. \qquad (12)$$

Equation (12) implies that the unconditional expected return from day t to $t+1$ is positively related to inventory, showing that knowledge of market maker positions is valuable to other traders. It can be shown that $(\zeta_1\gamma + \zeta_2) > 0$, so that the expected price change is inversely related to the anticipated order imbalance. Intuitively, equation (9) shows that if the expected order imbalance is positive, the profit maximizing strategy of the specialist is to raise prices. As order imbalances are transitory, the expected price on the following day is less than the current price, producing a negative expected return.

2.4 Comparative Statics

In this section, we examine the effects of changes in the model parameters on the specialist's price quotations and trades. From equation (8), 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. Recall that the stabilization parameter γ in equation (8) measures the fraction of the anticipated order imbalance that the specialist plans to accommodate. As $\gamma = -(1+\beta)/2$, lower values of β imply less stabilization and greater price fluctuations. Thus, the speed of mean reversion can be used to gauge the specialist's willingness to stabilize temporary order flow shocks.

From the solution to the specialist's optimization problem in the appendix, $\beta = -\delta\omega_s\sigma_v^2/(1+\delta\omega_s\sigma_v^2)$, where σ_v^2 represents the conditional variance of $\tilde{v}_t$ given the specialist's information set at time t. We derive a closed-form expression for this variance in the appendix and show it to be a complicated function of the variances of the error terms ξ and η. This expression for β shows that the speed of adjustment is a decreasing function of depth, δ, the specialist's degree of risk aversion, ω_s, and the conditional asset variance, σ_v^2. Intuitively, the greater the demand responsiveness, the disutility of carrying unwanted inventory or uncertainty about asset value, the greater the incentives to control inventories and the stronger the mean reversion.

The demand parameter δ in equation (2) measures the responsiveness of order flow to a change in specialist quotes; it is a metric for the liquidity or depth of the market. In a liquid or deep market, where δ is high, large trades can be accomplished with very little change in quoted prices. Depth, δ, increases with trading frequency or breadth, N, and with the likelihood of an information event, $(1-\rho)$, and decreases with greater risk aversion, ω, and asset risk, σ_η^2.

In the special case where δ or ω_s is near zero, β is near zero, and inventory changes approximate a random walk. Intuitively, if demand is completely unresponsive or 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.

Another special case occurs when traders and the specialist have the same degree of risk aversion. In this case, β is independent of the coefficient risk aversion.

The effect of inventories on price is measured by $\zeta_1 = \omega_s\sigma_v^2/(1+\delta\omega_s\sigma_v^2)$. The inventory effect diminishes with greater depth, δ, and increases with greater risk aversion, ω_s, and uncertainty, σ_v^2. Finally, the information effect, measured by the coefficient $\delta^{-1}\Omega$ in equation (11) is independent of the specialist's degree of risk aversion; it is inversely related to market depth and is positively related to the parameter Ω which represents the weight placed on the signal content of order flow. In turn, Ω increases with the variance of noise trading and decreases with the variance in fundamentals. This discussion shows that less liquid markets tend to be associated with stronger inventory and information effects.

3 The Data

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. (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 non-specialist 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 non-specialist transactions and quotes used in this study were obtained from the ISSM transaction files. The file contains transaction and quote data for NYSE and AMEX stocks, as well as transactions and quotes from other markets in the National Market System (NMS). Although the specialist data signs volume, the ISSM files do not indicate whether a transaction was buyer- or seller-initiated. These transactions were classified as buyer or seller-initiated using a method developed by Lee and Ready (1991) that compares the transaction price with the prevailing bid-ask quotations. Lee and Ready consider quotes that are eligible for inclusion in the National Market System and NASD Best Bid and Offer (BBO) quotes. The ISSM quote records include associated condition codes which allow users to identify BBO eligible quotes. A transaction is classified as a buy if the price is greater than or equal to the prevailing BBO ask, or closer to the ask than the bid, and as a sell if the price is less than or equal to the prevailing BBO bid, or closer to the bid than the ask. Trades for which there were no BBO eligible quotes, or which take place within the spread, are classified using the traditional tick test. The prevailing bid and ask quotations are quotes that are at time stamped at least five seconds before the reported time of the trade. Using the signed volume, we computed 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 non-specialist trading in the 16 stocks in the sample period. We formed a time-series of the specialist's inventory at the beginning of each trading day for the 16 stocks traded between February 1 to December 31, 1987 which had an average of at least 4 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 formed a time-series of opening bid and ask quotes and daily order imbalances in the National Market System (NMS). The ISSM file was 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 was 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. In computing closing inventories, the absolute value of the daily value was used. Changes in closing inventories are the absolute difference between the inventory values on a given day and on the previous trading day.[9] 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 inventories levels to be larger in more active stocks. Over all 16 stocks, the smallest average 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 ranged from $15,000 in stock 1 to $835,000 million in stock 12. The average ratio of inventory change to value of trading is 2.6%. 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% if the specialist were the only party on the other side of every trade. The average dollar participation ratio is 5.1% but the average transaction participation ratio is 52.7%. 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% of the trading value.

[9] In the case of stock splits, data measured in shares were adjusted to the pre-split levels.

4 Empirical Methods

4.1 Generalized Method of Moments Estimation

The econometric technique used throughout the paper is Hansen's (1982) Generalized Method of Moments (GMM). This technique is particularly appropriate here because the procedure is based on weak assumptions for the stochastic process generating the data. Using the results of Newey and West (1987), we can adjust the error term for serial correlation and conditional heteroskedasticity. Further, GMM provides a method to perform multivariate tests of stationarity in specialist inventories, which may be important because of cross-stock effects.

To illustrate the GMM approach, note that equation (8) can be written in the form:

$$u_t = I_t - I_{t-1} - \beta(I_{t-1} - I_d), \tag{13}$$

where $u_t = \epsilon_{t-1} + \gamma x_{t-1}$ represents the error term from the viewpoint of the econometrician, who does not observe the specialist's *ex ante* order imbalance signal. (As x_t has mean zero, this presents no difficulty.) If the model is correctly specified, u_t has a conditional mean of zero given information at time $t-1$, so that $E[u_t|Z_{t-1}] = 0$, where Z_{t-1} is a r-dimensional vector of instrumental variables that are in the information set. 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 then choosing parameter values for θ that minimize a criterion function based on these r orthogonality conditions.

Specifically, if the vector of instruments is $Z'_{t-1} = (1, I_{t-1})$, the moment conditions for a single stock are represented by the (2×1) vector:

$$g_T(\theta) = \frac{1}{T}\sum_{t=1}^{T}\left[(I_t - I_{t-1} - \beta(I_{t-1} - I_d) \otimes Z_{t-1}\right]. \tag{14}$$

These conditions are analogous to the OLS normal equations. With N stocks, there will be $2N$ such orthogonality conditions. Hansen (1982) shows that these orthogonality conditions can be used to estimate the unknown parameters of the model and test the restrictions implied by theory.

