



American Finance Association

The Trades of Market Makers: An Empirical Analysis of NYSE Specialists

Author(s): Joel Hasbrouck and George Sofianos

Source: *The Journal of Finance*, Vol. 48, No. 5 (Dec., 1993), pp. 1565-1593

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/2329060>

Accessed: 25/11/2009 11:39

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

The Trades of Market Makers: An Empirical Analysis of NYSE Specialists

JOEL HASBROUCK and GEORGE SOFIANOS*

ABSTRACT

This paper presents a transaction-level empirical analysis of the trading activities of New York Stock Exchange specialists. The main findings of the analysis are the following. Adjustment lags in inventories vary across stocks, and are in some cases as long as one or two months. Decomposition of specialist trading profits by trading horizon shows that the principal source of these profits is short term. An analysis of the dynamic relations among inventories, signed order flow, and quote changes suggests that trades in which the specialist participates have a higher immediate impact on the quotes than trades with no specialist participation.

THE IMPORTANCE OF LIQUIDITY in securities markets motivates strong interest in the trading behavior of dealers. Dealers seek to provide liquidity in an ongoing manner and they can influence the short-run dynamics of securities prices through their trading behavior. Not surprisingly, many academic studies examine dealers' trading activities and their role in price determination. Because of the difficulty in obtaining detailed data, most (but not all) of this work is theoretical. Several important empirical issues, therefore, remain unresolved. This paper analyzes inventory adjustment, price determination, and trading profits for one important class of dealers, New York Stock Exchange (NYSE) specialists.¹

* Hasbrouck is from the New York Stock Exchange and the Stern School of Business, New York University, and Sofianos is from the New York Stock Exchange. For comments on an earlier draft, we are indebted to René Stulz, the anonymous referee, David Chapman, Joseph Cherian, Eric Hughson, Ananth Madhavan, Eric Sirri, Seymour Smidt, Jim Shapiro, Deborah Sosebee, and seminar participants at the Board of Governors of the Federal Reserve System, Boston College, the Commodities Futures Trading Commission, Cornell University, the Federal Reserve Bank of New York, University of Texas at Austin, the Western Finance Association, and the European Finance Association. All errors are our own responsibility. The comments and opinions contained in this paper are those of the authors only. In particular, the views expressed here do not necessarily reflect those of the directors, members, or officers of the New York Stock Exchange, Inc.

¹ This paper uses the term "dealer" interchangeably with "market maker." Although this conforms to customary usage in the finance literature, the term "broker-dealer" in regulatory matters is considerably broader, encompassing all agents who trade directly on an exchange for their own account.

In the microstructure literature, there are two principal models of dealer trading behavior. Asymmetric information models focus on the adverse selection exposure of any trader who presents a bid or offer to a market consisting of indistinguishable informed and uninformed traders.² Inventory control models consider the problem of keeping securities holdings within certain bounds determined by the dealer's risk aversion. In practice, a dealer is likely to confront both problems. This study concentrates on the inventory control paradigm without ignoring the adverse selection problem faced by dealers.

The dealer's inventory control paradigm is based on the optimization models of classical economics. In microstructure analysis, however, the dealer controls his inventory not through the timing and quantity of a restocking order, but by actively adjusting his prices (quotes). In an early theoretical analysis, Garman (1976) succinctly asserts the necessity of active price control: in its absence a dealer who neutrally sets prices to balance average order inflow will suffer nonstationary inventories.³ In this literature, therefore, the dealer's inventory control problem becomes a pricing problem. Amihud and Mendelson (1980) solve the pricing problem of a monopolistic dealer, while Ho and Stoll (1983) solve the dealer's pricing problem when there are multiple dealers. O'Hara and Oldfield (1986) consider the effects of day boundaries on the dealers pricing problem.⁴

Perhaps due to the difficulty of obtaining detailed data, there are few empirical studies of the dealer's inventory control mechanism. In a U.S. Securities and Exchange Commission study (1971), Smidt presents average inventory levels and same-day inventory/price change correlations for a sample of NYSE specialists, but does not attempt to characterize the dynamic adjustment process. Stoll (1976) studies daily net dealer purchases and sales for a sample of National Association of Securities Dealers Automated Quotation (NASDAQ) stocks; Stoll (1985) estimates a similar specification for NYSE stocks. Ho and Macris (1984) provide the first transaction level study of dealer behavior.

In general, this early work neglects asymmetric information and attributes the entire dynamic relation between trades and prices to inventory control. At best, the asymmetric information effect is modeled as a static component of the bid-ask spread. This approach is unsatisfactory because in the short run both effects can have a similar impact on prices. The price increase following a buy from a dealer, for example, may be due to asymmetric information: the dealer revises the conditional expectation of the security value upwards based on the private information inferred from the trade. The initial price increase may also be the result of inventory control, due to the dealer's desire to induce a rebalancing sale.

² See Bagehot (1971), Copeland and Galai (1983), Kyle (1985), Glosten and Milgrom (1985), Easley and O'Hara (1987), Glosten (1987, 1989), and Admati and Pfleiderer (1988).

³ This point is implicit in many of the early studies of the dealer's inventory control problem.

⁴ Other relevant works include Bradfield (1979), Smidt (1979), Zabel (1981), and Son (1991).

Hasbrouck (1988) suggests that the two effects can be separated in a dynamic analysis. The inventory control component is transient: the price increase is matched by a decrease after the inventory rebalancing is complete. If the dealer inventory is initially at the optimal level, for example, an incoming buy order depletes it. The dealer raises the quotes to attract incoming sell orders. Once sell orders have arrived and the inventory position is restored, the quotes revert to their previous level. The asymmetric information component, on the other hand, is permanently impounded into the stock price. Consistent with the inventory control mechanism, Hasbrouck finds evidence that buyer-initiated trades are likely to be reversed by seller-initiated trades. These trade reversals, however, occur over long (multiday) horizons. The short-run trade autocorrelations suggest no trade reversal, and there is no evidence of the transient price overshooting predicted by the inventory control hypothesis. Stoll (1989), Madhavan and Smidt (1991, 1993), and Neuberger (1992) also try to distinguish asymmetric information and inventory control effects. The Madhavan and Smidt studies are particularly noteworthy because they develop theoretical models that incorporate both effects and they test these models using data sets that identify specialist inventories and trades. The earlier paper analyzes transaction data and finds no short-run price effect due to inventory control. Madhavan and Smidt (1993) analyze daily data and find that inventory changes do have an effect on prices, which in their specification is consistent with active inventory control.

This paper is organized as an exploratory data analysis that uses general statistical models to investigate the broad features of dealer trading behavior. The objective of the analysis is to establish empirical regularities against which the general predictions of the theoretical models can be assessed; the paper does not attempt to estimate any particular theoretical model. The paper first examines the time horizon of inventory mean reversion. Inventory autocorrelations are positive and persistent over long lags (in some cases, weeks) suggesting slow inventory adjustments. Evidence from various specialist activity measures and from incidents involving large exogenous inventory shocks, however, suggests that specialists are capable of rapid inventory adjustments. The observed long-term fluctuations in specialist stock holdings, therefore, are the result of changes in long-term investment positions. Madhavan and Smidt (1993) reach a similar conclusion and propose a time series model with random level shifts.

The paper then turns to the nature of specialist trading profits. Because of barriers to entry and various presumed informational advantages, one may infer that specialist activities are highly profitable. In fact, estimates of gross trading profits are found to vary considerably across stocks and are often statistically indistinguishable from zero. It appears that this low statistical power results from the long-term investment positions alluded to earlier. When profits are decomposed according to trading horizon, a more compelling picture emerges: the short- and medium-term components of trading profits

are strongly positive and often statistically significant. These findings suggest that NYSE specialists are good short-term traders but undistinguished long-term speculators.

Finally, the paper uses vector autoregressions to examine the joint behavior of quotes, inventories, and trade variables (functions of trade volume, signed as buyer or seller initiated on the basis of the prevailing quote). There is strong evidence that unexpected shocks to the specialist's inventory affect quotes, at least in the short run. Since inventory shocks are (inversely) correlated with the signed trade variables, this finding is consistent with asymmetric information effects: following a trade the dealer revises the expected value of the security based on the private information inferred from the trade. The classic inventory control mechanism also predicts transient quote overshooting. The paper does not find widespread evidence of this phenomenon.

