University of Southampton

Determinants of Prices and Spreads in Global Currency and Money Markets

by
Francis Joseph Anthony McGroarty

Thesis Submitted for the Degree of
Doctor of Philosophy

School of Management
Faculty of Law, Arts and Social Sciences

November, 2003

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF LAW, ARTS AND SOCIAL SCIENCES

SCHOOL OF MANAGEMENT

Doctor of Philosophy

DETERMINANTS OF PRICES AND SPREADS IN GLOBAL CURRENCY AND MONEY MARKETS

by Francis Joseph Anthony McGroarty

This thesis tackles two big questions. The first, from the macroeconomic literature is: what drives price? The second, from the market microstructure literature, is: what determines the bid-ask spread? The classification of these questions under these headings is conventional but is not strictly accurate. While macroeconomics has nothing to say about bid-ask spreads, market microstructure is concerned with price determination. Indeed, market microstructure views the two questions as so closely related that each is a linear function of the other in certain model settings. This duality arises because every transaction price differs from the mid-quote price by the amount of the half-spread. Since the price sequence itself is made up of a series of price innovations, it follows that the average price innovation consists of a half-spread and of mid-quote revisions due to public news releases which are assumed random. However, this relationship does not tell us why bid-ask spreads arise in the first place nor does it fully describe why prices change. Two additional factors which move prices are inventory and asymmetric information. Inventory describes the (usually temporary) imbalances between supply and demand which give rise to the bid-ask spread as a management cost but which nonetheless require price innovations (concessions) to be absorbed. Asymmetric information about future price innovations not only contributes to the bid-ask spreads because of adverse selection risk, but it also drives price. The remaining factor necessary to complete the picture of what determines prices and bid-ask spreads is termed 'microstructure effects'. These include price discreteness and price clustering. It is a simple fact that prices are not continuous but instead move in discrete units and that some prices are used more frequently than others. For the first time, this thesis reveals the percentage contribution of asymmetric information, inventory and of price clustering to bid-ask spreads in the order-driven inter-dealer spot FX and short term interest rate futures markets. It also quantifies the respective contributions of news and inventory in shaping prices in these markets. Independently, it proves that asymmetric information could not be the dominant driver of prices or of the bid-ask spreads, in both markets. Finally, it shows that the level of asymmetric information in spot FX rates fell precipitously after EMU.

Table of Contents

List of Tables

List of Figures

List of Equations

Acknowledgements

I dedicate this thesis to the memory of my late grandmother Mrs. Kathleen McGroarty whose steadfast support continues to sustain me.

I would like to thank my parents, John and Maura, whose faith in me and support of me, both moral and financial, was frequently more than I deserved.

I wish to thank my wife Paulina and our two sons, Leon and Anton, for their support and for putting up with the many long hours I spent away from home working on this thesis.

I thank my supervisors Professor Steve Thomas and Dr. Owain ap Gwilym for their guidance, encouragement, feedback and for providing me with the LIFFE futures data.

I am grateful to my friend Dr. William Killeen whose decision to undertake a part-time PhD while working at an investment bank spurred me to do the same, for our early collaboration on a research paper which led to the EBS data that forms a key foundation stone in both our dissertations and for many hours of helpful conversation on a topic that few others would have found interesting.

I am indebted to Tina Kane of EBS who took the time to provide the spot FX data for this study and to the staff at the various futures exchanges who supplied the STIR futures data.

Abbreviations

AMEX....................American Stock Exchange
API....................Application Programme Interface
APT....................Automated Pit Trading
BBA....................British Bankers Association
BIS....................Bank of International Settlements
BTP....................Italian Government Bond
CIP....................Covered Interest rate Parity
CET....................Central European Time
CHF....................ISO code for Swiss Franc
CME....................Chicago Mercantile Exchange
DEM....................ISO code for German Deutschemark
DTB....................Deutsche Terminbörse
ECU....................European Currency Unit
EBF....................European Bankers Federations
EBS....................Electronic Broking Systems
EMU....................European Monetary Union
EUR....................ISO code for the Euro
FRA....................Forward Rate Agreement
FRF....................ISO code for French Franc
FRP....................Forward Rate Parity
FTSE....................Financial Times-Stock Exchange
FX....................Foreign Exchange
GBP....................ISO code for British Pound (Sterling)
GLOBEX....................Global Exchange
GMM....................Generalised Method of Moments
GMT....................Greenwich Mean Time

IFE..International Fisher Effect
IID..Independent and Identically Distributed
IPO..Initial Public Offering
ISO.....................................International Organization for Standardization
JPY...ISO code for Japanese Yen
LIFFE...............................London International Financial Futures Exchange
MATIF...................................... Marché à Terme International de France
MDH...Mixture of Distributions Hypothesis
MEFF..Mercados Financieros
NASDAQ........National Association of Securities Dealers Automated Quotation
NYSE...New York Stock Exchange
PPP...Purchasing Power Parity
SIMEX....................................Singapore International Monetary Exchange
SOFFEX..............................Swiss Options and Financial Futures Exchange
SQL..Structured Query Language
STIR...Short Term Interest Rate
UIP...Uncovered Interest rate Parity
USD...ISO code for American Dollar

Chapter 1.

Introduction

1.1 Two Big Questions

This dissertation tackles two of the biggest questions in financial economics: 1) what drives price?, and 2) what determines the bid-ask spread? It is an empirical market microstructure study, set in two of the world's largest financial markets: the global spot foreign exchange (FX) market and the collective European markets for short term interest rate (STIR) futures. Both of these markets are under represented in the literature, largely because of lack of data. The parts played by the market participants and by "microstructure effects" of the market are both investigated. The work presented offers a deeper understanding of these two issues than has been available hitherto, primarily because the high-frequency data assembled for each market is broader, deeper and longer than what has been put together for any previous analysis of either market, to the best of my knowledge.

Although it is not immediately obvious, the questions of what drives price, and what determines the bid-ask spread, are equivalent in market microstructure theory. Price moves for one of three reasons: 1) mid-quote revision in light of public news, 2) private information transmitted via trading, i.e. adverse selection and 3) mean-reverting noise arising from inventory that market makers hold for short periods. In market microstructure models, the first driver is assumed to be random and is captured by an error term. The focus is on how private or "asymmetric" information and inventory together drive price innovations. On the other hand, the bid-ask spread is also deemed to have three components: 1) order processing cost, 2) a cost associated with adverse selection risk due to private information and 3) a cost associated with managing inventory. In empirical market microstructure studies, the first component is computed as a residual when the portions of the bid-ask spread due to the other two have been deduced. Factors two and three are the statistically identical for both price innovations and bid-ask spreads.

The microstructure effects alluded to above include price discreteness and price clustering. Each of these has independently been associated with exacerbating volatility and with fostering path dependence. Both of these outcomes affect both price innovation and bid-ask spread levels. I find that these microstructure effects constitute an important factor which has been omitted in much previous empirical microstructure research into both questions. In price innovations, these features induce a bias that is not captured by normally distributed random error. In the bid-ask spread, they provide a windfall gain to liquidity providers.

1.2 FX Bias

The reader will sense that the assessment of both the spot FX and the STIR futures markets is not evenhanded. This thesis displays a distinct bias towards the former. I offer four justifications for this. First, this dissertation draws heavily on my own previous experience as a finance practitioner and more of that experience is in the foreign exchange market. Second, this study focuses on international markets and exchange rates play a key role in the study of international assets, even though they are not the specific focus. Third, the history of exchange rate economics is broader, longer and richer than that of financial futures economics. Fourth, two of the most widely acknowledged puzzles in economics relate to the formation of spot exchange rates and the present analysis can shed new light on both of them. However, I contend that the inclusion of both markets in this study makes many of the findings and conclusions, particularly about order-driven markets, more general and more robust than they would be if only the spot FX analysis were presented.

The two spot FX puzzles are: 1) the determination puzzle and 2) the excess volatility puzzle. The first puzzle relates to why exchange rate movements appear so poorly related to the fundamentals that macroeconomic theory says should drive them. The second puzzle, which is actually very closely related to the first, is why are exchange rate movements so volatile, given how stable their macroeconomic fundamental drivers appear to be.

I find that excess volatility in exchange rates is largely caused by inventory, the magnitude of which turns out to be determined by the size of trade volume. Price clustering plays a supplementary role. A number of recent research papers have established a link between long-run exchange movements and fundamental drivers. This evidence supports a widely espoused view among FX market microstructuralists that order flow is the critical proximate determinant of exchange rates. I explain what drives order flow in terms of inventory, asymmetric information and news. The same explanation may also answer the determination puzzle but more information is required about the nature of sustained deviation of FX rates from fundamental fair value before this can be confirmed.

1.3 Two Sub-themes

Alongside the central theme described above, two sub-themes emerge. The first is: how have these markets changed since EMU? The second is: do market microstructure models which were developed for quote-driven markets apply to order-driven markets?

The first sub-theme arises because the data span the period during which EMU occurs. This enables analysis of futures contracts and exchange rates which are denominated in deutschemark (DEM) as well as in euro (EUR). Also, the number of available instruments and the volume levels are very different. This helps to tease out inter-dependencies and robust features.

The second sub-theme arises because the vast bulk of market microstructure theory originates in America and focuses on the equity market. By contrast, the futures markets I study are all in Europe and the global spot FX market has more of its volume traded in Europe than in America. However, the crucial difference is that these markets are not quote-driven, they are order-driven. To be strictly accurate, some of the earlier STIR futures data does come from floor-based, pit trading which is often called quote-driven, but I find this to be similarly different from US equities and to be in many ways similar to electronic order-driven markets.

1.4 Structure of the Analysis

Chapters 2, 3 and 4 set the scene for the subsequent analysis by addressing three key background issues. The first is a review of our current understanding of how prices and bid-ask spreads are arrived at. Second, the institutions of these markets and aspects of market practice are described. Third, details of and background to the data used in this study are given.

Chapter 5, is my first empirical chapter. It is a study of the links between price discreteness / price clustering and price innovation / spread size across all spot FX and STIR futures instruments, markets and sample periods. A new theory of price clustering is proposed and evidence in support of it is found. Much of this analysis is relevant to a currently topical debate about the reason for lower than expected spot FX volumes since EMU. Most pertinent to this, I provide the first exact measure of the bid-ask spread cost of changing from DEM to EUR and show that re-denomination alone increased the minimum bid-ask spread by 74%.

In chapter 6, I explore the much documented intra-day seasonal patterns in bid-ask spread, price change, volume and order flow and the relationship between them. A theory which can explain why these patterns collectively occur may help to explain what drives prices and bid-ask spreads. While intra-day patterns involving the first three of the factors mentioned have been studied before in a variety of market settings, the study of intra-day order flow is completely new. I synthesise the existing microstructure theory into a coherent, cogent and consistent latticework of hypotheses around a central theme of asymmetric information. In doing so, I uncover a major discrepancy in one of the key models, which I address. I use a correlation matrix structure to test these hypothesised relationships. Pre-EMU and post-EMU periods are compared. Finally, the importance of asymmetric information in forming the aforementioned intra-day patterns is disclosed.

Building on the findings of chapters 5 and 6, chapter 7 quantifies the respective contributions of price clustering, asymmetric information and inventory in the formation of prices and bid-ask spreads. This analysis uses a well established model from the US equity market called the trade indicator model. As well as applying the original, I develop a modified variant of the model for order-driven markets. The modified model produces a much better result than the original. A comparison of pre-EMU and post-EMU markets using the modified model reveals how the roles of these factors have changed since EMU.

Finally, chapter 8 summarises my findings, illustrates my contributions, distils their implications and highlights important areas for future research.

Chapter 2.

A Review of Relevant Theoretical Literature

2.1 Introduction

There are three aims in this chapter. The first it to trace the origins of and to summarise our current understanding of how prices are determined in currency and money markets, with the intention of exposing gaps in this understanding. The second to provide similar background to and snapshot of our current understanding of bid-ask spread formation. The third is to explain how these issues blend together and how they can be investigated as a single issue, which is the central theme of this thesis.

2.2 Price Determination

According to Markham(1987), futures trading in commodities could date back as far as 2000B.C., when "the merchants of what is now Bahrein Island took goods on consignment for barter in India.". However, Daigler(1993) points out that financial futures did not come about until the late 1970s, when high and volatile U.S. interest rates drew the interest of the entire financial community to the idea of hedging.

The "price" of a futures contract is a little misleading. In the context of a futures contract, the term price is not an immediate transaction price. Instead, it refers to the expected delivery price when the contract is settled at the end of the period. In the specific case of a 3-month STIR futures contract, the price refers to the value on settlement day of the 3-month discounted par value. For example, if the par value is €100, the price equates to the value of €100 three months prior to the par value date. This value, in turn, equates to the amount of money that I would have to place on deposit at the STIR rate of interest in order to get €100 at the end of 3 months, i.e. on the par date. The interest rate to which STIR futures relate is the inter-bank rate. This is closely associated with the headline 'base rate' which changes only infrequently and is accompanied with great media fanfare when it does change. The difference between the base rate and the inter-bank rate is a margin determined by the aggregate perceived risk of the banks that make up the official inter-bank rate- setting panel. The inter-bank rate is normally set once a day. Futures prices deviate around the price implied by the official rate because of order flow (= buy volume – sell volume). These order flow imbalances are often interpreted as revealing the aggregate change in market expectations about future interest rates.

Unless exchange rates are fixed, interest rates and STIR futures can not be compared internationally without first taking account of exchange rates. Indeed, as I discuss below, in a floating exchange rate regime, the relationship between two countries interest rates and their exchange rate is tightly inter-woven.

2.2.1 Parity Theories

Parity theory suggests that exchange rates are governed by the "law of one price" which states that identical goods, commodities or assets should have the same price in all markets. In the combined foreign exchange and money market relationship, the law of one price is expressed in four different ways, as two riskless arbitrage conditions and two "risky" arbitrage conditions. The two riskless arbitrage conditions are the spot exchange rate no-arbitrage-pricing condition and Covered Interest rate Parity (CIP). The two "risky" arbitrage conditions are Forward Rate Parity (FRP) and Purchasing Power Parity (PPP).

The spot exchange rate no-arbitrage-pricing condition theory tells us that every currency pair must be priced consistently with all other currency pairs. So, for any set of currencies A, B and C, the exchange rate A/B must equate to the combined exchange rate A/C multiplied by C/B. Similarly, the rates A/C and C/B must conform to combinations of the other two. If one exchange rate were to become mispriced, then combinations of the other two could be used to indirectly exchange the same currency pair. If a mispricing arises, it means that the direct and indirect rates do not match. By simultaneously buying at the cheaper rate and selling at the dearer, a riskless profit may be realised, while the pressure of supply and demand drive the direct and indirect rates back into line.

CIP was first proposed by Keynes(1930). CIP arbitrage says that it is not possible to make a risk free profit by borrowing at one risk free rate and investing at another risk free rate, while at the same time, hedging away any currency risk. If country A had an interest rate of 8% and country B had an interest rate of 10%, we might expect to borrow at 8% and lend at 10% and make a 2% profit for doing nothing. However, if country A and country B have floating exchange rates, any expected profit could be more than offset by exchange rate drift. In order for no riskless profit opportunities to exist, the cost of hedging must equal the potential profit that would exist if exchange rates were fixed (2% in this example). If the cost of hedging were higher, the counter-intuitive arbitrage opportunity would be to borrow at the higher rate and lend at the lower rate. This time, as a result of going from country B to country A, instead of paying for hedging, I receive the hedging payment. So, my loss on the interest rate difference (2%) is more than made up for by my gain on the forward (>2%).

In today's world, CIP is a riskless arbitrage proposition because traders in foreign exchange markets derive the forward exchange rate by applying an interest rate differential to the prevailing spot FX rate using the CIP relationship.

FRP relates the forward rate to the future spot rate. It says that the forward rate should not systematically over-estimate or under-estimate the future spot rate change. If it should systematically over-estimate the future spot rate, then funds should be borrowed at the lower interest rate and invested at the higher interest rate (without hedging the currency). This would exert upward pressure on the spot exchange rate and, respectively, upward and downward pressure on the low and high interest rates to restore equilibrium. On the other hand, if the forward rate should systematically under-estimate the future spot, the same apparently counterintuitive result arises as with CIP. Namely, that funds should be borrowed at the high interest rate and invested at the low interest rate. This strategy should be pursued because the combined low interest rate and currency gain would exceed the high interest rate borrowing cost.

PPP is arguable the best known parity theory. It is most closely linked with the name of Cassel(1916 and 1918) whose work popularised it. However, its origins go back to the mercantilist literature of the seventeenth century. Lahtinen(2003) states that the Spanish Salamaca school were probably the first to articulate the theory. An awareness of the relationship is clearly evident in the writing of Hume(1741) (pages 318-319). PPP says that goods, net of transport costs and taxes, should cost the same everywhere. If goods are cheaper or dearer in one place than in another, then market forces should act to erode the anomaly. This will happen because demand will increase for the cheaper goods, exerting upward pressure on the cheap price and the cheap currency. This will cause demand for the dearer good to fall, exerting downward pressure on the dear price and the dear currency.

Adding much confusion to the study of foreign exchange rate determination are two alternative theories which some writers describe as parity theories in their own rights. In fact, both are combinations of FRP and PPP. These other parity theories are: 1) the International Fisher Effect (IFE) and 2) Uncovered Interest rate Parity theory (UIP). IFE asserts that, in equilibrium, real interest rates are the same in all countries and so differences in nominal interest rates solely reflect the inflation rate differential. UIP states that the nominal interest rate equals the real interest rate plus compensation for expected inflation. Neither of these contributes anything more to our understanding of exchange rate dynamics beyond that already provided by the four parity theories discussed above. In fairness, the intellectual contribution in the FRP-IFE-UIP theory is Fisher's(1930) IFE. However, in the present work, I use the FRP theory because I think it is clearer.

The Mundell-Fleming model introduces foreign exchange market equilibrium into the Keynesian equilibrium IS-LM model. This model can be traced back to two separate papers: Mundell(1962) and Fleming(1962). The original Hicks-Hansen IS-LM model showed how the real and money economies relate to each other. Under floating exchange rates, this does not add much to our understanding of the relationship between interest rates and exchange rates. Rather, it implements interest rate parity theory within a broader macroeconomic model. The model assumes that, without capital restrictions, all domestic real interest rates equate to the world real interest rate. Therefore, any apparent difference in nominal rates must be associated with domestic inflation and must disappear in exchange rate drift. In other words, it imposes the above parity theories on the IS-LM framework.

The parity theories describe a world which is nice, neat and self-regulating, in which economic value is transferred from place to place and time-period to time-period smoothly and seamlessly. In the real world, things work less perfectly. The two riskless arbitrage conditions do indeed hold. Note that these two conditions are independent of each other, i.e. an anomaly arising in one would not necessarily cause the other to be mispriced. That FRP has not held is well documented in the empirical foreign exchange literature. Empirically, it has been observed that the forward rate consistently overestimates the magnitude of the subsequent spot rate change. In many instances it even persistently gets the direction wrong. Engel(1996) wrote a comprehensive review of the literature on the forward rate bias anomaly. PPP also fails empirically. The best defence that can be mounted for PPP by its proponents is that it is a very long run force and will eventually erode pricing anomalies. Engel(2000) produces a convincing dismissal of even that.

The theories of PPP and FRP make heavy demands on each other because both depend on the spot exchange rate as a transmission mechanism. If changes in underlying demands should drive either one from equilibrium, the pressure would fall on the spot exchange rate to restore equilibrium. If a perturbation appears in one measure while the other remains in equilibrium, one rule would require the spot exchange rate to change while the other requires it to stay the same. Another way to look at this is that if FRP were fixed, then for PPP to hold, domestic inflation would have to exactly absorb all exchange rate drift. On the other hand, if PPP were fixed, any exogenous rise in inflation would need to be simultaneously offset by an interest rate which aligned the forward rate with the spot exchange rate trajectory implied by the inflation differential. The big problem here is that inflation is a lagged variable, in that we only find out about after it has occurred, while interest rates are a leading measure since it has to be declared in advance. This means that the only way for the FRP-PPP relationship to hold exactly, through time, is if nominal interest rates perfectly predict future inflation, thereby holding the real interest rate constant. This is stronger that the FRP alone demands because errors in the expectation of inflation would still cause a breach of PPP.

Hartmann(1998) says that spot currency flows split into three components, related to investment activity, to trading and to government intervention, in decreasing order of importance. In the context of the FRP-PPP problem, this means that spot FX flow activity linked to FRP should dominate that of PPP. However, investment based flows are determined by the whole panoply of investment opportunities within a country, not just the short term interest rate assumed in the

FRP condition. For FRP to hold across all investment horizons and credit ratings, the ratio of returns would have to be identical for every time horizon and instrument type, across all countries. This is not what is observed empirically. Indeed, there is no shortage of evidence which shows that interest rate cuts often cause the equity market to rally because they are perceived as positive economic stimuli. Direct observation of FX dealers taught me that this higher short-term expected return in the equity market can more than outweigh lower expected return in the short-term interest rate. This is because it draws an inflow of capital into the equity market, thereby producing the opposite effect on the exchange rate, via order flow, than that predicted by FRP.

The FRP-PPP failure essentially permits the existence of the three most widely acknowledged major puzzles in exchange rate economics. If FRP-PPP held, it would be impossible for any of these puzzles to arise. As Lyons(2001) puts it, the three big FX puzzles are:

"*1. The determination puzzle: exchange rate movements are virtually unrelated to our best measures of fundamentals*

2. The excess volatility puzzle: exchange rate movements are excessively volatile relative to our best measures of fundamentals

3 The forward rate bias puzzle: excess returns in foreign exchange are predictable and inexplicable".

The present thesis contributes to the debate surrounding the first two puzzles.

2.2.2 Introducing Expectations

With the breakdown of Bretton Woods in the 1970s, came floating exchange rates. At the time, economists were greatly surprised by the volatility and directional behaviour of exchange rates. It was not consistent with what the dominant Mundell-Fleming theory had predicted. Dornbusch(1976) suggested an important modification to the model – to include expectations. He suggested that the initial response to a fiscal or monetary stimulus might be for the exchange rate to overshoot and then to reverse. Dornbusch(1976) also suggested that goods prices are sticky in the domestic market and that exchange rate overreaction will be more closely associated with asset market activity.

Meese and Rogoff(1983) showed that even with the Dornbusch(1976) insight, exchange rate economics had still failed to produce a model which fit the empirical data.

Frankel and Froot(1990) furthered the Dornbusch(1976) dogma of expectations in exchange rates. However, their work also signalled an important shift in academic thinking because they consider the behaviour of agents within the FX market instead of only the behaviour of general economic agents whose focus lay on goods or assets outside the market mechanism itself. They proposed a model of "chartist" speculators who, collectively, do not use any fundamental means to evaluate investment opportunities, but instead chase trends. For example, if a central bank were to raise short-term interest rates. According to the Dornbusch(1976) model, the domestic currency would first appreciate, but then would almost immediately go into a gradual reverse. Instead, Frankel and Froot(1990) suggest that after that initial jump, speculators exhibit herding behaviour and create a bubble. They observe that sterling has risen. In the belief that the trend will persist, they continue to invest in sterling. Their speculative behaviour pushes the price up which causes their prediction to fulfil itself. Thus, market participants are led to conclude that an interest rate increase is associated with future appreciation of the domestic currency.

2.2.3 Market Microstructure and Exchange Rate Determination

Market microstructure is a broad church and numerous facets of it contribute to an explanation for exchange rate evolution. Things have moved on considerably since Flood (1991) stated that, "little attention has been paid to the particular microstructure of the foreign exchange market". However, the microstructure of the foreign exchange market is still less researched than other markets. On the empirical side, as ap Gwilym and Sutcliffe(1999) point out, this is largely due to the lack of data.

The microstructure perspective studies the market mechanism which actually sets prices, rather than on how value imbalances between exogenously determined equilibria pressure individual exchange rates and asset prices into line. One way to think about this is as the "invisible hand" of the market being put under a microscope. Macroeconomists expect this to show how shifts in the macro level equilibria are smoothly translated at the level of the individual asset or exchange rate. However, the empirical side of the field finds that aggregate trading activity, in the form of order flow, is frequently and awkwardly evident at the macro level, in a way that does not snugly fit with the prevailing fundamentals.

The spotlight is drawn to the phenomenon of vehicle currencies. Vehicle currencies are a feature of the inter-dealer market. When a bank exchanges a currency with a customer, it ends up holding a position in one of the currencies that it may not want. The dealer will then use the inter-dealer market to exchange this currency for a currency that the bank does want to hold. However, if the volume of trade in this exchange rate is low, the dealer will first exchange the currency into a liquid currency and then from the liquid currency into the desired currency. The liquid currency in this scenario is called a vehicle currency. The use of vehicle currencies distorts the relationship that price determination has with volume and order flow because it permits the possibility that instances of either could be merely passing through and may not be in any way related to fundamentals, in the wider Lyons sense of the word. In addition, as Hartmann(1998) illustrates, vehicle currency trading should lower the bid-ask spread associated with the vehicle exchange rate. This is largely because it spreads out the fixed overhead costs associated with order processing.

Osler (2003a and 2003b) shows how features of the market practice and structure give rise to certain trading behaviours which, in turn, could give rise to considerable volatility and significant deviations of exchange rates from the long-run equilibrium levels. The main feature that her work highlights is the tendency of banks to have hidden limit orders at certain price levels. These are often placed by the exotic derivatives desks in these banks. The effect of these limit orders is to release an amount of volume that pushes in the same direction as prevailing trajectory of the exchange rate. These limit orders take the following forms: "If the rate should rise to X, then buy Y amount of the currency" or "If the rate should drop to X, then sell Y amount of the currency". This is because they are associated with "stop loss" and "take profit" activities. These features can result in "price cascades" when volatility is high.

Evans and Lyons(2002) was the first widely acknowledged paper to present an empirical work which shows that "order flow" helps considerably in predicting behaviour in spot FX rates. Order flow is defined as "buying volume – selling volume". (Actually, they use "number of buy orders – number of sell orders" as a proxy for order flow because they do not have actual volumes.) They find that the cumulative imbalances between buying and selling pressure are not mean-reverting. Instead cumulative order flow follows a random walk. They further show that this cumulative order flow is an important factor in determining the exchange rate, and that its significance is far greater that of the short-term interest rate.

Killeen, Lyons and Moore(2003) use the model from Evans and Lyons(2002) to address the excess volatility puzzle. They use daily aggregate inter-dealer data from EBS to explore the volatility in DEM/FRF before and after EMU. (Note that their data was not high frequency like the data I received from EBS and use here.) This currency pair and time period is particularly interesting because it affords the opportunity to compare volatility between fixed and floating exchange rate regimes and to identify its cause. They found that order flow and price are strongly cointegrated under the floating regime, whereas they are unrelated under the fixed regime. Furthermore, they found that order flow was exogenous and that it drove price rather than the other way around. As Lyons(2001) put it, "when a gap opens in the long-run relationship between cumulative order flow and the exchange rate, it is the exchange rate that adjusts to reduce the gap, not cumulative order flow.".

Lyons(2001) argues that exchange rate levels should be linked to discounted future payoffs and therefore that changes in exchange rates should be related to either changes in future payoffs or changes in the discount rate. He further argues that the discount rate comprises two distinct parts. One is related to transient inventory and the other to "portfolio balance" effects. The latter have a permanent effect on price but are not associated with changes in expected futures payoffs. Lyons(2001) suggests that asymmetric information about future accumulated inventory could drive price.

Flood and Rose(1995) suggest that asymmetric information is not credible as a significant driver of exchange rates when they state: "Intuitively, if exchange rate stability arises across regimes without corresponding variation in macroeconomic volatility, then macroeconomic variables will be unable to explain much exchange rate volatility. Thus, existing models, such as monetary models, do not pass our test; indeed, this is also true of any potential model that depends on standard macroeconomic variables. We are driven to the conclusion that the most critical determinants of exchange rate volatility are not macroeconomic."

Froot and Ramadorai(2001) use cross border institutional investors flows data from State Street Bank (my former employer), to explore movements in exchange rates. They find that these portfolio flows are strongly and positively related to movements in exchange rates. They also find that the flows are strongly autocorrelated. Furthermore, they find that these flows forecast a rise in the equity market into which the funds flow. This is evidence that asset returns, other than the short term interest rates associated with FRP, are important for determining exchange rates. Froot and Ramadorai(2001) go on to compare the performance of the equity market with funds which are invested in that market but are quoted on NYSE or AMEX (i.e. closed-ended country funds). They find that these two exhibit similar responses to flows of funds into the country but that the discount rate on the fund is uncorrelated with the flows. They interpret this as evidence in favour of the information driver over the inventory driver of exchange rate changes.

Froot and Ramadorai(2002) find that their investor flows data is not closely related to future fundamentals. By separating currency returns into temporary and permanent components, the authors show that flows appear important in understanding temporary and transitory deviations in currency returns. In other words, portfolio flows contribute to volatility. However, they find that portfolio flows are either not related to, or are negatively related to, permanent exchange rate shifts. Instead, over long horizons, they offer some evidence that future interest rate differentials are associated with the permanent part of exchange rate changes. These findings support the conclusions of Mark(1995), as well as those of Flood and Taylor(1996), about fundamentals driving price over the long run.

2.2.4 A Summary of Our Current Understanding of FX Rate Determination

Lyons(2001) describes the three ongoing puzzles in exchange rate economics as 1) determination, 2) excess volatility and 3) forward rate bias. I attribute all three to the combined failures of the FRP and PPP propositions. A critical problem with FRP is that it focuses on a very narrow definition of return, i.e. short term interest rate. In practice, investors are motivated by a wide range of available investments opportunities and, empirically, these do not rise monotonically as the maturity horizon rises.

I show that the market microstructure approach brings many potentially explanatory factors into focus. I show that a close link has been established between order flow and exchange rate movements. In addition, a firm link has been established between international institutional investor flows and exchange rate movements. The evidence so far affirms the role of flows as a proximate determinant of changes in FX rates, which helps to explain excess volatility. On the other hand, evidence is emerging which suggests that long-run exchange rate determination is more closely linked to fundamentals. However, this finding should not obscure the fact that, at any instant, an exchange rate will probably be some distance away from its fundamental fair value level. It remains to be resolved whether order flow is generated by uninformed inventory or are prompted by information about future fundamentals.

2.3 Bid-Ask Spread Theory

In contrast to the long history of price determination theory described above, the theory relating to how bid-ask spreads are formed only evolved over the past three decades. Most of the theory that has emerged is closely modelled on the US equity market and characterises the interaction between three types of economic agent: the market maker, the informed trader and the uninformed trader. This is the subject of the next section. In section 2.3.2, I present some research that relates to order-driven market models, which is the market structure that most closely resembles the markets that my data come from.

2.3.1 Quote-Driven Models

Quote-driven bid-ask spread models form two distinct groups. The first is made up of inventory models. The second group consists of asymmetric information models. Both classes of models are explored in detail in O'Hara(1995).

Demsetz(1968) was probably the first to address the question of what determines the bid-ask spread. Although he lacked the language of bid-ask spread analysis familiar to market microstructure researchers today, he squarely addressed how market microstructure affected transactions costs for NYSE stocks. He used the concepts of "immediacy", "externality" and "monopoly" to explore issues of efficiency and inventory management risk. However, it was Garman's(1976) work on temporary order flow imbalances that formally launched the field of market microstructure. Amihud and Mendelson(1980) re-formulated Garman's model to permit inventory to influence the setting of the bid and ask prices. The Garman(1976) and Amihud and Mendelson(1980) models assume that the market maker is a risk neutral monopolist. Stoll(1978) re-casts the market maker as a risk averse supplier of immediacy who must be rewarded for the service that s/he provides. It was Stoll(1978) who first defined the three components of the bid-ask spread as those linked to 1) order processing, 2) inventory holding and 3) adverse selection risk. Stoll(1978) also was first to show that the mid-quote should respond to the inventory level but that the bid-ask spread itself should not. Ho and Stoll(1981) extend Stoll's(1978) model to a multi-period framework and provide insight on how time horizon, transaction size, price change variance and relative risk aversion affect the relationship between inventory and the bid-ask spread. In general, inventory models only model the interaction of uninformed traders with market makers. Interaction involving informed traders is the focus of asymmetric information models.

A short paper by Jack Treynor, using the pseudonym Walter Bagehot(1971), is credited with first distinguishing between profit made from holding an asset and profit made or lost from trading it. Market makers lose money when they trade with informed traders. This insight sowed the seed for the concept of informed trading and how this affects bid-ask spreads. In the 1980s, a number of important papers were published which greatly expanded the theory on asymmetric information and bid-ask spreads. The first of these was Copeland and Galai(1983) who develop a model in which order flow is not exogenous but exploits information about future price. This shows that a bid-ask spread can arise for reasons other than inventory or fixed exogenous transaction costs. Glosten and Milgrom(1985) permit the market maker to deduce the information from informed trades and to adjust bid and ask quotes accordingly. This is achieved is through setting "regret-free" bid and ask prices which are conditioned upon the next buyer or seller being informed. The sequence of prices in this model forms a

Martingale. Glosten and Milgrom(1985) assume that trading occurs as a sequence of constant units and that price signals to informed traders happened in every period. In contrast, Easley and O'Hara(1987) show that varying the trade size impacts both bid-ask spread and the transaction price, and introducing price signal uncertainty affects the evolution of the mid-quote.

2.3.2 Order Driven Models

In an order-driven setting Cohen, Maier, Schwartz and Whitcomb(1981) conclude that the bid-ask spread has a different meaning to that in quote-driven markets. They assert that the size of the bid-ask spread is a measure of non-execution risk for self-interested long-term holders, rather than compensation for a risky intermediate holding service provided by obliging short-term holders. Cohen et al(1981) find that the size of the bid-ask spread is the crucial factor that determines whether a trader submits a market or a limit order, reflecting the fact that the advantage of price betterment lessens as the bid-ask spread falls. It follows from this that the bid-ask spread must be positive in equilibrium. A recent related study of bid-ask spreads in order-driven markets by Foucault, Kadan and Kandel(2001) supports these conclusions.

2.4 Conclusion

Heretofore, very different literatures have addressed the separate questions of price determination and bid-ask spread determination. The FX pricing question is firmly in the realm of macroeconomics. Although much recent work does take a market microstructure approach, the puzzles and questions remain macroeconomic ones. On the other hand, macroeconomics has nothing to say about the bid-ask spread. Interest in bid-ask spread determination is the sole preserve of market microstructuralists. However, market microstructuralists do also have an interest in price formation. Indeed, at the level of the individual trade, almost all microstructure models define price change as a function of the bid-ask spread or vice-versa. Important examples of this include Roll(1984) and Kyle(1985), as well as the bid-ask spread models described above. The close relationship between the bid-ask spread and price change at the elemental level means that the factors that determine one must also determine the other. This equivalence is clearest in the trade indicator model which is presented in chapter 7. This model has been used to derive the components of the bid-ask spread by Huang and Stoll(1997) and also to show the drivers of price innovations by Madhavan, Richardson and Roomans(1997). I exploit this duality in chapter 7 to jointly answer the two central questions of this thesis. Before that, the parts of this answer are explored separately in chapters 5 and 6. There are stark differences between order-driven and quote-driven bid-ask spreads. The trade indicator model introduced in chapter 7 draws heavily on its quote driven-inventory and asymmetric information heritage. I modify this model specifically to fit it to the order-driven market environment.

It is worth noting that market microstructure actually provides insight into two distinct sets of factors which shape prices and spreads. Only the first of these relates to the inter-action of market participants which is addressed above. The second set of factors is termed microstructure "effects". These relate to the market institutions and the denominational units of price. They do not involve behaviour or inter-action and so lack any formal theory. However, that does not mean they do not have an influence. My new derivation of the trade indicator model permits the inclusion of two of these microstructure effects: price clustering and price discreteness. The influence of these two features on price and the bid-ask spreads been has recognised in previous research papers, as I discuss in chapter 5. Although they have yet to be embraced by the core literature that purports to explain either price determination or bid-ask spread formation.

Chapter 3.

Market Practices and Structures

3.1 Introduction

Market microstructure thrusts the practices and structures of the market mechanism into the centre of the stage in the quest to unmask the key drivers of price and the bid-ask spread. Both of instrument types that I study here are traded within vast markets which have disjointed structures. In both cases, the data I use comes from the largest, most liquid and most transparent part of the broad market, the inter-dealer market. This chapter presents details of the practices in and the structure of spot FX and STIR futures inter-dealer market. It shows why they are so important. It also addresses problems that arise in trying to explain behaviour in these markets with theoretical models that come from a very different environment.

3.2 The View from a Dealing Desk

I have firsthand experience of the currency and money markets, as I spent several years working as a manager in the foreign exchange division of State Street Bank, who are one of the largest banks in both of these markets. This experience showed me that large banks dominate the foreign exchange business. They provide market-making services to fund managers, self-managed institutional investors, hedge funds corporate treasuries and smaller banks. Dealers also make money for their employer banks by speculating in the market. They do this by taking a long or short position in an instrument based on their expectations of the near future.

Lyons(2001) gives an excellent account of sub-markets that make up the foreign exchange market and how these inter-relate. Briefly, FX divisions in banks deal with customers on one side and with other dealers, via the inter-dealer market, on the other. The ostensible reason for dealing with other dealers is to manage inventory, i.e. to offload a currency position that a bank has received from a customer but does not want to hold. The customer bid-ask spreads are much larger than the inter-dealer bid-ask spread. This means that customer deals are generally profitable and inventory risk is low since it is cheap to dispose of.

Banks generally maintain separate spot and forward exchange rate desks. On a spot desk, individual traders specialise in certain liquid currency pairs. The 5 most liquid currency pairs in 1999 were believed to be, in order of decreasing liquidity: EUR/USD, USD/JPY, USD/CHF, GBP/USD and EUR/GBP (Killeen, McGroarty and Moore(2000)). From this and from Hartmann(1998), it is clear that the US dollar, the euro and the Japanese yen are by far the three most important currencies in the world. Of the remainder, Sterling and the Swiss franc are more important than the others.

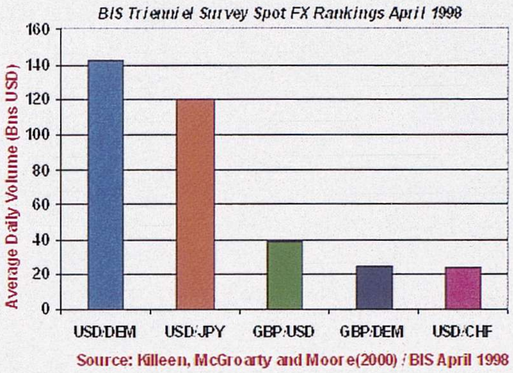

Figure 3.1: Volume per FX rate in 1998

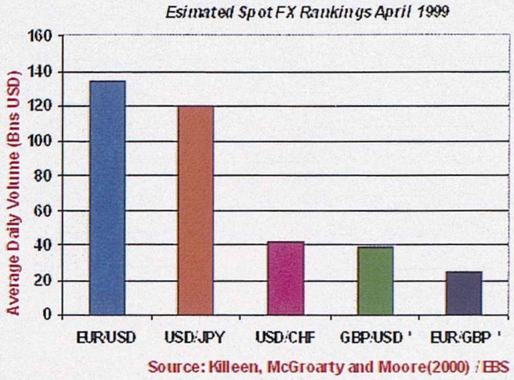

Figure 3.2: Volume per FX rate in 1999(est.)

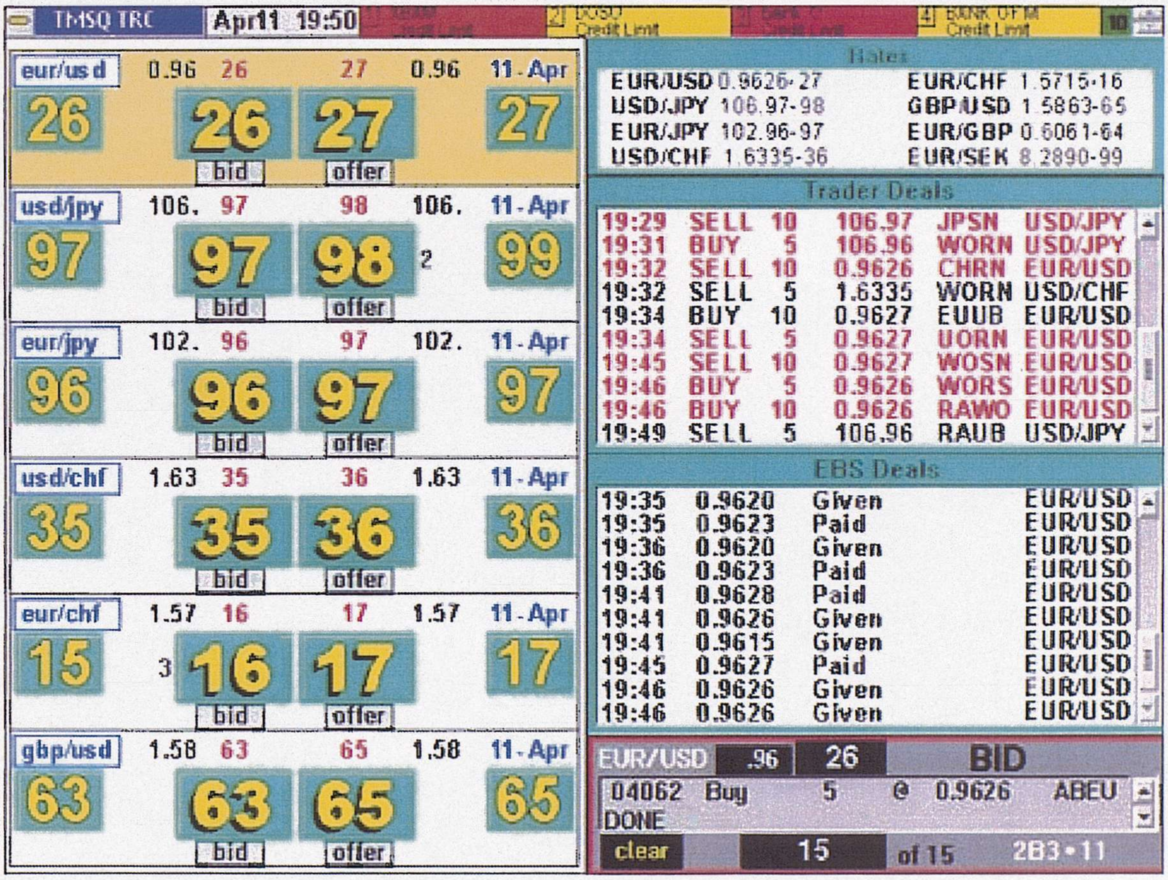

Figure 3.3: EBS Spot FX Dealing Screen

Most inter-dealer spot FX trading is now done over two electronic inter-dealer trading systems. Certain currency pairs trade on one system and other pairs trade on the other system. The first trading system is called Reuters D2000-2. This platform has liquidity in the commonwealth currencies and Scandinavia. The second system is EBS (Electronic Broking Systems), who provide the FX data for this study. EBS hosts larger volumes than Reuters and has liquidity in euro, US dollar, Japanese yen and Swiss franc currency pairs.

