What are the Origins of Foreign Exchange Movements?

Martin D. D. Evans*
Department of Economics
Georegetown University
Washington DC 20057
and the N.B.E.R
Email evansm1@gunet.georgetown.edu
Phone (202) 687-1570
Fax (202) 687 6102

(Second Draft: Comments Welcome)

October 15, 1999

Abstract

This paper uses a new transactions data set on the interbank foreign exchange market to examine the origins of spot exchange rate movements. The data provide a comprehensive picture of trading activity and allow me to examine the contribution of public news to spot rate dynamics over hours, days, and weeks. Contrary the presumption of macroeconomic exchange rates models, I find that public news only accounts for a fraction of exchange rate volatility over the whole frequency spectrum. In particular, I estimate that less that 50% of the variance of spot rate changes at very high frequencies is attributable to public news. At daily and weekly frequencies, changes in the spot rate *understate* the effects of public news by 20 to 40 percent because the cumulative effects of independent public and private news exert offsetting effects. These findings suggest one reason for the poor performance of macroeconomic exchange rate models; namely their exclusive focus on public news.

*I thank Richard Lyons for the many comments and suggestions he made on a earlier draft of this paper. Part of the research reported on here was conducted while I was visiting the Financial Markets Group at the L.S.E. and The Bank of England. I am grateful for the hospitality I received at both institutions and acknowledge the financial support of the Houblon-Norman Fund at The Bank of England. I would also like to thank the staff at The Bank and Reuters for allowing me to collect the data used in the study.

1 Introduction

The origins of nominal exchange rate dynamics remain elusive. In particular, there is no widely accepted explanation for the sizable short and medium-term movements in the dollar during the floating rate period. More generally, theoretical models relating exchange rates to macroeconomic fundamentals are still outperformed by simple time series models in forecasting spot rates over short and medium-term horizons [Frankel and Rose (1995)].

This paper proposes an explanation for the poor performance of macroeconomic exchange rate models based on the analysis of how the arrival of information affects spot rates. According to macro models, innovations in spot rates are attributable to the arrival of public news. More precisely, this approach assumes that (a) all information relevant for exchange rate determination is common knowledge; and (b) the mapping from information to equilibrium prices is also common knowledge. Thus, public news is characterized by the arrival of new information to all market participants simultaneously and their homogeneous interpretation of its implications for equilibrium prices.

In this paper, I use a new data set to examine the extent to which the arrival of new information in the FX market conforms to this definition of public news In the case of the DM/$, the most heavily traded currency, I find that public news only accounts for a fraction of exchange rate volatility over the whole frequency spectrum. In particular, I estimate that less that 50% of the variance of spot rate changes at very high frequencies is attributable to public news. At daily and weekly frequencies, changes in the spot rate *understate* the effects of public news by 20 to 40 percent. These findings suggest that the information structure of typical macro models, embodied in assumptions (a) and (b), is too restrictive. They also raise the possibility that macro models with less restrictive information structures may meet with more empirical success.

The data used in this study details trading activity in the spot FX market over a four month period, May 1 to August 1996.[1] These data are unique in that they provide information on trading between FX dealers around the world. In particular, they allow us to track the pattern of trade and spot rates in the direct interdealer market on a transaction-by-transaction basis. As such, I can identify the impact of public news very simply. Following Hasbrouck (1991), I identify public news as the arrival of information that affects transactions prices but not the pattern of trade. The rationale for this identification scheme is straightforward. Trade between FX dealers should only take place when there is a difference between their valuations of the asset (at the margin). In microstructure models of asset price determination, these different valuations arise from either asymmetric information [i.e., Kyle (1985) and Glosten and Milgrom (1985)] or heterogeneous

[1]Evans (1997) provides a detailed description of the data set focusing on the relationship between the transactions data and the indicative quote data that has been used in earlier research by Baillie and Bollerslev (1991), Goodhart and Giugale (1993), Guillaume et al (1994), and others.

inventories [i.e., Amihud and Mendelson (1980) and Ho and Stoll (1983)]. By definition, public news alters all dealer valuations *equally.* Differences between valuations are not affected so there is no impact on the pattern of trade.

Under this identification scheme there is no need to specify the news subject or its source. Thus, we need not confine our attention to the effects of announcements concerning particular macroeconomic fundamentals. This is clearly advantageous because there is little theoretical agreement over the appropriate set of fundamentals to consider. Similarly, there is no need to rely on a particular set of news announcements, say from a wire service. And, as a consequence, we avoid the recurrent question of whether wire service stories where anticipated in advance. Furthermore, the scheme makes no assumption about persistence in the impact of news on spot rates. As macroeconomic models show, public news may have temporary or permanent effects on spot rates depending upon the process for the fundamentals concerned.

In view of the strong presumption that public news accounts for most of the movements in exchange rates over days, weeks and months, some elaboration on possible alternative sources of information is in order. Specifically, what types of non-public information could affect exchange rates? Ito, Lyons and Melvin (1998) provide an answer in the form of a taxonomy; fundamental private information, and semi-fundamental private information. The former refers to information on the asset price that would obtain in an economy with symmetric information, while the latter refers to information about the future path of prices before they reach this value.

As an example of fundamental information, Ito et al. cite the FX orders dealers receive from outside the market derived from real international trade. Because these orders are received in advance of trade statistics [and are only known to dealers with which they are placed] they represent a private signal on the shift in demand for foreign currency due to trade flows [Lyons (1997)]. Foreign exchange intervention by central banks represents another example. Because few central banks deal directly in the interbank market, most interventions are made by placing orders with FX dealers at one or more commercial banks. Such an order would represent fundamental private information to the dealer [Pieres (1997)]. Notice that both these examples contradict the common knowledge assumption (a).

Disagreement amongst dealers about the meaning of a public announcement provides an example of semi-fundamental private information in the FX market. Given the lack of consensus by researchers around a theoretical exchange rate model, it would be very surprising if all dealers used the same model to discern the price-implications of a particular news item. In fact, there often appears to be disagreement over the direction the spot rate should take in response to a news item. Under these circumstances, as Ito et al. note, any dealer with superior information about the dispersion of beliefs across the market, would be in a position to more accurately forecast the path of prices as disagreements resolves through some learning process. This form of semi-fundamental information contradicts the homogeneous mapping assumption (b). Superior information about the distribution of dealer inventories is also a potential source of semi-fundamental private information. Insofar as inventory risk is priced, information about the distribution of inventories could allow a dealer to forecast the path of prices more accurately as inventory positions are unwound. Thus, assumption (a) is once again contradicted.

The research reported here has links to both the microeconomic and macroeconomic literatures on exchange rates. Microeconomic evidence on private information comes from comparing trading patterns under different regimes. In particular, Ito et al. (1988) and Covrig and Melvin (1998), found the differences in the high frequency behavior of Yen quotes before and after the introduction of lunch-time FX trading in Tokyo to be consistent with predictions of theoretical models in which traders have private information. Complementing these results, Payne (1999) finds evidence of asymmetric information in a week of FX trading on an electronic brokerage system.[2] One difference between these studies and the analysis here is that I examine the role of private information as a source of exchange rate movements over horizons of days and weeks. This leads to an important finding; the effects of private news do not vanish as we move to the lower frequencies considered by macroeconomic exchange rate models.

Engel and Frankel (1984) and Ito and Roley (1987) document significant effects of monetary announcements in the 1980s for the DM and Yen respectively, while Hardouvelis (1988) and Beck (1993) found evidence indicating that trade and fiscal announcements affected the dollar. In all these macro studies, the fraction of the exchange rate variance attributable to the announcements is very small. One interpretation of these findings suggested by my results is that macro announcements rarely constitute public news, at least not in the strict information theoretic sense of macro exchange rate models. While the announcement is heard simultaneously across the market, the homogeneous mapping assumption does not apply because market participants have differing views about its implications for equilibrium prices.[3]

The impact of public news is quantified using a VAR. Hasbrouck (1991) pioneered the use of VARs to identify the information content of trades in equity markets. Here I modify his approach to account for the decentralized nature of direct interdealer FX trading and the limits on trader's information sets. The modifications lead to a structural VAR that identifies the role of multiple trade-related shocks. These shocks represent the effects of private news hitting the FX market through different channels and have very different implications for the behavior of exchange rates.

The VAR provides impulse response functions and variance decompositions of exchange rate

[2]Other microeconomic studies pointing to the presence of private information include, Lyons (1995), and Yao (1997a) Bjonnes and Rime (1998), and Cheung and Wong (1998).

[3]Fleming and Remolona (1999) report evidence from the U.S. Treasury market consistent with this view. They find that trading volume surges for a prolonged period shortly after macro announcements and argue that this is due to "residual disagreement among investors about what precisely the announcement means for prices".

movements over different horizons. The estimates show that public news contributes less than 50% of the variance in high frequency exchange rate movements. The VAR estimates also allow us to decompose the spot rate into components that identify the cumulative effects of public and private news. Comparing these components against the actual behavior of the spot rate provides a simple yet dramatic measure of how public and private news contributed to exchange rate movements on a macroeconomic scale. The cumulative effects of public and private news have offsetting effects on the spot rate. As a consequence, daily and weekly changes in the spot rate are positively correlated with but significantly understate the changes implied by the arrival of public news.

The remainder of the paper is organized as follows. Section 2 describes the data set and presents some of its key statistical features. Section 3 presents a trading model that forms the basis for identifying public and private news in the data. This section also discusses the links between the trading model and macroeconomic exchange models. Section 4 describes the econometric model and presents the high frequency empirical results. The macroeconomic implications of the results are examined in Section 5. Section 6 concludes.

2 The Data

2.1 Sources and Definitions

The analysis below utilizes new micro data on activity in the DM/$ spot FX market over a four-month period, May 1 to August 31, 1996. The data set contains time-stamped tic-by-tic data on actual transactions taking place through the Reuters Dealing 2000-1 system via an electronic feed that was customized for the purpose by Reuters. This is the most widely used electronic dealing system. According to Reuters, over 90% of the world's direct interdealer transactions take place through the system. These transactions account for about 80% of total trading in major spot markets; the remaining 20% is between dealers and non-bank customers. Interdealer trading breaks into two transaction types, direct and brokered. Direct trading between dealers accounts for about 60% of interdealer trade and brokered trading accounts for about 40%.

Trades on the D2000-1 system take the form of electronic bilateral conversations. The conversation is initiated when a dealer calls another dealer on the system to request a quote. Users of the system are expected to provide a fast two-way quote with a tight spread, which is in turn dealt or declined quickly (i.e., within seconds). To settle disputes, Reuters keeps a temporary record of all the conversations on the system. This record is the source of the transactions data. Every time an electronic conversation on D2000-1 results in a trade, the Reuters feed provides a time-stamped record of the transactions price, a bought or sold indicator, and a measure of cumulative trading volume.

Two features of the data are particularly noteworthy. First, they provide transaction information for the whole interbank market over the full 24-hour trading day. This contrasts with earlier transaction data sets covering single dealers over some fraction of the trading day [Lyons (1995), Yao (1997a, 1997b), and Bjonnes and Rime (1998)]. The data set makes it possible, for the first time, to analyze trading patterns and prices at the level of "the market." The only other multiple-dealer data set in the literature covers brokered interdealer transactions [the electronic system examined by Goodhart, Ito and Payne (1996) and Payne (1999)]. The system they examine, however, accounts for only a small fraction of daily trading volume.[4]

Second, these market-wide transactions data are not observed by individual FX dealers on the system as they trade. Though dealers have access to their own transaction records, they do not have access to others' transactions on the system. The transactions data therefore represents a history of market activity that market participants could only infer indirectly. This feature has important implication for interpreting the results reported below.

Although the transactions data are the prime focus of the analysis below, I shall also utilize indicative quotes. These quotes, known as the DFX series, are the data reported on the Reuters information terminals. They constitute the most current source of information on the FX market that is available to the public. For each currency, the DFX series includes every time-stamped indicative bid and ask price posted by banks to Reuters.[5]

The analysis below concentrates on three variables; (i) the DM transaction price paid by dealers purchasing dollars, (ii) the indicative asking DM price of dollars quoted by banks on the Reuters information terminal and, (iii) aggregate order flow. This last variable provides a measure of the pattern of direct interdealer trading and plays a central role in the analysis below. Specifically, order flow measures the net of buyer-initiated orders and seller-initiated orders. Clearly, this measure is a variant of a key concept in economics – "net demand." The basic difference is that order flow keeps track of who initiates transactions, and net demand does not. Importantly, while net demand must be zero in equilibrium, order flow can differ from zero.