The paper is organized as follows. Section I discusses the institutional framework and the data. Section II analyzes the univariate time series properties of inventories. Section III presents estimates of trading profits and the allocation of these profits across different horizons. Section IV discusses the dynamic behavior of inventories, quotes, and trades. A brief summary concludes the paper in Section V. A summary of the spectral analysis techniques used in the paper is given in the Appendix.

I. The Institutional Framework and Data

The specialist is the NYSE member charged with making the market (posting bid and ask quotes) in a given security. There is only one specialist for each security. The Exchange imposes various affirmative and negative obligations on the specialist. These obligations generally require the specialist to maintain a market presence (post bid and ask quotes if no one else is willing to do so) and to promote a "fair and orderly market." In compensation for performing this function, the specialist is given limited authority for regulating the trading process, full knowledge of the limit order book, and first knowledge of the orders arriving on the computerized routing systems.⁵ Specialist performance is monitored and evaluated by the Exchange and (via questionnaires) other member brokers. Rule violations can result in fines and the Exchange considers the results of the evaluations in allocating new listings to particular specialists.

The market presence requirement effectively directs the specialist to maintain a meaningful (narrow) spread at all times. Specialists are required to maintain price continuity, which in practice limits most successive intraday price changes to one tick (one-eighth of a dollar for most stocks). They must

⁵ Visitors to the trading floor have long noted specialists' inclination to share information on the condition of the limit order book with floor traders. This sharing is now mandated by rule. Off-floor traders, however, have no direct access to this information.

trade in a stabilizing manner, buying on downticks and selling on upticks.⁶ Specialists must also yield to nonmember (public) orders: if the specialist is bidding and a public limit buy order arrives at the same price, the limit order must be filled before the specialist can purchase the stock.

On the NYSE, the specialist is the sole trader responsible for maintaining a market presence and is thereby distinguished from other liquidity providers such as public limit order traders. Together with the price continuity and stabilization requirements, this suggests that specialist exposure to sudden exogenous inventory shocks is large. On the other hand, the informational advantages inherent in full knowledge of the limit order book and of the incoming computerized order flow facilitate the process of inventory control. This combination of high exposure to shocks and advantageous control suggests that relative to dealers in various alternative market structures, NYSE specialists are most likely to conform to the inventory control paradigm.

As part of the surveillance process, specialists routinely report their activities to the Exchange. These reports constitute the basic data source for this paper. The NYSE maintains two specialist data files. The specialist equity trade (SPET) file contains specialist inventory positions and trades, reconciled with market data system information. The specialist equity trade summary (SPETS) file contains daily close positions and aggregate figures derived from the SPET file. Although it is not as detailed as the SPET file, the summary file is available for a longer historical period. We use both these files. A random sample (stratified by market capitalization) of 150 stocks was chosen in June 1990. By the time data collection commenced (in November 1990), 144 stocks remained in the sample. A few stocks with insufficient numbers of observations were dropped from the sample. Overall, 138 firms were used for the daily analysis and 137 firms for the transaction analyses. The transaction level (SPET) data for these stocks span all trading days in the three months November 1990 through January 1991. The daily summary (SPETS) data cover November 1988 through August 1990.

Various adjustments are made to the specialist inventory series. The two most common adjustments are odd-lot and "as of" adjustments. Odd lots are orders smaller than 100 shares. These are executed automatically against the specialist account. When the cumulative net position change exceeds a threshold level, the specialist receives an "odd-lot advisory." Partial round lots (e.g., the 60-share remainder of a 560-share order) are handled in the same fashion. "As of" adjustments are error corrections indicating that a trade previously thought to have taken place either did not take place, or took

⁶ The Exchange monitors each specialist's average performance on spread size, stabilization, and continuity (among other things). In 1991, 95.8 percent of all transactions occurred with no change or a one-eighth variation; 83.1 percent of specialist purchases (sales) were at prices below (above) the last different price, i.e., on a downtick or zero-downtick (uptick or zero-uptick); and the spread was one-fourth or less in 84.5 percent of the submitted quote records. (See the *New York Stock Exchange Fact Book* (1991).)

place at a different price. We leave these adjustments in place in the time sequence in which they are reported. For example, if an "as of" adjustment on day 5 indicates that a purchase of 10,000 shares from the specialist on day 1 was canceled, then the actual inventory during the intervening period was 10,000 shares larger than the record indicated. The specialist's beliefs, however, are correctly measured if the 10,000-share downward adjustment is made on day 5, and not before.⁷

II. The Time Series Characteristics of Dealer Positions

This section examines the univariate time series properties of dealer positions. The primary objective is determination of the time horizon over which inventory adjustment takes place. In the theoretical models, the speed of adjustment generally depends on the penalty of being away from the optimum and on the cost of adjustment. The penalty is a function of the riskiness of the security and the dealer's risk aversion. The cost of adjustment depends on the price elasticity of the order arrival function. With the possible exception of security risk, these parameters are not easily quantifiable. These models, therefore, offer few substantive predictions as to the speed of adjustment.

Figures 1, 2, and 3 present plots of the daily closing specialist inventory for three stocks.⁸ These stocks are not randomly chosen, but are selected because each distinctively illustrates a feature that appears to some degree in many of the other stocks. In Figure 1, the inventory series for stock A (a frequently traded issue with a large market capitalization) displays rapid reversion to a sample mean near zero. There is little evidence of persistence in positions across months. The sharp upward spike reflects specialist purchases in the October 1989 market break. But even this shock, which is roughly twice the amplitude of the next highest peak, is quickly eliminated. Stock B (Figure 2) also exhibits short-term variation. In this case, however, the variation is around a long-term level that shifts in a fashion that might be characterized as a discrete jump. The plot for stock C (Figure 3) also displays a combination of long- and short-term effects. In this case, however, the long-term effects are not discrete, but fluctuate slowly. Stocks B and C are infrequently traded issues with low capitalization. Most of the stocks in the sample display a mixture of these three features: short-run variation, discrete long-term variation, and smooth long-term variation.

Before proceeding further in analyzing the dynamic behavior of inventory levels, the stationarity of the series requires examination. In particular, the

⁷ In one respect at least, the data used in the present study are poorer than those used by Madhavan and Smidt (1991, 1993): the present study does not have access to the clearing records ("canceled checks" is their apt description) of the individual firms. On the other hand, Madhavan and Smidt use sixteen stocks from one (self-selected) specialist firm. The present sample is both more numerous and randomly selected.

⁸ The terms under which the NYSE collects these data preclude public identification of individual stocks or specialist units.




Figure 1. Daily closing specialist inventory in shares for "stock A." Source: NYSE Specialist Performance Evaluation Trade Summary file.

issue of whether the series possesses a unit root is extremely important because it bears on the specification of the empirical tests. If inventory levels contain a random walk (unit root) component, the level autocorrelations would not be meaningful and some degree of differencing might be appropriate. On the other hand, over-differencing of a series that is already stationary may also lead to statistical problems.⁹

Univariate Dickey-Fuller tests are conducted for the daily closing inventory series. For autoregressive specifications of lag 10, the null hypothesis that the inventory level possesses a unit root is rejected at the 0.10 level in 112 of 147 firms. With 50 lags, the rejection rate drops to 35/147. The null is rejected for stocks A, B, and C with 10 lags, but not at 50 lags. It is, however, notoriously difficult to ascertain statistically whether a time series that contains large long-run components possesses a unit root. Looking for a unit root for stock B is tantamount to asking if the jump process that generated the single apparent shift is mean reverting. Conventional unit-root tests possess little power in such applications: when long lags (relative to the sample size) are included in the autoregressive specification, the standard errors become so large that rejection of the null (unit root) hypothesis is difficult.

⁹ One cannot, for example, construct a convergent autoregressive representation for an over-differenced series.




Figure 2. Daily closing specialist inventory in shares for "stock B." Source: NYSE Specialist Performance Evaluation Trade Summary file.

