NBER WORKING PAPER SERIES

PORTFOLIO BALANCE, PRICE IMPACT, AND SECRET INTERVENTION

Martin D. D. Evans
Richard K. Lyons

Working Paper 8356
http://www.nber.org/papers/w8356

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2001

Portfolio Balance, Price Impact, and Secret Intervention
Martin D. D. Evans and Richard K. Lyons
NBER Working Paper No. 8356
July 2001
JEL No. F31, G12, G15

ABSTRACT

This paper tests the portfolio-balance approach to exchange rate determination in a new way. Past work on portfolio balance in foreign exchange falls into two groups: (1) tests using measures of asset supply and (2) tests using measures of central-bank asset demand. We address the demand side, but we use a broad measure of public demand, rather than focusing on demand by central banks. Under floating rates, changing public demand has no direct effect on interest rates, current or future. This provides an opportunity to test for portfolio-balance effects on price. We develop and estimate a micro portfolio-balance model that has both Walrasian and microstructure features. Portfolio-balance effects are clearly present: the immediate price impact of public trades is 0.44 percent per $1 billion (of which, about 80 percent persists indefinitely). This estimate is applicable to central-bank trades as well, as long as they are sterilized, secret, and provide no monetary-policy signal. Intervention of this type is most effective when the flow of macroeconomic news is strong.

Martin D. D. Evans
Department of Economics
Georgetown University
Washington, DC 20057
and NBER

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

Portfolio Balance, Price Impact, and Secret Intervention

Portfolio-balance models occupy an important place within exchange-rate economics.[1] They are still used, for example, as a basis for why central-bank intervention can be effective even when sterilized. Empirical evidence supporting these models is scant, however.[2] This paper tests the portfolio-balance model in a new, more powerful way and finds it strongly supported.

Work on portfolio balance in foreign exchange falls into two groups: (1) tests using measures of asset supply and (2) tests using measures of central-bank asset demand. Here, we address the demand side, but we examine demand by the public broadly, rather than focusing on demand by central banks. Under floating rates, changing public demand has no direct effect on interest rates, current or future. This provides an opportunity to test for portfolio-balance effects on price. Because data on public trades became available only recently (due to the advent of electronic trading), this strategy is feasible for the first time.

The discriminating power of our approach arises from avoiding difficulties inherent in past approaches. The asset-supply approach, for example, has low power because measuring supplies is notoriously difficult. First, one must determine which measure of supply is the most appropriate. (There is considerable debate in the literature about this issue; see, e.g., Golub 1989.) Then, for any given measure, the consistency of data across countries is a concern. Finally, these data are available only at lower frequencies (e.g., quarterly or monthly) and are rather slow-moving, making it difficult to separate the effects of changing supply from the many other forces moving exchange rates.

The central-bank-demand approach—an "event study" approach—may also have limited statistical power because central-bank trades in major markets are relatively few and are small relative to public trading. For example, the average US intervention reported by Dominguez and Frankel (1993b) is only $200 million, or roughly 1/1000[th] of the daily spot volume in either of the two largest markets. (Since then, US intervention has been larger, typically in the $300 million to $1.5 billion range, but market volume has been higher too; see Edison 1998.) Studies using this

---

[1] For theoretical work, see Kouri (1976), Branson (1977), Girton and Henderson (1977), Allen and Kenen (1980), and the survey by Branson and Henderson (1985).

[2] See Branson et al. (1977), Frankel (1982), Backus (1984), Lewis (1988), and the survey by Lewis (1995), among many others.

latter approach are more successful in finding portfolio-balance effects (e.g., Loopesko 1984, Dominguez 1990, Dominguez and Frankel 1993a). Nevertheless, results using this approach are not exclusively positive (e.g., Rogoff 1984) and the extent these event studies pertain to portfolio effects more broadly is not clear.

The "micro portfolio-balance model" we develop embeds both Walrasian features (as in the traditional portfolio-balance approach) and features more familiar to models from microstructure finance. Regarding the latter, the model clarifies the role played by order flow in conveying information about shifts in traders' asset demands.[3] Beyond this clarification, two analytical results in particular are important guideposts for our empirical analysis: (i) Order flow's effect on price is persistent as long as public demand for foreign currency is less than perfectly elastic (even when beliefs about future interest rates are held constant); and (ii) When central bank trades are sterilized, conducted secretly, and convey no policy signal, then the price impact of these trades is indistinguishable from that of public trades.[4] This latter result links our analysis directly to intervention operations of this type.

We establish three main results. First, testable implications of our model are borne out: we find strong evidence of price effects from portfolio balance. Thus, the portfolio-balance approach—with its rich past, but lack of recent attention—appears to warrant some fresh consideration. Second, we provide a precise estimate of the immediate price impact of trades: 0.44 percent per $1 billion (of which, about 80 percent persists indefinitely). Our third result speaks to intervention policy. (As noted above, our price-impact estimate is applicable to central-bank trades as long as they are sterilized, secret, and provide no signal.) Estimates suggest that central-bank intervention of this type is most effective at times when the flow of macroeconomic news is strong.

The remainder of the paper is in five sections. Section 1 introduces our trading-theoretic approach to measuring price impact. Section 2 presents our micro portfolio-balance model. Section 3 describes the data. Section 4 presents model estimates and discusses their implications (e.g., for central-bank intervention). Section 5 concludes.

---

[3] Order flow is not synonymous with trading volume. Order flow is a concept from microstructure finance that refers to *signed* volume. Trades can be signed in these microstructure models depending on whether the "aggressor" is buying or selling. (The dealer posting the quote is the passive side of the trade.) For example, a sale of 10 units by a trader acting on a dealer's quotes is order flow of −10. (Order flow is undefined in rational-expectations models of trading because all transactions are symmetric in that setting—an aggressor cannot be identified.)

[4] For more on how central banks intervene secretly, see, e.g., Hung (1997) and Dominguez and Frankel (1993b). Along the spectrum of intervention transparency, from secret to announced (in every detail), we are considering here only the secret end.

1. A Trading-Theoretic Approach to Portfolio Balance

This section links the traditional macroeconomic approach to exchange rates to microeconomic theories of asset trading. This is useful for two main reasons. First, theories of asset trading provide greater resolution on how trades 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 establishes that most channels through which trades—including intervention—affect price involve information asymmetry. Impounding dispersed information in price is an important function of the trading process (which our model is designed to capture).

Within macroeconomics, central bank (CB) currency demand affects price through two channels: imperfect substitutability and asymmetric information. Distinct modeling approaches are used to examine these two channels. For the first channel, imperfect substitutability, macro analysis is based on the portfolio-balance approach. Models within this approach are most useful for analyzing intervention that is sterilized and conveys no information (signal) about future monetary policy. Macro analysis of the second channel, asymmetric information, is based on the monetary approach. These models are most useful for analyzing intervention that conveys information about current policy (unsterilized intervention) or future policy (sterilized intervention with signaling). This channel captures the CB's superior information about its own policy intentions. Let us examine these two macro channels within a trading-theoretic approach.

Imperfect Substitutability

In contrast to macro models, which address imperfect substitutability at the marketwide level only, theories of asset trade address imperfect substitutability at two 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 adjustment is temporary because this risk premium is not necessary once positions are shared with the wider market. In trading-theoretic models, price effects from this channel are termed "inventory effects." These effects dissipate quickly in most markets because full risk sharing occurs rapidly (e.g., within a day).[5]

---

[5] For theoretical work on inventory effects on price, see Amihud and Mendelson (1980), Ho and Stoll (1983), and Vogler (1997) among many others. Empirical work on inventory effects in FX include Lyons (1995), Yao (1998), and Bjonnes and Rime (2000).

Within trading models, imperfect substitutability 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.[6] This too induces a risk premium, which elicits a price-level adjustment. Unlike price adjustment at the first level of imperfect substitutability, price adjustment at this second level is persistent (because risk is fully shared at this level). This is precisely the price adjustment that macro models refer to as a portfolio-balance effect.

The first of these two levels of imperfect substitutability is not present within the macro approach. Indeed, use of the term imperfect substitutability within that approach refers to the second level only. The logic among macroeconomists for addressing only the second level is that effects from the first level are presumed fleeting enough to be negligible at longer horizons. This is of course an empirical question—one that our trading-theoretic approach allows us to address in a rigorous way. Moreover, our modeling of this channel provides a more disciplined way to understand why part of intervention's effect on price is fleeting (and what determines the duration of this part of the effect).

Below we test empirically whether either or both of these two levels of imperfect substitutability are present. If the first level is present—the dealer level—then FX trades should have an impact on the exchange rate, but the effect should be temporary. We term this effect a "temporary portfolio-balance channel." If the second level is present—the marketwide level—then trades should have persistent impact. We term this effect a "persistent portfolio-balance channel."

