PRICE DISCOVERY IN CURRENCY MARKETS

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

Abstract

This paper makes three contributions to our understanding of the price discovery process in currency markets. First, it provides evidence that this process cannot be the familiar one based on adverse selection and customer spreads, since such spreads are inversely related to a trade's likely information content. Second, the paper suggests three potential sources for the pattern of customer spreads, two of which rely on the information structure of the market. Third, the paper suggests an alternative price discovery process for currencies, centered on inventory management strategies in the interdealer market, and provides preliminary evidence for that process. We suggest more broadly that the price discovery process will vary with market structure, and that our proposed mechanism may apply to liquid two-tier markets in general.

March 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, Anthony Neuberger, Paolo Pasquariello, Uday Rajan, Stefan Reitz, Dagfinn Rime, Erik Theissen, and Dan Weaver for insightful comment

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 this process cannot be the familiar one based on adverse selection and customer spreads, including the observation that the cross-sectional pattern of customer spreads is directly contrary to the predictions of that theory. Second, the paper suggests three potential sources for the pattern of customer spreads, two of which rely on the information structure of the market. Third, the paper suggests an alternative price discovery process for currencies, centered on inventory management strategies in the interdealer market, and provides preliminary evidence for that process.

The paper's message is in fact broader than the foreign exchange market. We 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 where dealers can trade with each other separately, such as FX. Our proposed price discovery process is potentially relevant to all liquid, two-tier markets, including the U.S. Treasury market, the U.S. corporate bond market, and the London Stock Exchange. Evidence shows that the cross-sectional pattern of customer spreads in those markets violates the predictions of adverse selection in much the same way it does in currency markets.

The adverse-selection-based price discovery process asserts that dealers build into their price quotes the potential information revealed by a given customer transaction. The theory, articulated most notably in Kyle (1985) and Glosten and Milgrom (1985), highlights the information advantage customers often have over dealers. Dealers will rationally protect themselves from this adverse selection risk, as shown in Glosten and Milgrom (1985), by quoting a spread such that both bid and ask prices incorporate the expected information content of each trade. In this way prices move, on average, in the direction required by the information coming into the market. 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, then spreads should also be higher for specific customers, or customer types, that typically have information.

Most FX microstructure papers draw on adverse selection as their primary interpretive framework. This originally gained support from Lyons (1995), which shows that trade size and the interbank spreads of a particular dealer were positively related in 1992. 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 FX 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 FX 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, 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. These reportedly informed customers are financial firms, meaning asset managers such as hedge funds or mutual funds; the other broad category of customers is commercial customers, meaning firms that import or export.[1]

If adverse selection doesn't drive customer spreads in FX, what does? The paper's second contribution is to outline three factors that seem likely to be important. Of necessity, these represent just a subset of the factors that influence dealers; we focus on these three because practitioners indicate they are particularly important. The first factor, fixed operating costs, can explain the negative relation between trade size and customer spreads if some costs are fixed, but cannot explain the cross-sectional variation across customer types. To explain why FX spreads are larger for commercial than financial customers we suggest that asymmetric information – in the broad sense of information that is held by some but not all market participants – may influence spreads through two channels distinct from adverse selection, one involving market power and a second involving strategic dealing.

The market power hypothesis suggests that firms gain market power, even in a market with hundreds of competitors like FX, from holding information. It can be costly for customer firms to search out the best available quotes in the FX market, so each individual dealer can exert a certain amount of market power despite the competition. As suggested in Green *et al.* (2004), dealers may quote the widest spreads when their market power is greatest, and market power in quote-driven markets depends on knowledge of current market conditions. In FX, commercial customers typically know far less about market conditions than financial customers so they might be expected to pay wider spreads, as they do.

The second channel through which asymmetric information might affect customer spreads in FX involves strategic dealing. Building on abundant evidence that customer order flow carries information (e.g., Evans and Lyons (2004), Daníelsson *et al.* (2002)), we argue that rational FX dealers might strategically vary spreads across customers, subsidizing spreads to informed customers in order to gain information which they can then exploit in upcoming interbank trades. In standard adverse-selection models, by contrast, dealers passively accept the information content of order flow. The idea that dealers strategically vary spreads to gather information was originally explored in Leach and Madhavan (1992, 1993). When applied to two-tier markets in Naik *et al.* (1999) it implies that customer spreads will be narrower for trades with information, consistent with the pattern in FX.[2]

The paper's third contribution is to outline a process through which information may become embedded in exchange rates, and more generally to outline a process for price discovery in liquid two-tier markets. In contrast to adverse selection, which focuses on spreads in the customer market, our suggested process focuses on dealers' inventory management practices in the interdealer market. The mechanism is this: After trading with an informed customer, a dealer's information and inventories provide strong incentives to place a market order in the interdealer market. An informed-customer buy would thus tend 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 will generate 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 more likely to generate liquidity in the interdealer market than to move exchange rates.[3]

This view of dealer behavior differs critically from that of the “portfolio shifts” model of the FX market (Evans and Lyons (2002)). In that model, there are three rounds of trading. In the first, dealers absorb inventory from end-users; in the second round dealers trade with each other; in the third round dealers sell their inventory to end-users and prices adjust to reflect information. We suggest, by contrast, that prices begin to reflect information earlier, during interbank trading.

Our view of dealer behavior predicts a number of the key stylized facts in FX 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 the dealer is responding 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 will be substantially permanent, consistent with evidence in Lyons (2001) and Bjønnes and Rime (2005).

We test the two most direct 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. The evidence provides encouraging support for both implications.

