WP/03/115
Revised: 11/13/06

IMF Working Paper

An Intraday Pricing Model of Foreign Exchange Markets

Rafael Romeu

INTERNATIONAL MONETARY FUND

WP/03/115
Revised: 11/13/06

IMF Working Paper

International Capital Markets Department

An Intraday Pricing Model of Foreign Exchange Markets

Prepared by Rafael Romeu[1]

Authorized for Distribution by Donald J. Mathieson

June 2003

Abstract

Dealers learn about asset values as they set prices and absorb portfolio flows. These flows causes inventory imbalances. Previous work argues that dealers deviate from their estimates of asset values to induce flows that unwind inventory imbalances. This study models dealer price-setting using multiple instruments to smooth inventory imbalances and update priors about asset values. This approach shows that canonical models in which price-setting is the only instrument for inventory control, and incoming order flow is the only source of asymmetric information, are misspecified. Thus, estimates of canonical models reject predicted asymmetric information and inventory effects because of omitted and extraneous variables. These estimations miss information from sources other than incoming order flow, and they overemphasize price shading in managing inventories. Estimates of the model presented support heretofore elusive inventory and asymmetric information effects. Price shading is found to have smaller role in inventory management and information effects are shown to be stronger than previously estimated. Additionally, this approach yields direct measures of the structural liquidity cost parameters in the model akin to Kyle's Lambda. For example, estimates presented suggest that a standard $10 million incoming purchase pushes price up by roughly one basis point, and dealers expect to immediately lay-off one-third of every incoming order.

JEL Classification Numbers: G15; F31
Keywords: Foreign Exchange, Microstructure, International Macroeconomics
Author's E-Mail Address: rromeu@imf.org

[1] I am grateful to Roger Betancourt, Michael Binder, Robert Flood, Carmen Reinhart, and John Shea for guidance; and also to Richard Lyons for guidance and data. Thanks also to Richard Adams, Juan Blyde, Alain Chaboud, Giovanni Dell'Ariccia, Andrew DePhillips, Larry Evans, Martin Evans, Stefan Hubrich, Andrei Kirilenko, Don Mathieson, José Pineda, Eswar Prasad, Mike Pries, Pedro Rodriguez, Francisco Vázquez, Citibank FX, seminar participants at the Board of Governors, the IMF Research Department, the University of Maryland Economics Department and R.H. Smith School of Business Finance Department, and the Bank of Canada for comments.

Contents Pages

Figures

Tables

I. Introduction

The evidence supporting a tight relationship between a market's absorption of portfolio flows and its assets' returns is mounting.[2] At the very least, it implies that asset returns depend on how dealers interact with each other and with end users. The question now is how long market trading affects asset returns. Assuming that asset fundamentals follow a random walk, there could be permanent effects if trading reveals new information. For example, dealers aggregating portfolio flows may also aggregate information dispersed in the economy. Conversely, the market's temporary indigestion from absorbing large portfolio shifts may imply transitory effects, as in microstructure inventory models. At the level of the individual dealer, however, there is surprisingly little (if any) evidence supporting theoretically predicted inventory effects. This paper presents a new model of asset trading that shows evidence of both information and inventory effects at the individual dealer level. The empirical results link portfolio flows to asset prices at the highest resolution, and provide direct estimates of the cost of liquidity, asymmetric information, and inventory effects. The results suggests that previous models have underestimated, if not missed or rejected these effects in markets with multiple dealers, such as bond or foreign exchange markets. An example illustrates why.

Consider a foreign exchange (FX) dealer who is trading U.S. dollar-Euro and watching the price of the currency fluctuate throughout the day. Assume that the dealer is constrained with a finite inventory (or, equivalently, inventory costs). If random-walk asset values drive incoming trades, she must respond with an inventory-management strategy or exhaust her supply. Past models suggest that this dealer divert her price away from the equilibrium full-information value to induce trades that compensate for inventory imbalances. But changing prices to induce trades equates to intentionally selling low or buying high. What if there is another way? In markets with multiple dealers she can call other dealers and unload her inventory imbalances on them. This allows the dealer another instrument for managing inventory and learning about asset values. In this example, the dealer's instruments are to change prices to induce incoming trades (i.e., incoming order flow), or to call others and use outgoing trades (i.e., outgoing order flow). Canonical single-dealer models fail to consider how this affects price formation.

Canonical modeling of dealer price-setting is grounded in the two general microstructure-pricing effects. The first is the inventory effect, in which the dealer must manage a finite stock of the asset against a demand that responds to a random-walk fundamental value.[3] In this situation, if the dealer passively fills orders, the probability of a

---

[2] Examples in equity markets include Froot, O'Connell and Seasholes (2001), and Froot and Ramadorai (2001a). Examples in foreign exchange markets include Evans and Lyons (2002), Froot and Ramadorai (2001b), and Rime (2001). Examples in bond markets include Massa and Simonov (2001).

[3] For example, Stoll (1978), Amihud and Mendelson (1980), Ho and Stoll (1981, 1983), O'Hara and Oldfield (1986) among others.

stock out is unity. Hence, inventory models argue that dealers changes prices away from the expected asset value to induce trades that unwind undesired positions. The second effect is the asymmetric information effect, where, for example, the dealer faces a market where some insiders have information about the asset's liquidation value.[4] Recognizing that incoming order flow partially reflects this information, the dealer changes her price accordingly.

When multiple increasing-marginal-cost instruments are available for managing inventory, as in the example, the dealer optimally spreads her inventory management across all of them. Furthermore, communication with other dealers through outgoing calls is as informative as communication through incoming trades. The dealer may use this information to update her prior beliefs about asset values and adjust inventory levels. Hence, part of observed inventory and price changes may be correlated with innovations in information, but be unrelated to either inventory carrying costs or incoming order flow. This paper models this phenomenon in the context of foreign exchange markets. In the model, the ability to make outgoing trades alters both inventory driven price changes, and learning about asset values. Ignoring outgoing orders leads to both neglecting the role of information learned from these orders and overemphasizing price changes in inventory management. Modeling price setting without considering these effects explicitly leads to misspecified tests of information and inventory effects.

While empirical evidence of asymmetric information based on canonical dealer pricing models abounds,[5] tests for inventory effects have failed. For example, Madhavan and Smidt (1991) and Hasbrouck and Sofianos (1993) reject expected inventory effects in equity and futures markets, respectively. Madhavan and Smidt (1993) only find evidence of unexpectedly long-lived effects by modeling inventory mean reversion with shifts in the desired inventory level. Manaster and Mann (1996) actually find robust effects opposite to theoretical predictions. Lyons (1995) extends microstructure models to foreign exchange markets and does find inventory effects; however, Romeu (2005) overturns the Lyons (1995) result supporting canonical models' inventory specifications – specifically, inventory and information effects are not simultaneously present in subsamples. Other studies of foreign exchange markets also fail to find inventory effects, and hence, the evidence supporting these is at best a mixed bag.[6]

[4] For example, Kyle (1985), Glosten and Milgrom (1985), Admati and Pfleiderer (1988), Easley and O'Hara (1987, 1992), among others.

[5] For example, Hasbrouck (1991 a, b), Hasbrouck (1988), Madhavan and Smidt (1991, 1993), Lyons (1995), Evans and Lyons (2002), Yao (1998), Bjonnes and Rime (2000), Ausubel and Romeu (2005), among others.

[6] In foreign exchange markets Yao (1998) and Bjonnes and Rime (2000) find no evidence of inventory effects. The former suggests that it is due to dealers' aversion to revealing their position (or private information) through inventory-induced bid shading, whereas the latter suggest that the introduction of electronic brokering is the cause. The model here suggests that misspecification is the cause. More generally, see O'Hara (1995) on the empirical difficulties of predicted inventory effects.

The model presented here nests canonical dealer pricing models, and demonstrates why they fail empirically. Previous models are misspecified insofar as they neglect both alternatives to controlling inventory through price-induced flows, and alternative sources of market information. The model presented uses decentralized markets with multiple dealers to underscore the impact of these alternatives on price setting. At its heart is the idea that dealers exploit every alternative when rebalancing portfolios, rather than relying solely on price-induced order flow to change their portfolio composition. As dealers face increasing marginal losses for inducing flows through price shading, they turn to other methods of unloading unwanted positions. Competitive dealer markets offer a clear opportunity to observe this phenomenon.

Previous work on price formation in decentralized markets, both at the dealer and at the market-level, support the model presented here. For example, in discussing inventory control, O'Hara (1995) singles out foreign exchange dealers' ability to lay off orders on one another. At the dealer level, the Ho and Stoll (1983) framework permits interdealer trading (although it does not arise in the model solution) which is the basis of the approach presented here. Moreover, Romeu (2005), Lyons (1995) and Mello (1996) all speculate that non-linearities in dealer pricing models related to inter-transaction time or multiple inventory control instruments may be present in canonical estimations of dealer behavior – both of which are central to the model presented here. At the general-equilibrium level, the "hot potato" model of Lyons (1997) favors dealer pricing with multiple instruments. In that framework, high trading volume in the FX market results from dealers passing on inventory imbalances.

Market makers in all types of markets have an incentive to minimize guaranteed losses from inducing trades via price changes, not just in FX. While laying off inventory on others is an alternative in multiple dealer settings such as FX or bond markets, there is evidence that similar phenomenon exist in more centralized markets as well. For example, Madhavan and Sofianos (1997) find that New York Stock Exchange (NYSE) specialists engage in selectively trading to balance inventory. Hence, previous equity market studies possibly overemphasize the role of prices in inventory management and miss other inventory effects. In addition, if previous models account perfectly for inventory costs, they still overlook price changes resulting from new information that alternative instruments yield. Accounting for both these effects presents more complex behavior, where the market maker is using multiple instruments to both manage inventory and update priors.

Empirical tests presented here support the model and offer several novel results. For example, asymmetric information effects driving price changes are likely twice as large as previously estimated – not only is the price response to order flow effect larger, but there are more instruments. One can graphically compare prices with the new information signals that the dealer sees. Inventory pressure on prices is lower, perhaps as low as one-fourth previous estimates. This makes sense since multiple instruments will keep inventory management costs at the lower end of an increasing marginal cost curve. After controlling for inventory and information effects, the base bid-ask spread is wider than previously estimated, and statistically indistinguishable from the market spread convention (3 pips).[7] When setting prices, the dealer plans to trade out about one-third of the difference between her current and the optimal inventory positions. A standard ($10 million) incoming trade moves the dealer's price less than 2 pips or $1,000, and the expected cost of executing an outgoing trade is about double that amount. Accordingly, the dealer is observed accepting incoming trades about nine times more often than outgoing trades, and five times more volume is handled through incoming trades.

A Federal Reserve intervention of $300 million in the data temporarily moves prices about 6.7 pips per $100 million.[8] This increases the asymmetric information impact of trades on price changes by fifteen percent, which suggests that order flow becomes more informative as the market learns of the intervention. That is, the estimates of how much our dealer shades her price in response to inventory imbalances is fairly robust to intervention. This, taken with the result on asymmetric information, suggests that the central bank intervention was transmitting information rather than inducing portfolio balance effects. Finally, the base spread tightens by five percent when the intervention is included in the estimation.

While both transitory and permanent effects are present in the data, the results suggest a stronger permanent impact of portfolio flows on prices. With multiple instruments, market participants share intraday inventory more efficiently. That is, dealers exhaust the gains from sharing a large inventory position more quickly and with less price impact in this model. As a result, the transitory effects of inventory imbalances are present, albeit less important in determining intraday price changes than estimated previously. Furthermore, multiple instruments facilitate a more efficient aggregation of the dispersed information embedded in order flow, which can be interpreted as favoring permanent price movements.

The paper is organized as follows. Section II describes the theoretical framework and the model solution, which is detailed in the Appendix I. Section III shows empirical estimates, tests of the model, and discusses intervention effects. Section IV concludes. Estimation details are in Appendix II.

II. Intraday Price Discovery In Markets with Multiple Dealers

This section generalizes the Madhavan and Smidt (1993) framework in which an uninformed market maker with inventory carrying costs sets prices in a market with informed agents. Optimally, the market maker extracts information from arriving order flow, and sets prices to induce inventory-balancing trades. The Madhavan and Smidt (1993) framework is

---

[7] A pip is the smallest price increment in a currency. The value depends on the currency pair. The data used here are dollar/deutsche mark, so a pip is DM 0.0001.

[8] This amount observed concords with studies of intervention, e.g. Evans & Lyons (1999) estimate 5 pips and Dominguez and Frankel (1993) estimate 8 pips per $100 million.

representative of the canonical microstructure hypothesis of price formation. In actuality, however, this abstraction may miss important alternatives available to dealers in competitive dealer markets, such as bond and FX markets. For example, an FX dealer only sets prices when she passively receives an order (i.e., another dealer initiates the trade).[9] This price-setting is the focus of this study. Besides setting prices, however, she can initiate interdealer bilateral dealer trades, initiate brokered dealer trades, or initiate IMM Futures trades, as well as receive information from these, or her sales and floor managers or fellow traders, among other sources. At no time does she set interdealer prices under any of these alternatives; however, they may indirectly affect her price setting. It is intractable to model all these alternatives explicitly.[10] Furthermore, the data available (inventory levels, incoming orders, and their corresponding prices) would limit empirical tests of any such model. These limitations withstanding, the dealer modeled here has two instruments for balancing inventory: inducing order flow through price changes, and initiating outgoing trades with others at their prices. She also has two instruments for updating priors: information reflected in incoming quantities, and information reflected in unplanned (at the time of price-setting) outgoing quantities. The optimal price updates priors from both information sources and spreads inventory costs across both instruments, hence the misspecification in canonical models.