Most economic considerations, however, militate in favor of stationarity. A series that possesses a unit root diverges over time. By reason of capital and regulatory constraints, however, dealer positions are bounded. Although over sufficiently long horizons these constraints and other parameters may themselves be nonstationary, such variation over the present sample is likely to be small. Inventory level series will henceforth be assumed stationary.¹⁰

Under the assumption that the inventory level always (eventually) reverts to its mean, the adjustment speed may be evaluated by examining the autocorrelations. For each stock in the sample, we compute the autocorrelations in closing inventories through the fiftieth lagged trading day. Let $n_{i,t}$ denote the number of shares of stock i in the specialist's inventory on the close of day t . Figure 4 plots the median autocorrelations in $n_{i,t}$ in subsamples constructed as quartiles on transaction frequency. The autocorrelations are typically large at low lags, and then decay. The lower trade frequency subsamples generally possess larger autocorrelations at all lags, indicating that adjustment is faster for frequently traded stocks. In all but the highest trade frequency subsample, the median autocorrelations are above 0.1 at a lag of 10 trading days (two calendar weeks). One way of interpreting this figure is that if on day 0 the closing inventory exceeds the optimum by 1,000

¹⁰ Madhavan and Smidt (1993) also assume stationarity in the inventory levels. The estimations reported in their Table II are functionally similar to a unit root test specification.




Figure 3. Daily closing specialist inventory in shares for "stock C." Source: NYSE Specialist Performance Evaluation Trade Summary file.

shares, the expected excess inventory (in the simple univariate regression sense) will be above 100 shares on day 10. For the least-frequently-traded subsample, the median autocorrelations are approximately 0.1 at a lag of about 42 days (roughly two calendar months).

The autocorrelations suggest that inventory levels in at least some stocks contain long-term (weekly or longer) components. Several plausible economic explanations for this characteristic may be conjectured. One possibility is that the long-term shifts arise from exogenous shocks to the specialist inventory. The adjustment may be slow due to the high cost of working off these positions. Alternatively, the long-term variation may arise from shifts in the desired level of holdings. Under this interpretation, the long-term components are more appropriately interpreted as speculative positions. Evidence on specialist trading activity and on the path of inventories subsequent to exogenous shocks indicates that the latter is the case.

First consider specialist trading activity relative to the total transaction flow. If trading volume is taken as exogenous (an admittedly crude assumption), and if the specialist's purchases and sales are large relative to this volume, one might reasonably conjecture that the trading volume is a binding constraint on inventory adjustment.

Table I presents descriptive statistics for the daily sample of inventory data. The average daily closing inventories for the stocks in our sample are 11,897 shares or \$160,650. Table I also presents three measures of specialist




Figure 4. Median autocorrelations in closing daily inventory levels. Figure depicts the median (across stocks) of $\rho_{i,k}$ where $\rho_{i,k}$ is the autocorrelation of the daily closing inventory level for stock i at lag k . The sample is 144 NYSE stocks from September 1988 through August 1990.

activity relative to the transaction flow. The first is the traditional specialist participation rate, defined as

$$\frac{(\text{Specialist Purchases}) + (\text{Specialist Sales})}{\text{Total Purchases and Sales}} \quad (1)$$

where purchases and sales are in share units, and the totals are computed over some appropriately long time period. The denominator is twice total volume, as conventionally reported. Over the entire sample, the participation rate averages 13 percent, although it is higher (19 percent) for the lowest trade frequency subsample. These estimates are far short of the 50 percent figure that would obtain if the specialist were counterparty to all trades.

Table I also presents average values of the ratio of average closing inventory to average daily volume. This ratio is similar to the inventory/sales ratio of classic financial analysis. For present purposes, it answers the question: if the specialist were the counterparty to all purchases (sales), how many days would it take to reduce the average positive (negative) inventory position to zero? This averages 0.84 days for all stocks, but even for the least-traded subsample, it is only 2.77 days. The final statistic relates the average daily inventory adjustment to the total trading volume. Like the participation rate, it suggests that the net inventory adjustment is small relative to the total volume. There is little evidence, therefore, that the long inventory adjustment lags are the result of volume constraints.

Table I
Summary Statistics

The table reports summary statistics for NYSE securities used in the study, from November 1988 through August 1990. Except for "Number of securities," table entries are sample means and (in parentheses) standard deviations. The participation rate for a given firm is defined as the sum of specialist purchases and specialist sales (over the sample period) divided by twice total volume.

$$\frac{\text{Avg. } |\text{Inventory}|}{\text{Avg. Daily Volume}} \text{ and } \frac{\text{Avg. } |\Delta \text{ Inventory}|}{\text{Avg. Daily Volume}}$$

are computed for a given firm using daily closing inventory (in shares), or the first difference of the daily closing inventory and the average daily trading volume for the firm. Statistics are derived from the NYSE's specialist equity trade summary file.

Variable	Total Sample	Quartile Subsamples by Average Daily Number of Transactions			
		1 (lowest)	2	3	4 (highest)
Number of Securities	138	29	36	37	36
Daily Closing Inventory (100-share lots)	118.97 (605.44)	106.66 (290.18)	71.33 (110.15)	48.68 (75.97)	248.78 (1,150.80)
Daily Closing Inventory (\$1,000)	160.54 (518.00)	100.16 (155.29)	99.91 (196.74)	83.67 (110.46)	348.81 (964.49)
Number of Transactions per Day	51.20 (96.95)	5.44 (1.90)	12.74 (3.44)	32.34 (8.20)	145.90 (154.43)
Daily Volume (100-share lots)	1,184.45 (2,518.98)	61.82 (38.67)	180.36 (127.74)	667.34 (655.15)	3,624.36 (3,986.36)
Participation Rate	0.13 (0.06)	0.19 (0.07)	0.15 (0.06)	0.11 (0.06)	0.10 (0.03)
Avg. $ \text{Inventory} $	0.84 (3.11)	2.77 (6.47)	0.65 (0.41)	0.20 (0.14)	0.13 (0.06)
Avg. $ \Delta \text{ Inventory} $	0.16 (0.16)	0.34 (0.27)	0.17 (0.06)	0.10 (0.05)	0.06 (0.04)

Next consider the response of specialists to sudden involuntary shocks to their holdings. While it is not in general possible to identify conclusively such shocks, some suggestive inferences may be gleaned from days in which there were large price changes. The price stabilization requirements discussed in the preceding section tend to force the specialist to buy on price declines and sell on price advances. Taking the position changes associated with large price movements to be approximately exogenous, the time required to reverse the shock can then be examined. The inventory paths in Figures 1 and 2, for example, display sharp rises on October 13, 1989. The fact that these shocks were quickly reversed suggests that the increase was not desired and that rapid correction was feasible. By implication, then, when similar levels

persist for long periods of time (as in the later portion of Figure 2), these levels are essentially voluntary.

For each stock in the sample we compute the standard deviation of the change in daily inventory levels and examine the day with the largest absolute percentage price change. If the change in the inventory is at least two standard deviations away from zero, we tabulate the number of days that elapse before the inventory is restored to its preshock level. Of the 40 occurrences that satisfy these criteria, 11 were on October 13, 1989. In 24 cases, full adjustment occurs by the end of the first day after the shock, and in 36 cases, full adjustment occurs by the end of the fourth day. This adjustment is much more rapid than that implied by the autocorrelations.

Finally, specialists may adjust inventory levels gradually because they hedge with offsetting positions in options, futures, and other stocks. Despite the apparent advantages that such strategies might permit, conversations with specialists indicate that they are not extensively used. The insurance afforded by options on individual stocks is regarded as being too expensive, and the market risk hedge implicit in an index futures contract is considered too inexact.

To investigate the third possibility, hedging using offsetting equity positions, we calculate the correlations of holdings within specialist units. If offsetting positions occur, there should be negative correlations between holdings of different stocks assigned to the same specialist. The data sample used for this analysis comprises all holdings for all specialist units for all days of 1991. For each specialist unit the bivariate correlations between the holdings of each pair of specialty stocks (i.e., those issues handled by the specialist unit) are computed. The representative (median) correlation is positive for 41 of the 50 specialist units. While this rough analysis does not rule out the possibility that some hedging occurs, it does suggest that such behavior is not pervasive.

The evidence presented to this point can be summarized as follows. For a typical security, if the mean inventory estimated over a long data sample is taken as the target level, adjustment toward this target takes place at a slow rate. Comparison of the adjustments to daily total volume suggests that more rapid adjustment is feasible. Furthermore, on days when the inventory shocks are arguably exogenous, more rapid adjustment does in fact occur. The simplest explanation that reconciles these findings is that short-term variation reflects classic dealer behavior, while the long-term variation stems from investment holdings. The distinction between short- and long-term behavior in the inventory and related series is a recurrent theme in subsequent sections.

