

# PROFITS AND POSITION CONTROL: A WEEK OF FX DEALING

Richard K. Lyons

U.C. Berkeley and NBER

This version: June 1997

## Abstract

This paper examines foreign exchange trading at the dealer level. The dealer we track averages \$100,000 in profits per day on volume of \$1 billion per day (or one basis point). The half-life of the dealer's position is only ten minutes, providing strong support for inventory models. A methodological innovation allows us to identify his speculative position over time. This speculative position determines the share of profits deriving from speculation versus intermediation: intermediation is much more important.

Forthcoming, *Journal of International Money & Finance*.

### Correspondence

Professor Richard K. Lyons  
Haas School of Business, U.C. Berkeley  
Berkeley, CA 94720-1900  
Tel: 510-642-1059, Fax: 510-642-4700  
E-mail: lyons@haas.berkeley.edu

---

\* I thank the following for helpful comments: Jeff Bohn, Frank Diebold, Bernard Dumas, Mark Flood, Dick Meese, Mark Ready, Andrew Rose, Patrik Säfvenblad, Matt Spiegel, Avanidhar Subrahmanyam, and seminar participants at Berkeley, UCLA, the Federal Reserve Board, the Federal Reserve Bank of New York, the NBER, and the Stockholm School of Economics. I also thank Jeff Bohn for valuable research assistance. Parts of this paper were written while visiting the IIES at Stockholm University and I thank them for their hospitality. Financial assistance from the National Science Foundation and the Berkeley Program in Finance is gratefully acknowledged.

# PROFITS AND POSITION CONTROL: A WEEK OF FX DEALING

Empirical work on FX microstructure is still in its early stages. The earliest work used futures data since that was available at high frequencies.<sup>1</sup> In FX, however, the futures market is both much smaller than spot and tightly linked to spot through arbitrage. Moreover, early data sets did not have sufficient granularity to capture agent heterogeneity, the hallmark of the microstructure approach. Work on the spot market itself grew in the early 90s with the availability of quotes on an intraday basis (specifically, the indicative quotes from Reuters called FXXF).<sup>2</sup> These quotes provide a quite accurate picture of price dynamics. More important, they also speak to heterogeneity issues since the names and locations of the quoting banks are also included. A number of interesting questions could thus be addressed that earlier data did not permit. The FXXF data did not, however, leave much room for direct testing of theory since they provide no direct measures of quantity (order flow) — and quantity's role in determining price is central to microstructure theory. As quantity data became available, more direct tests became possible.<sup>3</sup>

This paper extends earlier work by answering a number of key quantity-dependent questions. Answering some of these is rather straightforward. For example, how profitable is dealing in FX? And how rapidly do FX dealers dispose of risky inventory compared to, say, NYSE equity specialists? Other questions are less straightforward, requiring some methodological progress. For example, what share of dealer profits come from speculation versus intermediation? To answer this, one needs a method for measuring dealer speculation over time. This paper introduces a method for doing this.

Our answers to the more straightforward questions are striking. First consider dealer profits: the dealer we track averages \$100,000 profit per day (on volume of \$1 billion per day). By comparison, equity dealers average about \$10,000 profit per day (on volume of roughly \$10

---

<sup>1</sup> See for example Grammatikos and Saunders (1986).

<sup>2</sup> See Goodhart and Figliuoli (1991) and Bollerslev and Domowitz (1993), among many others. At the daily frequency, early work includes Glassman (1987) and Bossaerts and Hillion (1991).

<sup>3</sup> See Lyons (1995), Goodhart, Ito, and Payne (1996), and Yao (1996). For an emerging experimental literature see Flood et al. (1996). For work on information networks embedded in FX trading technologies see Zaheer and Zaheer (1995).

million per day).<sup>4</sup> Next, consider the pace of inventory management: the half-life for the dealer we track is only 10 minutes. This is remarkably short relative to half-lives for equity specialists of one week [Madhavan and Smidt (1993)]. Inventory theory is clearly essential to understanding trading in this market. Though striking, these results also present a challenge: Why do equity and FX markets look so different in these dimensions? That is a issue this paper raises, but does not resolve.

As stated above, addressing some of the less straightforward questions — e.g., the share of profits from speculation — requires some methodological progress. For perspective we turn to the relevant theory. Note that microstructure theory is composed of two distinct branches. The earlier branch of inventory models began in the 1970s.<sup>5</sup> The later branch of information models developed in the 1980s.<sup>6</sup> Underlying assumptions kept these branches distinct: inventory models include no information asymmetries, and information models include no inventory costs (e.g., no risk aversion). Consequently, a dealer's position in inventory models is purely nonspeculative, while in information models it is purely speculative. In the data, of course, positions include both speculative and nonspeculative components. Disentangling them is an important step for empirical work since theory is clear that the respective driving forces are different.

The methodological progress we make here is our means of disentangling speculative and nonspeculative positions. We are certainly not the first to effect this decomposition. The difference is that our approach draws from economics rather than being purely statistical [see Madhavan and Smidt (1993) and Hasbrouck and Sofianos (1993)]. Specifically, our method identifies the dealer's speculative position by projecting his total position on an information set, one that contains non-public information useful for speculation. A second, crucial ingredient is that this information set is selected such that its contents are orthogonal to the dealer's nonspeculative position. This instrumental-variables approach provides the identification we need.

In general, the information driving speculative positions comes from two key sources: incoming orders and incoming quotes. To drive speculation, it is important that this information not be public information, which it is not in this case: quantity and price information from FX

---

<sup>4</sup> We use the numbers from Hansch et al. (1994) for the London Stock Exchange because, unlike the NYSE, this is a pure dealership market and therefore more comparable to FX. They find dealers average a profit of roughly 10 basis points on each transaction. Though the authors do not provide an average turnover by dealer, they do provide data that allow a rough estimate. The average daily turnover for FTSE-100 stocks is about \$10 million (£6.9 million). This total turnover is divided among dealers, but active dealers make markets in many stocks. Given the market shares the authors report for the more active dealers, and given the number of stocks in which each makes markets, the estimated average turnover of \$10 million per dealer is about right.

