The Price Impact of Currency Trades: Implications for Intervention

Martin D. D. Evans*

Richard K. Lyons

This draft: 30'th June 2000

Abstract

Central-bank intervention always introduces asymmetric information. There are no plausible, common-knowledge alternatives, regardless of intervention type (e.g announced vs. unannounced), and regardless of efficacy channel (e.g., portfolio vs information). This paper recasts intervention analysis into a framework of asymmetric information. This allows us to address long-standing questions in a new way. We offer four main results. First, we find evidence of imperfect substitutability, a necessary condition for the efficacy of the portfolio channel. Second, we find evidence of an operative signaling channel. Third, we establish a lower bound on the unconditional price impact of sterilized trades: 5 basis points per $100 million. Fourth, we provide an explicit answer to an important unanswered question: "When should Central Banks intervene?"

Correspondence
Richard K. Lyons
Haas School of Business, UC Berkeley
Berkeley, CA 94720-1900
Tel: 510-642-1059, Fax: 510-643-1420
lyons@haas.berkeley.edu
www.haas.berkeley.edu/~lyons

---

* Respective affiliations are Georgetown University and NBER, and UC Berkeley and NBER. We thank the following for valuable comments: Rich Clarida. Lyons thanks the National Science Foundation for financial assistance.

The Price Impact of Currency Trades: Implications for Intervention

Central Bank intervention in the Foreign Exchange market always introduces asymmetric information. There are no plausible, common-knowledge alternatives, regardless of intervention type (e.g. announced vs. unannounced), and regardless of efficacy channel (e.g., portfolio vs. information). This paper recasts intervention analysis into a framework of asymmetric information. In doing so, we bring new perspective to two long-standing questions: (i) Why do interventions affect foreign currency prices? (ii) When is intervention particularly effective?

Central bankers would love an answer to the "when" question. But the right time to execute an intervention trade is not well understood. As an empirical matter, there is little work to draw upon. Dominguez and Frankel (1993b) make progress linking intervention trades directly to price movements. But their intervention data—though the best available—span only a few time intervals (e.g., 10-11 am EST) and provide little discriminating power in terms of market "states" (e.g., calm versus turbulent). In this paper, we utilize transactions data that permit us to examine how the price-impact of interventions change over the day, and over market states.

We draw on the economics of information to examine why intervention can affect prices. There are two basic types of asymmetric information that are relevant for this analysis. The first type is well recognized in the literature; asymmetric information between the central bank and the public.[1] The mere fact that the central bank controls several exchange-rate fundamentals (e.g., interest rates and inflation) gives it superior information about the paths of those fundamentals. Another source of central-bank information advantage is its access to privileged information (e.g., data that are not yet public). Models that focus on this "Type-1" information asymmetry play an important role in intervention theory. These models produce the well known "information channel", a distinct channel through which intervention can affect the exchange rate.

The second type of asymmetric information exists between private agents. This "Type-2" asymmetry receives little attention in theoretical work on intervention (notable exceptions include Bhattacharya and Weller 1997, Vitale 1997, and Montgomery and Popper 1998). Traditional models of intervention, which are based on asymmetric information of the first type, do not include the second. In those earlier models, all private agents learn about

[1] The intervention literature is surveyed by Dominguez and Frankel (1993b), Edison (1993), Almekinders (1995), and Sarno and Taylor (2000).

changes in fundamentals simultaneously, and the mapping from these new fundamentals to price is common knowledge.

Type-2 asymmetric information is in many respects more general than Type-1—at least as it relates to intervention—because it plays a role in determining intervention's effectiveness irrespective of the channel through which intervention works. Consider, for example, the information channel: Irrespective of whether intervention is sterilized, some agents observe the central bank's trade before others, and the market's understanding of this fact will affect the price response. This is also true for the portfolio-balance channel: The portfolio shift that the intervention trade is forcing on the public is not publicly observed, so price adjustment occurs via learning from the trading process. But what about the case when the intervention is announced? For an announced intervention to eliminate the information asymmetry it would have to (i) be fully credible, (ii) be fully informative regarding the size, direction, and timing of the trade; and (iii) occur in advance of or simultaneous with the trade. "Announced" interventions in the existing empirical literature do not fulfill these conditions.

The fact that intervention's effects always involve information asymmetry opens the door to novel theoretical and empirical strategies. On the theoretical side, we present a model that embeds intervention analysis in a setting with information asymmetry. The model is a variant of the Portfolio Shifts model in Evans and Lyons (1999). It clarifies the role played by interdealer order flow in conveying information about the actions of the central bank across the market. The model provides us with three key results: (i) Order flow moves prices if the public's demand for foreign currency is less than perfectly elastic. (ii) Unannounced sterilized interventions can only affect prices if they contributed to observed interdealer order flow. (iii) In the absence of Type-1 asymmetry, the price-impact of private trades is indistinguishable from the impact of sterilized interventions conveying no signal about future policy.

Our empirical analysis is guided by result (iii). If private trades are representative of intervention trades—in the sense of providing a reliable proxy for how price impact varies over time and states—then we can use data on private trades to address the state-dependence of intervention's efficacy.[2] Our empirical analysis uses four months of transac-

[2] This indirect approach is necessitated by the available data. If existing transactions data sets included intervention trades, we could employ the methodologies to detect the effects of asymmetric information used elsewhere in the exchange rate literature (e.g., Lyons 1995, Yao 1997, Covrig and Melvin 1998, Ito et al. 1998, Cheung and Wong 1998, Bjonnes and Rime 1998, Evans 1999, Naranjo and Nimalendran 1999, Payne 1999, and Evans and Lyons 1999). Moreover, there is no compelling evidence that central bank orders have greater price impact than other participants' orders. The two papers most useful for comparative purposes are Dominguez and Frankel (1993) and Evans and Lyons (1999). The Dominguez-Frankel estimate of price impact from intervention trades is very similar to that found by Evans and Lyons (1999) for interdealer trades—about 8 basis points per $100 million.

tion data from the spot DM/$ market that contain a truly vast number of private trades. As such, they provide us with the empirical power we need to address questions that hitherto have been difficult to answer. Examples include: Is there evidence of imperfect substitutability between different currency-denominated assets? Can we quantify the magnitude and duration of an intervention's price-impact? Is the price-impact of intervention state-dependent?

In addressing these questions we establish four main results. First, we find strong evidence of imperfect substitutability, a necessary condition for the efficacy of the portfolio channel. Second, we find evidence of persistent price effects consistent with an operative signaling channel. Third, we establish a lower bound on the price-impact of intervention trades: 5 basis points per $100 million. Fourth, we provide explicit guidance regarding when intervention trades are most effective in moving prices.

The remainder of the paper is composed of five sections. Section 1 introduces our trading-theoretic approach to the analysis of intervention. Section 2 presents our trading model that embeds intervention analysis in a setting with information asymmetry. Section 3 describes the data and presents our empirical results. Section 4 discusses the economic implications of our findings. Section 5 concludes.

1. A Trade-Theoretic Approach to Intervention

In this section we project the traditional macro approach to intervention onto theories of asset trading. This is useful for two main reasons. First, theories of asset trading provide greater resolution on how intervention can affect price. By greater resolution, we mean that individual channels within the macro approach can be broken into separate sub-channels. These sub-channels are themselves empirically identifiable. Second, a trading-theoretic approach makes it easy to establish that all channels through which intervention affects price involve some kind of information asymmetry. This point is not so clear within the macro approach.

Within the macro approach, intervention affects price through two channels: imperfect substitutability and asymmetric information. Imperfect substitutability is addressed within this approach using portfolio-balance models. These models are most useful for analyzing a particular type of intervention, namely sterilized intervention. The second channel—asymmetric information—is addressed within the macro approach using monetary models. These models are most useful for analyzing unsterilized intervention, or sterilized intervention that includes effects from information signaling. Let us consider how each of these two macro channels is manifested within a trading-theoretic approach.

Imperfect Substitutability

Theories of asset trading address imperfect substitutability at two different levels. The first level is the dealer level. Dealers—being risk averse—need to be compensated for holding positions they would not otherwise hold. This requires a temporary risk premium, which takes the form of a price-level adjustment. This price effect is temporary because the risk premium is not necessary once these positions are shared with the wider market. In trading-theoretic models, price effects from this channel are termed "inventory effects." In most markets, these effects fizzle quickly because wider risk sharing occurs rapidly (e.g., within a day). Figure 1 illustrates this. The short-run market (net) demand curve slopes up. Barring any other factors that produce upward slope, the longer-run demand curve is flat.[3]

Figure 1: Demand Under Inventory Effects Only

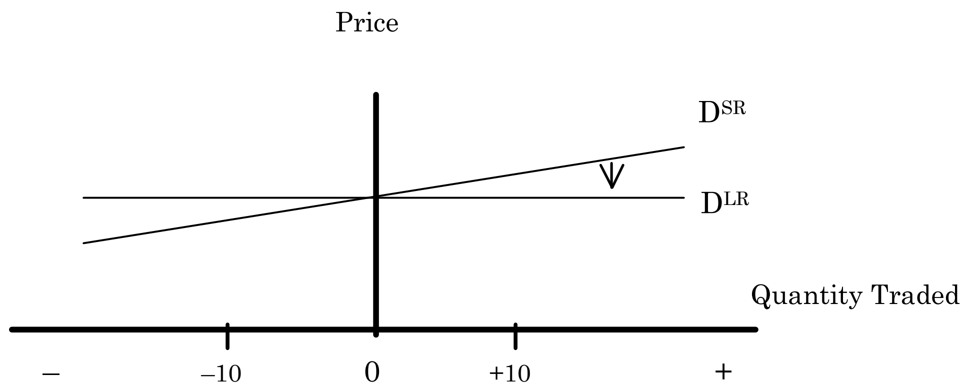

To read the figure, note that the effective spread for 10 units is the difference in price along the D^{SR} curve between –10 and +10.

