

The Microstructure Approach to Exchange Rates

Richard K. Lyons
U.C. Berkeley (faculty.haas.berkeley.edu/lyons)
MIT Press, published Dec. 2001

Figure and Table Caption List

- Figure 1.1: The Two Stages of Information Processing.
Figure 1.2: Four Months of Exchange Rates (solid) and Order Flow (dashed).

Figure 2.1: Supply Curves with only Transitory Inventory Effects.
Figure 2.2: Supply Curves With Inventory & Portfolio-balance Effects.
Figure 2.3: Supply Curves when Order Flow Conveys Information about both Payoffs and Discount Rates.
Figure 3.1: Three Types of Trades.
Figure 3.2: Dealer's net position (in \$ millions) over one trading week.
Table 3.1: Diagram of position sheet structure.
Table 3.2: Summary of DM/\$ dealer's trading and profits.
Table 3.3: Foreign exchange market turnover.
Table 3.4: Reported foreign exchange market turnover by currency pair.
Figure 3.3: The microstructure effects question.
Figure 3.4: The Accelerationist view of order-flow information.

Figure 4.1: A Bird's Eye View of Microstructure Models.
Figure 4.2: Timing of Rational Expectations Model.
Figure 4.3: Summary of Rational-Expectations Auction Model.
Figure 4.4: The Two Stages of Information Processing.
Figure 4.5: Timing of Kyle Model.
Figure 4.6: Summary of Kyle Auction Model.
Figure 4.7: Timing of a Single Trade in Sequential-Trade Model.
Figure 4.8: Summary of Sequential-Trade Model.
Figure 4.9: Probability of Different Trade Types—Sequential-Trade Model.
Figure 4.10: Timing in the Simultaneous-Trade Model.
Figure 4.11: Summary of Simultaneous Trade Model.
Figure 4.12: Negative exponential utility.

Figure 5.1: The Three Data Groupings.
Figure 5.2: Example of Dealing 2000-1 Communication.
Figure 5.3: Diagram of data structure.
Figure 5.4: An illustration of spread components.
Figure 5.5: Timing in each period of the DP model.
Figure 5.6: Dealer i's quote schedule in the DP model.

Table 5.1: Structural Model Estimates.

Table 5.2: Testing the hot potato hypothesis.

Table 5.3: Testing the hot potato hypothesis: Is order flow less informative when transactions follow in the same direction?

Figure 6.1: An illustration of overshooting.

Figure 6.2: Three approaches to exchange rates and their models.

Figure 6.3: The Issues Spectrum.

Figure 7.1: Spanning Macro and Microstructure Graphically.

Figure 7.2: Daily Timing in the Evans-Lyons Model.

Figure 7.3: Summary of Evans and Lyons (1999) Model.

Figure 7.4: Portfolio Balance Effects: One Period Example.

Table 7.1: Estimates of the Evans-Lyons model.

Figure 7.5: DEM-FRF Level and Cumulative Net Order Flow.

Figure 7.6: The Two Trading Regimes.

Figure 7.7: The Statistician's Perspective on Forward Bias.

Figure 7.8: The Practitioner's Perspective on Forward Bias.

Table 7.2: Sharpe ratios (annual basis) from pure currency strategy.

Figure 8.1: Fed balance sheet: Unsterilized purchase of \$100 million with yen.

Figure 8.2: Fed balance sheet: Sterilized purchase of \$100 million with yen.

Figure 8.3: Types of Intervention.

Figure 8.4: Intervention Transparency Spectrum.

Figure 9.1: The Trading Volume Pie.

Table 9.1: Customer Trades: Volumes and Order Flow.

Figure 9.2: Cumulative customer flow and exchange rates.

Figure 9.3: Cumulative customer flow and exchange rates over the Evans-Lyons sample.

Table 9.2: The Price Impact of Aggregate Customer Orders.

Table 9.3: The Price Impact of Disaggregated Customer Orders.

Figure 9.4, Panel A: Cumulative total customer flow and the Yen/\$ rate around the October 1998 collapse.

Figure 9.4, Panel B: Cumulative flow of leveraged financial institutions and the Yen/\$ rate around the October 1998 collapse.

Figure 9.4, Panel C: Cumulative flow of non-financial corporations and the Yen/\$ rate around the October 1998 collapse.

Figure 9.4, Panel D: Cumulative flow of unleveraged financial institutions and the Yen/\$ rate around the October 1998 collapse.

Figure 1.1

The Two Stages of Information Processing




Figure 1.2

Four Months of Exchange Rates (solid) and Order Flow (dashed)

May 1-August 31, 1996





Figure 2.1


Supply Curves with only Transitory Inventory Effects



The dotted region represents the transitory inventory effects. The effective spread faced by a customer for a 10-unit order is the difference in price along the short-run net supply curve S^{SR} between -10 and $+10$. If a customer wants to buy 10 British pounds from the dealer—an order of $+10$ —then he must pay the higher dollar price. If the customer wants to sell 10 pounds to the dealer—an order of -10 —then he will receive the lower dollar price. Over the longer run, however, the dealer unloads his position on the rest of the market at a price that does not include the transitory inventory effects. The market's net supply is perfectly elastic, by assumption, which corresponds to a longer-run supply curve S^{LR} slope of zero. The linear relationship shown along S^{SR} is a special case, which I adopt for simplicity.

Figure 2.2


Supply Curves With Inventory & Portfolio-balance Effects



The dotted region represents the transitory inventory effects. The gray region represents persistent portfolio-balance effects. Due to inventory effects, the short-run price impact of an incoming order is larger than the long-run impact. But the long-run impact is non-zero, due to imperfect substitutability; i.e., the long-run net supply curve S^{LR} now slopes upward. The linear relationships shown are a special case, which I adopt for simplicity.

Figure 2.3

**Supply Curves when Order Flow Conveys Information
about both Payoffs and Discount Rates**



The dotted region represents the transitory inventory effects. The cross-hatched region represents persistent payoff-information effects. The gray region represents persistent portfolio-balance effects. The figure therefore reflects all 3 of the information types that arise in microstructure theory. The long-run supply curve SLR^* reflects both the long-run effects from imperfect substitutability (SLR), plus an additional long-run effect due to the payoff information conveyed by order flow. The linear relationships shown are a special case, which I adopt for simplicity.