The GMM estimates of θ are found by minimizing the quadratic criterion function:

$$J_T(\theta) = g_T(\theta)'W_T(\theta)g_T(\theta),$$

where $W_T(\theta)$ is the weighting matrix. If the number of orthogonality conditions equals the number of parameters to be estimated, as is the case of the unrestricted model, $J_T(\theta) = 0$ for all choices of the weighting matrix $W_T(\theta)$. However, if there are overidentifying restrictions, Hansen (1982) shows that 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:

$$W_T(\theta) = S^{-1}(\theta), \tag{15}$$

where $S(\theta) = E[f_t(\theta)f_t'(\theta)]$ and $f_t(\theta) = u_t \otimes Z_{t-1}$. Denote by $S_0(\theta)$ an estimator of the covariance matrix $S(\theta)$. Then, the asymptotic covariance matrix for the GMM estimator is given by:

$$\frac{1}{T}(D_0'(\theta)S_0^{-1}(\theta)D_0(\theta))^{-1}, \tag{16}$$

where:

$$D_0 = E\left[\frac{\partial u_t}{\partial \theta} \otimes Z_{t-1}\right] \tag{17}$$

is the Jacobian matrix, which is evaluated at the GMM parameter estimates.

Hansen (1982) proves that the GMM estimates, $\hat{\theta}$, are consistent and asymptotically normally distributed.[10] 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 orthogonality conditions, less the number of parameters to be estimated. Hansen proposes this as a test statistic for the goodness-of-fit of the model. The asymptotic covariance matrix provides standard errors used to test the significance of the parameters separately.

[10]The asymptotic results require only that the distribution of the dependent variable is stationary and ergodic.

Newey and West (1987) show that the weighting matrix can be adjusted to account for serial correlation and conditional heteroskedasticity, allowing a more general specification of the error term. With this adjusted weighting matrix, multivariate hypothesis tests are straightforward. In general, consider a null hypothesis, $H_0 : \alpha(\theta) = 0$, where $\alpha(\theta)$ is a vector of model restrictions of dimension k, and denote by $J_T(\theta^*)$ the minimized objective function under the parameter restrictions implied by the null hypothesis. Then, the test statistic:

$$R = T[J_T(\theta^*) - J_T(\hat{\theta})] \tag{18}$$

is asymptotically distributed χ^2 with k degrees of freedom. A high value of this measure suggests some form of model misspecification.

5 Analysis of Inventories

5.1 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 later on. Table II presents the GMM estimates of equation (8), with t-ratios based on Newey-West autocorrelation-heteroskedasticity consistent standard errors, following the procedure described in the previous section.[11] For all 16 stocks, the estimate of β is negative. The average value of β is about -0.05, and this is significantly negative for half the stocks. These estimates are not consistent with inventory effects taking place within the day since this would imply that $\beta = -1$, i.e., that inventories I_t are 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.

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

[11]As autocorrelation may be a problem, we used lagged inventory over the previous three days as additional instruments. The standard errors reflect a correction for conditional heteroskedasticity and for serial correlation up to three lags.

of cross- stock inventory effects. Following the procedure described above, the statistic R has a value of 106.074 under the joint restriction. This statistic has a χ^2 distribution with degrees of freedom equal to the number of restrictions imposed. The corresponding p-value is below 0.001, so that the null hypothesis of non-stationarity is strongly rejected. This result is consistent with Garman (1976), who shows theoretically that if dealer inventories were non-stationary, a market maker with finite capital would eventually become bankrupt with certainty.

We also estimated 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. 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.

5.2 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 (8) are quite weak. Another way to assess these effects is to estimate equation (8) 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 -0.0316, and the corresponding t-ratio, using autocorrelation-heteroskedasticity consistent standard errors based on the Newey-West (1987) procedure, was -1.355. The implied desired inventory level was \$33.23 million, with a t-ratio of 5.841. We cannot reject the null hypothesis of a unit root for dollar inventory. Moreover, the fit is extremely poor; the regression R^2 is just 0.015.

The absence of significant cross-stock inventory effects may appear surprising at first glance, but is consistent with the comments of NYSE specialists. In particular, observers familiar with the actions of specialists note that the most efficient method for a specialist to hedge his portfolio is through the use of futures contracts, rather than by using his individual stock positions to offset market risk. Also, the value of this specialist firm's stock portfolio may not be all that large relative to the total assets of the firm's owners.

5.3 Shifts in Desired Inventory

The GMM estimation indicates that specialist inventories are jointly stationary, but the adjustment process appears to be very slow, and the model captures only a small fraction of the variance in specialist trades. Indeed, the R^2 for the regressions in Table II is low, about 2% on average, suggesting that lagged inventories by themselves account for only a small portion of the day-to-day variation in specialist trades. The model provides a possible explanation for these findings. In particular, specialist trades may reflect periodic revisions in the desired inventory level, I_d, reflecting long-term investment strategy. Similarly, specialist trades also reflect short-run speculation based on signals about impending order imbalances. 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. 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. These observations are apparent in Figure 1, which shows the time-series for dollar inventory and for the inventory for every third stock.

To obtain consistent estimates of the inventory effect, we must take into account the potential for shifts in the mean of the inventory process. It would be easy to estimate a modified version of equation (8) that embodied changes in desired inventory or abnormal order flow shocks using dummy variables if the days on which these events occurred were known. However, neither the dates nor the types of these events is observable. Fortunately, a technique, known as *intervention analysis*, has been developed to address precisely 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.[12]

To develop a formal intervention model, suppose that revisions in desired inventory holdings occur on s different days in the estimation period and that desired inventory changes by α_i shares on day t_i when the i^{th} event occurs. Similarly, suppose that the error term in equation (8) consists of the original white noise process and a process that produces an error drawn from a distribution with a large standard deviation on p different days. Let e_j be the abnormal error on day t_j when the j^{th} shock occurs. The latter distribution captures the effect on inventories of large but infrequent order imbalances (x_t) as well as inventory shocks arising from block positioning or stabilization of transitory imbalances. The revised model is:

$$I_t - I_{t-1} = \beta(I_{t-1} - I_{t-1}^d) + u_t \tag{19}$$

$$I_t^d = \alpha_0 + \sum_{i=1}^{s} \alpha_i S_{it} \tag{20}$$

$$u_t = \epsilon_t + \sum_{j=1}^{p} e_t P_{jt}, \tag{21}$$

where $S_{it} = 1$ if $t \geq t_i$ and 0 otherwise and $P_{jt} = 1$ if $t \neq t_j$ and 0 otherwise. The type of intervention illustrated in equation (20) 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. By contrast, equation (21) illustrates an intervention known as an *Additive Outlier* (AO). Unlike a LS-type intervention, an additive outlier is a 'pulse' that affects the series for only a single period.

The idea behind estimating an intervention model of the type described by equations (19)–(21) is to examine the residuals associated with specific observations, unlike most time-series analyses which focus on the overall patterns of serial correlation. Intuitively, the estimated residuals should not display systematic patterns if the model is correctly

[12] For a description of these techniques see, for example, Fox (1972), Box and Tiao (1975), and Tsay (1986).

specified. Various kinds of misspecifications will be reflected in systematic patterns in the estimated residuals. For example, suppose in equation (20), that $s = 1$ and $t_1 = 98$, $\alpha_0 = 0$ and $\alpha_1 = 18,000$, implying that from day 99 onward, desired inventory increased by 18,000 shares. In this case, if we estimated equation (8) assuming a constant mean, the estimated residuals from day 99 onwards would tent to be positive, while the estimated residuals on days 1 through 98 would tend to be negative. If α is large, the timing of the shift in the mean can be identified by analyzing the time series of the model's residuals.

