ASYMMETRIC INFORMATION AND THE CROSS-SECTION OF CURRENCY SPREADS

Alexander Mende, University of Hannover, Germany
Lukas Menkhoff, University of Hannover, Germany
Carol L. Osler, Brandeis University, USA *

Abstract

This paper shows that the spreads charged by currency dealers vary inversely with deal size and that they are wider for importers and exporters than for asset managers and other dealers. This pattern is the opposite of that predicted by standard models of market making under asymmetric information, given the information structure of currency markets. The paper suggests that private information gives certain customers market power relative to their dealers. Symmetrically, it suggests that dealers strategically quote narrower spreads to privately informed customers to increase their access to information. Finally, the paper suggests that dealers primarily seek information about transitory market developments rather than fundamentals.

October 2004

Corresponding author: Lukas Menkhoff, Department of Economics, University of Hannover, D-30167 Hannover, Germany, menkhoff@gif.uni-hannover.de; Alexander Mende, Department of Economics, University of Hannover, Königsworther Platz 1, D-30167 Hannover, Germany, mende@gif.uni-hannover.de; Carol Osler, Brandeis International Business School, Brandeis University, 415 South Street, Waltham, MA 02454, USA, cosler@brandeis.edu.
We are grateful for helpful comments from Alain Chaboud, William Clyde, Peter Eggleston, Thomas Gehrig, Valerie Kraus, Peter Nielsen, Dagfinn Rime, Erik Sirri, Erik Theissen, and Peter Tordo, along with participants at several seminars and the Stockholm Workshop on FOREX Microstructure and International Macroeconomics. We are deeply indebted to the bankers who provided and discussed the data with us.

Asymmetric Information and the Cross-Section of Currency Spreads

1. Introduction

This paper examines the cross-section of currency spreads. From the seminal papers on dealing under asymmetric information we learn that spreads compensate market makers for their losses to privately informed counterparties (Copeland and Galai, 1983; Glosten and Milgrom, 1985; Easley and O'Hara, 1987). Thus, spreads should rise with the likelihood that a given counterparty has private information, other things equal, a likelihood that depends on deal size and counterparty type. Spreads should vary positively with deal size because, as shown by Easley and O'Hara (1987), informed traders have an incentive to undertake larger trades. Spreads should be narrower for counterparty types considered less informed if dealers can discriminate among them. Since currency dealers generally do know their counterparty's type the narrowest currency spreads should therefore be enjoyed by importers and exporters ("commercial customers"), since they are considered less informed than asset managers and hedge funds (jointly "financial customers") or other dealers.

This paper shows that spreads charged by currency dealers conform to the opposite of these two predictions: They are inversely related to deal size and are wider for commercial customers than for financial customers and other market makers. The resulting variation in spreads is substantial. Average spreads vary from about 2 pips for large interbank deals to about 20 pips for small deals with commercial customers.[1] The paper suggests two mutually consistent explanations for this pattern, both based on asymmetric information. According to the first,

[1] A pip is the smallest unit in an exchange-rate quote. For USD/EUR one pip equals USD 0.0001.

private information gives certain customers market power relative to their dealers. According to the second, dealers strategically quote narrower spreads to privately informed customers to increase the chances of transacting with them, thereby gathering some of their information. The paper also notes that, even though the familiar models cited above do not explain how the absolute size of currency spreads varies cross-sectionally, they do explain why the share of the asymmetric information component of spreads is positively related to deal size and is lowest for commercial customers.

Our data comprise the entire USD/EUR transaction record of a bank in Germany over four months in 2001. These data have two advantages relative to other currency transaction data analyzed in the literature: they distinguish between financial and commercial customers, and they cover a longer time period. The earliest study of currency spreads using transaction data, Lyons (1995), concludes that spreads on interdealer transactions vary positively with deal size, as predicted by Copeland and Galai (1983), Glosten and Milgrom (1985), and Easley and O'Hara (1987) ("the standard models"). More recent studies generally find little relation between currency spreads and deal size (Yao, 1998; Bjønnes and Rime, 2003).

Though not previously identified in the literature, the negative relationship between currency spreads and deal size found here is not just widely known, it is deeply woven into the fabric of thought in the practitioner community. A similar relationship holds in the U.S. municipal bond market (Green, Hollifield, and Schurhoff, 2004; Harris and Piwowar, 2004). By contrast, the relationship between spreads and deal size is positive in U.S. equity markets (Peterson and Sirri, 2003).

The pattern of spreads in stock markets can presumably be explained by the adverse selection costs highlighted in the standard models. Indeed, real-world stock markets provided the inspiration for these models. To explain the pattern of spreads in municipal bond markets, Green, Hollifield and Schurhoff (2004) highlights that market's dealership structure and the associated dispersion of information. In opaque markets, agents with knowledge of current market conditions, like dealers, have market power. Green, Hollifield and Schurhoff (2004) hypothesizes that agents making small municipal deals are tend to be poorly informed about market conditions, a tendency that gives dealers the market power to extract wider spreads. This logic can be apply directly to explain the inverse relationship between spreads and deal size in currency markets, since small currency transactions also tend to be undertaken by uninformed agents.

Green, Hollifield and Schurhoff's (2004) analysis can also explain our finding that financial customers in currency markets pay lower currency spreads than commercial customers, since financial customers are generally well informed about market conditions while commercial customers are not. In addition, financial customers have private information about their own large trades and large stop-loss and take-profit orders (Osler, 2003).[2] Since deal flow is a key determinant of exchange-rate dynamics (Evans and Lyons, 2002), and financial deal flow in particular appears to be a key factor in extreme high-frequency moves (Fan and Lyons, 2003), information about large trades and large orders can be critical to market makers.[3] This private information enhances financial customers' market power relative to their dealers and allows them to demand narrower spreads.

---

[2] Stop-loss and take-profit orders are conditional market orders, where the conditioning variable is market price. A stop-loss order instructs a dealer to buy (sell) a specific amount at market prices if and only if the market price rises (falls) to a certain pre-specified level. A take-profit order instructs a dealer to sell (buy) a specific amount at market prices if and only if the market price rises (falls) to a certain pre-specified level. Orders are distinct from regular deals, in which a market maker provides a two-way quote and the counterparty chooses whether to deal at those prices.

[3] Deal flow is defined as the net of buy-initiated and sell-initiated deals over a given interval.

The benefits from learning financial customers' private information provide a strategic incentive for dealers to quote narrow spreads to such customers. When market makers compete with each other as in dealerships markets, the more attractive a dealer's quoted spreads the more likely customers are to transaction with that dealer, both now and in the future, and the more chances that dealer will have to learn about their trading activity. That is, currency dealers may not passively accept the information content of deal flow, as assumed in standard models, but may instead set prices strategically to increase their access to information. This strategic dealing hypothesis complements, and is fully consistent with, the market power hypothesis.

The possibility of strategic dealing was originally explored in theoretical papers by Gammill (1989) and Leach and Madhavan (1992, 1993), which show that dealers may rationally adjust prices in early transactions with the goal of becoming more informed, and thus more profitable, in later transactions. In the Leach and Madhavan models, specialists rationally quote wider spreads at the beginning of a trading session, driving informed trades *out*. In Gammill's model, dealers choose to take a loss on early trades, driving informed trades *in*. Our hypothesis is closest to Gammill's.

The hypothesis that currency dealers strategically subsidize deal flow with informed customers raises the question: Are currency dealers seeking to gain fundamental information or information about transient market developments? Consistent with the hypothesis that dealers seek fundamental information (Evans and Lyons, 2002, 2004), we show that interbank and financial-customer deal flow are both positively cointegrated with exchange rates. Nonetheless, we propose that dealers primarily seek information about transient market developments ("non-fundamental information"), and provide three justifications for this view. First, fundamentals are primarily relevant in the long run, but currency dealers generally close their positions by the end of the day. Second, dealers themselves characterize the information they seek as relating to transient developments. Third, there may be relatively little fundamental information to be gleaned from currency trades. Exchange-rate fundamentals are generally considered to be broad macroeconomic aggregates such as money supplies and price levels, information about which is in the public domain. By contrast, much market-relevant information about equity and municipal bond issuers is unearthed by equity analysts but never becomes public.

To illustrate the type of non-fundamental information currency dealers may seek we note that intraday currency deal flow is correlated, even though daily returns are not. As shown by Goodhart, Ito, and Payne (1996), there tend to be runs of buy orders and of sell orders. Presumably these runs often reflect large trades, which are typically divided into many smaller transactions. By subsidizing transactions with the largest customers, a market maker may raise the likelihood of participating in large trades and being informed about the associated runs.

The difference between spread determination in currency and equity markets can be summarized with a cost-benefit analysis of market making. Standard adverse selection costs may be lower for currency dealers because private fundamental information is less common in foreign exchange markets. Meanwhile, currency dealers may benefit from trading with customers who have private information about transient aspects of deal flow like large trades.

The bank from which this paper's data are derived is relatively small. Nonetheless, there are three reasons why our conclusions should generalize to the overall currency market. First, the intense competition in major currency markets means that any bank's pricing practices should accurately represent practices at all banks. Second, traders from large banks tell us that their pricing policies conform to those described here: supporting quotes from two market participants are provided below. Third, our small bank behaves similarly to large banks in many other dimensions. Indeed, a secondary contribution of the paper is to show this consistency between the behavior of small and large currency dealers.

The paper proceeds as follows: Section 2 describes our data and shows how our small bank's pricing and inventory management practices parallel those at large banks. Section 3 provides our core results that currency spreads vary inversely with deal size and tend to be larger for commercial customers than financial customers and other dealers. This section also shows that these results can be explained in terms of the market power and strategic dealing hypotheses. Section 4 discusses whether currency dealers strategically seek fundamental or non-fundamental information. Section 5 shows that the share of the asymmetric information component of spreads is largest when counterparties are most likely to be informed, consistent with standard models. Section 6 concludes.