The EBS Spot Dealing System (figure 3.3) is a screen-based, order driven, pre-trade anonymous, electronic trading system. After a dealer "takes" an existing limit order, identities are revealed and the counterparties settle directly with each other. The EBS screen displays the best bid and ask prices alongside "size" (in US$). Traders typically deal in unit values of either $5 million or $10 million. Market depth is not visible on the dealing screens, as amounts available for trading at prices "outside the touch" (lower than the highest bid or higher than the lowest offer) are not displayed. A dealer can monitor his own recent trades in a sub-window visible only on his own terminal. All recent trades executed on the EBS system are displayed in another sub-window but the counterparties are not shown. EBS pre-screens the credit of dealers so that they can only see prices for available trades with counterparties with whom they have credit limit agreements and, even then, only when the trade size is within their remaining available credit. Dealers can submit either limit orders or market orders to the system. To place a limit order, the dealer "makes" or posts a (bid or offer) price and quantity for a currency pair, at which he is willing to deal. To execute a market order, a dealer "takes" or accepts a price-quantity combination previously posted by another dealer. The dealer who executes a market order is deemed the initiator of the trade. A dealer who submits a limit order that is equal to or higher(lower) than the best available offer(bid) triggers immediate execution of an existing limit order and is deemed the initiator of the trade.

Notably, Lyons(2001) attributes only one-third of major market FX spot volumes to trades with customers. Therefore, two-thirds of spot FX volumes take place between dealers. He also points out that over half the total spot volume of the most liquid currency pairs is concentrated in the largest ten FX banks. Lyons explains this puzzling fact by "hot potato" trading. The latter arises in direct inter-dealer FX trading, where inventory imbalances are passed from one dealer to another, who in turn passes it on to another dealer, and so on. However, hot potato trading has steadily declined in recent years as brokered inter-dealer trading has steadily eroded the market share of direct inter-dealer trading. Hot potato trading is not a feature of brokered inter-dealer trading because limit order traders are not obliged to provide two-way prices. Therefore, a party who submits a limit order must actually want the position that results from the trade, removing the motivation for hot potato trading. According to Lyons(2001), "by the end of 2000, only about 10 percent of inter-dealer trades were direct".

Forward exchange rates are provided by banks to their customers. Forward rates are priced by applying relative interest rate differentials to the prevailing spot rate. This is how the CIP relationship is implemented to derive the forward rate from the current spot rate. Banks do not hedge these forward instruments by dealing in forward rates in the inter-dealer market. Instead, they hedge their spot component and their interest rate components separately. For this reason, every forward trade entails an immediate spot trade but the reverse is not true. This, in turn explains why the forward desk is always smaller than the spot desk. When it comes to hedging interest rate exposure, much of this is done bilaterally between banks using forward rate agreements(FRAs). However, to the best of my knowledge, there are no data available which show what proportion of total banks' interest rate risk is hedged bilaterally and what proportion is hedged in an open market. Forward traders also monitor a variety of money market instruments, including STIR futures, on systems such as LIFFE CONNECT. They use these both to hedge and to exploit arbitrage opportunities, if any should arise.

The LIFFE CONNECT system is a fully integrated electronic trading platform. However, instead of creating a standard interface, like EBS, LIFFE chose to build a core architecture for electronic trading. They collaborate with sixteen independent software vendors to build front ends with functionality customised to the needs of different groups of traders. Each of these front ends communicates with the LIFFE "Trading Host" via the LIFFE API (Application Programme Interface), which is a common software protocol. Like EBS, at the heart of LIFFE CONNECT lies a central order book. Dealers can submit bid or offer limit orders. Alternatively they can submit market orders which will take out existing limit orders. The LIFFE CONNECT facility also permits "Market on Open Orders". These are submitted at prior to market open for execution at market open. Like EBS, LIFFE CONNECT is pre-trade anonymous. In fact, traders don't get to see who their counterparty was until three days after the trade. This happens because LIFFE CONNECT counterparties don't settle directly with each other. Rather the market is settled centrally through the London Clearing House. A notable feature of the LIFFE CONNECT system is that it makes available real-time market depth information. This enables users to track and view all price and aggregate volumes for buy and sell orders of a specified contract month or series. Two pages from the LIFFE CONNECT service offered by Bloomberg are shown below.

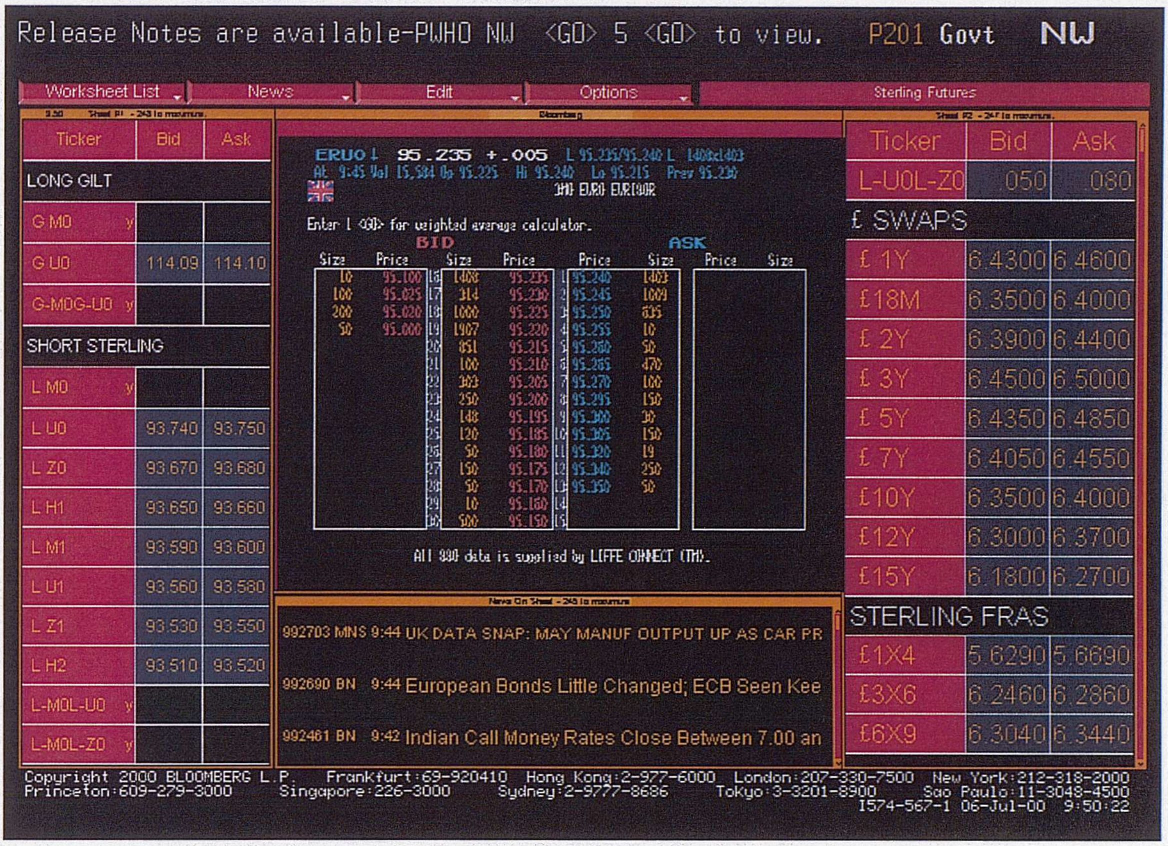

Figure 3.4(a): LIFFE CONNECT on Bloomberg

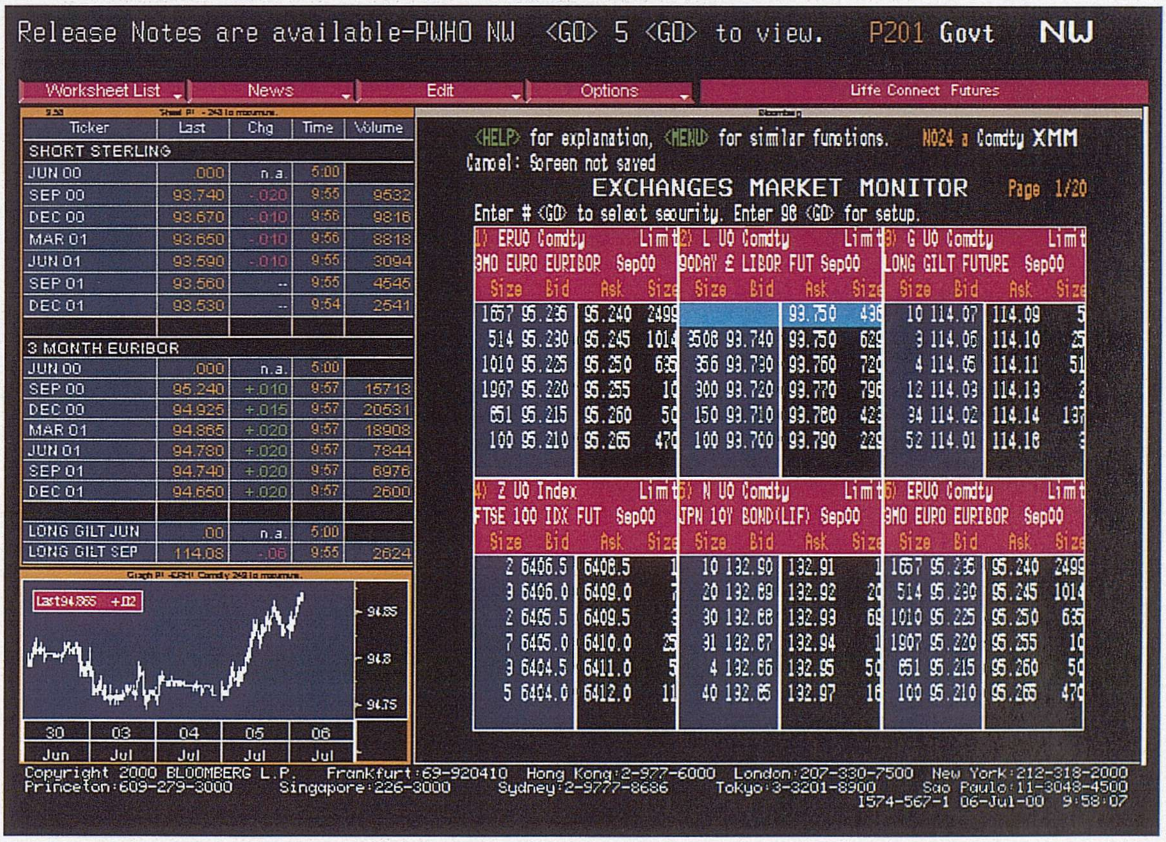

Figure 3.4(b): LIFFE CONNECT on Bloomberg

3.3 Inter-Dealer Price-Setting Hubs

The inter-dealer spot market plays an important role in the determination of price because it is transparent, while the other parts of the FX market are not. So, the inter-dealer market effectively sets the mid-quote level for the bilateral, opaque, customer-focused side of the FX business. Over the past decade, the market share of direct inter-dealer spot FX trading has steadily declined and brokered trading has risen to dominate this inter-dealer market. In contrast to BIS(1998) where only half of all inter-dealer trades were electronic, BIS(2001) showed that this rose to around 90% just three years later. The fact that brokered inter-dealer trades take place on market-wide electronic trading systems means that prices of executed trades are quickly disseminated across the market. Approximately 60% of total global spot FX volume takes place in the inter-dealer market. By contrast, customer trades are bilateral as each bank only sees its own customers' trades and not those of any of its rivals. This is why banks need a broader measure of how prices are moving and what their current levels are. The inter-dealer market provides these. In this sense, the spot FX inter-dealer market can be seen as a price-setting hub from which the rest of the spot FX market takes its cue.

In a similar way, the STIR futures market can be seen as a real-time interest-rate-setting hub for banks who wish to manage risk associated with future cash flows. Apart from liabilities arising from foreign exchange forward transactions, banks accumulate a wide range of other cash and short-term liabilities through various other business activities. Banks even offer Forward Rate Agreements (FRAs) directly to their customers. The money market instruments that banks use to hedge these exposures are characterised by a high degree of safety of principal and tend to be issued in very large denominations, e.g. €1,000,000. Cook and Laroche(1993) provide a detailed description of the different kinds of money market futures and other instruments used by banks. The main reason these money markets arose in the first place, is to bridge the time gap between receipts and expenditures. The largest, most liquid and most transparent market in which a bank can hedge these risks is the STIR futures market.

3.4 Order-Driven V Quote-Driven Markets

From the discussion of electronic trading above, it may be deduced that both the spot FX inter-dealer market and the STIR futures are order-driven markets. Actually, this is true only for the latter half of the data period studied. In the first half of the period, STIR futures trading was "floor based" (i..e. quote-driven), switching to electronic trading in 1999. Although, as I shall argue in chapter 7, many features of the quote-driven STIR futures market are more closely aligned with those of the order-driven market models than the particular type of US-equity based quote-driven on which most of the extant theory is currently based. In any case, the fact that most of the data comes from an order-driven system presents a problem for an empirical market microstructure analysis such as this. The problem is that most existing market microstructure theory assumes a radically different market environment.

"Quotes" in quote-driven models are indicative rather than binding. They are often said to be good "only as long as the dealer's breath is warm". By contrast, "quotes" in order-driven markets are referred to as firm or binding. These are limit order prices on declared quantities which a (market order) trader can choose to accept or not. Most existing US-equity-based theoretical models contain three key types of agent: market makers, uninformed traders and informed traders. The inter-action of these three is what determines price, returns, spreads and intra-day trading patterns. The market maker will raise his/her bid-ask spread when s/he anticipates increased risk of adverse selection, or if s/he expects to have to hold inventory for longer than usual. Traders are either informed or uninformed and choose whether or not to trade in the face of available prices and spreads. The inter-dealer market's lack of market-makers creates a problem. Without exogenous liquidity providing market makers, trades are made up of limit order makers and market order takers. Anyone wishing to effect a trade has the choice of executing an immediate market order or placing a limit order in the system. An increase in the bid-ask spread will make placing a limit order more attractive than executing a market order. However, a very important aspect of an order-driven regime is that each of the parties to a trade wants to arrive at the portfolio positions that the trade delivers them to. In other words, an order-driven system never generates unwanted inventory and there is nobody who can be adversely selected.

Note that the order-driven issue raised above do not present a problem for studies which focus on the individual dealer perspective within spot FX or STIR futures markets. For example, Lyons(1995) successfully implemented the Madhavan and Smidt(1991) model which was originally developed for the NYSE. The reason that the order-driven issue does not create a problem here is that individual dealers do have inventory they need to control and they do fear adverse selection by market participants who know more than they do – but not in their brokered inter-dealer trading! Since the electronic inter-dealer market does not require traders to quote two-way prices, there is no obvious reason why a trader would place a limit order that s/he did not genuinely want fulfilled. In a nutshell, the structure of the spot FX and STIR futures markets is best thought of as having an order-driven inter-dealer core and quote-driven market maker satellites.

It should not be construed from the preceding arguments that the market as a whole is immune to unwanted inventory or to the presence of informed traders. The aggregate market could act as a giant, monopolistic, market maker. As such, the insights of the quote-driven models may still prove useful. Although, some adjustments to the models may be appropriate. While theoretical work on order-driven markets is advancing steadily, the workhorse models for order-driven markets have yet to emerge. By workhorse, I mean alternatives to the models of Stoll(1978), Kyle(1985) and Glosten and Milgrom(1985), etc. which identify the protagonists, isolate and capture their motivations, as well as the general dynamics that rule the market, and which underlie so much subsequent theory. However, in the absence of such alternatives tailored to the order-driven markets, I believe it would be too drastic a move to set all the insights of the established models aside and start again with from the beginning. So, I continue to utilise the existing theory and models as the basis for the empirical work contained in the ensuing chapters, although substantial modifications are made where the assumptions underlying models do not accord with acknowledged features of the markets being studied.

3.5 Conclusion

This chapter describes the internal workings of the spot FX and STIR futures markets and shows where their inter-dealer markets fit in. Section 3.3 emphases the critical role played by the inter-dealer markets in setting the price for the general market. Finally, in section 3.4, I acknowledge the mismatch between the features of these markets and those usually assumed by the established market microstructure models but I defend the continued use of these established models as more sensible than the alternative of starting from scratch.

Chapter 4.

Two Unique Datasets

4.1 Introduction

I began this research in early 2000, and the catalyst for it was EMU. These two facts explains the choice of instruments studied, which are the two most directly affected by EMU. They also explain the period spanned by the data, 1997 to 2000. Although my FX data only spans two (separate) months during this period, they are particularly important because very little high frequency FX data have previously been made available. For example, two important papers, Goodhart, Ito and Payne(1996) use only one day, while Lyons(1995) uses only one week. While Evans'(1997) dataset covers a four month period, it comes from the smaller direct inter-dealer market and contains only transaction data, not quotes. Previously available STIR data have similar drawbacks. Much of it comes from US exchanges. These only provide sparse "time and sales" data, consisting of only trade price data with no volumes or no quotes. Furthermore, they do not even report all trade prices. Instead, they only report prices which differ from the immediately preceding trade price. Even the poorest European STIR series used in the present study contains all traded price and volume data.

This chapter describes the data in more detail, as well as the sample period from which they come. It also provides some background on the exchanges that supplied the data. Section 4.5 contains details of the specific STIR contract specifications. Finally, I give a brief overview of the work I did to structure and prepare the data for analysis.

4.2 The Data

The spot FX data is from the EBS Partnership, who handled roughly one-third of the total global spot foreign exchange transactions at the time from which the sample is taken. These data consist of two distinct months of per-second best bid and ask quotes, and trade price and side records, for eight currency pairs. Volume data is not provided. One month is from the pre-EMU period and the other is post-EMU data. The first half of the data spans the period 01/08/98 to 04/09/98 and the other is from 01/08/99 to 03/09/99. To the best of my knowledge, no other academic researcher has been given access to high frequency EBS data.

Tick data for the 3-month STIR futures contracts, which are by far the most liquid money market instruments, come from the five European futures exchanges that I discuss below. This STIR data spans the period 01/01/97 to 31/12/00. LIFFE provides quote and trade price data, while the other exchanges provide only trade price data. All exchanges provide volume data. Although each of these datasets has been available for a while, I am not aware that anybody has brought them together before. The amalgamated dataset is new.

The number of observations assembled for this study amounted to over 12 million in total. Approximately 8 million observations come from the STIR markets. The other 4 million come from the very liquid spot FX markets. I believe that no previous researcher has amassed so much spot FX data for a single study. The most brokered inter-dealer data that any previous high frequency FX study has had access to was one trading week of data from Reuters 2000-2 used by Payne(1999). Most high frequency FX studies to date have been based on customer indicative quotes data collected from the FXFX page on the Reuters news terminal which is the only readily available source of high-frequency and is collected and sold by Olsen and Associates. Surprisingly, Reuters do not store this data themselves. Research which make use of this dataset includes Goodhart and Figliuoli(1991), Dacarogna, Muller, Nagler, Olsen and Pictet(1993), Zhou(1996), Goodhart et al.(1996), Evans(1997), Hartmann(1998) and Andersen, Bollerslev and Das(2001).

The EBS data comprises anonymous, per-second quote and trade prices for 8 exchange rates in which they are the dominant inter-dealer venue: EUR/USD, USD/DEM, USD/JPY, USD/CHF, EUR/JPY, DEM/JPY, EUR/CHF and DEM/CHF. However, I treat these as 5 exchange rates because I treat the EUR as the linear successor of the DEM on the grounds that the DEM acted as a pan-European vehicle currency prior to the emergence of the EUR. The quotes provided are firm quotes, not indicative since they relate to limit orders entered by dealers into EBS's electronic trading system. EBS also reveal whether transactions were buyer or seller initiated. EBS did not supply volume data. In later work, I will use the number of trades as a proxy for volume.

The bulk of the STIR futures data comes from LIFFE and consists of quote and trade prices and volumes. Most of the LIFFE data was made available to me by my supervisors. Whether individual trades were buyer or seller initiated is not revealed. The identity of quote submitters or parties to a trade is also not revealed. The first part of the LIFFE quotes data comes from the floor-based market, when trading took place in trading pits. The second part consists of limit order prices from the electronic trading regime introduced in September 1999. Data from the EUREX, MATIF and MEFF exchanges contain transaction prices and volumes only, i.e. no quotes. Once again, whether the initiator was a buyer or a seller is not revealed. In all, 11 STIR instruments are examined in this study. There are, from LIFFE: Euribor, Euro Libor, Euromark, Eurolire, Short Sterling and Euroswiss, from EUREX: Euribor and Euromark, from MATIF: Euribor and PIBOR, and from MEFF: MIBOR. I treat these as 7 STIR contracts for two reasons. First, the Euro Libor was quickly absorbed into LIFFE's Euribor, so I combine them. Second, I treat the PIBOR and the two Euromark contracts as the linear antecedents to the Euribor, for two reasons. First, each country's STIR liquidity needs would be channelled directly into that country's national exchange – at least initially. Second, the leading role played by the DEM and Germany within Europe should have attracted more international hedgers to the Euromark than to the STIR contracts from other European countries, in advance of EMU. In every case, only the 4 most liquid delivery contracts are used. These are the March, June, September and December contracts.

4.3 The Sample Period

The data used in this study span a period which includes EMU. This is useful because EMU involved certain structural changes which directly revealed important aspects of the microstructure of these markets. Also, four of the ten STIR instruments underwent the transition from floor to electronic trading during the period studied. Their changeover days are depicted in figure 4.1. The timing of these switches is also arguably linked to EMU. In addition, four STIR instruments also experienced a change in minimum tick size which would probably not have occurred without EMU. Figure 4.1 places all these events in the context of a timeline. The STIR futures data spans the length of the X-axis. The EBS data is shown as two columns which span just over a month each. Changes in minimum tick sizes are indicated by coloured asterisks, while coloured dots show when certain instruments switched to electronic trading. The date on which the Euro Libor conceded defeat to the Euribor is depicted as a black square. The EMU convergence event is illustrated by a heavy dashed vertical line.

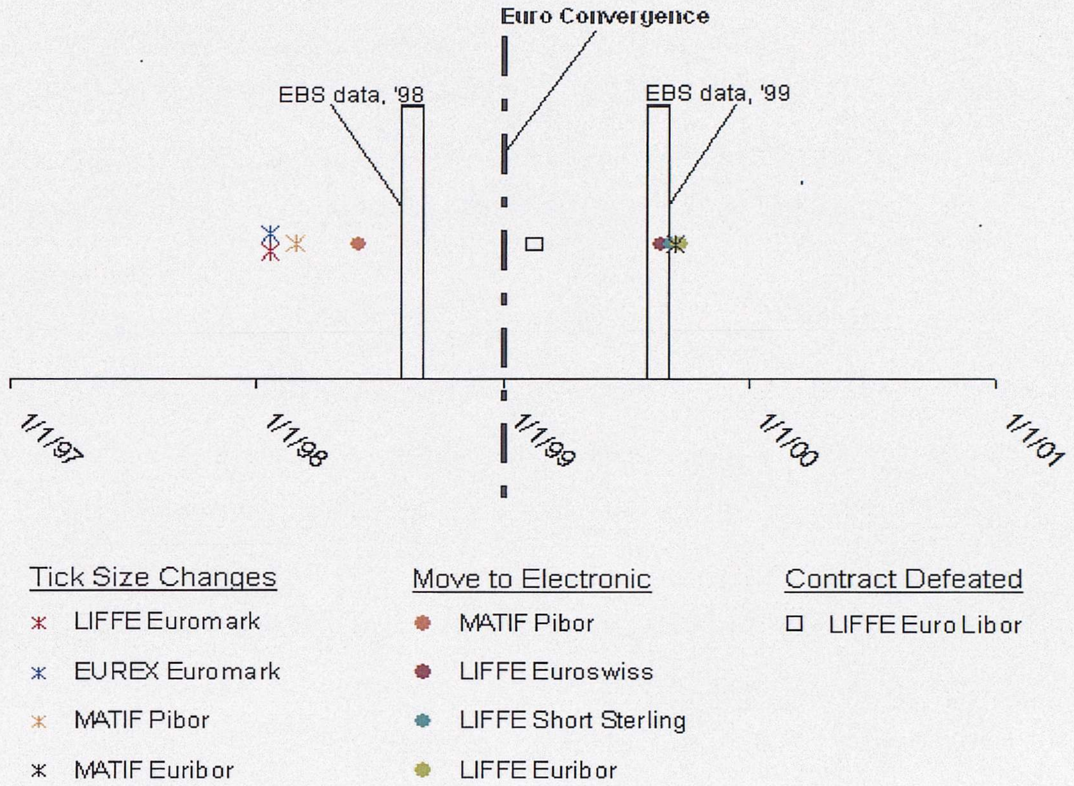

Figure 4.1: Timeline of Data and Events

% BREAKDOWN OF DAILY VOLUME IN 3-MONTH EURIBOR FUTURES PER EXCHANGE OVER 1999 AND 2000

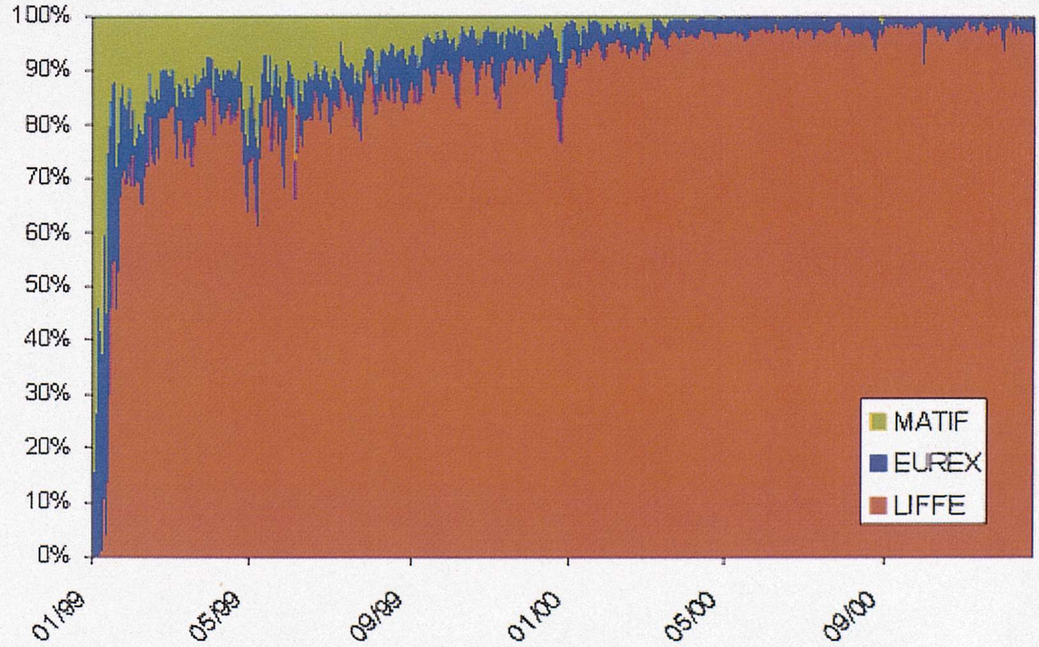

Figure 4.2: LIFFE's rise to dominance in the 3-month Euribor futures market.

Of the six contracts which are not shown to have converted to electronic trading, four were already electronically based during the data period. These are MEFF's MIBOR, MATIF's Euribor and EUREX's Euribor and Euromark. The remaining two, LIFFE's Euromark and Eurolire contracts, had ceased to exist before LIFFE moved over to electronic trading in 1999. So, their data is all floor based.

On 17/2/99, LIFFE signalled defeat for the Euro Libor by offering a zero-cost voluntary position conversion facility for transferring Euro Libor derivatives to the equivalent Euribor instrument. On 29/6/00, the Euribor finally vanquished the Euro Libor, absorbing the latter's liquidity.

The long STIR data series is split into four blocks over most of this study. These span the sub-periods: 01/01/97 to 19/01/98, 20/01/98 to 31/12/98, 01/01/99 to 22/08/99 and 20/09/99 to 31/12/00. The first break occurs because both Euromark contracts, trading on both LIFFE and EUREX, underwent a 50% reduction in minimum tick size on 20/01/98. In order to compare all instruments within contiguous time-blocks, a break across the whole dataset is introduced. The second break is for EMU on 01/01/99. The final break signifies the transition of STIR futures from floor to electronic trading on LIFFE. Since different contracts were converted over a period of successive weeks, approximately one month of data from the dataset is removed from the analysis.

In contrast to the relatively minor affect that euro convergence had on European equity and long term bond markets, its impact on the FX and STIR markets was dramatic. Currency convergence wiped out eleven European currencies and all cross-exchange-rate which had involved any of them. Also as a direct result of convergence, European STIR futures could no longer be denominated in any of those eleven European currencies. In their place, a new pan-European STIR futures contract emerged - the Euribor.

Global foreign exchange market volume dropped instantly when internal European cross-rates ceased to be tradable. Figure 4.2 reflects the tough competition from the start between futures exchanges in Frankfurt, London and Paris to win Euribor volume. A short time after convergence, in February 1999, LIFFE's fledgling Euro Libor contract gave way to LIFFE's competing Euribor contract. The latter went on to attract most of the volume from its continental rivals. The overall upheaval caused some inter-European cross rate dealers and smaller country money market dealers to be driven out. As figure 4.2 shows, LIFFE steadily gathered liquidity from the other exchanges throughout 1999 and the first quarter of 2000, when it levelled off at about 98% of the total global Euribor volume. Both MATIF and EUREX remained active in the contract after that, but liquidity fell to a trickle, especially on MATIF.

4.4 The Exchanges

The following subsections provide some information on the recent histories of the exchanges that have provided data for this study.

4.4.1 EBS

The EBS partnership was established in 1993 by 12 of the world's largest foreign exchange banks[1]. Since then, the EBS Spot Dealing System has grown to be the dominant intermediary for brokered inter-dealer spot FX trading. There is only one other intermediary for brokered inter-dealer spot FX trading. This is the Dealing 2000-2 system from Reuters. The Reuters system dominates in currency pairs that involve Sterling and Scandinavian currencies. EBS dominates in all other spot currency trades.

EBS claim average daily spot volume in excess of $80 billion. BIS(2001) estimates average daily total global FX volume at $1.17 trillion, down from $1.43 trillion in

[1] The banks who established the EBS Partnership were: ABN-Amro, Bank of America, Barclays, Citibank, Commerzbank, Credit Suisse First Boston, HSBC, JP Morgan, Lehman Brothers, Royal Bank of Scotland, SEB and UBS. The MINEX Corporation of Japan is the only non-bank member.

1998 (BIS(1998)). Of this, spot FX volume amounts to around 40% ($600 billion) in 1998 and 33% ($400 billion) in 2001. This credits EBS with between 13% and 20% of total global spot FX volume. However, even the higher number understates EBS's significance in the major currency pairs. Lyons(2001) states that, in the major markets, brokered inter-dealer trades account for a full third of spot volumes. Furthermore, Bank of England (2001) estimates show that the percentage of all FX trades which are transacted via electronic broking venues had doubled from their 33% share in 1998. The Federal Reserve Bank of New York (2001) provides an even more startling estimate of 71% of all FX trades. Killeen, Lyons and Moore(2003) credit EBS with the majority market share, which they claim is increasing. Furthermore, the BIS(1998) and BIS(2001) reports both state that brokered inter-dealer trading is growing more rapidly than either the direct inter-dealer or the dealer-customer category.

4.4.2 LIFFE

From 1982 to 1998, most derivatives trading on LIFFE was conducted by open outcry, i.e. face-to-face in the trading pits. From 1989, floor trading was supplemented with APT (Automated Pit Trading). APT enabled trading to be carried out after hours electronically, but this was an electronic form of open outcry and not the order book form of electronic trading which is prevalent today. In May 1998, as a result of the shock of loosing the 10 year German Bund contract to the electronic EUREX, LIFFE began developing an electronic order-driven trading platform of its own, LIFFE CONNECT. From the end of 1998 and through 1999, LIFFE implemented a phased transition of all futures and options contracts from the trading floor to their new electronic marketplace. STIRS were moved on to LIFFE CONNECT from late August, through September 1999.

An issue for the LIFFE data that does not affect the other exchanges is there are two kinds of prices in the LIFFE dataset. Prices quoted under the floor-trading regime are indicative only. Such quotes may be viewed as a form of advertising. They are not binding. In contrast, under the LIFFE CONNECT regime, quotes are binding. This is because they are entered into the system as limit orders. If another trader "takes" a posted limit order, that constitutes a trade.

According to press releases on the LIFFE website, LIFFE CONNECT is the world's second largest derivatives exchange and largest European exchange by value, with daily traded values exceeding $400 billion in 2001. STIR derivatives constitute about 97% of all LIFFE trading. So, it is not surprisingly that LIFFE is by far the largest STIR data source in the present study.

After much speculation and several failed courtships with other suitors, LIFFE was taken over by Euronext in 2002. Euronext have since ceased trading the Euribor on MATIF, which is also part of their stable, in order to concentrate all their STIR futures trading on LIFFE.

4.4.3 EUREX

DTB (Deutsche Terminbörse) and SOFFEX (Swiss Options and Financial Futures Exchange) had been early pioneers in providing electronic access to derivatives markets. On 04/09/1997, they collaborated to create a joint electronic platform for trading derivatives. In May 1998, they formally named this platform "EUREX". EUREX also provides an automated and integrated joint clearing house for products and participants. EUREX claims to be the largest futures exchange by the number of contracts traded.

4.4.4 MATIF

MATIF has offered financial derivative instruments since 1986. It switched to electronic trading over April and May 1997. It nurtured and made a success of the CME/Reuters spawned GLOBEX network, by being its biggest volume contributor. GLOBEX was the first global electronic trading network for futures and option contracts. As part of GLOBEX, the Euro GLOBEX agreement allowed members of MATIF, MEFF and Italy's MIF to trade each others' contracts from their own workstations across interconnected electronic trading platforms. The MATIF/MEFF Renta Fija trading link has been in operation since 05/02/1999 and the MATIF/MIF trading link started in June 1999.

In late 2000, MATIF merged with smaller exchanges in Amsterdam and Brussels to form Euronext. Subsequently, in 2002, Euronext acquired LIFFE.

4.4.5 MEFF

MEFF is the oldest electronic derivatives exchange in Europe, dating back to March 1989. From the beginning, it aimed to be a totally electronic solution, integrating trading, clearing and settlement in a single widely admired system. It offers a front office system to provide users with market information critical in decision making. In 1994, MEFF introduced their back office (MIBOS) system which provides direct access to house and client portfolio positions, allowing these to be monitored and controlled. MEFF claims its technology has attracted much interest from other derivatives exchanges. MEFF customers can trade on MATIF via the GLOBEX network.

4.5 STIR Contract Specifications

The following 3 month STIR futures contracts are used in this study: Short Sterling, Euroswiss, Eurolire, Euromark, Euribor, Euro Libor, Pibor and Mibor. These come from all four futures exchanges. Futures exchanges tend to have their own idiosyncratic ways of presenting contract specifications. For this reason, I use my own template as a standard and fill this with information provided by the exchanges. The first two contracts span the pre- and post-euro time periods, while the last six do not. The 11 tables below provide summary details of the specifications of the STIR contracts used.

	Contract Specifications
Contract	3-month Sterling (Short Sterling) Interest Rate Future
Exchange	LIFFE
Underlying Instrument	3-month Libor: London Interbank Offered rate on 3-month Sterling deposits, calculated by the British Bankers Association (BBA)
Size	GBP 500,000
Settlement	Cash settled.
Delivery Day	1st business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 2 decimal places.
Tick Size (Value)	0.01% (GBP 12.50)
Delivery Months	March, June, September, December, and 2 serial months, such that 22 delivery months are available for trading, with the nearest 3 being consecutive calendar months
Last Trading Day	3rd Wednesday of the delivery month
Final Settlement Price	Based on the BBA Libor Offered Rate for 3-month Sterling deposits at 11.00 (London time) (12:00 CET) on the Last Trading Day.
Trading Hours	08:05 - 18:00 (UK time)

Table 4.1: Short Sterling (LIFFE)

	Contract Specifications
Contract	3-month Euro Swiss Franc (Euroswiss) Interest Rate Future
Exchange	LIFFE
Underlying Instrument	3-month Libor: London Interbank Offered rate on 3-month Euroswiss Franc deposits, calculated by the British Bankers Association (BBA)
Size	CHF 1,000,000
Settlement	Cash settled.
Delivery Day	1st business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 2 decimal places.
Tick Size (Value)	0.01% (CHF 25)
Delivery Months	March, June, September, December, such that the next 8 quarterly delivery months are available for trading.
Last Trading Day	2 business days prior to the 3rd Wednesday of the delivery month
Final Settlement Price	Based on the BBA Libor Offered Rate for 3-month Euroswiss Franc deposits at 11.00 (London time) (12:00 CET) on the Last Trading
Trading Hours	08:10 - 17:55 (UK time)

Table 4.2: Euroswiss (LIFFE)

	Contract Specifications
Contract	3-month Euro Lire (Eurolire) Interest Rate Future
Exchange	LIFFE
Underlying Instrument	3-month Libor: London Interbank Offered rate on 3-month Eurolire deposits, calculated by the British Bankers Association (BBA)
Size	ITL 1,000,000,000
Settlement	Cash settled.
Delivery Day	1st business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 2 decimal places.
Tick Size (Value)	0.01% (ITL 25,000)
Delivery Months	March, June, September, December
Last Trading Day	2 business days prior to the 3rd Wednesday of the delivery month
Final Settlement Price	Based on the BBA Libor Offered Rate for 3-month Eurolire deposits at 11.00 (London time) (12:00 CET) on the Last Trading Day.
Trading Hours	07:55 - 18:00 (UK time)

Table 4.3: Eurolire (LIFFE)

	Contract Specifications
Contract	3-month Euro Deutschemark (Euromark) Interest Rate Future
Exchange	LIFFE
Underlying Instrument	3-month Libor: London Interbank Offered rate on 3-month Eurodeutschemark deposits, calculated by the British Bankers Association (BBA)
Size	DEM 1,000,000
Settlement	Cash settled.
Delivery Day	1st business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 2 decimal places.
Tick Size (Value)	0.01% (DEM 25)
Delivery Months	*March, June, September, December, and 2 serial months, such that* 18 delivery months are available for trading, with the nearest 3 being consecutive calendar months
Last Trading Day	2 business days prior to the 3rd Wednesday of the delivery month
Final Settlement Price	Based on the BBA Libor Offered Rate for 3-month Eurodeutschemark deposits at 11.00 (London time) (12:00 CET) on the Last Trading Day.
Trading Hours	07:30 - 18:00 (UK time)

Table 4.4: Euromark (LIFFE)

	Contract Specifications
Contract	3-month Euro (Euribor) Interest Rate Future
Exchange	LIFFE
Underlying Instrument	3-month Euribor: European Interbank Offered rate on 3-month Euro deposits, calculated by the European Bankers Federation (EBF/FBE)
Size	EUR 1,000,000
Settlement	Cash settled.
Delivery Day	1st business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 3 decimal places.
Tick Size (Value)	0.005% (EUR 12.50)
Delivery Months	March, June, September, December, and 4 serial months, such that 24 delivery months are available for trading, with the nearest 6 being consecutive calendar months
Last Trading Day	2 business days prior to the 3rd Wednesday of the delivery month
Final Settlement Price	Based on the EBF Euribor Offered Rate for 3-month Euro deposits at 11.00 (CET) (10:00 London time) on the Last Trading Day.
Trading Hours	07:00 - 18:00 (UK time)

Table 4.5: Euribor (LIFFE)

	Contract Specifications
Contract	3-month Euro (Euro Libor) Interest Rate Future
Exchange	LIFFE
Underlying Instrument	3-month Euribor: European Interbank Offered rate on 3-month Euro deposits, calculated by the British Bankers Federation (BBA)
Size	EUR 1,000,000
Settlement	Cash settled.
Delivery Day	1st business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 3 decimal places.
Tick Size (Value)	0.005% (EUR 12.50)
Delivery Months	March, June, September, December, and 4 serial months, such that 24 delivery months are available for trading, with the nearest 6 being consecutive calendar months
Last Trading Day	2 business days prior to the 3rd Wednesday of the delivery month
Final Settlement Price	Based on the BBA Libor Offered Rate for 3-month Euro deposits at 12.00 (CET) (12:00 London time) on the Last Trading Day.
Trading Hours	07:00 - 18:00 (UK time)

Table 4.6: Euro Libor (LIFFE)

	Contract Specifications
Contract	3-month Euro Deutschemark (Euromark) Interest Rate Future
Exchange	EUREX
Underlying Instrument	3-month Libor: London Interbank Offered rate on 3-month Eurodeutschemark deposits, calculated by the British Bankers Association (BBA)
Size	DEM 1,000,000
Settlement	Cash settled.
Delivery Day	1st business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 2 decimal places.
Tick Size (Value)	0.01% (DEM 25)
Delivery Months	March, June, September, December, and 2 serial months, such that 14 delivery months are available for trading, with the nearest 3 being consecutive calendar months
Last Trading Day	2 business days prior to the 3rd Wednesday of the delivery month
Final Settlement Price	Based on the BBA Libor Offered Rate for 3-month Eurodeutschemark deposits at 11.00 (London time) (12:00 CET) on the Last Trading Day.
Trading Hours	08:30 - 19:00 (CET)

Table 4.7: Euromark (EUREX)

	Contract Specifications
Contract	3-month Euro (Euribor) Interest Rate Future
Exchange	EUREX
Underlying Instrument	3-month Euribor: European Interbank Offered rate on 3-month Euro deposits, calculated by the European Bankers Federation (EBF/FBE)
Size	EUR 1,000,000
Settlement	Cash settled.
Delivery Day	1^{st} business day after the last trading day.
Quotation	100 *minus* the rate of *interest*(%), to 3 decimal places.
Tick Size (Value)	0.005% (EUR 12.50)
Delivery Months	March, June, September and December, such that the next 12 quarterly delivery months are available for trading.
Last Trading Day	2 business days prior to the 3^{rd} Wednesday of the delivery month
Final Settlement Price	Based on the EBF/ACI Euribor Offered Rate for 3-month Euro deposits at 11.00 (CET) (10:00 London time) on the Last Trading Day.
Trading Hours	08:30 - 19:00 (CET)

Table 4.8: Euribor (EUREX)

	Contract Specifications
Contract	3-month French Franc (Pibor) Interest Rate Future
Exchange	MATIF
Underlying Instrument	3-month Pibor: Paris Interbank Offered rate on 3-month French Franc deposits, calculated by AFB/Telerate
Size	FRF 5,000,000
Settlement	Cash settled.
Delivery Day	1^{st} business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 2 decimal places.
Tick Size (Value)	0.01% (FRF 125)
Delivery Months	March, June, September and December, such that the next 12 quarterly delivery months are available for trading.
Last Trading Day	2 business days prior to the 3^{rd} Wednesday of the delivery month
Final Settlement Price	Based on the interest rate supplied by AFB/Telerate on the Last Trading Day.
Trading Hours	07:00 - 22:00 (CET)

Table 4.9: Pibor (MATIF)

	Contract Specifications
Contract	3-month Euro (Euribor) Interest Rate Future
Exchange	MATIF
Underlying Instrument	3-month Euribor: European Interbank Offered rate on 3-month Euro deposits, calculated by the European Bankers Federation (EBF/FBE)
Size	EUR 1,000,000
Settlement	Cash settled.
Delivery Day	1st business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 3 decimal places.
Tick Size (Value)	0.002% (EUR 5)
Delivery Months	March, June, September, December, and 2 serial months, such that 22 delivery months are available for trading, with the nearest 3 being consecutive calendar months.
Last Trading Day	2 business days prior to the 3rd Wednesday of the delivery month
Final Settlement Price	Based on the EBF Euribor Offered Rate for 3-month Euro deposits at 11.00 (CET) (10:00 London time) on the Last Trading Day.
Trading Hours	07:45 - 22:00 (CET)

Table 4.10: Euribor (MATIF)

	Contract Specifications
Contract	MIBOR '90 Plus Interest Rate Future
Exchange	MEFF
Underlying Instrument	Madrid Interbank Offer Rate (MIBOR) interest rate on 90 day interbank deposit, published by the Bank of Spain.
Size	ESP 100,000,000
Settlement	Cash settled.
Delivery Day	1st business day after the last trading day.
Quotation	100 minus the rate of interest(%), to 2 decimal places.
Tick Size (Value)	0.01% (ESP 2,500)
Delivery Months	March, June, September and December, such that the next 8 quarterly delivery months are available for trading.
Last Trading Day	2 business days prior to the 3rd Wednesday of the delivery month
Final Settlement Price	Based on the the Bank of Spain's interbank Offered Rate for 90 day deposits C50on the Last Trading Day.
Trading Hours	09:00 - 17:15 (CET)

Table 4.11: Mibor (MEFF)

4.6 Data Acquisition, Management and Analysis

It took about 18 months to gather, prepare and marry together all of this data.