Our analysis, though it concerns the microstructure of trading in currency markets, is also relevant to the exchange-rate dynamics literature within macroeconomics – which is analogous to the “asset pricing” literature within finance. A critical recent contribution to exchange-rate economics is been the recognition that order flow is a key determinant of returns (Evans and Lyons 2002). Since this insight is inconsistent with the traditional “efficient markets” view that has dominated exchange-rate theory since the 1980s, it is particularly important for the evidence to have a solid conceptual foundation. So far, 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 information in order flow gets into to the exchange rate. Our work is designed to 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 FX: (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 our data. Section II shows that customer spreads in FX are narrowest for the trades most likely to carry information. Section III discusses how operating costs, market power, and strategic dealing can explain this pattern. Section IV presents our interpretation of the price discovery process in currency markets, along with supporting evidence. Section V concludes.

I. FX MARKET STRUCTURE AND DATA

The currency markets make all other markets look tiny. FX trading averages almost $2 trillion per day (B.I.S. (2004)) – over twenty times daily trading on all NYSE stocks. An active currency trades as much in a half hour as a high-volume stock on the Paris Bourse trades in a day. About half of this trading takes place in the interdealer market (B.I.S. (2004)), trading in which is now largely carried out on order-driven electronic exchanges. The customer market, by contrast, is quote-driven.[4] Hundreds of dealers compete in the euro-dollar market, which accounts for almost a fifth of all transactions (B.I.S. (2004)). There is no significant retail component to FX trading; virtually all trading is carried out by institutions. Since currencies are important in commerce as well as finance, however, the institutional customer base for FX includes non-financial as well as financial firms. (Bjønnes and Rime (2005) provides an excellent description of the market.) In the foreign exchange market it is accurate to treat the dealers as the only intraday suppliers of liquidity. Trading in euro-dollar and other major currency pairs is never constrained by a shortage of a particular bond or stock. Furthermore, during our sample period only dealing banks had access to the interbank market. In consequence, there is no need to consider “latent liquidity” (Chacko *et al.* 2006).

Our data comprise the complete USD/EUR transaction record of a bank in Germany over the 87 trading days from 11 July 2001 to 9 November 2001. Though the data technically refer to the overall bank, they are an accurate reflection of a single dealer's behavior because only one dealer was responsible for the bank's USD/EUR trading. For each transaction we have the following information: (1) the date and time;[5] (2) the direction (customer 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.

Though our sample period includes September 11, 2001, that day’s events will not distort the results. The foreign exchange customer market functioned quite smoothly that day, though trading volume was low. This no doubt stems in part from the wide geographical dispersion of dealers around the world – indeed, dealers are dispersed even within New York City. The interdealer market also functioned smoothly that day, due in part to wise planning by the two major electronic exchanges. Both Reuters and EBS, though based in the United Kingdom, have servers in multiple geographic locations around the world performing real-time replication of all functions.

We include outright forward trades, adjusted to a spot-comparable basis by the forward points, as recommended by Lyons (2001). The bank's inventory position is inferred by cumulating successive transactions.[6] Following Lyons (1995), we set the daily starting position at zero. This should not introduce significant distortions since our dealer keeps his inventory quite close to zero, as shown Figure 1. 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. Table I provides basic descriptive statistics.[7]

A preliminary comparison of our dealer with the large dealers described in the literature is provided in Table II. Table III provides information on the size distribution of our dealer’s transactions. (We would, of course, prefer to present statistics on spreads themselves; however, such figures cannot be calculated from transactions data.) The small size of our dealer is reflected in his total daily trading value, average transactions per day, average inventory position, and mean absolute price change between transactions.[8] Our dealer is comparable in size to a NOK/DEM dealer at the large dealing bank examined in Bjønnes and Rime (2004). Our bank is probably a reasonably good representative of the average currency dealing bank because small dealing banks are far more common than large ones (B.I.S. (2002)). Nonetheless, big banks dominate currency dealing.

Despite the small size of our bank, our main qualitative conclusions should generalize to the entire foreign exchange market for at least four reasons. First, the FX market is extremely competitive. Hundreds of banks deal in the major currency pairs and even the largest dealer's market share is only on the order of 10 percent. In such a market, the behavior of any (successful) dealer should accurately represent the behavior of all (successful) dealers. Second, surveys of currency dealers reveal that the primary determinant of currency spreads is the conventional level of such spreads (Cheung and Chinn (2001)). Third, market participants consistently confirm that the pattern we identify is correct.

Finally, our small bank's behavior should be representative because it is broadly consistent with that of large banks in many well-studied dimensions. The Appendix provides a detailed comparison of our bank's pricing and inventory management practices with those of large banks analyzed in earlier studies. This analysis suggests that the following statements about larger dealers are equally true for our dealer:

- The baseline spread for interbank trades is on the order of two pips (or equivalently two ticks) [9]
- 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 lightning fast relative to traders in equity and bond markets.

These parallels support the reasonableness of generalizing from this bank to the market.

II. THE CROSS-SECTIONAL PATTERN OF CURRENCY SPREADS

This section evaluates whether adverse selection is likely to drive price discovery in the FX currency markets by testing its implications for the cross-section of FX customer spreads. We analyze two different models of spreads, the Huang and Stoll model (1997) and the Madhavan and Smidt model (1991). The behavior of spreads fails to conform to the predictions of adverse selection for both models. Among the violations of adverse selection, we especially note that FX customer spreads are wider for small trades than for large trades and that they are wider for commercial customers than for financial customers.

A. Preliminary Analysis

Our transactions data cannot be used to calculate direct measures of spreads, since they only indicate one side of each quote. Nonetheless, we can extract indirect 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 equals the price change. Even if the market is volatile, any associated distortions should ultimately average to zero if there is no dominant trend.[10]

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:

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

P_t is the price and ΔP_t is the price change from period t-1 to t: $\Delta P_t = P_t - P_t$, measured in pips. D_t is the direction of trade [D_t = 1 (-1) if the counterparty is a buyer (seller)]. The coefficient β should represent half of the spread. To understand this, note that if D_t is at the ask and D_{t-1} is at the bid then $D_t - D_{t-1}$ is two. The price change in this case, however, should just equal the spread. Likewise, if D_t is at the ask and D_{t-1} is at the bid then $D_t - D_{t-1}$ is negative two and the price change should equal the 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.