Imperfect substitutability within trading models also operates at a second level, the market-wide level. At this level, the market as a whole—being risk averse—needs to be compensated for holding positions it would not otherwise hold. This too requires a risk premium that takes the form of a price-level adjustment. Price adjustment at this second level is not temporary because risk is fully shared at this level. Figure 2 illustrates this. The short-run market (net) demand curve slopes up, as before, reflecting the first level of imperfect substitutability—the dealer level. The longer-run demand also slopes up now, reflecting the persistent effect on price necessary to establish the appropriate risk premium.

3 Readers familiar with microstructure finance will recognize that the figure assumes there is no fixed component to the spread—i.e., the effective spread shrinks to zero as the quantity traded shrinks to zero.

Figure 2: Demand With Both Levels of Imperfect Substitutability

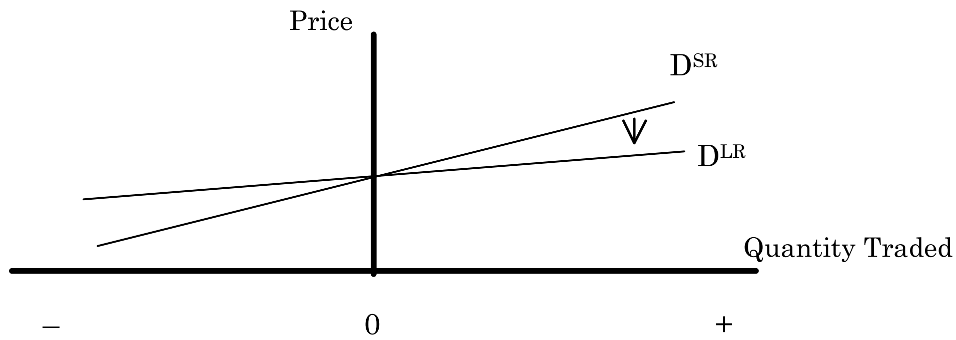

The first of these two levels of imperfect substitutability is not present within the macro approach. The term is used within the macro approach to refer to the second level only. The logic for addressing only the second level is that effects of the first are presumed fleeting enough to be negligible at horizons of one week or more. This is of course an empirical question—one that our trading-theoretic approach allows us to address in a rigorous way.

Our approach tests whether either of these two levels of imperfect substitutability is present in foreign exchange markets. If the first level is present—the dealer level—then intervention can be expected to have temporary effects on the exchange rate. We term these effects a "temporary portfolio-balance channel." If the second level is present—the market wide level—then intervention can be expected to have persistent effects on the exchange rate. We term these effects a "persistent portfolio-balance channel."

There is an important sense in which our strategy for measuring portfolio-balance effects from intervention differs from those in the literature. Traditionally, empirical estimation of a quantity role in portfolio balance relies on a quite narrow definition of quantities, namely, net supplies of government—i.e. public—bonds (see, e.g., Dominguez 1986 and Lyons 1987).[4] This is reasonable if (i) the net supply of private securities is zero (which is reasonable given public portfolios), and (ii) risk premia are a function of net supplies only. The last of these two conditions is an empirical matter, one we address in this paper. To wit: a finding that market order flow generates persistent portfolio-balance effects is inconsistent with the traditional empirical approach, and would appear to violate the second condition.

[4] Still more extreme is the assumption of Ricardian Equivalence, under which even net government assets are irrelevant to risk premia since private agents internalize the government's budget constraint and balance sheet.

Asymmetric Information

Theories of asset trading provide a third channel through which order flow affects price, the information channel (see., e.g., Kyle 1985 and Glosten and Milgrom 1985). If order flow conveys fundamental information, then it will have a persistent effect on price beyond the persistent portfolio-balance effect noted above. Mapped into the market demand curve, this channel adds additional slope to the longer-run schedule shown in Figure 2. Figure 3 provides an illustration.

Figure 3: Demand With Level-Two Imperfect Substitutability and Asymmetric Information (D^{LR*})

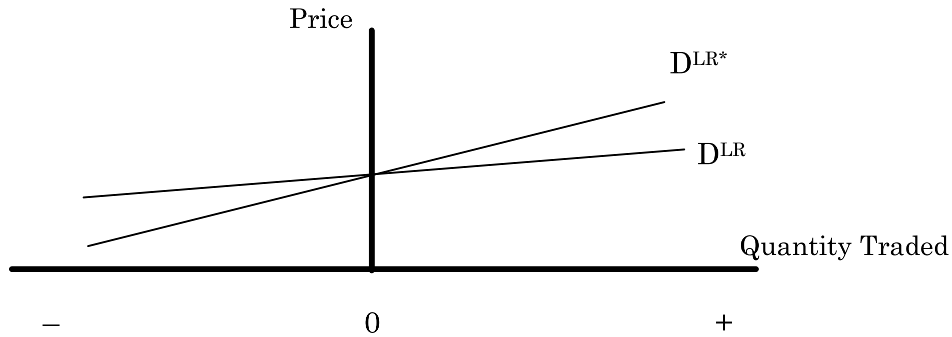

Unsterilized intervention trades are an obvious example of order flow that conveys fundamental information. Sterilized intervention trades that signal future fundamentals are another obvious example. Examples from the general public include trades that convey information about fundamentals like trade balances and capital flows (these trades are typically realized well before published macroeconomic statistics are available).

One manifestation of asymmetric information that is relevant for intervention is called “event uncertainty.” The concept of event uncertainty was introduced by Easley and O’Hara (1992). In an economy characterized by event uncertainty, *trades are more informative when trading intensity is high*. To understand why, consider an information structure in which new information may exist, but not necessarily. Suppose, for example, that with probability p some traders have observed some new information, and with probability (1-p) no new information exists. In the event of new information, the news can be either good or bad (with known probabilities). Easley and O’Hara demonstrate that if there is no trade at time t, then a rational dealer raises her subjective probability of the no-information event. If trading intensity is low, then, an incoming trade induces a smaller update in beliefs because it is less likely to be signaling news (i.e., it is less informative). On the flip side, trades occurring when trading intensity is high are more likely to be signaling news.

It is important to note that this "signaling channel" under event uncertainty is quite different from the standard signaling channel in the traditional intervention literature. Here, the signal conveyed by order flow resolves uncertainty about agents' expectations of future fundamentals. To link this to the discussion in the introduction about asymmetric information, the signaling channel here relates to Type-2 asymmetric information. The signaling channel in the traditional intervention literature relates instead to Type-1.

2. The Model

Overview

The model is designed to show how information conveyed by different types of sterilized intervention is learned through the trading process. In particular, it clarifies the role played by interdealer order flow in conveying information about the actions of the central bank (CB) across the market. The model accommodates the main intervention types and the main channels through which intervention can be effective. The four intervention types we examine appear in the following matrix:[5]

Figure 4: Sterilized Intervention Types

	No Signal	**Signal**
Unannounced		
Announced		

Intervention is identified by a foreign currency trade between a central bank and a foreign exchange dealer. Trades may be classified along two dimensions. First, they may or may not contain a signal about future CB (monetary) policy. CB trades that are uncorrelated with future interest rates are classified by our model as interventions containing no signal. Interventions conveying a signal are identified by trades that are correlated with future interest rates. Because we only consider sterilized interventions, CB trades are uncorrelated with current interest rates in the model.

[5] Our empirical work focuses in the upper-left cell of the matrix: sterilized intervention that is unannounced and conveys no signal about future CB policy. For this reason, we do not examine the impact of unsterilized interventions, (i.e., CB trades that are correlated with current interest rate changes) or the effects of coordinated intervention by more than one CB. However, our model could be extended in both these directions.

The second dimension concerns whether CB trades are announced or unannounced. As an empirical matter, announcements can be simultaneous or delayed. Sometimes announcements communicate the direction and exact trade size, but typically this is not the case. For our model, announced interventions occur simultaneously with the trade, and the announcement communicates the direction and exact size. An unannounced intervention in our model provides no information beyond that conveyed by the trade itself to the counter-party of that trade.

Order flow is a key forcing variable in the model, and it is important to understand why. The quick answer is that it conveys information. In the model there are two types of information conveyed by order flow: *portfolio-balance information* and *interest-rate information*. To understand portfolio-balance information, start with the basic structure of the model. At the beginning of the day, the public and central bank can place orders in the foreign exchange market. These orders are not publicly observable (except in the case of a central-bank order that is announced, i.e., an announced intervention). Initially, dealers take the other side of these trades—shifting their portfolios accordingly. At the end of the day, dealers share the risk of these positions with the public. Because we assume that public demand is not perfectly elastic (i.e., different-currency assets are imperfect substitutes), the sharing of these risky positions does not completely diversify away their price impact.[6] How much price impact is necessary to clear the market? That depends on two factors: the size of the initial portfolio shift and the elasticity of public demand. Order flow conveys information about the first of these two factors—the size of the initial portfolio shift.

Interest-rate information is conveyed by order flow only in the unannounced-sterilized-signal cell of the matrix (the upper left cell). In the announced-sterilized-signal cell (the lower left cell), order flow is correlated with future interest rates, but order flow does not convey any information about interest rates beyond what is conveyed by the announcement.

Specifics

Consider a pure exchange economy with two assets, one riskless and one risky, the latter representing foreign exchange. At the end of day T, foreign exchange earns an interest payoff R, which is composed of a series of random increments:

$$R = \sum_{t=1}^{T} r_t$$

[6] For evidence of imperfect substitutability across U.S. stocks, see Scholes (1972), Shleifer (1986) and Bagwell (1992), among others. Note that the size of the order flows the DM/$ spot market needs to absorb are on average more than 10,000 times those absorbed in a representative U.S. stock (e.g., the average daily volume on NYSE stocks in 1998 was $9.3 million, whereas the average daily volume in DM/$ spot was about $300 billion).

representing, for example, interest differentials. The increments r_t are i.i.d. Normal$(0,\Sigma_r)$ and are observed at the end of trading each day.

The foreign exchange market has three participant types. There are N dealers, indexed by i, a continuum of customers (the public), indexed by $z \in [0,1]$, and a central bank (CB). Dealers and customers all have identical negative exponential utility defined over end-of-day-T wealth. The CB's objective is to maximize the price impact of its trades, conditional on the type of intervention chosen.