<sup>5</sup> See Garman (1976), Amihud and Mendelson (1980), Ho and Stoll (1983), and O'Hara and Oldfield (1986), among others.

<sup>6</sup> See Copeland and Galai (1983), Kyle (1985), Glosten and Milgrom (1985), and Admati and Pfleiderer (1988), among others.

dealer trades are not available to the market at large. [For more on this see Lyons (1996).] The first of these two sources, incoming orders, is the standard source of dealer information in information models. The second source, incoming quotes, is a special feature of multiple-dealer markets [see Leach and Madhavan (1993)].<sup>7</sup> In fact, our position decomposition would not be possible without this feature: in single-dealer models there is typically a single determinant of both speculative and nonspeculative positions — namely order flow — and a single determinant cannot serve as an instrument for disentangling these two position components; with multiple-dealers, however, incoming quotes provide a second determinant of speculative position, one that may be uncorrelated with non-speculative position, and can therefore serve as an instrument.

Performing the decomposition allows us to pin down the degree to which profit is due to intermediation (i.e., the spread) versus speculative gains. We find that intermediation, at least for the dealer we track, is much more important, accounting for roughly 90 percent of his profits. To this we should add, however, that our dealer is described by his peers as a “liquidity machine,” by which they mean that his style focuses on high volume at competitive prices. In addition, more than 90 percent of our dealer’s trading is with other dealers (as opposed to non-dealer customers), whereas the average interdealer share is closer to 80 percent. Thus, in terms of the source of his profits, our dealer should not be viewed as representative of all dealers in this market.<sup>8</sup>

Though many empirical papers address dealer behavior, two are especially relevant here due to their focus on quantity data. Each examines implications of microstructure models using data on equity specialist inventories. Madhavan and Smidt (1993) use their inventory data to test whether inventories revert toward a target (as predicted by the theory). Their model of the target to which inventory reverts is statistical, based on systematic patterns in residuals. They find that allowing the target to vary intensifies mean reversion. Nevertheless, the shortest half-life they estimate is one week. Hasbrouck and Sofianos (1993) also focus on the inventory series itself. And their model of target inventory is also statistical, in this case driven by inventory’s time-series properties.<sup>9</sup> Like Madhavan and Smidt (1993), they find allowing the target to vary helps account for earlier findings that inventory control is weak.

Before proceeding we want to clarify our terms, in particular our use of “position

---

<sup>7</sup> It could be argued that limit orders play a similar role within specialist markets.

<sup>8</sup> See, for example, Yao (1996), who chronicles the trading of a commercial bank dealer who receives much more customer order flow than our dealer.

<sup>9</sup> The three time-series properties of inventory highlighted by Hasbrouck and Sofianos (1993) are short-term variation, discrete long-term variation, and smooth long-term variation (page 1570). Their view is that the short-term variation reflects classic dealer behavior (nonspeculation) whereas long-term variation of both types reflects investment holdings (speculation). The position series for an FX dealer displays only short-term variation by their standards since in general the daily closing position is zero. Thus, by their logic our position series must be devoid of any speculative components. Our discussions with FX dealers convince us otherwise.

control.” Here position control refers to management of both speculative and nonspeculative positions. The more traditional term “inventory control” is narrower, referring only to management of nonspeculative positions. Since this paper’s scope spans both, we adopt a more flexible terminology.

The paper is organized as follows: The next section outlines our method for disentangling speculative and nonspeculative positions; Section II describes the data; Section III disentangles the position components and estimates dealer profits; Section IV concludes.

## I. Decomposing Speculative and Nonspeculative Position Components

### I.A. Linking Target Position to Information

Consider the equation used by Madhavan and Smidt (1993) to estimate mean reversion in dealer positions (our appendix sketches the optimizing model underlying it):

$$(1) \quad I_{t+1} - I_t = \beta(I_t - I_t^*) + \varepsilon_{t+1}$$

where  $I_t$  is the dealer's total position, and  $I_t^*$  is the dealer's target position, both at time  $t$ . The position disturbance  $\varepsilon_{t+1}$  is realized at  $t+1$  and is uncorrelated with  $I_t - I_t^*$  (a property deriving from the assumption that quotes are regret-free in the sense of Glosten and Milgrom (1985)). The difficulty in estimating  $\beta$  comes from the unobservability of  $I_t^*$ . Since most models specify order flow as the only source of variation in both  $I_t$  and  $I_t^*$ , it is impossible to identify  $I_t^*$  on the basis of order flow alone, hence the past reliance on statistical models of  $I_t^*$ .

We need to be precise about the term target position. The dealer's target position can be decomposed into a speculative demand component  $I_t^s$  and a nonspeculative component  $I_t^{*n}$ .

$$(2) \quad I_t^* = I_t^s + I_t^{*n}$$

For an FX dealer, it is reasonable to presume that  $I_t^{*n}$  equals zero. In particular, it is unclear why a non-zero nonspeculative position would aid a major-market FX dealer in the provision of immediacy (e.g., an FX dealer does not worry about stock-out as he might in a less liquid market). Also supportive is the fact that FX dealers generally close each day with a position of zero (unlike NYSE specialists). Under the assumption that  $I_t^{*n}$  equals zero, we can therefore write:

$$(3) \quad I_{t+1} - I_t = \beta(I_t - I_t^s) + \varepsilon_{t+1}$$

Note that we are not assuming the nonspeculative component of total position  $I_t$  is zero, only that the nonspeculative component of the target position is zero. Accordingly, we decompose the total position at any time  $t$  into speculative and nonspeculative components:

$$(4) \quad I_t = I_t^s + I_t^n$$

where  $I_t^n$  denotes the time  $t$  nonspeculative component.

