



ELSEVIER

Journal of International Money and Finance
17 (1998) 97–115

Journal of
International
Money
and Finance

Profits and position control: a week of FX dealing

Richard K. Lyons^{*,†}

*U.C. Berkeley and NBER, Haas School of Business, U.C. Berkeley, Berkeley, CA 94720-1900,
USA*

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 10 min, 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. © 1998 Elsevier Science Ltd. All rights reserved.

JEL classifications: G15; F31

Empirical work on FX microstructure is still in its early stages. The earliest work used futures data since that was available at high frequencies.¹ 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 1990s with the availability of quotes on an

*Tel.: +1 510 6421059; fax: +1 510 6424700; 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.

¹See for example Grammatikos and Saunders (1986).

intraday basis (specifically, the indicative quotes from Reuters called FAFX).² 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 FAFX 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.³

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 million per day).⁴ Next, consider the pace of inventory management: the half-life for the dealer we track is only 10 min. This is remarkably short relative to half-lives for equity specialists of 1 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 an 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

²See 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).

³See Lyons (1995), Goodhart et al. (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).

⁴We 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.

in the 1970s.⁵ The later branch of information models developed in the 1980s.⁶ 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 non-speculative, while in information models it is purely speculative. In the data, of course, positions include both speculative and non-speculative 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 non-speculative 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; 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 non-speculative 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 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).⁷ 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 non-speculative 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% 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% of our dealer's trading is with other dealers (as opposed to non-dealer customers),

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

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

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

whereas the average interdealer share is closer to 80%. Thus, in terms of the source of his profits, our dealer should not be viewed as representative of all dealers in this market.⁸

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 theory predicts). 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 1 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.⁹ 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 control'. Here position control refers to management of both speculative and non-speculative positions. The more traditional term 'inventory control' is narrower, referring only to management of non-speculative positions. Since this paper's scope spans both, we adopt a more flexible terminology.

The paper is organized as follows: Sec. 1 outlines our method for disentangling speculative and non-speculative positions; Sec. 2 describes the data; Sec. 3 disentangles the position components and estimates dealer profits; Sec. 4 concludes.

1. Decomposing speculative and non-speculative position components

1.1. Linking target position to information

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

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

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

⁹The 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 (p. 1570). Their view is that the short-term variation reflects classic dealer behavior (non-speculation) 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.

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 ε_{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 β 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 non-speculative component I_t^{*n} .

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

For an FX dealer, it is reasonable to presume that I_t^{*n} equals zero. In particular, it is unclear why a non-zero non-speculative 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:

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

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

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

where I_t^n denotes the time t non-speculative component.

The speculative component is the dealer's speculative demand which we model in the usual way (see O'Hara, 1995, pp. 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:

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

where P_t is the price of the risky asset (FX) at time t , Ω_t is the dealer's information set at time t , and θ is a constant. The constant θ 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.

1.2. Consistent estimation of β

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

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

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 non-speculative position component from:

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

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

1.3. A consistent covariance matrix of $\hat{\beta}$

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

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

If we denote the variance of the disturbance from (Eq. (3)) as σ_ε^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 σ_ε^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)).

2. 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 5 trading days of the week August 3–7, 1992. Each of the 5 trading days begins at 08:30 h and ends on average at 13:30 h, Eastern Standard Time. The first part comprises the dealer's position

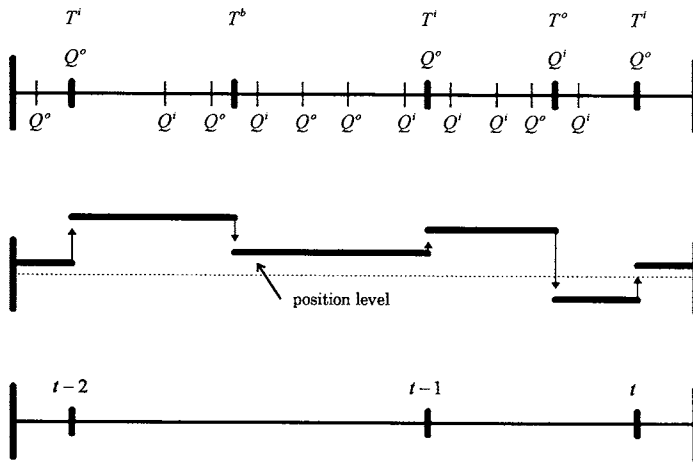


Fig. 1. Diagram of data structure. Definitions: Q^o is an outgoing interdealer quote (i.e. a quote made) and, if the quote is hit, T^i is the incoming dealer trade. Q^i is an incoming interdealer quote (i.e. a quote received) and, if the quote is hit, T^o is the outgoing trade. T^b is a brokered interdealer 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 subsamples; namely, periods are defined over incoming transactions.

