

Life in the Pits: Competitive Market Making and Inventory Control

Steven Manaster
University of Utah

Steven C. Mann
Texas Christian University

We use futures transaction data to investigate cross-sectional relationships between market-maker inventory positions and trade activity. The investigation documents strongly that traders control inventory throughout the trading day. Despite this evidence of inventory management, typical inventory control models are contradicted by our data. These inventory models predict that market-maker reservation prices are negatively influenced by inventory. Surprisingly, our evidence shows, as a strong and consistent empirical regularity, that correlations between inventory and reservation prices are positive. We interpret the evidence as consistent with active position taking by futures market floor traders.

We thank Chester Spatt (the editor), two anonymous referees, Jeff Coles, Avner Kalay, Ron Lease, Peter Locke, Uri Loewenstein, Fallaw Sowell, Elizabeth Tashjian, Jaime Zender, and participants at the 1993 Western Finance Association meetings, the 1994 UCLA/USC/NYSE Microstructure Conference, the 1994 Chicago Board of Trade Fall Research Seminar, the 1995 SWFA meetings, and seminar participants at the University of Utah, Texas Christian University, Indiana, Virginia Tech, and New Hampshire for helpful discussions and comments. We thank Zhiming Zhang for adding to our understanding of the Brownian bridge process. Appropriately, the authors retain sole responsibility for any errors. We thank Ron Hobson, Jonathan Smith, Hans Dutt, and Jay Huhman of the Commodity Futures Trading Commission for assistance in the provision of the data, and the Utah Supercomputing Institute (particularly Rich Reynolds) for technical support. Steven C. Mann acknowledges financial support from the Marriner S. Eccles Foundation and the Charles Tandy American Enterprise Center. Address correspondence to Steven C. Mann, M. J. Neeley School of Business, Texas Christian University, P.O. Box 32868, Fort Worth, TX 76129.

Two ideas dominate the theoretical market microstructure literature: inventory control¹ and adverse selection.² Inventory control models assume that market makers face exogenous demands to buy and sell. These market makers earn profits by selling at the ask price and buying at the bid price. The risk associated with this profitable trading is inventory risk. The models predict that market makers will manage risk and control inventory levels by adjusting bid and ask prices. When inventory is greater (less) than desired, prices are reduced (increased) to motivate sales (purchases). Adverse selection models analyze trading environments populated by agents with differential information. The models typically assume that market-maker inventory has no impact on prices in order to facilitate analysis of issues such as the information revealed by order flow.

This article uses detailed futures audit trail transaction data to provide evidence on cross-sectional inventory behavior and relationships between inventory and prices. The evidence shows that neither the adverse selection models nor the inventory control models provide an accurate description of market-maker behavior.

With respect to the inventory control literature, we find that market makers manage the levels of their inventories, but the predictions regarding price and inventory together are strongly contradicted in our data. For example, we find that when dealers hold positive inventory, they tend to become net sellers, consistent with inventory control theory. However, they tend to sell at higher prices, not lower prices, in direct contradiction of traditional inventory control theories.

With respect to adverse selection, its principal conclusion that price responds to the information content of order flow is verified empirically in our data. However, some simplifying assumptions employed to develop that conclusion are contradicted in our experiments. Our observation that market makers attempt to control their inventory levels contradicts the assumption of *inventory neutrality* implicit in most adverse selection models.³ We use the term inventory neutrality to describe market-maker behavior that is impervious to inventory levels [as in Admati and Pfleiderer (1988) and Kyle (1985)]. Inventory neutrality is a strong assumption, requiring risk-neutral utility for the market

¹ Inventory control articles include Amihud and Mendelson (1980), Biais (1993), Demsetz (1968), Garman (1976), Stoll (1976), Ho and Macris (1985), Ho and Stoll (1980, 1981, 1983), Laux (1995), Mildenstein and Schleef (1983), O'Hara and Oldfield (1986), and Zabel (1981).

² Some examples of adverse selection models include Copeland and Galai (1983), Kyle (1985), Glosten and Milgrom (1985), Admati and Pfleiderer (1988, 1989), Foster and Viswanathan (1990), and Benveniste, Marcus, and Wilhelm (1992).

³ Work such as Subrahmanyam (1991) suggests that the primary insights of the adverse selection literature may not require inventory neutrality.

makers, zero inventory carrying costs, and either no constraints on wealth or zero costs for default.

Interestingly, empirical attempts to contradict the assumption of inventory neutrality through examination of the time-series behavior of prices or of market-maker inventories are likely to be inconclusive. Some inventory neutral models, such as Kyle (1985), suggest time series for inventory and prices that are similar to predictions of inventory control models.⁴ Fortunately, cross-sectional tests are available to resolve the ambiguity, and our examination of the cross-sectional relationship between inventory positions and trade activity provides convincing evidence of inventory management. At a typical point in time, traders with long positions are the most active sellers, while traders with short positions are the most active buyers.

Additional price evidence contradicts the inventory neutral paradigm. We find that the price change in response to customer order flow (market depth) varies, conditioned on market-maker inventory levels. Inventory neutrality is inconsistent with market depth that varies with inventory.

Our empirical findings suggest avenues for future modeling design. From our investigation we infer that market makers have informational advantages that enable them to adjust inventory in anticipation of favorable price movements. Thus, this article's evidence provides support for the recent articles by Madhavan and Smidt (1993) and Spiegel and Subrahmanyam (1995), in which market makers have more information than "noise" or liquidity traders, but face adverse selection from traders with access to finer information partitions.

The rest of the article is organized as follows. Section 1 describes the data. Section 2 provides analysis of the time-series and cross-sectional behavior of trader inventory. Our investigation of the relationship between individual trader inventory and prices is presented in Section 3. Section 4 examines relationships between characteristics of aggregate pit inventory and market liquidity. Section 5 provides a concluding comment.

⁴ Kyle's dealer inventory can be shown to conform to the nonstandard Brownian bridge process

$$dI(t) = - \left[\frac{v - p_0}{(1-t)\lambda} + \frac{I(t)}{1-t} \right] dt - dU(t),$$

where net "noise" trades $dU(t) \sim N(0, \sigma_t^2)$; $\lambda > 0$ is the price response per unit net order flow; $I(t)$ is dealer inventory at the start of instant t ($0 < t < 1$); and $\bar{v}$, the ex post liquidation value of the asset, is normally distributed with mean p_0 . The informed trader observes v , the realization of $\bar{v}$. For $v = p_0$ inventory follows a standard Brownian bridge process, where if $I(0) = 0$, then $I(1) = 0$ with probability one. Although Kyle's model is inventory neutral the time series for inventory is in some cases indistinguishable from the predictions of inventory control models.

1. Data

The CFTC provided audit trail transaction records for all Chicago Mercantile Exchange (CME) futures trades for the first half of 1992. This data set consists of more than 12 million records that detail complete trade history for more than 2,000 individual floor traders in 16 different commodity pits. To protect trader privacy, the CFTC mapped each trader's audit trail identification (exchange badge numbers) to a randomly selected number unique to each trader. Therefore, the data provides a 6-month history of trade activity for each trader, but codes each trader's badge number.⁵ Besides the audit trail data, we also use Time and Sales data that records all price changes for the contracts and period represented in the sample.

