The Microstructure Approach to Exchange Rates

MIT Press, forthcoming

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

4 December 2000

Contents

Preface

This book addresses a growing interaction between two fields—exchange-rate economics and microstructure finance. Historically, these fields have progressed independently of one another. More recently, however, they have begun to interact, and that interaction has stimulated new perspective within both. In exchange-rate economics, that new perspective has given rise to a distinct approach—a microstructure approach.

A natural audience for this book includes people with expertise in only one of these two fields who are interested in learning more about the other. In the past this has not been easy—these areas have different intellectual traditions, so spanning them required a large investment of time. This book should facilitate that investment. To that end, I have tried to present material on both fields in an accessible way. By making both accessible, readers lacking specialization in either one should find the material within reach. Practitioners, too, should find most of the material accessible (the empirical work in chapters 7 through 9 are most relevant to that audience).

A notable feature of this book is its treatment of both theoretical and empirical work (c.f., O'Hara's 1995 theory text). Indeed, the book is organized to highlight their interplay. Chapter 4 is dedicated to the canonical frameworks in microstructure theory. Some of these frameworks were originally developed to address the New York Stock Exchange, a market with a single dealer in each stock. The circumstances under which they are appropriate for foreign exchange markets (FX) is an important topic addressed in that chapter. The following chapter, chapter 5, summarizes empirical microstructure frameworks, with emphasis on those employed in FX. These empirical frameworks draw heavily on chapter 4's theory. The last four chapters, chapters 7-10, are applications of tools presented in the first six. They bring the theoretical/empirical interplay into sharp focus.

A second notable feature of this book is its emphasis on information economics rather than institutions. People unfamiliar with microstructure finance typically believe its focus is institutional (e.g., the consequences of different trading mechanisms). This is an important part of the field. But in terms of this book, it is not the most important part. The focus of this book is the economics of financial information, and how microstructure tools clarify the types of information relevant to exchange rates. In keeping with this focus, I move immediately in chapter 2 to the economics of financial information, saving detailed institutional material on FX trading until chapter 3.

This book also has some notable features in terms of pedagogy. They are concentrated in chapters 4-6, the three survey chapters: microstructure theory, microstructure empirical frameworks, and exchange-rate theory, respectively. In chapter 4, I begin the survey of microstructure models with the “rational expectations” (RE) model. The chapter proceeds from the RE model's Walrasian auctioneer to the explicit auctioneer of the Kyle model. Conceptually, I consider this a more natural progression than the typical sequencing in microstructure books. To maintain focus on essentials, chapter 4 also presents the microstructure models in a common format. I close chapter 4 with an appendix that collects several important technical tools for easy reference. On the empirical side, chapter 5 opens with a survey of available data sets and information on Internet access to them (several are available from my web site: www.haas.berkeley.edu/~lyons). Data quality has increased substantially in recent years (in large part due to the shift to electronic trading), and this survey provides an easy reference for familiarizing oneself. Chapter 6 is the transition from the earlier, more micro-oriented chapters, to the later, more macro-oriented chapters. That chapter begins its survey of exchange-rate theory with a treatment of "valuation," which includes traditional definitions of exchange-rate "fundamentals." This should help readers whose perspective on valuation comes primarily from the dividend-discount models used in equity markets. Chapter 6 also highlights the parallels between exchange-rate models and microstructure models (such as the conceptual link between portfolio-balance models and inventory models).

This book also contains new work. For example, some of the material in chapters 7 and 8 is new to the literature, and most of the material in chapter 9 is new. Many of the arguments and points made throughout other chapters are new as well (e.g., perspectives on the microstructure approach's potential applications). Much of the material in chapters 7 to 9 comes from work I have done jointly with three co-authors: Martin Evans, Mintao Fan, and Michael Moore.

For teaching, there are three types of courses for which this book would be appropriate. The first is a PhD course in international macroeconomics. Of the ten chapters, the ones most appropriate for this type of course are chapters 1, 4, 6, 7, 8, and 9. A second course for which the book would be appropriate is a PhD course in market microstructure. The most appropriate chapters for this type of course are 1, 2, 3, 4, 5, and 10. A third course for which the book would be appropriate is a Masters course in international finance. Appropriate chapters for a course at this level include 1, 3, 6, 7, and 9.

Some caveats are in order. Like microstructure theory in general, my treatment of the microstructure approach emphasizes the role of order flow (signed transaction quantities) in price determination. This emphasis aligns closely with my own interests and my own work. Though much of the material I present is not my own, my work plays a more prominent role than it would in a balanced survey. To borrow the words of one reviewer, this book is more of a personal synthesis than a professional consensus. I apologize in advance to those whose work is underrepresented. (For those interested in a survey, let me recommend Sarno and Taylor 2000.)

A second caveat is my use of the term "microstructure approach." Though there is no consensus definition of what constitutes a microstructure model, it is safe to say that I adopt a rather broad definition. For me, the microstructure approach is not just a rich set of tools for addressing the issues, but also a way of framing those issues. Indeed, the framing *per se* is an important aspect of the approach's value. If I have done my job, the chapters that follow should make this point increasingly clear.

To these previous caveats I must add another. The tone of parts of this book may to some readers seem a bit missionary. This is particularly true of chapter 1, where I address the book's overarching themes. (Exchange-rate economists will find these themes more provocative than will people in microstructure finance.) At the risk of appearing to over-sell, I felt it best to address these themes early. Most of the evidence that supports these themes appears in later chapters.

I owe thanks to many people. The support and encouragement of Victoria Richardson and Elizabeth Murry at MIT Press are much appreciated. I also want to thank several of my colleagues who, as anonymous reviewers of early drafts, provided valuable comments. I am indebted to them. I am also indebted to Mike

Dennis, Petra Geraats, Philipp Hartmann, Jen Lyons, Michael Melvin, Michael Moore, Richard Portes, Hélène Rey, Dagfinn Rime, Andrew Rose, Patrik Safvenblad, and Mark Taranto for extensive comments and suggestions. A friend and colleague, Michael Klein, planted the seed for this book, and also provided many insights along the way. I thank the University of Toulouse and Stockholm University for providing generous visiting opportunities while I was writing. Financial support from the National Science Foundation is gratefully acknowledged; this book (and much of the research underlying it) is in large part the product of their investment. Finally, I'd like to thank Jim McCarthy, a long-time buddy and FX trader who introduced me to the world of FX trading.

Chapter 1: Overview of the Microstructure Approach

A friend of mine who trades spot foreign exchange for a large bank invited me to spend a few days at his side. That was ten years ago. At the time I considered myself an expert, having written my thesis on exchange rates. I thought I had a handle on how it worked. I thought wrong. As I sat there my friend traded furiously, all day long, racking up over $1 billion in trades each day ($U.S.). This was a world where the standard trade was $10 million, and a $1 million trade was a "skinny one." Despite my belief that exchange rates depend on macroeconomics, only rarely was news of this type his primary concern. Most of the time he was reading tea leaves that were, at least to me, not so clear. The pace was furious—a quote every 5 or 10 seconds, a trade every minute or two, and continual decisions about what position to hold. Needless to say, there was little time for chat. It was clear my understanding was incomplete when he looked over, in the midst of his fury, and asked me "What should I do?" I laughed. Nervously.

This book is an outgrowth of my subsequent interest in this area. It is principally concerned with the gap between what I knew before I sat down with my friend, and what I saw when I got there. In effect, this gap is the space between two fields of scholarship: exchange-rate economics on the one hand, and microstructure finance on the other. Exchange-rate economists use models of exchange-rate determination that are macroeconomic (i.e., rates are determined as a function of macro variables such as inflation, output, interest rates, etc.). These same exchange-rate economists are largely unfamiliar with microstructure models. Most microstructure scholars, in contrast, view foreign exchange as the purview of international economists, and are unfamiliar with macroeconomic exchange-rate models. Their traditional focus is the microeconomics of equity markets, particularly the New York Stock Exchange.

Though this book has several objectives, two deserve mention at the outset. The first is to lower entry barriers faced by scholars interested in this burgeoning area. Lowering barriers on both sides—exchange-rate economics and microstructure finance—will help this research domain realize its potential. A second objective is to channel past work into a more unified approach—a microstructure approach. In the 1990s, many authors have applied microstructure tools to foreign exchange (FX) markets. But existing work is still largely fragmented.

Does exchange-rate economics need a new approach? Yes. Exchange-rate economics is in crisis. It is in crisis in the sense that current macroeconomic approaches to exchange rates are empirical failures. In their recent survey in the *Handbook of International Economics*, Jeffrey Frankel and Andrew Rose (1995) put it this way:

> To repeat a central fact of life, there is remarkably little evidence that macroeconomic variables have consistent strong effects on floating exchange rates, except during extraordinary circumstances such as hyperinflations. Such negative findings have led the profession to a certain degree of pessimism vis-à-vis exchange-rate research.

In the end, it is my hope that his book might rouse a little optimism.

1.1 Three Approaches to FX: Goods, Assets, and Microstructure

Before the 1970s, the dominant approach to exchange-rate determination was the **goods market approach**. Under this approach, demand for currencies comes primarily from purchases and sales of goods. For example, an increase in exports increases foreign demand for domestic currency to pay for those exported goods. In this simple form, the implication is rather intuitive: countries with trade surpluses experience appreciation (which comes from the currency demand created by the surplus). Despite the approach's intuitive appeal, however, it fails miserably when one looks at the data: trade balances are virtually uncorrelated with exchange-rate movements in major-currency FX markets. This negative result is perhaps not surprising given that trade in goods and services accounts for only a small share of currency trading—less than 5 percent of the $1.5 trillion of FX traded daily (on average).

In the 1970s the **asset market approach** emerged. It built on the earlier approach by recognizing that currency demand comes not only from purchases and sales of goods, but also from purchases and sales of assets. For example, in order to purchase a Japanese government bond, a US investor first purchases the necessary yen. Going one level deeper, the investor's dollar return will depend on movements in the yen, so his demand for the bond depends in part on his desire to speculate on those currency movements. This shift in perspective brought a shift in modeling strategy. In particular, models began to conform to the notion of speculative "efficiency": exchange rates were modeled as efficient in that they incorporate all publicly available information, making public information useless for producing excess returns. This is a feature the goods market approach did not share.[1]

Disconcertingly, empirical work does not support the asset market approach either. The macroeconomic variables that underlie the approach do not move exchange rates as predicted. The classic reference is Meese and Rogoff (1983); they show that asset-approach models fail to forecast major-currency exchange rates better than a simple "no change" model, even when forecasts are allowed to use macro information that is realized over the forecasting period (up to one year). Thus, asset-approach models are not even consistently getting the direction right. In his later survey, Meese (1990) summarizes by writing, "The proportion of (monthly or quarterly) exchange rate changes that current models can explain is essentially zero." (The literature documenting this poor empirical performance is vast; for surveys see Frankel and Rose 1995, Isard 1995, and Taylor 1995.)

The FX market's enormous trading volume is also problematic for the asset approach. Explaining volume is difficult because trading is awarded no role in mapping macroeconomic variables into exchange-rate behavior. Rather, because all macroeconomic news is publicly available, when news occurs the exchange rate is presumed to jump to the new consensus level; the change in expectations that causes the jump does not require any trading. Differing beliefs is not a driver of trading under this approach either because the approach assumes homogeneous

---

[1] Another approach in exchange rate economics—the "flow" approach—is a variant of the goods market approach, so I do not include it separately in my three-approach taxonomy. In addition to currency demand from goods flows, the flow approach includes currency demand from the other main category of a country's balance-of-payments—the capital account. Like its goods-market cousin, the flow approach does not typically model the exchange rate in a way consistent with financial-market efficiency. I return to the parallels between the flow and microstructure approaches in chapter 7.

beliefs. These negative observations do not imply that the asset market approach is "wrong"; indeed, most agree that it is, in broad terms, appropriate. Rather, it appears the approach is missing some key features—features that matter for how exchange rates are actually determined.

This book presents a new approach to exchange rates, the **microstructure approach**. Under this approach, like the asset market approach, the demand for currencies comes from purchases and sales of assets. In this sense these approaches are complementary, not competing. What distinguishes the microstructure approach is that it relaxes three of the asset approach's most uncomfortable assumptions:[2]

(1) **Information**: microstructure models recognize that some information relevant to exchange rates is not publicly available.
(2) **Players**: microstructure models recognize that market participants differ in ways that affect prices.
(3) **Institutions**: microstructure models recognize that trading mechanisms differ in ways that affect prices.

People unfamiliar with microstructure believe its focus is on the third of these—the consequences of different trading mechanisms. *The focus of this book is resolutely on the first—the information economics.* (In keeping with this focus, the next chapter moves immediately to the economics of financial information; material on trader heterogeneity and trading mechanisms—points 2 and 3—is in chapter 3.)[3] Empirically, it is simply not true that all information used to determine market-clearing exchange rates is publicly available. The consequences of this can be analyzed—theoretically and empirically—using tools within the microstructure approach. The resulting analysis shows that the public-information assumption is not a good one: much of exchange-rate determination is missed.

Consider some examples that suggest that the microstructure approach is on target with respect to these three assumptions. Regarding non-public information, FX traders at banks regularly see trades that are not publicly observable. As I show in later chapters, this information forecasts subsequent exchange rates (e.g., seeing the demands of private participants or central banks before the rest of the market). Regarding differences across market participants, it is the regular occurrence that traders with common information will interpret that information differently. Another example of differences across participants is motives for trade: some traders are primarily hedgers whereas others are primarily speculators. Regarding trading mechanisms, consider a market where transparency is low, i.e., where trading outcomes are not generally observable (such as individual transaction sizes and prices). This can slow the updating of beliefs about appropriate prices.

From these examples one might describe the microstructure approach as taking an in-the-trenches, trading-room perspective. Given that exchange rates are actually determined in the trading room, this would seem a reasonable starting

[2] The field of macroeconomics is also moving in the direction of relaxing these three assumptions. See, for example, the literatures on asymmetric information (e.g., Gordon and Bovenberg 1996 and Morris and Shin 2000), non-representative-agent macro (e.g., Caballero, Engel, and Haltiwanger 1997), microfoundations of monetary policy (e.g., Rotemberg and Woodford 1997), and new open-economy macro (Obstfeld and Rogoff, 1996).

[3] I continue with this three-point characterization in chapter 4 (which reviews microstructure theory) by providing model summaries in three parts: Information, Players, and Institutions.

point. But can relaxing the three assumptions above help us understand exchange rates? Relaxing the corresponding assumptions for other asset classes has certainly deepened our understanding of these other markets. The final judgment, though, will be based on specific applications of microstructure tools. The latter half of this book presents a number of applications.

In advance of those applications, I urge the reader to bear in mind the overarching fact that traditional approaches are not consistent with the data. Indeed, this fact induces Flood and Taylor (1996) to conclude that:

> Given the exhaustive interrogation of the macro fundamentals in this respect over the last twenty years, it would seem that our understanding of the short-run behavior of exchange rates is unlikely to be further enhanced by further examination of the macro fundamentals. And it is in this context that new work on the microstructure of the foreign exchange market seems both warranted and promising.

1.2 Hallmarks of the Microstructure Approach

The previous section introduces microstructure in the context of exchange rates but does not define the term as used in domestic finance. Maureen O'Hara (1995) defines market microstructure as "the process and outcomes of exchanging assets under explicit trading rules." The definition I adopt here is consistent with hers. Because her definition is so broad, though, it may be helpful to clarify further.

When one moves from a macro approach to a micro approach, two variables that play no role in the macro approach take center stage. These variables are hallmarks of the micro approach. As hallmarks, they help to define microstructure. These variables are:

(1) **order flow**
(2) **spreads** (bid-ask)

If I labeled these "quantity" and "price" it would be clear that they are the old mainsprings of economics after all. These labels are a bit facile though. Describing them as quantity and price viewed through a magnifying glass is nearer the truth. Let me clarify by touching on each.

<u>Order Flow</u>

Understanding order flow is essential for appreciating how the microstructure approach to exchange rates departs from earlier approaches. First, it is important to recognize that transaction volume and order flow are not the same. Rather, order flow is transaction volume that is *signed*. For example, if you approach a dealer (marketmaker), and you decide to sell the dealer 10 units (shares, euros, etc.), then transaction volume is 10, but order flow is –10. You as the initiator of this transaction are on the sell side, so order flow takes on a negative sign. The quoting dealer is on the passive side of the trade. The trade is signed according to the active, or initiating side. Over time, order flow can then be measured as the sum of the signed buyer-initiated and seller-initiated orders. A negative sum means net selling pressure over the period.

This definition needs to be adjusted slightly for markets that do not have dealers. Some financial markets replace dealers with something known as a "limit order book."[4] Here is an example of a limit order: "buy 10 units for me when the market reaches a price of 50." Anyone can submit these limit orders. They are collected together in an electronic "book." The most competitive of the orders in the book define the best available bid and offer prices. For example, the limit order to buy with the highest buying price becomes the best bid in the market. If you entered the market, and wanted to sell 10 units immediately, you could sell at this best bid price, but no higher. (Think of these best limit orders as analogous to dealer bid and offer quotes in markets that have dealers.) The limit orders are the passive side of any transaction, just as the quoting dealer is always on the passive side in the example of the previous paragraph. When orders arrive that require immediate execution (e.g., an order to "sell 10 units now at the best available price), these orders—called market orders—generate the signed order flow. In the example above, as in the case of the previous paragraph, the execution of the market order to sell 10 units produces an order flow of –10.

Order flow, as used in microstructure finance, is a variant of a key term in economics, "excess demand." I refer to it as a variant of excess demand rather than a synonym because, in equilibrium, order flow does not necessarily equal zero. (By definition excess demand equals zero in equilibrium—there are two sides to every transaction.) In markets organized like foreign exchange, orders are initiated against a marketmaker. The marketmaker, if properly compensated, stands ready to absorb imbalances between buyers and sellers. These "uninitiated" trades of the marketmaker account for the wedge between the two concepts, excess demand and order flow.

A distinctive feature of microstructure models, across the board, is the central role played by order flow. This across-the-board property deserves emphasis because it expands the applicability of microstructure enormously. Recall that order flow plays no role in the asset approach. As viewed from the asset approach, then, order flow is not a variable that helps us understand exchange rates. That microstructure models across the board tell us this variable is important expands microstructure from the narrow concept of "institutional structures with price effects" to the broader concept of "a new lens for viewing markets." It instructs us that order flow deserves our attention. The question of order flow's importance in FX is distinct from—and in my judgment much larger than—the question of how specific FX institutions affect prices.

Consider a simple diagram that illustrates an important feature of microstructure models that relates directly to order flow. The diagram shows that information processing has two stages. The first stage is the analysis or observation of fundamentals by non-dealer market participants (mutual funds, hedge funds, individuals with special information, etc.). The second stage is the dealer's—i.e., the price setter's—interpretation of the first-stage analysis. The dealer's interpretation comes from reading the order flow. Dealers set price on the basis of this interpretation.[5]

---

[4] Markets with limit-order books include the Paris and Hong Kong stock exchanges, both of which are operated electronically. Chapter 3 provides much more detail along these lines.

[5] That dealers learn only from order flow is too extreme. It arises in standard microstructure models because these models assume that all information is private. Macro models go to the other extreme,

Figure 1.1

The Two Stages of Information Processing

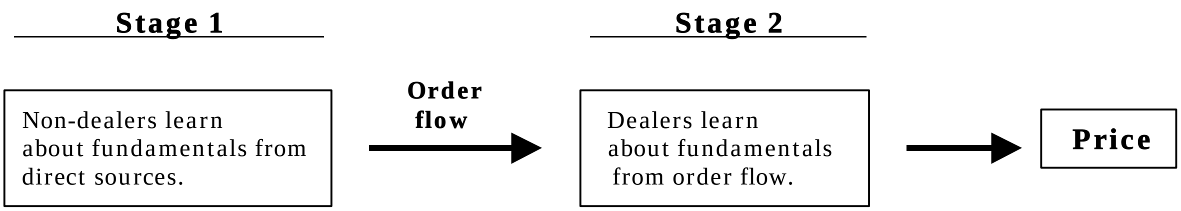

Order flow conveys information about fundamentals because it contains the trades of those who analyze fundamentals. In this sense it is a transmission mechanism. Naturally, though, these informative trades may be mixed with uninformative trades, making the task of "vote counting" rather complex. In standard microstructure models, the dealer learns nothing about fundamentals that she does not learn from order flow. As a practical matter, this is clearly too strong. The dealer's dependence on learning from order flow arises in these models because the information being learned is not publicly known. When information is publicly known, dealers do not need to learn from order flow. In practice, though some information relevant to FX is publicly known, some is not, so learning from order flow can be important.

<u>Spreads</u>

Spreads—the second hallmark variable of the micro approach—receive a lot of attention within the field of microstructure. There are three reasons for this, one scientific, one practical, and one historical. The scientific reason relates to data. Spread data are a core element of most data sets. As such, spreads are a ready target for testable hypotheses. This stands in contrast to other features of the trading environment that are important, but not as readily measurable (such as the flow of information, the dispersion of beliefs, and the flow of liquidity-motivated orders). The second reason spreads receive so much attention is practical. Practitioners are intensely concerned with the management of trading costs. This concern, and the resources devoted to it, has naturally influenced the course of research within microstructure. The third reason spreads receive so much attention is historical. From the beginning, the field of microstructure sought to separate itself from the literature on trading under rational expectations. Rational-expectations models abstract from trading mechanisms completely, the premise being that trading mechanisms have little effect on the relationship between underlying fundamentals and price. A natural means of distinguishing microstructure research was to orient it toward the following question: How does altering the trading

---

assuming all information is public. I return to this in chapter 7, where I present an empirical model that admits both public and private information.

mechanism alter price? This orientation led to a focus on the determination of real-world transaction prices—i.e., spreads.[6]

Though spreads receive a lot of attention, the sub-field of spread determination is but one branch of the broader field of microstructure. Many microstructure models, for example, do not even include a spread (such as Kyle models, presented in chapter 4). I raise the issue early because for some people the association between spreads and microstructure is very tight. That many microstructure models have no spread should loosen this association.