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

2.1. 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

Table 1

Diagram of position sheet structure, first 14 trades on Monday, August 3, 1992

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		
		-10	1.4797		
		-10	1.4797		
		-10	1.4797		
		-3	1.4797		
-2	1.4797				
		0.5	1.4794		
		0.75	1.4790		
		3	1.4791		
2	1.4791				
		-10	1.4797		
-8	1.4797				
		2	1.4799		
-6	1.4797				
		5	1.4805		
		-7	1.4810		
-8	1.4808				

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 15 transactions. The 'Position' column accumulates the individual trades in the 'Trade' column. Quantities are in millions 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 Deutschmarks 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).

amounting to \$7.0 billion. Table 1 provides a diagram of the position sheet's structure. The only part 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.

2.2. 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 FAFX 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 five of the following seven 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% of incoming quotes generated a trade. (In contrast, 34% of outgoing quotes generated a trade, an indication that our dealer is not the average dealer in this respect.)

2.3. The resulting position series

Figure 2 provides a plot of the dealer's position over the sample's 5 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 5 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 comparative 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.

3. Results: profits and position components

3.1. Dealer profits

A summary of the dealer's profitability over the week appears in Table 2. 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




Fig. 2. 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.

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 5-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 2. The profits imputed to intermediation correspond quite closely to total profits: the two differ by less than 10%. 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.

3.2. Position decomposition

Per Sec. 1, disentangling the speculative and non-speculative position compo-

Table 2
Summary of DM/\$ dealer's trading and profits from Monday, August 3 to Friday, August 7, 1992

	Transactions	Volume (mil)	Profit: actual	Profit: spread
Monday	333	\$1403	\$124 253	\$95 101
Tuesday	301	\$1105	\$39 273	\$74 933
Wednesday	300	\$1157	\$78 575	\$78 447
Thursday	328	\$1338	\$67 316	\$90 717
Friday	458	\$1966	\$198 512	\$133 298
Total	1720	\$6969	\$507 929	\$472 496

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/\$).

nents requires instruments that are correlated with the speculative component but uncorrelated with the non-speculative component. The instrument sets we consider rely exclusively on non-dealt, incoming quotes (specifically, lagged changes).¹⁰ 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 non-speculative 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 non-speculative position because dealers — being risk averse — set price to avoid non-speculative positions; though non-speculative 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 3 (rows 2–4) verifies that these instruments are indeed correlated with I_t .¹¹ The R^2 s indicate that \hat{I}_t^s , which is the projection of I on the information set Ω , accounts for about one-quarter of the variation in the dealer's total position. Put differently, Ω 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

¹⁰ In our data set, typically incoming quotes do not elicit a transaction: more than 80% 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. 5 min in Easley et al., 1997). An important strength of our dataset is that it includes non-dealt events explicitly.

¹¹ The last of the three information sets — row 4 of Table 3 — 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.

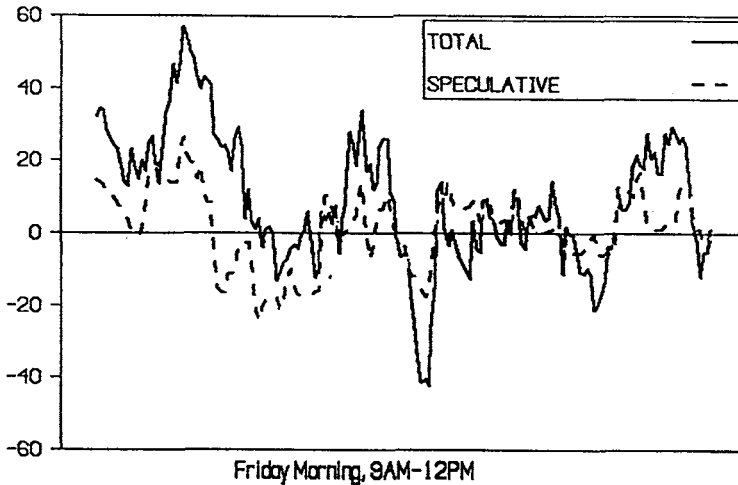


Fig. 3. Plot of DM/\$ dealer's total position I_t and speculative position \hat{I}_t^s (in millions of dollars) on Friday from 09:00 to 12:00 h, August 7, 1992. The total position I_t is constructed from the dealer's position sheets. Short dollar positions correspond to long DM positions, and vice versa. The speculative component \hat{I}_t^s is estimated as the projection of I_t on the information set Ω , with $\Omega = \{\Delta P_{t,5}\}$, where $\Delta P_{t,5}$ is the cumulative change in the non-dealt quotes received by the dealer over the period spanning five incoming trades prior to the incoming trade at t .