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

1. First, 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 pre-specified 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 i^{th} 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 re-estimated 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 applied Tsay's 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.[13] Conservative values were chosen for two reasons: first, we wished to estimate a parsimonious model, and second, as illustrated in Figure 1, the inventory process is subject to a few, relatively large shocks to inventory in the sample period.

Figure 2 illustrates the intervention totals for all stocks, grouped by month and type. Not surprisingly the rate at which the interventions occur increases greatly in the days immediately following the October crash. Of the 63 interventions identified for the individual stocks, 44 are of the LS-type, the remainder being AO-types. 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 finding 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 re-estimated equation (8), adding indicator variables based on the identified timings of the interventions. Table III presents the GMM estimates of the non-linear model:

$$I_t - I_{t-1} = \beta(I_{t-1} - \alpha_0 - \sum_{j=1}^{5} \delta_j D_{jt}) + u_t, \tag{22}$$

where D_{jt} is an indicator variable defined as follows. If outlier j is of the AO-type and occurred 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 occurred on day t_j, then $D_{jt} = 1$ for $t \geq t_j$ and 0 otherwise.

Table III should be compared with Table II. The mean-reversion coefficient β in equation (22) is negative in all 16 cases. The average estimate of β in Table III is -0.134 as opposed to an average of -0.05 for equation (8). The Newey-West standard errors of β are also much lower, which is consistent with our intuition that the regression equation (8) was misspecified.[14] Again, a multivariate χ^2 test of the null hypothesis $H_0 : \beta = 0$ is

[13] See Chang, Tiao, and Chen (1988) for further details of the implementation.

[14] With the inclusion of the D_{jt} indicator variables, we found no need to adjust for serial correlation. The standard errors do correct for conditional heteroskedasticity.

rejected; the p-value is below 0.001.

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. Most of the indicator variables have statistically significant coefficients. Fewer than five indicator variables are included in the regressions for some stocks if fewer than five outliers were identified. No interventions were estimated for stock 4. The average R^2 for this set of regressions is 0.18.

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. Another perspective on our results is provided by Hasbrouck and Sofianos (1992). Based on a spectral analysis of specialist profits for a sample of 144 stocks, they conclude that specialist positions 'are driven by rapid inventory adjustments towards targets that are themselves time-varying.'

5.4 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%, where:

$$h = -\frac{\ln(2)}{\ln(1+\beta)}. \tag{23}$$

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. It is worth noting that in the macroeconomic literature, commodity inventories are slow to adjust. The slow speed of adjustment for specialist inventories is puzzling because, unlike commodity inventories, the security positions can be rapidly changed at low cost, especially for the specialist.

6 Specialist Trades and Stock Returns

6.1 Information and Inventory Effects

Our approach to analyzing the specialist's behavior so far has focused on inventory dynamics. Equation (11) 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 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. The drawback to this approach is that it requires a model for the evolution of the specialist's beliefs over time. Specifically, equation (11) shows that the revision in prices is positively related to the change in the specialist's prior beliefs and inversely related to specialist trading. In turn, the revision in beliefs is positively related to the unanticipated order imbalance, which is an unobservable variable.

Let s_t denote the shock to order imbalances. An analysis of the estimated order 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 were pre-arranged in the so-called upstairs market and executed on the floor, and should be excluded from an analysis of specialist quote determination.

Accordingly, we constructed a time series of all intradaily large-block trades. We define a block buy (sell) for a particular stock as a buy order the top (bottom) 0.5% of the size distribution of intradaily buys (sells) for that stock over the entire 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 constructed 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:

$$q_t = \alpha_0 + \sum_{i=1}^{M} \beta_i q_{t-i} + \sum_{j=1}^{N} \gamma_j (p_{t-j} - p_{t-j-1}) + u_t. \tag{24}$$

In estimating these regressions (which are not reported here) we used $M = N = 3$. This procedure is suggested 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 we define order imbalances excluding large-block trades that are likely to have been negotiated in the upstairs market.

The regressions generally have low explanatory power, but using a F-test we can reject the hypothesis that the coefficients of the independent variables are zero. In general, there was small, but significant, positive autocorrelation in the q_t variable. The shock variable s_t is then defined as the residual in the regression, i.e., $s_t = q_t - E[q_t|\Phi_{t-1}]$. 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 directly from equation (11):

$$p_t - p_{t-1} = \beta_0 + \beta_1 s_t + \beta_2 (I_t - I_{t-1}) + \beta_3 \text{OCT19} + \varepsilon_t, \tag{25}$$

where p_t is the opening mid-quote 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, OCT19_t is a dummy variable that equals one on October 19, 1987 and 0 otherwise, and ε_t is the regression error term. Equation (11) implies that $\beta_0 = 0$,

$\beta_1 > 0$, and $\beta_2 < 0$. 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 $\beta_3 = 0$, the drop in prices on October 19 can be explained entirely by selling pressure on that day. The regression error term, ε_t, captures the unobservable term $x_t - x_{t-1}$ that reflects the specialist's signals regarding order imbalances.

We estimated the model using the GMM procedure, writing the moment conditions with instruments.[15] We also estimated 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 presents the GMM estimates of equation (25). The results provide strong support for the model. As hypothesized, the constant β_0 is generally close to zero. The information effect, as measured by β_1, is positive in 15 cases and is statistically significantly positive at the 5% level for 13 stocks using Newey-West standard errors. These results are especially strong because the innovation s_t is likely to be measured with error so that the coefficient β_1 is biased downward towards the null hypothesis.

The inventory parameter, as measured by β_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 is significant for 13 stocks, suggesting that selling pressure and inventory effects cannot entirely explain the drop in prices on October 19, 1987. The model fits remarkably well and the R^2 exceeds 0.25 for 10 stocks.

When the daily large-block volume is added to the regression equation, this variable has little or no significance suggesting that it is the non-block order flow that conveys information to the specialist. This appears surprising at first glance, but is consistent with the results of Madhavan and Smidt (1991) who find that the price impacts of large-

[15]The instruments include lagged values of inventory, order imbalances, and block trades. This procedure yields consistent parameter estimates when the explanatory variables may be correlated with the disturbance term.

block trades is bounded above. This result may also reflect the 'leakage' of information concerning the impending block transaction since negotiating a large-block trade in the upstairs market takes time. So, the information conveyed by large-block trades may already have been impounded in security prices before these trades occur. We turn now to a further investigation of the information available to the specialist and the uses to which this information is put.

6.2 Price Dynamics and the Specialist's Information

In this section, we examine whether the revision in quotes reflects information about future order imbalances, as suggested by our model. To test this hypothesis, we estimated the following regression:

$$p_t - p_{t-1} = \beta_0 + \beta_1 q_{t+1} + \beta_2 s_t + \beta_3(I_t - I_{t-1}) + \beta_4 \text{OCT19} + \varepsilon_t. \qquad (26)$$

This equation is identical to equation (25) except that it includes q_{t+1}, i.e., the leading order imbalance (net of large-block trades) as an independent variable. The regression estimates, again using the GMM procedure, are reported in Table VI. It is interesting to compare this table with Table V. First, our previous conclusions regarding the inventory and information effects are largely unaltered. Second, the coefficient of leading order imbalances is positive in all 16 cases and is significantly positive in 10 of these cases. This is consistent with the recurrent argument that the specialist's privileged position (which includes access to the limit order book) provides a short-run informational advantage relative to some traders.