III. Trading Profits

This section examines specialist trading profits, here defined as the gains (or losses) from purchases and sales of stock. Due to a lack of relevant data, these profits are not adjusted for the fees paid by the specialists to the NYSE,

financing charges, other operational expenses, and salary imputations. Specialist commission revenues (the fees received for acting as agent for orders) are also ignored. The focus is strictly on gross trading profits.

Trading profits can be measured either on a cash flow basis or a mark-to-market basis. Let p_t denote the price per share of the security at time t , and let n_t denote the number of shares held. It is assumed that the trade at time t takes place at price p_t and that n_t is net of all trading at time t . The mark-to-market profit is defined as the appreciation in the value of the holdings:

$$\Pi_t = n_{t-1}(p_t - p_{t-1}), \quad (2)$$

i.e., the capital gains on shares held at the beginning of the period. The cash flow arising from establishing or liquidating a position is determined by the change in holdings and the current price:

$$F_t = -p_t(n_t - n_{t-1}). \quad (3)$$

From a data sample for $t = 0, \dots, T$, the average mark-to-market and cash flow profits may be computed.

Investment returns are typically determined by marking positions to market. One criticism of this approach is that it assumes that entire positions can be purchased or liquidated at the prevailing market prices. This assumption of infinite elasticity is not plausible for large position changes over short intervals. In fact, it is demonstrated in Table I that changes in specialist positions tend to be small relative to the total order flow. Cash flow profits, on the other hand, are subject to the criticism that the profit is entirely attributed to the instant at which the position is liquidated.

Although F_t and Π_t will often differ dramatically for any particular t , the differences stem from the timing of profit recognition. In any given data sample, the averages of the two measures are likely to be close.¹¹ The mark-to-market approach is used here mainly for statistical reasons. It is the

¹¹ The average mark-to-market profit and its relation to the average cash flow profit are given by:

$$\begin{aligned} \bar{\Pi}_t &= \frac{1}{T}[\Pi_1 + \Pi_2 + \dots + \Pi_T] \\ &= \frac{1}{T}[n_0(p_1 - p_0) + n_1(p_2 - p_1) + \dots + n_{T-1}(p_T - p_{T-1})] \\ &= \frac{1}{T}[-n_0 p_0 - p_1(n_1 - n_0) - p_2(n_2 - n_1) - \dots - p_T(n_T - n_{T-1}) + n_T p_T] \\ &= \frac{1}{T}[-n_0 p_0 + F_1 + F_2 + \dots + F_T + n_T p_T] \\ &= \frac{-n_0 p_0}{T} + \bar{F}_t + \frac{n_T p_T}{T} \end{aligned}$$

Thus, $\bar{\Pi}_t$ and $\bar{F}_t$ are essentially similar, differing only by endpoint terms corresponding to initial acquisition and final liquidation.

product of two series that, subject to the caveats of the earlier discussion, are approximately stationary. The cash flow, in contrast, is the product of a stationary series (the first difference of n_t) and a nonstationary series (p_t). The stationarity of the mark-to-market profit components greatly facilitates the empirical analysis.

The definition of mark-to-market profits, equation (2), is modified to allow for nontrade adjustments to specialist holdings. For example, suppose that at the close of the day the share price is \$10 and that the specialist reports an "as of" purchase of 1,000 shares at \$11 (stemming from the resolution of a disputed trade). Based on the current liquidation value of the position, the specialist has incurred a loss of \$1,000. Lacking any further information about when the trade actually occurred, attributing the loss to the time when the adjustment was recognized seems preferable to ignoring it. Accordingly, a profit measure that incorporates adjustments is defined as

$$\Pi_t = n_{t-1} dp_t + A_t^b(p_t - p_t^b) - A_t^s(p_t - p_t^s). \quad (4)$$

where A_t^b and A_t^s are the number of shares reported at time t as having been bought and sold earlier, and p_t^b and p_t^s are the prices at which these adjustments were effected.

Using the three-month transactions data sample, the average trading profits per transaction are computed for all stocks in the sample.¹² For stock A, estimated profits are \$19.48 per transaction (with a standard error of 3.25); for stock C, \$10.99 (20.64). (Stock B was delisted by the start of the the transaction data sample.) The median across securities of the average profits per transaction (Table II, first column) is \$7.55. The median standard error, however, is 15.78, while the interquartile range is 19.55. These statistics imply large uncertainty about the individual point estimates and substantial variation across stocks. Overall, the null hypothesis of zero profits can be rejected at the 5 percent level for only 37 of the 137 stocks, suggesting that only a minority of specialty stocks have nonzero gross trading profits.

This surprising finding has a simple statistical explanation. The strong persistence in inventory levels documented in Section II suggests that average profits contain a large long-term component. Consider the extreme case of a specialist whose inventory is flat at some level $n > 0$ for the entire sample period. The average trading profits in this case are the product of n and the mean price change over the sample period. Mean price changes, however, computed over short samples are extremely noisy estimates of the corresponding population values (see Merton (1980)).

The large estimation error, therefore, most likely arises from the long-run component of profits and the short sample period. This feature motivates the decomposition of specialist profits by investment horizon. Such a decomposition is useful for two further reasons. The largest contribution to a dealer's profits is generally thought to arise from turning over a position between the bid and offer prices over relatively brief holding intervals. Furthermore, to

¹² All NYSE transactions are included, not just the ones in which the specialist participates.

Table II
Gross Trading Profits

For a given firm, let $\bar{\Pi}_t$ represent the average gross trading profit, in dollars per transaction. For profits over all horizons, $\bar{\Pi}_t$ is computed using equation (4). Within the total sample of firms (and the indicated transaction frequency subsamples), the table reports the median $\bar{\Pi}_t$, the median standard error (S.E.) of the $\bar{\Pi}_t$, and the interquartile range (I.Q.R.) of the $\bar{\Pi}_t$. (The median S.E. is the median of the estimated standard errors, not the standard error of the median.) Long-, medium-, and short-term profits per transaction are computed using spectral analysis. Spectral analysis resolves a time series into periodic components. The period of a component is the number of transactions required for the component to complete one full cycle. For a given firm, the short-term profit is estimated as the integral (over periods of 10 transactions or less) of the cospectrum between the first difference in the price per share and the lagged inventory. Medium-term profits (for periods between 10 and 100 transactions) and long-term profits (for periods over 100 transactions) are computed similarly. Standard errors are computed using spectral estimates smoothed with a Daniell window (see Bloomfield (1976)).

	Total Sample	Quartile Subsamples by Average Number of Daily Transactions			
		1 (low)	2	3	4 (high)
Number of securities	137	34	34	35	34
Median (average number of transactions per day)	19.73	5.15	12.17	31.22	85.17
Total profits (\$ per transaction)					
Median	7.55	5.49	7.15	10.42	7.26
Median S.E.	15.78	17.39	19.44	14.25	12.39
I.Q.R.	19.55	30.28	23.38	16.34	23.78
Long-term profits (\$ per transaction)					
Median	-7.02	-1.02	-6.93	-10.65	-11.06
Median S.E.	13.22	10.97	17.35	13.23	11.85
I.Q.R.	22.73	29.94	22.01	17.24	21.15
Medium-term profits (\$ per transactions)					
Median	6.08	3.57	4.18	7.88	8.17
Median S.E.	5.55	11.65	7.88	4.74	3.41
I.Q.R.	8.89	15.25	10.96	8.01	9.03
Short-term profits (\$ per transaction)					
Median	12.24	8.23	12.71	13.00	11.63
Median S.E.	2.81	4.97	3.09	1.86	1.63
I.Q.R.	10.34	13.12	9.21	8.77	9.57

the extent that profits are attributable to information, the advantage possessed by a dealer or any trader close to the market-making process is essentially short term (certainly intraday and possibly intraweek).

One approach to differentiating effects in time series over different horizons involves time domain techniques. The series is filtered or smoothed in order to extract a smooth "long-run" component and a residual "short-run" component. These techniques are sensitive to judgements regarding filter

design and estimation. A standard alternative approach is spectral (frequency domain) analysis. Although spectral estimates of a time series contain no more and no less information than the standard time domain techniques, they offer a more direct characterization of horizon effects. The formalities of the approach are addressed by a number of excellent references, but the intuition is simple.¹³ A realization of a time series may be expressed as a weighted sum of sinusoidal functions. Once the series has been expressed in this fashion, it is possible to decompose a moment or cross-moment (such as the expected profits) according to the contributions of different frequencies.