The following sections formalize this modeling approach. Subsection A describes the model setting: the market, inventory, capital, and information variables. Subsection B shows the optimal updating using multiple informative signals. Subsection C shows the optimal inventory management, and the model solution. Subsection D shows the model nesting previous work, and their misspecifications. Proofs are in the appendix.

A. The Market

Consider an economy where a dealer holds a portfolio of three assets. She only makes markets in the first, a risky asset with a full information value denoted by v_t, which evolves as a random walk. Write this value as:

$$v_t = v_{t-1} + \theta_t, \qquad \theta \sim N(0, \sigma_v^2) \quad (1)$$

[9] An extensive description of the Foreign Exchange (FX) market's institutional make-up can be found in Lyons (2001). FX is traded bilaterally, over–the–counter, and privately, via computer emailing systems called Reuters Dealing. There are also electronic brokers similar to bulletin boards, provided by Reuters or EBS. Most large trades are done via the Reuters Dealing system, and the spread is fixed by convention.

[10] That is, the return in economic insight to modeling competitive dealers is likely to be small relative to the cost of overcoming the intractability, particularly in terms of the necessary assumptions. See O'Hara (1995) on precisely this intractability.

The second is an exogenously endowed risky asset that is correlated with the first, and generates income y_t. The third is capital, the risk-free zero-return numeraire, denoted by K_t. The distribution of the two risky assets is:[11]

$$\begin{pmatrix} v_t \\ y_t \end{pmatrix} \equiv N\left(\begin{pmatrix} v_{t-1} \\ 0 \end{pmatrix}, \begin{bmatrix} \sigma_v^2 & \sigma_{vy} \\ \sigma_{vy} & \sigma_y^2 \end{bmatrix} \right). \qquad (2)$$

The dealer's total wealth is:

$$W_t = v_t I_t + K_t + y_t, \qquad (3)$$

With I_t being the dealer's inventory or risky asset position.

The market is open for $t = 1, 2, ..., T$ periods. The terminal date T is unknown, however, at the beginning every period $t = T$ with probability $(1-\rho)$. Hence, every period the probability that the market closes is $(1-\rho)$, at which time the dealer liquidates her position and pays a inventory carrying cost.[12] With probability ρ, $t \neq T$, so the dealer engages in trading activities, pays the inventory carrying cost, and goes on to the next period.

Figure 1 (page 24) depicts the timing of the model. The total change in the dealer's inventory from one event to the next occurs in two stages. In the first stage, the dealer faces an incoming order (denoted by q_{jt}) and knows her inventory (denoted by I_t). Part of q_{jt} comes from informed dealers who know the full information value (v_t). The informed part of q_{jt}, denoted by Q_t, is driven by differences between the dealer's price, denoted p_t, and the asset value v_t:

$$Q_t = \delta(v_t - p_t), \quad \delta > 0. \qquad (4)$$

The rest of the incoming order is an uninformed or liquidity component, denoted by X_t:

$$X_t \equiv N\left(0, \sigma_X^2\right). \qquad (5)$$

One can think of the uninformed as quantities demanded by parties not monitoring the markets or constrained to trade independent of price, for reasons not modeled here. The dealer only observes the aggregate order, (q_{jt}), and sets the price. Hence, the incoming order flow is:

$$q_{jt} = Q_t + X_t = \delta(v_t - p_t) + X_t. \qquad (6)$$

When our dealer sets her price at (incoming) trade t, she knows she can also call others and initiate outgoing trades (denoted q_t^{out}). These outgoing trades are depicted in the upper box of Figure 1. q_t^{out} indicates our dealer's desired outgoing quantity in expectation, and conditional on information available at the time of price setting. Because the dealer has

[11] Note that this is a one-period-ahead conditional distribution, as the unconditional distribution would have a time-varying variance.

[12] The inventory carrying cost, shown below, follows Madhavan and Smidt (1993). It is a cost proportional to the variance of the dealer's wealth.

this tool of outgoing trades (q_t^{out}) available, she does not control inventory solely through price induced order flow. In this sense q_t^{out} captures the planned amount the dealer prefers to lay off by initiating trades rather than by shading price to induce incoming trades.

The role of outgoing trades (q_t^{out}) in price formation is a departure from canonical dealer models. In considering multiple dealer markets, it is an empirical reality one typically has data on trade prices for only a subset of all dealer trades (this is particularly because they are relatively unregulated with much lower reporting requirements). We want to model the subset of available trades to the fullest extent possible, while at the same time recognizing the role of trades not in that subset. In this case, prices, inventories, and quantities traded are available only for incoming trades. Inventory, however, summarizes all quantities: incoming and outgoing. That is, we at least have quantity information for integrating the trades without price data into the analysis of the available data. Thus, we decompose the total change in inventory from one trade to the next into three components: the observed incoming trade (q_{jt}), the expected outgoing trade (q_t^{out}), and unexpected quantity shocks to inventory, as shown in Figure 1.

Denote the unexpected quantity shocks to inventory as γ_t. While our dealer is trading q_t^{out} , these exogenous quantity shocks change her inventory beyond the outgoing trade (q_t^{out}) planned at the time of setting prices. The source of these shocks can be unplanned trading with clients of our dealer's bank (her employer), other bank dealers, brokered trading, the trading floor manager, and so on. Accordingly, the total quantity $(q_{t-1}^{out} + \gamma_{t-1})$ will be the inventory change apart from the incoming trade (q_{jt-1}) from $t-1$ to t. Hence, last event's inventory (I_{t-1}), adjusted for the last incoming trade (q_{jt-1}), as well as the total realized outgoing quantity $(q_{t-1}^{out} + \gamma_{t-1})$, yields next event's inventory (I_t).

An example using actual dealer transactions helps motivate the key assumptions regarding q_t^{out} and γ_t. Table 1 (page 26) shows the first five incoming trades received by a NY based foreign exchange dealer on a given trading day (these data are discussed in detail below). The first column indexes the trades according to their order of arrival; the second column shows the price set by the dealer at each incoming trades. The next columns show incoming order flow, followed by the inventory at the beginning of the trade. The last column shows $q_t^{out} + \gamma_t$, which are observed jointly. Consider, for example, the third incoming trade, which was a sale to the dealer of \$28.5 million. At the time of the trade, the dealer was long \$1 million, as reflected in her inventory. Canonical models of price formation assume that incoming orders are the only instrument by which a dealer can adjust inventory levels and update prior information. If one assumes that this were the case, and since the dealer buys \$28.5 million, her inventory at entry four should be \$29.5 million long (the next incoming trade). Instead, the dealer is short \$1.5 million at entry four, which implies that her inventory declined by \$30.5 million between the third and the fourth trade. This decline is reflected in the last column, $q_t^{out} + \gamma_t$. It captures the inventory evolution that incoming order flow did not generate. This column is expressed as the sum of two components because q_t^{out} reflects the optimal amount that the dealer should trade given the information available at the time of the incoming trade. It is a first order condition. Any deviation from q_t^{out} must be a result of new information, and is reflected in γ_t. Therefore, the part of inventory changes not generated by incoming trades is the sum of planned and unplanned outgoing trades, $q_t^{out} + \gamma_t$.

Hence, new information and events occurring in the clock time between events *t-1* and *t* are assumed to be driving the quantity shocks (γ_{t-1}); The shock γ_{t-1} is informative because after the dealer chooses her outgoing quantity (q_{t-1}^{out}), she should trade this quantity and nothing else unless new information motivates a revision in the outgoing trade. That is, the choice made at *t-1* is optimal until new information (at the next incoming order, q_{jt}) arrives.[13] Hence, the only reason our dealer would deviate from the optimal outgoing quantity (q_{t-1}^{out}) between *t-1* and *t* is that new information is revealed. For this reason, the evolution of v_t can be inferred from γ_{t-1}, and the total outgoing quantity will reflect the desired quantity (q_{t-1}^{out}) plus the quantity driven by new information (γ_{t-1}). γ_{t-1} captures that information in the dealer's decision process beyond strictly what is derived from incoming order flow, while keeping the analysis tractable.[14]

In summary, the identity that describes the evolution of inventory is:

$$I_t \equiv I_{t-1} - \delta\left(v_{t-1} - p_{t-1}\right) - X_{t-1} + q_{t-1}^{out} + \gamma_{t-1} \qquad (7)$$

In contrast, at the time of setting prices, the dealer's expectation of next period's inventory is:

$$E[I_{t+1} \mid \Omega_t^j] = I_t - q_{jt} + q_t^{out}. \qquad (8)$$

Our dealer manages inventory because she pays a cost every period that is proportional to the variance of her portfolio wealth, which includes the cash value of the inventory. One can motivated this cost, for example, by risk aversion or marginally increasing borrowing costs. Assume that the dealer incurs a capital charge due to the γ shocks. That is, any gains (losses) entering into the dealer's wealth due to γ are subtracted

---

[13] An alternative to this interpretation of γ_{t-1} is that it could be an (uninformative) systematic factor missing in the analysis of available price changes. This would not introduce interesting alternative economics because dealers can anticipate this fully – there is no news in it.

[14] Although they include multiple informative signals, incoming order flow is the only source of private information in Madhavan and Smidt (1991) or Lyons (1995).

(added) from (to) the dealer's capital, K_t at a cost v_t.[15] Incorporating this charge, at trade t the dealer's wealth position is given by:

$$W_t = v_t\left(E[I_t \mid \Phi_{t-1}] + \gamma_{t-1}\right) + \left(E[K_t \mid \Phi_{t-1}] - v_t \gamma_{t-1}\right) + y_t\,. \tag{9}$$

This assumption implies that the dealer only pays the inventory carrying cost on the expected wealth, and the inventory carrying cost due to quantity shocks is canceled by the capital charge. The appendix shows that the inventory cost is a function of the deviations from the optimal hedge ratio of the risky assets, given by I^d. This hedge ratio optimally smoothes the dealer's wealth, and enters the inventory cost as:

$$c_t = \omega\left[\sigma_W^2\right] = \omega\left[\phi_0 + \phi_1\left(I_t - I^d\right)^2\right]. \tag{10}$$

B. The Information Structure

What is of interest is how the dealer sets prices, which occurs only in the event of an incoming trade. The incoming trade is, in part, based on the equilibrium asset value, v_t. The dealer wishes to learn this value, and she will estimate the full information value of the asset based on her trading history and any publicly available information. The appendix shows the solution to the dealer's learning problem modeled as a rational expectations consistent Kalman filter.[16] This section outlines the two sources of information available for learning v_t and updating prior beliefs in this model. Denote the dealer's expectation of the full information value of the risky asset as:

$$E\left[v_t \mid \Phi_t\right] = \mu_t\,. \tag{11}$$

The dealer has two ways of updating μ_{t-1} and learning about the full information value of the asset v_t. The first is the incoming trade, q_{jt}. From this incoming quantity the dealer extracts a signal of the asset value, v_t. Denote this signal by s_t. The second source of

---

[15] This assumption simply eases the exposition of the problem at hand, and keeps it in a discrete time framework. As discussed below, γ has a time-varying variance. This complicates calculating the variance of the portfolio – this would involve moving the entire model to a continuous time framework. Because of the discrete-time arrival process of incoming calls, this would make for a cumbersome solution with little added payoff in relation to the problem of how dealers set prices on incoming orders. It would not, however, change the model's conclusions regarding price setting with multiple instruments.