Figure 3.1

Three Types of Trades

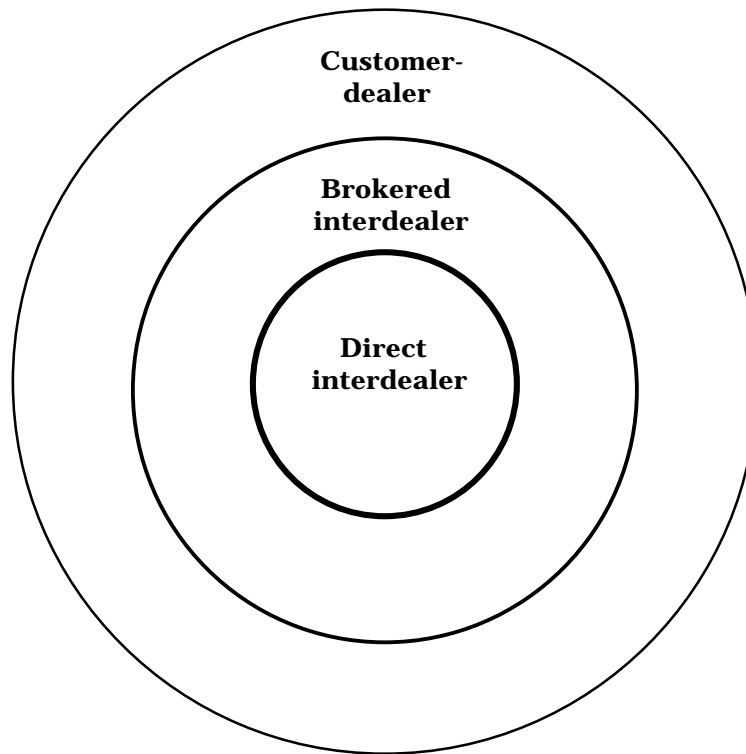
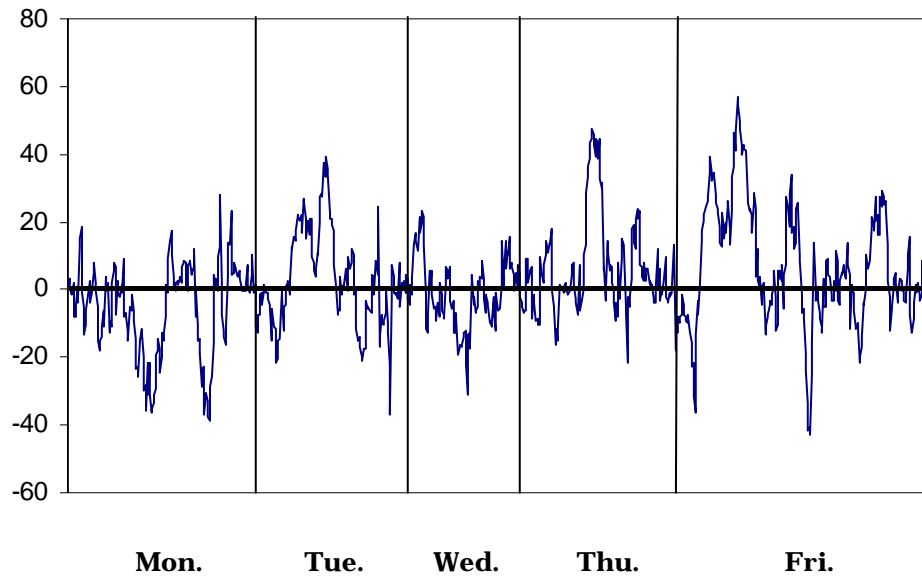


Figure 3.2

Dealer's net position (in \$ millions) over one trading week.



The vertical lines represent the overnight periods over which this dealer was not trading. The horizontal distance between those vertical lines is scaled by the number of transactions made by this dealer each trading day.

Table 3.1

Diagram of position sheet structure,
first fourteen trades on Monday, August 3, 1992

Trade date: 8/3					
Value date: 8/5					
Position	Position rate	Trade	Trade rate	Source	Time
		1	1.4794	r	8:30
		2	1.4797	r	
3	1.4796				
		28	1.4795	r	
		-10	1.4797	r	
		-10	1.4797	b	
		-10	1.4797	r	
		-3	1.4797	b	
-2	1.4797				
		0.5	1.4794	r	
		0.75	1.4790	r	
		3	1.4791	r	
2	1.4791				
		-10	1.4797	r	
-8	1.4797				
		2	1.4799	b	
-6	1.4797				8:38
		5	1.4805	b	
		-7	1.4810	r	
-8	1.4808				

The "Position" column accumulates the individual trades in the "Trade" column. Quantities are in millions of dollars. A positive quantity in the Trade column corresponds to a purchase of dollars. A positive quantity in the Position column corresponds to a net long dollar position. The "Trade Rate" column records the exchange rate for the trade, in deutschemarks per dollar. The "Position Rate" column records the dealer's estimate of the average rate at which he acquired his position. The Position and Position Rate are not calculated after every trade due to time constraints. The "Source" column reports whether the trade is direct over the Reuters Dealing 2000-1 system (r=Reuters) or brokered (b=Broker). All trades on this position sheet are interdealer.

Table 3.2

Summary of DM/\$ dealer's trading and profits
from Monday, August 3 to Friday, August 7, 1992.

	<u>Transactions</u>	<u>Volume (mil)</u>	<u>Profit: Actual</u>	<u>Profit: Spread</u>
Monday	333	\$ 1,403	\$ 124,253	\$ 95,101
Tuesday	301	\$ 1,105	\$ 39,273	\$ 74,933
Wednesday	300	\$ 1,157	\$ 78,575	\$ 78,447
Thursday	328	\$ 1,338	\$ 67,316	\$ 90,717
Friday	458	\$ 1,966	\$ 198,512	\$ 133,298
Total	1,720	\$ 6,969	\$ 507,929	\$ 472,496

The "Profit: Spread" column reports the profit the dealer would have realized if he had cleared one-third of his spread on every transaction. It is calculated as the dollar volume times one-third the median spread he quoted in the sample (median spread = 0.0003 DM/\$), divided by the average DM/\$ rate over the sample (1.475 DM/\$).

Table 3.3Foreign exchange market turnover (BIS 1999, Table A-1)¹

Daily averages in billions of US dollars

Category	April 1989	April 1992	April 1995	April 1998
Spot transactions ²	350	400	520	600
Outright forwards and forex swaps ²	240	420	670	900
Total “traditional” turnover	590	820	1,190	1,500
Memorandum item: Turnover at April 1998 exchange rates	600	800	1,030	1,500
¹ Adjusted for local and cross-border double counting. ² Includes estimates for gaps in reporting.				

Table 3.4

Reported foreign exchange market turnover by currency pair (BIS 1999, Table B-4)

Daily averages in billions of US dollars and percentage shares

April 1995					April 1998				
	Total	Spot	Outright forwards	Foreign exchange swaps		Total	Spot	Outright forwards	Foreign exchange swaps
	Amount	Percentage share				Amount	Percentage share		
USD/DEM	253.9	56	7	37	USD/DEM	290.5	49	8	43
USD/JPY	242.0	36	9	55	USD/JPY	266.6	45	10	44
USD/othEMS	104.3	19	8	73	USD/othEMS	175.8	14	7	79
USD/GBP	77.6	33	7	60	USD/GBP	117.7	33	9	59
USD/CHF	60.5	37	9	55	USD/CHF	78.6	30	7	62
USD/FRF	60.0	17	9	74	USD/FRF	57.9	16	8	76
DEM/othEMS	38.2	74	9	17	USD/CAD	50.0	25	6	68
USD/CAD	38.2	32	11	57	USD/AUD	42.2	33	8	59
DEM/FRF	34.4	86	4	9	DEM/othEMS	35.1	75	12	13
USD/AUD	28.7	31	7	63	DEM/GBP	30.7	79	10	11
DEM/JPY	24.0	79	12	9	DEM/JPY	24.2	77	14	9
DEM/GBP	21.3	84	6	10	DEM/CHF	18.4	85	7	8
DEM/CHF	18.4	86	6	7	USD/XEU	16.6	7	4	89
USD/XEU	17.9	11	7	82	USD/SGD	17.2	71	2	27
All currency pairs	1,136.9	43	9	48	All currency pairs	1,441.5	40	9	51

USD=U.S. dollar, DEM=Deutsche mark, JPY=Japanese yen, othEMS=other EMS (European Monetary System) currencies, GBP=British pound, CHF=Swiss franc, FRF=French franc, CAD=Canadian dollar, AUD=Australian dollar, XEU=European currency unit (a basket currency that includes all European Union members), and SGD=Singapore dollar.

Figure 3.3

The microstructure effects question: Does the trading process affect the mapping?