To further cement ideas, it is useful to consider an example of how market wide order flow is determined relative to a set of transactions. Suppose dealer A initiates a conversation with dealer B by asking for quotes on the DM for a \$10m trade. If after hearing the bid and ask quotes, A decides to buy, he pays the DM price of dollars specified in the asking price, and order flow is +1.

[4]There is also evidence that dealers attach more informational importance to direct interdealer order flow than to brokered interdealer order flow, [see Bjonnes and Rime (1998)].

[5]The DFX series differs from the FXFX quote series that has been used in other studies (see Guillaume et al., 1994 for a survey). Reuters transmits the FXFX and DFX data to its terminals through different networks. While the DFX data uses a high speed digital network that can keep up with the high frequency of quotes received from banks, the FXFX data uses an older system with a lower throughput. To keep the FXFX series "up-to-date", Reuters sends only a sample of the quotes it receives through this network. As the frequency of bank quoting rises, Reuters sends a smaller fraction of the quotes into the FXFX network. Thus, the FXFX series are a sample of the quotes being sent to Reuters whereas the DFX series records all the quotes. Evans (1997) provides a detailed analysis of the relationship between trading and the posting of indicative quotes.

At the same time, dealer C may be initiating a trade against dealer D's quotes by selling \$5m. The order flow from this transaction is -1. If there were no other interdealer transactions currently taking place, aggregate order flow would be 0. Notice that this measure takes no account of the differing transaction sizes. In principle such a measure could be even more informative, but there is insufficient information in the data to determine the dollar value of individual transactions.

2.2 Summary Statistics

Although the D2000-1 system runs 24 hours a day, the vast majority interdealer transactions in the DM/\$ are concentrated during the European trading. This institutional feature gives rise to recurrent intradaily patterns in the data. Exemplifying this phenomena, Engle, Ito and Lin (1990), Baillie and Bollerslev (1991), Bollerslev and Domowitz (1993), Goodhart and Giugale (1993), Payne (1997) and Andersen and Bollerslev (1998) have all studied the intradaily patterns in the volatility of indicative FX quotes.

Figure 1 shows intradaily patterns in the volatility of transactions prices and indicative quotes together with the patterns for trading and quote intensity. Volatility in each series is calculated from the standard deviation of the second-by-second observations over a five minute period. The intensity of trade and quotes is measure by the number of trades and quotes per five minute period. In each panel, the solid line plots the average value of the variable for the particular period over the 79 trading days in the sample. The lower and upper dashed lines show the 10% and 90% bounds of the variable's distribution across the sample period.

Panel I plots the volatility in ask quotes against British Summer Time (BST). On average, volatility rises over the trading day, peeking around 1600 hrs. BST. This pattern is similar to that found in earlier studies and is commonly attributed to market openings and closures around the world [Andersen and Bollerslev (1998), Dacorogna et al (1993), and Payne(1997)]. In support for this view, Panel II shows that quote activity is heavily concentrated between 800 hrs. and 1700 hrs. BST when the European markets are open.

Panels III and IV plot the volatility of asking prices and transaction intensity. Here we see some evidence of a "U-shaped" pattern in volatility between 800 hrs. and 1700 hrs. BST. Although the change in average volatility is small compared to the error bands, this seasonal pattern is more pronounced than in the quote data. It is also worth noting that the average volatility of transactions prices is higher that the volatility of quotes, particularly around 800 hrs. BST. The "U-shaped" volatility pattern corresponds to the seasonal in trading intensity. Panel IV shows that, on average, trading intensity rises after 800 hrs. BST, peeks around 1000 hrs., and then declines towards mid-day before peeking again in the afternoon before 1600 hrs. BST.

Overall, Figure 1 confirms that the majority of trade and quote activity in the DM/\$ is concentrated within a time period associated with European Trading. In my analysis below, I shall therefore focus on the dynamics of transactions prices, quotes and order flow during "the trading day" defined as the period between 700 and 1900 hrs. BST. Of course, these times are somewhat arbitrary. However, changing the opening and closing times by an hour in either direction has no material effect on the results presented below.

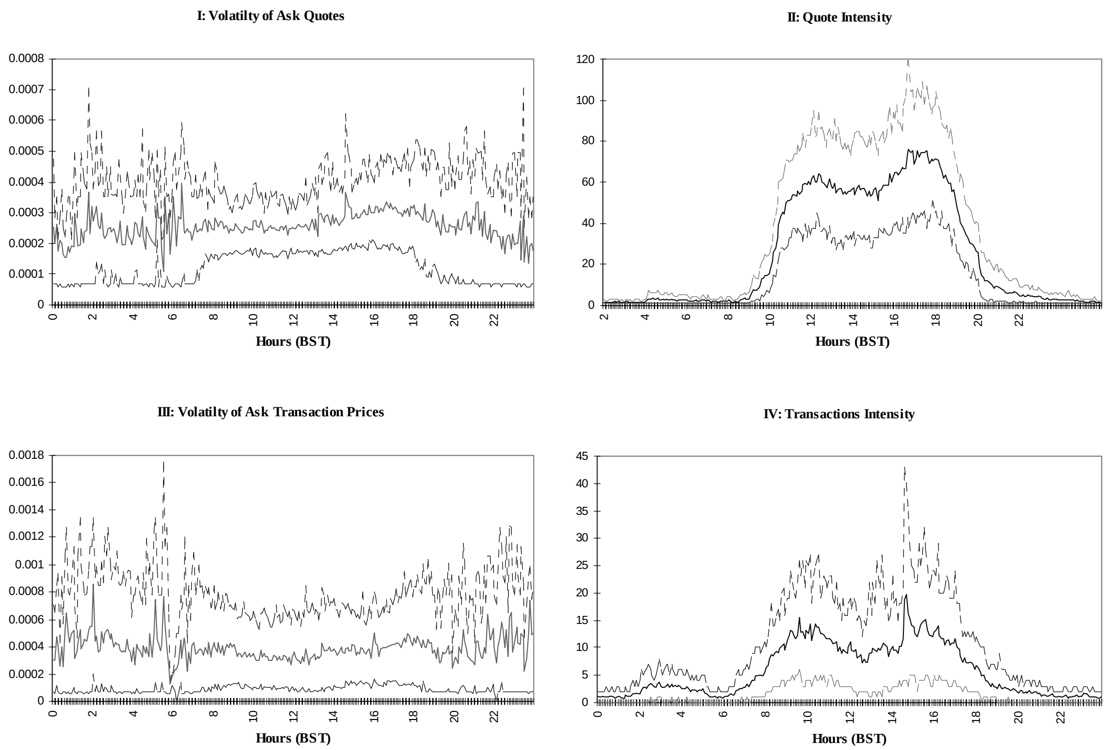

Figure 1

Table 1 below reports summary statistics on; the change in log ask transaction prices over a five minute period, $\Delta \ln p$, aggregate order flow, x, measured as the difference between the number of buyer-initiated and seller-initiated transactions over the period, and the difference between the logs of last indicative ask quote and the last ask transaction price each period, $\ln q - \ln p$. For comparison, the table reports statistics calculation over the 24 hour and trading day samples.

Several of the statistics are noteworthy. First, the unconditional distributions for all three variables appear highly non-normal. The distribution for transaction prices changes, $\Delta \ln p$, is skewed to the left and is highly leptokuritic. The distributions for order flow, x, and the quote/price spread, $\ln q - \ln p$, are also fat-tailed but are skewed to the right. Second, the average of returns measured over the trading day is -1.121% per day, but only -0.087% measured over the whole day - a difference implying positive overnight returns. Third, the statistics imply that the average value

Table 1: Summary Statistics						
	24 Hour Sample 13664 Observations			700-1900 Sample 10403 Observations		
	$\Delta \ln p$	x	$\ln q - \ln p$	$\Delta \ln p$	x	$\ln q - \ln p$
Mean	-0.087	0.147	0.020	-0.121	0.020	0.020
Std. Deviation.	14.344	6.108	0.035	14.029	6.784	0.033
Skewness	-0.254	0.038	0.095	-0.341	0.037	0.102
Kurtosis	7.838	11.052	8.029	8,963	9.177	8.788
ρ_1	-0.262	0.233	0.034	-0.268	0.229	0.026
ρ_3	-0.005	0.091	-0.001	-0.009	0.093	-0.003
ρ_6	0.011	0.060	-0.002	0.006	0.063	-0.003
Stability Tests						
P-Values						
ρ_1	0.010	0.058	0.208	0.051	0.224	0.125
ρ_3	0.071	0.229	0.169	0.095	0.223	0.422
ρ_6	0.665	0.241	0.803	0.224	0.163	0.840
$\rho_1 - \rho_6$	<0.001	0.025	0.860	0.212	0.132	0.630

Notes: $\Delta \ln p$ is the change in the log ask transactions prices over 5 minutes multiplies express in daily percent (i.e., multiplied by 28800). x is orderflow measured as the difference between the number of buyer-initiated and seller-initiated transactions over five minutes. $\ln q - \ln p$ is the difference between the last log ask quote and last log ask price at the end of the five minute period multiplied by 100. The asymptotic standard errors for the skewness and kurtosis statistics are 0.021 and 0.042 in the full sample, and 0.024 and 0.048 in the sub-sample. ρ_i denotes the autocorrelation at lag i. The lower portion of the table reports p-values for the null hypothesis that there is no variation in autocorrelation ρ_i over periods starting at 8:00, 10:00, 12:00 14:00 and 16:00 hrs. The last line reports p-values for the null hypothesis of joint stability in autocorrelations at lags 1 through 6.

for order flow outside normal trading hours is 0.552. Dealers initiating transactions outside normal trading hours appear on average to be far more likely to buy than sell dollars. By comparison, the average of 0.02 for x during the day indicates much more balanced between buying and selling. Finally, the positive average values for $\ln q - \ln p$ are consistent with the notion that indicative ask quotes are on average "less competitive" than the actual ask transactions prices available contemporaneously in the market.

As noted above, several studies have found evidence of recurrent intradaily patterns in the volatility of indicative FX quotes and these seasonal volatility patterns appear in Figure 1. In the analysis below, I will use a VAR to examine the dynamics of order flow and price levels rather than price volatility. So it is important to gauge the extent to which seasonal patterns appear present in the autocorrelation structure of, $\Delta \ln p$, x, and $\ln q - \ln p$. To this end, the lower portion of Table 1 reports the results from stability tests performed on the sample autocorrelations for each variable. In particular, I estimate the autocorrelation over two hour periods beginning at 800, 1000, 1200, 1400, and 1600 hrs. and test the null hypothesis that all these estimates are equal to the autocorrelation estimated over all hours in the sample.[6] As the table shows, evidence of seasonal instability appears most strongly in the case of the first-order autocorrelations for $\Delta \ln p$ and x in the 24 hour sample. In the trading day sample, none of the statistics are significant at the 5 percent level. Based on this evidence, there do not appear to be pronounced seasonal patterns in the autocorrelation structure of our variables between 700 and 1900 hrs.

3 Information and FX Trading

This section considers the impact of public and private information on the behavior of spot rates and order flow. I first present a simply trading model, based on Glosten and Milgrom (1985), that provides the formal basis for identification the effects of innovations in public and private information on transaction prices and aggregate order flow. Next, I consider how the arrival of information affects the relationship between indicative quotes and transaction prices. I then discuss the link between the trading model and traditional macroeconomic exchange rate models.

3.1 A Trading Model

The market is populated by a large number of dealers that trade dollars and marks with one another. We are interested in the problem of how transactions prices and order flow are determined as a result of bilateral conversations between dealers across the market. As noted earlier, a conversation is initiated when one dealer calls another dealer requesting a quote for a trade of a certain size. By convention, the dealer receiving the call must respond immediately with a *firm* bid and ask quote (with a tight spread). These quotes are of the "take-it-or-leave-it" variety and represent the DM prices at which the dealer is willing to buy or sell dollars to the caller. On hearing the quotes, the calling dealer must immediately decide to either initiate a trade (by "hitting" one of the quotes), or pass. The model focuses on how *firm* quotes are set, and on the decision to initiate trade.

Suppose there are $2N$ dealers in the market that request and set quotes according to an exogenous matching mechanism that results in N conversations each period. Consider the problem of how dealer i sets quotes to buy or sell a known amount when receiving a request from dealer j. Let

[6] Formally, to test for stability in ρ_i for variable z_t, I estimate the regression

$$z_t = \rho + \rho_i z_{t-i} + \sum_{j=1}^{k} D_t(j)\beta_j z_{t-i} + u_t$$

where $D_t(j)$ is a dummy variable equal to one when t falls within the j'th. 2 hour period, and zero otherwise. I then perform a Wald test for the joint significance of the β_j coefficients allowing for heteroskedasticity. To test for the joint stability in $\rho_1, \rho_2, ...\rho_6$, I estimate a system of regressions like the one above by GLS allowing for contemporaneous correlations between the residuals. I then test for the significance of the estimated $\beta_j' s$ across the whole system.