Another perspective on this issue is provided by examining the causal relation between quotation revisions and specialist trades. Granger causality is a purely statistical concept — x is said to *Granger cause* y if knowledge of past values of x enables better predictions about y, other things equal. Using a Wald test proposed by Geweke, Meese, and Dent (1983), we reject the null hypothesis that price changes do not Granger cause inventory changes in 12 of the 16 cases. Repeating the Wald tests in the other direction we rejected the null hypothesis that inventory does not cause quote revisions in only two of the 16 cases. The results indicate that quote revisions anticipate future inventory movements, but not vice versa. We interpret the result that price changes appear to Granger-Sims cause inventory changes as additional support for our model where the specialist's information about future imbalances affects current price quotations. The apparent lack of causality running from inventory changes to subsequent quote changes may reflect long-term speculation by the specialist that is not detected at short horizons. These results are also consistent with other explanations, such as the specialist altering quotations to induce informative trade, thereby expediting the process of price discovery.[16]

7 Conclusions

This paper analyzes dealer behavior using data obtained from a NYSE specialist. We develop a theoretical model of trading where all agents take actions to maximize their expected utility. The model differs from previous work in that it incorporates formally both asymmetric information and inventory control effects. The specialist is shown to act not just as a market maker, but also as an active investor. Finally, the specialist's information about liquidity-based trading is made explicit and is shown to have an important influence on price formation.

The model yields a number of testable hypotheses. We focus on two representations of the specialist's optimal trading strategy. The first representation shows that inventories in the stocks assigned to the specialist unit may exhibit mean reversion to target levels that themselves may shift in response to the changes in the composition and risk characteristics of other assets or liabilities owned by the specialist. The specialist's trades also reflect forecasts of short term order imbalances. The second representation shows that stock price changes are positively related to the innovation in order flow and are inversely related to specialist trades. Further, any information on future order imbalances should be reflected in a current quote revision. While these representations are equivalent theoretically, they

[16]See, e.g., Leach and Madhavan (1992). Unlike other asymmetric information models such as Kyle (1985), the price experimentation hypothesis implies that quote revisions are associated with nonzero expected order flow.

differ from an empirical viewpoint, and we use both approaches. Throughout, we use Hansen's (1982) Generalized Method of Moments as our estimation technique.

Beginning with inventories, when specialist desired inventories are assumed constant, we find weak support for the mean reversion hypothesis. On this assumption it appears to take over 49 trading days, on average, for an imbalance in inventory to be reduced by 50%. Desired inventories are uniformly positive, and a multivariate GMM test suggests the system is jointly stationary. 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. 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 the predictions of our model. We find that unanticipated (non-block) order imbalances convey signals to the specialist regarding future price movements. 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 paper suggests that the specialist plays a far more complex role in price formation than previously thought. As expected, the specialist is an market maker who adjusts quoted prices to control fluctuations in inventory. The specialist is also an active investor for his own account, a role that was not suspected until recently. As an active investor, specialists may hold large positions, and may periodically adjust the size of these positions based on relatively long-term considerations that reflect his whole portfolio, including personal assets that are not part of the specialist firm.

Our analysis raises several questions that lie beyond the scope of this paper: What considerations underlie the shifts in target inventory levels, and do these movements anticipate future changes in fundamentals? Do predictable patterns in security prices induced by mean reversion in inventory give rise to possible arbitrage opportunities or provide a rationale for the persistence of technical trading strategies? What is the nature of the specialist's information about impending order flow, and how much of this information is public? These are topics for future research.

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.

Forster, M. M., and T. J. George, 1991, "Anonymity in Securities Markets," Working paper, The Ohio State University.

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.

Geweke, J., R. Meese, and W. Dent, 1983, "Comparing Alternative Tests of Causality in a Temporal System: Analytical Results and Experimental Evidence," *Journal of Econometrics*, 21, 161–194.

Glosten, L., and L. Harris, 1988, "Estimating the Components of the Bid-Ask Spread," *Journal of Financial Economics*, 21, 123–142.

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, United Kingdom.

Hasbrouck, J., 1988, "Trades, Quotes, Inventories and Information," *Journal of Financial Economics*, 22, 229–252.

Hasbrouck, J., 1991, "Measuring the Information Content of Stock Trades," *Journal of Finance*, 46, 178–208.

Hasbrouck, J., and G. Sofianos, 1992, "The Trades of Market-Makers: An Analysis of NYSE Specialists," Working paper, New York Stock Exchange.

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.

Kyle, A., 1985, "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1335.

Leach, J. C., and A. Madhavan, 1992, "Price Experimentation and Security Market Structure," Working paper, University of Pennsylvania.

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.

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.

Stoll, H., 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, IL.

Working, H., 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, IL.

Appendix

Derivation of Equation (2):

For trader i, who trades z_{it} at price p_t in period t, the final period wealth is given by:

$$\widetilde{W}_\tau = \tilde{v}_\tau(z_{it} + E_{it}) + C_{it} - p_t z_{it}, \tag{A.1}$$

where E_{it} represents the trader's endowment of the risky asset, C_{it} the initial cash holding, and $\tilde{v}_\tau$ is the value of the asset at date τ. Since the liquidation date is random, we need to derive an expression for the variance of final period wealth to solve the maximization problem. Observe that at time t, given that an informed trader observes the fundamental price at the time of the trade, the asset's value at time τ is a random variable, $\tilde{v}_\tau$, where $\tilde{v}_\tau = v_t + \tilde{\eta}_{t+1} + \cdots + \tilde{\eta}_\tau$. Thus, at the time of the trade at time t, the conditional variance of $\tilde{v}_\tau$ is:

$$\sigma^2[\tilde{v}_\tau] = \sum_{k=1}^{\infty} \sigma^2[\tilde{v}_{t+k}|\tau = k] \Pr[\tau = t+k] = \sum_{k=1}^{\infty}(1-\rho)\rho^{k-1}k\sigma_\eta^2 = \frac{\sigma_\eta^2}{1-\rho}. \tag{A.2}$$

For trader i at time t, $E[\tilde{\eta}_{t+j}] = 0$, for $j > 0$, implying that the conditional expectation of v_τ at time t is just v_t. Using (1), investor's quantity, z_{it}, solves:

$$\max_{\{z_{it}\}} \left\{ E[v_\tau|\Phi_{it}](z_{it} + E_{it}) + C_{it} - p_t z_{it} - \frac{\omega\sigma_\eta^2}{(1-\rho)}(z_{it} + E_{it})^2 \right\}. \tag{A.3}$$

The first-order conditions for equation (A.3) yield:

$$z_{it} = \frac{(1-\rho)(v_t - p_t)}{2\omega\sigma_\eta^2} - E_{it}. \tag{A.4}$$