Asymmetric Payoff Information

Theories of asset trading provide a third channel through which trades affect price—asymmetric payoff information (see., e.g., Kyle 1985 and Glosten and Milgrom 1985).[7] If trades convey future payoff information (sometimes referred to as "fundamentals" in exchange-rate economics), then they will have a second persistent effect on price beyond the persistent portfolio-balance effect noted above. (For example, in equity markets managers of firms have inside information about earnings, and their trades can convey this information.) Unsterilized interventions are an example of currency trades that convey payoff information (i.e., information

---

[6] Relevant theory includes rational expectations models like Grossman and Stiglitz (1980) and Kodres and Pritsker (1998). Though not a fully rational model, another recent paper that includes price effects from market-wide imperfect substitutability is Kyle and Xiong (2001). Empirical work on imperfect substitutability across stocks at the market-wide level includes Scholes (1972), Shleifer (1986), Bagwell (1992), and Kaul et al. (2000), among others.

[7] The word "payoffs" in trading-theoretic models refers to the cash flows that accrue to the security's holder (e.g., dividends in the case of a stock).

about current interest rates). Another example is sterilized intervention that signals future interest rate changes.

In foreign-exchange markets, however, trades by market participants other than central banks (the public) do not in general convey payoff information: under floating, public trades have no direct effect on interest rates. For these trades, then, the payoff-information channel is not operative. This presents an opportunity to use public trades to test for the presence of the two types of portfolio-balance effect.

2. A Micro Portfolio-Balance Model

The model is designed to show how the trading process reveals information contained in order flow. At a micro level, it is the flow of orders between dealers that is particularly important: public trades are not observable market-wide, but are subsequently reflected in interdealer trades, which *are* observed market-wide. Once observed, this information is impounded in price. This information is of two types, corresponding to the two portfolio-balance effects outlined in the previous section: information about temporary portfolio-balance effects and information about persistent portfolio-balance effects.

To understand these different portfolio-balance effects, consider the model's basic structure. At the beginning of each day, the public and central bank place orders in the foreign exchange market. (These orders are stochastic and are not publicly observed.) Initially, dealers take the other side of these trades—shifting their portfolios accordingly. To compensate the (risk-averse) dealers for the risk they bear, an intraday risk premium arises, producing a temporary portfolio-balance effect on price. The size of this price effect depends on the size of the realized order flow. This is the first of the two information types conveyed by order flow.

To understand the second, first note that at the end of each day, dealers pass intraday positions on to the public (consistent with empirical findings that dealers end their trading day with no position; see Lyons 1995 and Bjonnes and Rime 2000). Because the public's (non-stochastic) demand at the end of the day is not perfectly elastic—that is, different-currency assets are imperfect substitutes in the macro sense—beginning-of-day orders have portfolio-balance effects that persist beyond the day. Thus, the price impact of these risky positions is not diversified away even when they are shared marketwide.[8] The size of this price effect, too, is a function of

---

[8] 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 individual NYSE stocks in 1998 was about $9 million, whereas the average daily volume in DM/$ spot was about $300 billion).

the size of the beginning-of-day order flow. This is the second of the two information types conveyed by order flow.

<u>Specifics</u>

Consider an infinitely lived, pure-exchange economy with two assets, one riskless and one risky, the latter representing foreign exchange.[9] Each day, foreign exchange earns a payoff R, publicly observed, which is composed of a series of random increments:

(1)
$$R_t = \sum_{i=1}^{t} \Delta R_i$$

The increments ΔR are i.i.d. Normal$(0, \sigma_R^2)$. We interpret the increments as the flow of public macroeconomic information (e.g., interest rate changes).

The foreign exchange market has three participant types: dealers, customers, and a central bank. The N dealers are indexed by i. There is a continuum of customers (the public), indexed by $z \in [0,1]$. Dealers and customers all have identical negative exponential utility defined over periodic wealth. Central bank trades are described below.

Within each day t there are four rounds of trading:

Round 1: Dealers trade with the central bank and public.
Round 2: Dealers trade among themselves (to share inventory risk).
Round 3: R_t is realized and dealers trade among themselves a second time.
Round 4: Dealers trade again with the public (to share risk more broadly).

The timing of events within each day is shown in Figure 1, which also introduces some notation clarified below.

---

[9] The structure of the model is similar to that of Evans and Lyons (1999). Because the published version of Evans and Lyons (1999) does not include the model's derivation and specifics—being focused instead on empirical results—this paper now contains the original modeling contribution. Other papers that have adopted the model developed here include Rime (2000).

Figure 1

Daily Timing

Round 1		Round 2			Round 3				Round 4	
dealers quote	trades: public and CB	dealers quote	inter-dealer trade	order flow observed	dealers quote	payoff realized	inter-dealer trade	order flow observed	dealers quote	trades: public
P_1^i	C_1^i, I	P_2^i	T_2^i	X_2	P_3^i	ΔR	T_3^i	X_3	P_4^i	C_4^i
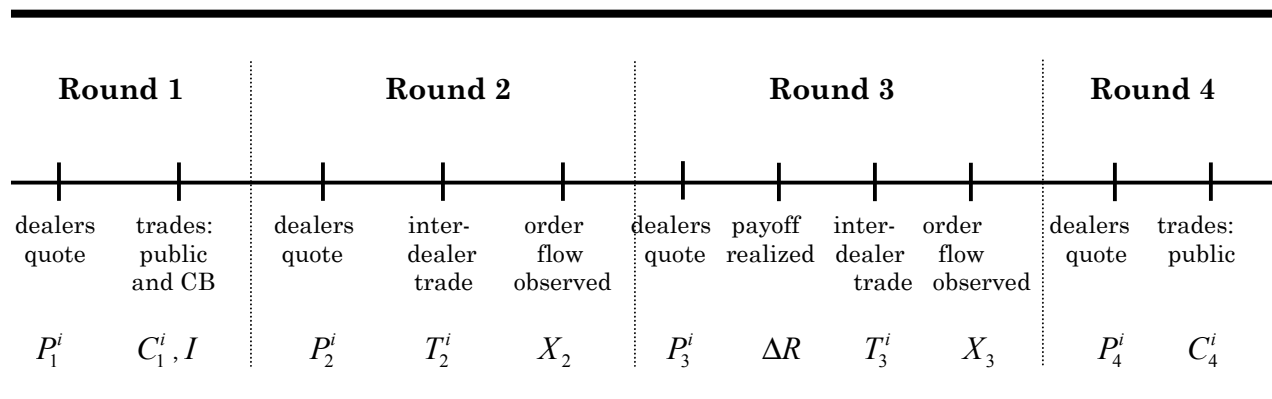
Central Bank Trades

To accommodate analysis of intervention, we include trades by a central bank. The intervention we consider is of a particular type, equivalent in its features to public trades: intervention that is sterilized, secret (anonymous and unannounced), and conveys no signal of future monetary policy.[10]

More specifically, each day, one dealer is selected at random to receive an order from the central bank. To maintain anonymity, the CB order is routed to the selected dealer via an agent. 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 $(0, \sigma_I^2)$.[11] Because the CB trade is sterilized and conveys no signal, I_t and the daily interest increments ΔR_t are uncorrelated (at all leads and lags). Secret intervention insures that only the dealer who receives the CB trade observes its size (though not its source). A CB trade is, under these circumstances, indistinguishable from other customer orders.[12]

---

[10] The model can be extended to accommodate all the main intervention types and the main channels through which intervention can be effective. (For example, sterilized versus unsterilized intervention can be modeled by admitting correlation between CB trades and the current periodic payoff R_t. A signaling channel can be introduced by admitting correlation between CB trades and future R_{t+1}.)

[11] We treat the CB trade as having an expected value of zero. This should be viewed as a normalization around an "expected" intervention trade, should any non-zero expectation exist.

[12] That intervention in practice often does not take this form in no way negates the fact that this intervention strategy is indeed available.

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.[13] We denote this round-1 price of dealer i as P_1^i. (We suppress unnecessary notation for day t; as we shall see, it is the within-day rounds—the subscripts—that capture the model's economics.) This price is conditioned on all information available to dealer i.