Within each day t there are three rounds of trading:

Round 1: Dealers trade with the central bank and public.

Round 2: Dealers trade among themselves to share risk.

Round 3: Dealers trade again with the public to share risk more broadly.

The timing of events within each day is shown in Figure 2.

Figure 5: Daily Timing

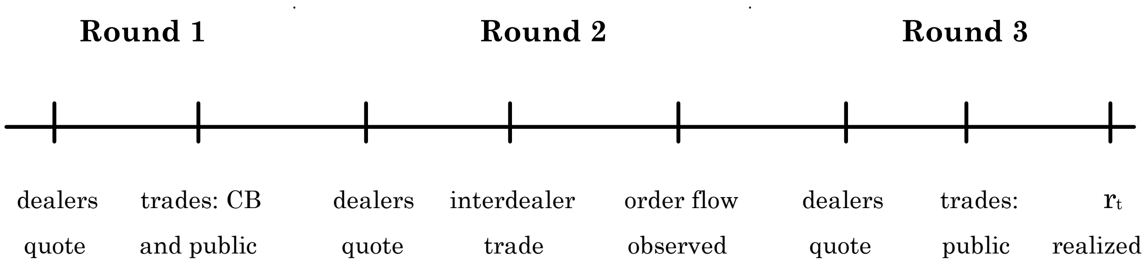

<u>Intervention Trades</u>

Each day, one of the dealers is selected at random to receive the order from the central bank. Let I_t, denote the intervention on day t, where $I_t<0$ denotes a CB sale (dealer purchase). The central-bank order arrives with the public orders at the end of round 1. The CB trade is distributed normally: $I_t \sim$ Normal $(0,\Sigma_I)$.[7] The CB trade I_t and the end-of-day interest increment r_t are correlated, with correlation coefficient ρ. The case of $\rho=0$ corresponds to sterilized intervention without signaling. (The CB trade I_t is uncorrelated with all other components of R.) An announced intervention is publicly observable. For unannounced interventions, only the dealer who receives the CB trade observes it.

[7] For now, we treat the CB trade as exogenous, with expected value zero. This might be viewed as a normalization around an "expected" intervention trade.

Trading Round 1

At the beginning of each day t, each dealer simultaneously and independently quotes a scalar price to the public and central bank.[8] We denote this round-1 price of dealer i as P_{i1}. (To ease the notational burden, we suppress the period subscript t when clarity permits.) This price is conditioned on all information available to dealer i.

Each dealer then receives a net "customer-order" from the public, c_{i1}, that is executed at his quoted price P_{i1}, where $c_{i1}<0$ denotes a net customer sale (dealer i purchase). Each of these N customer-order realizations is distributed normally, and they are independent across dealers: $c_{i1} \sim$ Normal $(0,\Sigma_{c1})$. We assume that customer orders are orthogonal to the interest increment r_t (though this is not hard to relax). These orders represent public portfolio shifts. Their realizations are not publicly observed. At the time the customer orders are received, one dealer also receives the intervention trade.

Trading Round 2

Round 2 is the interdealer trading round. Each dealer simultaneously and independently quotes a scalar price to other dealers at which he agrees to buy and sell (any amount). These interdealer quotes are observable and available to all dealers in the market. Each dealer then simultaneously and independently trades on other dealers' quotes. Orders at a given price are split evenly between dealers quoting that price.

Let T_{i2} denote the (net) interdealer trade initiated by dealer i in round two. At the close of round 2, all agents observe a noisy signal of interdealer order flow from that period:

(1) $$x = \left(\sum_{i=1}^{N} T_{i2} \right) + e \quad ,$$

where the noise term e is distributed Normal $(0,\Sigma_\varepsilon)$.

The model's difference in transparency across trade types corresponds well to institutional reality. Customer-dealer trades in major foreign-exchange markets are not generally observable, whereas interdealer trades do generate (noisy) signals of order flow than can be observed publicly.[9]

---

[8] Introducing a bid-offer spread (or price schedule) in round one to endogenize the number of dealers is a straightforward—but distracting—extension of our model. The simultaneous-move nature of the model is in the spirit of simultaneous-move games more generally (versus sequential-move games).
[9] The screens of interdealer brokers (such as EBS) are an important source of these interdealer order-flow signals.

Trading Round 3

In round 3, dealers share overnight risk with the non-dealer public. Unlike round 1, the public's trading in round 3 is non-stochastic and purely speculative. Initially, each dealer simultaneously and independently quotes a scalar price P_{i3} at which he agrees to buy and sell any amount. These quotes are observable and available to the public at large.

The mass of customers on the interval [0,1] is large (in a convergence sense) relative to the N dealers. This implies that the dealers' capacity for bearing overnight risk is small relative to the public's capacity. Dealers will optimally set prices such that the public willingly absorbs dealer inventory imbalances, and each dealer ends the day with no net position. These round-3 prices are conditioned on the round-2 interdealer order flow x. The interdealer order flow informs dealers of the total size of the position that the public needs to absorb to bring the dealers back to a position of zero.

Knowing the total size of the position the public needs to absorb is not sufficient for determining round-3 prices. Dealers also need to know the risk-bearing capacity of the public. We assume it is less than infinite. (In a multiple risky-asset setting, this would be equivalent to assuming that foreign-currency assets and domestic-currency assets are not perfect substitutes, in much the same spirit as traditional portfolio balance models.) Specifically, given negative exponential utility, the public's total demand for the risky asset in round-3, denoted c_3, is proportional to its expected return conditional on public information:

$$c_3 = \gamma(E[P_{3,t+1} \,|\, \Omega_3] - P_{3,t}), \tag{2}$$

where the positive coefficient γ captures the aggregate risk-bearing capacity of the public ($\gamma = \infty$ is infinitely elastic demand), and Ω_3 is the public information available at the time of trading in round 3.

Equilibrium

The dealer's problem is defined over four choice variables, the three scalar quotes P_{i1}, P_{i2}, and P_{i3}, and the dealer's interdealer trade T_{i2} (the latter being a component of x, the observed interdealer order flow). The appendix provides details of the model's solution. Here we provide some intuition.

Consider the three quotes P_{i1}, P_{i2}, and P_{i3}. No arbitrage ensures that all dealers quote a common price at any given time: because quotes are available to all dealers, any difference would provide an arbitrage opportunity. Hereafter, we write P_1, P_2, and P_3 in lieu of P_{i1}, P_{i2}, and P_{i3}. It must also be the case that if all dealers quote a common price, that price must be conditioned on common information only. If the CB intervention I_t is announced publicly at the time of trading in round 1, the common price quoted by dealers for period-2 trading, P_2, will reflect that information. Otherwise, order flow is not observed until the end of round 2. The price for round-3 trading, P_3, reflects the information in both observed order flow and any announcements of intervention.

<u>Model Solution</u>

Table 1 shows the equilibrium trading and pricing rules for each type of intervention we consider.

Table 1: Model Solution

No Signaling		**Signaling**	
Unannounced	**Announced**	**Unannounced**	**Announced**
Case 1	Case 2	Case 3	Case 4
$P_{2t} - P_{1t} = 0$	$P_{2t} - P_{1t} = \delta_2 I_t$	$P_{2t} - P_{1t} = 0$	$P_{2t} - P_{1t} = \delta_4 I_t$
$P_{3t} - P_{2t} = \lambda_1 x_t$	$P_{3t} - P_{2t} = \lambda_2 x_t - \delta_2 I_t$	$P_{3t} - P_{2t} = \lambda_3 x_t$	$P_{3t} - P_{2t} = \lambda_4 x_t$
$P_{1t+1} - P_{3t} = r_t$	$P_{1t+1} - P_{3t} = r_t$	$P_{1t+1} - P_{3t} = r_t - \delta_3 x_t$	$P_{1t+1} - P_{3t} = r_t - \delta_4 I_t$
$T_{j2t} = \alpha_1(c_{j1t})$	$T_{j2t} = \alpha_2(c_{j1t})$	$T_{j2t} = \alpha_3(c_{j1t})$	$T_{j2t} = \alpha_4(c_{j1t})$
$T_{i2t} = \alpha_1(c_{i1t}+ I_t)$	$T_{i2t} = \alpha_2(c_{i1t}+ I_t)$	$T_{i2t} = \alpha_3(c_{i1t}) + \beta_3 I_t$	$T_{i2t} = \alpha_4(c_{i1t}+ I_t)$

Note: We use the subscript j to denote all dealers not receiving the CB order, and subscript i for the one dealer who does receive the CB order. The appendix provides the details of the model solution.

To understand the intuition behind these sets of pricing and trading rules, consider Case 1 where intervention is unannounced and coveys no signal. The price change from round 1 to round 2 in that case is zero because no addition public information is observed from round-1 trading (customer trades are never publicly observable, and in this case the CB trade is not publicly observable either—it's unannounced). The change in price from round 2 to round 3, $\lambda_1 x_t$, is driven by public observation of the interdealer order flow x_t. The value $\lambda_1 x_t$ reflects the price adjustment required to induce the public to absorb the total portfolio shift from round 1, which equals $\Sigma_i c_{i1}+I_t$. (The CB's portfolio shift—in the form of I_t—must be absorbed by the public as well.) The parameter λ_1 insures that, at the round-3 price P_{3t}, $c_3+\Sigma_i(c_{i1}+b_{i1}) =0$. The price change from round 3 to round 1 of the following day is r_t in this case: this is the only new public information, and price should reflect it one-for-one (given that the final interest payoff R is the sum of the increments r_t).

The trading strategies in Case 1 are based on the private information that each dealer has, namely, his customer order (c_{j1t}) if he did not receive the intervention trade, and ($c_{i1t}+ I_t$) if he did. The parameter α_1 is the same in the two trading rules because intervention is indistinguishable from public order flow in the case of sterilized intervention that conveys no signal.