Although there are certain commonalities, spectral analysis should not be confused with harmonic analysis. The aim of the latter is detection of regular cyclical components in a data sample. These cycles are held to be deterministic components of the series in question. Sunspots, agricultural yields, and airline traffic are generally considered to possess deterministic cycles. In spectral analysis, the cyclic components found in a sample of data are assumed to be estimates of the unknown underlying components. The distinction is important because financial return series are generally assumed to be free of deterministic components. This paper does not suggest otherwise.

Since the trading profit defined in (2) is a cross-product, basic spectral techniques may be applied directly. First, Fourier analysis is used to extract the sinusoidal components of the price changes and inventory levels at various frequencies. At a particular frequency, if the price change and inventory level are in phase (i.e., have peaks and troughs that match), then the contribution to the covariance (trading profits) is positive. If the two series are out of phase, then the contribution is negative.

As an example, the component of the highest observable frequency corresponds to a pattern that completes a full cycle every two transactions. Suppose that the bid and ask quotes are set symmetrically about an implicit efficient price that follows a random walk. Suppose further that transaction prices alternate between the bid and the ask, and that the specialist is a party to all transactions (buying at the bid and selling at the ask). This short-horizon behavior will show up in the spectral analysis as a source of profits. A sequence consisting of two trades at the bid followed by two trades at the ask would correspond to a component with a period of four transactions, and so on. More complicated (and realistic) patterns can be generated by combinations of sinusoidal functions.

Application of spectral analysis requires covariance stationarity of the component series. While one aspect of this question (the unit root problem) has already been addressed in connection with the daily data (Section II), this and the subsequent section rely on intraday transaction data. The subscript t indexes transactions, and the covariance stationarity is assumed to hold in transaction time. This approach is likely to be preferable to the alternative assumption of stationarity in natural time. There is substantial evidence that return variances per unit time, for example, are substantially elevated at the

¹³ See Fuller (1976), Bloomfield (1976), and Granger and Engle (1983).

beginning and end of trading sessions. Furthermore, Harris (1987) suggests that most daily return moments are functions of the number of transactions.

The paper does not report complete estimations of the spectra and cross-spectra for the inventory and price change series. It is econometrically much simpler, and for the present purposes adequate, to consider components of variances and covariances attributable to broad frequency bands. The approach is essentially the one suggested in Engle (1974, 1978). A summary of the computational details and an example is given in the Appendix.

For the intraday analysis, the bands are selected to correspond to periodic components of the following lengths. Greater than 100 transactions (long horizon); 100 to 10 transactions (medium); and under 10 transactions (short). Average profits per transaction over each of these bands are computed for each security. Table II presents the medians (across securities) for the total sample and subsamples. For any given stock, the frequency decomposition is an identity: the short-, medium-, and long-term profits sum to the total.¹⁴

The spectral profit estimates display a striking pattern. In the short run, they are positive (median value \$12.24 per transaction) with relatively low standard errors (median \$2.81). Positive profits persist in the medium term, although the median standard errors are somewhat higher. Over the long horizon, however, while the median profits are negative, the median standard errors are large. This suggests that the noisy estimates found for total profits are primarily a consequence of the long-term variation in inventory levels and returns. Over short and medium horizons, the profit estimates are more positive and distinct. These findings strongly confirm that specialist trading profits are due to short- and medium-term trading activity.¹⁵ The estimates of the long-term components are not statistically reliable, and so cannot support any conclusions regarding long-term specialist profitability.

Two likely sources of short-run profits are inventory turnover at the bid and ask prices and trading on the basis of brief informational advantages. One way of resolving these two components is to consider a modified definition of trading profits with the bid-ask spread removed:

$$\Pi_t^q = n_{t-1}(q_t - q_{t-1}) = n_{t-1} dq_t \quad (5)$$

where q_t is the midpoint of the bid and ask quotes prevailing at the time the trade was initiated.

Table III presents quote midpoint profits (net of adjustments) with a horizon decomposition identical to that given for actual profits. The long-term component is virtually identical to that of actual profits. This reflects the fact that long-term price components are little affected by the bid-ask spread. Over short and medium horizons, however, the median profits are substan-

¹⁴ The decompositions in Tables II and III do not sum exactly because the entries are sample and subsample medians.

¹⁵ For stock A, the short-, medium-, and long-term profits per transaction are (with standard errors) 11.74 (.83), 9.45 (1.46), and -1.71 (2.78); for stock B, 8.97 (2.68), 9.30 (5.92), and 7.28 (19.59).

Table III

Gross Quote Midpoint Trading Profits

For a given firm, let $\bar{\Pi}_t$ represent the average gross trading profit, in dollars per transaction. For profits over all horizons, $\bar{\Pi}_t$ is computed using equation (4), but with prevailing quote midpoints used in lieu of the actual transaction price. Within the total sample of firms (and the indicated transaction frequency subsamples), the table reports the median $\bar{\Pi}_t$, the median standard error (S.E.) of the $\bar{\Pi}_t$, and the interquartile range (I.Q.R.) of the $\bar{\Pi}_t$. (The median S.E. is the median of the estimated standard errors, not the standard error of the median.) Long-, medium-, and short-term profits per transaction are computed using spectral analysis. Spectral analysis resolves a time series into periodic components. The period of a component is the number of transactions required for the component to complete one full cycle. For a given firm, the short-term profit is estimated as the integral (over periods of 10 transactions or less) of the cospectrum between the first difference in the price per share and the lagged inventory. Medium-term profits (for periods between 10 and 100 transactions) and long-term profits (for periods over 100 transactions) were computed similarly. Standard errors were computed using spectral estimates smoothed with a Daniell window (see Bloomfield (1976)).

	Total Sample	Quartile Subsamples by Average Number of Daily Transactions			
		1 (low)	2	3	4 (high)
Number of securities	137	34	34	35	34
Number of transactions per day	19.73	5.15	12.17	31.22	85.17
Total profits (\$ per transaction)					
Median	-3.30	-1.26	-4.40	-5.56	-2.60
Median S.E.	15.15	15.57	18.17	14.19	12.17
I.Q.R.	20.48	44.05	16.68	16.38	27.94
Long-term profits (\$ per transaction)					
Median	-7.18	-1.61	-7.28	-10.87	-10.63
Median S.E.	13.13	10.99	16.48	13.24	11.71
I.Q.R.	20.97	18.14	21.81	18.79	22.11
Medium-term profits (\$ per transaction)					
Median	3.86	1.66	1.20	4.65	6.09
Median S.E.	5.44	9.79	6.96	4.55	3.34
I.Q.R.	7.84	14.04	9.52	6.56	6.87
Short-term profits (\$ per transaction)					
Median	0.39	0.35	0.13	0.36	0.82
Median S.E.	1.98	3.94	2.65	1.45	1.18
I.Q.R.	5.12	8.53	4.05	6.70	5.09

tially reduced, although they remain positive. The strongest evidence of this lies in the highest trade frequency subsample. This may be indicative of an ability to adjust positions based on information about local market conditions.

Summarizing, specialist trades are profitable, but only in the short and medium terms. Most of these profits vanish if the trades are forced to be at the midpoint of the prevailing bid and offer quotes. Nevertheless, the small positive short- and medium-term components of quote midpoint profits suggest that there are price trends over this horizon that generate trading profits absent the spread.

IV. Trades, Inventories, and Quote Revisions

This section examines the joint behavior of specialist inventory levels, quote revisions, and trade variables. The objective is to characterize the relation between specialists' trading and the dynamic behavior of prices. Theory and empirical realities suggest that this relation reflects several diverse effects. According to the inventory control model, dealers set the quotes to elicit orders of the desired sign. By this mechanism, a trade causes a transient change in the quotes. Asymmetric information complicates matters. If an uninformed dealer faces a population of traders, some of whom possess superior information, the buyer-initiated trade will have positive information content. This will lead to a permanent revision in the quote. Assuming as an approximation that the inventory and asymmetric information effects on the security price are additive, a buyer-initiated trade will lead to an immediate increase in the quote, reflecting both information and immediate inventory effects. When the inventory imbalance is corrected, quotes should revert to the level consistent with the information content of the trade.

If, on the other hand, it is the dealer who possesses an informational advantage, an inventory depletion should precede a *decline* in the price (quotes). The duration of this effect depends on the nature of the information. Position adjustments in anticipation of transient price movements (caused perhaps by transient shocks in the liquidity-motivated component of the order flow) need not be associated with permanent price effects. Superior information concerning the terminal cash flows of the security or the informed order flow, on the other hand, would have permanent effect. The evidence on quote midpoint profits discussed in the last section is consistent with superior information over the medium term.