Each dealer then receives a net "customer-order" from the public, C_1^i, that is executed at his quoted price P_1^i, where $C_1^i < 0$ denotes a net customer sale (dealer i purchase). Each of these N customer-order realizations is distributed normally, C_1^i ~ Normal $(0, \sigma_C^2)$. They are uncorrelated across dealers and uncorrelated with the payoff R. These orders represent portfolio shifts by the public, coming, for example, from changing hedging demands, changing transactional demands, or changing risk preferences. 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 first of two interdealer trading rounds. Each dealer simultaneously and independently quotes a scalar price to other dealers at which he agrees to buy and sell (any amount), denoted P_2^i. 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_2^i 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:

$$(2) \qquad X_2 = \sum_{i=1}^{N} T_2^i + \nu$$

where $\nu \sim$ Normal$(0, \sigma_\nu^2)$, independently across days. The model's difference in transparency across trade types corresponds well to institutional reality: customer-dealer trades in major foreign-exchange markets (round 1) are not generally

---

[13] 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).

observable, whereas interdealer trades do generate signals of order flow than can be observed publicly.[14]

Trading Round 3

Round 3 is the second of the two interdealer trading rounds. At the outset of round 3 the payoff increment ΔR_t is realized and the daily payoff R_t is paid (both observable publicly). Like in round 2, each dealer then simultaneously and independently quotes a scalar price to other dealers at which he agrees to buy and sell (any amount), denoted P_3^i. 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_3^i denote the net interdealer trade initiated by dealer i in round 3. At the close of round 3, all agents observe interdealer order flow from that period:

$$(3) \qquad\qquad X_3 = \sum_{i=1}^{N} T_3^i$$

This specification with noiseless observation of round-three order flow captures dealer learning and the increasing precision of their order-flow beliefs.

Trading Round 4

In round 4, dealers share overnight risk with the non-dealer public. Unlike round 1, the public's trading in round 4 is non-stochastic. Initially, each dealer simultaneously and independently quotes a scalar price P_4^i at which he agrees to buy and sell any amount. These quotes are observable and available to the public.

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. With this assumption, dealers set prices optimally such that the public willingly absorbs dealer inventory imbalances, and each dealer ends the day with no net position (which is common practice among actual spot foreign-exchange dealers). These round-4 prices are conditioned on the interdealer order flow X_3, described in equation (3). We shall see that this interdealer order flow informs dealers of the size of the total position that the public needs to absorb to bring the dealers back to a position of zero.

---

[14] The screens of interdealer brokers (such as EBS) are an important source of these interdealer order-flow signals.

To determine the round-4 price—the price at which the public willingly absorbs the dealers' aggregate position—dealers need to know (1) the size of that aggregate position and (2) the risk-bearing capacity of the public. We assume the latter is less than infinite. Specifically, given negative exponential utility, the public's total demand for foreign exchange in round-4 of day t, denoted C_4, is proportional to the expected return on foreign exchange conditional on public information:

(4)
$$C_4 = \gamma\left(E[P_{4,t+1} + R_{t+1} \mid \Omega_{4,t}] - P_{4,t}\right)$$

where the positive coefficient γ captures the aggregate risk-bearing capacity of the public ($\gamma = \infty$ is infinitely elastic demand), and $\Omega_{4,t}$ includes all public information available for trading in round 4 of day t.

Equilibrium

The dealer's problem is defined over six choice variables, the four scalar quotes P_1^i, P_2^i, P_3^i, and P_4^i, and the two dealer's interdealer trades T_2^i and T_3^i. Appendix A provides details of the model's solution. Here we provide some intuition.

Consider the four quotes P_1^i, P_2^i, P_3^i, and P_4^i. No arbitrage ensures that at any given time all dealers will quote a common price: quotes are executable by multiple counterparties, so any difference across dealers would provide an arbitrage opportunity. Hereafter, we write P_1, P_2, P_3, and P_4 in lieu of P_1^i, P_2^i, P_3^i, and P_4^i. It must also be the case that if all dealers quote a common price, then that price must be conditioned on common information only. Common information arises at three points: at the end of round 2 (order flow X_2), at the beginning of round 3 (payoff R), and at the end of round 3 (order flow X_3). The price for round-4 trading, P_4, reflects the information in all three of these sources.

The following optimal quoting rules specify when the common-information variables (X_2, X_3, ΔR) are impounded in price. These quoting rules describe a linear, Bayes-Nash equilibrium.

Proposition 1: Dealers in our micro portfolio-balance model choose the following quoting rules, where the parameters λ_2, λ_3, δ, and ϕ are all positive:

$$P_2 - P_1 = 0$$

$$P_3 - P_2 = \lambda_2 X_2$$

$$P_4 - P_3 = \lambda_3 X_3 + \delta\Delta R - \phi\left(P_3 - P_2\right)$$

For intuition on these quoting rules, note that the price change from round 1 to round 2 is zero because no additional public information is observed from round-1 trading (neither customer trades nor the CB trade are publicly observed). The change in price from round 2 to round 3, $\lambda_2 X_2$, is driven by public observation of the interdealer order flow X_2. X_2 here serves as an information aggregator. Specifically, it aggregates dispersed information about privately observed trades of the public and CB. The value $\lambda_2 X_2$ is the price adjustment required for market clearing—it is a risk premium that induces dealers to absorb the round-1 flow from the public and CB (that round-one flow equaling $\Sigma_i C_1^i + I$). The price change from round 3 to round 4 includes both pieces of public information that arise in that interval: X_3 and ΔR. The second round of interdealer flow X_3 conveys additional information about round-1 flow (from the public and CB) because it does not include noise. The payoff increment ΔR will persist into the future, and therefore must be discounted into today's price. The third component of the price change from round 3 to 4 is dissipation of a temporary portfolio-balance effect that arose between rounds 2 and 3. Specifically, part of $\lambda_2 X_2$ is a temporary risk premium that arises to induce dealers to hold risky positions intraday. The end-of-day price P_4 does not include this because dealers hold no positions overnight.

The persistent portion of the portfolio-balance effect arises in this model because interdealer order flow informs dealers about the portfolio shift ($\Sigma_i C_1^i + I$) that must be absorbed at day's end by the public. If the end-of-day public demand were perfectly elastic, order flow would still convey information about the portfolio shift, but the shift would not affect the end-of-day price. This persistent portfolio-balance effect is the same in the model regardless of whether the initial order flow came from the public or the central bank. Thus, CB trades of the type we consider here have the same effect on price as a customer order of the same size—both induce the same portfolio shift at day's end by the public.

3. Empirical Analysis

3.1 Data

The dataset contains time-stamped, tick-by-tick observations 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.[15] 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.)

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 is therefore measured as the difference between the *number* of buyer-initiated and seller-initiated trades.[16]

The variables in our empirical model are measured hourly. We take the spot rate, as the last purchase-transaction price (DM/$) in hour h, P_h. (With roughly 1 million transactions per day, the last purchase transaction is generally within a few seconds of the end of the hour. Using purchase transactions eliminates bid-ask bounce.) Order flow, X_h, is the difference between the number of buyer- and seller-initiated trades (in hundred thousands, negative sign denotes net dollar sales) during hour h. We also make use of three further variables to measure the state of the market: trading intensity, N_h, measured by the gross number of trades during hour h; price dispersion, σ_h, measured by the standard deviation of all transactions prices during hour h, and the number of macroeconomic announcements, A_h. These announcements comprise all those reported over the Reuter's News service that relate to macroeconomic data for the U.S. or Germany. The source is Olsen Associates (Zurich); for details, see, e.g., Andersen and Bollerslev (1998).

Although trading can take place on the D2000-1 system 24 hours a day, 7 days a week, the vast majority of transactions in the DM/$ take place between 6 a.m. and 6 p.m. London time, Monday to Friday. Though the results we report below are based on this sub-sample, they are quite similar to results based on the 24-hour

---

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

[16] This is common in the literature; see, for example, Hasbrouck (1991). See also Jones et al. (1994) for analysis suggesting that trade size conveys no additional information (beyond that conveyed by the number of buys minus sells).

trading day (as noted below). This sub-sample still leaves us with vast number of trades, providing us with considerable power to test for effects from portfolio balance.

3.2 The Empirical Model

Our model is specified with each day split into 4 trading rounds. We now develop an empirical implementation for examining the model's implications in hourly data.

Let $\rho_j(y_h)$ denote the probability that the market will move from round j to $j+1$ between the end of hours h and $h+1$, when the state of the market at the end of hour h is y_h.[17] Given these transition probabilities, the probability that the market will be in round j at the end of hour h, $\pi_j(Y_{h-1})$, is defined recursively as

$$(5) \qquad \pi_j(Y_{h-1}) = \rho_{j-1}(y_{h-1})\pi_{j-1}(Y_{h-2}) + [1 - \rho_j(y_{h-1})]\pi_j(Y_{h-2})$$

where $Y_h = \{y_h, y_{h-1}, ...\}$ denotes current and past states of the market.

According to Proposition 1, prices change when the market moves from rounds 2 to 3, and from rounds 3 to 4. Let ΔP_h and ΔR_h respectively denote the change in price and the flow of macroeconomic information between the end of hours h-1 and h. With the aid of the probabilities $\rho_j(y_{h-1})$ and $\pi_j(Y_{h-1})$, we can derive the probability distribution of hourly price changes as shown in Table 1.