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 distinguish 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 large (*Lg*), medium (*Md*), and small trades (*Sm*).

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)). Since our data operates on transaction time and involves trades in which the dealer sets the price, our dependent variable is the sequence of prices on transactions initiated by customers. We exclude the few transactions over $25 million because such trades essentially represent a distinct market: customers hire dealers to manage such trades by breaking them up into smaller interbank transactions.[11] Interbank and customer trades may not strictly be comparable, given the structural differences between quote- and order-driven markets, so we exclude interbank transactions.

The results of this analysis on 1,640 customer trades, shown below, show a negative relation between trade size and spreads, inconsistent with adverse selection:

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

The half-spread is, on average, 8.4 pips for small customer trades, 6.9 pips for medium-sized trades, and only 3.2 pips for large trades. All coefficients are strongly statistically significant. The Wald test statistic for the difference between the coefficients on large and medium-sized trades is a highly-significant 7.04, indicating that spreads are indeed narrower for large trades than for medium-sized trades, on average. The Wald statistic for the large-small distinction is a similarly-significant 34.87. The Wald statistic for small versus medium-sized trades is only 1.46, which is not statistically significant. This may reflect a lack of power due to the paucity of medium-sized trades in our sample (see Table 3). We note in passing that the constant term is positive and statistically significant. This reflects the incompleteness of our sample of trades: when interdealer trades are included the constant becomes economically small and statistically insignificant.

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

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

As before, the coefficients are strongly statistically significant. The Wald statistic of 12.38 for the difference between coefficients on financial and commercial trades is highly significant, indicating that financial customers pay narrower spreads, on average, than commercial customers. Adverse selection predicts the opposite.

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 a small bank. Thus the counterparty-based tiering identified here should be roughly accurate for our bank.

Trade Size and Counterparty Type: Since commercial trades tend to be smaller than financial trades, it is possible that the commercial trades' smaller spreads simply reflect their 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, at least for small and medium-sized trades. Commercial spreads are wider larger than financial spreads for both small and medium-sized trades, though the spreads are roughly the same for large trades. Wald tests confirm the broad outlines of this pattern though they do not indicate statistical significance for all the pairwise differences across coefficients. This could reflect a lack of power due to the low number of observations for large commercial-customer trades and for medium-sized trades for both customer types (see Table 3).

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. For this reason, we 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 foreign exchange ─ 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}, a 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 we assume to be zero since this is generally the case for FX dealers (Bjønnes and Rime 2005). The model permits dealers to manage existing inventories by shading prices to customers (e.g., quote lower prices when his inventory is high), which implies $\theta > 0$.

Dealer *i* updates his expectation of the asset's fundamental value in light of the private information revealed by the direction of the previous trade as well as public news: $\mu_{it} - \mu_{it-1} = (\lambda S/2)D_{t-1} + \varepsilon_t$. The term

$\lambda S/2$ captures the information effect of trade direction and is thus a direct manifestation of adverse selection. The shock ε_t is a serially uncorrelated and is assumed to reflect public information. 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 $\{Lg, Md, Sm\}$ and counterparty type $\{FC, CC\}$. We again use GMM with Newey-West standard errors.

The results, shown in the first two columns of Table V, broadly confirm our two preliminary findings. First, baseline spreads are wider for small trades than for large trades. Second, baseline spreads are wider for the small and medium-sized trades of commercial customers than for the equivalently-sized trades of financial customers. As noted above, these results are not implied by adverse selection.

The estimated adverse-selection coefficients, λ, also provide no support for the influence of adverse selection in this market. Four of these are insignificant, and the patterns implied by the two significant coefficients do not conform to adverse selection. While the theory predicts that λ should be larger for financial customers than commercial customers, the two significant coefficients both apply to commercial customers. Furthermore, the theory predicts that λ rise monotonically with trade size, but the point estimates of λ vary non-monotonically with trade size for both commercial and financial customers. If we instead take statistical significance as our guide, the estimates of λ imply that small and medium-sized commercial trades have wider spreads than large commercial trades, the opposite of what we would expect under adverse selection.

We note in passing that the results suggest that inventory levels are not relevant to FX customer spreads: 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 shading prices to customers.

As shown in Table V, we test the robustness of these results in three ways. First, we rerun the regressions excluding inventories, which appear to have no influence. Second, we rerun the regressions using only spot transactions. Forward transactions account for 20 percent of all trades, so their inclusion could impede direct comparisons with earlier papers, which focus exclusively on spot trades. Finally, 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. Our results consistently prove robust.

2. 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.[12] For this reason, we also estimate the Madhavan and Smidt model (1991), in which trade size is informative. This is perhaps used more commonly in FX microstructure than any other model of spreads (see, for example, Lyons (1995) and Bjønnes and Rime (2005)).

The Madhavan and Smidt model (1991) model assumes that agent *j* calls dealer *i* requesting a quote on amount Q_{jt} which is determined as follows: $Q_{jt} = \xi(\mu_{jt} - P_{it}) + X_{jt}$. The term μ_{jt} represents agent *j*'s expectation of the asset's true value, conditional on a noisy private signal of the asset's true value and on a noisy public signal. X_{jt} represents agent *j*'s liquidity demand. The coefficient ξ is positive, so demand increases with the gap between the true value and the quoted price; it is because of this relation that trade quantity conveys information.

Dealer *i*'s regret-free price, P_{it}, is determined as $P_{it} = \mu_{it} + \zeta(I_{it} - I_i^*) + \chi D_t$, where variables are defined as above. If dealers shade prices to manage existing inventories, $\zeta < 0$. After solving for conditional expectations and taking first differences, one arrives at the following expression for the price change between incoming transactions, $\Delta P_{it} = P_{it} - P_{it-1}$:

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

The intercept, α, should be zero if the dealer's desired inventory is zero. If the dealer shades prices in response to inventories then $\gamma_2 = |\gamma_1|/\phi > |\gamma_1| > 0 > \gamma_1$. 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 of the parameters in this model. First, it could influence the values of β_2, the coefficient on lagged direction, which according to the model is the negative of the baseline half-spread. Under adverse selection these 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, Q_t: under adverse selection this 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 could also reflect 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. Adverse selection and this inventory effect, which we will refer to as a "prospective" inventory effect, are observationally equivalent here.