The intuition for Case 2 is broadly similar. One key difference is that the announcement of the intervention in round 1 causes price to adjust between rounds 1 and 2. The logic is the same as the logic in Case 1 for the $\lambda_1 x_t$ effect on price: it is a pure portfolio balance effect, in this case coming wholly from the CB's portfolio shift. Given the adjustment in price between rounds 1 and 2 in this case, the portfolio-shift information in x_t needs to be corrected so that only the incremental information is added to price. Like in Case 1, the parameter α_2 is the same in the two trading rules because intervention is indistinguishable from public order flow in the case of sterilized intervention that conveys no signal.

In Cases 3 and 4 we have signaling, so information about order flow is now correlated with the realization of r_t. In the case of unannounced intervention, Case 3, it is the interdealer order flow that conveys this incremental information about r_t. The parameter λ_3 is unambiguously greater than λ_1, its counterpart under Case 1. In the case of announced intervention, Case 4, it is the announcement itself that conveys the information. In both cases, the end-of-day realization of r_t does not convey as much incremental information is it did without signaling.

Several aspects of these results are worth emphasizing. First, in all four cases observed interdealer order flow conveys information about the aggregate of customer and CB orders received by dealers. As such, it informs dealers about the portfolio shift that must be absorbed by the public. Observed order flow moves prices because dealers recognize that the public's demand is less that perfectly elastic. If demand were perfectly elastic, order flow would still covey information about the portfolio shift, but it would not move prices. Thus, imperfect substitutability between different currency-denominated assets is a necessary condition for order flow to move prices. Second, unannounced interventions can only affect prices if they contributed to observed interdealer order flow. And, even then, the interventions will only affect prices if the public demand is imperfectly elastic. Third, interventions in the form of unannounced trades conveying no signal – a common form undertaken by the U.S. Federal Reserve[10] – have the same effect on prices as a customer order of the same size. This can be seen explicitly if we combine (1) with the trading rules from Case 1 to give

$$x_t = a_1\left(\sum_{i=1}^{N} c_{i,t} + I_t\right) + e_t \,.$$

[10] One might argue with this on the grounds that the Fed's intervention—which is sterilized as a matter of course—does sometimes have signal content. This is an open empirical question. Announced intervention (by our definition here) is quite rare.

Observed order flow depends only on the sum of customer and CB orders. In this case customer trades and CB trade are indistinguishable in terms of their price impact. A fact that we shall exploit in the empirical analysis that follows.

3. Empirical Analysis

Our empirical analysis aims to answer three questions. (i) Is there evidence of imperfect substitutability between different currency-denominated assets? (ii) Can we quantify both the magnitude and duration of an intervention's price-impact? (iii) Is the price-impact of intervention state-dependent? In other words, can we identify market conditions under which CB interventions will be particularly effective or ineffective?

To address these questions, we make use of intraday data on transactions prices and interdealer order flow for the spot DM/$. These data (described in detail below) do not contain direct information on CB trades. However, as our model demonstrates, they do contain a rich vein of information about the effects of unannounced no-signal interventions because these CB trades are theoretically indistinguishable from customer trades in terms of their price-impact. Because our data span a period during which there are a truly vast number of customer trades, this indirect approach provides us with the statistical power necessary to answer the questions above.

<u>Data</u>

Our dataset contains time-stamped, tick-by-tick data on actual transactions for the largest spot market – DM/$ – over a four-month period, May 1 to August 31, 1996. These data are the same as those used by Evans (1997), and we refer readers to that paper for additional detail. The data were collected from the Reuters Dealing 2000-1 system via an electronic feed customized for the purpose. Dealing 2000-1 is the most widely used electronic dealing system. According to Reuters, over 90 percent of the world's direct interdealer transactions take place through the system.[11] All trades on this system take the form of bilateral electronic conversations. The conversation is initiated when a dealer uses the system to call another dealer to request a quote. Users 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 bilateral conversations. This record is the source of our data. (Reuters was unable to provide the identity of the trading partners for confidentiality reasons.)

---

[11] Interdealer transactions account for about 75 percent of total trading in major spot markets. This 75 percent from interdealer trading breaks into two transaction types—direct and brokered. Direct trading accounts for about 60 percent of interdealer trade and brokered trading accounts for about 40 percent. For more detail on the Reuters Dealing 2000-1 System see Lyons (1995) and Evans (1997).

For every trade executed on D2000-1, our data set includes a time-stamped record of the transaction price and a bought/sold indicator. The bought/sold indicator allows us to sign trades for measuring order flow. This is a major advantage: we do not have to use the noisy algorithms used elsewhere in the literature for signing trades. One drawback is that it is not possible to identify the size of individual transactions. For model estimation, order flow x is therefore measured as the difference between the number of buyer-initiated trades and the number of seller-initiated trades.

Three features of the data are especially 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 1998, and Bjonnes and Rime 1998). Second, individual dealers do not observe our marketwide transactions data. Though dealers have access to their own transaction records, they cannot observe others' transactions on the system. Our data therefore represent activity that, at the time, participants could only infer indirectly. This is one of those rare situations where the researcher has more information than market participants themselves (at least in this dimension). Third, our data cover a relatively long time span in comparison with other micro data sets—four months.

The variables in our model are all measured hourly. The change in the spot rate (DM/$), Δp_t, is the log change in the purchase transaction price. When a purchase transaction does not occur precisely at the hour, we use a linear interpolation between the bracketing transaction prices (with roughly 1 million transactions per day, the bracketing transactions are generally within a few seconds of the hour). Order flow, x_t, is the difference between the *number* of buyer-initiated trades and the number of seller-initiated trades (in hundred thousands, negative sign denotes net dollar sales). We use two variables to measure the state of the market: trading intensity, n_t, measured by the number of trades during hour t; and price dispersion, s_t measured by the standard deviation of all the log transactions prices during hour t.

Empirical Models

To address the questions posed above, we need an empirical model of the relationship between transactions price and interdealer order flow. We develop the model in three steps. First, we consider linear models to establish the basic relationships between the various variables. Next, we study the relationship between prices and order flow in the context of a nonparametric (kernel) regression model. Finally, we take the results from this analysis to develop a parametric nonlinear model that well-describes the behavior of prices and order flow.

We begin by considering linear models of the form:

(3) $$\Delta p_t = b_1 x_t + b_2 \Delta p_{t-1} + b_3 s_{t-1} + b_4 n_{t-1} + h_t .$$

All variables are measured at the hourly frequency. The dependent variable Δp_t is the change in the log spot rate (DM/$), which we shall simply refer to as the exchange-rate return. The regressors include x_t, the interdealer order flow, measured contemporaneously with Δp_t (negative for net dollar sales), n_t, the number of transactions during hour t and s_t the standard deviation of all the log transactions prices during hour t.

Table 2 reports the OLS estimates of (3) together with t-statistics corrected for conditional heteroskedasticity. As the table shows, contemporaneous order flow is extremely significant, with t-statistics in excess of 10 in every model. Model I includes only contemporaneous order flow. In this case, the strong serial correlation and heteroskedasticity are suggestive of misspecification. Model II—a simple AR1 in exchange-rate returns—has very little explanatory power in terms of R^2, but the coefficient on lagged return is significant, and serial correlation is no longer significant. The negative value of that coefficient implies high-frequency return reversals.[12] Model III shows that when both contemporaneous order flow and lagged return are included, the point estimates and t-statistics all increase (in absolute value), as does the fit of the model as measured by R^2. The model is now accounting for 20 percent of the variation in hourly returns, versus only 2 percent when contemporaneous order flow is excluded. Models IV and V show that once we control for contemporaneous order flow and lagged returns, lagged order flow, lagged trading intensity, and lagged price dispersion are all insignificant.

Although these results point to order flow having a significant price-impact, none of these linear specifications appear to adequately capture the relationship. The right hand columns of the table show that residual serial correlation is present in all the models containing contemporaneous order flow, despite the striking significance of that variable. As serial correlation remains a problem in models that add one of two further lags of order flow and returns (unreported to conserve space), this diagnostic seems to point to the presence of a nonlinear relationship.

To explore this possibility, we consider nonparametric regressions of the form:

(4) $$\Delta p_t = m(z_t) + h_t ,$$

where $m(.)$ is an arbitrary fixed but unknown nonlinear function of the variables in the vector z_t, and h_t is a mean zero i.i.d. error. We estimate the $m(.)$ function by kernel

[12] The negative serial correlation in returns does not come from the bid-ask spread as is often the case in high-frequency price series. Returns are calculated from the sequence of DM prices dealers paid for Dollars and, as such, each is equal to the ask quote they received from another dealer immediately prior to the transaction taking place.

Table 2: Linear Models

$$\Delta p_t = b_1 x_t + b_2 \Delta p_{t-1} + b_3 s_{t-1} + b_4 n_{t-1} + h_t$$

Model	x_t	Δp_{t-1}	x_{t-1}	n_{t-1}	s_{t-1}	Diagnostics R^2	Serial	Hetero
I	1.430					0.13	<0.01	<0.01
	(11.422)						<0.01	<0.01
II		-0.137				0.02	0.79	<0.01
		(2.883)					0.10	<0.01
III	1.742	-0.251				0.20	0.05	<0.01
	(13.588)	(6.023)					0.04	<0.01
IV	1.738	-0.254	0.036			0.20	0.02	<0.01
	(13.343)	(5.802)	(0.355)				0.02	<0.01
V	1.742	-0.253		-0.001	-0.001	0.20	0.03	<0.01
	(13.578)	(6.317)		(0.662)	(0.004)		0.03	<0.01

The dependent variable Δp_t is the hourly change in the log spot exchange rate (DM/$). The regressor x_t is the hourly interdealer order flow, measured contemporaneously with Δp_t (negative for net dollar sales, in hundred thousands). s_t is the standard deviation of all the log transactions prices during hour *t*. n_t is the number of transactions during hour t, in hundreds. T-statistics are shown in parentheses and are calculated with standard errors corrected for the presence of heteroskedasticity. The sample spans four months (May 1 to August 31, 1996), which is 89 trading days. The Serial column presents the p-values of LM tests for first-order (top row) and sixth-order (bottom row) serial correlation in the residuals. The Hetero column presents the p-values of LM tests for first-order (top row) and sixth–order (bottom row) ARCH in the residuals.

regression utilizing the multivariate Gaussian kernel with a bandwidth parameter chosen by cross-validation (see the appendix for details). Conditions for consistency and asymptotic normality of kernel estimates in a time series context have been derived by Bierens (1983) and Robinson (1983) and appear to be satisfied here.

