PRICE DISCOVERY IN CURRENCY MARKETS

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

Abstract

This paper makes three contributions to our understanding of the price discovery process in currency markets. First, it provides evidence that price discovery cannot follow the familiar process based on adverse selection in customer transactions, since customer spreads are inversely related to a trade's likely information content. Second, the paper highlights three hypotheses from the literature that may explain the pattern of customer spreads, two of which rely on asymmetric information. Third, the paper suggests an alternative price discovery process for currencies that incorporates the pivotal distinction between the customer market and the interdealer market, and provides preliminary evidence for that process. [*JEL F31, G14, G15. Keywords: Bid-ask spreads, foreign exchange, asymmetric information, microstructure, price discovery, interdealer, inventory, market order, limit order]*

October 2007

Corresponding author: Carol Osler, cosler@brandeis.edu or Brandeis International Business School, Brandeis University, Mailstop 32, Waltham, MA 02454, USA. Tel. (781) 736-4826. Fax (781) 736-2269. We are deeply grateful to the bankers who provided the data and to William Clyde, Pete Eggleston, Keith Henthorn, Valerie Krauss, Peter Nielsen, Peter Tordo, and other bankers who discussed dealing with us. Special thanks go to Clyde, Eggleston, and Tordo, for reading and commenting on drafts of the paper. We thank, without implicating, Geir Bjønnes, Alain Chaboud, Yin-Wong Cheung, Joel Hasbrouck, Thomas Gehrig, Michael Goldstein, Rich Lyons, Albert Menkveld, Bruce Mizrach, Anthony Neuberger, Paolo Pasquariello, Uday Rajan, Tarun Ramadorai, Stefan Reitz, Dagfinn Rime, Gideon Saar, Erik Theissen, Dan Weaver, and participants at the Second Annual Microstructure Conference in Ottawa and at the NBER microstructure meeting for insightful comments.

Price Discovery in Currency Markets

This paper makes three contributions to our understanding of the price discovery process in currency markets. First, it provides evidence that price discovery cannot follow the familiar process based on adverse selection in customer transactions, since customer spreads are inversely related to a trade's likely information content. Second, the paper highlights three hypotheses from the literature that may explain the pattern of customer spreads, two of which rely on asymmetric information. Third, the paper suggests an alternative price discovery process for currencies that incorporates the pivotal distinction between the customer market and the interdealer market, and provides preliminary evidence for that process.

The dominant model of price discovery in finance focuses on the information advantage customers often have over dealers (Kyle (1985); Glosten and Milgrom (1985)). In one-tier markets, where all trades are between customers and dealers, dealers rationally protect themselves from this adverse selection risk by quoting a spread (Glosten and Milgrom (1985)). Equilibrium bid and ask prices should incorporate the expected information content of each trade and in consequences prices move in the direction required by the customers' information. In this framework, spreads should be higher when customers are more likely to have information, such as when they undertake large trades (Glosten and Milgrom (1985), Easley and O'Hara (1987), Glosten (1989)). If dealing is not anonymous, spreads should also be higher for customers that are typically informed.

Most foreign exchange microstructure papers draw on adverse selection as their primary interpretive framework. The relevance of this framework was initially suggested in Lyons (1995), which shows that, for a particular dealer in 1992, trade size and the interbank spreads of were positively related. The implicit adoption of adverse selection is clear in Marsh and O'Rourke (2005), for example, which estimates Easley, Kiefer, and O'Hara's (1996, 1997) adverse-selection-based measure of private information on daily foreign exchange customer data. Similarly, Payne (2003) estimates a VAR decomposition of interdealer trades and quotes and interprets the results, following Hasbrouck (1991), through the lens of adverse selection.

Our evidence indicates, however, that the behavior of customer spreads in foreign exchange is inconsistent with adverse selection. This holds whether we estimate the Huang and Stoll (1997) model or the Madhavan-Smidt model (1991). Among other violations of the theory, we find that customer spreads are widest for the trades least likely to carry information. More specifically, customer spreads are inversely related to trade size and are narrower for the customers that dealers consider most informed. Dealers consider financial customers such as hedge funds and other asset managers to be more informed than the other broad category of customers, firms that participate in international trade ("commercial firms").[1]

If adverse selection doesn't drive customer spreads in foreign exchange, what does? The paper's second contribution is to highlight three factors taken from the broader microstructure literature that seem likely to be important. The first factor, fixed operating costs, can explain the negative relation between trade size and customer spreads, but cannot explain variation across customer types. The customer-based variation could be explained by two hypotheses that involve asymmetric information. The first of these suggests that firms in dealership markets gain transitory "market power" from their information about current market conditions. It can be costly for customer firms to find the best quotes, so individual dealers has market power during the moment of communicating with the customer, even in a market with hundreds of competitors. In foreign exchange, this means dealers may extract wider spreads from commercial than financial customers, since commercial customers typically know less about market conditions. The second hypothesis proposes that dealers strategically vary spreads to optimize information flow, rather than passively accepting the information flow as assumed in adverse-selection models. Building on theoretical results in Naik *et al.* (1999) and evidence that customer order flow does carry information (e.g., Evans and Lyons (2004), Daníelsson *et al.* (2002)), we suggest that foreign exchange dealers strategically set narrower spreads to informed customers to gain information which they exploit in later interdealer trading. We also provide empirical support for this hypothesis.

The paper's third contribution is to outline a process through which information may ultimately become embedded in exchange rates. The proposed mechanism focuses on dealers' order choice in the interdealer market when they unload inventory accumulated in customer trades. After trading with an informed customer, a dealer's information and inventories provide strong incentives to place market orders. [2] An informed-customer buy thus tends to trigger market buys in the interdealer market and thus higher interdealer exchange rates. In this way the information brought to the market by informed customers generates information-consistent changes in interdealer prices. By contrast, after trading with an uninformed customer a dealer has only weak incentives to place market orders. Thus dealer transactions with uninformed customers may be most likely to generate limit orders and liquidity in the interdealer market.

Our view of dealer behavior predicts a number of the key stylized facts in foreign exchange microstructure. First, it predicts the positive relation between interdealer order flow and exchange-rate returns documented in Lyons (1995), Payne (2003), Evans (2002), Evans and Lyons (2002), and Daníelsson *et al.* (2002), *inter alia*. If dealers respond to fundamental information, it also predicts that the relation should be substantially permanent, consistent with evidence presented in Killeen *et al.* (2006) and Bjønnes *et al.* (2005). In addition, our view of dealer behavior predicts the positive relation between exchange rates and financial order flow documented in Evans and Lyons (2004), Bjønnes *et al.* (2005), and Marsh and O'Rourke (2005). Finally, our view predicts that the response of exchange rates to financial order flow is substantially permanent, consistent with evidence in Lyons (2001) and Bjønnes and Rime (2005).

We test two further implications of our proposed price discovery mechanism. First, dealers should be more likely to make outgoing trades after financial-customer trades than after commercial-customer trades. Second, dealers should be more likely to make outgoing trades after large incoming trades than after small ones, controlling for inventory level. The evidence provides support for both implications.

The paper’s message is potentially broader than the foreign exchange market. We essentially suggest that the price discovery process in any market depends on the market’s structure. The standard adverse selection process assumes a one-tier market, and thus it may not be relevant in markets like foreign exchange with active interdealer trading. Our proposed price discovery process may thus be relevant to other liquid two-tier markets, such as the U.S. Treasury market, the U.S. corporate bond market, and the

London Stock Exchange. Evidence shows that the relation between customer spreads and trade size in those markets fits the pattern we identify for currency markets.

Though our analysis concerns the trading process it is also potentially relevant to the exchange-rate dynamics literature within international macroeconomics – which is analogous to the "asset pricing" literature within finance. A critical empirical development in exchange-rate economics has been the recognition that order flow strongly influences returns (Evans and Lyons 2002). Theory has shown that order flow matters in part because it contains information (e.g., Evans and Lyons 2002, 2004), but has not articulated the mechanism through which the customers' information gets into to exchange rates. Our proposed price discovery framework and supporting empirical evidence help fill that gap.

Our data comprise the entire USD/EUR transaction record of a single dealer at a bank in Germany during four months in 2001. These data have two advantages relative to most other tick-by-tick transactions datasets in foreign exchange: (i) they distinguish between financial and commercial transactions, and (ii) they cover a longer time period.

The rest of the paper has four sections and a conclusion. Section I describes the foreign exchange market and our data. Section II shows that customer spreads in foreign exchange do not conform to the predictions of adverse-selection theory, in part because they are narrowest for the trades most likely to carry information. Section III discusses how this pattern may be explained by operating costs, market power, and strategic dealing, and provides empirical support for the strategic dealing hypothesis. Section IV articulates our proposed price discovery process for currency markets and presents supporting evidence. Section V concludes.

I. FOREIGN EXCHANGE MARKET STRUCTURE AND DATA

The foreign exchange market is geographically dispersed around the world so trading takes place almost round the clock. As documented in the most recent (2004) survey from the Bank for International Settlements (B.I.S.), there are roughly two thousand dealers in total who compete for business in spot, forward, swap, and derivative contracts. Dealers can be found in almost every country, but the three major trading centers – Tokyo, London, and New York – account for roughly half of each day's trading. Trading is relatively concentrated across currencies: the euro, which we study here, is involved in almost one fifth of all trades. Trading is not highly concentrated across dealers, however: the market share of the top five firms was recently estimated at 54 percent (Euromoney 2007). Even banks outside the top five are quite active since the market is immense – daily trading was last formally estimated at almost $2 trillion (B.I.S. 2004) and has reportedly grown rapidly since (Euromoney 2007). The foreign exchange market – or the "FX" market, as it is called by market participants – is intensely competitive. Banks compete in terms of spreads, of course, and also in terms of transaction speed, pricing consistency, trading strategies and ideas, electronic products, and overall relationship quality (Euromoney 2007).[3]

In contrast to equity markets, where individuals can account for half or more of all trading, trading by individuals is almost invisible in foreign exchange, and the overwhelming majority of trading is carried out by institutions.[4] Since currencies are important in commerce as well as finance, the institutional customer base for foreign exchange includes non-financial firms as well as financial firms. Foreign exchange dealers can be accurately viewed as the only intraday suppliers of liquidity and there is no need to consider "latent liquidity" (Chacko *et al.* 2006) in part because, during our sample period, access to the brokerages was restricted to dealing banks. More broadly, by their nature banks create deposits, so trading in the major currency pairs is never constrained by supply.

The customer market in foreign exchange is quote-driven. The interdealer market includes both "direct" trading in the quote-driven market and "indirect" trading on order-driven exchanges. The share of interdealer trading has dropped markedly in recent years, possibly in response to the increasing dominance of electronic brokerages.[5] Trading is now split almost evenly between the interdealer and customer markets (B.I.S. 2004).

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. Like all significant foreign exchange dealers, the bank offers the full range of foreign exchange products, including forwards and derivatives; it serves the full range of customers, including financial as well as commercial customers; and it participates in the interdealer market as both liquidity provider and liquidity demander. While the bank desires anonymity, we can state that the bank was included among the banks rated in Euromoney's annual FX poll (2007), which places it among the top ten percent of all dealing banks worldwide. Furthermore, in that same poll the bank was among those considered "best" in euro-dollar trading. Nonetheless, the bank is not among the ten immense banks that top those lists.

The data technically refer to the overall bank, but 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;[6] (2) the direction (counterparty buys or sells); (3) the quantity; (4) the transaction price; (5) the type of counterparty – dealing bank, financial customer, commercial customer, preferred customer; (6) the initiator; and (7) the forward points if applicable. Table I provides basic descriptive statistics.[7]

We include outright forward trades, adjusted to a spot-comparable basis by the forward points, as recommended by Lyons (2001). 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). Like all foreign exchange dealers, ours manages his own inventory (subject to position and loss limits);[8] we infer this inventory by cumulating successive transactions.[9] Following Lyons (1995), we set the daily starting position at zero. This should not introduce significant distortions since our dealer, like most FX dealers, keeps his inventory quite close to zero. Figure 1 plots the complete inventory record. The dealer's average inventory position is EUR 3.4 million during the trading day and only EUR 1.0 million at the end of the day.

Our sample period includes September 11, 2001, but that day's events do not materially affect our conclusions. It may surprise those familiar with the extreme disruption of U.S. money, bond, and stock markets around that time to learn that the foreign exchange market functioned smoothly overall. Volatility, trading volume, and spreads were high on September 11, 2001 itself, but this would be expected in an efficient market subject to major news shocks. Furthermore, volatility and spreads were indistinguishable from their normal levels by the next day (Mende 2006). This resilience no doubt stems in part from the large number of dealers and their wide geographical dispersion – dealers are even dispersed within New York City. In addition, Reuters and EBS, the two major electronic brokerages, both have servers in multiple locations around the world performing real-time replication of all functions.