no first stage projection, and the R^2 is not applicable.) Figure 3 plots the speculative demand and total position corresponding to row 2 of the table (we include a three-hour interval for greatest resolution).

Table 3 also presents estimates of the half-life of non-speculative positions.¹² A half-life of 10 min is extremely short relative to 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 10 min 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.¹³

3.3. 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

¹²Note 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, pp. 345–346.

¹³Though not as short as FX, other multiple-dealer markets also exhibit shorter half-lives than the NYSE. See Reiss and Werner (1997) and Hansch et al. (1994).

Table 3

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

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

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 Ω . The information sets we consider appear in column 1. $\Delta P_{t,5}$ is the cumulative change in the non-dealt 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 ε_{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 Ω . Then, we use \hat{I}_t^s to estimate β from the above equation using OLS. Column 2 presents the R^2 s from the projection of I_t on Ω . 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 min) 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.

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 1% 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)):¹⁴

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

where the proxy ' $\%Day_t$ ' is defined over each day as (observation #, day n) divided by (total # observations, day n), with $n = 1, \dots, 5$.¹⁵ The results appear in Table 4. 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

¹⁴Note 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.

¹⁵We refer to this as a proxy because the denominator — total transactions for the day — is not known to the dealer in advance.

Table 4
Time-dependent mean reversion in dealer positions

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

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 min
$\Omega = \{\Delta P_{t,5}, \Delta P_{t-5,5}\}$	0.03 (0.46)	-0.32 (-3.13)	16	7	5 min

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 Ω . The information sets we consider appear in column 1. $\Delta P_{t,5}$ is the cumulative change in the non-dealt 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 ε_{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 Ω . Then, we use \hat{I}_t^s to estimate β 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 min), 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.

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

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

where the dummy variable ' $HighVol_t$ ' 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 5. The coefficient $\hat{\beta}_2$ is significant at the 5% level but not at 1%. The size and sign of $\hat{\beta}_2$ imply that mean reversion is lower when volatility is high. This suggests that when volatility

Table 5

The impact of volatility on time-dependent mean reversion in dealer positions

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

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 min
$\Omega = \{\Delta P_{t,5}, \Delta P_{t-5,5}\}$	-0.37 (-4.72)	0.21 (2.11)	13	6 min

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 Ω . The information sets we consider appear in column 1. $\Delta P_{t,5}$ is the cumulative change in the non-dealt 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 dummy variable ' $HighVol_t$ ' is equal to one if $(\Delta P_{t-1})^2$ is higher than its median, zero otherwise. The error term ε_{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 Ω . Then, we use \hat{I}_t^s to estimate β 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 intertransaction time (1.8 min), 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 ' $\%Day_t$ '.

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

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.

3.4. Profits from speculative component

Though Sec. 3.1 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 2, 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 1 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. Nevertheless, our findings are consistent with those from the much lower frequency analysis of Ammer and Brunner (1997). Nevertheless, our findings are consistent with those from the much lower-frequency analysis of Ammer and Brunner (1997).

4. 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 min. This is remarkably short compared to half-lives for equity specialists of 1 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 a much more important source of profits, at least for the dealer we track over the week we are able to track him.

Appendix 1

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 cost (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 τ will be the liquidation time. The asset's fundamental value follows a random walk. This value is known by the informed trader.

