

ASYMMETRIC INFORMATION AND THE FOREIGN-EXCHANGE TRADES OF GLOBAL CUSTODY BANKS

Carol Osler

Brandeis International Business School

Tanseli Savaser

Williams College

Thang Nguyen

Brandeis International Business School

Abstract

This paper provides the first rigorous empirical analysis of markups on custodial foreign exchange trades. It finds that they substantially exceed relevant benchmarks such as interbank half-spreads. We trace this to an information asymmetry: custodial bank dealers know more about their prices and bid-ask spreads than their client funds. We also examine the asset managers' continued heavy reliance on this high-cost approach to trading when alternatives are available with lower markups. We provide evidence that this choice does not reflect ignorance of the cost differential. Analysis relies on the complete foreign exchange trading record of a mid-sized global custody bank during calendar year 2006.

June 21, 2011

Osler: Cosler@Brandeis.edu. Savaser: Tanseli.Savaser@Williams.edu. Nguyen: TanThang@Brandeis.edu. The authors are grateful for the insightful comments of Jim Angel, Dan Bergstrasser, Blake LeBaron, Lukas Menkhoff, Dagfinn Rime, Elif Sisli, seminar participants at Bilkent University, Brandeis University, Wesleyan University, and Williams College, and participants at the Sixth Annual Central Bank Workshop on the Microstructure of Financial Markets (2010). We take responsibility for remaining errors.

ASYMMETRIC INFORMATION AND THE FOREIGN-EXCHANGE TRADES OF GLOBAL CUSTODY BANKS

This paper examines the foreign-currency trades between institutional asset managers and their dealers at global custody banks. Global custody banks serve client funds by tracking asset values, repatriating income, settling trades, and investing funds as directed. The reputedly high markups of custodial dealers on foreign-exchange trades has “been a vexing problem for investors for decades” (Galanek, 2010). The issue came to public prominence in 2009 when the California Public Employees’ Retirement System and the California State Teachers’ Retirement system sued State Street Bank, a top global custodian, alleging that State Street set unreasonably high markups.¹ Our analysis provides the first rigorous empirical investigation of this topic. It thus follows Christie and Schultz (1994a, 1994b) in providing evidence relevant to a topic of current legal concern.

Our evidence indicates that markups on custodial foreign exchange trades greatly exceed relevant benchmarks such as interbank half-spreads. We trace this to an information asymmetry: custodial bank dealers know more about their prices and bid-ask spreads than their client funds. The paper also examines the strong revealed preference of asset managers for this high-cost approach to trading despite the availability of low-cost alternatives. We provide evidence that this choice does not reflect ignorance of the cost differential and suggest that it is, instead, a rational response to the asset managers’ incentives and constraints. The constraints include a second information asymmetry: the funds know more about their execution costs than their investors.

When a client fund decides to trade an equity or bond it typically instructs its “fund accountant” at the custodian to trade foreign currency as required. With such “non-negotiated” transactions, a client learns the details some time later. According to the client’s preferences this could be the next day or as part of a regular summary of recent trading activity received after some days or weeks. These reports historically specified only the transaction’s date and price but not the time or the bid-ask spread (Kothare and Raynor, 2010). This information asymmetry is an unintended by-product of the very reason that custody banks exist: client funds prefer to outsource administrative burdens. In effect, global custody

¹ More recently, State Street has been sued by the Arkansas Teachers Retirement Fund and the Attorneys General of Virginia and Florida are involved in similar suits against Bank of New York/Mellon.

banks serve as “headache relief” for their client funds, for whom foreign-exchange trading – with associated back-office processing and the potential for out-trades – is often viewed as a headache.

The client funds are not forced to trade this way: alternatives with known low execution costs are available. Client funds can contact regular OTC currency dealers, who stand ready during business hours to quote bid and ask prices in response to customer inquiries. After considering the quoted prices a client fund decides whether to trade and, if so, in which direction (buy/sell). The client funds can also call dealers at global custody banks who stand ready to trade at their own quoted prices. With such “negotiated” trades the asset managers and the dealers are roughly equally informed: they both immediately know the trade’s date, time, and price, as well as the bid-ask spread. From the perspective of the client funds, however, the information advantage from negotiating a price is accompanied by the disadvantage of higher administrative costs, since they must not only contact the dealers but also arrange for confirmation. Further, these costs require higher investor fees, and fees are easier for the investors to monitor than bid-ask spreads.

We examine the consequence of these asymmetries on custody bank dealers and their client funds using the complete foreign exchange (FX or forex) trading record of a mid-sized global custody bank during calendar year 2006. The sample comprises between 75,000 and 100,000 transactions in 27 currencies vis-à-vis the U.S. dollar (USD) with total value between \$50 and \$100 billion. Like most global custodian banks, this one trades a broad range of currencies; it trades with a broad range of asset managers and trust funds; and it provides a broad range of services to its clients. Dealers at this custody bank focus primarily on client service: in this they differ from regular OTC dealers and also from dealers at some large custody banks, for whom speculative position-taking in the interbank market is a critical component of foreign exchange desk profitability (Bjønnes and Rime, 2005; Ramadorai, 2006). The investigation employs panel regressions in which standard errors cluster simultaneously along the dimensions of time and currency and are robust to heteroskedasticity and autocorrelation. All results are supported by numerous robustness tests.

Our analysis focuses first on the custody-bank dealers. We hypothesize that their information advantage on non-negotiated trades allows them to set higher markups. This “fog effect” is closely related to Green et al.’s (2007) “market power” hypothesis that OTC dealers quote wider bid-ask spreads when they have a stronger information advantage relative to specific customers (Angel, 1996). The information

asymmetry driving Green's market power hypothesis implies variation in spreads across customers according to their familiarity with market conditions. The fog effect implies, instead, variation in spreads according to the opacity of different trade types (negotiated vs. non-negotiated), and not necessarily variation across client funds. We test the fog effect using regressions, but simple averages tell the story. Our custody bank's average markup on non-negotiated transactions, 22.4 basis points, substantially exceeds the 3.4 basis-point average markup on negotiated transactions.

The possibility that custodial traders are aware of the opacity of the custodial relationship, and keen to preserve its advantages, is raised by a statement that forms part of the allegation from the original complaint filed against State Street Corporation:

When discussing inquiries by the Pension Funds about providing "transparency" in foreign exchange execution costs, one Senior Vice President with State Street California commented to other State Street executives that, "[i]f providing execution costs will give [CalPERS] any insight into how much we make off of FX [foreign exchange] transactions, I will be shocked if [a State Street V.P.] or anyone would agree to reveal the information." (Complaint page 3, paragraph 4).

We examine whether custody banks attempt to preserve the naturally-occurring ambiguity about markups on non-negotiated trades by avoiding extreme prices. If the prices paid are far beyond the day's interbank trading range, custody clients could infer that the custodial markup is substantial. Custody banks may thus reduce markups when they send a trade through a sub-custodian – as required for currencies subject to exchange controls, for example – since the sub-custodian charges its own markup. Likewise, markups should rise with return volatility because the day's range – and thus a safe range for pricing – tends to be wider for more volatile currencies. A positive relation between markups and exchange-rate volatility is also predicted for regular OTC dealers because volatility increases inventory risk (Ho and Stoll, 1981). This prediction is not relevant to custodial dealers on non-negotiated trades, however, because such trades are covered in the interbank market before being priced so they bear no inventory risk.

The ambiguity preservation hypothesis gains consistent empirical support from our regression analysis and, like the fog effect, it is equally apparent in less sophisticated summary statistics. The average markup on emerging-market currencies that involve sub-custodians, 12.8 basis points, is less than half the 39.5 basis-point average markup associated with other emerging market currencies. The influence of volatility is illustrated by the Hong Kong dollar, which is so tightly fixed to the US dollar that asset managers can easily infer their transaction costs. The ambiguity preservation hypothesis would predict that markups on non-negotiated Hong-Kong dollar transactions would be relatively small, and indeed the

0.5 basis-point average markup on such transactions is tiny compared with the 22.4 basis-point average across all currencies.

The paper also notes that custodial foreign exchange trades tend to be executed towards the end of the day, which could reflect two motivations. The custodial dealers may wait to observe the day's trading range before setting prices, with the goal of maximizing their markups. Alternatively, the dealers may wait to net small transaction amounts into larger trades to exploit transaction economies of scale. We find that larger transactions tend to be executed more swiftly, which supports netting as the dominant motivation for late-day pricing.

Our analysis of client fund behavior examines why only a minority of custodial trades are negotiated: in our sample this fraction is just three percent. According to some practitioners, the infrequency of negotiated custodial trades suggests that asset managers have not "been doing basic best-execution due diligence" (Galanek 2010) and are unaware of either the cost difference or that custodians trade currencies as a principal (custodians trade equities as agents, a service for which they are compensated with a commission).

Most of the funds in our sample seem unlikely to be ignorant, however. U.S. regulators consider any institution with over \$10 million in assets under management to be "sophisticated" (McGeehan, 2010) and the vast majority of funds in our sample exceed this benchmark. Sophisticated asset managers would have heard of the reputed high costs of custodial foreign exchange trades because the issue has been broadly discussed in the financial press and at foreign exchange conferences for about a decade.