[16] In the empirical estimation, this study uses total incoming orders (rather than the unexpected component) as signals, as in Lyons (1995), Madhavan and Smidt (1991), Yao (1998) and others. The s_t represents a function that reflects the information in incoming order flow, and $\kappa(\gamma_{t-1})$ represents a function reflecting the information in inventory shocks, consistent with the approach Hasbrouck (1991a), Madhavan and Smidt (1993), and others. Generally, estimations are robust to either approach, as is the case here.

information about v_t is the information learned while executing the outgoing trade, which is reflected in a function of the inventory shock, $\kappa(\gamma_{t-1})$. While both $\kappa(\gamma_{t-1})$ and s_t are used to update μ_{t-1}, assumed that the variance of $\kappa(\gamma_{t-1})$ is increasing in the real time (i.e., clock time) elapsed between incoming trades. That is, assume that $\text{var}(s_t) = \sigma_w^2$ and $\text{var}(\kappa(\gamma_{t-1})) = \sigma_w^2 \Delta\tau$, with $\Delta\tau$ being the clock time elapsed between events *t-1* and *t*. As the appendix shows, this gives an updating as a function of:

$$\mu_t - \mu_{t-1} = \left(\Delta\tau / 1+\Delta\tau\right) s_t + \left(1 / 1+\Delta\tau\right) \kappa(\gamma_{t-1}). \tag{12}$$

In equation (12), as elapsed inter-transaction time gets larger ($\Delta\tau \to \infty$) the dealer places the majority of the weight on the incoming order's information, s_t. The longer the time in between trades, the less relevant is the information from that time in relation to the incoming trade's information. Intuitively, (12) says that the moment the dealer is setting p_t, s_t has just arrived because it is based on the incoming order itself (q_{jt}). The quantity shock signal ($\kappa(\gamma_{t-1})$) also serves to signal the new innovation, but it arrives between *t-1* and *t*, and hence it is not assumed to have the same precision as s_t. Instead it is assumed that $\kappa(\gamma_{t-1})$'s precision decreases (i.e., variance increases) as the clock-time elapsed from event $t-1$ to t increases. As more time has passed in between trades, $\kappa(\gamma_{t-1})$ has more noise.[17]

Finally, the appendix shows that the estimate of the full-information asset value, μ_t, generates an unbiased estimate of the liquidity trade, X_t. We denote this statistic as $E[X_t \mid \Omega_t] = x_t$.

C. The Dealer's Optimization

Here the problem is set up as a stochastic dynamic programming problem; ~ denote random variables, and the solution is given in the appendix. The dealer solves:

$$J(I_t, x_t, \mu_t, K_t) = \max_{p_t, q_t^{out}} E\left\{(1-\rho)\left[\tilde{v}_t I_t + K_t + y_t - c_t\right] + \rho J(\tilde{I}_{t+1}, \tilde{x}_{t+1}, \tilde{\mu}_{t+1}, \tilde{K}_{t+1})\right\}, \tag{13}$$

subject to the following evolution constraints:

Inventory: $$E\left[\tilde{I}_{t+1} \mid \Phi_t^i\right] = I_t - \delta\left(\mu_t - p_t\right) - x_t + q_t^{out}, \tag{14}$$

Noise Trading: $$E\left[\tilde{x}_{t+1} \mid \Phi_t^i\right] = 0, \tag{15}$$

---

[17] One might argue that as $\Delta\tau \to 0$, the dealer has less time to carry out planned transactions, but she can always elect to not answer the incoming calls until the part of planned transactions she wants done are satisfied. Furthermore, the increasing frequency of incoming calls and shortening of inter-transaction time would itself be a source of new information for the dealer, as suggested by Easley and O'Hara (1992). Indeed, Lyons (1995) finds evidence supporting that longer inter-transaction clock times increases the informativeness of incoming order flow, as interpreted in this study.

Information: $$E\left[\tilde{\mu}_{t+1} \mid \Phi_t^i\right] = \mu_t, \tag{16}$$

Capital: $$E\left[\tilde{K}_{t+1} \mid \Phi_t^i\right] = K_t + p_t\delta\left(\mu_t - p_t\right) + p_t x_t - (\mu_t + \alpha q_t^{out})q_t^{out} - c_t, \tag{17}$$

Equations (10), and (13) through (17) comprise the optimization problem. (14) constrains inventory evolution. (15) constrains liquidity trades to be zero in expectation. (16) constrains the asset to a random walk. (17) constrains the capital evolution, and specifies that when the dealer trades q_t^{out}, she expects to pay a price centered on the full-information value, and with a price impact $(\mu_t + \alpha q_t^{out})$. α captures the price impact of a marginal increase in her outgoing quantity. Hence the dealer, while not a monopolist in the interdealer market, does face a downward sloping demand curve in her trades. Assuming that the dealer faces α when trading out is similar to assuming that there is marginal declining revenue from selling to an informed agent (recall that revenue from the sale is $p\delta(\mu - p)$). Modeling outside prices explicitly requires a general equilibrium framework that normally mutes dealer level pricing effects.[18] The appendix shows the model solution to be:

$$p = \mu + \beta\left(\alpha/_{(1+\delta\alpha)}\right)\left(I - I^d\right) + \left(\tfrac{1+\delta\alpha(1-\beta)}{2\delta(1+\delta\alpha)}\right)x; \tag{18}$$

$$q^{out} = \left(A_1/_{A_1-\alpha}\right)\left[-(I - I^d) + \delta(\mu - p) + x\right]; \tag{19}$$

$$I' = I + \beta(I - I^d) - \tfrac{(1+\beta)}{2}x, \tag{20}$$

$$\Delta p_t = \psi\eta_t q_{jt} + \beta\left(\alpha/_{(1+\delta\alpha)}\right)(q_{t-1}^{out} + \gamma_{t-1} + q_{t-1}) + \psi(1-\eta_t)\gamma_{t-1} + \left(\tfrac{1+\delta\alpha(1-\beta)}{2\delta(1+\delta\alpha)}\right)\Delta x_t \tag{21}$$

Equation (18) shows the price of the dealer as a function of the estimated asset value, (μ_t), the deviation from optimal inventory, ($I_t - I^d$), and the liquidity shocks (x_t). In (19) the outgoing quantity shows that as the price impact of outgoing trades goes to zero, i.e., $\alpha \to 0$, outgoing trades fully adjusts inventories to the optimal level (in the appendix, A_1<0 is shown). In this case, the price will depend only on the estimate of *v* and the liquidity demand. In equation (21), s_t is the information from incoming order flow (q_{jt}) and the elapsed time is measured by $\eta = \Delta\tau/_{1+\Delta\tau}$. This equation shows that the increment in dealer price contains information-driven components from both the current incoming order (ηs_t), and the previous inventory shock ($(1-\eta_t)\gamma_{t-1}$), both weighted by the Bayesian updating term, ψ. The ($q_{t-1}^{out} + \gamma_{t-1} + q_{t-1}$) term captures component of the price change attributable to inventory pressure – it is the change in the inventory. Finally, the dealer changes her price due to the noise-trading component (Δx_t).

---

[18] For example, the Evans and Lyons (2002) assumes that dealers submit bids simultaneously and transparently, which in equilibrium implies that prices be based on common information only. This paper avoids such restrictions because the focus is on interdealer price dynamics, but this comes at the expense of the market-wide price determination of such models.

Intuitively, the dealer would like to maintain inventory at the optimal level, but as a market maker she must accept incoming orders that constantly disturb her inventory position. As incoming orders arrive, she tries to restore balance to her inventory with q_{t-1}^{out} and price changes. Adjusting back to the optimal level I^d via q_{t-1}^{out} implies absorbing the costs from the outgoing order's price impact (α). Adjusting inventories via price induced orders implies absorbing the certain loss to the informed dealers, via $\delta\left(\mu_t - p_t\right)$. The coefficients in (21) reflect the balance between these competing losses. Furthermore, the price is centered on the best guess of v_t, which is derived from two information sources, s_t and $\kappa(\gamma_{t-1})$. The respective coefficients reflect the information extraction, which involves weighing these signals by the time elapsed between events.

D. A Comparison with Existing Models

This section shows how the model presented nests the previous dealer-level frameworks. Restricting the model to no outgoing trades, and consequently no inventory shocks, the solution would be (22). This is the Madhavan and Smidt (1993) pricing behavior for an equity market specialist;

$$\Delta p_t = s_t + \zeta_1(I_t - I^d) + \zeta_2 x_t \quad \Leftrightarrow \quad \gamma_t \equiv q_t^{out} \equiv 0 \quad \forall t \leq T . \tag{22}$$

This model suggests, however, that these restrictions may shut down other avenues of inventory management available to specialists. That is, as NYSE specialists face increasing marginal costs to inventory management through price changes, they optimally spread these costs across different avenues available. For example, Madhavan and Sofianos (1997) find evidence supporting this. Hence, restrictions that yield (22) would lead to biased estimates of inventory effects since they overemphasize the role of changing prices to manage inventory. Romeu (2005), Bjonnes and Rime (2000), Yao (1998), Lyons (1995) and Madhavan and Smidt (1991) postulate that prices are set according to:

$$p_t = \mu_t - \alpha(I_t - I^d) + \gamma D_t \tag{23}$$

Equation (23) yields the price change as:

$$\Delta p_t = \beta_0 + \beta_1 q_{jt} + \beta_2(I_{t-1} - q_{j,t-1} + q_{t-1}^{out} + \gamma_{t-1}) + \beta_3 I_{t-1} + \beta_4 D_t + \beta_5 D_{t-1} \tag{24}$$

With the data used here, Romeu (2005) shows that estimates of (24) are misspecified. Breaks present in the data coincide with systematic differences in the length of inter-transaction time ($\Delta\tau$). Previous studies using canonical dealer pricing models have indeed noted that inter-transaction times imply changes in the precision of incoming order flow, however, there are, in fact, changes in both informative variables (q_{jt}, γ_t). The model presented here shows why inter-transaction times would cause breaks. Rewriting (24) consistent with this paper's data generation process, note the omitted term in brackets weighed by $(1-\eta_t)$ below:

$$\Delta p_t = \varphi_0 + \varphi_1 q_{jt} + \varphi_2(-q_{jt-1} + q_{t-1}^{out} + \gamma_{t-1}) + \underbrace{(\varphi_3 - \varphi_2) I_{t-1}}_{\text{extraneous term}} + \varphi_5 \Delta x_t + \underbrace{(1-\eta_t)[\varphi_4 \kappa(\gamma_{t-1}) - \varphi_1 q_{jt}]}_{\text{omitted term}}$$

The data generating process under the hypothesis of multiple instruments places zero weight on lagged inventory (the extraneous term), which would tend to bias $(\varphi_3 - \varphi_2)$ toward zero. However, the estimated coefficient φ_2 captures not only the inventory effect, but it partially reflects information from γ_{t-1} which is contained in the inventory term. Thus, the omitted term would normally transmit information from γ_{t-1} to prices, but its absence drives the inventory term to partially reflect this information. Hence, the variation in the informativeness of γ_{t-1} will affect the inventory term. When inter-transaction times are long ($\Delta\tau \to \infty$ and $\left(\frac{\Delta\tau}{1+\Delta\tau}\right) \equiv \eta \to 1$), the omitted term should be irrelevant. At such times, one should expect the incoming order flow coefficient (φ_1) to be significant, and $\text{var}(\kappa(\gamma_{t-1})) \to \infty$, hence γ_{t-1} will be mostly noise, and uncorrelated to price changes. This would in turn make φ_2 less correlated with the information effect in Δp, since the inventory term picks up the information in γ_{t-1} in lieu of the omitted term. Hence, one would expect to see the inventory effect dampened at these times. When inter-transaction times are short ($\Delta\tau \to 0$, and $\eta \to 0$), one would see the order flow coefficient (φ_1) become less significant, whereas the coefficients on the inventory terms would be more significant, and pick up the inventory effect more clearly. Hence, canonical models fail to find inventory effects because they are confounded with information effects, or they include extraneous variables that are assigned the inventory role.

III. Data Considerations

This section discusses the data sources employed in testing the model, and then presents the data graphically to motivate both the new inventory and the asymmetric information effects predicted here, as well as those predicted by canonical models.

The data set consists of one week of a New York based foreign exchange dealer's prices, incoming order flow, inventory levels, and transaction clock times. Hence, p_t, q_{jt}, I_t, and $\Delta\tau$ (and η) come directly from the recordings of a *Reuters Dealing* trading system. Out of the 843 transactions, four overnight price changes are discarded since the model at hand deals exclusively with intraday pricing. A few measurement errors are present in transaction clock times, and these are treated with a dummy variable in the estimation.[19] Table 2 (page 26) presents some descriptive statistics. One observes that the dealer keeps the average inventory at $2.1 million, however, it deviates as much as ±$50 million. Given a median

[19] The data are for the dollar/DM market from August 3–7, 1992. See Lyons (1995) for an extensive exposition of this data set. The transaction clock time measurement errors show up when the sequential order of the trades is not consistent with the clock-times, e.g. trade 2 cannot have occurred earlier than trade 1.

incoming order of roughly $3 million, reversing a one standard deviation swing in inventory necessitates about five sequential incoming trades. This suggests alternative measures of inventory management other than inducing incoming trades are at work, which is also suggested by other FX studies.[20]

Table 3 (page 26) shows the observed incoming trades received by the dealer, as well as bilateral trades that our dealer initiates with other dealers in the FX market. The table shows on average 20 outgoing trades per day initiated by our dealer. These, however, are conceptually different from q^{out}, which represents an outgoing quantity planned at the time of price-setting that captures alternatives to shading the incoming transaction price for inventory control. Thus q_t^{out} is unobservable in that it represents the dealer's commitment to an outgoing trade *at the moment of price setting only*. At this moment she commits irreversibly to trading at a price who's optimality depends on being able to trade q_t^{out}; one of the messages of this model is that the price set by the dealer would be different if q_t^{out} were not available for inventory control. Observed outgoing quantities differ from the planned q_t^{out} because the dealer reoptimizes in response to unanticipated information, frictions, or differences in the trading venues utilized to execute the outgoing trade. For example, at each incoming trade, because a price is set, there necessarily exists an expected outgoing trade. However, the dealer may not execute an outgoing trade before then next incoming trade is observed in the sample. Although they are unobservable, the model solution provides equations which allow estimation of q_t^{out} and γ_t. Table three shows that the spread on both incoming and outgoing trades is tightly maintained at the market's convention of 3 pips. Diverging from this spread is frowned upon by others in the market, as it is interpreted as failing to provide predictable over the counter liquidity. Hence, point estimates of the model that imply widening or narrowing the spread should be interpreted as theoretical constructs that in practice manifest themselves in other ways, e.g. as shifts in the midpoint of the spread.