Of the two hallmarks—order flow and spreads—this book focuses much more on order flow. In those instances where I do address spreads, my focus is on their information-theoretic implications. These implications can be substantial. For example, in chapter 2 I present evidence—gleaned from spread behavior alone—that implies that order flow forecasts price movements. This is important because it violates the premise of traditional models that all FX market participants are equally well informed: some participants are better informed because they observe more order flow. Thus, on the basis of spread behavior alone, one reaches a rather profound conclusion: contrary to the asset approach, exchange-rate determination is not wholly a function of public news.[7]

1.3 Overarching Themes

Readers familiar with exchange-rate economics are unlikely to be familiar with microstructure. For these readers, it may be helpful at the outset to address some of this book's overarching themes. Introducing them early provides valuable advance perspective on applications in later chapters. There are four themes in particular that exchange-rate economists may find provocative.

<u>Theme 1: Order flow matters for exchange-rate determination.</u>

Let me offer some perspective on this assertion since it is one I will revisit. Consider an example, one that clarifies how economist and practitioner worldviews differ. The example is the timeworn reasoning used by practitioners to account for price movements. In the case of a price increase, practitioners will assert, "there were more buyers than sellers." Like other economists, I smile when I hear this. I smile because in my mind the expression is tantamount to "price had to rise to balance demand and supply." These phrases may not be equivalent however. For economists, the phrase "price had to rise to balance demand and supply" calls to mind the **Walrasian auctioneer**. The Walrasian auctioneer is an abstract way to

---

[6] My own view is that this central question—How does altering the trading mechanism alter price?—has unwittingly narrowed the scope of questions to which microstructure tools have been applied. I return to this issue in section 3.4, after more groundwork is laid.

[7] Regarding spreads, let me communicate an experience that may speak to the reader. I recall years ago presenting a paper on exchange rates that was firmly embedded in the macro-orientation of the asset approach. Someone in the audience asked me a question about bid-ask spreads. What went through my mind—though not my lips—was "I couldn't care less about bid-ask spreads." To me, at the time, bid-ask spreads were simply a nuisance parameter, bad manners next to an otherwise elegant approach. At the time, spreads meant microstructure to me, and microstructure meant spreads. I no longer subscribe to this view.

think about how price adjusts to a market-clearing level. The Walrasian auctioneer collects “preliminary” orders, which he uses to find the market-clearing price. All actual trades, however, occur at this price—no trading occurs in the process of finding it. (Readers familiar with the rational-expectations model of trading will recognize that in that model, this property is manifested by all orders being conditioned on a market-clearing price.)[8]

Practitioners seem to have a different model in mind. In the practitioners’ model there is a dealer instead of an abstract auctioneer. The dealer acts as a buffer between buyers and sellers. The orders the dealer collects are actual orders, rather than preliminary orders, so trading does occur in the transition to the new price.[9] Crucially, the dealer then determines new prices from information about demand and supply that is embedded in the order flow (as suggested in the “two-stage processing” diagram above).

Can the practitioner model be rationalized? Not at first blush, because it appears that trades are occurring at disequilibrium prices (prices at which the Walrasian auctioneer would not allow trading). This suggests irrational behavior. But this misses an important piece of the puzzle: whether these trades are out-of-equilibrium depends on the information available to the dealer. If the dealer knows at the outset that there are more buyers than sellers, eventually pushing price up, then it is unclear why that dealer would sell at a low interim price. If the buyer/seller imbalance is not known, however, then rational trades can occur through the transition. (Put differently, in setting prices the dealer cannot condition on all the information available to the Walrasian auctioneer.) This is precisely the story developed in standard microstructure models. Trading that would be irrational if the dealer knew as much as the Walrasian auctioneer can be rationalized in models with more limited—and more realistic—conditioning information.

<u>Theme 2: Microstructure implications can be long-lived</u>.

It is common to associate “microstructure” with “high frequency.” The association is natural, but deceptive. It is true that empirical work in microstructure is in general high frequency. But this does not imply that microstructure tools are irrelevant to lower-frequency, resource relevant phenomena. Indeed, there are ample tools within the microstructure approach for addressing lower-frequency phenomena. And new tools continue to emerge, thanks in part to recognition within the broader microstructure literature that resource allocation warrants greater attention. The later chapters of this book are dedicated to examples of lower-frequency relevance, particularly chapter 7.

Regarding long-lived effects, the most important point to recognize is that when order flow conveys information, its effect on price *should* be long-lived. Indeed, a common empirical assumption for distinguishing information from pricing errors is that information’s effects on price are permanent, whereas pricing errors are transitory (French and Roll 1986, Hasbrouck 1991; chapter 2 provides details).

---

[8] Walrasian-style mechanisms do exist in real-world financial markets: they are often used to open trading (see, e.g., work on the pre-opening of the Paris Bourse by Biais, Hillion, and Spatt 1999). Indicative orders are collected for a period before actual trading begins, and these orders are used to determine the opening price (the orders become executable at the opening, but typically can be retracted before the opening).

[9] Actual orders come in different types. These different types are covered in chapter 3.

These long-lived effects are borne out in the data, in equity markets, bond markets, and FX markets. In FX, for example, Evans (1997), Evans and Lyons (1999), Payne (1999), and Rime (2000) show that order flow has significant effects on exchange rates that persist. Indeed, statistically these effects appear to be permanent. Among microstructure's long-lived implications, this "information" channel is definitely the most fundamental.

Let me touch on another source of lower-frequency relevance: multiple equilibria that depend on microstructure parameters. Certain parameters' values can determine whether multiple equilibria within a given model are possible, and if so, which are more likely to be selected (e.g., Portes and Rey 1998, Hau 1998, Hartmann 1998a, and Rey 1999). These different equilibria apply in some models to the exchange rate's level, and in other models to the exchange rate's volatility (multiple volatility states are the focus of Jeanne and Rose 1999 and Spiegel 1998). Either way, multiple equilibria that depend on microstructure parameters open the door to price effects that are long-lived (long-lived because a given equilibrium is by nature persistent).[10]

An analogy may be helpful. Microstructure can speak to longer-horizon exchange rates in much the same way that microscopes speak to pathologies with macro impact. In medicine, microscopes provide resolution at the appropriate level—the level at which the phenomenon emerges. This is true irrespective of whether the phenomenon also has macro impact. Resolution at this level is the key to our understanding. Similarly, microstructure tools provide resolution at the level where its "phenomenon" emerges—the level where price is determined. What information do dealers have available to them, and what are the forces that influence their decisions? (Whether we like it or not, it is a stubborn fact that in the major currency markets, *there is no exchange rate other than the price these people set*.) Answering these questions does indeed help explain exchange rates over longer horizons. I provide evidence of this in section 7.1 and elsewhere.

Theme 3: Microstructure is relevant to macroeconomists.

In 1987 stock markets crashed around the world, an event that most people consider macro relevant. What analytical tools did the profession use to address the Crash? The tools were microstructure tools (see, for example, Grossman 1988, Gennotte and Leland 1990, Jacklin et al. 1992, and Romer 1993). These leading papers on the Crash are explicit about relaxing all three of the asset-approach assumptions noted above. In particular, the richness of these models comes from (1) information structure: which participants knew what; (2) heterogeneity: what types of participant were active and what were their motives for trading; and (3) institutions: what role did each participant play in the trading process and what trading information did each have available. The microstructure approach certainly bore fruit in this case.

Microstructure tools have also been applied by macroeconomists to understand exchange-rate collapses during the 1990s financial crises in Asia. These papers also introduce information structures, heterogeneity, and institutional

[10] Andersen and Bollerslev (1997) establish another link between microstructure and lower-frequency relevance. They show that forecasts of low-frequency volatility are more precise when based on high-frequency data.

factors that are not in general present within the traditional macro approach (see Chen 1997, Calvo 1999, Corsetti et al. 1999, and Carrera 1999).

Theme 4: Exchange-rate economics merits an information-theoretic perspective.

In many ways, this theme follows from the first three. The information economics of traditional exchange-rate models is rather simple. Macroeconomic news is announced publicly, so everybody learns new information at the same time. This news can then be impounded in price directly. The aggregation of bits of information that are not publicly known is presumed to play no role. The question is, of course, whether this captures the essence of price determination, or whether it neglects something important. This question is addressed extensively in the next chapter.

1.4 Applying Microstructure Tools to Exchange-Rate Puzzles

I turn now to puzzles within exchange-rate economics and how microstructure helps to resolve them. Though later chapters address this in detail, let me offer some initial thoughts. Consider first the puzzle of the FX market's enormous trading volume ($1.5 trillion per day, by far the highest of any financial market). In fact, the microstructure approach has made considerable progress on this puzzle. I have in mind here recent findings on the volume-amplification effects of the so-called **hot potato**. Hot potato trading is the passing of unwanted positions from dealer to dealer following an initial customer trade. In the words of Burnham (1991):

> [When hit with an incoming order, a currency dealer] seeks to restore his own equilibrium by going to another marketmaker or the broker market for a two-way price. A game of "hot potato" has begun ... It is this search process for a counterparty who is willing to accept a new currency position that accounts for a good deal of the volume in the foreign exchange market.

Thus, the passing of unwanted positions is a consequence of dealer risk management. Under the asset approach, in contrast, volume is attributed to speculation.

Understanding the causes of high volume is not as important as understanding price determination, but it is still important. Three reasons come immediately to mind. First, misunderstanding the causes of high volume can lead to bad policy. Consider the issue of transaction taxes—an issue that has attracted much attention among exchange-rate economists. Proponents of levying transaction taxes tend to associate high volume with excessive speculation. If instead much of this volume reflects dealer risk management, then a transaction tax would—unintentionally—impede that risk management. Second, high volume can impede order flow's information role. As suggested above, and detailed in the next chapter, order flow conveys information. The precision of that information is a function of the underlying causes of order flow. It is important to understand whether those causes contribute to, or detract from, informational precision. Third, misunderstanding the causes of order flow can lead to bad theory. The mere existence of the volume puzzle is an indication that the asset approach is not addressing key features of the FX market. The features it is missing may not be important, but to be confident that they are unimportant requires a tremendous leap of faith.

What about the big puzzles in exchange-rate economics? Chapter 7 addresses this question directly. The three biggest puzzles are:[11]

(1) **The determination puzzle**: exchange-rate movements are virtually unrelated to macroeconomic fundamentals (at least over periods of less than about two years);
(2) **The excess volatility puzzle**: exchange rates are excessively volatile relative to our best measures of fundamentals; and
(3) **The forward bias puzzle**: excess returns to speculating in foreign exchange are predictable and inexplicable.

The microstructure approach links these puzzles to one-another, the link being expectations formation—i.e., how market participants form their expectations of future fundamentals. It makes this link without departing from the asset-approach convention of rational expectations. Rather, the microstructure approach grounds expectations formation more directly in a richer, information-economic setting. The focus is on information *types* (such as public versus private) and *how* information maps into expectations (e.g., whether the aggregation of order flow “votes” is efficient). The issues of information type and mapping to expectations are precisely where microstructure tools provide resolving power.

Chapter 7 addresses the three big puzzles and shows that the microstructure approach has already made empirical progress. Section 7.1 addresses the first puzzle; it reviews the work of Evans and Lyons (1999), who find that exchange-rate movements can be explained—they are largely a function of order flow. Section 7.3 addresses the excess volatility puzzle; the focus of that section is recent work on the sources of FX volatility (e.g., Killeen et al. 2000a, Hau 1998, Jeanne and Rose 1999). Section 7.4 offers a microstructure-based explanation of the third puzzle—forward discount bias. It would be wrong for me to suggest that these three big puzzles have been put to rest by the above-mentioned work. But progress is being made.[12]

1.5 Spanning the Micro-Macro Divide

As noted above, a core distinction between the microstructure approach and the asset approach is the information role of trades. Under the asset approach,

[11] Within international finance more broadly, there are four main puzzles, the three listed above plus the “home bias” puzzle, which is that investors under-invest internationally. I do not include the home bias puzzle above because it is not clear it stems from the exchange rate, whereas the other three puzzles involve the exchange rate directly. Nevertheless, microstructure tools have already proven valuable for addressing home bias (e.g., the trading model of Brennan and Cao 1997 and the empirical results of Kang and Stulz 1997, among others).

These three exchange-rate puzzles have analogues in other markets. For equities, papers recognizing the three puzzles include Roll (1988), Shiller (1981), and Mehra and Prescott (1985), respectively. (However, the equity-market version of the forward bias puzzle—the so-called equity premium puzzle—is a much looser analogue than the other two: the large risk premium on equity is rather stable over time and remains positive, whereas the large risk premium in FX changes over time, including frequent changes in sign.) Microstructure tools are just beginning to be applied to those major equity puzzles (see, for example, Easley et al. 1999).

[12] Of course, the microstructure approach also has its drawbacks (e.g., lack of publicly available data over long periods). I consider those drawbacks in later chapters.

trades play no role (macroeconomic information is publicly announced), whereas in microstructure models trades are the driving force. It is instructive to frame this distinction with a bird's eye view of structural models used by empiricists within these two approaches.

Structural Models: Asset Approach

Equations of exchange-rate determination within the asset approach take the form:

$$\Delta P_t = f(i, m, z) + \varepsilon_t. \tag{1.1}$$

where ΔP_t is the change in the nominal exchange rate over the period, typically one month. The driving variables in the function f(i,m,z) include current and past values of home and foreign nominal interest rates i, money supply m, and other macro determinants, denoted here by z.[13] Changes in these public-information variables are presumed to drive price without any role for order flow. If any price effects from order flow should arise, they would be subsumed in the residual ε_t. Though logically coherent and intuitively appealing, a long literature documents that these macro determinants account for only a small portion (less than 10 percent) of the variation in floating exchange rates (see the surveys by Frankel and Rose 1995, Isard 1995, and Taylor 1995).

Structural Models: Microstructure Approach

Within the microstructure approach, equations of exchange-rate determination are derived from the optimization problem faced by the actual price setters (the dealers).[14] These models are variations on the following specification:

$$\Delta P_t = g(X, I, Z) + \varepsilon_t \tag{1.2}$$

where now ΔP_t is the change in the nominal exchange rate between two transactions, versus the monthly frequency of the macro model in equation (1.1). The driving variables in the function g(X,I,Z) include order flow X (signed so as to indicate direction), a measure of dealer net positions (or inventory) I, and other micro determinants, denoted by Z. It is interesting to note that the residual in this case is the mirror image of the residual in equation (1.1) in that it subsumes any price changes due to the public-information variables of the asset approach.

The key to spanning the micro-macro divide is the role of signed order flow X. Microstructure models predict a positive relation between price and signed order flow. This is because order flow communicates information that is not publicly available, and once communicated it is reflected in price. Empirical estimates of this

---

[13] The precise list of determinants depends on which model within the larger asset approach is selected. Here our interest is simply a broad-brush contrast between the asset and microstructure approaches. I return to precise specifications of the asset approach in chapters 6 and 7.

[14] Work using structural models includes Glosten and Harris (1988), Madhavan and Smidt (1991), and Foster and Viswanathan (1993), all of which address the NYSE; structural models in a multiple-dealer setting include Snell and Tonks (1995) for stocks, Lyons (1995) for currencies, and Vitale (1998) for bonds.

relation between ΔP and X are uniformly positive and significant for a number of markets (including stocks, bonds, and foreign exchange). It is noteworthy that these empirical estimates have been possible only a relatively short time: the switch to electronic trading means that we now have detailed records of order flows. What used to be a black box is no longer.

A Hybrid Approach

To establish the link between the micro and macro approaches, I investigate in chapter 7 equations with components from both approaches:

$$\Delta P_t = f(i, m, z) + g(X, I, Z) + \varepsilon_t. \tag{1.3}$$

These equations are estimable at frequencies corresponding to the asset approach through the use of time-aggregated measures of order flow X. The time-aggregated measures of X span much longer periods than those typically employed in empirical microstructure.

Estimates of this equation show that time-aggregated order flow has much more explanatory power than macro variables. In fact, chapter 7 shows that regressing daily changes in log DM/$ rates on daily order flow produces an R^2 statistic greater than 60 percent.[15] Figure 1.2 below provides a convenient summary of this explanatory power. The solid lines represent the spot rates of the DM and yen against the dollar over the four-month sample of the Evans (1997) dataset (described in chapter 5). The dashed lines represent cumulative order flow for the respective currencies over the same period. Order flow, denoted by X, is the sum of signed trades since the beginning of the sample period between foreign exchange dealers worldwide.[16] Cumulative order flow and nominal exchange-rate levels are strongly positively correlated (price increases with buying pressure). This result is intriguing. Order flow appears to matter for exchange-rate determination, and the effect appears to be persistent (otherwise the exchange rate's level would reflect only concurrent or very recent order flow and not *cumulative* order flow). This persistence is an important property, one that I examine more closely in later chapters. For order-flow to be helpful in resolving big exchange-rate puzzles, its effects have to persist over horizons that match those puzzles (monthly, at a minimum).[17]

[15] I use the notation "DM/$" here because I am referring specifically to the value of an exchange rate, and this exchange rate is quoted in the marketplace as the DM price per dollar. When referring in general terms to an FX market, practitioners typically list the dollar first: orally, they refer to this market as "dollar-mark," and this is typically written as the "$/DM market," or simply as "$/DM." (the Bank for International Settlements also follows this dollar-first convention in its FX-market surveys; see section 3.2). In this book, I remain true to these differing conventions. When describing an actual *rate*, precision requires me to write it as it is traded (e.g., DM/$, ¥/$, and $/£). When referring in general terms to a particular *market*, I always list the dollar first ($/DM, $/¥, and $/£). Finally, when I use the symbol "P" in equations to denote the exchange rate as a price, P always denotes the dollar price of the other currency ($/other).

[16] Because the Evans (1997) data set does not include the size of every trade, this measure of order flow is in fact the *number* of buys minus sells. That is, if a dealer initiates a trade against another dealer's DM/$ quote, and that trade is a $ purchase (sale), then order flow is +1 (–1). These are cumulated across dealers over each 24-hour trading day (weekend trading—which is minimal—is included in Monday).

[17] Readers familiar with the concept of co-integration will recognize that it offers a natural means of testing for a long-run relationship. In chapter 7, I present evidence that cumulative order flow and

Figure 1.2

Four Months of Exchange Rates (solid) and Order Flow (dashed)
May 1-August 31, 1996

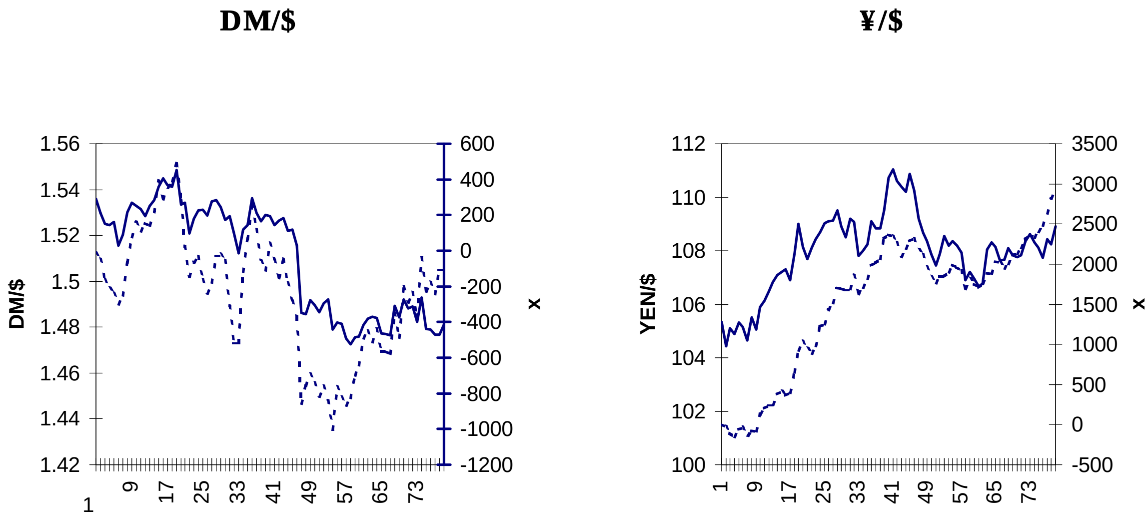

That order flow matters for exchange-rate determination does not imply that order flow is the **underlying cause** of exchange-rate movements. Order flow is a **proximate cause**. The underlying cause is information. How, specifically, can one identify the information that determines order flow? The notion of order flow as an intermediate link suggests several strategies to answering this question, and I address these further in later chapters (particularly in chapters 7 and 9).

One strategy for linking order flow to underlying determinants starts by decomposing order flow. (That it can be decomposed is one of its nice properties.) In chapter 9 I test whether all parts of the aggregate order flow have the same price impact. They do not: the price impact of FX orders from financial institutions (e.g., mutual funds and hedge funds) is significantly higher than the price impact of orders from non-financial corporations. This suggests that order flow is not just undifferentiated demand. Rather, the orders of some participants are more informative than the orders of others. Analyzing order flow's parts gives us clues as to the underlying information structure.

A second strategy for linking order flow to underlying determinants is based on the view that order flow measures changing market expectations. As a measure of expectations, it reflects a willingness to back one's beliefs with money—the backed-by-money expectational votes, if you will. Expectations measured from

---

the level of the exchange rate are indeed cointegrated (i.e., the relationship between order flow and price is not limited to high frequencies). I also present models in chapter 7 that show why a long-run relationship of this kind is what one should expect.

macro data, on the other hand, are slow moving and imprecise. If order flow is serving as an expectation proxy, then it should forecast surprises in important macroeconomic variables (like interest rates). New order-flow data sets that cover up to six years of FX trading provide enough statistical power to test this.

Chapter 2: The Economics of Order-Flow Information

As noted in chapter 1, an important focus of this book is the economics of order-flow information. The objective of this chapter is to provide a foundation for that focus. How do we know order flow conveys information? What type of information does order flow convey? I answer both of these questions here.