A preliminary comparison of our dealer with other dealers described in the literature is provided in Table II. Table III provides information on the size distribution of our dealer's transactions. (Statistics on the spreads themselves cannot be calculated from transactions data.) In terms of daily trading value, average transactions per day, average inventory position, and mean absolute price change between transactions our dealer is comparable in size to a NOK/DEM dealer at the large dealing bank examined in Bjønnes and Rime (2004), and more generally it is about average sized (B.I.S. 2002). Since it is less specialized than the dealer considered in Lyons (1995), it is more likely to be representative.

Our main qualitative conclusions should generalize to the entire foreign exchange market, including the biggest banks, because foreign exchange is an intensely competitive market and agents in competitive markets are price takers. In FX, the product is liquidity and the price is the bid-ask spread. Dealers confirm that they are price takers with respect to spreads when they claim, in surveys, that the primary factor they consider in setting spreads is the spreads' conventional level (Cheung and Chinn 2001). Logically, therefore, the behavior of any (successful) dealer should accurately represent the behavior of other (successful) dealers.

To underscore the representativeness of our bank, we show that its behavior is broadly consistent with that of the biggest banks in many well-studied dimensions. Appendix A provides a detailed comparison of our bank's pricing and inventory management practices with those of the immense banks analyzed earlier. This analysis suggests that the following statements about immense dealers are equally true for our average-sized dealer:

- The baseline spread for interbank trades is on the order of two pips (or equivalently two ticks) [10]
- The baseline spread for customer trades is a few times larger than the spread on interbank trades
- Existing inventories are not statistically related to quoted prices

- The dealer typically brings his inventory back to zero by the end of the trading day
- The dealer tends to bring inventory back to zero in a matter of minutes, a speed that is comparable with that of futures traders and extremely fast relative to traders in equity and bond markets.

II. THE CROSS-SECTIONAL PATTERN OF CURRENCY SPREADS

This section evaluates whether the cross-sectional pattern of customer spreads in foreign exchange is consistent with the implications of adverse selection. After presenting some simple descriptive results we analyze two structural models of spreads, the Huang and Stoll model (1997) and the Madhavan and Smidt model (1991). We find no consistency between the predictions of adverse selection and the behavior of customer spreads.

A. Preliminary Analysis

We extract measures of spreads from a statistical analysis of successive price changes. Consider a simple market where everyone pays the same spread and the spread never changes. If the market price is stable then prices only change if trading moves from the bid to the ask or vice versa so the spread can be estimated from the price changes. In the absence of price stability, the same conclusion still applies so long as there is no dominant trend.[11]

We begin our analysis with crude estimates of how price changes are related to trade size and counterparty type. These preliminary estimates are based on the following equation:

$$\Delta P_{it} = \alpha + \beta(D_t - D_{t-1}) + \eta_t \quad . \tag{1}$$

P_t is the price and $\Delta P_t \equiv P_t - P_{t-1}$ is measured in pips. D_t is the direction of trade [$D_t \equiv 1$ (-1) if the counterparty is a buyer (seller)]. The coefficient β should represent half of the spread.[12]

Trade size: Market participants tell us that they informally divide normal-sized customer transactions into three categories: regular trades, which vary from €1 million to about €25 million; modest trades; and tiny trades. Though the line between the latter two categories is ambiguous, their treatment can vary substantially: tiny trades are often spread by formula rather than by dealers' discretion and on such trades a one percent spread is not considered unreasonable. For estimation purposes we distin guish the following size ranges: Large trades: $\{|Q_t| \in [€1 \text{ million}, € 25 \text{ million})\}$; medium trades: $\{|Q_t| \in [€0.5 \text{ million}, €1 \text{ million})\}$; and small trades: $\{|Q_t| \in (€0, € 0.5 \text{ million})\}$. To capture the influence of trade size we interact the change-in-direction with dummies for small (*Sm*), medium (*Md*), and large trades (*Lg*).

We follow standard practice and use generalized method of moments (GMM) with Newey-West correction for heteroskedasticity and autocorrelation (e.g., Yao (1998); Bjønnes and Rime (2005)). Our dependent variable is the sequence of prices on transactions initiated by customers; thus our regression operates on transaction time rather than clock time. We exclude the few transactions over \$25 million because such trades essentially represent a distinct market: customers hire dealers to manage such trades.[13] We also exclude interbank transactions because they may not strictly be comparable with customer trades, given the structural differences between quote- and order-driven markets.

Though adverse selection predicts that trade size and spreads are positively related, the results of this analysis, shown below, suggest they are negatively related:

$$\Delta P_t = \underset{(0.296)}{0.979} + \underset{(0.503)}{8.403}\,(D_t - D_{t-1}) \, x \, Sm + \underset{(1.188)}{6.859}\,(D_t - D_{t-1}) \, x \, Md + \underset{(0.724)}{3.181}\,(D_t - D_{t-1}) \, x \, Lg \qquad Adj\ R^2 = 0.266.$$

The estimated half-spreads, which are all highly significant, fall monotonically as trade size increases, from 8.4 pips for small customer trades to 6.9 pips for medium-sized trades to 3.2 pips for large trades. This overall pattern provides no support for the positive relation predicted by adverse selection; indeed, the reverse relation gets support from Wald tests indicating that two of the three differences are statistically significant.

Customer Type: To examine the bilateral relationship between spreads and customer type we the change-in-direction variable with dummies for trades with financial customers (*FC*) and commercial customers (*CC*). Adverse selection predicts that financial customers pay higher spreads than commercial customers, but the results, shown below, indicate the reverse: the half-spread is 4.0 pips, on average, for financial customers and roughly twice as large, 7.9 pips, for commercial customers:

$$\Delta P_t = \underset{(0.298)}{1.069} + \underset{(1.005)}{4.010}\,(D_t - D_{t-1})\ x\ FC + \underset{(0.454)}{7.878}\,(D_t - D_{t-1})\ x\ CC \qquad Adj\ R^2 = 0.329.$$

As before, the coefficients are strongly statistically significant and the Wald statistic supports the hypothesis that the more informed financial customers pay narrower spreads than the commercial customers.

According to our correspondents at large dealing banks, the correct customer disaggregation is between small commercial customers, on the one hand, and financial customers and large multinational (commercial) corporations, on the other. Though we cannot technically distinguish large multinationals from other commercial customers, large multinationals are unlikely to do much business with our bank. Thus the counterparty-based tiering identified here should be roughly accurate in the present context.

Trade Size and Counterparty Type: Since commercial trades tend to be smaller than financial trades (Table 3), it is possible that the small spreads on commercial trades simply reflect the trades' small size. We test whether counterparty type has an independent influence by interacting the change-in-direction variable with dummies for both trade size and counterparty type. The results of this analysis, reported in Table IV, indicate that currency spreads are influenced by counterparty type as well as trade size. Importantly, the results continue to provide no support for adverse selection. Commercial spreads are larger than financial spreads for both small and medium-sized trades and roughly the same size for large trades. Spreads are not positively related to trade size within customer categories.

B. Structural Analysis

In reality, of course, spreads vary for a number of reasons, so the univariate regression above need not provide reliable estimates of the influence of adverse selection. We next estimate the determinants of spreads using two standard structural models, those of Huang and Stoll (1997) and Madhavan and Smidt (1991). Though each model parameterizes the influence of adverse selection differently, we find no evidence for that influence.

1. The Huang and Stoll Model

Huang and Stoll (1997) observes that trade size is relatively unimportant for pricing in markets ─ like the foreign exchange interbank market ─ where large trades are routinely broken up into multiple smaller transactions. Even in such markets, however, the risk of trading with a better informed counterparty remains. Huang and Stoll's model analyzes the pricing decision of a representative dealer in a competitive market whose counterparties have private information that is revealed by their trade direction of (buy or sell). Agents are fully rational. The model assumes that dealer *i*'s quote is determined by the dealer's expected true value of the asset, μ_{it}, the trade's direction, and the dealer's existing inventory, as follows:

$$P_{it} = \mu_{it} + \frac{S}{2} D_t - \theta \frac{S}{2}\left(I_{it} - I_t^*\right) + \upsilon_t \quad . \tag{2}$$

The baseline half-spread ─ meaning the spread that would apply before adjustment for existing inventories ─ is $S/2$. I_{it} is dealer *i*'s inventory at the beginning of period *t*; I^*_i is his desired inventory, which appears to be zero for our dealer (see Figure 1). The model permits dealers to manage existing inventories by shading prices to customers (e.g., quote lower prices when inventory is high), which implies $\theta > 0$. The term υ_t is a mean-zero random disturbance.

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, ε_t: $\mu_{it} - \mu_{it-1} = (\lambda S/2)D_{t-1} + \varepsilon_t$. The term $\lambda S/2$ captures the information effect of trade direction and is thus a direct manifestation of adverse selection. The public news shock ε_t is a serially uncorrelated. Combining the pricing and updating rules gives the following expression for price changes between customer transactions:

$$\Delta P_{it} = \frac{S}{2}(D_t - D_{t-1}) + \lambda \frac{S}{2} D_{t-1} - \theta \frac{S}{2} \Delta I_{it} + e_t \, , \tag{3}$$

where $e_t \equiv \varepsilon_t + \Delta\upsilon_t$. We follow Huang and Stoll (1997) in estimating separate coefficients for trades in our various size and customer-type categories, which we achieve by interacting the key right-hand-side variables with dummies for both transaction size $\{Sm, Md, Lg\}$ and counterparty type $\{FC, CC\}$. We again use GMM with Newey-West standard errors.

The results, shown in Table V, are again inconsistent with the predictions of adverse selection and instead broadly confirm the reverse pattern highlighted previously. Within customer categories spreads are inversely related to trade size. Point estimates for the baseline commercial customer spreads on small, medium, and large trades, for example, are 13.5 pips, 11.6 pips, and 3.8 pips, respectively. For small and medium trades these values exceed corresponding estimates for financial customers.

Adverse selection also fails to get support from the estimated adverse-selection coefficients, λ. The point estimates of λ imply that information content varies non-monotonically with trade size, while adverse selection predicts a monotonic relation. If we instead take statistical significance as our guide, the estimates of λ imply that small and medium-sized commercial trades carry more information than large commercial trades or equal-sized financial trades, the opposite of what dealers report.

We note in passing that the results suggest that inventory levels are not relevant to FX customer spreads, since none of the six inventory coefficients is significant at standard significance levels. This presumably reflects the general preference among FX dealers for managing inventory via interbank trades (Bjønnes and Rime (2005)), rather than by shading prices to customers.

Table V presents three robustness tests. First, we rerun the regressions excluding inventories, which appear to have no influence. Second, we rerun the regressions using only spot transactions, to facilitate comparisons with earlier papers. Third, we rerun the regressions including interdealer as well as customer trades. This provides comparability with Bjønnes and Rime (2001), where customer transactions (as a single category) and interbank transactions are included in the main regressions. We also carried out two unreported robustness tests. We also ran the regressions excluding September 11, 2001 and (separately) using a different cutoff between small and medium trades, (€300,000 rather than €500,000).[14] None of these alternative specifications suggest any change to our qualitative conclusions.

2. The Madhavan and Smidt Model

FX dealers consistently report that they consider large customer trades to be more informative than small ones, so the Huang and Stoll (1997) model's assumption that trade size is uninformative may not be valid in the FX customer market.[15] Trade size is informative in the Madhavan and Smidt model (1991), which has frequently been applied in FX microstructure research (see, for example, Lyons (1995) and Bjønnes and Rime (2005)). In that model, agent *j* calls dealer *i* requesting a quote and chooses an amount Q_{jt} according to his perceived gap between his expected value of the asset, μ_{jt}, and the price quoted by the dealer: $Q_{jt} = \xi(\mu_{jt} - P_{it}) + X_{jt}$. X_{jt} represents agent *j*'s liquidity demand; it is assumed that $\xi > 0$. In setting his regret-free price, P_{it}, dealer *i* considers his own expected value of the asset, μ_{jt}, his inventory, and the direction of the trade: $P_{it} = \mu_{it} + \zeta(I_{it} - I_i^*) + \chi D_t$. If the dealer shades prices to manage existing inventories, $\zeta < 0$. Dealer *i* rationally considers customer's desired trade size in determining μ_{jt}.

After solving for conditional expectations and taking first differences, one arrives at the following expression for the price change between dealer *i*'s incoming transactions:

$$\Delta P_{it} = \alpha + \beta_1 D_t + \beta_2 D_{t-1} + \gamma_1 I_{it} + \gamma_2 I_{it-1} + \delta Q_{jt} + \eta_t \tag{4}$$

The intercept, α, should be zero if the dealer's desired inventory is zero. If the dealer shades prices in response to inventories then $\gamma_2 > |\gamma_1| > 0 > \gamma_1$. Our estimates of the Huang and Stoll model suggest that both γ_1 and γ_2 should be zero.