Table 1: Distribution of Hourly Price Changes

	ΔP_h: Hourly Price Change	Probability
I	0	$\theta_I(Y_{h-1})$
II	$\lambda_2 X_h$	$\theta_{II}(Y_{h-1})$
III	$\lambda_3 X_h - \phi \Delta P_{h-k} + \delta \Delta R_h$	$\theta_{III,k}(Y_{h-1})$

Lines II and III of Table 1 identify the price change associated with the market moving into round 3 and into round 4 between the end of hours h-1 and h, respectively. In the former case, the price change is proportional to order flow during the hour. In the latter, the price change depends on order flow and macroeconomic information during the hour, and a lagged price change ΔP_{h-k}, for $k > 0$. The length of the lag k equals the number of hours the market spends in round 3 trading before

---

[17] We assume that the market remains in a round for a minimum of one hour. Dropping this assumption complicates the calculations needed to find the distribution of hourly price changes, but it does not alter the basic structure of the empirical model; see Appendix B for details.

moving to round 4. The probabilities in the right-hand column are complicated functions of $\rho_j(y_{h-l})$ for $j = 1, 2, 3, 4$, and $l > 0$ and so depend on the past states of the market, $Y_{h-1} = \{y_{h-1}, y_{h-2}, ...\}$ (see appendix B for details). In the special case where the probability of moving from round 3 to round 4, $\rho_3(y_{h-1})$, equals one, k must also equal one, and the probabilities simplify to

$$
\begin{aligned}
\theta_I(Y_{h-1}) &= 1 - \theta_{II}(Y_{h-1}) - \theta_{III,1}(Y_{h-1}), \\
\theta_{II}(Y_{h-1}) &= \rho_2(y_{h-1})\pi_2(Y_{h-1}), \\
\theta_{III,1}(Y_{h-1}) &= \rho_2(y_{h-2})\pi_2(Y_{h-2}).
\end{aligned}
$$

Our empirical model is derived from the distribution of hourly price changes. Specifically, let $\Omega_h = \{X_h, Y_{h-1}, \Delta P_{h-1}, \Delta P_{h-2}, ...\}$ denote the information set spanned by current order flow, past states of the market, and past hourly price changes. The observed hourly price change can be written as

$$
(6) \qquad \Delta P_h = E[\Delta P_h \mid \Omega_h] + \eta_h,
$$

where η_h is the expectational error in hour h. Since the flow of macroeconomic information in hour h, ΔR_h, is orthogonal to Ω_h, this error includes ΔR_h. To complete the empirical model, we need the conditional expectation from the distribution of hourly price changes. For the special case noted above where $\rho_3(y_{h-1}) = 1$, this expectation is given by

$$
(7) \qquad E[\Delta P_h \mid \Omega_h] = \beta_1(Y_{h-1})X_h + \beta_2(Y_{h-1})\Delta P_{h-1}
$$

with $\beta_1(Y_{h-1}) = \lambda_2\theta_{II}(Y_{h-1}) + \lambda_3\theta_{III,1}(Y_{h-1})$ and $\beta_2(Y_{h-1}) = -\phi\theta_{III,1}(Y_{h-1})$. Hourly price-change dynamics can therefore be represented by

$$
(8) \qquad \Delta P_h = \beta_1(Y_{h-1})X_h + \beta_2(Y_{h-1})\Delta P_{h-1} + \eta_h
$$

In the more general case where $\rho_3(y_{h-1}) \leq 1$, the equation for price changes contains more than one lag of past price changes on the right-hand side (see appendix B for details). These lags are not statistically significant in our data. We therefore focus attention on equation (8), which takes the form of a regression with state-dependent coefficients.

4. Results and Implications

Estimation of our micro portfolio-balance model allows us to answer three key questions. First, is there support for portfolio-balance in the data? Though existing negative results have led to the view that portfolio-balance theory is moribund, past work may suffer from low power (as noted in the introduction). Second, do trades have both temporary and persistent portfolio-balance effects? Third, does the price-impact of trades depend on the state of the market? This last question is central to identifying states in which intervention is most effective.

4.1 Model Estimates

Our estimation strategy proceeds in two stages. First, we estimate a constant-coefficient version of equation (11) and test for state dependency in the coefficients. As we shall see, the coefficients in this model accord with portfolio-balance predictions in terms of sign and significance. The estimated coefficients also accord with our model in that they are indeed state dependent. This latter result motivates the second stage of our strategy, namely, estimation of the precise nature of this state dependency (using non-parametric kernel regressions).

Table 2 presents results from the first stage of our estimation: the constant-coefficient model. Both contemporaneous order flow X_h and lagged price change ΔP_{h-1}—the two core variables in our model—have the predicted signs and are significant. (Though constants do not arise in our derivation, for robustness we also estimate the model with constants; they are insignificant.) A coefficient on order flow X_h of 0.26 translates into price impact of about 0.44 percent per $1 billion.[18] (The magnitude is similar when we use log price change as the dependent variable, as can be seen in Table A1 of the appendix.) A coefficient on lagged price change of –0.2 implies that 1/1.2, or 83 percent of the impact effect of order flow persists indefinitely. Thus, we are finding evidence of both types of portfolio-effect noted in section I: the temporary portfolio-balance channel and the persistent portfolio-balance channel. Though the former is clearly present, the latter accounts for the lion's share of order flow's impact effect.

---

[18] Order flow X_h is measured as the net number of dollar purchases (in 10,000s). With an average trade size in the sample of $3.9 million, this implies that $39 billion of positive order flow raises price by 0.26 DM/$, which is 17 percent of the average spot rate of 1.5 DM/$. 17 percent divided by $39 billion is approximately 0.44 percent per $1 billion.

This estimate is slightly lower than that reported in Evans and Lyons (1999) for the price-impact of order flow in *daily* data. The difference stems from order flow being positively auto-correlated 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.) It is also slightly less than the 8 basis points per $100m reported by Dominguez and Frankel (1993b).

Rows II-VI of Table illustrate the sensitivity of the estimates to various departures from our derived specification. Note from row III, for example, that returns are not negatively autocorrelated unconditionally, as one might expect in an efficient market; only when order flow is included does the negative relationship emerge. Note too from rows V and VI that when both our core variables are included, lags of these variables do not enter significantly (which accords with the model). Row IV shows that lagged order flow proxies for lagged price change, albeit imperfectly, when the latter is excluded (as one would expect). The reduced explanatory power of the proxy may arise for two reasons: our order flow measure does not capture all order flow in the market (e.g., it does not include brokered interdealer flow) and the transparency of order flow in practice is lower than the transparency of price.

The bottom panel of Table 2 presents results from the state dependency tests. These tests address whether the two core coefficients are affected by variables describing the state of the market, which include time-of-day, the number of trades, the volatility of prices, and the number of macroeconomic announcements (the latter three measured over the previous hour). The results of these tests are clear: the coefficients are indeed state dependent.

To measure the effects of changing state variables, we turn to the second stage of our estimation: non-parametric kernel estimation (see Bierens 1983 and Robinson 1983). This is a flexible and intuitive means of estimating precisely how the state variables affect the two core coefficients. These regressions take the form:

$$\Delta P_h = m(z_h) + \eta_h$$

where $m(.)$ is an arbitrary fixed but unknown nonlinear function of the vector $z_h = [X_h \ \Delta P_{h-1} \ \sigma_{h-1} \ N_{h-1} \ A_{h-1}]$, which includes the two core variables and three state variables. The error η_h is i.i.d mean zero. Intuitively, this method estimates $m(.)$ by taking a weighted average of the ΔP's from other observations with similar z's. The weights are Normally distributed, attaching greater weight to the most similar z's. (See appendix B for a formal definition of the weighting procedure.) The derivatives of the estimated kernel $\hat{m}(z_h)$ with respect to our two core variables can then be regressed on the state variables, providing a precise measure of how they vary by state.[19]

---

[19] Including the state variables in estimating the kernel $m(z_h)$ means that we are not restricting the partial derivatives of the kernel with respect to these variables to equal zero. Imposing that restriction in estimating the kernel would have transformed our subsequent tests on the kernel derivatives into joint tests.