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 that captures 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 interact the key variables with dummies corresponding to trade size and counterparty type. We include the same robustness tests used when estimating the Huang-Stoll model plus one more. The new test excludes the quantity variables but not the inventory variables. As before, our results consistently prove robust.[13]

The results of this analysis, presented in Table VI, do provide no more support for adverse selection in the FX customer market than our earlier results. We begin with the estimated baseline half-spreads. For financial customers, 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. Wald tests are not able to detect statistically significant differences within the financial-customer spreads, which could be due to a lack of power since our bank primarily serves commercial customers. However, Wald tests indicate that financial-customer spreads are indeed smaller than commercial-customer spreads for small trades (marginal significance 0.0002) and medium trades (marginal significance 0.03). Wald tests detect no significant difference between spreads on large trades for financial and commercial customers, which is not surprising since the spreads themselves were insignificant in this size category for both types of customers.

The other two potential sources of evidence for adverse selection are equally unsupportive. The point estimates of the coefficient on trade size, Q_t, are positive for both commercial and financial customers but neither is significant at standard significance levels. Finally, we turn to the ratio between the coefficients on lagged and current direction, $\phi = |\beta_2|/\beta_1$. Though adverse selection predicts these will vary inversely with trade size, the estimates rise with deal size for the financial customers — the customers to which the theory most likely applies. Similarly, though adverse selection predicts that ϕ is smaller for financial than commercial customers, the reverse is true for medium-sized and large trades.

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.

C. Discussion

This section has provided evidence inconsistent with the hypothesis that price discovery in FX follows the standard adverse selection model. That is, dealers do not appear to adjust customer spreads to protect themselves against the trades' likely information content. 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 banks, to whom we sent the paper for a reality check, all confirmed both this qualitative pattern as well as the specific magnitudes we find for spreads during our sample period. Indeed, they assert that the pattern just identified approximates common knowledge within the FX market: the pattern is known by virtually everyone who trades, and virtually everyone who trades knows that virtually everyone else who trades knows it, etc. Only rank beginners might find the pattern unfamiliar, they claim.[14]

We stress that these results apply only to the *customer* FX market. Existing evidence suggests that interdealer spreads have a different relation to both counterparty type and trade size. Interdealer spreads are likely 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.[15]

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 we find that spreads are positive for both commercial and financial customers. Indeed, negative spreads do not appear to happen in this market even in the extreme: the authors have had extensive conversations with many FX dealers, most of whom work at large banks and none of whom were aware of negative spreads.[16] Marsh and O'Rourke (2005) also suggest that the negative relationship may instead reflect feedback trading, and provide evidence that commercial-customer trades are indeed influenced by lagged returns. Our evidence provides a further indication that the feedback-trading interpretation is more likely to be correct.

How do our results compare with results from other markets? Adverse selection has a mixed record of success in explaining spreads elsewhere. 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. As shown by Theissen (2003), those specialists provide price improvement when adverse selection costs are lowest and adjust quotes the most after trades that did not get price improvement.

Despite this evidence and the general acceptance of the adverse selection hypothesis in the FX microstructure literature, there are a number of markets where adverse selection cannot explain the behavior of spreads. A negative relationship between trade size and spreads has been observed in the 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)). We argue below that this may reflect the lack of anonymity in this market coupled with asymmetries in bargaining power among different customers, properties that are also shared with FX. A negative relationship between spreads and trade size has also been documented in the U.S. corporate bond market (Goldstein *et al.* (2006)) and the London Stock Exchange (Hansch *et al.* (1999)).[17]

We argue below that the negative relationship between trade size and spreads in these markets may reflect their two-tier structure coupled with their high liquidity, properties they share with FX.

III. OPERATING COSTS, MARKET POWER, AND STRATEGIC DEALING

This section examines possible alternative explanations for the pattern of FX customer spreads documented in the previous section. We begin by considering the components of the standard paradigm beyond adverse selection, which are: monopoly power, inventory risk, and operating costs. Pure monopoly power is unlikely to be important in FX, where hundreds of dealers compete intensely. Inventory risk can also be ruled out as a determinant of our pattern, since the prospective inventory effect implies a positive relationship between spreads and trade size (Ho and Stoll (1981)), rather than the negative relation we observe. In addition, inventory risk is invariant across customers, so this element cannot explain the relationship between spreads and customer type.

The remaining component of the standard paradigm is operating costs. In discussing the negative relationship between spreads and trade size on the London Stock Exchange, Angel (1996) and Hansch *et al.* (1999) note that such a relationship could arise if per-unit processing costs are smaller for large trades. This occurs with fixed costs, which certainly exist in FX and, in conversation, foreign exchange dealers themselves suggest that they are relevant. However, the relationship between costs and customer type seems unable to explain the smaller spreads paid by financial 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. In short, the standard paradigm can explain the relationship between FX customer spreads and trade size but not the relationship between spreads and customer type.

One might wonder whether the customer-based differences in spreads could result merely from differences in the intraday pattern of trading. If commercial customers trade more intensely during hours when spreads are widest, they will naturally pay larger spreads on average. Interdealer spreads (the only spreads for which intraday patterns are available) are widest during the London morning and during the

FX market's "overnight" period, which lasts for a few hours after about 5 pm London time (Payne (2003)). Financial-customer trades tend to be concentrated during the London morning hours (Figure 2A,B), while commercial-customer trades are more evenly distributed across the trading day. 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 have more success in explaining why FX spreads vary across counterparty types. One theory suggests that information about current market conditions provides market power which, in turn, affects spreads. The other theory suggests that dealers strategically vary spreads across customers in an attempt to gather private information about near-term exchange-rate returns.[18] It is our view that both of these information-based forces operate simultaneously with operating costs.