In many cases, establishing a link with the right contact proved critical is getting the data. In most cases, formal written requests for data also had to be made. These forced me to be clear and concise about the data I needed, so that the data provider would spend the minimum amount of time putting it together. In most cases, I also had to sign confidentiality agreements and assurances that data would not be used for commercial purposes and had to have these declarations stamped by the university. This bureaucracy took several months, but eventually I acquired several CDs full of data, from which I extracted many enormous files in ASCII format.

The main tool that I used to store and manipulate the data was Microsoft ACCESS. I re-formatted all STIR data from the continental exchanges so that it had the same table structure and layout as the LIFFE data. This enabled me to later write data analysis programmes only once and apply them to several tables in succession. In the process of doing this, my STIR database hit the upper size limit for ACCESS databases. I had to develop a more elaborate distributed database structure with several satellite databases acting as data depositories. These were connected, via "linked" tables, to a central database which I used for manipulation and analysis.

It was not obvious initially whether all the data spanning dozens of instruments and exchange rates could be utilised in my study. After assessment of several summary statistics, many instruments had to be rejected because they were too illiquid. I whittled the viable set down to the 8(5) spot exchange rates and 11(7) STIR futures contracts listed above.

The minimum tick size change events described above were also far from obvious in the beginning. They first emerged as strange discontinuities in the data. Only upon investigation could they be confirmed as being due to changes in the minimum tick size brought about by a change in the exchange's policy.

Inevitably there were data errors and much early work went into "cleaning" the data. With low frequency a key method for assessing data integrity is to graph it. With high frequency data this is not possible. It is difficult to find graphics routines which can handle massive numbers of observations. When one does eventually find capable software, the data emerges as a concentrated indistinguishable, amorphous blob. This is because there is simply too much of it. Focusing on narrow time segments does not provide a very satisfactory solution because the amount of data is so great that the number of segments at any meaningful resolution is still vast. The best that can be done is to bound each data series within known or clearly defined limits. For example, prices for specific futures contracts are restricted to dates before the contract expires. Also, misplacement-of-decimal-point errors are excluded by limiting data values to an arbitrary range around the mean value. In spite of these steps, certain anomalies arose which had to be investigated individually and the data amended or excluded. An example was where the USD/JPY bid-ask spread appeared negative over a few days in early August 1999. Close inspection of the data revealed that the ask price was constant over three days. This was not credible because the USD/JPY is a very liquid market in which prices change every few seconds and, besides, the bid price exhibited considerable variation over the same period. In this example, all USD/JPY observations over these days were excluded from the analysis.

Data transformation and statistical analysis were carried out using a combination of SQL queries and Visual Basic programmes which are both supported in the Microsoft ACCESS software. I benefited enormously from the ACCESS experts on the "tek-tips" website (www.tek-tips.com), who helped me solve many difficult SQL query and programming problems. The size of the data tables meant that computations frequently had to be run overnight, even when using an up-to-date high specification computer.

I used Microsoft EXCEL for producing graphs and some low frequency data analysis. For regression analysis, I used TSP.

Along the way, I spent a lot of time and energy investigating several other software packages for various aspects of this project, only to decide that these were not appropriate because, for example, they could not handle the size of the series or they could not implement some necessary constraints.

4.7 Conclusion

This chapter describes the two exceptional datasets that form the foundation for this study and illustrates why they are such an important advance in empirical finance. It also discusses the time period and sources from which these dataset came. Finally, the work that I put in, and the tools that I use, to obtain, format, prepare and analyse these data are briefly outlined.

Chapter 5.

Microstructure Effects

5.1 Introduction

Microstructure effects refer to aspects of prices which arise because of structural features of the market or of the pricing unit rather than from economic agent' behaviour. Examples of microstructure effects include price discreteness, price clustering and non-synchronous trading. This chapter emphasises the links that the first two of these have with price innovation and the bid-ask spread but also briefly addresses the effect of non-synchronicity. It establishes the presence and explores the cause of price clustering which is evident in the data for spot FX and STIR futures market instruments. I develop two new test statistics based around established price clustering explanations. I propose a brand new theory to capture a type of price clustering behaviour not previously addressed in the literature. My findings contribute to a debate about changes in bid-ask spreads and volumes in the spot FX market since EMU, by providing more detail on both the pre-EMU and post-EMU markets than had been available previously. I also offer some insight into why the observed changes occurred. Finally, I produce the first precise measure of the true bid-ask spread cost of re-denomination from DEM to EUR.

One reason that the present study is particularly noteworthy is that price clustering studies are rare in both the FX and the futures markets. Studies of price clustering in STIR futures seem especially rare. Previous FX clustering studies have only used quote price cluster patterns, whereas my analysis also uses trade price data.

EMU has brought widespread change to financial markets. Much of this change is directly due to the re-denomination of certain instruments from DEM to EUR. Since these currency units are of different values, the nature of the price discreteness affecting instruments which are now denominated in EUR will be different from what it was under DEM. This point is exemplified by the fact that the smallest sized bid-ask spread and smallest price increment for the EUR are both 74% greater than that for the DEM, independent of any currency drift.

It is an acknowledged fact that volumes have decreased in the inter-bank spot FX market since EMU. It is becoming increasingly accepted that the EUR/USD bid-ask spread rose at the same time. Recently, Goodhart, Love, Payne and Rime(2002) attracted much interest when they argued that the change in denomination from DEM to EUR is sufficient to account for observed increases in bid-ask spreads, post-EMU. Their "price granularity" hypothesis combines the difference in price discreteness between the DEM and the EUR, with trader inertia in setting bid-ask spreads. The authors argue that the fall in volume is a coincidence primarily due the increased use of electronic inter-dealer broking and to banking industry consolidation. This model was put forward in response to controversial work by Hau, Killeen and Moore(2000 and 2002), who suggested that lower FX trading volumes and higher bid-ask spreads since EMU, are both due to "market transparency". The latter hypothesis centres on the idea that the availability of fewer FX currency pairs after EMU, makes risk management harder to do. Hau et al. suggest this causes market makers to charge higher bid-ask spreads, which result in lower volumes. In the same volume as Goodhart et al(2002), Detken and Hartmann(2002) used a wider but lower frequency dataset and reached conclusions which Goodhart et al(2002).

Until very recently, price discreteness and clustering studies on FX markets were rare. Goodhart and Curcio(1991) was the only widely cited paper. Besides

Goodhart et al(2002), Sopranzetti and Datar(2002) also produced a price clustering study for the FX market, using data from the Federal Reserve Bank of New York. Using Royal Bank of Scotland data, Osler(2003) found that the clustering of "stop loss" and "take profit" orders at certain price points explained why certain technical analysis forecasts have predictive power. The principal reason for the paucity of research is lack of FX data.

Since more data has been available for futures markets for a longer period, one might expect that the price discreteness and price clustering aspects of futures markets to be better researched. In fact, they do not appear to be. The main papers in this area include: Ball, Torous and Tschoegl(1985), Brown, Laux and Schachter(1991), ap Gwilym, Clare and Thomas(1998a and 1998b) and ap Gwilym and Alibo(2003). However, price clustering in STIR futures contracts appears not to have been studied before.

The vast bulk of the work done on price discreteness and clustering focuses on the equity markets. The most notable contributions include: Harris(1991), Christie and Schultz(1994), Christie, Harris and Schultz(1994), Kleidon and Willig(1995), Aitken, Brown, Buckland, Izan and Walter(1996), Grossman, Miller, Fischel, Cone and Ross(1997), Woodward(1998), Weston(2000), Kandel, Sarig and Wohl(2001) and Brown, Chua and Mitchell(2002).

5.1.1 How Does Price Discreteness Affect the Bid-Ask Spread and Price Innovation?

Prices move in discrete units. The precise resolution of these discrete units may be imposed by a regulator or an exchange, or it may arise as a market convention. Price discreteness is clearly important for bid-ask spreads because the minimum tick size places a lower bound on how low the (non-zero) bid-ask spread can be. It also determines the increments by which it can increase.

The significance of price innovation effects is harder to penetrate. Campbell, Lo and MacKinley(1997), (page 113), use a sequence of stock return plots with progressively finer scales to illustrate graphically how discreteness imposes structure on the return generating process. Gottlieb and Kalay(1985) found that "...the variance and...the higher order moments of the rate of return of stocks are upward biased due to the discreteness of observed stock prices.". Harris(1990) found similar. Osler(2003) suggested that discreteness in FX rates could contribute to excess kurtosis.

Hausman, Lo and MacKinley(1992) link price discreteness to path dependence when their ordered probit analysis concludes that the conditional mean of price changes is determined by the sequence of preceding order-flow. Rime(2000), Lyons(2001), Evans(2002) and Evans and Lyons(2002) all provide strong evidence that FX rates rise in fall in line with aggregate order flow, suggesting that FX rates are similarly path dependent. Osler(2003) found evidence of links between FX rate path dependence, certain "round number" price points and technical analysis.

Price clustering effects bid-ask spreads and price innovation in the same way as price discreteness. Harris(1990) discusses in detail how these two features interact.

5.1.2 What causes Price Clustering?

"Price clustering" refers to traders' tendency to not use the full range of price points uniformly, but rather to gravitate towards some numbers and avoid others.

Yule(1927) observed that numerical clustering arose systematically from errors that people made when asked to read numbers from a scale. Osborne(1962) was the first to address clustering in the context of financial prices. He was followed by Niederhoffer(1965). Niederhoffer(1966) pointed out that price clustering could be at odds with market efficiency. Niederhoffer and Osborne(1966) find profitable trading rules based on cluster frequencies..

Following Harris(1991), Goodhart and Curcio(1991) compare "attraction" with "price resolution" as possible explanations of price clustering in the FX market. The attraction hypothesis focuses on people's fondness for round numbers, which is the same kind of rounding behaviour that Yule(1927) identified. Aitken et al(1996) found evidence of this type of clustering in the Australian stock market. Kandel et al.(2001) found the same effect in the Israeli IPO market. ap Gwilym et al(1998) also find it for international bond futures. The resolution hypothesis, put forward by of Ball et al(1985), asserts that clustering is the natural result when the market had reached "the optimal degree of price resolution". Harris(1991) and Grossman et al.(1997) both use the example of the housing market to convey the idea of price resolution. They point out that it would be inefficient, in terms of search and negotiations cost, to fine tune house prices to the nearest penny. Similarly, in the case of the spot foreign exchange market, rates could theoretically contain an infinite number of digits. However, at some point, the marginal cost of having to deal with an extra digit will exceed the marginal benefit of a slightly more accurate price. Goodhart and Curcio(1991) find evidence in favour of price resolution in their price (=FX rate) data and of attraction in their bid-ask spreads data. Harris's(1991) own results for US stocks are consistent with price resolution hypothesis. As are those Grossman et al(1997) from a variety of global markets. Most notably, Ball et al(1985) found this effect at work in the gold market.

Harris(1991) and Brown et al(1991) both suggest negotiation could produce a two-tier price system, which would give the appearance of clustering when combined. This negotiation hypothesis asserts that large trades result in harder bargaining and progressively more finely tuned trade price, while small trades make do with cruder pricing from a reduced set of prices. The negotiation hypothesis is an interpretation of the price resolution hypothesis, rather than an alternative to it.

Most recent papers to address the subject of price clustering either start with, or make early reference to Christie and Schultz(1994). In the early 1990s, these researchers caused a stir with a price clustering study, when they suggested that widespread NASDAQ market makers avoidance of odd-eights quotes could amount to tacit collusion to maintain wider bid-ask spreads. They could not explain this observation in terms of stock price levels and volatility, or in terms of Harris's negotiation hypothesis. So, they conclude that they "...are unable to envision any scenario in which 40 to 60 dealers who are competing for order flow would simultaneously and consistently avoid using odd-eighth quotes without an implicit agreement to post quotes only on the even price fractions", and they consider this evidence of an "...apparent lack of competitiveness of the NASDAQ market". This collusion hypothesis has gathered huge support, mainly because the use of odd-eight quotes increased following the publication of their results and increased again after subsequent rule changes.

Several competent and convincing rebuttals of the Christie and Schultz(1994) result have emerged since their paper was published. These include Kleidon and Willig(1995), Grossman et al(1997) and Woodward(1998). On the other hand, other sources like the US Securities and Exchange Commission (1996) and Weston(2000) seem to support the Christie and Schultz(1994) case. The debate is still considered open.

Grossman et al(1997) show that price clustering is a both a common and a variable feature across global financial markets and that, instead of NASDAQ being out of step with NYSE, it is actually NYSE that is anomalous for its lack of clustering compared with other markets. They go on to show that a wide variety of factors contribute to price clustering. Market structure can play a part. Also, whether the quotes are binding or not matters. Most importantly, they link price clustering to the inventory and information costs of market making, suggesting that clustering should be high when these costs are high. They also expect clustering to be negatively related to volume and positively related to volatility. In contrast, Sopranzetti and Datar(2002) found a positive empirical relationship between FX rate clustering and volume. Grossman et al(1997) borrow heavily from Harris(1991) who had asserted that "clustering increases with price level and volatility, and decreases with capitalization and transaction frequency". I will refer to this set of predicted associations as the "cost of market making" hypothesis. Note that these conclusions are not restrictive in relation to final-digit price clustering pattern. They could be entirely consistent with either the attraction or resolution hypotheses described above.

5.1.3 The Price Concentration Hypothesis

Most of what has been written on price clustering in recent years adopts a perspective best summarised by the question: "Is the uneven use of all available final digits evidence of market inefficiency?". Opinion is divided on the answer. I would like to introduce an entirely new perspective on the debate by asking the question: "Could the minimum tick size be set too high to permit 'normal' price clustering to occur?". The argument here is that price could be prevented from taking on attraction- or resolution-type patterns because the tick size is overly restrictive. This perspective leads one to contemplate an entirely different form of price clustering from what has been discussed hitherto – price concentration. By price concentration, I mean the localised concentration of final digits in one part of the available range because certain prices have a very high frequency. One interpretation of such tightly bunched prices is that their price resolution is not high enough. In other words, a higher price resolution (lower minimum tick size) would open up a greater array of prices and a broader dispersion of final digit values. I refer to these ideas as the price concentration hypothesis.

The price concentration hypothesis can best be illustrated using an instrument that does not actually exhibit price concentration in its final digit pattern at all. As is shown in figure 5.1(e), the EUR(DEM)/USD exchange rate has the usual kind of final digit pattern in both 1998 and 1999, with more price observations ending in 0 and 5 than in the other digits. Figures 5.1(a) to 5.1(d), I graph the frequency of observations at each digit to the left of the final digit, starting at the leftmost digit. The five graphs reflect the five significant digits normally used in spot FX rates.

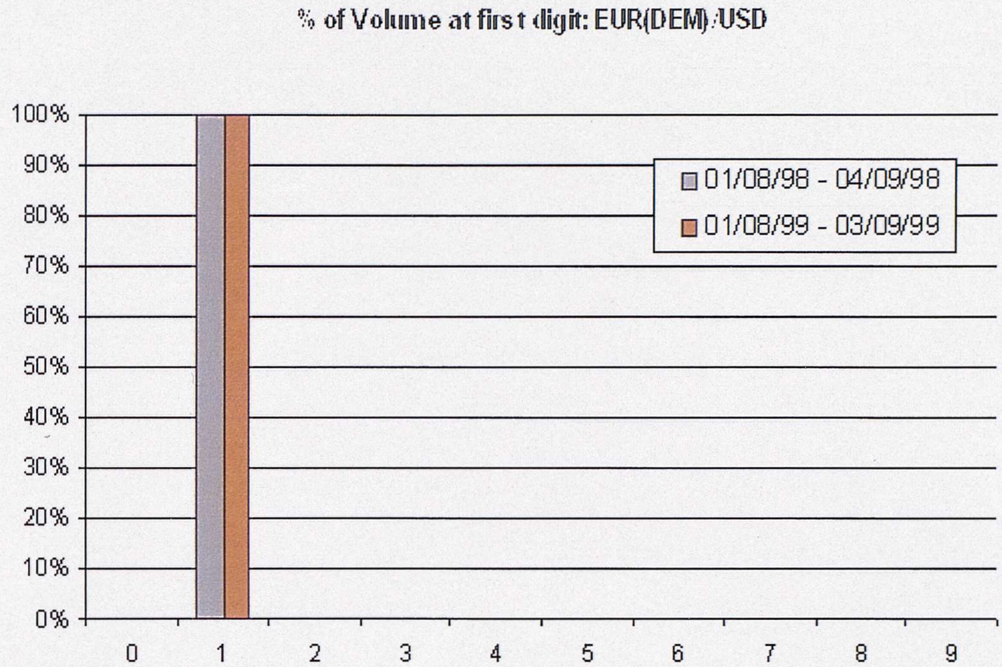

Figure 5.1(a): All rates for both EUR/USD and USD/DEM start with the number 1

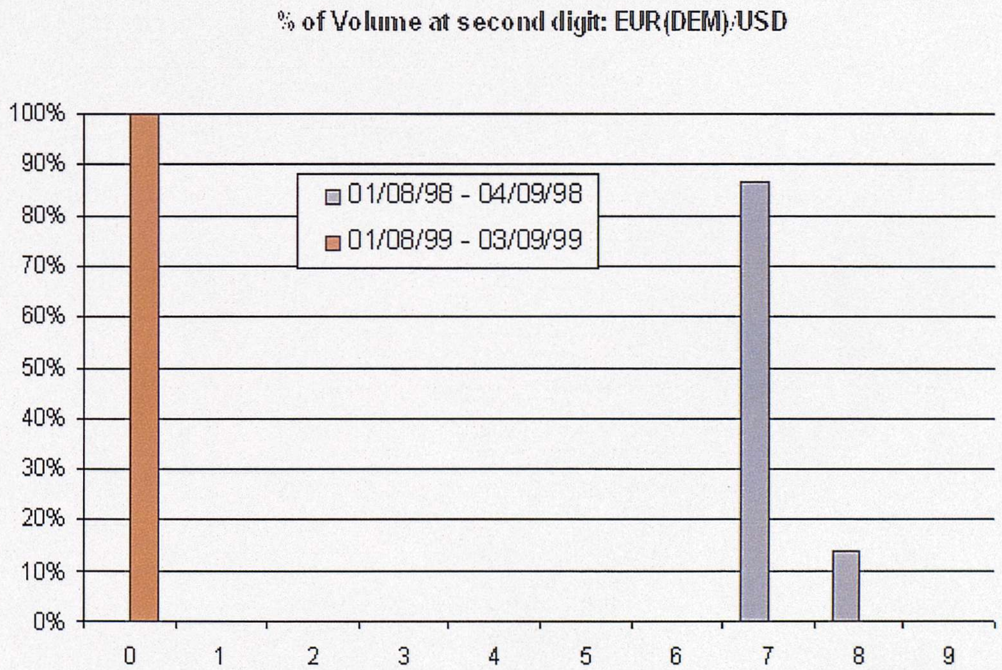

Figure 5.1(b): Still all of EUR/USD concentrated at a single digit(0), while some dispersion emerges in USD/DEM

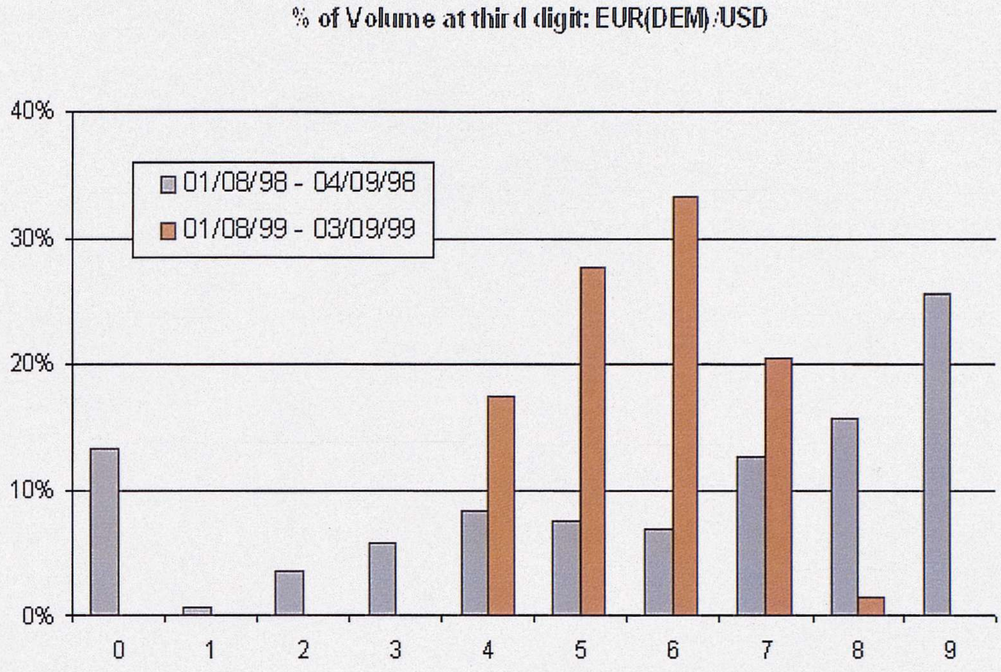

Figure 5.1(c): Wide dispersion in USD/DEM, less dispersion in EUR/USD

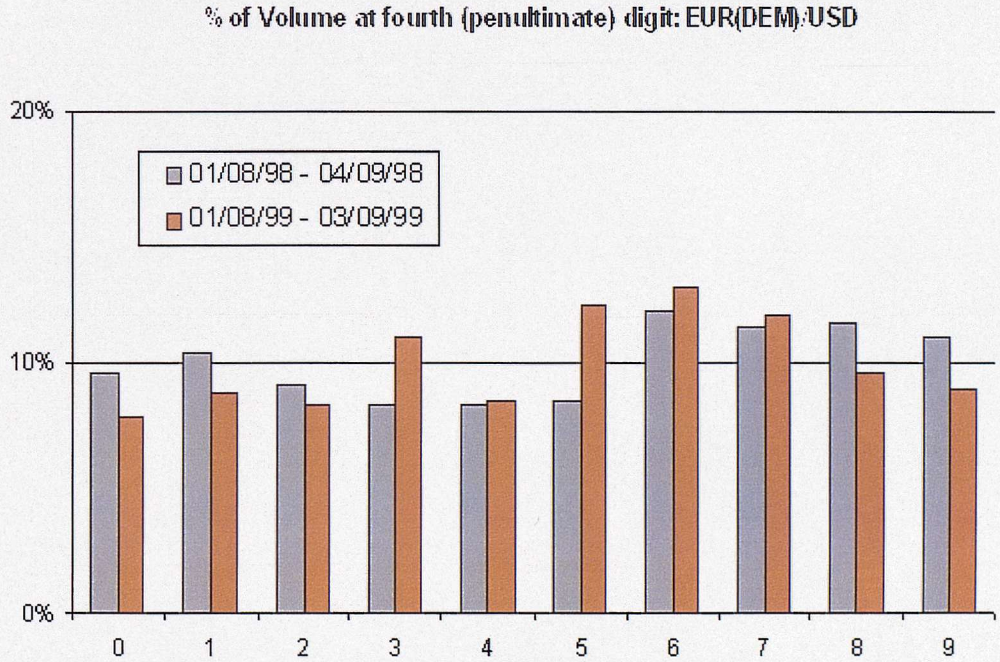

Figure 5.1(d): Bunching to the right in penultimate digit.

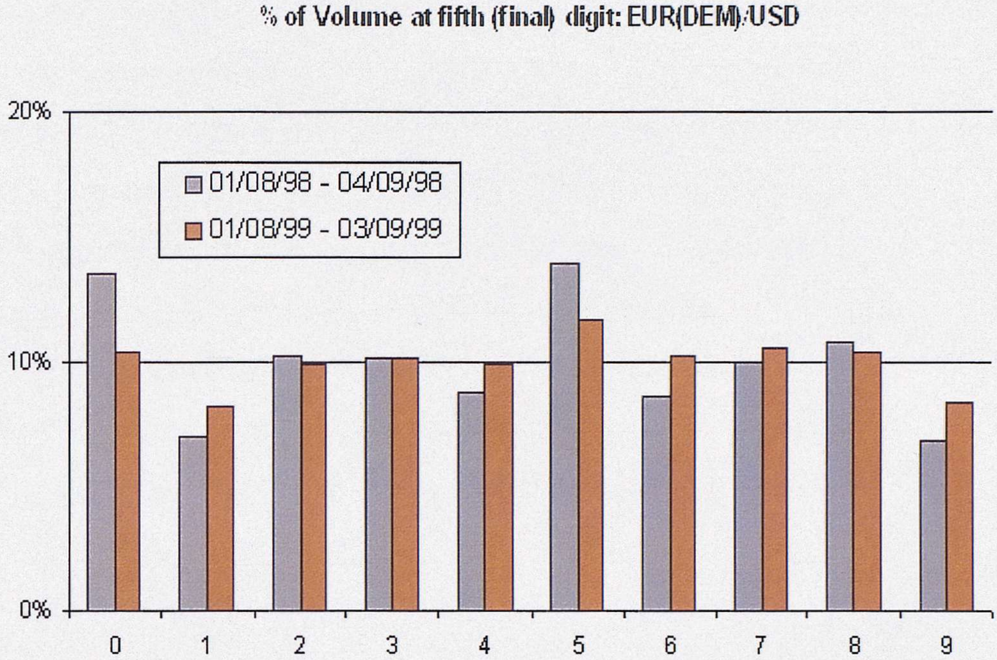

Figure 5.1(e): Resolution type price clustering with more observations at 0 and 5.

Neither the attraction pattern nor the resolution pattern is remotely apparent before figure 5.1(e). However, none of the pre-final number distributions are uniform. They exhibit successively decreasing levels of price concentration. The most important graph here is figure 5.1(d) because this represents the kind of price concentration that one would most expect to find empirically. There is a pronounced bunching in the upper half of the set of available digits for both the USD/DEM and the EUR/USD. The price concentration hypothesis implies that if this pattern were found in the final digits of an instrument, then a lower minimum tick size would probably permit something like the attraction or resolution final digit pattern to emerge.

This remainder of this chapter is divided into three main parts. The next section lays out the statistical methods and hypotheses that I use to probe the price discreteness and price clustering properties of both prices and bid-ask spreads and describes the data to be used. The third part shows my findings. The final section contains my conclusions.

5.2 Statistical Measures and Testable Hypotheses

This section takes its lead from Goodhart et al(2002). However, the datasets available to me are both wider and deeper than those available to the latter. Like Goodhart et al(2002), I compute series for percentage bid-ask spreads and for 'pip' bid-ask spreads. The term 'pip' is commonly used in the foreign exchange market in place of the word 'tick'. It may be worth acknowledging the distinction that pips arise as a matter of convention, whereas ticks are formally enforced, usually by an exchange. However, in the present work I use the two terms inter-changeably. A pip usually refers to the incremental value in the fifth non-zero digit position from the left. Note that it is not related to the position of the decimal point. For example, one pip in a USD/JPY value of 113.57 would be 0.01, while one pip for EUR/USD of 1.0434 would be 0.0001. The fact that the decimal place does not occupy a fixed position necessitates the introduction of a 'scaling factor' whose job is to bring the pip to the left of the decimal point. For example, the scaling factor for the USD/JPY is 100 and that for the EUR/USD is 10,000.

My computation method differs a little from that of Goodhart et al(2002). First, I use the occurrence of a trade to mark which times are of interest. Then I compute the ambient bid-ask spread at the time of that trade by selecting the most recently preceding bid and ask prices, but only where these are less than 1 minute old[2]. Trades that do not have both bid and ask value less than one minute old are excluded. My pip bid-ask spread is simply: *pip = (ask – bid)*scaling factor*. My percentage bid-ask spread formula is: *%_bid-ask spread = 100*(ask – bid)/ traded price*. I take issue with one feature of Goodhart et al's(2002) computations. They removed non-positive bid-ask spreads "which primarily represent the matching of market orders on the Reuters 2000-2 system." However, I contend that in an order-driven regime, zero-spreads are neither erroneous nor irrelevant. Indeed, a theoretical microstructure model for order-driven markets by Cohen, Maier, Schwartz and Whitcomb(1981) found that, in the absence of exogenous market order transactions costs, the order-driven bid-ask spread should collapse to zero. The spot FX market is order-driven and so does not depend on exogenous market makers. For this reason, I do not remove zero-spreads from my data series. As a result, many of my results appear very different from those arrived at by Goodhart et al(2002). Also, I emulate Goodhart et al.'s(2002) practice of restricting datasets to between 06:00 and 16:00 GMT.

Goodhart et al(2002) compute 9 summary statistics, encompassing both the unadjusted and time-weighted average bid-ask spreads, which I follow here. The summary statistics are: average bid-ask spread in basis points(AS), time weighted average bid-ask spread in basis points(TWAS), average bid-ask spread in pips(ASPIP), time weighted average bid-ask spread in pips(TWASPIP), total number of trades(TRAD), absolute imbalance between number of buy trades and

[2] Huang and Stoll(1997) discuss using a 2 minute interval, although they use an alternative method in the end. Since, my volume far exceeds theirs in every case, I decided a shorter interval would be appropriate. However, the choice of 1 minute is essentially arbitrary.

number of sell trades(ABIM), return volatility over 5-minute intervals (VOLAT) which uses the method outlined by Andersen, Bollerslev, Diebold and Labys(2001), the standard deviation of the bid-ask spread in basis points(STDSP), and the standard deviation of the bid-ask spread in pips(STDSPPIP).

The standard χ^2 goodness-of-fit test statistic can be applied to the observed set of final digits to detect the presence of price clustering. If clustering is absent, we would expect to see an equal number of observations at each available final digit. The formula for the χ^2 statistic used here is:

$$\chi^2 = \sum_{i=1}^{10} \frac{\left(n_i - \frac{\sum_{j=1}^{10} n_j}{10} \right)^2}{\frac{\sum_{j=1}^{10} n_j}{10}}$$

where

$n_{i(j)}$ = number of observations at final digit i(j)

Equation 5.1: χ^2 test applied to final-10-digit data

The χ^2 critical value at the 1% significance level with 9 degrees of freedom is 21.7.

Grossman et al.(1997) proposed the "standard range" as a measure of level of clustering for comparison across different instruments. The formula for the standard range is:

$$SR = (Max(\omega_i) - Min(\omega_i)) / x_i$$

where
ω_i = percentage of observations at final digit i
Min( ) = minimum value of set
Max( ) = maximum value of set
x_i = percentage at each final digit i, if no clustering

Equation 5.2: Standard Range

No clustering would give a standard range of 0. For the ten final digit data range analysed in the present study, 100% concentration would give an SR value of 10.

Goodhart and Curcio (1991) provide two ordered final digits groupings which are consistent with the Attraction hypothesis and the Price Resolution hypothesis respectively. For Attraction, the following final-digits should occur in descending order of frequency: 0, 5, {7=3}, {8=2}, {4=6} and {1=9}. In other words, if the attraction hypothesis is correct, 0 and 5 should be the most frequent, followed by 7 and 3, 8 and 2, 4 and 6, in last place, 1 and 9. I compare the differences between these groupings with the differences within each group, thereby developing this expected ordering into a test statistic:

$$A = \frac{Min\Big(\big(Av(\phi_{3,7}) - Av(\phi_{2,8})\big), \big(Av(\phi_{2,8}) - Av(\phi_{4,6})\big), \big(Av(\phi_{4,6}) - Av(\phi_{1,9})\big)\Big)}{Av\big(|\phi_3 - \phi_7|, |\phi_2 - \phi_8|, |\phi_4 - \phi_6|, |\phi_1 - \phi_9|\big)}$$

where
ϕ_i = number of observations at final digit i
$\phi_{i,j}$ = set of numbers of observations at final digits i and j
Min( ) = minimum value of set
Av( ) = mean value of set
| | = absolute value

Equation 5.3: Attraction test

Equation 5.3 is first conditioned upon 0 and 5 being the two most frequent final digits. The numerator of Equation 5.3 then takes the average of the observations for each sequential pair and then detects the minimum difference between each of these averaged pairs. The denominator calculates the average of the absolute difference within each ordered pair. The smallest difference between the pairs is divided by the average difference within the pairs. Note that the denominator will always be non-negative, so that a negative number on the numerator can never be made positive by the denominator.

An attraction test value greater than 1 denotes strong evidence in favour of the attraction hypothesis, for it shows that, not only are the sets of numbers of observations in the required order, but the difference between the set means dominates the internal differences within each set. To put it another way, the magnitude of the statistic measures the degree to which the between groups difference dominates the within group difference. Positive values below 1 also identify series with the appropriate rank ordering but differences within at least one of the groupings dominates the difference between the groups. The test is conditioned such that a value of zero results where 0 and 5 and also where the minimum difference between the group averages is exactly 0. Negative values denote that the set means are in the wrong order for the hypothesis. In other words, non-positive test values show the hypothesis is rejected, while positive value show increasing levels of acceptance.

The second hypothesis I consider is Price Resolution hypothesis. Once again, Goodhart and Curcio (1991) provide an ordered set of final digits groupings. In descending order the resolution hypothesis groupings are: 0, 5, {2=3=7=8} and {1=4=6=9}. Again, I develop a test statistic from their expected ordering:

$$R = \frac{Av\left(\phi_{2,3,7,8}\right) - Av\left(\phi_{1,4,6,9}\right)}{Max\left(Max\left(\phi_{2,3,7,8}\right) - Min\left(\phi_{2,3,7,8}\right), Max\left(\phi_{1,4,6,9}\right) - Min\left(\phi_{1,4,6,9}\right)\right)}$$

where

ϕ_i = number of observations at i

$\phi_{i,j,k,l}$ = set of numbers of observations at final digits i, j, k and l

Min( ) = minimum value of set

Max( ) = maximum value of set

Av( ) = mean value of set

Equation 5.4: Resolution test

Like Equation 5.3, Equation 5.4 depends on 0 and 5 being the two most frequent. The numerator of Equation 5.4 subtracts the average of the set that should have the lower rank from the higher ranked set. The denominator computes the largest deviation between values within each of the two sets and selects the larger one. As with the previous test, the denominator will always be non-negative, so it can never offset a negative numerator.

As for the attraction test, a resolution test value greater than 1 denotes strong evidence in favour of the resolution hypothesis. Positive values below 1 suggest the right ordering but a weak fit. Once again, non-positive test values show the hypothesis is rejected, while a positive number shows acceptance. In short, the higher the number, the better the fit!

In the original Goodhart and Curcio(1991) specification, 0 was required to be strictly greater than 5. I relaxed this requirement, but still require that 0 and 5 both be greater than the other final digit numbers.

I have not yet been able to come up with a satisfactory comprehensive statistical test for the price concentration hypothesis.

5.2.1 Data

This study uses both the spot FX best bid and ask quote price data and trade price data, per second, from EBS, as well as the European STIR tick data series that I collected from LIFFE(UK), EUREX(Germany), MATIF(France) and MEFF(Spain). Only EBS and LIFFE supply quotes data, so these are the only ones that can be used for studying bid-ask spreads.

The specific series that I focus on are: from EBS; EUR(DEM)/USD, USD/JPY, USD/CHF, EUR(DEM)/JPY and EUR(DEM)/CHF, from LIFFE; Euribor, Euromark, Eurolire, Euroswiss and Short Sterling, from EUREX; Euribor and Euromark, from MATIF; Euribor and Pibor, and from MEFF; Mibor.

The STIR data spans the period from 01/01/97 to 31/12/00. During this period, almost all STIR futures contracts switched from floor based trading to electronic. In addition, the following instruments and dates relate to minimum tick size changes: LIFFE Euromark(20/01/98), EUREX Euromark(20/01/98), MATIF Pibor(27/02/98) and MATIF Euribor(14/09/99). In the first three instances a half-point price increment was introduced in place of the preceding whole point price increments, while MATIF Euribor progressed from a minimum tick size of 0.005 to one of 0.002. For analysing bid-ask spreads, I use only the front month STIR futures contracts. I use all contracts for studying price clustering.

The spot FX data consists of two samples. The first covers the period 01/08/98 to 04/09/98. The second covers 01/08/99 to 03/09/99. Although there are no explicit structural changes to adjust for, it should be borne in mind that EBS's share of global spot FX volumes grew hugely during this period. BIS(2001) estimates that between 85% and 95% of inter-bank trading was occurring over electronic trading systems by 2000. This compares with only 50% in 1998 from BIS(1998). All currency pairs used here consist of 5 significant digits throughout, as is the convention in the spot foreign exchange market. In the second sample period, all rates use the full final digit range of 0 to 9. However, in the first sample period, two rates, DEM/JPY and DEM/CHF, only use 0 and 5 in the final digit position. This makes it difficult to apply the Goodhart and Curcio(1991) price clustering pattern tests to all currency pairs and sample periods. However, I found that both these pairs exhibited pronounced attraction/resolution type clustering in their penultimate digits. This was not evident for any of the other spot FX rates. Therefore, in the price clustering analysis below, the "final-digit" patterns for DEM/JPY and DEM/CHF are actually penultimate, or fourth significant digit, patterns. Where STIR prices have minimum tick sizes of 0.005 or 0.002, the penultimate digit, which can span all ten possible digits, is similarly used. This means that all of the price clustering discussed in the context of fitting the hypotheses outlined above, is 10-digit price clustering.

In the tables below, I adopt the convention of displaying my own data in red when it comes from an order-driven regime and in blue when it comes from a quote-driven regime. Statistics quoted from other sources are displayed in black.

5.3 Results

5.3.1 EMU, Price Granularity, Non-Synchronicity and Bid-Ask Spreads

The tables below show the summary statistics used by Goodhart et al.(2002) applied to my spot FX data from EBS and to my STIR futures data from LIFFE. The fact that my spreads appear much smaller that those of Goodhart et al(2002) is primarily due to the fact that I included zero-spreads and they did not. The fact that our studies use different sample periods also contributes to these differences. However, comparing the rates of change in the EUR(DEM)/USD, it is notable that both our results show a marked increase in average percentage bid-ask spread but a very small increase in average pip bid-ask spread. Both studies indicate a big drop in volume, as approximated by the average number of trades (TRAD) and a big drop in absolute order flow(ABIM). The picture relating to EUR(DEM)/USD volatility is a lot less clear. In contrast to Goodhart et al(2002), who found that EUR(DEM)/USD volatility increased by 22%, I find that volatility for all currency pairs has actually fallen since EMU. This finding also conflicts with the majority of other studies, including Hau et al.(2000 and 2002), Galati and Tsataronis(2001), and the BIS(2001). It may simply be a feature of the particular sample periods.

The USD/JPY result provides a telling insight. It shows that the average bid-ask spread has risen by more than a quarter, but the average pip bid-ask spread exhibits no change at all. The change in the average bid-ask spread is entirely explained by a 21% fall in USD/JPY between the sample periods. The fact that the pip bid-ask spread is completely unmoved by such a large change in exchange rate suggests that the size and use of pip bid-ask spreads may be rigid.

Volume (TRAD) is lower for all exchange rates except USD/CHF which shows an increase of over 60%. Absolute order flow (ABIM) appears drastically lower except for the USD/CHF. However, absolute order flow as a percentage of volume (ABIM/TRAD) shows that USD/CHF is no exception. There is a lot less order flow generally after EMU than there was before. Volatility is also lower, very sharply in the case of the EUR(DEM)/CHF.

On the basis of previous evidence, the fall in EUR(DEM)/JPY volume since EMU could have been attributed to vehicle currencies, in that the DEM/JPY may have seen as a viable currency pair for direct trading, while the EUR/JPY may have been considered more cost effective to transact using the USD as a vehicle. However, the picture that emerges from these table presented below is that for both of these currency pairs, indirect trading via the USD was more cost effective than direct trading. That said, the saving on indirect trading the EUR/JPY was twice that for DEM/JPY.

Tables 5.3 and 5.4 shows these same statistics except ABIM applied to LIFFE STIR futures. EMU appears to have remarkably little direct impact on these markets. Two important findings emerge. First, the reduction of the Euromark minimum tick size by 50% brings an almost exact 50% drop in both the average percentage bid-ask spread and the average pip bid-ask spread. In fact, this move was reversed in absolute terms when the Euromark gives way to the Euribor, because the EUR/USD currency unit is almost exactly twice the size of the USD/DEM. However, this effect is masked in these tables by the fact that Euromark is shown denominated in DEM and Euribor in EUR. The second notable finding is that Euribor and Short Sterling volumes increase hugely when they move to electronic trading.