We hypothesize that asset managers generally do not negotiate trades with their custodial dealers because this approach is more directly costly to themselves. While the administrative burdens of non-negotiated trades are largely borne by custodians, negotiated trades burden the asset managers with additional administrative expense: the client funds must hire dedicated traders (or allocate some of the portfolio manager's time for trading) along with dedicated staff to handle back-office processing and compliance. Asset managers are also constrained by a difference in the investors' perception of costs between negotiated and non-negotiated trades. Bid-ask spreads on non-negotiated trades are extremely difficult for underlying investors to monitor.² In contrast, the high administrative costs associated with negotiated trades must be passed on to investors as higher fees, which are relatively easy to monitor.

² Transactions cost analysis (TCA), standard for equity and bond trades, is not well established in foreign exchange.

We suggest that asset managers weigh their own advantage against their clients' disadvantage, balancing their own fixed administrative costs of negotiated transactions against the variable opportunity costs to investors' from non-negotiated transactions. Those opportunity costs rise with transaction size, so asset managers aware of the cost differential will be more likely to negotiate a price and forego the convenience of non-negotiated transactions when transacting large amounts.

We analyze the determinants of transaction type (negotiated, non-negotiated) using probit regressions in which the key independent variable is transaction size. Consistent with our hypothesis that asset managers are aware that negotiated transactions have smaller markups, the estimated likelihood that a transaction is negotiated rises strongly with trade size. With effects evaluated at sample means, a \$10 million transaction is more than twice as likely to be negotiated as a transaction worth \$10 thousand.

Overall, our analysis shows that custody-bank dealers and asset managers both fulfill their primary professional responsibility, which is to enhance shareholder returns subject to the constraints and incentives built into the trading situation. The custody-bank dealers do this by trading judiciously with clients. When a custody-bank dealer carefully assesses whether a given customer is willing to pay a higher markup, it acts like a dealer in any market, whether it be the OTC currency market, the municipal bond market, or the market for automobiles. The institutional asset managers enhance shareholder returns by outsourcing administrative services to their custodians and avoiding investor backlash from higher fees.

The rest of this paper has five sections. Section I, which follows, describes our data. Section II presents our analysis of custody-bank markups on foreign-exchange trades. Section III extends our analysis of custodians to consider the timing of their trades. Section IV examines the client funds, focusing on the choice between negotiated and non-negotiated transactions. Section V concludes.

I. DATA

Custody banks administer assets worth about \$100 trillion worldwide (Institutional Investor, 2007). The business has expanded rapidly in recent decades, propelled by growth in the asset management industry overall and by a 1974 U.S. law requiring that pension funds separate investment management and custody management services. Global custody banks, which administer internationally-owned assets, handle roughly one third of the industry's total (Institutional Investor, 2007). Foreign exchange trading is one of the services provided by global custodians, and it has been highly lucrative in

recent years. Extrapolation from our database suggests that the custodial industry's global profits from foreign exchange trading are on the order of \$14 billion annually.³ Despite its immense size, the custody business has rarely been examined by academicians and never by researchers in microstructure.

After a massive consolidation during the past decade, the industry is now dominated by a few large banks; the top fifteen custodian banks manage over eighty percent of total assets (Institutional Investor, 2007). This consolidation presumably reflects, at least in part, the industry's intense computational requirements and consequent economies of scale. As shown in Schmeidel et al. (2006), doubling the assets under management of a mid-sized European custody bank would increase its total costs by only 70 percent. Cullinan et al. (2005) reach a similar conclusion.

We analyze the complete record of foreign-exchange transactions between a mid-sized global custody bank and its client funds during calendar year 2006. For each transaction the data include: (i) the two currencies, (ii) the amount traded in both currencies, (iii) the transaction price, (iv) the time the trade was requested and the time the trade was carried out, (v) a variable indicating whether the trade was for a special purpose such as repatriating investment income or corporate action (meaning, for example, the exercise of warrants or participation in a tender offer), (vi) a code distinguishing each specific fund (note: individual funds may be part of large fund families but these are not explicitly identified), (vii) the fund's assets under management on December 31, 2006, and (viii) the custody bank's income from the transaction measured in US dollars.

We exclude all transactions that do not involve the US dollar (e.g., transactions between yen and euros), trust-fund transactions, transactions of funds listed as having assets below \$1,000, currencies that account for less than one percent of all transactions, and corporate-action transactions (which likewise account for only one percent of transactions). The final sample includes between 75,000 and 100,000 transactions in 27 currencies, worth in aggregate between \$50 and \$100 billion (see Table 1).⁴ About 20 percent of the transactions are intended to repatriate investment income; the rest are presumably intended for asset allocation. Transaction sizes range widely: the income-repatriation transactions are sometimes tiny while the largest single trade was worth almost \$500 million.

³ In our database, foreign exchange trading volume represented 18 percent of fund net asset value and average markups were 20.4 basis points. Applying these fractions to the estimated net asset value at global custody banks in 2007, \$37 trillion, we infer that global custodians may have earned on the order of \$14 billion.

⁴ We leave some figures ambiguous to protect the anonymity of the custody bank.

Our sample of client funds includes on the order of 500 individual asset managers with mean assets under management of about \$700 million. Of these, roughly 40 percent had net asset value below \$100 million, one third had net asset value between \$100 and \$600 million, and one fifth had net asset value in excess of \$600 million.

Only a small minority of transactions in our sample – 3 percent – were negotiated directly with the dealers. These negotiated transactions tend to be large, averaging \$2.6 million. By contrast, non-negotiated transactions for asset-allocation purposes average \$0.8 million; non-negotiated transaction for income repatriation average \$0.3 million. Even though negotiated transactions tend to be large, many large trades are not negotiated. Defining a “large” transaction to be worth \$5 million or more, we find that only 13.5 percent of large transactions are negotiated.

The funds that make negotiated transactions tend to be big and active. The average such fund has roughly \$2.0 billion of assets under management and traded about \$1.6 billion in 2006; the average fund that does not make negotiated transactions has only \$600 million in assets under management and traded \$200 million over the year. Negotiated transactions are more common for the three major currencies than for other developed-country currencies and least likely for the emerging-market currencies that do not involve sub-custodians.

Custodial banks typically aggregate individual transaction requests for a given currency into a single sum for pricing. For example, if one fund intends to purchase Thai shares worth 1,100 baht while another intends to liquidate 100 baht of dividend income, the quantities would be offset. The dealer would purchase 1,000 baht on the interbank market and the markup would be set on the basis of the 1,000 baht amount. Netting is handled automatically by the custodian’s information systems and “the FX manager simply converts the ... aggregate/net positions to/from the base currency. ... Settlement is no more than book entry transfer between the currencies concerned” (Joseph, 2005). For clarity we label a purchase or sale of the netted amount as a “trade” and we label the client funds’ individual requests to buy or sell currency as “transactions.”

When we analyze custodian behavior our main interest is in the proportionate markup as perceived by the dealer when s/he prices each trade, a markup quoted on the basis of the netted amount. To reconstruct these netted amounts we identify clusters of transactions sharing a currency and an exact price that occur within five minutes of each other. For each cluster we assign the trade amount to be the

net of the component amounts; dollar income is the sum of income from the individual transactions; the markup is that income divided by the trade amount. The associated fund is identified as the one that made the largest component transaction.

After clustering, the sample includes roughly 25,000 trades with average trade size of about \$1.7 million. This is comparable to standard trade sizes in the electronic foreign-exchange interbank market (EBS and Reuters Matching), but it is large relative to most equity trades. Jones and Lipson (2003), for example, report a mean trade size of \$39,000 for their NYSE sample. The custody-bank's income-repatriation trades, after clustering, average around \$600,000 each.

Roughly 20 percent of (clustered) trades were in euro-dollar. Dollar-yen and sterling-dollar each represent a further 10 to 12 percent of trades. The Hong Kong dollar and Korean won each represent between seven and eight percent of trades. The currencies of Australia, Canada, and Switzerland each represent four to five percent of trades; all other currencies represent less than three percent of trades. As shown in Figure 1A, trades tend to be largest for the major currencies.

We formally measure the custody bank's markups as the proportionate gap between the price it charges the customer (P_C) and the price at which it covers the trade. For 16 of the 27 currencies in our sample, the trade is covered in the interbank market at price P_{IB} , so the markup is $|(P_C - P_{IB})/P_C|$.⁵ Trades for the remaining 11 currencies involve a sub-custodian, which is essentially a bank located in the currency's sponsoring country. The sub-custodian purchases the currency in the local interbank market and charges the main custodian a marked-up price, P_{SC} , in which case the main custodian's markup is $|(P_C - P_{SC})/P_C|$. A sub-custodian is involved whenever a country imposes exchange controls and also in a few other emerging markets. In Brazil, for example, the involvement of a local sub-custodian is essential because formal government approval is required which, in turn, requires the submission of proof that trades are associated with specific underlying investment transactions.

On *trades* that are not negotiated, markups should never be negative because dealers cover their positions before setting prices for the client. On *transactions* that are not-negotiated, by contrast, recorded spreads can be negative if a transaction in one direction is netted into a trade in the other direction. Indeed, 8.4 percent of the original transactions have negative markups. As expected, after clustering only 0.2 percent of trades have negative markups, and almost all of those correspond to negotiated trades. The

⁵ The decision to use the customer price as the scale factor was the custodian's.

0.2 percent share seems plausible because negotiated trades are subject to inventory risk. For comparison, Green et al. (2007) find an overall 1.4 percent loss frequency for (negotiated) municipal bond trades.