The speculative component is the dealer's speculative demand which we model in the usual way [see O'Hara (1995), pages 156-158]. The demand for the risky asset in a competitive two-asset economy, one riskless and one risky, with utility defined over end-of-period wealth is:

$$(5) \quad I_t^s = \theta E[P_{t+1} - P_t | \Omega_t]$$

where  $P_t$  is the price of the risky asset (FX) at time  $t$ ,  $\Omega_t$  is the dealer's information set at time  $t$ , and  $\theta$  is a constant. The constant  $\theta$  depends on the coefficient of absolute risk aversion and the conditional variance of the risky asset price. For now, we assume a constant conditional variance, though we return to this issue below.

#### *I.B.. Consistent Estimation of $\beta$*

The usual difficulty in extracting  $I_t^s$  from  $I_t$  is due to the correlation between  $I_t^s$  and  $I_t^n$ . To see this, consider the canonical specialist who draws inferences about the risky asset's value from the incoming order flow. In this context, incoming order flow is driving both  $I_t^n$  and  $I_t^s$ , the latter due to order flow's information content. Consequently, order flow alone is insufficient for disentangling these two components of  $I_t$ . Instruments are needed for  $I_t^s$ : variables correlated with  $I_t^s$  but uncorrelated with  $I_t^n$ .

A richer specification of  $\Omega_t$  provides the needed instruments. Partition  $\Omega_t$  into three components: (i) a set of public signals  $\Omega_t^P$ , (ii) a set of non-public signals correlated with  $I_t^n$ , or  $\Omega_t^N$ , and (iii) a set of non-public signals orthogonal to  $I_t^n$ , or  $\Omega_t^{N\perp}$ :

$$(6) \quad \Omega_t = \{ \Omega_t^P, \Omega_t^N, \Omega_t^{N\perp} \}$$

The least-squares projection of  $I_t$  on  $\Omega_t^{N\perp}$  is thus purged of any correlation with inventory changes due to  $I_t^n$ . Let  $\hat{I}_t^s$  denote this projection. We can then consistently estimate the half-life of the nonspeculative position component from:

$$(7) \quad I_{t+1} - I_t = \beta(I_t - \hat{I}_t^s) + \varepsilon_{t+1}$$

since  $\hat{I}_t^s$  is by construction orthogonal to  $\varepsilon_{t+1}$ . This is the mean reversion we associate with inventory control in the classic models.

### *I.C. A Consistent Covariance Matrix of $\hat{\beta}$*

Though an OLS estimate of from Eq. (7) is consistent, the OLS standard errors are incorrect. To see this, define the projection error of  $\hat{I}_t^s$  as  $\eta_t$ .

$$(8) \quad \hat{I}_t^s = I_t^s + \eta_t$$

If we denote the variance of the disturbance from Eq. (3) as  $\sigma_\varepsilon^2$ , then the variance of the disturbance in Eq. (7) is  $\sigma_\varepsilon^2 + \beta^2 \sigma_\eta^2$ . The probability limit of the OLS estimate of the residual covariance matrix from Eq. (7) is  $\sigma_\varepsilon^2$ , which clearly understates the asymptotic standard errors. Pagan (1984) shows that Two-Stage Least Squares (2SLS) produces a consistent estimator of the variance of  $\hat{\beta}$ . This is the estimator we use to calculate standard errors for estimates of Eq. (7).

## **II. Data**

The dataset is an augmented version of that used in Lyons (1995), the principal addition being the non-dealt quotes we use to identify the dealer's speculative position. Here we provide a broad overview, with detail only where necessary to augment that in Lyons (1995).

The dataset consists of two linked parts, each covering the activity of a DM/\$ dealer at a major New York bank. The sample spans the five trading days of the week August 3-7, 1992. Each of the five trading days begins at 8:30 A.M. and ends on average at 1:30 P.M., Eastern Standard Time. The first part comprises the dealer's position cards. His position cards include all his trades, with transaction quantities and prices. The second part is narrower but deeper: narrower because it contains only a subset of his trades, his direct (meaning non-brokered) inter-

dealer trades; deeper because it contains all his direct quotes, irrespective of whether dealt. Figure 1 provides a diagram of the data flow through time.

### *II.A. Dealer Data: Position Cards*

Because the position cards include all of the dealer's transactions, they are sufficient for constructing the dealer's position  $I_t$  through time. An average day consists of approximately 20 cards, each with about 15 transaction entries. Each card includes the following information for every trade:

- (1) the signed quantity traded (which determines  $I_t$ )
- (2) the transaction price
- (3) the counterparty, including whether brokered.

Note that the bid/ask quotes at the time of the transaction are not included on the position sheets. Note also that each entry is not time-stamped, though the dealer does record the time at the outset of every card (to the minute), and occasionally within the card too. This part of the dataset includes 1720 transactions amounting to \$7.0 billion. Figure 2 provides a diagram of the position sheet's structure. The only dimension of an actual position sheet that is not in the diagram is the names of the counterparties. Counterparty names are considered highly confidential to this dealer's institution.

### *II.B. Dealer Data: Direct Quotes and Trades*

The second part of the dataset includes the quotes, transaction quantities, and transaction prices for all the non-brokered interdealer interactions. These data derive from bank dealing records from the Reuters Dealing 2000-1 system. This system is different from the system that produces the Reuters FFX indications used elsewhere in the literature [see, for example, Goodhart and Figliuoli (1991)]. It allows dealers to communicate quotes and trades bilaterally via computer rather than verbally over the telephone. This dealer uses Dealing 2000-1 for more than 99% of his non-brokered interdealer trades.

Each record from the Dealing 2000-1 system includes the first 5 of the following 7 variables. The last two are included only if a trade takes place.

- (1) The time the communication is initiated (to the minute, with no lag)
- (2) which of the two dealers is requesting the quote
- (3) the quote quantity
- (4) the bid quote
- (5) the ask quote
- (6) the quantity traded
- (7) the transaction price.

This part of the dataset includes 952 transactions amounting to \$4.1 billion. Though it does not include the dealer's brokered transactions, it has two key advantages over the data from the position cards: it provides quotes rather than just transaction prices and it identifies which counterparty is the aggressor. It also provides a means of error-checking the data from the position cards. Perhaps most important for our purposes, this part of the dataset provides the incoming quotes that did not generate a trade. We use these non-dealt quotes to identify the dealer's speculative position. Over our sample week, only 5 percent of incoming quotes generated a trade. (In contrast, 34 percent of outgoing quotes generated a trade, an indication that our dealer is not the average dealer in this respect.)