Hourly	USD/JPY			USD/CHF			EUR(DEM)/USD			EUR(DEM)/JPY			EUR(DEM)/CHF			Goodhart et al. : EUR(DEM)/USD		
Average	98	99	%Δ	98	99	%Δ	98	99	%Δ	98	99	%Δ	98	99	%Δ	97	99/00	%Δ
AS	0.72(0.1513)	0.91(0.1689)	26%	1.58(0.456)	1.15(0.2646)	-27%	0.42(0.0964)	0.66(0.081)	57%	1.32(0.3796)	2.09(0.5914)	58%	0.76(0.219)	0.57(0.1411)	-25%	1.62(0.4783)	2.77(1.2823)	71%
TWAS	0.81(0.2363)	0.92(0.2304)	14%	1.75(0.6984)	1.24(0.4316)	-29%	0.41(0.107)	0.65(0.1099)	59%	1.42(0.5364)	2.27(0.9489)	60%	0.77(0.398)	0.59(0.2482)	-23%	1.45(0.531)	2.67(1.4189)	84%
ASPIP	1.03(0.2082)	1.03(0.1833)	0%	2.35(0.6651)	1.73(0.4003)	-26%	0.75(0.1678)	0.69(0.0847)	-8%	1.06(0.2982)	2.5(0.7)	136%	0.63(0.1791)	0.91(0.2256)	44%	2.84(0.8339)	2.82(1.2843)	-1%
TWASPIP	1.16(0.3281)	1.04(0.2549)	-10%	2.59(1.0205)	1.87(0.6546)	-28%	0.73(0.1864)	0.69(0.1159)	-5%	1.14(0.4255)	2.72(1.1293)	139%	0.64(0.3303)	0.94(0.397)	47%	2.53(0.9274)	2.71(1.4305)	7%
TRAD	846(520)	507(305)	-40%	148(115)	240(150)	62%	1367(672)	933(465)	-32%	332(193)	117(84)	-65%	237(149)	96(56)	-59%	602.74(313.279)	289.55(146.127)	-52%
ABIM	67.18(56.4131)	30.37(24.1556)	-55%	17.59(16.5925)	15.98(13.2052)	-9%	89.65(58.7582)	33.03(25.3438)	-63%	26.98(24.0536)	16(14.5088)	-41%	22.08(20.5351)	9.9(9.0706)	-55%	51(47)	30(27)	-41%
VOLAT	0.0456(0.0264)	0.037(0.018)	-19%	0.0363(0.0211)	0.0325(0.0135)	-10%	0.0324(0.018)	0.0308(0.0126)	-5%	0.0422(0.0203)	0.0408(0.0193)	-3%	0.0186(0.0096)	0.0065(0.003)	-65%	0.0318(0.05996)	0.0387(0.058)	22%
STDSP	0.89(0.1829)	0.92(0.1846)	3%	1.29(0.2064)	1.06(0.1628)	-18%	0.49(0.1129)	0.64(0.0854)	31%	1.53(0.4376)	1.56(0.253)	2%	0.88(0.3264)	0.49(0.1479)	-44%	1.11(0.0566)	2(1.7955)	80%
STDSPPIP	1.27(0.2432)	1.04(0.2016)	-18%	1.92(0.2859)	1.6(0.2461)	-17%	0.87(0.1955)	0.68(0.0897)	-22%	1.23(0.3414)	1.87(0.2935)	52%	0.73(0.268)	0.79(0.2366)	8%	1.93(0.9888)	2.03(1.8531)	5%

Table 5.1: Summary statistics (Hourly data)

Daily	USD/JPY			USD/CHF			EUR(DEM)/USD			EUR(DEM)/JPY			EUR(DEM)/CHF			Goodhart et al. : EUR(DEM)/USD		
Average	98	99	%Δ	98	99	%Δ	98	99	%Δ	98	99	%Δ	98	99	%Δ	97	99/00	%Δ
AS	0.71(0.0951)	0.9(0.0876)	27%	1.48(0.2579)	1.09(0.0892)	-26%	0.42(0.0627)	0.65(0.0344)	55%	1.25(0.2431)	1.93(0.1997)	54%	0.74(0.1248)	0.55(0.0381)	-26%	1.63(0.2917)	2.77(0.6592)	70%
TWAS	0.81(0.1245)	0.92(0.0901)	14%	1.75(0.3446)	1.24(0.1423)	-29%	0.41(0.0556)	0.65(0.0363)	59%	1.42(0.3198)	2.27(0.356)	60%	0.77(0.1613)	0.59(0.0688)	-23%	1.44(0.2474)	2.69(0.6654)	87%
ASPIP	1.01(0.121)	1.01(0.0877)	0%	2.19(0.3591)	1.64(0.1409)	-25%	0.74(0.1058)	0.69(0.0345)	-7%	1.01(0.1845)	2.3(0.2152)	128%	0.61(0.0987)	0.88(0.0606)	44%	2.86(0.4997)	2.82(0.6264)	-1%
TWASPIP	1.16(0.1584)	1.04(0.0911)	-10%	2.6(0.4772)	1.87(0.2198)	-28%	0.73(0.0933)	0.69(0.0368)	-5%	1.14(0.2472)	2.72(0.4109)	139%	0.64(0.1302)	0.94(0.1097)	47%	2.53(0.4246)	2.74(0.6379)	8%
TRAD	8493(3197)	5108(1635)	-40%	1496(676)	2395(655)	60%	13830(3995)	9369(2155)	-32%	3323(971)	1171(472)	-65%	2375(954)	960(280)	-60%	6027(914)	2864(735)	-52%
ABIM	671.84(283.4645)	303.68(94.1602)	-55%	175.88(61.0535)	159.84(30.6236)	-9%	896.48(349.3857)	330.28(90.225)	-63%	269.84(85.3511)	160(56.4816)	-41%	220.8(92.6058)	99(28.2297)	-55%	86.6(129.009)	104.16(88.4)	20%
VOLAT	0.0456(0.0197)	0.037(0.0108)	-19%	0.0363(0.0156)	0.0325(0.0058)	-10%	0.0324(0.0132)	0.0308(0.0057)	-5%	0.0422(0.0149)	0.0408(0.0118)	-3%	0.0186(0.0068)	0.0065(0.0017)	-65%	0.3178(0.2181)	0.3888(0.293)	22%
STDSP	0.89(0.1426)	0.94(0.1202)	6%	1.31(0.1445)	1.06(0.0796)	-19%	0.5(0.093)	0.65(0.0479)	30%	1.5(0.3234)	1.59(0.1088)	6%	0.9(0.2076)	0.51(0.0649)	-43%	1.25(0.411)	2.58(1.2311)	106%
STDSPPIP	1.27(0.1795)	1.05(0.1265)	-17%	1.94(0.1845)	1.6(0.1201)	-18%	0.89(0.1591)	0.69(0.0499)	-22%	1.21(0.2452)	1.9(0.1079)	57%	0.75(0.1666)	0.81(0.1036)	8%	2.19(0.7133)	2.63(1.2682)	20%

Table 5.2: Summary statistics (Daily data)

Summary statistics for EBS spot FX with the data first aggregated into hourly(Table 5.1) and then daily(Table 5.2) units. Standard deviations are shown in parentheses. Statistics are computed in the same manner as Goodhart et al.(2002), using the same 10 hour time window (07:00 to 17:00, London time). The original Goodhart et al.(2002) findings are shown in black at the end. Since all the data comes directly from the original source (EBS), there are no missing data intervals due to system crashes as some previous studies have had to allow for. Average bid-ask spread and volatility measures are shown in basis points.

	Pre				Post			
Short Sterling		%Δ		%Δ		%Δ		%Δ
AS			1.06(0.0766)		1.04(0.0763)	-2%	1.09(0.1437)	5%
ASPIP			0.99(0.0708)		0.99(0.0722)	0%	1.03(0.1353)	4%
TWAS			1.07(0.0698)		1.05(0.0449)	-2%	1.07(0.1119)	2%
TWASPIP			1(0.0646)		1(0.0425)	0%	1.01(0.1052)	1%
TRAD			11(10)		11(10)	0%	22(24)	100%
VOLAT			0.0032(0.0038)		0.0035(0.0039)	9%	0.0032(0.0034)	-9%
STDSP			0.08(0.1583)		0.08(0.1743)	0%	0.19(0.2524)	138%
STDSPPIP			0.07(0.1471)		0.07(0.1651)	0%	0.18(0.237)	157%
Euroswiss								
AS			1.04(0.1014)		1.01(0.1036)	-3%	1.27(0.465)	26%
ASPIP			1.02(0.0996)		1(0.1024)	-2%	1.23(0.4552)	23%
TWAS			1.04(0.1008)		1.02(0.0817)	-2%	1.19(0.4693)	17%
TWASPIP			1.02(0.0991)		1.01(0.0807)	-1%	1.16(0.459)	15%
TRAD			16(16)		12(10)	-25%	11(11)	-8%
VOLAT			0.0046(0.0044)		0.0037(0.0041)	-20%	0.0045(0.0049)	22%
STDSP			0.18(0.2032)		0.13(0.1963)	-28%	0.44(0.4032)	238%
STDSPPIP			0.17(0.1998)		0.13(0.1938)	-24%	0.42(0.3932)	223%
Eurolire								
AS			1.09(0.202)					
ASPIP			1.02(0.1891)					
TWAS			1.09(0.1786)					
TWASPIP			1.03(0.1666)					
TRAD			27(24)					
VOLAT			0.0051(0.0196)					
STDSP			0.21(0.2762)					
STDSPPIP			0.2(0.2599)					
Euromark								
AS	1.03(0.0541)		0.51(0.0559)	-50%	0.51(0.0656)	0%	0.55(0.0812)	8%
ASPIP	0.99(0.0522)		0.49(0.0539)	-51%	0.49(0.0636)	0%	0.52(0.079)	6%
TWAS	1.03(0.0363)		0.52(0.0552)	-50%	0.52(0.056)	0%	0.54(0.1094)	4%
TWASPIP	1(0.0349)		0.5(0.0533)	-50%	0.5(0.0542)	0%	0.51(0.1058)	2%
TRAD	9(8)		13(10)	44%	11(9)	-15%	60(51)	445%
VOLAT	0.0031(0.0036)		0.0018(0.0019)	-42%	0.0021(0.0019)	17%	0.0024(0.0017)	14%
STDSP	0.04(0.1214)		0.05(0.0947)	25%	0.06(0.0944)	20%	0.14(0.1385)	133%
STDSPPIP	0.04(0.1172)		0.05(0.0913)	25%	0.06(0.0918)	20%	0.13(0.1335)	117%

Table 5.3: Summary statistics (Hourly data)

	Pre				Post			
Short Sterling		%Δ		%Δ		%Δ		%Δ
AS			1.05(0.0464)		1.04(0.0456)	-1%	1.1(0.0914)	6%
ASPIP			0.98(0.0423)		0.98(0.0431)	0%	1.03(0.0866)	5%
TWAS			1.07(0.0273)		1.05(0.015)	-2%	1.08(0.0524)	3%
TWASPIP			1(0.0246)		1(0.014)	0%	1.01(0.0497)	1%
TRAD			91(50)		90(51)	-1%	216(133)	140%
VOLAT			0.0031(0.0017)		0.0035(0.0015)	13%	0.0032(0.0014)	-9%
STDSP			0.18(0.1234)		0.18(0.1573)	0%	0.31(0.1931)	72%
STDSPPIP			0.17(0.1147)		0.17(0.1489)	0%	0.29(0.1817)	71%
Euroswiss								
AS			1.04(0.0444)		1.01(0.0304)	-3%	1.26(0.2195)	25%
ASPIP			1.03(0.0436)		1(0.0301)	-3%	1.22(0.2176)	22%
TWAS			1.04(0.0373)		1.02(0.0253)	-2%	1.19(0.2197)	17%
TWASPIP			1.02(0.0366)		1.01(0.0249)	-1%	1.16(0.2177)	15%
TRAD			143(87)		111(64)	-22%	113(61)	2%
VOLAT			0.0045(0.0018)		0.0036(0.0016)	-20%	0.0044(0.0021)	22%
STDSP			0.29(0.0959)		0.23(0.108)	-21%	0.61(0.2745)	165%
STDSPPIP			0.28(0.0943)		0.23(0.1066)	-18%	0.6(0.2692)	161%
Eurolire								
AS			1.08(0.0738)					
ASPIP			1.02(0.0682)					
TWAS			1.09(0.0628)					
TWASPIP			1.03(0.0559)					
TRAD			284(150)					
VOLAT			0.0049(0.0059)					
STDSP			0.34(0.1974)					
STDSPPIP			0.32(0.1855)					
Euromark								
AS	1.02(0.0256)		0.51(0.0217)	-50%	0.52(0.103)	2%	0.55(0.0447)	6%
ASPIP	0.99(0.0245)		0.49(0.021)	-51%	0.5(0.0996)	2%	0.53(0.0445)	6%
TWAS	1.03(0.0114)		0.52(0.0146)	-50%	0.53(0.0985)	2%	0.54(0.0612)	2%
TWASPIP	1(0.0107)		0.5(0.0141)	-50%	0.51(0.0952)	2%	0.52(0.0599)	2%
TRAD	89(52)		126(63)	42%	106(55)	-16%	639(349)	503%
VOLAT	0.003(0.0015)		0.0018(0.0008)	-40%	0.002(0.0007)	11%	0.0024(0.0006)	20%
STDSP	0.11(0.1067)		0.1(0.08)	-9%	0.12(0.062)	20%	0.18(0.1135)	50%
STDSPPIP	0.11(0.1031)		0.1(0.0772)	-9%	0.11(0.0602)	10%	0.17(0.1099)	55%

Table 5.4: Summary statistics (Daily data)

Summary statistics for LIFFE STIR futures with the data first aggregated into hourly(Table 5.3) and then daily(Table 5.4) units. Standard deviations are shown in parentheses. Data from order-driven trading regimes are shown in red, while data from quote-driven regimes are in blue. On the right hand side of the tables, the red column to the right of the blue column shows where the electronic market superseded the trading floor. On the left hand side, there are two columns for Euromark. The first column relates to the period 01/01/97 to 19/01/98 when this market used a minimum tick size of 1. The second column covers the period 20/01/98 to 31/12/98, when market used a minimum tick size of 0.005, ending at EMU. Average bid-ask spread and volatility measures are shown in basis points.

The following charts show the percentage volume at each bid-ask spread level (in ticks). For the spot FX data, the 1999 data is shown alongside the 1998 data. In the spot FX data, the highest bid-ask spread, 6 ticks, contains not only the observations at 6 ticks but also all those all those observed at tick values greater than 6. For the STIR futures, up to four periods are shown. The latter represent not only each side of EMU, but also the shift in minimum tick size in the Euromark and the shift to electronic trading in mid-1999.

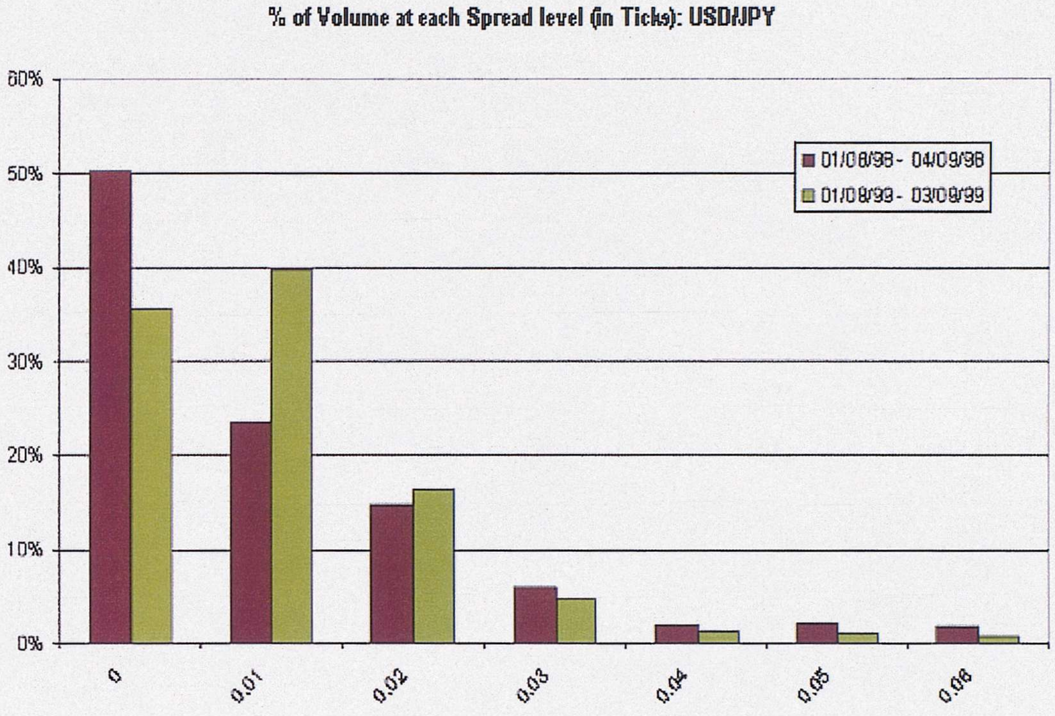

Figure 5.2: USD/JPY

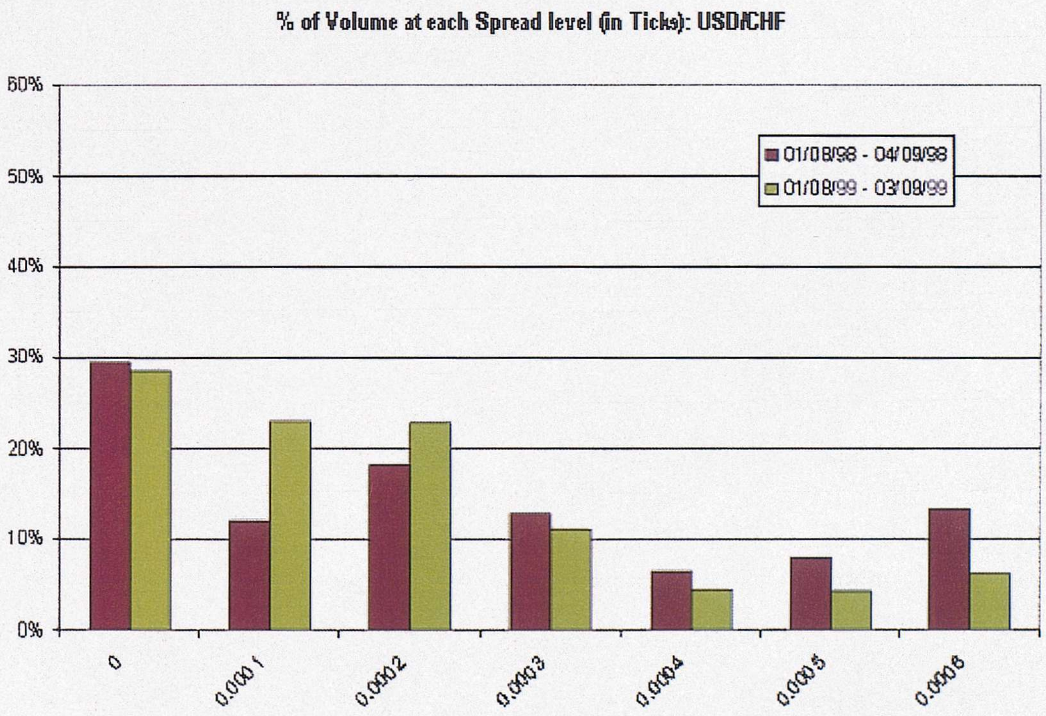

Figure 5.3: USD/CHF

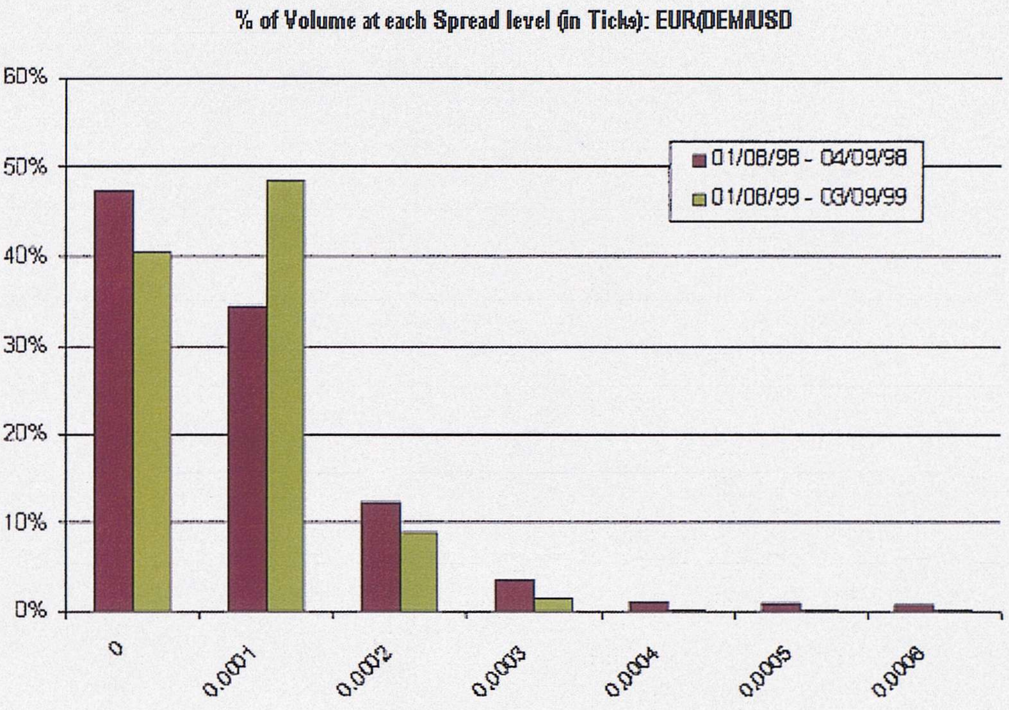

Figure 5.4: EUR(DEM)/USD

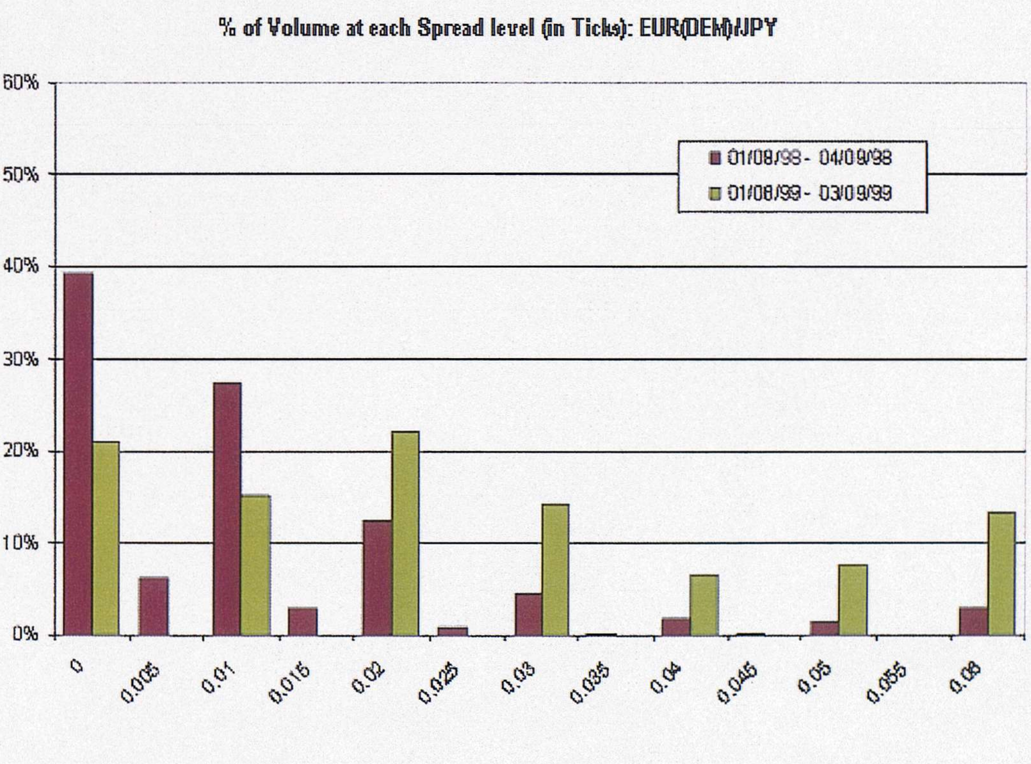

Figure 5.5: EUR(DEM)/JPY

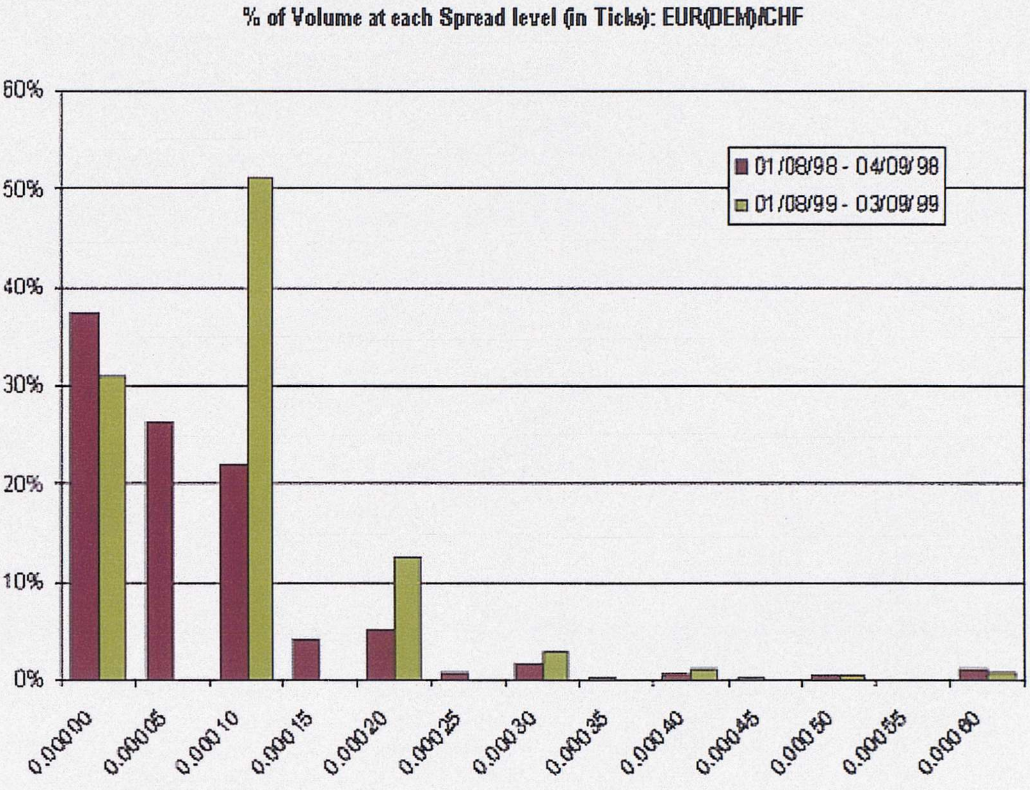

Figure 5.6: EUR(DEM)/CHF

Figure 5.2 shows that in spite of the fact that the mean USD/JPY bid-ask spread has not changed, the mode has changed. Most trades in this currency pair now execute at a bid-ask spread of 1 pip, whereas pre-EMU, most trades were done at a bid-ask spread of zero pips. Figure 5.3 shows a greater concentration of EUR/CHF trading activity in the first three ticks. After EMU, 85% of trading in this currency pair take place within 3 ticks, as opposed to 72% before EMU. Figure 5.4 shows another shift in the mode from zero to one pip, this time for the EUR(DEM)/USD. Pre-EMU, 47% of EUR(DEM)/USD trades were executed at a zero bid-ask spread compared with post-EMU, when 49% of trades coincided with a one pip bid-ask spread.

The most remarkable thing about the EUR/JPY and EUR/CHF charts is that they show clear usage of an additional final digit, while the DEM/JPY and DEM/CHF only used 0 and 5. In other words, in practice, DEM/JPY and DEM/CHF had a minimum tick of a half. EUR/CHF made more extensive use of this half-point than did EUR/JPY. This fact does not seem to have been previously noted in the empirical literature. In common with the findings of the previous paragraph, EUR(DEM)/CHF shows a mode switch from zero to one pip after EMU. Pre-EMU, there is a clear declining order of bid-ask spread tick usage for EUR(DEM)/JPY, but an apparent avoidance of half-point usage. Post-EMU, most EUR/JPY trades go through at either zero or two ticks, with what could be interpreted as an avoidance of the one-pip bid-ask spread.

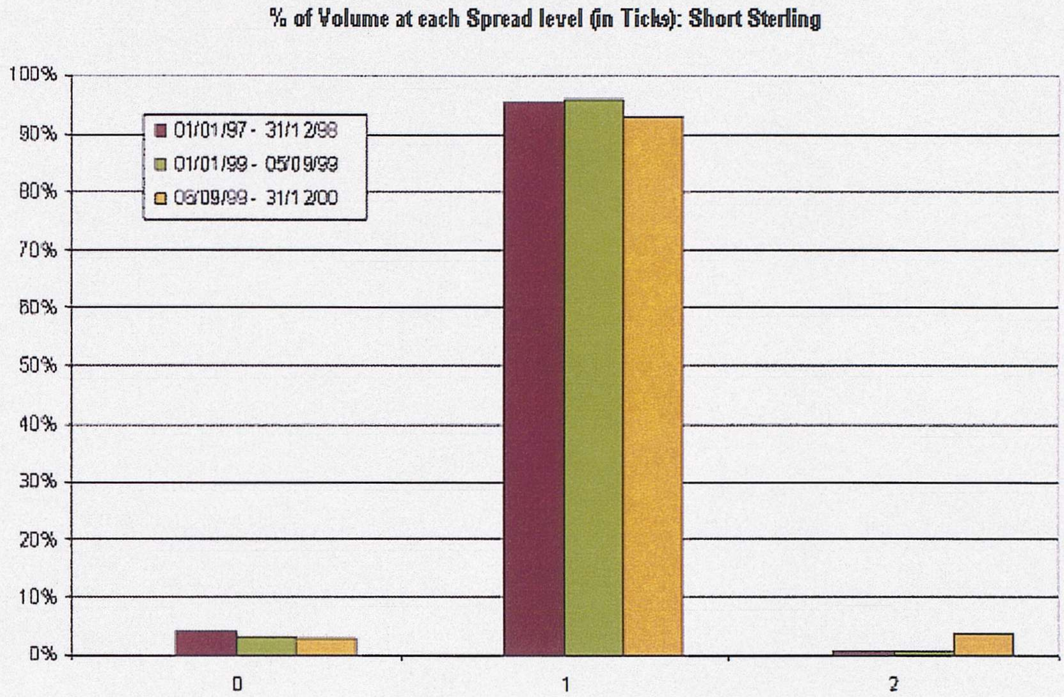

Figure 5.7: Short Sterling

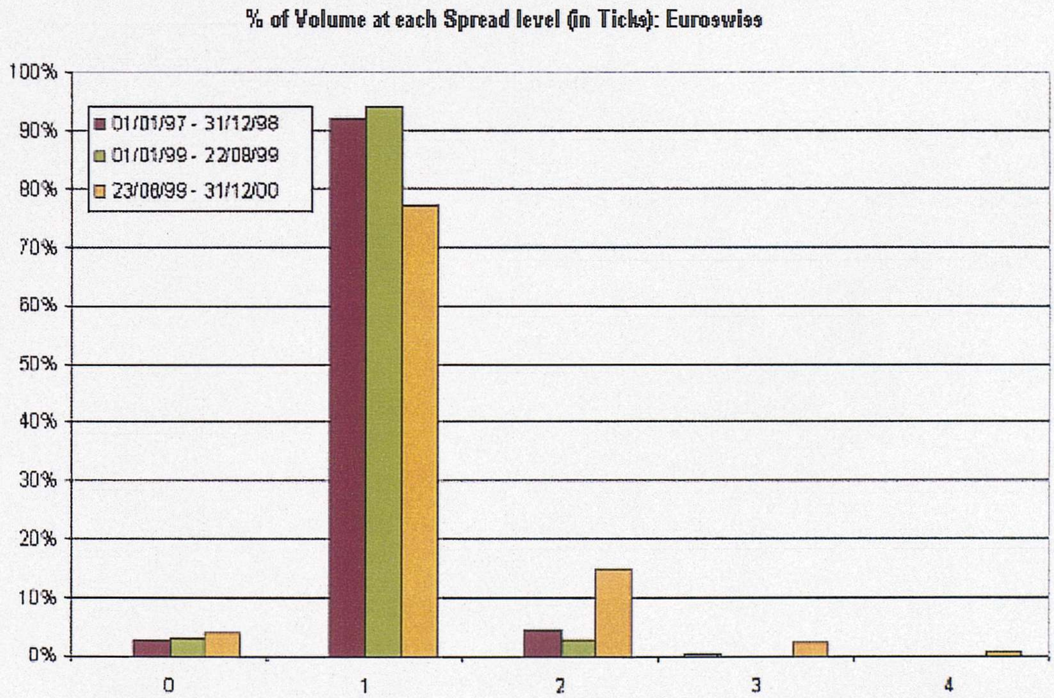

Figure 5.8: Euroswiss

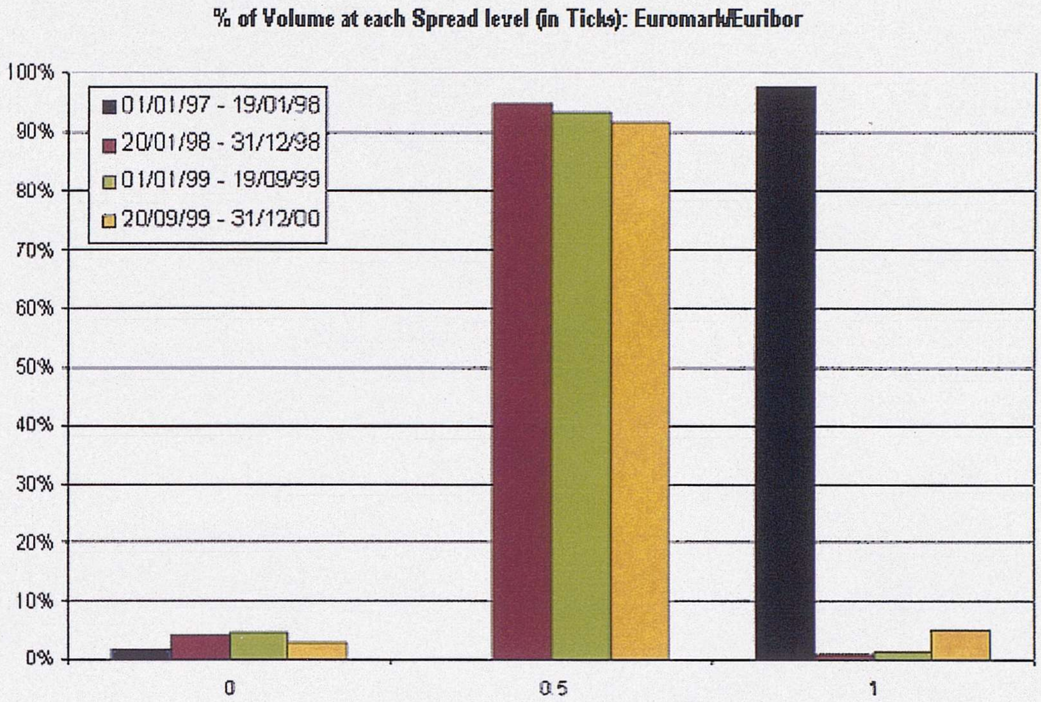

Figure 5.9: Euromark/Euribor

The STIR graphs all clearly show that almost all trading activity is transacted at the minimum tick size. In the case of Euromark/Euribor, when the minimum tick size was 0.01, almost all trades were transacted at 0.01. After the minimum tick size was dropped to 0.005, almost all trades were transacted at 0.005 ticks. Since the advent of electronic trading in these markets, some small amount of trading activity appears to go through at higher bid-ask spreads, most notably for the Euroswiss. Then again, this comparison needs to be considered carefully since quoted bid-ask spreads may not be directly comparable to limit order bid-ask spreads.

The fact that there are almost no observations at the zero-tick level, particularly in light of the spot FX results, warrants some comment. In the floor-based regime, it is conceivable that the scalpers could have such a grip on the market that they ensure that everybody pays the bid-ask spread. However, there is no obvious reason why this should hold when the markets move to electronic trading. Perhaps, the fact that volume does not reach spot FX proportions means that there is a longer time-gap than is tolerable by market order traders who would rather pay the small spread.

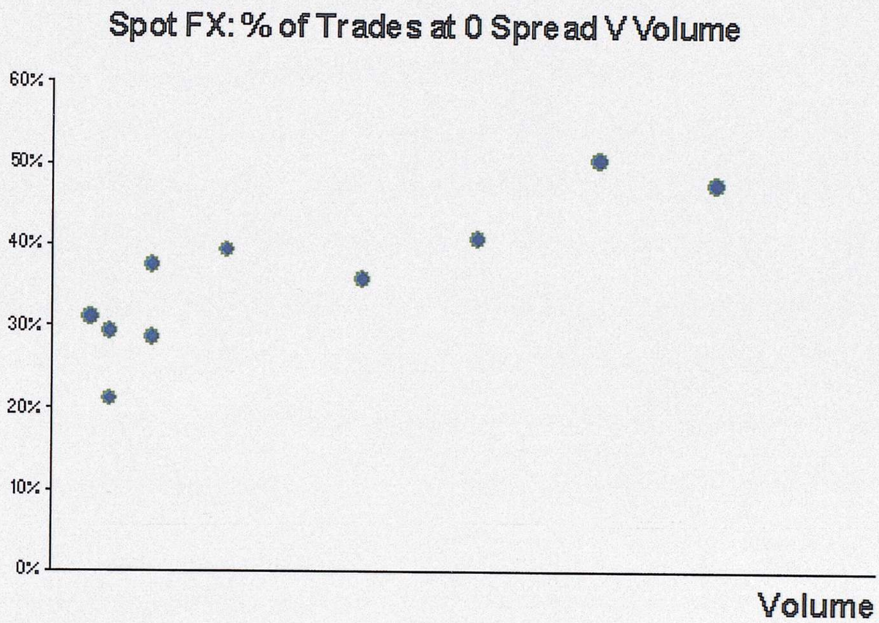

Figure 5.10: Each FX rate / sample period's percentage of trades at 0 spread plotted against its volume.

Figure 5.10 shows the average percentage of trades executed at zero spread for each currency-pair, per sample periods plotted against volume. It shows an upward trend, indicating that as volume increases, a greater proportion of trades are executed at zero-spread, as Cohen et al.(1981) predicted.

Insofar as the average spot FX inter-dealer bid-ask spread has risen, an important factor is the widespread shift in the mode from the zero-spread to the one-tick spread. However, there is a technical microstructure explanation for this – non-synchronous price revision. In an order-driven market, the bid-ask spread is deduced from successive bid and ask prices. The large fall in the number of trades from 1998 to 1999, over time segments of the same length, as shown in tables 5.1 and 5.2, means that the time gap between trades has increased. As Lo and MacKinley(1999) demonstrate, this tends to increase the average gap between successive prices. This happens because the process driving one of the price series of interest, the ask price, is the same as the process driving the other series of interest, the bid price. If the underlying process is dynamic, the likelihood that two successive sample prices drawn from the same series will be unchanged decreases as the time-gap increases. This will lead to the widening of the dispersion of values around the mean. However, this is not simply a measurement error issue. At one extreme, a time-induced price increment could drive the ask price below the bid price, or equivalently, the bid price above the ask price. Either move would turn the bid-ask spread negative, opening up an arbitrage opportunity which would be reversed quickly because, in an order driven market, these are real prices on real limit orders. At the other extreme, the time-gap induced drift simply drives the bid and ask prices apart. Consequently, an increase in the time-gap alone is sufficient to increase the average gap between the bid and ask prices, and so to increase the bid-ask spread.

Goodhart et al(2002) broadly conclude that the "price granularity" resulting from rigid "pip" bid-ask spread quoting practices combined with the re-denomination of the currency pair, is what is responsible for the sharp increase in the bid-ask spread that they observed in the EUR(DEM)/USD. My results lend support to this conclusion. The inertia observed in the pip bid-ask spread for the USD/JPY in tables 5.1 and 5.2, in spite a 20% weakening in the currency is compelling evidence of rigidity in pip quoting practices. This is supported in Figure 5.2. Similar evidence from the EUR(DEM)/USD itself, in tables 5.1, and figure 5.3, lends even more strength to this argument.

Normally, if the pip bid-ask spread is fixed, then the only way the average percentage bid-ask spread can change is through exchange rate fluctuations. However, the EUR introduces another source of change – re-denomination, and the EUR/USD has a further complication – inversion. In order to explore what effect these two factors have independent of exchange rate fluctuations, I control specifically for fluctuations. I divide the minimum tick size by the concurrent average exchange rate gives a measure of the percentage value of the minimum tick for each currency pair. However, I want to compare this to the pre-EMU percentage value of one tick for each (DEM) currency pair, but without any currency fluctuations. I achieve this by converting all my average EUR exchange rate into DEM at the official fixed conversion rate of 1.95583. For EUR/JPY and EUR/CHF, I divide the average EUR rate by 1.95583. In the case of the EUR/USD, I divided 1.95583 by the rate, because this rate is inverted compared to its predecessor. The result is a set of DEM exchange rates that have no drift in value compared to the EUR rates. I divide the minimum tick size, associated with the DEM and the EUR currency pairs respectively, by its "constant price" exchange rate. The results are displayed in the table 5.5

	98				99		
(A) EUR & DEM	Min Tick	Av. Rate*	Basis Pts		Min. Tick	Av. Rate	Basis Pts
USD/DEM	0.0001	1.8430	0.5426	EUR/USD	0.0001	1.0612	0.9423
DEM/JPY	0.005	60.957	0.8203	EUR/JPY	0.01	119.22	0.8388
DEM/CHF	0.00005	0.8181	0.6112	EUR/CHF	0.0001	1.6001	0.6250
(B) non-EUR(DEM)							
USD/JPY	0.01	112.74	0.8870	USD/JPY	0.01	112.74	0.8870
USD/CHF	0.0001	1.5075	0.6634	USD/CHF	0.0001	1.5075	0.6634

Table 5.5: Value (in basis points) of the minimum tick associated with each currency pair at constant (99) prices.

** - 1998 rates are adjusted to 1999 prices to remove fluctuations in the FX rate.*

Table 5.5(A) shows that, for the EUR(DEM)/CHF and EUR(DEM)/JPY, re-denomination of these rates from DEM to EUR increased the bid-ask spread by a mere 2%, independent of exchange rate fluctuations. If the USD/DEM been quoted as DEM/USD, whereby the constant price average FX rate would have been 0.54260 rather than 1.8430, and if traders had only used final-digits 0 and 5 in the fifth digit position, this re-denomination change would also have been 2%. However, the USD/DEM was quoted as USD/DEM and the isolated impact of re-denomination on this exchange rate is a 74% increase in the value of the bid-ask spread. Currency drift has a minor offsetting affect, as it lowers the increase to 67%.

In table 5.5(B), note that the size of the one-tick-spread for the non-EUR(DEM) USD based currency pairs is unchanged in 1999 compared with 1998, in spite of big changes in trading volumes and the market environment. This helps to justify the use above of 1998 DEM ticks with what are effectively 1999 prices and market conditions.