$p^b_{i,t}$ and $p^a_{i,t}$ denote the bid and ask quotes at which dealer i is willing to buy and sell a fixed number of dollars. The trading decision of dealer j depends on these quotes. In particular, I assume that

$$\text{dealer } j \quad \begin{cases} \text{buys if} & r_{j,t} > p^a_{i,t}, \\ \text{sells if} & r_{j,t} < p^b_{i,t} \end{cases}, \tag{1}$$

where $r_{j,t}$ is dealer j's reservation price. The resulting order flow is given by

$$x_{i,t} = \begin{cases} 1 & \text{if } r_{j,t} > p^a_{i,t} \\ -1 & \text{if } r_{j,t} < p^b_{i,t} \end{cases},$$

and the transaction price is $p_{i,t} = p^a_{i,t}\mathcal{I}(x_{i,t} = 1) + p^b_{i,t}\mathcal{I}(x_{i,t} = -1)$, where $\mathcal{I}(.)$ is an indicator function [equal to one when the condition in $\mathcal{I}(.)$ holds, and zero otherwise]. The data set contains the cross section of transactions prices and order flows that arise from dealer conversations across the market, i.e., the set of variables $\{p_{i,t}, x_{i,t} | x_{i,t} \neq 0\}_{i=1}^N$. Aggregate order flow is defined as the sum of bilateral order flows across the market, $x_t \equiv \sum_{i=1}^N x_{i,t}$.

Following Glosten and Milgrom (1985), dealer i sets quotes to maximize the expected profit from each quote request:

$$E\left\{ \left[p^a_{i,t} - p^*_t\right] \mathcal{I}(r_{j,t} > p^a_{i,t}) + \left[p^*_t - p^b_{i,t}\right] \mathcal{I}(r_{j,t} < p^b_{i,t}) \middle| \Omega_{i,t} \right\}, \tag{2}$$

where $\Omega_{i,t}$ represents dealer i's information set. This function measures trading profits relative to a benchmark value, p^*_t. In standard trading models, this benchmark is identified as the terminal value of the asset. A more natural interpretation in the FX context is that p^*_t defines the DM price of dollars that would obtain in an economy with symmetric information amongst dealers. We shall see below that this definition creates a natural link between the trading model and traditional general equilibrium macro models of the exchange rate. Here $p^a_{i,t} - p^*_t$ represents dealer i's profit from selling dollars at the ask quote while $p^*_t - p^b_{i,t}$ represents the profit from buying them at the bid in a world of symmetric information.[7]

Maximizing expected trading profits gives the following equations for optimal quotes:

$$\begin{aligned} p^a_{i,t} &= E\left[p^*_t | \Omega_{i,t}\right] + (1 - G(p^a_{i,t}))/g(p^a_{i,t}), \\ p^b_{i,t} &= E\left[p^*_t | \Omega_{i,t}\right] - G(p^b_{i,t})/g(p^b_{i,t}), \end{aligned} \tag{3}$$

[7] Note that dealers choose quotes to maximize these expected profits without regard for the level of their inventories. This is consistent with the behavior of an individual FX dealer reported by Yao (1997a).

where $G(.)$ and $g(.)$ are the distribution and density functions for dealer j's reservation price, $r_{j,t}$, conditioned on dealer i's information [see Appendix A for details].

Optimal quotes are equal to the dealer's prior expectation, $E\left[p_t^*|\Omega_{i,t}\right]$, plus or minus an adjustment factor that depends on the subjective distribution of reservation prices. These factors are equal to the inverse of the hazard function associated with the distribution. In the first equation, the hazard identifies the probability that $p_{i,t}^a$ is equal to the highest asking price at which the caller is willing to buy. The hazard rate in the second equation identifies the probability that $p_{i,t}^b$ is the lowest bid price at which the caller is willing to sell. The intuition behind these adjustment factors is clear: Both factors depend on how dealer i perceives the distribution of reservation prices. The dealer sets a higher asking price when (i) the right hand tail of this distribution rises, and/or (ii) when the shrinkage in the distribution's tail from given price increase is smaller. Similarly, the bid price will be lowered when (i) the left hand tail of the perceived distribution increases and/or (ii) rate of shrinkage in the tail falls.

Equation (3) provides an implicit expression for optimal quotes given a particular form for the conditional distribution of reservation prices, $G(.)$. In a rational expectations equilibrium, this distribution must be consistent with the actual distribution of reservation prices in the market. Formally, this implies that $G(y) = \Pr(r_{j,t} < y|\Omega_{i,t})$ where $\Pr(.|\Omega_{i,t})$ represents the conditional probability based on the actual distribution of the reservation prices. In general equilibrium, these prices are determined endogenously as part of the solution to the dealer's optimization problem. Optimal reservation prices should reflect customer orders from outside the market, and inventory imbalances that result from unintended trade with other dealers.

In the analysis below, I assume that $G(.)$ conforms with rational expectations, but I leave the actual form of distribution unspecified. That is to say, I treat quotes satisfying (3) as rational expectations equilibrium quotes. The cost of this simplification is that we cannot relate changes in the actual or conditional distribution of reservation prices to specific economic (structural) shocks. Importantly, it does not affect the distinction between the effects of public and private information on transactions prices and order flow, which is the focus of the analysis.

3.2 Price and Order Flow Innovations

Three factors can change the pattern of transactions prices and order flow in the model. Changes in $E\left[p_t^*|\Omega_{i,t}\right]$ and the conditional distribution of reservation prices, $G(.)$, alter quotes and hence the prices at which any transactions take place. In addition, order flow also depends on the actual distribution of reservation prices. Although changes in the conditional and actual distribution of reservation prices can take many forms, for expositional clarity I shall consider changes in the conditional mean, $E\left[r_{i,t}|\Omega_{i,t}\right]$, and the true mean μ_t.

Let ∇y_t denote the innovation in a variable y that takes place at the beginning of period t.

Writing the explicit solution for quotes implied by (3) as $p_{i,t}^n = \mathcal{F}_n(E[p_t^*|\Omega_{i,t}], E[r_{j,t}|\Omega_{i,t}])$, for $n = \{a,b\}$, and linearizing $\mathcal{F}_n(.,.)$ using the implicit function theorem, we obtain:

$$\nabla p_{i,t}^n \simeq \beta_i^n \nabla E[p_t^*|\Omega_{i,t}] + (1-\beta_i^n)\nabla E[r_{j,t}|\Omega_{i,t}] \qquad n = \{a,b\}, \tag{4}$$

where $\beta_i^n > 0$.[8] Equation (4) shows that innovations in optimal quotes can approximately be written as a weighted average of the innovations in $E[p_t^*|\Omega_{i,t}]$ and $E[r_{j,t}|\Omega_{i,t}]$. Notice that changes in the actual distribution of reservation prices can only affect quotes insofar as they change dealers expectations, $E[r_{j,t}|\Omega_{i,t}]$. Changes in $E[r_{j,t}|\Omega_{i,t}]$ have no effect on quotes when $\beta_i^n = 1$, a case that only arises when the conditional distribution of reservations prices has a constant hazard rate [see (3) above].

Next, we need to consider innovations in order flow. Let $H(.)$ denote the true distribution of reservation prices relative to their market-wide average, μ_t. We can now write the discrete distribution for order flow between dealers i and j as

$$x_{i,t} = \begin{cases} 1 & \text{with probability } 1 - H(p_{i,t}^a - \mu_t) \\ 0 & \text{with probability } H(p_{i,t}^a - \mu_t) - H(p_{i,t}^b - \mu_t) \\ -1 & \text{with probability } H(p_{i,t}^b - \mu_t) \end{cases} . \tag{5}$$

Using least squares theory, the innovations in order flow can be written as

$$\nabla x_{i,t} = \phi_i^n \nabla(p_{i,t}^n - \mu_t) + \xi_{i,t}^n \qquad n = \{a,b\}, \tag{6}$$

where $\xi_{i,t}^n$ is uncorrelated with the innovations in $p_{i,t}^n - \mu_t$. Appendix B shows that the regression coefficients are negative and are well-approximated by

$$\phi_i^a \simeq \frac{-h(p_t^a - \mu_t)E[\nabla(p_{i,t}^a - \mu_t)|\nabla(p_{i,t}^a - \mu_t) > 0]}{[1 - H(p_t^a - \mu_t)]\, Var\left(\nabla(p_{i,t}^a - \mu_t)\right)}$$

and

$$\phi_i^b \simeq \frac{h(p_t^b - \mu_t)E[\nabla(p_{i,t}^b - \mu_t)|\nabla(p_{i,t}^b - \mu_t) < 0]}{H(p_t^b - \mu_t) Var\left(\nabla(p_{i,t}^b - \mu_t)\right)}.$$

[8]Specifically,

$$\beta_i^n = \frac{g(p_{i,\tau}^n - E[r_{j,\tau}|\Omega_{i,\tau}])}{2g(p_{i,\tau}^n - E[r_{j,\tau}|\Omega_{i,\tau}]) + (p_{i,\tau}^n - E[p_\tau^*|\Omega_{i,\tau}])g'(p_{i,\tau}^n - E[r_{j,\tau}|\Omega_{i,\tau}])}$$

where τ denote the start of period t, (i.e., before the innovation takes place). Appendix A shows that the second order conditions from maximizing expected profits imply $\beta_i^n > 0$.

Finally, combining (4), and (6) with the definition of aggregate order flow, we obtain

$$\nabla x_t = \sum_{i=1}^{N} \phi_i^n \left(\nabla E\left[p_t^* | \Omega_{i,t}\right] - \nabla \mu_t\right) + \sum_{i=1}^{N} \phi_i^n (1-\beta_i^n) \left(\nabla E\left[r_{j,t} | \Omega_{i,t}\right] - \nabla E\left[p_t^* | \Omega_{i,t}\right]\right) + \xi_t^n \qquad (7)$$

where $\xi_t^n \equiv \sum_{i=1}^{N} \xi_{i,t}^n$, $n = \{a, b\}$. Equation (7) decomposes innovations in aggregate order flow into a linear combination of the innovations to $E\left[p_t^* | \Omega_{i,t}\right]$, $E[r_{j,t} | \Omega_{i,t}]$ and μ_t, that can be used in conjunction with (4) to examine the effects of public and private news.

3.3 Public and Private News

Public news is defined as the arrival of new price-relevant information to all market participants simultaneously and their homogeneous interpretation of its implications for equilibrium prices. Such news is identified in the model by a homogeneous revision in dealer reservation prices and their valuations of the dollar across the market. This implies that $\nabla E\left[p_t^* | \Omega_{i,t}\right] = \nabla \mu_t$ for all i. Rational dealers also recognize this fact so that $\nabla E\left[r_{j,t} | \Omega_{i,t}\right] = \nabla E\left[p_t^* | \Omega_{i,t}\right]$ for all i. Public news is not only heard all dealers, but it is also perceived by each one as being received and interpreted homogeneously across the whole market.

Under these circumstances, the innovations to aggregate order flow and quotes from (4) and (7) are

$$\nabla x_t = 0, \qquad \nabla p_{i,t}^n = \nabla E\left[p_t^* | \Omega_{i,t}\right] \qquad \vee\, i \text{ and } n = \{a, b\}. \qquad (8)$$

Thus, aggregate order flow remains unchanged while quotes move one-for-one with the common revaluation of the dollar across the market. The first result follows from the fact that there is no change in the distribution for $x_{i,t}$ for any dealer pairs in the market. This means that the market-wide distribution of transactions prices (for both buyer-initiated purchases and sales) shifts in line with the distribution of quotes. Thus, in total, the arrival of public news induces a shift in the distribution of transactions prices with no change in aggregate order flow,

Private information can arrive in several forms. Type I private news to dealer i is represented by an innovation in $E\left[p_t^* | \Omega_{i,t}\right]$ that is not matched by the innovations in *both* $E\left[r_{j,t} | \Omega_{i,t}\right]$ and μ_t. The degree to which $E\left[r_{j,t} | \Omega_{i,t}\right]$ varies depends on the extent to which dealer i believes that other dealers are privy to the news. At one extreme, dealer i may believe that he is the sole recipient of the news so that $\nabla E\left[r_{j,t} | \Omega_{i,t}\right] = 0$. At the other, i may believe the news to be public so that $\nabla E\left[p_t^* | \Omega_{i,t}\right] = \nabla E\left[r_{j,t} | \Omega_{i,t}\right]$. Changes in μ_t depend on the extent to which the news is actually received across dealers in the market. Once again, if i is the only dealer privy to the news, $\nabla \mu_t = 0$, but if all other dealers receive the news $\nabla E\left[p_t^* | \Omega_{i,t}\right] = \nabla \mu_t$.