### *II.C. The Resulting Position Series*

Figure 3 provides a plot of the dealer's position over the sample's five trading days. Mean reversion is evident. Note too the range of the position: rarely did open positions rise above \$40 million (\$40 million is well below the *intraday* position limits imposed on senior dealers at major banks, typically in the \$100-\$150 million range). Each of the five days the dealer closed his trading day with a position of zero. (A common *overnight* limit is about \$75 million; most dealers, however, close their day at zero. Two common rationales for closing at zero are that (1) carrying an open position requires monitoring it through the evening and (2) a dealer's compara-

tive advantage in speculation is when she is seated at her desk observing order flows and quotes.) The relative width of the days is a reflection of the relative trading activity. Monday and Friday are the highest volume days.

### **III. Results: Profits and Position Components**

#### *III.A. Dealer Profits*

A summary of the dealer's profitability over the week appears in Table 1. Note that the daily volume is well over \$1 billion each day. Profits for the week total \$507,929, Friday being the most profitable day. Friday is also the highest volume day. Of course, if profits derive primarily from speculative positions, then volatility is more important than volume. Friday also has the highest volatility in the five-day sample.

We turn now to the question of whether profits tend to come from intermediation or speculation. As a first step, we impute the intermediation component of profits using an assumption that the dealer himself felt is reasonable. Specifically, we assume that the cost of laying off an incoming order is one-third of the dealer's bid-ask spread. For example, if the dealer quotes his median spread of three pips at DM 1.4750/\$ - DM 1.4753/\$ and the incoming order is a purchase of dollars at DM 1.4753/\$, then on average the dealer must provide price improvement on the bid side of DM 1.4751 to induce an offsetting customer sale. This translates into an average profit of one pip on every transaction.

The results appear in the "Profit: Spread" column of Table 1. The profits imputed to intermediation correspond quite closely to total profits: the two differ by less than 10 percent. Given the dealer himself feels the assumption driving this comparison is reasonable, intermediation appears to be a much more important source of profits. Still, to get a more complete picture, we need to examine whether his speculative positions are profitable.

#### *III.B. Position Decomposition*

Per section I, disentangling the speculative and nonspeculative position components requires instruments that are correlated with the speculative component but uncorrelated with the nonspeculative component. The instrument sets we consider rely exclusively on non-dealt,

incoming quotes (specifically, lagged changes).<sup>10</sup> These quotes are not public information, so they are arguably relevant for speculative position taking. On the other hand, they do not produce a transaction, so there is no direct relation to the nonspeculative component of our dealer's position. Of course, no direct relation does not rule out an indirect relation. One rationale for ruling out even an indirect relation is embedded in the model of Lyons (1997). In that model, past price change is orthogonal to nonspeculative position because dealers — being risk averse — set price to avoid nonspeculative positions; though nonspeculative positions still arise, optimal use of past information brings the correlation with price to zero.

Nevertheless, incoming quote changes cannot identify the speculative component  $I_t^s$  if they are uncorrelated with the dealer's total position. Column two of table 2 (rows 2-4) verifies that these instruments are indeed correlated with  $I_t$ .<sup>11</sup> The  $R^2$ 's indicate that  $\hat{I}_t^s$ , which is the projection of  $I_t$  on the information set  $\Omega$ , accounts for about one-quarter of the variation in the dealer's total position. Put differently,  $\Omega$  contains information useful to the dealer in determining his speculative demand. The first row of the table is included as a benchmark. With an empty information set, it is equivalent to assuming that target position  $\hat{I}_t^s$  equals zero. (Hence, there is no first stage projection, and the  $R^2$  is not applicable.)

Table 2 also presents estimates of the half-life of nonspeculative positions.<sup>12</sup> A half-life of ten minutes is extremely short relative to the position half-lives estimated in equity markets; for example, the lowest estimate for equity specialists in Madhavan and Smidt (1993) is 7 days. Our finding of ten minutes is clear evidence of aggressive inventory control. Of course, the FX market is not the NYSE. Why they trade so differently, however, is still an open question. Results like this reinforce the fact that the FX market is distinctive at the microstructural level.

---

<sup>10</sup> In our data set, typically incoming quotes do not elicit a transaction: more than 80 percent of incoming quotes lapse without a deal. This large share of “non-dealt” quotes suggests that dealers gather quotes because they are informative, not simply because they provide immediacy.

For lack of data, previous work has had to identify no-trade “events” as the elapsing of a given amount of time without a transaction [e.g., five minutes in Easley, Kiefer, and O'Hara (1997)]. An important strength of our data set is that it includes non-dealt events explicitly.

<sup>11</sup> The last of the three information sets — row 4 of Table 2 — includes a squared lagged price change in addition to the signed lagged price changes. Though this can only improve the first-stage fit, the orthogonality of this variable is more difficult to defend on theoretical grounds.

<sup>12</sup> Note that these regressions are estimated in transaction time rather than real time. For a discussion of this issue in the context of the FX market see Lyons (1995), page 345-346.

### III.C. Regression Diagnostics

Some regression diagnostics are in order. First, there is no evidence of serial correlation (Durbin-Watson and Ljung-Box Q tests) or heteroskedasticity (White test). A natural next step is to examine structural stability, in particular, with respect to time-of-day. To do so, we split the sample into two subsamples, early and late, where early is defined as all observations occurring before the midpoint of the trading day (median transaction time). A Wald test of the equality of  $\hat{\beta}$  across the two subsamples is rejected at the one percent level. Clearly, time-of-day is related to inventory control, which is not surprising given the dealer starts and ends each day with a zero net position.