These various effects are investigated using vector autoregression (VAR) models and the associated impulse response computations. Consistent with this paper's approach, VARs impose few theoretical restrictions on the estimated model and permit the consideration of broad hypotheses in a framework of robust statistical generality. Hasbrouck (1991a, 1991b, 1993) discusses microstructure applications of VARs.

The empirical specification is based on the following framework. At the beginning of time t , the prevailing quote midpoint is that set at the close of time $t - 1$, denoted q_{t-1} . A trade occurs, defined as $x_t = |\text{trade volume}|$ if the transaction price is greater than q_{t-1} (a buyer-initiated trade), $x_t = -|\text{trade volume}|$ if the price is less than q_{t-1} (a seller-initiated trade), and $x_t = 0$ if the price is equal to q_{t-1} .¹⁶ Following this trade, the specialist sets new quotes, with quote midpoint q_t ; the quote revision is $r_t = q_t - q_{t-1}$; the

¹⁶ Classifying trades into buyer and seller initiated by comparing the trade price to the midquote is, however, a rough approximation. Midquote trades cannot be classified and are assigned a zero value. Tick-sensitive orders and stopped orders may be misclassified. A sell-plus order, for example, will trade at the ask and will be misclassified as a buy order. Madhavan and Smidt (1993) use a similar approach in the construction of their order flow variable.

specialist's inventory at the close of time t is n_t , which is net of x_t if the specialist is the counterparty to the trade. In order to capture nonlinear effects of trade size on quote changes, let the signed trade variable of order k be denoted $x_t^k = \text{sign}(x_t)|x_t|^k$. The column vector of the five variables included in the VAR is $z_t = [r_t, x_t^0, x_t^{1/2}, x_t^1, n_t]'$. A five-lag structure is adopted, so the VAR model may be summarized as:

$$z_t = A_0 z_t + A_1 z_{t-1} + \cdots + A_5 z_{t-5} + u_t \quad (6)$$

where u_t is the column vector of residuals and the A_i are conformable coefficient matrices. The contemporaneous term, $A_0 z_t$, reflects the recursion at time t : x_t and n_t are determined jointly and prior to r_t .¹⁷ This system was estimated for all securities with at least 1,000 transactions in the three-month sample period (73 securities). Of the three stocks considered in Section II, only stock A meets this criterion.

A key feature of this specification is that it includes signed trade variables as well as the inventory level. This permits a distinction between trades in which the specialist participates and those (the majority) that take place without participation. The differential effect of participation can thereby be analyzed. The asymmetric information effect suggests that a trade will have an impact on the price irrespective of the counterparty's identity. An inventory control effect, however, should be evident, only in those instances in which the specialist participates.

This hypothesized resolution of inventory and information effects is undermined by the limitations of the data. First, a substantial proportion (typically 40 percent) of the trades take place at prices at the midpoint of the prevailing quotes, and cannot be signed. The inventory change in these cases provides additional information about the direction of the trade. This burdens the inventory variable with an additional proxy role. Second, specialist participation is frequently associated with other characteristics of the state of the market, leading to a type of selection bias. For example, if perceived informational asymmetries (and price impacts of trades) are high, it is likely that alternative suppliers of liquidity, such as public limit order traders, will withdraw. The specialist must maintain a market presence, however, and so will participate in a disproportionate share of the trades that occur during these times.

Both inventory control and asymmetric information effects suggest that inventory shocks should affect quotes. In addition, if the specialist actively changes the quotes to elicit orders of the desired sign, quote shocks should affect inventories. Tests of Granger-Sims causality were overwhelmingly consistent with bidirectional causality between inventory levels and quote revisions.¹⁸

A more refined picture of the inventory/quote-revision interactions is obtained from the impulse response functions of the system. The VAR (6) may

¹⁷ Since r_t is the first equation in the system, the only nonzero elements of A_0 are in the first row, $[A_0]_{1,2}$ through $[A_0]_{1,5}$.

¹⁸ Madhavan and Smidt (1993) find Granger-Sims causality running from quote revisions to inventories, but not the reverse. Their estimations employ only daily data, however, and so may neglect effects operating over shorter intervals.

be rearranged as

$$z_t = (I - A_0)^{-1} A_1 z_{t-1} + \dots + (I - A_0)^{-1} A_5 z_{t-5} + (I - A_0)^{-1} u_0 \quad (7)$$

Taking all lagged values in (7) as zero, the initial impact of a shock u_0 may be computed as $E[z_0 | u_0] = (I - A_0)^{-1} u_0$, and by recursively substituting this in for z_0 , the expectation of next period's realization may be computed as $E[z_1 | u_0]$ and so on. The first element of z_t is the quote revision. Accordingly, the expected cumulative quote midpoint change through transaction m is

$$\alpha_m(u_0) = \sum_{t=0}^m E[r_t | u_0].$$

For each stock, the initial shock was chosen to correspond to a buy order of size equal to the 90th percentile of all trade sizes. Two separate impulse response computations were made, one assuming no specialist participation and the other assuming that 10 percent of the sell volume comes out of the specialist inventory (and the remaining 90 percent from other traders). Let $\alpha_{m,i}^{no\ part}$ denote the cumulative quote revision through transaction m for stock i assuming that the specialist does not participate. Similarly, $\alpha_{m,i}^{part}$ denotes the cumulative quote revision assuming participation.

In both participation cases, the α 's were generally positive over all horizons through $m = 50$, indicating that a buyer-initiated trade conveys positive information whether or not the specialist participates. A buy order with specialist participation corresponds to a depletion of inventory. The subsequent rise in the quote midpoint implied by the impulse response functions does not support the view that the specialist possesses a significant informational advantage.

The differential effect of specialist participation is investigated by considering the difference between the two implied quote revisions,

$$\Delta_{m,i} = \alpha_{m,i}^{part} - \alpha_{m,i}^{no\ part}.$$

Across the 73 securities, this difference exhibits substantial variation and is frequently statistically indistinguishable from zero. Some inferences regarding the sign and significance can be drawn from estimates pooled across stocks.

The pooling is performed as follows. For each stock i and horizon m , a standard error of estimate for $\Delta_{m,i}$, denoted $SE(\Delta_{m,i})$ is computed using Monte Carlo simulations.¹⁹ A standardized variate is then constructed as

$$z_{m,i} = \frac{\Delta_{m,i}}{SE(\Delta_{m,i})}.$$

¹⁹ The Monte Carlo standard errors are computed as follows. Let A denote the column vector of all coefficients for the VAR (6). Regression analysis provides both estimates of the parameters A and the associated covariance matrix $\text{Var}(A)$. The estimates are asymptotically normally distributed. The distribution of a function of the A , such as the impulse response function, can be approximated by computing the function over a set of random coefficient draws from the multivariate normal distribution with mean A and variance $\text{Var}(A)$. The standard errors of the cumulative quote revision were computed in this fashion using 30 random draws.

These standardized variates are then averaged across all securities:

$$\bar{z}_m = \frac{1}{73} \sum_{i=1}^{73} z_{m,i}.$$

These averages and associated two-standard-error confidence bounds are plotted in Figure 5.

These standardized measures of the specialist participation effect suggest that over all horizons, trades in which the specialist participates are associated with larger quote revisions. This behavior is statistically significant in the short run, but not beyond the first few transactions subsequent to the trade. The larger apparent information impact of trades in which the specialist participates may arise from either or both of the study's limitations described above. That is, in consideration of the market maintenance obliga-




Figure 5. The differential effect on quotes of specialist participation in a buyer-initiated trade. The incremental impact of a buy order on the quote midpoint conditional on the knowledge that the specialist is the counterparty seller for a tenth of the order. The plot is based on impulse response functions computed from estimated vector autoregressions using the five variables consisting of inventories, quote revisions, and signed trade variables. The plot is based on results pooled from estimates for the 73 firms in the sample that possessed at least 1,000 transactions. For each firm i , an initial shock was assumed to consist of a buy order of size equal to the 90th percentile of all trade sizes. The implied impact on the quote midpoint at m transactions subsequent to this initial shock is defined as $\alpha_{m,i}^{part}$ if the specialist participated in the trade and $\alpha_{m,i}^{no\ part}$ otherwise. The incremental difference is $\Delta_{m,i} = \alpha_{m,i}^{part} - \alpha_{m,i}^{no\ part}$. The standardized incremental difference is $z_{m,i} = \Delta_{m,i}/SE(\Delta_{m,i})$ where $SE(.)$ is the standard error. The values plotted in the figure are the mean standardized incremental differences (across firms), with two-standard-error confidence bands.

tion, the specialist may be forced to participate in a relatively high proportion of “difficult” high-impact trades. Alternatively, specialist participation may contain additional information about the direction of trades with prices equal to the quote midpoint.