The custody bank's markups vary widely across currencies, as shown in Figure 1B. Average markups range from a minimum of 0.5 basis points for trades in Hong Kong dollars to 51 basis points for trades in (New) Turkish lira. The average markup overall is 20.4 basis points, which would imply an average bid-ask spread over twice as large. Consistent with the fog effect the average markup on non-negotiated trades is 22.4 basis points, well above the 3.4 basis point average markup on negotiated trades. Consistent with the ambiguity preservation hypothesis, the custodian's average markup on emerging-market currencies that involve sub-custodians, 12.8 basis points, is significantly lower than the 39.5 basis-point average markup associated with other emerging market currencies.

For comparison, our average estimated markup is substantially larger than the estimated markups on financial customer transactions reported in Osler et al. (forthcoming), which range up to about 6 basis points. Indeed, in 2006 interdealer half-spreads in euro-dollar averaged about 1 basis point (Osler and Sherman, 2009), and even a retail trader "with a \$500 margin account [could] get a three-pip [basis point] price in euro, 24 hours a day" (Euromoney, 2006).

Our average estimated markup is also substantially larger than the 9 basis-point "average transaction cost" on foreign exchange trades by pension funds estimated by Russell Investment Group (Collie, 2004). The Russell dataset, however, does not distinguish transactions with custody banks from regular OTC transactions. Further, those data include only the date and the traded price so execution costs are estimated by comparing each traded price to the day's average price or the price range. Our data focus exclusively on custodial trades and they allow us to reconstruct the exact price at which each trade was covered in the interbank market.

II. CUSTODY BANK BEHAVIOR: WHAT DRIVES MARKUPS?

To examine the ways in which asymmetric information influences custody-bank markups, we regress the markup for each trade t , $Markup_t$, on variables that capture potential manifestations of the information asymmetry between the custodian and its client funds, X_t , plus other relevant determinants suggested in the literature, Z_t :

$$Markup_t = \alpha + \beta X_t + \gamma Z_t + \eta_t \quad (1)$$

Since the residuals may be correlated across currencies and within a given day, statistical significance is evaluated using standard errors that are allowed to cluster simultaneously along the dimensions of currency and time. Standard errors are also robust to heteroskedasticity and autocorrelation. The independent variables included in X_t and Z_t are described below; Table 1 provides descriptive statistics.

A. Fog and Ambiguity-Preservation Hypotheses

The fog effect hypothesizes that custodial dealers charge wider spreads on the trades about which the client funds are least well informed. This could be achieved via a simple fixed effect, whereby markups on non-negotiated trades are automatically widened by a certain amount. Thus the asymmetric information variables, X_t , include a dummy to distinguish negotiated from non-negotiated trades, taking negotiated trades as the baseline.

The ambiguity-preservation hypothesis suggests that custodial banks attempt to maintain uncertainty about markups by ensuring that the price charged to the client, P_C , is not extreme. To formalize this, begin by supposing asset managers generally consider a markup of x percent or larger to be unreasonable.⁶ That is, if the price on a trade falls outside the day's reported interbank price range by x percent or more, the asset managers would tend to infer an excessive markup. Custody-bank dealers may thus strategically ensure that on day d the range of prices quote to customers, P_C , is no wider than this:

$$\text{MaxCustodyPriceRange}_d = \ln[P^H_d(1+x)/P^L_d(1-x)] \approx \ln(P^H_d/P^L_d) + 2x/100 \quad (2)$$

where P^H_d (P^L_d) is the day's high (low) interbank price.

The custodian's markup is determined not only by prices charged to customers but also by the prices at which trades are covered, P_{IB} or P_{SC} . For currency C , the main custodian's maximum sustainable markup consistent with preserving ambiguity is thus:

$$\text{MaxMarkup}_C \approx \text{Mean}[\ln(P^H_d/P^L_d)]/2 + x/100 - \text{AvgSubCustMarkup}_C - \text{AvgInterBkSprd}_C/2 \quad (3)$$

This expression highlights three factors of potential relevance for custodial markups under the ambiguity preservation hypothesis: 1) a currency's volatility; 2) the markups imposed by sub-custodians; 3) the interbank spread.

We calculate a currency's current (realized) volatility as the square root of the sum of five lags of squared daily returns (data from Global Insight); its annual volatility is the square root of the sum of

⁶ Foreign-exchange traders at regular dealing banks as well as custody banks have mentioned, anecdotally, that they are discouraged from spreading prices beyond the day's range plus/minus three percent.

squared daily returns over calendar year 2006. Annual realized volatility for all 27 currencies, shown in Figure 2A, is generally between 5% and 15% and tends to be highest for emerging markets without exchange controls. The regressions capture cross-sectional variation in volatility by including $CSVol_C$, currency C 's annual realized volatility divided by the (unweighted) average of annual volatility over the 27 currencies. The regression captures each currency's volatility history, $TSVol_C$, as its current realized volatility annualized and divided by its annual volatility.

The inclusion of volatility as a determinant of spreads is fairly standard in the literature, but the interpretation of this variable must be specific to the custodial situation and, within that, to the type of transaction. In the regular OTC currency market volatility can bring wider spreads when it reflects the adverse-selection risk associated with the arrival of fundamental information (Foucault, 1999). However, traders at small and mid-sized custody banks are not exposed to adverse-selection risk, since they generally do not take speculative positions. (At some larger custody banks, by contrast, foreign exchange dealers behave more like interbank traders, holding their own currency inventory and taking positions at will (Ramadorai 2006).) In the regular OTC currency market volatility also heightens dealer inventory risk and should thus raise bid-ask spreads (Ho and Stoll, 1981). This might be a concern for negotiated custodial trades, but it should not be of concern for non-negotiated custodial trades which carry zero inventory risk.⁷ To capture the potentially differing implications of volatility for negotiated and non-negotiated trades, we include not just $CSVol_C$ and $TSVol_C$, but also each variable interacted with the dummy for non-negotiated trades.

The second factor suggested by the ambiguity preservation hypothesis is the sub-custodial markup. Since we do not have data on these markups, we include a dummy variable for currencies that are traded through sub-custodians. The dummy enters by itself and also interacted with the dummy for non-negotiated trades. For thoroughness, we also include dummies for two other currency groups: (i) emerging-market countries that do not involve a sub-custodian and (ii) developed countries excluding the three majors. The majors – euro-dollar, dollar-yen, and sterling-dollar – are the base group. We also interact these currency dummies with the dummy for non-negotiated trades. The ambiguity preservation

⁷ Custodians also face negligible credit risk on foreign exchange trades because they hold the client's assets, including cash. To make doubly sure they may also place a lien on the assets (Joseph, 2005).

hypothesis predicts a negative coefficient on the interaction term between the sub-custodial currency and non-negotiated trade dummies, but otherwise predicts nothing for the other currency dummies.

The final ambiguity-preservation factor is the interbank half-spread. We calculate this as the (log) difference between daily exchange rates in Global Insight, which are the average of intraday interbank *mid*-quotes, and daily exchange rates from Olsen and Associates, which are the average of intraday interbank *ask* quotes (exchange-rate data cover 2006). Since these spreads are noisy, we include only the sample average for each currency. Average half-spreads are a few basis points for the currencies from developed economies and are generally between 10 and 30 basis points for emerging-market currencies (Figure 2B). We include the half-spread itself plus the half-spread interacted with the non-negotiated trade dummy.

The total extra markup associated with non-negotiated trades will be the sum of the fixed effect plus up to three variable components, specifically those associated with volatility, sub-custodial markups, and interbank half-spreads.

B. Traditional and Other Influences on Bid-Ask Spreads

The literature identifies at least three additional influences on bid-ask spreads that we include in Z_i : fund size, market liquidity, and trade size.

Our first measure of fund size is trading volume in billions of dollars. In OTC markets a dealer's most active customers, who generate the most trading profits, can generally expect a volume discount (Bernhardt et al., 2005). Indeed, one practitioner asserts that "the amount of spread [charged by custodians] is dependant upon the global custodians 'rating' of the client's importance to their business through segmentation exercises. Key factors are annual turnover and the nature of the overall relationship" (Joseph, 2005). We also include (log) fund net asset value in millions of dollars as of December 31, 2006, which could capture either fund size or fund sophistication. To capture "the nature of the overall relationship," we later add individual fund dummy variables.

To capture aspects of market liquidity beyond bid-ask spreads, we include a currency's overall trading volume in April, 2007 as a percent of global trading volume (Bank for International Settlements, 2007). In the model of Roşu (2009) traders are sensitive to execution delays, so spreads narrow in markets with high activity and low delays. It would be preferable to capture time-series dimensions of liquidity by including daily trading activity, but such data are not available in currency markets. As a crude alternative

we include day-of-the-week dummies, with Monday as the baseline day, to capture the tendency of market liquidity to be low on Friday afternoons in the active currencies.⁸

We include (log) trade size because microstructure theory suggests many reasons why this variable could influence spreads. Nonetheless, a quick review indicates that those reasons may not apply to custodial trades. Theory suggests that bid-ask spreads *rise* with trade size because larger trades bring higher inventory risk (Ho and Stoll, 1981). But inventory risk should only be important for the minority of trades that are negotiated, as discussed earlier. Alternatively, bid-ask spreads could rise with trade size because such trades bring higher adverse-selection risk (Glosten, 1989; Easley and O'Hara, 1987). But adverse-selection risk is unlikely to be a significant concern to dealers at small and mid-sized custodians because speculative position-taking is not one of their responsibilities; indeed, some global custodians prohibit it. Finally, bid-ask spreads could *decline* with trade size in liquid two-tier markets like foreign exchange, where dealers have a strategic incentive to attract large trades for their information value (Naik et al., 1997; Ramadorai 2008; Osler et al., forthcoming). However, information has little value to traders who do not speculate.