The audit data records each trade twice, once for each party to each trade. The trade is timed to the nearest minute. Traders report time in 15-minute brackets. An exchange algorithm known as computerized trade reconstruction (CTR) uses each trader's independently reported sequence of trades in conjunction with the time and sales data to time each trade to a specific minute of the bracket. While some timing errors are likely, the timing of the trade is a critical element in the use of the audit trail data in internal (exchange) and external (CFTC enforcement) investigations of legal trading practices.

In addition to trade time, the audit trail record details price, quantity, the contract, and the traders' identification. Unique to this data, the record also specifies the trade direction (whether the trade was a buy or a sell) and a classification of the customer types for both sides of the trade. There are four customer type indicators (CTI), labeled 1 through 4. The CTI 1 trades are market-making trades for personal accounts, and account for 46.6% of total CME volume for the period. CTI 2 trades (6.8% of volume) are trades executed for the account of the trader's clearing member. CTI 3 trades (6.2% of volume) are trades executed for the account of any other exchange member.⁶ CTI 4 trades are the trades of outside customers, and account for 40.4% of volume. For example, if a commercial clearing member's floor trader executes a trade for the firm, it will be a CTI 2 trade, but trades for any other member's account will be CTI 3 trades. The most frequent CTI combination is a customer order (CTI 4) filled by a market maker (CTI 1), or local.⁷

⁵ Fishman and Longstaff (1992) use a smaller sample of audit trail data (15 days for one pit at the Chicago Board of Trade) in their investigation of dual trading.

⁶ Only a subset of exchange members are also members of the exchange clearinghouse. Floor traders that are not clearing members must maintain margin accounts with their clearing member in order to clear trades.

⁷ The term "local" is futures market jargon for "professional market maker." While overall local

2. Inventory Management

2.1 Tracking inventory changes

The audit data times trades to the nearest minute. We therefore track each local's (CTI 1) inventory changes minute by minute. For each commodity, we define each trader's change in inventory as the difference between the number of contracts of all delivery months bought and sold during the minute. This definition treats all contract delivery months as essentially the same security, as price movements on different delivery months are strongly correlated.

Tracking intraday inventory changes is straightforward. However, tracking inventory over longer periods is problematic. About 6.2% of total CME volume for the period is listed as CTI 3 trade, or trades executed on behalf of other members other than the trader's clearing member. The data does not identify the other member, so that if some proportion of CTI 3 volume is executed on behalf of the account of market makers, observed inventory changes may not provide a complete history.⁸ The CME's mutual offset system with SIMEX (Singapore) also poses potential problems, as traders can unwind positions on SIMEX overnight. SIMEX trades are not included in the audit trail data. Due to CTI 3 trades, SIMEX, and potential data errors, trader inventory positions that we calculate at the end of that day may not precisely correspond to actual opening positions for the next trading day. If traders pervasively carry inventory overnight, then the audit data is not well suited for tracking inventory changes over periods longer than a day. For the empirical analysis we assume that each trader begins each day with zero inventory.

There is strong anecdotal support for the conjecture that traders rarely take positions home with them, preferring to end the day with zero inventory, or a "sleeping position."⁹ We provide evidence regarding the propensity of traders to close out positions by the end of the day by examining daily net inventory changes. Figure 1 shows the

volume is higher than customer volume, locals are more likely to trade with customers than each other. Concurrently, customers are much more likely to trade with a local than another customer. Overall, 48.7% of total volume in our sample represents customers trading against locals.

⁸ Our conversations with floor traders lead us to believe that delta hedging by futures options traders accounts for a large part of CTI 3 volume. The conversations indicated that the CTI 3 order flow is unidirectional, as the futures floor traders do not use the futures options to hedge.

⁹ Kolb (1991) categorizes floor traders as either scalpers, day traders, or position traders. Kolb (p. 158) writes, "The overwhelming majority of speculators are either scalpers or day traders, which indicates just how risky it can be to take a position home overnight." Trader studies by Silber (1984) and Working (1967) confirm this "market lore." Working provides empirical evidence on the distribution of floor trader inventory positions, and reports that 44% carried no overnight inventory, another 23% carry less than 10% of their daily volume home for either night, and of the remaining 33%, a "substantial" number carried no more than 1 contract home on either night. Trader reluctance to carry positions overnight is influenced by margin funding costs.




Figure 1
Trader daily inventory changes

Inventory changes are standardized across pits by dividing the ending number of contracts by the pit median trade size. The figure represents 135,015 trades/day, except for 1,528 (1.13%) inventory changes less than -25 and 1,456 (1.07%) inventory changes greater than 25.

frequency distribution of the calculated daily inventory changes. To facilitate comparison across pits, changes are scaled by the median trade size for each pit.¹⁰ Well over half the 135,015 daily trader inventory changes are less than or equal in size to one median trade. For raw (nonscaled) inventory changes, more than 50% of daily changes were no more than two contracts.

The evidence shows that daily inventory changes are concentrated about zero. If traders are starting the day flat, they are generally ending the day flat. Of the observed nonzero daily inventory changes, many are likely due to CTI 3 trades, errors, and/or pending SIMEX offsets. Based on the foregoing observations we believe that our procedure of assuming that all traders begin the day with a zero inventory position provides the most accurate estimate of market-maker inventory that is attainable with the available data.

2.2 Inventory management over time

Consider a simple model of minute-by-minute inventory time-series behavior consistent with inventory control models:

$$\Delta I_t = \alpha + (\rho - 1)I_t + \epsilon_t. \quad (1)$$

¹⁰ In our sample, the median trade is one contract for the S&P 500 index, hogs, pork bellies, lumber, and feeder cattle. The median trade is two contracts for the live cattle, Swiss franc, pound, S&P Midcap 400, and the Canadian and Australian dollar. Three contracts is the median trade for the duetschemark, yen, and the Nikkei index. Interest rate contracts trade in higher quantity, as median trades are four contracts for T-bills, 10 contracts for LIBOR, and 15 contracts for Eurodollars.

We define I_t as inventory at the start of minute t , and define ΔI_t as the net inventory change during the minute t . Inventory control implies that market makers adjust inventory toward a desired level, inducing mean reversion. A method of testing for inventory control is to examine whether a given inventory series follows a random walk. If $\rho = 1$, then inventory is a random walk; mean reversion predicts that $\rho < 1$.

A random walk inventory hypothesis seems to be an easy target. Stoll (1976) presents evidence of NASDAQ net daily inventory changes consistent with mean reverting inventory. However, other than Stoll's groundbreaking article, random walk inventory has been a remarkably stalwart "straw man." Hasbrouck (1988), proxying inventory with estimated order flow, finds evidence of inventory mean reversion for only the lowest decile NYSE stocks. Madhavan and Smidt (1993) formally test specialist inventory for a unit root, but reject the random walk for less than half the 16 stocks in their sample. Hasbrouck and Sofianos (1993) examine inventory autocorrelations for a sample of 144 NYSE stocks and estimate that inventory adjustment takes place at a "very slow rate."