Adverse selection, if operative, could influence three parameters. First, it could influence β_2, the coefficient on lagged direction, which according to the model is the negative of the baseline half-spread. Under adverse selection this would be bigger (in absolute value) for large trades and for financial trades. This same effect could also be reflected in δ, the coefficient on trade size: under adverse selection this coefficient should be positive. Large trades can reflect a big gap between the asset's true value and the dealer's quote, so a rational dealer in the model increases the spread with trade size. Unfortunately, the interpretation of a positive coefficient on trade size is inherently ambiguous, since it is observationally equivalent to an inventory effect noted in Ho and Stoll (1981). Larger trades leave market makers with higher inventory and thus greater inventory risk, so larger trades should carry wider spreads.

Finally, adverse selection should influence the relation between β_1 and β_2. The model implies that $\beta_1 = |\beta_2|/\phi > 0 > \beta_2$, where $0 < \phi < 1$ is a model-derived parameter capturing the extent to which dealers rely on their priors rather than the current trade in updating their estimate of the currency's true value. Under adverse selection, estimates of ϕ should be smaller (farther below unity) for large trades and for financial trades, since dealers consider such trades to be relatively informative.

As before, we estimate the model using generalized method of moments (GMM) with Newey-West correction for heteroskedasticity and autocorrelation and we interact the key variables with dummies corresponding to trade size and counterparty type. The results of this analysis, presented in Table VI, provide no more support for adverse selection in the FX customer market than our earlier results.

The estimated baseline half-spreads once again indicate that spreads vary inversely with the likely information content of each trade, the reverse of the prediction from adverse selection. For financial customers, the estimated baseline half-spreads are a statistically significant 6.6 pips on small trades and roughly half that size – and insignificantly different from zero – for medium and large trades. Commercial customer spreads on small and medium trades are estimated to be 12.7 and 7.2 pips, respectively, roughly twice the corresponding estimates for financial customers. The baseline half-spread on large commercial trades are two pips, far below those on smaller trades, and insignificant. Wald tests indicate that financial-customer spreads are indeed smaller than commercial-customer spreads for small and medium trades and that commercial customer spreads are negatively, rather than positively, related to trade size.

The other two potential sources of evidence for adverse selection are equally unsupportive. The point estimates of the coefficient on trade size, δ, are positive for both commercial and financial customers, consistent with adverse selection, but neither is significant. The ratio between the coefficients on lagged and current direction, $\phi = |\beta_2|/\beta_1$, which should vary negatively with trade size under adverse selection, varies positively with deal size for the financial customers, the customers to which the theory most likely applies: for small, medium, and large financial trades the point estimates are 0.63, 0.76, and 1.51, respectively (the coefficients should also be unity or less under adverse selection). Similarly, though adverse selection predicts that ϕ is smaller for financial than commercial customers, the reverse is true for medium-sized and large trades: point estimates for medium and large commercial trades are 0.57 and 0.44, respectively.

The other results from this model are generally unsurprising. The constant term is insignificant, implying that our dealer's preferred inventory level is indeed zero. The coefficients on inventory are significant for commercial customers but have signs opposite those predicted by the model. This can be traced to one particular trade; when that trade is excluded the coefficients become statistically insignificant. The other inventory coefficients are also insignificant, implying that our dealer does not manage inventories by shading prices to customers, consistent with the results from the Huang and Stoll model.

We once again test whether the results are robust to excluding inventories, excluding forward transactions, including interbank trades, excluding September 11, and splitting small and medium trades at €300,000 rather than €500,000. In addition, we test whether the results are affected by excluding the quantity variable. None of these alternative specifications, most of which are reported in the table, bring noticeable changes to the results.[16]

C. Discussion

This section has provided evidence that price discovery in FX does not follow the standard adverse selection model. Dealers do not appear to adjust customer spreads to protect themselves against the likely information content of customer trades. Indeed, our analysis has shown that the cross-sectional pattern of FX customer spreads is the opposite of that predicted by adverse selection: spreads are widest for the smallest trades and they are wider for the least informed customers, the commercial customers. Market participants at large FX dealing banks, to whom we sent the paper for a reality check, all confirmed this qualitative pattern and also the magnitudes we find for spreads during our sample period. Representative comments are included in Appendix B. They affirm that the pattern just identified approximates common knowledge within the FX market: the pattern is known by virtually all traders, and virtually all traders know that virtually all other traders know it, etc.[17]

We stress that these results apply only to the *customer* FX market. Quoted interdealer spreads should be invariant to counterparty type, since the interdealer market is anonymous. The true qualitative relation between interdealer spreads and trade size is unclear, but it does not seem to be negative. The earliest study using interbank transaction data, Lyons (1995), finds a positive relationship between trade size and spreads, and subsequent studies find little or no relationship (Yao (1998), Bjønnes and Rime (2005)). The absence of any such relationship in recent years presumably reflects the fact that interbank trades are consistently small, in part because large trades are split into many small trades, so the size of an interdealer trade is unlikely to carry much information.[18]

Our results indicate that caution should be applied when interpreting correlation results at lower frequencies. For example, Lyons (2001) and Marsh and O'Rourke (2005), which both analyze daily data, suggest that the observed negative correlation between commercial order flow and exchange rates might reflect a negative price impact of commercial trades. If so, our analysis suggests that the price impact is certainly not instantaneous, since spreads appear to be positive for all customers.

Adverse selection has a mixed record of success in explaining customer spreads in other financial markets. It successfully explains the relation between spreads and trade size on the NYSE (see, for example, Harris and Hasbrouck (1996); Bernhardt and Hughson (2002); Peterson and Sirri (2003)). Not only are NYSE spreads wider for larger trades, but some stock brokers pay for order flow from retail (uninformed) customers (Easley *et al.* 1996). The theory also works well in explaining the pattern of price discrimination among specialists on the non-anonymous Frankfurt Stock Exchange (Theissen 2003). There are, however, a number of other markets where adverse selection cannot explain the behavior of customer spreads. A negative relationship between spreads and trade size has been documented in the U.S. corporate bond market (Goldstein *et al.* (2006)) and the London Stock Exchange (Hansch *et al.* (1999)).[19] A negative relationship between trade size and spreads has also been observed in the relatively illiquid U.S.

municipal bond market, where spreads average 2.23 percent for small trades and 0.10 percent for large trades (Harris and Piwowar (2004)).

While our results indicate that adverse selection does not dominate FX customer spreads, they do not prove that adverse selection has zero influence on customer spreads. They also do not point to specific alternative determinants of FX customer spreads. Nonetheless, at least three hypotheses have been suggested in the broader microstructure literature that are consistent with the patterns documented above.

III. Operating Costs, Market Power, and Strategic Dealing

This section highlights three explanations for the pattern of FX customer spreads taken from the broader microstructure literature. The first, fixed operating costs, is suggested by Angel (1996) and Hansch *et al.* (1999) as a source of the negative relationship between spreads and trade size on the London Stock Exchange, since per-unit processing costs become smaller for larger trades. Fixed operating costs certainly exist in FX, so this factor may explain size-based variation in customer spreads. However, operating costs are unlikely to explain the gap between spreads paid by financial and commercial customers. Fixed costs do not vary strongly by customer type and marginal costs are, if anything, higher for asset managers, who often require the proceeds of a large trade to be "split" among numerous individual funds.

Perhaps commercial traders pay wider spreads simply because they trade at times of day when spreads are wider. During our bank's regular trading hours, interdealer spreads (the only spreads for which intraday patterns are available) are widest during the London morning (Payne (2003)). Financial-customer trades tend to be concentrated during the London morning hours, while commercial-customer trades are more evenly distributed across the trading day (Figure 2). In consequence, intraday trading patterns predict variation in customer spreads opposite to that just documented.

The rest of this section highlights two mutually consistent theories of dealing under asymmetric information that might explain why FX spreads vary across counterparty types. One theory suggests that information about current market conditions provides dealers with transitory market power relative to uninformed customers, allowing dealers to charge wider spreads. The other theory suggests that dealers strategically vary spreads across customers in an attempt to gather private information.[20] These information-based forces are mutually consistent and we hypothesize that they operate simultaneously with operating costs. The theories we highlight in this section do not exhaust the long list of factors dealers consider in setting spreads – though a longer list of theories could exhaust the patience of our readers.

A. Market Power

Green *et al.* (2007) point out that dealership markets are opaque due to the dispersion of trading, so current market conditions – meaning real-time mid-quotes, spreads, volatility and the like – can be hard to ascertain. Asymmetric information about these conditions may create market power for dealers. As Angel (1996) describes it, "a dealer knows that an unsophisticated individual … may have higher search costs per share and is not in a good position to monitor the quality of a broker's execution. The broker has little incentive to spend time negotiating or shopping around for a better deal for a small order. Thus, a dealer may take advantage of this by quoting a wider market …" (p. 4). Duffie *et al.* (2004) develops this insight into a formal model and shows that bargaining power in dealership markets partly reflects the alternatives to trading immediately, alternatives that are influenced by the costs and benefits of further search.

The market-power hypothesis can be applied directly to explain why commercial FX customers pay wider spreads than financial customers. Currency markets are also dealership markets with dispersed information, and dealers consider commercial customers less sophisticated than financial customers. Commercial firms typically trade currencies in order to purchase imports necessary for production. Since the firms' currency transactions are scattered across time and sometimes across currencies, the firms do not hire professional traders. Instead, the individuals that carry out FX trading at such firms are typically administrators with many other responsibilities. Because they trade infrequently it is difficult for these individuals to gain an intuitive understanding of any one market, a difficulty compounded by lack of real-time market information, which commercial firms rarely purchase. Further, these individuals are rarely evaluated on execution quality, so they have no incentive to achieve better spreads. Thus traders at commercial firms perceive high costs to search and low benefits, making them ripe targets for wide spreads.[21]

By contrast, FX traders at financial firms are often professionals, have plenty of real-time information, and are often evaluated on their execution quality. These traders get better spreads because, perceiving lower costs and greater benefits to search, they are more likely to keep searching until they are quoted a narrow spread. Financial firms may also gain market power from their tendency to undertake large trades (see Table III). As shown in Bernhardt *et al.* (2004), customers who regularly provide a dealer with substantial amounts of business may receive better spreads as dealers compete for their business.

B. Strategic Dealing

The gap between commercial and financial spreads may also reflect strategic dealing, in which dealers adjust their pricing so as to extract private information from customer trades. Order flow at large banks includes information about upcoming high-frequency currency returns, as documented by Daníelsson *et al*. (2002), Evans and Lyons (2007), and Rime, Sarno, and Sojli (2007). Evidence from equity markets confirms that access to real-time order flow information can provide an informational advantage (Anand and Subramanyam (2005)). The two-tier structure of FX makes it especially logical that FX dealers try to capture informative customer order flow, since they can exploit the information in subsequent interdealer trading.[22]

The insight that market makers might strategically manipulate spreads to increase the information value of order flow was originally explored in Leach and Madhavan (1992, 1993). Those papers use equity-market inspired models to demonstrate that market makers in one-tier markets may adjust prices early in a trading session to enhance later profitability. Flood *et al.* (1999) present evidence for this type of strategic dealing in an experimental market similar to the FX interdealer market.

Our evidence, however, concerns cross-sectional variation in spreads rather than variation across time. Naik *et al.* (1999) presents an equity-inspired model of strategic dealing in a two-tier market. They conclude that customer spreads will be narrower for more informed customers, consistent with the pattern we document for FX. However, the Naik *et al.* model also concludes that customer spreads vary positively with trade size, while our data fit the opposite pattern.

Our bank's order flow need not be hugely informative for strategic dealing considerations to influence its customer spreads. Since dealers are largely price takers when it come to spreads (Cheung and Chinn (2001)), strategic dealing will indirectly influence spreads at moderate-sized banks so long as it directly influences spreads set at the largest banks.[23] It is also noteworthy that dealers need not know exactly which customers are informed for strategic dealing to be influential. Strategic dealing can arise even if, as is true in FX, dealers discriminate only according to a customer's likelihood of being informed. Nonetheless, we show in the next section that our dealer's order flow does carry information and that the information patterns are consistent with the strategic dealing hypothesis outlined above.

C. Empirical Support for Strategic Dealing

This section provides evidence consistent with the strategic dealing hypothesis, structured around three properties of currency order flow consistently asserted by FX dealers: (1) order flow carries information; (2) order flow carries more information when it comes from financial customers than when it comes from commercial customers; (3) large trades are more informative than small trades. Further evidence for strategic dealing in FX is presented in Ramadorai (2006), who analyzes the transactions of asset managers and finds that spreads are narrower for managers that produce higher (risk-adjusted) FX returns.