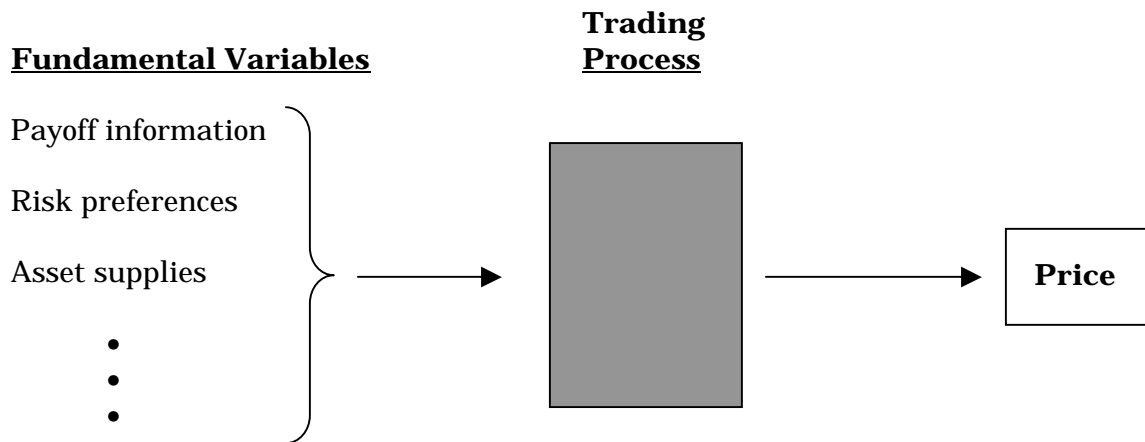



Figure 3.4

The Accelerationist view of order-flow information



The solid line shows a hypothetical price path for a stock under the assumption that price responds only to a higher-than-expected public earnings announcement at time t . The dotted line shows the price path under the assumption that an insider is trading in advance of the announcement, and the information in the insider's buy orders are pushing price up.

Figure 4.1

A Bird's Eye View of Microstructure Models

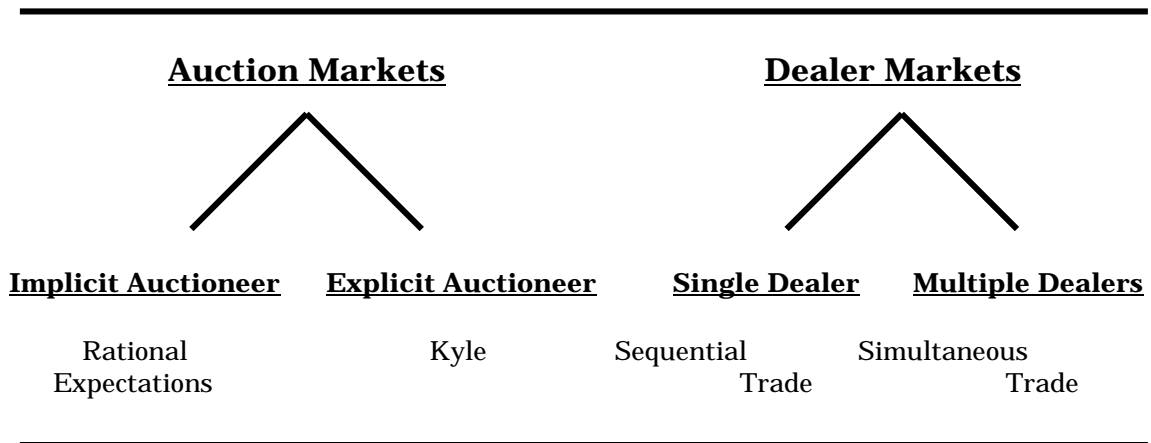


Figure 4.2

Timing of Rational Expectations Model




Figure 4.3

Summary of Rational-Expectations Auction Model

Players:	<ul style="list-style-type: none">• 1 informed trader (risk averse, perfect competitor)• 1 uninformed trader (risk averse, perfect competitor)
Information:	<ul style="list-style-type: none">• final payoff V of risky asset distributed $\text{Normal}(0, \sigma_V^2)$• informed trader has private information about V• uninformed sees only market-clearing price P
Institutions:	<ul style="list-style-type: none">• single trading period• batch clearing: all trades cleared at single price• consistency of conjectured and actual pricing rule

Figure 4.4

The Two Stages of Information Processing




Figure 4.5

Timing of Kyle Model




Figure 4.6

Summary of Kyle Auction Model

Players:	<ul style="list-style-type: none">• 1 marketmaker (risk neutral)• 1 informed trader (risk neutral)• many uninformed traders (non-strategic)
Information:	<ul style="list-style-type: none">• end-of-period value V of risky asset distributed $\text{Normal}(0, \sigma_V^2)$• informed trader sees V• informed trader does not see uninformed orders (but knows distribution)• marketmaker only sees total orders, not the individual traders' components
Institutions:	<ul style="list-style-type: none">• single trading period• batch clearing: all trades cleared at single price• marketmaker pricing such that expected profit equals zero

Figure 4.7

Timing of a Single Trade in Sequential-Trade Model




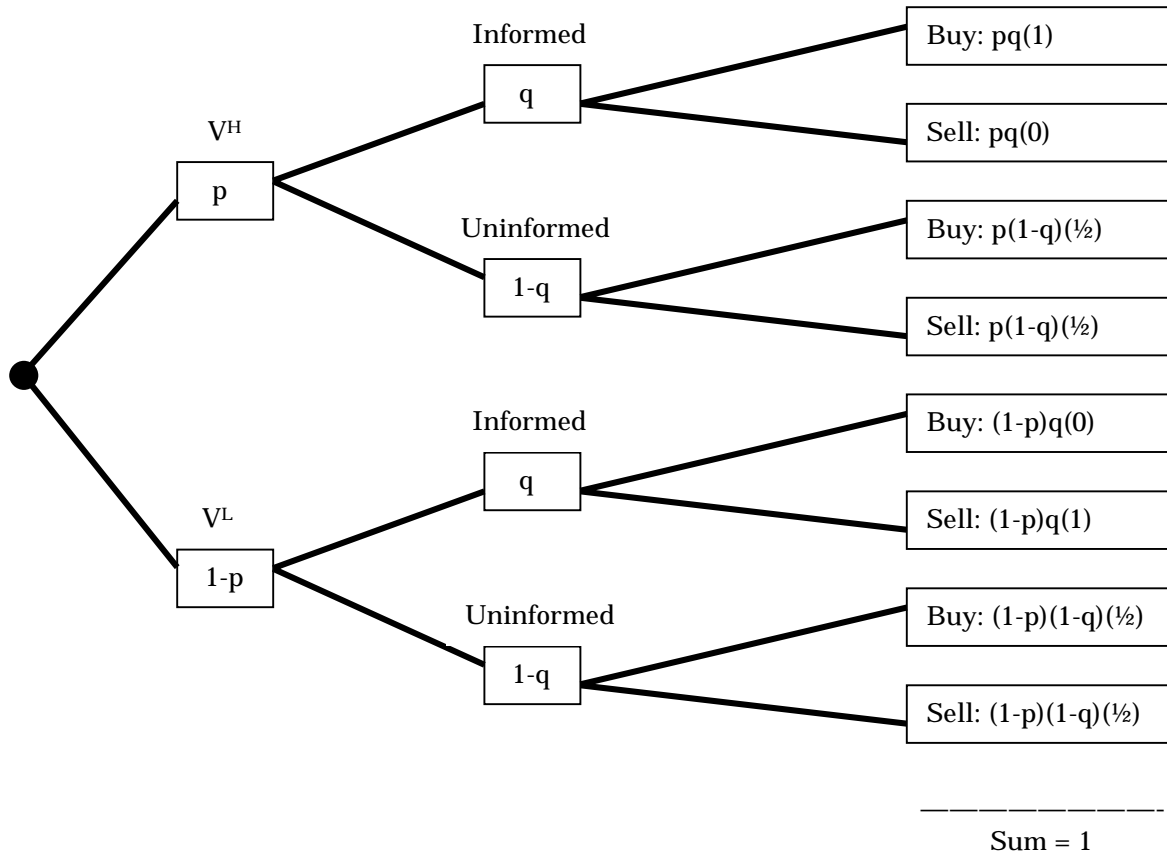
Figure 4.8

Summary of Sequential-Trade Model

Players:	<ul style="list-style-type: none">• 1 dealer (risk neutral)• many informed traders (risk neutral and non-strategic)• many uninformed traders (non-strategic)
Information:	<ul style="list-style-type: none">• terminal value V of risky asset is either high V^H or low V^L• all informed traders know whether value is V^H or V^L• dealer knows unconditional probability of V^H (p)• dealer knows probability that next trader is informed (q)• dealer sees sequence of incoming orders
Institutions:	<ul style="list-style-type: none">• sequence of trading periods, 1 trade maximum per period• dealer participates in all trades• trade size limited to one unit• potential trader randomly selected from pool each period• dealer presents bid and offer price to potential trader• dealer sets prices such that expected profit equals zero

Figure 4.9


Probability of Different Trade Types—Sequential-Trade Model



There are 8 possible trade types and the probability of each occurring appears in the far-right column (the sum of the eight probabilities equals 1). Each of these probabilities has three components. First, nature produces either a high payoff value V^H or a low payoff value V^L , with probabilities p and $(1-p)$ respectively. Then a trader is selected from a pool who is either informed or uninformed, with probabilities q and $(1-q)$ respectively. Informed traders know whether the realized value is V^H or V^L . Finally, the selected trader chooses to buy or sell. If the selected trader is uninformed he buys with probability $\frac{1}{2}$ and sells with probability $\frac{1}{2}$. If the selected trader is informed he buys with probability 1 if payoff value is high and sells with probability 1 if payoff value is low.