We directly estimate Equation (1) for daily inventory time-series of individual traders. Table 1 reports summary statistics on the estimates of $\hat{\rho}$ for all daily individual inventory series for which a trader traded at least 20 times. In marked contrast to evidence from equity markets, floor trader inventory exhibits strong mean reversion over a short time horizon. Hasbrouck and Sofianos (1993) and Madhavan and Smidt (1993) report that it takes a typical equity specialist a matter of weeks to reduce an inventory imbalance by 50%. Table 1 shows that the median S&P 500 index trader reduces inventory by almost 50% in the next trade!¹¹ Table 1 also reports the percentage of daily series that Phillips–Perron [Perron (1988)] tests reject the null hypothesis that $\rho = 1$.¹² For the S&P 500 index pit, 79.6% of the Phillips–Perron statistics reject the random walk hypothesis at the 5% level. Floor trader inventory does not follow a random walk, but instead exhibits very rapid mean reversion.

¹¹ Our construction of inventory, assuming that traders start flat with zero inventory, has potential impact on the results reported in Table 1. If any traders begin trade with nonzero inventory, then our measured inventory deviates from true inventory by a constant for each day of trading. However, as each regression summarized in Table 1 represents one day of trading, the effect is that of adding a constant to inventory, which has no impact on the slope coefficient. The impact will be reflected in the regression intercept. It may be more precise to state that for the median S&P 500 trader, the deviation of absolute inventory from desired inventory is reduced by 50% in one trade.

¹² The Phillips–Perron test reported ($Z(t_{\alpha}^*)$), is used in conjunction with the Newey and West (1987) adjustment procedure with three lags.

Table 1**Floor trader inventory time-series: estimates of the model $\Delta I_t = \alpha + (\rho - 1)I_t + \epsilon_t$**

Pit	Number of series	Median number of trades	Median $\hat{\rho}$	Percentage of tests rejecting $\rho = 1$
S&P 500 index	20,490	49	0.51	79.6
Eurodollars	10,518	32	0.55	62.5
Deutschemark	7,681	48	0.43	85.5
Swiss franc	4,945	47	0.54	75.9
Yen	3,656	46	0.48	83.8
Live cattle	3,402	41	0.70	50.6
Pound sterling	2,869	49	0.51	79.9
Live hogs	2,126	40	0.70	43.6
Pork bellies	2,060	38	0.60	58.8
Canadian dollar	972	37	0.64	55.1
Feeder cattle	914	36	0.75	31.6
T-bills	584	26	0.63	41.1
Lumber	455	31	0.74	27.7
Nikkei index	179	27	0.55	53.1
LIBOR	147	23	0.70	17.7
S&P Midcap 400	22	23	0.51	54.5
Australian dollar	14	24	0.78	14.3

The table summarizes the results of estimating the model $\Delta I_t = \alpha + (\rho - 1)I_t + \epsilon_t$ (where I_t is trader inventory at the start of minute t and ΔI_t is the inventory change during minute t), for all trader daily inventory series with at least 20 trade/minutes (minutes with at least one trade). The last column reports the percentage of Phillips-Perron [Perron (1988)] unit root tests ($Z(\hat{\rho})$) that reject the null that $\rho = 1$ at the 5% level. The Phillips-Perron tests incorporate the Newey and West (1987) modification procedure with three lags.

2.3 Inventory management in the cross section

Competitive dealer models [e.g., Biais (1993), Ho and Macris (1985), Ho and Stoll (1983), and Laux (1995)] make strong predictions about the cross-sectional relationship between relative inventory positions of market makers and contemporaneous trade activity. In competitive dealer models, the trader with the longest inventory position is the most likely to sell, and the shortest trader is the most likely to buy. This prediction has been untested due to lack of cross-sectional trader inventory data. The futures audit data provides a unique opportunity to investigate cross-sectional inventory behavior.

Table 2 examines the contemporaneous relation between trades and inventory for all 1-minute periods with reasonable activity. We report Spearman rank correlations between trader inventory (I_{it}) rank (at the start of the minute) and subsequent inventory change (ΔI_{it}) rank for all minutes in which at least 10 locals traded for personal account (five locals for the lower volume pits).¹³ Table 2 reports the number of valid minutes for each pit (e.g., 41,734 for the S&P 500), the median rank correlation, and the percent of negative correlations. For example, 85.6% of the correlations between start-of-minute inventory

¹³ We report the nonparametric Spearman rank correlation rather than the more familiar Pearson statistic due to the Spearman statistics' superior small sample properties with potentially nonnormal data.

Table 2
Cross-sectional inventory management: Spearman rank correlations (r_s)

Pit	Number of minutes	Percent of $r_s < 0$	Median r_s	Median number of observations	% negative with p -value < 0.25	% positive with p -value $< .025$
Panel A: Spearman rank correlation statistics calculated for all minutes that at least 10 locals traded for personal account						
S&P 500 index	41,734	85.6	-0.26	25	30.4	0.3
Eurodollars	19,508	83.5	-0.26	19	27.7	0.5
Deutschemark	19,421	84.5	-0.32	15	27.4	0.4
Swiss franc	10,709	82.0	-0.31	13	23.1	0.4
Live cattle	7,884	71.8	-0.18	13	12.9	0.9
Yen	5,948	83.1	-0.33	12	25.4	0.5
Pound sterling	3,069	79.8	-0.31	11	24.2	0.8
Pork bellies	2,466	76.5	-0.22	12	18.8	0.9
Live hogs	2,359	67.1	-0.15	12	14.3	1.5
Panel B: Due to lower trading volume, we report statistics for the following pits based on all minutes that at least 5 locals traded for personal account						
Feeder cattle	2,275	59.6	-0.13	6	11.6	2.4
Canadian dollar	2,187	69.3	-0.30	5	13.9	1.2
T-bills	1,785	59.0	-0.11	6	15.9	3.2
Lumber	1,510	63.5	-0.21	6	15.8	2.0
LIBOR	280	59.4	-0.06	6	18.2	2.0
Nikkei index	209	74.6	-0.39	5	25.4	1.7
S&P Midcap 400	43	70.5	-0.34	5	13.6	2.3

The Spearman correlations (r_s) are cross-sectional correlations of trader inventory (I_{it}) ranks at the start of each minute and subsequent trader inventory change (during the minute) (ΔI_{it}) ranks. P -values come from treating the statistic calculated as $r_s \sqrt{n-2} / \sqrt{1-(r_s)^2}$ as coming from a t -distribution with $n-2$ degrees of freedom.

and subsequent trade are negative for the S&P 500 pit, and the median correlation is -0.26 . For the CME as a whole, generally 70–80% of active trading minutes have negative correlations.

The last two columns of Table 2 report the proportion of negative and positive rank correlation estimates (transformed into t -statistics) that are classically significant at the two-sided 5% level. For the S&P 500 pit, 30.4% of the estimated correlations are negative and significant, while 0.3% are positive and significant.

The evidence reported in Table 2 provides much support for competitive dealer models. As the models predict, the most active sellers are traders with long positions, and the most active buyers are the traders with shorter inventory positions.

3. Individual Inventory and Relative Trade Prices

Inventory control models suggest that trades increasing the absolute level of excess dealer inventory are made on terms that are favorable to the dealer. Favorable terms compensate dealers for the increased

cost and/or risk exposure associated with increases in absolute inventory. Conversely, the models suggest that trades reducing dealer inventory are made on terms that are unfavorable to dealers.