Table 2: Estimates of Micro Portfolio-Balance Model (constant coefficients)

$$\Delta P_h = \beta_1 X_h + \beta_2 \Delta P_{h-1} + \eta_h$$

	X_h	ΔP_{h-1}	X_{h-1}	ΔP_{h-2}	R²	Serial	Hetero
						Diagnostics	
I	0.258 (13.205)	-0.203 (4.103)			0.212	0.437 0.287	0.071 0.020
II	0.225 (12.297)				0.173	<0.001 <0.001	0.070 0.011
III		-0.061 (1.359)			0.003	0.150 0.311	0.271 0.016
IV	0.234 (12.083)		-0.041 (2.551)		0.180	<0.001 <0.001	0.067 0.009
V	0.258 (13.016)	-0.202 (3.857)	-0.001 (0.081)		0.212	0.205 0.282	0.071 0.020
VI	0.260 (13.213)	-0.200 (3.878)		0.044 (0.729)	0.220	0.823 0.381	0.086 0.028

Tests for State-Dependency
p-values

State variables	$\beta_1(.)$	$\beta_2(.)$
τ: time of day	<0.001	0.652
N_h: number of trades	0.002	0.051
σ_h: std. dev. of prices	<0.001	0.239
A_h: number of announcements	0.176	0.006

OLS estimates are based on hourly observations from 6:00 to 18:00 BTS from May 1 to August 31, 1996, excluding weekends. T-statistics are shown in parentheses and are calculated with asymptotic standard errors corrected for heteroskedasticity. ΔP_h is the hourly change in the spot exchange rate (DM/$). X_h is the hourly interdealer order flow, measured contemporaneously with ΔP_h (negative for net dollar sales, in thousands). The Serial column presents the p-value of a chi-squared LM test for first-order (top row) and fifth-order (bottom row) serial correlation in the residuals. The Hetero column presents the p-value of a chi-squared LM test for first-order (top row) and fifth–order (bottom row) ARCH in the residuals. The lower panel presents the p-values for a heteroskedasticity-consistent Wald test of the null hypothesis that $\beta_i(.)$ does not vary with the state variable shown in the left-hand column. σ_h is the standard deviation of all the transactions prices, N_h is the number of transactions, and A_h is the number of macroeconomic announcements, all during hour h. τ is a vector of three dummy variables, $[\tau_1 \ \tau_2 \ \tau_3]$. τ_1 equals one for hours between 6:00am and 7:59am, zero otherwise; τ_2 equals one for hours between 8:00am and 11:59am, zero otherwise; and τ_3 equals one for hours between 12:00pm and 1:59 pm, zero otherwise.

Table 3 presents these dependencies of the two core coefficients on movements in the state variables. The first panel presents the effects of state variables on the derivative of the kernel with respect to order flow X_h. For example, the positive and significant coefficient on the number of announcements A_{h-1} implies that the price impact of order flow is 0.024 higher per additional macro announcement (for perspective, recall that the unconditional estimate of that price-impact coefficient in Table 2 is 0.26). Note too from this first panel that there does not appear to be non-linearity in the order-flow/price-change relation (per the insignificant impact of X_h on the derivative of the estimated kernel with respect to X_h, reported in column three).

The effects of the state variables are more broadly present in the case of the lagged price-change coefficient (panel two of Table 3). Column four, for example, shows that there is some non-linearity in this case: the larger the lagged price change, the greater the negative auto-correlation in price (i.e., the greater the transitory portion of portfolio-balance effects). Columns five, six, and seven show additional state dependence, in these cases due to volatility, trading activity (numbers of trades), and announcements. The negative (and significant) coefficient on announcements in this case implies that the greater the number of announcements, the greater the negative auto-correlation in price.

The third panel of Table 3 shows the effects of the state variables on the volatility (absolute change) of the estimated residual from the kernel regressions. The most striking result in this panel is the strong intraday seasonality in this unexplained volatility. This is evident from the estimated coefficients on τ_h (the number of hours since midnight at the end of hour h) and τ_h^2 in columns eight and nine. These estimates indicate a U-shape in volatility over the trading day, consistent with findings elsewhere (e.g., Andersen and Bollerslev 1998). In addition, the second and third rows of panel three indicate that conditional heteroskedasticity is tied to the flow of announcements.

Table 3: Nonparametric (Kernel) Regressions

$$\Delta P_h = m(z_h) + \eta_h$$

Dependent Variable (1)	Const. (2)	z_h					τ_h (8)	τ_h^2 (9)	Diagnostics		
		X_h (3)	ΔP_{h-1} (4)	σ_{h-1} (5)	N_{h-1} (6)	A_{h-1} (7)			R^2	Serial	Hetero
$\hat{m}_x(z_h)$	0.134	<0.001	<0.001	-0.014	0.001	0.024			0.071	0.866	0.556
	(15.249)	(0.342)	(0.997)	(1.225)	(0.264)	(3.788)				0.231	0.315
	0.077			-0.013	-0.001	0.023	0.010	<0.001	0.071	0.943	0.531
	(1.989)			(0.977)	(0.261)	(3.558)	(1.493)	(1.366)		0.198	0.323
	0.136			-0.018	0.001	0.024			0.069	0.968	0.536
	(14.719)			(1.432)	(0.352)	(3.803)				0.225	0.313
$\hat{m}_{\Delta p}(z_h)$	-0.050	0.001	-0.002	-0.072	0.016	-0.028			0.066	0.116	<0.001
	(3.120)	(1.552)	(2.851)	(3.287)	(2.560)	(2.626)				0.230	<0.001
	-0.005			-0.055	0.018	-0.026	-0.007	<0.001	0.044	0.294	<0.001
	(0.072)			(1.915)	(2.490)	(2.340)	(0.545)	(0.201)		0.147	<0.001
	-0.061			-0.052	0.014	-0.027			0.038	0.279	<0.001
	(3.662)			(1.958)	(2.215)	(2.534)				0.195	<0.001
$\left|\hat{\eta}_h\right|$	4.488	-0.011	-0.027	13.780	-1.096	0.566			0.184	0.973	0.999
	(5.042)	(0.662)	(1.048)	(9.757)	(3.145)	(1.207)				0.222	0.217
	29.214			12.201	-0.297	0.946	-4.458	0.176	0.200	0.493	0.661
	(4.952)			(7.049)	(0.737)	(2.089)	(4.436)	(4.364)		0.209	0.232
	31.046			11.698		0.910	-4.807	0.189	0.200	0.414	0.590
	(5.815)			(7.442)		(2.073)	(5.326)	(5.126)		0.249	0.255

The table presents estimates from the kernel regression shown in the heading, where z_h is the vector of five conditioning variables and $\hat{m}(z_h)$ is the nonparametric estimate of $E[\Delta P_h \mid z_h]$. $\hat{m}_x(z_h)$ and $\hat{m}_{\Delta p}(z_h)$ are the derivatives of the estimated Kernel with respect to X_h and ΔP_{h-1} respectively. See Table 2 for other variable definitions. Estimated using OLS, with t-statistics shown in parentheses (calculated with standard errors corrected for heteroskedasticity). τ_h is the number of hours since midnight at the end of hour h. The Serial column presents the p-value of a chi-squared LM test for first-order (top row) and fifth-order (bottom row) residual serial correlation. The Hetero column presents the p-value of a chi-squared LM test for first-order (top row) and fifth–order (bottom row) ARCH in the residuals.

<u>Implications for Intervention</u>

What are the implications of these empirical results for the 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 happens only when public demand for foreign currency is less than perfectly elastic. The data provide solid evidence of imperfect substitutability, a necessary condition for the efficacy of intervention that is sterilized, secret, and conveys no policy signal.

Our results also provide a guide to the size of an intervention's price-impact. From Table 2, $1 billion of net dollar purchases increases the DM price of dollars by 0.44 percent, with about 80 percent of this persisting indefinitely. When linking this estimate to intervention, however, the finding needs to be interpreted with care: a dollar of net interdealer flow is not equivalent to a dollar of public flow. To take an extreme case, a dealer trading with a central bank could decide to retain the resulting inventory indefinitely, so that intervention has no impact whatsoever on interdealer order flow. However, this isn't an optimal strategy for the risk-averse dealers in our model, and we doubt it would be optimal more generally. Moreover, empirical evidence in Lyons (1995) and Bjonnes and Rime (2000) indicates that dealers unwind inventory positions rather quickly. Only under extreme and counterfactual assumptions would interdealer order flow be unaffected by an intervention trade.

It is also possible for intervention's effect on interdealer flow to be magnified. For example, suppose the inventory position caused by the CB's trade is quickly passed from dealer to dealer over a succession of trading rounds, a phenomenon market participants refer to as "hot potato" trading (Lyons 2000). This process could amplify the effects on interdealer order flow considerably, thereby amplifying the subsequent impact on price. With this possibility of hot potato trading, then, it is appropriate to consider the 0.44-percent-per-$1-billion estimate as a lower bound.

The model estimates may also shed light on 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."

Estimates in Table 3 suggest that the price impact of order flow is not significantly affected by variables commonly associated with disorder, namely, the

volatility of transaction prices σ_h and the intensity of trading (proxied by the number of trades per unit time, N_h). This suggests that—at least in terms of mean effects—intervention retains its efficacy even during times of higher volatility and trading intensity. The one state variable that clearly affects price impact is the flow (number) of macroeconomic announcements. In this case, a stronger flow of announcements is associated with greater price impact.

5. Conclusion

This paper measures portfolio adjustment where it actually occurs—within the trading process. The resulting flow of transacted orders provides a powerful means of testing portfolio-balance models. Until recently, this strategy was not feasible due to data limitations. The advent of electronic trading, and the data it provides, has made it feasible for the first time.