Goodhart et al(2002) consider whether the market practice relating to pip quoting could have been changed to add a decimal place and so facilitate a reduction in bid-ask spreads. They conclude that this could have introduced complications as other researchers found that trade size and market depth fell when policy induced smaller tick sizes were imposed (Jones and Lipson(2001), Goldstein and Kavajecz(2000)). However, Goodhart et al(2002) only considered the case of using all 10 final digits in an additional decimal place. It is evident from my data that the market had frequently used half-points in the past. The introduction of a half-point (0 and 5 only in the sixth decimal place) may have dissipated the increase in bid-ask spreads without introducing the complications outlined by Goodhart et al(2002).

The probable reason that this half-point option for the EUR may not have been considered is the spot FX market convention of having prices which consist of "5 significant digits". So, USD/JPY which has three digits to the left of the decimal point, is quoted to two decimal places. DEM/CHF had zero to the left of the decimal point and so could have five decimal places. The half-point option for the EUR/USD would have entailed the use of a sixth significant digit. By contrast, the DEM/CHF and DEM/JPY usage of half-points had occurred in the fifth decimal place. If the half-point had been introduced for the EUR/USD, the percentage one-tick bid-ask spread would have dropped by 13% according to my constant prices methodology. This is in stark contrast to the 74% increase that was actually recorded.

The question remains as to whether inversion plays a significant separate role in the steep change evident in the bid-ask spread. At first blush, one is tempted to conclude that it does not, on the grounds that the value of EUR/USD has hovered around one and inverting it will still produce a number near one. However, if values below one cause an additional digit in the fifth decimal place to kick in, the resulting bid-ask spreads could be substantially lower. This also means that the high bid-ask spread phenomenon may be intermittent, since the value of EUR/USD has dipped below parity in periods subsequent to my sample period. As the EUR/USD rate goes below 1, a fifth significant digit of "1" will go from being worth one basis point of a number beginning with 1, to being worth one-tenth of a basis point of a number beginning with 9. Alternatively, a final digit of 5, with a value of half a basis point, may be a more reasonable companion for a high first significant digit. The evidence from table 5.5(A) shows that DEM/JPY and DEM/CHF had high first significant digits of 6 and 8 respectively and both used a final digit set of 0 and 5.

One tick is also the minimum increment by which successive prices can deviate from their predecessors. Similar to the one-tick bid-ask spread, the EUR/USD minimum (one tick) price increment is also greater than the USD/DEM minimum increment by 74%. This means that a change in the tick size drives the bid-ask spread and price change volatility in the same direction. The positive effect of discreteness on volatility was predicted by Gottlieb and Kalay(1985), Harris(1990) and Osler(2003). For very similar reasons, the low-volume induced non-synchronous pricing effect discussed above will increase price change volatility, for the same reason that it increased the bid-ask spread.

5.3.2 The Level of Price Clustering Before and After EMU

Tables 5.6 and 5.7 below reveal the use of final digits in STIR trade and quotes prices, where the minimum tick size is below 1. These tables show a persistent but moderate tendency for the whole number to be used in preference to the half, i.e. 0 is preferred to 5. On 14/9/99, MATIF moved from using a minimum tick size of 0.005 to 0.002 for the 3-month Euribor future. This gives a set of five possible final digits. If this level of price resolution was unnecessarily fine, one would expect to see very little deal flow at 0.002 and 0.008. However, both of these price points exhibit significant levels of demand. Furthermore, MATIF Euribor does not show any evidence of having to move to adversely small trade sizes. Only 6% of its trades were done at the smallest volume category of 1 between 14/9/99 and 31/12/00. This compares with 8% for EUREX Euribor trades and 15% for LIFFE Euribor trades, respectively. 29% of MATIF Euribor trades consisted of 10 contracts or less. This compares with 21% on EUREX and 34% of LIFFE. These findings suggest the possibility that a minimum tick size lower than the current 0.005 for the Euribor on LIFFE and on EUREX would probably reduce bid-ask spreads and may even increase volume.

% at 0	Pre	Post	
LIFFE Euromark/Euribor	54%	52%	51%
EUREX Euromark/Euribor	55%	54%	
MATIF Pibor/Euribor	53%	53%	

Table 5.6(A): Percentage of STIR trade prices with 0 in fifth place

% at even ticks	Post	Digit
MATIF Euribor	42%	0
post 14/9/99	13%	2
	16%	4
	19%	6
	10%	8

Table 5.6(B): Breakdown (%) for MATIF's 0.002 minimum tick size for STIR futures.

% at 0	Pre	Post	
LIFFE Euromark/Euribor	51%	52%	50%

Table 5.7: Percentage of STIR quote prices with 0 in fifth place

% at 0	Pre
DEM/JPY	93%
DEM/CHF	74%

Table 5.8: % of spot FX trades at final-digit 0.

% at 0	Pre
DEM/JPY	91%
DEM/CHF	73%

Table 5.9: % of spot FX quotes at final-digit 0.

% Odd	Pre	Post	%Δ
USD/JPY	47%	48%	2%
USD/CHF	46%	47%	1%
EUR(DEM)/USD	48%	49%	2%
EUR(DEM)/JPY	48%	46%	-4%
EUR(DEM)/CHF	49%	49%	1%

Table 5.10: % at odd final digits - spot FX trade prices

% Odd	Pre	Post	%Δ
USD/JPY	48%	49%	2%
USD/CHF	47%	47%	0%
EUR(DEM)/USD	49%	50%	2%
EUR(DEM)/JPY	49%	46%	-5%
EUR(DEM)/CHF	49%	49%	-1%

Table 5.11: % at odd final digits - spot FX quote (limit order) prices.

DEM/JPY and DEM/CHF only used 0 and 5 in the fifth significant digit position. Tables 5.8 and 5.9 show that both currency pairs did not use the final-digit 5 very much, but the DEM/JPY used it even less than DEM/CHF. Tables 5.10 and 5.11 expose the odd versus even number usage among full price points in spot FX markets. There is persistent evidence that even numbers are preferred to odd. The STIR markets consistently use 50% each of even and odd final digits and so are not shown.

χ^2(No.of Obs.)	Pre		Post	
LIFFE Short Ster		300 (238,662)	322 (99,513)	1386 (351,986)
LIFFE Euroswiss		321 (189,334)	96 (47,585)	111 (98,464)
LIFFE Eurolire		313 (379,889)		
LIFFE Euromark/Euribor	112 (136,735)	553 (152,947)	703 (132,931)	3843 (1,005,060)
EUREX Euromark/Euribor	57 (7,020)	65 (3,866)	116 (60,354)	
MATIF Pibor/Euribor	138 (33,515)	537 (47,422)	408 (46,665)	121 (10,971)
MEFF Mibor		909 (117,351)		

Table 5.12: χ^2 and number of observations for STIR trade prices

χ^2(No.of Obs.)	Pre		Post	
LIFFE Short Ster		115 (537,357)	282 (223,287)	9715 (1,507,640)
LIFFE Euroswiss		219 (509,679)	129 (126,483)	347 (234,217)
LIFFE Eurolire		303 (877,031)		
LIFFE Euromark/Euribor	227 (300,577)	1211 (387,585)	622 (303,884)	20324 (4,012,947)

Table 5.13: χ^2 and number of observations for STIR quote prices

χ^2(No.of Obs.)	Pre	Post	%Δ
USD/JPY	79,729 (399,187)	15,195 (225,825)	-81%
USD/CHF	15,009 (42,952)	13,783 (72,939)	-8%
EUR(DEM)/USD	24,658 (484,006)	2,809 (310,300)	-89%
EUR(DEM)/JPY	6,025 (128,064)	12,046 (42,743)	100%
EUR(DEM)/CHF	464 (73,898)	317 (29,654)	-32%

Table 5.14: χ^2 and number of observations for spot FX trade prices

χ^2(No.of Obs.)	Pre	Post	%Δ
USD/JPY	77,074 (819,673)	17,116 (633,963)	-78%
USD/CHF	90,466 (251,624)	66,874 (399,722)	-26%
EUR(DEM)/USD	20,535 (771,614)	3,567 (540,868)	-83%
EUR(DEM)/JPY	21,846 (495,771)	103,619 (329,584)	374%
EUR(DEM)/CHF	1,516 (232,362)	1,552 (121,166)	2%

Table 5.15: χ^2 and number of observations for spot FX quote (limit order) prices.

Std. Range	Pre		Post	
LIFFE Short Ster		0.08	0.18	0.23
LIFFE Euroswiss		0.12	0.12	0.12
LIFFE Eurolire		0.09		
LIFFE Euromark/Euribor	0.10	0.19	0.26	0.18
EUREX Euromark/Euribor	0.25	0.39	0.15	
MATIF Pibor/Euribor	0.24	0.31	0.34	0.33
MEFF Mibor		0.26		

Table 5.16: Standard range for STIR trade prices.

Std. Range	Pre		Post	
LIFFE Short Ster		0.07	0.19	0.30
LIFFE Euroswiss		0.09	0.15	0.11
LIFFE Eurolire		0.09		
LIFFE Euromark/Euribor	0.09	0.18	0.23	0.22

Table 5.17: Standard range for STIR quote prices

Std. Range	Pre	Post	%Δ
USD/JPY	1.34	0.79	-41%
USD/CHF	1.81	1.34	-26%
EUR(DEM)/USD	0.71	0.34	-52%
EUR(DEM)/JPY	0.69	1.65	139%
EUR(DEM)/CHF	0.28	0.38	39%

Table 5.18: Standard range for spot FX trade prices

Std. Range	Pre	Post	%Δ
USD/JPY	0.93	0.52	-44%
USD/CHF	1.82	1.22	-33%
EUR(DEM)/USD	0.53	0.29	-44%
EUR(DEM)/JPY	0.62	1.73	177%
EUR(DEM)/CHF	0.29	0.41	44%

Table 5.19: Standard range for spot FX quote (limit order) prices.

Tables 5.12 and 5.13 show that χ^2 statistics are above the critical level (21.7) for all STIR contracts. This firmly rejects the null hypothesis of no clustering for every STIR contract and time period. For the spot FX market, the rejection of the no clustering case is even more emphatic. In both markets, the same story is borne out in both the trades and quotes data.

The standard range measure of Grossman et al(1997) uncovers much more pronounced price clustering in the spot FX markets than in the STIR futures markets. The level of price clustering seems to have fallen post-EMU, in both trade and quote data, for currency pairs using the USD, compared with those using the EUR(DEM). This may be linked to vehicle currency status, since the volumes evident in both CHF and JPY suggest there has been a shift in spot FX volume from EUR(DEM) to USD.

The following graphs reveal the precise frequency of final digit usage for the spot FX and STIR futures instrument. All graphs use trade price data. Where available, the quote price patterns proved visually indistinguishable from these graphs and so are not shown.

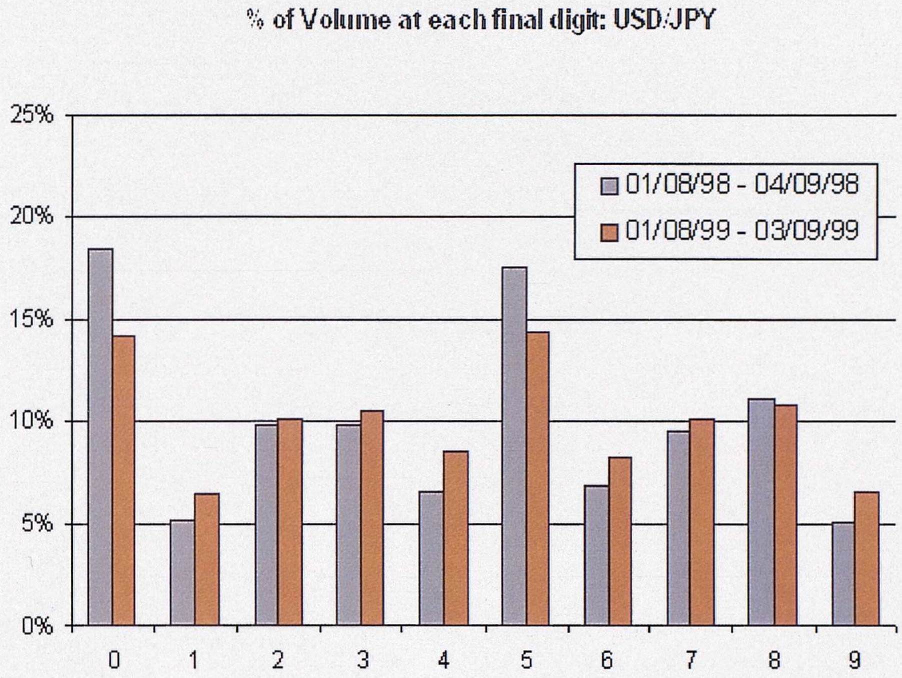

Figure 5.11: USD/JPY

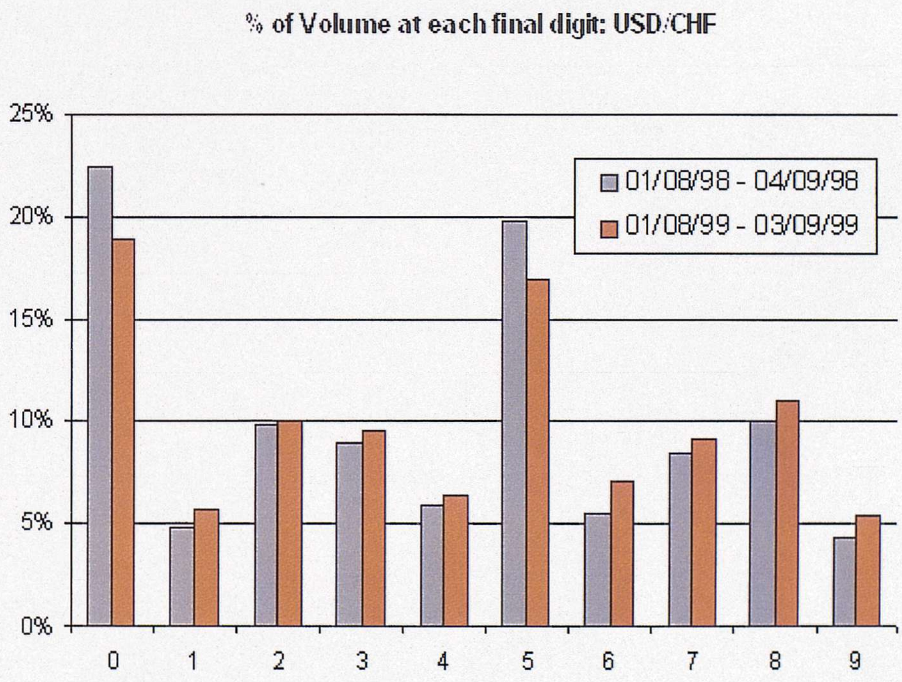

Figure 5.12: USD/CHF

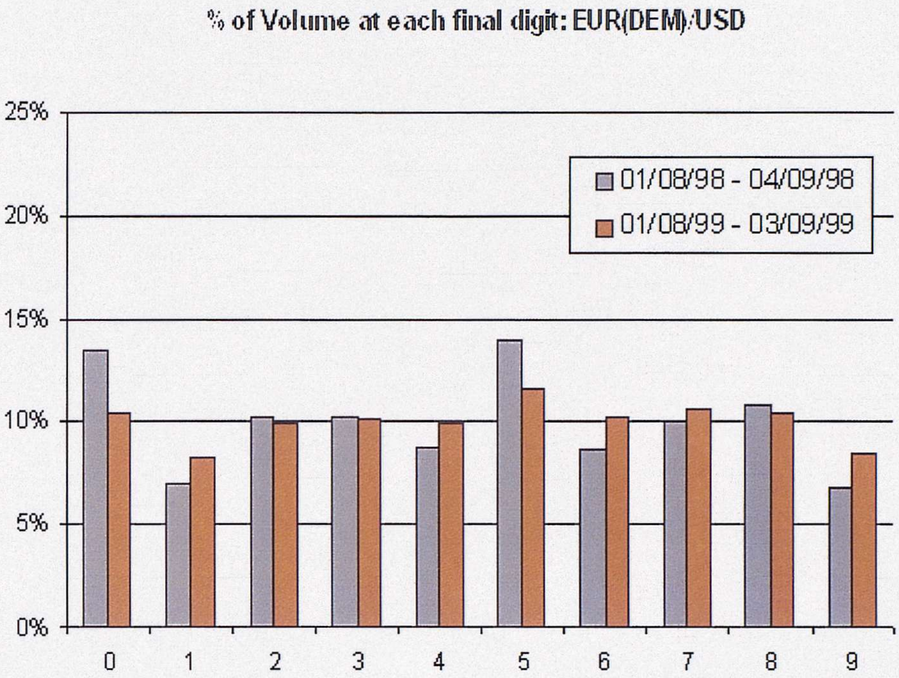

Figure 5.13: EUR(DEM)/USD

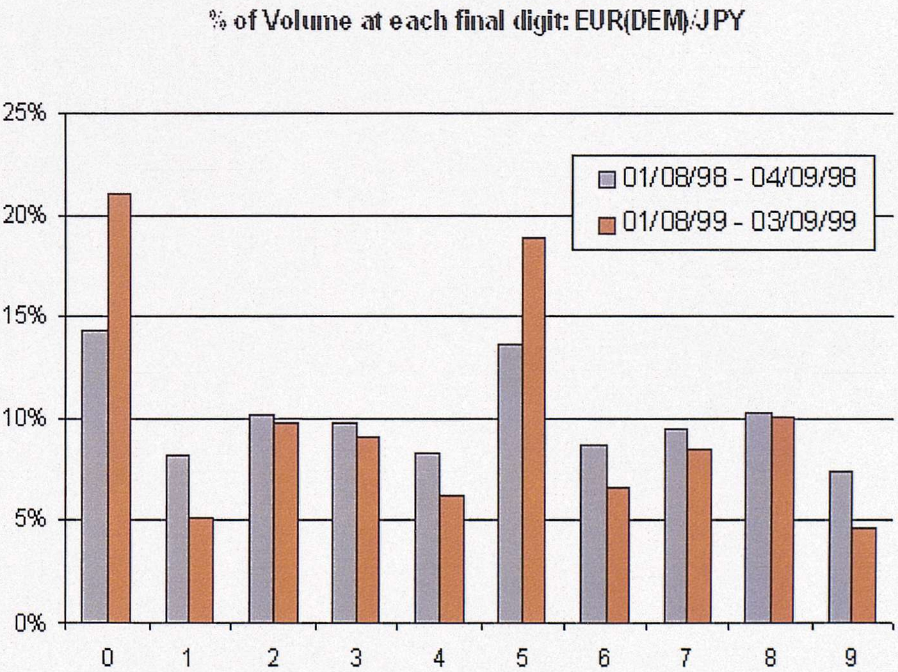

Figure 5.14: EUR(DEM)/JPY

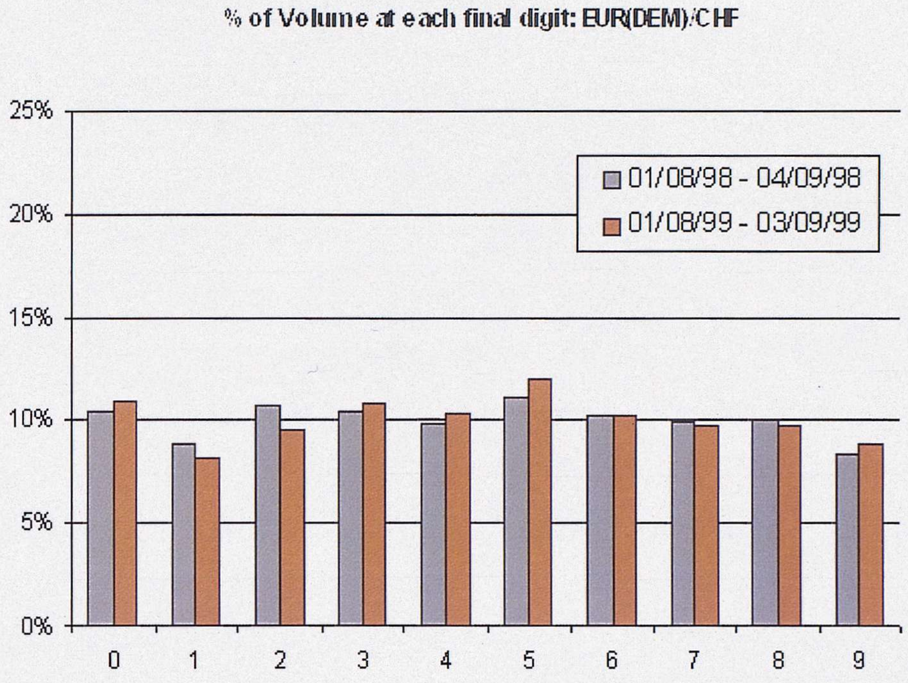

Figure 5.15: EUR(DEM)/CHF

Visual inspection of the USD/JPY final digit frequencies shows that the dominance of 0 and 5 over the rest of the numbers has declined since EMU. USD/CHF experienced a similar post-EMU decline in 0-5 dominance.
The same result is evident in EUR/USD, but here there also appears to be a reduction in the general dispersal pattern also. Conversely, EUR(DEM)/JPY exhibits an increase in 0-5 dominance and a more pronounced cluster pattern post-EMU. However, it must be remembered that in the post-EMU period, the price clustering is in the fifth significant digit position. In contrast, the pre-EMU period shows clustering in the fourth digit position and is accompanied by an extra half-point price increment and a high first position digit. The visual evidence for EUR(DEM)/CHF is weaker in both samples than for the other currency pairs. Although, for the same reasons as EUR(DEM)/JPY, price clustering is more pronounced after EMU.

5.3.3 The Price Concentration Hypothesis

Note that the Y-axis scale is smaller for the STIR graphs in figures 5.16 to 5.22 below than for the spot FX graphs above. In comparison with the spot FX results, all STIR futures in all time periods look much nearer to the uniform distribution. Unlike spot FX, 0 and 5 most clearly do not dominate the other final digits. Instead, what is apparent is a bunch of moderately higher observations in one half of the range than in the other. This is precisely the description of the penultimate digit pattern discussed in section 5.1.3 above. The fact the price concentration behaviour is evident in every single STIR contract and sample period strongly suggests that the tick size, which the exchanges have collectively decided upon, could be too high and that a finer price resolution might lower spreads and consequently, perhaps, increase volume. There is no suggestion of price concentration in the spot FX graphs above.

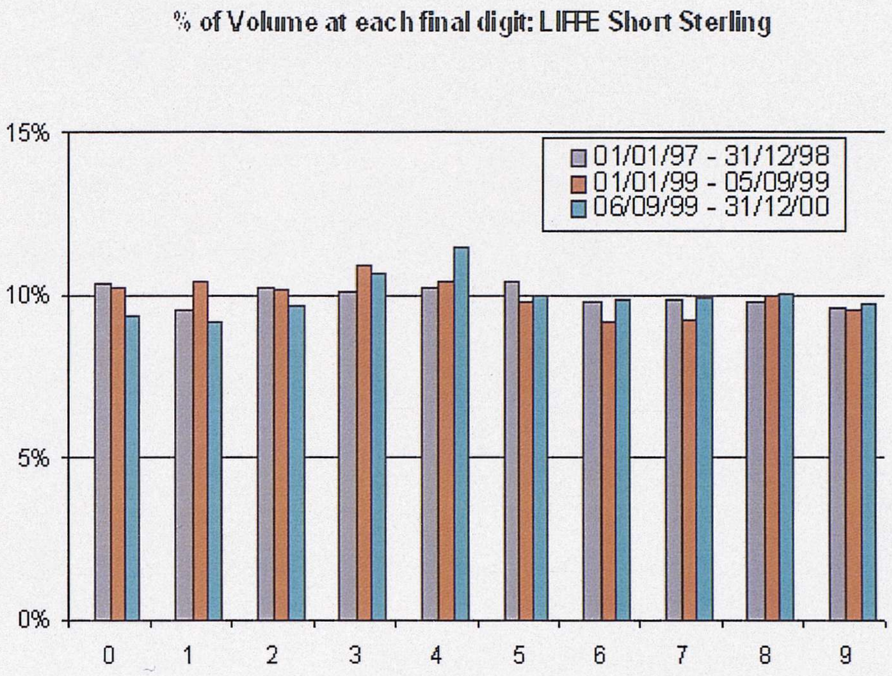

Figure 5.16: LIFFE Short Sterling

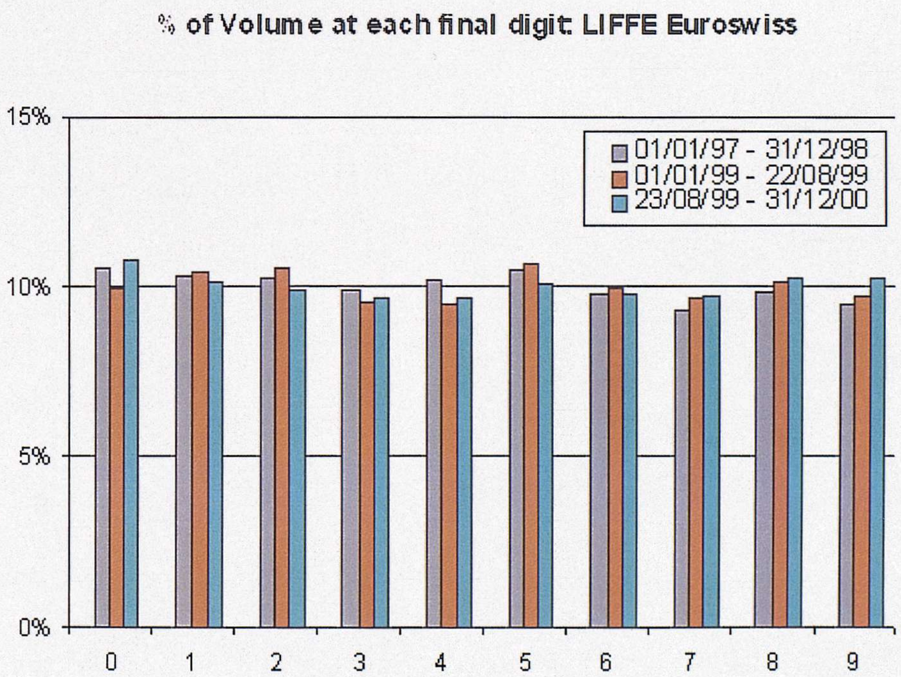

Figure 5.17: LIFFE Euroswiss

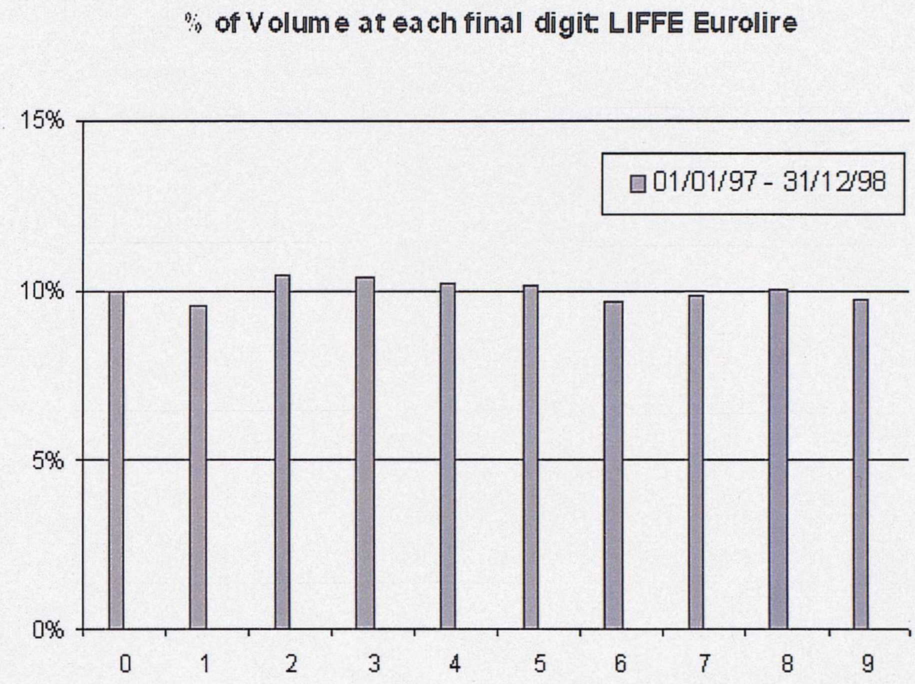

Figure 5.18: LIFFE Eurolire

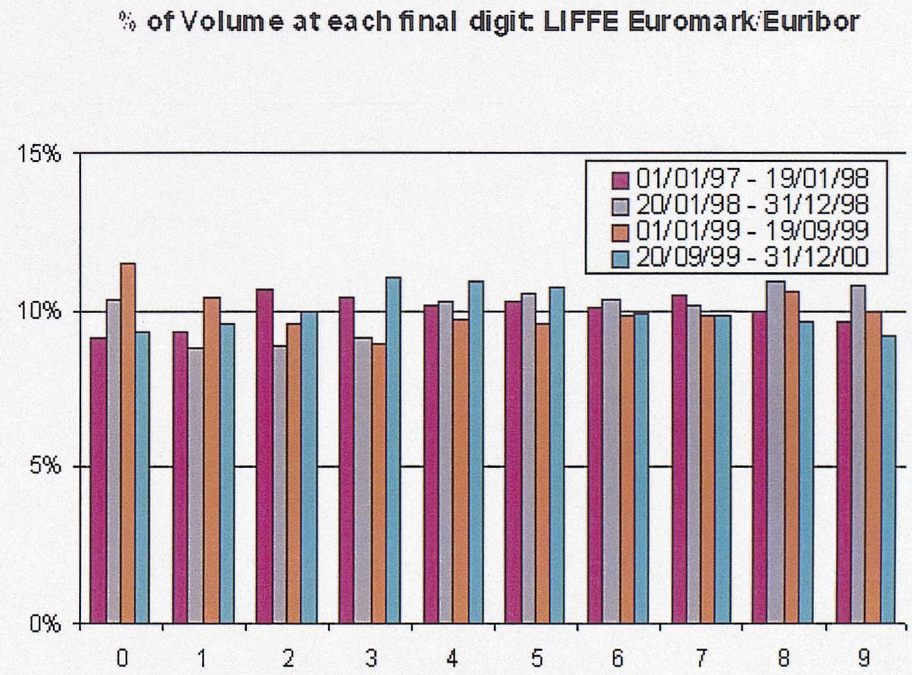

Figure 5.19: LIFFE Euromark/Euribor

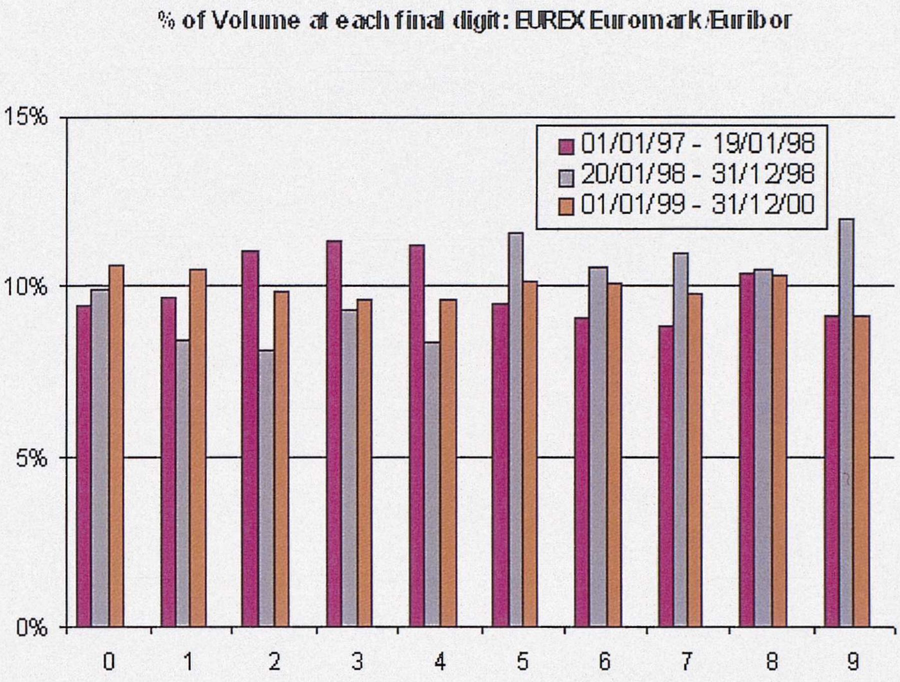

Figure 5.20: EUREX Euromark/Euribor

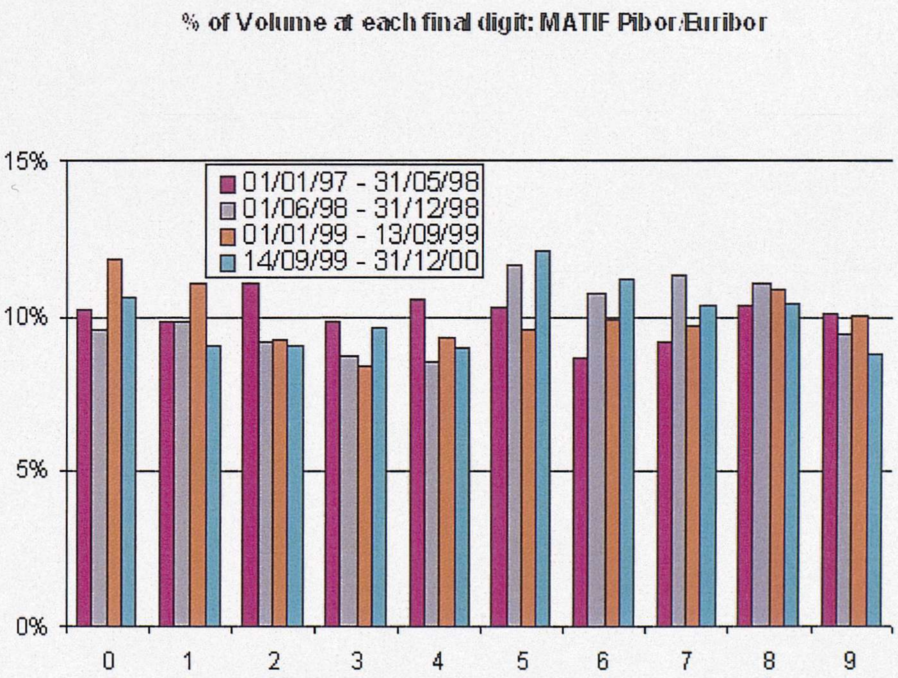

Figure 5.21: MATIF Pibor/Euribor

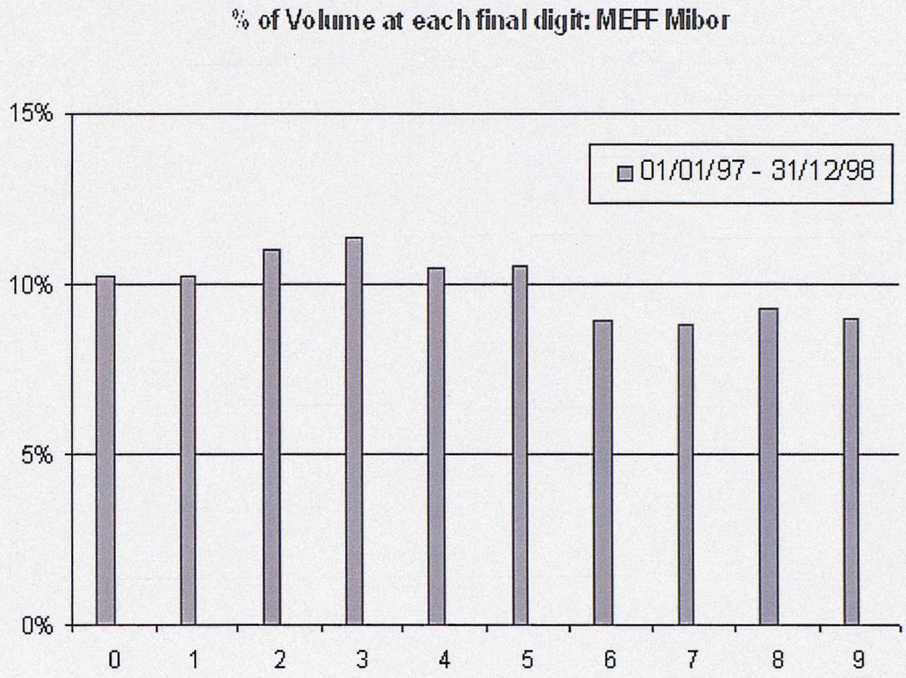

Figure 5.22: MEFF Mibor

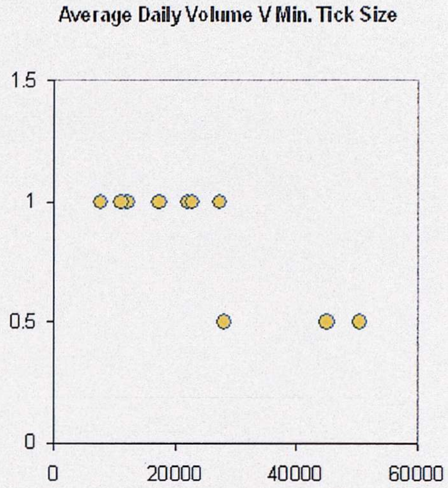

Figure 5.23(a)

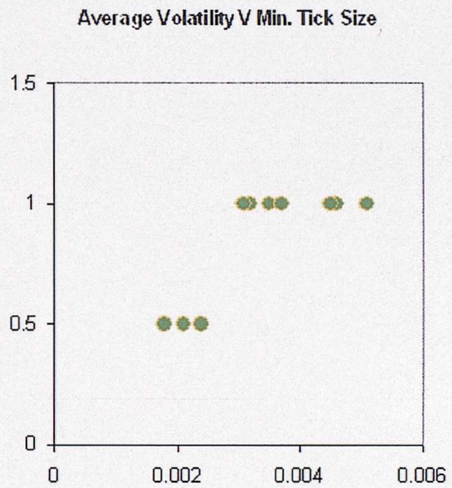

Figure 5.23(b)

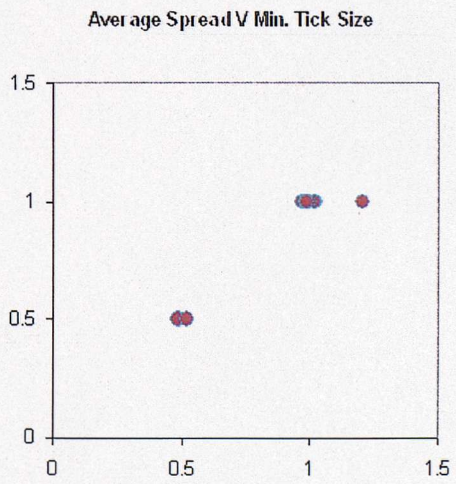

Figure 5.23(c)

Figures 5.23(a) to 5.23(c) show the minimum tick size from all LIFFE STIR contract and sample periods plotted against volume, volatility and bid-ask spread, respectively. It shows very clearly that lower tick size is associated with higher volume, lower volatility and lower bid-ask spreads. The concentrated data points in figure 5.23(c) suggest that the minimum tick size seems particularly binding on the size of the bid-ask spread.

5.3.4 The Attraction and Resolution Hypotheses

The Attraction Hypothesis was completely rejected for all instruments and sample periods and so is not displayed. The Resolution Hypothesis was rejected in all the STIR data and this also not shown. However, the Resolution Hypothesis fits the final digit usage pattern in the spot FX data. This is evident in both the trade and quotes data.

Resolution	Pre	Post	%Δ
USD/JPY	2.23	1.47*	-34%
USD/CHF	2.49	2.03	-19%
EUR(DEM)/USD	1.33*	0.50*	-62%
EUR(DEM)/JPY	1.34	1.78	33%
EUR(DEM)/CHF	0.52*	0.24*	-54%

Table 5.20: Resolution statistic - spot FX trade prices

Resolution	Pre	Post	%Δ
USD/JPY	2.06	1.60*	-22%
USD/CHF	2.27	1.97	-13%
EUR(DEM)/USD	1.48*	0.52*	-65%
EUR(DEM)/JPY	1.86	2.05	10%
EUR(DEM)/CHF	0.85*	0.57*	-32%

Table 5.21: Resolution statistic - spot FX quote (limit order) prices.

* - *these values go to zero if the restriction of volume(0) > volume(5) is enforced*

The tables above reveal that resolution type price clustering has decreased since EMU, except for EUR/JPY which has increased. In isolation, the sharp reduction of resolution-type clustering in the EUR(DEM)/USD by almost two-thirds, evident in both the trades and the quotes data, could be interpreted as counterbalancing the 74% increase in tick size. However, the large shifts apparent in the other currency pairs suggest that this is not the whole story. In the case of the EUR/CHF, the apparent 54% fall in resolution type price clustering is a decrease in a series that exhibited weak resolution characteristics to begin with. The other results tell the same story as the standard range.

5.3.5 The Cost of Market Making Hypothesis

The graphs and tables above provide little support for the cost-of-market-making hypothesis. The latter predicts a negative relationship between the price level and the level of clustering. In the context of the spot FX market with its 5 significant digit convention, all the EUR rates have effectively experienced a fall in price level. However, the associated bid-ask spread picture is mixed. EUR/USD and EUR/JPY show an increase in bid-ask spread, while EUR/CHF shows a fall. A positive association is predicted between bid-ask spreads and price clustering. While the non-CHF rates show a rise in bid-ask spreads, they show a fall in two out of the three clustering measures. The CHF rates show a fall in bid-ask spreads and a fall in clustering. Furthermore, the widespread falls in volume and in resolution-type price clustering conflict with Grossman et al(1997) and with Harris(1991). However, my findings are corroborated by Sopranzetti and Datar's(2002) recent empirical findings for the spot FX market.

5.3.6 The Negotiation Hypothesis

The data available for the spot FX market does not permit testing of the negotiation hypothesis. However, there is appropriate data to test this hypothesis for the STIR futures markets. I divided the data for each instrument and time period up into different trade sizes, using various hurdle rates. If the negotiation hypothesis is valid, larger trade sizes should exhibit less clustering than smaller trades. However, numerous different trade size classifications could not produce any consistent evidence in support of the negotiation hypothesis.

5.4 Conclusion

In both the spot FX and STIR futures markets, the influence of price discreteness on the bid-ask spread is clearly evident, not least because bid-ask spreads are so tight in these inter-dealer markets and the greatest number of trades localise around either the zero-tick and one-tick levels. This leaves little room for price clustering to play a part.

The link with price determination is less clear cut. Increases in tick size and in the time-gap between prices point to increased price changes post-EMU. Furthermore, I find tentative evidence that price clustering falls in spot FX rates in order to dissipate the positive impact of tick size on price change. In addition, the connection between price discreteness and path dependence indicated by Hausman et al.(1992) and supported in the spot FX market by the research of Rime(2000), Lyons(2001), Evans(2002), Evans and Lyons(2002) and Osler(2003) can not be ignored. Osler(2003) also reveals features of the FX market structure which would cause path dependent behaviour beyond what could be justified by economic fundamentals.