Evidence presented in Section 2 shows that futures floor traders manage their inventory continually throughout each trading day. Does the desire to manage inventory impact trade performance? If bid and ask prices never changed, traders accumulating positive inventory due to buying at the bid would be able to wait until they could sell at the ask, thus earning the bid-ask spread while managing inventory costlessly. But prices do change. Consider a trader buying at the bid one moment and attempting to sell at the ask the next moment. Frequently the market ask has changed between the moment of purchase and the moment of attempted sale. Indeed, as adverse selection models show, the price change is likely to be correlated with the direction of customer order flow and therefore have a negative expected impact on the trader's inventory. However, even if order flow has no information content, pure uncertainty about the income from this market-making round trip is "costly" to a risk-averse market maker.

In fact, most futures markets exhibit considerable short-term volatility, and market makers buy at varying prices in brief intervals. The existence of stable bid and ask prices seems unlikely given rapid price movement in a competitive pit trading environment. Indeed, as Stoll (1992, p. 76) writes, "In pit trading, as on futures markets, many bilateral negotiations occur simultaneously, and slightly different prices may be agreed upon at the same time."

Table 3 provides evidence on short-term variation in market-maker purchase prices for each pit's most active daily contract (usually the nearby). The table reports distributional percentiles for the number of different buy prices for CTI 1 (personal account) trades during 1-, 5-, and 15-minute intervals. For example, Table 3 shows that market makers buy the most active S&P 500 contract at eight different prices in the median 5-minute interval. As the tick is \$25 for the S&P 500 contract, the price range of market-maker purchases is at least \$175 during the median 5-minute period.¹⁴

Traders have the opportunity to transact at a variety of prices during short time frames. In a typical period, some floor traders buy contracts at relatively low prices and some buy at relatively high prices. Given a decision to buy, executing the trade at the lowest price available is clearly desirable. We introduce trade execution skill as a measure designed to indicate an ability to transact at "desirable prices."

¹⁴ While \$175 is the minimum price range for eight different prices with a \$25 minimum price variation (tick), larger price ranges are possible if some ticks are skipped. Market-maker sell-price variability, not reported, is virtually identical to buy-price variability.

Table 3
Price variability: number of different prices for local (CTI 1) buy trades (purchases)

Pit	1-minute intervals			5-minute intervals			15-minute intervals		
	Med	Q3	P90	Q1	Med	Q3	Q1	Med	Q3
S&P 500 index	1	2	3	6	8	10	10	14	18
Eurodollars	1	1	2	1	2	2	2	2	3
Deutschemark	1	2	3	3	5	6	6	8	11
Swiss franc	1	2	3	3	5	7	6	9	13
Yen	1	2	2	2	4	5	4	6	10
Pound	1	2	2	3	4	6	6	8	12
Live cattle	1	1	2	2	3	4	4	6	8
Hogs	1	1	2	2	2	3	3	5	6
Canadian dollar	1	1	2	1	2	2	2	3	4
Pork bellies	1	1	2	2	3	4	4	6	8
Feeder cattle	1	1	2	1	1	2	2	3	4
90-day T-bills	1	1	1	1	1	1	1	1	2
Lumber	1	1	1	1	2	2	2	3	5
LIBOR	1	1	1	1	1	1	1	1	1
S&P 400 Midcap	1	1	1	1	1	1	1	1	2
Australian dollar	1	1	1	1	1	1	1	1	2

The number of different prices for local buy trades per period for the daily most active contract, Chicago Mercantile Exchange, 1/1/92–6/30/92.

Q1 = lower quartile, Q3 = upper quartile, Med = median.

For 1-minute intervals the lower quartile is not reported and is replaced by the 90th percentile. Price variability for local sell trades is virtually identical.

We define execution skill for buy trades (purchases) as

$$\pi_{ti}^b = \bar{p}_\tau^b - p_{ti}^b, \quad (2)$$

where $\bar{p}_\tau^b$ is the volume-weighted mean buy price for all floor trader personal account trades during interval τ , and p_{ti}^b is trader i 's buy price paid during minute t where $t \in \tau$.¹⁵ We define execution skill for sales as

$$\pi_{ti}^s = p_{ti}^s - \bar{p}_\tau^s. \quad (3)$$

Therefore, trading skill is measured for each minute of buying and each minute of selling by comparing the trader's price with the benchmark buy or sell price. Benchmark prices are computed for 5-minute intervals, so that a 15-minute bracket has three 5-minute local buy-price benchmarks and three local sell-price benchmarks.

Execution skill associates good trades with positive execution. Positive selling execution shows that a trader sells at a higher price than the average CTI 1 (market maker) sales price, and positive buy execution represents a purchase at a lower price than the average CTI 1 purchase.

¹⁵ If a trader buys the same contract more than once during a minute, the trader's volume-weighted mean price for the minute is used. We report results using 5-minute intervals. Mann (1994) provides evidence using both 1- and 5-minute intervals for the above specification of trading skill as well as alternative skill measures.

Inventory control models imply that relative trade execution will be affected by relative inventory. Models such as Biais (1993), Ho and Macris (1985), Ho and Stoll (1983), and Laux (1995), assuming homogeneous information and homogeneous risk aversion, predict that reservation prices of traders are negatively correlated with inventory. In theory, traders offer prices that discourage inventory increasing trades and encourage trades that reduce inventory exposure. Inventory's predicted impact on trade execution depends on the trade direction. Inventory models suggest that traders with relatively long positions will have good buy execution (charging a premium to increase inventory), but poor sell execution (making a concession to reduce inventory). Conversely, traders with relatively short positions are predicted to have good sell execution, but poor buy execution. We label this predicted trader behavior "price concessions."

Relative inventory is determined by subtracting the average inventory position of the traders in the pit ($\bar{I}_{t,pit}$) from each trader's inventory position ($I_{t,i}$) for the beginning of each minute.¹⁶ As inventory's predicted effects on trade execution depend on the trade direction, we define orthogonal vectors for each trader's relative inventory, I^B (for buys) and I^S (for sells), specified as

$$\begin{aligned} I_{ti}^B &= I_{ti} - \bar{I}_{t,pit} && \text{for buy trades of trader } i \text{ at time } t, \\ &= 0 && \text{for sell trades of trader } i \text{ at time } t, \end{aligned} \quad (4)$$

and similarly,

$$\begin{aligned} I_{ti}^S &= I_{ti} - \bar{I}_{t,pit} && \text{for sell trades of trader } i \text{ at time } t, \\ &= 0 && \text{for buy trades of trader } i \text{ at time } t, \end{aligned} \quad (5)$$

so that the elements of each vector are either relative inventory or zero. To test for inventory effects on trade execution, the following model is estimated once for each trader for the entire sample period:

$$\pi_{ti} = \alpha + \beta(I_{ti}^B - I_{ti}^S) + \epsilon_t, \quad (6)$$

where π_{ti} is execution skill (either buy or sell), α is a constant, and ϵ_t is an error term. The term β measures the effect of relative inventory on subsequent relative trade prices. If traders make price concessions, then we expect β to be positive. Table 4 reports summaries for each

¹⁶ The average inventory position for traders in the pit for a minute is defined as the aggregate pit inventory (the summed current inventory of all traders) divided by the number of traders that have inventory positions. A trader is defined as having an inventory position for a minute if the trader executes personal account trades in the pit prior to or inclusive of each minute. If the pit is net long 400 contracts, and 100 traders have positions, then the average position is four contracts. A trader short three contracts at that time has defined relative inventory of negative seven contracts (short seven contracts).

pit of regression estimates of Equation (6) for each trader with at least 126 trades over the 6-month (126 day) sample period.