The excess demand in period t is obtained by adding the demands of the N participating traders, i.e., $z_t = \sum_i^N z_{it}$. From equation (A.4), we see that z_t takes the form given by equation (2) where:

$$\delta = \frac{N(1-\rho)}{2\omega\sigma_\eta^2} \tag{A.5}$$

$$X_t = -\sum_i^N E_{it}. \tag{A.6}$$

Derivation of Equation (8):

Using equation (5), the specialist's wealth at the start of period t is a random variable with conditional expectation:

$$E[\widetilde{W}_t] = \mu_t I_t + K_{t-1} + p_{t-1}z_{t-1}, \tag{A.7}$$

where we assume, without loss of generality, that $E[\tilde{y}] = 0$. Similarly, the variance of the specialist's wealth at the start of trading on day t can be expressed as:

$$\sigma^2[\widetilde{W}_t] = \sigma_v^2 I_t^2 + \sigma_y^2 + 2I_t\sigma_{yv}, \tag{A.8}$$

where σ_y^2 is the variance of $\tilde{y}_t$, σ_{yv} 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. We will derive a closed-form representation for this variance later on. It follows that:

$$\sigma^2[\widetilde{W}_t] = \phi_0 + \phi_1(I_t - I_d)^2, \tag{A.9}$$

where:

$$\phi_0 = \sigma_y^2 - \left(\frac{\sigma_{yv}}{\sigma_v}\right)^2 \tag{A.10}$$

$$\phi_1 = \sigma_v^2 \tag{A.11}$$

$$I_d = \frac{-\sigma_{yv}}{\sigma_v^2} \tag{A.12}$$

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

Applying the Bellman principle, the functional equation equivalent to the optimality equation associated with the maximization problem in equation (7) is:

$$\begin{aligned} J(I_t, x_t, \mu_t, K_t) &= \max_{\{p_t\}} E\Big\{(1-\rho)\Big[E[\widetilde{W}_t] - \omega_s(\phi_0 + \phi_1(I_t - I_d)^2)\Big] + \\ &\quad \rho J(\tilde{I}_{t+1}, \tilde{x}_{t+1}, \tilde{\mu}_{t+1}, \widetilde{K}_{t+1}) \;\Big|\; \Phi_t^s\Big\}. \end{aligned} \tag{A.13}$$

Intuitively, the functional equation consists of two parts: The first term in brackets in the functional equation is the expected utility if the event occurs in the current period, i.e., at time t, (note that by convention there is no trading in period t itself if $\tau = t$) times 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 the functional equation is the expected utility if the event does not occur at time t, multiplied by the probability that the event does not occur in the current period, and represents the continuation value if there is trading in period t.

For now, we suppress subscripts to minimize the notational burden. Define by $\widetilde{I}'$ the inventory position in the next period, and similarly define $\tilde{x}', \tilde{\mu}', \widetilde{K}'$. Before deriving the solution to the maximization problem, we need to formally express the transition equations governing the evolution of the state variables, i.e., inventories, beliefs, order imbalance signals, and capital:

$$E[\tilde{I}'|\Phi^s] = I - \delta(\mu - p) - x \tag{A.14}$$

$$E[\tilde{x}'|\Phi^s] = 0 \tag{A.15}$$

$$E[\tilde{\mu}'|\Phi^s] = \mu \tag{A.16}$$

$$E[\widetilde{K}'|\Phi^s] = K + p(\delta(\mu - p) + x). \tag{A.17}$$

These transition equations are interpreted as follows: Equation (A.14) states that expected inventories are current inventories less expected demand, equation (A.15) states that the expected order imbalance is zero, equation (A.16) is the Law of Iterated Expectations and implies that the revision in beliefs is an innovation, and equation (A.17) states that capital increases with expected trading income. The specialist's price choice affects his utility in future periods through its effect on future inventory and on his capital position.

Let I' denote the specialist's expected inventory position where, from equation (A.14), I' is linearly related to the specialist's quoted price, p. It is easiest to solve the dynamic programming problem by finding the specialist's optimal choice of I'. This solution yields a policy function, denoted by $g(I, x, \mu, K)$, that relates the expected inventory to the state variables. Once we have obtained the function g, we can then derive the implied optimal price quotation function using equation (A.14).

The first step to solving the problem is to substitute the transition relations (A.14)–(A.17) into equation (A.13), taking care to express the choice variable in terms of expected inventory

I', not the price p. This is straightforward as $p = \mu - \delta^{-1}(I - I' - x)$, and $\tilde{I}' = I' + \epsilon$, where ϵ is a mean zero error term.

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 (A.13) with respect to I'. This yields:

$$E[J_1(\tilde{I}', \tilde{x}', \tilde{\mu}', \widetilde{K}')] - E\left[J_4(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}')\left(\mu + \frac{2(I' - I) + x}{\delta}\right)\right] = 0. \qquad (A.18)$$

The envelope conditions are found by differentiating the value function (A.13) with respect to each of the four state variables. These yield:

$$\begin{aligned}
J_1(I, x, \mu, K) &= (1 - \rho)(\mu - 2\omega_s\phi_1(I - I_d)) + \\
&\quad \rho E[J_4(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{W}')(\mu + \delta^{-1}(2(I' - I) + x))] && (A.19)\\
J_2(I, x, \mu, K) &= -\frac{\rho}{\delta}E[J_4(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}')(I' - I)] && (A.20)\\
J_3(I, x, \mu, K) &= (1 - \rho)I + \rho E[J_3(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}')] + \\
&\quad \rho E[J_4(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}')(I - I')] && (A.21)\\
J_4(I, x, \mu, K) &= (1 - \rho) + \rho E[J_4(\tilde{I}', \tilde{x}', \tilde{\mu}', \tilde{K}')]. && (A.22)
\end{aligned}$$

The linear-quadratic structure of the maximization problem suggests that $J(I, x, \mu, K)$ has a quadratic form given by $J(I, x, \mu, K) = A_0 + \mu I + K + A_1(I - I_d)^2 + A_2 xI + A_3 x^2 + A_4 x$, where A_i $(i = 0, \ldots, 4)$ are constants.

Equations (A.19) and (A.20) then imply that the optimal policy $g(I, x, \mu, K)$ is a linear function given by $g(I, x, \mu, K) = I + \beta(I - I_d) + \gamma x$, so that J indeed has the conjectured form. It remains to show that β and γ exist. Differentiating J, we obtain $J_1 = \mu + 2A_1(I - I_d) + A_3 x$. Using the fact that $E[\mu'] = \mu$ and $E[x'] = 0$, we see that $E[J_1] = \mu + 2A_1(I' - I_d)$. Using this expression and the conjectured policy function in equation (A.18) and equating the coefficients of $(I - I_d)$ and x respectively, we obtain $A_1 = \beta/[(1 + \beta)\delta]$ and $A_1 = (2\gamma + 1)/(2\delta\gamma)$. Similarly, setting $J_1 = \mu + 2A_1(I - I_d) + A_3 x$ in equation (A.20), and equating coefficients yields $A_1 = \rho\beta\delta^{-1} - (1 - \rho)\omega_s\phi_1$. Combining these expressions for A_1, we obtain:

$$\begin{aligned}
\beta &= \frac{-\delta\omega_s\phi_1}{1 + \delta\omega_s\phi_1} && (A.23)\\
\gamma &= \frac{-(1 + \beta)}{2}. && (A.24)
\end{aligned}$$

So, $\beta \in (-1, 0)$ and $\gamma \in (-\frac{1}{2}, 0)$. In the special case where all agents have the same coefficient of risk aversion (i.e., $\omega = \omega_s$), $\beta = \frac{-1}{1+\left(\frac{2\sigma_\eta^2}{N(1-\rho)\sigma_v^2}\right)}$, which is independent of the risk aversion parameter.