The upper panel of Table 3 reports kernel regression results based on three sets of conditioning variables. The pseudo R^2 statistics are larger in models that include trade intensity and price dispersion. While far form a formal test, this suggests that these "state-of -the-market" variables play some role despite their insignificance in the linear models.

Table 3: Nonparametric (Kernel) Regressions

$$\Delta p_t = m(z_t) + h_t, \quad \widehat{m}(z_t) = a' z_t + w_t$$

Model	Dependent Variable	z_t				Diagnostics	
		x_t	Δp_{t-1}	n_{t-1}	s_{t-1}	R^2	Serial
I	Δp_t	yes	Yes			0.174	0.034
							0.014
II	Δp_t	yes	Yes	Yes		0.195	<0.001.
							<0.001
II	Δp_t	yes	Yes	Yes	yes	0.219	<0.001
							<0.001
I	$\widehat{m}(z_t)$	3.433	-2.094			0.698	0.108
		(22.541)	(6.264)				0.604
II	$\widehat{m}(z_t)$	3.696	-1.746	-0.170		0.673	<0.001
		(21.514)	(11.688)	(1.722)			0.003
III	$\widehat{m}(z_t)$	3.487	-2.054	0.145	-0.154	0.569	0.003
		(19.187)	(6.421)	(0.786)	(1.059)		0.121

See Table 2 for the definitions of the variables. $m(z_t)$ is the kernel regression function where z_t is the vector of conditioning variables. T-statistics are shown in parentheses and are calculated with standard errors corrected for the presence of heteroskedasticity. The Serial column presents the p-values of LM tests for first-order (top row) and sixth-order (bottom row) serial correlation in the residuals.

The lower panel of Table 3 reports the results from regressing the fitted kernel estimates, $\hat{m}(z_t)$, on the standardized values of z_t (i.e., each element of z_t is divided by its sample standard deviation). These regressions summarize the relationship between the value of the $m(.)$ function and the conditioning variables. Consistent with the linear models, order flow has a strong positive effect and lagged returns a strong negative effect on $m(.)$, while the "state-of-the-market" variables remain insignificant.

To examine the form of nonlinearity in $\hat{m}(z_t)$, we calculated the partial derivatives of the estimated kernel for model III in Table 3 with respect to x_t and Δp_t evaluated at the values of x_t and Δp_t for every observation in the sample. The resulting series of estimated partial derivatives, $\hat{m}_x(z_t)$ and $\hat{m}_p(z_t)$, were then used as the dependent variables in a series of regressions.

The upper panel of Table 4 reports the results for the order flow derivative. The first two rows show results for $\hat{m}_x(z_t)$ regressed on a constant and the conditioning variables measured in deviations from their respective sample means. The constant is positive and highly significant in both cases. Thus, the nonparametric estimates strongly support the hypothesis that order flow affects prices under average market conditions. The estimates also show that the price-impact of order flow varies significantly with trading intensity during the previous hour. Interestingly, there is no evidence that the price-impact of order flow varies with the size of order flow or lagged returns. The coefficients on both variables are insignificantly different from zero. The next two lines report estimates from regressions that include dummy variables to capture time-of-the-day effects. These estimates indicate that the price-impact of order flow is greatest during European trading (i.e., 600 – 1800 hrs. British Summer Time, BST).

The lower panel of Table 4 report results with $\hat{m}_p(z_t)$ as the dependent variable. As above, the estimated constants are consistent with the results from the linear models: returns are significantly negatively related to lagged returns implying the presence of return reversals. Moreover, the strength of these reversals appears to depend on both "state-of-the-market" variables. The coefficients on trading intensity and price dispersion during the previous hour are both statistically significant even when the time-of-day dummies are included in the regressions.

The results above indicate that (i) order flow has a significant price-impact that varies with trading intensity, and (ii) lagged returns exert an influence on current returns that varies with past trading intensity and price dispersion. To explore these features further, we now consider a parametric nonlinear model.

Table 4: Kernel Regression Tests For State-Dependency

$$\hat{m}_j(z_t) = \delta z_t + u_t$$

										Diagnostics	
	00:01 06:00	06:01 12:00	12:01 18:00	18:01 24:00	Const.	x_t	Δp_{t-1}	n_{t-1}	s_{t-1}	R^2	Serial
$j = x$					2.140	0.250	0.241			0.015	<0.001
					(26.525)	(1.296)	(1.419)				
					2.140	0.262	0.260	0.889	0.191	0.081	0.161
					(27.463)	(1.339)	(1.325)	(6.829)	(1.072)		
	1.049	2.931*	3.018	0.822*						0.099	0.392
	(15.414)	(20.927)	(14.672)	(9.365)							
	1.238	2.840*	2.827	1.015*				0.270		0.103	0.483
	(8.873)	(20.715)	(15.421)	(6.169)				(1.553)			
$j = p$					-1.894	0.260	-0.373			0.016	0.360
					(24.994)	(1.156)	(2.718)				
					-1.894	0.273	-0.471	0.283	-0.500	0.038	0.363
					(25.277)	(1.220)	(3.059)	(2.335)	(4.099)		
	-1.739	-1.839*	-2.073	-1.865*						0.002	0.533
	(33.992)	(10.033)	(11.464)	(27.757)							
	-1.383	-2.123*	-2.374	-1.436*				0.574	-0.494	0.026	0.494
	(11.632)	(11.353)	(14.716)	(11.266)				(3.504)	(3.271)		

See table 2 for variable definitions. The dependent variables $\hat{m}_x(z_t)$ and $\hat{m}_p(z_t)$ are the derivatives of the estimated kernel $m(z_t)$ with respect to x_t and Δp_t respectively. The regressions estimated using OLS. The regressors, x_t, Δp_{t-1}, n_{t-1}, and σ_{t-1} are measured in deviation form. T-statistics are shown in parentheses and are calculated with standard errors corrected for the presence of heteroskedasticity. The Serial column presents the p-value of a chi-squared LM test for first-order serial correlation in the residuals. A "*" indicates that the estimated time dummy is statistically different from the value of the dummy in the earlier time period at the 5% significance level. The time intervals defining the dummies are British Summer Time (BST).

The model we estimate is a version of the smooth transition regression (STR) model due to Terasvirta (1994) and Granger and Terasvirta (1993) and incorporates features (i) and (ii) noted above. The form of the model is

(5) $$\Delta p_t = f_x(Y_{t-1})x_t + f_p(Y_{t-1})\Delta p_{t-1} + w_t,$$

where w_t is an i.i.d. mean zero random variable. The transition functions $f_x(.)$ and $f_p(.)$ are given by

$$f_z(Y_{t-1}) = a_z + b_z\left(1 - \exp\left(-(Y_{t-1} - m_z)\Omega_z(Y_{t-1} - m_z)'\right)\right),$$

where $Y_{t-1} = [n_{t-1}, s_{t-1}]$. The parameters of transition function are α_z, β_z and the elements of the vector m_z and the symmetric matrix Ω_z. The model in (5) allows the relation between returns, order flow and lagged returns to vary with Y_{t-1} via the transition functions which are bounded between α_z and $\alpha_z + \beta_z$. Because the price-impact of order flow appears to vary with trade intensity alone, we set the lower element of m_x and all but the upper left hand element of Ω_x equal to zero when estimating the model.

Table 5 reports model estimates obtained by nonlinear least squares using the full data sample, and a sub-sample that only includes observations between 6:00 and 18:00 hrs. BST. The majority of the parameter estimates in both models are statistically significant. The estimates imply similar bounds on $f_x(.)$ which measures the price-impact of order flow. The upper bound of approximately 0.2 is higher that the estimate of 0.17 obtained from the linear models and just below the estimate of 0.21 implied by the kernel regressions. The estimated lower bound is approximately 0.05 and is statistically insignificant. The estimated bounds for $f_p(.)$ are somewhat different across the two models, with the full sample estimates appearing far more negative than those found in the sub–sample. The next rows report tests for the null hypothesis of no state-dependency, or equivalently, $\beta_z = 0$. Testing this hypothesis is complicated by the fact that μ_z and Ω_z are not identified under the null. To circumvent this problem, we employ the LM-type test developed by Escribano and Jorda (1999) which is based on a Taylor series expansion of the transition function under the null (see appendix for details). As the table shows, the null hypothesis is strongly rejected for both transition functions in each model. Model diagnostics are shown at the bottom of the table. While there is evidence of residual serial correlation and heteroskedasticity in the full sample estimates, the sub-sample estimates appear to be relatively free of misspecification. We shall therefore focus our attention on these estimates.

Table 5: STR Models of State-Dependency

$$\Delta p_t = f_x(Y_{t-1})x_t + f_p(Y_{t-1})\Delta p_{t-1} + w_t$$
$$f_z(Y_{t-1}) = a_z + b_z\left(1-\exp\left(-(Y_{t-1}-m_z)\Omega_z(Y_{t-1}-m_z)'\right)\right)$$
$$m_z = [m_{z1} m_{z2}], \Omega_z = [w_{zi,j}]$$

	Full Sample		Sub-Sample	
a_x	0.209	(14.579)	0.205	(13.058)
b_x	-0.157	(-2.298)	-0.164	(-2.414)
m_{x1}	3.308	(8.757)	2.984	(8.376)
w_{x11}	0.103	(1.208)	0.120	(1.246)
a_p	-0.925	(-4.131)	-0.591	(-4.856)
b_p	0.708	(3.173)	0.522	(4.238)
m_{p1}	0.782	(19.075)	0.821	(9.001)
m_{p2}	2.541	(15.462)	2.587	(14.502)
w_{p11}	58.259	(1.565)	4.466	(1.481)
$w_{p12} = w_{p21}$	-9.057	(-1.249)	-0.763	(-0.818)
w_{p22}	3.205	(1.666)	0.997	(1.840)
Range for $f_x(.)$	0.209	0.052	0.205	0.041
Range for $f_p(.)$	-0.217	-0.925	-0.069	-0.591
State-Dependency Tests	Statistic	p-value	Statistic	p-value
$f_x(.)$	27.489	<0.001	19.007	0.001
$f_p(.)$	46.682	<0.001	16.589	0.002
R^2	0.211		0.243	
Serial	0.042		0.748	
Hetro.	<0.001		0.156	

NLLS parameter estimates with t-statistics shown in parentheses. The sub-sample contains hourly observations from 600 to 1800 hrs. BST. State-Dependency tests are LM tests for the null hypothesis that the transition function is constant. The p-values for LM tests for residual first-order serial correlation and residual first-order ARCH are reported in the rows marked Serial and Hetro.