In the first section, I provide background on the information economics of traditional asset-approach models. In section two, I review recent empirical evidence on order-flow information. The evidence supporting an information role for order flow comes from many sources, and is documented using several different methodologies. In section three, I provide a working definition of the term private information; if order flow is going to convey information that is not public, we need a clear sense of what that information is. The concluding section provides a more theoretical view of the types of information that order flow conveys. Although these types share certain characteristics, they also have important differences. These differences have implications for theory and empirical work.

This chapter does not extend to the mechanics of *how* order flow is impounded in price. That topic is addressed in chapter 4, which presents the key theoretical frameworks. It is only in the context of the theory that the details become meaningful. The goal of this chapter is more modest: to establish a comfort level with the idea that order flow is a transmission mechanism, a means of transmitting information to price.

2.1 Background

To many macroeconomists, the idea that order flow conveys incremental information relevant to exchange rates is controversial. In traditional models, macroeconomic news is announced publicly, and can therefore be impounded in price directly, without any role for order flow.[18] In light of this common belief, let me provide a bit more background before examining evidence that order flow does indeed play an information role.

As noted in chapter 1, models of exchange-rate determination within the asset approach take the form:

$$\Delta P_t = f(i, m, z) + \varepsilon_t.$$

where ΔP_t is the change in the nominal exchange rate over the period, typically one month. The driving variables in the function f(i,m,z) include current and past values of home and foreign nominal interest rates i, money supplies m, and other macro determinants (e.g., home and foreign real output), denoted here by z.

Now consider the possibility of private information. It is difficult to imagine circumstances in which agents would have private information about interest rates.

[18] A related, less extreme view is that order flow may convey some non-public information, but the information it conveys is likely to become public soon. In this case, order flow advances the impounding of information in price by only, say, a few minutes. This less extreme view is not consistent with the data, however. To a first approximation, our best measures of public-information flow are uncorrelated with the direction of exchange-rate movements (at annual or higher frequencies), whereas order flow *is* correlated with exchange-rate movements. I return to this issue in chapter 3.

Perhaps somebody had an enlightening conversation with the Chairman of the Federal Reserve Board (Fed) that morning. Doubtful. Maybe somebody has inside information about the next Fed vote regarding monetary policy. Again, doubtful.[19] The natural presumption is that all agents have the same information about current—and future—interest rates. How about private information on money supplies, or real output? For these variables, too, it is natural to presume that agents are symmetrically informed. When money-supply or real-output data are publicly announced, all agents learn new information at the same time. This is not a recipe for speculative activity based on information advantage.

With a slight shift in perspective, however, an information role for order flow emerges. This shift in perspective is perfectly consistent with the ideas of the last paragraph: even if exchange-rate determination is based wholly on public information, this is not sufficient to rule out an information role for order flow. To understand why, recognize that there are in fact two crucial assumptions in these macro-asset models that disconnect order flow from price. These two assumptions are:

(1) all information relevant for exchange rates is publicly known; and
(2) the mapping from that information to prices is also publicly known.

If either assumption is relaxed, then order flow can convey information that is relevant for prices. As an example of relaxing the second, suppose you and I are FX traders for large banks and we have both seen the same macro announcement. If I do not know how you will interpret the announcement's price implications, then I need to watch your trade to learn about your interpretation. Your trade—order flow—will both influence price and teach me something.

As an empirical matter, relaxing the assumption that everybody knows the mapping from public information to price should not be controversial. There is no existing consensus on the "right" macroeconomic model: extant macro-asset models fit the data poorly, as noted in chapter 1. In his book "Exchange Rate Economics," Isard (1995) puts it this way:

> Thus, economists today still have very limited information about the relationship between equilibrium exchange rates and macroeconomic fundamentals. Accordingly, it is hardly conceivable that rational market participants with complete information about macroeconomic fundamentals could use that information to form precise expectations about the future market-clearing levels of exchange rates.

This fact has important implications for the role that order flow plays in mapping information to price.

Relaxing the all-agents-know-the-mapping assumption is not, however, the only way to restore a role for order flow. There are many types of information that do not conform to the first of the two macro-asset model assumptions—that all information relevant for exchange rates is publicly known. Section 2.3 reviews these

[19] I have in mind here the major FX markets. This kind of private information is more reasonable in emerging markets. For example, the IMF (1995) reports that Mexican investors were the first to flee Mexico in the period immediately prior to the December 1994 devaluation. It is not unreasonable to believe that certain people close to the devaluation decision had inside information and were able to act on it.

types of information. In anticipation of that material, I offer a recent quote from Rubinstein (2000) that presages some of these information types:

> Perhaps the most important missing generalization in almost all work on asset prices thus far has been uncertainty about the demand curves (via uncertainty about endowments or preferences) of other investors. This injects a form of endogenous uncertainty into the economy that may be on a par with exogenous uncertainty about fundamentals.

Before reviewing these information types, let me first provide some empirical evidence on order flow's information role.

2.2 Empirical Evidence that Order Flow is Informative

Four distinct empirical methodologies have been used to generate evidence of an information role for order flow. That these different methodologies produce similar results is an indicator of the evidence's robustness. The methodologies look in different places for indications of order flow's information effects. The four places are: (1) persistent effects of order-flow on price; (2) adverse-selection components of bid-offer spreads; (3) volatility responses to trading halts; and (4) survey data from FX dealers. In this section, I briefly review some of the papers that fit in each of these four categories. I keep it brief here because it is too early to present these empirical methodologies in detail. The detailed material appears in chapter 5, which is dedicated to the empirical frameworks.[20]

Methodology 1: Persistent Order-Flow Effects on Price

One methodology used to show that order flow is informative focuses on order flow's price effects and their persistence. It is common in the literature to distinguish between order flow that has transitory effects on price and order flow that has permanent effects on price. When order flow has transitory effects on price—sometimes called "indigestion" or "inventory" effects—these effects are often referred to as pricing errors. When order flow has permanent effects on price, however, these effects are taken to reflect underlying fundamental information. (This identification scheme is used by French and Roll 1986, for example, in their celebrated paper on information arrival and stock return volatility.)

In empirical microstructure, the standard way to implement this idea is to estimate vector auto-regression models (VAR) and test whether innovations in order flow have long-run effects on price (Hasbrouck 1991). When applied to data from major FX markets, one finds that order flow innovations do indeed have long-run effects on price, indicating that they convey bona fide information. Examples of findings along these lines include Evans (1999) and Payne (1999).

There is a second method for testing whether order flow effects are persistent that has been applied to FX markets. This method uses time-aggregated order flow to explain price movements. That is, rather than asking whether a single trade has an impact on price, this work asks whether trades aggregated over time (say a day)

[20] This section emphasizes empirical work that uses order flow data, as opposed to empirical work that uses price data only.

have an impact on price. The idea is that if single trades have only fleeting effects on price, then order flow aggregated over the day will not be closely related to daily price movements. When applied to the FX market, one finds that daily order flow does remain strongly positively related to daily price changes. Examples of findings along these lines include Evans and Lyons (1999) and Rime (2000).

Methodology 2: Spreads and Informative Order Flow

Though not obvious to people new to microstructure, bid-offer spreads provide a means of testing whether order flow is informative. This is rather striking. It implies that purely on the basis of spread behavior, one can learn something quite important about the FX market's information structure. To wit, the information structure is not one of public information with a publicly known mapping of that information to price, but one where individual orders convey non-public information.

To understand why spreads provide a test, it is important to understand something about how spreads are determined. I show in chapter 5 that spreads exist because of three costs faced by dealers. One of these costs—typically referred to as an **adverse selection** cost—results from asymmetric information. Dealers know that when they trade with someone who is better informed, they can expect to lose money from the trade. If one could identify better-informed traders before trading, then this would not be a problem—dealers could choose not to trade, or adjust price appropriately. But dealers typically cannot identify those who are better informed. Given this, one way to protect against losses is to increase the width of the spread quoted to all potential counter-parties—informed and uninformed alike. That way, there is more room for prices to move against the dealer before he begins to lose money. When a dealer protects himself this way, we say that he has included an adverse selection effect (or component) in his spread. If we find this adverse-selection component of the spread empirically, this indicates that dealers believe that some traders' orders reflect better information (despite the dealers' inability to identify exactly who those traders are).

Empirical findings show that an adverse selection effect is indeed present: dealers are increasing spreads to protect themselves against informative incoming order flow (Lyons 1995, Yao 1998b, and Naranjo and Nimalendran 1999).[21] In Lyons (1995), for example, I find that the FX dealer I track protects himself from adverse selection by increasing the width of his spread by about one pip (or 0.0001 DM/$) for every $5 million increase in the size of the incoming order. (This estimate is a pure adverse selection effect on the spread, i.e., the model controls for other costs that can affect spread width.) Yao (1998b) also finds an adverse selection effect on the spread. In addition, he finds that the dealer he tracks profits from trades executed after observing order flows. Unlike the Lyons and Yao papers, which look at market trading in general, Naranjo and Nimalendran (1999) focus on central-bank trading—i.e., intervention. They find that the adverse selection effect on the spread is positively related to the variance of unexpected intervention trades.

[21] Hartmann (1998) finds that FX spreads widen with unexpected volume, which is consistent with an adverse selection component.

Methodology 3: Volatility Responses to Trading Halts

The third methodology used to show that order flow is informative focuses on price volatility over periods that contain a sub-period during which trading is halted. (Think, for example, of the volatility of prices over weeks in which trading is halted on Wednesday.) The trick in this case is to identify trading halts that are unrelated to the flow of public information. If the trading halt is related to the flow of public information—holidays for example—then changes in volatility that occur because of the halt could easily be due to changes in that public-information flow. If one is confident that the flow of public information has not changed, however, then changes in volatility that occur because trading is halted must be attributed to something other than public information. In particular, a finding of *lower* volatility over periods that contain trading halts would indicate that either (1) informative order flow was not reaching the market during the halt, (2) the lack of trading during the halt reduced pricing errors, or (3) some combination of both.

French and Roll (1986) were the first to apply this methodology. They applied it to the trading of stocks. In their case, they identified a set of days—mostly Wednesdays—on which the New York Stock Exchange (NYSE) was closed due to order-processing backlogs. These were regular weekdays, not economy-wide holidays, so the closure had no impact on the underlying firms, nor on the economy as a whole. The only difference is that the NYSE was not trading. Thus, there is no reason to believe the flow of public information changed (save, perhaps, firms' choosing not to release information due to the closure). French and Roll then measured the volatility of returns from Tuesday to Thursday in weeks when there was trading on Wednesdays, and compared this to the volatility of returns from Tuesday to Thursday in weeks with no trading on these special Wednesdays. They found that volatility was significantly lower in weeks with the Wednesday closures. This, together with other evidence they provide, leads them to conclude that the main source of the volatility reduction is possibility (1) above: informative order flow was not reaching the market during the Wednesday halts.

Ito, Lyons and Melvin (1998) analyze an analogous experiment in the FX market. Their experiment was made possible by a change in the trading rules in the Tokyo FX market. Until December 1994, banks were restricted from trading in Tokyo over a lunch period (from 12 to 1:30 local time). In December 1994, the restriction was lifted. There were no other changes in policy or other shifts in the flow of macroeconomic information (e.g., announcement dates and times), so the flow of public information was not affected. Ito et al. find that after lifting the restriction, volatility in the $/yen market over the lunch period doubled.[22] Was this due to informative order flow or pricing errors? The increase in volatility does not, in itself, allow one to distinguish between them. But because it is unlikely that the increase is due *wholly* to pricing errors, this result also supports an information role for order flow.

---

[22] In the case of FX, volatility over the lunch period can still be calculated because trading continues in other trading locations, such as Singapore and Hong Kong. Andersen, Bollerslev, and Das (2000) verify the significance of this volatility increase over lunch. They do find, however, that some other results in the Ito et al. paper are sensitive to an outlier in the data set. The analysis of Covrig and Melvin (1998) of the same Tokyo experiment excludes the outlier observation, but still finds evidence that Japanese banks are relatively well informed, corroborating the Ito et al. interpretation. The Covrig-Melvin evidence is based on price leadership by Japanese banks.

Methodology 4: Surveys of Foreign Exchange Dealers

If we think dealers might believe order flow is informative, why don't we just ask them? This is the tack taken by Cheung and Wong (1999, 2000), Cheung and Chinn (1999a,b), and Cheung, Chinn, and Marsh (1999). These papers survey foreign exchange dealers in major trading centers around the world (London, New York, and Tokyo). Fifty percent of the dealers surveyed by Cheung and Wong believe that large players in the FX market have a competitive advantage that derives from "better information" and "a large customer base." The latter is often described as a source of advantage for dealers because it gives them privileged information about their customers' orders, and they base their speculative trades on this information (see, e.g., the Yao 1998b paper cited above).

Other authors report results from discussions with dealers that are similar, though these results are less formal than the above-noted surveys. For example, Goodhart (1988) writes, "A further source of informational advantage to the traders is their access to, and trained interpretation of, the information contained in the order flow." Similarly, based on interviews with nine FX dealers in London, Heere (1999) reports that the dealers emphasize the importance of asymmetric information. The dealers she interviews state that information asymmetry is based on both order flows and the identities of the institutions behind those order flows.

2.3 Defining Private Information

In this book, information is **private information** if:

(1) it is not known by all people; and
(2) it produces a better forecast of price than public information alone.[23]

This is a natural definition. I should say, though, that it is a bit broader than some people have in mind. For example, under this definition, if a dealer has privileged order-flow information, and that information aids him in forecasting prices, then this constitutes private information. (He could be expected to take speculative positions based on it.) For understanding motives for speculative trade, this definition is quite useful.

To add concreteness to the definition, let me identify two sub-categories of private information, and provide some examples from FX. Consider a simple two-period trading model in which trading occurs initially at a price P_0, again at t=1 at a price P_1, and then a terminal payoff value V (e.g., an unknown final dividend) is realized at t=2. The first type of private information that can arise in this setting is private information about the size of the terminal payoff V. Information that alters expectations of the payoff V will clearly be relevant for the prices P_0 and P_1. Let me offer two examples of private information in the FX market that are arguably of this type. The first example is private information about central bank intervention: the dealer who receives a central bank's order has also received private information (Peiers 1997). The second example is private information about particular compo-

[23] By "better" I mean lower conditional variance.

nents of exchange-rate fundamentals, e.g., the trade balance: real trade in goods and services generates FX orders that provide information to dealers about trade balances long before published statistics are available (Lyons 1997a).

In contrast to this first type of private information, a second type of private information is unrelated to the payoff value V, but is related to interim prices P_0 and P_1. The function that determines P_0 and P_1 includes many variables beyond expectations of the payoff V. An example is traders' risk preferences. Other examples include traders' trading constraints, the supply/distribution of the risky asset, and other features of the trading environment. Insofar as these affect P_0 or P_1 without altering expectations of V, superior knowledge of them is private information of this second type.

Consider some examples from the FX markets of private information of this second type. A first example begins with the fact that even when traders have common information about V, they may still disagree on the meaning of this information, thereby affecting P_0 or P_1. (Consider the disagreement—despite access to the same data—among financial analysts about the direction of the stock market.) Theory identifies the source of this disagreement as different prior beliefs, different models, or both. Whatever the source, if one adds superior knowledge about others' beliefs—perhaps because their trades have been observed—then it is possible to forecast interim price more accurately than the market at large.

Another example of private information of this second type is superior knowledge about dealer inventories (i.e., the distribution of risky-asset holdings). Because the transparency of order flow in FX is low, dealers always have superior knowledge of their own inventory, and often have better information than the general public about other dealers' inventories as well. If inventory risk earns a risk premium, as many microstructure models predict, then superior knowledge of this kind allows a dealer to forecast interim price more accurately than the market at large.[24] As aggregate inventory across dealers becomes known, this induces a change in the risk premium and an attendant change in price, even though terminal-payoff expectations remain unchanged.[25] Discussions with FX dealers indicate that this type of private information is indeed relevant.

This book tends to highlight the second of these two broad types of private information—information about interim prices. There are several reasons for this. First, in my judgment this type of private information is especially relevant for the FX market. Second, people who believe that private information does not exist in the FX market typically have in mind only the first of the two types—private information about the size of the payoff, V (which in the FX market translates to private information about future interest differentials). Highlighting the second of the two broad types broadens that perspective. Third, previous literature tends to neglect this category of private information. Information-theoretic models of trading are specified with private terminal-payoff information. Empirical models follow suit. But this makes interpretation of the empirical results difficult: should one interpret evidence of private information as reflecting the first type or the second type? The answer is not clear.

---

[24] Empirical evidence of inventory effects on FX prices is provided in Lyons (1995).

[25] To understand why the change in the risk premium—i.e., the expected return—changes price, think of a pure discount bond: when the market interest rate—the expected return—changes, the price of the bond must change, even though the cash flow at maturity does not change.

A Comment on the Term "Fundamentals"

The term **fundamental** means different things to different people. For example, one might be tempted to consider the second of my two private-information types as non-fundamental. The quote from Rubinstein (2000) that closed section 2.1 is suggestive of this narrower definition of fundamentals. In that quote, he distinguishes uncertainty about fundamentals from uncertainty about agents' preferences and endowments. But it is not clear that preferences and endowments should be distinguished from fundamentals. My choice to put the two broad types of private information on equal footing recognizes the joint, complementary nature of these two categories of fundamentals. The issue is more than semantic; it affects the way we frame our thinking about price determination.

My use of the term "fundamental" to refer to information of both types is not so broad that it is no longer meaningful. The examples of private information above are all bona fide determinants of price in optimizing, well-specified models. None of the examples presented here rely on "bubbles," "greater-fool" behavior, or irrationality.[26]

2.4 Extending the Taxonomy of Information Types

Let us shift to a more theoretical approach to identifying the types of information that order flow might convey. This section extends the last section by adding more granularity to the earlier two-category breakdown of private information. The danger in adding more granularity—and grounding it in theory—is that readers new to microstructure will find this section tougher going. Readers who do find this section more difficult should rest assured that the previous section is an ample introduction to the basic information-theoretic issues that arise later in the book.

Let us begin by recognizing an important connection between order flow and information within microstructure models. As we will see in chapter 4 (which reviews microstructure theory), order flow is the proximate determinant of price in all the standard models. This "all-models" property is important: it ensures that the causal role played by order flow is not dependent on the specification of market structure.

To understand the specific types of information that order flow can convey, one needs to understand the specific channels through which order flow has price impact. Theory provides a disciplined approach for achieving this. At the broadest level, the information conveyed by order flow can be any information that currency markets need to aggregate (e.g., differential interpretation of news, shocks to liquidity demands, shocks to hedging demands, etc.). Within this broad class of information, theory provides a taxonomy of different types.

To set the stage for those different types, consider the following highly simplified view of asset pricing. Let us write the equilibrium price of a risky asset in a one-period setting as:

[26] This logic is also behind my choice of "V" in this book to denote payoffs, rather than the more customary "F"; the symbol F is too easily interpreted as the whole of fundamentals, an interpretation that is too narrow in my judgment.

(2.1) $$P_0 = \frac{E[V_1]}{1+d}$$

where P_0 is the price at the beginning of the period, $E[V_1]$ is the expected value of the risky asset's payoff, and d is the market-clearing discount rate. In the case of a stock, where V_1 is a dividend, this equation is the familiar dividend discount model. (The whole stream of discounted dividends would be included in a multiple-period setting.) In the background of such a pricing equation is a market-clearing condition, which equates market demand with market supply (supply is typically assumed fixed). Any factor—other than the numerator $E[V_1]$—that affects market demand/supply will affect price through the market-clearing discount rate d.

Now we are ready to outline the types of information that order flow can convey. There are three key types that arise within microstructure theory:

(1) Payoff information
(2) Discount-rate information—inventory effects
(3) Discount-rate information—portfolio-balance effects

In the case of a stock, payoff information refers to information about future dividends—the numerator $E[V_1]$ in our dividend discount model. In the case of foreign exchange, payoffs take the form of current and future interest differentials (foreign minus domestic; see section 6.1 for more detail). These represent the cash flows that accrue to holders of money-market instruments denominated in foreign exchange—akin to the dividends that accrue to holders of a stock. (FX speculators who buy foreign exchange do not hold actual currency, which bears no interest, but instead invest their holdings in short-term, interest-bearing instruments.) We will see that private information about payoffs is the basis for a class of microstructure models known as **information models** (reviewed in chapter 4).

Let me provide examples of how order flow, per se, might convey private information about payoffs. The simplest example—though not so common in the major FX markets—is information about future interest rates conveyed in the orders of a central bank (intervention). A second example that is likely to operate on a more regular basis is information about people's *expectations* of future interest differentials.[27] To understand this example, recognize that in reality, people do not all share the same $E[V_1]$. Instead, each of us has our own expectation about the direction of future interest rates, based on the millions of bits of information we use to form this view. This can be described using the numerator $E[V_1 \mid \Omega_i]$, where Ω_i denotes the information that market participant i uses to form her expectation. Because participant i's orders depend on $E[V_1 \mid \Omega_i]$, observing her orders conveys information about that expectation. Thus, order flow is serving as a proxy for people's expectations about future payoffs, and the information embedded therein. These orders are the backed-by-money expectational votes that the market "counts" when determining price.[28]

---

[27] This expectations story is also applicable to information about discount rates, in which case it would fall into type (2) or type (3).