Figure 4.10

Timing in the Simultaneous-Trade Model



Notation

- S: common signal received by all dealers.
 - S_i : private signal received by dealer i .
 - P_{i1} : dealer i 's quote in period one.
 - C_i : net customer order received by dealer i .
 - T_{i1} : dealer i 's net outgoing order to other dealers in period one.
 - X : net interdealer order flow in period one.
 - P_{i2} : dealer i 's quote in period two.
 - T_{i2} : dealer i 's net outgoing order to other dealers in period two.
 - V : payoff on the risky asset.
-

Figure 4.11

Summary of Simultaneous Trade Model

Players:	<ul style="list-style-type: none">• a continuum of “customers” (risk-averse and non-strategic)• N dealers (risk-averse and strategic)
Information:	<ul style="list-style-type: none">• terminal value V of risky asset distributed $\text{Normal}(0, \sigma_V^2)$• each dealer receives a signal S_i distributed $\text{Normal}(V, \sigma_{S_i}^2)$• all dealers receive a signal S distributed $\text{Normal}(V, \sigma_S^2)$• each dealer i receives customer orders that aggregate to C_i, distributed $\text{Normal}(0, \sigma_C^2)$• after trading dealers observe a signal of interdealer order flow X
Institutions:	<p><u>Quoting</u></p> <ul style="list-style-type: none">• dealer quoting is simultaneous, independent, and required• quotes are available to all dealers• a quote is a single price at which the dealer agrees to buy and sell any amount <p><u>Trading</u></p> <ul style="list-style-type: none">• trading is simultaneous and independent• trading with multiple partners is feasible

Figure 4.12

Negative exponential utility




Figure 5.1

The Three Data Groupings




Figure 5.2

Example of Dealing 2000-1 Communication


```
From CODE  FULL NAME HERE  *1250GMT 030892 */1080
Our Terminal: CODE  Our user: DMK
      SP DMK 10
# 8891
      BUY

# 10 MIO AGREED
# VAL 6AUG92
# MY DMK TO FULL NAME HERE
# TO CONFIRM AT 1.5891 I SELL 10 MIO USD
#
      TO CONFIRM AT 1.5891 I SELL 10 MIO USD
      VAL 6AUG92
      MY USD TO FULL NAME HERE AC 0-00-00000
      THKS N BIFN
#
#      #END LOCAL#
#
##WRAP UP BY DMK DAMK 1250GMT 3AUG92
#END#
```

The opening word “From” establishes this as an incoming quote request (outgoing quote requests begin with “To”); this information is crucial for signing trades. The caller’s four-digit code and institution name follow; “GMT” denoted Greenwich Mean Time; the date follows, with the day listed first; the “1080” at the end of line one is simply a record number. “SP DMK 10” identifies this as a request for a spot DM/\$ quote for up to \$10 million; “8891” denotes a bid of 88 and an offer of 91. Only the last two digits are quoted because it involves fewer keystrokes; dealers are well aware of the first digits of the price—sometimes called the “handle.” From the confirmation that follows, one can see that the earlier bid quote was in fact 1.5888 DM/\$ and the offer quote was 1.5891 DM/\$. The confirmation also provides the transaction price and verifies the transaction quantity; “THKS N BIFN” is shorthand for “thanks and bye for now.”

Figure 5.3

Diagram of data structure



Definitions: Q^o is an outgoing interdealer quote (i.e., a quote made) and, if the quote is hit, T^i is the incoming direct dealer trade. Q^i is an incoming interdealer quote (i.e., a quote received) and, if the quote is hit, T^o is the outgoing direct trade. T^b is a brokered interdealer trade. Brokered trades do not align vertically with a quote because the data for brokered trades in the Lyons (1995) data set come from the dealer position sheets, and the broker-advertised quotes at the time of the transaction are not recorded. “**|**” appears whenever a trade occurs; “|” appears whenever a non-dealt quote occurs. The disjoint segment below the top time-line presents a hypothetical path of the dealer’s position over the same interval; it changes with trades only. The time-line at the bottom clarifies the definition of “periods” within the Lyons (1995) analysis: incoming trades define an event, not all trades (that model is presented in chapter 5).