Notation

- ν_t : fundamental value at time t
- μ_t : dealer's conditional expectation of ν_t
- Ω_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.

$$Q_t(\nu_t, P_t) = \delta(\nu_t - P_t). \quad (\text{A1})$$

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

$$D_t = \delta\nu_t + X_t. \quad (\text{A2})$$

When divided by δ , this variable D_t , provides an unbiased signal of the fundamental value ν_t (since X_t is mean zero):

$$D_t/\delta = \nu_t + X_t/\delta \quad (\text{A3})$$

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

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

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

$$MAX_{P_i} \sum_{j=1}^{\infty} E[W_j | \Omega_j] Pr[j = \tau]. \quad (A5)$$

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

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

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

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

References

- Admati, A., Pfleiderer, P., 1988. A theory of intraday patterns: volume and price variability. *Rev. Financial Stud.* 1, 3–40.
- Amihud, Y., Mendelson, H., 1980. Dealership markets: market making with inventory. *J. Financial Econ.* 8, 31–53.
- Ammer, J., Brunner, A., 1997. Are Banks market timers or market makers? Explaining foreign exchange trading profits. *J. Int. Financial Markets, Institutions & Money* 7, 43–60.
- Bessembinder, H., 1994. Bid-ask spreads in the interbank foreign exchange market. *J. Financial Econ.* 35, 317–348.
- Bollerslev, T., Domowitz, I., 1993. Trading patterns and prices in the interbank foreign exchange market. *J. Finance* 48, 1421–1444.
- Bollerslev, T., Melvin, M., 1994. Bid-ask spreads and volatility in the foreign exchange market: an empirical analysis. *J. Int. Econ.* 36, 355–372.
- Bossaerts, P., Hillion, P., 1991. Market microstructure effects of government intervention in the foreign market. *Rev. Financial Stud.* 4, 513–541.
- Copeland, T., Galai, D., 1983. Information effects on the bid-ask spread. *J. Finance* 38, 1457–1469.
- Easley, D., Kiefer, N., O'Hara, M., 1997. One day in the life of a very common stock. *Rev. Financial Stud.* (in press).
- Flood, M., Huisman, R., Koedijk, K., 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. *J. Financial Econ.* 3, 257–275.
- Glassman, D., 1987. Exchange rate risk and transactions costs: evidence from bid ask spreads. *J. Int. Money Finance* 6, 479–490.
- Glosten, L., Milgrom, P., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed agents. *J. Financial Econ.* 14, 71–100.
- Goodhart, C., Figliuoli, L., 1991. Every minute counts in financial markets. *J. Int. Money Finance* 10, 23–52.
- Goodhart, C., Ito, T., Payne, R., 1996. One day in June 1993: a study of the working of the Reuters 2000-2 electronic foreign exchange trading system. In: Frankel, J., Galli, G., Giovannini, A. (Eds.), *The Microstructure of Foreign Exchange Markets*. University of Chicago Press, Chicago, IL, pp. 107–179.

- Grammatikos, T., Saunders, A., 1986. Futures price variability: a test of maturity and volume effects. *J. Bus.* 59, 319–330.
- Hansch, O., Naik, N., Viswanthan, S., 1994. Trading profits, inventory control and market share in a competitive dealer market. Duke University typescript, November, forthcoming *J. Finance*.
- Hasbrouck, J., Sofianos, G., 1993. The trades of market makers: an empirical analysis of NYSE specialists. *J. Finance* 48, 1565–1593.
- Ho, T., Stoll, H., 1983. The dynamics of dealer markets under competition. *J. Finance* 38, 1053–1074.
- Kyle, P., 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315–1335.
- Leach, J., Madhavan, A., 1993. Price experimentation and security market structure. *Rev. Financial Stud.* 6, 375–404.
- Lyons, R., 1995. Tests of microstructural hypotheses in the foreign exchange market. *J. Financial Econ.* 39, 321–351.
- Lyons, R., 1996. Optimal transparency in a dealer market with an application to foreign exchange. *J. Financial Intermediat.* 5, 225–254.
- Lyons, R., 1997. A simultaneous trade model of the foreign exchange hot potato. *J. Int. Econ.* 42, 275–298.
- Madhavan, A., Smidt, S., 1993. An analysis of changes in specialist inventories and quotations. *J. Finance* 48, 1595–1628.
- O'Hara, M., 1995. *Market Microstructure Theory*. Blackwell Business, Cambridge, MA.
- O'Hara, M., Oldfield, G., 1986. The microeconomics of the market making. *J. Financial Quantit. Anal.* 21, 361–376.
- Pagan, A., 1984. Econometric issues in the analysis of regressions with generated regressors. *Int. Econ. Rev.* 25, 221–247.
- Reiss, P., Werner, I., 1997. Interdealer trading: Evidence from London. Stanford Business School Research Papers No. 1430, February.
- Yao, J., 1996. Market making in the interbank foreign exchange market. New York University typescript, December.
- Zaheer, A., Zaheer, S., 1995. Catching the wave: alertness, responsiveness, and market influence in global electronic networks. University of Minnesota typescript, December.