Our analysis provides three main results. First, our model's implications are borne out: we find strong evidence of price effects from portfolio balance, both temporary and persistent. This contrasts with a common belief that these effects (from intervention or otherwise) are too small to be detectable. Not only are they detectable, they are also economically significant, leading us to conclude that portfolio-balance theory is more applicable than many believe. The second result pertains to the economic significance noted above. Specifically, we establish the (unconditional) price impact of trades of about 0.5 percent per \$1 billion. With gross flows in the largest spot markets at about \$300 billion per day, this level of price impact is potentially quite important.[20] Our third result clarifies how this unconditional price impact varies with the state of the market. The most important state variable for the size of the price-impact coefficient is the flow of macroeconomic announcements. (It may be, for example, that order flow is the variable that market participants use to resolve uncertainty about how these announcements are interpreted.) Whatever the reason, our estimates imply that trades have the most price impact when the flow of macroeconomic news is strong. This applies to intervention trades as well, provided they are sterilized, secret, and provide no policy signal (i.e., as long as they mimic private trades).

---

[20] 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: the average intervention of \$200 million reported by Dominguez and Frankel (1993b) translates into an exchange-rate movement of only 0.10 percent, an amount easily swamped by movements due to other factors.

Finally, we offer some thoughts on future application of our trading-theoretic approach to intervention. Market data now coming available allow precise tracking of how the market absorbs actual CB trades, and any information in them. CB's with precise knowledge of their own trades—announcements, timing, stealth level, etc.—can estimate the impact of these various "parameter" settings. Consider, for example, the type of data used by Payne (1999), which includes the order book of an electronic interdealer broker. A CB with these data, over a sufficiently large number of intervention trades, can learn exactly how the "book" is affected, including the process of price adjustment, liquidity provision on both sides, and transaction activity. It is something like a doctor who has a patient ingest blue dye to determine how it passes through the system. The whole process becomes transparent. Such is the future of empirical work on this topic.

References

Allen, P., and P. Kenen, 1980, *Asset Markets, Exchange Rates, and Economic Integration* (Cambridge University Press: New York).

Amihud, Y., and H. Mendelson, 1980, Dealership markets: Marketmaking with inventory, *Journal of Financial Economics*, 8: 31-53.

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.

Backus, D., 1984, Empirical models of the exchange rate: Separating the wheat from the chaff, *Canadian Journal of Economics*, 17: 826-846.

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

Bierens, H., 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, 2000, 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.

Branson, W., 1977, Asset markets and relative prices in exchange rate determination, *Sozialwissenschaftliche Annalen*, 1: 69-89.

Branson, W., H. Halttunen, and P. Masson, 1977, Exchange rates in the short run: The Deutschemark rate, *European Economic Review*, 10: 303-324.

Branson, W., and D. Henderson, 1985, The specification and influence of asset markets, in R. Jones and P. Kenen (eds.), *Handbook of International Economics*, Volume 2, North-Holland: Amsterdam.

Chaboud, A., and B. LeBaron, 1999, Foreign exchange market trading volume and Federal Reserve intervention, typescript, Brandeis University, July, forthcoming in the *Journal of Banking and Finance*.

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.

Dominguez, K., 1990, Market responses to coordinated central bank intervention, *Carnegie-Rochester Series on Public Policy*, 32: 121-163.

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., 1998, On foreign exchange intervention: An assessment of the U.S. experience, typescript, Board of Governors of the Federal Reserve System, February.

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, *Journal of Political Economy*, forthcoming.

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.

Girton, L., and D. Henderson, 1977, Central bank operations in foreign and domestic assets under fixed and flexible exchange rates, in P. Clark, D. Logue, and R. Sweeney (eds.) *The Effects of Exchange Rate Adjustment*, (Department of the Treasury, Washington, D.C.) 151-179.

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.

Golub, S., 1989, Foreign currency government debt, asset markets, and the balance of payments, *Journal of International Money and Finance*, 8: 285-294.

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

Grossman. S., and J. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review*, 70: 393-408.

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.

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

Hasbrouck, J., 1991, Measuring the information content of stock trades, *Journal of Finance*, 46: 179-207.

Ho, T., and H. Stoll, 1983, The dynamics of dealer markets under competition, *Journal of Finance* 38, 1053-1074.

Hung, J., 1997, Intervention strategies and exchange rate volatility: A noise trading perspective, *Journal of International Money and Finance*, 16: 779-793.

Jones, C., G. Kaul, and M. Lipson, 1994, Transactions, volume, and volatility, *Review of Financial Studies*, 7: 631-651.

Kaul, A., V. Mehrotra, and R. Morck, 2000, Demand curves for stock do slope down: New evidence from an index weights adjustment, *Journal of Finance*, 55: 893-912.

Kodres, L., and M. Pritsker, 1998, A rational expectations model of financial contagion, working paper, International Monetary Fund.

Kouri, P., 1976, The exchange rate and the balance of payments in the short run and in the long run: A monetary approach, *Scandinavian Journal of Economics*, 78: 280-304.

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

Kyle, A., and W. Xiong, 2001, Contagion as a wealth effect, typescript, Duke University.

Lewis, K., 1988, Testing the portfolio balance model: A multilateral approach, *Journal of International Economics*, 7: 273-288.

Lewis, K., 1995, Puzzles in international financial markets, in G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, vol. 3, North Holland: Amsterdam.

Loopesko, B., 1984, Relationships among exchange rates, intervention, and interest rates: An empirical investigation, *Journal of International Money and Finance*, 3: 257-277.

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.

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

Rogoff, K., 1984, On the effects of sterilized intervention: An analysis of weekly data, *Journal of Monetary Economics*, 14: 133-150.

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., 1999, Sterilized central bank intervention in the foreign exchange market, *Journal of International Economics*, 49: 245-267.

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.

Appendix A

<u>Model Solution</u>

Each dealer determines quotes and speculative demand by maximizing a negative exponential utility function defined over periodic (daily) wealth. Within a given day t, let W_j^i denote the end-of-round j wealth of dealer i, using the convention that W_0^i denotes wealth at the end of day t-1. (We suppress notation to reflect the day t where clarity permits.) With this notation, and normalizing the gross return on the riskless asset to one, we can define the dealers' problem over the six choice variables described in section 2, namely, the four scalar quotes P_j^i, one for each round j, and the two outgoing interdealer trades, T_2^i and T_3^i:

(A1)
$$\underset{\{P_1^i, P_2^i, P_3^i, P_4^i, T_2^i, T_3^i\}}{Max} \quad E\left[-\exp(-\theta W_4^i \mid \Omega^i)\right]$$

s.t.

$$
\begin{aligned}
W_4^i = \ & W_0^i \ + \ C_1^i(P_1^i - \tilde{P}_2^i) \ + \ \tilde{T}_2^i\left(P_2^i - \tilde{P}_3^i\right) \ + \ \tilde{T}_3^i\left(P_3^i - P_4^i\right) \\
& + \ \left(T_2^i - C_1^i\right)\left(\tilde{P}_3^i - \tilde{P}_2^i\right) \ + \ \left(T_3^i + T_2^i - C_1^i - \tilde{T}_2^i\right)\left(P_4^i - \tilde{P}_3^i\right)
\end{aligned}
$$

Dealer i's wealth over the four-round trading day is affected by positions taken two ways: incoming random orders and outgoing deliberate orders. The incoming random orders include the public order C_1^i and the incoming interdealer orders $\tilde{T}_2^i$ and $\tilde{T}_3^i$ (tilde distinguishes incoming interdealer orders and prices from outgoing). The outgoing deliberate orders are the two interdealer trades T_2^i and T_3^i. $\tilde{P}_j^i$ denotes an incoming interdealer quote received by dealer i in round j. As an example, the second term in the budget constraint reflects the position from the public order C_1^i received in round one at dealer i's own quote P_1^i and subsequently unwound at the incoming interdealer quote $\tilde{P}_2^i$ in round-two. (Recall that the sign of dealer i's position is opposite that of C_1^i, so a falling price is good for dealer i if the public order is a buy, i.e., positive. The dealer's speculative positioning based on information in C_1^i is reflected in the final two terms of the budget constraint.) Terms three and four reflecting the incoming random dealer orders are analogous.

Terms five and six of the budget constraint reflect the dealer's speculative and hedging demands. The outgoing interdealer trade in round 2 has three components:

(A2)
$$T_2^i = \ C_1^i \ + \ D_2^i \ + \ E[\tilde{T}_2^i \mid \Omega_{2T}^i]$$

where D_2^i is dealer i's speculative demand in round 2, and $E[\tilde{T}_2^i|\Omega_{2T}^i]$ is the dealer's hedge against incoming orders from other dealers (this term is zero in equilibrium given the distribution of the C_1^i's). The dealer's total demand can be written as follows:

$$T_2^i - C_1^i = D_2^i + E[\tilde{T}_2^i|\Omega_{2T}^i]$$

which corresponds to the position in term five of the budget constraint. The sixth term in the budget constraint is analogous: the dealer's total demand in round three is his total trade in round three (T_3^i) plus his total demand in round two ($T_2^i - C_1^i$) less the random interdealer order he received in round two ($\tilde{T}_2^i$).