The fundamental question of interest is how dealers set prices, i.e. equation (21). Its estimation requires decomposing the inventory change so as to get at the outgoing orders, q_t^{out} and inventory shocks. Because γ_t is driven by new information, the model solution reflects this information in our dealer's estimate of the liquidation value of the asset. That is, price changes depend on updating priors using two sources of information: the incoming order flow, and the unexpected outgoing order flow (γ_{t-1}). Canonical models typically employ incoming order flow as a source of information; however, the use of γ_{t-1} as a source of information is new. To get a feel for this variable, Figure 2 (page 25) superimposes

[20] For example, Lyons (1995) finds evidence that observed outgoing bilateral interdealer trades and brokered dealer trading are used to control inventory in the context of a canonical dealer pricing model. These do not include a small amount of brokered trading (which occurs at 5 percent of the sample) which the dealer also engages in.

cumulative daily unexpected order flow on the price, and Figure 3 does the same for cumulative daily inventory shocks (i.e., cumulative daily γ_{t-1}).

The vertical lines represent the end of each day of the five-day sample (Monday through Friday). The correlation of two signals with price seems to vary. For example, on Monday and Wednesday, incoming order flow appears to be a more precise signal of price than inventory shocks, whereas on Friday the opposite seems to be true. In the model, elapsed clock-time affects the relative precision between these signals. Table 4 (page 26) reports the daily correlations and average inter-transaction clock-time. Although these are cumulative signals, Friday gives an example of short inter-transaction clock-time, and higher correlation in the (cumulative) inventory shocks than (cumulative) order flow shocks. Hence, these signals seem to compliment each other and are weighted by inter-transaction time in the model.

IV. ESTIMATION

The framework presented provides sufficient identifying relationships so as to permit an almost direct system estimation of the model solution. Only leveling constants, an autoregressive error on the inventory equation, and bid-ask bounce dummies on the pricing equation are added. Table 5 (page 27) lays out the system of equations given in the model solution (the first column), with the empirical implementation of the solution (the second column), and the parameters recovered from each equation (third column). The first equation in the system, the inventory evolution, yields the optimal inventory level. The second equation identifies the optimal outgoing order q_t^{out} and γ_t. This is simplified as:

$$\hat{q}_{t-1}^{out} = c_3\left(-I_{t-1} + \hat{I}^d + q_{jt-1}\right),\ \ with\ \ \hat{I}^d = \frac{c_1}{(1-c_2)}\ \ and\ \ (q_{t-1}^{out} + \gamma_{t-1}) \equiv (\Delta I_t + q_{jt-1}) \quad (25)$$

Solving for γ_{t-1} by adding and subtracting c_3I_t, yields:

$$(\Delta I_t + q_{jt-1}) - c_3\left(-I_{t-1} + \hat{I}^d + q_{jt-1}\right) = (1-c_3)(\Delta I_t + q_{jt-1}) + c_3\left(I_t - \hat{I}^d\right) \quad (26)$$

Hence, the transformation of (26) allows the estimation of the proportion of incoming trade that is expected to be traded out, c_3, as a moving average of the net outgoing order flow $\left(\Delta I_t + q_{jt-1}\right)$, and the deviation from target inventory $\left(I_t - \hat{I}^d\right)$. Moreover, in the pricing equation (the third row of Table 5), removing expected outgoing trade, as well as the incoming trade, from the inventory change identifies the outgoing trade shock $\hat{\gamma}_{t-1}$. However, since (26) is a function of terms such as ΔI_t that are already present in the pricing equation, it is necessary to transform it so as to eliminate multicollinearity. Thus, (26) is simplified for the pricing equation to:

$$(1-c_3)(\Delta I_t + q_{jt-1}) + c_3\left(I_t - \hat{I}^d\right) = (1-c_3)\Delta I_t + q_{jt-1} + c_3\overbrace{\left(I_t - \hat{I}^d - q_{jt-1}\right)}^{(-\hat{q}_{t-1}^{out})} \quad (27)$$

one can express (27) in a more conceptual way using $\hat{q}_{t-1}^{out}$:

$$(1-c_3)\Delta I_t + q_{jt-1} + c_3\left(I_t - \hat{I}^d - q_{jt-1}\right) = (1-c_3)\Delta I_t + (q_{jt-1} - \hat{q}_{t-1}^{out}) \qquad (28)$$

Equation (28) identifies $\hat{\gamma}_{t-1}$ as a weighted function of the inventory change which the dealer did not trade, less the part of the last incoming order that the dealer did not trade out. Grouping the terms on ΔI_t in (28) with the inventory effect permits estimation of the system without multicollinearity in the pricing equation.

In estimating the incoming order flow's information content canonical models use either order flow or its unexpected component. This study uses order flow directly in the price equation, so as to maintain comparability to FX market studies, such as Lyons (1995), however, estimation is robust to either measure.[21] In addition, the model predicts that the only difference in the informativeness of incoming and outgoing order flow is due to the clock time between trades, η. Thus, the solution allows the identification of the information effect from the different components of (28) since the inter-transaction times are observed. Hence, since the model solution predicts identical coefficients on these terms, the components of γ_{t-1} outlined above are accordingly constrained to have the same coefficient as incoming order flow after accounting for η.[22] Two direction-of-trade dummy variables are included to capture the fixed costs such as order processing costs, and pick up the base spread for quantities close to zero. These variables equal unity if the incoming order is a purchase (i.e., the caller buys), and negative one if the incoming order is a sale (i.e., the caller sells). The elapsed time in between transactions is measured to the minute, and estimates are robust to monotonic transformations of η.[23] Finally, scaling constants are included in all three equations, and the first equation is estimated with an AR(1) error to control for autocorrelation. The system is estimated simultaneously using Seemingly Unrelated non-linear least squares. Table 6 (page 28) shows the estimations of the model. Below, Table 7 presents canonical model estimates of the same data as a basis for comparison.

---

[21] For example, Hasbrouck (1991) and Madhavan and Smidt (1993) use the unexpected component of incoming order flow, and estimate this measure as a residual of a vector autoregression. In the case of the FX data used here, these autoregressions tend to have little explanatory power, making the residual almost identical to the incoming order flow.

[22] Estimating the model with independent information coefficients on incoming order flow and gamma is possible, and support the restriction imposed here. However, under such estimations some inventory terms cannot be grouped as presented here, and collinearity prevents satisfactory estimations of the inventory effect, hence these estimable forms are not used.

[23] Some measurement error in the time stamps leads to the inclusion of a dummy interacted with the absolute value of the clock time (which turns out to be insignificant).

The estimations in Table 6 indicate that the model fits the data fairly well. The main results are the very significant and properly signed coefficients on the information and inventory effects, c_{11} and c_{12}, as well as the predicted inventory evolution and outgoing trade estimates, c_1, c_2, and c_3. Canonical model estimates are presented in Table 7 as a basis for comparison. Note that the canonical estimates are not robust to subsample estimation. Specifically, canonical model predictions of inventory effects are rejected in the first half of the sample, and similarly, predicted information effects are rejected in the second half of the sample. [24] The model presented here is robust to subsample estimation, notwithstanding the lower p-values of estimated coefficients in the first sub-sample. Moreover, all three equations in the system are jointly significant as predicted, and the estimates fail to reject any of the testable restrictions. Hence, this model rejects canonical model point estimates of asymmetric information and inventory effects. The model predicts that the dealer plans to trade out roughly one-third of each incoming trade ($\hat{c}_3$=0.34) each time she quotes a price. Additionally, the model estimates the dealer's target inventory at about two million ($\hat{I}^d$ =2.09). From Table 2 the average inventory is 2.16, which is statistically indistinguishable from our dealer's observed average.[25]

Asymmetric information

The asymmetric information component (c_{11}) is significant and larger than canonical model estimates given by β_1 in Table 7 (10^5 multiply the pricing equation coefficients). One way to interpret the estimates is that the dealer widens her spread by 3.5 pips per \$10 million of incoming order flow or inventory shocks (twice c_{11}, since orders are quoted based on absolute size). These estimates indicate a more intense asymmetric information effect than previously estimated; not just because of the higher estimated effects, but because there are two sources of private information – both incoming and outgoing order flow – both pushing price changes. In terms of economic significance, the estimates suggest that the marginal \$1 million dollar order pushes the dealer's price by about 2 basis points, given the average exchange rate in the sample of roughly DM 1.5 per US dollar, or 2 percentage points per excess US\$1 billion traded. This is higher than market-wide estimates of the price impact of US\$ 1 billion of excess order flow, which fluctuate around half a percent.[26] However, these latter these estimates are not comparable because of the inherent difficulties of linearly interpolating one dealer's behavior to the market-wide equilibrium. These difficulties are particularly acute since the dealer generating these data predominantly provides interdealer liquidity, not end-user liquidity. The hot potato hypothesis of Lyons (1997) would suggest

[24] Note that Romeu (2005) documents evidence of model misspecification and structural breaks present in these estimates of the canonical dealer pricing model used here for comparison.

[25] A Wald test fails to reject equality of the mean to the target with a p-value of 0.94.

[26] See Evans and Lyons (2002) or Chaboud, et. Al. (2006).

that this dealer pushes prices in response to excess order flow more than others who have access to end-users that absorb order imbalances.[27] Put crudely, an FX position is like a hot potato. Liquidity providers such as our dealer pass it around, pushing prices until an end user is found who is willing hold the off-setting position.

One may consider why previous work underestimates the information component. Even if pure inventory pressures were perfectly explained by previous models, there is a component of inventory change driven by new information. Inventory theory cannot explain this information-driven inventory component. This component is one of multiple signals that, according to the model, vary in precision depending on elapsed clock-time. This suggests that incoming order flow can be relatively less informative at different times, and should be weighed accordingly. Previous estimations assign all information-driven price changes to the (at times, noisy) incoming order flow that mute its true informative impact.

Inventory effects

Turning to inventory effects, comparing coefficient estimates of the canonical model and the model presented here is unsatisfactory because the dealer's pricing decision is affected differently by inventory. Instead, it is more useful to compare estimates of the structural parameters that reflect the dealer's bid-shading in response to inventory pressure. Canonical models' inventory specification depend crucially on the linear price relationship $p_t = \mu_t - \alpha(I_t - I^d) + \gamma D_t$, as shown in equation (23) (Section II.D, page 14).[28] That pricing assumption yields two inventory terms:

$$\beta_2 I_t + \beta_3 I_{t-1} \equiv \beta_2 (I_{t-1} - q_{jt-1} + q_{t-1}^{out} + \gamma_{t-1}) + \beta_3 I_{t-1} \equiv \underset{<0}{\beta_2}(-q_{jt-1} + q_{t-1}^{out} + \gamma_{t-1}) + \underset{<0;\ |\beta_2|>\beta_3}{(\beta_2 + \beta_3)} I_{t-1} . \qquad (29)$$

The estimate in Table 7 (page 28) of $\hat{\beta}_3 = 0.72$ from (29) is the canonical model's (absolute) structural price adjustment per one-million dollar deviation from the desired inventory level (i.e. $\hat{\beta}_3$ is the empirical estimate of the canonical model parameter α in equation (23)). In the model presented here, the analogous relationship is given in the first order conditions specified by equation (18), where $\beta(\alpha/_{(1+\delta\alpha)})$ is our structural inventory effect on prices. A direct estimate of our model's parameter β (the inventory evolution parameter in equation (20)) is $\hat{\beta} = (\hat{c}_2 - 1)$, as shown in Table 5 (page 27). This yields $\hat{\beta} = -0.34$. Moreover, $(\alpha/_{(1+\delta\alpha)}) < 1$ for the range of $\alpha > 0$ and $\delta > 0$ consistent with our model. Hence, the total inventory effect in our model is β multiplied by a factor that approaches unity from below. That is, to arrive at the equivalent measure of the canonical inventory effect in

---

[27] Lyons (1996) describes this dealer as a "liquidity machine" in reference to the interdealer market.

[28] For example, this pricing relationship forms the basis of Madhavan and Smidt (1991) or Lyons (1995).

equation (18), one must multiply $\hat{\beta} = (\hat{c}_2 - 1)$ by a factor of at most, one. Hence, in comparing the price impact per million dollar deviation from the desired inventory level in equation (23) against (18), canonical model estimates of inventory costs are at least two to three times larger than the estimates presented here. Ignoring multiple instruments will overweigh the inventory component because price changes are empirically assigned such an important role in inventory management. The model presented here suggests that price is but one of multiple instruments used to control inventory costs. As a result, inventory accumulation is not as important in explaining price changes.