To examine the effects of Type I news, denote θ as the set of dealers that receive the information and requests for quotes. Also, let $\delta_i^r \equiv \nabla E\left[r_{j,t} | \Omega_{i,t}\right] / \nabla E\left[p_t^* | \Omega_{i,t}\right]$ and $\delta_i^\mu \equiv \nabla \mu_t / \nabla E\left[p_t^* | \Omega_{i,t}\right]$.

Substituting these expressions into (4) and (7) gives

$$\begin{aligned}\nabla x_t &= \sum_{i\in\theta}\phi_i^n\left[(\delta_i^r-\delta_i^\mu+\beta_i^n(1-\delta_i^r))\right]\nabla E\left[p_t^*|\Omega_{i,t}\right]+\xi_t^n \\ \nabla p_{i,t}^n &= \left[\delta_i^r+\beta_i^n(1-\delta_i^r)\right]\nabla E\left[p_t^*|\Omega_{i,t}\right] \qquad i\in\theta\end{aligned} \tag{9}$$

Equation (9) shows that Type I private news leads to an innovation in both aggregate order flow and quotes under most circumstances. Exceptions occur when either the innovations in $x_{i,t}$ cancel out across the market, or when $\delta_i^r-\delta_i^\mu+\beta_i^n(1-\delta_i^r)=0$ for all i. Notice that this latter condition is always met for public news because in the this case $\delta_i^r=\delta_i^\mu=1$ for all i. We cannot rule out this possibility here without further restrictions.

Private information may also arrive in a form that leads dealer i to change $E\left[r_{j,t}|\Omega_{i,t}\right]$ but not $E\left[p_t^*|\Omega_{i,t}\right]$. I shall refer to this as Type II private news. Changes in $E\left[r_{j,t}|\Omega_{i,t}\right]$ may or may not be accompanied by changes in the actual distribution of reservation prices. For example, $E\left[r_{j,t}|\Omega_{i,t}\right]$ may change with out any variation in μ_t if dealer i learns more about the distribution of reservation prices in the market. Alternatively, there may be changes in the inventory positions of dealers across the market that alter the distribution of reservation prices. These changes could be due to interdealer trade or trade between dealers and customers outside the market. In either case, since they cannot be directly observed by any single dealer, it is unlikely that the innovations in μ_t and $E\left[r_{j,t}|\Omega_{i,t}\right]$ will match for any i.

Type II private news is characterized by innovations in μ_t and/or $E\left[r_{j,t}|\Omega_{i,t}\right]$ with $\nabla E\left[p_t^*|\Omega_{i,t}\right]=0$. The innovations to aggregate order flow and quotes under these circumstances are

$$\begin{aligned}\nabla x_t &= \sum_{i\in\theta}^{N}\phi_i^n(1-\beta_i^n)\nabla E\left[r_{j,t}|\Omega_{i,t}\right]-\sum_{i=1}^{N}\phi_i^n\nabla\mu_t+\xi_t^n \\ \nabla p_{i,t}^n &= (1-\beta_i^n)\nabla E\left[r_{j,t}|\Omega_{i,t}\right] \qquad i\in\theta\end{aligned} \tag{10}$$

where θ is the set of dealers that receive Type II news and requests for quotes. Here the news affects both aggregate order flow and quotes except in the special case where $\nabla\mu_t=(1-\beta_i^n)\nabla E\left[r_{j,t}|\Omega_{i,t}\right]$.

In summary, the arrival of public news always leads to a change in quotes and transactions prices without any affect on the pattern of interdealer trade, and hence aggregate order flow. By contrast, private news of either type affects the pattern of trade between dealers and aggregate order flow under all but the most exceptional of circumstances. This is the basis for the identification scheme employed in the empirical model below.

3.4 Indicative Quotes and Transaction Prices

To this point I have focused on how the arrival of information affects trading between FX dealers. Information will also affect the indicative quotes posted by dealers to news services like Reuters. These quotes represent the public's most current information on interdealer transactions and they provide a signal about the prices at which dealers will trade with the public (i.e., dealer-customer trades).

Let $q^b_{i,t}$ and $q^a_{i,t}$ denote the indicative bid and ask quotes posted by dealer i, indicating the price at which the dealer will try to buy and sell a fixed number of dollars from the public. Because there may be unexpected changes in market conditions between the time a customer hears the quotes and places an order, and the time that the dealer can fulfill it, market convention allows for the actual purchase or sale price to differ from the quotes. Thus, $q^b_{i,t}$ and $q^a_{i,t}$ are "indicative" of the prices at which dealer-customer trades actually take place.[9]

To understand how the arrival of information affects indicative quotes, we need to first consider what motivates dealers to post them. Broadly speaking, the aim must be to generate customer orders. These orders are important to dealers for two reasons. First, they represent a significant source of a dealer's profits because the spreads on dealer-customer trades are much larger than in the interdealer market.[10] This means, for example, that dealer i could fill a customer purchase order for \$5m by buying from dealer j at $p^a_{j,t}$, and selling the dollars on at $p^a_{j,t} + v \simeq q^a_{i,t}$ (i.e., close to the indicative quote) with markup v. Second, orders from customers provide an important channel through which dealers can manage their inventories. This is particularly true at the end of the trading day when dealers try (or are forced) to eliminate their inventory imbalances. With very few traders willing or able to carry overnight positions, customers have to be induced to absorb much of the aggregate dealer imbalance.[11] The posting of indicative quotes represents the principle means by which dealers can affect the flow of customer orders. Consequently, dealers have a strong incentive to maintain a reputation for posting indicative quotes that are viewed as unbiased estimates of the actual prices at which they will trade with the public.

How should the arrival of public news affect these quotes? Consider the posting decision of dealer i when public news arrives raising interbank transactions prices by 100 pips (i.e., 0.01DM/\$) If he raises his indicative quotes by less than 100 pips, and he has a reputation for unbaisedness,

[9] As Lyons (1996b) points out, indicative quotes represent real-time information targeted at the public rather than dealers. The reason is that dealers have access to live quotes (either voice or electronic) from brokered trading between dealers which are more informative about the firm quotes being made in direct interdealer trading.

[10] In his single dealer study, Yao (1997b) estimates that 75% of the dealer's profit came from dealer-customer transactions even though these trades account for only 14% of his total trading volume.

[11] Although trading can take place on a 24 hour basis, in reality trading activity sharply falls at the end of European trading around 1900 hrs BST [see Figure 1 above]. So while individual dealer imbalances can be passed on to a large number of other dealers earlier in the day, by this stage a large portion of any aggregate imbalance held by dealers must be absorbed by the public.

there will be a large increase in customer orders to purchase dollars because everyone recognizes that dealer i has become a relatively cheaper source. Since the (opportunity) cost of purchasing dollars in the interbank market has risen by 100 pips, dealer i now faces the prospect of either filling the orders at the usual markup and reneging on the indicative ask quote, or lowering the markup so as to keep his reputation. Alternatively, now suppose dealer i posts new quotes that are more than 100 pips higher. In this case, there will be a large increase in customer orders to sell dollars. Once again the dealer must choose between filling the orders at a lower markup, or reneging on the indicative bid quote. The dealer faces none of these choices when he posts 100 pip higher quotes. In this case there should be no change in the flow of customer orders because the competitive position of dealer i relative to others remains the same in the eyes of the public. Moreover, since the public recognizes that transactions prices in the interbank market have risen by 100 pips, the dealer can fill orders at his usual markup without the loss of reputation.

This example illustrates why indicative quotes should move one-for-one with interbank transactions prices when public news arrives. A dealer following an alternative quoting policy could only maintain his reputation for unbiased quoting at the cost of lowering (or eliminating) his markup, υ, and hence his profit from dealer-customer trading.

The arrival of private news need not have the same affect on quotes. Returning to our example, suppose dealer i raises his firm quotes $\{p^a_{i,t}, p^b_{i,t}\}$ by 100 pips on the basis of either Type I or Type II private news. Here there is little danger of loosing his reputation for unbiased quoting if he chooses not to increase his indicative quotes by the same amount. Unlike the case of public news, there is no consensus about the changing value of the dollar so a failure to fill an order at the indicative quote can be credibly attributed to an unexpected change in market conditions.

In summary, the arrival of public news should lead indicative quotes to move one-for one with interbank transactions prices otherwise dealers risk loosing their reputation for unbiased quoting. This reputation is important because indicative quotes are the principle means by which dealers affect customer order flow, and dealer-customer trades are an important source of dealer profits. The reputational incentives are far weaker when private news arrives. Under these circumstances, indicative quotes may or may not move one-for-one with transaction prices in the interbank market.

3.5 Exchange Rate Fundamentals

For perspective on the results that follow, it is useful to relate the trading model to more familiar macroeconomic models. In these models exchange rates are determined by fundamentals such as interest rates, money supplies, output, and other macroeconomic variables. At the foundation of general equilibrium macro models lies a no-arbitrage condition

$$1 = E\left[m_{t+1} z_{t+1} | \Omega_t\right], \tag{11}$$

where z_{t+1} is the nominal return on any traded asset, and m_{t+1} is the nominal pricing kernel. In models were there is a representative agent, m_{t+1} is equal to the nominal intertemporal marginal rate of substitution, and Ω_t is the agent's information set. In models where there are agents with differing information sets, say $\Omega_{i,t}$, the pricing kernel is a Ω_t-measurable non-negative random variable and $\Omega_t = \bigcup_{i,=1}^{N} \Omega_{i,t}$ [see Duffie (1992)].

The implications of (11) for the behavior of the exchange rate follow from the fact that the condition applies to both foreign and domestic risk-free bonds with returns, r_t^f and r_t^d. If p_t^* denotes the domestics price of foreign currency, and m_{t+1} is the pricing kernel measured in terms of domestic currency, (11) implies that

$$E\left[m_{t+1}|\Omega_t\right] r_t^d = E\left[m_{t+1}p_{t+1}^*/p_t^*|\Omega_t\right] r_t^f.$$

The equilibrium exchange rate is found by solving this stochastic difference equation for p_t^*. Iterating forward, with the long-run level of p_t^* normalized to unity, we get

$$p_t^* = E\left[\prod_{j=0}^{\infty} r_{t+j}^f \psi_{t+j}/r_{t+j}^d \middle| \Omega_t\right] \tag{12}$$

where $\psi_t \equiv 1 + Cov(m_{t+1}, p_{t+1}^*|\Omega_t)/E[m_{t+1}|\Omega_t]E[p_{t+1}^*|\Omega_t])$.[12]

In representative agent models, ψ_t^{-1} identifies the foreign exchange risk premium [i.e. the expected excess return on investing in foreign bonds relative to the domestic risk-free rate.] In this case, (12) shows the equilibrium spot rate to be equal to the expected product of future interest differentials adjusted for the risk premium. Here changes in p_t^* *must* originate from news about current and future interest rates and/or movements in the risk premium; news about money supplies, output, or other macroeconomic variables can only affect spot rates insofar as they affect the forecasts of r_{t+j}^f, r_{t+j}^r or ψ_{t+j}. Moreover, (12) implicitly rules out any role for private information in the determination of spot rates. By construction, all the relevant news in these models is macroeconomic and public. This leads to the strong prediction that changes in equilibrium spot rate takes place without trade.

We can also use (12) to consider the behavior of equilibrium exchange rates in macro models with heterogenous agents. In this case, Ω_t is the pooled set of information available individuals at

[12]The complete forward solution is

$$p_t^* = E\left[\lim_{k\to\infty} p_{t+k}^* \prod_{j=0}^{k-1} r_{t+j}^f \psi_{t+j}/r_{t+j}^d \middle| \Omega_t\right].$$

In a fully specified model, the limiting behavior of p_t^* is typically tied down by a no-arbitrage condition in goods markets (i.e., PPP) and other long run conditions (i.e., money neutrality). For simplicity, I assume in (12) that the long run equilibrium value of the exchange rate is non-stochastic and normalized to one. Without this assumption, the current spot rate would be affected by news concerning the long run equilibrium exchange rate in addition to the factors cited below.

t, and p_t^* identifies the equilibrium spot exchange rate in the economy in the special case where agents have the same information sets. When agents have different information, $\Omega_{i,t} \subset \Omega_t$, so by iterated expectations,

$$E\left[p_t^* | \Omega_{i,t}\right] = E\left[\prod_{j=0}^{\infty} r_{t+j}^f \psi_{t+j} / r_{t+j}^d \middle| \Omega_{i,t}\right]. \tag{13}$$

Equation (13) shows how individual's estimates of p_t^* are related to fundamentals in the context of heterogenous agent macroeconomic exchange rate model. In particular, it shows that traditional macroeconomic variables like money supplies and output will affect $E\left[p_t^* | \Omega_{i,t}\right]$ insofar as they impact upon individual forecasts of the interest rates and the risk premium.[13] Within this context, news concerning macro variables will be public if it leads all individuals to the same revision in the forecasts for $r_{t+j}^f \psi_{t+j} / r_{t+j}^d$ for all $j \geq 0$. By contrast, forecast revisions that differ across individuals constitute private news. Such news may concern advanced information about macro variables or monetary policy. If could also concern information about the many other state variables that affect the equilibrium behavior of interest rates or the pricing kernel in a heterogeneous agent model.