2. Small banks and large banks

This section describes our transactions data and provides a preliminary comparison of our bank's pricing and inventory management practices and those at large banks. It also presents the model on which we base our core results, which are presented in the next section. We find that our bank is indeed small relative to others examined in the literature, but that its behavior is nonetheless consistent with that of large banks in recent years. Readers familiar with the model and willing to trust that small and large banks behave similarly can safely skip to Section 3, which presents our central results.

2.1. Data

Our data comprise the complete USD/EUR transaction record of a bank in Germany over the 87 trading days from 11 July, 2001 to 9 November, 2001. Though our data technically refer to the overall bank, they are an accurate reflection of a single dealer's behavior because only one dealer was responsible for the bank's USD/EUR trading. For each transaction we have the following information: (1) the date and time;[4] (2) the direction (customer buys or sells); (3) the quantity; (4) the transaction price; (5) the type of counterparty: bank, financial customer, commercial customer, preferred customer; (6) the initiator; (7) the forward points if applicable. Table 1 provides basic descriptive statistics.[5]

We include outright forward trades, adjusted to a spot-comparable basis by the forward points, as recommended by Lyons (2001). Since forward transactions account for 20 percent of all trades, their inclusion could impede direct comparisons between our results and those of most earlier papers, which focus exclusively on spot trades. Reassuringly, our main qualitative conclusions are sustained when forward transactions are excluded.

The bank's inventory position is inferred by cumulating successive transactions. Following Lyons (1995), we set the daily starting position at zero. This should not introduce significant distortions since our dealer keeps his inventory quite close to zero. As shown Figure 1, which charts the dealer's inventory over the sample period, the average inventory position is EUR 3.4 million during the trading day and only EUR 1 million at the end of the day.

Our ability to distinguish among customer types is almost unique in currency transaction data. Lyons (1995) only uses data on interbank trading; Yao (1998) uses customer trade data but does not generally distinguish among customer types; Bjønnes and Rime (2004) have insufficient customer transactions to perform a detailed analysis; Carpenter and Wang (2003) have such information but lack inventory data to test the relevant models; finally, Lyons (2001) and Fan and Lyons (2003) can distinguish among customer types but only in daily data.

[4] The time stamp indicates the time of data entry and not the moment of trade execution, which will differ slightly. Nevertheless, there is no allocation problem because all trades are entered in a strict chronological order.

[5] We exclude trades with "preferred customers", typically commercial customers with multi-dimensional relationships with the bank, because these customers' spreads may reflect cross-selling arrangements and because their trades are typically very small (average size EUR 0.18 million). We also exclude a few trades with tiny volumes (less than EUR 1,000) or with apparent typographical errors.

A preliminary comparison of our dealer with the other large dealers described in the literature is provided in Table 2. The dimensions in which our dealer is small include total daily trading value, average transactions per day, and average inventory position. Our dealer is comparable in size to a NOK/DEM dealer employed by a large dealing bank examined in Bjønnes and Rime (2004). Small dealing banks are far more common than large ones (B.I.S., 2002), so our bank is probably a reasonably good representative of the average currency-dealing bank. Nonetheless, big banks are thought to dominate such dealing.

The small size of our bank is also reflected in the prominence of customer deals, especially those with commercial customers (Table 1). Our bank's customer business is 23 percent of its spot trading value. Though this does not differ much from the 33 percent share of customer business at all foreign exchange banks (B.I.S., 2002), it greatly exceeds the customer shares reported for bigger dealers, which range from zero percent (Lyons, 1995) to 14 percent (Yao, 1998). Commercial customers generate roughly twice the business of financial customers at our bank, by value. By contrast, commercial customers do roughly half the business of financial customers (B.I.S., 2002) in the foreign exchange market overall.

The small size of our bank is also reflected in the large mean absolute change in transaction price between successive deals, 10.7 pips. This presumably reflects the relative infrequency of transactions at our small bank as well as the high proportion of small commercial customer deals, which tend to have wide spreads (as we show in Section 3). Table 3 provides information on the size distribution of our dealer's transactions.

2.2. Pricing and inventory management practices

Despite the small size of our bank, there are a number of reasons to believe our qualitative conclusions generalize to the entire currency market. First, currency markets are extremely competitive. Hundreds of banks deal in the major currency pairs and even the largest dealer's market share is only on the order of 10 percent. In such a market, the behavior of any agent should accurately represent the behavior of all agents. Second, market participants consistently confirm that the patterns we identify are correct. Third, our small bank's pricing and inventory management strategies are generally consistent with those documented in recent years for large banks, as we show next.

2.2.1. Pricing

Lyons' (1995) version of the Madhavan-Smidt model of market making (1991) has been widely used in subsequent studies of currency dealers (e.g., Yao, 1998; Bjønnes and Rime, 2004). The model assumes a representative dealer in a competitive market whose counterparty has private information about the asset's fundamental value. The model's agents are fully rational and there is a detailed informational setting. Agent *j* calls dealer *i* requesting a quote on amount Q_{jt}; that amount is determined as follows:

$$Q_{jt} = \xi(\mu_{jt} - P_{it}) + X_{jt} \tag{1}$$

X_{jt} represents agent *j*'s nonspeculative need for currency, which constitutes agent *j*'s private information. The term μ_{jt} represents agent *j*'s expectation of the asset's true value, conditional on X_{jt} and on public information. P_{it}, dealer *i*'s regret-free price, is determined as follows:

$$P_{it} = \mu_{it} + \zeta(I_{it} - I_{it}^{*}) + \gamma D_t\,. \tag{2}$$

Here, I_{it} is dealer *i*'s inventory at the beginning of period *t*, I^*_{it} is his desired inventory, and D_t is the direction of trade [$D_t = 1$ (-1) if agent *j* is a buyer (seller)].

After solving for conditional expectations and taking first differences, one arrives at the following expression for the price change between incoming transactions, $\Delta P_{it} = P_{it} - P_{it-1}$:

$$\Delta P_{it} = \alpha + \beta_1 D_t + \beta_2 D_{t-1} + \gamma_1 I_{it} + \gamma_2 I_{it-1} + \delta Q_{jt} + \eta_t \tag{3}$$

The model predicts that $\beta_1 > | \beta_2| > 0 > \beta_2$ and that $| \beta_2|$ equals the baseline half-spread, meaning the half-spread that would apply before adjustment for deal size or existing inventories. The model assumes that dealers shade prices to help manage existing inventory (e.g., dealers lower prices in response to high inventory), implying $\gamma_2 > 0 > \gamma_1$.

According to the model, the coefficient on deal size should be positive, reflecting adverse selection considerations: spreads should be wider for larger deals because they are more likely to be undertaken by privately informed agents (Easley and O'Hara, 1987). However, a positive coefficient on deal size could also capture a second type of inventory concerns: as shown in Ho and Stoll (1981), larger deals leave market makers with higher inventory and thus greater inventory risk, so they should carry wider spreads. We call this a "prospective" inventory effect, since it concerns inventories that may arrive if the counterparty decides to deal at the current quote and because we need to distinguish it from the effect of existing inventory captured by γ_1 and γ_2. The adverse selection and prospective inventory effects both predict $\delta > 0$ and are observationally equivalent in this setting.

The model is typically estimated using generalized method of moments with Newey-West correction for heteroskedasticity (e.g., Yao 1998; Bjønnes and Rime, 2004). We first estimate Equation (3) without discriminating among counterparties (Table 4A, column 1), and then re-estimate it distinguishing interdealer transactions and customer transactions by interacting dummy variables for each with the direction, inventory, and deal size variables (Table 4B, column 1). Since existing inventories appear to have no influence we re-run both regressions excluding inventories. We also re-run the regressions using only spot transactions. The results are robust to these changes, as shown in columns 3 and 4 of each panel of the table.

We compare these results with the results of similar regressions using large-bank data reported in earlier studies. This comparison indicates that dealer behavior is consistent across dealers of all sizes in three dimensions: baseline spreads, the influence of existing inventories, and the relationship between deal size and spreads. We discuss each dimension in turn.

Baseline spreads: Our bank's average baseline half-spread for interbank transactions is about 1.5 pips (Table 4B), similar to estimates from other studies. For example, Goodhart et al. (2002) finds that the average spread for USD/EUR transactions on Electronic Brokerage Service (EBS, one of the two major electronic brokerage systems for interbank trading) was 2.8 pips about one year after the euro was introduced. Our bank's average half-spread for customer deals, 9 pips, is much higher than its interdealer spread. Bjønnes and Rime's (2001) NOK/DEM dealer also makes a sharp distinction between dealers and customers.

Influence of existing inventories: Our results indicate that existing inventories have no influence on the prices our dealer quotes to other dealers, consistent with recent studies of large banks (Yao, 1998; Bjønnes and Rime, 2004). By contrast, Lyons (1995) provides evidence that his dealer did engage in inventory-based price shading towards other dealers in 1992. This may reflect the unusual character of Lyons' dealer who, as a jobber, dealt exclusively with other dealers at extremely high frequency. Yao (1998) claims that his dealer avoided such shading because it would reveal information about his inventory position.

Bjønnes and Rime (2004) argue that the apparent shift away from inventory-based price shading over the 1990s may reflect the way the interbank market shifted rapidly to a heavy reliance on electronic brokerages after their introduction in the mid-1990s (Melvin and Wen, 2003).[6] Indeed, our dealer reports that for interbank trades he generally uses EBS because it is less expensive and faster than direct interbank dealing.[7] Together, this these observations imply that our dealer controls inventories via interbank trading instead of price shading, a conclusion we support empirically later in this section.

The estimates seem to provide slight evidence of price shading with respect to customers, but the shading seems to go the "wrong" way. Reassuringly, this can be traced to one trade carried out in the first month of our sample period. When that month is excluded, both coefficients are insignificant.