To capture this time-of-day effect, we include an additional term in the model of Eq. (1):<sup>13</sup>

$$(9) \quad I_{t+1} - I_t = \beta_1(I_t - \hat{I}_t^s) + \beta_2(\% Day_t)(I_t - \hat{I}_t^s) + \varepsilon_{t+1}$$

where the proxy “%  $Day_t$ ” is defined over each day as (observation #, day n) divided by (total # observations, day n), with  $n=1, \dots, 5$ .<sup>14</sup> The results appear in Table 3. Note that the time-independent component of mean reversion is insignificantly different from zero ( $\hat{\beta}_1$ ). In contrast, the time-dependent component,  $\hat{\beta}_2$ , now does all the work. The estimates imply that the time-dependent component of mean reversion moves from zero at the start of each day to about -0.3 at the end. The implied half-lives at three points within the dealer's day appear in the last columns. On the whole, this time profile squares nicely with the fact that FX dealers typically close their positions at the end of each trading day.

Another important consideration is the impact of price volatility on the pace of mean reversion. It is not a priori clear, however, how volatility should affect mean reversion. On one hand, the higher volatility, the greater the benefit of inventory control. On the other hand, the higher volatility, the greater the cost of inventory control (spreads at brokers widen, as do direct spreads quoted by other dealers; for some evidence from indicative quotes, see Bessembinder

---

<sup>13</sup> Note that while Eq. (1) is derived from an optimizing model, the same cannot be said of our adjustments to it. The technical challenge of doing so is formidable, and the payoff, in our judgment, marginal.

<sup>14</sup> We refer to this as a proxy because the denominator—total transactions for the day—is not known to the dealer in advance.

(1994) and Bollerslev and Melvin (1994)). To capture this, we estimate the following linear model:

$$(10) \quad I_{t+1} - I_t = \beta_1(\% Day_t)(I_t - \hat{I}_t^s) + \beta_2(HighVol_t)(\% Day_t)(I_t - \hat{I}_t^s) + \varepsilon_{t+1}$$

where the dummy variable “*HighVol<sub>t</sub>*” is equal to one if  $(\Delta P_{t-1})^2$  is above its median, zero otherwise. The first regressor is the same as the second regressor in the model of Eq. (9) (recall that the first regressor in Eq. (9) produced an insignificant coefficient). The results appear in Table 4. The coefficient  $\hat{\beta}_2$  is significant at the five percent level but not at one percent. The size and sign of  $\hat{\beta}_2$  imply that mean reversion is lower when volatility is high. This suggests that when volatility increases, the added cost of inventory control dominates the added benefit. The implied half-lives at high and low volatility appear in the last columns (calculated at %Day=50%). The half-life is about twice as long with high volatility than with low volatility.

#### *III.D. Profits From Speculative Component*

Though section III.A. provides estimates of profit from intermediation, our measure of the dealer’s speculative position allows us to estimate speculative profit directly. We proceed in two stages. First, we determine whether the dealer’s speculative position forecasts future price changes. In particular, we test whether speculative position Granger causes price changes. Though not a sufficient condition for profitable speculation, the result provides useful perspective. Then, in stage two, we mark the speculative position to market using the midpoint of the incoming quote closest to the measured speculative position change. The accumulated changes in net value provide an estimate of speculative profit. (Unlike our exact measure of total profit in Table 1, it is not possible to calculate speculative profit exactly since our measured changes in speculative position are not linked one-for-one with market transactions.)

On both counts, our results do not evince speculative profits. We find no evidence of Granger causality running from speculative position to prices. Further, even using quote midpoints, the speculative position does not generate profits; in fact, over the week the dealer suffered negligible losses on the speculative position. It is important to keep in mind, however,

that we are working with only one week of data, which may not be representative with respect to speculative profits. To be sure, speculative profits are much more volatile than profits from intermediation.

#### **IV. Conclusions**

We are able to address a number of questions that have not yet been examined in empirical work on FX microstructure. And our answers to these questions are striking. First the dealer we track is quite profitable: he averages \$100,000 profit per day (on volume of \$1 billion per day). By comparison, equity dealers average about \$10,000 profit per day (on volume of roughly \$10 million per day). Second, we also find quite rapid inventory management: the half-life for the dealer we track is only 10 minutes. This is remarkably short compared to half-lives for equity specialists of one week. Why do equity and FX markets generate such different results at the microstructural level? This is an important topic for future research.

The paper also makes methodological progress. Specifically, we develop a method for measuring speculative position-taking over time. Unlike past — purely statistical — methods for doing this, our measure is based on the information flows that drive speculation. We use this measure of speculative position to pin down the extent to which dealer profits derive from speculation. In the end we find that intermediation is much more important source of profits, at least for the dealer we track over the week we are able to track him.

## Appendix

### Sketch of Madhavan and Smidt (1993) Derivation of Eq. (1)

Agents: (1) Risk neutral informed trader  
 (2) Liquidity traders  
 (3) Risk neutral dealer, with inventory carrying costs (proportional to the variance of wealth)

Preferences: The dealer and informed trader maximize the expected value of terminal wealth. Liquidity trader preferences are not modeled (their trading can be generated as optimizing in the context of endowment shocks).

Market: There is a single risky asset that trades in a series of auctions. There is a constant probability  $(1-\rho)$  that a given time  $\tau$  will be the liquidation time. The asset's fundamental value follows a random walk. This value is known by the informed trader.