1. Order flow carries information

Earlier studies confirm that the order flow of individual large dealing banks carries information (Evans and Lyons 2002, Daniélsson *et al*. 2002, Evans and Lyons 2004). Current research shows that large banks have more information than small ones. But the question remains: Does the limited order flow of a modest dealing bank carry any information? To address this question we estimate cointegrating relationships between exchange rates and our three types of cumulative incoming order flow:

$$P_t = \omega_i + \phi_{it} trend + \kappa_i CumOF_{it} + \nu_{it}, \qquad (3)$$

where $i \in \{FC, CC\}$. If incoming order flow of type i is associated with a currency appreciation, κ_i will be positive. The results, shown in Table VII, show that the residuals are stationary as required for cointegration. We conclude that order flow can be informative at banks of modest size, as well as immense banks.

2. Information and Counterparty Type

Dealers claim that, on average, commercial order flow carries less fundamental information than financial order flow. Empirical support for this claim is presented in Carpenter and Wang (2003), which finds that dealers widen their spreads in the interbank market after transactions with financial customers but leave those spreads unchanged after transactions with commercial customers.

The idea that financial order flow is relatively informative may seem counterintuitive, however, based on the literature. Pasquariello, Yun, and Zhu (2006), for example, show that large commercial firms sometimes exhibit timing foresight when issuing ADRs in emerging market currencies, which suggests that commercial firms sometimes do have important information about currencies. In the major currencies that dominate currency markets, however – including the euro, which we study here – commercial firms did not appear to have such timing information. And of course the vast majority of commercial trades have no connection to major capital events such as ADR issues, and instead reflect mundane real-side concerns like a firm's need to import intermediate inputs. Such import decisions are likely to be dominated by microeconomic concerns such as demand for a firm's products and the cost of this input, rather than by exchange-rate forecasts.

The idea that an "informed" agent in FX is a firm involved in real economic activity is also consistent with Evans and Lyons (2004) insight that aggregate order flow conveys information about economic fundamentals that is dispersed among myriad agents. Their illustrative model focuses on information about the real economy, which could naturally come from a commercial firm. Information may in fact be carried by commercial order flow, but so far the evidence suggests that the information is only relevant to long-run exchange rate dynamics. Fan and Lyons (2003) shows that non-financial order flow seems more strongly related to multi-year exchange-rate dynamics than financial order flow while financial trades that carry more information with respect to short-run exchange rates. In their words, "extreme exchange-rate movements at high frequency are generally associated with large net flows from financial institutions" but not from commercial institutions (p. 160). Additional evidence that financial institutions' currency trades provide information about short-run exchange rates is provided in Froot and Ramadorai (2005).

Because currency dealers typically end the day with minimal inventory (Appendix B), they are of necessity focused solely on short horizons. Dealers themselves report that the information they seek concerns short-run phenomena (Gehrig and Menkhoff 2004, Oberlechner 2004). Overall, the evidence to date is consistent with the dealers' claim that financial order flow is most informative for their purposes (as illustrated in Appendix A).

The results of Table VII provide further support for the hypothesis that financial order flow is more informative than commercial order flow. The coefficient on cumulative financial order flow in Table VII is positive, suggesting that this order flow is informative in the familiar way: net buying demand from such customers is associated with an appreciation of the commodity currency. (These results are statistically weak, but later analysis shows that this reflects the mix of informative large trades and uninformative small trades.) The coefficient on commercial order flow is negative, however, which seems to suggest that such order flow is not only informative, contrary to the views of dealers, but informative in the "wrong" way. This statistical result is not unique to our dataset: Qualitatively consistent results are reported in three studies of large banks -- Lyons (2001), Evans and Lyons (2004), and Marsh and O'Rourke (2005) – as well as a study based on comprehensive trading data for the Swedish krone (Bjønnes *et al.* 2005).

The literature suggests a logical reason why commercial order flow would not be considered informative despite its negative relationship with exchange rates. Specifically, it appears that commercial order flow reacts to exchange rates and provides liquidity, on average, rather than propelling exchange rates by demanding liquidity (Bjønnes *et al.* 2005, Osler 2006). Suppose that financial customers drive short-run exchange-rate returns. If so, some other set of participants must be providing the necessary liquidity – if financial customers are buying and driving the price up, say, some other group must be selling. Dealers are obviously the immediate source of that liquidity. But dealers would not readily provide immediate liquidity if they did not anticipate that they could quickly pass their inventory to counterparties outside their own circle. Whoever they are, these ultimate liquidity providers must, by definition, have cumulative order flow that is negatively correlated with exchange rates, and furthermore information about their order flow will have zero incremental value once one knows about financial customers' order flow. In short, if commercial customers are the ultimate providers of liquidity in FX, their trades would simultaneously have a statistically significant negative relationship with exchange rates and yet have no incremental information value for dealers.

Direct support for this interpretation of the market comes from Bjønnes *et al*. (2005), which carefully examines the structure of overnight liquidity using ten years of comprehensive data for the market between Swedish krone and the euro. They first show that market-wide cumulative financial (commercial) order flow is positively (negatively) cointegrated with exchange rates, consistent with our finding. In addition, they show that "changes in net positions of non-financial customers are forecasted by changes in net positions of financial customers." They conclude that "non-financial customers are the main liquidity providers in the overnight foreign exchange market" (p. 1). Further evidence that commercial customers are induced to provide liquidity by exchange-rate changes comes from Marsh and O'Rourke's (2005) analysis of daily customer flows from the Royal Bank of Scotland, which finds strong negative relationship between commercial order flow and lagged returns in most major currency pairs.[24]

In practice, how do commercial customers provide liquidity in FX? We highlight two relevant institutional practices. First, commercial customers are relatively heavy users of take-profit orders, conditional market orders in which dealers are instructed to buy (sell) a specific amount of currency at the market rate immediately after its value falls (rises) to a certain level (Osler 2003, 2005). Of all the euro-dollar, dollar-yen, and sterling-dollar orders placed by commercial customers at the Royal Bank of Scotland between June 2001 and September 2001, 72 percent were take-profits. This structure generates the quick negative feedback trading found empirically. Second, large exporters with inventories of foreign currency are often alert to intraday exchange-rate movements and will sell when the rate reaches intraday targets. The market has historically been particularly aware of this with respect to "Japanese exporters."

3. *Information and Deal Size*

Are there systematic differences in the information content of large and small transactions? That is, could the strategic-dealing hypothesis provide a second explanation, beyond fixed operating costs, for the inverse relationship between deal size and currency spreads? The existing evidence is scarce and inconclusive. Since trades over roughly $25 million are usually broken up into smaller transactions, large regular-sized deals need not carry more information than small ones. Indeed, Chakravarty (2001) shows that the most informative trades on the NYSE are not large but medium-sized. On the other hand, Biais *et al.* (1995) finds that large trades on the Paris Bourse carry more information than small ones. Similarly, Kurov and Lasser (2004) provides evidence that large futures trades are especially informative.

The relative information value of large versus small individual FX transactions is then an empirical question. We focus on financial transactions – since they are considered informative – which we partition into two size categories, small and medium-and-large, as suggested by the results in Table VI. Cointegration tests (Table VII, columns 3 and 4) show that the link between cumulative financial-customer order flow and exchange-rate levels is significant for medium-and-large trades but insignificant for small trades. This is consistent with our hypothesis that medium-and-large financial trades are quoted tighter spreads than small financial trades because they carry more information.

IV. PRICE DISCOVERY IN FOREIGN EXCHANGE

The evidence presented so far suggests that spreads in the FX customer market are inversely related to a trade's information content, the opposite of the pattern predicted by adverse selection. But if adverse selection is not the basis for price discovery in currency markets, what is? This section proposes an alternative price discovery mechanism for FX and provides evidence in support of that mechanism. Asymmetric information is the centerpiece of our story, as it must be. We suggest, however, that information influences inventory management and order choice in the interdealer market rather than spreads in the customer market. Our proposed mechanism thus reflects the foreign exchange market's two-tiered structure rather than the one-tier structure assumed in adverse-selection models.

Any formalization of our price discovery hypothesis will have to wait for the resolution of technical constraints in the literature. The dynamics of order-driven markets with asymmetric information are so complex that financial models have yet to provide closed-form solutions. Our hypothesis involves the interaction of an order-driven (interdealer) market under asymmetric information with a quote-driven (customer) market, also under asymmetric information.

A. The Mechanism

According to the evidence presented above, price discovery must happen in the interbank market, rather than the customer market, since a given trade's information content is not embedded in prices at the customer level.[25] Interdealer markets are crucially important for inventory management in FX (Lyons 1996) as in other two-tier markets (Manaster and Mann (1996), Reiss and Werner (1998)). Our proposed price discovery mechanism involves dealers' trading behavior in the interbank trading after customer trades. As we will show, our price discovery mechanism applies regardless of whether FX dealers choose "indirect" trading in the order-driven broker market or "direct" trading in the regular quote-driven market.

Consider a dealer whose inventory rises abruptly in response to a customer decision to sell. Since FX dealers prefer to have zero inventory, this dealer will most likely try to offload the new inventory to another dealer. Assume for now that our dealer chooses to trade through an interdealer broker, in which case he must decide whether to submit a market sell or a limit sell. Harris (1998) and Foucault (1999) highlight a central trade-off: market orders provide immediate execution with certainty while limit orders provide better prices but uncertain execution. Since FX dealers can identify their customers, this order choice could depend on the customer providing the inventory (Reiss and Werner (2004)).

Suppose the customer is considered informed. In this case the dealer has three incentives to exploit the immediacy offered by market orders: (i) he has information, (ii) he has inventory with its inherent risk, and (iii) his information indicates that his inventory could soon bring a loss. Our dealer therefore seems likely to place a market sell order and earn the lower bid price. The price will fall. Suppose instead the customer is uninformed. In this case the dealer has only one incentive to place a market order: the inherent riskiness of his inventory. Thus our dealer might be more likely to place a limit order which, if executed, would earn him the higher offer price. The limit order would provide liquidity rather than move the price. In short, we suggest that dealers will have a stronger tendency to place market orders after informed customer trades than after uninformed customer trades.[26] The connection to price discovery is direct: brokered interdealer prices will therefore tend to move in the direction indicated by informed trades. (Note that our discussion of price discovery does not assume a priori that informed dealers place outgoing/market orders, but instead derives that outcome.)[27]

If our dealer chooses to trade directly, this analysis changes only superficially. Calling another dealer produces a quick, certain trade at a relatively undesirable price, like placing a market order; waiting for someone else to call could bring a better price but could instead bring no trade at all, like placing a limit order. Thus, a dealer has strong incentives to call another dealer after trading with an informed customer, in which case he trades at the lower bid price and the market price declines. After trading with an uninformed customer the dealer may be more likely to wait for incoming calls.

This set of motivations could conceivably lead to a no-trade equilibrium if (i) customer identity were the only factor determining whether a dealer makes an outgoing trade and (ii) customer identity were a reliable indicator of whether the customer is informed at a given point in time. In this case all market orders would be placed by dealers with information, so uninformed agents would always be at an informational disadvantage when placing limit orders and the market would cease to exist. These conditions do not hold, however. The literature identifies many factors besides information that influence order choice; we shortly confirm the importance of one of those factors, inventories. Furthermore, customer identity is imperfectly correlated with a given customer's private information at any point in time.[28]

B. Explaining the Stylized Facts

Our proposed price discovery mechanism predicts a number of the stylized facts in FX microstructure. For example, it predicts that financial order flow, which dealers assert is relatively informed, will be positively related to exchange-rate returns. Evidence for this positive relationship is provided in Evans and Lyons (2004), Bjønnes *et al.* (2005), and Marsh and O'Rourke (2005). If the information in question is permanent, then our analysis also predicts that this relationship between financial order flow and exchange rates is substantially permanent, evidence for which is provided in Lyons (2001) and in Bjønnes *et al.* (2005).

Our proposed price discovery mechanism also predicts a positive and largely permanent relationship between exchange rates and interdealer order flow. (In the order-driven or brokered portion of the interdealer market this is market buys minus market sells; in the quote-driven or direct dealing portion of that market the initiator is the dealer that calls out.) Consistent with this prediction, substantial evidence indicates a strong and positive contemporaneous correlation between interdealer order flow and exchange-rate returns at the daily and weekly horizons (see Lyons (1995), Payne (2003), Evans (2002), Evans and Lyons (2002), Killeen, Lyons, and Moore (2002), and Daníelsson *et al.* (2003), *inter alia*). Furthermore, a substantial portion of this relationship is permanent (Evans and Lyons (2002), Payne (2003), Killeen *et al.* (2005), Bjønnes *et al.* 2005).