C. Results

The results of running equation (1) on the full sample provide substantial support for the fog and ambiguity preservation hypotheses (Table 2, column 1). The estimated total extra markup on non-negotiated trades is 20.9 basis points for a currency of average volatility and average interbank spreads. Of this, roughly half is driven by the fixed effect: the coefficient on the dummy for non-negotiated trades is highly significant at 11.9. By itself this easily exceeds interbank half-spreads on the currencies in our sample, which average just eight basis points (trading-volume weighted).

Consistent with the ambiguity-preservation hypothesis, volatility in the cross section has a large, positive, and highly significant effect on markups for non-negotiated trades. For a currency with average cross-sectional volatility, markups on non-negotiated trades are 9.0 basis points wider than for a currency with essentially zero volatility like the Hong Kong dollar. This difference accounts for almost half the 19.9 basis-point difference between markups on Hong Kong dollar trades and the overall average markup on all non-negotiated trades. Time-series variation in volatility has no influence on markups for non-

⁸ Regular OTC dealers report that anyone trading late on Fridays should expect to pay wider spreads.

negotiated trades. The estimated effects of both cross-sectional and time-series volatility on the markups for negotiated trades are positive but economically small and statistically insignificant. This suggests that inventory risk is not a major concern among custodial foreign exchange dealers even when they trade as OTC dealers.

The coefficient on the interbank half-spread itself is positive and significant, implying that markups respond positively to interbank spreads for negotiated trades. The coefficient on the interbank half-spread interacted with the non-negotiated dummy is negative and significant, consistent with the ambiguity-preservation hypothesis. Since both coefficients are below 0.1 in absolute value, their sum – which represents the overall effect of interbank spreads on markups for non-negotiated trades – is tiny. Consistent with another aspect of the ambiguity-preservation hypothesis, custodial markups decline by an estimated 14.8 basis points when a sub-custodian is involved in the trade. Otherwise, markups do not vary by currency group for either negotiated or non-negotiated trades.

The coefficients on the traditional determinants of bid-ask spreads generally have the expected sign and most are significant, though their economic influence is modest. Turning first to the effect of fund characteristics, we find a negative and significant coefficient on fund trading volume, suggesting that large funds do indeed get a volume discount. However, this variable's estimated influence is economically modest, since it suggests that a one-standard deviation rise in annual trading volume reduces the average markup by only 0.4 basis points. The negative coefficient on fund NAV, which is only significant at the 10 percent level, suggests that a one standard deviation rise in (log) fund size likewise reduces markups by 0.4 basis points. Since funds have difficulty ascertaining their custodial execution costs, there may be little competitive advantage to be gained from providing discounts based on trading volume or net asset value.

As predicted by theory, custody-bank markups are narrower for the most actively traded currencies, and the coefficient implies that markups between the least and most liquid currencies differ by 3.5 basis points, other things equal. The estimated coefficients hint that markups might be a bit high on Fridays, when liquidity is low, but the difference is not significant. (Later results show that the Friday effect is significant for the most actively traded currencies.) The results also indicate that markups rise slightly with trade size, though the effect is modest: a rise of one standard deviation raises markups by less than one basis point.

How much are funds paying their custodian in foreign exchange execution costs? These estimates imply that an average-sized fund that trades an average amount (with other variables likewise at sample means) would have paid on the order of \$200,000 annually in custodial markups. By contrast, a fund that was one standard deviation higher in terms of trading volume and (log) NAV would have paid on the order of \$2 million annually.

D. Robustness

We test the robustness of our results by applying different estimation approaches. To verify that the small minority of trades with negative markups do not distort the results, we run a censored regression on all trades with positive markups. The estimated marginal effects at sample means, reported in Table 2, second column, differ little from the coefficients in the previous regression. Notable but minor changes include the newly full significance of cross-sectional volatility on negotiated trades and the negative effect of market liquidity on all trades. Both changes bring our results more in line with established theory.

To allow for differences in “the nature of the overall relationship” (Joseph, 2005), we next add individual fund dummies. Some funds prefer to compensate their custodian by paying high up-front fees and low execution costs while others prefer the reverse, and these arrangements can vary in other ways. *F* statistics (unreported) indicate that the fund dummies are jointly statistically significant. Nonetheless, as shown in Table 2, third column, the inclusion of hundreds of fund dummy variables has little effect on the estimated coefficients. The coefficient on the dummy for non-negotiated trades, at 10.9; differs little from its previous value of 11.9; cross-sectional volatility still accounts for a large share of the cross-sectional variation in markups; and the presence of a sub-custodian still reduces markups by over 10 basis points. The coefficient on fund trading volume becomes insignificant: since the coefficient was negative in the regression without fund dummies, this suggests that active funds tend to compensate their custodians with smaller markups and higher fees, which seems logical. Trade size also becomes insignificant, which suggests that funds with higher average trade sizes wisely make the same choice.

We next examine the implications of constraining the sample, first by including only trades by funds that negotiate transactions. The decision to negotiate presumably reveals relatively strong concern about execution costs, so such funds might be treated differently. We run the regression both with and without fund dummies. As shown in the first and second columns of Table 3, the results are very close to

the results obtained with the full sample. The one minor difference is that the coefficient on fund NAV becomes insignificant. This could reflect the sample constraint, since it tends to be the largest firms that negotiate transactions. We infer that the higher markups on non-negotiated trades apply equally regardless of a fund's concern about execution quality, consistent with the fog effect.

We finish these robustness tests by running separate regressions for the most liquid currencies and for the emerging-market currencies. The currencies of the six countries in the most-liquid group – Australia, Canada, Europe, Japan, Switzerland, and the U.K. – account for about half of the trades in our sample and a bit more than 60 percent of trade value. The currencies of the 15 countries in the emerging-market group – Brazil, Chile, Hungary, India, Indonesia, Israel, Malaysia, Mexico, the Philippines, Poland, South Africa, South Korea, Taiwan, Thailand, and Turkey – account for 30 (12) percent of trades (trade value).

The estimated coefficients, shown in the third through sixth columns of Table 3, highlight subtle but interesting differences between the two currency groups. First, the coefficients indicate that the well-known market illiquidity on Fridays is an issue for the most liquid currencies but not for emerging-market currencies. In addition, time-series volatility does matter for negotiated trades in emerging-market currencies, though not for the liquid currencies. This asymmetry may reflect the higher turnover in liquid currencies, combined with the custodian's tendency to net trades. When the dealer negotiates a trade in a liquid currency it may already know it can cover the accumulated inventory with a transaction request that has not yet been executed. With less-liquid currencies, however, this may often be impossible.

Turning to our central hypotheses, we note that the total extra markup on non-negotiated trades is lower for liquid currencies than non-sub-custodial emerging-market currencies. In the regressions without fund dummies, the total extra markup is 16.1 for the liquid currencies and 31.0 for the non-sub-custodial emerging-market currencies. Similar conclusions hold when fund dummies are included in the regression.

The drivers of the extra markup on non-negotiated trades also differ across currency groups. For the liquid currencies, the extra markup is largely driven by volatility: indeed, both time-series and cross-section variation in volatility are significant while the fixed effect is essentially insignificant. The extra markup for emerging-market currencies, by contrast, comes primarily from a substantial fixed effect – the point estimate is as large as 24.2 basis points in the regression without fund dummies. For these

currencies, cross-sectional variation in volatility has a significant but more-modest effect than it has for the liquid currencies and time-series variation in volatility has no influence.

Despite the nuanced differences across regressions just discussed, these robustness tests strongly support the fog and ambiguity-maintenance hypotheses: Markups are substantially higher on non-negotiated trades than negotiated trades, much of this effect works through volatility, and the presence of a sub-custodian reduces custody-bank markups.

III. CUSTODIAL DEALER BEHAVIOR: TRADE TIMING

This section examines the timing of trades relative to the timing of transaction requests. On non-negotiated trades there is a delay between the time a transaction request is sent to the trading floor and the time it is priced. As shown in Figure 3A, these delays average about six hours for currencies without exchange controls and range from 6 hours to 3 days for currencies with exchange controls (average = 49.2 hours). Notably, the average delays are not especially high for the three currencies free of exchange controls that nonetheless involve sub-custodians, specifically Hong Kong, Israel, and the Philippines. This highlights the role of exchange controls, per se, in the longest delays.

The delays on liquid currencies seem to be a natural reflection of our custody bank's tendency to execute most trades towards the end of the day. Almost two thirds of transaction requests arrive between 8:00 am and 11:00 am; in contrast, that same share of trades are executed between 2:00 and 4:00 pm (Figure 3B). Most trading must be completed by 4:00 pm so the back-office staff can cross-check all trades for amount, direction, and price before they leave for the day.

The tendency for trading to peak towards the end of the day has been interpreted by some market practitioners as a reflection of the dealers' desire to know the day's price range when setting prices, which presumably helps them maximize markups. "[T]he custodian has a free option to set the price at the end of the day. At this point it already knows the full trading range and is able, if it chooses, to pick any price from this range" (Collie, 2004, p. 3).