Trade size and spreads: The coefficient on deal size is statistically insignificant for interbank trades, suggesting that neither information asymmetries nor prospective inventories cause large interbank deals to be priced less attractively than small deals. This is consistent with recent empirical depictions of interbank trading at large banks. Bjønnes and Rime (2004) finds that spreads are independent of deal size for the brokered trades that now dominate such trading. It also finds that spreads rise with deal size for direct interbank transactions, a distinction that makes economic sense. Dealers have limited control over the relationship between deal size and spread for brokered transactions, but they have full control for direct deals. Notably, the earliest studies of currency dealers (Lyons, 1995; Yao, 1998), which do not control for the distinction between direct and brokered trades, found that interbank spreads do rise with deal size,

---

[6] In a direct interbank trade, one bank calls another and asks for a two-way quote for a specific amount. Electronic brokers take limit orders from dealers and post the best bid and ask prices. Limit orders are then crossed with dealers' market orders.

[7] This preference is supported by the transactions data. Our dealer's mean interbank transaction size was only EUR 1.42 million (Table 1), the maximum interbank trade size was only EUR 16 million, and the standard deviation of these trade sizes was only 1.42. These small values are consistent with heavy use of EBS, where the mean USD/EUR transaction size in August 1999 was EUR 1.94 million and the standard deviation of (absolute) transaction sizes was 1.63 million. By contrast, interbank deals averaged closer to USD 4 million prior to the emergence of electronic brokerages (Lyons, 1995).

consistent with standard models. This could reflect the fact that interbank trading was mostly carried out through direct transactions until the late 1990s.

The coefficient on deal size is also insignificant for customers in our baseline regression. We note in passing that this coefficient is negative and significant when inventories are excluded. Section 3 shows that the overall relationship between spreads and transactions sizes is indeed negative for customer transactions.

2.2.2. Inventory management

Our dealer's tendency to keep inventories close to zero (Figure 1) is itself similar to inventory management practices at large banks. As Table 1 shows, currency dealers of all sizes tend to keep minimal inventories. A more rigorous description of our dealer's approach to inventory management comes from the following regression:

$$I_{it} - I_{it-1} = \varpi + \rho I_{it-1} + \varepsilon_t \ . \qquad (4)$$

If the dealer instantly eliminates unwanted inventories, $\rho \approx$ -1. If the dealer allows his inventory to change randomly, $\rho = 0$.

Results from estimating Equation (4), once again using GMM with Newey-West correction for heteroskedasticity, are presented in Table 5. They confirm that our small bank does actively keep inventories close to zero. The negative and statistically significant coefficient on lagged inventory implies that our dealer typically brings inventories halfway back to zero within 19 minutes of an inventory shock. This is quite close to the 18-minute median inventory half-life for Bjønnes and Rime's (2004) NOK/DEM dealer. By contrast, the median inventory half-lives of that bank's DEM/USD dealers are only 0.7 to 3.7 minutes.

Since our dealer does not seem to use price shading to control inventory, it seems likely he controls it through the interbank market, instead. This would be consistent with at least two of the dealers at the large bank examined in Bjønnes and Rime (2004), including the NOK/DEM dealer. To examine this possibility, we carry out a probit analysis of the probability that a given trade is outgoing:

$$Prob(Trade_t\!=\!IB^{out}) = P(|I_{it}|,\ I_{it}^{2},\ |Q_{jt}|,\ IB|Q_{jt-1}|,\ FC|Q_{jt-1}|,\ CC|Q_{jt-1}|)\ . \qquad (5)$$

If dealers are more aggressive in eliminating large inventories than small ones the coefficient on the absolute value of inventory, $|I_{it}|$, will be positive. If dealers automatically eliminate inventory, this coefficient will be insignificant. We include squared inventory, I_{it}^2, to capture nonlinearities in this relationship. The absolute transaction size, $|Q_{jt}|$, may capture technical aspects of dealing described below. The variable $IB|Q_{jt-1}|$ is an interaction term between absolute transaction size and a dummy that equals unity if the previous transaction was an incoming interbank (*IB*) deal and zero otherwise; $FC|Q_{jt-1}|$, and $CC|Q_{jt-1}|$ are defined accordingly for financial customer (*FC*) and commercial customer (*CC*) transactions. Coefficients on these variables should be positive if outgoing transactions are customarily used to eliminate unwanted inventory. We allow the coefficients to vary according to counterparty type because one might expect dealers to be relatively aggressive in eliminating inventory accumulated in deals with more informed counterparties.

The results of estimating Equation (5), shown in Table 6, indicate that the likelihood of an outgoing deal rises with the absolute amount of existing inventory, and that the relationship is convex. The positive overall relationship implies that our dealer does rely on outgoing transactions to manage his inventory, consistent with the large dealers of major currency pairs analyzed by Bjønnes and Rime (2004).

We find a positive relationship between absolute deal size, $|Q_{jt}|$, and the likelihood that the deal itself is outgoing, which indicates that the transactions submitted to the brokers tend to be larger than the dealer's average incoming transaction. We infer that our dealer sometimes collects inventory from small customer transactions and then squares his position by submitting one relatively large outgoing order. Like our earlier results, this is consistent with behavior at large banks (Bjønnes and Rime, 2004).

As expected, the likelihood of turning to the brokers appears to be influenced by information considerations: outgoing trades are more likely when the previous trade was an incoming interbank or financial customer trade. This supports dealers' statements that they consider transactions with financial customers and other banks to carry more information than transactions with commercial customers.

Overall, this section shows that the dealer from which we take our data is small by international standards but behaves much like large dealers. On this basis it seems reasonable to generalize from the additional observations of this bank’s pricing practices provided below to those of the currency market as a whole.

3. Cross-sectional variation in currency spreads

This section provides our core results concerning the cross-section of currency spreads. Our empirical analysis shows that currency spreads are wider for commercial customers than for financial customers and other banks, and wider for small deals than for large deals. Together these results imply that spreads are narrowest for the deals most likely to be with informed counterparties. Though standard models of dealing under asymmetric information do not predict this pattern, it can be explained in terms of market power and strategic dealing.

3.1. Empirical analysis

Customer type: We first incorporate the distinction between commercial and financial customers into the Lyons-Madhavan-Smidt model presented in Section 2. The results, shown in

Table 7, indicate that the baseline half-spread on commercial customer transactions is 10 pips, while the baseline half-spread on financial customer transactions is only 2 pips. In fact, the baseline half-spread on financial customer transactions is only significant at the ten percent level, and differs little from baseline half-spread for other dealers. As shown in the table, these qualitative conclusions are robust: they do not change if we exclude inventory terms, or if we consider only incoming spot trades.[8]

The standard models of market making under asymmetric information do not imply these results. At a superficial level those models imply the complete absence of counterparty-based variation in spreads because dealers in the models, like real-world equity specialists, know little about individual counterparties. Under a more nuanced interpretation the models suggest that spreads for informed counterparties would be wider, not narrower, if dealers could identify them.

Trade size: We now modify the Lyons-Madhavan-Smidt model to reflect some potentially important nonlinearities in the relationship between deal size and currency spreads. According to our correspondents in the market, normal-sized customer transactions (those below about 25 EUR or USD millions) are informally divided into three categories: regular deals, which vary from EUR 1 million to about EUR 25 million, modest deals, and tiny deals. Though the line between the latter two categories is ambiguous, their treatment can vary substantially: tiny deals are often spread by formula, rather than dealers' discretion (and three percent is not considered unreasonable). For estimation purposes we distinguish the following three size ranges for customer transactions: Large deals: $\{|Q_{jt}| \in [\text{EUR 1 million, EUR } \infty)\}$; Medium: $\{|Q_{jt}| \in [\text{EUR 0.5 million, EUR 1 million})\}$; and Small: $\{|Q_{jt}| \in (\text{EUR 0, EUR 0.5 million})\}$.

[8] Our correspondents at dealing banks indicate that the correct counterparty disaggregation is between small commercial customers, on the one hand, and financial customers and large multinational (commercial) corporations, on the other. Though we cannot distinguish large multinational corporations from other commercial customers, our small bank is unlikely to do much business with large multinational corporations. Thus the tiering by counterparty we show here should be accurate.

For interbank deals, market participants suggest a single pricing nonlinearity at EUR 1 million, since the brokered deals that now dominate interbank trading are only permitted in integer multiples of EUR 1 million. Thus, we distinguish two categories of interbank deals: Large: $\{|Q_{jt}| \in [\text{EUR 1 million, EUR } \infty)\}$; and Small: $\{|Q_{jt}| \in (\text{EUR 0, EUR 1 million})\}$. We do not create separate categories for deals above EUR 25 million because our dataset includes few.

The results, shown in Table 8, indicate a negative relationship between baseline spreads and deal size, while continuing to indicate an independent influence of counterparty type. The tiering of spreads by deal size is most pronounced for commercial customers, for whom the baseline half-spread is estimated to be 11.4 pips on small deals, 8.4 pips on medium-sized deals, and 2.9 pips on large deals. For financial customers, the baseline half-spread is estimated to be 5.5 pips on small deals and is smaller, and insignificantly different from zero, for medium-sized and large deals.[9] For interbank deals, the baseline half-spread is estimated to be 2.9 pips on small deals and 1.3 pips on large deals. We note in passing that spreads on large interbank deals are estimated to exceed those on large financial customer deals; we return to this below.

These qualitative conclusions are generally robust to the exclusion of inventories from the right-hand-side or the exclusion of spot trades from the sample. The one exception is the negative relationship between spreads and deal size for interbank trades, which seems to reverse when we consider only spot deals. This change is largely due to the new insignificance of baseline half-spreads on small interbank deals; the very small number of such deals (Table 3) makes these results sensitive to changes in our sample.

We also find that, conditional on the baseline spread for its size category, the spread charged on commercial customer deals rises with deal size. As noted in Section 2, this positive

---

[9] It may seem surprising that that spreads could be zero. However, currency dealers have complained for over a decade that good financial customers are sometimes given a single price and allowed to choose a direction.