The additional 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 would doubtless exhaust the patience of our readers. For example, dealers acknowledge that conforming with standard market practice, presumably in order to maintain a strong reputation, is the single most important determinant of spreads (Cheung and Chinn 2001). This does not, however, help us understand the issue on which we focus: why market practice itself involves wider spreads for smaller trades and for commercial customers. We choose to focus on the market power and strategic dealing theories based on the view that these theories reveal the factors most influential in determining this pattern. Our readers in the market have found no reason to complain about this choice.

A. Market Power

Green *et al.* (2004) shows that variations in market power between dealers and their customers may explain why spreads are inversely related to trade size in the U.S. municipal bond market. That paper points 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 – are hard to ascertain. The customers who make smaller municipal bond trades tend to know the least about current market conditions, so they have the least market power relative to the dealers and are charged the widest spreads.

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. What little market information is available to customers is expensive. Financial customers typically purchase real-time information and hire professional traders who know how to interpret it. By contrast, most commercial customers do not purchase that information and do not hire sophisticated traders, so their traders are usually considered relatively uninformed about market conditions.

Information about market conditions is not the only potential source of financial customers' market power. In discussing the NYSE, Angel (1996) notes that

> a dealer knows that an unsophisticated individual who places a small order 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 for small orders (p. 4).

Duffie *et al.* (2004) develops this insight into a formal model and shows that bargaining power in OTC markets partly reflects the alternatives to trading immediately, alternatives that are determined by the relative costs and benefits of further search. In currency markets, the benefits to search are smaller for commercial customers than for financial customers.[19] FX traders at commercial firms are not always rewarded for finding better prices; for them, trading is typically just one of many administrative responsibilities. By contrast, FX traders at financial customers are often explicitly evaluated on execution quality. Since FX traders at financial firms perceive greater benefits to search, they are more likely to keep at it until they find a narrow spread. Knowing this, dealers may not even try to quote them a wide spread. Financial customers' market power may also come 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 counterparty-based tiering of currency spreads may also reflect “strategic dealing,” in which the dealers adjust their pricing so as to extract private information from their customers. Order flow at large banks includes information about upcoming high-frequency currency returns, as documented by Evans and Lyons (2004) and Daníelsson *et al*. (2002). Evidence from equity markets confirms that access to real-time order flow information can provide an informational advantage (Anand and Subramanyam (2005)). Thus it seems logical that FX dealers might try to capture a larger share of the most informative order flow, since the information could help increase returns and/or lower risk through better inventory management, better pricing on upcoming trades, and better speculative positioning.[20]

Our own small bank’s order flow need not be hugely informative for strategic dealing considerations to influence its customer spreads. As noted earlier, dealers’ dominant concern when setting FX spreads is conforming to standard practice (Cheung and Chinn (2001)), so strategic dealing will (at least) indirectly influence spreads at small banks so long as it directly influences spreads set at large banks.[21] 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.

The insight that market makers might strategically manipulate spreads to increase the information value of order flow is not new. Leach and Madhavan (1992, 1993) use equity-market inspired models to demonstrate that market makers may adjust prices early in a trading session to enhance later profitability. This general insight motivates the empirical tests of Hansch and Neuberger (1997), which “provide[s] evidence that dealers [on the London Stock Exchange] do act strategically, and that they deliberately accept losses on some trades in order to make superior revenues on others” (p. 1). Evidence for this type of strategic dealing in an experimental market which shares many properties with the FX interdealer market is presented in Flood *et al.* (1999).

Our evidence, however, concerns cross-sectional variation in spreads rather than variation across time. An equity-inspired strategic dealing hypothesis that overlaps more substantially with our own is presented in Naik *et al.* (1999), whose analysis of a two-tier market indicates that customer spreads will be narrower for more informed customers, consistent with the pattern we document for FX. The motivation for this conclusion is similar to the first two outlined above: after gleaning the information included in the current customer trade, dealers can profit more in subsequent trading. However, the Naik *et al.* model also concludes that customer spreads vary positively with trade size, while our data fits the opposite pattern.

Consistent with the Naik *et al.* (1999) model, Reiss and Werner (2004) report that "[d]uring the period of our sample, London [Stock Exchange] dealers were known to solicit large customer orders, even if the terms were unfavorable. The explanation most often given for this behavior was that dealers were 'purchasing' information …" (p. 625). In the context of the FX market we hypothesize that information content is not only inversely related to trade size, as on the London Stock Exchange, but it is typically higher for financial customers than commercial customers. Evidence that financial transactions in FX carry information for high-frequency exchange-rate returns is provided in Froot and Ramadorai (2005). Evidence that the information in financial transactions tends to exceed the information in commercial transactions, at least information about high-frequency dynamics, is provided in Fan and Lyons (2003) and Carpenter and Wang (2003). Related evidence is presented in Ramadorai (2005), whose buy-side view of the phenomenon supports our sell-side view. His analysis of the FX transactions of a large set of asset managers finds that spreads are narrower for the managers that produce higher (risk-adjusted) FX returns.

There is a sense in which the strategic dealing hypothesis mirrors the market power hypothesis. In the game between dealers and their commercial customers, dealers gain market power from their knowledge of market conditions, on the basis of which they extract wider spreads. In the game between dealers and their financial customers, both sides are well informed about market conditions but financial customers also have private information relevant to near-term exchange-rate dynamics. Financial customers view themselves as exploiting the market power associated with their private information to extract smaller spreads. Dealers simultaneously view themselves as strategically setting small spreads to increase their business with privately informed customers and learn their information. Both sides are right.