Another natural source of delay would be transactions netting, which in turn may be motivated by economies of scale in interbank trading. For trades below roughly \$25 million, spreads quoted by regular OTC dealers to their customers tend to be inversely related to trade size and the differences have historically been substantial. Osler et al. (forthcoming) find that in 2001 financial customers paid negligible average spreads for trades over \$1 million while their average spread on a transaction below

\$0.5 million ranged from 14 to 28 basis points. In our sample, 52 percent of transactions in euro-dollar are below \$0.5 million, and of course the share for emerging market currencies, where transactions tend to be smaller, ranges up to 63 percent. This is not unusual: average equity trade sizes are around \$40,000 in the Paris Bourse sample of Biais et al. (1995) and \$39,000 in the NYSE sample of Jones and Lipson (2003).

The custody dealer's tendency to aggregate many individual transactions into larger, single trades suggests by itself that scale economies represent a significant motivation for the delays between transaction requests and trade executions. Without scale economies such aggregation can be costly not only because it relies on special software but also because it increases the share of transactions on which the custodian loses money to customers. When a buy transaction is aggregated into a large trade dominated by sell transactions, the buying customer pays the same price earned by the selling customers. Consequently, the buying customer makes money on the transaction and the custody bank loses money. As noted earlier, customers profited from 8.4 percent of transactions in our dataset.

We evaluate the markup-maximization interpretation of intraday trading patterns by considering the relationship between trade size and delay. If the dominant motivation for trading late in the day is to maximize the markup – and thus the profits per transaction – the incentive to delay transactions could be positively related to trade size and we would find longer delays for larger transactions; at the extreme, if every trade is delayed to maximize the markup, the delay would be unrelated to trade size. By contrast, if the late-day executions reflect netting and the benefits of scale economies, the tendency for delay should be more pronounced for small trades than large trades. Figure 3C compares, for each of our 27 currencies, delays on transaction requests for amounts above and below \$1 million, where \$1 million is the minimum trade size for brokered interdealer trades and thus the natural minimum target for trade sizes. The average delay is 107 minutes shorter for the larger trades; the average proportionate reduction in delay is 32 percent. Delays are longer for the smaller transactions in 25 of the 27 currencies. Under the null hypothesis that trade size and delay are unrelated, the number of currencies with shorter delays for larger trades has a binomial distribution with $n = 27$ and $p = 0.5$. The figure of 25 out of 27 is significant at the level of $1.0E-5$. This suggests that the dominant motive for trading late in the day is netting.

For the subset of currencies facing exchange controls, delays are fairly insensitive to transaction size. The average proportionate delay is only 11 percent higher for small trades than large trades in such countries, versus a 37 percent difference for countries without exchange controls, a difference that is

highly statistically significant. This further confirms the importance of exchange controls in determining delays in emerging market countries. The nature of the controls is also important. Thai baht trades do not experience high delays even though the currency is subject to exchange controls. Those controls do not involve extra steps, however, as in Brazil, but instead simply involve restrictions on the types of bank accounts that can be used to fund asset trading.

IV. CLIENT FUND BEHAVIOR: WHICH TRANSACTIONS ARE NEGOTIATED?

In light of the high markups on non-negotiated trades, it is remarkable that negotiated transactions represent only three percent of the total. This section investigates why asset managers do not take greater advantage of the option to call and negotiate prices with custodial dealers. Practitioners have suggested what we call the “ignorance hypothesis”: client funds are not aware of the high markups on non-negotiated trades. Either the funds are unable to assess those costs due to the opacity of the trade situation or they are unaware that foreign exchange trading involves dealers who trade as principals. Whatever the reason, practitioners report that at some asset managers, “currency execution is viewed as an administrative function and delegated to staff with little market experience” (McGeehan, 2010). Equity investors, in particular, might be confused on this point since equity trades typically involve a broker paid by commission; fixed-income trades, like foreign exchange trades, typically involve a dealer compensated through the bid-ask spread. The ignorance hypothesis is thus consistent with the finding, from a Russell Investment Group study of 24 investment programs, that equity managers pay substantially more per foreign exchange trade than currency managers or fixed-income managers (reported in *Euromoney*, 2006).

Nonetheless, the ignorance hypothesis seems unlikely to apply broadly among the funds in our sample. As noted in the introduction, there has been industry-wide concern about the high markups on custodial foreign exchange trades for at least a decade. Substantial information about the situation has been available in the financial media at least since 2003. Firms offering to help reduce those costs, including Record Currency Management and Russell Investment Group, have publicized the situation at conferences, in practitioner magazines, and in practitioner-oriented research journals. Normatively speaking, funds *should* not be ignorant about execution costs according to the best practices for pension funds listed in the Myner’s Report (2001) of the U.K. That report asserts that trustees of pension funds “should have a full understanding of the transaction-related costs they incur.”

We hypothesize, instead, that client funds are generally aware of the higher costs of non-negotiated trades, but their incentives and constraints lead them to rationally choose this approach nonetheless. The strongest incentive is the administrative cost of negotiated trading and the fact that recovering these costs requires raising up-front fees. To enable active trading with third parties, such as electronic exchanges or regular OTC dealers, an institutional investor must create an in-house trading operation. This requires it to hire traders and, to control the risk that a trader incurs losses that destroy the bank, to assign someone besides the trader to handle confirmation and settlement. In addition, compliance officers must be assigned to ensure that everyone follows established processes, and all these staff members require substantial technology and associated technology support. For the average individual fund it would not be cost-effective to create an in-house trading operation of this nature, since the cost would exceed our earlier estimate of \$200,000 spent annually in transaction costs when trading foreign exchange with the custodian. However, many of these individual funds are presumably members of fund families and it might well be cost effective for them to pool their resources with other members. Further, the costs of negotiating trades with the custodial dealers – as opposed to a regular OTC dealer – are substantially lower, since settlement and some compliance can all be handled by the custodial staff.

An asset manager dissatisfied with its current custodian can certainly shift assets to a different custodian. But switching custodians is itself an extremely complicated process that takes over a year of concerted staff effort, and the benefits would be difficult to ascertain due to the same opacity that triggered the initial dissatisfaction.

Funds that incur these costs must either absorb them and report lower returns to shareholders or raise fees. But fees, as explicit costs, are relatively easy for underlying investors to monitor. The bid-ask costs associated with non-negotiated transactions, by contrast, are implicit – funds buy at higher prices and sell at lower prices, thus incurring opportunity costs – and thus difficult for investors to monitor. A unsatisfied investor would not have sufficient information to disentangle the extent to which his/her low returns reflect the fund's poor asset allocation decisions or the fund's poor management of its execution costs. Since the vast majority of asset managers customarily primarily rely on non-negotiated foreign exchange trades, any asset manager confident in its asset-picking capability would have felt confident of its ability to compete in the eyes of investors.

A. Which Transactions Are Negotiated?

We evaluate the ignorance hypothesis by examining the determinants of the choice between negotiated and non-negotiated transactions. We note that the costs of trading are unrelated to transaction size while the benefits from lower spreads rise with transaction size. This suggests that, if funds are aware of the cost differential and its implications for investor returns, they will be more likely to negotiate prices for larger transactions.

We use probit regressions in which the dependent variable is a dummy coded 1 if a transaction is negotiated and 0 otherwise. We return to the original transactions data, rather than data on netted trade amounts, because the funds' trade-type decisions are made on a transaction-by-transaction basis. The key independent variable is the (log) trade size. Since anticipated investor savings are also influenced by the expected difference in markup between negotiated and non-negotiated trades, we include the average of this difference for a given currency over 2006.

It is possible that direct calls are more likely from funds that are relatively sophisticated in the management of transaction costs. We use (log) fund NAV and (log) fund average trade size as measures of fund sophistication. As before, we include three currency-group dummy variables; we allow the residuals to cluster along the two dimensions of currency and trade date; and we use heteroskedasticity- and autocorrelation-consistent standard errors.

The results of this analysis do not support the ignorance hypothesis. Client funds are increasingly likely to call a dealer directly as transaction size increases and the effect is economically substantial (Table 4). With marginal effects evaluated at sample means, a transaction worth \$10 million is 4.8 percentage points more likely to be negotiated than a transaction worth \$10 thousand, more than doubling the 3.0 percent unconditional likelihood of such transactions. The coefficient on the average gap between markups on negotiated and non-negotiated transactions is positive, as expected, but not statistically significant. This could reflect the difficulty funds have in actually identifying their markups on non-negotiated transactions.

The point estimate suggests that funds with larger NAV are more likely to negotiate transactions, as anticipated, but the effect is insignificant. This suggests that the tendency for larger funds to rely relatively heavily on direct transactions, noted in the descriptive statistics, primarily reflects those funds' larger transaction sizes, rather than the size or sophistication of the funds per se.

Negotiated transactions are substantially less likely for emerging market currencies, other things equal. According to the marginal effects, for emerging-market currencies involving custodians the likelihood of a negotiated transaction is reduced by 1.3 percentage points – roughly one third of the unconditional likelihood of 3 percent – while for emerging-market currencies that do not involve custodians the effect is reduced even more, by 2.6 percentage points.

We evaluate the robustness of these results with three additional regressions. In the first, the (log) transaction size is replaced by a transaction's size relative to the average transaction size of the fund making the transaction. This adjustment allows for the possibility that funds define a large transaction relative to their normal trade size rather than in absolute terms. The regression results once again indicate that the size of the transaction in question is positively and significantly related to the likelihood that the asset manager decides to negotiate the price. In addition, the results again indicate that the funds most likely to negotiate prices are those that typically trade large amounts, and that such transactions are relatively unlikely for emerging-market currencies.