Notation:  $v_t$  : fundamental value at time  $t$   
 $\mu_t$  : dealer's conditional expectation of  $v_t$   
 $\Omega_t$  : dealer's information set  
 $W_t$  : dealer's wealth  
 $I_t$  : dealer's position  
 $P_t$  : price  
 $z_t$  : excess demand, a function of  $P_t$   
 $Q_t$  : demand of informed trader, a function of  $P_t$   
 $X_t$  : demand of liquidity traders (aggregate), not a function of  $P_t$   
 $D_t$  : intercept of demand schedule, informed plus uninformed

The informed trader's demand is proportional to the gap between fundamental value and price.

$$(A1) \quad Q_t(v_t, P_t) = \delta(v_t - P_t)$$

It is helpful to define the intercept of the excess demand schedule as the following:

$$(A2) \quad D_t = \delta v_t + X_t$$

When divided by  $\delta$ , this variable  $D_t$  provides an unbiased signal of the fundamental value  $v_t$  (since  $X_t$  is mean zero).

$$(A3) \quad D_t/\delta = v_t + X_t/\delta$$

Now, the updating equation can be expressed as a weighted average of this signal and the dealer's prior expectation.

$$(A4) \quad \mu_t = \phi(D_t/\delta) + (1-\phi)\mu_{t-1}$$

This evolution of the dealer's expectation can then be integrated with the dealer's stochastic programming problem.

$$(A5) \quad \underset{P_t}{\text{Max}} \quad \sum_{j=1}^{\infty} E[W_j | \Omega_j] \Pr[j = \tau]$$

The solution to this stochastic programming problem yields Eq. (1) in the text:

$$(1) \quad I_{t+1} - I_t = \beta(I_t - I_t^*) + \varepsilon_{t+1}$$

where the error term in the regression corresponds to the dealer's estimate of liquidity-trader demand  $X_t$  in the model:

$$(A6) \quad \varepsilon_{t+1} \equiv -\left(\frac{1+\beta}{2}\right)E[X_t | \Omega_t]$$

## References

- Admati, A. and Pfleiderer, P. (1988) A theory of intraday patterns: Volume and price variability. *Review of Financial Studies* **1**, 3-40.
- Amihud, Y. and Mendelson, H. (1980) Dealership market: Market making with inventory. *Journal of Financial Economics* **8**, 31-53.
- Ammer, J. and Brunner, A. (1994) Are banks market timers or market makers? Explaining foreign exchange trading profits. International Finance Discussion Paper #484, Board of Governors of the Federal Reserve System.
- Bessembinder, H. (1994) Bid-ask spreads in the interbank foreign exchange market. *Journal of Financial Economics* **35**, 317-348.
- Bollerslev, T. and Domowitz, I. (1993) Trading patterns and prices in the interbank foreign exchange market. *Journal of Finance* **48**, 1421-1444.
- Bollerslev, T. and Melvin, M. (1994) Bid-ask spreads and volatility in the foreign exchange market: An empirical analysis. *Journal of International Economics* **36**, 355-372.
- Bossaerts, P. and Hillion, P. (1991) Market microstructure effects of government intervention in the foreign exchange market. *Review of Financial Studies* **4**, 513-541.
- Copeland, T. and Galai, D. (1983) Information effects on the bid-ask spread. *Journal of Finance* **38**, 1457-1469.
- Easley, D., Kiefer, N. and O'Hara, M. (1997) One day in the life of a very common stock. *Review of Financial Studies*, forthcoming.
- Flood, M. (1992) Market structure and inefficiency in the foreign exchange market. *Journal of International Money and Finance* **13**, 131-158.
- Flood, M., Huisman, R., Koedijk, K. and Mahieu, R. (1996) Price discovery in multiple-dealer financial markets: The effect of pre-trade transparency. Concordia University typescript, December
- Garman, M. (1976) Market microstructure. *Journal of Financial Economics* **3**, 257-275.
- Glassman, D. (1987) Exchange rate risk and transactions costs: Evidence from bid-ask spreads. *Journal of International Money and Finance* **6**, 479-490.
- Glosten, L. and Milgrom, P. (1985) Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents. *Journal of Financial Economics* **14**, 71-100.
- Goodhart, C. and Figliuoli, L. (1991) Every minute counts in financial markets. *Journal of International Money and Finance* **10**, 23-52.

- Goodhart, C., Ito, T. and Payne, R. (1996) One day in June 1993: A study of the working of the Reuters 2000-2 electronic foreign exchange trading system. In *The Microstructure of Foreign Exchange Markets*, eds. J. Frankel, G. Galli and A. Giovannini, pp. 107-179. University of Chicago Press, Chicago, IL.
- Grammatikos, T. and Saunders, A. (1986) Futures price variability: A test of maturity and volume effects. *Journal of Business* **59**, 319-330.
- Hansch, O., Naik, N. and Viswanathan, S. (1994) Trading profits, inventory control and market share in a competitive dealer market. Duke University typescript, November.
- Hasbrouck, J. (1988) Trades, quotes, inventories, and information. *Journal of Financial Economics* **22**, 229-252.
- Hasbrouck, J. and Sofianos, G. (1993) The trades of market makers: An empirical analysis of NYSE specialists. *Journal of Finance* **48**, 1565-1593.
- Ho, T. and Stoll, H. (1983) The dynamics of dealer markets under competition. *Journal of Finance* **38**, 1053-1074.
- Kyle, P. (1985) Continuous auctions and insider trading. *Econometrica* **53**, 1315-1335.
- Leach, J. and Madhavan, A. (1993) Price experimentation and security market structure. *Review of Financial Studies* **6**, 375-404.
- Lyons, R. (1995) Tests of microstructural hypotheses in the foreign exchange market. *Journal of Financial Economics* **39**, 321-351.
- Lyons, R. (1996) Optimal transparency in a dealer market with an application to foreign exchange. *Journal of Financial Intermediation* **5**, 225-254.
- Lyons, R. (1997) A simultaneous trade model of the foreign exchange hot potato. *Journal of International Economics* **42**, 275-298.
- Madhavan, A. and Smidt, S. (1991) A bayesian model of intraday specialist pricing. *Journal of Financial Economics* **30**, 99-134.
- Madhavan, A. and Smidt, S. (1993) An analysis of changes in specialist inventories and quotations. *Journal of Finance* **48**, 1595-1628.
- O'Hara, M. (1995) *Market Microstructure Theory*. Blackwell Business, Cambridge MA.
- O'Hara, M. and Oldfield, G. (1986) The microeconomics of market making. *Journal of Financial and Quantitative Analysis* **21**, 361-376.
- Pagan, A. (1984) Econometric issues in the analysis of regressions with generated regressors. *International Economic Review* **25**, 221-247.
- Yao, J. (1996) Market making in the interbank foreign exchange market, New York University typescript, December.