Our proposed price discovery mechanism also answers a natural question regarding the strategic dealing hypothesis: If dealers narrow spreads to attract informed customers, how do the dealers benefit from that information? We answer: they benefit via enhanced interdealer trading. The information permits them to reduce their inventory risk and to position themselves to profit from anticipated returns.

C. Additional Evidence

Our proposed price discovery mechanism has four additional testable implications. First, it predicts that interdealer prices are the best measure of "the market" at any instant. Abundant institutional evidence confirms this implication. Most critically, dealers universally base their customer quotes on the interdealer market's current best bid and offer. In a large dealing room, salespeople construct the quote given to a customer from a preliminary quote provided at that moment by an interdealer trader. Those preliminary quotes are in turn anchored on the best bid and offer in the interdealer market. In electronic communication networks (e.g., Currenext, FXAll) the connection between interdealer prices and customer quotes is programmed directly into the pricing algorithm.

Second, our proposed price discovery mechanism predicts that dealers with the most customers should be best informed and should profit the most from interdealer trading. Concurrent research by Bjønnes *et al.* (2007) supports both implications. Trades by the largest banks are positively correlated with each other but negatively correlated with trades by smaller banks, which suggests implicitly that dealers can be divided into two size categories, big and small. Large-bank (small-bank) trades are positively (negatively) correlated with returns, indicating that large banks are relatively informed.

The third and fourth testable implications of our proposed price discovery mechanism concern the likelihood of outgoing interbank transactions. Under our proposed price discovery mechanism, dealers should be more likely to place interdealer market orders after trades with financial customers than after trades with commercial customers. Similarly, dealers will be more likely to place interdealer market orders after larger trades than after small ones, even after controlling for inventory. If our hypothesis is incorrect and information from individual customers is not critical to the dealers' "make or take" decisions, then customer type and trade size should not matter once we control for the dealer's inventory level and other institutional factors.

We test these implications via a probit analysis of the conditional probability that a given transaction is outgoing in the interbank market:

$$Prob(Trade_t{=}IB^{out}) = P(FC_{t-1}, CC_{t-1}, 10mio_{t-1}, |I_{it}|, I_{it}^{2}, |Q_{jt}|) \quad . \tag{5}$$

Our hypothesis concerns the first three variables: dummy variables indicating that the previous trade was with a financial customer, FC_{t-1}, or with a commercial customer trades, CC_{t-1}, and a dummy set to one if the previous transaction was worth €10 million or more, $10mio_{t-1}$. Our conjecture suggests that the coefficient on the financial dummy will be higher than the coefficient on commercial dummy and the coefficient on $10mio_{t-1}$ will be positive.

The last three terms in equation (5) capture other factors relevant to the decision to place a market order. The coefficient on absolute inventory, $|I_{it}|$, should be positive since higher inventory brings higher inventory risk.[29] Following Bjønnes and Rime (2005) we include squared inventory, I_{it}^2, to capture potential nonlinearities in this relationship. The absolute size of the current transaction, $|Q_{jt}|$, is included because our dealer's customer transactions are often smaller than the $1 million minimum size for brokered trades. Since our dealer prefers to carry out interbank trades on EBS, a broker, rather than by dealing directly, he seems likely to collect inventory from small customer transactions and then square his position by submitting one relatively large market order.

The results of estimating Equation (5), shown in Table VIII, support our view that the likelihood of an outgoing interbank transaction is higher when the previous transaction is considered informed. Outgoing interbank transactions are statistically significantly more likely when the previous transaction involves a financial customer than when it involves a commercial customer. They are also statistically significantly more likely after big trades, meaning those over €10 million. The results are economically meaningful, as well. After a moderate-sized commercial trade the estimated probability of an outgoing interbank transaction is 9.5 percent; after a similarly-sized financial trade that probability is roughly twice as large, at 18.5 percent. After commercial trade over €10 million the probability of an outgoing interbank transaction is 25.4 percent. After a similarly-sized financial trade this probability reaches 40.2 percent. (In these calculations, all other independent variables are taken at sample means.) As indicated by the three robustness tests, these results, like our earlier results, are not sensitive to whether inventories are included as an independent variable or to whether the data include spot trades or interdealer trades. Unreported results show that the results are also not sensitive to the inclusion of September 11, 2007.

The rest of the results from estimating Equation (5) also make sense. The likelihood of an outgoing trade rises with the absolute value of existing inventory and the relationship is concave. As noted above, the influence of inventory level on dealer order choice helps the market avoid no-trade equilibria. The positive relationship between absolute trade size and the likelihood that the trade itself is outgoing indicates that outgoing brokered transactions tend to be larger than the dealer's average incoming transaction, as expected.[30]

V. CONCLUSIONS

This paper's overall message is that the standard adverse selection model of price discovery may not apply in the foreign exchange market. Instead, we propose a new price discovery process for this liquid two-tier market. Our data comprise the complete USD/EUR trading record of a bank in Germany over four months in 2001. The paper first provides evidence that adverse selection does not dominate the pattern of customer spreads. While adverse selection predicts wider spreads for large than for small trades, and wider spreads for financial than commercial customers, the reverse appears to be true. Other implications of adverse selection are also not supported for this market. This evidence implies that adverse selection in the customer market cannot dominate the price discovery mechanism in FX.

The paper then highlights three hypotheses from the broader microstructure literature that help explain the cross-sectional pattern of currency spreads. We first note that operating costs are partially fixed in FX, which could help explain the negative relationship between trade size and spreads. The customer-based variation in spreads could be explained in part by Green *et al.*'s (2007) market-power hypothesis, which asserts that spreads in quote-driven markets vary positively with a dealer's market power relative to a given customer, and that such market power derives in part from knowledge of market conditions. Commercial customers tend to know the least about current market conditions, so this theory predicts they pay the widest spreads, as they do. The customer-based variation in spreads could also reflect dealers' attempts to strategically gather information (Leach and Madhavan (1992), (1993), Naik *et al.* (1999)). Dealers may narrow spreads to attract informed customers and extract information from their trades, information from which they hope to profit in subsequent interdealer trades. Dealers consider financial order flow to be relatively informative, so financial customers pay the narrowest spreads.

The paper finishes by proposing a new price discovery process that incorporates the foreign exchange market's two tiers and high liquidity. We first note that, since customers' information is not immediately reflected in the prices they pay, price discovery must take place in the interdealer market. The key mechanism behind our suggested price discovery process involves the way dealers trade in the interdealer market after individual customer trades. We suggest that after trading with informed customers dealers will tend to make parallel outgoing interdealer trades – placing a market buy order after an informed customer buy, for example – motivated by their inventory as well as by their newly-acquired information. In this way the information from customer trades will be reflected in interdealer prices. After trading with uninformed customers, by contrast, dealers will be more likely to place parallel limit orders or to wait for incoming calls, leaving price relatively unaffected.

We provide evidence that our dealer is more likely to place outgoing interdealer trades after informed customer trades, consistent with our proposed price discovery mechanism. Our theory also predicts some key stylized facts in FX: the positive and substantially permanent relation between cumulative interdealer order flow and exchange rates, as well as the positive and substantially permanent relation between financial order flow and exchange rates.

Customer spreads are known to vary inversely with trade size in other liquid two-tier markets, including the U.S. Treasury market, the U.S. corporate bond markets, and the London Stock Exchange. Our proposed price discovery mechanism may thus apply in these markets as well as FX, since the mechanism incorporates the interplay of highly liquid trading between two tiers. In future research it would be valuable to test the relevance of our proposed price discovery process in these other markets.

REFERENCES

Anand, Amber, and Avanidhar Subramanyam, 2005, Information and the intermediary: Are market intermediaries informed traders in electronic markets? Mimeo, Syracuse University.

Angel, James, 1996, Who Gets Price Improvement on the NYSE? Georgetown University Working Paper FINC-1377-12-694.

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

Bank for International Settlements, 2004, Preliminary global results, Triennial central bank survey, foreign exchange and derivatives market activity in 2004, (B.I.S., Basel).

Bernhardt, Dan, and Eric Hughson, 2002, Intraday trade in dealership markets, *European Economic Review* 46, 1697-1732.

Bernhardt, Dan, Vladimir Dvoracek, Eric Hughson, and Ingrid Werner, 2005, Why do larger orders receive discounts on the London Stock Exchange? *Review of Financial Studies* 18, 1343-1368.

Biais, Bruno, Pierre Hillion, and Chester Spatt, 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *Journal of Finance* 50, 1655-1689.

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

Bjønnes, Geir Høidal and Dagfinn Rime, 2005, Dealer behavior and trading systems in foreign exchange markets, *Journal of Financial Economics* 75, 571-605.

Bjønnes, Geir Høidal, Dagfinn Rime, and Haakon O.Aa. Solheim, 2005, Liquidity provision in the overnight foreign exchange market, *Journal of International Money and Finance* 24, 175-196.

Bjønnes, Geir Høidal, Carol Osler, and Dagfinn Rime, 2007. Asymmetric Information in the Foreign Exchange Interdealer Market," Mimeo.

Carpenter, A., and J. Wang 2003, Sources of private information in FX trading. Mimeo, University of New South Wales.

Chacko, George, Sriketan Mahanti, Gaurav Mallik, Marti Subrahmanyam, 2005. The Determinants of Liquidity in the Corporate Bond Markets: An Application of Latent Liquidity, Working Paper.

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.

Christie, William G., Paul H. Schultz, 1994, Why do NASDAQ market makers avoid odd-eighth quotes? *Journal of Finance* 49, 1813-1840.

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

Daníelsson, Jon, Richard Payne, Jinhui Luo, 2002, Exchange rate determination and inter-market order flow effects, Typescript, Financial Markets Group, London School of Economics.

Duffie, Darrell, Nicolae Garleanu, and Lasse Heje Pedersen, 2004, Over-the-counter markets. *NBER Working Paper* 10816.

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

Easley, David, Nicholas M. Kiefer, and Maureen O'Hara, 1996, Cream skimming or profit sharing? The curious role of purchased order flow, *Journal of Finance* 51, 811-833.

Easley, David, Nicholas M. Kiefer, and Maureen O'Hara, 1997, One day in the life of a very common stock, Review of Financial Studies 10, 805-835.

Euromoney, May 2007. Euromoney FX Poll 2007: http://www.euromoney.com/article.asp?PositionID=polls&ArticleID=1330478

Evans, Martin, 2002, "FX trading and exchange rate dynamics," *Journal of Finance* 57, 2405-2448.

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

Evans, Martin D.D. and Richard K. Lyons, 2004, Exchange-rate fundamentals and order flow, Mimeo, U.C. Berkeley and Geortown University.

Evans, Martin D.D. and Richard K. Lyons, 2005, Meese-Rogoff Redux: Micro-Based Exchange Rate Forecasting. *American Economic Review* 95, 405-14.

Evans, Martin D.D. and Richard K. Lyons, 2006, Understanding Order Flow, *International Journal of Finance and Economics* 11, 3-23.

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.

Federal Reserve Bank of New York (2005). http://www.ny.frb.org/.

Flood, Mark D., Ronald Huisman, Kees G. Goedijk, and Ronald J. Mahieu, 1999, Quote disclosure and price discovery in multiple-dealer financial markets, *Review of Financial Studies* 12, 37-59.

Foucault, Thierry, 1999, Order flow composition and trading costs in a dynamic limit order market, *Journal of Financial Markets* 2, 99-134.

Froot, Kenneth. A., and Tarun Ramadarai (2005), Currency returns, intrinsic value, and institutional investor flows, *Journal of Finance*, forthcoming.

Gehrig, Thomas and Lukas Menkhoff, 2004, The use of flow analysis in foreign exchange: exploratory evidence, *Journal of International Money and Finance* 23, 573-594.

Glosten, Lawrence, 1989, Insider trading, liquidity, and the role of the monopolist specialist, *Journal of Business* 62, 211-235.

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

Goettler, Parlour, and Rajan, 2005, Equilibrium in a dynamic limit order market, *Journal of Finance* 60, 2149-2192.

Goldstein, Michael A., Edith S. Hotchkiss, Erik R. Sirri, 2006, Transparency and liquidity: A controlled experiment on corporate bonds, Forthcoming, *Review of Financial Studies.*

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, 2007, Financial intermediation and the costs of trading in an opaque market, *Review of Financial Studies,* 2007, 20: 275-314.

Hansch, Oliver, and Anthony Neuberger, 1997, Strategic trading by market makers on the London Stock Exchange, Mimeo, London Business School.

Hansch, Oliver, Narayan Naik, and S. Viswanathan, 1998, Do inventories matter in dealership markets? Some evidence from the London Stock Exchange, *Journal of Finance* 53, 1623-1656.

Hansch, Oliver, Narayan Naik, and S. Viswanathan, 1999, Preferencing, internalization, best execution, and dealer profits, *Journal of Finance* 54, 1799-1828.