Table 4 provides little evidence that traders make price concessions. Roughly 30% of the regressions have an estimated $\hat{\beta}$ that is positive, as suggested by inventory models. Overall, only 2% or 3% of the estimates are positive and significant by classical standards (two-sided 5% level). On the other hand, about two-thirds of the β estimates are negative, and overall, about 20% of the estimates are negative and significant.

Although individual $\hat{\beta}$ s range widely, taken as a whole, inventory's impact on execution skill appears to be contrary to the implications of inventory models. There is little evidence that traders make price concessions for trades reducing absolute inventory, or extract better prices for trades that increase inventory exposure. In fact, most traders appear to obtain better prices for trades that decrease absolute inventory. Thus, while long traders are the most active sellers, the selling activity is not due to price concessions, as they seem to sell at higher, not lower, prices.

We directly compare contemporaneous trade prices of short and long traders in Table 5, classifying traders as either short or long based on existing inventory at the start of each trade minute. We define buy price differences for 5-minute intervals as the mean purchase price secured by buying long traders ($\bar{p}_L^b$) less the mean purchase price secured by buying short traders ($\bar{p}_S^b$). Selling price differences are also calculated for each interval by subtracting the mean sell price obtained by short traders from the mean sell price obtained by long traders. Inventory control models, predicting negative correlations between reservation prices and inventory, imply that $\bar{p}_L - \bar{p}_S < 0$, both for buying and selling price differences.¹⁷

Table 5 summarizes price differences by pit.¹⁸ The table reports the median of buy- and sell-price differences for each pit ($\bar{p}_L - \bar{p}_S$), a t -statistic testing a mean zero null hypothesis, and the percent positive and negative (some price differences are zero). For example, the median price difference for the S&P 500 pit is \$3.25, the mean is reliably positive ($t = 28.7$), and 60.2% of the calculated price differences are positive. For all 12 pits with valid observations, the median price difference is positive. The mean price difference is significantly positive

¹⁷ To conserve space we combine buying and selling price differences in Table 5. The distributions of buying and selling price differences are very similar. See Mann (1994).

¹⁸ Only 5-minute intervals with at least five short traders and five long traders buying for personal account are included. Intervals with no local personal account price variation (e.g., all locals buy at the same price for the entire period) are dropped. Analogous criteria are used to report sell-price differences.

Table 4
Inventory and trade execution: distribution of the inventory coefficient $\hat{\beta}$

Pit	Number of traders	Percent of $\hat{\beta} \geq 0$	Percent of $t(\hat{\beta}) \geq 1.96$	Percent of $t(\hat{\beta}) \leq -1.96$	Median number of observations for each regression	<i>t</i> -statistic for the mean of all regression <i>t</i> -statistics
S&P 500 index	370	30	2	20	3,272	-10.20
Eurodollars	329	34	3	14	1,609	-7.82
Deutschemark	145	36	3	17	2,929	-4.65
Live cattle	116	30	2	20	924	-5.01
Swiss franc	100	32	3	25	1,969	-6.02
Yen	71	27	0	27	2,816	-4.57
Live hogs	65	31	3	23	755	-3.36
Pound	52	27	4	12	3,336	-2.65
Pork bellies	52	23	6	27	1,641	-4.45
Lumber	25	36	4	4	810	-1.62
T-bills	22	55	0	14	1,610	-0.38
Canadian dollar	21	33	0	57	2,493	-3.91
Feeder cattle	23	43	0	13	1,115	-1.08
LIBOR	13	69	8	0	1,001	1.23
Australian dollar	2	0	0	0	1,286	-1.15
S&P Midcap 400	8	75	0	0	402	1.51

The table summarizes for each pit the results from estimating the following regression model for each trader with at least 126 trades over the sample period: $\pi_{it} = \alpha + \beta(I_{it}^B - I_{it}^S) + \epsilon_{it}$, where π_{it} is the execution skill for trader i for trade at time t , and the orthogonal vectors I^B and I^S represent trader i 's relative inventory at time t ; I_{it}^B is trader i 's start-of-minute inventory less average pit start-of-minute inventory if the trade is a buy; $I_{it}^S = 0$ if the trade is a sell; I_{it}^S is the trader's start-of-minute inventory less average pit start-of-minute inventory if the trade is a sale; $I_{it}^S = 0$ if the trade is a buy.

The statistic $t(\hat{\beta})$ for each regression is the t -statistic associated with testing whether the parameter estimate $\hat{\beta}$ is significantly different from zero.

Table 5
Price differences for long versus short traders ($\bar{p}_L - \bar{p}_S$)

Pit	Number of intervals	Median price difference $\bar{p}_L - \bar{p}_S$	<i>t</i> -statistic for difference	Percent negative	Percent positive	Mean customer execution spread
S&P 500 index	19,479	\$ 3.25	28.7	39.7	60.2	\$ 4.33
Eurodollars	9,122	0.73	10.7	43.0	54.3	1.55
Deutschemark	12,558	0.93	10.5	44.3	55.4	3.30
Swiss franc	8,545	1.91	12.1	42.5	57.3	5.30
Yen	5,436	1.82	10.6	41.6	57.9	3.56
Live cattle	3,536	1.00	6.6	42.6	56.6	1.75
Pound sterling	3,629	1.83	3.2	44.6	55.1	4.44
Pork bellies	947	2.26	3.3	42.0	56.6	7.71
Live hogs	564	0.83	3.0	45.9	52.8	3.97
Canadian dollar	100	1.67	1.8	43.0	56.0	3.06
Feeder cattle	16	4.90	0.6	31.3	68.8	6.58
T-bills	8	5.68	0.7	37.5	62.5	3.62

Comparisons of contemporaneous (5-minute brackets) buy or sell prices obtained by traders with long (positive) or short (negative) inventory positions. Every minute, each active trader's mean buy and/or sell price is calculated. Price differences ($\bar{p}_L - \bar{p}_S$) for buy trades and sell trades for each interval are defined as the mean price secured by long traders ($\bar{p}_L$) less the mean price secured by short traders ($\bar{p}_S$). Each price difference compares either long trader buy prices to short trader buy prices or long trader sell prices to short trader sell prices. If traders with long positions trade at prices less than traders with short positions, the differences would be negative. Intervals without at least five long and five short traders with transactions are dropped. Each 5-minute interval may have two price differences, one for buys and one for sells, given enough trades by both long and short traders. Four pits had no intervals that satisfied the criteria.

for 9 of the 12 pits. Consistent with the results of Table 4, long traders are trading at higher prices than short traders. To provide a benchmark for the economic significance of the price difference, mean 5-minute customer execution spreads for the most active contract are also reported in Table 5. The customer execution spread is defined as the mean customer buy price less the mean customer sell price, and is a pure measure of the realized bid-ask spread.¹⁹ Comparing the S&P 500 pit's mean execution spread (\$4.33) to the median price difference (\$3.25), it appears that long traders are trading at economically as well as statistically higher prices than short traders. Evidently, long traders obtain higher selling prices than contemporaneous traders with short positions, but they also pay more to purchase contracts.