Derivation of Equation (11):

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 (z_t, x_t) is equivalent to observing w_t where

$$w_t = \frac{z_t - x_t + \delta p_t}{\delta} = v_t + \frac{\xi_t}{\delta}.$$

Using this definition and equation (3), we see that $w_t = v_t + \xi_t/\delta$. Since ξ_t has a zero mean, w_t is an unbiased signal about the unobserved fundamental price. Let σ^2 denote the variance of this signal, i.e., the variance of $\tilde{\xi}_t/\delta$. Let $\theta = \sigma_\eta^2/\sigma^2$ represent the signal-to-noise ratio; a higher value of θ implies more order flow 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 information before the market opens on day t and recall that σ_v^2 denotes the conditional variance associated with this forecast. It is convenient to express this conditional variance in the form: $\sigma_v^2 = \sigma^2\Omega$, where σ^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:

$$\mu_t = \mu_{t-1} + \left[\frac{\Omega + \theta}{\Omega + \theta + 1}\right](w_{t-1} - \mu_{t-1}), \tag{A.25}$$

where:

$$\Omega = \frac{-\theta + \sqrt{\theta^2 + 4\theta}}{2}. \tag{A.26}$$

Equation (A.25) shows that expected value of the security can be represented in 'error correction' form, while equation (A.26) provides an expression for Ω, which represents the conditional variance of this 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. Then, it is straightforward to show that equation (A.25) can be written as:

$$\mu_t = \Omega w_{t-1} + (1 - \Omega)\mu_{t-1}. \tag{A.27}$$

Thus, Ω also has an interpretation as the steady-state weight placed on the signal about the unobservable state. Intuitively, as the signal-to-noise ratio, θ, increases, Ω increases and less weight is placed on prior beliefs.

Using equation (A.27) and the definition of w_t, we can write the revision in beliefs as: $\mu_t - \mu_{t-1} = \Omega(v_{t-1} - \mu_{t-1} + \xi_{t-1}/\delta)$. This decomposition shows that the revision in beliefs is proportional to the unanticipated component of order flow in the last trading round. Then, equation (11) follows directly from (10).

Table I

Descriptive Statistics on Specialist and Market Transactions for 16 NYSE Stocks

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

Stock Number	Closing Inventories		Change in Specialist's Inventories		Value of Transactions		Number of Transactions	
	$1,000's	Round Lots	$1,000's	Round Lots	Specialist $1,000's	Market $1,000's	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

Table II

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

Generalized Method of Moments Estimates of the model:

$$I_t - I_{t-1} = \beta(I_{t-1} - I_d) + \epsilon_t$$

where I_t represents the opening inventory on day t. The figures in parentheses are t-statistics based on autocorrelation-heteroskedasticity consistent standard errors computed using the procedure of Newey and West (1987).

Stock	$\hat{\beta}$	$\hat{I}_d \times 10^{-3}$	$\bar{R}^2$
1	−0.017 (−1.20)	7.77 (1.68)	0.010
2	−0.008 (−0.46)	198.43 (3.62)	0.001
3	−0.020 (−2.01)	72.51 (3.32)	0.006
4	−0.030 (−1.82)	0.25 (0.01)	0.000
5	−0.115 (−3.46)	10.81 (3.44)	0.064
6	−0.035 (−1.78)	211.91 (10.03)	0.019
7	−0.006 (−0.81)	166.74 (2.07)	0.005
8	−0.076 (−2.96)	14.53 (2.38)	0.037
9	−0.019 (−1.90)	225.98 (7.56)	0.005
10	−0.012 (−0.69)	91.52 (0.66)	0.005
11	−0.069 (−2.35)	20.45 (2.88)	0.028
12	−0.065 (−2.58)	58.23 (2.65)	0.027
13	−0.047 (−2.33)	24.75 (3.43)	0.021
14	−0.036 (−2.50)	97.20 (5.11)	0.022
15	−0.002 (−0.23)	790.94 (0.24)	0.000
16	−0.125 (−3.21)	39.30 (2.85)	0.055

Table III

Mean Reversion in Inventories Corrected for Interventions

The table presents Generalized Method of Moment estimates of the model:

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

where I_t is the opening inventory on day t (scaled by 10^{-5}), and D_{jt} is an indicator variable for intervention j on day t_j. The figures in parentheses are t-values computed from standard errors obtained using the Newey and West (1987) procedure.

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-type intervention; all the remaining are LS-type interventions.

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:

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

where β is the mean-reversion parameter. Two figures are reported, based on the the Generalized Method of Moments 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-type 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	7.334
σ	78.16	5.200

Table V

Estimates of the Model of Quote Revisions Using Generalized Method of Moment Estimators

The table presents coefficient estimates of the model:

$$p_t - p_{t-1} = \beta_0 + \beta_1 s_t + \beta_2(I_t - I_{t-1}) + \beta_3 \text{OCT19}_t + \varepsilon_t,$$

where, on day t, p_t is the opening mid-quote, s_t is the unanticipated order imbalance from the previous day, computed as the residual from a VAR of non-block order imbalances on lagged order imbalances and lagged quote revisions , I_t is the opening inventory, OCT19_t is a dummy variable that equals 1 on October 19, 1987 and 0 otherwise. Figures in parentheses are t-statistics using Newey-West standard errors.

Stock	$\hat{\beta}_0$	$\hat{\beta}_1 \times 10^5$	$\hat{\beta}_2 \times 10^5$	$\hat{\beta}_3$	$\bar{R}^2$
1	−0.02 (−1.05)	6.82 (3.22)	−6.11 (−2.40)	−0.94 (−2.98)	0.288
2	−0.04 (−1.90)	1.50 (2.74)	−2.42 (−3.56)	−1.33 (−4.22)	0.199
3	−0.01 (−0.50)	1.43 (3.09)	−0.64 (−1.47)	−4.10 (−7.58)	0.321
4	0.01 (0.15)	−0.20 (−0.21)	−5.57 (−4.31)	0.18 (0.32)	0.186
5	−0.00 (−0.03)	0.10 (0.98)	−1.80 (−3.68)	−0.98 (−2.27)	0.099
6	−0.03 (−1.91)	0.22 (2.32)	−0.54 (−4.35)	−0.26 (−1.02)	0.130
7	−0.00 (−0.10)	3.84 (4.40)	−4.84 (−5.65)	−8.55 (−8.76)	0.509
8	0.06 (0.67)	4.32 (4.75)	−7.33 (−4.90)	−6.27 (−4.48)	0.341
9	0.09 (1.88)	2.68 (5.24)	0.82 (1.21)	−7.72 (−10.60)	0.470
10	0.11 (1.90)	0.45 (2.49)	0.18 (0.64)	−3.67 (−4.36)	0.106
11	−0.02 (−0.86)	1.48 (6.23)	−0.63 (−2.07)	0.08 (0.22)	0.226
12	−0.01 (−0.19)	1.09 (3.48)	−1.37 (−3.84)	−3.89 (−4.41)	0.371
13	−0.00 (−0.57)	0.91 (3.91)	−0.62 (−1.76)	−1.33 (−5.46)	0.267
14	−0.00 (−0.26)	0.45 (1.57)	−1.26 (−3.09)	−3.49 (−7.92)	0.320
15	−0.02 (−0.35)	0.80 (2.97)	−0.20 (−0.91)	−7.33 (−7.52)	0.334
16	−0.02 (−0.74)	0.28 (3.62)	−0.43 (−4.09)	−2.04 (−4.68)	0.258