Combining (13) with (8) in the case of public news, gives the following equation for the innovation in transaction prices

$$\nabla p_{i,t}^n = \nabla E\left[\prod_{j=0}^{\infty} r_{t+j}^f \psi_{t+j} / r_{t+j}^r \middle| \Omega_{i,t}\right] \qquad n = \{a, b\}, \tag{14}$$

where $\nabla E\left[. | \Omega_{i,t}\right]$ is the same across all individuals, i. Here transactions prices immediately reflect the full change in individuals' expectations concerning fundamentals and so exhibit semi-strong-from efficiency.

Transactions prices respond to public news in the trading model in the same way as equilibrium macroeconomic exchange models predict. While macro models with heterogeneous agents admit the possibility that private news could also affect exchange rates, this possibility has not been emphasized in the literature. Rather the focus has been on the role of public news. With this focus, equation (14) represents the canonical macro view of how transactions prices in the FX market are determined. By combining transactions prices and order flow, the empirical model presented below allows us to assess the validity of this view without the need to identify individual information, $\Omega_{i,t}$, or the equilibrium risk premium, ψ_t.

[13] Strickly speaking, there will not be a single foreign exhange risk premium in a heterogeneous agent model because risk tolerence will differ across agents. Nevertheless, the pricing kernel and ψ_t are well-defined.

4 Results

4.1 The Econometric Model

The effects of public and private news are estimated with a structural Vector Autoregression, VAR. Let $Y_t = [\Delta \ln p_t, x_t, \ln q_t - \ln p_t]$ be a vector containing the difference between the logs of the last transaction prices in periods t and $t-1$, aggregate order flow during period t, and the difference between the logs of the last indicative quote and transaction price during period t. Based on the statistic in Table 1, I assume that Y_t follows a stationary time series process with the Wold representation,

$$Y_t = \delta + \Psi(L)\varepsilon_t, \tag{15}$$

where ε_t is a 3×1 vector of one-step-ahead forecast errors in Y_t given information on lagged values of Y_t. $\Psi(L) = \sum_{i=0} \Psi_i L^i$ is a matrix polynomial in the lag operator with Ψ_0 normalized to the identity matrix. The forecast errors are serially uncorrelated with mean zero and covariance matrix Σ_ε. Equation (15) is a reduced form representation for the dynamics of FX trading that can be estimated by a VAR from the data.

We are interested in the structural model consistent with the trading model that leads to (15). This model has the form

$$Y_t = \delta + \Gamma(L)\eta_t, \tag{16}$$

where $\eta_t \equiv [u_t, v_{1,t}, v_{2,t}]'$ is a vector of serially uncorrelated structural disturbances with mean zero and diagonal covariance matrix Σ_η. Assuming that these disturbances lie in the space spanned by current and lagged values of Y_t, (15) is a reduced form for (16) with $\varepsilon_t = \Gamma_0 \eta_t$ and $\Psi(L) = \Gamma(L)\Gamma_0^{-1}$. The structural model is identified by placing restrictions on the impact matrix Γ_0 assuming that η_t comprises public news, u_t, and two private news shocks, $v_{1,t}$ and $v_{2,t}$:

$$\Gamma_0 = \begin{bmatrix} \gamma_{11} & \gamma_{12} & 0 \\ 0 & \gamma_{22} & \gamma_{23} \\ 0 & \gamma_{32} & \gamma_{33} \end{bmatrix}. \tag{17}$$

The first column of Γ_0 shows that impact of the public news. As in the trading model, this shock affects transactions prices but not aggregate order flow. Public news also induces indicative quotes to move in line with transactions prices so there is no change in the log spread, $\ln q - \ln p$.[14] The remaining columns show the impact of the two private news shocks, $v_{1,t}$, and $v_{2,t}$. The center

[14] Strictly speaking, the analysis above implies that public news has no affect on the spread, $q-p$, because indicative quotes move one-for-one with interbank transactions prices. However, since $q^a_{i,t} - p^a_{i,t}$ and $q^b_{i,t} - p^b_{i,t}$ are on the order of 3 to 4 pips, changes in transactions prices that are matched one-for-one by indicative quotes have no measurable affect on the log spread, $\ln q^n_{i,t} - \ln p^n_{i,t}$, $n = \{a, b\}$.

column shows that $v_{1,t}$ affects transactions prices and aggregate order flow as in trading model. $v_{2,t}$ represents the effect of a shift the distribution of reservation prices that goes unrecognized by dealers making quotes. This is a form of private news that affects aggregate order flow but not transactions prices. Since there is no strong theoretical presumption about how private news affects the relation between indicative quotes and transactions prices, I do not restrict how $v_{1,t}$, and $v_{2,t}$ affect innovations in $\ln q_t - \ln p_t$.

There are two important differences between the structure of this empirical model and the VARs employed by Hasbrouck (1991) and Payne (1999). First, the model uses transactions prices rather than the firm quotes of individual dealers. Payne, for example, uses a VAR to estimate how quotes on an electronic brokerage system are affected by innovations in order flow. Since there is no information on the sequence of quotes made by each dealer in my data, this is not a feasible method for estimating the impact of private news here. Second, the empirical model above allows some forms of private news to contemporaneously affect both transactions prices and order flow. This would not be possible in a centralized market because firm quotes are set before trade takes place. Here we are considering prices and order flow across a decentralized market where private news is distributed across some dealers initiating conversations and some making quotes [see (9) and (10) above].

The structural model is estimated from a VAR for Y_t as follows. Estimates of the impact matrix, Γ_0, are found as the solution to $\hat{\Sigma}_\varepsilon = \hat{\Gamma}_0 \Sigma_\eta \hat{\Gamma}_0'$ where Σ_η is normalized to the identity matrix and $\hat{\Sigma}_\varepsilon$ is the estimated covariance matrix of the VAR innovations. The remaining elements of $\Gamma(L)$ are derived as $\hat{\Gamma}(L) = \hat{\Psi}(L)\hat{\Gamma}_0$ where $\hat{\Psi}(L)$ is the moving average polynomial in L implied by the VAR estimates.

With these estimates in hand we can examining how u_t shocks contribution the variability of exchange rate movements. Strictly speaking, such estimates will provide an upper bound on the contribution of public news because it is possible for some forms of private news to be treated as an u_t shock under the identification scheme in (17). From this perspective, the results presented below will be biased towards the traditional macroeconomic view that public news is the dominant source of exchange rate movements.

4.2 Model Estimates

Estimates of the structural model in (16) and (17) are derived from estimating a VAR for Y_t using data sampled every five minutes during the trading day. In view of the seasonal patterns in Figure 1, the trading day is defined by the period between 700 and 1900 hrs. BST. The VARs include a set of dummy variables to allow for the effect on non-continuous observations, such as between trading days. The results reported below are derived from estimates with 6 lags.

Table 2 reports summary statistics from the VAR. Panel I shows the marginal significance

Table 2: Summary VAR Results			
Granger Causality Tests			
	Equations		
Variables	$\Delta \ln p_t$	x_t	$\ln q_t - \ln p_t$
$\Delta \ln p_t$	<0.001	<0.001	0.474
x_t	<0.001	<0.001	0.210
$\ln q_t - \ln p_t$	<0.001	<0.001	0.242
II: Diagnostics			
Q_3	0.972	0.999	0.999
Q_6	0.418	0.999	0.999
Seasonals	0.282	0.111	0.530
Eigenvalues	0.713	0.579	0.579

Notes: Aggregate order flow x_t is the difference between the number of dollar purchases and sales, $\Delta \ln p_t$ is the change in the log purchase (ask) price, and $\ln q_t - \ln p_t$ is the difference between the logs of the last ask quote and the last purchase price, each calculated over a five minute interval.

levels from the Granger Causality tests. We can reject the null of no Granger Causality from any of the variables in the equations for transactions price changes, $\Delta \ln p_t$, and aggregate order flow, x_t, at the one percent level. By contrast, there is little evidence that any variable predicts the difference between quote and transactions prices, $\ln q_t - \ln p_t$. Panel II reports significance levels for Box-Pierce Q-statistics for 3'rd. and 6'th. order serial correlation in the VAR residuals, under Q_3 and Q_6 respectively. The line marked Seasonals shows significant levels from LM tests for seasonal variation in the residuals over periods starting at 800, 1000, 1200 1400 and 1600 hrs. None of these tests are significant at the 10 percent level. The table also reports the modulus of the three largest eigenvalues of the VAR companion matrix. All these statistic suggest that $\Delta \ln p_t, x_t$ and $\ln q_t - \ln p_t$ jointly follow a stationary process which has dynamics that are well-represented by the estimated VAR.

The results in table 2 may be surprising to readers familiar with the difficulties of forecasting exchange rate movements at lower frequencies. However, the predictability implied here is consistent with the unpredictable changes in spot rates over weeks, months and longer. It is also worth noting that the predictability of $\Delta \ln p_t$ doesn't *necessarily* imply market inefficiency. Even though dealers have access to their own transaction records, they do not have access to information on the trades between other dealers. The lagged values of $\Delta \ln p_t$ and x_t represents a history of market activity that dealers can only infer indirectly. If the market is transparent so these inferences are precise, the predictability of $\Delta \ln p_t$ implied by the aggregate data is representative of how well individual dealers can predict price changes. In this case, the evidence in Table 2 points to the presence of unexploited profit opportunities. Alternatively, if low transparency in the market makes inferences about the lagged values of $\Delta \ln p_t$ and x_t very imprecise, price changes may appear completely unpredictable to dealers. In this case, the evidence in the Table 2 points to low transparency in the market rather than inefficiency.

Table 3: VAR Variance Decompostions

		Equation					
		$\Delta \ln p_t$		x_t		$\ln q_t - \ln p_t$	
horizon	source shock	%	(Std. Err.)	%	(Std. Err.)	%	(Std. Err.)
5 mins	public u	43.427	(9.518)	0.000	-	0.000	-
	private v_1	56.523	(9.518)	0.757	(0.106)	99.924	(0.405)
	private v_2	0.000	-	99.243	(0.106)	0.076	(0.405)
1 hour	public u	43.406	(9.500)	56.196	(10.814)	42.553	(8.391)
	private v_1	56.569	(9.499)	9.386	(16.917)	0.620	(31.753)
	private v_2	0.025	(0.060)	34.417	(15.236)	56.827	(31.822)
5 hours	public u	43.406	(9.500)	61.542	(11.994)	61.543	(11.991)
	private v_1	56.569	(9.499)	11.058	(17.202)	11.059	(17.358)
	private v_2	0.025	(0.060)	27.399	(14.773)	27.398	(14.936)

Notes: The table reports the percentage of the variance in the variable heading each column attributable to each of the three shocks. Standard errors are based on Monte Carlo simulations of the VAR with 1000 replications. $\Delta \ln p_t$ is the change in the log purchase (ask) price, x_t is aggregate order flow, and $\ln q_t - \ln p_t$ is the difference between the logs of the last ask quote and the last purchase price, each calculated over a five minute interval.

Table 3 reports the variance decompositions for the three variables at the 5 minute, 1 hour, and 5 hour horizons. The standard error associated with each statistic are calculated from 1000 Monte Carlo simulations based on the VAR estimates.

The most striking result in this table concerns the relative contribution of public and private shocks to the variance of transactions price changes. In the left hand column of the table we see that public shocks account for only 43% of the variance in price changes at all three horizons. This finding contrasts strongly with the traditional macro view that public news is the predominant source of exchange rate movements. Private news, in the form of the v_1 shock contributes most to the variance $\Delta \ln p_t$ at all horizons. As the center column shows, this shock accounts for a small but statistically significant fraction of the variance of order flow at the 5 minute horizon. The remainder of the variance in x_t is due to the v_2 shock. At longer horizons, public shocks become the dominant source of variance in aggregate order flow. In the right hand column we see that v_1 shocks account for almost 100% of the variance in $\ln q_t - \ln p_t$ within the 5 minute interval, but public shocks account for the majority of the variance as the horizon increases.