[28] Other proxies for expectations—such as time-series measures like ARIMA or VAR models—do not share this backed-by-money, information-encompassing property. A quote by Frankel and Rose (1995) provides some perspective on these time-series measures. In their words, "To use an ARIMA or VAR

Turning to discount-rate information, microstructure theory emphasizes two distinct causes of time-variation in discount rates. Crucially, order flow is related to both of these causes. The first of these causes is **inventory effects.** (These arise in the class of microstructure models known as **inventory models**; see chapters 4 and 5 for details.) The idea here is that risk-averse dealers will require compensation for absorbing transitory mismatches in supply and demand over time. The larger the mismatch, the greater the risk the dealer must assume, and the greater the compensation the dealer requires. Suppose, for example, that the mismatch is such that the dealer needs to absorb market sell orders (i.e., the dealer needs to buy). The dealer may be willing to absorb a small amount at only a slightly lowered price, but he would require a significantly lowered price to absorb a large amount. Dealers thus earn a transitory risk premium for providing liquidity. These effects on price last only as long, on average, as the mismatches in market supply and demand. Once the supply-demand mismatch is remedied, the dealer no longer holds a position (inventory), so the effect on price dissipates—the risk is diversifiable. This type of order-flow effect on price is what people have in mind when they assert, "microstructure effects fizzle quickly." To summarize, these effects arise because risk is not perfectly and instantaneously spread throughout the whole market; instead, dealers bear disproportionate risk in the short run, and this affects price in the short run.[29]

The second cause of time-variation in discount rates—the third of our three information categories above—is **portfolio-balance effects**. To distinguish these effects from inventory effects, the idea in this case is that even after risky positions are spread through the economy as a whole, order flow's effect on price will not disappear completely. In other words, the risk that drives the inventory effect is, ultimately, diversifiable, whereas the risk that drives the portfolio-balance effect is undiversifiable. Of course, to distinguish this from the first of our information types—information about payoffs—it must be shown to be the case that order flow is not conveying information about $E[V_1 | \Omega_i]$.

Let me offer two types of FX order flow that are unrelated to $E[V_1 | \Omega_i]$, but may be large enough to have persistent portfolio-balance effects. These are orders—or the aggregation of orders—that arise from liquidity demand (e.g., importing and exporting) or from hedging demand. Consider a specific example. Suppose IBM sells equipment worth $3 billion in the UK, is paid in pound sterling, and then dumps those pounds on the spot market to exchange them for dollars. Somebody must be willing to step up and hold those pounds for the indefinite future. (The word indefinite is important here: to keep the example simple I am assuming that IBM is not going to reverse its decision to sell the pounds sometime in the future—it is a permanent portfolio shift. Also, I am assuming that the equipment buyer makes not adjustment in its portfolio other than the pound payment.) If elasticity of market demand for those pounds is less than infinite, then the dollar price of a pound (the

process as a measure of what agents expect, is to ascribe to them simultaneously not enough information, and too much. It does not ascribe to them enough information, because it leaves out all the thousands of bits of information that market investors use … It ascribes to them too much information … because it assumes that they know the parameters of the statistical process from the beginning of the sample period." Also relevant is the discussion in Engel (1996), where he describes "peso problems" in exchange rates as a case where the market had more information than empiricists, and "learning" as a case where the market had less information than empiricists (ex post).

[29] Though this discount-rate effect on price emerges in most models as a risk premium, for technical convenience sometimes models are specified with risk-neutral dealers who face some generic "inventory holding cost," which produces similar results.

\$/£ rate) must fall to induce people to purchase the pounds. Algebraically, for the market to clear we need:

$$S_{£} = D_{£}(P_{\$/£}, IBM)$$

where I will assume that the aggregate supply of pounds, $S_{£}$, is fixed by the bank of England, and the aggregate demand for pounds depends negatively on the current spot rate, $P_{\$/£}$ (with less than infinite elasticity), and positively on the liquidity demand for pounds from IBM. Though quite simple, this example is in the same spirit as the portfolio-balance effects from sterilized intervention by central banks—an example familiar to international economists; the only difference is that instead of the central bank forcing a (payoff-unrelated) position on the public, in my example it is a subset of the public forcing a (payoff-unrelated) position on the rest of the public. (See chapter 8 for more on central-bank intervention, including the distinction between sterilized and unsterilized intervention.)

Let me relate the above discussion to the previous section's introduction of private information. In effect, the discussion above splits one of the earlier private-information types into two parts. The previous section introduced a type of private information that is unrelated to the payoff value V, but relevant to interim prices P_0 and P_1. This is what I am now calling private information about discount rates, of which I have established two distinct categories, inventory effects and substitutability effects. By splitting the earlier type in two, this section provides more granularity, and links the earlier description more tightly to microstructure theory. It is helpful to distinguish between the two sub-categories because their different properties allow us to isolate them empirically: inventory effects are transitory, but portfolio-balance effects persist.

A Graphic View of These Information Types

To clarify these three types of order-flow information further, I turn to a graphic interpretation. Consider, for example, the second type of information outlined above—information about transitory inventory effects on discount rates. These effects are presumed to dissipate quickly because dealers are not holding these positions for long; the spreading of risky positions to non-dealer participants occurs rapidly (within a day in foreign exchange). Figure 2.1 provides a qualitative illustration. The short-run market (net) supply curve, denoted S^{SR}, slopes up. (Think of this supply curve as a willingness to sell to accommodate incoming order flow, rather than as changes in the physical supply of foreign exchange.) If order flow is not conveying either of the other two information types whose effect on price persists beyond the short run, then longer-run (net) supply curve S^{LR} is flat.[30]

[30] Readers familiar with microstructure theory will recognize that this figure assumes that there is no fixed component to the spread—i.e., the effective spread shrinks to zero as the size of the incoming order shrinks to zero. This is a detail that need not concern the more macro-oriented reader.

Figure 2.1

Supply Curves with only Transitory Inventory Effects

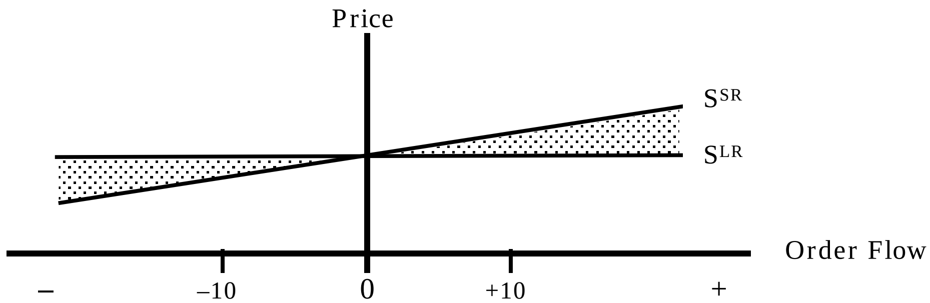

* The dotted region represents the transitory inventory effects. The effective spread faced by a customer for a 10-unit order is the difference in price along the short-run net supply curve S^{SR} between –10 and +10. If a customer wants to buy 10 British pounds from the dealer—an order of +10—then he 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.

Now let us add persistent effects from imperfect substitutability.[31] Doing so implies that the market as a whole—being risk averse—needs to be compensated for holding a position it would not otherwise hold. This requires an enduring risk premium, which takes the form of a price-level adjustment (per the IBM example). This price adjustment is not temporary because the risk premium for holding this position must be sustained. Figure 2.2 illustrates this. The short-run market (net) supply curve still slopes up, but now it reflects both the transitory effect of inventory at the dealer level and the longer-run effect from imperfect substitutability. It is, as a result, more steeply sloped than the short-run supply curve from inventory effects alone illustrated in Figure 2.1. To understand why this short-run effect goes in the same direction as the longer-run effect, think of the underlying dealer behavior. An individual dealer will buy pound sterling over the short run only at a relatively discounted \$/£ price. The market as a whole will take those pounds off the dealer's hands at a slightly discounted price, but not as discounted as was required by the dealer in doing the trade initially with the customer. The dealer, knowing this cost of laying off his inventory has increased, will pass this on to the customer in his initial quotes.

[31] For evidence of imperfect substitutability across U.S. stocks, see Scholes (1972), Shleifer (1986), Bagwell (1992), and Kaul et al. (2000), among others. For at least two reasons, imperfect substitutability may be more applicable to currency markets than to markets in individual stocks. First, note that the size of the order flows that the \$/euro 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). Second, there are far more individual stocks that can substitute for one another in portfolios than there are individual currencies (particularly major currencies).

Figure 2.2

Supply Curves With Inventory & Portfolio-balance Effects

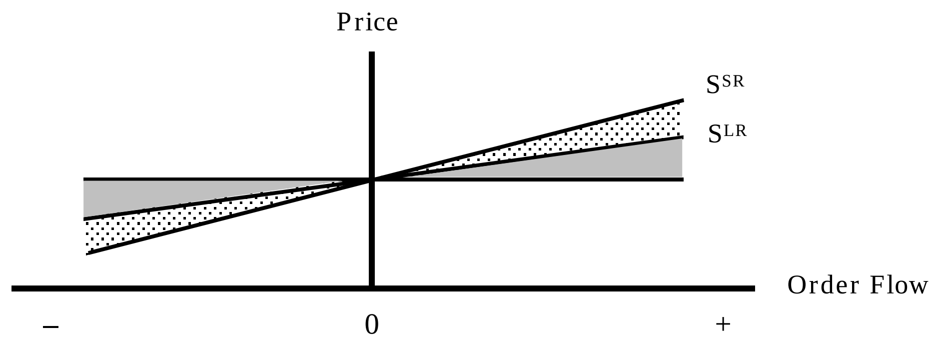

* The dotted region represents the transitory inventory effects. The gray region represents persistent portfolio-balance effects. Due to inventory effects, the short-run price impact of an incoming order is larger than the long-run impact. But the long-run impact is non-zero, due to imperfect substitutability; 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.

Now I will allow order flow to convey information about expected future payoffs. Like in the case of imperfect substitutability, order-flow effects on price from this channel will persist (see, e.g., French and Roll 1986, Hasbrouck 1991). Mapped into the market supply curve, this channel adds additional slope to the longer-run schedule shown in Figure 2.2. Figure 2.3 provides an illustration of these longer-run supply curves. Note that the short-run supply curve is more steeply sloped than either of the two longer-run curves.

Figure 2.3

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

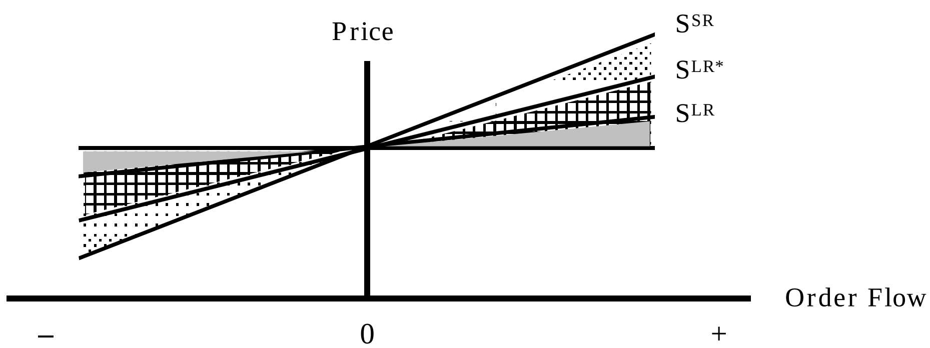

* The dotted region represents the transitory inventory effects. The crosshatched region represents persistent payoff-information effects. The gray region represents persistent portfolio-balance effects. The figure therefore reflects all 3 of the information types that arise in microstructure theory. The long-run supply curve S^{LR*} reflects both the long-run effects from imperfect substitutability (S^{LR}), 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.

In Figure 2.3 there is now a new longer-run supply curve, S^{LR*}. This new longer-run supply curve reflects both the longer-run effects from imperfect substitutability (S^{LR}), as before in Figure 2.2, plus an additional longer-run effect due to the payoff information conveyed by order flow.

In later chapters, I address the slopes of these net supply curves empirically. At this stage, it is worth bearing one point in mind: order flows in the FX market are enormous relative to other asset markets. In the figures, this corresponds to being far to the left or far to the right of the order-flow-equals-zero point. Thus, even if the slopes of these supply curves are nearly zero, large order flow can still produce substantial price adjustment.

Concluding Thoughts

To conclude this chapter, it is worth stepping back to reflect on an important, overarching point. The microstructure tools applied in this chapter are useful for addressing a rather deep question:

What is the nature of the information this market is processing?

By focusing attention on order flow, these tools help to characterize which types of information are relevant, and how this information is aggregated.

In terms of financial markets' economic role, the aggregation of information is of profound conceptual importance. Nobel laureate Friedrich Hayek (1945) provides an early and powerful articulation of this point. He writes:

> The peculiar character of the problem of rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form, but solely as disbursed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess. The economic problem of society is thus a problem of the utilization of knowledge not given to anyone in its totality.

Relative to traditional exchange-rate approaches, the information-theoretic perspective offered here is qualitatively different. As we shall see in Chapter 6 (where I survey exchange-rate models), exchange-rate economics may warrant a richer information-theoretic perspective.

Chapter 3: The Institutional Setting

Chapter 1 began with an overview of the microstructure approach, making the point that the approach relaxes three of the asset approach's most uncomfortable assumptions. Those assumptions are that (1) all FX-relevant information is publicly available, (2) all market participants are alike, and (3) trading mechanisms are inconsequential for prices. Chapter 2 addressed the first of these by providing frameworks for thinking about information that is both relevant and not publicly available. This chapter addresses the other two assumptions: trader heterogeneity and the role of trading mechanisms.

In the first section, I provide an overview of the market participants, and how they differ from one another. The first section also describes the trading mechanisms used in the major FX markets, and includes comparisons with other financial markets. (See also Luca 2000 for a great deal of institutional information, including an historical account of FX market development.) Section 2 introduces an important source of institutional information—the triennial central-bank surveys summarized by the Bank for International Settlements (BIS 1999). These survey data provide institutional perspective not available from any other source. Section 3 addresses market transparency. Transparency is a market feature that is crucial to understanding how order-flow information is conveyed. In markets that are highly transparent, order flow is observed by all participants, thereby affecting expectations rapidly and precisely. In opaque markets (the FX market is relatively opaque), order flow is not widely observed, so that any information it conveys is impounded into prices more slowly. The concluding section, section 4, provides reflections on a common and powerful association—the association between institutions and the field of microstructure.

One should be more careful than I have been thus far when referring to "the" FX market. Many people understand this term to mean spot markets in the major floating exchange rates, such as $/euro and $/¥, the two largest spot markets. In its broadest sense, though, the term includes markets other than spot and rates other than the major floaters. FX markets other than spot include the full array of derivative instruments (forwards, futures, options, and swaps). Rates other than the major floaters include those in smaller markets, such as emerging markets, and those in pegged regimes, such as Western Europe before the euro. When people quote the daily trading volume in FX at $1.5 trillion, that statistic applies to the broadest definition of the market.[32]

Nevertheless, the essence of the FX market is the spot market. In 1998 the spot market accounted for 40% of total turnover across all FX instrument categories ($600 billion out of $1.5 trillion). Though this share has been trending downward—it was 59% in the BIS survey of 1989—the falling spot share is not due to lower spot turnover in absolute terms. Rather, the derivatives markets have grown up around the spot market.

For the purposes of this book, however, a vital point must be understood about the previous paragraph's market-share figures: of the $900 billion of daily

[32] This statistic is from the Bank for International Settlements, BIS (1999). I examine BIS survey data in more detail in section 2.

volume that is not from the spot market, $734 billion of this is **FX swaps**,[33] and *FX swaps have no order-flow consequences in the FX market*. To understand why, one first needs to know what these swaps are. An FX swap bundles two FX transactions that go in opposing directions. For example, I agree to buy 100 million euros today for dollars (spot), and at the same time I agree to sell 100 million euros for dollars for settlement in one month (forward). This example is called a spot-forward swap. (One can also do forward-forward swaps, in which case the first of the two transactions is a nearer-dated forward transaction than the other.) Note that the two orders in this example are of equal size, but opposite sign, so the net order-flow impact is zero. Readers familiar with covered interest parity will recognize that this contract is a means of locking in an interest differential, and market participants use them for this purpose (whether hedging or speculating).[34] The net demand impact is therefore mainly on relative short-term interest rates, not on the FX market. This point is borne out in the behavior of banks: bankers tell me that when they design in-house models for forecasting exchange rates using order-flow, they exclude FX-swap transactions from their order-flow measures. Bottom line: of the transaction activity that corresponds to the order-flow models of this book, the spot market accounts for about $600 billion of $766 billion, or 78 percent.

Nevertheless, 78 percent is not 100 percent, so these statistics still highlight a tension in defining the scope of this book. On the one hand, defining the FX market broadly to include derivative instruments is consistent with official definitions like that of the BIS; moreover, arbitrage relationships link these sub-markets tightly, suggestive of one market rather than many. On the other hand, these sub-markets do not share the same market structure, particularly in the case of futures, which are often traded in pits using face-to-face open outcry.

To avoid these difficulties, henceforth I focus attention explicitly on the spot market, in particular the major floating-rate spot markets. (Fixed-rate spot markets are much smaller in terms of trading volume than floating-rate spot markets.) Unless otherwise noted, my use of the term "the FX market" corresponds to spot markets like $/euro and $/¥. Work thus far on FX microstructure is heavily concentrated on spot markets; broadening the scope to derivatives would bring us into uncharted terrain. I do not, however, exclude work on the derivatives segment—there are many notes and references to these related areas.[35]

3.1 Features of Spot FX Trading

Before digging deeper into institutional detail, it will be helpful to start at a general level by reviewing the basic institutional forms. Though in practice, most

---

[33] BIS (1999), page 17.

[34] Briefly, **covered interest parity** is a no-arbitrage condition that implies that $F_{\$/£}/P_{\$/£} = (1+i_{\$})/(1+i_{£})$, where $F_{\$/£}$ is today's one-period forward rate ($/£), $P_{\$/£}$ is today's spot rate, $i_{\$}$ is today's one-period nominal interest rate in dollars, and $i_{£}$ is today's one-period nominal interest rate in pound sterling. In economic (and order-flow) terms, taking offsetting positions on the left-hand side of this equation is equivalent to taking offsetting positions on the right-hand side (in markets free of capital controls).

[35] One fascinating branch of related work on derivatives involves derivative and spot market interaction (as opposed to analysis of a derivatives market in isolation). The question is whether introducing a derivatives market can rectify specific market failures. See, for example, Brennan and Cao (1996) and Cao (1999), among many others.

market structures are hybrids of these forms (or involve concurrent use of more than one form), appreciation of this fact requires familiarity with those forms. The three basic forms of market structure are:

(1) **auction markets**
(2) **single-dealer markets**
(3) **multiple-dealer markets.**

Naturally, within each of these forms there are further refinements. Since this is not intended as a survey, I offer a few words about each and provide some examples.

In an auction market (in particular, a "two-sided" auction market), a participant can submit a buy order, a sell order, or both. These orders can be market orders (buy X units now at the best available price) or limit orders (buy X units when the market reaches a price of Y). In a pure auction market, there is no explicit dealer, so the most competitive limit orders define the best available bid and offer prices. Examples of auction markets include the Paris Bourse and the Hong Kong Stock Exchange, both of which are operated electronically.[36]

In a single-dealer market, a lone dealer stands ready to buy at his bid quote and sell at his offer quote, thereby defining the best available price.[37] In this setting, incoming orders from customers are necessarily market orders—a customer either buys at the dealer's offer, sells at the dealer's bid, or chooses not to trade. Though some consider the "specialist system" of the New York Stock Exchange a single-dealer market, this is not accurate. In reality, the NYSE is a hybrid system, with both auction and single-dealer features. Each specialist (marketmaker) maintains a collection of customer limit orders—the limit-order book. If a market order to buy arrives, the specialist can either match it with the best (i.e., lowest priced) sell limit order, or, if he offers an even lower price himself, then he can take the other side. (In the parlance, he can step in front of the limit-order book.) Thus, the specialist must work within the parameters of the best buy and sell limit orders when trading for his own account. In this sense, the specialist is forced to compete against the limit-order book. This constrains the specialist's ability to exercise monopoly power. Pure examples of single-dealer markets are rare. Examples include FX markets in some developing countries where all orders must be routed through a single dealer—the central bank (developing-country markets of this type tend to be fixed-rate markets).

Multiple-dealer markets come in two main varieties—centralized and decentralized. In both cases, competition is provided via multiple competing dealers, rather than via limit orders as is the case in auction and hybrid specialist systems. In a centralized market, quotes from many dealers are available in a consolidated format, such as on a single screen (like the US's NASDAQ), or in a single physical place (like a futures trading pit). In a decentralized market, there is generally some

[36] There is a large literature on limit-order auction markets. See, for example, Glosten (1994), Biais et al. (1995), Chakravarty and Holden (1995), Harris and Hasbrouck (1996), Handa and Schwartz (1996), and Seppi (1996).
[37] The terms "dealer" and "marketmaker" are typically used interchangeably in the academic literature. When a distinction is made, the term dealer is used for dealership-market settings (like the FX market) and marketmaker is used for hybrid auction-dealership settings (like the New York Stock Exchange). My use of these terms will be consistent with this distinction throughout the book. I should also note that FX practitioners typically use the term "trader" to describe a dealer. Because the term trader is rather general, I opt for more specific terms when possible.

degree of fragmentation because not all dealer quotes are observable. One result of this fragmentation is that simultaneous transactions can occur at different prices.

The spot foreign exchange market is best described as a **decentralized multiple-dealer market**. (This is also true of forwards, options, and swaps markets in major currencies worldwide.) There is no physical location—or exchange—where dealers meet with customers, nor is there a screen that consolidates all executable dealer quotes in the market.[38] In this way, it is quite different from most equity and futures markets. In its structure, the spot FX market is perhaps most similar to the U.S. government bond market (bond markets have only recently attracted attention in the microstructure literature).[39]

Three characteristics in particular are often cited as distinguishing trading in FX from that in other markets:

(1) trading volume is enormous,
(2) trades between dealers account for most of this volume, and
(3) trade transparency is low.

Volume in the spot $/euro market alone is about $200 billion per day, dwarfing that of any other single financial instrument. Remarkably, interdealer trading currently accounts for roughly two-thirds of this volume, a much higher share than in other multiple-dealer markets (the remaining one-third is between dealers and non-dealer customers).[40] Finally, the FX market has an uncommon information structure. Specifically, order flow in FX is not as transparent as in other multiple-dealer markets: in most national markets—whether equity or bonds—trades must be disclosed within minutes by law. FX trades have no disclosure requirement. As a consequence, trades in this market are not generally observable. From a theoretical perspective, this feature is quite important because order flow conveys information about fundamentals. If order flow is not generally observed, then the trading process will be less informative, and the information reflected in prices will be reduced.