(but now second-order) relationship could reflect either adverse selection costs or a prospective inventory effect. To us, it seems more likely to reflect inventory concerns because dealers consider the information content of commercial customer transactions to be small.

Once again, it is worth contrasting these results with those implied by standard models. The model of Easley and O'Hara (1987) indicates that spreads should vary positively with deal size, and empirical evidence supports this prediction for U.S. equities (Peterson and Sirri, 2003). We basically find the reverse in currency markets.

3.2. Explanations: market power and strategic dealing

Though this pattern of currency spreads is not predicted by standard models, it is familiar to all active currency market participants. The Appendix provides two quotes from market participants who have, or have had, significant managerial responsibility for trading activity at large banks. These individuals confirm the patterns identified above and suggest reasons for it. We elaborate on these and other potential explanations below.

A negative relationship between spreads and deal size would be predicted by simple economies of scale: large spreads help small deals cover the fixed administrative costs of dealing. However, our results indicate that spreads vary widely across counterparty type even within size categories. Thus economies of scale cannot completely explain the cross-sectional pattern of currency spreads. To fill the remaining gap, we highlight two mutually consistent theories of dealing under asymmetric information. One theory suggests that information provides market power, which in turn determines spreads; the other suggests that dealers strategically vary spreads across counterparties in an attempt to gather information.

3.2.1. Market Power

Green, Hollifield, and Schurhoff (2004) shows that variations in market power between dealers and their customers may explain why spreads are inversely related to deal size in the U.S. municipal bond market. They point out that dealership markets, like the municipal and currency markets, are relatively opaque due to the dispersion of trading. Indeed, it takes effort in such markets even to learn the current market price, and "comparison shopping [is] relatively costly" (p. 1). Empirically, the customers who make larger municipal bond deals tend to know more about current market conditions. According to the theory, this knowledge gives them market power relative to their dealers, allowing them to demand narrower spreads.

Some sources of the information dispersion in municipal bond markets are not at work in currency markets. Unlike municipal bonds, currencies are not issued with specific maturities, and few (if any) are traded exclusively at obscure dealers. Nonetheless, information is sufficiently dispersed in currency markets that commercial customer are considered generally uninformed about current market conditions. These customers' transactions are often carried out by clerks for whom trading is only one among many administrative responsibilities. By contrast, financial customers typically hire professional traders, provide them with real-time trading information, and track their performance carefully. Thus the market power hypothesis of Green, Hollifield, and Schurhoff (2004) can be applied directly to explain why commercial customers pay wider currency spreads than financial customers and other market makers.

3.2.2. Strategic Dealing

The strategic dealing hypothesis of Leach and Madhavan (1991, 1992) and Gammill (1989) can also shed light on the counterparty-based tiering of currency spreads. In its broadest form, this hypothesis asserts that market makers strategically adjust prices to influence the information content of their deal flow. Gammill's (1989) version of this hypothesis, which suggests that dealers sometimes subsidize transactions to gain information to increase trade with informed counterparites, could explain why financial customers are offered narrower currency spreads. The remainder of this section provides evidence that currency deal flow does carry information, even at a small bank, and that the information content of large financial customer deals may be especially high, as required to explain the pattern we observe.

Theoretical sources for the information value of deal flow are not difficult to discern. Trading is costly and the vast majority of customer transactions are too small to influence markets noticeably, so it is usually not sensible for customers to speculate on the basis of the information content of their own trades. When such trades are aggregated across many counterparties, however, their overall information value could be non-negligible since aggregate deal flow does affect exchange rates (Evans and Lyons, 2002) and does appear to forecast exchange-rate fundamentals (Evans and Lyons, 2004).[10]

One might presume that large banks gather the most information from deal flow. Indeed, a small bank's relatively limited deal flow might might not be informative at all. To examine the information value of our small bank's deal flow, we estimate a cointegrating relationship between exchange rates and our three types of cumulative deal flow:

$$P_t = \omega + \phi trend + \kappa_{IB} CumDealFlow_{IBt} + \kappa_{FC} CumDealFlow_{FCt} + \kappa_{CC} CumDealFlow_{CCt} + v_t \quad (6)$$

The residuals, v_t, are stationary, as required.[11] More importantly, from the dealers' perspective, the coefficient on the residual in the associated error correction equation is negative (-0.010) and statistically significant (t-value -3.99). This implies that our small bank's deal flow has

[10] Note that this logic implies that dealers, rather than their customers, are the ones with private fundamental information in currency markets – the opposite scenario from that postulated in standard microstructure models.
[11] The ADF statistic for this test is -2.93. The Philips-Peron statistic is -3.58.

predictive power for exchange rates. The estimated coefficient implies that the half-life of this predictive power is about fifteen hours.

For the strategic dealing hypothesis to help explain why currency spreads vary accross counterparties, dealers must consider certain types of counterparties to have more information than others. Dealers certainly claim that transactions with financial customers and other dealers are more informative than transactions with commercial customers (see Appendix). Our dealer's behavior is consistent with this claim, insofar as he is more aggressive in eliminating inventories when the counterparty is a financial customer than when it is a commercial customer (as shown in Section 2). However, it is still worth inquiring: Does empirical evidence confirm that financial customer deal flow is more informative than commercial customer deal flow?

One empirical indication of the high information value of financial customer deal flow comes from Fan and Lyons (2003), which shows that "extreme exchange-rate movements at high frequency are generally associated with large net flows from financial institutions" (p. 160), but not from commercial institutions.

A comparison of our own bank's financial and commercial customer deal flow provides suggestive results. We estimate separate cointegrating relationships between exchange rates and our three categories of cumulative incoming deal flow:

$$P_{it} = \omega_i + \phi_i trend + \kappa_i CumDealFlow_{it} + \nu_{it} . \qquad (7)$$

Here i represents the counterparty type, $i \in \{IB, FC, CC\}$. If demand-driven deal flow of type i tends to bring a currency appreciation, κ_i will be positive.

As shown in Table 9, the USD/EUR exchange rate is positively cointegrated with cumulative incoming interbank deal flow and with cumulative financial customer deal flow, but negatively cointegrated with cumulative commercial customer deal flow. Superficially this seems consistent with the idea that transactions with financial customers and other dealers, for which the relationship has the "right" sign, are more informative than transactions with commercial customers. Nonetheless, it is possible that commercial customer deal flow is informative despite the "wrong" sign of the cointegrating relationship. To us, however, it seems more likely that commercial customer business simply responds to exchange rates, just as predicted by traditional international economics: a stronger currency discourages exports and encourages imports, generating a negative relationship between deal flow and exchange-rate levels if certain elasticity conditions are satisfied (as is usually assumed).[12,13] In short, our small bank's transaction record provides some support for the idea that dealers strategically subsidize the transactions of financial customers and other market makers in order to gather market-relevant information.[14]

The strategic dealing hypothesis can also help explain why currency spreads are narrower for large financial customer transactions than for similarly sized interdealer transactions. This might not be predicted by the currency microstructure literature (e.g. Lyons, 2001), which stresses that dealers gain private information from deal flow. In turn, this implies that dealers, as the most informed agents in the market, would get the narrowest spreads. However, market sources tell us that the largest asset managers gather information from many dealers, and are thus more informed than individual dealers. This puts the largest asset managers at the pinnacle of the

[12] The traditional "Marshall-Lerner-Robinson" condition would have to be amplified to reflect exchange-rate pass-through and related matters studied in recent research.

[13] It is also worth noting that total deal flow in the market must sum to zero: the amount of sales must equal the amount of purchases. For some types of traders – like financial customers or banks as a group – to be influential, or to successfully anticipate the market, they must find counterparties who are not influential or who fail to anticipate the market. It appears that commercial customers fulfill this role.

[14] In addition to quoting financial customers narrow spreads, dealers seek the business of large financial customers by entertaining them lavishly: real-world examples include dude ranch visits in the U.S. west, NY harbor cruises followed by fireworks, and multicourse feasts at elegant establishments.

currency information pyramid, a position that maximizes their market power and enables them to demand the smallest spreads.

Could dealers' strategic incentives also contribute to the negative relationship between deal size and currency spreads? To investigate the role of deal size in communicating information, we focus on financial customer transactions, taking at face value dealers' claim that commercial customer transactions are largely uninformative regardless of size (see Appendix). As suggested by our earlier analysis of spreads (Table 8), we partition these trades into two categories, small and medium-and-large. Cointegration tests, reported in Table 9, show that the link between cumulative financial customer deal flow and exchange-rate levels is significant for medium-and-large deals but insignificant for small deals (again, the residuals appear stationary). This is consistent with the hypothesis that medium-and large-financial customer deals are quoted the tightest spreads in part because they have the highest information value to dealers.

This section has shown that spreads charged by currency dealers vary inversely with deal size and are wider for commercial customers than for financial customers and other dealers. We explain this pattern in terms of economies of scale and two mutually consistent theories of market making under asymmetric information. First, spreads may be positively related to dealers' market power relative to their customers. Second, spreads may reflect dealers' strategic attempts to increase their access to information. In currency markets, financial customers and other dealers view themselves as exploiting the market power associated with their private information to demand smaller spreads. The quoting dealers view themselves as strategically setting small spreads to increase their business with informed customers.[15] They are all right.

---

[15] A brief anecdote underscores the seriousness with which banks pursue informative deal flow. About a decade ago, a dispute arose between the salespeople and interbank traders at a significant foreign exchange dealing bank. The salespeople were "spreading the quotes" to customers wider than the interbank traders preferred. The salespeople

4. *Information in currency markets*

If it is indeed strategically wise for currency dealers to gather information by subsidizing transactions with informed customers, what kind of information do they seek? Or equivalently, what kind of private information do financial customers possess that gives them so much market power relative to their dealing banks? Two alternatives roughly span the space: fundamental information and information about transient market developments. We consider each in turn.

4.1. *Fundamental information*

In standard microstructure models, private information concerns the asset's fundamental value and is possessed by customers. This corresponds well to the reality of equity and bond markets where asset managers regularly investigate issuers to gain private information about factors such as executive talent. Equity and bond dealers rarely have that luxury.