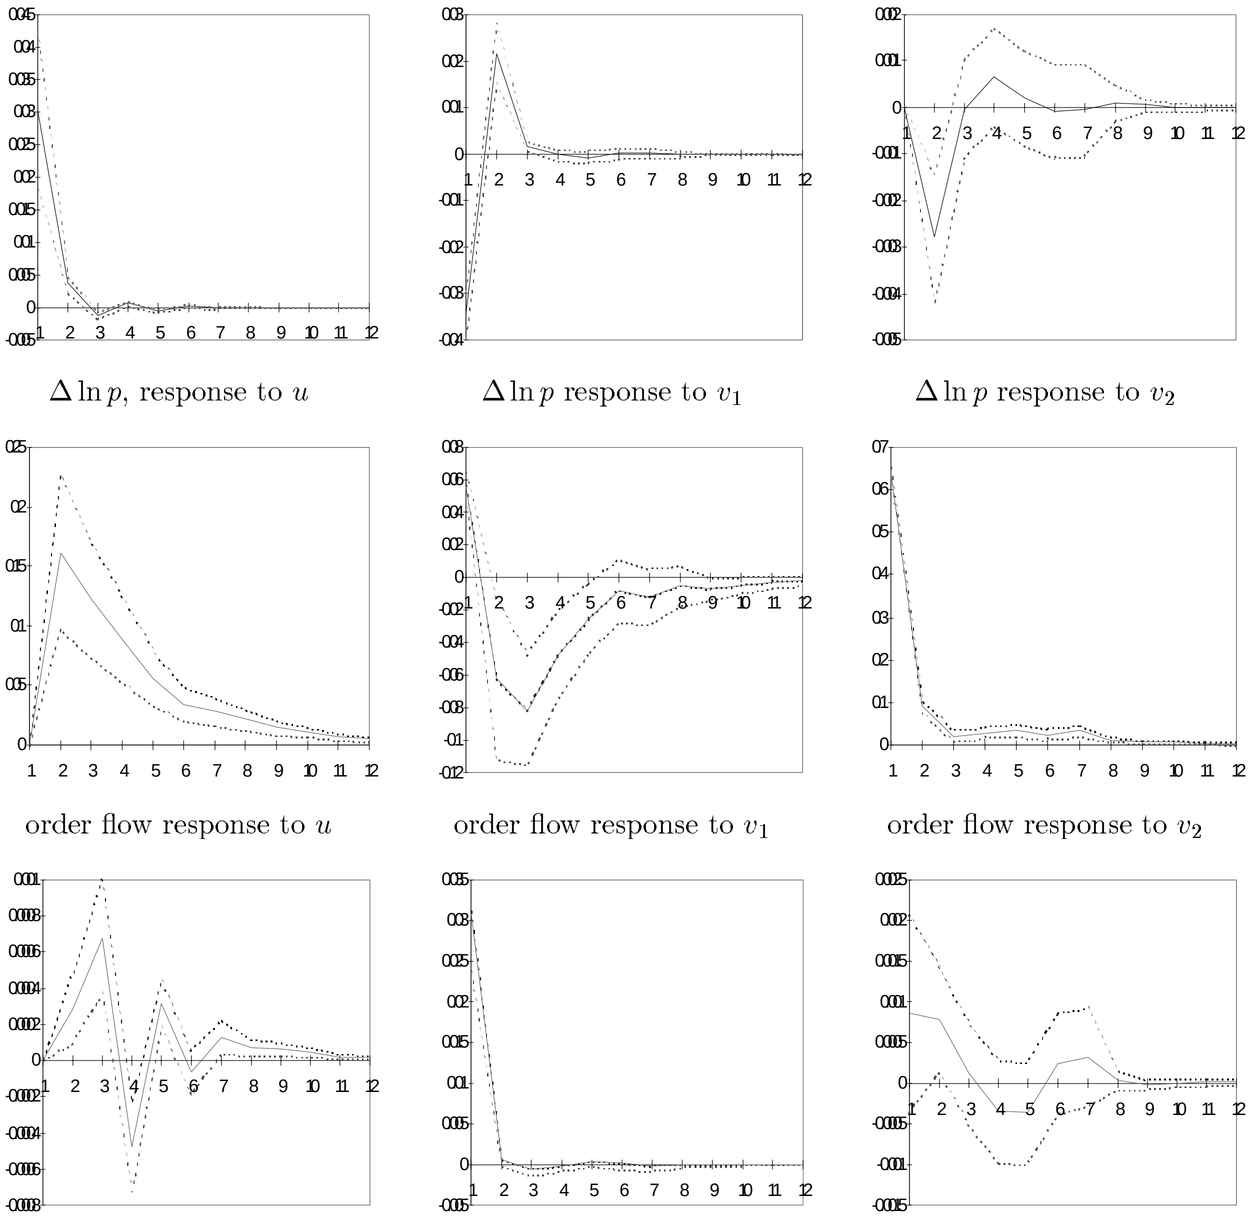

$\Delta \ln p$, response to u — $\Delta \ln p$ response to v_1 — $\Delta \ln p$ response to v_2

order flow response to u — order flow response to v_1 — order flow response to v_2

spread response to u — spread response to v_1 — spread response to v_2

Figure 3

The impulse responses shown in Figure 3 provide complementary evidence on the role of public and private news shocks. The graphs show the response patterns to a one standard deviation shock over one hour (12 five minute periods) with the dashed lines denoting the 95% confidence band calculated from the Monte Carlo experiments The left hand columns show the effects of the public news shock, u. Here we see that $\Delta \ln p$ rises significantly for just two periods. This means that the full effect of public news on the *level* of transactions prices is reached within 10 minutes and that the effect thereafter appears to be permanent. The effects on aggregate order flow are large and persist for as long as one hour. Thus, while public news appears to be quickly reflected in the level of prices, it has a long-lasting impact on the pattern of trade across the market. Public news shocks also appear to have persistent but alternating effects on the spread between quotes and transactions prices, $\ln q_t - \ln p_t$. While statistically significant, these variations are very small in magnitude. Thus, indicative quotes closely mirror the response of transactions prices to public news.

The center column shows that impact of the private news shock, v_1. This shock lowers prices and raises aggregate order flow generating a negative contemporaneous covariance between x_t and p_t. We can interpret this finding using the trading model. In particular, equations (4) and (6) imply that $Cov(\nabla p_{i,t}^n, \nabla x_{i,t}) = \phi_i^n \left[Var(\nabla p_{i,t}^n) - Cov(\nabla p_{i,t}^n, \nabla \mu_t) \right]$ where $\phi_i^n < 0$. Hence innovations in quotes and order flow will be negatively correlated when $Var(\nabla p_{i,t}^n) > Cov(\nabla p_{i,t}^n, \nabla \mu_t)$, a condition implying that dealers "over-adjust" quotes relative to the actual change in reservation prices. In other words, when private news leads to a 1% fall in the average reservation price, dealers lower their quotes buy more than 1%. Interestingly, this overreaction of quotes is followed by a correction. As the impulse response path for $\Delta \ln p$ shows, after the initial period, transactions prices rise quickly to their new long term level. By contrast, aggregate order flow falls below zero after the first period and remains there for some time. Indicative quotes do not mirror the response of transactions prices to the v_1 shocks. Initially, positive v_1 shocks raise the spread between quotes and transactions prices. This increase almost mirrors the fall in transactions prices so there is little impact on indicative quotes. In the next period, the spread remains unchanged so indicative quotes rise with transaction prices.

The effects of the v_2 private news shock are displayed in the right hand column. This shock identifies the effect of private news that shifts the distribution of reservation prices without affected quotes. Thus, it has no initial effect on transactions prices by definition. As the figures show, a positive v_2 shock temporarily raises aggregate order flow and leads to a small fall in transactions prices. This response pattern is consistent with an initial fall in average reservation prices that leads dealers to subsequently revise the quotes downwards. v_2 shocks also raise the spread between indicative quotes and transactions prices but the effect is small and statistically insignificant.

To summarize, the results above clearly demonstrate that public news is not the dominant source of high frequency exchange rate movements. Rather, the estimates suggest that these shocks contribute 50% at most to the variance of spot rates over horizons ranging from 5 minutes to 5 hours. The remaining variance is attributable to private news that leads to innovations in both transactions prices and aggregate order flow representing the pattern of interdealer trade. The results also show that while prices respond promptly to public news, the effects on aggregate order flow persist for as much as an hour. There is also evidence that prices overreact to private news.

5 Macro Implications

Although these results contrast with the prevailing macro view that exchange rate variations are primarily attributable to public news, their relevance for the analysis of exchange rate behavior over periods of weeks, months and longer is as yet unclear. To consider this issue, I now turn to examine the low frequency implications of the VAR results.

Figure 4 provides a simple yet informative way to examine how public and private news shocks impact upon spot exchange rates. The graphs decompose the movements in the log spot rate and cumulative aggregate order flow into three components that correspond to each of the three news shocks. Specifically, the components are

$$\begin{aligned} \ln \hat{p}_{i,t} &= (1-L)^{-1}\hat{\Gamma}_{1,i}(L)\hat{\eta}_{i,t} \\ \hat{x}^c_{i,t} &= (1-L)^{-1}\hat{\Gamma}_{2,i}(L)\hat{\eta}_{i,t} \end{aligned} \tag{18}$$

for $i = 1, 2, 3$ where $\hat{\Gamma}_{j,i}(L)$ denotes the j, i'th. element of the estimated matrix polynomial $\Gamma(L)$ in (16), $\hat{\eta}_{i,t}$ is the i'th element in the estimated shock vector $\hat{\eta}'_t \equiv [\hat{u}_t, \hat{v}_{1,t}, \hat{v}_{2,t}]$, and $x^c_t = \sum_{j=0}^{t} x_j$. Thus, $\ln \hat{p}_{1,t}$ and $\hat{x}^c_{1,t}$ represent estimates of the spot rate and cumulative order flow if the only news hitting the foreign exchange market during the sample period is public. The figure plots the $\ln \hat{p}_{i,t}$ and $\hat{x}^c_{i,t}$ components with solid lines and the series for $\ln p_t$ and x^c_t with dashed lines.

The graphs in the top row compare $\ln p_t$ with $\ln \hat{p}_{1,t}$, and x^c_t with $\hat{x}^c_{1,t}$. In the spot rate case, the two series follow broadly the same path over the first 35 days of the sample. Thereafter, the $\ln \hat{p}_{1,t}$ series appreciates (falls) by approximately 5% while the actual exchange rate appreciates by only 2%. The resulting 3% gap between the series continues for the next 20 days before slowly closing over the remainder of the sample. Thus, viewed over the whole sample, the behavior of $\ln \hat{p}_{1,t}$ and $\ln p_t$ differ considerably. In the case of order flow, the differences between x^c_t and $\hat{x}^c_{1,t}$ are much less pronounced.