Harris, Lawrence, 1998, Optimal dynamic order submission strategies in some stylized trading problems, *Financial Markets, Institutions, and Instruments* 7, 1-75.

Harris, Lawrence, 2003, *Trading and Exchanges: Market Microstructure for Practitioners* (Oxford University Press, New York).

Harris, Lawrence, and Joel Hasbrouck, 1996, Market vs. limit orders: The SuperDOT evidence on order submission strategy, *Journal of Financial and Quantitative Analysis* 31, 213-231.

Harris, Lawrence E. and Michael S. Piwowar, 2004, Report on Transactions in Municipal Securities, Securities and Exchange Commission.

Hasbrouck, Joel, and George Sofianos, 1993, The trades of market makers: An empirical analysis of NYSE specialists, *Journal of Finance* 48, 1565-1593.

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.

Kaniel, Ron, and Hong Liu, 2004, So what orders do informed traders use? Working Paper, Fuqua School of Business.

Killeen, William, Richard K. Lyons, and Michael Moore, 2005. "Fixed versus flexible: Lessons from EMS order flow," Forthcoming, *Journal of International Money and Finance*.

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.

Leach, J. Chris and Ananth N. Madhavan, 1993, Price experimentation and security market structure, *Review of Financial Studies* 6, 375-404.

Lo, Ingrid, and Stephen G. Sapp, 2005, Price aggressiveness and quantity: How are they determined in a limit order market? Presented at the Norges Bank/Norwegian School of Management Conference on Equity and Currency Market Microstructure, September 2005.

Luo, Jinhui, 2002, Market conditions, order flow, and exchange rate determination, Mimeo, Financial Market Group, London School of Economics.

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

Lyons, Richard K., 1997, A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics* 42, 275-98.

Lyons, Richard K., 2001, The microstructure approach to exchange rates (Cambridge: MIT Press).

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

Madhavan, Ananth N. and Seymour Smidt, 1993, An analysis of changes in specialist inventories and quotations, *Journal of Finance* 48, 1595-1628.

Manaster, Stevern, and Steven C. Mann, 1996, Life in the pits: Competitive market making and inventory control, *Review of Financial Studies* 9, 953-975.

Marsh, Ian, and Ceire O'Rourke, 2005, Customer order flow and exchange rate movements: Is there really information content? Presented at the Norges Bank/Norwegian School of Management Conference on Equity and Currency Market Microstructure, September 2005.

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

Mende, Alexander, 2006, 09/11 on the USD/EUR exchange market, *Applied Financial Economics* 16: 213-222.

Naik, Narayan Y., Anthony Neuberger, and S. Viswanathan, 1999, Trade disclosure regulation in markets with negotiated trades, *Review of Financial Studies* 12, 873-900.

New York Stock Exchange, 2004. Annual Report 2003. New York, NY.

Payne, Richard, 2003, Informed trade in spot foreign exchange markets: an empirical investigation, *Journal of International Economics* 61, 307-329.

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

Ramadorai, Tarun, 2006, Persistence, performance, and prices in foreign exchange markets, Oxford University Working Paper.

Reiss, Peter C., and Ingrid M. Werner, 1998, Does risk sharing motivate interdealer trading? *Journal of Finance* 53, 1657-1703.

Reiss, Peter C., and Ingrid M. Werner, 2004, Anonymity, adverse selection, and the sorting of interdealer trades, *Review of Financial Studies* 18, 599-636.

Rime, Dagfinn, 2003, New Electronic Trading Systems in Foreign Exchange Markets, in *The New Economy Handbook* (Elsevier Science, USA), 469-504.

Rime, Dagfinn, Lucio Sarno, and Elvira Sojli (2007) "Exchange-rate Forecasting, Order Flow, and Macro Information," *Norges Bank Working Paper* 2007-2.

Theissen, Erik, 2003, Trader anonymity, price formation and liquidity, *European Finance Review* 7, 1-26.

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 A: Banks of Different Sizes Banks Behave Similarly

Microeconomic theory implies that the intense competition of the foreign exchange market should leave all banks as “price takers.” Banks confirm that they are price takers with respect to bid-ask spreads when they assert that the main determinant of their spreads is the spreads’ “conventional level” (Cheung and Chinn 2001). This appendix further documents that FX dealing banks behave similarly. We document that our moderate-sized bank dealer behaves very similarly to large-bank dealers in terms of pricing and inventory management. The analysis is based on the Madhavan-Shmidt model outlined in Section II, with customers aggregated into one category for comparability with earlier studies.

Baseline spreads: As shown in Table AI, our bank's average baseline half-spread for interbank transactions is about 1.5 pips, which is similar to estimates from other studies. For example, Goodhart *et al.* (2002) finds that the average spread for USD/EUR transactions on the 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 trades, 9.2 pips, is much higher than its average interdealer spread of 1.6 pips. Customers are also quoted sharply higher spreads than other dealers by Bjønnes and Rime's (2001) NOK/DEM dealer. These figures imply that currency spreads average less than 0.1 percent; for comparison, average municipal bond spreads were 180 basis points in 2003 (Harris and Piwowar (2004)) and average spreads on the London Stock Exchange were 110 basis points in 1991 (Reiss and Werner (2004)).

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 (2005)). Survey-based evidence confirms that inventories are of minimal importance when dealers set spreads, and that the dominant concern is whether spreads conform to market convention (Cheung and Chinn (2001)). 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 (2005) argue that any shift away from inventory-based price shading in recent years may reflect the interbank market’s rapid shift to a heavy reliance on electronic brokerages after their introduction in the mid-1990s (Melvin and Wen (2003)). Our dealer reports that for interbank trades he generally uses EBS because it is less expensive and faster than direct interbank dealing.[31] Together, these observations imply that our dealer controls inventories via interbank trading instead of price shading, a conclusion we support empirically later in this section. Studies from other markets also show that dealers in two-tier markets with access to brokerage services prefer to manage their inventory through interdealer transactions (Reiss and Werner (1998)).

The estimates in Table AI seem to provide slight evidence of inventory-based price shading in the “wrong” direction with respect to transactions with customers. Reassuringly, this can be traced to one trade carried out in the first month of our sample period. If that trade is excluded, the coefficients on inventory are insignificant.

Trade size and spreads: The coefficient on trade size is statistically insignificant for interbank trades, suggesting that neither information asymmetries nor prospective inventories cause large interbank trades to be priced less attractively than small ones. This is consistent with the large dealing bank examined in Bjønnes and Rime (2004), for which spreads on brokered interbank transactions seem independent of trade size. That paper also finds that spreads rise with trade size for direct interbank transactions, a distinction that makes economic sense. Dealers have limited control over the relationship between trade size and spread for brokered transactions, but they have full control for direct trades. Notably, the earliest studies of currency dealers (Lyons (1995), Yao (1998)), which did not control for the distinction between direct and brokered trades, found that interbank spreads do rise with trade size, 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 trade size is also insignificant for customers in our baseline regression. Note that this coefficient is negative and significant when inventories are excluded: Section II showed that the overall relationship between spreads and trade size is indeed negative for customer transactions.

2. Inventory Management

Our dealer's tendency to keep inventories close to zero (Figure 1) is itself similar to inventory management practices at the largest banks. As Table I 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 estimating the following regression:

$$I_t - I_{t-1} = \omega + \rho I_{t-1} + \varepsilon_t. \tag{A1}$$

If the dealer instantly eliminates unwanted inventories, then $\rho \approx -1$. If the dealer allows his inventory to change randomly, then $\rho = 0$. The time subscript corresponds to transaction time, and only incoming transactions, for which our dealer quotes the price, are included (giving 2,858 observations). Results from estimating Equation (A1), once again using GMM with Newey-West standard errors, confirm that our small bank strives to keep inventories close to zero. Our point estimate of $\rho = -0.20$ has a standard error of 0.008 and is thus highly statistically significant. The dealer on average eliminates 20 percent of an inventory shock in the next trade, which implies a median inventory half-life of 19 minutes.

Our estimated inventory half-life is quite close to the 18-minute median inventory half-life for Bjønnes and Rime's (2004) NOK/DEM dealer. The speed of adjustment is faster in futures markets, where dealers eliminate almost half of any inventory shock in the next trade (Manaster and Mann 1994). Adjustment speeds are also faster of the large DEM/USD dealers at the bank studied by Bjønnes and Rime, for which inventory half-lives range from 0.7 to 3.7 minutes. Nonetheless, our dealer's adjustment speed is lightning fast, and differs little from the others just reported, when compared with inventory adjustment lags in other markets. On the NYSE these lags average over a week (Madhavan and Smidt (1993)) and can extend beyond a month (Hasbrouck and Sofianos (1993)). Even on the London Stock Exchange, which is a dealership market like FX, inventory half-lives average 2.5 trading days (Hansch *et al.* (1998)).

Overall, this analysis shows that the dealer from which we take our data behaves much like the largest dealers.

Appendix B: Representative Comments From Market Participants

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 … provid[e] a greater facility for price discovery than smaller customers…. 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 then moved on to be 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 I. 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 trades, over the 87 trading days between July 11th, 2001 and November 9th, 2001.

	All Transactions	Interbank	Customer		
			All	Financial	Commercial
Number of Transactions	3,609	1,919	1,690	171	1,519
(percent)	(100)	(44)	(56)	(5)	(42)
Of Which, Forward	646	114	532	60	472
Value of trades (€ mil.)	4,335	2,726	1,609	405	1,204
(percent)	(100)	(61)	(39)	(9)	(28)
Of Which, Forward	999	87	912	226	686
Mean Size (€ mil.)	1.20	1.42	0.95	2.37	0.79
Mean Size, Forwards (€ mil.)	1.55	0.76	1.71	3.77	1.45

Table II. Comparison of small bank studied here with larger banks studied in other papers.

The table shows the complete USD/EUR trading activity of a 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 trades.

	Bank in Germany	B.I.S. (2002) per Bank	Lyons (1995)	Yao (1998)	Bjønnes and Rime (2005): Four Dealers, Range	Bjønnes and Rime (2005): DEM/USD Dealer	Bjønnes and Rime (2005): 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		
Transactions per Day	**40 (51)**	---	267	181	58 - 198	198	58
Transaction value per Day (in $ millions)	**39 (52)**	50 - 150	1,200	1,529	142 - 443	443	270
Value per Transaction ($ mil.)	**1.0**	---	4.5	8.4	1.6 - 4.6	2.2	4.6
Customer Share of Transaction value (in percent)	**23 (39)**	33	0	14	0 – 18	3	18
Average Inventory Level (in € or $ millions)	**3.4**		11.3	11.0	1.3 – 8.6	4.2	8.6
Average Transaction Size (in € or $ 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 III. Size distribution of individual trades

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			
Below € 0.1 million	7%	15%	54%
€ 0.1 – 0.5 million	9%	26%	32%
€ 0.5 – 1.0 million	7%	14%	5%
€ 1.0 – 20 million	77%	44%	8%
€ 20 million and above	0%	1%	1%

Table IV: We estimate this equation: $\Delta P_{it} = \alpha + \beta(D_t - D_{t-1}) + \eta_t$. The dependent variable is the change in price between two successive customer trades measured in pips. D_t is an indicator variable picking up the direction of the deal: D_t is +1 for buy-initiated trades and –1 for sell-initiated trades. The change-in-direction variable is interacted with dummy variables for two customer types, financial customers (*FC*) and commercial customers (*CC*), and with dummies for three trade size categories, large trades (*Lg*), meaning those worth $1 million or more; medium trades (*Md*), meaning those worth $500,000 to $1 million, and small trades (*Sm*), meaning those smaller than $500,000. Data include all incoming customer USD/EUR spot and forward trades of a small bank in Germany, except those with preferred customers, during 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.

	Coefficient	**Standard Error**
Constant	0.894‡	0.296
FC x Sm x ΔD_t	5.619‡	2.420
CC x Sm x ΔD_t	8.618‡	0.512
FC x Md x ΔD_t	2.821†	1.414
CC x Md x ΔD_t	9.967‡	1.458
FC x Lg x ΔD_t	3.365‡	1.100
CC x Lg x ΔD_t	3.060‡	0.929
Adj. R^2	0.271	
No. Obs.	1,640	

Table V. 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 customer trades measured in pips. D_t is +1 for buy-initiated trades and –1 for sell-initiated trades. I_{it} is the dealer's inventory, measured in EUR millions. These variables are interacted with dummy variables for trades with financial customers (*FC*) and trades with commercial customers (*CC*). They are also interacted with dummies for trade size: *Lg*. = $\{|Q_{jt}| \in [1,\infty)\}$; *Md*. = $\{|Q_{jt}| \in [0.5,1)\}$; *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. Estimates of the baseline half spread are highlighted in bold.