I introduce a brand new hypothesis for price clustering which I call the "price concentration" hypothesis. For the first time, I reveal the price clustering patterns evident in STIR futures data and find that this evidence supports the price concentration hypothesis. The price concentration hypothesis suggests that the minimum tick size might be set too high in the STIR futures markets. Ball and Chordia(2001) make a similar claim about the tick size of the largest NYSE stocks.

A lower minimum tick size would probably lower the average bid-ask spread for STIR futures and may well increase volume. Evidence from France's MATIF exchange, where the Euribor's minimum tick size was cut to 0.002 from 0.005, shows that additional price points are not redundant. Furthermore, there is nothing in the available evidence that would suggest any detrimental effects would result from lowering the minimum tick size of 3 month STIR futures contracts.

Using new test statistics that I developed based on their theoretical work, I confirm earlier findings of Goodhart and Curcio(1991) that clustering among spot FX rates fits the resolution pattern predicted by Ball et al(1985). I found no evidence in favour of the attraction hypothesis, the cost of market making hypothesis or the negotiation hypothesis.

My empirical findings reveal that the widespread fall in FX volume post-EMU is not ubiquitous and that it is not matched by a widespread rise in bid-ask spreads. Insofar as bid-ask spreads have risen on average, non-synchronous pricing induced by lower-volume seems the most likely contributor. However, in the specific case of the EUR(DEM)/USD bid-ask spread, I concur with the price granularity hypothesis of Goodhart et al(2002) which asserts that redenomination is the most likely cause of an increase in the spread for this exchange rate. I go on to show that, if FX traders only use five significant digits, then re-denomination alone is enough to cause a 74% increase in the one-tick EUR(DEM)/USD bid-ask spread. This increase did not occur in the other EUR FX bid-ask spreads. If the market had been willing to embrace an additional half point price increment by using 0 and 5 in the sixth digit position, the one-tick EUR/USD bid-ask spread would have been 13% lower than for USD/DEM.

Chapter 6.

Can Asymmetric Information Explain Observed Intra-Day Patterns in Bid-Ask Spreads, Price Changes, Volume and Order Flow?

6.1 Introduction

Peculiar recurring patterns have been widely observed in intra-day bid-ask spread, price change and volume data. These have been found in a wide variety of financial markets. In this chapter, I introduce a new intra-day pattern – order flow. Order flow and volume are widely acknowledged as having close links with prices and with bid-ask spreads. A theory which can explain the relationships between these intra-day patterns should be able to shed some light on what drives prices and bid-ask spreads. Many theories have been put forward to explain both the observed phenomena and the relationships between them. The most common theme across these is asymmetric information / informed trading. For the first time, I bring this group of theories together and unify them into a coherent and internally-consistent network of hypotheses. In doing so, I find that the impact of informed trading on volatility is misspecified in one of the core models and has subsequently been widely misunderstood. I develop new theory which defines the correct informed trading - volatility relationship. I propose a novel use for the correlation matrix as a means of testing multiple contemporaneous hypotheses. I examine whether these relationships have changed since EMU. Finally, I reveal how important asymmetric information really is in explaining observed intra-day patterns in the four key variables mentioned above.

The present analysis is important because both spot FX and STIR futures instruments, have previously been underrepresented in the intra-day pattern literature. This is largely because of the lack of data. FX data in particular is very hard to come by. For STIR markets, many previous studies have been done on the US markets and they do not provide volumes or quote data. The LIFFE STIR data used here, has been available in high frequency form only since the mid-1990s.

Understanding intra-day patterns is important for market participants, regulators and researchers. It can help traders identify the most/least advantageous times of the day to trade. It is important for policy makers and enforcers to understand what is happening in the market, when and why, if they are to formulate and implement effective regulation. Empiricists need to take account of seasonal effects, which can bias other analysis.

In the following sections, I review previous theoretical work and subsequently the empirical work relevant for this topic. In the following section, I draw together the various strands of theory and expound a set of broad hypotheses about market relationships. The third section addresses my empirical analysis. This includes a description of the data used here, the exact definitions of variables analysed, the intraday patterns in these data are shown and the method for evaluating the hypotheses is detailed. The final section is the conclusion.

6.1.1 Theoretical Background

As I discuss below, many empiricists have found a pronounced U-shape in the patterns of bid-ask spreads and volumes in a variety of markets. What could cause such a pattern? Two broad sets of explanations are proposed for the periodic patterns that emerge in intra-day bid-ask spreads and volumes. One argues that asymmetric information is the key. The other says it is a by-product of the market structure. The first says that patterns arise as market agents strategically optimise their trading behaviour to minimise trading costs and market impact of their trades. The second says that these patterns are incidental and occur because of longer horizon strategic behaviour of traders.

The asymmetric information case can be traced back to Admati and Pfleiderer(1988) who extend the model of Kyle(1984) to explain intra-day phenomena. Their central argument is that volume patterns emerge because informed and uninformed traders choose to trade at the same time to minimise transactions costs. Rejecting the Admati and Pfleiderer(1988) explanation as insufficient to fully explain empirical observations, Brock and Kleidon(1992) make the case that traders have different optimal holding portfolios when the market is closed from when the market is open. They argue that volumes are larger at the open and close because of portfolio rebalancing. While the Admati and Pfleiderer(1988) and Brock and Kleidon (1992) models focus on different explanations of volume, their conclusions do not actually contradict each other. The drivers identified in both models could combine in the overall result.

Brock and Kleidon(1992) predict a U-shaped intra-day pattern in bid-ask spreads if market makers have some degree of monopoly power. They argue that this is a natural response to increased order flow at the open and close. Actually, it is two different natural responses. At the open, market makers maintain high bid-ask spreads because they fear they could be adversely selected before they can get a firm estimate of the level of the true price. At the close, they maintain high bid-ask spreads in an attempt to avoid exposing themselves to the risk of holding unwanted inventory positions over the closed period. The monopoly power assumption is necessary because in a perfectly competitive market, market makers would always compete the bid-ask spread down to the minimum level.

Admati and Pfleiderer(1988) address the issue of transaction costs indirectly, in the form of the Kyle-λ, which measures market makers' price sensitivity (i.e. aversion) to order flow. Their model shows that the Kyle-λ is expected to be lower at times of high volume. It could be interpreted from this that market makers will put up bid-ask spreads when they are more averse to order flow but Admati and Pfleiderer(1988) do not make this explicit. Kyle(1984) defines λ as the inverse of market depth. Empirical work by Lee at al(1993), Ahn et al(1999) and Danielsson and Payne(2001) all show that bid-ask spreads and depth are negatively related. While it is true that other literature like Harris(1994), Jones and Lipson(2001) and Goldstein and Kavajecz(2000) find that narrower spreads coincide with less depth, those findings all relate to markets in which the price resolution or minimum tick size has changed. Absent changes in price granularity, the evidence indicates that the Kyle-λ should be positively related to bid-ask spreads, which implies that the bid-ask spread should fall as volume rises. For markets that close overnight, this means that a U-shaped volume pattern should be accompanied by and inverted-U-shaped spread pattern. Brock and Kleidon(1992) take issue with this prediction because it does not match empirical observations.

Subrahmanyam(1991) extends the Admati and Pfleiderer(1988) model by allowing informed traders, who had been risk neutral, to become risk averse and so, enables high volumes and high bid-ask spreads to co-exist. However, in so doing, he loses motivation for the volumes to bunch together in the first place. This is because risk averse informed traders would trade more during high volume periods than risk neutral informed traders. Increased informed trading increases adverse selection risk causing high volume trading costs to rise with the result that discretionary liquidity traders no longer wish to trade alongside informed traders.

The conventional measure of price change in intra-day market microstructure studies is price change volatility or, more precisely, the across-day variance of successive price changes over each time segment. This volatility measure concentrates on the magnitude of price changes associated with a particular time of day because price changes can be both positive and negative. So, high price change volatility in a particular time segment reveals the presence of extreme price moves at that time of day. Admati and Pfleiderer(1988) make the case that as informed trading takes kicks in, price changes rise. Subsequently, most empirical researchers ascribe to the Admati and Pfleiderer(1988) model that volume and volatility move in the same direction. However, as I discuss in detail below, this interpretation of their model is questionable.

Other researchers contend that asymmetric information explains the observed relationship of volume and volatility. Copeland(1976) and Jennings, Starks and Fellingham(1981) both develop models based on sequential information arrival. In these models, an individual trader receives a signal ahead of the market and trades on it, thereby creating volume and moving price (=creating volatility). Hence, volatility and volume move in the same direction. In a supporting empirical contribution, French and Roll(1986) consider three alternative explanations for the observed positive relationship between volume and volatility. The three alternative explanations are summarised as follows: (1) Relevant public announcements are made primarily during trading hours and so affect price at that time; (2) errors in pricing rise linearly with volume; and (3) some traders may be trading on private information which is either not available or can not be exploited in quite times. French and Roll's(1986) analysis focused on periods of unscheduled closure rather on the regular intra-day high-low or open-close pattern. However, that does not detract in any way from the relevance of their contribution to the present study. By ruling out the first two explanations, they conclude that informed trading must be the source of the excess volatility they observed in the periods when the market was continuously open. Ito, Lyons and Melvin(1998) applied this model to the Tokyo spot FX market over a period during which a ban on lunchtime trading was lifted. Echoing French and Roll(1986), they found that volatility doubled when trading was permitted at lunchtime.

Brock and Kleidon(1992) do not say anything about volatility. However, other researchers use non-asymmetric information arguments to explain why volatility is observed to have the same intra-day pattern as volume. The central theory here is Clark's(1973) Mixture of Distributions Hypothesis(MDH), which argues that volatility and volume move together in response to a common unobservable external stimulus, deemed to be information flow. The arrival of news pushes both volume and volatility (a measure of absolute price adjustment) in the same direction. Later researchers have elaborated on this idea. Epps and Epps(1976) link intra-day volatility to disparate opinions among traders following a price signal. Tauchen and Pitts (1983) develop the disparate opinion among traders model more formally. They propose a bivariate mixture model in which volume and price change are jointly distributed due to the presence of a latent variable. This model shows the covariance between volume and price change as zero, while the covariance between volume and price change volatility is positive, which is what has been commonly observed empirically.

Other models based on asymmetric information, which relate the bid-ask spread to volume and volatility, exist outside of the specifically intra-day pattern models. A significant contribution from Easley and O'Hara(1992) suggests volume is in itself important for price and bid-ask spread determination. In their own words "absent abnormal volume prices do not move". Their central message is that "no-trades" convey information too. They convey to market makers that an information event has probably not occurred, thereby decreasing the uncertainty of the expected price. This in turn should decrease bid-ask spreads. The upshot of this is that unanticipated volume, which monotonically reveals the level of informed trading, should be positively related to bid-ask spreads. While according to Cornell(1978), anticipated volume should be negatively related to the bid-ask spread because of economies of scales, competition among market makers and inventory management opportunities.

An important factor that I utilise below is order flow or, more specifically, the across-day volatility of order flow. Order flow is defined as buyer volume minus seller volume. I was unable to find any theoretical literature, which addressed intra-day order flow directly.

One feature to be borne in mind is that most of the above models are based in markets which can be modelled in the Kyle(1984, 1985) tradition. Specifically, there were designated market makers involved in the price setting and bid-ask spread setting processes. Even in futures markets, the open outcry system enabled dedicated traders known as "scalpers" to fulfil a market making function. Under electronic order driven systems, however, there are no designated market makers. Any trader can choose to execute his trade via a limit or a market order. As a result, price and bid-ask spread behaviour may be very different. However, for the purposes of exposition, I continue to use the term "market maker" throughout this chapter. For the moment, in the context of order driven markets, I interpret the term as "the abstract, nebulous means by which liquidity is provided to the market". This function should still cause bid-ask spreads to increase when there is a risk of adverse selection. However, inventory imbalances should be less of a problem. Furthermore, if informed traders can choose between market orders and limit orders, and do so in response to environmental conditions, it may be difficult to distinguish buys from sells in the ex-post order flow. Although the trader does not switch between being a buyer or a seller, it is the behaviour of the aggressor in a trade that determines whether the trade is a buy or a sell. This problem could make any model, which is reliant on signed order flow, difficult to evaluate.

Models specified on order driven markets are important for the present study because all of the EBS data and roughly half of the STIR data to be analysed below occur under order driven regimes. There is a general dearth of theoretical research devoted to order driven markets. However, a number of important papers have emerged in recent years. Glosten(1994) develops a model where limit order traders are uninformed and risk-neutral. Market order traders are comprised of both risk-neutral uninformed traders and risk-averse informed traders. He finds that the bid-ask spread is positively related to the level of informed trading because large orders are more likely to come from informed traders. Harris(1998) examines order placement in a variety of market conditions and finds that informed traders preference for market over limit orders is positively related to volume and negatively to bid-ask spreads. Foucault(1999) presents a sequential, dynamic, one-period model of limit and market order placement in which market participants have diverse opinions about asset valuations. He finds that the propensity to place market rather than limit orders decreases with volatility. However, in volatile periods, many of these limit orders arrive at uncompetitive prices, resulting in a high proportion of these being unfilled. He also finds that bid-ask spreads are positively related to volatility.

In summary, there are different but not always contradictory theories for observed intra-day volume patterns. On the one hand, it may be the case that traders rebalance their portfolio when switching between different market states, e.g. open and closed. On the other hand, the observed patterns may result from informed and uninformed traders trading alongside each other. Intra-day bid-ask spread patterns may arise because market makers exercise monopoly power in the face of higher volumes. Alternatively, bid-ask spreads may increase in response to unanticipated volumes and fall in response to anticipated volumes. Price change volatility may be positively related to informed trading or it may be linked to disagreement between traders about the true price. Unanticipated volume is associated with informed trading, as is volatility, but in different ways. The market making function is very different in order-driven markets, compared with quote driven ones. Most notably, traders can switch from the market side to the limit side of an order without switching their buy or sell intention. Recent work has shown that bid-ask spreads should still rise in response to informed trading under order driven regimes and also that bid-ask spreads should still rise with volatility. In addition, theorists have found that both high volatility and high bid-ask spreads cause informed traders to choose limit orders over market orders, whereas high volume has the opposite effect.

6.1.2 Previous Empirical Evidence

An overwhelming number of empirical papers have documented U-shaped patterns in intra-day data on a wide number of variables. Wang et al.(1994) noted a U-shaped pattern in bid-ask spreads for S&P500 futures. Abhyanker et al (1995) finds similar U-shaped intra-day bid-ask spreads in the FTSE100 futures. Ma et al(1992) finds a U-shaped bid-ask spread pattern for US Treasury bond futures. Franses et al(1997) find a flat distribution of bid-ask spreads over the day for the LIFFE Bund future. Ding(1999) finds the U-shaped bid-ask spread pattern in foreign exchange futures. Abhyankar et al.(1997) finds that UK equity bid-ask spreads are U-shaped over the day. As do Levin and Wright(1999) and Werner and Kleidon(1996). The latter also find that NYSE stocks have U-shaped bid-ask spreads. This is confirmed by McInish and Wood(1992), Brock and Kleidon (1992), Lee et al.(1993), Chan, Chung and Johnson(1995) and Madhavan et al.(1997). Shifting the focus to order driven markets, Chan, Christie and Schultz(1995) find Nasdaq bid-ask spreads are flat throughout the day and tail off significantly at the close. However, Lehmann and Modest(1994) find the U-shaped pattern again in equity bid-ask spreads on the Tokyo Stock Exchange. Brockman and Chung(1998) finds the same U-shape on the Hong Kong Stock Exchange. Ahn et al(1999) also find the U-shape in Hong Kong stocks. Danielsson and Payne(2001) find a W-shaped pattern in USD/DEM bid-ask spreads for the 24-hour inter-dealer spot FX market.

Many studies address the issue of intra-day volumes or number of trades. Ekman(1992) finds a U-shaped pattern in the number of trades for S&P500 futures. DeJong and Donders(1998) find the same result for Dutch AEX futures. Abhyankar et al(1995), Buckle et al(1998) and Tse(1999) all find a U-shaped intra-day volume pattern for the FTSE100 future. ap Gwilym et al(1999) also find this volume pattern for the FTSE100 future, as well as for the Short Sterling and Long Gilt futures. Gannon(1994) finds the same for the Australian All Ordinaries future. Franses et al(1997) finds U-shaped volume for the LIFFE Bund contract. Ap Gwilym et al(1996) find the same result for the LIFFE Bund, BTP and Long

Gilt contracts. ap Gwilym et al(1999) find the same U-shape for Long Gilt futures volumes. Buckle et al(1998) find the same U-shape for Short Sterling futures. Piccinato et al.(1998) confirm this U-shape in a variety of CME, LIFFE and SIMEX STIR contracts, including EuroDollar, EuroSwiss, EuroMark, Short Sterling and ECU, using trade and quote data. Jain and Joh(1988) find a U-shaped pattern in NYSE stock volumes and this finding is confirmed by Stephan and Whaley(1990), Gerety and Mulherin(1992), Lee et al(1993), Foster and Viswanathan(1993), Atkins and Basu(1995), Chan, Chung and Johnson(1995), Madhavan et al(1997) and Werner and Kleidon(1996). The latter find the same pattern for UK stocks. Abhyankar et al(1997) find an M-shaped volume pattern for UK stocks. McInish and Wood(1990) find U-shaped intra-day volumes on the Toronto Stock Exchange. In relation to order driven markets, Chan, Christie and Schultz(1995) found U-shaped volume for Nasdaq stocks. Vila and Sandmann(1995) find the same result for Nikkei futures. Lehmann and Modest(1994) find U-shaped volumes on the Tokyo Stock Exchange. Niemeyer and Sandas(1993) find a U-shaped volume pattern for Swedish stocks. Benos and Rockinger(2000) find the same pattern for French stocks. Danielsson and Payne(2001) find an M-shaped volume pattern for USD/DEM spot exchange rate on Reuters global inter-dealer FX trading platform.

Researchers have found compelling empirical evidence of the same U-shaped pattern in the intra-day volatility of price changes. Kawaller et al(1990) found this pattern for S&P 500 futures. So did Froot et al(1990), Cheung and Ng(1990), Chan et al(1991), Ekman(1992), Daigler(1997), Lee and Linn(1994), Wang et al(1994), Chang et al(1995) and Kofman and Martens(1997). The latter also found the same pattern for FTSE 100 futures. ap Gwilym et al(1999) find this U-shape in the FTSE100 future, and again in the Short Sterling and Long Gilt futures. Anderson and Bollerslev(1997) found a W-shaped pattern for S&P 500 futures. Gannon(1994) found a U-shaped pattern in Australian All-Ordinaries futures. Buckle et al(1998) found the same pattern for FTSE 100 futures. As did Lequeux(1999) and Tse(1999). Franses et al(1997) found a U-shaped pattern for the Bund futures on both LIFFE and DTB. Buckle et al(1998) found a U-shape in

Short Sterling futures price change volatility. Lequeux(1999) found the same U-shape for the Long Gilt future. So did ap Gwilym et al(1999) and Becker et al(1993). The latter found the same shape in Treasury Bond and Eurodollar futures. Daigler(1997) found the U-shaped pattern again for Treasury Bond futures. Lequeux and Acar(1996) found the same pattern in Bund and BTP futures. Kawaller et al(1994) also found this pattern in Eurodollar futures. McInish and Wood(1990) find the same pattern in US equities. Chan et al(1995) the U-shape again in NYSE equities. Gerety and Mulherin(1992) find the U-shape for the Dow-Jones. Ito and Lin(1992) find the same U pattern in S&P 500 and Nikkei 225 futures. Werner and Kleidon(1996) find that volatility for both US and UK stocks is U-shaped also. Abhyankar et al(1997) finds US stocks have U-shaped price change volatility. In order driven markets, Hiraki et al(1995) found a reverse L shape for the Nikkei futures. Lehmann and Modest(1994) find the U-shaped pattern for price change volatility on the Tokyo Stock Exchange. Chan, Christie and Schultz(1995) find the U-shape on Nasdaq. In the spot FX market, Baillie and Bollerslev(1990) find that volatility for the major currency pairs peaks twice during the day – when London and New York open. Low and Muthuswamy(1996) find peaks in price change volatility for three major currency pairs when London and New York open and close. Hseih and Kleidon(1996) and Docking et al(1999) find the same result for the USD/DEM.

Like the theoretical literature, the empirical literature on intra-day studies neglects the role of order flow.

Aside from the spot foreign exchange examples, the predominant finding from the above is that bid-ask spreads, volumes and price change volatility all exhibit a U-shaped pattern. Both bid-ask spreads and volumes appear to rise at the open and the close of the market. Such patterns might not be found in this study. The EBS spot FX data is from a 24-hour global market and so does not have an open or close per se. Although, peaks may coincide with the London/New York Open/Close, it is likely that spot FX research findings will be more relevant for EBS analysis than more general U-shaped observation. Evidence from previous

STIR futures empirical work leads me to expect a U-shaped volume pattern, at least in the floor-based trading regime, but there is little specific guidance on bid-ask spreads. Broader futures markets results suggest that both floor-based bid-ask spreads and volumes should be U-shaped.

Bessembinder(1994) tests the Easley and O'Hara(1992) idea that expected and unexpected volume affect bid-ask spreads in opposite ways on FX futures. Jorion(1996) carries out a similar exercise also on FX futures. Hartmann(1998) tests the same relationships on inter-dealer spot USD/JPY volumes. Danielsson and Payne(2001) apply the expected-unexpected split to high-frequency inter-dealer spot USD/DEM volumes. All find evidence supporting the Easley and O'Hara(1992) argument.

6.2 A Unified Theoretical Framework

The theoretical approach underlying the present analysis is based on the asymmetric information explanation for intra-day empirical regularities. The Brock and Kleidon(1992) approach would not be appropriate for analysing the spot FX market because this market is 24-hour/global and therefore does not have daily open and close events per se. Consequently, the initial premise that these researchers tackle, the U-shaped pattern, does not exist in this market. We know this from Danielsson and Payne(2001). We also know, from the same paper, that bid-ask spreads and volumes have opposing patterns, while their patterns are usually aligned in other markets.

Another justification for concentrating on the asymmetric information explanation over the market institution scenario is that inventory risk is observed to be generally much lower in both the spot FX and the futures markets, compared with the more widely studied equity markets. Lyons(2001) witnessed that a "large bank dealer in the USD/DEM market that [he] tracked in 1992 finished his trading day with no net position within 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 position and zero was only ten minutes." Manaster and Mann(1996) observed that the average trader of S&P500 index futures could reduce their inventory holdings by 50% in a single trade. By contrast, both Hasbrouck and Sofianos(1993) and Madhavan and Smidt(1993) find that it takes NYSE specialists a full week to achieve the same result. For reasons discussed above, the inventory factor should be less important under order driven regimes.

Drawing the various strands of theory together, but leaning particularly on Admati and Pfleiderer(1988), I form a number of hypotheses set out below. In doing this, I came across a problem with Admati and Pfleiderer's(1988) volume-volatility relationship. Most empirical work (e.g. Abhyankar et al(1997), Buckle et al(1998), ap Gwilym and Sutcliffe(1999)), which draws on this theory, ascribes a positive volume-volatility relationship to it. However, what Admati and Pfleiderer(1988) actually say about price change volatility is that it: A) rises at the transition point when informed trading volume kicks in, B) falls at the transition point when informed trading subsides and C) is constant everywhere else. In their own words, "When the number of informed traders is greater in the later period, [price change volatility rises]. This is because more information is revealed in the later period than in the earlier one. When the number of informed traders decreases from one period to the next, [price change volatility falls], since more information is revealed in the earlier period.". Their conclusion on the non-transition trading periods is revealed in Proposition 3, where they state that, "...the variance of price changes is the same when *n* informed traders trade in each period as it is when there is no informed trading.....With some informed traders, the market gets information earlier than it otherwise would, but the overall rate at which information comes to the market is unchanged." However, this idea that order flow has no impact on price change volatility just because it does not reduce the total amount of information is flawed. It is only true if foresight is both perfect and free, which the authors do not assume. Indeed, such an assumption would be inconsistent with any variation in the number of informed traders which is central to their model. At the core of this issue, is the conceptual relationship between order flow and the price change process. Admati and Pfleiderer(1988) define their price innovation process as a Martingale. Therefore, it follows that successive price innovations can have opposing signs. This allows order flow to counteract the prevailing price innovation and reduce the price change. As I illustrate below, the relationship between volume and volatility actually implied by this model, turns out to be the opposite of that usually inferred from the authors' findings.

Admati and Pfleiderer(1988) use the price generating process from Kyle(1984):

$$P_t = P_0 + \sum_{\tau=1}^{t} \delta_\tau + \lambda_t \omega_t$$

Equation 6.1: The Kyle price formation equation

The current price, P_t is made up of a starting value P_0, the sum of price innovations, δ, since the start and current order flow, ω_t, multiplied by a coefficient, λ, which reflects market maker aversion to order flow. λ is always positive. Order flow is driven by private information about the next price innovation δ_{t+1}, and also by the expected transaction cost, given by λ. Price changes in this model are defined as:

$$R_t = P_t - P_{t-1} = \delta_t - \lambda_{t-1}\omega_{t-1} + \lambda_t \omega_t$$

Equation 6.2: Price changes in the Kyle model

In this price change definition, the first component of the price change, δ is the price innovation for the current trading period. The second term captures earlier price movement caused by prior predictions of δ. These two terms together comprise the residual or unexploited current price innovation. The last term is a price disturbance predicting the next price innovation.

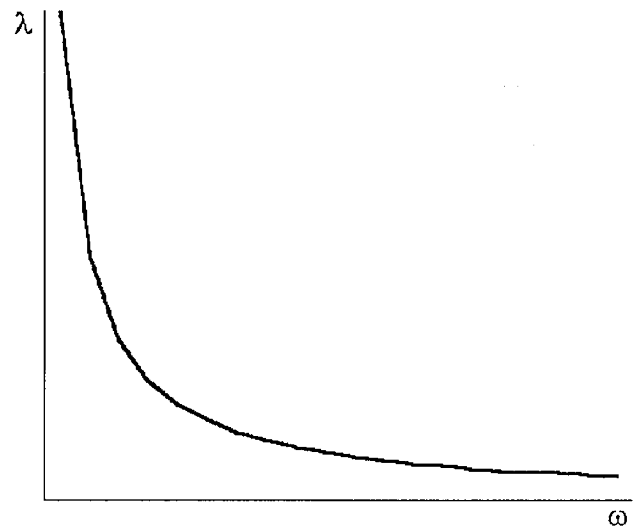

Figure 6.1: The relationship between λ and ω, at a given level of price informativeness

The Kyle(1984) framework contains an implicit relationship, which is important for understanding the relationship between volume and price change volatility - price change does not simply depend on either λ or ω, but, rather, on how these two interact. For a given price signal, λ and ω are solely and inversely determined by the expected value of the other. If λ is anticipated to rise, ω should fall, such that just enough ω is put through to exhaust the price signal. If increased volume should cause the general level of ω to rise, it is accompanied by a fall in λ, such that no more than the price signal feeds through to price change. If each fully anticipates the other, then any point, $\lambda\omega$, on the curve in the first graph above will fully exhaust the price signal and fully convey it to price change. Given this, the focus shifts to how informative the price signal actually is. If the price signal carries no information then ω is zero and price change volatility is simply the variance of δ, var(δ). At the other extreme, if foresight is perfect and costless, then, for any given λ, order flow rises to fully capture each price innovation. This causes the first two terms in the price change equation above to cancel out and the third term to equal δ_{t+1}. In other words, the price innovation is shifted in time by one period and price change volatility is again var(δ). However, var(δ) is only preserved in these extremes. As the following equation shows, order flow will erode var(δ) at all intermediate points:

$$\sigma_r^2 = \sigma^2_{\delta_t - \lambda_{t-1}\omega_{t-1}} + \sigma^2_{\lambda_t\omega_t} + 2\rho_{\delta_t - \lambda_{t-1}\omega_{t-1}, \lambda_t\omega_t}\sigma_{\delta_t - \lambda_{t-1}\omega_{t-1}}\sigma_{\lambda_t\omega_t}$$
$$= (1-\varphi)^2 \sigma^2_{\delta_t} + \varphi^2\sigma^2_{\delta_{t+1}} + 2\varphi(1-\varphi)\rho_{\delta_t,\delta_{t+1}}\sigma_{\delta_t}\sigma_{\delta_{t+1}}$$

where

σ_r^2 = Return volatility

$\sigma^2_{\delta_t - \lambda_{t-1}\omega_{t-1}}$ = Variance of unexploited δ_t

$\sigma^2_{\lambda_t\omega_t}$ = Variance of $\lambda_t\omega_t$ (= forecast part of δ_{t+1})

$\rho_{\delta_t - \lambda\omega_{t-1}, \lambda\omega_t}$ = Correlation between unexploited δ_t and $\lambda_t\omega_t$

φ = Percentage of δ_{t+1} revealed by ω_t

$\sigma^2_{\delta_t} = \sigma^2_{\delta_{t+1}}$ = Variance of δ

$\rho_{\delta_t,\delta_{t+1}}$ = 1st order autocorrelation of δ

Equation 6.3: Price change volatility with informed trading in the Kyle model

This equation shows that price change volatility is composed of the residual, or unexploited, component of δ_t (=δ-$\lambda\omega$), the forecast part of δ_{t+1} (=$\lambda\omega$), plus the interaction between the two. Since λ is constant within each regime, the variance of the forecast part of δ_{t+1} can also be modelled as the forecast part of the variance, where φ is the portion of δ_{t+1} revealed by ω_t. This relationship is exposed in the following identity:

$$\sigma^2_{\lambda_t\omega_t} \equiv \sigma^2_{\varphi\delta_{t+1}} \equiv \varphi^2\sigma^2_{\delta_{t+1}}$$

Equation 6.4: The variance of the part of δ_{t+1} revealed at time t by informed trading

Similarly, the residual part of the variance can be modelled as $(1-\varphi)^2\sigma^2_{\delta_t}$. This form of the equation reveals that the interaction part depends on the first order autocorrelation of the δ time series, $\rho_{\delta_t,\delta_{t+1}}$. However, Admati and Pfleiderer(1988) assume that δ is independent and identically distributed (IID). In other words, under the Admati and Pfleiderer's(1988) assumptions, price change volatility looks like this:

$$\sigma^2 = \sigma^2_{\delta-\lambda\omega} + \sigma^2_{\lambda\omega}$$
$$= (1-\varphi)^2\sigma^2_{\delta_t} + \varphi^2\sigma^2_{\delta_{t+1}}$$

Equation 6.5: Admati and Pfleiderer(1988) assume $\rho_{\delta_t,\delta_{t+1}}$ *is 0.*

In an example where the price signal captures 50% of the next price innovation (i.e. φ=0.5), order flow will erode 50% of the current price innovation, thus reducing that component of price change volatility. However, contemporaneous order flow relating to δ_{t+1} provides an additional source of price change volatility. As these two are uncorrelated, there is no interaction component in the aggregate price change. If var(δ) is equal to 1, price change volatility will equal 0.5^2, from the unforecast part of δ_t, plus 0.5^2, from the forecast part of δ_{t+1}, which together sum to 0.5. In fact, in the absence of autocorrelation, 50% price informativeness produces the minimum possible price change volatility of var(½δ). All other levels of informativeness nudge price change volatility towards one of the extremes of no information or full information, where it reverts to var(δ). The following graph gives a visual representation of this example:

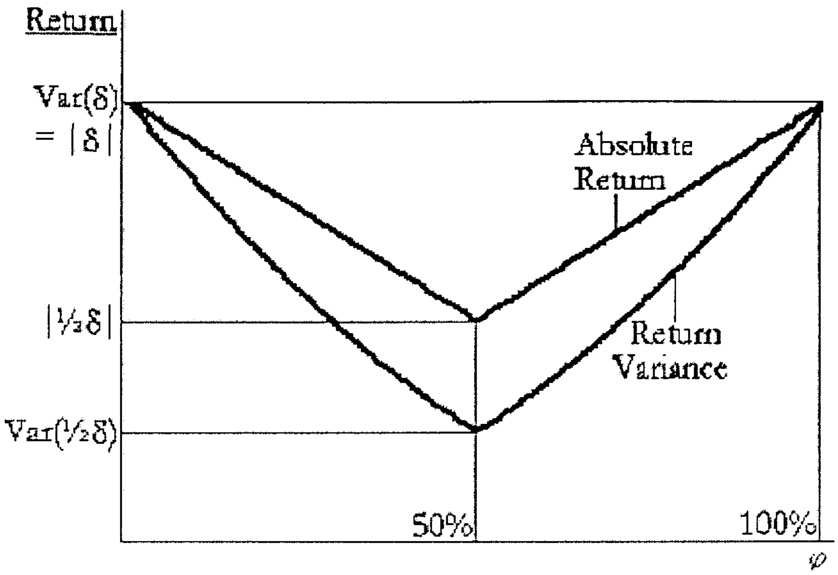

Figure 6.2: How informativeness of price signal (φ) affects price change volatility, where $\delta = 1$

Admati and Pfleiderer(1988) introduce the assumptions of a fixed information acquisition cost and only two states for λ. λ is assumed to be low during high volume periods and high during low volume periods. Information is now only acquired when it can be fully exploited, i.e. when λ is low. So, when λ is high, ω should go to zero and price change volatility reverts to var(δ). In the model, the latter should happen during low volume trading periods. Since foresight is probably less than perfect, price change volatility should fall below var(δ) when informed traders are active. Even if foresight is perfect, the cost of information acquisition will require that informed traders make a profit, in which case they will not fully exhaust the price signal, again indicating that price change volatility should be below var(δ). In reality, it is hard to envisage circumstances where price would consistently convey more than 50% of δ_{t+1}. As such, the negative volume-volatility relationship described by the left half of the curve in Figure 6.2 seems far more plausible than the positive one described by the right hand half. In any case, in the Admati and Pfleiderer(1988) model, informed traders are only active during high volume trading periods. Therefore, price change volatility is the full var(δ) in the low volume period and should fall when volume rises. This flatly contradicts both the conclusion reached by Admati and Pfleiderer(1988) and the assertions of empiricists drawing on their work. I adopt the expectation of a negative relationship between volume and volatility in the hypotheses below.

This negative volume-volatility prediction is controversial. It goes beyond just re-interpreting Admati and Pfleiderer(1988). It rejects the volume-volatility predictions of Copeland(1976) and Jennings et al(1981). However, the central point of those two papers was to show that sequential information arrival encompasses volume (= order flow) in the same direction as the price change, as part of the change process. This compares with the alternative tâtonnement process which requires no volume in order to revise price. Neither paper gives any consideration to how a future price innovation would be affected by having information about it released early, which is the key to my argument. My negative volume-volatility conclusion has another important implication. It implies that informed trading is not the additional source of exogenous volatility that French and Roll(1986) had argued. Instead, it suggests that informed trading serves to reduce exogenous volatility by dispersing and mixing price reactions to news.

The Easley and O'Hara(1992) theory poses no opposition to the negative volume-volatility prediction. Their main conclusion on the relationship between volume and price changes is that, in the absence of informed trading and unusual volumes, price change is equal to δ_t. However, unusual volume, whether motivated by information or not, will disturb price changes. They further show that price will move in the direction of whichever quote is hit. So, if an informed trader finds out that the next price innovation is downwards and initiates a sell order now, the current price will be driven down by his order flow. This will close the gap between the price now and that predicted at the end of the next trading period, reducing the price change, as predicted above.

It is clear from the preceding paragraphs that order flow, ω, is a very important factor in determining both price changes and price change volatility. As a result I include intra-day order flow in this investigation. Taking a lead from Hartmann(1998), I utilise an order flow volatility measure. The volatility measure of order flow has the benefit of neatly getting around the signed order flow problem in order driven markets that I identified earlier. The most obvious question is, how should order flow volatility and price change volatility be related? As figure 6.1 above shows, if λ is correctly anticipated, the relationship between order flow and price change depends solely on the average information content of the price signal. As the intervening paragraphs show, the relationship between their variances is similarly dependent on this information content. However, since information is assumed to carry a fixed cost, ω should be absent when λ is high, i.e. when volume is low. In the Admati and Pfleiderer(1988) framework, order flow is the result of informed trading activity and so order flow volatility is positively linked to volume.

In the process of trying to reconcile the conclusions of the various theoretical papers with each other, a deep-rooted apparent inconsistency came to light. This time it was between the Admati and Pfleiderer(1988), and Easley and O'Hara(1992) and relates to their respective conclusions on the bid-ask spread - volume relationship. Admati and Pfleiderer(1988) say that bid-ask spreads should fall as volume rises, partly because much of the volume increase will be uninformed and partly because informed traders compete with each other. Easley and O'Hara(1992) say that bid-ask spreads should rise as volume rises because volume is inversely related to the number of no-trade events. The last point appears tautological but it is not. It says that if there is no volume, there can not be any informed volume. Hence, market makers can not be adversely selected. Conversely, adverse selection risk must rise linearly with volume. Adapting Easley and O'Hara(1992) to the intra-day case means that the excess or absence of volume, relative to what is normal at that time of day, indicates the strength or weakness of a price signal. Therefore, bid-ask spreads should rise when volume is relatively high and fall when it is relatively low.

The Admati and Pfleiderer(1988) framework does not cater for this kind of variation in the price signal discussed in the preceding paragraph. Admati and Pfleiderer's(1988) focus is purely on when traders should choose to dispatch trades, given inter-temporal bid-ask spreads and constantly available price signals. The conflict can be resolved, while preserving the insights of both models, by recasting volume into two parts, expected and unexpected. This refinement means that market makers should now drop their bid-ask spreads in high volume periods partly because they *expect* a high number of informed traders. But competition among these informed traders erodes the adverse selection risk that each one would pose if acting alone. To put it another way, market makers find informed traders more tolerable in high volume periods because they are accompanied by high uninformed volumes and competition among informed traders makes their order flow less damaging. Now, in both high and low volume regimes, where trading deviates from the expected level, variations in the price signal can be inferred and bid-ask spreads can rise or fall as Easley and O'Hara(1992) predict. A hypothesised negative relationship between bid-ask spreads and expected volume supports Admati and Pfleiderer(1988). While, the Easley and O'Hara(1992) expectation is manifested in the hypothesis that bid-ask spreads and unexpected volume are positively related. Expected volume and unexpected volume should be unrelated to each other.

In order to preserve the relationships established previously within a coherent structure, it is also necessary to split order flow volatility into expected and unexpected. Expected order flow volatility should be closely and positively aligned with expected volume, since order flow should be highest during high volume periods. These two should be negatively related to price change volatility. This follows directly from the discussion of the volume-volatility relationship above, since the Admati and Pfleiderer(1988) model did not allow for variations in the price signal. Since the latter do occur in actual data, the expected value is more appropriate for testing this model. Like expected volume, expected order flow volatility should be negatively related to bid-ask spreads. This is because, according to Admati and Pfleiderer(1988), market makers find order flow more tolerable in high volume periods because it accompanies high uninformed volumes and competition among informed traders makes this order flow less damaging. Unexpected order flow volatility and expected order flow volatility should be unrelated to each other.

Both unexpected volume and unexpected order flow volatility are believed to capture deviations in the relative participation rate of informed traders and should be closely and positive aligned with each other. In some market microstructure models (e.g. Diamond and Verrecchia(1987)) the absence of informed trading activity is perceived as bad news because of restrictions on short selling. The latter does not apply here. Short selling restrictions are not believed to be a problem in either the spot FX or STIR futures markets. Therefore, in the present analysis, the level of both unexpected volume and unexpected order flow volatility are seen as indicative of the strength of the price signal. Since both unexpected volume and unexpected order flow volatility represent abnormal adverse selection risk, both should also be positively linked to bid-ask spreads.

Unexpected order flow should increase price change volatility. This is because the former is inversely related to the latter by the value of λ determined by expected order flow. Therefore, its impact on price change is larger when it rises and smaller when it falls. The association between unexpected order flow and unexpected volume means that the latter should also increase price change volatility as it rises.

The combined Admati and Pfleiderer(1988) and Easley and O'Hara(1992) model predicts that bid-ask spreads should be positively related to price change volatility. This follows because both are expected to fall as expected volume and expected order flow rise. Similarly, both are expected to rise in response to increases in unexpected volume and unexpected order flow volatility. The predicted bid-ask spread-price change volatility relationship also accords with intuition, as one expects the bid-ask spread to rise when the level of price change becomes more volatile.

As alluded to at the beginning of this section, the primary purpose of the set of hypotheses below is to explore the relationship between bid-ask spreads/price innovations and the timing of informed and uninformed trading decisions. While bid-ask spreads and price changes may be measured directly, trading volume can not easily be split into informed and uninformed. However, a number of variables that are closely associated with informed trading activity are directly measurable. These are: unexpected volume, expected order flow volatility and unexpected order flow volatility. Drawing from Easley and O'Hara(1992), unexpected volume depicts informed trading by encapsulating both its presence (positive values) and its absence (negative values). Expected order flow volatility illustrates the normal level of informed trading. Like unexpected volume, unexpected order flow volatility encapsulates both the presence and absence of informed trading. Expected and unexpected informed trading both contribute to the level of price change volatility. The remaining variable, expected volume, comprises trading from both informed and uninformed traders. The relationships between all six variables implied by the theory above are now encapsulated in the following fifteen hypotheses.

Hypothesis 6.1: The bid-ask spread is positively related to unexpected order flow volatility.

Hypothesis 6.2: The bid-ask spread is negatively related to expected order flow volatility.

Hypothesis 6.3: The bid-ask spread is positively related to unexpected volume

Hypothesis 6.4: The bid-ask spread is negatively related to expected volume

Hypothesis 6.5: The bid-ask spread is positively related to price change volatility

Hypothesis 6.6: Price change volatility is positively related to unexpected order flow volatility

Hypothesis 6.7: Price change volatility is negatively related to expected order flow volatility

Hypothesis 6.8: Price change volatility is positively related to unexpected volume

Hypothesis 6.9: Price change volatility is negatively related to expected volume

Hypothesis 6.10: Expected volume is not related to unexpected order flow volatility

Hypothesis 6.11: Expected volume is positively related to expected order flow volatility

Hypothesis 6.12: Expected volume is not related to unexpected volume. *(should hold by construction)*

Hypothesis 6.13: Unexpected volume is positively related to unexpected order flow Volatility

Hypothesis 6.14: Unexpected volume is not related to expected order flow volatility

Hypothesis 6.15: Unexpected order flow volatility is not related to expected order flow volatility *(should hold by construction)*

6.3 Empirical Analysis

	UO	EO	UV	EV	RV
BA	+	-	+	-	+
RV	+	-	+	-	
EV	0	+	0*		
UV	+	0			
EO	0*				

Key

BA – Bid-ask spread

RV – Return (Price Change) Volatility

EV – Expected Volume

UV – Unexpected Volume

EO – Expected Order Flow Volatility

UO – Unexpected Order Flow Volatility

* - *should hold by construction*

This correlation matrix representation is used to present the results for various instruments and market regimes (i.e. floor or electronic) below.