Table VI

Regressions of Daily Quote Revisions on Leading (Non-Block) Order Imbalances, Unanticipated Order Flow, and Specialist Trades

Generalized Method of Moments estimates (with t-values based on Newey-West standard errors in parentheses) of the regression equation:

$$p_t - p_{t-1} = \beta_0 + \beta_1 q_{t+1} + \beta_2 s_t + \beta_3 (I_t - I_{t-1}) + \beta_4 \mathrm{OCT19}_t + \varepsilon_t,$$

where p_t is the opening mid-quote on day t, q_{t+1} is the order imbalance (net of block trades) over day t, s_t is the unanticipated order imbalance over day $t-1$, I_t is the opening inventory on day t, $\mathrm{OCT19}_t$ is a dummy variable for October 19, 1987.

Stock	$\hat{\beta}_0$	$\hat{\beta}_1 \times 10^5$	$\hat{\beta}_2 \times 10^5$	$\hat{\beta}_3 \times 10^5$	$\hat{\beta}_4$	$\bar{R}^2$
1	−0.02	3.25	7.24	−5.16	−1.03	0.303
	(−0.94)	(2.22)	(3.44)	(−2.02)	(−3.26)	
2	−0.02	1.00	1.38	−2.36	−1.27	0.211
	(−1.21)	(1.96)	(2.51)	(−3.49)	(−4.04)	
3	−0.02	0.67	1.30	−0.68	−4.08	0.328
	(−0.61)	(1.69)	(2.77)	(−1.56)	(−7.57)	
4	−0.01	1.46	−0.48	−5.60	0.28	0.203
	(−0.24)	(1.72)	(−0.48)	(−4.38)	(0.51)	
5	−0.00	0.01	0.10	−1.79	−0.97	0.095
	(−0.02)	(0.10)	(0.95)	(−3.63)	(−2.23)	
6	−0.02	0.22	0.18	−0.60	−0.19	0.153
	(−1.13)	(2.44)	(1.89)	(−4.79)	(−0.75)	
7	0.02	1.69	3.60	−4.62	−8.16	0.519
	(0.35)	(2.22)	(4.14)	(−5.41)	(−8.32)	
8	0.04	1.81	3.84	−7.23	−5.83	0.355
	(0.43)	(2.22)	(4.14)	(−4.88)	(−4.17)	
9	0.05	1.63	2.16	0.66	−7.36	0.513
	(1.11)	(4.13)	(4.26)	(1.01)	(−10.45)	
10	0.01	0.89	0.38	0.46	−3.25	0.257
	(0.22)	(6.21)	(2.30)	(1.76)	(−4.22)	
11	−0.01	0.36	1.39	−0.63	0.05	0.234
	(−0.76)	(1.76)	(5.73)	(−2.08)	(0.14)	
12	−0.00	0.37	0.99	−1.47	−3.64	0.376
	(−0.15)	(1.56)	(3.09)	(−4.07)	(−4.07)	
13	−0.00	0.50	0.79	−0.77	−1.31	0.285
	(−0.56)	(2.39)	(3.38)	(−2.17)	(−5.47)	
14	0.00	0.70	0.45	−1.17	−3.47	0.343
	(0.14)	(2.76)	(1.61)	(−2.90)	(−8.01)	
15	−0.00	0.32	0.75	−0.17	−7.18	0.338
	(−0.07)	(1.47)	(2.78)	(−0.80)	(−7.35)	
16	−0.00	0.26	0.24	−0.41	−2.11	0.313
	(−0.25)	(3.97)	(3.28)	(−4.01)	(−5.01)	