IV. Price Discovery in Foreign Exchange

The evidence presented so far shows that spreads in the FX customer market are inversely related to a deal'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, relevant to FX and other liquid two-tier markets, and provides evidence in support of that proposal. Asymmetric information is the centerpiece of our story, as it must be, but we suggest that information influences inventory management and order choice in the interdealer market rather than spreads in the customer market. Our proposed mechanism thus reflects institutional features of the FX market, such as its two-tiered structure and the importance of the interdealer market for inventory management, that distinguish FX from the simpler market structures assumed in adverse-selection models.

Our proposed price discovery mechanism differs in a key way from the familiar "portfolio shifts" model of the FX market articulated in Evans and Lyons (2002). In that model, dealers first absorb inventory from end users, then trade that inventory among themselves, and finally sell the inventory to other end users. The exchange rate moves to reflect information only during the customer trading of round three. If one were to graft our price discovery framework to the Evans and Lyons model, however, one would conclude that the exchange rate moves to reflect information during the interbank trading of round two. Nonetheless, our proposal creates a coherent picture from disparate stylized facts from FX microstructure.

A. The Mechanism

Our proposed price discovery mechanism involves dealers' interbank trading in response to customer trades. We focus on the interbank market because the evidence presented above implies that a given trade's potential information content is not embedded in customer prices. We infer that price discovery does not happen in the customer market and must therefore happen in the interdealer market.[22] 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), Lyons (1997)).

Consider a dealer whose inventory rises abruptly in response to an incoming customer call. Since FX dealers prefer to have zero inventory (this is documented for our dealer in the Appendix and for large dealers in Bjønnes and Rime (2005)), our dealer will most likely try to offload the new inventory to another dealer. In FX the dealer must choose between "indirect" trading in the order-driven broker market or "direct" trading in the regular quote-driven market.

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 with 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 informed. In this case the dealer has three incentives to exploit the immediacy offered by market orders: He has information, he has inventory with its inherent risk, and 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. 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. In short, we suggest that dealers using the brokers market to manage inventory will have a stronger tendency to place market orders after informed customer trades than after uninformed customer trades.[23] The connection to price discovery is direct: brokered interdealer prices will tend to move in the direction indicated by informed trades.

If our dealer chooses to trade directly, a modified version of this cost-benefit analysis still applies. 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 who chooses the direct interdealer market has strong incentives to call another dealer after trading with an informed customer and may be more likely to wait for incoming calls after trading with an uninformed customer.

The overall conclusion is consistent regardless of whether a dealer chooses to manage his inventory via brokered or direct trades. After trades with informed customers, a dealer will be more likely to make a (parallel) outgoing trade than after trades with uninformed customers. As a result, interdealer prices will tend to move in the direction required by the information contained in customer 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.)[24]

In equilibrium, trading might cease entirely if (a) customer identity were the only factor determining whether a dealer makes an outgoing trade and (b) customer identity were a reliable indicator of whether the customer is informed at a given point in time. Under this combination of circumstances dealers would only place market orders after informed-customer trades, so placing a limit order would be a recipe for losing money and the market might cease to exist. In reality, however, customer identity is imperfectly correlated with a given customer's private information at any point in time. Furthermore, the decision to make an outgoing trade depends on more than just customer identity, as shown below.[25]

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). 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, which is defined as buy-initiated interdealer transactions minus sell-initiated transactions. (In the order-driven or brokered portion of the interdealer market, the initiator of a transaction is considered to be the dealer placing the market order; in the quote-driven or direct dealing portion of that market, the initiator is the dealer that calls out. In both cases the initiator makes an "outgoing trade.") 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 this natural question regarding the strategic dealing hypothesis: If dealers subsidize the trades of their informed customers in order to buy information (in effect), 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 profit from anticipated high-frequency exchange-rate moves.

C. Other 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 actually given to a customer from a preliminary quote provided at that moment by the relevant 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 banks with the most customers are positively correlated with each other but negatively correlated with trades by small 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. Likewise, when a large dealer is accumulating speculative positions via interdealer market orders, his counterparties banks tend to be small, which suggests that other large dealers have information that helps them avoid picking-off risk.

The most direct testable implications of our proposed price discovery mechanism concern the likelihood of outgoing interbank transactions. Under our proposal, dealers should be more likely to place interdealer market orders after trades with financial customers than after trades with commercial customers, since financial customers are considered more informed. Similarly, dealers will be more likely to place interdealer market orders after larger trades than after small ones, even after controlling for inventory, since large trades are considered relatively informative.

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

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

Our hypothesis concerns the first three variables, dummy variables for lagged financial-customer trades, FC_{t-1}, lagged 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.[26] 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 VII, support our view that the likelihood of an outgoing interbank transaction is higher when the most recent 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 a lofty 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.

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 importance of inventory level for dealer order choice helps the market avoid no-trade equilibria and maintain low interbank spreads by reducing the signal/noise ratio associated with outgoing interbank trades. Importantly, the significance of inventory levels eliminates one alternative possible explanation for the influence of trade size on order choice. Specifically, it appears that the influence of large trades in our regression does not reflect the inventory risk they bring, since the influence of the dealer’s inventory level per se is already accounted for. 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.[27]

To summarize: This section proposes a mechanism through which price discovery may occur in FX. We first note that price discovery must happen in the interdealer market since customer spreads vary inversely with a trade's likely information content. We then show both conceptually and empirically that dealers are more likely to make outgoing interbank trades after trading with informed customers than after trading with uninformed customers. This could be the force that drives interdealer prices in the direction implied by the information customers bring to the market.

V. CONCLUSIONS

This paper's overall message is that the standard adverse selection model of price discovery may not apply in liquid two-tier markets. Instead, we propose a new price discovery process relevant to such markets. Our data comprise the complete USD/EUR trading record of a bank in Germany over four months in 2001. The paper first shows that adverse selection in the customer market cannot be the mechanism through which price discovery happens in FX. Spreads on normal-sized currency trades vary inversely with trade size and are wider for commercial customers than for financial customers. Both components of this pattern are inconsistent with adverse selection, since FX dealers consider large trades to be more informative than small trades and financial customers to be more informed than commercial customers.