4. Discussion

What are the implications of these empirical results for the effectiveness and efficacy of intervention? First, and foremost, our results indicate that order flow has a significant price-impact under normal conditions. Recall from the theoretical model that this only happens when the public's demand for foreign currency is less than perfectly elastic. Our data provide indirect but strong evidence of imperfect substitutability, which is a necessary condition for (sterilized) unannounced no-signal intervention to move prices.

Our results also provide a guide to the size of an intervention's price-impact. According to the sub-sample estimates in Table 5, the upper bound on $f_x(.)$ is approximately 0.2, which implies that 1000 more Dollar purchases than sales increases the DM price of the Dollar by 2 percent. Given a trade size of \$3.9m (the average in our data), \$100m of net Dollar purchases would therefore increase the DM price of Dollars by approximately 5 basis points.[13]

This finding needs to be interpreted with some care. In principle, a \$100m trade by the central bank could move the exchange rate by much less or much more than 5 basis points. In the extreme case, the dealer trading with the central bank could decide to keep all the resulting inventory, so that intervention would have no impact on interdealer order flow at all. In effect, the dealer would be taking a speculative position against the central bank. This isn't an optimal trading strategy for the risk-averse dealers in our model (the α_i's are always positive) and we doubt that it would be more generally. Moreover, the empirical evidence in Lyons (1995) indicates that dealers unwind inventory imbalances rather quickly. It seems very unlikely, therefore, that interdealer order flow could be unaffected by an intervention trade.

Alternatively, a \$100m intervention could have a bigger effect on interdealer order flow if the inventory imbalance caused by the central bank's trade is quickly passed from dealer to dealer. This is not possible in our theoretical model because there is only one round of interdealer trading. However, in a richer model with multiple trading rounds, dealers will trade among themselves until the inventory imbalance is spread across the market or can be absorbed through trade with the public. This process could create an interdealer order flow several times the size of the initiating central bank trade as the inventory imbalance is passed from one dealer to another, a phenomenon termed "hot potato" (Lyons 1997). Thus, to the extent that central bank intervention generates "hot potato" trading, the effects of a

[13] This figure is lower that the estimate of 7 basis reported in Evans and Lyons (1999) for the price-impact of order flow in daily data. The difference stems from the fact that order flow is positively autocorrelated at the hourly frequency so that swings in order flow persistently move prices from one hour to the next. See Evans (1999) for more on the dynamics of order flow.

\$100m intervention could be far greater than the 5 basis points estimate based on our order flow data. For example, suppose dealers past on 90 percent of the inventory imbalance. The total interdealer order flow generated by a \$100m Dollar purchase by the Fed. would then be (1/(1-0.9)) × \$100m. = \$1bn. and the DM price of Dollars would rise by 0.5 percent. With this perspective, the 5 basis point estimate seems a rather conservative lower bound. It is less than the estimate of 8 basis points per \$100m reported by Dominguez and Frankel (1993).

Our estimates of the STR model also provide information on the duration of the price-impact of order flow. In particular, we can calculate the long-run impact of order flow as

$$g_x(Y_{t-1}) = \frac{f_x(Y_{t-1})}{1 - f_p(Y_{t-1})}.$$

Using our sub sample estimates, the upper bound on $g_x(.)$ is 0.191 with a standard error of 0.016, and the lower bound is 0.026 with a standard error of 0.043. On the basis of these estimates, order flow appears to have a permanent price-impact in some states of the market. Referring to our discussion in Section 1, this finding points to an operative persistent portfolio-balance channel.

Next, we turn to the question of state-dependence. The results from the kernel regressions in Table 4 indicate that the price-impact of order flow is on average significantly larger during European trading. Evans (1997) documents a strong "seasonal pattern" in trade intensity across the 24 hour day with most trading occurring during European trading hours. With the perspective of Easley and O'Hara's event-uncertainty model, this seasonal pattern in the price-impact of order-flow is consistent with the presence of a signaling channel.

Further evidence on state-dependency is provided by Figure 6, which plots the transition functions from the STR model estimated with the sub-sample of our data. Panel I plots $f_x(.)$, the transition function for order flow, against a standardize measure of trade intensity, $(n_t - \bar{n})/s_n$; where $\bar{n}$ and s_n are the sample average and standard deviation of the number of trades per hour. Here we see that the price-impact of order flow rises and falls as trade intensity increases, with the peek impact occurring when trade intensity is approximately half one standard deviation above the sample mean. This pattern is consistent with both the event-uncertainty and hot-potato views of trading. As trading intensity rises from a low level, dealers view order flow as being increasingly informative about a market-wide event, such as a portfolio shift, so that the price-impact increases. However, beyond a certain intensity point, dealers view trades as being increasing driven by other factors, like the hot-potato passing of unwanted inventory, so that the information about market-wide conditions conveyed by order flow falls. Notice also that the peek price-impact occurs in “normal” states of the market (i.e., when trading intensity is close to its sample average) so our estimate of 5 basis points per $100m of order flow applies under these conditions.

Figure 6: Transition Function Estimates

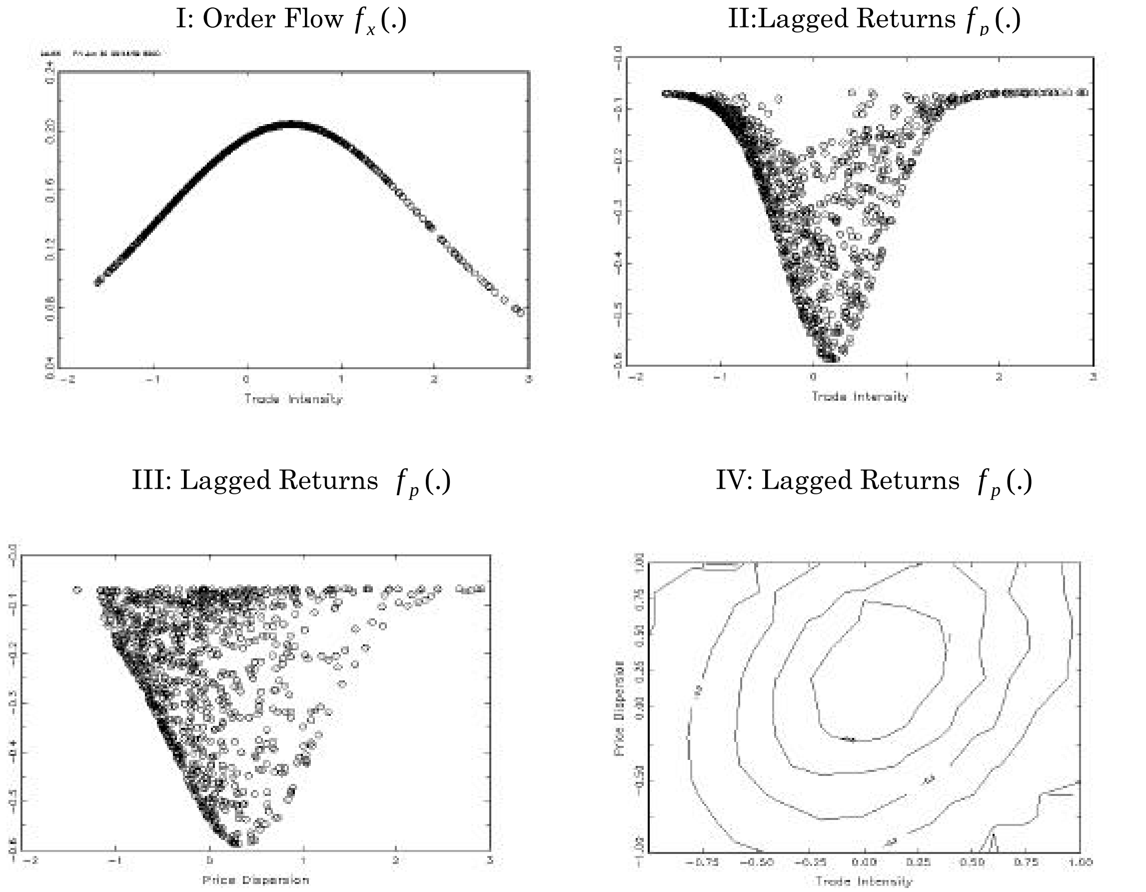

Panels II-IV show how the transition function on lagged returns, $f_p(.)$, varies with trade intensity and price dispersion. In panel II we see that, unconditioned on price dispersion, $f_p(.)$ is most negative when trading intensity is close to the sample average. Under these conditions, the persistent price-impact of order flow is approximately 1/(1+0.6) = 0.62 times the immediate impact. When trading intensity is unusually high or low, $f_p(.)$ is close to zero. Under these conditions, both the immediate and persistent price-impact of order flow is small. Price dispersion has a similar impact on the transition function. Panel III shows $f_p(.)$ plotted against a standardize measure of price dispersion, $(s_t - \bar{s})/s_s$;

where $\bar{s}$ and s_s are the sample average and standard deviation of price dispersion over the hour. Here we see that $f_p(.)$ is most negative when price dispersion is close to its sample average, and close to zero when price dispersions is unusually high or low. Panel IV shows the contour map of $f_p(.)$. The transition function takes the form of a distorted bowl when mapped against the two "state-of-the-market" variables.