Zaheer, A. and Zaheer, S. (1995) Catching the wave: Alertness, responsiveness, and market influence in global electronic networks. University of Minnesota typescript, December.

**Table 1**

Summary of DM/\$ dealer's trading and profits  
from Monday, August 3 to Friday, August 7, 1992.

The "\$ Profit: Spread" column reports the profit the dealer would have realized if he had cleared one-third of his spread on every transaction. It is calculated as the dollar volume times one-third the median spread he quoted in the sample (median spread = DM 0.0003/\$), divided by the average DM/\$ rate over the sample (DM 1.475/\$).

---

	<b><u>Transactions</u></b>	<b><u>Volume (mil)</u></b>	<b><u>Profit: Actual</u></b>	<b><u>Profit: Spread</u></b>
<b>Monday</b>	333	\$ 1,403	\$ 124,253	\$ 95,101
<b>Tuesday</b>	301	\$ 1,105	\$ 39,273	\$ 74,933
<b>Wednesday</b>	300	\$ 1,157	\$ 78,575	\$ 78,447
<b>Thursday</b>	328	\$ 1,338	\$ 67,316	\$ 90,717
<b>Friday</b>	458	\$ 1,966	\$ 198,512	\$ 133,298
<b>Total</b>	1,720	\$ 6,969	\$ 507,929	\$ 472,496

---

**Table 2**

## Mean Reversion in Dealer Positions

$$I_{t+1} - I_t = \beta(I_t - \hat{I}_t^s) + \varepsilon_{t+1}$$

**Definitions:**  $I_{t+1}$  is the dealer's position following the incoming trade at time  $t+1$ , where each incoming trade in the sample defines a single period (i.e., transaction time).  $\hat{I}_t^s$  is an estimate of the dealer's speculative position at time  $t$ , defined as a projection on an information set  $\Omega$ . The information sets we consider appear in column one.  $\Delta P_{t,5}$  is the cumulative change in the nondealt quotes received by the dealer over the period spanning five incoming trades prior to the incoming trade at  $t$ . Note that the first of the information sets implies  $\hat{I}_t^s = 0$ . The error term  $\varepsilon_{t+1}$  represents disturbances to the dealer's position from the provision of immediacy.

**Estimation:** First, we estimate  $\hat{I}_t^s$  as the OLS projection of  $I_t$  on  $\Omega$ . Then, we use  $\hat{I}_t^s$  to estimate  $\beta$  from the above equation using OLS. Column two presents the  $R^2$ 's from the projection of  $I_t$  on  $\Omega$ . In rows 2-4, the t-statistics in parentheses are calculated from 2SLS standard errors to account for the use of a generated regressor. The implied half-life is calculated as the mean intertransaction time (1.8 minutes) times the half-life in transactions implied by  $\hat{\beta}$  [defined as  $-\ln(2)/\ln(1+\hat{\beta})$ ].

**Sample:** 838 observations from Monday, August 3 to Friday, August 7, 1992.

Information set	$R^2$ of projection	$\hat{\beta}$	Implied half-life
$\Omega = \{ \}$	N.A.	-0.10 (-5.63)	12 minutes
$\Omega = \{ \Delta P_{t,5} \}$	0.24	-0.12 (-5.73)	10 minutes
$\Omega = \{ \Delta P_{t,5}, \Delta P_{t-5,5} \}$	0.26	-0.12 (-5.64)	10 minutes
$\Omega = \{ \Delta P_{t,5}, \Delta P_{t-5,5}, (\Delta P_{t,5})^2 \}$	0.29	-0.13 (-5.97)	9 minutes

**Table 3**

## Time-Dependent Mean Reversion in Dealer Positions

$$I_{t+1} - I_t = \beta_1(I_t - \hat{I}_t^s) + \beta_2(\% Day_t)(I_t - \hat{I}_t^s) + \varepsilon_{t+1}$$

**Definitions:**  $I_{t+1}$  is the dealer's position following the incoming trade at time  $t+1$ , where each incoming trade in the sample defines a single period (i.e., transaction time).  $\hat{I}_t^s$  is an estimate of the dealer's speculative position at time  $t$ , defined as a projection on an information set  $\Omega$ . The information sets we consider appear in column one.  $\Delta P_{t,5}$  is the cumulative change in the nondealt quotes received by the dealer over the period spanning five incoming trades prior to the incoming trade at  $t$ . The variable “%  $Day_t$ ” is defined over each day as (observation #, day  $n$ )/(total # observations, day  $n$ ),  $n=1,\dots,5$ . The error term  $\varepsilon_{t+1}$  represents disturbances to the dealer's position from the provision of immediacy.

**Estimation:** First, we estimate  $\hat{I}_t^s$  as the OLS projection of  $I_t$  on  $\Omega$ . Then, we use  $\hat{I}_t^s$  to estimate  $\beta$  from the above equation using OLS. The t-statistics in parentheses are calculated from 2SLS standard errors to account for the use of a generated regressor. The three half-life columns report the implied half-life when the dealer's day (% $Day$ ) is 25% done, 50% done, and 75% done. The implied half-life is calculated as the mean intertransaction time (1.8 minutes), times the half-life in transactions implied by  $\hat{\beta}_2$  [defined as  $-\ln(2)/\ln(1+\hat{\beta}_2)$ ], times “%  $Day_t$ ”.