The conditioning information Ω^i at each decision node (4 quotes and 2 outgoing orders) is summarized below (see also the daily timing in figure 1).

$$\Omega_{1P}^i = \left\{ \sum_{k=1}^{t-1}\left\{\Delta R_k, X_{2k}, X_{3k}, P_{1k}, P_{2k}, P_{3k}, P_{4k}\right\}\right\}$$

$$\Omega_{2P}^i = \left\{\Omega_{1P}^i, P_{1t}, C_1^i\right\}$$

$$\Omega_{2T}^i = \left\{\Omega_{2P}^i, P_{2t}\right\}$$

$$\Omega_{3P}^i = \left\{\Omega_{2T}^i, X_{2t}\right\}$$

$$\Omega_{3T}^i = \left\{\Omega_{3P}^i, P_{3t}, \Delta R_t\right\}$$

$$\Omega_{4P}^i = \left\{\Omega_{3T}^i, X_{3t}\right\}$$

At this stage it is necessary to treat each of the prices in these information sets as a vector that contains the price of each individual dealer i (though in equilibrium each of these prices is a scalar, as shown below).

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 beliefs.

To solve for the symmetric BNE, first consider optimal quoting strategies.

PROPOSITION A1: A quoting strategy is consistent with symmetric BNE only if quotes within any single trading round are common across dealers.

PROPOSITION A2: A quoting strategy is consistent with symmetric BNE only if $P_1 = P_2$ and these prices are equal to the final round price P_4 from the previous day.

PROPOSITION A3: A quoting strategy is consistent with symmetric BNE only if the common round-three quote is:

$$P_3 = P_2 + \lambda_2 X_2$$

where the constant λ_2 is strictly positive and X_2 denotes the signal of round-two interdealer order flow.

PROPOSITION A4: A quoting strategy is consistent with symmetric BNE only if the common round-four quote is:

$$P_4 = P_3 + \lambda_3 X_3 + \delta \Delta R - \phi(P_3 - P_2)$$

where the constants λ_3, δ and ϕ are strictly positive and X_3 denotes round-three interdealer order flow.

Propositions A1 through A4

The proof of proposition A1 is straightforward: That all dealers post the same quote in any given trading round is required to eliminate risk-free arbitrage. (Recall from section 2 that all quotes are scalar prices at which the dealer agrees to buy/sell any amount, and trading with multiple partners is feasible.)

The proof of proposition A2 is straightforward as well: Common prices require that quotes depend only on information that is commonly observed. In round one, this includes the previous day's round-four price. Because there is no new information that is commonly observed between round four and round two quoting the following day, the round-four price is not updated. (Recall that public trading in round four is a deterministic function of round-four prices and therefore conveys no information.) Thus, dealers' round-two quotes are not conditioned on individual realizations of C_1^i.

Propositions A3 and A4 require equations that pin down the levels of the four prices. Per above, these equations are necessarily functions of public information. Naturally, they also embed the equilibrium trading rules of dealers and customers. The equations are the following:

$$(A3) \qquad E\left[C_1 + I \middle| \Omega_{1P}\right] + E\left[ND_2^i(P_1)\middle|\Omega_{1P}\right] = 0$$

$$(A4) \qquad E\left[C_1 + I \middle| \Omega_{2P}\right] + E\left[ND_2^i(P_2)\middle|\Omega_{2P}\right] = 0$$

$$(A5) \qquad E\left[C_1 + I \middle| \Omega_{3P}\right] + E\left[ND_3^i(P_3)\middle|\Omega_{3P}\right] = 0$$

$$(A6) \qquad E\left[C_1 + I \middle| \Omega_{4P}\right] + E\left[C_4(P_4)\middle|\Omega_{4P}\right] = 0$$

where C_1 denotes the sum of C_1^i over all N dealers. The first three equations state that for

each round j (j=1,2,3), at price P_j dealers willingly absorb the estimated demand from customers and the central bank (realized at the beginning of the day, but not observed publicly). The fourth equation states that at price P_4 the public willingly absorbs the estimated demand from customers and the central bank. These equations pin down equilibrium prices because any price other than that which satisfies each generates irreconcilable demands in interdealer trading in rounds two and three (e.g., if price is too low, all dealers know that on average dealers are trying to buy from other dealers, which is inconsistent with rational expectations; see Lyons 1997 for a detailed treatment in another model within the simultaneous trade approach).

From these equations, $P_2-P_1=0$ follows directly from two facts: (1) the expected value of C_1+I conditional on public information Ω_{1P} or Ω_{2P} is zero and (2) expected dealer demand D_2^i is also zero at this public-information-unbiased price. To be more precise, this statement *postulates* that the dealer's demand D_2^i has this property; derivation of the optimal trading rule shows that this is the case.

That $P_3-P_2=\lambda_2 X_2$ with $\lambda_2>0$ follows from two facts: (1) interdealer order flow X_2 is the only public information revealed in this interval and (2) X_2 is positively correlated with—and therefore provides information about—the morning portfolio shift C_1+I. The positive correlation arises because each of the dealer orders T_2^i of which X_2 is composed is proportional to the C_1^i received by that dealer (and proportional to C_1^i+I in the case of the dealer receiving the central bank order, per proposition A5 below). A positive expected C_1+I induces an increase in price because it implies that dealers—having taken the other side of these trades—are short and need to be induced to hold this short position with an expected downward drift in price.

The exact size of this downward drift in price depends on where price is expected to settle at the end of the day. Per proposition A4, $P_4-P_3 = \lambda_3 X_3 + \delta\Delta R - \phi(P_3-P_2)$. This price change depends positively on the two pieces of public information revealed in this interval, X_3 and ΔR.[21] The logic behind the positive X_3 effect is the same as that behind the positive X_2 effect in round two: a positive average T_3^i implies that the market's estimate of C_1+I from X_2 was too low; absorption of the additional short position requires price increase. (That a positive average T_3^i implies this is clear from the derivation of T_3^i.) The term $\delta\Delta R$ is the perpetuity value of the change in the daily payoff R_t. Finally, the drift term $-\phi(P_3-P_2)$ is the equilibrium compensation to dealers for having to absorb the morning portfolio shift through the interval in which ΔR (and the associated price risk) is realized. This is an intraday price effect that dissipates by the end of the day.

---

[21] Interdealer order flow X_3 is observed without noise, which means it reveals the value of C_1+I fully. The price in period four must therefore adjust such that equation (A6) is satisfied exactly.

Equilibrium Trading Strategies

An implication of common interdealer quotes is that in rounds two and three each dealer receives a share 1/(N-1) of every other dealer's interdealer trade. These orders correspond to the position disturbances $\widetilde{T}_2^{\,i}$ and $\widetilde{T}_3^{\,i}$ in the dealer's problem in equation (A1).

Given the quoting strategy described in propositions 1-4, the following trading strategy is optimal and corresponds to symmetric linear equilibrium:

PROPOSITION A5: The trading strategy profile:
$$T_2^i = \alpha C_1^i$$
for dealers not receiving the central bank order and
$$T_2^i = \alpha \left(C_1^i + I \right)$$
for the one dealer receiving the central bank order, with $\alpha > 0$, conforms to Bayesian-Nash equilibrium.

PROPOSITION A6: The trading strategy profile:
$$T_3^i = \kappa_1 C_1^i + \kappa_2 X_2 + \kappa_3 \tilde{T}_2^i$$
for dealers not receiving the central bank order and
$$T_3^i = \kappa_1 \left(C_1^i + I \right) + \kappa_2 X_2 + \kappa_3 \tilde{T}_2^i$$
for the one dealer receiving the central bank order, conforms to Bayesian-Nash equilibrium.

Sketch of Proofs for Propositions A5 and A6

As noted above, 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 day. Because there are only N dealers, however, each dealer acts strategically in the sense that his speculative demand depends on the impact his trade will have on subsequent prices.

Propositions A5 and A6 are special cases of the analysis in Lyons (1997), which is also set in the context of a simultaneous-trade game with two interdealer trading rounds. Accordingly, we refer readers to that analysis for details on the derivation of optimal trading rules in this setting.[22] Two differences warrant note here. First, the Lyons (1997) analysis also includes private and public signals (denote s_i and s in that paper) that are not present in our specification, i.e., they equal zero. Second, our model includes a central bank trade. However, because the central bank trade is sterilized, secret, and provides no monetary policy signal, the dealer receiving that trade treats it the same as any other customer trade.

---

[22] Available from the authors on request.