Tables 4 and 5 show that long traders, while active sellers, buy and sell at high relative prices. Over time (Table 4), long traders have good skill at selling (selling higher than other traders) and poor skill at buying (paying a higher price to buy). At a particular point in time (Table 5), long traders sell and buy at higher prices than other traders. Conversely, short traders, while active buyers, buy and sell at low relative prices. These results are inconsistent with received inven-

¹⁹ The execution spread is used by the CFTC (1989) and Chang and Locke (1993).

tory models. However, extant competitive dealer models assume both homogenous information and homogenous risk aversion. These assumptions probably do not hold in the trading pits. Development of a formal model incorporating heterogeneous risk aversion and information is beyond our scope. However, the results from Tables 4 and 5 do provide empirical guidance for future modeling efforts. To be consistent with our evidence future models must incorporate heterogeneous information among market makers. Differences in risk aversion alone might produce differences in reservation prices conditional on inventory, but differences in risk aversion alone will never cause market makers with long positions to have reservation prices greater than market makers with short positions. We conjecture that this empirical regularity can only be supported by a model that contains differential information among market makers.

4. Market Liquidity and Aggregate Pit Inventory

In this section we examine predicted relationships between characteristics of aggregate pit inventory and market liquidity characteristics, specifically spreads and depth.

4.1 Inventory dispersion and customer spreads

Biais (1993), Ho and Macris (1985), Ho and Stoll (1980, 1983), and Laux (1995) argue that competitive dealer inventory positions will not diverge substantially, as dealer price concessions act as an equilibrating mechanism for inventory. If inventory does diverge, the models imply that the market bid-ask spread will be reduced. According to the theory, the spread is squeezed as long traders reduce ask prices and short traders increase bids.

Besides potential inventory effects, trade volume and volatility may also affect the spread. Brock and Kleidon (1992) argue that spreads increase as volume (demand for liquidity) increases. Models of the spread from both the inventory control and asymmetric information literature imply that increased price volatility widens the spread. Measuring volume is straightforward. However, a wider spread will affect most price volatility estimates. Therefore, we employ the detail of the audit trail data to calculate a volatility measure that eliminates bid-ask bounce. For each 5-minute period, we compute one quantity weighted standard deviation for customer buy trade prices and one for customer sell trade prices. We define the volatility measure as the maximum of the customer buy-price or sell-price standard deviation. The maximum is used to avoid potential problems such as no customer trade on the buy side despite considerable activity on the sell side. This one-sided volatility measure eliminates bid-ask bounce by

exclusively using either buy or sell prices for each 5-minute interval.

We examine the impact of inventory dispersion on the spread by estimating the following regression model:

$$Spread_t = \alpha + \beta_1 \sigma_t + \beta_2 Volume_t + \beta_3 Range(I^*)_t + \epsilon_t, \quad (7)$$

where the dependent variable is the 5-minute customer execution spread (recall that the execution spread is the mean customer buy price less the mean customer sell price), *volume* is measured as contemporaneous total two-sided (buys and sells) customer volume, and σ_t is the bid-ask bounce-free measure of contemporaneous price volatility. To measure inventory dispersion, we define $Range(I^*)_t$ as an interfractile inventory range of trader inventory positions at the start of a 5-minute period. We calculate the range by ranking trader inventory positions at the start of each minute for all traders active in the commodity that day, before or inclusive of the minute. The inventory range is calculated as an interfractile range of the inventory positions at the start of the first minute of each 5-minute period, where the fractiles are varied depending upon the number of traders. For less than 20 traders, we use the interquintile range (80th less 20th percentile). For at least 20 but less than 100 traders, we use the interdecile range, and for 100 or more traders, we use the intervigintile range (95th less 5th percentile).

Table 6 reports estimates of Equation (7) for each pit. As predicted by all models, volatility is associated with wider spreads. Of 16 estimated coefficients on volatility, all are positive, and only one (Swiss franc) is not classically significant at the 1% level. Contemporaneous volume has a more ambiguous impact; loose interpretation suggests that increased volume is associated with wider spreads for high-volume pits, but reduced spreads in lower volume pits.

The evidence regarding the relationship between inventory ranges and realized spreads is difficult to interpret. The coefficients for inventory range should be negative if, as available theories predict, inventory dispersion narrows the spread. However, only three pits (Eurodollars, yen, and the S&P Midcap 400) have estimated inventory range coefficients that are negative and significant at a classical 5% level. The estimated inventory range coefficient is negative for only 8 of 16 regressions. Our interpretation is that inventory range does not appear to affect the spread in a manner consistent with available theory. However, based on the evidence we report, there is room for other interpretations. For example, if market-maker inventory is held in anticipation of favorable price movements [as described in Ho and Stoll (1981, Section 8)], then at a point in time locals may be content with their divergent inventories and have no incentives to narrow the

Table 6
Inventory dispersion and the spread

Pit	Number of observations	Intercept	Standard deviation of one-sided customer prices	Customer volume	Interfractile inventory range
S&P 500 index	10,325	-0.38 (-0.5)	0.017 (4.8)	0.004 (6.3)	0.083 (2.2)
Eurodollars	8,034	1.40 (5.1)	0.023 (11.1)	0.000 (0.7)	-0.005 (-3.1)
Deutschemark	9,946	0.98 (4.2)	0.050 (15.8)	0.001 (2.9)	-0.010 (-1.5)
Live cattle	5,776	0.30 (1.0)	0.059 (7.38)	0.001 (1.13)	0.012 (1.2)
Swiss franc	9,744	3.41 (8.1)	0.005 (1.3)	0.009 (10.2)	0.007 (0.3)
Yen	9,460	2.62 (8.9)	0.041 (10.0)	0.000 (0.6)	-0.036 (-2.5)
Live hogs	5,092	3.13 (8.8)	0.084 (7.3)	-0.003 (-2.4)	-0.010 (-0.9)
Pound	9,226	1.94 (4.9)	0.075 (17.5)	-0.002 (-1.9)	-0.036 (-1.5)
Pork bellies	4,854	4.79 (7.6)	0.089 (6.6)	0.010 (2.4)	0.030 (0.8)
Lumber	2,466	9.86 (7.9)	0.224 (10.9)	-0.081 (-2.7)	0.073 (0.6)
T-bills	2,213	2.81 (7.9)	0.145 (13.0)	-0.011 (-3.9)	0.001 (0.1)
Canadian dollar	6,093	2.78 (14.2)	0.064 (8.2)	-0.009 (-9.7)	0.013 (1.0)
Feeder cattle	2,983	3.77 (7.3)	0.213 (12.3)	-0.036 (-4.4)	0.050 (1.7)
LIBOR	901	0.94 (3.2)	0.052 (7.7)	-0.003 (-2.1)	-0.007 (-1.0)
Australian dollar	375	3.91 (3.1)	0.305 (6.8)	-0.093 (-2.9)	-0.020 (-0.4)
S&P Midcap 400	908	7.63 (5.0)	0.098 (3.7)	-0.098 (-2.6)	-0.322 (-2.9)

The table reports estimates of the model $S_t = \alpha + \beta_1 \hat{\sigma}_t + \beta_2 Volume_t + \beta_3 Range(I^*)_t + \epsilon_t$, where S_t is a 5-minute customer execution spread (mean customer buy price less mean customer sell price) for each day's most active contract, $\hat{\sigma}_t$ is contemporaneous one-sided customer price standard deviation, volume is aggregate (buy plus sell) customer volume, and $Range(I^*)_t$ is an interfractile range of trader inventory positions at the start of each 5-minute period (t -statistics in parentheses).

spread.