In major currency markets, by contrast, real-world sources of private fundamental information are harder to identify, and the possessors of such information probably include the dealers themselves. Currency fundamentals are generally thought to be macroeconomic aggregates like prices, interest rates, and money supplies, all of which are public information. Information about foreign exchange interventions and other large anticipated currency trades might qualify as private, but intervention intentions are extremely closely guarded for policy reasons, while anyone planning a large trade has a strong profit-based incentive to keep the plan secret. In theory, asymmetric information about fundamentals could also arise from private insights about how to interpret public information. However, exchange-rate forecasts by

---

were simply responding to the structure of their bonuses, which depended exclusively on their own profits. The interbank traders, by contrast, wanted narrower spreads on certain types of deals, to encourage informative deal flow (they were explicit about the need to learn about the market from deal flow). The interbank traders won this dispute, and bonus formulae were revised to ensure that everyone was motivated by the dealing room’s total profit pool. Spreads narrowed and deal flow picked up immediately.

professionals tend to be biased and inefficient, among other difficulties (MacDonald, 2000), so this type of private information may also be in short supply.

A more promising source of private fundamental information in currency markets is each bank's deal flow (Lyons, 2001). As shown above, even a small bank's deal flow provides exchange-rate relevant information, and this information could be fundamental since deal flow and exchange rates are cointegrated. Even more importantly, deal flow predicts exchange-rate fundamentals (Evans and Lyons, 2004).

4.2. Information about transitory market developments

While it is possible that dealers seek fundamental information when subsidizing certain types of transactions, it is our view that dealers are more interested in information about transitory market developments. Three arguments support this view. First, dealers generally close their positions by the end of the trading day. Since fundamental information has essentially no forecasting power for short-run exchange rate dynamics (Meese and Rogoff, 1983), it is not obvious how dealers could use fundamental information for pricing or position-taking at such a high frequency. Second, the information dealers can gather from deal flow appears to be primarily valuable for a relatively short time period: as shown earlier, deal flow seems to lose half its predictive value in fifteen hours. This does not support the idea that deal flow is valuable to dealers because of its relationship to fundamentals.

Third, dealers themselves report that the information they seek is not fundamental. Gehrig and Menkhoff's (2003) survey of foreign exchange dealers and fund managers in Germany and Austria finds that information about deal flow is important and that it has become more so over time. More critically, the survey indicates that these agents use deal flow less to learn about fundamentals than to exploit transitory market developments. Similar responses are reported in Cheung and Chinn (2001) for the U.S.

Further direct evidence of how highly dealers value non-fundamental information comes from a 2002 survey of North American currency traders (Oberlechner, 2004). The survey asked each trader to rate, on a scale of 1 to 6, the importance to their profitability of short-term, medium-term, and long-term price forecasts. Sixty-eight percent of the 400-plus respondents indicated that short-term forecasts were in the top two categories of importance, while only 27 percent of respondents made that claim for long-term forecasts. Similarly, only six percent of respondents indicated that short-term forecasts were in the lowest two categories of importance, while 28 percent indicated that long-term forecasts were in those lowest categories. Since economic fundamentals are poor at forecasting short-run exchange rate dynamics (Meese and Rogoff, 1983), it seems reasonable to assume that the relatively-important short-term forecasts are based on transitory market developments.

One type of non-fundamental information that dealers might seek is whether a customer currently has a large need for currency. Large trades matter because of the way they are managed. The well-organized upstairs market for block trades on the New York Stock Exchange has no equivalent in currency markets. Instead, a customer typically breaks a large trade into many smaller transactions, hoping to get better prices through counterparty banks' ignorance of the trade's total size. This presumably helps explain why there tend to be runs of buy transactions and runs of sell transactions (Goodhart, Ito, and Payne, 1996). Since deal flow affects exchange rates (Evans and Lyons, 2002), knowing the direction in which a large customer is trading should help predict rates. Such information is not fundamental since its relevance is unlikely to exceed a few days. Dealers report that it rarely takes more than one day to fill customers' large currency needs; consistent with this, daily exchange-rate changes are not correlated despite the tendency of buy and sell transactions to come in runs (Goodhart, Ito, and Payne, 1996).

Even if a financial customer is not currently involved in a large trade, a dealer may still benefit from quoting him a narrow spread. In dealership markets a good customer can call myriad dealers. Quoting attractive spreads to such customers presumably increases the odds of being called when the customer does undertake future large trades. Ideally, a dealer will be asked to manage those large trades, a privilege that provides the maximum information advantage.

Dealers also seek information on the placement of stop-loss and take-profit orders. Indeed, such information is one of the major commodities traded in the informal market for information among dealers and their better customers. Stop-loss orders generate positive-feedback trading that can propagate short-term price trends and can sometimes contribute to price cascades (Osler, 2004) or, more extremely, liquidity black holes (Morris and Shin, 2003). Individual orders can exceed EUR 500 million, and even the small ones can become important since they have a tendency to cluster at certain exchange rate levels (Osler, 2003). Take-profit orders generate negative-feedback trading that can stop or reverse trends.

5. The asymmetric-information share of currency spreads

We have suggested that adverse selection costs may be lower in currency markets than equity markets because private fundamental information is relatively scarce in currency markets. This need not imply, however that adverse selection costs in currency markets are zero. Currency dealers do worry, for example, that some other dealer is managing the latest FDI transaction. This section provides evidence that adverse selection costs do affect currency spreads even though they do not dominate them.

We implement an approach to estimating spread components suggested in Huang and Stoll (1997) and modified in Bjønnes and Rime (2004). Huang and Stoll's analysis begins with the observation that trade size is relatively unimportant for pricing in limit order markets (such as EBS) because deal sizes are standardized and large transactions are divided into small ones. Nonetheless, the risk of trading with a better informed counterparty remains. Thus, Huang-Stoll's model assumes that that the critical determinants of price are a deal's direction and the market maker's existing inventories.

In the model dealer i sets his price, P_{it}, as follows:

$$P_{it} = \mu_{it} - \theta \frac{S}{2} I_{it} + \frac{S}{2} D_t + \upsilon_t \,. \tag{8}$$

Once again, μ_{it} represents dealer i's conditional expectation of the asset's fundamental value. The baseline half-spread is $S/2$ and the contribution of existing inventory I_{it} to the spread through price shading is $\theta S/2$. Dealer i updates his expectation of the asset's fundamental value in light of the private information revealed by the direction of the previous trade as well as public news:

$$\mu_{it} - \mu_{it-1} = \lambda \frac{S}{2} D_{t-1} + \varepsilon_t \,. \tag{9}$$

Here, $\lambda S/2$ captures the information effect of trade direction and λ is monotonically related to the share of the spread determined by asymmetric information; ε_t is a serially uncorrelated public information shock. Combining these expressions gives

$$\Delta P_{it} = \frac{S}{2}(D_t - D_{t-1}) + \lambda \frac{S}{2} D_{t-1} - \theta \Delta I_{it} + e_t \,, \tag{10}$$

where $e_t = \varepsilon_t + \Delta \upsilon_t$. We follow Huang and Stoll (1997) in estimating the model separately for various size categories. We also disaggregate deals according to counterparty, as in Section 3.

Once again we use generalized method of moments with Newey-West correction for heteroskedasticity. The results are shown in Table 10. We first note that our earlier findings are generally confirmed: as before, spreads are inversely related to deal size and they are wider for commercial customers than for financial customers and other banks. We again find that spreads are little influenced by existing inventory.

The estimates generally indicate that, within size categories, the adverse selection coefficient λ – and thus the share of the adverse selection component of spreads – is smaller for commercial customers than for financial customers. For example, λ is 0.10 on small commercial customer deals but 0.39 on small financial customer deals. Similarly, the estimates generally suggest that the share of the asymmetric information component rises with deal size, after controlling for counterparty type. For example, point estimates of λ for small, medium, and large financial customer deals are 0.39, 0.80, and 1.97, respectively. The main ambiguity here concerns the adverse selection component of spreads on the largest deals. While the point estimates are consistent with the pattern we present, those for large deals are measured imprecisely and are not statistically significant.

In short, these results suggest that the asymmetric information share of spreads is highest for counterparties that are most likely to have private information. We infer that direct costs to dealing with privately informed counterparties are not zero in currency markets and that these costs do affect currency spreads.

6. Conclusions

This paper provides evidence that currency spreads vary inversely with deal size and are wider for commercial customers than for financial customers and other dealers. These results are not predicted by standard models of market making under asymmetric information (Copeland

and Galai, 1983; Glosten and Milgrom, 1985; Easley and O'Hara, 1987). They are, however, consistent with evidence of a negative relationship between spreads and deal sizes in another dealership market, the U.S. market for municipal bonds (Green, Hollifield, and Schurhoff, 2004; Harris and Piwowar, 2004). They are also consistent with the hypothesis advanced in Green, Hollifield, and Schurhoff (2004) that spreads vary positively with a dealer's market power relative to different counterparties. Our results are also consistent with the idea of strategic dealing presented in Gammill (1989) and Leach and Madhavan (1992, 1993), whereby dealers strategically manipulate prices to increase the information content of their deal flow.

We also provide evidence that adverse selection costs are not zero, even in currency markets. As standard models predict, the share of the adverse selection component of spreads varies positively with deal size and is highest for the most informed counterparties. We infer that adverse selection costs are smaller in currency markets than in equity markets and are outweighed by the information value gained through dealing with privately informed counterparties. The information gleaned from this pricing strategy can include the direction and magnitude of current and future trades by large customers as well as the location of major stop-loss orders

The data used in this study comprise the complete USD/EUR trading record of a small bank in Germany over four months in 2001. They have the unusual advantage of distinguishing between financial and commercial customer transactions. To demonstrate that our results should apply to the foreign exchange market overall, despite our bank's small size, we show that other dimensions of our bank's pricing and inventory management practices are congruent with such practices at large banks. This consistency of pricing across dealers makes economic sense, given the intensely competitive nature of currency markets.