In the final robustness tests we run probit regressions separately on the liquid and emerging-market currencies. The results indicate that relative transaction size only influences the decision to negotiate for the more actively traded currencies. For emerging-market currencies the only significant determinant of whether a trade is negotiated seems to be a fund's average transaction size and the effect remains economically small.

B. Simultaneity?

Our analysis of the asset managers' choice between negotiated and non-negotiated transactions may raise questions about whether markups and transaction type are simultaneously determined. If causality runs in both directions on a transaction-by-transaction basis, the econometric analysis of markups presented in Section II would be mis-specified. Suggestive evidence that simultaneity is not an issue comes from juxtaposing two results introduced previously: (1) the insignificance of the coefficient on markups in the probit regressions concerning the decision to negotiate, and (2) the significance of transaction-type dummy in the regressions concerning markups. This suggests that the causality operates from transaction-type to markup and not vice versa. The inference is not reliable, however, because simultaneity could be biasing the coefficients. This section highlights two more robust reasons why simultaneity should not be an issue.

First, the client's decisions are made for each *transaction* while the custodial dealers' markup decisions are made for each (clustered) *trade*. Second, the decision whether to negotiate is based on the *expectation* of a markup difference, not any realized markup difference. Suppose the client fund decides to negotiate a transaction with its custodian and is quoted a relatively unattractive price. The fund might nevertheless have no incentive to switch to a non-negotiated transaction because, given our earlier analysis, the fund can expect to get an even less attractive price if it sends the instruction through the fund accountant. Suppose instead the client fund decides not to negotiate and sends the transaction instruction to its fund accountant. Since it will only learn the actual price paid a few days or weeks later, it has no incentive to change its mind.

There certainly may be some *long-run* influence of realized markups on the negotiation decision. A larger expected difference in markups between negotiated and non-negotiated transactions should eventually increase the tendency of asset managers to negotiate prices. A straightforward analysis of the direction of simultaneity bias, however, shows that it should be towards smaller estimated effect. This suggests that the very substantial effect of trade type on markup reported earlier is, if anything, an underestimate of the long-run connection.

V. CONCLUSIONS

This paper examines how asymmetric information influences the nature and price of foreign exchange trades between dealers at global custody banks and their client funds. Our data comprise the complete trading record of a mid-sized global custody bank during calendar year 2006. The magnitude of markups on such trades is the subject of a high-profile legal dispute and ours is the only rigorous academic analysis to date of this type of trading. We highlight two key layers of information asymmetries. First, the custodians typically know more about their prices and costs than their client funds. Second, the client funds know more about their execution costs than their ultimate investors.

Most asset managers take the cost of foreign-exchange trading as a fixed cost of doing business overseas and simply assign the custodian to take care of it. As a result, the asset managers typically learn the prices paid belatedly in periodic activity reports. This arrangement, though entered into willingly by both parties, gives the custodian's currency dealers a considerable information advantage relative to their clients. Our evidence suggests that custodians set wider markups when the ambiguity or "fog"

surrounding their prices is highest. It also suggests that custodians avoid setting extreme prices so as to preserve that ambiguity.

The client funds have lower-cost trading options, one of which is to call the custodial dealers directly and negotiate a price. In our sample, the markups on negotiated trades average 3.4 basis points, distinctly less than the 22.4 basis point average markup for non-negotiated trades. Nonetheless, client funds chose to negotiate prices for only three percent of transactions.

Our analysis shows that clients are significantly more likely to negotiate a price for larger transaction sizes, from which we infer that the client funds are aware that negotiated transactions are relatively inexpensive. We hypothesize that the funds may choose not to negotiate transactions because it is more administratively burdensome. If they were to negotiate all foreign-exchange transactions, they would eventually have to charge their investors higher fees, and fees are relatively easy for investors to monitor. By contrast, the higher markups associated with non-negotiated transactions directly lower the investors' returns, but these costs are implicit and are thus difficult to monitor.

Since the administrative cost of negotiated trades is fixed while the execution costs rise with trade size, client funds are aware of the cost difference between negotiated and non-negotiated trades the latter would tend to negotiate their larger trades, consistent with our evidence. This evidence cannot be considered conclusive, however, so further clarification of the motives for non-negotiated trades between client funds and their custodians could be a fruitful direction for future research.

REFERENCES

- Angel, James (1996), Who Gets Price Improvement on the NYSE? Georgetown University Working Paper FINC-1377-12-694.
- Bank for International Settlements (2007), Triennial central bank survey, foreign exchange and derivatives market activity in 2007, (B.I.S., Basel).
- 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 S. 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 (2005), Dealer Behavior and Trading Systems in Foreign Exchange Markets. *Journal of Financial Economics* 75: 571-605.
- Christie, William G., Jeffrey H. Harris, and Paul H. Schultz (1994a). Why Do NASDAQ Market Makers Avoid Odd-Eighth Quotes? *Journal of Finance* 49: 1813-1840.
- Christie, William G., Jeffrey H. Harris, and Paul H. Schultz (1994b). Why Did NASDAQ Market Makers Stop Avoiding Odd-Eighth Quotes? *Journal of Finance* 49: 1841-1860.
- Collie, Bob (2004). It's Time for More Choice in FX. Russell Investment Group Viewpoint, December 2004.
- Complaint, People of the State of California versus State Street Corporation, http://ag.ca.gov/cms_attachments/press/pdfs/n1823_october_20.pdf
- Cullinan, Charles P., and Dennis M. Blin (2005). An Analysis of Mutual Fund Custodial Fees. *Journal of Applied Business Research* 21 (1): 13-22.
- Easley, David, and Maureen O'Hara (1987), Price Trade Size and Information in Securities Markets. *Journal of Financial Economics* 19: 69-90.
- Euromoney (2006), The Growing Pains of Foreign Exchange (March).
- Foucault, Thierry (1999), Order Flow Composition and Trading Costs in a Dynamic Limit Order Market. *Journal of Financial Markets* 2: 99-134.
- Galanek, John (2010). The Search for Clarity: exploring New Frontiers in FX TCA. Currency Investor Magazine. www.currency-investor.net/transaction+cost+analysis.html.
- Glosten, Lawrence (1989), Insider Trading, Liquidity, and the Role of the Monopolist Specialist. *Journal of Business* 62: 211-235.
- Green, Richard C., Burton Hollifield, and Norman Schurhoff (2007), Financial Intermediation and the Costs of Trading in an Opaque Market. *Review of Financial Studies* 20: 275-314.
- Ho, T., and H. Stoll (1981). Optimal Dealer Pricing Under Transaction and Return Uncertainty. *Journal of Financial Economics* 9: 47-73.
- Institutional Investor Magazine (2007), Global Custody Rankings, <http://www.iimagazinerankings.com/globalcustody/GlobalCustodyRanking.asp>
- Jones, Charles M. and Marc L. Lipson (2003), Are Retail Orders Different? Working Paper, Columbia University.

- Joseph, Paul (2005), Custody Vs. 3rd Party Foreign Exchange Execution. Global Money Management. <http://www.globalmoneymanagement.com/Article.aspx?ArticleID=1073746>.
- Kothare, Shashank, and Lloyd Raynor (2010). Are Your FX Fees Too High? Ways to Close the 'Responsibility' Gap. Russell Investments Forum (February): 1-8.
- McGeehan, Jimmy (2010), What Can Usain Bolt Teach Us About FX Transaction Costs Analysis? www.fxtransparency.com/2010/04/what-can-usain-bolt-teach-us-about
- Myner's Report (2001), <http://www.icaew.com/en/library/subject-gateways/corporate-governance/codes-and-reports/myners-report>.
- Naik, Narayan Y., Anthony Neuberger, and S. Viswanathan (1999). Trade Disclosure Regulation in Markets with Negotiated Trades, *Review of Financial Studies* 12: 873-900.
- Osler, C.L., Alexander Mende, and Lukas Menkhoff (2009). "Price Discovery in Currency Markets," *Journal of International Money and Finance* forthcoming.
- Osler, C.L., and Rimma Sherman (2009). Intraday Dynamics of Foreign Exchange Interdealer Spreads. Mimeo, Brandeis University.
- Ramadorai, Tarun (2006), What Determines Transaction Costs in Foreign Exchange Markets? *International Journal of Finance and Economics* 13: 14-25.
- Roşu, Ionni (2009). A Dynamic Model of the Limit-Order Book. *Review of Financial Studies* Forthcoming.
- Russell Investment Group (2006). Pushing for An FX Revolution.
- Schmeidel, Heiko, Markku Malkamaki, and Juha Tarkka (2006). Economies of Scale and Technological Development in Securities Depository and Settlement Systems. *Journal of Banking and Finance* 30: 1783-1806.

Figure 1A: Average Trade Sizes

Sample includes the complete trade record for a mid-sized custody bank during calendar year 2006.