4.2 Pit inventory and market depth

Inventory-neutral models that focus on price discovery assume that market-maker inventory has no impact on the price response to order flow (market depth). In contrast, models that incorporate both price discovery and inventory control, such as Chordia and Subrahmanyam (1992), Diamond and Verrecchia (1991), Laux (1993), and Madhavan and Smidt (1991, 1993), predict that market-maker inventory does affect the price response to order flow.

Our empirical analysis of the relationship between the price response to order flow and aggregate pit inventory is based on a generic

implied relationship between the price response to order flow and inventory:

$$p_{t+1} - p_t = f(\omega_t, I_{pit,t}), \quad (8)$$

where ω_t is net customer order flow (positive for net customer buys) during minute t , p_{t+1} is the price at the end of minute t , and $I_{pit,t}$ is total pit inventory at the beginning of minute t . The common prediction of all models is that $\partial p / \partial \omega > 0$.

The models diverge regarding the effect of inventory, as inventory-neutral models imply that $f(\omega, I) = g(\omega)$ [for example, in Kyle (1985), $g(\omega) = \lambda\omega$]. In contrast, models incorporating price discovery and inventory control imply that market-maker inventory does affect market depth, or that $\partial f(\omega, I) / \partial I \neq 0$. In these models, inventory has an effect because dealers use price to manage inventory. In essence, the models predict that market depth is greater (lesser price response) for order flow that moves dealer inventory to a preferred position, and that order flow that increases market-maker inventory exposure is associated with larger price impacts.

To examine inventory's impact on market depth, we estimate the following regression:

$$p_{t+1} - p_t = \alpha + \beta_1 \omega_t + \beta_2 \delta_t \omega_t + \epsilon_t \quad (9)$$

for 1-minute price changes, net customer (CTI 4) order flow, and a dummy variable δ_t set equal to one if order flow increases absolute aggregate pit inventory, and zero otherwise.²⁰ All microstructure models predict that the coefficient on order flow will be positive. The inventory management models suggest that the coefficient on the dummied order flow will also be positive. Intuition for the prediction can be developed by considering the case of a market maker with long inventory. Customer sales (negative order flow) to the market maker lead to inference that "true" price is lower [as in Kyle (1985)], and this is captured in the term $\beta_1 \omega_t$. However, selling to a long market maker also increases the market maker's inventory exposure, which, according to the theory, further lowers the price at which the market maker is willing to make a purchase.

²⁰ The 1-minute price change is taken from the time and sales data records, which record all price changes to the nearest second. We define the last price for the contract prior to the start of minute t as p_t . If any price changes are recorded for minute t then the last price for the minute is defined as p_{t+1} . If no price changes are recorded, price change is designated as zero. Net customer order flow is calculated as the summed quantity of customer buys of the contract less the total sells for the minute. Price changes and order flow are for the most active contract for each day (usually the nearby, or closest to maturity, contract).

Table 7
Inventory and the price impact of order flow

Pit	Number of observations	Intercept	Contract order flow	Inventory dummy times order flow
S&P 500 index	50,816	0.18 (0.6)	0.200 (14.6)	-0.030 (-1.5)
Eurodollars	33,956	0.13 (1.4)	0.005 (6.9)	-0.007 (-6.7)
Deutschemark	47,618	0.08 (0.5)	0.048 (14.4)	-0.004 (-1.0)
Live cattle	25,024	0.13 (1.5)	0.061 (10.9)	-0.028 (-3.7)
Swiss franc	45,258	0.08 (0.5)	0.101 (13.6)	-0.007 (-0.7)
Yen	42,404	0.03 (0.2)	0.056 (12.0)	0.008 (1.3)
Live hogs	20,239	0.08 (0.7)	0.129 (13.7)	-0.043 (-3.2)
Pound	41,388	0.10 (0.5)	0.100 (8.7)	0.010 (0.7)
Pork bellies	19,798	0.13 (0.7)	0.573 (19.2)	-0.144 (-3.3)
Lumber	9,498	0.29 (0.5)	1.920 (13.0)	-1.620 (-7.3)
T-bills	9,922	-0.06 (-0.4)	0.069 (8.9)	-0.029 (-2.7)
Canadian dollar	24,474	-0.12 (-1.4)	0.087 (10.5)	-0.009 (-0.8)
Feeder cattle	11,449	0.20 (0.9)	0.359 (10.2)	-0.031 (-0.6)
LIBOR	3,362	0.30 (0.9)	0.063 (5.5)	-0.053 (-3.44)
Australian dollar	2,178	0.13 (0.1)	1.613 (5.8)	-1.339 (-4.4)
S&P Midcap 400	4,451	-1.30 (-1.2)	0.285 (1.5)	-0.104 (-0.4)

Model: $p_{t+1} - p_t = \alpha + \beta_1 \omega_t + \beta_2 \delta_t \omega_t + \epsilon_t$. Price (p_{t+1}) is end-of-minute t -price for the most active contract, ω_t is net customer (CTI 4) order flow (most active contract), and δ_t is a dummy variable equal to one if customer order flow increases absolute pit inventory, and equal to zero if customer order flow reduces pit absolute inventory. Pit inventory is aggregate pit trader inventory at the start of minute t (t -statistics in parentheses).

Table 7 reports estimates of Equation (9) for daily most active contracts for each pit. As predicted, order flow moves price. The coefficient for net customer order flow is positive for all 16 pits, and significant at the 1% level for 15 of the 16. However, inventory's predicted impact is not evidenced in the data. Only 2 of the 16 pits (pound and yen) have a positive estimated coefficient on inventory-increasing order flow (neither is significant). In fact, of the 14 negative estimated coefficients for inventory-increasing order flow, 8 are significantly negative at the classical 1% level. As with results for individual trader inventory, the results for aggregate pit inventory seem opposite to the predictions of inventory models.

5. Conclusion

The empirical investigations reported in the previous sections reveal several strong and consistent empirical regularities for the CME futures trading activities. Two principal conclusions deserve emphasis.

First, we conclude that CME locals aggressively manage their inventories. This conclusion is based on the observation that individual market-maker inventories are rapidly mean reverting (Table 1), and that trade direction is negatively correlated with inventory (Table 2). Simply put, long market makers are more likely to sell, short market makers are more likely to buy.²¹ Compared with other studies, which are based on stock trading data, our empirical results show that the speed of inventory adjustment for futures markets is much greater than for equity markets. We do not regard any of these results as surprising.