The type of information dealers seek by subsidizing informed deal flow is an open question. We provide evidence that deal flow may provide fundamental information even at a small bank, but highlight two reasons to question whether dealers are primarily seeking this type of information. First, currency dealers usually close their positions within a day, while the horizon over which fundamental information is useful for exchange-rate prediction exceeds a year. Second, there may be less private fundamental information in currency markets than in stock markets, since currency fundamentals are mostly macroeconomic aggregates and are therefore known publicly.

We suggest that currency dealers are more interested in information about transitory market developments. Indeed, dealers themselves stress the relative importance of non-fundamental information, when asked. In particular, we highlight the importance of information on large trades. By subsidizing customers who tend to have large currency needs, dealers may increase their chances of learning about large trades and potentially even managing such trades, thereby gaining valuable information. Our analysis brings into focus the need for further research into the nature of information in currency markets.

References

Bank for International Settlements, 2002. Triennial central bank survey, foreign exchange and derivatives market activity in 2001. Basel.

Bjønnes, Geir Høidal and Dagfinn Rime, 2001. Customer trading and information in foreign exchange markets. Working paper, Stockholm Institute for Financial Research.

________________________________, 2003. Dealer behavior and trading systems in foreign exchange markets. Norges Bank Working Paper, November, Oslo. Forthcoming, Journal of Financial Economics.

Carpenter, Andrew and Jianxin Wang, 2003. Sources of private information in FX trading. Working paper, University of New South Wales, Australia.

Cheung, Yin-Wong and Menzie D. Chinn, 2001. Currency traders and exchange rate dynamics: a survey of the U.S. market. Journal of International Money and Finance 20, 439-471.

Copeland, Thomas E. and Dan Galai, 1983. Information effects on the bid-ask spread. Journal of Finance 31, 1457-1469.

Easley, David and Maureen O'Hara, 1987. Price trade size and information in securities markets. Journal of Financial Economics 19, 69-90.

Evans, Martin D.D. and Richard K. Lyons, 2002. Order flow and exchange rate dynamics. Journal of Political Economy 110, 170-180.

__________________________________, 2004. Do markets learn about fundamentals from order flow? Mimeo, U.C. Berkeley and Geortown University.

Fan, Mintao and Richard K. Lyons, 2003. Customer trades and extreme events in foreign exchange, in: Paul Mizen (ed.), Monetary history, exchange rates and financial markets: essays in honour of Charles Goodhart, Northampton: Edward Elgar, 160-179.

Gammill, James F., 1989. The organization of financial markets: competitive versus cooperative market mechanisms. Harvard University Mimeo.

Gehrig, Thomas and Lukas Menkhoff, 2003. The use of flow analysis in foreign exchange: exploratory evidence. Journal of International Money and Finance, forthcoming.

Glosten, Larry and Paul Milgrom, 1985. Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents. Journal of Financial Economics 14, 71-100

Goodhart, Charles, Takatoshi Ito, and Richard Payne, 1996. One day in June 1993: a study of the working of the Reuters 2000-2 electronic foreign exchange trading system, in: Frankel, Jeffrey A., Giampaolo Galli, and Alberto Giovannini (eds.), The microstructure of foreign exchange markets, University of Chicago Press, Chicago 1996, 107-179.

Goodhart, Charles, Ryan Love, Richard Payne, and Dagfinn Rime, 2002. Analysis of spread in the Dollar/Euro and Deutsche-Mark/Dollar foreign exchange markets. Economic Policy 17:35, 536-552.

Green, Richard C., Burton Hollifield, and Norman Schurhoff, 2004. Financial intermediateion and the costs of trading in an opaque market. Mimeo, Graduate School of Industrial Administration, Carnegie Mellon University.

Harris, Lawrence E. and Michael S. Piwowar, 2004. Municipal bond liquidity. Working paper, Securities and Exchange Commission.

Ho, Thomas S.Y. and Hans R. Stoll, 1981. Optimal dealer pricing under transactions and return uncertainty. Journal of Financial Economics 9, 47-73.

Huang, Roger D. and Hans R. Stoll, 1997. The components of the bid-ask spread: a general approach. Review of Financial Studies 10, 995-1034.

Killeen, William P., Richard K. Lyons, and Michael Moore, 2001. Fixed versus flexible: lessons from EMS order flow. NBER Working Paper, No.8491.

Leach, J. Chris and Ananth N. Madhavan, 1992, Intertemporal price discovery by market makers: Active versus passive learning, Journal of Financial Intermediation 2: 207-235.

__________________________________, 1993. Price experimentation and security market structure. Review of Financial Studies 6, 375-404.

Lyons, Richard K., 1995. Tests of microstructural hypotheses in the foreign exchange market. Journal of Financial Economics 39, 321-351.

________________, 2001. The microstructure approach to exchange rates. Cambridge: MIT Press.

MacDonald, Ronald, 2000. Expectations formation and risk in three financial markets: surveying what the surveys say. Journal of Economic Surveys 14: 69-100.

Madhavan, Ananth N. and Seymour Smidt, 1991. A Bayesian model of intraday specialist pricing. Journal of Financial Economics 30, 99-134.

Meese, Richard, and Kenneth Rogoff, 1983, Empirical exchange-rate models of the seventies: do they fit out of sample? Journal of International Economics 14: 3-24.

Melvin, Michael and Lin Wen, 2003. The choice of direct dealing or electronic brokerage in foreign exchange trading. Mimeo, ASU.

Morris and Shin, 2003. Liquidity black holes. Review of Finance 8: 1-18.

Oberlechner, Thomas, The psychology of the foreign exchange market, John Wiley & Sons.

Osler, Carol L., 1998. Stop-loss orders and price cascades in currency markets, Federal Reserve Bank of New York Staff Paper No. 150.

____________, 2003. Currency orders and exchange-rate dynamics: an explanation for the predictive success of technical analysis. Journal of Finance 58: 1791-1819.

Peterson, Mark A. and Erik R. Sirri, 2003. Order preferencing and market quality on U.S. equity exchanges, Review of Financial Studies 16:2, 385-415.

Yao, Jian M., 1998. Market making in the interbank foreign exchange market. Working paper S-98-3, Stern School of Business, New York University.

Appendix

While writing this paper we corresponded frequently with currency market participants. There was no question among our correspondents that our broad conclusions accurately represent the cross-sectional pattern of currency spreads. We provide comments from two individuals.

Peter Nielsen is currently Global Head of Foreign Exchange, Currency Options, Equities & Futures at the Royal Bank of Scotland, the world's largest dealer in U.K. pounds and one of the larger foreign exchange dealing banks overall. He states:

> "Large customers tend to get better prices than smaller customers as they generally have more banking relationships, thereby providing a greater facility for price discovery than smaller customers who may only have one banking counterparty. In addition, in general, larger transactions are quoted with tighter spreads than smaller transactions, although the large customers tend to receive best pricing for all business due to the buying power associated with their overall size and volume of business." (Personal correspondence, April 8, 2004)

William Clyde, Ph.D., was Vice President and Manager of overnight trading at First Chicago Corp. and is now Professor of Finance at Quinnipiac University. He states:

> Banks will want to make good quotes on large, potentially information-bearing amounts for two reasons. First, it gets them better access to the current information: in addition to getting the directional information won by being dealt on, the caller will sometimes share a little additional information with the bank. With this information you don't get caught out and you can make better trading decisions. Second, it ensures that institutions with large amounts continue to call whenever they have something going on.
>
> Small trades, no matter what the source, do not contain much information. They are valuable only for either relationship building (which could result in very tight spreads – I've even quoted zero spreads on small trades to important relationships), or as sources of profit due to large spreads. In fact, it is common for someone asking for a price on a small trade to 'give up their side' and only ask for the bid or the offer (the one they want), in which case the spread implied by the price could be quite large without actually being quoted as a large spread.
>
> Financial customers tend to get better spreads because their trades reflect their view of the market, and their views are often shared with other asset managers. So when you see a lot of financial institutions doing one thing you're sometimes getting a sense of a broad opinion. With corporates you're just seeing their core business activities – car building or whatever. Almost all of them will tell you 'we're not in the business of speculating.' And the trades they're executing now don't tell you much about what other corporates are doing because their current trades reflect business deals done a long time ago, driven by lots of different things. (Personal correspondence, August 18, 2004)

Table 1. Descriptive statistics, currency dealing at a small bank in Germany

The table shows the complete USD/EUR trading activity of a small bank in Germany, except preferred customer deals, over the 87 trading days between July 11th, 2001 and November 9th, 2001.

A. All Business

	All Transactions	Interbank	Customer
All Transactions			
Number of Transactions	3,609	1,919	1,690
(percent)	(100)	(44)	(56)
Of Which, Forward	646	114	532
Value of deals (EUR mil.)	4,335	2,726	1,609
(percent)	(100)	(61)	(39)
Of Which, Forward	999	87	912
Mean Size (EUR mil.)	1.20	1.42	0.95
Mean Size, Forwards (EUR mil.)	1.55	0.76	1.71

B. Customer Business

	Customer Transactions	Financial Customers	Commercial Customers
All Transactions			
Number of Transactions	1,690	171	1,519
(percent)	(100)	(7)	(61)
Of Which, Forwards	532	60	472
Value of Transactions (EUR mil.)	1,609	405	1,204
(percent)	(100)	(23)	(69)
Of Which, Forwards	912	226	686
Mean Size (EUR mil.)	0.95	2.37	0.79
Mean Size, Forwards (EUR mil.)	1.71	3.77	1.45

Table 2. Comparison of small bank studied here with larger banks studied in other papers. The table shows the complete USD/EUR trading activity of a small bank in Germany, except preferred customer trades, over the 87 trading days between July 11th, 2001 and November 9th, 2001. For comparison purposes we focus on statistics based exclusively on the small bank's spot deals.