My objective here is to examine how these variables actually fit together, compared with how the theory above says they should fit together. The correlation matrix approach side-steps the whole issue of causality. ap Gwilym et al(1999) found strong evidence of bi-directional causality between volume and volatility in three LIFFE futures contracts. In addition, the matrix method enables me to evaluate all elements in this lattice of hypotheses simultaneously. If a particular relationship does not conform to that hypothesised, there are three possible explanations. First, a variable may be a poor proxy for the trading behaviour it is supposed to be linked with. The magnitude of the correlation between the three variables supposedly linked with informed trading will hopefully expose any rogue proxies for that variable. Second, the underlying behavioural premise that the theory projects may be flawed. Third, a pair of variables may be driven by an external factor in such a way that their natural relationship is overwhelmed. The pattern of relationships in the correlation matrix should help to explain what is going on. I only explore concurrent behaviour and not leading or lagging relationships, since this is what the theory above addresses. As throughout this thesis, I compare the cases pre and post EMU to explore if and how these relationships have changed.

6.3.1 Data

The data I use for this analysis are 5-minute observations sampled from my large EBS and STIR data sets. The large data sets themselves have been discussed in detail in an earlier chapter. I chose to use 5-minute samples for four reasons. First, because some instruments are less heavily traded and so, come with less data than others. 5-minute intervals should capture a good representation from all instruments that I am interested in. Second, I have so much data for some instruments that it necessary to condense it in some way in order to extract any meaningful insight from it. Third, data which is evenly spaced in time makes it easier to use conventional time series methods. Fourth, the 5-minute interval is used in many previous empirical research papers, including Buckle et al(1998), ap Gwilym et al(1996), ap Gwilym et al(1999), Payne(1997) and Andersen and Bollerslev(1998).

A correlation value is computed for each pair of factors sampled at 5-minute intervals, over the full length of each sample period. In the case of spot FX, each correlation coefficient uses around 7,000 observations. Since the STIR data spans a longer period, the numbers of observations per sample are usually greater, in spite of the fact that their trading day is shorter. The number of observations for each STIR contract sample ranges between 7,000 and 40,000.

Hypotheses involving the bid-ask spread and also involving Order Flow require quotes data. Since, only LIFFE and EBS provide quotes, only instruments from these markets are used in this analysis. Furthermore, only the front (i.e. nearest to maturity) STIR contract is used for this analysis. There are 9 instruments that have sufficient liquidity to be used in this analysis. These are, from EBS: USD/JPY, USD/CHF, EUR(DEM)/USD, EUR(DEM)/JPY, EUR(DEM)/CHF, and from LIFFE: Short Sterling, Euroswiss, Eurolire and Euribor/Euromark,

The bid-ask spread and price change volatility, are both calculated using the difference in log prices. This is the same method used by Buckle et al(1998) and gives results for bid-ask spreads similar to those used by Abhyankar et al(1997) and Werner and Kleidon(1996), for example. Bid-ask spreads use the last bid and ask prices, from series of best quote prices, in each 5-minute interval. Price change is calculated using the last trade prices in the interval. For testing, the absolute value of price change is used as a proxy for price change volatility, which is consistent with the method proposed by Andersen, Bollerslev, Diebold and Labys(2001). The other four variables, expected and unexpected order flow, and expected and unexpected volume are either in units of number of STIR futures contracts traded, or the number of individual spot FX transactions. The expected and unexpected components are derived from the total order flow and total volume series respectively. Total volume comprises the total sell volume over the interval plus the total buy volume. Total order flow is arrived at by subtracting the total buy from the total sell volumes. Similar to price change, the absolute value of this total order flow is used as a proxy for order flow volatility in time series analysis.

Since EBS do not provide volume data, the number of trades is used as a proxy. Fortunately, EBS do provide the side of each trade. On the other hand, LIFFE do not provide the side of each trade. Here, I use a method described in Stoll(2000) and also in Huang and Stoll(1997) for ascribing initiating side to trades. Where a trade price is above the preceding mid-quote, it is labelled a "buy". If it is below the mid-quote, it is labelled a "sell". When the trade price exactly equals the mid-quote, it is deemed indeterminate. In every case, the preceding bid and ask prices are the nearest preceding in time, up to one-minute old. Observations that do not have both bid and ask prices which are under one-minute old are excluded from the calculation of order flow.

In low frequency empirical analysis, expected volume is often derived using an ARIMA model (e.g. Hartmann(1998), Bessembinder(1994), Jorion(1996)). However, I follow an alternative approach, pioneered by Danielsson and Payne(2001). They found that using the repetitive intra-day volume pattern directly, which is measured as the across-day average of volume for each time segment, worked at least as well for high frequency intra-day FX data. Unexpected volume is calculated as the difference between the actual volume and this expected measure. Similarly, expected order flow volatility is computed as the across-day standard deviation of order flow per time segment. While, unexpected order flow volatility is the difference between the absolute value of actual order flow and the aforementioned expected measure. It is important to note that both unexpected variables will contain both positive and negative numbers and that, for each time segment, their across-day averages will sum to zero. For this reason, neither of the "unexpected" variables is depicted in the graphs below. Formally, the variables are defined as follows:

$$s_{i,\xi,\varepsilon,d,t} = \log(a_{i,\xi,\varepsilon,d,t}) - \log(b_{i,\xi,\varepsilon,d,t})$$

Equation 6.6: Bid-ask spread

$$\left|r_{i,\xi,\varepsilon,d,t}\right| = \left|p_{i,\xi,\varepsilon,d,t} - p_{i,\xi,\varepsilon,d,t-1}\right|$$

Equation 6.7: Return Volatility

$$EV_{i,\xi,\varepsilon,t} = \sum_{d=1}^{D} V_{i,\xi,\varepsilon,d,t} / D$$

Equation 6.8: Expected Volume

$$UV_{i,\xi,\varepsilon,d,t} = V_{i,\xi,\varepsilon,d,t} - EV_{i,\xi,\varepsilon,t}$$

Equation 6.9: Unexpected Volume

$$EO_{i,\xi,\varepsilon,t} = \sum_{d=1}^{D} \left|O_{i,\xi,\varepsilon,d,t}\right| / D$$

Equation 6.10: Expected Order Flow Volatility

$$UO_{i,\xi,\varepsilon,d,t} = \left|O_{i,\xi,\varepsilon,d,t}\right| - EO_{i,\xi,\varepsilon,t}$$

Equation 6.11: Unexpected Order Flow Volatility

In the equations above, the measures are shown with five subscripts, i which represents the instrument, ξ which denotes the regime (floor or electronic), ε which indicates pre-EMU or post-EMU, d which represents the day and t which represents the time of day (to 5-minute accuracy).

The intra-day patterns for each of these variables, across each market/regime, for pre-EMU and post-EMU periods, are shown in the graphs below. The purpose of these graphs is to show the pattern of activity throughout the day and to directly compare the pre-EMU and post-EMU periods. In all cases, the Y-axis intersects the X-axis at 0. The X-axis shows time of day. For the spot FX data, this is shown in GMT. For the LIFFE data, it is shown in UK time which adjusts from summer to winter. The spot FX intra-day data spans 24-hours, while the LIFFE data is shorter. The time period displayed for the LIFFE data was determined by the official trading period, e.g. for the Euribor it is 07:30 to 18:00, UK time. The data is sampled at 5 minute intervals, so the first data point on the graph is at 07:35. For the LIFFE data, the instruments are examined in four separate time samples. The first time-break marks the change in Euromark minimum tick size. The second coincides with EMU. The third represents the move from floor to electronic trading regimes. Splitting all the instruments into these groups means that comparisons of instruments can be made over consistent time frames. Correlation matrices showing how factors relate to each other under different regimes are shown underneath the intra-day pattern charts. I use my own convention of displaying numbers relating to electronic markets in red and numbers relating to 'floor' markets in blue.

6.3.2 Results

USD/JPY

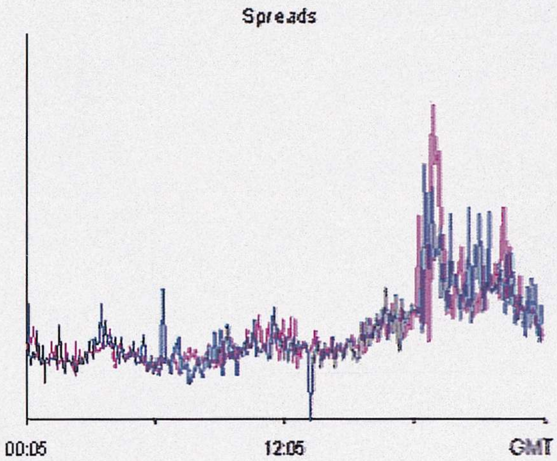

Figure 6.3

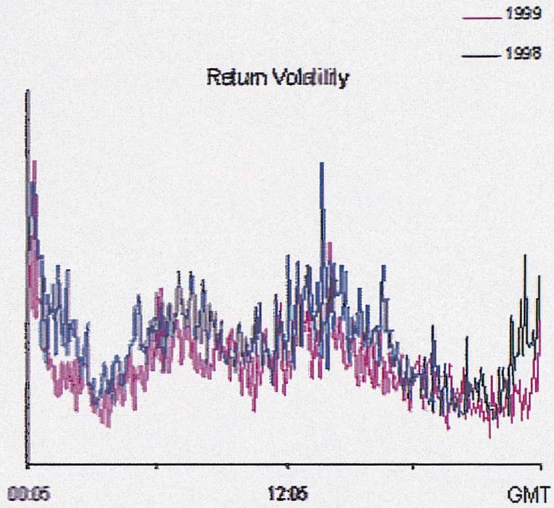

Figure 6.4

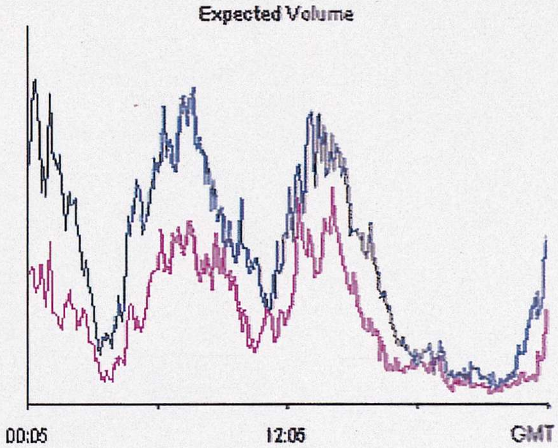

Figure 6.5

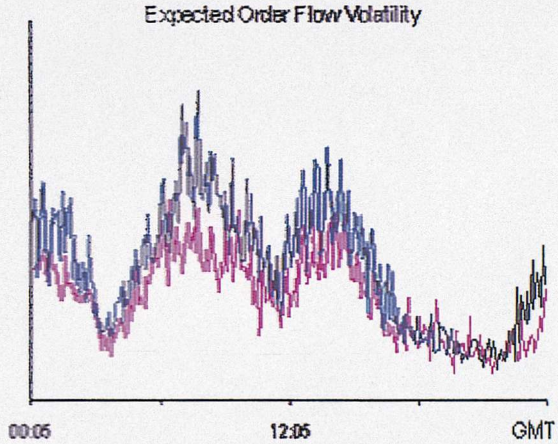

Figure 6.6

USD/JPY, pre-EMU

	UO	EO	UV	EV	RV
BA	-1%	-22%	-1%	-23%	5%
RV	35%	19%	53%	21%	
EV	0%	89%	0%		
UV	38%	0%			
EO	0%				

Table 6.1

USD/JPY, post-EMU

	UO	EO	UV	EV	RV
BA	-3%	-20%	-3%	-20%	4%
RV	39%	18%	52%	19%	
EV	0%	90%	0%		
UV	41%	0%			
EO	0%				

Table 6.2

USD/CHF

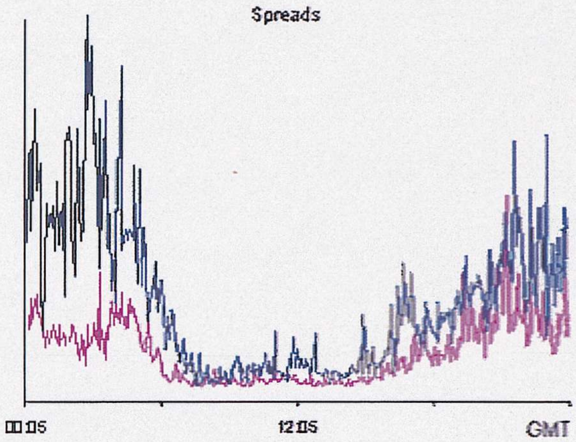

Figure 6.7

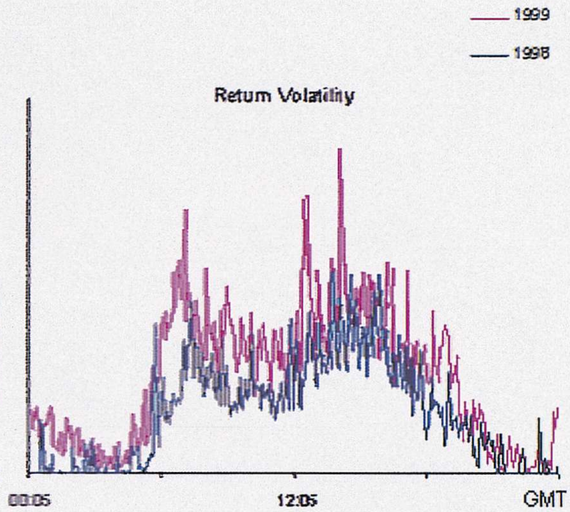

Figure 6.8

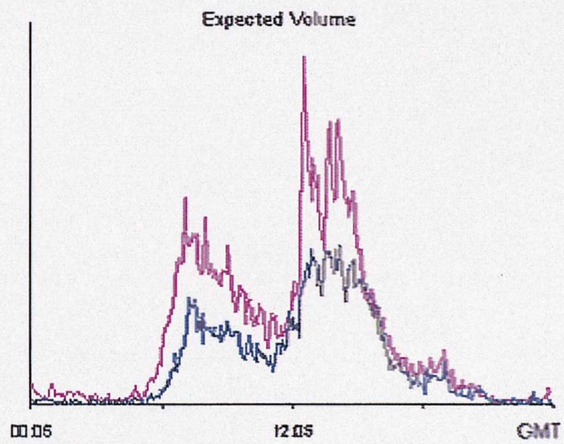

Figure 6.9

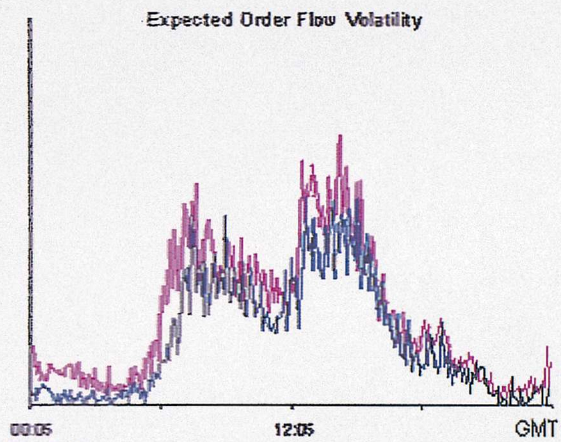

Figure 6.10

USD/CHF, pre-EMU

	UO	EO	UV	EV	RV
BA	-2%	-40%	-2%	-36%	8%
RV	33%	7%	42%	8%	
EV	0%	93%	0%		
UV	54%	0%			
EO	0%				

Table 6.3

USD/CHF, post-EMU

	UO	EO	UV	EV	RV
BA	-3%	-33%	-2%	-30%	1%
RV	37%	19%	43%	19%	
EV	0%	95%	0%		
UV	45%	0%			
EO	0%				

Table 6.4

EUR(DEM)/USD

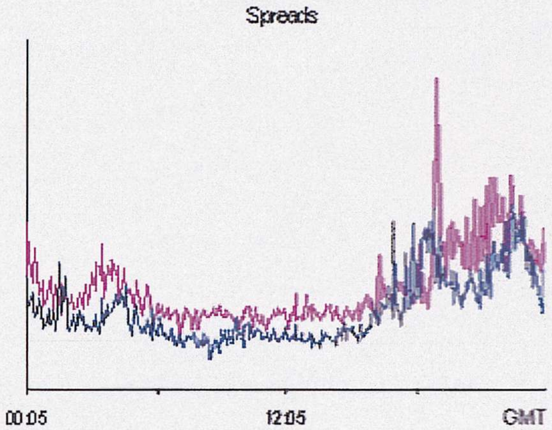

Figure 6.11

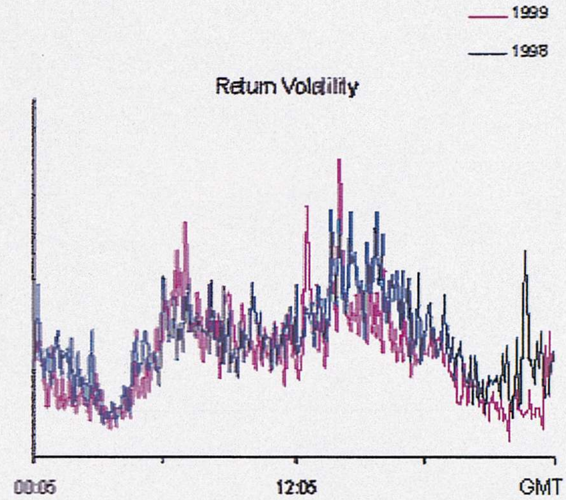

Figure 6.12

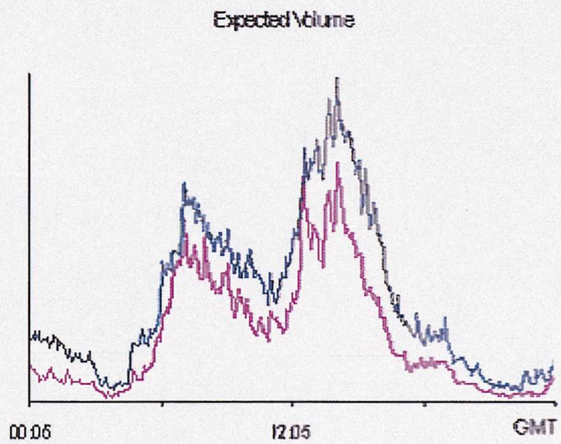

Figure 6.13

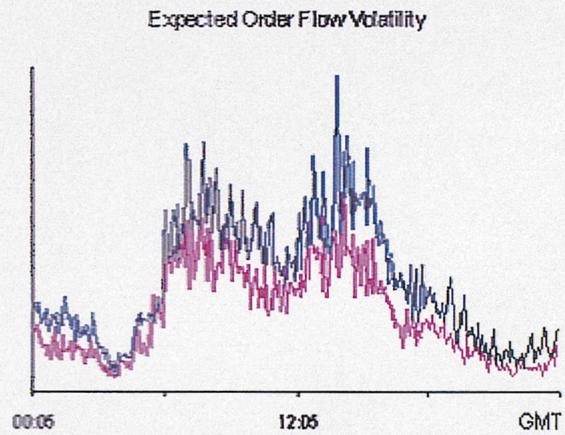

Figure 6.14

USD/DEM, pre-EMU

	UO	EO	UV	EV	RV
BA	-1%	-20%	-1%	-21%	5%
RV	35%	19%	52%	20%	
EV	0%	91%	0%		
UV	35%	0%			
EO	0%				

Table 6.5

EUR/USD, post-EMU

	UO	EO	UV	EV	RV
BA	-3%	-25%	-4%	-24%	-6%
RV	38%	29%	48%	29%	
EV	0%	93%	0%		
UV	33%	0%			
EO	0%				

Table 6.6

EUR(DEM)/JPY

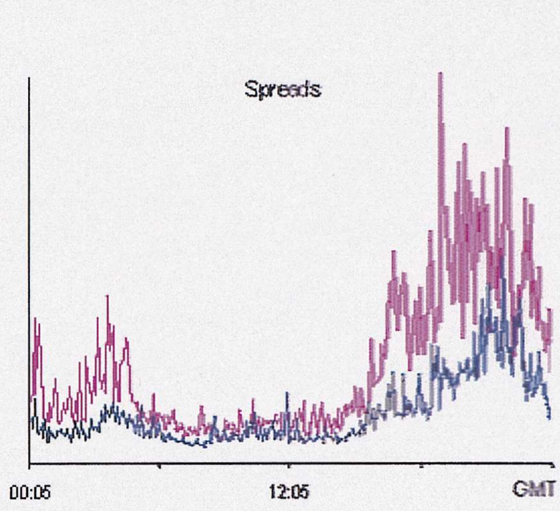

Figure 6.15

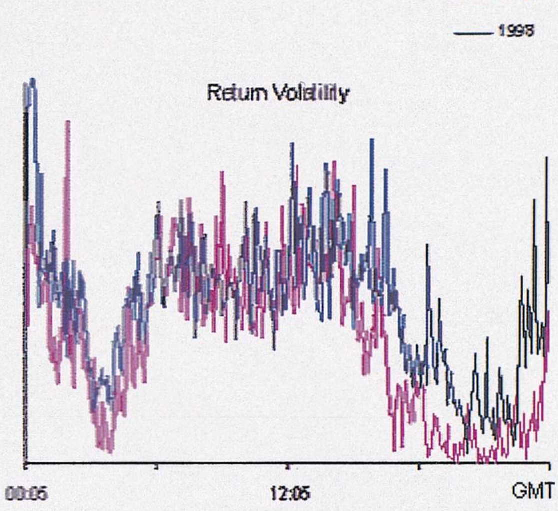

Figure 6.16

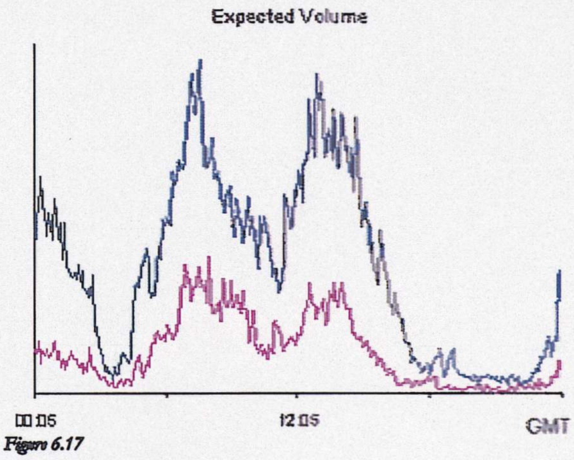

Figure 6.17

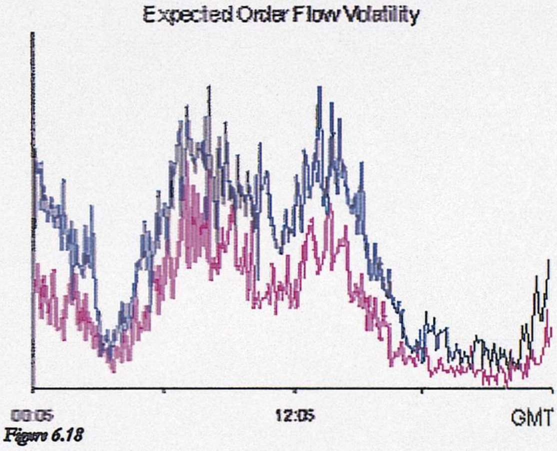

Figure 6.18

DEM/JPY, pre-EMU

	UO	EO	UV	EV	RV
BA	2%	-31%	0%	-30%	6%
RV	36%	14%	38%	13%	
EV	0%	94%	0%		
UV	46%	0%			
EO	0%				

Table 6.7

EUR/JPY, post-EMU

	UO	EO	UV	EV	RV
BA	-2%	-33%	-2%	-32%	5%
RV	41%	12%	39%	10%	
EV	0%	94%	0%		
UV	62%	0%			
EO	0%				

Table 6.8

EUR(DEM)/CHF

1999
1998

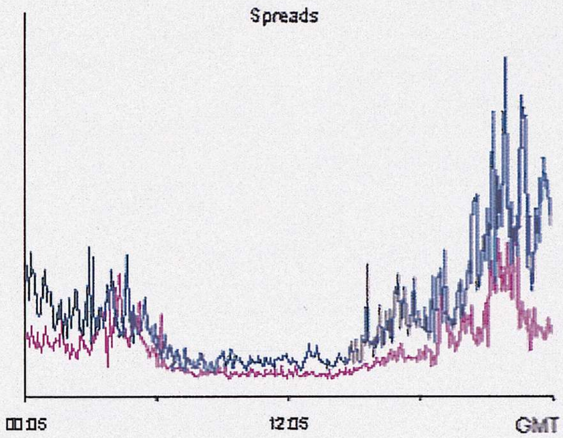

Figure 6.19

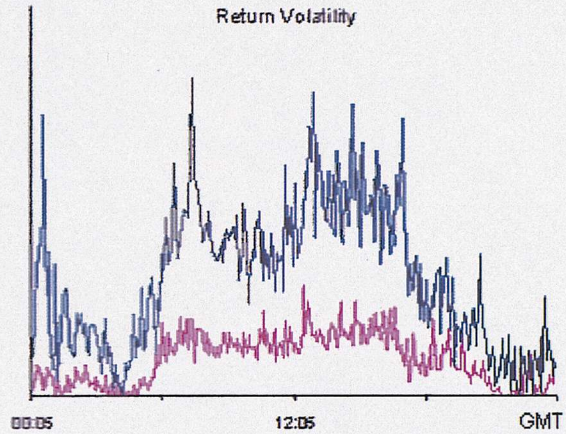

Figure 6.20

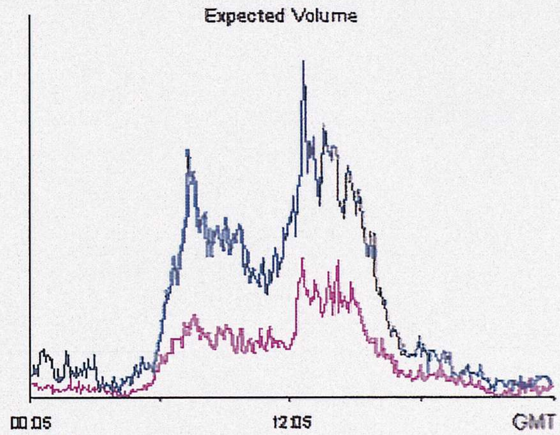

Figure 6.21

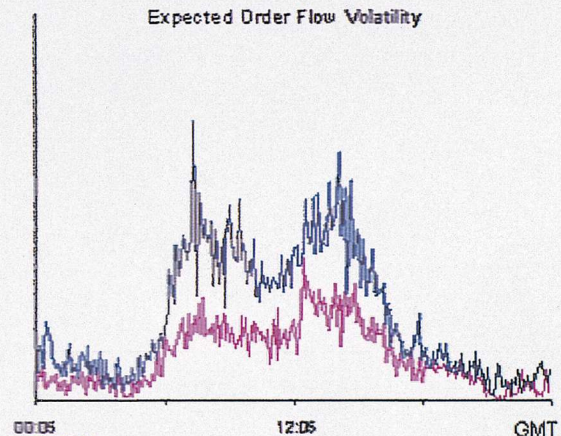

Figure 6.22

DEM/CHF, pre-EMU

	UO	EO	UV	EV	RV
BA	1%	-28%	3%	-27%	9%
RV	38%	12%	46%	12%	
EV	0%	94%	0%		
UV	51%	0%			
EO	0%				

Table 6.9

EUR/CHF, post-EMU

	UO	EO	UV	EV	RV
BA	-2%	-36%	-2%	-35%	8%
RV	34%	4%	41%	4%	
EV	0%	94%	0%		
UV	52%	0%			
EO	0%				

Table 6.10

Short Sterling

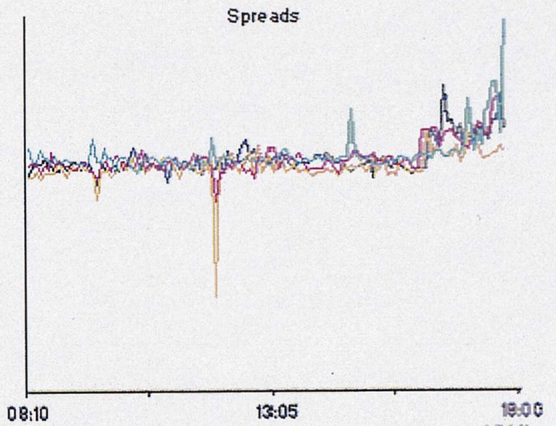

Figure 6.23

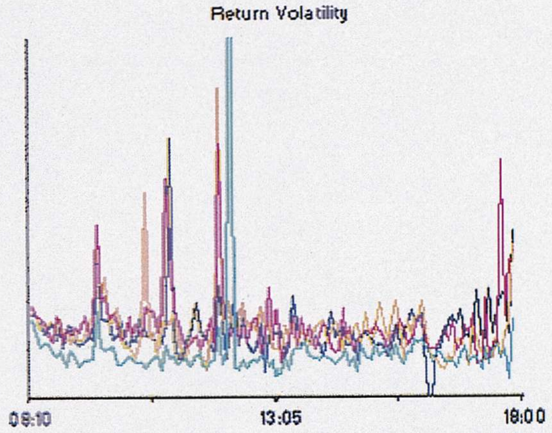

Figure 6.24

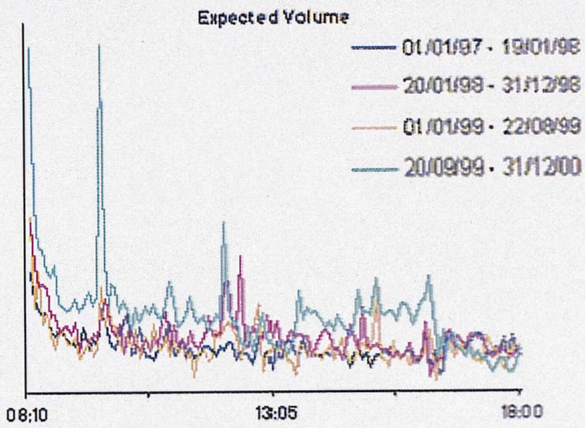

Figure 6.25

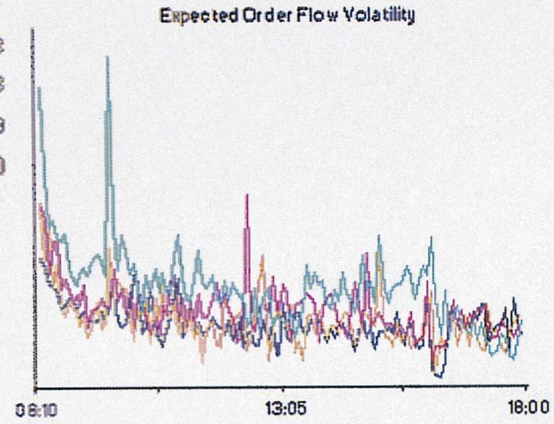

Figure 6.26

Short Sterling, 01/01/97 - 19/01/98

	UO	EO	UV	EV	RV
BA	0%	-4%	-1%	-5%	2%
RV	1%	15%	1%	15%	
EV	0%	94%	0%		
UV	88%	0%			
EO	0%				

Table 6.11

Short Sterling, 20/01/98 - 31/12/98

	UO	EO	UV	EV	RV
BA	-2%	-6%	-4%	-7%	1%
RV	0%	11%	0%	13%	
EV	0%	90%	0%		
UV	86%	0%			
EO	0%				

Table 6.12

Short Sterling, 01/01/99 - 22/08/99

	UO	EO	UV	EV	RV
BA	1%	-4%	-2%	-5%	-3%
RV	0%	22%	0%	20%	
EV	0%	96%	0%		
UV	92%	0%			
EO	0%				

Table 6.13

Short Sterling, 20/09/99 - 31/12/00

	UO	EO	UV	EV	RV
BA	0%	1%	0%	1%	1%
RV	0%	9%	0%	10%	
EV	0%	95%	0%		
UV	86%	0%			
EO	0%				

Table 6.14

Euroswiss

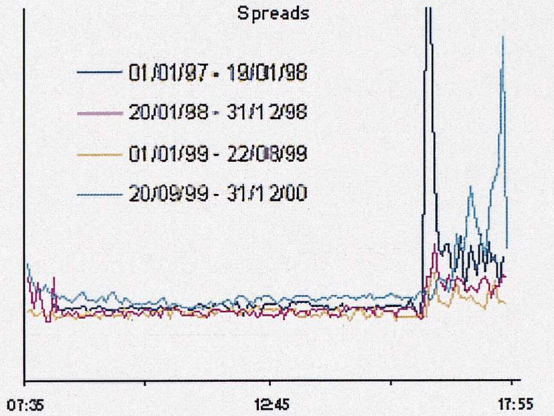

Figure 6.27

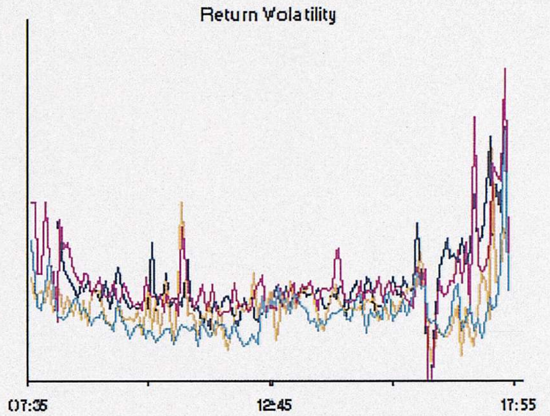

Figure 6.28

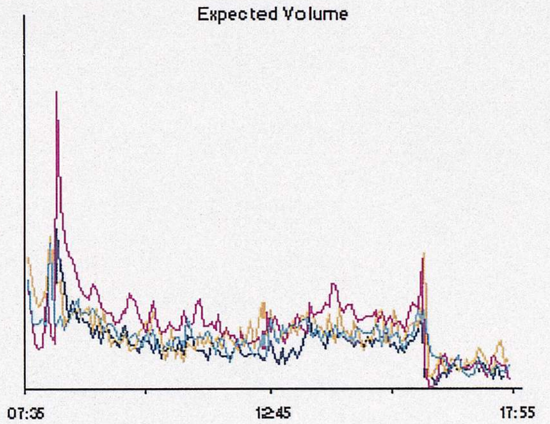

Figure 6.29

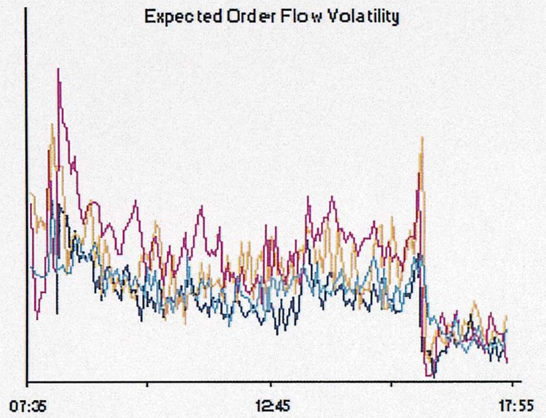

Figure 6.30

Euroswiss, 01/01/97 - 19/01/98

	UO	EO	UV	EV	RV
BA	-4%	-21%	-4%	-15%	12%
RV	0%	-3%	0%	11%	
EV	0%	93%	0%		
UV	86%	0%			
EO	0%				

Table 6.15

Euroswiss, 20/01/98 - 31/12/98

	UO	EO	UV	EV	RV
BA	-3%	-20%	-3%	-14%	7%
RV	0%	8%	0%	20%	
EV	0%	93%	0%		
UV	88%	0%			
EO	0%				

Table 6.16

Euroswiss, 01/01/99 - 22/08/99

	UO	EO	UV	EV	RV
BA	-1%	-9%	0%	-8%	-2%
RV	0%	19%	0%	21%	
EV	0%	96%	0%		
UV	87%	0%			
EO	0%				

Table 6.17

Euroswiss, 20/09/99 - 31/12/00

	UO	EO	UV	EV	RV
BA	-1%	-7%	0%	-6%	4%
RV	0%	26%	0%	35%	
EV	0%	94%	0%		
UV	84%	0%			
EO	0%				

Table 6.18

Eurolire

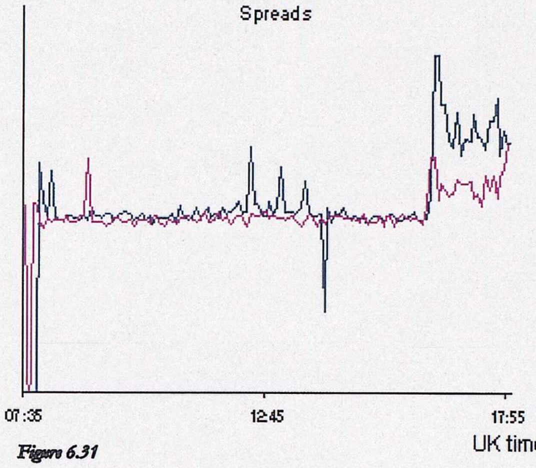

Figure 6.31

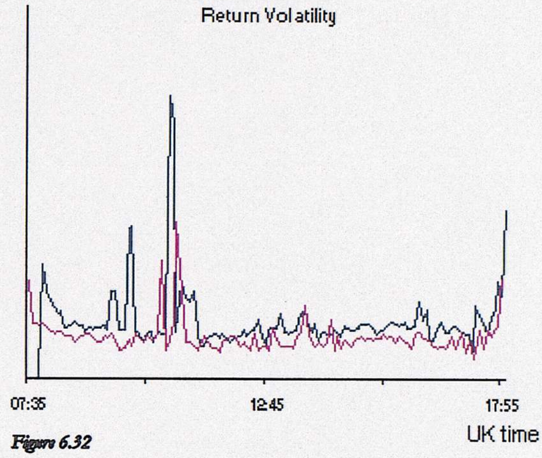

Figure 6.32

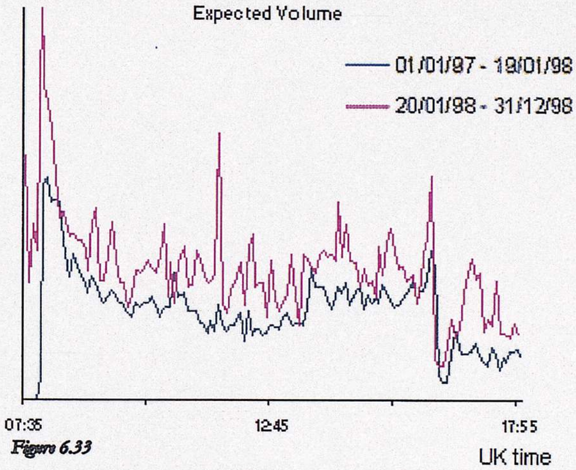

Figure 6.33

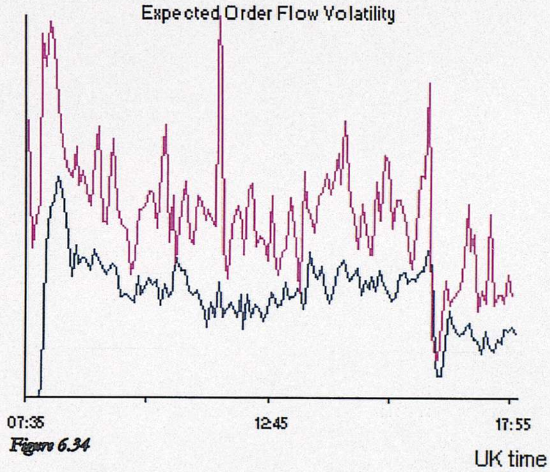

Figure 6.34

Eurolire, 01/01/97 - 19/01/98

	UO	EO	UV	EV	RV
BA	-1%	-9%	-2%	-8%	0%
RV	0%	11%	0%	12%	
EV	0%	94%	0%		
UV	81%	0%			
EO	0%				

Table 6.19

Eurolire, 20/01/98 - 31/12/98

	UO	EO	UV	EV	RV
BA	0%	-5%	0%	-4%	1%
RV	0%	13%	0%	18%	
EV	0%	93%	0%		
UV	92%	0%			
EO	0%				

Table 6.20

Euromark/Euribor

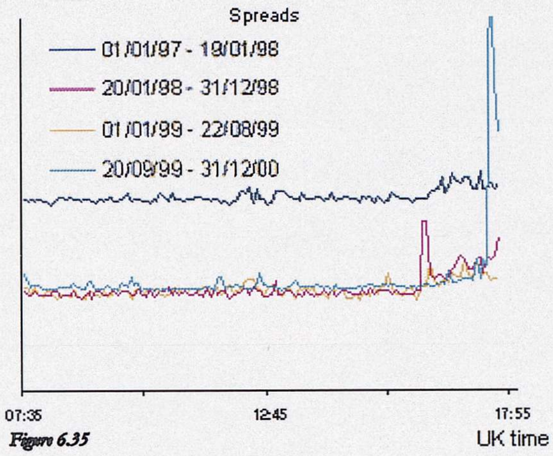

Figure 6.35

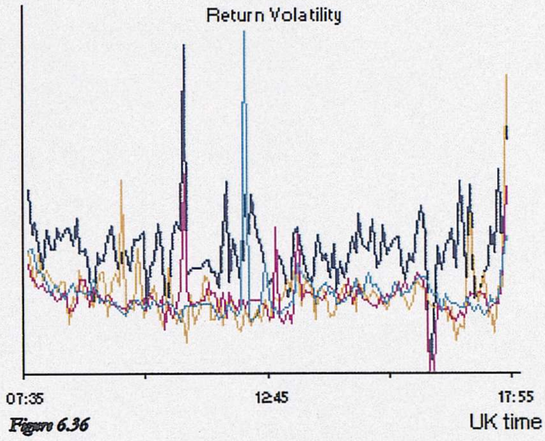

Figure 6.36

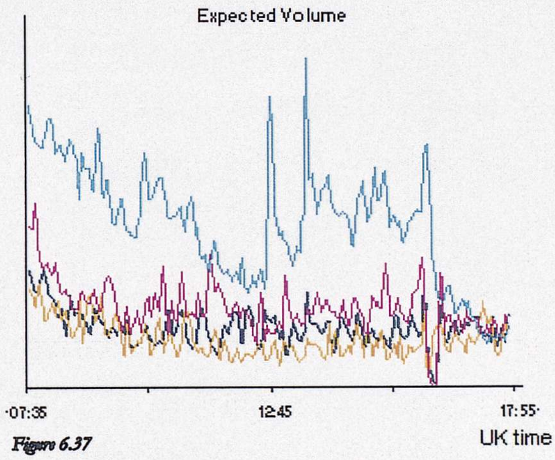

Figure 6.37

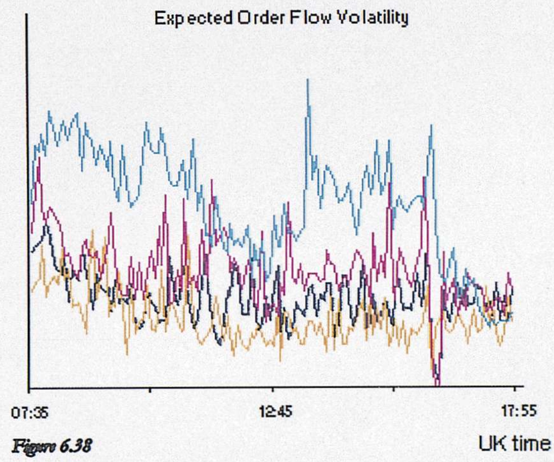

Figure 6.38

Euromark, 01/01/97 - 19/01/98

	UO	EO	UV	EV	RV
BA	2%	-2%	1%	-2%	4%
RV	0%	18%	0%	21%	
EV	0%	93%	0%		
UV	90%	0%			
EO	0%				

Table 6.21

Euromark, 20/01/98 - 31/12/98

	UO	EO	UV	EV	RV
BA	-2%	-12%	-3%	-10%	7%
RV	0%	13%	0%	16%	
EV	0%	91%	0%		
UV	92%	0%			
EO	0%				

Table 6.22

Euribor, 01/01/99 - 22/08/99

	UO	EO	UV	EV	RV
BA	-1%	-4%	0%	0%	2%
RV	-1%	9%	0%	19%	
EV	0%	90%	0%		
UV	88%	0%			
EO	0%				

Table 6.23

Euribor, 20/09/00 - 31/12/00

	UO	EO	UV	EV	RV
BA	0%	-7%	0%	-6%	1%
RV	0%	5%	0%	12%	
EV	0%	94%	0%		
UV	79%	0%			
EO	0%				

Table 6.24

6.3.2.1 Observations on Patterns

The most obvious regularity that is evident across the LIFFE STIR futures is that at 16:25, bid-ask spreads rise, while volumes and volatility fall. This represents the opening of APT trading in the 'floor' regime. These prices are also used for settlement which means that this patterns carries on after the demise of the floor trading regime. The Euroswiss contract offers the clearest evidence of the fabled U-shaped pattern, displaying a clear U-shape in bid-ask spreads, price change volatility, volume and order flow volatility with peaks at 08:10 and 04:25. Short Sterling shows spikes of activity just after 09:30 and just after 12:00, which are the times of official UK national statistics news releases and of Bank of England announcements, respectively. Curiously, this pattern is associated with low bid-ask spreads in the 'floor' trading period and with high bid-ask spreads under electronic trading. Short Sterling volume is clearly higher in the electronic period. The huge surge in Euribor volume when the contract goes electronic is very clear in the volume graph. This new pattern of volume is reminiscent of the sharp-peak-followed-by-distributed-peak pattern evident throughout the spot FX volume graphs, perhaps suggesting a link with the spot FX market what was not visible previously. Before the minimum tick size was changed, the Euromark bid-ask spread was twice the level of the post-change Euromark and Euribor bid-ask spreads. The higher price change volatility and associated bid-ask spread peak which occurs after 16:25 for the Euromark, disappears for the Euribor. However, the volume behaviour over the day for the periods before the introduction of electronic trading, is very similar before and after EMU, as is consistent with the notion that the early Euribor was the linear successor of the legacy Euromark. After the introduction of electronic trading, the structure of the Euribor trading day looks different. This period coincides with the LIFFE Euribor's rise to dominance over rival contracts in Frankfurt and Paris that was illustrated in figure 4.2.