	Baseline Regression		**Robustness 1: No Inventories**	**Robustness 2: Spot Trades Only**	**Robustness 3: Interbank Trades Included**
	Coefficient	**Std. Error**	**Coefficient**	**Coefficient**	**Coefficient**
Half-Spread, *S*/2					
S/2 x *FC* x *Sm*.	**10.538‡**	*2.55*	**10.606‡**	**7.807‡**	**9.304‡**
S/2 x *FC* x *Md*.	**5.354†**	*2.39*	**4.125**	**2.763**	**4.918†**
S/2 x *FC* x *Lg*.	**4.202†**	*1.94*	**4.214†**	**0.998**	**1.597**
S/2 x *CC* x *Sm*.	**13.478‡**	*0.59*	**13.436‡**	**11.346‡**	**12.805‡**
S/2 x *CC* x *Md*.	**11.621‡**	*2.74*	**12.298‡**	**13.561‡**	**12.963‡**
S/2 x *CC* x *Lg*.	**3.804†**	*1.65*	**3.480†**	**6.505†**	**4.478‡**
S/2 x *IB* x *Sm*.+ *Md*.					**0.817**
S/2 x *IB* x *Lg*.					**3.934‡**
Adverse Selection					
λ x *FC* x *Sm*.	0.319	*0.21*	0.333*	0.529*	0.391†
λ x *FC* x *Md*.	0.457	*0.52*	0.330	-0,395	0.802*
λ x *FC* x *Lg*.	0.266	*0.57*	0.346	-3.360	1.965
λ x *CC* x *Sm*.	0.056†	*0.02*	0.048†	0.197‡	0.101‡
λ x *CC* x *Md*.	0.393†	*0.18*	0.426‡	0.614‡	0.348†
λ x *CC* x *Lg*.	0.513	*0.46*	0.534	0.489	0.364
λ x *IB* x *Sm*.+ *Md*.					-2.729
λ x *IB* x *Lg*.					0.717‡
Inventory					
θ x *FC* x *Sm*.	0.038	*0.18*		0.116	0.18
θ x *FC* x *Md*.	-0.512	*0.42*		-1.315	0.42
θ x *FC* x *Lg*.	0.003	*0.05*		0.152	0.05
θ x *CC* x *Sm*.	-0.078*	*0.04*		-0.002	0.04
θ x *CC* x *Md*.	0.081	*0.27*		-0.003	0.27
θ x *CC* x *Lg*.	-0.011	*0.02*		-0.017	0.02
θ x *IB* x *Sm* +*Md*.					4.814
θ x *IB* x *Lg*.					-0.077
Adjusted *R*²	0.33		0.33	0.35	0.23
Observations	1,651		1,651	1,129	2,859

Table VI. Spread variation across trade sizes and counterparty types

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

The dependent variable is the change in price between two successive incoming trades, measured in pips. 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); I_{it} is the dealer's inventory at time *t,* and Q_{jt} is order flow measured in millions of euros. These variables are interacted with dummy variables for financial customers (*FC*) and commercial customers (*CC*). They are also interacted with dummies for trade size: *Lg*. = $\{Q_{jt} \in [1,\infty)\}$; *Md*. = $\{Q_{jt} \in [0.5,1)\}$; *Sm*. = $\{Q_{jt} \in (0,0.5)\}$. Data include all incoming customer 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. Estimates of the (negative of the) baseline half spread are highlighted in bold.

	Baseline Regression		Robustness Tests			
			No Inventory	No Quantity	Spot Trades Only	Interbank Trades Included
	Coeff.	***Std. Error***	**Coeff.**	**Coeff.**	**Coeff.**	**Coeff.**
Constant	0.031	*0.32*	0.159	-0.047	0.718*	-0.597†
Direction						
FC x Sm x D_t.	10.456‡	*2.58*	10.419‡	7.880‡	12.924‡	9.034‡
FC x Sm x D_{t-1}	**-6.615‡**	*2.39*	**-6.935‡**	**-3.230**	**-13.236‡**	**-5.420‡**
FC x Md. x D_t	3.921	*2.69*	3.905	3.161	5.574	3.364
FC x Md. x D_{t-1}	**-2.972**	*2.99*	**-2.930**	**-2.174**	**-4.679**	**-0.895**
FC x Lg. x D_t	2.397	*2.93*	2.788	5.117‡	4.013	-0.164
FC x Lg. x D_{t-1}	**-3.622***	*2.02*	**-3.100**	**-1.308**	**-0.065**	**0.343**
CC x Sm. x D_t	13.329‡	*0.61*	13.327‡	11.766‡	11.403‡	12.934‡
CC x Sm. x D_{t-1}	**-12.681‡**	*0.64*	**-12.729‡**	**-10.138‡**	**-11.100‡**	**-11.469‡**
CC x Md. x D_t	12.618‡	*1.56*	12.473‡	12.914‡	13.945‡	14.570‡
CC x Md. x D_{t-1}	**-7.199‡**	*1.86*	**-7.161‡**	**-7.267‡**	**-5.607‡**	**-8.492‡**
CC x Lg. x D_t	4.682†	*2.31*	4.721†	5.759‡	1.010	6.296‡
CC x Lg. x D_{t-1}	**-2.064**	*1.76*	**-1.715**	**-4.010‡**	**0.001**	**-3.189†**
IB x Md.+Sm.x D_t						2.027
IB x Md.+Sm.x D_{t-1}						**-3.757†**
IB x Lg. x D_t						3.450‡
IB x Lg. x D_{t-1}						**-1.122†**
Inventory						
FC x I_{it}	-0.464	*0.59*		0.049	-0.234	1.119
FC x I_{it-1}	0.365	*0.60*		-0.135	0.169	-1.180
CC x I_{it}	1.052†	*0.41*		0.144	0.029	1.012†
CC x I_{it-1}	-1.087‡	*0.42*		-0.143*	-0.036	-1.097‡
IB x I_{it}						-0.263
IB x I_{it-1}						0.198
Trade size						
FC x Q_{jt}	0.121	*0.73*	0.435		-0.263	1.597
CC x Q_{jt}	0.773*	*0.47*	-0.240		0.311	0.522
IB x Q_{jt}						-0.347
Adjusted R^2	0.33		0.33	0.34	0.32	0.24
Observations	1,640		1,640	1,640	1,125	2,848

Table VII. Exchange Rates and Cumulative Deal Flow

Table presents ordinary least squares estimates of the cointegrating relationships between exchange rates and cumulative order flow:

$$P_t = \omega_i + \phi_{it} trend + \kappa_i CumOF_{it} + \nu_{it},$$

where $i \in \{FC, CC\}$. If incoming order flow of type i is associated with a currency appreciation, κ_i will be positive. Preliminary statistical tests indicate that the variables are not stationary, so t-values for coefficients are not reliable and are not reported. The ADF tests is a standard augmented Dickey-Fuller test on the regression residuals. PP test is a Phillips-Perron test on those residuals. The number of lags 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. Significanct at the 1, 5, and 10 percent levels is indicated by ‡, †, and *, respectively. Flow and trend coefficients are multiplied by 10^3.

	Commercial Customers	Financial Customers: All Trades	Financial Customers	
			Med. & Lg. Trades	Small Trades
Cumulative Order Flow	-0.301‡	0.150*	0.255†	4.330
Constant	0.884‡	0.885‡	0.891‡	0.875‡
Trend	0.008‡	0.167‡	0.167†	0.553‡
ECM Coefficient	-0.120‡	-0.051†	-0.098†	-0.132†
Adjusted R^2	0.59	0.30	0.20	0.48
ADF-test	-2.72‡	-1.63*	-1.64†	-1.55
PP-test	-3.10‡	-1.46†	-2.62‡	-2.40†

Table VIII. Probit regression of choice of outgoing interbank trades

We estimate this equation, $Prob(Trade_t = IB^{out}) = P(FC_{t-1}, CC_{t-1}, |I_{it}|, I_{it}^2, |Q_{jt}|)$, as a probit regression.

Incoming (outgoing) interbank trades are coded 0 (1). FC_{t-1} is a dummy coded 1 if the previous counterparty was a financial customer, CC_{t-1} and IB_{t-1} are defined similarly for commercial customers and other banks. I represents inventories, in millions of euros; $|Q_{jt}|$ represents the absolute size of the current deal, measured in EUR millions; $10\ mio_{t-1}$ is a dummy set to one if the size of the previous transaction was €10 million or larger. Significance at the 1, 5 and 10 percent levels indicated by ‡, † and *, respectively.

	Baseline Regression			Robustness Tests: Spot Trades Only	Robustness Tests: Interbank Trades Included
	Coefficient	**Std. Error**	***z*-Statistic**	**Coefficient**	**Coefficient**
Information Variables					
FC_{t-1}	-0.116	*0.116*	-1.00	-0.091	-0.256*
CC_{t-1}	-0.531‡	*0.055*	-9.60	-0.409‡	-0.672‡
IB_{t-1}					-0.214‡
$10\ mio_{t-1}$	0.650‡	*0.190*	3.43	0.770‡	0.657‡
Control Variables					
$\lvert I_{it}\rvert$	0.030‡	*0.011*	2.85	0.051‡	0.028‡
I_{it}^2	-0.001‡	*0.000*	-2.64	-0.002‡	-0.001†
$\lvert Q_{jt}\rvert$	0.029‡	*0.008*	3.58	0.070‡	0.028‡
Constant	-0.875‡	*0.044*	-19.92	-0.893‡	-0.728‡
McFadden's R^2		0.041		0.044	0.044
Observations		3,534		2,894	3,534

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. The horizontal axis is transaction-time. Vertical lines indicate the end of each calendar week.

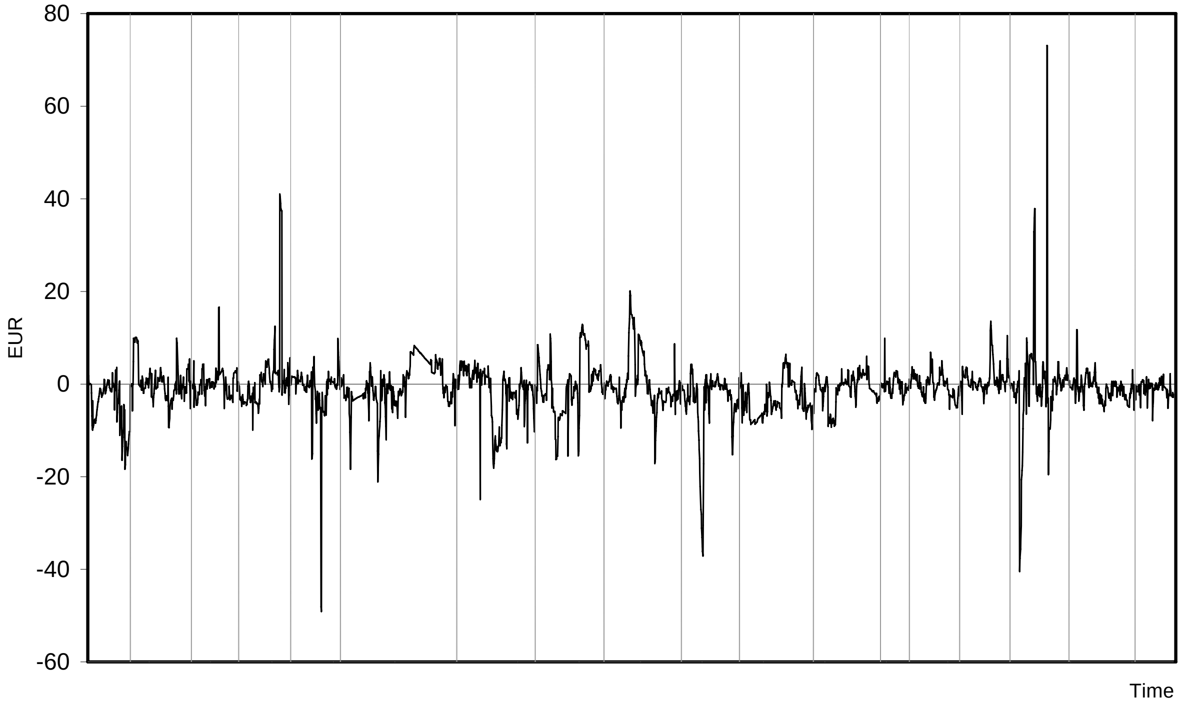

Figure 2: Intraday distribution of trades
The charts below show the average number of trades during each five-minute period of the trading day. Data come from a small bank in Germany and include all USD/EUR spot and forward trades during four months in 2001.