The variates plotted in Figure 5 are standardized across firms, not transaction time. That the magnitude at transaction fifty is much lower than that at transaction one is not due to the fact that the $\Delta_{m,i}$ are generally lower at $m = 50$, but rather to the fact that the associated standard errors are generally much higher. The evidence in Figure 5 suggests that trades in which the specialist participates have a larger quote impact, but does not constitute evidence of quote reversion. A pooled analysis similar to the one implemented above for the $\Delta_{m,i}$ finds no statistically significant evidence of quote reversion.

The power of the analysis to detect such reversion, however, is small. In the usual inventory control models, the horizon over which the quote reversion operates is comparable to that of the mean reversion in inventories. That is, the quote reversion should be complete only when the inventory has been restored to the target level. In transaction time, the first-order autocorrelation in inventory levels for these firms is typically close to unity (often in the neighborhood of 0.99), implying that reversion in inventories (and quotes) occurs over long transaction horizons. Over such horizons, the inventory control component of the quote changes is likely to be dwarfed by changes in quotes due to the arrival of new information.

The quote revision specification estimated here is similar to Madhavan and Smidt's (1993) equation (17). In both specifications signed order flow and inventory levels (in their case, changes in the inventory level) appear separately, and the estimated inventory coefficients are statistically significant. In their specification, this suffices to establish the presence of active inventory control. The limitation that inventory changes may be proxying for order flow information is common to both studies, however. This proxy relation could explain both Madhavan and Smidt's nonzero inventory coefficients and our finding of a differential effect of trades in which the specialist participates.

In summary, the results of the dynamic analyses confirm a role for inventory levels in explaining subsequent quote movements. Examination of the differential impact of specialist participation suggests that trades in which the specialist participates have a larger impact on the quotes than trades with no specialist participation. This may reflect institutional constraints requiring participation in high-impact trades, or a proxy role for inventories. There is no compelling evidence of the quote overshooting and reversion pattern predicted by the traditional inventory control mechanism.

V. Conclusions

This study examines a comprehensive sample of quote, trade, and inventory data for New York Stock Exchange specialists. The most salient univari-

ate characteristic of specialist holdings is the long persistence of deviations from the mean values. If the mean is taken as a good estimate for a time-invariant optimum target level, this finding suggests that the conventional inventory control mechanisms operate with adjustment lags of days or weeks. This paper, however, provides evidence that specialists are capable of rapidly adjusting their positions. Moreover, the paper finds no evidence that specialists react slowly to inventory shocks because they are hedged with offsetting positions in other stocks. The most plausible explanation, therefore, for the long persistence in inventory levels is that specialists adjust inventory levels toward time-varying targets.

The study then presents a novel analysis of specialist trading profits, using spectral decompositions to attribute components of these profits to particular horizons. Most of the profits arise over short- and medium-term holding periods (fewer than 100 transactions). Furthermore, these profits are almost entirely a consequence of the bid-ask spread. When specialist trades are assumed to take place at the midpoint of the prevailing quotes, the implied profits are sharply reduced. The presence of small positive components in the quote midpoint profits over the short and medium terms (under 100 transactions) does suggest, however, that specialists may be able to anticipate price reversals over these horizons. Estimates of long-term profit components are too noisy to permit inference.

Finally, the analysis turns to dynamic modeling of the interactions among quotes, trades, and inventories. There is overwhelming evidence that shocks to inventories lead at least to transient effects on quotes. The impulse response analyses suggest that the impact of a trade is higher if the specialist is the counterparty to the trade than otherwise. The larger impact may arise because specialist market maintenance obligations require participation in trades during times of high informational asymmetries. Alternatively, specialist participation may be serving as a proxy variable for the direction of trades that occur at the midpoint of the quotes, and cannot therefore be signed. There is little support for the classic inventory control mechanism, which predicts that in response to a trade the quotes initially overshoot and then revert as the inventory is rebalanced.

Appendix: The Computational Details of the Spectral Analysis

This appendix provides a summary of the spectral techniques used in the paper and gives an illustrative example. The discussion closely parallels Engle (1974). For an equally spaced time series of length N , the Fourier frequencies are given by $\omega_k = 2\pi k/N$, for $k = 0, \dots, N-1$. The period (length of cycle) corresponding to a frequency $\omega > 0$ is given by $2\pi/\omega$, so the lowest positive frequency in the data corresponds to a component that cycles once over the full sample. The Fourier transform of the data is

$$J_x(\omega_k) = \frac{1}{N} \sum_{t=1}^N x_t e^{-i\omega_k t}. \quad (\text{A1})$$

$J_x(\omega_k)$ is the Fourier component of x_t at frequency ω_k . The data may be recovered using the inverse transform

$$x_t = \sum_{k=0}^{N-1} J_x(\omega_k) e^{i\omega_k t}. \quad (\text{A2})$$

This expresses x_t as the sum of the components at various frequencies. The usual estimate of the crossproduct between two series, x_t and y_t , is

$$\hat{M} = \frac{1}{N} \sum_{t=1}^N x_t y_t. \quad (\text{A3})$$

This estimate is computationally equal to the one formed from the Fourier transforms:

$$\hat{M} = \sum_{k=0}^{N-1} J_x(\omega_k) \overline{J_y(\omega_k)} \quad (\text{A4})$$

where the summation is over all Fourier frequencies and the overbar denotes complex conjugation. Written in this fashion, the contributions to the covariance from different frequencies are clearly visible.

A particular subset of frequencies, such as the ones corresponding to a certain horizon, may be denoted by K , a subset of $\{k: k = 0, \dots, N-1\}$. Engle shows that there are two computationally equivalent ways of estimating the covariance over this frequency subset. The first is simply to restrict the summation in (A4):

$$\hat{M}_K = \sum_{k \in K} J_x(\omega_k) \overline{J_y(\omega_k)}. \quad (\text{A5})$$

The second is to transform the data so that they reflect only the contributions from the desired frequencies. That is, compute the restricted inverse transform

$$x_t^* = \sum_{k \in K} J_x(\omega_k) e^{i\omega_k t} \quad (\text{A6})$$

and similarly compute y_t^* . This corresponds to a selective reconstruction of the data. Then apply the usual formula:

$$\hat{M}_K = \frac{1}{N} \sum_{t=1}^N x_t^* y_t^*. \quad (\text{A7})$$

Some minor implementation details remain. Despite the fact that the Fourier frequencies are in the interval $[0, 2\pi)$, the highest frequency that is actually resolvable in the data is π . Components at higher frequencies are folded into the interval $[0, \pi)$ in the phenomenon known as aliasing. Computationally, this means that if frequency ω_k is included in the band K , then frequency ω_{N-k} should also be included. Standard errors were computed using spectral estimates smoothed with a Daniell window (see Bloomfield

(1976)). Finally, the expected trading profit definition in (2) is actually a cross-product with one of the variables lagged. Although it is possible to compute this cross-covariance from the Fourier transforms of the original data, it is computationally simpler merely to shift the original series by one observation.

As an example of this analysis, consider the following model for security prices and holdings. The stock price per share is

$$p_t = m_t + s_t \quad (\text{A8})$$

where m_t is the expected end-of-trading value of the stock conditional on all time- t information (the implicit "efficient price"), and s_t is a serially uncorrelated disturbance with variance σ_s^2 . In the present context s_t may be thought of as arising from the bid-ask spread. The efficient price follows a random-walk: $m_t = m_{t-1} + w_t$, where the w_t are serially uncorrelated increments that are also uncorrelated with the s_t . At the close of time- t trading, the trader's holdings are n_t shares.