Expected cost of outgoing trades and the base spread

The use of multiple increasing–marginal–cost instruments to manage inventory requires having an expected cost of the outgoing trade at the time of price setting. This expected cost is estimated at $\hat{\alpha}$ =0.35 pips. This measure reflects the dealer's expected marginal cost of trading out an extra million dollars, i.e. the dealer's opportunity cost of changing the spread in response to a \$1 million incoming trade. In principle, the dealer's alternative is to change the price to offset the inventory carrying cost, estimated to be at most 0.34 pips per million, as discussed above. Hence, the estimates suggest that trading out excess inventory has a higher marginal cost for the dealer than accepting incoming trades, and the estimated proportion of excess inventory that is traded out, C_3, is 0.33, meaning that for each incoming dollar, the dealer expects to trade out one third. Finally, c_4 measures the effective spread for q_{jt} close to zero. It suggests that after having controlled for information and inventory effects, the baseline spread is roughly 2.5-2.8 pips (twice c_4 times 10^{-5}). Note that these estimates are approximately equal the median interdealer spread observed in the FX market of 3 pips.

Fed Intervention

The last five percent of recorded trades that occurred while the Fed intervened to support the dollar. In Figure 2 (page 25), the sharp appreciation on the last day reflects the market reaction to the intervention. It perhaps succeeded in slowing the slide of the dollar, but was unsuccessful in sustaining a reversal. The market closed down on the day, and down from its high after the start of intervention. It involved multiple dollar purchases totaling \$300 million after the close of European markets. The Fed does not reveal the exact start time and there are too few observations to meaningfully estimate the intervention in isolation.[29] Wald tests fail to reject equality between estimates of the model with and without the intervention period (i.e., 95 percent of the sample, versus 100 percent).

[29] Quoting the Wall Street Journal, August 10, 1992: "The Federal Reserve Bank of New York moved to support the U.S. currency... as the dollar traded at 1.4720." This is the most precise documentation available of the intervention start, and that price corresponds to 12:32 pm. Other times are selected because of reports of a mid-day start (hence, 12:02 pm), and at 12:26 pm the price jumps 36 pips, suggesting a possible intervention start at that point.

Table 8 (page 29) shows the impact of the intervention on the estimated parameters. The intervention increases the asymmetric information effect of incoming order flow (c_{11}) by over 8 percent, while the change in the estimated inventory effect (c_{12}), as well as in other model parameters, is negligible. The dealer price appreciation recorded during the Fed intervention period, which presumably would be induced by Fed purchases of dollars, serves as a rough check on market wide studies of market liquidity. While the exact start time is not revealed, the $300 million intervention moved the market price between 20 and 32 pips before falling back. At the lower end of the range, this concords with estimates of between 5 and 8 pips per 100 million from Evans and Lyons (5 pips per $100 million), and Dominguez and Frankel (8 pips per $100 million). At the higher end, 12 pips per $100 million implies a market-wide elasticity closer to the estimates of dealer costs in this study.

V. Conclusions

The model presented incorporates the realistic options available to market makers for absorbing portfolio flows. In canonical models making markets entails moving prices away from the full information value to induce trades that compensate inventory imbalances. But these models constrain the dealer behavior to either paying inventory costs, or intentionally selling low and buying high. This paper suggests that there are multiple ways to control inventory costs.

One clear example is that in foreign exchange markets, the dealer has the ability to call others in the market and unload her unwanted inventory on them. Of course, this is not to suggest that outgoing orders are a panacea for inventory control, so these are modeled with price impact (i.e., increasing marginal costs). However, at the margin, she will equate the loss of trading unwanted inventory to incoming calls with the marginal price impact (i.e., the loss of trading unwanted inventory in outgoing calls) and with the marginal loss of the inventory imbalance (i.e., the marginal inventory carrying cost).

In addition, these outgoing calls do not occur in a vacuum. As long as events transpire during the outgoing call period, the dealer will learn through trading at those times and update her beliefs. These updates bring about price changes that neither inventory costs nor incoming order flow can explain. And FX dealers are just one example of market makers that smooth costs over multiple instruments. This paper argues that one should consider where dealers or specialists might be substituting away from conventional inventory costs when modeling price setting. Moreover, price-induced order flow is one of a multiplicity of informative instruments available to market makers.

The estimations support the proposed model and provide several novel empirical results. Generally, these indicate that previous studies overemphasize the role of price changes in inventory management, since no other instruments are considered. This omission biases downward the role of information in price changes, can make inventory effects appear insignificant, and tightens the bid-ask spread. The data generating process modeled here suggests that information effects are also biased downward in canonical estimations, since the dealer infers asset values from multiple signals which vary in their precision. Canonical estimates fail to correct for the varying precision of the informative flows, and hence, the information effect is biased downward as it overweighs the signals at uninformative times. The estimates also suggest that at the time of price setting, planned outgoing trades are one-third of the difference between dealer's current and optimal inventory positions, and a Fed intervention increases the informativeness of order flow, and lowers the cost of liquidity for the dealer. It also lowers inventory costs and tightens the spread.

Finally, the model addresses the broader relation between portfolio flows and asset prices. The presence of inventory effects suggests that part of observed price changes is transitory. However, with multiple instruments, dealers exhaust the gains from sharing a large inventory position with less price impact. As a result, the transitory component of price changes is less important than the information components from the multiple instruments. Hence, while both transitory and permanent effects are present in the data, the model favors a permanent impact of portfolio flows on prices.

Figure 1. The Timing of the Model

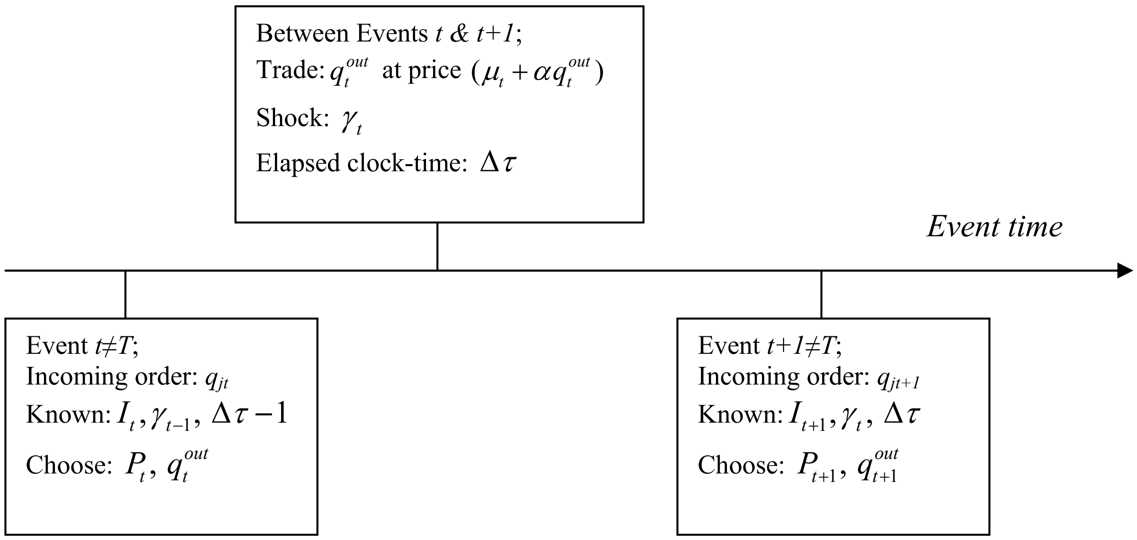

The figure above describes the timing of the model. At every event:

1. if $t \neq T$, the dealer knows her current inventory (denoted I_t), and a new incoming trade (one source of information for updating priors) occurs. The incoming quantity is q_{jt}.
2. The dealer decides her price (denoted by P_t) and plans her outgoing trade (denoted by q_t^{out}). These are the alternate methods available for offsetting inventory disturbances caused by the incoming trade.
3. Between events, the dealer executes the planned outgoing trade (q_t^{out}), and faces a quantity shock, (denoted by γ_t). This is another source of information for updating priors.
4. In addition, the dealer observes time elapsed between trades (denoted by $\Delta\tau$).
5. At the next event ($t+1$), the dealer uses the new incoming trade q_{jt+1} as well as the quantity shock between trades and the time elapsed between trades to update priors on the evolution of the asset value, and set prices.

Figure 2. Canonical Models' Information Effect: Incoming Order Flow and Price

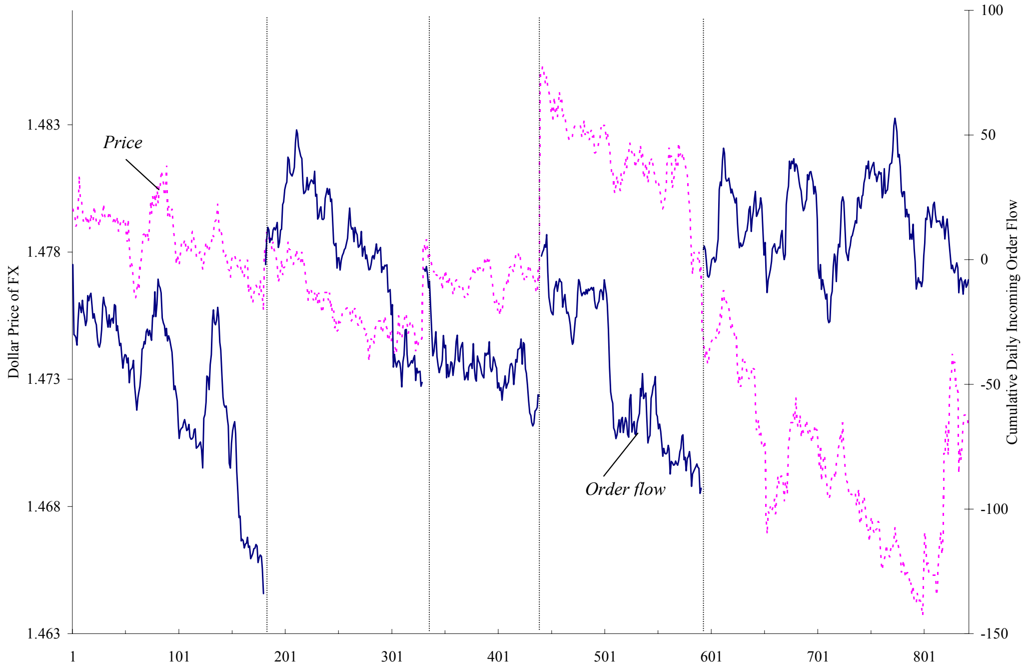

Figure 2 superimposes price on cumulative incoming order flow, August 3-7, 1992.

Figure 3. New Information Effect: Cumulative Inventory Shocks and Price

Figure 3 superimposes price on cumulative *unexpected* inventory shocks, August 3-7, 1992.

Table 1. Inventory Control: First Five Entries of Lyons (1995) Dataset

entry	p_{it}	q_{jt}	I_t	$q_t^{out} + \gamma_t$
1	1.4794	-1	1	1
2	1.4797	-2	3	-4
3	1.4795	-28	1	-30.5
4	1.4794	-0.5	-1.5	0.25
5	1.479	-0.75	-0.75	...

Notes: Table 1 shows the first five entries of the price (second column), incoming order flow (third column), and inventory (fourth column) variables from the data set. The last column captures the part of inventory evolution that is not due to incoming order flow, which reflects the optimal outgoing trade (q^{out}), and deviations driven by new information (γ). Lyons (1995) data: NY based dollar/DM dealer, August 3–7, 1992.

Table 2. Descriptive Statistics

	Inventory	Order flow	Order flow (absolute value)
Mean	2.16	-0.4	3.8
Median	0.7	0.5	2.5
Maximum	56.8	20.0	28.0
Minimum	-42.7	-28.0	0.0
Std. Dev.	15.4	5.2	3.6
Observations	838	838	838

Table 2 shows descriptive statistics for the dealer's inventory and incoming order flow.

Table 3. Observed Incoming Order Flow and Outgoing Trades

Observed Trades	Daily No. (mean)	Size (median)	Spread (median)
Incoming	170	3	0.0003
Outgoing	20	5	0.0003

Table 3 shows observed trades made by the dealer (not including a small amount of brokered trades). Note that outgoing refers to trades that the dealer is observed initiating. This is conceptually different from q^{out}, which represents an outgoing quantity planned at the time of price-setting that captures alternatives to shading the incoming transaction price for inventory control. Lyons (1995) data: NY based dollar/DM dealer, August 3–7, 1992.

Table 4. Information Effect: Daily Correlation of Order Flow Variables with Price

	Order Flow	**Unexpected Order Flow**	**Inventory Shocks**	**Mean Elapsed Time***	**End of Day Observation**
Monday	0.83	0.83	-0.54	1.77	181
Tuesday	0.69	0.71	0.58	1.86	330
Wednesday	0.53	0.48	0.03	2.44	440
Thursday	0.82	0.81	0.66	2.01	592
Friday	-0.03	-0.02	0.71	1.31	843

Table 4 shows the daily correlation between price and the order flow variable used to update priors. The first column shows incoming unexpected order flow and the second inventory shocks correlations for each day, August 3-7, 1992. The last column shows daily mean elapsed inter-transaction time.
* Reporting errors imply mean absolute value transaction time.