Let me clarify the players in the spot foreign exchange market. In addition to providing context, this will help to classify trades into types depending who are the counterparties. (The classification of trades into types is relevant to the material in later chapters.) The three main categories are:

(1) **dealers**
(2) **customers**
(3) **brokers (strictly interdealer)**

---

[38] There is a screen called Reuters FXFX that displays dealer quotes, but these quotes are not firm. (I provide more detail on the FXFX screen in chapter 5.) Though quotes on the electronic brokerage screens that I describe later in this section are firm, these brokerage quotes reflect only a subset of firm quotes in the market at any given time. For currency options, though some trading occurs on exchanges, most occurs in a decentralized, multiple-dealer setting. For currency futures, trading around the world tends to be centralized in various futures exchanges.

[39] Though the spot FX market and the U.S. government bond market currently share a similar structure, the way the bond market trades is evolving toward a more centralized auction structure (like the Paris Bourse and Hong Kong Stock Exchange, as described above). See section 10.3 for more detail. Other countries' bond markets have already moved to a centralized auction format (e.g., the Italian government bond market).

[40] See BIS (1999). The interdealer share of volume for NASDAQ and SEAQ (London's Stock Exchange Automated Quotations System) is less than 40% (see Reiss and Werner 1995).

Dealers provide two-way prices to both customers and other dealers. In major spot markets (like $/euro and $/yen), each dealer trades only a single currency pair. Though the number of banks that have a dealer in any major spot market is large (greater than 100 worldwide), the top ten banks handle the lion's share of the order flow, and concentration is increasing over time: over the past ten years, the combined market share of the top ten dealers has risen from around 40 percent to around 50 percent (BIS 1999, page 15; see also *Financial Times*, Survey: Foreign Exchange, 5 June 1998).

The customer category includes many institution types, such as non-financial corporations, financial firms/managers, and central banks. Chapter 9 provides some more detailed analysis of the trades of these individual customer categories. Also, chapter 8 provides analysis of the central bank category, and the role of these trades in intervention. (Because it is natural to consider central banks as having superior information, this customer type receives special attention in the macro literature.)

The term "broker" is confusing to people more familiar with equity markets than the FX market. Brokers in equity markets trade for both their customers and for themselves. FX brokers do not trade for themselves; they only facilitate trades between dealers. This facilitation role is important in the spot market. To understand the FX brokers' role, note that there are two methods for dealers to trade with one another. One way is for a dealer to call another for a quote and either sell at the quoted bid, or buy at the quoted offer. This is often referred to as a **direct** interdealer trade. The other method is to trade indirectly through a broker. (In 1998, about half of all interdealer trades in the largest spot markets were direct, which implies that the total-volume pie split rather cleanly into one-third customer-dealer, one-third direct interdealer, and one-third brokered interdealer.)

Think of brokers as a bulletin board. Brokers do not make prices themselves. They gather firm prices from dealers, and then communicate those prices back to dealers. A dealer might want to post a price through a broker because he prefers not to reveal his identity before the trade is executed (revealing one's identity before the trade is a necessary consequence of trading directly). For example, one dealer may post with the broker a limit order to buy a $5 million at a price of 10. Another dealer may post a limit order to sell $3 million at a price of 12. If these are the best prices the broker has received on either side, then the broker will advertise a two-way price of 10-12, and will do so without identifying the dealers behind those prices. A third dealer can choose to trade at one of those prices through the broker. (If so, after the transaction the broker reveals the counterparty, and settlement occurs directly between the counterparties; both pay the broker a small commission.) Thus, brokers are pure matchmakers—they do not take positions of their own, they only connect dealers that might not otherwise find each other. In the parlance of the three basic market structure forms above, brokers are running an interdealer auction market that operates concurrently with the multiple-dealer market. In this way, brokers provide a degree of centralization in an otherwise decentralized FX market.[41]

[41] One of the big changes in the FX market over the last few years is the shift away from voice-based brokers, where prices are advertised over intercoms at dealers' desks, and toward electronic brokers, where prices are advertised over a screen. Indeed, there is (informal) evidence that the electronic brokers have also taken market share away from direct interdealer trading. (I expect the 2001 BIS survey results to bear this out.) Currently, the dominant electronic broker in $/euro and $/yen is EBS. The other major electronic broker is Reuters 2000-2. For more on this shift to electronic brokerage see

These three categories of market participants give rise to three basic types of trades. We can illustrate these trade types using three concentric rings, shown below in Figure 3.1.

Figure 3.1

Three Types of Trades

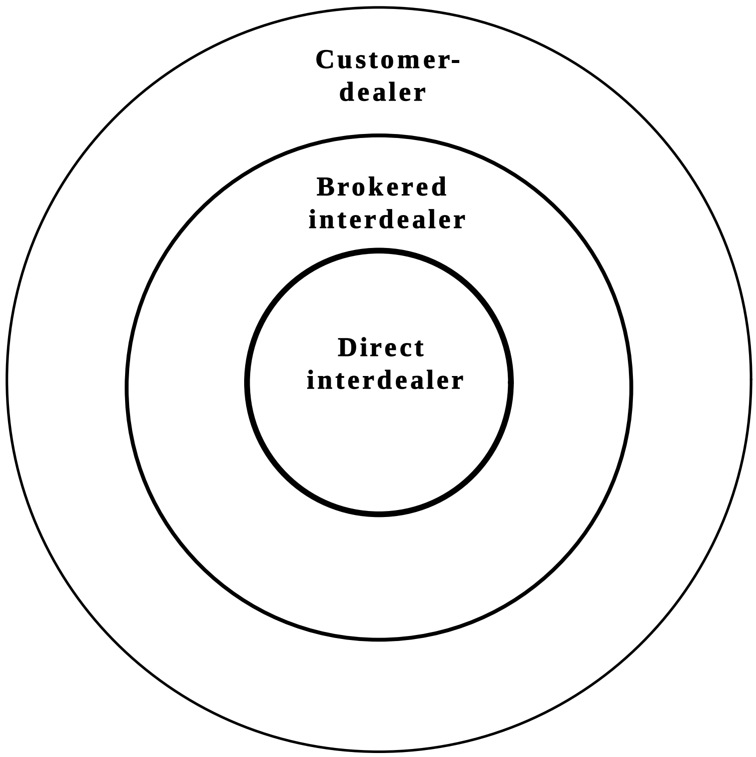

The inner ring represents direct interdealer trading, the most liquid part of the market. In the $/euro market, current spreads in this inner ring are 1-2 basis points (one basis point equals 1/100th of one percent) for $10 million trades (the standard size) between large banks during active trading hours (the full London trading day plus the morning hours in New York). Dealers typically choose direct trading for larger interdealer trades (above $10 million). The second ring represents brokered interdealer trading. The effective spread for a $10 million trade in this ring is roughly 2-3 basis points during active trading. I add the word "effective" here because the inside—i.e. lowest—spread in the brokered part of the market can be less than 2-3 basis points, but that inside spread often applies to trade sizes less than $10 million; any remainder must be executed at less attractive prices, such that a $10 million trade will have price impact beyond the initial inside spread. The third ring represents customer-dealer trading. Dealers tell me that current spreads

the *Financial Times* "Survey: Foreign Exchange," 5 June 1998. Interdealer brokers are also quite important in U.S. bond markets.

for a $10 million trade are in the 3-7 basis point range for "good" customers. ("Good" to most dealers means high volume.)

Visually, Figure 3.1 reflects a common metaphor used for the foreign exchange market, namely that the market is like a pool of water, with stones being thrown in the center, where the action is most intense. The stones are the customer orders. Direct interdealer trading lies at the center. Stones landing in that center send ripple effects through the brokered interdealer trading, and, ultimately, back to the customers themselves. Why back to customers themselves? Because dealers tend not to hold positions for very long in this market, as we shall see below. This metaphor also clarifies the typical order's "life cycle."[42]

One might ask why a dealer would use a broker if direct prices are tighter (brokers also charge a commission). Part of the answer is that smaller banks often do not have access to the tighter direct spreads among large banks. Large banks, too, have incentives to use brokers. From the large-bank perspective, providing a broker with a limit order provides a wider advertisement of a willingness to trade than bilateral direct quoting provides. (Keep in mind that a large bank that provides a limit order to buy, for example, is still buying at the bid price, which is below the offer price. If a second, smaller bank hits that bid, it is the second bank that sells at this lower price.) Another reason banks—large and small alike—may choose to trade via brokers is that they provide pre-trade anonymity (as noted above). In a direct interdealer trade, the dealer providing the quote knows the identity of the other dealer. (For more detail on incentives to use particular FX trading systems, see Luca 2000.)

Features for Modeling

With this more-complete picture of FX-market institutions, let us now consider other features of FX microstructure that influence modeling strategies. Among many cited in the literature, three in particular deserve note:

- **Dealers receive information from their customer orders**. As described by Citibank's head of FX in Europe "If you don't have access to the end user, your view of the market will be severely limited" (*Financial Times*, 4/29/91). In a similar spirit, Goodhart (1988) writes: "A further source of informational advantage to the traders is their access to, and trained interpretation of, the information contained in the order flow....Each bank will also know what their own customer enquiries and orders have been in the course of the day, and will try to deduce from that the positions of others in the market, and overall market developments as they unfold." Note that banks have little information regarding the customer

[42] I should mention two important variations on this story. First, banks do their best to "internalize" as much customer order flow as possible. That is, they want to match the incoming customer buy orders with incoming customer sell orders. When successful, any net flows that get passed on to the interdealer market are much reduced. Second, the order-flow life cycle does not really "end" once non-dealers reabsorb net balances. Rather, that reabsorption moves through a "chain" of liquidity providers, the first non-dealer link being hedge funds and banks' proprietary trading desks, and the last link being so-called "real money accounts." The real money accounts, such as mutual funds and pension funds, are institutions that absorb positions over longer horizons. Chapter 9 on customer trading provides more perspective.

orders of other banks. Consequently, insofar as this order-flow information helps forecast prices, it is private information (by the definition of chapter 2).

- **Dealers learn about marketwide order-flow largely from brokered interdealer trades**. Because dealers do not observe one-another's customer orders, they need to gather order-flow information from interdealer trading. As noted above, though, direct interdealer trading does not provide order-flow information to anyone other than the counter-parties. Brokered interdealer trading, on the other hand, does provide some order-flow information beyond the counter-parties. This is important: of the three trade types (customer-dealer, direct interdealer, and brokered interdealer), the brokered interdealer trades are the only order-flow information communicated to all dealers. The broker systems, which are now electronic, typically communicate this information by indicating whether incoming market orders are executed at the bid or offer side (indicated with a colored highlight, which provides the order flow's sign), and by providing information on how the incoming market order has changed the quantity available on the bid or offer side (information on the size of the order flow). Though there is noise in this order-flow measure, on a market-wide basis it is the best measure that is available to dealers.[43]

- **Large dealer positions are frequent and non-trivial**. They are a natural consequence of marketmaking in a fast-paced market with tight spreads (less than 2 basis points in the interdealer $/euro market).[44] FX dealers manage these large positions intensively. The large-bank dealer in the $/DM market that I tracked in 1992 (see Lyons 1995) finished his trading day with no net position each of the five days in the sample, despite trading over $1 billion each day. Within the day, the half-life of the gap between his current net position and zero was only ten minutes (Lyons 1998). From the plot of that dealer's net position in Figure 3.2 below, the strong reversion toward zero is readily apparent.

[43] One important source of noise comes from the fact that brokered-trading systems do not indicate the size of all individual transactions in real time, so it is not possible from information about transactions' signs to construct an exact order flow measure. (For example, on the D2000-2 system dealers only see an "r" on both the bid side and the offer side if the quantity available at the best price is $10 million or more, otherwise they see the quantity available. If, after a trade, the screen is still showing "r" on both the bid and offer sides, then one cannot infer the size. Often, though, the sizes drop below $10 million, or the quantity on one side is exhausted, in which case dealers do receive information about size. EBS, the other major electronic broker, is similar. Both systems also provide a high-frequency (but not complete) listing of deals as either "paid" or "given," where paid indicates buyer-initiated and given indicates seller-initiated.

Research data sets are a bit different. (In chapter 5 I review the data sets that cover brokered interdealer trading.) These data sets, constructed from records well after actual trading, do provide exact measures of order flow. But these exact ex-post measures were not in the dealers' information sets while trading.

[44] For comparison, a stock selling at $50 per share with a spread of only 10 cents would still translate into a spread of 20 basis points—10 times that in $/euro.

Figure 3.2

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

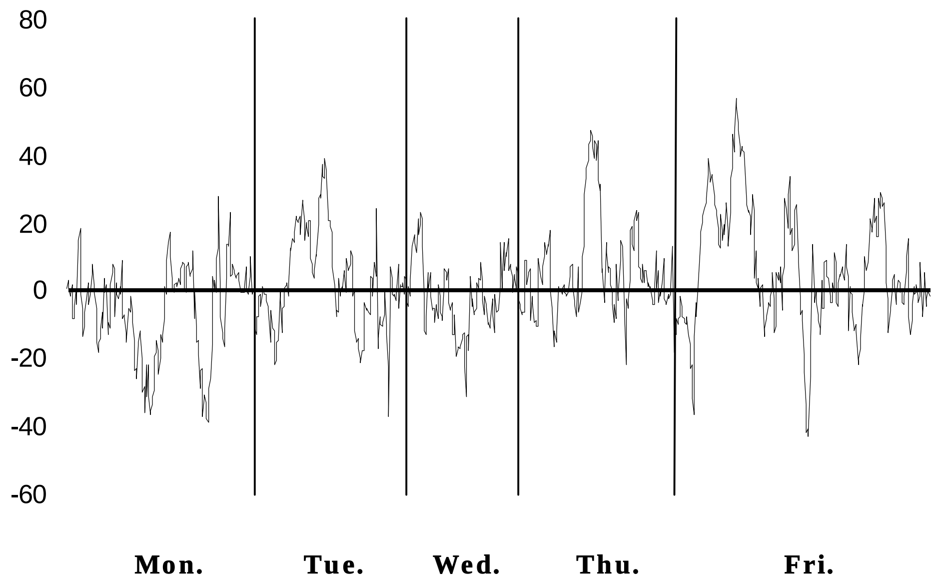

* The vertical lines represent the overnight periods over which this dealer was not trading. The horizontal distance between those vertical lines is scaled by the number of transactions made by this dealer each trading day.

Though the three features noted above are the most important from a modeling-strategy perspective, let me provide a bit more perspective on what the life of a dealer in a major FX market is like.[45] Two aspects of a dealer's life are particularly telling: his position sheets and his profitability. Tables 3.1 and 3.2 below provide some perspective on each of these two aspects. The position sheet is the source for the dealer position shown in Figure 3.2.

Table 3.1 provides a representative position sheet from the dollar-mark dealer I tracked through one week of trading in August 1992 (see Lyons 1998 for details). The dealer's position sheets provide the dealer with a running record of his net position and the approximate cost of that position. The dealer fills it in by hand as he trades (though these days it is mostly electronic and automatic). Each sheet (page) covers about fifteen transactions. The "Position" column accumulates the individual trades in the "Trade" column. Quantities are in millions of dollars. A

[45] There are two nice multimedia resources that bring still more perspective on the life of an FX dealer: (1) Goodhart and Payne (1999) and (2) Citibank's "Bourse Course." The former provides a visual account of trading in an electronic interdealer broker system (Reuters Dealing 2000-2). The latter is a simulated trading game designed to replicate the FX market; it includes trading of all three types depicted in figure 3.1. (Contact Citibank directly for more information.)

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 DM per dollar (which is the way dealers quoted prices in this market). 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 average inter-transaction time of this dealer is 1.8 minutes over this trading week, and during especially active periods it was much shorter). The “Source” column reports whether the trade is direct over the Reuters Dealing 2000-1 system[46] (r=Reuters) or brokered (b=Broker). One dimension of actual position sheets not shown in the diagram is the names of the counterparties. Most banks are reluctant to provide these names since they are considered confidential.

There are several take-aways. First, as is common for New York dealers, the trading day starts around 8:30 am, New York time (an important fact for empirical researchers needing to identify when New York comes “on line”). Second, the distinction between direct and brokered trades is evident in the way this dealer structures his sheet, with “r” capturing his direct trades and “b” capturing his brokered trades. (Beyond tracking position, one needs to track brokered trades because they involve a commission payment.) Third, trading is hectic enough that this dealer does not have time to update all information following every trade. (His average daily volume is over $1 billion, as we shall see in Figure 3.3.) Fourth, each entry does not record the time, though the dealer does record the time at the beginning of every card (to the minute), and occasionally within the card, as was the case here. Finally, and perhaps most important, one needs to remember that this is but one dealer. In many respects he appears to be representative; in other respects, he clearly is not. One important way in which he is not representative is that roughly 95 percent of his trades are interdealer (taking all his position cards together), compared to only about 80% of marketwide volume at the time (1992—the interdealer share has fallen to about two-thirds since then). Thus, the only types of trades that appear on this position sheet are either direct interdealer (r) or brokered interdealer (b) trades—no customer trades are listed (see the detailed discussion in chapter 5 of how these data compare to data on other dealers).

[46] Chapter 5 describes the Reuters Dealing 2000-1 system in detail, including the rich data it generates for empirical work.

Table 3.1

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

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.

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 second aspect of a dealer's livelihood that is particularly telling is her profitability. It is illuminating to get a sense for the source and size of a dealer's profits (particularly given the common, and not unreasonable view, that FX trading is a zero-sum game). Table 3.2 provides some perspective on these profits, based on the same dealer whose position sheet appears in Table 3.1. From the "Profit: Actual" column, we see that this dealer averages about $100,000 profit per day (on volume of about $1 billion per day). By comparison, equity dealers average about $10,000 profit per day (on volume of roughly $10 million per day).[47] So, even though this FX dealer's profit as a percentage of his volume is only one-tenth that of the average equity dealer, because his volume is 100 times as high, each day he earns 10 times as much.

To determine where these profits come from, first recognize that there are two possible sources: speculation and intermediation. Speculative profits come from being long dollars, on average, when the DM price of dollars goes up (and vice versa). Intermediation profits come from the bid-offer spread: buying at the lower price (bid) and selling at the higher price (offer). The column "Profit: Spread" is an attempt to impute the profit the dealer would have earned each day purely from intermediation. It is calculated under the assumption that the dealer earns one third of his spread on every transaction. Specifically, it is his daily dollar volume times one-third of his median quoted spread (0.0003 DM/$), divided by the average DM/$ rate over the sample (1.475 DM/$).[48]

Let me explain the rationale behind this assumption that the dealer earns one third of his spread on every transaction. Suppose the dealer has no net position, and quotes bid and offer prices of 1.4750 and 1.4753 DM/$, respectively (quotes apply to a standardized amount in this market, at the time $10 million). If the counter-party chooses to sell $10 million at 1.4750, then the dealer is long $10 million after the transaction. If the market has not moved, and another potential counter-party calls for a quote, this particular dealer will typically shade his price to induce the counter-party to buy—relieving the dealer of his long position. For example, this dealer would probably quote bid and offer prices to the next potential counter-party of 1.4749 and 1.4752 DM/$ (versus the original 1.4750 and 1.4753). Relative to the first pair of quotes, the new quote is attractive on the offer side—the 1.4752—but unattractive on the bid side. If the caller goes for the attractive offer quote, and buys, then the dealer will have sold his $10 million position at 1.4752 DM/$. The net result of both transactions is that he cleared two-thirds of his spread (two ticks) on two transactions—or one-third of his spread on each transaction.

From the last line of Table 3.2 we can see that, under this assumption, this dealer makes most all of his profit from the spread. Of the $507,929 he made over the week, our estimate of the amount that came from intermediation is $472,496. (If

---

[47] For the equity dealer comparison, I use the numbers from Hansch et al. (1999) for the London Stock Exchange because, unlike the NYSE, the LSE was a pure dealership market and therefore more comparable to FX. The authors find that dealers make a profit of roughly 10 basis points on the average transaction. Though the authors do not provide an average turnover by dealer, they do provide data that allow a rough estimate. The average daily turnover for FTSE-100 stocks is about $10 million (£6.9 million). This total turnover is divided among dealers, but active dealers make markets in many stocks. Given the market shares the authors report for the more active dealers, and given the number of stocks in which each makes markets, the estimated average turnover of $10 million per dealer is about right.

[48] The median quoted spread in the sample of 0.0003 DM/$ is the mode as well: that spread size accounts for about three-quarters of all the dealer's bilateral interdealer quotes.

I had assumed that the dealer makes half his spread from intermediation, then the profit from intermediation would have been higher than his total profit, indicating that he suffered speculative losses.) It appears that this dealer makes his money from intermediation. This is broadly consistent with the idea that—in terms of speculative profits—this market is a **zero-sum game**. The market need not be a zero-sum game in terms of intermediation profits, however. In that case, customers in this market are paying the dealers, on average, for the liquidity that the dealers provide.[49] Is that compensation inordinate? The data in Table 3.2 do not allow us to answer that question because these trades are mostly interdealer trades. $100,000 per day is not a bad day's work though. At least not where I come from.[50]

Table 3.2

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

	Transactions	Volume (mil)	Profit: Actual	Profit: Spread
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/$).

[49] The story is a bit subtler here, though, because this dealer does not have much customer business. In my judgment, the right way to think about this dealer is that he was supplying liquidity and inventory-management services to other dealers that have more customer business. So, in effect, there is a kind of "tiering" in the interdealer market—FX dealers are not a homogeneous lot. My understanding is that dealers of this type are much less profitable now that electronic interdealer brokers play such an important role (Table 3.2 corresponds to a trading week in 1992).
[50] An important determinant of whether $100,000 per day is "large" is the amount of bank capital this dealer ties up when trading. If the capital required were $1 billion per day—equal to his total volume—then $100,000 would be rather small. In reality, the capital required to support a dealing operation like this one, which involves only intraday positions, is far smaller than $1 billion.