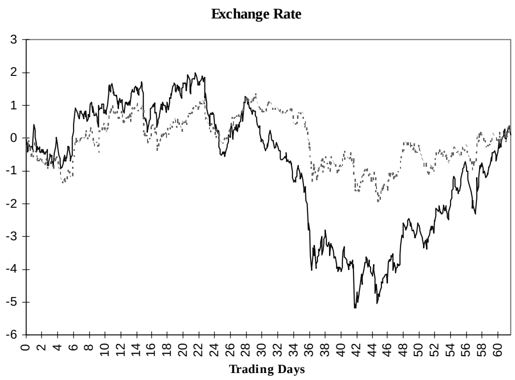

$\ln\hat{p}_{1,t}$ (solid line), $\ln p_t$ (dashed line)

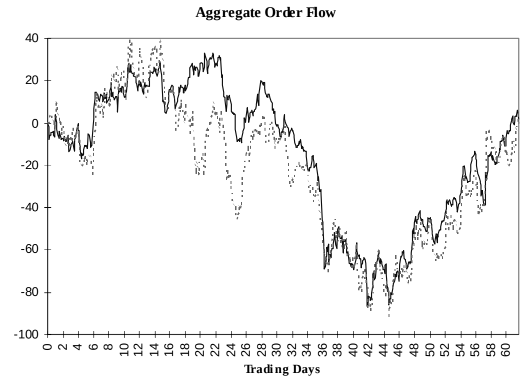

$\hat{x}^c_{1,t}$ (solid line), x^c_t (dashed line)

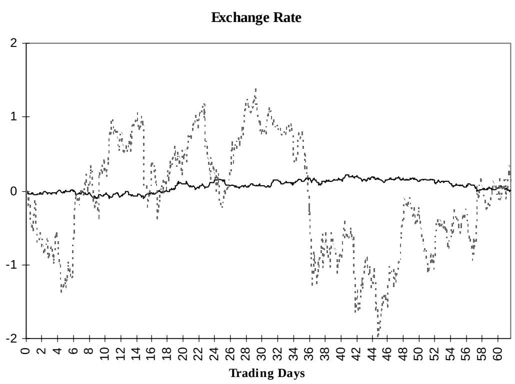

$\ln\hat{p}_{2,t}$ (solid line), $\ln p_t$ (dashed line)

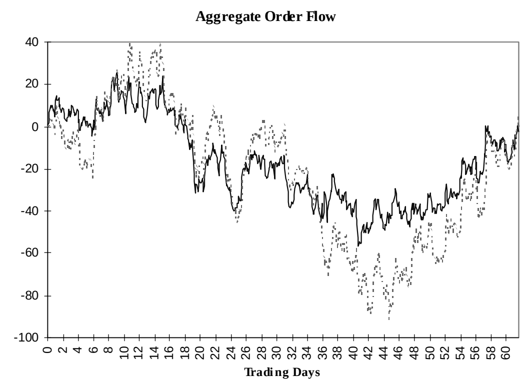

$\hat{x}^c_t$ (solid line), x^c_t (dashed line)

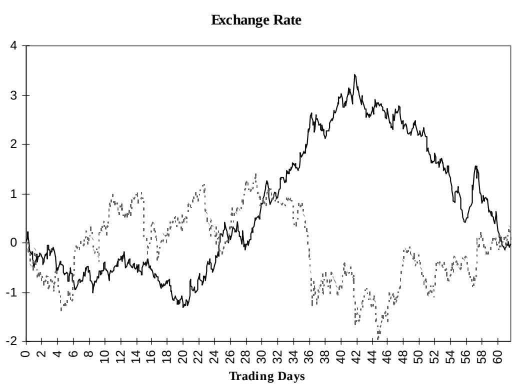

$\ln\hat{p}_{3,t}$ (solid line), $\ln p_t$ (dashed line)

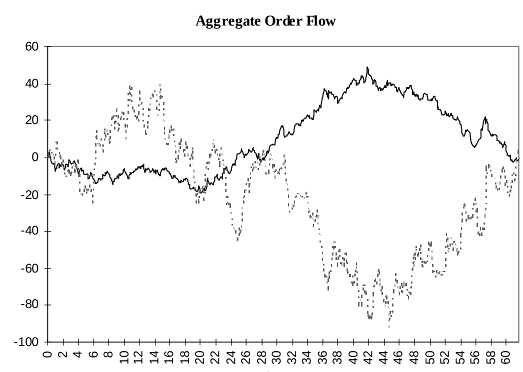

$\hat{x}^c_t$ (solid line), x^c_t (dashed line)

Figure 4

The center row of graphs compare $\ln p_t$ with $\ln\hat{p}_{2,t}$, and x^c_t with $\hat{x}^c_{2,t}$. Recall that the v_1 private news shocks contributed more than 50% of the variance of spot rate changes at high frequencies. Here we see that these effects almost completely dissipate as we move to lower frequencies. The plot for $\ln \hat{p}_{2,t}$ is essentially flat and displays none of the fluctuations in $\ln p_{2,t}$. The reason for this difference can be found in the impulse response functions. There we saw that spot rates overreacted to v_1 shocks, initially falling and then rising. These movements effectively cancel out when we consider the low frequency movements in exchange rate. The effects of v_1 on order flow also partially cancel out; $\hat{x}^c_{2,t}$ is less variable over the sample than x^c_t.

The effects of the v_2 shock are shows in the bottom row of graphs. Although these private news shocks contribute very little to the variance of spot rate changes at high frequency, they do have a significant cumulative effect on spot rates. As in the case of public news, there is a considerable divergence between the plot for $\ln \hat{p}_{3,t}$ and the actual exchange rate. In the absence of any other shocks, the spot rate would have depreciated by approximately 4% during the middle third of the sample, and depreciated by about 3% in the last third. Similarly, there is a significant difference between the paths for x^c_t with $\hat{x}^c_{3,t}$.

Since the spot rate is the sum of its estimated components, $\ln p_t = \ln \hat{p}_{1,t} + \ln \hat{p}_{2,t} + \ln \hat{p}_{3,t}$ we can decompose the variance of a k-period change in $\ln p_t$, $\Delta^k \ln p_t$, as

$$Var(\Delta^k \ln p_t) = Cov(\Delta^k \ln \hat{p}_{1,t}, \Delta^k \ln p_t) + Cov(\Delta^k \ln \hat{p}_{2,t}, \Delta^k \ln p_t) + Cov(\Delta^k \ln \hat{p}_{3,t}, \Delta^k \ln p_t).$$

If public news is the predominant source of spot rate variations over weeks, months and longer, we should expect to see the second and third covariance terms disappear as k becomes large. Alternatively, if these covariance terms remain different from zero as k rises, private news must contribute significantly to the low frequency variations in shot rates.

Table 4 reports estimates of the three covariance terms as a fraction of $Var(\Delta^k \ln p_t)$ for spot rate changes over one to fifteen trading days. The estimates are calculated using the values of $\ln p_t$,and $\ln \hat{p}_{i,t}$ at 1200 hrs. each day. The table reports the slope coefficient from the regression of $\Delta^k \ln \hat{p}_{i,t}$ on $\Delta^k \ln p_t$ with these daily data, together with standard errors that allow for conditional heteroskedasticity and MA$(k-1)$ error structure induced by the overlapping observations.

As column III of the table shows, private news in the form of v_2 shocks contribute significantly to spot rate changes over 1 to 15 days. Moreover, since the absolute size of the statistics rises with k, these results suggest that the contribution of private news may even increase as we move to consider spot rate dynamics at lower frequency. The results in column I compliment these findings. Here we see that the covariance between changes in the spot rate and the public news component, $\Delta^k \ln \hat{p}_{1,t}$, are greater than the variance of spot rate changes, and increase with k. Thus, as we move to lower frequencies, changes in the spot rate tend to increasingly *understate* the changes purely due to public news. The reason is that the $\ln \hat{p}_{1,t}$ and $\ln \hat{p}_{3,t}$ components move in different directions during much of the sample period. As a consequence, the effects of public and private news tend to partially cancel out over long horizons.

Table 4: Variance Decompositions By News Source

	I	II	II
k (days)	$Cov(\Delta^k \ln \hat{p}_{1,t}, \Delta^k \ln p_t)$	$Cov(\Delta^k \ln \hat{p}_{2,t}, \Delta^k \ln p_t)$	$Cov(\Delta^k \ln \hat{p}_{3,t}, \Delta^k \ln p_t)$
1	1.282	-0.020	-0.261
	(0.058)	(0.007)	(0.057)
2	1.239	-0.016	-0.223
	(0.059)	(0.009)	(0.060)
3	1.242	-0.015	-0.227
	(0.088)	(0.010)	(0.087)
4	1.286	-0.020	-0.266
	(0.101)	(0.009)	(0.100)
5	1.364	-0.017	-0.347
	(0.122)	(0.008)	(0.122)
10	1.615	-0.033	-0.582
	(0.169)	(0.005)	(0.168)
15	1.857	-0.029	-0.827
	(0.200)	(0.008)	(0.191)

Notes: The table reports estimates of the covariances as a fraction of $Var(\Delta^k \ln p_t)$.
Standard errors are reported in parentheses.

Is private news an important source of exchange rate movements over horizons of weeks, months, and longer? On the basis of the evidence above, the short answer is yes. If public news was the predominant source of low frequency exchange rate movements, there should not have been any significant difference between the paths for $\ln \hat{p}_{1,t}$ and $\ln p_t$, the plots for $\ln \hat{p}_{2,t}$ and $\ln \hat{p}_{3,t}$ should have both been flat, and the statistics in columns II and III of Table 4 should have been insignificantly different from zero. The results above clearly contradict these predictions.

Quantify the importance of private news is more difficult once we recognize that the data cover a relatively short time span. Even though the standard errors in Table 4 correct for overlapping observations when estimating the variance components, it is well-known that these corrections may be unreliable. The span of the data set is simply insufficient to reliably pin down the size of the private news effects over several weeks. That said, the statistics in Table 4 are nevertheless impressive. In particular, the statistics in column I indicate that daily changes in the spot rate understate the effects of public news by approximately 22% (1-1/1.282). While changes over three weeks understate the effects of public news by as much as 46% (1-1/1.857) [15]

[15] This does not mean that the variance of daily spot rate changes would have been at least 20% larger if public news was the sole source of exchange rate movements. The estimates of $\ln \hat{p}_{1,t}$, are susceptible to the Lucas critique and so should not be vewied as the level of the exchange rate in a world with only public news shocks. Recall that dealers have relatively little information about what is going on across the whole market. It would be surprising if dealers' actions were unaffected by the knowledge that all the information they individually receive was simultaneously being

6 Conclusion

This paper has used a new data set to investigate the origins of exchange rate movements. The key feature of this analysis is that it uses the joint behavior of transactions prices and order flow to identify the effects of public and private news shocks without the need to identify their source. This simple approach allowed me to quantify the impact of public and private news with a structural VAR estimated in very high frequency data.

The VAR estimates were used to examine the contribution of public and private news to spot rate movements over minutes, hours, days and weeks. The key finding to emerge from this analysis is that private news makes a significant contribution to spot rate movements at all frequencies. In particular, I found that the cumulative effects of both public and private news display low frequency variations in excess of the variations in the spot rate. I also found that the cumulative effects of public and private news appear to have offsetting effects on the spot rate. As a consequence, daily and weekly changes in the spot rate are positively correlated with, but significantly understate, the changes implied by the arrival of public news. These findings suggest one reason for the poor performance of existing macroeconomic exchange rate models; namely, their exclusive focus on public news.

received by all other dealers. Thus, is unlikely that the reduced form coefficients in $\Gamma_{1,i}(L)$ are invariant to the presence of private news.

References

Amihud Y. and H. Mendelson, 1980, "Dealership Markets : Market Making with Inventory", *Journal of Finance.*

Andersen, T. and T. Bollerslev, 1998, "Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic annoucements, and longer run dependencies", *Journal of Finance*, 53, 219-266..

Baillie, R. and T. Bollerslev, 1991, "Intra day and inter market volatility in foreign exchange rates", *Review of Economic Studies*, 58, 565-585.

Beck, S. 1993, "The Ricardian equivalence proposition: Evidence from foreign exchange markets", *Journal of International Money and Finance,* 12, 154-169.

Bjonnes, G., and D. Rime, 1998, "FX Trading ... Live: Impact of New Trading Environments", typescript, Norwegian School of Management, December.

Bollerslev, T. and I. Domowitz, 1993, "Trading patterns and prices in the interbank foreign exchange market", *Journal of Finance*, 48, 1421-1443.

Cheung, Y. and C Wong, 1998, "A survey of Market Practitioners' View on Exchange rate Dynamics", *Journal of International Economics,* forthcoming.

Covrig, V., and M. Melvin, 1998, "Asymmetric Information and Price Discovery in the FX Market: Does Tokyo Know More About the Yen?" typescript, Arizona State University.

Dacorogna, M., Muller, U., Nagler, R., Olsen, R. and O. Pictet, 1993, "A geographical model for the daily and weekly seasonal volatility in the FX market", *Journal of International Money and Finance*, 12, 413-438.

Duffie, D., 1992, *Dynamic Asset Pricing Theory* (Princeton University Press, Princeton NJ).

Engel, C. and J. Frankel, 1984, "Why interest rates react to money announcements: An answer of the foreign exchange market", *Journal of Monetary Economics*, 13, 31-39.

Engle, R., Ito, T. and W-L. Lin, 1990, "Meteor showers or heat waves? Heteroskedastic intradaily volatility in the foreign exchange market", *Econometrica*, 58, 525-542.

Evans, M. D. D. "The Microstructure of Foreign Exchange Dynamics" typscript, Georgetown University, presented at the NBER Market Microstructure Conference, December 4 1997.

Fleming, M. J. and E. M. Remolona 1999 "Price Formation and Liquidity in the U.S. Treasury Market: The Response to Public Information", *Journal of Finance,* forthcoming.

Frankel J, and A.K. Rose, 1995 "Empirical Research on Nominal Exchange Rates", in Grossman and Rogoff (eds.), *Handbook of International Economics* Vol 3, North Holland.

Glosten L.and P. Milgrom, 1985, "Bid, Ask and Transactions Prices in a Specialist Market with Heterogeneously Informed Agents", *Journal of Financial Economics.*

Goodhart, C. and M. Giugale, 1993, "From hour to hour in the foreign exchange market", *The Manchester School*, 61, 1-34.