The paper then highlights three hypotheses that help explain the cross-sectional pattern of currency spreads. We first note that operating costs are largely fixed in FX, which could help explain the negative relationship between trade size and spreads. The customer-based variation in spreads could be explained by Green *et al.*'s (2004) market-power hypothesis. This hypothesis 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 will pay the widest spreads, as they do. The customer-based variation in spreads could also reflect dealers' attempts to strategically gather information about near-term returns (Leach and Madhavan (1992), (1993), Naik *et al.* (1999)). Dealers may subsidize trades with informed customers in order to learn the information embedded in their trades, 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 three hypotheses we highlight here do not exhaust the long list of factors dealers consider in setting spreads; we focus on these because we consider them most influential, and leave the others for future research.

The paper finishes by proposing a new price discovery process relevant to liquid two-tier markets like FX. This proposal creates a coherent picture of the FX price discovery process by fusing existing empirical evidence on FX microstructure, including our own, with insights from mainstream microstructure. We first note that, since customers' information is not immediately reflected in the prices they pay, price discovery must take place entirely in the interdealer market. We focus our analysis, therefore, on dealer behavior in the interdealer market, a market that is important for inventory management (Lyons 1997). The key mechanism behind our suggested price discovery process involves the dealer's response to individual customer trades. We suggest that after transactions with informed customers, dealers will tend to make parallel outgoing interdealer trades – placing a market order in the order-driven component of the market, 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 transactions with uninformed customers, by contrast, dealers will be relatively likely to place parallel limit orders or to wait for incoming calls.

Our proposed mechanism implies that dealers should be more likely to place outgoing interdealer trades after informed customer trades, and we provide evidence that this is true for our dealer. 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, as in FX, in 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 relies solely on the existence of two tiers and high liquidity. 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.

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, 2004, Financial intermediation and the costs of trading in an opaque market, Graduate School of Industrial Administration, Carnegie Mellon University, Working Paper 2004-11.

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.

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.

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: Small Banks and Large Banks Behave Similarly

This Appendix documents that our small-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.[28] 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 large 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 large dealers despite his small volume.

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 small bank in Germany, except preferred customer trades, over the 87 trading days between July 11th, 2001 and November 9th, 2001. For comparison purposes we focus on statistics based exclusively on the small bank's spot trades.

	Small Bank in Germany	**B.I.S. (2002) per Bank**	**Lyons (1995)**	**Yao (1998)**	**Bjønnes and Rime (2005)**		
					Four Dealers, Range	**DEM/USD Dealer**	**NOK/DEM Dealer**
	87 Trading Days in 2001[a]	April 2001	5 Trading Days in 1992	25 Trading Days in 1995	5 Trading Days in 1998		
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. (Modified) Huang and Stoll (1997) model

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

ΔP_{it} is the change in price between two successive 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. 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	Interbank Trades Included
	Coefficient	Std. Error	z-Statistic	Coefficient	Coefficient
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‡
$\lvert I_{it}\rvert$	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.

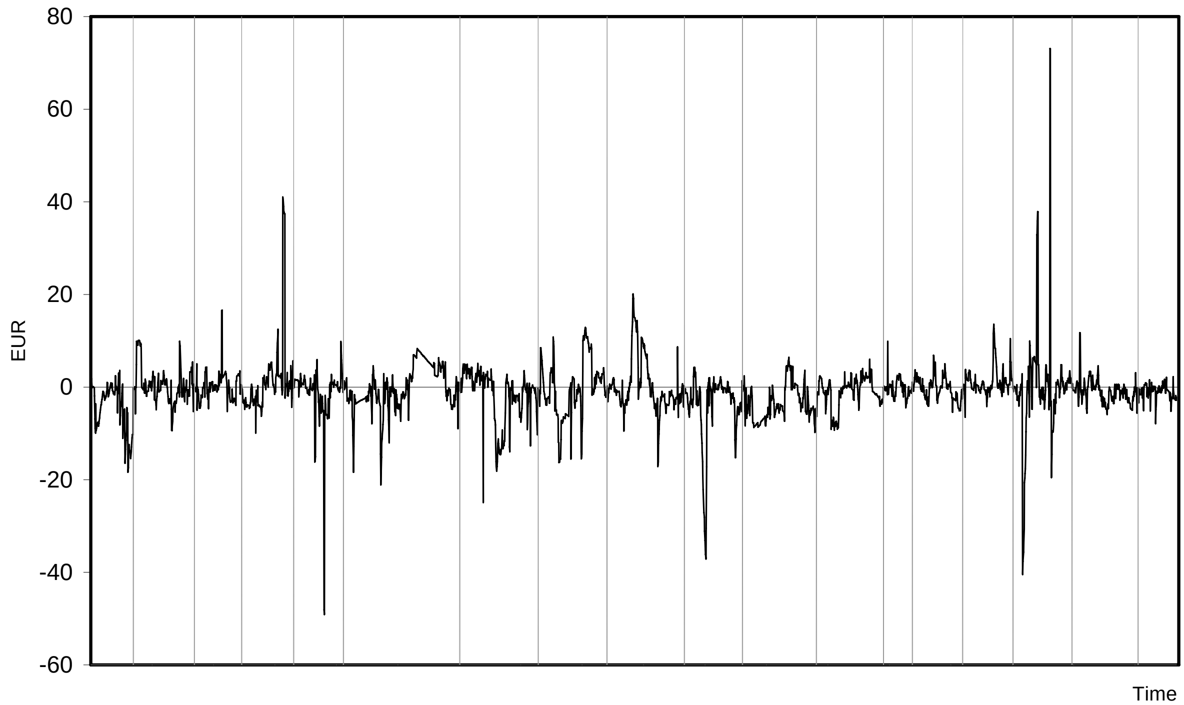

Figure 2: Intraday distribution of trades
The charts below show the average number of trades during each five-minute period of the trading day. Data come from a small bank in Germany and include all USD/EUR spot and forward trades during four months in 2001.