3.2 Descriptive Statistics: The BIS Surveys

Unlike equity markets, because FX trades are not reported in most countries, marketwide volume in foreign exchange is in general not available. Every three years, however, individual central banks survey their financial institutions regarding FX trading activity (for a single month, typically April) and a snapshot becomes available. The latest triennial survey was in April 1998 (see the summary in BIS 1999—available at www.bis.org—and also individual banks' findings, e.g., Bank of England 1998 and Federal Reserve Bank of New York 1998).[51] This 1998 survey was the fifth triennial survey. Forty-three countries' central banks participated. Because these survey data provide institutional perspective not available from any other source, they warrant attention here.

Let us begin with the first table in the BIS (1999) report, which is reproduced below as Table 3.3. This table shows that the $1.5 trillion in daily volume in April 1998 is composed of $600 billion in spot trading and $900 billion in trading of outright forwards and forex swaps. As I noted earlier in this chapter, however, forex swaps have no order-flow consequences in the FX market; the net demand impact from this category of FX trading is mainly on relative short-term interest rates, not on the FX market. Of the transaction activity that corresponds to the order-flow models of this book, the spot market accounts for $600 billion of $766 billion, or 78 percent. This is an important point that is easily missed by readers of the BIS report.

[51] Beyond the BIS, there are several institutions that serve as semi-official coordinators in establishing FX trading practices and serving as forums for debate. In the U.S., that institution is the Foreign Exchange Committee of the Federal Reserve Bank of New York. For more information, see their web site at www.ny.frb.org/fxc.

Table 3.3

Foreign 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.

When comparing data across surveys, it is important to keep in mind that the coverage of the survey has changed quite a bit since the first survey in 1986. Only four countries participated in that first survey. In 1989, the number rose to 21 (but some countries did not provide all types of information). In 1992, 26 countries participated, including all countries with larger FX markets. In 1995 the number of countries remained the same, but the coverage of financial activity was expanded significantly to include FX-related financial derivatives. In 1998, the number of reporting countries increased to 43 and the coverage of derivatives activity was further expanded.[52]

Subject to the caveat of the previous paragraph, Table A-1 shows that spot turnover increased by 14% from 1989-92, 30% from 1992-95, and 15% from 1995-98. Notice, though, how much the total turnover statistics change when measured at

[52] The role of derivatives in determining market resilience and efficiency was a topic of increasing public policy concern in the 1990s. The considerable emphasis afforded derivatives in the BIS survey is in keeping with this policy concern. One should not lose sight, however, of the fact that in FX it is the spot markets that generate most of the order flow (per my point in the text about the largest of the FX derivative markets—that for forex swaps—generating lots of turnover, but no order flow).

the constant (April 1998) exchange rates. The big change is in the 1995 total turnover, which falls considerably when measured at the April 1998 rates (from 1,190 to 1,030 billion dollars). This is because the dollar was quite weak in April of 1995, particularly against the Japanese yen. Indeed, at the end of April 1995 the yen/$ rate fell to an unprecedented low of about 80. The translation to April 1998 exchange rates means that each dollar of 1995 turnover is scaled down to reflect that a dollar in April 1995 was worth less than a dollar in April 1998 (when measured against other currencies). If one applies the same constant-rate correction to the spot growth statistics, one finds that the three growth rates cited above—14%, 30%, and 15%—change to 10%, 15%, and 33%, respectively.[53] This changes the picture from slowing growth to accelerating growth. It will be interesting to see whether the statistics for spot turnover in the April 2001 survey continue this trend of accelerating growth. Currently, many market participants are predicting a slowing of spot-market growth, due to two factors: (1) the collapsing of many European cross markets into the euro and (2) the more efficient inventory management that is resulting from the dominance of electronic interdealer brokers (which can dampen the hot potato process described in chapter 1).

Another key table in the BIS report is Table B-4, which breaks down the turnover statistics by currency pair (reproduced in Table 3.4). As noted earlier in this chapter, the dollar-DM and dollar-yen spot markets are a good deal larger than any other. Although the dollar versus all other EMS currencies is listed third, most of this trading is in foreign exchange swaps, not spot. (These foreign exchange swap trades—being in fact trades on interest differentials—were likely unusually high in April 1998 due to speculation on the convergence of interest rates in the run-up to the January 1999 launch of the euro.) Note that markets in currencies against the dollar are the largest.[54] These major dollar exchange rates are floating rates (i.e., are market determined, with little or no intervention on the part of central banks). Not until the DEM/othEMS line does one find rates that are officially pegged (though not rigidly so—these rates were allowed to vary within pre-set bands). Now that trading in the euro has been introduced, there will be no need for a DEM/othEMS line in the 2001 survey table (the largest market, USD/DEM, will be replaced by the USD/Euro).

[53] The statistic 10% for 1989-92 is (390/356)–1, with 390 being the corrected spot turnover for 1992 of 400(800/820) and 356 being the corrected spot turnover for 1989 of 350(600/590). Corrected spot turnovers for the other years are calculated similarly. Note that this correction should be viewed as approximate for the spot turnover because the mix of currencies that are traded spot does not match exactly the mix of currencies in total turnover.

[54] Part of this large amount of dollar trading (but only part) is due to use of the dollar as a vehicle currency. When the dollar is used as a vehicle, then when trading two non-dollar currencies one does so indirectly, by going through the dollar first. With the launch of the Euro, the dollar's role as an international currency has received much attention recently. See, for example, Hartmann (1998) and Alogoskoufis, Portes, and Rey (1998). Much of the analysis in this area turns on microstructural matters; it is a natural application of microstructure tools to what has traditionally been a macro topic.

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.

Table E-1 from the Statistical Annex of BIS (1999)—not reproduced here—provides statistics on spot counterparty types. The table shows an interdealer share of 60% (347,689/577,737). Recall that in section 3.1 I described the interdealer share as roughly two-thirds. The central bank surveys upon which the BIS draws tend to underestimate the interdealer share in total trading because the category “other financial institutions” includes some non-reporting investment banks, some of which are important in dealing. (Central banks’ role in supervision and regulation applies mainly to commercial banks, so commercial banks are more thoroughly represented in the surveys.) The evidence that dealers are included in this “other financial institutions” comes from the fact that this category includes significant brokered trading; FX brokers are strictly interdealer, so these trades belong in the interdealer category. It is difficult to know how much this biases the survey-measured share downward; my adjustment from 60% to two-thirds is an educated guesstimate.

Section 7 of the BIS (1999) report provides some information on the share of interdealer trading that is brokered, and the degree to which these brokered trades are handled by electronic brokers, rather than the traditional voice-based brokers.[55] Because this section is not linked to specific tables, one needs to be cautious in interpreting the data. For example, one needs to take care to distinguish statistics that apply to total FX turnover, as opposed to spot turnover. (This is important throughout the BIS report.) One particularly useful sentence in that section is the following: "Electronic brokers now handle almost one quarter of total spot transactions in the UK market."[56] That is nearly one half of all interdealer spot transactions (per above, the survey finds that interdealer transactions are roughly 60% of total spot transactions). Although the BIS report goes a long way toward this end, to arrive at a more complete picture of the share and type of brokered trading—electronic versus voice-based—one needs to piece together data from the individual central bank reports.

3.3 Transparency of Order Flow

Any model that includes order flow as a proximate determinant of price must also specify who observes that order flow. In microstructure research, this issue is called **transparency**. The term transparency is broader than just the observability of order flow, however. It is defined to encompass the full array of information types that the trading process might transmit. The three primary categories include:

(1) Pre-trade versus post-trade information,
(2) price versus order flow (quantity) information, and
(3) public versus dealer information.

Applying this three-part taxonomy to the chapter-2 discussion of order-flow information, we see that the material of that chapter relates most directly to the (1) post-trade (2) order-flow information that is (3) available to dealers. In equity markets, this issue—what post-trade order-flow information is seen, and by whom—is at the heart of current policy debates.[57] Theoretically, too, post-trade order-flow information is the most relevant since it is the main communicator of shifts in asset demand. When interpreting this information as shifts in asset demand, however, one needs to be precise. One cannot infer the *sign* of a shift in demand from the information that, say, 10 units just traded. One needs to know whether the trade represents buying or selling pressure. The trade needs to be signed—it needs to be converted from trading volume to order flow.

---

[55] Do not be misled by the expression "transacted by brokers" that appears in the first sentence of section 7 (BIS 1999): brokers do not transact themselves; they merely facilitate the transactions of dealers (for a fee).
[56] The UK market is the biggest in terms of spot trading, accounting for about 28% of the worldwide spot total (Table E-9).
[57] There is a second major policy issue in equity-market transparency that is more relevant to pre-trade information: is the public entitled to see all the price quotes that the dealers observe?

Actual markets differ radically in terms of order-flow transparency. In equity markets, the transparency regime is typically imposed (e.g., regimes are imposed on the London Stock Exchange, the NYSE, and NASDAQ). On the London Stock Exchange, for example, the price and size of smaller trades must be disclosed within three minutes, while disclosure of the largest trades can be delayed up to five business days. FX markets, in contrast, have no disclosure requirements. For this reason, FX is particularly interesting since its degree of transparency has arisen without regulatory influence.[58] With no disclosure requirements, it is perhaps not surprising that most FX trades do not generate public order-flow information. But, as described earlier in this chapter, some trades do generate widely available order-flow information. Interestingly, these trades—the brokered interdealer trades—produce a level of transparency that arises as a by-product of dealer's selective use of this trading method. To summarize, the FX market is certainly not a case of purposeful transparency-regime design, as is true for most equity markets.

The only other financial markets similar to FX in terms of low transparency are other non-equity OTC markets. (OTC, or over-the-counter, simply means not traded on a centralized exchange.) These include the U.S. bond markets and much of the trading in derivatives. With the advent of centralized electronic trading in these other markets, however, they are on the way to becoming more transparent than the FX market.[59]

Now that we have a better sense from section 3.1 for how transparency arises in the FX market, we can examine the impact of this transparency on price determination. In markets that are highly transparent, order flow is observed by all participants, thereby affecting expectations—and prices—rapidly and precisely. In opaque markets, order flow is not widely observed, so the information it conveys may be impounded in price more slowly.[60] The FX market is opaque with respect to customer-dealer order flow. As noted, however, interdealer FX transactions are not completely opaque. One of the models I present in chapter 4 captures this differential dissemination of order-flow information, depending on order-flow type.

For market design, determining which participants see what, and when, is central. At a broad level, the key trade-off that concerns policymakers is the following: though greater transparency can accelerate revelation of information in price, it can also impede dealers' ability to manage risk. Full transparency may therefore not be optimal, and this must be considered when designing a transparency regime. Board and Sutcliffe (1995) make the point this way:

> The purpose of a transparency regime is to allow marketmakers to off-set inventory risk by trading before the market as a whole is aware of the large trade.

[58] I should add, though, that transparency in centralized exchange markets (stock and futures exchanges) was historically imposed by the members on themselves, prior to government regulation.

[59] In June of 2000, three investment banks (Goldman Sachs, Merrill Lynch, and Morgan Stanley Dean Witter) announced that they are launching a centralized electronic system for the US bond markets (government and corporate bonds). This system, if successful, represents a fundamental transformation of that market, not just in terms of increased transparency, but also in terms of access and the cost of liquidity.

[60] As an aside, there is an interesting argument why increasing transparency of the trading process might *reduce* the information in prices: increased transparency might reduce incentives to invest in information production. This effect does not arise in standard trading models because the amount of private and public information is assumed fixed.

It is unclear, however, whether the current low level of transparency of the FX market is socially optimal. Because low transparency has arisen without regulatory influence, a reasonable premise is that low transparency serves the interests of dealers. (Lyons 1996a provides a model in which this is true; basically, dealers prefer low transparency because it slows the pace at which price reflects information, enabling dealers to better manage risk—in keeping with the logic of the Board and Sutcliffe quote above.)

But if low transparency is an "equilibrium" outcome in the FX market, how is that equilibrium maintained? It is unlikely to result from collusion; collusion is difficult to maintain in competitive, decentralized markets. More likely is that equilibrium low transparency arises as a kind of externality—a by-product of dealers' individual decisions to trade using brokers (i.e., to trade using a transparent method). A fact that is broadly consistent with this view is that actual transparency levels produced by brokered trading are quite similar across the world's trading centers: the share of total trading that is brokered is in the 20 to 40 percent range in all the major trading centers (BIS 1996). Though it is true that dealers have other reasons for using brokers, such as anonymity, brokered trading does determine transparency, so dealers are indeed choosing transparency *de facto*. From a welfare perspective this is an important issue in institution design.

3.4 Moving Beyond Institutions

In many people's minds, there is a powerful association between the field of microstructure and the study of institutions. The association is natural, but also a bit deceptive. It is true that institution design is one of the "poles" within the microstructure literature; the material presented in this chapter is suggestive of this pole. But there is a second pole—the economics of financial-market information. This book is aligned more with the second pole. To emphasize this alignment, I presented the information framework in chapter 2, before this chapter on institutions. Later chapters, too, reflect primarily this second pole: they present models and methods that characterize how, in reality, the FX market aggregates dispersed information. The questions addressed within this second pole are of a broader nature than institution design.

Let me provide some more concrete examples of how microstructure's "information" pole extends beyond its "institutions" pole. The first example relates to order flow, and the sense in which order flow's role is not about institutions. Within microstructure, order flow is an information transmission mechanism, and, crucially, it operates regardless of market-structure type. Given this rather general property of order flow, it would be a mistake to attribute its information role to specific institutional configurations. Pushing further, for reasons apart from institutions per se, recognition of order flow's information role has the potential to realign thinking—at least within exchange-rate economics. To understand why, recall from chapter 1 that within exchange-rate economics, order flow's role in transmitting information was not considered. This omission is evident from surveys of macro-style empirical work. Consider, for example, Meese's (1990) survey, where he writes that:

> Omitted variables is another possible explanation for the lack of explanatory power in asset market models. However, empirical researchers have shown considerable imagination in their specification searches, so it is not easy to think of variables that have escaped consideration in an exchange rate equation.

It is hard to argue with Meese's point from a macro perspective. But the macro perspective considers only variables within the traditional macroeconomic set. From the microstructure perspective, there is indeed a variable that escaped macro consideration—order flow. Microstructure has opened macroeconomists' minds to the idea that order flow can serve as a precise, real-time measure of changing fundamentals.

Let me illustrate further why order flow's role has little to do with institutions per se. Consider the following thought experiment: suppose spot FX markets began to trade in a completely centralized auction format (akin to the Paris Bourse or Hong Kong Stock Exchange, as described above).[61] Would the microstructure approach still be useful? Yes, because order flow would still be a determinant of prices. True, one would need to measure order flow in a different way, because the passive (non-initiating) side of each trade would now be a limit order rather than a dealer's quote (see chapter 1). But with this changed measure of order flow, one could then produce the same analysis of how order flow determines exchange rates that I present in later chapters. It is unlikely that this would change the main results in a qualitative way: the underlying information structure of this market has more to do with the properties of the asset being traded—foreign exchange—than it does with the market structure per se, particularly at lower frequencies.

Analysis of crashes and collapses is a second example of how microstructure addresses questions of a broader nature than institution design. The global stock-market crash of 1987 attracted tremendous research attention, much of it set within microstructure models (see, e.g., Gennotte and Leland 1990 and Jacklin et al. 1992). Note, however, that equity prices collapsed across a host of different market-structure types (e.g., the specialist-market NYSE, the dealer-market NASDAQ, and several auction-type stock markets around the world as well). Because the crash was common to markets with different structures, one might argue that microstructure cannot help us understand the crash. But this would be too extreme: by providing a disciplined approach to complex information problems, microstructure models provide a useful way to understand crashes and collapses. This kind of analysis moves beyond narrow institutional concerns such as how auction and dealer markets differ.

Microstructure Effects versus Microstructure Approach

The term **microstructure effects** is very commonly used—particularly among people who do not work in microstructure. The concept deserves attention in this particular chapter because its connection to institutions is quite close. People typically use the term to refer to *temporary* effects on prices that arise from specific institutional features. As such, the term pertains more to the institutions pole of

[61] I consider this possibility further in chapter 10.

microstructure than the information pole. I place emphasis on the word "temporary" because the term is most often used to describe fleeting effects on asset prices that, at lower frequencies, are unlikely to be significant.

Figure 3.3 provides a graphic representation of microstructure effects and where they arise in the process of price determination. The left-hand column lists certain fundamental variables that might drive equilibrium prices. These variables do not, however, translate directly into price. Rather, they are inputs to the trading process, represented by the gray box. The microstructure effects "question" is whether the trading process alters the mapping from fundamental variables to price, and if so, for how long. The presumption is that it probably does, but that these effects are short lived.

Figure 3.3

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

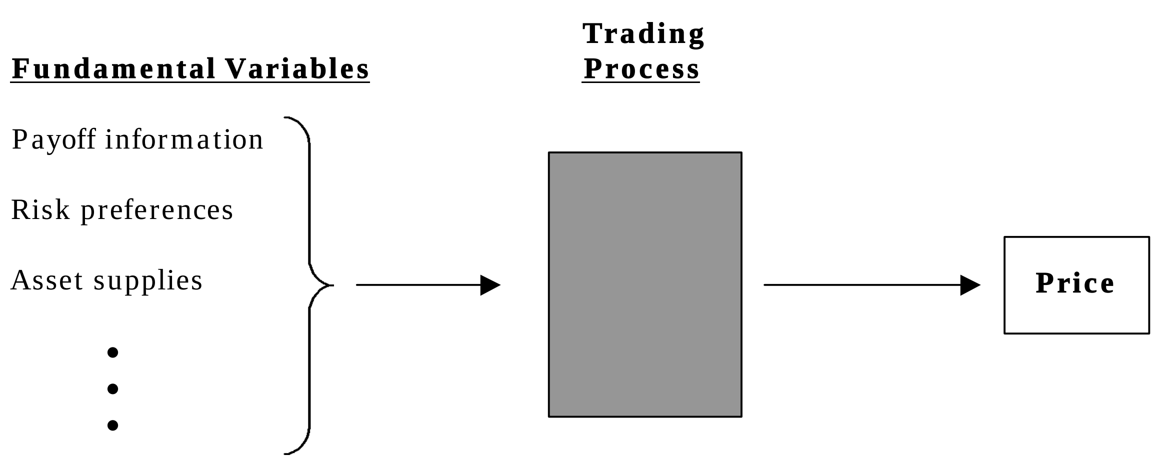

What is it that people really have in mind when they use the term microstructure effects? Though there is no explicit definition, in my judgment people use the term to refer to two particular types of effects. The first type is the temporary inventory effects described in chapter 2 (and diagrammed in Figure 2.1). As an empirical matter, inventory effects in FX are indeed likely to be short lived: as noted in section 3.1, the half-life of a dealer's inventory in FX is very short (ten minutes for the dealer whose inventory appears in Figure 3.2).

The second type of microstructure effect that people often have in mind is the price impact from order flow that accelerates the impounding of payoff information. The word "accelerate" is important here for understanding why this effect, too, might be only temporary. Consider the example of a pending earnings release by a firm and the possibility that insider trading prior to the announcement might hint at the quality of the earnings contained in the release. Figure 3.4 provides an illustration. The solid line is the price path under the assumption that the market responds to the (positive) public earnings announcement only. The dotted line shows how the price path would deviate if, in addition, the market were to respond to informative order flow occurring in advance of the public announcement. The difference in the two paths is only temporary here because the announcement reveals all the information contained in the prior order flow (and then some).

Figure 3.4

The Accelerationist view of order-flow information

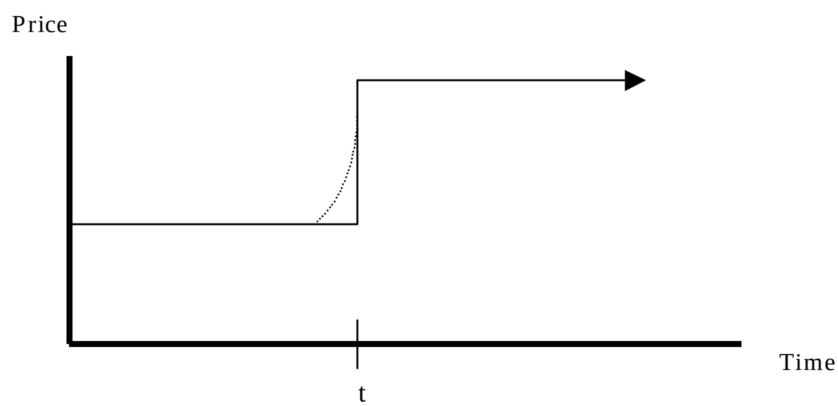

The solid line shows a hypothetical price path for a stock under the assumption that price responds only to a higher-than-expected public earnings announcement at time t. The dotted line shows the price path under the assumption that an insider is trading in advance of the announcement, and the information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

If, as an empirical matter, order flow were conveying only information that is on the verge of public announcement, then one would be justified in treating its effects as temporary, in the sense portrayed in Figure 3.4. This is certainly not the case for the FX market, however. Order flow's important role in determining exchange rates—documented in later chapters—is virtually unrelated to macroeconomic announcements that arrive within the subsequent year. The tight relation between order flow and exchange rates is not, therefore, simply the result of short-run acceleration of public information flow. Over the longer run, whether order flow conveys information about public news has yet to be determined. This is an active topic of ongoing research, one that I return to in chapter 7.

I raise the notion of microstructure effects because depicting the field of microstructure in too limited a way can affect research strategies. More specifically, it discounts microstructure's apparent potential, which can be counterproductive.

Let me provide an example embedded in the discussion above about analyzing crashes and collapses. In that crash discussion, I made the point that although the crash occurred across several different market-structure types, microstructure analysis was still fruitful for specifying information problems that can lead to a crash. More generally, people less familiar with microstructure are prone to assert that microstructure cannot resolve any puzzle that (1) is common to markets with different structures or (2) is not common to markets with the same structure. As an example of the former, it might inappropriately be argued that while (apparent) excess volatility is a property of both equity and FX markets, because the NYSE and FX markets have different structures, microstructure cannot help to resolve the excess volatility puzzle. This reasoning, in my judgment, is too oriented toward the institutions pole of microstructure. Effective use of information models within microstructure may indeed help to resolve puzzles in the FX market, even if the same puzzles occur in other, differently structured markets. Chapters 7 through 9 make this case.