	Small Bank in Germany	B.I.S. (2002) per Bank	Lyons (1995)	Yao (1998)	Bjønnes and Rime (2004)			Carpenter and Wang (2003), AUD/USD Dealer
					Four Dealers, Range	DEM/USD Dealer	NOK/DEM Dealer	
	87 Trading Days in 2001[a]	April 2001	5 Trading Days in 1992	25 Trading Days in 1995	5 Trading Days in 1998			45 Trading Days in 2002
Transactions per Day	**40 (51)**	---	267	181	58 - 198	198	58	203
Transaction value per Day (in USD millions)	**39 (52)**	50 - 150	1,200	1,529	142 - 443	443	270	213
Value per Transaction (USD mil.)	**1.0**	---	4.5	8.4	1.6 - 4.6	2.2	4.6	1.1
Customer Share of Transaction value (in percent)	**23 (39)**	33	0	14	0 – 18	3	18	11
Average Inventory Level (in EUR or USD millions)	**3.4**		11.3	11.0	1.3 – 8.6	4.2	8.6	---
Average Transaction Size (in EUR or USD millions)	**1.2**		3.8	9.3	1.5 – 3.7	1.8	3.7	---
Average Price Change Btwn. Transactions (in pips)	**11**		3	5	5 - 12	5	12	---

[a] Values in parentheses refer to the data set including outright-forward transactions.

Table 3. Size distribution of individual deals

The table shows the size distribution of all USD/EUR spot and forward transactions, except those for preferred customers, at a small bank in Germany over the period July 11, 2001 through November 9, 2001.

	Interbank Trades	Financial Customer Trades	Commercial Customer Trades
Number	1,872	171	1,492
Share (%)			
1 below EUR 0.1 million	7	15	54
2 EUR 0.1 – 0.5 million	9	26	32
3 EUR 0.5 – 1.0 million	7	14	5
4 EUR 1.0 – 20 million	77	44	8
5 EUR 20 million and above	0	1	1

Table 4. Baseline Lyons-Madhavan-Smidt model

The estimations below correspond to the following equation:

$$\Delta P_{it} = \alpha + \beta_1 D_t + \beta_2 D_{t-1} + \gamma_1 I_{it} + \gamma_2 I_{it-1} + \delta Q_{jt} + \varepsilon_t.$$

The dependent variable is the change in price between two successive incoming trades measured in pips. Q_{jt} is order flow measured in EUR millions, I_{it} is the dealer's inventory at time *t,* and D_t is an indicator variable picking up the direction of the trade, positive for purchases (at the ask) and negative for sales (at the bid). In Panel B these variables are interacted with dummy variables for the two counterparty groups, other dealers (*IB*) and all customers (*CU*). Data include all incoming USD/EUR spot and forward trades of a small bank in Germany, except those with preferred customers, over the period July 11, 2001 through November 9, 2001. Estimation uses GMM and Newey-West correction. Significance at the 1, 5 and 10 percent levels indicated by ‡, † and *, respectively.

A: All Incoming Deals	**Baseline Regression**		**Robustness 1: No Inventories**	**Robustness 2: Spot Trades Only**
	Coefficient	**Standard Error**	**Coefficient**	**Coefficient**
Constant	-0.599†	*0.23*	-0.477†	-0.519*
Direction				
D_t	7.377‡	*0.39*	7.250‡	6.613‡
D_{t-1}	-5.347‡	*0.34*	-5.300‡	-4.866‡
Inventory				
I_{it}	0.503*	*0.26*		-0.035
I_{it-1}	-0.595	*0.26*		-0.018
Deal Size, Q_{jt}	0.522	*0.27*	-0.343‡	-0.339†
Adjusted R^2	0.13		0.13	0.13
Observations	2,859		2,859	2,219

B: Interbank vs. Customer Deals	**Baseline Regression**		**Robustness 1: No Inventories**	**Robustness 2: Spot Trades Only**
	Coefficient	**Standard Error**	**Coefficient**	**Coefficient**
Constant	-0.498†	*0.22*	-0.376*	-0.458*
Direction				
$IB \times D_t$	2.783‡	*0.69*	2.728‡	0.717
$IB \times D_{t-1}$	-1.524‡	*0.48*	-1.513‡	-1.207†
$CU \times D_t$	10.946‡	*0.47*	10.887‡	11.155‡
$CU \times D_{t-1}$	-9.065‡	*0.44*	-9.048‡	-8.996‡
Inventory				
$IB \times I_{it}$	-0.253	*0.35*		-0.202
$IB \times I_{it-1}$	0.142	*0.35*		0.191
$CU \times I_{it}$	1.072‡	*0.41*		-0.053
$CU \times I_{it-1}$	-1.154‡	*0.41*		-0.016
Deal Size				
$IB \times Q_{jt}$	-0.144	*0.40*	0.064	0.543
$CU \times Q_{jt}$	0.748*	*0.42*	-0.353‡	-0.335†
Adjusted R^2	0.23		0.23	0.22
Observations	2,859		2,859	2,219

Table 5. Mean reversion in dealer's EUR inventory.
The table shows results for the following inventory change regression,

$$I_{it} - I_{it-1} = \varpi + \rho I_{it-1} + \varepsilon_t ,$$

where I_{it} represents the dealer's inventory level in EUR millions. Data come from a small bank in Germany and include all USD/EUR spot and forward deals, except those with preferred customers, over the period July 11, 2001, through November 9, 2001. Significance at the 1 percent levels indicated by ‡.

Sample:	ρ	Standard Error	Mean Half-Lives (Minutes)	Median Half-Lives (Minutes)	Obs.
All Deals	-0.17‡	0.006	83	24	3,534
Incoming Deals	-0.20‡	0.008	66	19	2,858
Incoming Deals With:					
Other Banks	-0.27‡	0.029	48	14	1,195
Customers	-0.33‡	0.020	38	11	1,663
Commercial Customers	-0.35‡	0.044	35	10	1,492
Financial Customers	-0.47‡	0.107	24	7	170

Table 6. Probit regression of choice of incoming versus outgoing deal

Probit regression of incoming/outgoing interbank trade decision. Incoming interbank deals are coded 0, while outgoing interbank deals are coded 1. *I* represents inventories, in millions of euros ; $|Q_{jt}|$ represents the absolute size of the current deal, measured in EUR millions; $|Q_{jt-1}|$X*IB* represents the absolute value of the previous transaction if it was an interbank deal; $|Q_{jt-1}|$X*FC* represents the absolute size of the previous deal if it was with a financial customer; $|Q_{jt-1}|$X*CC* represents the absolute value of the previous deal if it was with a commercial customer. Data come from a small bank in Germany and include all USD/EUR spot and forward deals, except those with preferred customers, over the period July 11, 2001, through November 9, 2001; 3,534 observations. Significance at the 1, 5 and 10 percent levels indicated by ‡, † and *, respectively.

$$Prob(Trade_t=IB^{out}) = P(|I_{it}|, I_{it}^2, |Q_{jt}|, IB|Q_{jt-1}|, FC|Q_{jt-1}|, CC|Q_{jt-1}|)$$

	Coefficient	Standard Error	*z*-Statistic
Constant	-1.117‡	*0.04*	-28.25
$\|I_{it}\|$	0.032‡	*0.01*	3.21
I_{it}^2	-0.001‡	*0.00*	-2.91
$\|Q_{jt}\|$	0.032‡	*0.01*	4.14
IB X $\|Q_{jt-1}\|$	0.088‡	*0.02*	3.62
FC X $\|Q_{jt-1}\|$	0.034*	*0.02*	1.84
CC X $\|Q_{jt-1}\|$	-0.004	*0.01*	-0.35
McFadden's R^2		0.01	

Table 7. Spread variation across counterparty types

The estimation below corresponds to the following equation:

$$\Delta P_{it} = \alpha + \beta_1 D_t + \beta_2 D_{t-1} + \gamma_1 I_{it} + \gamma_2 I_{it-1} + \delta Q_{jt} + \varepsilon_t.$$

The dependent variable is the change in price between two successive incoming deals measured in pips. Q_{jt} is order flow measured in EUR millions, I_{it} is the dealer's inventory at time *t,* and D_t is an indicator variable picking up the direction of the deal, positive for purchases (at the ask) and negative for sales (at the bid). These variables are interacted with dummy variables for the three counterparty groups, other dealers (*IB*), financial customers (*FC*), and commercial customers (*CC*). Data include all incoming USD/EUR spot and forward deals of a small bank in Germany, except those with preferred customers, over the period July 11, 2001, through November 9, 2001. Estimation uses GMM and Newey-West correction. Significance at the 1, 5 and 10 percent levels indicated by ‡, † and *, respectively.

	Baseline Regression		**Robustness 1: No Inventories**	**Robustness 2: Spot Trades Only**
	Coefficient	**Standard Error**	**Coefficient**	**Coefficient**
Constant	-0.522†	*0.23*	-0.419*	-0.473*
Direction				
$IB \times D_t$	2.952‡	*0.69*	2.927‡	0.775
$IB \times D_{t-1}$	-1.517‡	*0.48*	-1.520‡	-1.225†
$FC \times D_t$	5.171‡	*1.46*	5.155‡	5.865‡
$FC \times D_{t-1}$	-2.114*	*1.17*	-2.037*	-4.261‡
$CC \times D_t$	11.906‡	*0.52*	11.861‡	11.746‡
$CC \times D_{t-1}$	-10.027‡	*0.49*	-10.027‡	-9.555‡
Inventory				
$IB \times I_{it}$	-0.270	*0.35*		-0.199
$IB \times I_{it-1}$	0.175	*0.35*		0.192
$FC \times I_{it}$	1.297	*1.06*		-0.267
$FC \times I_{it-1}$	-1.422	*1.07*		0.284
$CC \times I_{it}$	1.000†	*0.43*		-0.056
$CC \times I_{it-1}$	-1.079†	*0.43*		-0.017
Deal Size				
$IB \times Q_{jt}$	-0.206	*0.40*	0.021	0.579
$FC \times Q_{jt}$	1.179	*1.08*	-0.180	-0.258
$CC \times Q_{jt}$	0.679	*0.44*	-0.347‡	-0.337†
Adjusted R^2	0.23		0.24	0.23
Observations	2,859		2,859	2,219

Table 8. Spread variation across counterparties and trade sizes

We estimate this equation: $\Delta P_{it} = \alpha + \beta_1 D_t + \beta_2 D_{t-1} + \gamma_1 I_{it} + \gamma_2 I_{it-1} + \delta Q_{jt} + \varepsilon_t$.