Figure 2A: Financial-customer trades

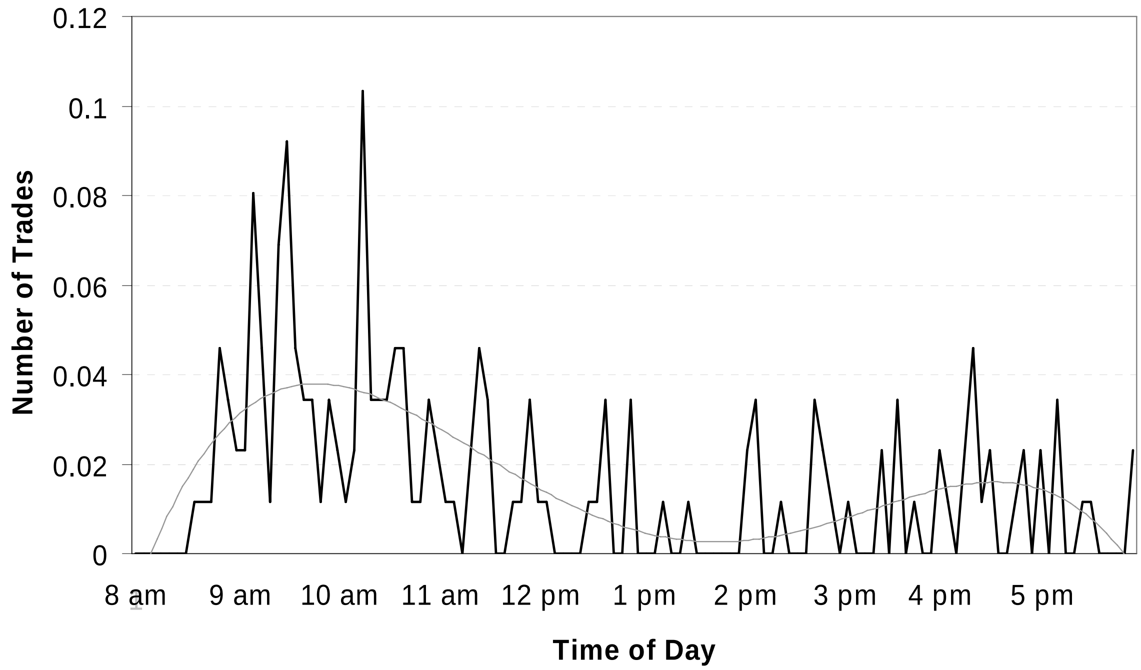

Figure 2B: Commercial-customer trades

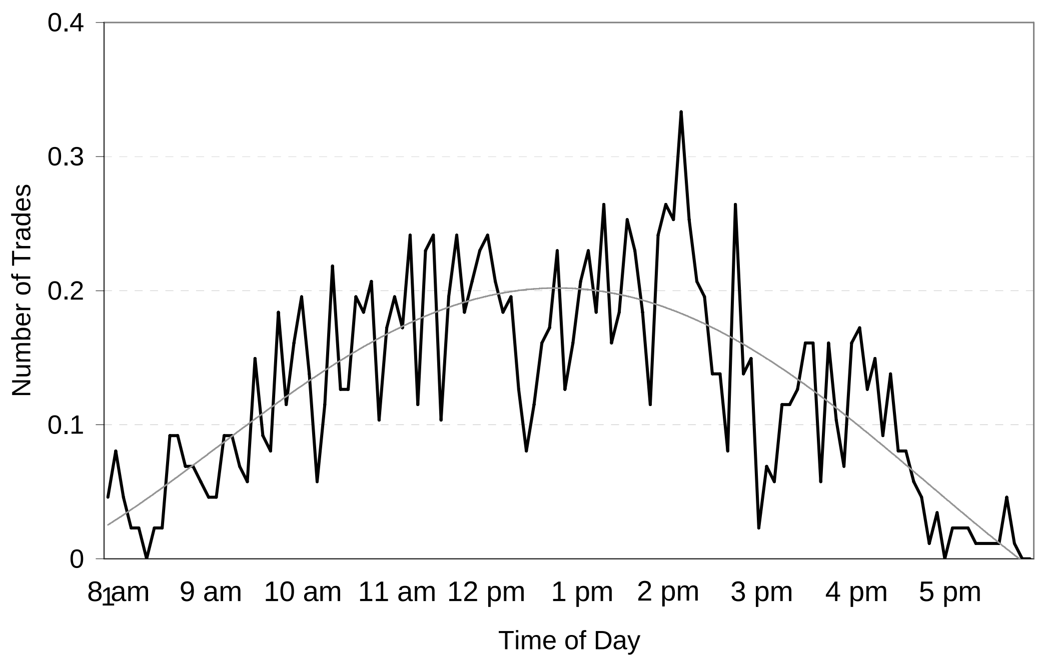

Table AI. Baseline Madhavan-Smidt model

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] Buy-side evidence for strategic dealing of this sort is provided in concurrent work by Ramadorai (2006).

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

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

[5] 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.

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

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

[8] The large mean absolute change in transaction price between successive trades, 10.7 pips, presumably reflects the relative infrequency of transactions at our bank as well as the high proportion of small commercial customer trades, which tend to have wide spreads (as we document below).

[9] 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.

[10] 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.* (2004)), where trades can be identified by the amount traded, as in FX, and also by the particular bond.

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

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

[13] 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 x Q_t and Fin x LG x D_t, and the qualitative conclusions from our baseline analysis are sustained when quantity variables are excluded.

[14] 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.

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

[16] Negative spreads, or the equivalent, are sometimes observed, such as the U.S. treasury market during the late 1980s.

[17] 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)).

[18] 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.

[19] 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.

[20] 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.

[21] 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.

[22] 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.

[23] 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)).

[24] 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.

[25] 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.

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

[27] 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.

[28] 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)).