References

Admati, Anat, and Paul Pfleiderer, 1988, A theory of intraday patterns: Volume and price variability, *Review of Financial Studies* 1, 3-40.

Allen, F. and G. Gorton, 1992, Stock price manipulation, market microstructure and asymmetric information, *European Economic Review*, 36, 624-630.

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

Ammer, J. and Brunner, A. 1997. Are banks market timers or market makers? Explaining foreign exchange trading profits, *Journal of International Financial Markets, Institutions & Money*, 7, 43-60.

Andersen, T., 1996, Return volatility and trading volume: An information flow interpretation of stochastic volatility, *Journal of Finance,* 51: 169-204.

Andersen, T., and T. Bollerslev, 1997, Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns, *Journal of Finance,* 52: 975-1006.

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.

Andersen, T., T. Bollerslev, and A. Das, 2000, Variance-ratio statistics and high-frequency data: Testing for changes in intraday volatility patterns, typescript, Northwestern University, *Journal of Finance*, forthcoming.

Bagehot, W., 1971, The only game in town, *Financial Analysts Journal*, 27: 12-22.

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

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

Baldwin, R., 1990, Reinterpreting the failure of foreign exchange market efficiency tests: Small transactions costs, large hysteresis bands, NBER Working Paper 3319.

Bank for International Settlements, 1996, Central bank survey of foreign exchange market activity in April 1995, publication of the Monetary and Economics Department, BIS, March.

Bank for International Settlements, 1999, Central bank survey of foreign exchange market activity in April 1998, publication of the Monetary and Economics Department, BIS, May (available at www.bis.org).

Bank for International Settlements, 1999b, A review of financial market events in autumn 1998, publication of the Committee on the Global Financial System, October.

Bank of England, 1998, The UK foreign exchange and over-the-counter derivatives markets in April 1998, September (available at www.bankofengland.co.uk).

Becker, T., B. Chadha, and A. Sy, 2000, Foreign exchange bid-ask spreads in the Asian crisis, typescript, International Monetary Fund, April.

Bessembinder, H. 1994. Bid-ask spreads in the interbank foreign exchange market. *Journal of Financial Economics* 35, 317-348.

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

Biais, B., 1993, Price formation and equilibrium liquidity in fragmented and centralized markets, *Journal of Finance*, 48: 157-184.
Biais, B., P. Hillion, and C. Spatt, 1995, An empirical analysis of the limit order book and the order flow in the Paris bourse, *Journal of Finance*, 50: 1655-1689.
Biais, B., P. Hillion, and C. Spatt, 1999, Price discovery and learning during the preopening period in the Paris Bourse, *Journal of Political Economy*, 107: 1218-1248.
Bishop, M., 1996, "Corporate Risk Management," *The Economist*, 2/10/96.
Bjonnes, G., and D. Rime, 2000, FX trading ... live: Dealer behavior and trading systems in foreign exchange markets, typescript, Norwegian School of Management, August.
Blanchard, O., and S. Fischer, 1989, *Lectures on Macroeconomics*, MIT Press: Cambridge, MA.
Blume, L., D. Easley, and M. O'Hara 1994. "Market Statistics and Technical Analysis: The Role of Volume, *Journal of Finance*, 49: 153-182.
Board, J., and C. Sutcliffe, 1995, The effects of trade transparency in the London stock exchange, report commissioned by the London Stock Exchange, January.
Bollerslev, T., and I. Domowitz, 1993, Trading patterns and prices in the interbank foreign exchange market, *Journal of Finance*, 48: 1421-1443.
Bollerslev, T., and M. Melvin, 1994, Bid-ask spreads and volatility in the foreign exchange market: An empirical analysis, *Journal of International Economics*, 36: 355-372.
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., 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.
Brennan, M., and H. Cao, 1996, Information, trade, and derivative securities, *Review of Financial Studies*, 9, 163-208.
Brennan, M., and H. Cao, 1997, International portfolio investment flows, *Journal of Finance*, 52, 1851-1880.
Brock, William, and Allan Kleidon, 1992, Periodic market closure and trading volume: A model of intraday bids and asks, *Journal of Economic Dynamics and Control* 16, 451-489.
Burnham, J. 1991. "Current Structure and Recent Developments in Foreign Exchange Markets," in S. Khoury, ed., *Recent Developments in International Banking and Finance*, Amsterdam: Elsevier North Holland.
Caballero, R., E. Engel, and J. Haltiwanger, 1997, Aggregate employment dynamics: Building from microeconomics, *American Economic Review*, 87: 115-137.
Cai, J, Y. Cheung, R. Lee, and M. Melvin, 1999, "Once in a generation" yen volatility in 1998: Fundamentals, intervention, or order flow? Typescript, Arizona State University, July.
Calvo, G., 1999, "Contagion in Emerging Markets: When *Wall Street* Is a Carrier," University of Maryland working paper.

Cao, H., 1999, The effect of derivative assets on information acquisition and price behavior in a rational expectations equilibrium, *Review of Financial Studies*, 12, 131-164.
Cao, H., and R. Lyons 1998. “Inventory Information” Working Paper, Haas School of Business, U.C. Berkeley, August.
Carlson, J., 1998, Risk aversion, foreign exchange speculation and gambler’s ruin, *Economica*, 65: 441-453.
Carlson, J., and C. Osler, 1999a, Determinants of currency risk premiums, typescript, New York Federal Reserve Bank, February.
Carlson, J., and C. Osler, 1999b, Rational Speculators and exchange rate volatility, *European Economic Review*, forthcoming.
Carrera, J., 1999, Speculative attacks to currency target zones: A market microstructure approach, *Journal of Empirical Finance*, 6, 555-582.
Chakravarty, S., and C. Holden, 1995, An integrated model of market and limit orders, *Journal of Financial Intermediation*, 4: 213-241.
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.
Chen, Z., 1997, Speculative market structure and the collapse of an exchange rate mechanism, working paper, London School of Economics.
Cheung, Y., and C. Wong, 1999, Foreign exchange traders in Hong Kong, Tokyo, and Singapore: A survey study, U.C. Santa Cruz Department of Economics Working Paper #425, January.
Cheung, Y., and C. Wong, 2000, A survey of market practitioners’ views on exchange rate dynamics, *Journal of International Economics*, 51, 379-400.
Cheung, Y., and M. Chinn, 1999a, Traders, market microstructure, and exchange rate dynamics, NBER Working paper 7416, November., *Journal of International Money and Finance*, forthcoming.
Cheung, Y., and M. Chinn, 1999b, Macroeconomic implications of the beliefs and behavior of foreign exchange traders, NBER Working paper 7417, November.
Cheung, Y., M. Chinn, and I. Marsh, 1999, How do UK-based foreign exchange dealers think their market operates?, CEPR Discussion Paper 2230, September.
Chinn, M., 1991, Some linear and non-linear thoughts on exchange rates, *Journal of International Money and Finance*, 10: 214-230.
Chinn, M., and R. Meese, 1994, Banking on currency forecasts, *Journal of International Economics*, 38: 161-178.
Clyde, W., 1993, Intraday foreign exchange market anomalies and microstructure, working paper, Quinnipiac College, August.
Copeland, T., and D. Galai, 1983, Information effects and the bid-ask spread, *Journal of Finance*, 38: 1457-1469.
Corsetti, G., S. Morris, and H. Shin, 1999, Does one Soros make a difference? The role of a large trader in currency crises, typescript, Yale University, March.
Covrig, V., and M. Melvin, 1998, Asymmetric information and price discovery in the FX market: Does Tokyo know more about the yen?, typescript, Arizona State University.
Dacorogna, M., U. Muller, R. Nagler, R. Olsen, and O. Pictet, 1993, A geographical model for the daily and weekly seasonal volatility in the FX market, *Journal of International Money and Finance*, 12: 413-438.

Danielsson, J., and R. Payne, 1999, Real trading patterns and prices in spot foreign exchange markets, typescript, London School of Economics, February.

Degennaro, R., and R. Shrieves. 1997. “Public Information Releases, Private Information Arrival, and Volatility in the Foreign Exchange Market,” *Journal of Empirical Finance*, 4: 295-315.

De Jong, F., R. Mahieu, P. Schotman, I. van Leeuwen, 1999, “Price Discovery on Foreign Exchange Markets with Differentially Informed Traders,” CEPR Discussion Paper No. 2296, November.

De Jong, F., T. Nijman, and A. Roell, 1995, Price effects of trading and components of the bid-ask spread on the Paris bourse, Discussion Paper 207, Financial Markets Group, London School of Economics.

Demsetz, H., 1968, The cost of transacting, *Quarterly Journal of Economics*, 82, 33-53.

Devereux, M., and C. Engel, 1999, The optimal choice of exchange-rate regime: Price-setting rules and internationalized production, NBER Working Paper 6992.

Diamond, D. and R. Verrecchia 1981. “Information Aggregation in a Noisy Rational Expectations Economy,” *Journal of Financial Economics*, 9: 221-235.

Dominguez, K., 1986, Are foreign exchange forecasts rational? New evidence from survey data, *Economic Letters*, 21: 277-281.

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

Dominguez, K., 1999, The market microstructure of central bank intervention, NBER Working Paper 7337, September.

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.

Domowitz, Ian, 1990, The mechanics of automated trade execution, *Journal of Financial Intermediation*, 1: 167-194.

Dooley, M., and P. Isard, 1982, A portfolio-balance rational-expectations model of the dollar-mark rate, *Journal of International Economics*, 12: 257-276.

Dornbusch, R., 1976, Expectations and exchange rate dynamics, *Journal of Political Economy*, 84: 1161-1176.

Dow, James, and Rohit Rahi, 1997, Should speculators be taxed? *Journal of Business*, January 2000.

Dumas, B., 1996, Comment on Jorion, in J. Frankel et al. (eds.), *The Microstructure of Foreign Exchange Markets*, University of Chicago Press: Chicago, 37-40.

Easley, D., S. Hvidkjaer, and M. O’Hara, 1999, Is information risk a determinant of asset returns? Typescript, Cornell University, November.

Easley, D., and M. O'Hara, 1987, Price, trade size, and information in securities markets, *Journal of Financial Economics*, 19: 69-90.

Easley, D., and M. O’Hara, 1992, Time and the process of security price adjustment, *Journal of Finance*, 47: 577-606.

Easley, D., N. Kiefer, M. O’Hara, and J. Paperman, 1996, Liquidity, information, and infrequently traded stocks, *Journal of Finance*, 51, 1405-1436.

Edison, H., 1993, The effectiveness of central-bank intervention: A survey of the literature after 1982, *Special Papers in International Economics*, No. 18, Princeton University.

Eichenbaum, M., and C. Evans, 1995, Some empirical evidence on the effects of shocks to monetary policy on exchange rates, *Quarterly Journal of Economics*, 110: 975-1009.

Eijffinger, S., and W. Verhagen, 1997, The advantage of hiding both hands: Foreign exchange intervention, ambiguity, and private information, Center for Economic Research Working Paper No. 9730, Tilburg University, January.

Engel, C., 1996, The forward discount anomaly and the risk premium: A survey of recent evidence, *Journal of Empirical Finance*, 3, 123-192.

Engle, R., and C. Granger, 1987, Cointegration and error correction: Representation, estimation and testing, *Econometrica*, 55: 251-276.

Engle, Robert, Takatoshi Ito, and Wen-Ling Lin, 1990, Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market, *Econometrica* 58, 525-542.

Engle, R., and J. Russell, 1998, Autoregressive conditional duration: A new model for irregularly spaced transaction data, *Econometrica*, 66, 1127-1162.

Euromoney, 2000, A true exchange for forex, July.

Evans, G., 1986, A test for speculative bubbles in the sterling-dollar exchange rate, *American Economic Review*, 76: 621-636.

Evans, Martin, 1997, The microstructure of foreign exchange dynamics, Georgetown University typescript, November.

Evans, Martin, 2000, FX Trading and Exchange Rate Dynamics, typescript, Georgetown University, October.

Evans, Martin, and Richard Lyons, 1999, Order flow and exchange rate dynamics, NBER Working Paper 7317, August, forthcoming in the *Journal of Political Economy*.

Evans, Martin, and Richard Lyons, 2000, The price impact of currency trades: Implications for Secret Intervention, U.C. Berkeley typescript, June, presented at the NBER Summer Institute, July 2000.

Fan, M., and R. Lyons, 2000, Customer-dealer trading in the foreign exchange market, typescript, U.C. Berkeley, July.

Federal Reserve Bank of New York, 1998, Foreign exchange and interest rate derivatives markets survey: Turnover in the United States, September (available at www.ny.frb.org).

Fieleke, Norman, 1981, Foreign-currency positioning by U.S. firms: Some new evidence, *Review of Economics and Statistics*, 63: 35-43.

Fleming, M., and E. Remolona, 1999, Price formation and liquidity in the U.S. treasury market: The response to public information, *Journal of Finance*, 54: 1901-1916.

Flood, M. 1991. "Microstructure Theory and the Foreign Exchange Market," *Review*, Federal Reserve Bank of St. Louis, November/December, 73: 52-70.

Flood, M. 1994. "Market Structure and Inefficiency in the Foreign Exchange Market," *Journal of International Money and Finance*, 13: 131-158.

Flood, M., R. Huisman, K. Koedijk, and R. Lyons, 1998, Search costs: The neglected spread component, presented at the HFDF-II conference, Zurich, April 1998.

Flood, M., R. Huisman, K. Koedijk, and R. Mahieu, 1999, Quote disclosure and price discovery in multiple-dealer financial markets, *Review of Financial Studies*, 12, 37-60.

Flood, M., R. Huisman, K. Koedijk, R. Mahieu, and A. Röell 1997. "Post-Trade

Transparency in Multiple-Dealer Financial Markets," typescript, Maastricht University, March.

Flood. R., and R. Hodrick, 1990, On testing for speculative bubbles, *Journal of Economic Perspectives*, 4: 85-101.

Flood, R., and A. Rose, 1995, Fixing exchange rates: A virtual quest for fundamentals, *Journal of Monetary Economics*, 36: 3-37.

Flood, R., and A. Rose, 1996, Fixes: Of the forward discount puzzle, *Review of Economics and Statistics*, 78: 748-752.

Flood, Robert, and Mark Taylor, 1996, Exchange rate economics: What's wrong with the conventional macro approach?, in Frankel, J., G. Galli, and A. Giovannini, (eds.) *The Microstructure of Foreign Exchange Markets* (Chicago: The University of Chicago Press), 261-294.

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

Foster, Douglas, and S. Viswanathan, 1993, Variations in trading volume, return volatility, and trading costs: Evidence on recent price formation models, *Journal of Finance* 48, 187-211.

Frankel, J., 1982, A test of perfect substitutability in the foreign exchange market, *Southern Economic Journal*, 46: 406-416.

Frankel, J., 1996, How well do markets work: Might a Tobin tax help?, in M. ul Haq et al. (eds.), *The Tobin Tax: Coping with Financial Volatility* (Oxford University Press: New York).

Frankel, J., and C. Engel, 1984, Do asset demand functions optimize over the mean and variance of real returns? A six-currency test, *Journal of International Economics*, 17: 309-323.

Frankel, J., and K. Froot, 1987, Using survey data to test standard propositions regarding exchange rate expectations, *American Economic Review*, 77: 133-153.

Frankel, J., and K. Froot, 1990, Chartists, fundamentalists, and trading in the foreign exchange market, *American Economic Review*, 80: 181-185.

Frankel, J., G. Galli, and A. Giovannini (eds.), 1996, *The Microstructure of Foreign Exchange Markets* (Chicago: The University of Chicago Press).

Frankel, J., G. Galli, and A. Giovannini, 1996, Introduction, in *The Microstructure of Foreign Exchange Markets* (Chicago: The University of Chicago Press), 1-15.

Frankel, Jeffrey, and Andrew Rose, 1995, Empirical research on nominal exchange rates, in G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, Elsevier Science: Amsterdam, 1689-1729.

French, Kenneth, and Richard Roll, 1986, Stock return variance: The arrival of information and the reaction of traders, *Journal of Financial Economics* 17, 99-117.

Frenkel, J., 1976, A monetary approach to the exchange rate: Doctrinal aspects and empirical evidence, *Scandinavian Journal of Economics*, 78, 200-224.

Froot, K., and E. Dabora, 1999, How are stock prices affected by the location of trade? *Journal of Financial Economics*, 53: 189-216.

Froot, K., and K. Rogoff, 1995, Perspectives on PPP and long-run real exchange rates, in G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, Elsevier Science: Amsterdam, 1647-1688.

Froot, K., and R. Thaler, 1990, Anomalies: Foreign exchange, *Journal of Economic*

Perspectives, 4: 179-192.

Galati, G., 2000, Trading volumes, volatility, and spreads in FX markets: Evidence from emerging market countries, typescript, Bank for International Settlements, July.

Garman, M. 1976. Market microstructure. *Journal of Financial Economics* 3, 257-275.

Gehrig, T., and L. Menkhoff, 2000, The use of flow analysis in foreign exchange: Exploratory evidence, typescript, University of Freiburg, Germany, August.

Gennotte, G., and H. Leland, 1990, Market liquidity, hedges, and crashes, *American Economic Review*, 80: 999-1021.

Gersbach, H., and K. Vogler, 1998, Learning from informed and uninformed traders and the role of interdealer trading, University of Heidelberg typescript, March.

Glassman, D. 1987. Exchange rate risk and transactions costs: Evidence from bid-ask spreads. *Journal of International Money and Finance* 6, 479-490.

Glosten, L., 1989, Insider trading, liquidity, and the role of the monopolist specialist, *Journal of Business*, 62: 211-236.

Glosten, L., 1994, Is the electronic open limit order book inevitable? *Journal of Finance*, 49: 1127-1162.

Glosten, L., and L. Harris, 1988, Estimating the components of the bid-ask spread, *Journal of Financial Economics*, 21, 123-142.

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

Goldberg, L., 1993, Foreign exchange markets in Russia: Understanding the reforms, *International Monetary Fund Staff Papers*, 40: 852-864.

Goldberg, L., and R. Tenorio, 1997, Strategic trading in a two-sided foreign exchange auction, *Journal of International Economics*, 42: 299-326.

Goldstein, M., and K. Kavajecz, 2000, Liquidity provision during circuit breakers and extreme market movements, typescript, Babson College, November.

Goodhart, C. 1988. "The Foreign Exchange Market: A Random Walk with a Dragging Anchor," *Economica*, 55: 437-460.

Goodhart, C. and Figliuoli, L. 1991. Every minute counts in financial markets. *Journal of International Money and Finance* 10, 23-52.

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

Goodhart, C., Ito, T. and Payne, R. 1996. One day in June 1993: A study of the working of the Reuters 2000-2 electronic foreign exchange trading system. In *The Microstructure of Foreign Exchange Markets*, eds. J. Frankel, G. Galli and A. Giovannini, pp. 107-179. University of Chicago Press, Chicago, IL.

Goodhart, C., and M. O'Hara, 1997, High frequency data in financial markets: Issues and applications, *Journal of Empirical Finance*, 4: 73-114

Goodhart, C., and R. Payne, 1996., Micro-structural dynamics in a foreign exchange electronic broking system, *Journal of International Money and Finance*, December.

Goodhart, C., and R. Payne, 1999, The foreign exchange market: A multimedia representation, compact disk, distributed by En-terprise LSE Ltd (www.enterprise-lse.co.uk).

Goodhart, C., and M. Taylor, 1992, Why don't individuals speculate in the forward foreign exchange market? *Scottish Journal of Political Economy*, 39: 1-13.

Gordon, R., and L. Bovenberg, 1996, Why is capital so immobile internationally? Possible explanations and implications for capital income taxation, *American Economic Review*, 86: 1057-1075.

Grammatikos, T. and A. Saunders, 1986, Futures price variability: A test of maturity and volume effects, *Journal of Business*, 59, 319-330.

Grossman, S., 1977. "The Existence of Futures Markets, Noisy Rational Expectations and Informational Externalities," *Review of Economic Studies*, 44: 431-449.

Grossman, S., 1988, An analysis of the implications for stock and futures price volatility of program trading and dynamic hedging strategies, *Journal of Business*, 61: 275-298.

Grossman, S., and M. Miller, 1988, Liquidity and market structure, *Journal of Finance*, 43: 617-633.

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

Grossman, S. and L. Weiss, 1983, A transactions-based model of the monetary transmission mechanism, *American Economic Review*, 73: 871-880.

Grundy, B., and M. McNichols 1989. "Trade and the Revelation of Information Through Prices and Direct Disclosure," *Review of Financial Studies*, 2: 495-526.

Guillaume, D., M. Dacorogna, R. Dave, U. Muller, R. Olsen, and O. Pictet, 1995, From the bird's eye to the microscope: A survey of new stylized facts of the intradaily foreign exchange markets, *Finance and Stochastics*, 1, 95-129.

Habermeier, K., and A. Kirilenko, 2000, Securities transaction taxes and financial markets, typescript, International Monetary Fund, June.

Hamao, Y., and J. Hasbrouck, 1995, Securities trading in the absence of dealers: Trades and quotes on the Tokyo stock exchange, *Review of Financial Studies*, 8, 849-878.

Hamilton, J., 1994, *Time Series Analysis*, Princeton University Press: Princeton, NJ.

Handa, P., and R. Schwartz, 1996, Limit order trading, *Journal of Finance*, 51: 1835-1861.

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.

Hansch, O., N. Narayan, and S. Viswanathan, S, 1999, Preferencing, internalization, best execution, and dealer profits, *Journal of Finance*, 54, 1799-1828.

Harris, L., 1986, A transaction data study of weekly and intradaily patterns in stock returns, *Journal of Financial Economics*, 16: 99-117.

Harris, L., and J. Hasbrouck, 1996, Market versus limit orders: The SuperDOT evidence on order submission strategy, *Journal of Financial & Quantitative Analysis*, 31, 213-231.

Harris, M., and A. Raviv, 1993, Differences of opinion make a horse race, *Review of Financial Studies* 6, 473-506.

Hartmann, P., 1998a, *Currency competition and foreign exchange markets: The dollar, the yen, and the euro*, Cambridge University Press: Cambridge.