Table 5. System of Estimable Equations

<table>
<tr><th>Model Solution</th><th>Empirical Implementation</th><th>Testable Restrictions</th></tr>
<tr><td>Equation (20), inventory:
$I' = I + \beta(I - I^d) + \left(\frac{(1+\beta)}{2\alpha}\right)x$</td><td>$I_t = c_1 + c_2 I_{t-1} + \varepsilon_{1t}$</td><td>$\hat{I}^d = \frac{\hat{c}_1}{(1-\hat{c}_2)}$
$\hat{c}_2 = (1+\hat{\beta}) > 0$</td></tr>
<tr><td>Equation (19), outgoing trade:
$q^{out} = \left({}^{A_1}\!/_{A_1-\alpha}\right)\left[-(I - I^d) + \delta(\mu - p) + x\right]$</td><td>$\varepsilon_{2t} = (1-c_3)\left(\Delta I_t + q_{jt-1}\right) + c_3\left(I_t - \frac{c_1}{1-c_2}\right)$</td><td>$\hat{\gamma}_{t-1} = \hat{\varepsilon}_{2t}$
$\hat{q}^{out}_{t-1} = \hat{c}_3\left(-I_{t-1} + \hat{I}^d + q_{jt-1}\right)$
$\hat{c}_3 = \left({}^{A_1}\!/_{A_1-\alpha}\right) > 0$</td></tr>
<tr><td>Equation (21), price change:
$\Delta p_t = \psi\eta_t s_t + \beta({}^{\alpha}\!/_{(1+\delta\alpha)})(q^{out}_{t-1} + \gamma_{t-1} + q_{t-1}) +$
$+\psi(1-\eta_t)\gamma_{t-1} + \left(\frac{1+\delta\alpha(1-\beta)}{2\delta(1+\delta\alpha)}\right)\Delta x_t$</td><td>$\Delta p_t = c_{10} + c_{11}\eta_t q_{jt} +$
$+[c_{12}c_2 + c_{11}(1-\eta_t)(1-c_3)]\Delta I_t +$
$+c_{11}(1-\eta_t)c_3[-q_{jt-1} + (I_t - \frac{c_1}{1-c_2})] +$
$+c_{11}(1-\eta_t)q_{jt-1} + c_{13}D_t + c_{14}D_{t-1} + c_{15}D^{time}_t + \varepsilon_{3t}$</td><td>$\hat{c}_{12}\hat{c}_2 = \beta({}^{\alpha}\!/_{(1+\delta\alpha)}) < 0,$
with $(1+\beta) = \hat{c}_2,$
$\hat{c}_{12} < 0,$
$\hat{c}_{11} > 0$</td></tr>
<tr><td colspan="3">Outgoing Trade, Expected Marginal Cost: $\hat{\alpha} = {}^{\hat{c}'_{12}(\hat{c}_3-1)}\!/_{\hat{c}_3}$</td></tr>
</table>

Table 5 compares the algebraic solution to the model (in the first column) with the estimable equations these imply (in the second column). The final column shows testable restrictions on the model parameters. Row (1) shows inventory evolution, row (2) shows outgoing quantity, and row (3) shows price changes. The bottom shows the structural parameter measuring the expected cost of liquidity at the time of price setting, $\hat{\alpha}$. The system is estimated simultaneously using seemingly unrelated non-linear least squares.

Table 6. Price Formation with Multiple Instruments

$$\left\{\begin{array}{l} I_t = c_1 + c_2 I_{t-1} + \varepsilon_{1t} \\ \varepsilon_{2t} = (1-c_3)\left(\Delta I_t + q_{jt-1}\right) + c_3\left(I_t - \tfrac{c_1}{1-c_2}\right) \\ \Delta p_t = c_{10} + c_{11}\eta_t q_{jt} + [c_{12}c_2 + c_{11}(1-\eta_t)(1-c_3)]\Delta I_t + \\ + c_{11}(1-\eta_t)c_3[-q_{jt-1} + (I_t - \tfrac{c_1}{1-c_2})] + c_{11}(1-\eta_t)q_{jt-1} + c_{13}D_t + c_{14}D_{t-1} + c_{15}D_t^{time} + \varepsilon_{3t} \end{array}\right\}$$

Model	C_1	C_2	C_3	C_{10}	C_{11}	C_{12}	C_{13}	C_{14}	C_{15}	adj R^2	α	I^d
											(pips)	*US$ M*
Full sample	0.72	0.66	0.34	0.00	1.71	-1.79	11.85	-9.40	-0.40	0.20	0.35	2.09
Sample: 2 838	*0.03*	*0.00*	*0.00*	*0.26*	*0.00*	*0.00*	0.00	0.00	*0.92*			
First half	-0.71	0.58	0.42	0.00	1.16	-0.70	15.07	-9.07	-0.23	0.30	0.09	-1.70
Sample: 2 460	*0.09*	*0.00*	*0.00*	*0.28*	*0.05*	*0.10*	0.00	0.00	*0.96*			
Second half	2.35	0.69	0.29	0.00	2.21	-3.22	9.93	-11.42	5.54	0.29	0.78	7.51
Sample: 460 796	*0.00*	*0.00*	*0.00*	*0.07*	*0.01*	*0.00*	0.00	0.00	*0.40*			

Table 6 estimates the system of equations imposing all identifying restrictions. Estimation is robust over subsamples of this dataset, including around the approximate break points found in previous canonical model estimations. α measures the expected price impact of augmenting the planned outgoing trade by $1 million in pips, and I^d measures the implicit optimal inventory level used by the dealer. All estimates multiplied by 10^5, p-values in italics, Lyons (1995) dataset.

Table 7. Canonical Model Estimates

$$\Delta p_t = \beta_0 + \beta_1 q_{jt} + \beta_2 I_t + \beta_3 I_{t-1} + \beta_4 D_t + \beta_5 D_{t-1} + ma(1)$$

	β_0	β_1	β_2	β_3	β_4	β_5	Adj. R^2
Full Sample	-1.34	1.47	-0.91	0.72	10.30	-9.12	0.22
Sample (adjusted): 2 838	*0.32*	*0.00*	*0.00*	*0.01*	*0.00*	*0.00*	
First half	-1.87	1.34	-0.45	0.21	12.44	-8.76	0.34
Sample (adjusted): 2 460	*0.11*	*0.00*	*0.10*	*0.43*	*0.00*	*0.00*	
Second half	-2.99	1.13	-1.99	1.82	10.00	-10.50	0.28
Sample: 460 796	*0.18*	*0.11*	*0.00*	*0.00*	*0.00*	*0.00*	
Memorandum:							
Break Tests	*Observation*	*F-statistic*	*Log likelihood ratio*				
	460	*0.03*	*0.03*				
	796	*0.00*	*0.00*				

Table 7 reproduces canonical microstructure estimates using the Lyons (1995) dataset. All estimates multiplied by 10^5. Estimating over the two halves of the sample reveals that the simultaneous presence of inventory and information effects predicted by canonical models are significantly not different from zero (See Romeu (2005)). Hence, while inventory and information appear to be present in the data, canonical model predictions are overturned as predicted by Section II.D.

Table 8. The Impact of a Federal Reserve Intervention

Model	C_1	C_2	C_3	C_{10}	C_{11}	C_{12}	C_{13}	C_{14}	C_{15}	adj R^2	α	I^d
											(pips)	*US$ M*
Full Sample	0.72	0.66	0.34	0.00	1.71	-1.79	11.85	-9.40	-0.40	0.20	0.35	2.09
	0.03	*0.00*	*0.00*	*0.26*	*0.00*	*0.00*	0.00	0.00	*0.92*			
No Fed Intervention	0.72	0.67	0.33	0.00	1.48	-1.78	12.99	-9.84	2.44	0.28	0.36	2.20
	0.03	*0.00*	*0.00*	*0.03*	*0.00*	*0.00*	0.00	0.00	*0.51*			

Table 8 shows cost comparisons for the $300 million Fed intervention on August 7, 1992. The exact start time and sequence of the intervention is unknown. * Wall Street Journal, August 10, 1992: "The Federal Reserve Bank of New York moved to support the U.S. currency... as the dollar traded at 1.4720." This is the most precise documentation available of the intervention start, and that price corresponds to 12:32 pm. Other times selected because of reports of a mid-day start, and because between 12:26 and 12:32 pm, the price jumped 36 pips, suggesting a possible intervention start there.

APPENDIXES

I. MODEL SOLUTION

Inventory Carrying Cost

From equations (2) and (3) the variance of the dealer's portfolio is

$$\sigma_{W_t}^2 = \sigma_V^2 I_t^2 + \sigma_y^2 + 2I_t\sigma_{vy}\,. \tag{30}$$

Add and subtract $\left(\frac{\sigma_{vy}}{\sigma_v^2}\right)^2$ into (10) to get:

$$c_t = \omega\left[\sigma_y^2 - \left(\frac{\sigma_{vy}}{\sigma_v^2}\right)^2 + \sigma_V^2 I_t^2 + 2I_t\sigma_{vy} + \left(\frac{\sigma_{vy}}{\sigma_v^2}\right)^2\right] = \omega\left[\left(\sigma_y^2 - \left(\frac{\sigma_{vy}}{\sigma_v^2}\right)^2\right) + \sigma_v^2\left(I_t - \frac{-\sigma_{vy}}{\sigma_v^2}\right)^2\right] \tag{31}$$

Which is the right-hand-side of (10) with coefficients:

$$I^d = \left(\frac{-\sigma_{vy}}{\sigma_v^2}\right) \quad \phi_1 = \sigma_v^2 \quad \phi_0 = \sigma_y^2 - \left(\frac{\sigma_{vy}}{\sigma_v^2}\right)^2. \tag{32}$$

Dealer's Beliefs

Given market demand q_{jt}, the dealer creates a statistic based on the intercept of the demand curve, which is independent of her price. Denote this statistic as D_t.

$$D_t = q_{jt} + \delta p_t = \delta(v_t - p_t) + X_t + \delta p_t = \delta v_t + X_t\,. \tag{33}$$

From the signal of market demand D_t the dealer forms two statistics. The first is an innovation in the full information value of the risky asset, which shall be denoted as s_t. The second is a signal of the liquidity demand, which is denoted as (lower case) x_t, and will depend on the estimate of full information value, μ_t.

$$w_t = \delta^{-1}D_t = v_t + \tfrac{X_t}{\delta}; \quad E[w_t] = v_t \tag{34}$$

$$x_t = D_t - \delta\mu_t, \quad E\left[x_t\right] = X_t\,. \tag{35}$$

Consistent with rational expectations, assume that the dealer's previous estimate, μ_{t-1} is the steady state distribution over the true asset value v_t, and that the variance of μ_t is proportional to the variance of w_t. Hence, one can write $\sigma_\mu^2 = \Omega\sigma_w^2$. Given the variance of w_t, form a signal to noise ratio given by:

$$\Upsilon = \frac{\sigma_v^2}{\sigma_w^2}, \quad \textit{with } \sigma_w^2 = \delta^{-2}\sigma_x^2\,. \tag{36}$$

The dealer uses the recursive updating of a Kalman filter to form the expectations over v_t. This implies that she updates the prior belief μ_{t-1} using the current order flow w_t. The resulting posterior, μ_t, converges to a steady-state distribution whose time varying mean is an unbiased estimate of the true value of v_t. The recursive equations to generate this estimate are given by:

$$\Omega = \frac{-\Upsilon + \sqrt{\Upsilon^2 + 4\Upsilon}}{2}, \quad (37)$$

Hence, if the dealer had only information based on the incoming order, she would use the following estimate, which is denoted as μ_t^Z, as the estimate of v_t :

$$\mu_t^Z = \Omega w_t + (1-\Omega)\mu_{t-1}. \quad (38)$$

Note, however, that the dealer also receives information for updating μ_{t-1} through a linear function of the inventory shock which is denoted by $\kappa(\gamma_{-1})$. Given $\kappa(\gamma_{-1})$, an unbiased estimate of v_t is given by:

$$\mu_t^\gamma = \Omega[\mu_{t-1} + \kappa(\gamma_{t-1})] + (1-\Omega)\mu_{t-1} = \mu_{t-1} + \Omega\kappa(\gamma_{t-1}), \quad (39)$$

where the same Kalman filter algorithm as defined above is used. Hence there are two signals of v_t at the time of setting the price. Given the assumption, the variance of μ_t^γ is a linear function of the variance of μ_t^Z. That is,

$$\mathrm{var}(\mu_t^Z) = \sigma_Z^2, \quad \mathrm{var}(\mu_t^\gamma) = \sigma_Z^2 * \Delta\tau, \quad (40)$$

where $\Delta\tau$ is the elapsed clock time between incoming order *(t-1)* and *t*. The optimal signal for the dealer is then:

$$\mu_t = \eta\mu_t^Z + (1-\eta)\mu_t^\gamma = \eta[\Omega w_t + (1-\Omega)\mu_{t-1}] + (1-\eta)[\mu_{t-1} + \Omega\kappa(\gamma_{t-1})]. \quad (41)$$

with $\eta = \left({}^{\Delta\tau}\!/_{1+\Delta\tau}\right)$. Now grouping and rearranging:

$$\mu_t - \mu_{t-1} = \eta\Omega(w_t - \mu_{t-1}) + (1-\eta)\Omega\kappa(\gamma_{t-1}) = \eta\Omega(\delta^{-1}D_t - \mu_{t-1}) + (1-\eta)\Omega\kappa(\gamma_{t-1}) \quad (42)$$