The ability of the trader to buy at the bid and sell at the ask (on average) is represented by the following relation:

$$n_t = -\beta s_t + u_t \quad (\text{A9})$$

where β is a positive constant and u_t is a disturbance uncorrelated with w_t and s_t . The force of this relation is that when the price disturbance s_t is positive, the trader is generally short. The expected trading profits per transaction are:

$$E\Pi_t = E n_{t-1} dp_t = E n_{t-1} (p_t - p_{t-1}) = \beta \sigma_s^2. \quad (\text{A10})$$

The spectral representation of the expected profits is developed as follows. The summand in (A5), $J_x(\omega_k) \bar{J}_y(\omega_k)$, gives the contribution to the sample cross-product at frequency ω_k . The population analog to this quantity is the cospectrum. To construct the cospectrum, first define the cross-covariance-generating function:

$$g(z) = \sum_{k=-\infty}^{+\infty} \text{Cov}(dp_t, n_{t-1-k}) z^k = \beta \sigma_s^2 (1 - z^{-1}). \quad (\text{A11})$$

The cospectrum between dp_t and n_{t-1} at frequency ω is the real part of the cross-covariance-generating function evaluated at $z = e^{-i\omega}$:

$$Co(\omega) = \text{Re}\{\beta \sigma_s^2 (1 - e^{i\omega})\} = \beta \sigma_s^2 (1 - \cos \omega). \quad (\text{A12})$$

The expected profits (expected cross-product) given in (A10) can be recovered by integrating

$$E\Pi_t = E dp_t n_{t-1} = \frac{1}{2\pi} \int_{-\pi}^{+\pi} Co(\omega) d\omega = \beta \sigma_s^2. \quad (\text{A13})$$

The cospectrum explicitly measures the contributions to profits from various frequencies. It reaches a maximum at π . This corresponds to the most rapidly

varying component of the spectrum (the component with the shortest period). In moving to lower frequencies (longer periods) the cospectrum falls off rapidly. This attenuation is a reflection of the fact that the profits are concentrated at the short horizons.

To further amplify this point, consider a "short-run" profit estimate in which only frequencies with absolute values greater than x are considered. This corresponds to restricting the integration in (A13):

$$\begin{aligned} E\Pi_t(\text{short-run}) &= \frac{1}{2\pi} \left(\int_{-x}^{-\pi} Co(\omega) d\omega + \int_x^{\pi} Co(\omega) d\omega \right) \\ &= \frac{1}{\pi} \beta \sigma_s^2 (\pi - x + \sin x) \end{aligned} \quad (\text{A14})$$

At the cutoff frequency that corresponds to a period of 20 transactions, $x = 2\pi/20$, $E\Pi_t(\text{short-run}) = 0.9984\beta\sigma_s^2$. Clearly, this captures virtually all of the trading profits.

The advantage of using the short-run spectral profit estimate is that it retains substantial power when the data are corrupted by long-term disturbances. Suppose, for example, that (A9) is modified to incorporate a component that might be interpreted as long-term speculative holdings:

$$n_t = n_t^* - \beta s_t + u_t \quad (\text{A15})$$

where n_t^* exhibits large variance and strong persistence (like the specialist inventory levels). We may model this as:

$$n_t^* = \alpha n_{t-1}^* + v_t \quad (\text{A16})$$

where v_t is uncorrelated with $\{s_t, w_t, u_t\}$. Due to this lack of correlation, this modification leaves the cospectrum function unchanged.

The statistical properties of the model are greatly affected, however. If we attempt to estimate trading profits by computing the average cross-product, the estimate will be subject to large errors introduced by long persistence in the long-term speculative holdings. The short-term spectral estimate is substantially more powerful. Consider the results of 100 simulations each with a sample size of 512 observations, $\text{Var}(w_t) = \text{Var}(s_t) = \text{Var}(u_t) = 1$, $\text{Var}(v_t) = 10$, $\beta = 1$ and $\alpha = 0.99$ (roughly the first-order autocorrelation in the actual inventory data). The mean of the unrestricted profit estimates was 1.249 with a standard deviation of 1.007. The mean of the short-term profit estimates was 1.027 with a standard deviation of 0.258. Thus, use of the restricted spectral profit estimate leads to a substantial reduction in the standard error of the profit estimate.

REFERENCES

- Admati, A. R., and P. Pfleiderer, 1988, A theory of intraday patterns: Volume and price variability, *Review of Financial Studies* 1, 3-40.
 Amihud, Y., and H. Mendelson, 1980, Dealership markets: Market making with inventory, *Journal of Financial Economics* 8, 31-53.

- Bagehot, W. 1971, The only game in town, *Financial Analysts Journal* 38, 50–59.
- Bloomfield, P., 1976, *Fourier Analysis of Time Series: An Introduction* (John Wiley and Sons, New York).
- Bradfield, J., 1979, A formal dynamic model of market making, *Journal of Financial and Quantitative Analysis* 14, 275–291.
- Copeland, T., and D. Galai, 1983, Information effects on the bid-ask spread, *Journal of Finance* 38, 1457–1469.
- Easley, D., and M. O'Hara, 1987, Price trade size, and information in securities markets, *Journal of Financial Economics* 19, 69–90.
- Engle, R., 1974, Band spectrum regression, *International Economic Review* 15, 1–11.
- , 1978, Testing price equations for stability across spectral frequency bands, *Econometrica* 46, 869–881.
- Fuller, W. A., 1976, *Introduction to Statistical Time Series* (John Wiley and Sons, New York).
- Garman, Mark, 1976, Market microstructure, *Journal of Financial Economics* 3, 257–275.
- Glosten, L. R., 1987, Components of the bid-ask spread and the statistical properties of transaction prices, *Journal of Finance* 42, 1293–1307.
- , 1989, The role of the monopolist specialist, *Journal of Business* 62, 211–235.
- and P. R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.
- Granger, C. W. J., and R. Engle, 1983, Applications of spectral analysis in econometrics, in D. R. Brillinger and P. R. Krishnaiah, eds.: *Handbook of Statistics*, Vol. 3 (Elsevier Science Publishers, New York).
- Harris, L. E., 1987, Transaction data tests of the mixture of distributions hypothesis, *Journal of Financial and Quantitative Analysis* 22, 127–141.
- Hasbrouck, J., 1988, Trades, quotes, inventories and information, *Journal of Financial Economics* 22, 229–252.
- , 1991a, Measuring the information content of stock trades, *Journal of Finance* 46, 179–208.
- , 1991b, The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies* 4, 571–595.
- , 1993, Assessing the quality of a security market: A new approach to transaction cost measurement, *Review of Financial Studies* 6, 191–212.
- Ho, T. S. Y., and R. G. Macris, 1984, Dealer bid-ask quotes and transaction prices: An empirical study of some AMEX options, *Journal of Finance* 39, 23–45.
- Ho, T. S. Y., and H. R. Stoll, 1983, The dynamics of dealer markets under competition, *Journal of Finance* 38, 1053–1074.
- Kyle, A. S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.
- Madhavan, A., and S. Smidt, 1991, A Bayesian model of intraday specialist pricing, *Journal of Financial Economics* 30, 99–134.
- , 1993, An analysis of changes in specialist inventories and quotations, *Journal of Finance* 48, 1595–1628.
- Merton, R. C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.
- Neuberger, A., 1992, An empirical examination of market maker profits, Working paper, London School of Business.
- New York Stock Exchange, 1991, *New York Stock Exchange Fact Book* (New York Stock Exchange, New York).
- O'Hara, M., and G. S. Oldfield, 1986, The microeconomics of market making, *Journal of Financial and Quantitative Analysis* 21, 361–376.
- Smidt, S., 1979, Continuous vs. intermittent trading on auction markets, *Journal of Financial and Quantitative Analysis* 14, 837–866.
- Son, G., 1991, Dealer inventory position and intraday patterns of price volatility, bid/ask spreads and trading volume, Working paper, Department of Finance, Kellogg School, Northwestern University.

- Stoll, H. R., 1976, Dealer inventory behavior: An empirical investigation of NASDAQ stocks, *Journal of Financial and Quantitative Analysis* 11, 359–380.
- , 1985, The stock exchange specialist system: An economic analysis, Monograph 1985-2, Salomon Brothers Center, Stern School, New York University.
- , 1989, Inferring the components of the bid-ask spread: Theory and empirical tests, *Journal of Finance* 44, 115–134.
- U. S. Securities and Exchange Commission, 1971, *Institutional Investor Study Report* (Arno Press, New York).
- Zabel, E., 1981, Competitive price adjustment without market clearing, *Econometrica* 49, 1201–1221.