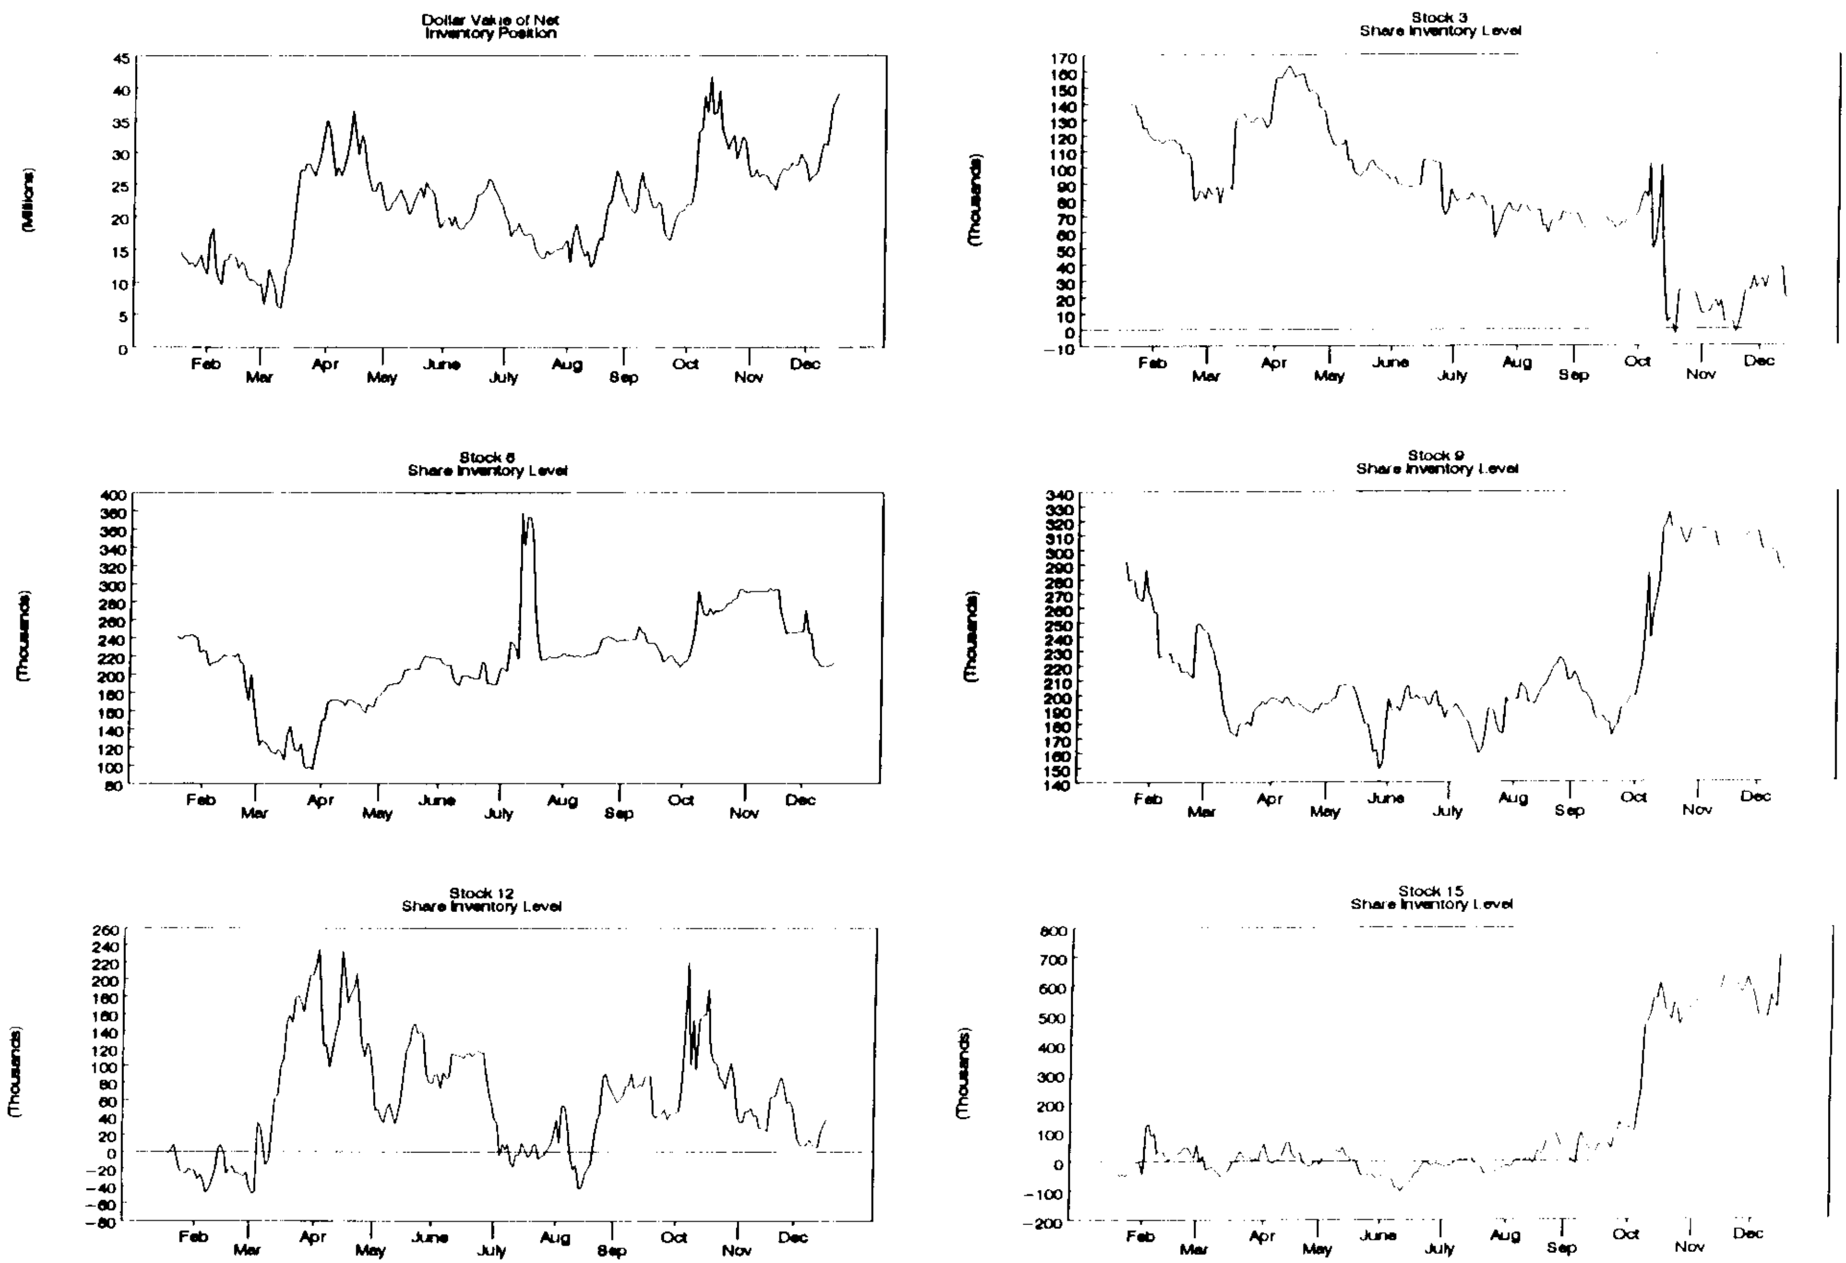

Figure 1: Inventory Levels for Dollar Inventory and Selected Individual Stocks, February-December, 1987.

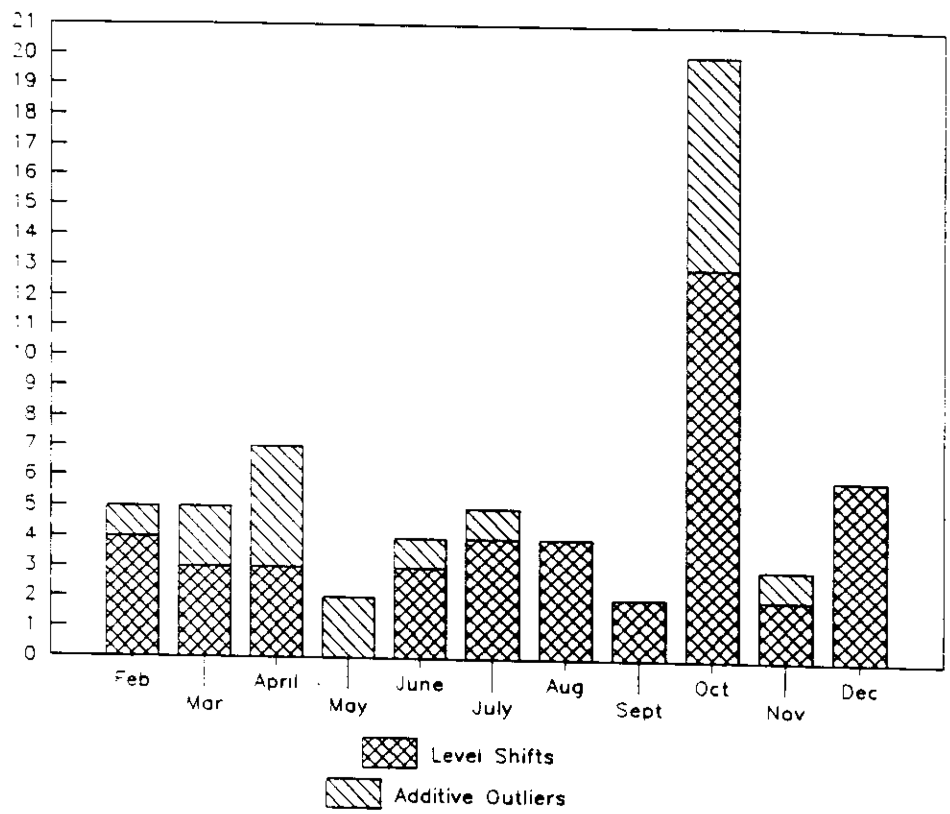

Figure 2: Intervention Totals for All Stocks, by Month and Type