In view of its empirical significance, the price dispersion measure deserves some further comment. Recall that price dispersion is identified by the standard deviation of log transaction prices over an hour. As such it contains both a time-series and cross-section dimension. The former comes from the fact that an individual dealer typically engages in many transactions per hour at different prices. If we could identify this price sequence, we could construct measures of price volatility for each dealer. The cross-section dimension arises from the fact that trade takes place simultaneously between pairs of dealers. This means, for example, that two dealers may be purchasing Dollars at different DM prices at the same time. Because transaction prices are only known to the dealers involved in the trade, the existence of price differences does not constitute an exploitable arbitrage opportunity. Rather it is a reflection of the heterogeneous information and inventory positions held by dealers at a point in time. It is hard to say how much each dimension contributes to the standard deviation of prices calculated over an hour. However, our data contains a very large number of instances where multiple dollar purchases (sales) occur during the same *second* with different transaction prices (see Evans 1997 for details). This suggests that our price dispersion measure contains a significant cross-section component. And, as a consequence, high price dispersion states are almost certainly associated with a higher than usual degree of dealer heterogeneity.

The STR model estimates may also shed some light on the issue of whether intervention can help maintain "an orderly market". Theory provides little guidance along these lines because—without market failure of some kind—it is unclear why markets should be any less orderly than required by market efficiency. Nevertheless, central banks do cite the maintenance of orderly markets as a distinct objective when articulating their policies. As an empirical matter, then, it is useful to consider whether intervention trades might affect exchange rates differently depending on whether the market is "orderly" or "disorderly." Figure 6 can help in this regard if we take the pragmatic approach of identifying a disorderly market as central bankers appear to: a state of high price dispersion that cannot be attributed to identifiable public news. Under this definition, Figure 6 shows that all the price–impact of order flow will be immediate because the $f_p(.)$ transition will be close to zero. If we go further to identify a disorderly market with both high trade intensity and price dispersion, the immediate price-impact of order flow will be very small. Thus, Figure 6 provides no support for the idea that intervention can calm a disorderly market. On the contrary, our model estimates suggest that the circumstances creating the disorderly market will make interventions particular ineffective in moving prices.[14]

5. Conclusion

In essence, this paper is a response to the following question: What do recently available transactions data and related theoretical innovations teach us about foreign exchange intervention? Several empirical papers have made progress along these lines (e.g., Bossaerts and Hillion 1991, Goodhart and Hesse 1993, Chang and Taylor 1998, Dominguez 1999, and Naranjo and Nimalendran 2000). What distinguishes our paper from this earlier work is that our data are not limited to high-frequency prices, but also include, crucially, high-frequency order flow. Order flow is central to the relevant theory, and measuring it is important for determining how trades impact prices.

Our analysis has several implications. At the broadest level, one implication is that it redirects attention within the intervention literature to order flow, a variable that plays little role in past work. This realignment of perspective is applicable in many contexts, including many we have not addressed explicitly. For example, how does the central bank know when current fundamentals are consistent with a fixed peg? Central banks take their cues from the micro level: *fundamentals are consistent with the peg when there are no longer significant imbalances between buyers and sellers in the market.* Real-time monitoring of movements in fundamentals is actually done at the microstructural level, even by central banks. This point is not obvious to someone who believes that the only source of asymmetric information relevant to intervention is the asymmetry between the central bank and the collective public (Type-1 asymmetric information).

More concretely, we present four main results, two that bear directly on theory and two that bear primarily on policy. On the theory side, we find evidence of imperfect substitutability, a necessary condition for the efficacy of the portfolio channel. This is important—there is something of a consensus in the profession that portfolio-balance effects from intervention are too small to be readily detectable. This has led some to dismiss portfolio-balance theory as irrelevant. Our approach harnesses additional statistical power that allows us to detect the effects. Portfolio-balance theory may be more relevant than was thought.

Also relevant for theory is our second finding of an operative signaling channel. In this case the implication for theory is different, however, because the signaling channel we detect is not the same signaling channel stressed in the literature. We find evidence

[14] There is no statistical evidence that order flow is related to trade intensity or price dispersion so it seems unlikely that intervention trades can affect the state-of-the-market directly.

consistent with signaling of Type-2 asymmetric information, whereas previous work focuses on signaling information of Type 1. Distinguishing these two types makes room for channels of intervention effects that are largely unrecognized. Further development of theory in this direction seems warranted.

The third of our main results bears primarily on policy: we establish a lower bound on the unconditional price impact of major-currency intervention trades of 5 basis points per $100 million. This responsiveness of price to order flow is measured in the existing literature only crudely. Our estimates bring substantial new statistical power. More work is certainly required before this (semi) elasticity can claim the status of "stylized fact." But there is little doubt that its size is important to exchange-rate economics.[15]

Our fourth result is also of direct relevance to policy-makers: we provide explicit guidance for ensuring that intervention trades have maximal price impact (which we associate with efficacy). This guidance includes both time-dependent rules and state-dependent rules. Our time-dependent rules indicate that order flow has the largest positive price impact during European trading (600-1800 hrs. BST). During our sample, New York begins significant trading at about 1500 hrs. BST (8 am EST), so this interval includes the first three hours of the New York trading day (which is, incidentally, the interval within which most Fed. intervention occurs). Our state-dependent rules indicate that order flow has the largest positive price impact when trading intensity is slightly above normal. The price-impact under unusually low or high trading intensity is negligible. Strikingly, our estimates indicate that attempts to stabilize an especially turbulent market are likely to be ineffective.

The existing empirical literature on intervention is vast, but as yet provides no direct analysis of how and why quantities affect prices. This paper provides that analysis. The source of our statistical power is our maintained hypothesis that private trades are indistinguishable from intervention trades. While we justified this hypothesis on both theoretical and empirical grounds, we cannot point to any one piece of evidence that establishes the validity of our maintained hypothesis in a strict sense. Thus, in the end, what we learn from this indirect approach is a matter of judgment. Because detailed quantity data are only recently available, our knowledge of how and why quantities affect prices is woefully underdeveloped. Given this state of our knowledge, we believe our indirect approach teaches us quite a lot.

Finally, we offer some thoughts on where this trading-theoretic approach to intervention is headed. The type of market data that are now coming available allow tracking—very precise tracking—of how the market absorbs actual CB trades, and any information in

[15] An immediate example of this fact's value is its ability to help us understand why portfolio-balance effects from sterilized intervention are so hard to detect: as we noted above, even a $1-billion intervention is likely to move the exchange rate by less than one percent.

them. CB's with precise knowledge of their own trades—announcements, timing, stealth level, etc.—will be able to estimate the impact of these various "parameter" settings. Consider, for example, the type of data used by Payne (1999), which provides the order book of an electronic interdealer broker. A CB with these data, over a sufficiently large number of intervention trades, could learn exactly how the "book" is affected, including liquidity provision on both sides, transaction activity, and the process of price adjustment. It is something like a doctor who has a patient ingest blue dye to determine how it passes through the system—throat, stomach, intestines, liver, bladder, and bowel. The whole process becomes transparent. Such is the future of empirical work on this topic.

References

Almekinders, G., 1995, *Foreign exchange intervention*, London: Edward Elgar.

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

Bagwell, L., 1992, Dutch auction repurchases: An analysis of shareholder heterogeneity, *Journal of Finance*, 47: 71-105.

Bhattacharya, U., and P. Weller, 1997, The advantage of hiding one's hand: Speculation and foreign exchange intervention in the foreign exchange market, Journal of Monetary Economics, 39: 251-277.

Bierens H.J., 1983, Uniform consistency of kernel estimators of a regression function under general conditions, *Journal of the American Statistical Association*, 77 699-707.

Bjonnes, G., and D. Rime, 1998, FX trading ... live: Impact of new trading environments, typescript, Norwegian School of Management, December.

Bossaerts, P. and Hillion, P. 1991, Market microstructure effects of government intervention in the foreign exchange market, *Review of Financial Studies* 4, 513-541.

Chaboud, A., and B. LeBaron, 1999, Foreign exchange market trading volume and Federal Reserve intervention, typescript, Brandeis University, July.

Chang, Y., and S. Taylor, 1998, Intraday effects of foreign exchange intervention by the Bank of Japan, *Journal of International Money and Finance*, 18: 191-210.

Cheung, Y., and M. Chinn, 1998, Macroeconomic implications of the beliefs and behavior of foreign exchange traders, typescript, U.C. Santa Cruz, November.

Cheung, Y., and C. Wong, 1998, A survey of market practitioners' views 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.

Dominguez, K., 1998, Central bank intervention and exchange rate volatility, *Journal of International Money and Finance*, 18:161-190.

Dominguez, K., 1999, The market microstructure of central bank intervention, University of Michigan typescript, August.

Dominguez, K., and J. Frankel, 1993a, Does foreign-exchange intervention matter? The portfolio effect, *American Economic Review*, 83: 1356-1369.

Dominguez, K., and J. Frankel, 1993b, *Does foreign-exchange intervention work?* Institute for International Economics, Washington, D.C.

Edison, H., 1993, The effectiveness of central-bank intervention: A survey of the literature after 1982, Special papers in international economics, No. 18, Princeton University.

Escribano and Jorda (1999) Improved Testing and Specification of Smooth Transition Regression Models, in N*onlinear Time Series Analysis of Economic and Financial Data*, Rothman (Ed.) Kluwer Academic Press.

Evans, M., 1997, The microstructure of foreign exchange dynamics, Georgetown University typescript, November, presented at the NBER Market Microstructure Conference, December 4, 1997.

Evans, M., 1999, What are the origins of foreign exchange movements?, Georgetown University typescript, March.

Evans, M., and R. Lyons, 1999, Order flow and exchange rate dynamics, NBER working paper 7317, August.

Foster, D., and S. Viswanathan, 1990, A theory of interday variations in volumes, variances, and trading costs in securities markets, *Review of Financial Studies*, 3: 593-624.