**Sample:** 838 observations from Monday, August 3 to Friday, August 7, 1992.

Information set	$\hat{\beta}_1$	$\hat{\beta}_2$	Implied half-life			
			25%	50%	75%	
$\Omega = \{\Delta P_{t,5}\}$	0.02 (0.29)	-0.30 (-3.06)	16	8	5	minutes
$\Omega = \{\Delta P_{t,5}, \Delta P_{t-5,5}\}$	0.03 (0.46)	-0.32 (-3.13)	16	7	5	minutes

**Table 4**

The Impact of Volatility on Time-Dependent  
Mean Reversion in Dealer Positions

$$I_{t+1} - I_t = \beta_1(\% Day_t)(I_t - \hat{I}_t^s) + \beta_2(HighVol_t)(\% Day_t)(I_t - \hat{I}_t^s) + \varepsilon_{t+1}$$

**Definitions:**  $I_{t+1}$  is the dealer's position following the incoming trade at time  $t+1$ , where each incoming trade in the sample defines a single period (i.e., transaction time).  $\hat{I}_t^s$  is an estimate of the dealer's speculative position at time  $t$ , defined as a projection on an information set  $\Omega$ . The information sets we consider appear in column one.  $\Delta P_{t,5}$  is the cumulative change in the nondealt quotes received by the dealer over the period spanning five incoming trades prior to the incoming trade at  $t$ . The variable “% Day<sub>*t*</sub>” is defined over each day as (observation #, day  $n$ )/(total # observations, day  $n$ ),  $n=1,\dots,5$ . The dummy variable “HighVol<sub>*t*</sub>” is equal to one if  $(\Delta P_{t-1})^2$  is higher than its median, zero otherwise. The error term  $\varepsilon_{t+1}$  represents disturbances to the dealer's position from the provision of immediacy.

**Estimation:** First, we estimate  $\hat{I}_t^s$  as the OLS projection of  $I_t$  on  $\Omega$ . Then, we use  $\hat{I}_t^s$  to estimate  $\beta$  from the above equation using OLS. The  $t$ -statistics in parentheses are calculated from 2SLS standard errors to account for the use of a generated regressor. The two half-life columns report the implied half-life when the dealer's day is 50% done and volatility is high ( $HighVol_t=1$ ) and low ( $HighVol_t=0$ ). The implied half-life is calculated as the mean inter-transaction time (1.8 minutes), times the half-life in transactions implied by  $\hat{\beta}_1$  and  $\hat{\beta}_2$  [defined as  $-\ln(2)/\ln(1+\hat{\beta})$ , where  $\hat{\beta}=\hat{\beta}_1$  with low volatility and  $\hat{\beta}=\hat{\beta}_1+\hat{\beta}_2$  with high volatility], times 0.5 to account for “% Day<sub>*t*</sub>”.


**Sample:** 838 observations from Monday, August 3 to Friday, August 7, 1992.

Information set	$\hat{\beta}_1$	$\hat{\beta}_2$	Implied half-life		
			High Vol.	Low Vol.	
$\Omega = \{\Delta P_{t,5}\}$	-0.36 (-4.56)	0.21 (2.09)	13	6	minutes
$\Omega = \{\Delta P_{t,5}, \Delta P_{t-5,5}\}$	-0.37 (-4.72)	0.21 (2.11)	13	6	minutes

**Figure 1**

Diagram of data structure

Definitions:  $Q^o$  is an outgoing interdealer quote (i.e., a quote made) and, if the quote is hit,  $T^i$  is the incoming dealer trade.  $Q^i$  is an incoming interdealer quote (i.e., a quote received) and, if the quote is hit,  $T^o$  is the outgoing trade.  $T^b$  is a brokered interdealer trade. Brokered trades do not align vertically with a quote because the data for brokered trades come from the dealer position sheets, and the broker-advertised quotes at the time of the transaction are not recorded. “ “ appears whenever a trade occurs; “ | ” appears whenever a non-dealt quote occurs. ■ The disjoint segment below the top time-line represents the dealer's position over the same interval. The time-line at the bottom clarifies the timing relevant to the regression subscripts; namely, periods are defined over incoming transactions.



**Figure 2**

Diagram of position sheet structure,  
first fourteen trades on Monday, August 3, 1992

The dealer's position sheets provide the dealer with a running record of his net position and the approximate cost of that position. This dealer fills it in by hand as he trades. Each sheet (page) covers about fifteen transactions. The "Position" column accumulates the individual trades in the "Trade" column. Quantities are in millions of dollars. A positive quantity in the Trade column corresponds to a purchase of dollars. A positive quantity in the Position column corresponds to a net long dollar position. The "Trade Rate" column records the exchange rate for the trade, in deutschemarks per dollar. The "Position Rate" column records the dealer's estimate of the average rate at which he acquired his position. The Position and Position Rate are not calculated after every trade due to time constraints. The "Source" column reports whether the trade is direct over the Reuters Dealing 2000-1 system (r=Reuters) or brokered (b=Broker).

Trade date: 8/3 Value date: 8/5					
Position	Position rate	Trade	Trade rate	Source	Time
		1	1.4794	r	8:30
		2	1.4797	r	
3	1.4796				
		28	1.4795	r	
		-10	1.4797	r	
		-10	1.4797	b	
		-10	1.4797	r	
		-3	1.4797	b	
-2	1.4797				
		0.5	1.4794	r	
		0.75	1.4790	r	
		3	1.4791	r	
2	1.4791				
		-10	1.4797	r	
-8	1.4797				
		2	1.4799	b	
-6	1.4797				
		5	1.4805	b	
		-7	1.4810	r	
-8	1.4808				

**Figure 3**

Plot of DM/\$ dealer's position  $I_t$  (in millions of dollars)  
from Monday, August 3 to Friday, August 7, 1992.

These positions are constructed from the dealer's position sheets. Short dollar positions correspond to long DM positions, and vice versa. The vertical lines designate the four overnight periods through which the dealer did not trade. The relative width of the days is a reflection of the relative trading activity. The two horizontal lines bracketing zero are two standard deviations from the mean long position of two million dollars.

**Figure 4**

Plot of DM/\$ dealer's total position  $I_t$  and  
speculative position  $\hat{I}_t^s$  (in millions of dollars)  
on Friday from 9AM to 12PM, August 7, 1992.

The total position  $I_t$  is constructed from the dealer's position sheets. Short dollar positions correspond to long DM positions, and vice versa. The speculative component  $\hat{I}_t^s$  is estimated as the projection of  $I_t$  on the information set  $\Omega$ , with  $\Omega = \{\Delta P_{t,5}\}$ , where  $\Delta P_{t,5}$  is the cumulative change in the nondealt quotes received by the dealer over the period spanning five incoming trades prior to the incoming trade at  $t$ .