Hartmann, P., 1998b, Do Reuters spreads reflect currencies' differences in global trading activity? *Journal of International Money and Finance*, 17: 757-784.

Hartmann, P., 1998c, Trading volumes and transaction costs in the foreign exchange market: Evidence from daily dollar-yen spot data, *Journal of Banking and Finance*, forthcoming.
Hasbrouck, J. 1988. Trades, quotes, inventories, and information. *Journal of Financial Economics* 22, 229-252.
Hasbrouck, J., 1991a, Measuring the information content of stock trades, *Journal of Finance*, 46: 179-207.
Hasbrouck, J., 1991b, The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies*, 4, 571-595.
Hasbrouck, J., 1992, Using the TORQ database, typescript, New York University (copy available at jhasbrouck.stern.nyu.edu/main).
Hasbrouck, J. 1997, One security, many markets, typescript, New York Univeristy.
Hasbrouck, J. and Sofianos, G. 1993. The trades of market makers: An empirical analysis of NYSE specialists. *Journal of Finance* 48, 1565-1593.
Hau, H., 1998, Competitive entry and endogenous risk in the foreign exchange market, *Review of Financial Studies*, 11: 757-788.
Hau, H., 2000, Information and geography: Evidence from the German stock market, typescript, Insead, March.
Hau, H., and A. Chevallier, 2000, Evidence on the volatility effect of a Tobin tax, typescript, Insead, April.
Hau, H., W. Killeen, and M. Moore, 2000, The euro as an international currency: Explaining the first evidence, typescript, Insead, April.
Hayek, F., 1945, The use of knowledge in society, *American Economic Review*, September.
He, H., and D. Modest, 1995, Market frictions and consumption-based asset pricing, *Journal of Political Economy*, 103: 94-117.
Heere, E., 1999, Microstructure theory applied to the foreign exchange market, dissertation, Maastricht University, Faculty of Economics and Business Administration, February.
Hellwig, M. 1980. On the aggregation of information in competitive markets, *Journal of Economic Theory*, 22: 447-498.
Hellwig, M., 1982, Rational expectations equilibrium with conditioning on past prices: A mean variance example, *Journal of Economic Theory*, 26, 279-312.
Hicks, J., 1939, *Value and Capital*, Clarendon Press: Oxford.
Ho, Thomas, and Hans Stoll, 1983, The dynamics of dealer markets under competition, *Journal of Finance* 38, 1053-1074.
Hodrick, R., 1987, *The empirical evidence on the efficiency of forward and futures foreign exchange markets*, Harwood.
Holden, Craig, and Avanidhar Subrahmanyam, 1992, Long-lived private information and imperfect competition, *Journal of Finance* 47, 247-270.
Hong, H., and J. Wang 1995. "Trading and Returns Under Periodic Market Closures," MIT working paper, July.
Houthakker, H., 1957, Can speculators forecast prices? *Review of Economic Studies*, 39: 143-151.
Hsieh, David, 1988, The statistical properties of daily foreign exchange rates: 1974-1983, *Journal of International Economics*, 24: 129-145.
Hsieh, David, and Allan Kleidon, 1996, Bid-ask spreads in foreign exchange markets: Implications for models of asymmetric information, in Jeffrey

Frankel et al. eds. *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago) 41-65.

Huang, R., and R. Masulis, 1999, FX spreads and dealer competition across the 24 hour trading day, *Review of Financial Studies*, 12, 61-94.

Huang, R., and H. Stoll, 1997, The components of the bid-ask spread: A general approach, *Review of Financial Studies*, 10, 995-1034.

Huisman, R., K. Koedijk, C. Kool, and F. Nissen, 1998, Extreme support for uncovered interest parity, *Journal of International Money and Finance*, 17: 211-228.

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

Institutional Investor, 2000, Trading meets the millennium, January.

Isard, P., 1995, *Exchange Rate Economics*, Cambridge University Press: Cambridge, England.

Ito, Takatoshi, Michael Melvin, and Richard Lyons, 1998, Is there private information in the FX market? The Tokyo experiment, *Journal of Finance*, 53: 1111-1130.

Ito, Takatoshi, and V. Vance Roley, 1990, Intraday yen/dollar exchange rate movements: News or noise? *Journal of International Financial Markets, Institutions and Money*, 1: 1-31.

Jacklin, C., A. Kleidon, and P. Pfleiderer, 1992, Underestimation of portfolio insurance and the crash of October 1987, *Review of Financial Studies*, 5: 35-63.

Jeanne, O. and A. Rose, 1999, Noise trading and exchange rate regimes, NBER Working Paper #7104, April, *Quarterly Journal of Economics*, forthcoming.

Johansen, S., 1992, Cointegration in partial systems and the efficiency of single equation analysis, *Journal of Econometrics*, 52: 389-402.

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

Jordan, J., and R. Radner, 1982, Rational expectations in microeconomic models: An overview, *Journal of Economic Theory*, 26, 201-223.

Jorion, P., 1996, Risk and turnover in the foreign exchange market, in J. Frankel et al. (eds.), *The Microstructure of Foreign Exchange Markets*, University of Chicago Press: Chicago, 19-36.

Kaminsky, G., and K. Lewis, 1996, Does foreign exchange intervention signal future monetary policy? *Journal of Monetary Economics*, 37: 285-312.

Kandel, Eugene, and Neil Pearson, 1995, Differential interpretation of public signals and trade in speculative markets, *Journal of Political Economy* 103, 831-872.

Kang, J., and R. Stulz, 1997, Why is there a home bias? An analysis of foreign portfolio equity ownership in Japan, *Journal of Financial Economics*, 46, 3-28.

Karpoff, J., 1986, A theory of trading volume, *Journal of Finance*, 41: 1069-1087.

Karpoff, J., 1987, The relation between price changes and volume: A survey, *Journal of Financial and Quantitative Analysis*, 22, 109-126.

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.

Keynes, J., 1936, *The General Theory of Employment, Interest, and Money*, Harcourt Brace: New York (see especially chapter 12).

Killeen, W., R. Lyons, and M. Moore, 2000a, Fixed versus floating exchange rates: Lessons from order flow, typescript, U.C. Berkeley, July.

Killeen, W., R. Lyons, and M. Moore, 2000b, The puzzle of the euro, typescript, U.C. Berkeley, July.

Kim, O., and R. Verrecchia, 1991, Trading volume and price reactions to public announcements, *Journal of Accounting Research* 29, 302-321.

Kirilenko, A., 1997, Endogenous trading arrangements in emerging foreign exchange markets, working paper, University of Pennsylvania (Wharton).

Klein, M., 1993, The accuracy of reports of foreign exchange intervention, Journal of International Money and Finance, 12: 644-653.

Kouri, P., and M. Porter, 1974, International capital flows and portfolio equilibrium, *Journal of Political Economy*, 82: 443-467.

Krugman, Paul, and Marcus Miller, 1993, Why have a target zone? *Carnegie-Rochester Conference Series on Public Policy*, 38: 279-314.

Kurz, M., 1997, Foreign exchange rates under alternative expectational paradigms: A resolution of the excess volatility puzzle, working paper, Stanford University.

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

Kyle, Albert, 1989, Informed speculation with imperfect competition, *Review of Economic Studies* 56, 317-356.

Leach, J., and A. Madhavan 1993. "Price Experimentation and Security Market Structure," *Review of Financial Studies*, 6: 375-404.

Lee, C., and M. Ready, 1991, Inferring trade direction from intradaily data, *Journal of Finance*, 46, 733-746.

Leland, H., 1992, Insider trading: Should it be prohibited? *Journal of Political Economy*, 100, 859-887.

Lewis, K., 1988, Testing the portfolio balance model: A multilateral approach, *Journal of International Economics*, 24: 109-127.

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

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

Luca, C., 2000, *Trading in the global currency markets*, second edition, Prentice Hall (New York Institute of Finance): New York.

Lui, Y., and D. Mole, 1998, The use of fundamental and technical analyses by foreign exchange dealers: Hong Kong evidence, *Journal of International Money and Finance*, 17: 535-545.

Lyons, R. 1995. "Tests of Microstructural Hypotheses in the Foreign Exchange Market," *Journal of Financial Economics*, 39: 321-351.

Lyons, R. 1996a. "Optimal transparency in a dealer market with an application to foreign exchange," *Journal of Financial Intermediation* 5, 225-254.

Lyons, R. 1996b. "Foreign Exchange Volume: Sound and Fury Signifying Nothing?" in J. Frankel et al. (eds.), *The Microstructure of Foreign Exchange Markets*, University of Chicago Press: Chicago, 183-201.

Lyons, R. 1996c. "Comment on A Study of the Reuters D2000-2 Foreign Exchange Trading System," in J. Frankel et al. (eds.), *The Microstructure of Foreign Exchange Markets*, University of Chicago Press: Chicago, 183-201.

Lyons, R. 1996d. "Microstructure of the Foreign Exchange Market," in T. Ito and D. Folkerts-Landau, eds., *International Capital Markets*, International Monetary Fund, September 1996, 125.

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

Lyons, R. 1997b. "Explaining Trading Volume in Foreign Exchange: Lessons From Tokyo," *Economic Letter*, Federal Reserve Bank of San Francisco, #97-38, December 26.

Lyons, R. 1998. "Profits and position control: A week of FX dealing," *Journal of International Money and Finance*, 17: 97-115.

Lyons, R., and A. Rose. 1995. "Explaining Forward Exchange Bias ... Intraday," *Journal of Finance* 50: 1321-1329.

Madhavan, A. 1992, Trading Mechanisms in Securities Markets, *Journal of Finance*, 47: 607-642.

Madhavan, A. 1996. Security prices and market transparency, *Journal of Financial Intermediation*, 5, 255-283.

Madhavan, A., 2000, Market microstructure: A survey, *Journal of Financial Markets*, 3: 205-258.

Madhavan, A. and Smidt, S. 1991. A bayesian model of intraday specialist pricing. *Journal of Financial Economics* 30, 99-134.

Madhavan, A. and Smidt, S. 1993. An analysis of changes in specialist inventories and quotations. *Journal of Finance* 48, 1595-1628.

Madrigal, Vicente, 1996, Non-fundamental speculation, *Journal of Finance* 51, 553-578.

MAR/Hedge, 2000, Managed Account Reports, Inc.: New York (www.marhedge.com), monthly issues

Mark, N., 1995, Exchange rates and fundamentals: Evidence on long-horizon predictability, *American Economic Review*, 85: 201-218.

Meese, R., 1986, Testing for bubbles in exchange markets, *Journal of Political Economy*, 94: 345-373.

Meese, R., 1990, Currency fluctuations in the post-Bretton Woods era, *Journal of Economic Perspectives*, 4: 117-134.

Meese, R., and K. Rogoff, 1983a, Empirical exchange rate models of the seventies, *Journal of International Economics*, 14: 3-24.

Meese, R., and K. Rogoff, 1983b, The out-of-sample failure of empirical exchange rate models, in J. Frenkel (ed.), *Exchange Rate and International Macroeconomics*, University of Chicago Press: Chicago.

Mehra, R., and E. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics*, 15: 145-161.

Melvin, M., and B. Peiers, 1997, The global transmission of volatility in the foreign exchange market, working paper, Arizona State University.

Melvin, M., and X. Yin, 2000, Public information arrival, exchange rate volatility, and quote frequency, *Economic Journal*, July.

Mendelson, Haim, 1987, Consolidation, fragmentation, and market performance, *Journal of Financial and Quantitative Analysis* 22, 189-208.

Menkhoff, L., 1998, The noise trading approach—Questionnaire evidence from foreign exchange, *Journal of International Money and Finance*, 17: 547-564.
Montgomery, J., and H. Popper, 1998, Information sharing and central bank intervention in the foreign exchange market, typescript, Santa Clara University, December.
Moore, M., and M. Roche, 1999, Less of a puzzle: A new look at the forward forex market, typescript, Queen’s University of Belfast, June.
Morris, S., and H. Shin, 2000, Rethinking multiple equilibria in macroeconomic modeling, *NBER Macroeconomics Annual 2000*, forthcoming.
Mussa, M., 1976, The exchange rate, the balance of payments, and monetary and fiscal policy under a regime of controlled floating, *Scandinavian Journal of Economics*, 78, 229-248.
Muth, J., 1960, Optimal properties of exponentially weighted forecasts, *Journal of the American Statistical Association*, 55: 290-306.
Naik, N., A. Neuberger, and S. Viswanathan, 1998, Trade disclosure regulation in competitive dealer markets, working paper, London Business School.
Naik, N., and P. Yadav, 2000, Do market intermediaries hedge their risk exposure with derivatives? Typescript, London Business School, February.
Naranjo, A., and M. Nimalendran, 1998, Government intervention and adverse selection costs in foreign exchange markets, working paper, University of Florida, October, forthcoming, *Review of Financial Studies*.
New York Federal Reserve Bank, 1995: “The Foreign Exchange Committee Annual Report, 1995”.
Obstfeld, M., and K. Rogoff, 1995, Exchange rate dynamics redux. *Journal of Political Economy*, 103: 624-660.
Obstfeld, M., and K. Rogoff, 1996, *Foundations of International Macroeconomics*, MIT Press: Cambridge, MA.
Obstfeld, M., and A. Taylor, 1997, Non-linear aspects of goods-market arbitrage and adjustment: Heckscher’s commodity points revisited, *Journal of the Japanese and International Economies*, 11: 441-479.
O'Hara, M. 1995. *Market Microstructure Theory*. Blackwell Business, Cambridge MA.
O’Hara, M., 1999, Making market microstructure matter, *Financial Management*, 28, 83-90.
O’Hara, Maureen, and George Oldfield, 1986, The microeconomics of market making, *Journal of Financial and Quantitative Analysis* 21, 361-376.
Osler, C., 1995, Exchange rate dynamics and speculator horizons, *Journal of International Money & Finance*, 14: 695-720.
Osler, C., 1998, Short-term speculators and the puzzling behavior of exchange rates, *Journal of International Economics*, 45: 37-57.
Osler, C., 2000, Support for Resistance: Technical Analysis and Intraday Exchange Rates, *Economic Policy Review*, 6: 53-68.
Pagano, M. 1989, Trading volume and asset liquidity, *Quarterly Journal of Economics*, 104: 255-274.
Pagano, M., and A. Roell, 1996, Transparency and liquidity: a comparison of auction and dealer markets, *Journal of Finance*, 51: 579-612.
Payne, R. 1999, Informed trade in spot foreign exchange markets: An empirical investigation, typescript, London School of Economics, January.

Payne, R., and P. Vitale, 2000, A transaction level study of the effects of central bank intervention on exchange rates, typescript, London School of Economics, July.

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.

Perold, A., and E. Schulman, 1988, The free lunch in currency hedging: Implications for investment policy and performance standards, *Financial Analysts Journal*, May-June, 45-50.

Perraudin, W., and P. Vitale, 1996, Interdealer trade and information flows in a decentralized foreign exchange market, in Jeffrey Frankel et al. eds. *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago) 73-106.

Pithyachariyakul, P., 1986, Exchange markets: A welfare comparison of market-maker and walrasian systems, *Quarterly Journal of Economics* 101, 69-84.

Portes, R., and H. Rey, 1998, The emergence of the euro as an international currency, *Economic Policy*, 26, 307-343.

Reiss, P., and I. Werner 1995. "Transaction Costs in Dealer Markets: Evidence From the London Stock Exchange," in A. Lo (ed.), *Industrial Organization and Regulation of the Securities Industry*, University of Chicago Press: Chicago.

Reiss, P., and I. Werner 1997. "Interdealer Trading: Evidence From London," Stanford Business School Research Paper No. 1430, February.

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

Reiss, P., and I. Werner, 1999, Adverse Selection in Dealers' Choice of Interdealer Trading System, Dice Working Paper No. 99-7, Ohio State University, June.

Rey, H., 1999, International trade and currency exchange, CEPR Discussion Paper No. 2226, September, forthcoming, *Review of Economic Studies*.

Rime, D., 2000, Private or public information in foreign exchange markets? An empirical analysis, typescript, University of Oslo, March (www.uio.no/~dagfinri).

Rochet, Jean-Charles, and Jean-Luc Vila, 1995, Insider trading without normality, *Review of Economic Studies* 61, 131-152.

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

Roll, R., 1988, R^2, *Journal of Finance*, 43: 541-566.

Romer, D. 1993. "Rational Asset-Price Movements Without News," *American Economic Review*, 83: 1112-1130.

Rotemberg, J., and M. Woodford, 1997, An optimization-based econometric framework for the evaluation of monetary policy, in B. Bernanke and J. Rotemberg (eds.), *NBER Macroeconomics Annual*, MIT Press: Cambridge, MA.

Rubinstein, M., 2000, Rational markets: Yes or No? The affirmative case, U.C. Berkeley Research Program in Finance Working Paper RPF-294, June (www.haas.berkeley.edu/finance/WP/rfplist.html), forthcoming, *Financial Analysts Journal*.

Saar, G., 1999, Demand uncertainty and the information content of the order flow, typescript, Cornell University, July.

Sarno, L, and M. Taylor, 2000a, The microstructure of the foreign exchange market, *Special Papers on International Economics*, Princeton University, forthcoming (available at http://www.princeton.edu/~ies).
Sarno, L., and M. Taylor, 2000b, Official intervention in foreign exchange markets, typescript, University of Oxford, April.
Seasholes, M., 2000, Smart foreign traders in emerging markets, typescript, Harvard University, January.
Seppi, D., 1996, Liquidity provision with limit orders and a strategic specialist, *Review of Financial Studies*, 10: 103-150.
Shiller, R., 1981, Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71: 421-436.
Shleifer, A., and R. Vishny, 1997, The limits to arbitrage, *Journal of Finance*, 52, 35-56.
Siegel, J., 1972, Risk, interest rates, and the forward exchange, *Quarterly Journal of Economics*, 86: 303-309.
Slezak, Steve, 1994, A theory of the dynamics of security returns around market closures, *Journal of Finance* 49, 1163-1211.
Smith, R., and B. Madigan, 1988, Exchange market management and monetary policy in the United States, in Bank for International Settlements (ed.), *Exchange Market Intervention and Monetary Policy*, Basel: BIS.
Snell, A., and I. Tonks, 1995, Determinants of price quote revisions on the London stock exchange, *Economic Journal*, 105, 55-73.
Spiegel, M., 1998, Stock price volatility in a multiple security overlapping generations model, *Review of Financial Studies*, 11, 419-447.
Spiegel, Matthew, and Avanidhar Subrahmanyam, 1995, On intraday risk premia, *Journal of Finance* 50, 319-339.
Stein, J. 1987. “Informational Externalities and Welfare-Reducing Speculation,” *Journal of Political Economy*, 95: 1123-1145.
Stein, J., 1989, Cheap talk and the Fed: A theory of imprecise policy announcements, *American Economic Review*, 79: 32-42.
Stockman, A., 1987, The equilibrium approach to exchange rates, *Federal Reserve Bank of Richmond Economic Review*, 73, no. 2, 12-31.
Stoll, H., 1978, The supply of dealer services in securities markets, *Journal of Finance*, 33: 1133-1151.
Subrahmanyam, A. 1991. “Risk Aversion, Market Liquidity, and Price Efficiency,” *Review of Financial Studies*, 4: 417-441.
Summers, Lawrence, 1986, Does the stock market rationally reflect fundamental values? *Journal of Finance* 41, 591-601.
Suvanto, Antti 1993. “Foreign Exchange Dealing: Essays on the Microstructure of the Foreign Exchange Market,” ETLA – The Research Institute of the Finnish Economy (Taloustieto Oy: Helsinki).
Taylor, M., 1995, The economics of exchange rates, *Journal of Economic Literature*, 83, 13-47.
Taylor, M., and H. Allen, 1992, The use of technical analysis in the foreign exchange market, *Journal of International Money and Finance*, 11: 304-314.
Varian, Hal, 1985, Divergence of opinion in complete markets, *Journal of Finance* 40, 309-317.
Viswanathan, S., and J. Wang, 1998, Why is inter-dealer trading so pervasive in financial markets?, working paper, Duke University, February.

Vitale, P., 1998, Two months in the life of several gilt-edged market makers on the London Stock Exchange, *Journal of International Financial Markets, Institutions & Money*, 8: 301-326.
Vitale, P., 1999, Sterilized central bank intervention in the foreign exchange market, *Journal of International Economics*, 49, 245-267.
Vitale, P., 2000, Speculative noise trading and manipulation in the foreign exchange market, working paper, London School of Economics, forthcoming in the *Journal of International Money and Finance.*
Vogler, K., 1997, Risk allocation and inter-dealer trading, *European Economic Review*, 41: 417-441.
Wei, S., 1994, Anticipations of foreign exchange volatility and bid-ask spreads, NBER Working Paper 4737.
Wei, S., and J. Kim, 1997, The big players in the foreign exchange market: Do they trade on information or noise?, NBER Working Paper 6256, November.
Werner, Ingrid, 1997, A double auction model of interdealer trading, Stanford Business School Research Paper no. 1454.
Weston, J., 1999, Competition on the Nasdaq and the impact of recent market reforms, *Journal of Finance,* forthcoming.
Wolinsky, A., 1990, Information revelation in a market with pairwise meetings, *Econometrica*, 58: 1-23.
Wood, Robert, Thomas McInish, and Keith Ord, 1985, An investigation of transaction data on NYSE stocks, *Journal of Finance* 40, 723-741.
Working, H., 1953, Futures trading and hedging, *American Economic Review*, 43: 314-343.
Yao, J. 1998a. Market making in the interbank foreign exchange market, New York University Salomon Center Working Paper #S-98-3.
Yao, J., 1998b, Spread components and dealer profits in the interbank foreign exchange market, New York University Salomon Center Working Paper #S-98-4.
Zaheer, A. and Zaheer, S. 1995. Catching the wave: Alertness, responsiveness, and market influence in global electronic networks. University of Minnesota typescript, December.
Zhou, Bin, 1996, High frequency data and volatility in foreign exchange rates, *Journal of Business and Economic Statistics* 14, 45-52.