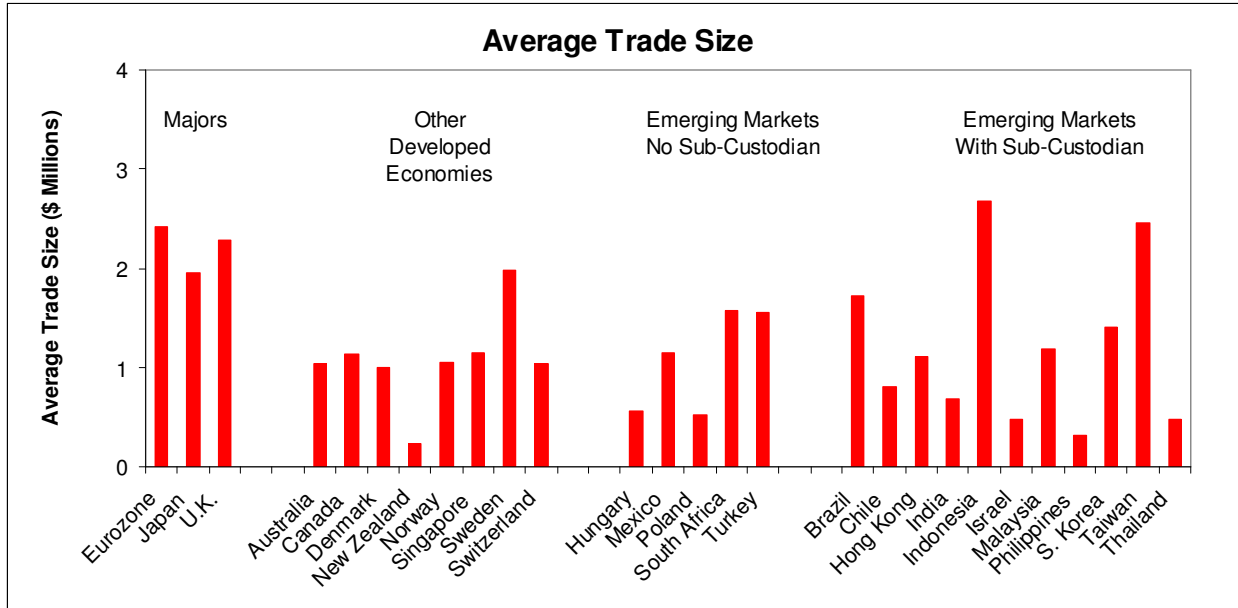


Figure 1B: Mean Custody-Bank Markups by Currency

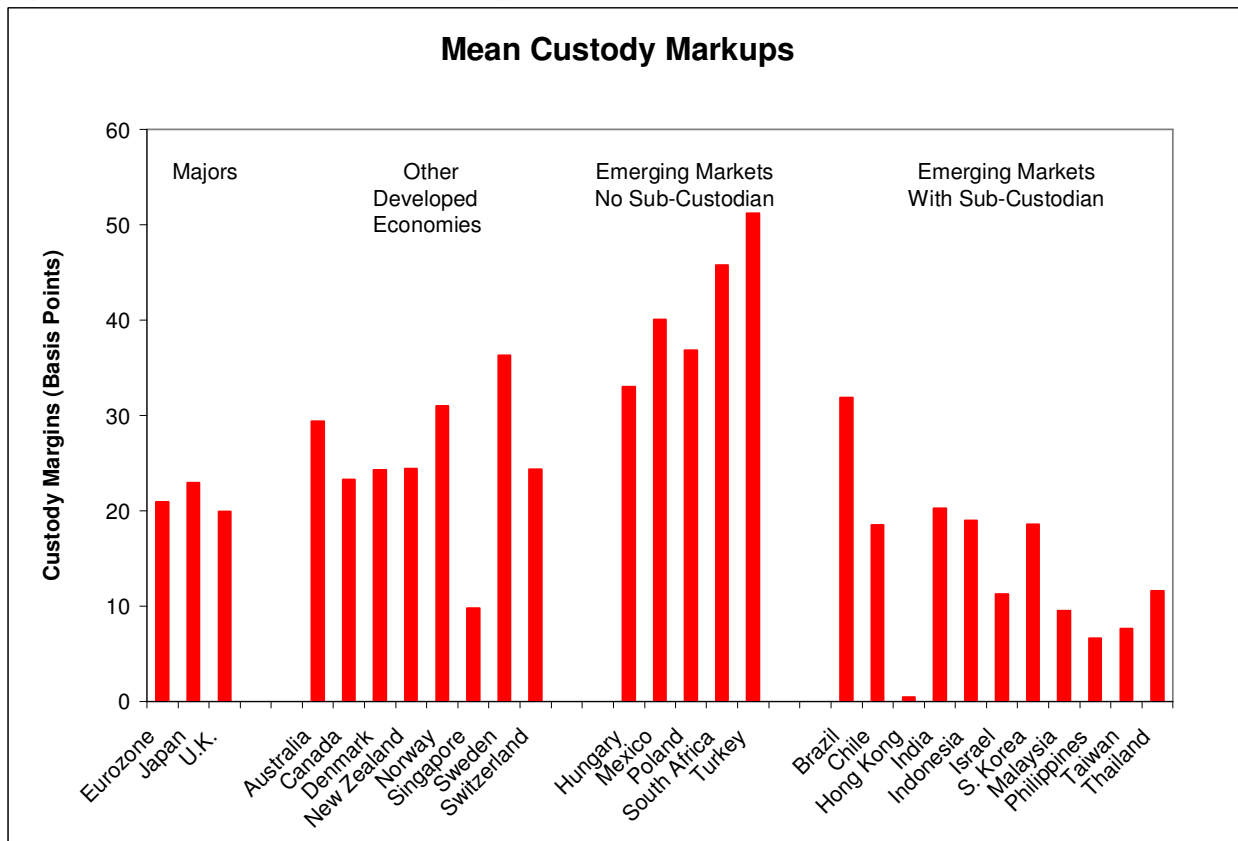



Figure 2A: Currency Volatility

Square root of sum of squared daily returns during calendar year 2006. Returns based on daily interbank mid-quotes from Global Insight.



2B: Mean Interbank Bid-Ask Spreads

Figure shows the mean (log) difference between interbank ask prices provided by www.Oanda.com and interbank mid-quotes provided by Global Insight. Daily data cover calendar year 2006. Units are basis points.





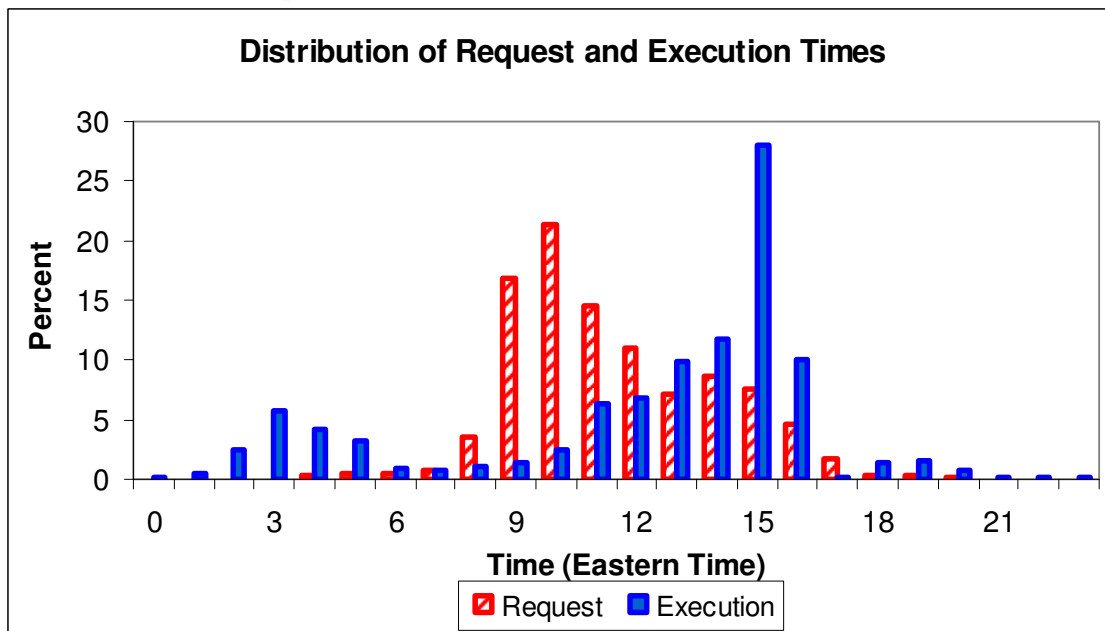
Figure 3. Trade Request and Execution Times

Transaction request refers to the time at which the fund accountant sends the order to the trading floor. Sample includes the complete record of non-negotiated trades for a mid-sized custody bank during calendar year 2006.

3A. Time from Transaction Request to Execution



3B. Distribution of Request and Execution Times



3C. Delays on Small and Large Transactions




Table 1. Descriptive Statistics

Sample comprises custody-bank markups taken from the complete transaction record for a mid-sized custody bank during 2006. Interbank half-spreads are calculated as the mean (log) difference between average intraday interbank ask prices provided by www.Oanda.com and average intraday interbank mid-quotes provided by Global Insight using weekday data for 2006.