Second, despite the strong evidence that market makers attempt to control their inventory risk exposure, the influence of inventory on price is contrary to the predictions of inventory control theory. Overall, we find that locals make inventory-reducing trades on more favorable terms (to the local) than inventory-increasing trades (Table 4), we find that price changes are less responsive to order flow that increases (absolute) pit inventory (Table 7), and we find no conclusive evidence that inventory dispersion reduces customer execution spreads (Table 6). These are surprising results. Furthermore, our experiments show that these results are not attributable to differential levels of risk aversion among market makers (Table 5). From this we infer that market makers are not merely passive order fillers, as depicted in some microstructure models, but are active profit-seeking individuals with heterogeneous levels of information and/or trading skill.

Heterogeneous market-maker information and the management of market-maker inventory are important issues. These issues, however, are not simultaneously included in the most widely cited theoretical models of market microstructure. We hope that the evidence we have presented will motivate model building that will abandon the implicit assumption of inventory neutrality and integrate the joint influences of inventory management and asymmetric information on individual and market behavior more precisely.

References

Admati, A. R., and P. Pfleiderer, 1988, "A Theory of Intraday Patterns: Volume and Price Variability,"

²¹ Additional evidence from logit regressions (not reported) shows that the probability that a market maker's next trade will be a sell (buy) increases the longer (shorter) his/her inventory. See Mann (1994).

Review of Financial Studies, 1, 3–40.

Admati, A. R., and P. Pfleiderer, 1989, "Divide and Conquer: A Theory of Intraday and Day-of-the-Week Mean Effects," *Review of Financial Studies*, 2, 189–223.

Amihud, Y., and H. Mendelson, 1980, "Dealership Market: Market-Making With Inventory," *Journal of Financial Economics*, 8, 31–53.

Benveniste, L. M., A. J. Marcus, and W. J. Wilhelm, 1992, "What's Special About the Specialist?," *Journal of Financial Economics*, 32, 61–86.

Biais, B., 1993, "Price Formation and Equilibrium Liquidity in Fragmented and Centralized Markets," *Journal of Finance*, 48, 157–185.

Brock, W. A., and A. W. Kleidon, 1992, "Periodic Market Closure and Trading Volume," *Journal of Economic Dynamics and Control*, 16, 451–489.

CFTC (U.S. Commodity Futures Trading Commission), 1989, "Economic Analysis of Dual Trading on Commodity Exchanges," Report prepared by the Division of Economic Analysis.

Chordia, T., and A. Subrahmanyam, 1992, "Off-Floor Market Making, Payment-for-Order-Flow, and the Tick Size," working paper, University of California, Los Angeles.

Chang, E. C., and P. R. Locke, 1993, "The Performance and Market Impact of Dual Trading: CME Rule 552," working paper, Commodity Futures Trading Commission.

Copeland, T., and D. Galai, 1983, "Information Effects and the Bid- Ask Spread," *Journal of Finance*, 38, 1457–1469.

Demsetz, H., 1968, "The Cost of Transacting," *Quarterly Journal of Economics*, 82, 33–53.

Diamond, D., and R. Verrecchia, 1991, "Disclosure, Liquidity, and the Cost of Capital," *Journal of Finance*, 46, 1325–1359.

Fishman, M., and F. Longstaff, 1992, "Dual Trading in Futures Markets," *Journal of Finance*, 47, 643–671.

Foster, F. D., and S. Viswanathan, 1990, "A Theory of the Interday Variations in Volume, Variance, and Trading Costs in Securities Markets," *Review of Financial Studies*, 3, 593–624.

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

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

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

Hasbrouck, J., and G. Sofianos, 1993, "The Trades of Market-Makers: An Analysis of NYSE Specialists," *Journal of Finance*, 48, 1565–1593.

Ho, T. S. Y., and R. G. Macris, 1985, "Dealer Market Structure and Performance," in Y. Amihud, T. S. Y. Ho, and R. Schwartz (eds.), *Market Making and the Changing Structure of the Securities Industry*, Lexington Books, Lexington, Massachusetts.

Ho, T. S. Y., and H. R. Stoll, 1980, "On Dealer Markets Under Competition," *Journal of Finance*, 35, 259–267.

Ho, T. S. Y., and H. R. Stoll, 1981, "Optimal Dealer Pricing Under Transactions and Return Uncertainty," *Journal of Financial Economics*, 9, 47–73.

- Ho, T. S. Y., and H. R. Stoll, 1983, "The Dynamics of Dealer Markets Under Competition," *Journal of Finance*, 38, 1053–1074.
- Kolb, R. W., 1991, *Understanding Futures Markets*, Kolb Publishing, Miami.
- Kyle, A. S., 1985, "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1336.
- Laux, P. A., 1993, "Intraday Price Formation in the Stock Index Futures Market," working paper, University of Texas at Austin.
- Laux, P., 1995, "Dealer Market Structure, Outside Competition, and the Bid-Ask Spread," *Journal of Economic Dynamics and Control*, 19, 683–710.
- Madhavan, A., and S. Smidt, 1991, "A Bayesian Model of Intraday Specialist Pricing," *Journal of Financial Economics*, 30, 99–134.
- Madhavan, A., and S. Smidt, 1993, "An Analysis of Changes in Specialist Inventories and Quotations," *Journal of Finance*, 48, 1585–1628.
- Mann, S. C., 1994, "Competitive Marketmaking on Futures Exchanges," Ph.D. dissertation, David Eccles School of Business, University of Utah.
- Mildenstein, E., and H. Schleef, 1983, "The Optimal Pricing Policy of a Monopolistic Marketmaker in the Equity Market," *Journal of Finance*, 38, 218–231.
- Newey, W., and K. West, 1987, "A Simple Positive Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–715.
- O'Hara, M., and G. Oldfield, 1986, "The Microeconomics of Market Making," *Journal of Financial and Quantitative Analysis*, 21, 361–376.
- Perron, P., 1988, "Trends and Random Walks in Macroeconomic Time Series," *Journal of Economic Dynamics and Control*, 12, 297–332.
- Silber, W., 1984, "Marketmaker Behavior in an Auction Market: An Analysis of Scalpers in Futures Markets," *Journal of Finance*, 39, 937–953.
- Spiegel, M., and A. Subrahmanyam, 1995, "On Intraday Risk Premia," *Journal of Finance*, 50, 319–339.
- Stoll, H. R., 1976, "Dealer Inventory Behavior: An Empirical Investigation of NASDAQ/NMS Stocks," *Journal of Financial and Quantitative Analysis*, 11, 359–380.
- Stoll, H. R., 1992, "Principles of Trading Market Structure," *Journal of Financial Services Research*, 6, 75–106.
- Subrahmanyam, A., 1991, "Risk Aversion, Market Liquidity, and Price Efficiency," *Review of Financial Studies*, 4, 417–441.
- Working, H., 1967, "Tests of a Theory Concerning Floor Trading on Commodity Exchanges," *Food Research Institute Studies: Supplement*.
- Zabel, E., 1981, "Competitive Price Adjustment Without Market Clearing," *Econometrica*, 49, 1201–1221.