Figure 5.4

An illustration of spread components

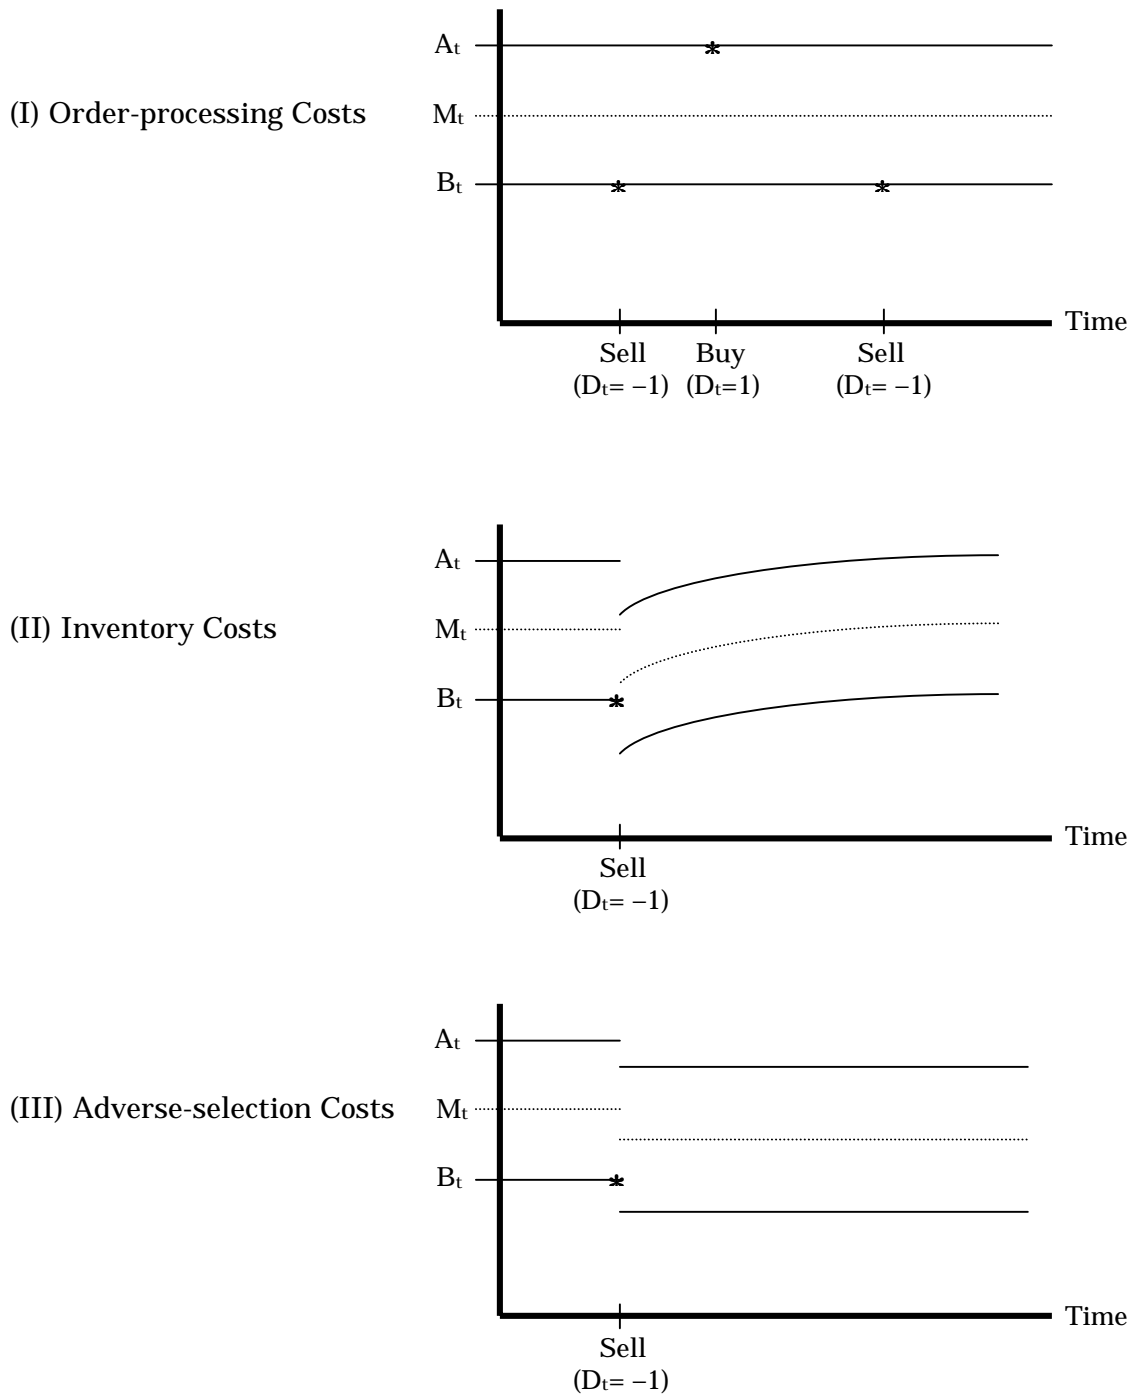


Figure 5.5

Timing in each period of the DP model





Figure 5.6

Dealer i's quote schedule in the DP model



The slope of the quoted price schedule is determined by β_1 , and reflects the information conveyed by order flow X_{jt} . Inventory is a shift variable: the larger I_{it} is relative to the desired position I_i^* , the lower the price schedule throughout (to induce inventory decumulating purchases by counterparties). The bid-offer spread at quantities near zero is pinned down by the parameter γ , which multiplies the direction-indicator variable D_t in the pricing rule of equation (5.18).

Table 5.1

Structural Model Estimates

$$\Delta P_{it} = \beta_0 + \beta_1 X_{it} - \beta_2 I_{it} + \beta_3 I_{it-1} + \beta_4 D_t - \beta_5 D_{t-1} + \beta_6 B_t + \beta_7 v_{it-1} + v_{it}$$

β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	R^2
-1.30 (-0.96)	1.44 (3.10)	-0.98 (-3.59)	0.79 (3.00)	10.15 (4.73)	-8.93 (-6.12)	0.69 (2.21)	-0.09 (-2.55)	0.23
-1.34 (-0.99)	1.40 (3.03)	-0.97 (-3.56)	0.78 (2.95)	10.43 (4.86)	-9.16 (-6.28)		-0.09 (-2.61)	0.22
	>0	<0	>0	>0	<0	>0	<0	

T-statistics in parentheses. The last row indicates the signs predicted by the structural model. ΔP_{it} is the change in the incoming transaction price (DM/\$) from t-1 to t. X_{it} is the incoming order transacted at dealer i's quoted prices, positive for purchases (i.e., effected at the offer) and negative for sales (at the bid). The units of X_{it} are such that $\beta_1=1$ implies an information effect on price of DM0.0001 for every \$10 million. I_t is dealer i's inventory at the end of period t. D_t is than indicator variable with value 1 if the incoming order is a purchase and value -1 if a sale. B_t is the net quantity of third-party brokered trading over the previous two minutes, positive for buyer-initiated trades and negative for seller-initiated trades. All quantity variables are in \$ millions. All coefficients are multiplied by 10^5 . Sample: August 3-7, 1992, 839 observations.

Table 5.2

Testing the hot potato hypothesis:
Is order flow less informative when inter-transaction time is short?

$$\Delta P_{it} = \beta_0 + \beta_1 s_t X_{jt} + \beta'_1 l_t X_{jt} - \beta_2 I_{it} + \beta_3 I_{it-1} + \beta_4 D_t - \beta_5 D_{t-1} + \varepsilon_{it}$$

	β_1 (short)	β'_1 (long)	Fraction short	$\beta_1 = \beta'_1$ P-value
Inter-transaction time short if:				
Less than 1 minute	-0.01 (-0.01)	2.20 (3.84)	262/842	0.000
Less than 2 minutes	0.76 (1.63)	2.60 (3.40)	506/842	0.009

T-statistics in parentheses. The coefficient β_1 measures the information effect of orders for which the time from the previous transaction is short ($s_t=1$ and $l_t=0$ in the equation in the heading), where short is defined in the first column. The coefficient β'_1 measures the information effect of those orders for which the time from the previous transaction is long ($s_t=0$ and $l_t=1$), where long is defined as not short. The “Fraction short” column presents the fraction of observations satisfying the corresponding definition of short inter-transaction times. In each case, the remaining observations fall into the long category. The P-value column presents the significance level at which the null $\beta_1 = \beta'_1$ can just be rejected. ΔP_{it} is the change in the incoming transaction price (DM/\$) from $t-1$ to t . X_{it} is the incoming order transacted at dealer i 's quoted prices, positive for purchases (i.e., effected at the offer) and negative for sales (at the bid). The units of X_{it} are such that $\beta_1=1$ implies an information effect on price of DM0.0001 for every \$10 million. I_t is dealer i 's inventory at the end of period t . D_t is than indicator variable with value 1 if the incoming order is a purchase and value -1 if a sale. Sample: August 3-7, 1992. Estimated using OLS with autocorrelation consistent (first order) standard errors.

Table 5.3

Testing the hot potato hypothesis:
Is order flow less informative when transactions follow in the same direction?

$$\Delta P_{it} = \beta_0 + \beta_1 s_t X_{jt} + \beta'_1 o_t X_{jt} + \beta''_1 l_t X_{jt} - \beta_2 I_{it} + \beta_3 I_{it-1} + \beta_4 D_t - \beta_5 D_{t-1} + \varepsilon_{it}$$

β_1 (short and same)	β'_1 (short and opposite)	β''_1 (long)	Fraction short and same	Fraction short and opposite	$\beta_1=\beta'_1$ P-value
-0.06 (-0.11)	1.90 (3.01)	2.64 (3.46)	276/842	230/842	0.009

T-statistics in parentheses. The coefficient β_1 measures the information effect of orders that have (i) short inter-transaction times, defined as less than the median of two minutes, and (ii) the same sign (direction) as the previous order ($s_t=1$, $o_t=0$, and $l_t=0$ in the equation in the heading). The coefficient β'_1 measures the information effect of orders that have (i) short inter-transaction times, defined as less than the median of two minutes, and (ii) the opposite sign (direction) of the previous order ($s_t=0$, $o_t=1$, and $l_t=0$). The coefficient β''_1 measures the information effect of orders that have long inter-transaction times, defined as greater than or equal to the median of two minutes ($s_t=0$, $o_t=0$, and $l_t=1$). The “Fraction short and same” column presents the fraction of observations satisfying the corresponding definition of short and same (similarly for the “Fraction short and opposite” column). The remaining observations fall into the long category. The P-value column presents the significance level at which the null $\beta_1=\beta'_1$ can just be rejected. ΔP_{it} is the change in the incoming transaction price (DM/\$) from $t-1$ to t . X_{it} is the incoming order transacted at dealer i 's quoted prices, positive for purchases (i.e., effected at the offer) and negative for sales (at the bid). The units of X_{it} are such that $\beta_1=1$ implies an information effect on price of DM0.0001 for every \$10 million. I_t is dealer i 's inventory at the end of period t . D_t is than indicator variable with value 1 if the incoming order is a purchase and value -1 if a sale. Sample: August 3-7, 1992. Estimated using OLS with autocorrelation consistent (first order) standard errors.