Appendix B

Kernel Regression

To examine whether there is state-dependency in the relation between aggregate order flow and price changes we consider nonparametric regressions of the form

$$\Delta p_h = m(z_h) + \eta_h$$

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

$$\widehat{m}(z_h) = \frac{\sum_{j=0, j\neq h}^{H} K_\vartheta(z_h - z_j)\Delta p_j}{\sum_{j=0, j\neq h}^{H} K_\vartheta(z_h - z_j)}$$

where the kernel function $K(u) \geq 0$, $\int K(u)du = 1$, and $K_b(u)=b^{-1}K(u/b)$, where b is the bandwidth parameter. In this application, we use the multivariate Gaussian kernel $K(u) = (2\pi)^{-d/2}\exp(-u'u/2)$, where $d = \dim(u)$. The bandwidth parameter, b, is chosen by cross-validation. That is to say, b minimizes

$$\frac{1}{H}\sum_{h}^{H}\left(\Delta p_h - m(z_h)\right)^2 w_h$$

where w_h 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 3, z_h contains $\{X_h, \Delta P_{h-1}, \sigma_{h-1}, N_{h-1}, A_{h-1}\}$. We follow the common practice of including the standardized value of each of these variables in the Gaussian kernel (i.e., each element of z_h is divided by its sample standard deviation).

Asymptotic theory for kernel regressions in the time series context appear 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 $H\sum_{h}^{\infty}\alpha(k)^{1-2/\delta} = O(1)$ and $E \mid \Delta p_h \mid^{\delta} < \infty$, $\delta > 2$.

The Empirical Model

In the general case where $\rho_3(y_t^{h-1}) \leq 1$, the distribution of hourly price changes is

$$\Delta P_h = \begin{cases} 0 & ; \quad \theta_I(Y_{h-1}) \\ \lambda_2 X_h & ; \quad \theta_{II}(Y_{h-1}) \\ \lambda_3 X_h - \phi\Delta P_{h-k} + \delta\Delta R_h & ; \quad \theta_{III,k}(Y_{h-1}) \end{cases}$$

where

$$\theta_I(Y_{h-1}) = \pi_1(Y_{h-1}) + (1 - \rho_2(y_{h-1}))\pi_2(Y_{h-1}) + \pi_4(Y_{h-1}) + (1 - \rho_3(y_{h-1}))\pi_3(Y_{h-1})$$
$$\theta_{II}(Y_{h-1}) = \rho_2(y_{h-1})\pi_2(Y_{h-1})$$

$$\theta_{III,k}(Y_{h-1}) = \rho_3(y_{h-1})\prod_{j=2}^{k}\{1 - \rho_3(y_{h-j})\}\rho_2(y_{h-k-1})\pi_2(Y_{h-k-1})$$

The expected change in prices is given by

$$E[\Delta P_h | \Omega_h] = \beta_0(Y_{h-1})\Delta X_h + \sum_{i=1}^{\infty}\beta_i(Y_{h-1})\Delta P_{h-i}$$

where

$$\beta_0(Y_{h-1}) = \lambda_2\theta_{II}(Y_{h-1}) + \lambda_3\sum_{k=1}^{\infty}\theta_{III,k}(Y_{h-1})$$

and

$$\beta_k(Y_{h-1}) = -\phi\theta_{III,k}(Y_{h-1}).$$

Appendix C

<u>Log Price Changes</u>

Our model derivation implies a dependent variable in the form of changes in the level of price, rather than changes in the log of price (despite the latter being customary in exchange-rate economics). For robustness, we estimate the model using log changes. The results corresponding to text Tables 2 and 3 appear below as Tables A1 and A2. Qualitatively, the results are quite similar.

Table A1: Estimates of the Micro Portfolio-Balance Model

Log Price changes

$$\Delta p_h = \beta_1 X_h + \beta_2 \Delta p_{h-1} + \eta_h$$

| | X_h | Δp_{h-1} | X_{h-1} | Diagnostics | | |
				R^2	Serial	Hetero
I	0.171	-0.204		0.211	0.430	0.074
	(13.098)	(4.053)			0.252	0.022
II	0.149			0.171	<0.001	0.074
	(12.210)				<0.001	0.012
III		-0.062		0.003	0.156	0.276
		(1.380)			0.303	0.018
IV	0.155		-0.027	0.179	<0.001	0.070
	(11.993)		(2.542)		<0.001	0.010
V	0.171	-0.203	-0.001	0.211	0.199	0.074
	(12.907)	(3.810)	(0.082)		0.247	0.022

Tests for State-Dependency

p-values

State variables	$\beta_1(.)$	$\beta_2(.)$
τ: time of day	<0.001	0.675
N_h: number of trades	0.002	0.051
σ_h: std. dev. of prices	<0.001	0.241
A_h: number of announcements	0.183	0.007

OLS estimates are based on hourly observations from 6:00 to 18:00 BTS from May 1 to August 31, 1996, excluding weekends. T-statistics are shown in parentheses and are calculated with asymptotic standard errors corrected for heteroskedasticity. Δp_h is the hourly change in the log spot exchange rate (DM/\$). X_h is the hourly interdealer order flow, measured contemporaneously with Δp_h (negative for net dollar sales, in thousands). The Serial column presents the p-value of a chi-squared LM test for first-order (top row) and fifth-order (bottom row) serial correlation in the residuals. The Hetero column presents the p-value of a chi-squared LM test for first-order (top row) and fifth–order (bottom row) ARCH in the residuals. The lower panel presents the p-values for a heteroskedasticity-consistent Wald test of the null hypothesis that $\beta_i(.)$ does not vary with the state variable shown in the left-hand column. σ_h is the standard deviation of all the transactions prices, N_h is the number of transactions, and A_h is the number of macroeconomic announcements, all during hour h. τ is a vector of three dummy variables, $[\tau_1\ \tau_2\ \tau_3]$. τ_1 equals one for hours between 6:00am and 7:59am, zero otherwise; τ_2 equals one for hours between 8:00am and 11:59am, zero otherwise; and τ_3 equals one for hours between 12:00pm and 1:59 pm, zero otherwise.

Table A2: Nonparametric (Kernel) Regressions

Log Price Changes

$$\Delta p_h = m(z_h) + \eta_h$$

Dependent Variable	Const.	z_h							Diagnostics		
		X_h	Δp_{h-1}	σ_{h-1}	N_{h-1}	A_{h-1}	τ_h	τ_h^2	R^2	Serial	Hetero
$\hat{m}_x(z_h)$	0.089	<0.001	<0.001	-0.009	0.001	0.016			0.069	0.873	0.577
	(15.213)	(0.326)	(0.966)	(1.215)	(0.284)	(3.741)				0.238	0.340
	0.051			-0.009	-0.001	0.015	0.007	0.000	0.069	0.948	0.552
	(1.980)			(0.969)	(0.246)	(3.515)	(1.487)	(1.356)		0.204	0.347
	0.090	-0.012	0.001	0.016					0.067	0.971	0.557
	(14.707)	(1.421)	(0.369)	(3.757)						0.233	0.338
$\hat{m}_{\Delta p}(z_h)$	-0.049	0.001	-0.002	-0.073	0.016	-0.030			0.069	0.116	<0.001
	(3.073)	(1.542)	(2.879)	(3.359)	(2.540)	(2.749)				0.184	<0.001
	-0.010			-0.054	0.018	-0.027	-0.006	0.000	0.046	0.300	<0.001
	(0.155)			(1.915)	(2.404)	(2.476)	(0.545)	(0.201)		0.121	<0.001
	-0.059	-0.052	0.014	-0.029					0.041	0.286	<0.001
	(3.653)	(1.973)	(2.189)	(2.656)						0.158	<0.001
$\left|\hat{\eta}_h\right|$	2.914	-0.008	-0.028	9.262	-0.732	0.378			0.186	0.926	0.978
	(4.859)	(0.697)	(1.051)	(9.660)	(3.143)	(1.201)				0.165	0.216
	19.213			8.232	-0.204	0.629	-2.937	0.116	0.202	0.440	0.633
	(4.878)			(6.977)	(0.754)	(2.072)	(4.381)	(4.301)		0.158	0.232
	20.472			7.887		0.604	-3.177	0.125	0.201	0.364	0.559
	(5.761)			(7.374)		(2.053)	(5.290)	(5.078)		0.192	0.255

The table presents estimates from the kernel regression shown in the heading, where z_h is the vector of conditioning variables and $\hat{m}(z_h)$ is the nonparametric estimate of $E[\Delta p_h \mid z_h]$, where Δp_h is the hourly change in the log spot exchange rate (DM/\$). $\hat{m}_x(z_h)$ and $\hat{m}_{\Delta p}(z_h)$ are the derivatives of the estimated Kernel with respect to X_h and Δp_{h-1} respectively. Estimated using OLS, with t-statistics shown in parentheses (calculated with standard errors corrected for heteroskedasticity). τ_h is the number of hours since midnight at the end of hour h. See Table A1 for other variable definitions. The Serial column presents the p-value of a chi-squared LM test for first-order (top row) and fifth-order (bottom row) residual serial correlation. The Hetero column presents the p-value of a chi-squared LM test for first-order (top row) and fifth–order (bottom row) ARCH in the residuals.