Frankel, J., 1982, In search of the exchange risk premium: A six currency test assuming mean-variance optimization, *Journal of International Money and Finance*, 1: 255-274.

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

Goodhart, C., and T. Hesse, 1993, Central bank forex intervention assessed in continuous time, *Journal of International Money and Finance*, 12: 368-389.

Granger C.W.J., and T. Terasvirta, 1993, *Modelling Nonlinear Relationships*. New York: Oxford University Press.

Hardle W., 1990 *Applied Nonparametric Regression*, Econometric Society Monograph, Cambridge University Press.

Hansch, O., N. Naik, and S. Viswanathan 1998. "Do Inventories Matter in Dealership Markets? Evidence from the London Stock Exchange," *Journal of Finance*, 53: 1623-1656.

Hsieh, D., and A. Kleidon, 1996, Bid-ask spreads in foreign exchange markets: Implications for models of asymmetric information, in Jeffrey Frankel et al. eds.: *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago) 41-65.

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.

Kyle, A., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315-1335.

LeBaron, B., 1996, Technical trading rule profitability and foreign exchange intervention, NBER working paper 5505, March, forthcoming in the *Journal of International Economics*.

Luukkonen, R.,P. Saikkonen and T. Terasvirta, 1988, Testing linearity against smooth transition autoregressive models, *Biometrika*, 75 491-499.

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

Lyons, R. 1997, A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics* 42, 275 – 298.

Lyons, R., 2000, *The Microstructure Approach to Exchange Rates*, MIT Press (chapters at www.haas.berkeley.edu/~lyons/NewBook.html), forthcoming.

Montgomery, J., and H. Popper, 1998, Information sharing and central bank intervention in the foreign exchange market, typescript, Santa Clara University, December.

Naranjo, A., and M. Nimalendran, 2000, Government intervention and adverse selection costs in foreign exchange markets, *Review of Financial Studies*, 13, 453-477.

Payne, R. 1999, Informed trade in spot foreign exchange markets: An empirical investigation, typescript, London School of Economics, January.

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.

Reiss, P., and I. Werner, 1998, Does risk sharing motivate interdealer trading? *Journal of Finance*, 53: 1657-1704.

Rime, D., 2000, Private or public information in foreign exchange markets? Typescript, Department of Economics, University of Oslo, March.

Robinson P.M., 1983. Nonparametric Estimators for Time Series, *Journal of Time Series Analysis*, 4, 185-207.

Sarno, L., and M. Taylor, 2000, Official intervention in the foreign exchange market, typescript, Department of Economics, University of Oxford, March.

Scholes, M., 1972, The market for securities: Substitution versus price pressure and the effect of information on share price, *Journal of Business*, 45: 179-211.

Shleifer, A., 1986, Do demand curves for stocks slope down? *Journal of Finance*, 41: 579-590.

Terasvirta, T., 1994, Specification, estimation and evaluation of smooth transition autoregressive models, *Journal of the American Statistical Association*, 89, 208-218.

Vitale, P., 1997, Sterilized central bank intervention in the foreign exchange market, working paper, London School of Economics, forthcoming in the *Journal of International Economics*.

Vogler, K., 1997, Risk allocation and interdealer trading, *European Economic Review* 41, 417-441.

Yao, J. 1998. Market making in the interbank foreign exchange market, New York University Salomon Center Working Paper #S-98-3.

Model Solution Appendix (Incomplete)

Each dealer determines quotes and speculative demand by maximizing a negative exponential utility function defined over terminal wealth. Because returns are independent across periods, with an unchanging stochastic structure, the dealers' problem collapses to a series of independent trading problems, one for each period. Within a given period t, let $W_{i\tau}$ denote the end-of-round τ wealth of dealer i, where we use the convention that W_{i0} denotes wealth at the end of period t-1. (To ease the notational burden, we suppress the period subscript t when clarity permits.) With this notation, and normalizing the gross return on the riskless asset to one, we can write the dealers' problem as:

$$\text{(A1)} \qquad \underset{\{P_{i1},\, P_{i2},\, P_{i3},\, T_{i2}\}}{Max} \qquad E\left[-\exp(-qW_{i3} \mid \Omega_i)\right]$$

$$\text{s.t.} \qquad W_{i3} = W_{i0} + (c_{i1} + I)(P_{i1} - P'_{i2}) + \left(D_{i2} + E[T'_{i2} \mid \Omega_{i2}]\right)\left(P_{i3} - P'_{i2}\right) - T'_{i2}\left(P_{i3} - P_{i2}\right)$$

$P_{i\tau}$ is dealer i's round-τ quote and a $'$ denotes an interdealer quote or trade received by dealer i. The dealers' problem is defined over four choice variables: the three scalar quotes P_{i1}, P_{i2}, and P_{i3}, and the dealer's outgoing interdealer trade in round 2, T_{i2}. The last three terms in W_{i3} capture capital gains/losses from round-1 orders from the public and CB (c_{i1}+I), round-2 speculative demand D_{i2}, and the round-2 position disturbance from incoming interdealer orders T'_{i2}. (Recall that b_{i1} is equal to zero for all but one of the dealers.) The dealer's own outgoing interdealer trade in round 2 has three components:

$$\text{(A2)} \qquad T_{i2} = D_{i2} + (c_{i1} + I) + E[T'_{i2} \mid \Omega_{Ti2}]$$

The first component D_{i2} is dealer i's speculative demand. The last two components arise from hedging: to arrive at his desired speculative demand, the dealer must offset the effect on his position of orders from the public and CB (c_{i1}+I), and must also hedge against incoming orders from other dealers $E[T'_{i2} \mid \Omega_{Ti2}]$. (In equilibrium, the latter term $E[T'_{i2} \mid \Omega_{Ti2}]$ is equal to zero.) The conditioning information Ω_i at each decision node (3 quotes and 1 outgoing order) is summarized below.

$$\Omega_{Pi1} \equiv \left\{ {}_{k=1}^{t}\{r_k\},\ {}_{k=1}^{t-1}\{x_k\} \right\}$$
$$\Omega_{Pi2} \equiv \left\{ \Omega_{Pi1},\ c_{i1},\ b_{i1} \right\}$$
$$\Omega_{Ti2} \equiv \left\{ \Omega_{Pi2} \right\}$$
$$\Omega_{Pi3} \equiv \left\{ \Omega_{Pi2},\ \Delta x_t \right\}$$

Conditional Variances

This appendix repeatedly uses several conditional return variances. These variances do not depend on conditioning variables' realizations (e.g., they do not depend on realizations of c_{i1} and I). These conditional variances are therefore known in period one. (It is a convenient property of the normal distribution that realizations of conditioning variables affect the conditional mean but not the precision of the condition mean.) This predetermination of conditional variances is key to the derivation of optimal quoting and trading rules.

Equilibrium

The equilibrium concept we use is Bayesian-Nash Equilibrium, or BNE. Under BNE, Bayes rule is used to update beliefs and strategies are sequentially rational given those beliefs.

Empirical Appendix

.

Kernel Regressions

We consider kernel regressions of the form

$$\Delta p_t = m(z_t) + h_t$$

where $m(.)$ is an arbitrary fixed but unknown nonlinear function of the variables in the vector z_t, and h_t is a mean zero i.i.d. error. An estimate of the $m(.)$ function is estimated by

$$\hat{m}(z_t) = \frac{\sum_{j=0, j\neq t}^{T} K_h(z_t - z_j)\Delta p_j}{\sum_{j=0, j\neq t}^{T} K_h(z_t - z_j)},$$

where $K_h(u)=h^{-1}K(u/h)$ with $K(x) \geq 0$ and $\int K(u)du = 1$. In this application, we use the multivariate Gaussian kernel $K(j) = (2p)^{-d/2}\exp(-j'j/2)$ where $d = \dim(j)$. The bandwidth parameter, h, is chosen by cross-validation. That is to say, h minimize

$$\frac{1}{T}\sum_t^T (\Delta p_t - m(z_t))^2 w_t ,$$

where w_t is a weighting function that cuts off 5% of the data at each end of the data interval as in Hardle (1990, p. 162). For the regressions in Table 2, z_t, contains $\{x_t, Dp_{t-1}, n_{t-1}, s_{t-1}\}$. We follow the common practice of including the standardize value of each of these variables in the Gaussian kernel (i.e., each element of z_t is divided by its sample standard deviation).

Asymptotic theory for kernel regressions in the time series context appears in Bierens (1983) and Robinson (1983). Robinson shows that consistency and asymptotic normality of the estimator can be established when the data satisfy α-mixing with mixing coefficients α(k) that obey the condition $T\sum_t^{\infty} a(k)^{1-2/d} = O(1)$ and $E\,|\,\Delta p_t\,|^{d} < \infty, \delta > 2$.

<u>Testing for State Dependency in STR Models</u>

Our model has the form of an exponential STR model that may be written as

$$y_t = pw_t + qw_t F(z_{t-d}, g, c) + u_t$$

where y_t is scalar, $w_t = (1, y_{t-1}.x_t)' = (1, \tilde{w}_t)$ and

$$F(z_{t-d}, \Omega, c) = 1 - \exp\{-(z_{t-d} - c)\Omega(z_{t-d} - c)'\}\,.$$

Testing linearity against STR-type nonlinearity implies testing the null hypothesis $q = 0$. However, the elements of Ω and c are not identified under the null. Luukkonen et al (1988a) proposed a solution to this problem in which *F(.)* is replaced by a suitable Taylor series approximation around Ω=0. Under the null, the LM test is shown to have the standard asymptotic distribution. Escribano and Jorda (1999) propose computing the LM tests by an auxiliary regression of the form

$$y_t = d_0 + dx_t + b_1\tilde{w}_t z_{t-d} + b_2\tilde{w}_t (z_{t-d})^2 + b_3\tilde{w}_t (z_{t-d})^3 + b_4\tilde{w}_t (z_{t-d})^4 + v_t$$

and testing the null of $\beta_1=\beta_2=\beta_3=\beta_4=0$ with an F-test. Table 5 report statistics computed in this manner where z_t is equal to trade intensity and price dispersion with d=1.