Figure 6.1

An illustration of overshooting

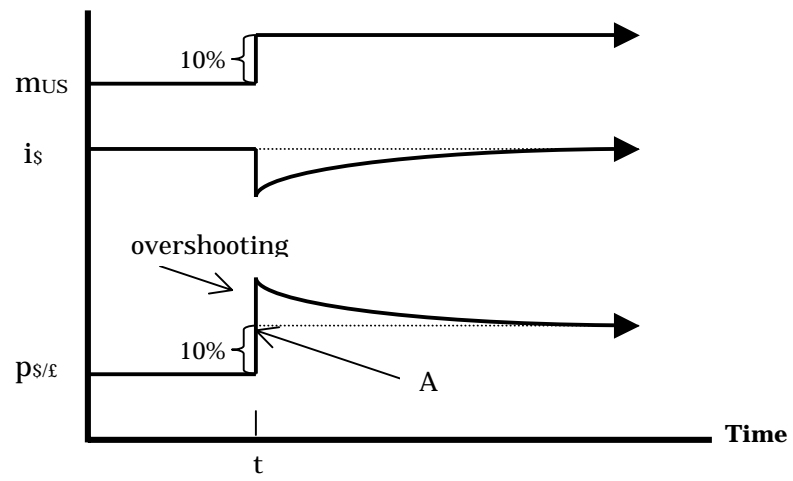


Figure 6.2

Three approaches to exchange rates and their models

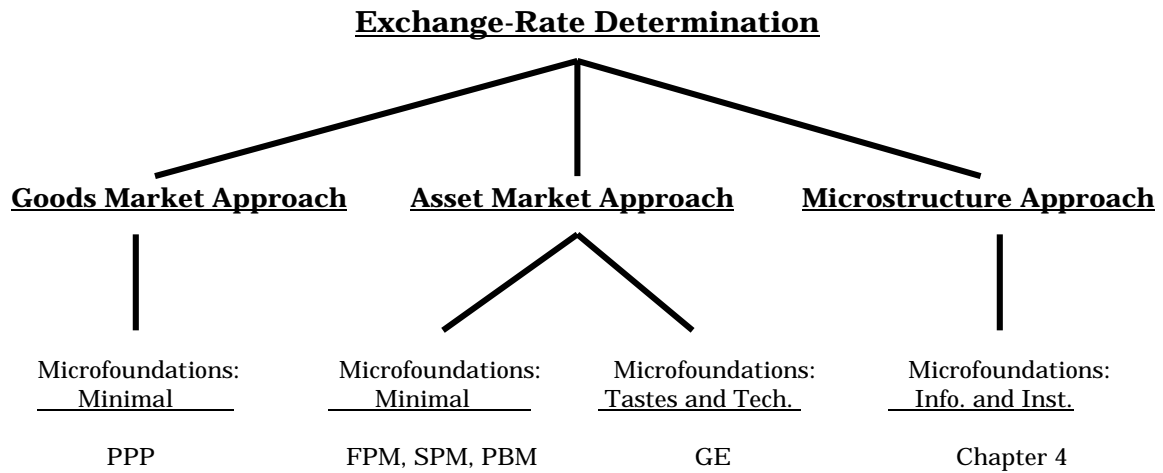


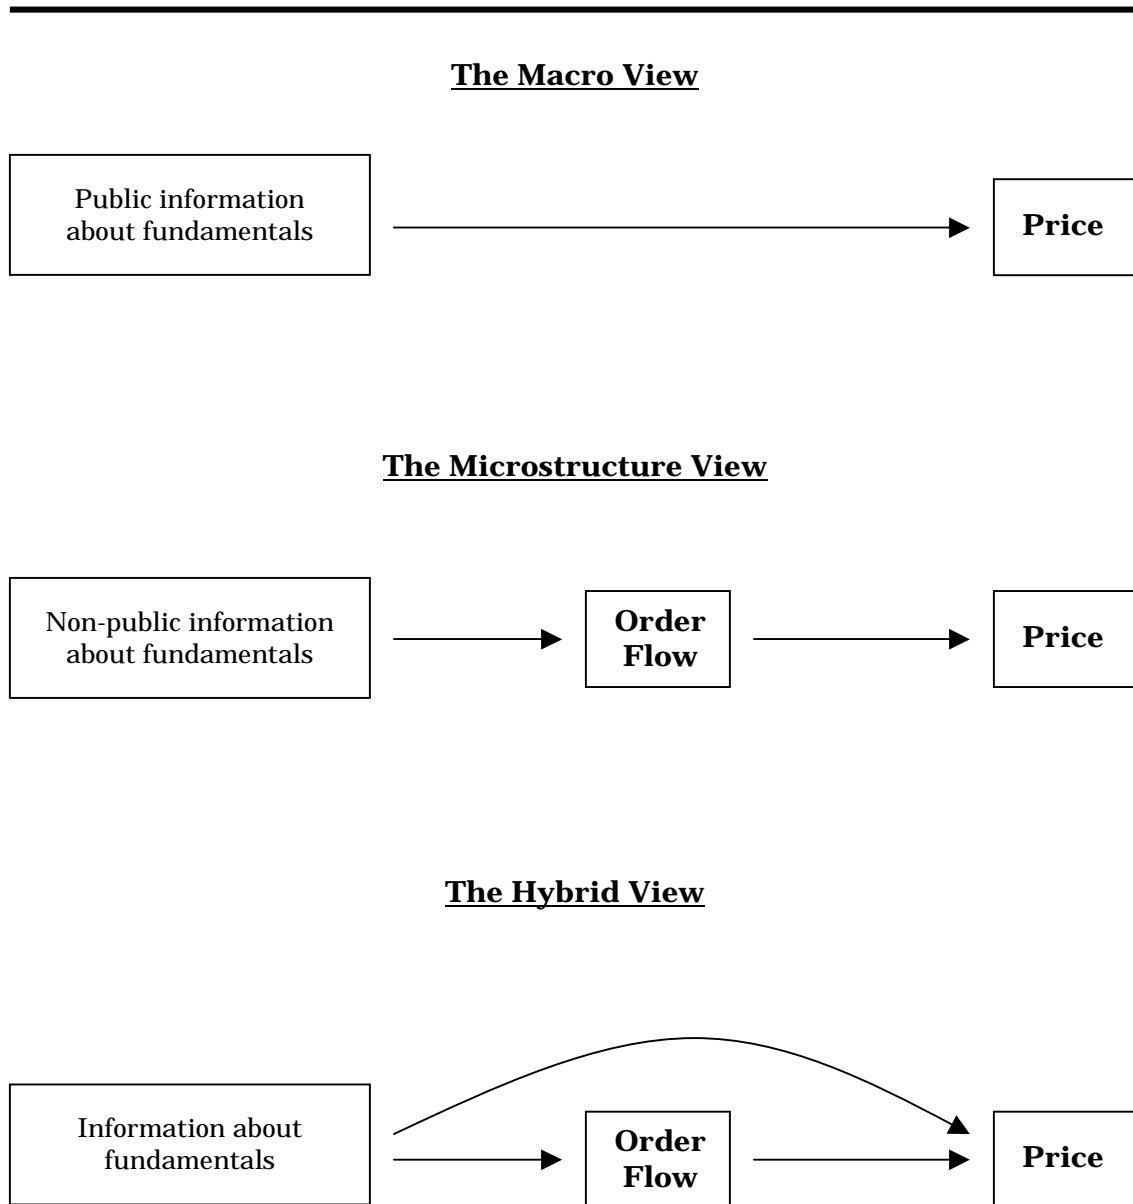
Figure 6.3

The Issues Spectrum



Figure 7.1

Spanning Macro and Microstructure Graphically



The top panel illustrates the connection between fundamentals and price under the traditional macro view (i.e., as reflected in the models of chapter 6): information about fundamentals is public, and so is the mapping to price, so price adjustment is direct and immediate. The middle panel shows the traditional microstructure view (as reflected in the models of chapter 4). The

focus in that case is fundamental information that is not publicly known. This type of information is first transformed into order flow, which becomes a signal to the price setter (e.g., dealer) that price needs to be adjusted. Actual markets include both, which is illustrated in the bottom panel—the hybrid view.