Figure 2A: Financial-customer trades

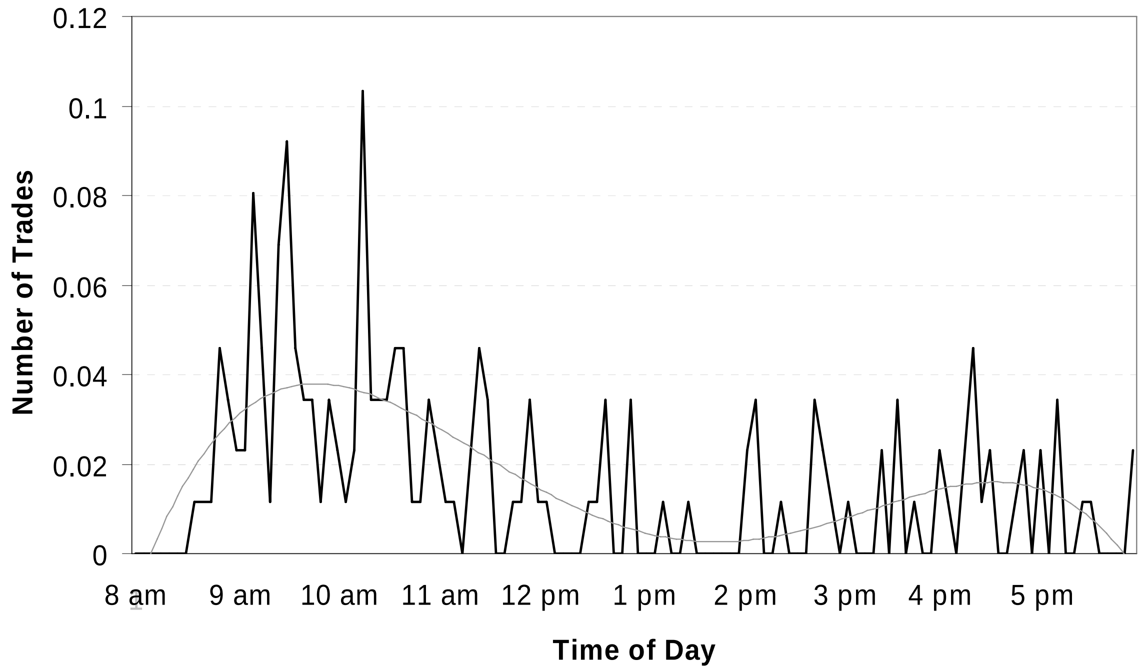

Figure 2B: Commercial-customer trades

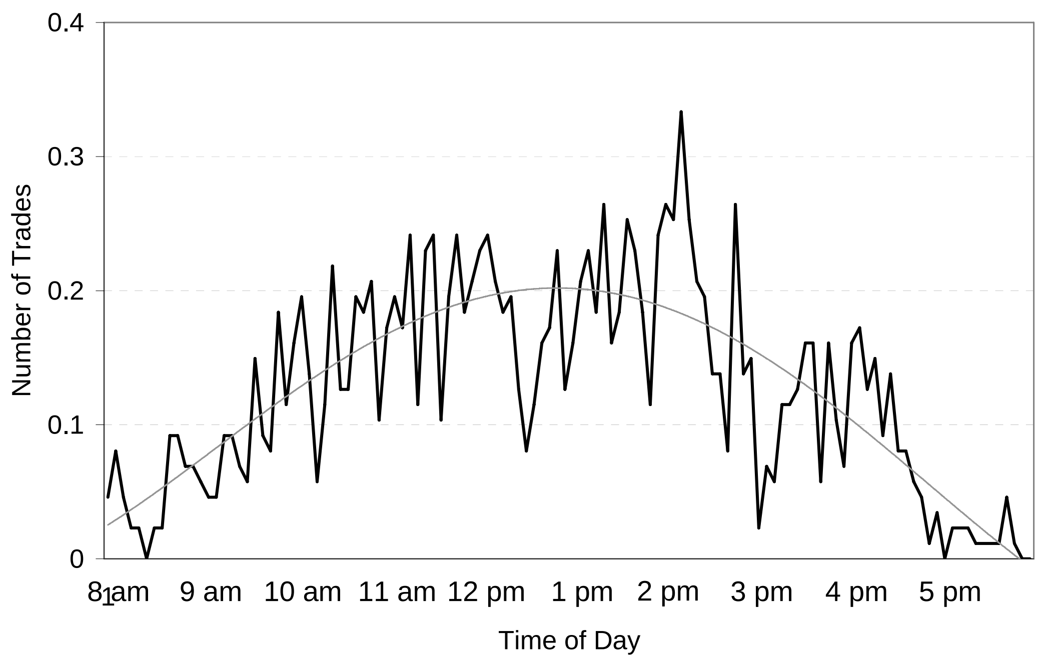

Table AI. Baseline Madhavan-Smidt model

We estimate this equation: $\Delta P_{it} = \alpha + \beta_1 D_t + \beta_2 D_{t-1} + \gamma_1 I_{it} + \gamma_2 I_{it-1} + \delta Q_{jt} + \varepsilon_t$.
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). These variables are interacted with dummy variables for the two counterparty groups, other dealers (*IB* for "interbank") and all customers (*CU*). Data include all incoming customer 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. Numbers in bold can be interpreted as the (negative of the) baseline half-spread.

	Baseline Regression		Robustness Tests		
			No Inventories	Spot Trades Only	Interbank Trades Excluded
	Coefficient	*Std. Error*	Coefficient	Coefficient	Coefficient
Constant	-0.590†	*0.23*	-0.426*	-0.383	0.070
Direction					
$CU \times D_t$	11.467‡	*0.50*	11.327‡	10.988‡	11.548‡
$CU \times D_{t-1}$	**-9.206‡**	*0.45*	**-9.186‡**	**-8.864‡**	**-10.025‡**
$IB \times D_t$	2.817‡	*0.69*	2.753‡	0.706	
$IB \times D_{t-1}$	**-1.579‡**	*0.48*	**-1.555‡**	**-1.025†**	
Inventory					
$CU \times I_{it}$	1.125‡	*0.38*		-0.064	0.855†
$CU \times I_{it-1}$	-1.264‡	*0.38*		-0.046	-0.974†
$IB \times I_{it}$	-0.259	*0.35*		-0.191	
$IB \times I_{it-1}$	0.133	*0.35*		0.187	
Trade size					
$CU \times Q_{jt}$	0.126	*0.39*	-1.001‡	-0.840‡	-0.001
$IB \times Q_{jt}$	-0.152	*0.40*	0.055	0.590	
Adjusted R^2	0.23		0.23	0.23	0.32
Observations	2,848		2,848	2,212	1,640

NOTES

[1] Our definition of a "customer" follows the market definition as any counterparty that is not another dealer.

[2] We show in Section IV that a similar analysis applies if the dealer uses direct trades to unwind his inventory.

[3] Bjønnes and Rime (2005) provides an excellent description of the market.

[4] This may change over time, but electronic foreign exchange trading platforms accessible to individuals are still relatively new.

[5] Electronic brokerages were not introduced until the early 1990s, so their dominance dates only from the late 1990s.

[6] 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.

[7] We exclude a few trades with tiny volumes (less than EUR 1,000) or with apparent typographical errors.

[8] Foreign exchange dealing institutions do not at present keep track, on an intraday basis, of their aggregate inventory across dealers and currencies.

[9] Inventory calculations are based on all trades for all tests, including those in which our statistical analysis is restricted to subsets of the data.

[10] A pip is equivalent to a tick: one unit of the smallest significant digit in an exchange rate as conventionally quoted. In the euro-dollar market, where the exchange rate averaged \$1.1128/€ during our sample period; a one-pip change from that level would bring the rate to \$1.112***9***/€. In this market one pip is approximately one basis point, since the exchange rate is near unity.

[11] It is also not possible to estimate spreads from matched pairs of trades. This technique is commonly used in analyzing bond markets (e.g., Goldstein *et al.* (2006), Green et al. (2007)), where trades can be identified by the amount traded, as in FX, and also by the particular bond.

[12] To understand this, note that if D_t is at the ask (bid) and D_{t-1} is at the bid (ask) then $D_t - D_{t-1}$ is two (negative two), but the price change should just equal the spread (negative of the spread). If both transactions are at either the bid or the ask, the price change and $D_t - D_{t-1}$ are both zero.

[13] Fewer than ten of the customer trades in our sample exceeded \$25 million. These trades were not excluded when calculating inventory levels.

[14] Results available upon request.

[15] Trade size is more likely to be irrelevant in the interdealer FX market, where trades are almost always either \$1 or \$2 million.

[16] One might naturally wonder about collinearity among our instruments. We are confident that this is not a problem, since the only pair of instruments with non-trivial correlation is $Fin \times Q_t$ and $Fin \times LG \times D_t$, and the qualitative conclusions from our baseline analysis are sustained when quantity variables are excluded.

[17] The market participants that checked our paper cautioned, however, that the magnitude of spreads on commercial trades has changed since 2001, even though the qualitative pattern identified here survives. In particular, intensified competition since 2001 associated with FXAll and other electronic communication networks has brought a compression in spreads to commercial customers.

[18] According to market participants, interbank trades on the electronic brokerages that now dominate that market are almost always \$1, \$2, \$3, or \$5 million.

[19] Spreads for BBB-rated corporate bonds average \$2.37 per \$100 face value for trades involving ten bonds or less but only \$0.37 per \$100 face value for trades involving over 1,000 bonds (Goldstein *et al.* (2006)). On the London Stock Exchange, average quoted spreads range from 165 basis points for the smallest stocks to 112 basis points for the largest stocks (Hansch *et al.* (1999): similar results are provided in Bernhardt *et al.* (2004)).

[20] Huang and Stoll (1997) propose yet another explanation for the negative relationship between adverse selection costs and transaction size in their analysis of equity market spreads. We pass over this explanation since it relies on the special properties of block trades. We exclude all trades over \$25 million from our regression analyses, so this

explanation cannot explain our results. Further, the management of large trades is carried out quite differently in FX than in equity markets.

[21] As interpreted here, asymmetric information has two roles in the Duffie et al. (2004) model. First, dispersed/asymmetric information about current prices generates the need to search in OTC markets. Second, information asymmetries determine the agency relationships within customer firms, between management and their traders, that in turn determine whether execution is rewarded.

[22] Strategic dealing may be more relevant in FX than the municipal or corporate bond markets, since most such bonds trade relatively infrequently so the information value of any trade may be negligible.

[23] This pre-occupation with standard practice may bring to mind the issues of collusion on the NASDAQ raised in Christie and Schultz (1994). However, since there are literally hundreds of dealers in the major currency pairs, and they are spread across the globe, it seems highly unlikely that collusion could maintain FX spreads for decades.

[24] The positive sign of the cointegration coefficient for cumulative interbank order flow provides further support for our hypothesis that commercial customers are net liquidity providers. Cumulative order flow from the bank's ultimate liquidity suppliers must have a negative relationship with the exchange rate, and only commercial order flow satisfies this requirement.

[25] We are not the first to note that some price discovery happens in the interdealer market (Evans and Lyons 2006), but to our knowledge we are the first to note that price discovery *cannot* happen in the customer market, and that therefore *all* price discovery must happen in the interdealer market.

[26] The choice between limit and market orders will also hinge on market conditions, such as the width of the bid-ask spread and the depth of the book (Biais *et al.* (1995), Goettler *et al.* (2005), Lo and Sapp (2005)).

[27] Our conclusion that dealers will place outgoing/market orders after trading with "informed" customers is consistent with the finding of Bloomfield et al. (2005) that informed traders "take (provide) liquidity when the value of their information is high (low)." In their experimental setting information is most valuable when it is new. In FX markets, information is newest right after a dealer trades with an informed customer, which corresponds to the time we suggest the dealer will place the outgoing/market order.

[28] Though it would be ideal to develop a formal model of this price discovery mechanism, space constraints preclude presenting a fully articulated model in this paper. Indeed, the influence of information on order choice has only begun to be analyzed theoretically (Kaniel and Liu (2004)), in part because such models are of necessity extremely complex. These complexities will multiply when information is incorporated into a two-tier market structure.

[29] A more general framework would replace $|I_{it}|$ with $|I_{it}\text{-}I^*_t|$, the gap between actual and desired inventory. However, currency dealers' desired inventory is usually zero.

[30] These inventory management practices are consistent with practices at large banks (Bjønnes and Rime (2004)). Further extensive parallels between our bank's behavior and that of large banks are documented in the Appendix.

[31] This preference is supported by the transactions data. Our dealer's mean interbank transaction size was only €1.42 million (Table 1), the maximum interbank trade size was only € 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 €1.94 million and the standard deviation of (absolute) transaction sizes was €1.63 million. By contrast, interbank trades averaged closer to $4 million prior to the emergence of electronic brokerages (Lyons (1995)).