ΔP_{it} is the change in price between two successive incoming trades measured in pips. Q_{jt} is deal size and I_{it} is the dealer's inventory, both measured in EUR millions. D_t is +1 for buy-initiated deals and –1 for sell-initiated deals. These variables are interacted with dummy variables for interbank transactions (*IB*), transactions with financial customers (*FC*), and transactions with commercial customers (*CC*). They are also interacted with dummies for deal size: $Lg. = \{Q_{jt} \in [1,\infty)\}$; $Med. = \{Q_{jt} \in [0.5,1)\}$; $Sm. = \{Q_{jt} \in (0,0.5)\}$. Data include all incoming USD/EUR spot and forward deals of a small bank in Germany, except those with preferred customers, over the period July 11, 2001, through November 9, 2001. Estimation uses GMM and Newey-West correction. Significance at 1, 5 and 10 percent levels indicated by ‡, † and *, respectively.

	Baseline Regression		Robustness 1: No Inventories	Robustness 2: Spot Trades
	Coefficient	**Std. Error**	**Coefficient**	**Coefficient**
Constant	-0.257	*0.22*	-0.188	-0.260
Direction				
$IB \times D_t \times Lg.$	3.869‡	*0.74*	3.896‡	2.600‡
$IB \times D_{t-1} \times Lg.$	**-1.282‡**	*0.51*	**-1.300‡**	**-1.753‡**
$IB \times D_t \times Med.+Sm.$	1.429	*1.22*	1.378	-4.123‡
$IB \times D_{t-1} \times Med.+Sm.$	**-2.883‡**	*1.33*	**-2.856†**	**0.867**
$FC \times D_t \times Lg.$	1.208	*1.94*	1.102	1.219
$FC \times D_{t-1} \times Lg.$	**0.509**	*1.81*	**0.953**	**-1.943**
$FC \times D_t \times Med.$	3.610	*2.46*	3.786	9.638
$FC \times D_{t-1} \times Med.$	**-0.772**	*2.39*	**-0.959**	**-1.049**
$FC \times D_t \times Sm.$	9.247‡	*2.49*	9.180‡	7.638‡
$FC \times D_{t-1} \times Sm.$	**-5.461‡**	*1.67*	**-5.679‡**	**-7.095‡**
$CC \times D_t \times Lg.$	4.142‡	*1.73*	4.372‡	5.376*
$CC \times D_{t-1} \times Lg.$	**-2.947‡**	*1.27*	**-2.669†**	**-3.528***
$CC \times D_t \times Med.$	14.043‡	*1.59*	13.914‡	13.463‡
$CC \times D_{t-1} \times Med.$	**-8.410‡**	*1.64*	**-8.354‡**	**-6.187‡**
$CC \times D_t \times Sm.$	12.811‡	*0.57*	12.821‡	11.815‡
$CC \times D_{t-1} \times Sm.$	**-11.447‡**	*0.57*	**-11.459‡**	**-10.005**
Inventory				
$IB \times I_{it}$	-0.223	*0.36*		-0.182
$IB \times I_{it-1}$	0.163	*0.36*		0.182
$FC \times I_{it} \times$	0.996	*0.96*		-0.238
$FC \times I_{it-1}$	-1.135	*0.96*		0.260
$CC \times I_{it}$	**0.993†**	*0.41*		0.003
$CC \times I_{it-1}$	**-1.013†**	*0.41*		-0.024
Deal Size				
$IB \times Q_{jt}$	-0.448	*0.39*	-0.263	-0.091
$FC \times Q_{jt}$	1.085	*0.98*	0.027	-0.022
$CC \times Q_{jt}$	**0.916†**	*0.46*	-0.082	-0.088
Adjusted R^2	0.24		0.24	0.24
Observations	2,859		2,859	2,219

Table 9. Tests of cointegration between exchange rates and cumulative deal flow

Table reports ordinary least squares estimates of the following cointegrating relationship between exchange rates and cumulative deal flow:

$$P_{it} = \omega_i + \phi_i trend_t + \kappa_i \sum_{j=0}^{t} CumulativeDealFlow_{ij} + \nu_{it},$$

where i represents the counterparty type, $i \in \{IB_{Incoming}, CC_{All}, FC_{All}, FC_{Small}, FC_{Med\text{-}Lg}\}$. Preliminary statistical tests indicate that the variables are not stationary, so t-values on the coefficients are not reliable and are not reported. ADF-test is a standard augmented Dickey-Fuller test on the regression residuals. PP-test is a Phillips-Perron test on the regression residuals. The number of lags included is calculated from the sample size (Newey-West automatic truncation lag selection). The tests do not include a constant since a constant is included in the original regression equation. Significance at the 1, 5 and 10 percent levels is indicated by ‡, † and *, respectively. Flow and trend coefficients are multiplied by 10^3.

	Commercial Customers	Incoming Interbank	Financial Customers: All Deals	Financial Customers: Small Deals	Financial Customers: Med. & Lg. Deals
Constant	0.884	0.871	0.884	0.875	0.891
Cumulative Deal Flow	-0.291	0.417	0.152	4.330	0.255
Trend	0.008	0.010	0.167	0.553	0.167
ADF-test	(-3.02)‡	(-2.38)†	(-2.31)†	(-1.58)	(-2.00)†
PP-test	(-3.77)‡	(-2.77)‡	(-2.45)†	(-2.58)†	(-2.66)‡
Observations	1,492	1,269	171	70	101

Table 10. Modified Huang and Stoll (1997) model

We estimate this model: $\Delta P_{it} = \frac{S}{2}(D_t - D_{t-1}) + \lambda \frac{S}{2} D_{t-1} - \theta \frac{S}{2} \Delta I_{it} + e_t$.

ΔP_{it} is the change in price between two successive incoming trades measured in pips. I_{it} is the dealer's inventory, measured in EUR millions. D_t is +1 for buy-initiated trades and –1 for sell-initiated trades. These variables are interacted with dummy variables for interbank trades (*IB*), trades with financial customers (*FC*), and trades with commercial customers (*CC*). They are also interacted with dummies for trade size: *Lg*. = $\{|Q_{jt}| \in [1,\infty)\}$; *Med*. = $\{|Q_{jt}| \in [0.5,1)\}$; *Sm*. = $\{|Q_{jt}| \in (0,0.5)\}$. Data include all incoming USD/EUR spot and forward trades of a small bank in Germany, except those with preferred customers, over the period July 11, 2001, through November 9, 2001. Estimation uses GMM and Newey-West correction. Significance at 1, 5 and 10 percent levels indicated by ‡, † and *, respectively. Constant term suppressed.

	Baseline Regression		Robustness 1: No Inventories	Robustness 2: Spot Trades Only
	Coefficient	Std. Error	Coefficient	Coefficient
Half-Spread, *S*/2				
S/2 X *IB* X *Lg.*	3.934‡	*0.75*	3.426‡	2.403‡
S/2 X *IB* X *Med. +Sm.*	0.817	*1.31*	1.138	-4.438‡
S/2 X *FC* X *Lg.*	1.597	*1.88*	1.690	1.642
S/2 X *FC* X *Med.*	4.918†	*2.31*	3.994	8.273
S/2 X *FC* X *Sm.*	9.304‡	*2.44*	9.424‡	7.850‡
S/2 X *CC* X *Lg.*	4.478‡	*1.65*	3.597‡	5.592†
S/2 X *CC* X *Med.*	12.963‡	*2.60*	13.767‡	13.046‡
S/2 X *CC* X *Sm.*	12.805‡	*0.57*	12.774‡	11.698‡
Adverse Selection				
λ X *IB* X *Lg.*	0.717‡	*0.13*	0.647‡	0.287
λ X *IB* X *Med. +Sm.*	-2.729	*5.57*	-1.395	0.757†
λ X *FC* X *Lg.*	1.965	*1.80*	1.685	-0.195
λ X *FC* X *Med.*	0.802*	*0.46*	0.749	0.825‡
λ X *FC* X *Sm.*	0.391†	*0.19*	0.385†	0.102
λ X *CC* X *Lg.*	0.364	*0.32*	0.293	0.446
λ X *CC* X *Med.*	0.348†	*0.15*	0.384‡	0.509‡
λ X *CC* X *Sm.*	0.101‡	*0.02*	0.094‡	0.136‡
Inventory				
θ X *IB* X *Lg.*	-0.077	*0.06*		0.077
θ X *IB* X *Med. +Sm.*	4.814	*8.83*		0.006
θ X *FC* X *Lg.*	0.003	*0.12*		0.130
θ X *FC* X *Med.*	-0.433	*0.42*		0.146
θ X *FC* X *Sm.*	-0.046	*0.20*		0.066
θ X *CC* X *Lg.*	-0.021*	*0.01*		-0.016
θ X *CC* X *Med.*	0.090	*0.23*		-0.002
θ X *CC* X *Sm.*	-0.072*	*0.04*		-0.002
Adjusted R^2	0.23		0.25	0.24
Observations	2,859		2,859	2,219

Figure 1. Overall inventory position (EUR millions)

Plot shows the evolution of a currency dealer's inventory position in EUR millions over the period July 11, 2001 through November 9, 2001. Data come from a small bank in Germany and include all USD/EUR spot and forward trades, except those with preferred customers. The horizontal axis is transaction-time. Vertical lines indicate the end of each calendar week.