Figure 7.2

Daily Timing in the Evans-Lyons Model




Figure 7.3

Summary of Evans and Lyons (1999) Model

- Players:**
- N dealers (risk-averse and strategic)
 - a continuum of “customers” (risk-averse and non-strategic), whose collective risk-bearing capacity is less than infinite
- Information:**
- terminal value V of risky asset is the sum of daily increments R_t , $t=1, \dots, T+1$, with each R_t distributed $\text{Normal}(0, \sigma_R^2)$
 - all participants observe R_t at the beginning of day t
 - in round one of each day, each dealer i receives a customer order C_{it}^1 , distributed $\text{Normal}(0, \sigma_C^2)$
 - after round two of each day, all dealers observe that day’s interdealer order flow X_t
- Institutions:**
- there are T trading days before V is realized
 - each trading day, there are three rounds of trading: customer-dealer, then interdealer, then another customer-dealer
 - dealers end each day with no net position (i.e., there are no daily frequency inventory effects)

Quoting


- dealer quoting is simultaneous, independent, and required
- quotes are available to all dealers
- a quote is a single price at which the dealer agrees to buy and sell any amount

Trading

- trading is simultaneous and independent
 - trading with multiple partners is feasible
-

Figure 7.4

Portfolio Balance Effects: One Period Example



The market-clearing gap $E[V] - P_0$ is a function of the risky asset's net supply. In traditional portfolio balance models, variation in gross supply is the driver. In the Evans-Lyons model, gross supply is fixed, but net supply is moving over time due to shifts in demand that are unrelated to $E[V] - P_0$. These demand shifts are the exogenous realizations of c_{it}^1 . In contrast to the dissipation of the portfolio-balance effect on price in the one-period example, the price effects do not dissipate in the Evans-Lyons model because payoff uncertainty is resolved smoothly over time.

Table 7.1

Estimates of the Evans-Lyons model

$$\Delta p_t = \beta_1 \Delta(i_t - i_t^*) + \beta_2 X_t + \eta_t$$

	β_1	β_2	R^2
DM	0.52 (1.5)	2.10 (10.5)	0.64
Yen	2.48 (2.7)	2.90 (6.3)	0.45

T-statistics are shown in parentheses. (In the case of the DM equation, the t-statistics are corrected for heteroskedasticity; there is no evidence of heteroskedasticity in the Yen equation, and no evidence of serial correlation in either equation.) The dependent variable Δp_t is the change in the log spot exchange rate from 4 pm GMT on day t-1 to 4 pm GMT on day t (DM/\$ or ¥/\$). The regressor $\Delta(i_t - i_t^*)$ is the change in the one-day interest differential from day t-1 to day t (* denotes DM or ¥, annual basis). The regressor X_t is interdealer order flow between 4 pm GMT on day t-1 and 4 pm GMT on day t (negative for net dollar sales, in thousands of transactions). Estimated using OLS. The sample spans four months (May 1 to August 31, 1996), which is 89 trading days. (Saturday and Sunday order flow—of which there is little—is included in Monday.)

Figure 7.5

Figure 1.

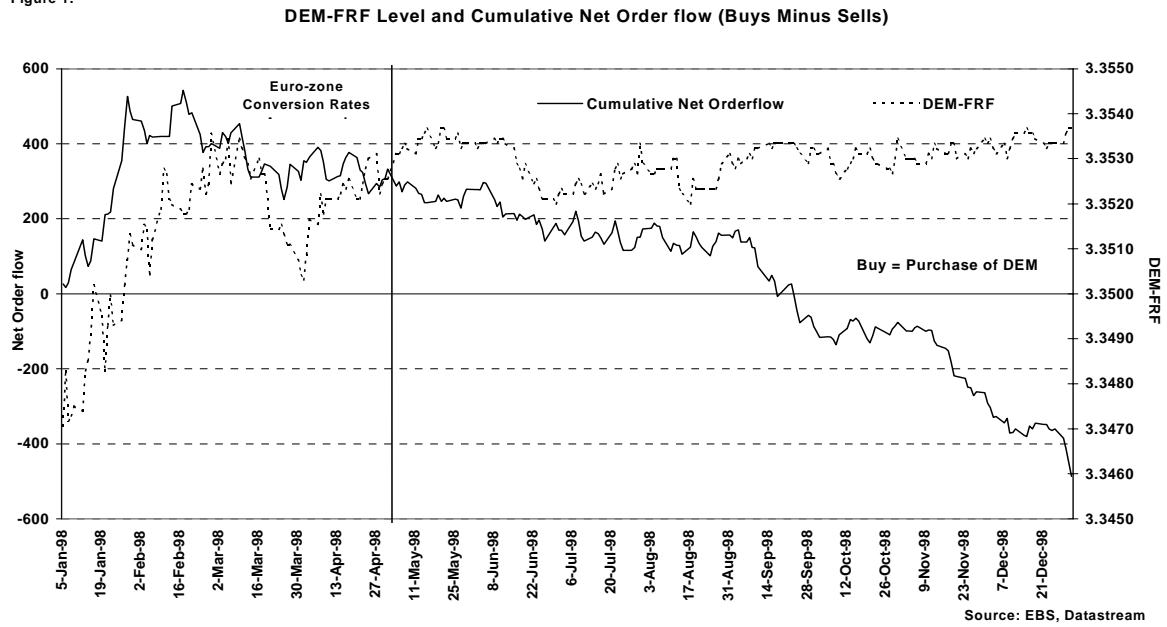
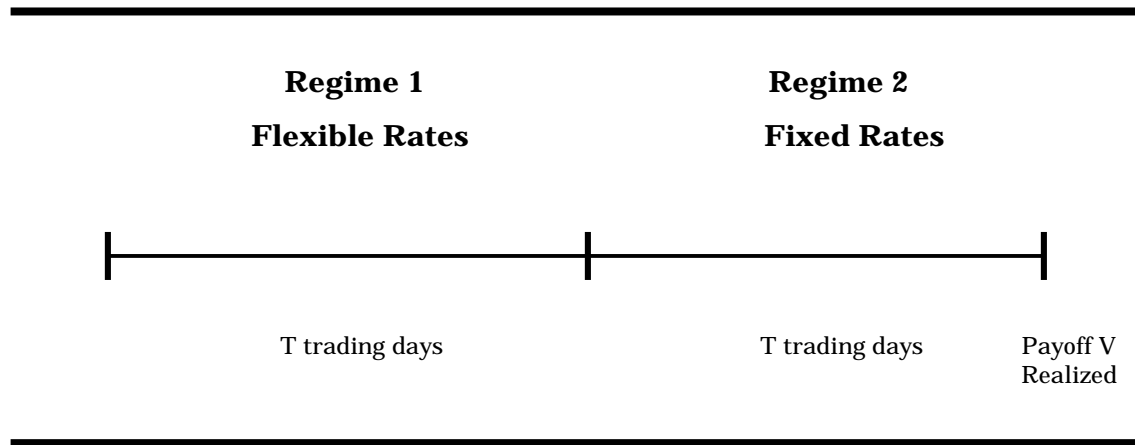


Figure 7.6

The Two Trading Regimes



Under the flexible-rate regime, payoff increments R_t are distributed Normally, with mean zero and variance Σ_R . On the first morning of the fixed-rate regime, the central bank (credibly) commits to pegging the exchange rate at $P_T + k$, where P_T is the previous day's closing price and k is a final draw from the distribution $\text{Normal}(0, \Sigma_R)$. It does so by setting R_{T+1} such that $V = P_T + k$ and maintaining $R_t = 0$ from day $T+2$ to day $2T$.