Since $w_t = \delta^{-1}D_t = v_t + \dfrac{X_t}{\delta}$,

$$\mu_t - \mu_{t-1} = \eta\Omega(\delta^{-1}(q_{jt} + \delta p_t) - \mu_{t-1}) + (1-\eta)\Omega\kappa(\gamma_{t-1}) \quad (43)$$

Add and subtract $\delta\mu$ to get:

$$\mu_t - \mu_{t-1} = \eta\Omega\delta^{-1}[q_{jt} - \delta(\mu_t - p_t) + \delta(\mu_t - \mu_{t-1})] + (1-\eta)\Omega\kappa(\gamma_{t-1}) \quad (44)$$

Solving for $(\mu_t - \mu_{t-1})$ yields,

$$(\mu_t - \mu_{t-1})[1-\Omega\eta] = \eta\Omega\delta^{-1}[q_{jt} - \delta(\mu_t - p_t)] + (1-\eta)\Omega\kappa(\gamma_{t-1}) \quad (45)$$

Which gives the final relationship for the updating:

$$\Delta\mu_t = \xi_1 s_t + \xi_2\kappa(\gamma_{t-1}), \quad (46)$$

Where $s_t = q_{jt} - \delta(\mu_t - p_t)$ is the unexpected order flow, and

$$\xi_1 = \frac{\eta\Omega}{\delta(1-\Omega\eta)} \quad \& \quad \frac{\partial\xi_1}{\partial\eta} > 0; \quad \xi_2 = \frac{(1-\eta)\Omega}{\delta(1-\Omega\eta)} \quad \& \quad \frac{\partial\xi_2}{\partial\eta} < 0. \quad (47)$$

Hence, ξ_1 and ξ_2 are inversely related with respect to η, and as inter-transaction time is longer, more weight is placed on the unexpected incoming order flow signal s_t. Here, $\kappa(\gamma_{t-1})$ is assumed to be some simple linear function: $\kappa(\gamma_{t-1}) = \omega_0 + \omega_1\gamma_{t-1}$, where ω_0 may be assumed zero if desired.

The Dealer's Problem

The dealer's problem is reproduced here:

$$J(I_t, x_t, \mu_t, K_t) = \max_{p_t, q_t^{out}} E\left\{(1-\rho)\left[\tilde{v}_t I_t + K_t - c_t\right] + \rho J(\tilde{I}_{t+1}, \tilde{x}_{t+1}, \tilde{\mu}_{t+1}, \tilde{K}_{t+1})\right\}, \tag{48}$$

subject to the following evolution constraints:

$$E\left[\tilde{I}_{t+1} \mid \Phi_t^i\right] = I_t - \delta\left(\mu_t - p_t\right) - x_t + q_t^{out}, \tag{49}$$

$$E\left[\tilde{x}_{t+1} \mid \Phi_t^i\right] = 0, \tag{50}$$

$$E\left[\tilde{\mu}_{t+1} \mid \Phi_t^i\right] = \mu_t, \tag{51}$$

$$E\left[\tilde{K}_{t+1} \mid \Phi_t^i\right] = K_t + p_t\delta\left(\mu_t - p_t\right) + p_t x_t - (\mu_t + \alpha q_t^{out}) q_t^{out} - c_t, \tag{52}$$

For expositional simplicity, in what follows the expectation operators on the evolution equations and the time subscripts are dropped, and a forward lag is denoted by a 'superscript.' The first order conditions are given by:

$$p:\ \delta E\left[J_I(I', x', \mu', K')\right] + (\delta\mu - 2\delta p + x) E\left[J_K(I', x', \mu', K')\right] = 0, \tag{53}$$

$$q^{out}:\ \ E\left[J_I(I', x', \mu', K')\right] - (\mu + 2\alpha q^{out}) E\left[J_K(I', x', \mu', K')\right] = 0. \tag{54}$$

Substituting (54) into (53), and assuming for now that $E\left[J_K(I', x', \mu', K')\right] \neq 0$ (I confirm this later), price is:

$$p = \mu + \frac{x}{2\delta} + \alpha q^{out}. \tag{55}$$

Denote from here on the value function without its arguments for notational simplicity, maintaining the convention that $J(')$ is the forward lag of $J()$. Furthermore, in what follows a subscript denotes the derivative of the function with respect to that argument. The envelope conditions for this problem are:

$$J_I() = (1-\rho)\mu + \rho E[J_I(')] - 2\omega\phi_1(I - I^d)\left[(1-\rho) + \rho E[J_K(')]\right]; \tag{56}$$

$$J_x() = -\rho\left(E[J_I(')] - pE[J_K(')]\right); \tag{57}$$

$$J_\mu() = (1-\rho)I - \delta\rho E[J_I(')] + \rho E[J_\mu(')] + \rho(\delta p - q^{out}) E[J_K(')]; \tag{58}$$

$$J_K() = (1-\rho) + \rho E[J_K(')]; \tag{59}$$

Based on the envelope conditions, it is conjectured that the value function takes on the functional form:

$$\tilde{J}(I, x, \mu, K) = A_0 + \mu I + K + A_1(I - I^d)^2 + A_2 x(I - I^d) + A_3 x^2. \tag{60}$$

Using the conjecture, and the evolution equations, taking the derivatives with respect to *I* and *K* updating:

$$E[\tilde{J}_I(')] = E[\mu' + 2A_1(I' - I^d) + A_2 x'] = \mu + 2A_1(I' - I^d). \tag{61}$$

$$E[\tilde{J}_K(')] = E[1] = 1. \tag{62}$$

Plugging (61) and (62) into (54) yields the optimal outgoing quantity:

$$\frac{A_1}{\alpha}(I' - I^d) = q^{out}. \tag{63}$$

Substituting (63) into (55) for q^{out} yields the pricing equation:

$$p = \mu + A_1(I' - I^d) + \frac{x}{2\delta}. \tag{64}$$

Taking the evolution equation for inventory, (49), one can substitute (64) in for *p* and solve for I' to get:

$$I' = I + \beta(I - I^d) - \tfrac{(1+\beta)}{2}x, \tag{65}$$

with

$$\beta = \left(\frac{A_1(1+\delta\alpha)}{\alpha - A_1(1+\delta\alpha)}\right) \Leftrightarrow A_1 = \left(\frac{\beta\alpha}{(1+\beta)(1+\delta\alpha)}\right). \tag{66}$$

Given the inventory evolution of (65), one can solve for the optimal pricing policy function, recognizing that relationship in (66) simplifies the implicit function of A_1 multiplied by $(1+\beta)$ to $A_1(1+\beta) = \beta(\alpha/_{(1+\delta\alpha)})$, and substituting:

$$p = \mu + \beta(\alpha/_{(1+\delta\alpha)})(I - I^d) + \left(\tfrac{1+\delta\alpha(1-\beta)}{2\delta(1+\delta\alpha)}\right)x. \tag{67}$$

Taking first differences of (67), and substituting in:

$$\Delta p = \Delta\mu - \beta(\alpha/_{(1+\delta\alpha)})q_{jt-1} + \beta(\alpha/_{(1+\delta\alpha)})(q_{t-1}^{out} + \gamma_{t-1}) + \left(\tfrac{1+\delta\alpha(1-\beta)}{2\delta(1+\delta\alpha)}\right)\Delta x. \tag{68}$$

Substituting the relationship for the updating of the μ_t given by (46) yields:

$$\Delta p = \xi_1 s_t + \xi_2\kappa(\gamma_{t-1}) - \beta(\alpha/_{(1+\delta\alpha)})q_{jt-1} + \beta(\alpha/_{(1+\delta\alpha)})(q_{t-1}^{out} + \gamma_{t-1}) - \left(\tfrac{1+\delta\alpha(1-\beta)}{2\delta(1+\delta\alpha)}\right)\Delta x \tag{69}$$

Next the conjectured functional form of (60) is confirmed. Begin by taking the envelope condition for *x*, (57), and solve for coefficients A_2 and A_3 of the conjectured functional form's derivative, which is:

$$\tilde{J}_2 = A_2(I - I^d) + 2A_3x \tag{70}$$

Substituting the optimal policy functions into (57), as well as the updated derivatives of the conjectured functional form which are given by (61) and (62) yields: $A_2 = -\rho A_1(1+\beta)$, and $A_3 = \rho(1+\delta A_1(1+\beta))/_{4\delta}$. Continuing, the envelope condition on *I* in (60) can be solved with the conjectured functional form's derivative, which is given in (61). This yields $A_1 = \left[(-\omega\phi_1)/_{1-\rho(1+\beta)}\right]$. An economically sensible solution requires $A_1<0$, hence, using the definition for A_1, it is required that:

$$\beta + \frac{\omega\phi(1+\beta)(1+\delta\alpha)}{1-\rho(1+\beta)} = 0. \tag{71}$$

This implies $\beta \in (-1,0)$. As $\beta \to -1$, the right-hand-side of (71) goes to negative one. As $\beta \to 0$, the right-hand-side of (71) is positive. Hence, since (71) is a continuous function, by the Mean Value Theorem $\exists \beta \in (-1,0) \therefore$ (71) holds.

The Informed Trader's Problem

This section shows that the conjectured behavior of the informed trader is optimal given the dealer's optimal solution for price setting. This proof adapts the Madhavan and Smidt (1993) proof that conditions exist such that any deviation from the conjectured result would be suboptimal. The informed maximizes her terminal wealth after observing the liquidation value of the asset, and facing the same stochastic probability of a trading event occurring as the dealer of the previous section. Hence, prior to trading at time t, the informed faces a probability $(1-\rho)$ of no trade occurring, in which she keeps her expected wealth, $v_t B_t + C_t$, with *B* and *C* representing the endowments of risky asset and capital, respectively. In the alternative, the informed trades, and updates her stocks to $B_t + Q_t$ and $C_t - p_t Q_t$, respectively. We show that for Δ different from zero, $Q_t = \delta(v_t - p_t) + \Delta_t$ is suboptimal. The informed observes the dealer's price, which is a function of her order through its effect on the dealer's inventory and information. Taking the information effect first, using $w_t = \delta^{-1} D_t$, the dealer's signal, with $D_t = \delta(v_t - p_t) + X_t + \Delta_t + \delta p_t$, the introduction of a non-zero deviation yields a distorted signal, $w_t' = w_t' + \delta^{-1}\Delta_t$. This, in turn, yields price as an increasing function of the deviation:

$$P_t(\Delta_t) = p_t + \lambda\Delta_t, \quad with\,\lambda = \Omega\eta\delta^{-1} \tag{72}$$

Where p_t would be the price prevailing if $\Delta_t = 0$ held. As in the case of the dealer, denote by $V(v_t, p_t, B_t, C_t)$ the maximum expected wealth given the state, represented by the price, asset liquidation value, and the capital and inventory stocks. The informed trader chooses the optimal quantity for the order, which, by construction, allows the problem to be expressed as:

$$V(v_t, p_t, B_t, C_t) = \max_{\Delta_t} E[(1-\rho)(\upsilon_t B_t + C_t) + \rho V(v_t, p_t, B_t, C_t)] \tag{73}$$

With transitional equations, $E[v_{t+1}] = v_t$, $B_{t+1} = B_t + Q_t$, and $C_{t+1} = C_t - P_t Q_t$, with $P_t = p_t + \lambda\Delta_t$, and $Q_t = \delta(v_t - p_t) + \Delta_t$. Turning to the transitional equation for the notational base price, note that the price next trade depends on the trader's current quantity through information and inventory effects, which in turn in a function of Δ_t. Hence, we can restate the dealer's solution consistent with (67) as:

$$p_{t+1} = \mu_{t+1}(\Delta_t) + \zeta_1(I_{t+1} - I^d) + \zeta_2 x_{t+1} \tag{74}$$

Where, $E[\mu_{t+1}(\Delta_t)] = \mu_t(\Delta_t)$ by iterated expectations, and from (41), $\mu_t(\Delta_t) = \mu_t + \lambda\Delta_t$, which implies that the dealer's expectations of the liquidation value are adjusted by the non-zero Δ_t "excess" trade if the informed deviates. From (67), we can rewrite the expectation of μ_t as $\mu_t = p_t - \zeta_1(I_t - I^d) - \zeta_2 x_t$. Using the expression derived for $\mu_t(\Delta_t)$ and (74), we have that price evolves by $p_{t+1} = p_t + \lambda\Delta_t + \zeta_1(I_{t+1} - I_t) + \zeta_2(x_{t+1} - x_t)$. Taking expectations, and using (19) for q^{out}, we have:

$$E[p_{t+1}] = p_t + \lambda\Delta_t + (\tfrac{\alpha}{A_1-\alpha})\zeta_1 q_{jt} - \zeta_1(\tfrac{\alpha}{A_1-\alpha})(I_t - I^d) + \zeta_2(x_{t+1} - x_t) \tag{75}$$