	Mean	Median	Std. Dev.	Skew	Kurtosis	Min	Max
Dependent Variable							
Markup in basis points	20.4	10.8	28.3	3.7	38.8	-129.9	727.0
Independent Variables							
Cross-Sec Volatility	0.83	0.79	0.59	1.97	8.44	0.00	3.22
Time-Series Volatility	0.97	0.72	1.01	5.72	81.21	0.00	28.43
Interbank Half-Spread	20.6	6.5	42.8	3.1	12.2	0.5	251.5
Fund Trading Volume	\$0.1 bn	\$0.0 bn	\$0.9 bn	15.3	269.0	\$0.0 bn	\$16.4 bn
Fund NAV	\$0.7 bn	\$0.1 bn	\$2.1bn	7.7	73.2	\$0.0bn	\$25 bn
(Log) Fund NAV	4.1	4.9	3.1	-1.3	4.7	-6.7	10.1
Market Liquidity	11.4%	0.042	13.3%	1.1	2.7	0.1%	37.0%
(Log) Market Liquidity	-3.3	-3.2	1.8	-0.5	2.1	-6.9	-1.0
Trade Value	\$1.7 mn	\$0.3 mn	\$6.8 mn	26.5	1,303.9	\$0.0 mn	\$479 mn
(Log) Trade Value	12.3	12.6	2.4	-0.8	4.1	0.0	20.0
Transaction Value	\$0.8 mn	\$0.1 mn	\$9.7 mn	43.2	2191.8	\$0.0 mn	\$631 mn
(Log) Transaction Value	10.8	11.0	2.5	-0.4	3.6	0.0	20.3

Table 2. Baseline Regressions: Determinants of Custody-Bank Markups

Table shows the results from running the following regression:

$$Markup_t = \alpha + \beta X_t + \gamma Z_t + \eta_t$$

where X_t comprises variables relevant to the information asymmetry between custodians and their client funds and Z_t comprises the control variables. Column 2 reports marginal effects from a censored regression on trades with positive markups. Columns 1 and 3 report coefficients from panel estimation with standard errors that are robust to simultaneous clustering by time and currency. Entries represent basis points. Coefficients highlighted in bold are statistically significant. Shaded rows correspond to non-negotiated trades. Sample comprises custody-bank markups taken from the complete transaction record for a mid-sized custody bank during 2006. *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Independent Variable: Markup	Baseline	Baseline Censored	With Fund Dummies
Fog and Ambiguity Preservation			
Non-negotiated (+)	11.87***	12.67***	10.86***
Cross-Sec Volatility	1.86**^{cc}	2.41**^{cc}	0.02 ^c
Cross-Sec Vol * Non-neg. (+)	10.84***	10.28***	12.86***
Time-Series Volatility	1.23 ^c	1.21** ^{c*}	0.94
Time-Series Vol * Non-neg (+)	1.66	1.69	1.97
Interbank Half-Spread	0.09***[*]	0.08***[*]	0.10***[*]
IBk Half-Spread * Non-neg (-)	-0.05***	-0.05**	-0.07***
Subcustodian	3.24 ^{c*}	4.28** ^{c*}	0.89** ^{c*}
Sub-custodian * Non-neg (-)	-14.76***	-15.79***	-11.84***
Other Emerging Markets	-3.93** ^{c*}	-4.48** ^{c*}	-0.09** ^{c*}
Other Ems * Non-neg	2.26	2.82	-2.62
Other Developed Countries	-0.01** ^{c*c*}	0.33** ^{c*c*}	0.62** ^{c*c*}
Other DC * Non-neg	1.03	0.70 ⁻	-0.14
Traditional			
Fund Trading Volume (-)	-0.49***^e	-0.49***^e	-0.08 ^e
Fund NAV (-)	-0.14*^e	-0.14*^e	-4.33***^e
Market Liquidity (-)	-0.60** ^v	-0.60***^v	-0.42*** ^v
Tuesday	0.61** ^{c*}	0.63** ^{c*}	0.85** ^{c*}
Wednesday	-0.55** ^{c*}	-0.54** ^{c*}	-0.36** ^{c*}
Thursday	0.58** ^{c*}	0.59** ^{c*}	0.79** ^{c*}
Friday (-)	1.46** ^{c*}	1.51*^{c**}	1.54** ^{c**}
Trade Value	0.33**[*]	0.33**[*]	-0.08 [*]
Constant	-5.40 ^{cccc}	^c	24.54***^c
Adjusted (or pseudo) R^2	0.16 ^{cccc}	0.02 ^{cccc}	0.20 ^{cccc}

Table 3. Regressions with Constrained Samples

We run this regression: $Markup_t = \alpha + \beta X_t + \gamma Z_t + \eta_t$, where X_t comprises variables relevant to the information asymmetry between custodians and their client funds and Z_t comprises control variables. Standard errors are clustered by currency and time and are robust to simultaneous correlation along both dimensions. Entries represent basis points. Coefficients highlighted in bold are statistically significant. Shaded rows correspond to non-negotiated trades. Sample comprises custody-bank markups on trades in 27 currencies vis-à-vis the US dollar from the complete transaction record of a mid-sized custody bank during 2006. *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Indep Variable: Markup Individual Fund Dummies?	Funds That Call Directly		Most Liquid Currencies		Emerging-Market Currencies	
	No	Yes	No	Yes	No	Yes
Fog and Ambiguity Preservation						
Non-negotiated (+)	10.64***	10.86***	6.18	10.98*	24.20***	13.39***
Cross-Sec Volatility	1.75 * cc	-0.05 * cc	-1.29 * cc	-2.12 cc	3.57* * cc	1.84 * cc
Cross-Sec Vol * Non-neg. (+)	10.28***	12.33***	17.20***	15.42**	8.14***	10.18***
Time-Series Volatility	1.23 c	0.94 c	-0.50 c	-0.50 c	7.78* c	7.63*** c
Time-Series Vol * Non-neg. (+)	2.14	2.49	1.85***	1.70***	-3.83	-3.74
Interbank Half-Spread	0.08*** *	0.10*** *	0.01 *	-0.15 *	0.33 *	-0.15 *
IBk Half-Spread * Non-neg. (-)	-0.05***	-0.07***	-0.03	0.13	-0.03	-0.13
Subcustodian	2.46 c *	-0.19 c *	NA	NA	14.02*** *	0.60 c *
Sub-custodian * Non-neg. (-)	-13.02***	-11.47***	NA	NA	-25.43***	-11.00
Other Emerging Markets	-4.64 ** c *	-1.44 ** c *	NA	NA	NA	NA
Other Ems * Non-neg.	4.29*	-1.12	NA	NA	NA	NA
Other Developed Countries	-0.39 c c *	-0.12 * c c *	2.40** c *	2.56** c *	NA	NA
Other DC * Non-neg.	0.72	-1.15	1.99*	0.82	NA	NA
Traditional						
Fund Trading Volume (-)	-0.36* c	-0.80*** e	-0.68*** e	-1.32*** e	0.81* c	-0.48 c
Fund NAV (-)	-0.16 c	0.21 c	-0.12 c	3.10*** e	-0.39 c	0.86** e
Market Liquidity (-)	-0.80 *** v	-0.84 *** v	1.49 *** v	1.40 *** v	-2.83 *** v	-2.35 *** v
Tuesday	0.86 ** c *	1.15 ** c *	0.68 ** c *	0.93 ** c *	-0.42 ** c *	0.02 ** c *
Wednesday	-0.55 ** c *	-0.26 ** c *	-0.89 ** c *	-0.70 ** c *	-0.42 ** c *	-0.23 ** c *
Thursday	0.52 * c * *	0.76 * c * *	0.32 * c * *	0.51 * c * *	1.82 * c * *	2.38 * c * *
Friday (-)	1.24 * c * *	1.31 * c * *	2.94*** *	3.01*** *	-0.82 * c * *	-0.11 * c * *
Trade Value	0.32* *	-0.04 *	0.12 *	-0.46*** *	0.66* *	0.58* *
Constant	-5.38 cccc	-4.89 cccc	4.92 cccc	-12.03** cc	-24.49 cccc	-14.10 cccc
Adjusted R ²	0.19 cccc	0.22 cccc	0.10 cccc	0.19 cccc	0.20 cccc	0.20 cccc

Table 4. Which Transactions Are Negotiated?

Table shows the results from a probit regression examining the decision whether to negotiate a given transaction (1= negotiated). Residuals are permitted to cluster by currency and transaction day and are robust to simultaneous correlation along both dimensions. In column 3, the log trade size is replaced with the ratio of the trade size to the average trade size for that particular fund. Sample comprises custody-bank markups for 27 currencies vis-à-vis the US dollar taken from the complete transaction record of a mid-sized custody bank during 2006. All coefficients $E+2$ except the constant term; all marginal effects $E+3$. *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Independent Variable: Trade Type Negotiated = 1	Baseline Regression		Robustness (Coefficients)		
	Coefficient	Marginal Effect	Relative Trade Size	Liquid Currencies	Emerging-Mkt Currencies
(Log) Trade Size (+)	10.7***	6.9***			
Relative Trade Size (+)			1.7***	2.1***	4.0
(Log) Fund NAV (+)	0.9	0.6	0.7*	1.2**	-0.8
(Log) Fund Average Trade Size (+)	14.5***	9.4***	15.1***	16.3***	8.8***
Non-Neg. Markup – Neg. Markup (+)	0.4	0.3	0.3	3.2***	1.7
Currency Dummies					
Other Developed Countries	1.9	1.2	0.2	-11.3**	
Emerging Markets, No Sub-Custodian	-40.8***	-26.5***	-25.6***		-1.6
Emerging Markets, Sub-Custodian	-19.6*	-12.7*	-13.3**		
Constant	-5.0***				
Pseudo-R ²	0.11		0.09	0.10	0.06