Figure 7.7

The Statistician's Perspective on Forward Bias




Figure 7.8

The Practitioner's Perspective on Forward Bias




Table 7.2**Sharpe ratios (annual basis) from pure currency strategy**

	Strategy 1: <u>Equal Weighted</u>	Strategy 2: <u>> Median Discount</u>	Strategy 3: <u>< Median Discount</u>
Sharpe Ratio: No Costs	0.48	0.46	0.49
Sharpe Ratio: With Costs	0.37	0.39	0.41

Strategies for profiting from forward bias entail selling foreign currency forward when $f_{t,1} > e_t$ and buying foreign currency forward when $f_{t,1} < e_t$. The three strategies shown are implemented using the six largest currency markets: USD/DEM, USD/JPY, USD/GBP, USD/CHF, USD/FRF, and USD/CAD (see table 3.4). The “Equal Weighted” strategy has an equal position weight each month in each of the six forward markets. The “> Median Discount” strategy only takes a position in a forward market in a given month if that month’s forward discount is greater than the median forward discount for that currency over the sample (weights are equal across forward positions taken). The “< Median Discount” strategy only takes a position in a forward market in a given month if the month’s forward discount is less than the median forward discount for that currency over the sample (weights are equal across forward positions taken). The Sharpe ratio estimate with costs assumes a cost of ten basis points per transaction (includes price impact of trade). The sample is monthly data, from January 1980 to December 1998 (1980 is about the time when forward bias was first documented in the literature). Source: Datastream.

Figure 8.1

Fed balance sheet: Unsterilized purchase of \$100 million with yen

Assets	Liabilities
FXR: ↓ \$100 million DGB: no change	MS: ↓ \$100 million

FXR is foreign-exchange reserves, DGB is domestic government bonds, and MS is money supply—i.e., money in circulation (currency plus monetary base).

Figure 8.2

Fed balance sheet: Sterilized purchase of \$100 million with yen

Assets	Liabilities
FXR: ↓ \$100 million (1)	MS: ↓ \$100 million (1)
DGB: ↑ \$100 million (2)	MS: ↑ \$100 million (2)

FXR is foreign-exchange reserves, DGB is domestic government bonds, and MS is money supply—i.e., money in circulation (currency plus monetary base). Transaction (1) is the unsterilized intervention; transaction (2) is the offsetting sterilization.

Figure 8.3

Types of Intervention

	Sterilized: no signal	Sterilized: signal	Unsterilized
Announced			
Unannounced, but partially revealed			
Secret			

Figure 8.4

Intervention Transparency Spectrum




Figure 9.1

The Trading Volume Pie




Table 9.1

Customer Trades: Volumes and Order Flow


Billions of euros for Euro and billions of dollars for Yen.

	Euro			Yen		
	Total Trading Volume	Cumul. Order Flow	Daily Standard Deviation	Total Trading Volume	Cumul. Order Flow	Daily Standard Deviation
Non-financial Corporations	539	-25.7	0.09	259	3.3	0.07
Leveraged Financial	667	2.5	0.16	681	16.1	0.16
Unleveraged Financial	507	11.8	0.13	604	-1.8	0.15
Total	1,713	-11.4	0.13	1,544	17.6	0.14

Euro denotes the USD/EURO market. Yen denotes the USD/YEN market. The sample for both currencies is January 1993 to June 1999. (Before the launch of the euro in January 1999, volume and order flow are constructed from trading in the euro's constituent currencies.) Positive order flow in the case of the euro denotes net demand for euros (following the convention in that market of quoting prices in dollars per euro). Positive order flow in the case of the yen denotes net demand for dollars (following the convention in that market of quoting prices in yen per dollar). Daily standard deviation measures the standard deviation of daily order flow.

Figure 9.2


Cumulative customer flow and exchange rates



The plots show the spot exchange rate and cumulative customer order flow received by the source bank. The sample for the \$/euro plot is January 1993 to June 1999. The sample for the yen/\$ plot is January 1996 to June 1999 (the January 1993 to December 1995 period is not included due to the lack of Tokyo-office data). The spot exchange rate is expressed on the left-hand scale. The cumulative customer order flow is expressed on the right-hand scale (in billions of euros for the \$/euro plot and in billions of dollars for the yen/\$ plot).

Figure 9.3

Cumulative customer flow and exchange rates over the Evans-Lyons sample



The plot shows the cumulative customer order flow in the USD/EURO market received by the source bank from May 1 to August 31, 1996, and the spot exchange rate over the same period. The spot rate is expressed in DM/\$ on the left-hand scale. The cumulative customer order flow is expressed in millions of euros on the right-hand scale (positive for net dollar purchases).

Table 9.2

The Price Impact of Aggregate Customer Orders

$$\Delta p_t = \beta_0 + \beta_1(\text{Aggregate Customer Flow})_t + \varepsilon_t$$

	β_1	R^2
<u>Monthly Data</u>		
Euro	0.8 (3.7)	0.15
Yen	1.1 (3.6)	0.15

T-statistics are shown in parentheses. The dependent variable Δp_t is the monthly change in the log spot exchange rate (the \$/euro rate and the yen/\$ rate, respectively). The order-flow regressors are measured over the concurrent month (in billions of euros for the euro equation and billions of dollars in the yen equation). Estimated using OLS (standard errors corrected for heteroskedasticity). The sample is January 1993 to June 1999. Constants (not reported) are insignificant in both equations.

Table 9.3

The Price Impact of Disaggregated Customer Orders

$$\Delta p_t = \beta_0 + \beta_1(\text{Unlev. Fin. Flow})_t + \beta_2(\text{Lev. Fin. Flow})_t + \beta_3(\text{Non-fin. Corp. Flow})_t + \varepsilon_t$$

	β_1	β_2	β_3	R^2
<u>Monthly Data</u>				
Euro	1.5 (4.3)	0.6 (1.9)	-0.2 (-0.39)	0.27
Yen	1.1 (2.6)	1.8 (4.5)	-2.3 (-2.8)	0.34

T-statistics are shown in parentheses. The dependent variable Δp_t is the monthly change in the log spot exchange rate. The three order-flow regressors are the order flows from unleveraged financial institutions, leveraged financial institutions, and non-financial corporations. Order flows are measured over the concurrent month (in billions of euros for the euro equation and billions of dollars in the yen equation). Estimated using OLS (standard errors corrected for heteroskedasticity). The sample is January 1993 to June 1999. Constants (not reported) are insignificant in both equations.

Figure 9.4

Panel A: Cumulative total customer flow and the Yen/\$ rate around the October 1998 collapse



Figure 9.4

Panel B: Cumulative flow of leveraged financial institutions
and the Yen/\$ rate around the October 1998 collapse




Figure 9.4

Panel C: Cumulative flow of non-financial corporations
and the Yen/\$ rate around the October 1998 collapse





Figure 9.4

Panel D: Cumulative flow of unleveraged financial institutions
and the Yen/\$ rate around the October 1998 collapse



The four plots show the cumulative customer order flow of various types received by the source bank from September 1 to October 31, 1998, and the spot yen/\$ rate over the same period. The spot rate is expressed in yen/dollar on the left-hand scale. The cumulative customer order flow is expressed in billions of dollars on the right-hand scale.