Goodhart, C., Ito, T. and R. Payne, 1996, "One Day in June 1993: A study of the working of Reuters 2000-2 electronic foreign exchange trading system", in Frankel, Galie and Giovannini (eds.) *The Microstructure of Foreign Exchange Markets*, Chicago.

Guillaume, D., Dacorogna, M., Dave, R., Muller, U., Olsen, R. and O. Pictet, 1994, "From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets", A discussion paper by the O&A research group, Zurich, October.

Hardouvelis, G., 1988, "Economic news, exchange rates, and interest rates", *Journal of International Money and Finance,*7, 23-25.

Hasbrouck, J., 1991 "Measuring the information content of stock trades", *Journal of Finance* 41(1), 179-202.

Ho. T. and H. Stoll, 1983 "The Dynamics of Dealer Markets Under Competition", *Journal of Finance.*

Ito, T., R. Lyons, and M, Melvin, 1998 "Is there private information in the FX market? The Tokyo experiment", *Journal of Finance,* 53, 1111-1130.

Ito, T. and V Roley, 1987 "News from the U.S. and Japan: Which moves the yen/dollar exchange rate", *Journal of Monetary Economics,* 19, 255-277.

Kyle, A., 1985 "Contintuous Auctions and Insider Trading", *Econometrica.*

Lyons, R., 1995 "Tests of microstructure hypotheses in the foreign exchange market", *Journal of Financial Economics*

Lyons, R., 1996a "Foreign exchange volume: sound and fury signifying nothing?" in Frankel, Galie and Giovannini (eds.) *The Microstructure of Foreign Exchange Markets*, Chicago

Lyons, R., 1996b Comment on "One Day in June 1993: A study of the working of the Reuters 2000-2 Electronic Foreign Exchange Trading System" in Frankel, Galie and Giovannini (eds.) *The Microstructure of Foreign Exchange Markets*, Chicago

Lyons, R.,1997 "A simultaneous trade model of the foreign exchange hot potato", *Journal of International Economics*, May, 275-298.

Payne, R., 1997 "Announcement effects and seasonality in the intra-day foreign exchange market", LSE Financial Markets Group Discussion Paper, No. 238, March.

Payne, R. 1999 "Informed trade in spot foreign exchange markets: An empirical investigation", typscript, LSE.

Peiers, B., 1997 "Informed Traders, Intervention, and Price Leadership: A Deeper View of the Microstructure of the Foreign Exchange Market", *Journal of Finance,* 52 1589-1614..

Yao, J., 1997a "Market making the interbank foreign exchange market", Working Paper S-98-3, Salomon Center, New York University

Yao, J., 1997b "Spread Components and Dealer Profits in the Interbank Foreign Exchange Market,' Working Paper S-98-4, Salomon Center, New York University

Appendix A

This appendix describes the deviation of the optimal quotes in equation (3). First, I rewrite expected profits in (2) as

$$\left(p_{i,t}^a - E[p_t^*|\Omega_{i,t}, r_{j,t} > p_{i,t}^a]\right)(1 - G(p_{i,t}^a)) - \left(p_{i,t}^b - E[p_t^*|\Omega_{i,t}, r_{j,t} < p_{i,t}^b]\right)G(p_{i,t}^b).$$

Next, note that

$$\begin{aligned} E[p_t^*|\Omega_{i,t}, r_{j,t} &> p_{i,t}^a] = \int_{p_{i,t}^a}^{\infty} E\left[p_t^*|\Omega_{i,t}, r_{j,t}\right] \frac{g(r_{j,t})}{1 - G(q_{i,t}^a)} dr_{j,t} \\ E[p_t^*|\Omega_{i,t}, r_{j,t} &< p_{i.t}^b] = \int_0^{p_{i,t}^b} E\left[p_t^*|\Omega_{i,t}, r_{j,t}\right] \frac{g(r_{j,t})}{G(p_{i,t}^b)} dr_{j,t} \end{aligned}$$

where the support of the reservation price distribution is in the positive orthant. Substituting these expressions into the expected profit function, gives

$$\begin{aligned} \pi(p_{i,t}^a, p_{i,t}^b) &\equiv p_{i,t}^a(1 - G(p_{i,t}^a)) - \int_{p_{i,t}^a}^{\infty} E\left[p_t^*|\Omega_{i,t}, r_{j,t}\right] g(r_{j,t}) dr_{j,t} \\ &+ \int_0^{p_{i,t}^b} E\left[p_t^*|\Omega_{i,t}, r_{j,t}\right] g(r_{j,t}) dr_{j,t} - p_{i,t}^b G(p_{i,t}^b). \end{aligned}$$

The first and second-order conditions for expected profit maximization are respectively

$$\begin{aligned} \frac{\partial \pi(p_{i,t}^a, p_{i,t}^b)}{\partial p_{i,t}^a} &= (1 - G(p_{i,t}^a) - p_{i,t}^a g(p_{i,t}^a) + E[p_t^*|\Omega_{i,t}, p_{i,t}^a] g(p_{i,t}^a) = 0 \\ \frac{\partial \pi(p_{i,t}^a, p_{i,t}^b)}{\partial p_{i,t}^b} &= -G(p_{i,t}^b) - p_{i,t}^b g(p_{i,t}^b) + E[p_t^*|\Omega_{i,t}, p_{i,t}^b] g(p_{i,t}^b) = 0 \end{aligned}$$

and

$$\begin{aligned} \frac{\partial \pi^2(p_{i,t}^a, p_{i,t}^b)}{\partial p_{i,t}^a \partial p_{i,t}^a} &= -2g(p_{i,t}^a) - (p_{i,t}^a - E\left[p_t^*|\Omega_{i,t}, p_{i,t}^a\right]) g'(p_{i,t}^a) < 0 \\ \frac{\partial \pi(p_{i,t}^a, p_{i,t}^b)}{\partial p_{i,t}^b \partial p_{i,t}^b} &= -2g(p_{i,t}^b) - (p_{i,t}^b - E\left[p_t^*|\Omega_{i,t}, p_{i,t}^b\right]) g'(p_{i,t}^b) < 0 \end{aligned}$$

The optimal quote equations are derived from the first-order conditions using the fact that $E[p_t^*|\Omega_{i,t}, p_{i,t}^a] = E[p_t^*|\Omega_{i,t}, p_{i,t}^b] = E[p_t^*|\Omega_{i,t}]$ because $\{p_{i,t}^a, p_{i,t}^b\} \in \Omega_{i,t}$.

Appendix B

This appendix shows that the regression coefficients in (6) are negative and can approximated as shown in the text. To begin, note that $E\left[\nabla(p^n_{i,t}-\mu_t)\right]=E\left[\nabla x_{i,t}\right]=0$. We can therefore write,

$$\begin{aligned}Cov\left(\nabla x_{i,t},\nabla(p^n_{i,t}-\mu_t)\right) &= E\left[\nabla x_{i,t}\nabla(p^n_{i,t}-\mu_t)|\nabla(p^n_{i,t}-\mu_t)>0\right]\Pr\left(\nabla(p^n_{i,t}-\mu_t)>0\right)\\ &+E\left[\nabla x_{i,t}\nabla(p^n_{i,t}-\mu_t)|\nabla(p^n_{i,t}-\mu_t)\le 0\right]\Pr\left(\nabla(p^n_{i,t}-\mu_t)\le 0\right).\end{aligned}$$

To evaluate the terms above, let τ denote the start of period t before the news hits the market, so $\nabla x_{i,t}\equiv x_{i,t}-x_{i\tau}$ and $\nabla(p^n_{i,t}-\mu_t)\equiv(p^n_{i,t}-\mu_t)-(p^n_{i,\tau}-\mu_\tau)$ for $n=\{a,b\}$.

To evaluate $Cov\left(\nabla x_{i,t},\nabla(p^a_{i,t}-\mu_t)\right)$, consider the case where $\nabla(p^a_{i,t}-\mu_t)>0$. Here $\nabla x_{i,t}$ will be zero except when $x_{i,t}=0$ and $x_{i\tau}=1$. This occurs with probability

$$\Phi^a(+)\equiv\int_{p^a_{i,\tau}-\mu_\tau}^{p^a_{i,t}-\mu_t}\frac{h(r)}{1-H(p^a_{i,\tau}-\mu_\tau)}drr,$$

so

$$E\left[\nabla x_{i,t}\nabla(p^n_{i,t}-\mu_t)|\nabla(p^n_{i,t}-\mu_t)>0\right]=-\Phi^a(+)E\left[\nabla(p^n_{i,t}-\mu_t)|\nabla(p^n_{i,t}-\mu_t)>0\right].$$

For the case where $\nabla(p^a_{i,t}-\mu_t)\le 0$, $\nabla x_{i,t}$ will be zero except when $x_{i,t}=1$ and $x_{i\tau}=0$. This occurs with probability

$$\Phi^a(-)\equiv\int_{p^a_{i,t}-\mu_t}^{p^a_{i,\tau}-\mu_\tau}\frac{h(r)}{1-H(p^a_{i,t}-\mu_t)}dr,$$

so

$$E\left[\nabla x_{i,t}\nabla(p^a_{i,t}-\mu_t)|\nabla(p^a_{i,t}-\mu_t)\le 0\right]=\Phi^a(-)E\left[\nabla(p^a_{i,t}-\mu_t)|\nabla(p^a_{i,t}-\mu_t)\le 0\right].$$

Combining these components we have

$$\begin{aligned}Cov\left(\nabla x_{i,t},\nabla(p^a_{i,t}-\mu_t)\right) &= \Phi^a(-)E\left[\nabla(p^a_{i,t}-\mu_t)|\nabla(p^a_{i,t}-\mu_t)\le 0\right]\Pr\left(\nabla(p^a_{i,t}-\mu_t)\le 0\right)\\ &-\Phi^a(+)E\left[\nabla(p^a_{i,t}-\mu_t)|\nabla(p^a_{i,t}-\mu_t)>0\right]\Pr\left(\nabla(p^a_{i,t}-\mu_t)>0\right)<0.\end{aligned}$$

Consequently, the regression coefficient, $\phi^a_i\equiv Cov\left(\nabla x_{i,t},\nabla(p^a_{i,t}-\mu_t)\right)/Var\left(\nabla(p^a_{i,t}-\mu_t)\right)$, must also be negative.

By analogous argument, we also have

$$\begin{aligned}Cov\left(\nabla x_{i,t},\nabla(p^b_{i,t}-\mu_t)\right) &= \Phi^b(-)E\left[\nabla(p^b_{i,t}-\mu_t)|\nabla(p^b_{i,t}-\mu_t)\le 0\right]\Pr\left(\nabla(p^b_{i,t}-\mu_t)\le 0\right)\\ &-\Phi^b(+)E\left[\nabla(p^b_{i,t}-\mu_t)|\nabla(p^b_{i,t}-\mu_t)>0\right]\Pr\left(\nabla(p^b_{i,t}-\mu_t)>0\right)<0,\end{aligned}$$

where

$$\Phi^b(+) \equiv \int_{p^b_{i,\tau}-\mu_\tau}^{p^b_{i,t}-\mu_t} \frac{h(r)}{H(p^b_{i,\tau}-\mu_\tau)} dr, \qquad \text{and} \qquad \Phi^b(-) \equiv \int_{p^b_{i,t}-\mu_t}^{p^b_{i,\tau}-\mu_{\tau t}} \frac{h(r)}{H(p^b_{i,t}-\mu_t)} dr.$$

So the regression coefficient, $\phi^b_i \equiv Cov\left(\nabla x_{i,t}, \nabla(p^b_{i,t}-\mu_t)\right)/Var\left(\nabla(p^b_{i,t}-\mu_t)\right)$, must also be negative.

In the case where the distribution of innovations, $\nabla(p^n_{i,t}-\mu_t)$ is symmetric and has a small variance, the expressions for the covariance terms above can be well-approximated by

$$\begin{aligned} Cov\left(\nabla x_{i,t}, \nabla(p^a_{i,t}-\mu_t)\right) &\simeq -\frac{h(p^a_{i,t}-\mu_t)}{1-H(p^a_{i,t}-\mu_t)} E\left[\nabla(p^a_{i,t}-\mu_t)|\nabla(p^n_{i,t}-\mu_t)>0\right] < 0, \\ Cov\left(\nabla x_{i,t}, \nabla(p^a_{i,t}-\mu_t)\right) &\simeq \frac{h(p^b_{i,t}-\mu_t)}{H(p^b_{i,t}-\mu_t)} E\left[\nabla(p^b_{i,t}-\mu_t)|\nabla(p^b_{i,t}-\mu_t)\leq 0\right] > 0. \end{aligned}$$

For clarity, the regression coefficients reported in the text are based on these approximations.