Here, we can assume without loss of generality that the informed trader does not have a priori knowledge about our dealer's inventory levels.[30] However, this equation shows the full impact of a deviation affects the future price both through changes in the dealer's expectation, and through her inventory pressure. Note that in the event that the marginal cost of trading out to other dealers is zero (i.e. $\alpha = 0$), only the information channel is relevant, as the inventory adjustment is complete, illustrating the dichotomy between multiple instruments in this approach and canonical models. Omitting time subscripts, using superscripts to denote one-period ahead, the first order condition for (73) is:

$$EV_B' - (\lambda\delta(v-p) + p + 2\lambda\Delta)EV_C' + (\lambda + \zeta_1)EV_p' = 0 \tag{76}$$

Taking the envelope conditions:

$$V_v = (1-\rho)B + \rho EV_v' + \rho\delta EV_B' - \rho EV_v'(p\delta + \lambda\Delta\delta) + \rho EV_p'\zeta_1\delta \tag{77}$$

$$V_p = -\rho\delta EV_B' + \rho(1-\zeta_1\delta)EV_p' + \rho(-\delta(v-p) + \delta p(1+\lambda\Delta) - \Delta)EV_C' \tag{78}$$

$$V_B = (1-\rho)v + \rho V_B' \tag{79}$$

$$V_B = (1-\rho)v + \rho V_C' \tag{80}$$

These suggest a conjectured functional form for the value function of:

$$\tilde{V} = vB + C + A(v-p)^2 \tag{81}$$

With derivatives, $\tilde{V}_v = B + 2A(v-p)$, $\tilde{V}_p = -2A(v-p)$, $\tilde{V}_B = v$, and $\tilde{V}_C = 1$.

The transitional equations yield $E(v'-p') = (v-p)(1 - (\tfrac{\alpha}{A_1-\alpha})\zeta_1\delta) - \Delta((\tfrac{\alpha}{A_1-\alpha})\zeta_1 + \lambda)$, which we can substitute into the first order condition, and set the deviation to zero, which in turn gives a condition for A:

$$A = \frac{(1-\lambda\delta)}{2(\zeta_1 + \lambda)(1 - (\tfrac{\alpha}{A_1-\alpha})\zeta_1\delta)}. \tag{82}$$

Taking the envelope condition for $\tilde{V}_v$, and substituting in the expected values with the use of the evolution equations, a second condition is imposed on A:

---

[30] This is consistent with other inventory models and evidence from financial markets. In the alternative, it is straightforward to show that including a non-zero expectation on either of the last two terms in (75) leaves the pricing equation unaltered.

$$A = \frac{\rho\delta}{1-\rho(1-\zeta_1\delta)(1-(\frac{\alpha}{A_1-\alpha})\zeta_1\delta)}. \tag{83}$$

Note that conditions (82) and (83) are analogous to the restricted case presented in Madhavan and Smidt (1993), where differences will appear in both the wedge associated with the inventory adjustment due to q^{out}, in this case, $(\frac{\alpha}{A_1-\alpha})$, and the scaling of the updating coefficient, Ω by the elapsed time fraction. As indicated in Section II.D, the model presented would yield the informed trader of Madhavan and Smidt (1993) if the aforementioned effects are restricted away.

Since $\delta\lambda = \Omega\eta$, we can express the conditions imposed by (82) and (83) as finding a $\delta \in (0,\infty)$ such that the function below satisfies:

$$\frac{(1-\Omega\eta)}{2\rho(1-(\frac{\alpha}{A_1-\alpha})\zeta_1\delta)} - \frac{\delta\zeta_1+\Omega\eta}{1-\rho(1-\zeta_1\delta)(1-(\frac{\alpha}{A_1-\alpha})\zeta_1\delta)} = 0 \tag{84}$$

Equation (84) represents a continuous function in δ, directly and indirectly through both Ω and β. We can express $\delta\zeta_1 = \delta A_1(1+\beta) = {}^{\beta(\delta\alpha)}\!/_{(1+\delta\alpha)}$, and it is straightforward to show that as $\delta \to 0$, $\Omega \to 0$, and $\delta\zeta_1 \to 0$, and we can express $(\frac{\alpha}{A_1-\alpha})\zeta_1\delta = {}^{-\delta\alpha\beta}\!/_{(1+\delta\alpha(1+\beta))}$. Hence, as $\delta \to 0$, (84) is positive, and converges to $1/_{2\rho} > 0$. Moreover, as $\delta \to \infty$, $\Omega \to 1$, $\beta \to -1$, and $\delta\zeta_1 \to -1$, and $(\frac{\alpha}{A_1-\alpha})\zeta_1\delta \to \infty$. Applying L'Hôpital's rule, it can be shown that (84) becomes negative, and hence, by the mean value theorem, $\exists\delta \in (0,\infty) \therefore (84)$ holds.

REFERENCES

Admati, A., and P. Pfleiderer, 1988, "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies*, Vol. 1 (Spring), pp. 3–40.

Amihud, Y., and H. Mendelson, 1980, "Dealership Market: Marketmaking with Inventory," *Journal of Financial Economic*s, Vol. 8 (March), pp. 31–53.

Andrews, D.W.K., 1993, "Testing for Parameter Instability and Structural Change with Unknown Change Point," *Econometrica,* Vol. 61 (July), pp. 821–56.

Ausubel, L. and R. Romeu, 2005, "Bidder Participation and Information in Currency Auctions." (unpublished; IMF Working Paper No. 05/157, International Monetary Fund).

Bai, J. and P. Perron, 1998, "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica,* Vol. 66 (January), pp. 47–78.

Bjonnes, G., and D. Rime, 2000, "FX Trading ... Live: Impact of New Trading Environments," (unpublished; Oslo: Norwegian School of Management).

Blitz, J., 1992, "Dollar Below Intervention level," *Financial Times* (London: August 11, 1992), p. 23.

Cao, H. H., M. D. D. Evans, and R. K. Lyons, 2006, "Inventory Information," *The Journal of Business*, Vol. 79, pp. 325–364.

Dominguez, K., and J. Frankel, 1993a, "Does Foreign-Exchange Intervention Matter? The Portfolio Effect," *American Economic Review*, Vol. 83 (December), pp.1356–69.

———, 1993b, "Does Foreign-Exchange Intervention Work?" (Washington: Institute for International Economics).

Easley, D., and M. O'Hara, 1987, "Price, Trade Size, and Information in Securities Markets," *Journal of Financial Economics*, Vol. 19 (September), pp. 69–90.

———, 1992, "Time and the Process of Security Price Adjustment," *Journal of Finance*, Vol. 47 (June), pp. 577–605.

Evans, M. D. D. and R. K. Lyons, 2002. "Order Flow and Exchange Rates Dynamics." *Journal of Political Economy*, Vol. 110 (February), pp. 170–180.

"US Fed Intervenes as Dollar Nears Low Against D-Mark," *Financial Times* (London: August 8, 1992), p.1.

Flood, R., and A. Rose, 1995, "Fixing Exchange Rates: A Virtual Quest for Fundamentals," *Journal of Monetary Economics*, Vol. 36 (December), pp 3–37.

Flood, R., and M. Taylor, 1996, "Exchange Rate Economics: What's Wrong with the Conventional Macro Approach?" *The Microstructure of Foreign Exchange Markets,* ed. by J. Frankel, G. Galli, and A. Giovannini (Chicago: The University of Chicago Press), pp. 261–294.

Frankel, J., G. Galli, and A. Giovannini, 1996, *The Microstructure of Foreign Exchange Markets,* (Chicago: University of Chicago Press).

Frankel, J., and A. Rose, 1995, "Empirical Research on Nominal Exchange Rates," *Handbook of International Economic*s, ed. By G. Grossman and K. Rogoff (Amsterdam: Elsevier Science) pp.1689–1729.

Froot, K. A., P. O'Connel, and M. Seasholes, 2001, "The Portfolio Flows of International Investors," *Journal of Financial Economics*, Vol. 59 (February). pp.151–93.

Froot K. A. and T. Ramadorai, 2001a. "The Information Content of International Portfolio Flows," NBER Working Papers 8472, National Bureau of Economic Research, Cambridge, MA.

———, Froot, K., and T. Ramadorai, 2002. "Currency Returns, Institutional Investor Flows, and Exchange Rate Fundamentals." NBER Working Paper 9101, National Bureau of Economic Research, Cambridge, MA.

Glosten, L., and P. Milgrom, 1985, "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Agents," *Journal of Financial Economic*s, Vol. 14 (March), pp. 71–100.

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

———, 1991a, Measuring the Information Content of Stock Trades," *Journal of Finance*, Vol. 46 (March), pp.179–207.

———. 1991b, The Informativeness of Stock Trades: An Econometric Analysis," *Review of Financial Studies*, Vol.4 (No. 3), pp. 571–595.

Ho, T., and H. Stoll, 1981, "Optimal Dealer Pricing Under Transactions and Return Uncertainty," *Journal of Financial Economic*s, Vol. 9 (March), pp.47–73.

———, 1983, "The Dynamics of Dealer Markets under Competition," *Journal of Finance* Vol. 38 (September), pp. 1053–74.

Ito, T., R. Lyons, and M. Melvin, 1998, "Is There Private Information in the FX Market? The Tokyo experiment," *Journal of Finance,* Vol. 53 (June), pp. 1111–30.

Kyle, A., 1985, "Continuous Auctions and Insider Trading," *Econometrica,* Vol. 53 (November), pp. 1315–35.

Linton, C., "Dollar Likely to Fall Though Banks May Try to Slow Currency's Descent," *Wall Street Journal* (New York: Aug 10, 1992), p. C13.

Lyons, R., 1995, "Tests of Microstructural Hypotheses in the Foreign Exchange Market," *Journal of Financial Economics*, Vol. 39 (October), pp. 321–51.

———, 1996, "Optimal Transparency in a Dealer Market with an Application to Foreign Exchange," *Journal of Financial Intermediatio*n, Vol. 5 (July) pp. 225–54.

———, 1997, "A Simultaneous Trade Model of the Foreign Exchange Hot Potato," *Journal of International Economics*, Vol. 42 (May), pp. 275–98.

———, 1998, "Profits and Position Control: A Week of FX Dealing," *Journal of International Money and Finance*, Vol. 17 (February), pp. 97–115.

———, 2001, *The Microstructure Approach to Exchange Rate*s, (Cambridge and London: MIT Press).

Madhavan, A., and G. Sofianos, 1998, "An Empirical Analysis of NYSE Specialist Trading," *Journal of Financial Economics* Vol. 48 (May), pp. 189–210.

Madhavan, A., and S. Smidt, 1991, "A Bayesian Model of Intraday Specialist Pricing," *Journal of Financial Economics*, Vol. 30 (November), pp. 99–134.

———, 1993, "An Analysis of Daily Changes in Specialist Inventories and Quotations." *Journal of Finance*, Vol. 48 (December), pp.1595–1648.

Manaster, S. and S. C. Mann, 1996, "Life in the Pits: Competitive Market Making and Inventory Control," *Review of Financial Studies* Vol. 9 (Fall), pp. 953–75.

Massa, M. and A. Simonov, 2003, "Reputation and Interdealer Trading. A Microstructure Analysis of the Treasury Bond Market." Journal of Financial Markets, Vol. 6, Number 2, (April), pp. 99-141.

Meese, R., and K. Rogoff, 1983, "The Out-of-sample Failure of Empirical Exchange Rate Models," *Exchange Rate and International Macro-economic*s, ed. by J. Frenkel (Chicago: University of Chicago Press).

Mello, A., 1996, Comment–Chapter 5. *The Microstructure of Foreign Exchange Markets,* ed. by J. Frankel, G. Galli, and A. Giovannini (Chicago: The University of Chicago Press).

O'Hara, M.; Oldfield, G. S., 1986 "The Microeconomics of Market Making," *Journal of Financial and Quantitative Analysis*, Vol. 21 (December), pp. 361–76.

Rime, D., 2000, "Private or public information in foreign exchange markets? An empirical analysis," (unpublished; Oslo: Norwegian School of Management).

Romeu, R. B., 2005, "Why Are Asset Markets Modeled Successfully, But Not Their Dealers?" *IMF Staff Papers*, Vol. 52, No. 3 (December), pp. 369–86.

Stoll, Hans R., 1978, "The Supply of Dealer Services in Securities Markets," *Journal of Finance* Vol. 33 (September), pp. 1133–51.

Yao, J., 1998, "Market Making in the Interbank Foreign Exchange Market" (unpublished; New York: Salomon Center Working Paper No. S-98-3, New York University).