The M-shape pattern previously documented by other FX researchers is clear in the spot FX volume and volatility graphs. It is also evident in the order flow volatility graphs. Spot FX bid-ask spreads display a U-shaped dip during the heavy trading part of the day, i.e. when London or New York are active. The general fall in volume, price change volatility and order flow volatility, and the rise in bid-ask spreads in 1999 from 1998 as previously noted in chapter 5 of this thesis are revealed in intra-day detail. USD/CHF did not conform to this pattern of change. For EUR(DEM)/CHF, everything is lower in 1999. Although, the bid-ask spread has fallen noticeably less than the other variables.

6.3.2.2 Does the Theory Fit?

The evidence from the both markets provides generally strong support for the hypotheses laid out above. However, the most contentious hypotheses, concerning the relationship between volume and volatility are overwhelmingly rejected. It is clear from visual comparison of the intra-day average patterns, particularly for spot FX, that expected volume and price change volatility are strongly positively correlated. The correlation matrices confirm this and, for spot FX, reveal further that unexpected volume has an even stronger positive link to price change volatility. While this finding is broadly in line with previous empirical findings, it flatly contradicts the negative volume-volatility relationship implied by the Admati and Pfleiderer(1988) framework that I demonstrated above.

A strong negative relationship between bid-ask spreads and volume is evident in the spot FX charts. The correlation matrices divulge that expected volume is significantly negatively correlated with bid-ask spreads, confirming the asserted relationship by Admati and Pfleiderer(1988). The evidence from the STIR market is weaker on this relationship but is still supportive. Contrary to the expectations of Easley and O'Hara(1992), the correlation matrices show that the link between bid-ask spreads and unexpected volume is very weak in both markets.

No strong relationship between bid-ask spreads and volatility is discernable in either instrument type. Although, all but one of the correlation tables exhibit the expected positive sign.

The expected relationship between order flow volatility and volume is strongly confirmed in both the spot FX data and STIR futures data. Expected order flow volatility is very highly correlated with expected volume, which is consistent with the notion that informed traders choose to trade at the same time as uninformed traders, as Admati and Pfleiderer(1988) conjectured. However, the graphs also clearly show that order flow is not dormant in the low volume period. In the spot FX market, the relationship between the unexpected components is weaker than that of the expected components but is still positive and far from insubstantial, which indicates that order flow is not the only source of unexpected volume. There are four permutations whereby an expected part of either volume or order flow volatility could be correlated with the unexpected part of either itself or the other. As predicted, all four were found to have correlation coefficients of zero for every instrument and sample period.

In short, there is strong evidence from the spot FX market for 10 of the hypotheses (6.2, 6.4, 6.6, 6.8, 6.10, 6.11, 6.12, 6.13, 6.14 and 6.15), weak evidence or no relationship in 3 cases (6.1, 6.3 and 6.5) and strong contrary evidence in 2 cases (6.7 and 6.9). The STIR futures market provides strong positive evidence for 6 hypotheses (6.10, 6.11, 6.12, 6.13, 6.14 and 6.15), weaker evidence for 2 more (6.2 and 6.4), while there is little or no evidence for a further 5 hypotheses (6.1, 6.3, 6.5, 6.6, and 6.8) and strong contrary evidence in 2 cases (6.7 and 6.9). In all cases, visual comparison of the charts validates the numerical finding in the correlation matrix.

6.3.2.3 What Do the Results Mean?

The positive relationship observed between volume and price change volatility raises a serious question about the ability of asymmetric information based models to completely explain the observed intra-day patterns. However, it should be remembered that the central model considered here, Admati and Pfleiderer(1988), contains an implicit assumption that the variance of the price innovation is constant. If relevant macroeconomic and company news were more likely to occur during trading hours, then this assumption would not hold. In that case, perhaps a combined flow-of-news and asymmetric information might fare better. Even then, there are two pieces of evidence which suggests that this answer falls short. First is the empirical evidence investigated by French and Roll(1986). These data included days where the market experienced unscheduled closures. There is no reason to believe that the amount of news was any less but volatility proved much lower than the open market data led one to expect. The same point applies to Ito et al.'s(1998) application of French and Roll's(1986) model to the spot FX market. The second piece of contrary evidence is contained in the pattern charts above. The patterns in spot FX data show how price change volatility and order flow volatility closely follow the peaks and troughs of volume over the trading day. The USD/JPY is particularly telling. After Tokyo closes and before New York opens, a large volume of London trading can be discerned. It is hard to believe that much news important for the USD/JPY occurs during this period. Yet, price change volatility and order flow volatility are shown to be high in this period. In other words, the magnitude of price change seems closely aligned with order flow and with volume, even when these occur at times when relevant news is unlikely to be released.

This close association between volume and price change volatility fits particularly well with Clark's(1973) original MDH model, in which volatility in daily price change composed of n successive individual price change increments within the day, increases as n increases. The number of trades, n, is interpreted as a proxy for volume.

The U-shaped bid-ask spreads evident in the STIR futures charts seem most consistent with Brock and Kleidon's(1992) explanation. This is further supported by the fact that the spot FX data reveal an intra-week U-shaped bid-ask spread pattern for all five currency pairs studied (not displayed). However, Admati and Pfleiderer's(1988) negative relationship between volume and spreads is still widely supported in the correlations and in the 24-hour spot FX intra-day patterns

6.4 What Else Could Drive Order Flow?

The order-flow-erodes-return argument in conjunction with the positive order flow volatility - return volatility empirical finding rule out asymmetric information about future price innovations as a main driver of order flow and therefore of prices and bid-ask spreads. The question arises as to what the alternative driver(s) of order flow could be.

Existing microstructure theory considers only one rival force to asymmetric information as capable of creating order flow – inventory. Inventory is normally defined as a temporary imbalance between supply and demand which a market maker is willing to hold for a short period. Buyers and sellers with different motives and/or different views co-exist in the market at the same time. Garman's(1976) seminal paper on order flow permitted demand and supply probability distributions to be identical but independent. This assumption would allow inventory to appear lumpy. However, this explanation does not obviate the need for each inventory increment to demand a price concession in order to be absorbed by the market. This simple argument could account for the observed positive order flow volatility-return volatility relationship. In addition, it is consistent with Clark's(1973) MDH. Furthermore, it appears to go some way to explaining one of the big FX puzzles: why exchange rates are excessively volatile.

In recent years, several papers have come to the fore which suggest that observed exchange rate changes can not be adequately explained as a random walk. More specifically, papers from Mark(1995), Flood and Taylor(1996) and Froot and Ramadorai(2002) have found a link between exchange rate changes and fundamentals, but only over the very long run. The question then arises as to whether the residual in these models dissipate quickly or whether some part of it has persistence. If it can be demonstrated that the residuals dissipate quickly, then an inventory plus fundamentals explanation should suffice to fully characterise how exchange rates are determined. If, however, some part of the error is shown to have persistence beyond a few minutes, then two facts make the inventory explanation appear inadequate. First, except on vary rare occasions when central banks intervene, the aggregate amount of a currency held is constant. Second, we know (e.g. from Lyons(1995)) that the intermediaries in the spot FX market dispatch unwanted inventory within minutes, which means that all units of the currency must be held by long term holders. These facts can comfortably accommodate inventory imbalances lasting for several minutes but not much more. If a more sustained factor is necessary, a different explanation would be required.

6.5 Conclusion

The fact that volume and (price change) volatility are shown to be strongly positively related firmly rejects asymmetric information as the dominant explanation of observed intra-day patterns and relationships. The reason for this is that order-flow volatility and price change volatility should be negatively related. This follows because high levels of order flow will erode price changes by breaking up the price impact of an independent news event and merging it with price impact from uncorrelated news events. The merging of two uncorrelated series lowers variance, and so lowers the average (price change) deviation. Allowing the size of price innovations to vary with news does not fully mitigate this. This order flow erosion insight also undermines the French and Roll(1986) conclusion, as well as that of Ito et al.(1998), that the cause of observed excess volatility during normal trading hours must be due to informed trading.

The pattern of intra-day order flow volatility, a variable hitherto unaddressed in the intra-day pattern literature, is revealed to be very closely linked to that of volume. Previous empirical research had only linked order flow to path dependence. My work reveals order flow's relationship with bid-ask spreads and price change volatility, as well as volume, for the two markets studied.

No empirical evidence was found to support the bid-ask spread – unanticipated volume predictions of Easley and O'Hara(1992). By contrast, the bid-ask spread – expected volume relationship predicted by Admati and Pfleiderer(1988) was borne out.

Echoing some of the findings of the previous chapter, the intra-day charts show that EMU has brought lower FX volumes, price change volatility and order flow volatility, and higher spreads, in most cases. STIR futures contracts appear to have been much more perturbed by the switch to electronic trading than by EMU per se. However, both these changes appear to have had little impact on the fundamental relationships between bid-ask spreads, price change volatility, volume and order flow volatility in these markets.

The Brock and Kleidon's(1992) explanation seems to account for why the bid-ask spread is U-shaped between the open and close. That aside, what is required now to explain the other observed intra-day patterns is an hypothesis which encompasses Clark's(1973) volume-volatility conclusion, but also retains Admati and Pfleiderer's(1988) insight about the negative intra-day link between the bid-ask spread and volume. To be consistent with my empirical evidence, it should also predict a weak relationship between the bid-ask spread and price change volatility and a very strong connection between the magnitude of order flow and volume.

I suggest inventory as a possible candidate as an alternative driver for order flow. If inventory were driven by random imbalances between buyer and seller volume, it would help explain not only the observed positive order flow volatility – return volatility relationship, but also could shed some light on why exchange rates appear to be excessively volatile. If deviations from fundamental fair value can be shown not to have persistence of more than a few minutes, an inventory plus fundamentals explanation may be sufficient to explain the spot FX determination puzzle as well.

Chapter 7.

The Relative Importance of Information, Inventory and Price Clustering

7.1 Introduction

This chapter quantifies the respective contributions of asymmetric information, inventory and price clustering in the formation of prices and bid-ask spreads. While previous components of bid-ask spread studies have used the asymmetric information and inventory factors extensively, price clustering is new to this arena. To apportion credit to each of these factors, I use a trade indicator model. In addition to an established version of this model, I introduce a brand new variant, which I call the "modified" trade indicator model. This is a model which has been specifically adapted for the order driven market environment and which proves to fit the empirical data much better than the original model, for both spot FX and STIR futures markets.

In previous research, trade indicator models have primarily been used to identify the components of the of the bid-ask spread, e.g. Glosten and Harris(1988), Huang and Stoll(1997). However, the model can also be used to explain what drives stock prices, e.g. Madhavan, Richardson and Roomans(1997). This is because of the close inter-dependence between price increments and bid-ask spreads evident in the theoretical work of Kyle(1985), Harris(1986), Glosten and Milgrom(1985) and Glosten(1987), as well as of Glosten and Harris(1988).

The work presented in this chapter is noteworthy in its own right for three distinct reasons. First, the trade indicator bid-ask spread model is applied to two financial instruments which have not previously been analysed in this way, i.e. spot FX and STIR futures. Second, the trade indicator model has not previously been adapted to electronic order-driven markets, as both of these markets are. Third, in line with a pervasive theme of this thesis, the model is engaged in a comparative analysis of the pre-EMU and post-EMU market environments.

The "original" trade indicator model that I use follows closely the derivation of the model used by Huang and Stoll(1997). They used this model to split the bid-ask spreads into three components for NYSE stocks. My second ("modified") trade indicator model is a variant of the original.

Glosten and Harris(1988) produced the seminal trade indicator model. They used the trade indicator sequence to separate price changes into two parts: 1) the "transitory" part which encompasses order-processing and inventory management costs and 2) the "adverse-selection" component which are permanent price shifts associated with informed trading. Both the Madhavan et al(1997) and Huang and Stoll(1997) models are direct descendents of the Glosten and Harris(1988) model. Stoll(1978) first identified three components of the bid-ask spread, linked to: 1) order processing, 2) inventory management and 3) adverse selection risk. To date, Huang and Stoll(1997) is the only model that is widely acknowledged as having successfully decomposed empirical bid-ask spreads into all three components. In addition, Huang and Stoll(1997) reconcile all the other notable bid-ask spread models, including trade indicator models within their framework. This includes Madhavan et al(1997) study. The latter harness the trade indicator model in an intra-day analysis, which enables them to study the changing structure of the bid-ask spread over the day and to explore the nature of price discovery.

Both the Huang and Stoll(1997) and the Madhavan, et al(1997) analyses use the Generalised Method of Moments (GMM) regression method to estimate parameters for the model. Both use New York Stock Exchange (NYSE) data obtained from the Institute for the Study of Securities Markets (ISSM). Both use a full year of tick data – Madhavan et al.(1997) use 1990, while Huang and Stoll(1997)use 1992. Indeed, both were published in the same journal, at the same time. Furthermore, each makes reference to the other. However, they reach surprisingly different conclusions. Madhavan, Richardson and Roomans(1997) find that asymmetric information comprises between 36% and 51% of the bid-ask spread. By contrast, Huang and Stoll(1997) find that order processing is responsible for about 90% of bid-ask spreads, ranging from 96.7% down to 57% across stocks. In a slightly more elaborate model, the latter isolate the asymmetric information component as 9.6% of the bid-ask spread on average, with a range of between 1.4% and 22%.

On the basis of previous research work and of acknowledged facts about the markets, one might expect that the bid-ask spread for spot FX and STIR futures will have quite a different composition from that of equities. Inventory risk is not nearly as significant a problem in the former markets as it is in the latter (Manaster and Mann(1996). Chartism is more widely practised and more generally accepted in both the FX (Allen and Taylor(1992), Liu and Mole(1998)) and futures markets (Schwager(1984)) than it is in the equity market. This may mean that shared beliefs are a factor in driving prices and that in turn could lead to greater amounts of certain pricing behaviour than markets where Chartism is less accepted. Furthermore, the lack of formal prohibition of insider trading in the unregulated trans-national FX market, coupled with the known practice, particularly in the spot FX market, of dealers exploiting their customer order-flow to their own advantage in the inter-dealer markets, could result in a significant adverse selection component which is linked to future accumulated inventory. Additionally, the phenomenon of central bank intervention has no parallel in the equity market, nor probably in any other financial market. Intervention, at least in its unsterilised form, can be interpreted as a naked attempt to use accumulated inventory to drive price. Finally, the tight price clustering and relatively stationary intra-day bid-ask spreads observed in an earlier chapter of this thesis may lead one to expect that the order-processing factor might be strongly dominant for STIR futures data.

7.2 The Components of Bid-Ask Spreads and Price Innovations

The conventional bid-ask spread decomposition model defines the components of the bid-ask spreads as being associated with: 1) adverse selection risk, 2) inventory management and 3) order processing. The latter are usually attributed to the administration costs of processing orders, but are actually more of a catch-all residual variable. Madhavan et al(1997) also include price discreteness as an additional explanatory variable. From my analysis in chapter 5 of this thesis, it is clear that price clustering could have an influence that supersedes and encompasses that of price discreteness alone. Furthermore, in an order-driven market, there is no obvious reason why a limit order trader would incur more administration costs or be associated with high fixed overhead cost than a market order trader. This makes price clustering a much more credible candidate as the third bid-ask spread component than order processing cost.

Since Madhavan et al(1997) focus on price innovation, they do not interpret price discreteness(clustering) in the context of the bid-ask spread. My own interpretation is that it would provide a windfall gain to the provider of liquidity. For example, suppose that the minimum tick size is 1 unit. If adverse selection risk warrants 0.4, inventory warrants 0.3 and there are no other costs, how large will the bid-ask spread be? While it would be mathematically convenient for the bid-ask spread to be 1 for 70% of the time and 0 for 30% of the time, a permanent rounding up to 1 is a possible result. If that happened, price discreteness alone would net a 30% windfall premium for the non-initiating side.

The interpretation of adverse selection risk also needs to be reconsidered here, in light of Lyons'(2001) suggestion that asymmetric information in spot FX markets could be about future accumulated inventory rather than about future payoffs and discount rates. This raises the question of whether information about future inventory should be counted as an inventory or an information effect. In my view, if inventory can distort price, then inventory should be viewed as a rival for future valuations and payoffs as a price driver, not to asymmetric information. If the Martingale assumption underlying the price innovation process were the same when inventory drives the innovations as when news drives it, then whether the driver is called "news" or "inventory" is irrelevant. Asymmetric information relates to future price innovation, regardless of its cause. This means that informed trading based on future inventory should be indistinguishable from informed trading based on future fundamentals. It follows from this that the presence of informed trading reveals nothing about whether the informed traders are good at reading fundamentals or whether they see future inventory before it arrives. Either way, informed traders profit from private information and reduce the price gap between the prevailing and the next equilibrium level.

7.3 The Huang and Stoll Model

Trade indicator models relate price changes to the "side" of trades, i.e. whether a given trade is transacted near the prevailing bid quote or the prevailing ask quote. Huang and Stoll's(1997) form of the model seems the most general, which is why theirs is the notation that I follow here. In this model, the trade indicator variable, denoted as Q_t, can take on only three distinct values, +1 when the transaction is initiated by the buyer (i.e. where the transaction price is above the mid-quote), -1 when the transaction is initiated by the seller (i.e. where the transaction price is below the mid-quote) and 0 where neither party can be identified as the initiator (i.e. where the transaction price exactly equals the mid-quote). The prevailing bid and ask quotes, which comprise the mid-quote, are defined as those which pre-exist each trade and must be no more than one minute old[3].

Huang and Stoll(1997) develop three different models and each model, in turn, has a variant which accommodates the trade size dimension. Only the first two of their models are utilised here as the third model uses data which can not be paralleled in this study. The second model is an expansion of the first.

In the first Huang and Stoll(1997) model, they identify the order processing (/ price clustering) component from the combined adverse selection and inventory management components. The regression model is expressed as follows:

$$\Delta P_t = \tfrac{S}{2}(Q_t - Q_{t-1}) + \lambda \tfrac{S}{2} Q_{t-1} + e_t$$
$$(= \tfrac{S}{2}(Q_t - Q_{t-1}) + (\alpha + \beta - 1)\tfrac{S}{2} Q_{t-1} + e_t)$$

Equation 7.1: The basic trade indicator model

The left hand side variable is the change in traded price. On the right hand side, the $S/2$ is the constant half-spread and λ represents the combined adverse selection and inventory management components. Both λ and $S/2$ are coefficient variables computed by the regression. The error term combines public news releases which change prices and random deviations in the bid-ask spread.

The second (expanded) Huang and Stoll(1997) model adds an additional lag of the trade indicator variable to the above equation as follows:

$$\Delta P_t = \tfrac{S}{2} Q_t + (\alpha + \beta - 1)\tfrac{S}{2} Q_{t-1} - \alpha \tfrac{S}{2}(1 - 2\pi) Q_{t-2} + e_t$$

Equation 7.2: The extended trade indicator model

[3] Huang and Stoll(1997) discuss using a 2 minute interval, although they use an alternative method in the end. Since, my volume far exceeds theirs in every case, I decided a shorter interval would be appropriate. However, the choice of 1 minute is essentially arbitrary.

(α plus β) equals λ in the previous equation. $(1\text{-}2\pi)$ indicates the conditional expectation of Q_{t-1} given Q_{t-2}, where π is the probability that a trade is of the opposite sign to the previous one. The second model also requires that $(1\text{-}2\pi)$ be simultaneously estimated along with the extended regression equation, using the following additional regression equation:

$$Q_t = (1 - 2\pi)Q_{t-1} + u_t$$

Equation 7.3: The second part of the extended trade indicator model, which is estimated jointly with Equation 7.2.

The above outlines the two original Huang and Stoll(1997) models which I implement below.

As Huang and Stoll(1997) explain, three separate but co-existing variables linked to price innovation underpin these models. These are 1) the fundamental valuation, "V_t", 2) the mid-quote value between the bid and the ask, "M_t", and 3) the transaction price, "P_t". These variables are linked together by the time variable, t, such that t refers to the values of these variables at the instant a transaction occurs. Therefore, V_t and M_t must exist prior to the transaction occurring and P_t is brought into existence by the transaction event. As I will show below, the generating process for each of these variables can be linked to the Trade (side) indicator, Q_t, and this fact enables the derivation of the structural models above.

From the theoretical work of Copeland and Galai(1983) and Glosten and Milgrom(1985), the relationship between V and Q may be defined as:

$$V_t = V_{t-1} + \alpha \tfrac{S}{2} Q_{t-1} + \varepsilon_t$$

Equation 7.4: How order flow drives fundamental value.

This says that the current value of the theoretical construct, V_t, equals the previous value, V_{t-1}, plus a value related to the side of the previous trade, Q_{t-1} and an error term, ε_t, which captures randomly occurring news events. S, which forms part of the coefficient on Q_{t-1}, is the bid-ask spread which is assumed constant over time. The other variable that makes up the coefficient is a which denotes the portion of the half-spread that is due to informed trading. To put it another way, the change in fundamental valuation over time is determined by both informed trading and the random release of public information.

In the inventory model literature, it can be inferred from Stoll(1979) and Ho and Stoll(1981), that M is related to V and Q as follows:

$$M_t = V_t + \beta \tfrac{S}{2} \sum_{i=1}^{t-1} Q_i$$

Equation 7.5: How the market maker's mid-quote reflects his accumulated inventory.

This equation says that the mid-quote equals the fundamental value plus a value related to the sum of the previous trade side indicator series. The latter is closely linked to accumulated inventory. The coefficient on the trade indicator series sum consists of S, which is the bid-ask spread as before, and β which is the portion of the half-spread due to inventory holding costs. This equation may be interpreted as saying that the mid-quote deviates from the fundamental value by the amount of accumulated inventory.

The traded price, P, is linked to M and Q in the following way:

$$P_t = M_t + \tfrac{S}{2} Q_t + \eta_t$$

Equation 7.6: How the mid-quote and the bid-ask spread together produce the transaction price.
The above equation says that the difference between the traded price and the mid-quote will be the half-spread multiplied by the side of the trade. So, transactions at the ask price will equal mid-quote plus the half-spread, while trades at the bid price will equal mid-quote minus the half-spread. Expanding this equation slightly reveals the power and simplicity of this model:

$$P_t = V_t + \beta \tfrac{S}{2} \sum_{i=1}^{t-1} Q_i + \tfrac{S}{2} Q_t + \eta_t$$

Equation 7.7: The transaction price equation expanded to show its value and inventory components.

This expression shows that this model relates price to two interesting components and two relatively uninteresting components. The relatively uninteresting components are the half-spread and the error term. For the purpose of this analysis, the interesting components are the fundamental value and the accumulated inventory. While it is common to think of inventory as being transitory with no long-run effects, Lyons(2001) shows that accumulated inventory can also directly affect price through "portfolio balance" inventory effects.

In the preceding equations, note that P encompasses M and that, in turn, M encompasses V. Taking the first difference of equation 7.7 allows the components of the bid-ask spread to be revealed as a function of a sequence of trade indicator variables:

$$
\begin{aligned}
\Delta P_t &= \Delta V_t + \tfrac{S}{2}\Delta Q_t + \beta \tfrac{S}{2} Q_{t-1} + \Delta \eta_t \\
&= \tfrac{S}{2}\Delta Q_t + \alpha \tfrac{S}{2} Q_{t-1} + \varepsilon_t + \beta \tfrac{S}{2} Q_{t-1} + \Delta \eta_t \\
&= \tfrac{S}{2}\Delta Q_t + (\alpha + \beta) \tfrac{S}{2} Q_{t-1} + (\Delta \eta_t + \varepsilon_t) \\
&= \tfrac{S}{2}\Delta Q_t + \lambda \tfrac{S}{2} Q_{t-1} + e_t
\end{aligned}
$$

Equation 7.8: How the change in the transaction price relates to the sign indicator sequence of past trades.

The final form of this equation is the one I used to introduce the first Huang and Stoll(1997) model in equation 7.1 above.

As mentioned above, this first model is unable to separate the adverse selection component from inventory management component. In order to achieve this separation, an additional lag of the trade indicator variable and another equation are introduced. The two new additions enable the estimation of a new coefficient, π, which makes it possible to identify *a* separately from β.

Although the two models appear very similar in structure, the second model represents a major shift. It actually embodies a blend of the two separate traditions of return decomposition models. The new term π denotes the probability that each successive trade will be of the opposite sign to its predecessor. In other words, it uses the serial correlation of the trade flow to reveal the components of the bid-ask spread. This is the method used in a different class of microstructure bid-ask spread models, referred to as "covariance" models, which originate with Roll(1984). Exploiting the Martingale assumption of Glosten and Milgrom(1985), Roll's(1984) model relates the bid-ask price to the serial covariance of price change. In other words, it infers the bid-ask spread from the bid-ask bounce. However, this model is valid only where the bid-ask spread entirely consists of order processing costs, i.e. it can not accommodate either inventory cost or adverse selection risk costs. Choi, Salandro and Shastri(1988) expand the covariance model framework to allow serial covariance. George, Kaul and

Nimalendran(1991) develop a covariance model which permits informed trading, where they find that time variation in the price change is an important factor in computing the traded bid-ask spread. Adjusting for the latter, they show the adverse selection component is smaller than had been claimed previously. All covariance models depend on the probabilities of change in trade direction, while trade indicator models only depend on the trade direction of incoming orders.

π is introduced into the relationships between V, M and P by tweaking our definition of V, as follows:

$$V_t = V_{t-1} + \alpha \tfrac{S}{2}(Q_{t-1} - (1-2\pi)Q_{t-2}) + \varepsilon_t$$

Equation 7.9: How order flow drives fundamental value when order flow is (negatively) serially correlated.

It is important to emphasise that this additional term is no way associated with price pressure generated by informed trading. The relationship between informed trading and V_t is already fully characterised by equation 7.4. Rather, this exploits what Huang and Stoll(1997) argue is a naturally occurring negative serial correlation in the mid-quote price. This negative serial correlation is unrelated the negative serial correlation in trade prices which is associated with bid-ask bounce. Negative serial correlation in the mid-quote occurs because of how market makers react to receiving or losing a unit of inventory. When they receive a unit of inventory from a seller initiated trade which they do not want to hold, they shade the mid-quote price down to make their ask price more attractive than their bid price. This improves the probability that the next trade will be at the ask price (i.e. buyer initiated), thereby enabling them to offload the unwanted inventory. This behaviour will ensure that the probability of a sell trade is highest just after a buy trade and the probability of a buy trade is highest just after a sell trade, which is sufficient to ensure negative serial correlation.

At first blush, it seems that this negative serial correlation from time t-2 should impact both informed trading and market maker reaction, implying that π should link to both α and β. However, π is not linked to β at t-2 because expectations of the next trade do not matter to market makers. Only the realisation of the next trade matters. To put it another way, market makers want to consummate the next trade, receive their half-spread and only then shade their price to get back to their desired inventory level. This fact means that π only relates to α.

By supplanting the definition of V in the sequence of equations above and deriving ΔP_t in the same way, we get:

$$\Delta P_t = \frac{S}{2} Q_t + (\alpha + \beta - 1)\frac{S}{2} Q_{t-1} - \alpha \frac{S}{2}(1 - 2\pi) Q_{t-2} + e_t$$

Equation 7.10: Equation 7.2.

This is the form of the second Huang and Stoll(1997) model that I introduced in equation 7.2 above. The second model also contains the supplementary equation shown in equation 7.3 which exposes the conditional expectation of the trade indicator Q_t, at time t-1 that facilitates the identification of α independently of β. Furthermore, a restriction of $\pi \geq 0.5$ is imposed which corresponds to negative serial correlation in the trade indicator seqeunce. Values of $\pi < 0.5$ would correspond to positive serial correlation.

Huang and Stoll(1997) also introduce a trade size dimension into their analysis. They divide trades into 3 different size categories. In the analysis below, I explore both this 3 trade size model and a simpler 2 trade size model. The equation for the latter is shown below:

$$\Delta P_t = \frac{S_s}{2} Q_{s,t} + \frac{S_l}{2} Q_{l,t}$$
$$+(\alpha_{ss} + \beta_{ss} - 1)\frac{S_s}{2} Q_{s,t-1} - \alpha_{ss}\frac{S_s}{2}(1-2\pi_{ss})Q_{s,t-2}$$
$$+(\alpha_{sl} + \beta_{sl} - 1)\frac{S_s}{2} Q_{s,t-1} - \alpha_{sl}\frac{S_s}{2}(1-2\pi_{sl})Q_{l,t-2}$$
$$+(\alpha_{ls} + \beta_{ls} - 1)\frac{S_l}{2} Q_{l,t-1} - \alpha_{ls}\frac{S_l}{2}(1-2\pi_{ls})Q_{s,t-2}$$
$$+(\alpha_{ll} + \beta_{ll} - 1)\frac{S_l}{2} Q_{l,t-1} - \alpha_{ll}\frac{S_l}{2}(1-2\pi_{ll})Q_{l,t-2} + e_t$$

Equation 7.11: The further extended trade indicator model which accounts for trade size.

7.4 The Modified Model

The negative serial correlation restriction, $\pi>0.5$, is the only element of the Huang and Stoll(1997) model that imposes transitory behaviour on the inventory component of the bid-ask spread. The authors justify this restriction by acknowledging the need for market makers to "recover inventory holding costs from trade and quote reversals". In order-driven markets, this need does not exist, as there are no exogenous liquidity providing market makers who risk their own capital to create a market and who requires compensation for doing so. Even in quote-driven futures markets, "scalpers", who are usually described as futures market makers, appear to have only small and brief inventory exposure. This combined with the fact they are not obliged to provide two-way quotes suggests that scalpers are probably more accurately described as "brokers" than as market "makers", where the latter term conveys the meaning of someone who puts up risk capital. As a result of these points, I propose to relax the $\pi>0.5$ assumption for all instruments and sample periods covered in my datasets.

There are several acknowledged features of the foreign exchange and futures markets that would lead one to expect prices to be positively serially correlated.

For example, stop loss orders, which are often linked to hedging of exotic derivatives, are a feature in both markets (Osler(2003b)). These accelerate demand in the existing direction when trigger-prices are crossed. Currency Risk Management services, used by many institutional funds, dynamically hedge with precisely the same effect. The latter are often automated with a computer triggering trades when a live price feed produces a rate above or below a pre-set level. In addition, in contrast to the equity market, widespread shared beliefs in Chartist phenomena, like trading bands, resistance levels, break-out levels and head-and-shoulders, at the very least cause traders to not counteract these markets trends and to possibly actively contribute to them (Osler(2003a)).

Rime(2000), Lyons(2001), Evans(2002) and Evans and Lyons(2002) showed that foreign exchange rates are disturbed by accumulated order flow. Lyons(2001) calls this a "portfolio balance effect". This assertion follows from earlier work by Scholes(1972), Shleifer(1986), Harris and Gurel(1986), Bagwell(1992) and Kaul, Mehrotra and Morck(2000), among others, who showed that aggregate demand is less than perfectly elastic in equity markets, which means that markets require price concessions to absorb inventory imbalances. In rudimentary microeconomics terms, portfolio balance type inventory imbalances equate to exogenous supply shifts. In the absence of any other changes, if demand curves are downward sloping, a rightward supply shift will drive price down. Reversion to "fundamental" equilibrium requires that supply curve and/or demand curve shifts, which can take a lot of time.

Lyons(2001) suggests that, in the trade indicator model, a should fully capture both informed trading associated with future fundamentals and trading linked to future accumulated inventory, and that β will capture only transitory inventory effects. His conclusion is predicated on the belief that a and β are perfectly aligned with permanent and transitory price effects respectively. However, this alignment does not allow inventory itself to accumulate or to have a permanent impact. In market microstructure theory, a is linked to the behaviour of informed traders and β with uninformed trading. As noted above, the restriction of π>0.5 in the Huang and Stoll(1997) analysis is the only element in the whole trade indicator framework that actually imposes transitory behaviour on β. If this restriction is relaxed, then both a and β can be associated with permanent price shifts. This permits accumulated inventory to have a determining influence on price. In order-driven markets, the absence of market makers rules out the possibility of transitory inventory price effects. In other words, in an order-driven market, if β is not associated with permanent price effects, it must be zero.

In market-maker-less order-driven markets, liquidity, defined as limit orders against which market orders can be executed, is endogenous. It is negatively related to the bid-ask spread and positively related to the risk of non-execution. Cohen, Maier, Schwartz and Whitcomb(1981) show that more limit orders will be submitted when bid-ask spreads are high and when the risk of non-execution is low, implying that the bid-ask spread is also endogenous because there are no market makers. Too narrow a bid-ask spread will elicit excess market orders which will widen the bid-ask spread. Too wide a bid-ask spread will draw in more keenly priced limit orders, narrowing the bid-ask spread. This means that the exogenous factor determining both liquidity and the bid-ask spread is the risk of non-execution. This is negatively related to volume and positively related to volatility. In short, in an order-driven market, not executing is unlikely when volume is high and volatility is low. In addition, Parlour(1998) finds that high depth at the best price also increases the risk of non-execution because there is a risk that a new limit order will be crowded out. Confirming the link between bid-ask spreads and non-execution risk, Foucault, Kadan and Kandel(2001) find that bid-ask spreads and times-to-execution are jointly determined in equilibrium. However these papers, like other key papers in the order-driven market microstructure literature, including Rock(1990), Glosten(1994) and Seppi(1997) rely on a crucial but flawed assumption. They all assume that informed traders would choose to submit market orders in preference to limit orders.

Recent experimental work by Bloomfield, O'Hara and Saar(2003) finds that informed traders are actually more likely to submit limit orders than market orders. The authors argue that this is because only informed traders know the true value of the underlying asset and they can extract profit using this knowledge to sell high and buy low around the true value. As a former practitioner in the FX market, I recognise features of the spot FX inter-dealer system in that description. I observed that larger (and arguably better informed) banks predominantly place limit orders on EBS, while small banks are more likely to place close out positions via market orders.

The Bloomfield et al(2003) insight utterly redefines one of the key fundamentals of the trade indicator model, the determination of fundamental value, *V*. If informed traders are setting prices, the idea of the uninformed market maker learning by vote-counting no longer applies. Instead, the evolution to *V* would be solely determined by public information shocks, ε:

$$V_t = V_{t-1} + \varepsilon_t$$

Equation 7.12: How fundamental value evolves if informed traders do not adversely select market makers.

In an order-driven market context, the definition of the mid-quote, *M*, is misleading.

$$M_t = V_t + \beta \tfrac{S}{2} \sum_{i=1}^{t-1} Q_i$$

Equation 7.13: Equation 7.5.

Equation 7.13 contains the implied suggestion that price-setting market makers adjust their mid-quote to accommodate inventory imbalances. This can not happen in order-driven markets since there are no market makers. However, there is no dispute that aggregate inventory imbalances will disturb V, insofar as downward sloping aggregate demand curves require price concessions for the excess to be held. As such, an interim variable representing the disturbed value of V seems more consistent with the mechanisms of order-driven markets. I propose the term V^* to represent V disturbed by an inventory imbalance.

$$V_t^* = V_t + \beta \tfrac{S}{2} \sum_{i=1}^{t-1} Q_i$$

Equation 7.14: Fundamental value plus accumulated inventory.

The mid-quote, M, can now be defined as a function of V^*. However, recall that Bloomfield et al(2003) said that informed traders set M in order-driven markets. The information they release can be captured by the following relationship:

$$M_t = V_t^* - \alpha \tfrac{S}{2} Q_t$$

Equation 7.15: How private information is reflected in the mid-quote.

Previously, in the quote-driven model Q_t acted as a vote counter, registering aggressive market order trading from informed traders. Here, things work a little differently. In an order-driven market trade indicator model, Q_t is the first opportunity to record the information released at M_t in the trade flow. In order-driven regimes, liquidity based trading endowments are exogenous. The choice for every trader is whether to submit a limit order or a market order. Using the Bloomfield et al(2003) insight, what had been aggressive buying or selling by an informed trader in a quote-driven context, will translate to the submission of an aggressive limit order. This will narrow the existing bid-ask spread and entice a trader on the opposite side to submit a market order in preference to a limit order.

For this reason an upward price revision will trigger a sell, thus producing a negative relationship between M_t and Q_t.

These new fundamental relationships produce the following price equation:

$$
\begin{aligned}
P_t &= M_t + \tfrac{S}{2} Q_t + \eta_t \\
&= V_t^* - \alpha \tfrac{S}{2} Q_t + \tfrac{S}{2} Q_t + \eta_t \\
&= V_t + \beta \tfrac{S}{2} \sum_{i=1}^{t-1} Q_i - \alpha \tfrac{S}{2} Q_t + \tfrac{S}{2} Q_t + \eta_t
\end{aligned}
$$

Equation 7.16: How the transaction price relates to inventory and information in an order-driven regime with informed traders submitting limit orders.

This results in a price change equation that is identical to the original one in every detail but one:

$$
\begin{aligned}
\Delta P_t &= \beta \tfrac{S}{2} Q_{t-1} - \alpha \tfrac{S}{2} Q_t + \alpha \tfrac{S}{2} Q_{t-1} + \tfrac{S}{2} Q_t - \tfrac{S}{2} Q_{t-1} + e_t \\
&= (1-\alpha) \tfrac{S}{2} Q_t + (\alpha + \beta - 1) \tfrac{S}{2} Q_{t-1} + e_t
\end{aligned}
$$

Equation 7.17: How the change in the transaction price relates to the sign sequence of past trades in an order-driven regime with informed traders submitting limit orders.

Now, order flow relating to P_{t+1} is a component of Q_t and is revealed by $-a$.

It is clearly evident that the modified model no longer needs π to identify a. However, there are still three parameters (a, β and S) to be estimated and only two explanatory variables(Q_t and Q_{t-1}). Given that there are quote prices available for all of the datasets that I use in this study, I employ the quoted bid-ask spread time series in place of the parameter S in the modified trade indicator model. As noted previously, the quoted bid-ask spread is derived from the nearest preceding bid and ask prices, if both of these are under one minute old. If the nearest bid or ask is older than that, the bid-ask spread is left blank.

In general, trade indicator models relate the price (return) on the left to demand (Q) on the right. In fact, quote revision and trade execution (= vote counting) are only two channels through which inventory and information can influence price. In quote driven regimes, inventory drives the mid-quote and information is revealed through executed trades. In an order driven regime, this is reversed. Information drives the mid-quote, while inventory impacts price via trade execution. Stoll's(1978) assumption of risk averse inventory holders, coupled with the knowledge that informed traders have an information advantage, ensures that inventory motivated trades will be effected solely via the trade execution channel. The most important point in all this is that, even though inventory and information swap channels, equation 7.17 shows that the underlying relationships that inventory and information have with price are preserved.

It seems intuitive that informed trading should feed through to Q_t. After all, why should order flow linked to P_{t+1} have any different relationship with Q_t than order flow linked to P_t had with Q_{t-1}? This leads me to wonder why a had been omitted from the coefficient of Q_t in the original quote-driven model. If informed trading activity is present at time, t, could it introduce distortion in the demand, compared with other times, t, where informed trading is not present?

The reason that quote-driven Q_t does not have an a term in the original models can be traced back to the Glosten and Milgrom(1985) assumption of regret-free quoted prices. If market makers condition their bid and ask prices on the possibility of the next trader being informed, in either direction, there can be no role for a at time t. However, the idea of regret-free prices relies on several critical assumptions. If trade size is variable and price signals are uncertain as they are in Easley and O'Hara(1987), the regret-free assumption can not hold. Furthermore, regret-free quote prices pre-suppose that the market maker is someone who could experience long-term capital exposure and finds unloading inventory difficult.

While this is evidently true of equity market makers, it is not a good description of the observed very rapid inventory turnover of market makers in quote-driven futures markets, as evidenced by Manaster and Mann(1996). Moreover, this notion of a market maker who can widen his bid-ask spread in the face of a surfacing adverse selection risk hinges on the assumption of monopoly power. If bid-ask spreads are kept tight by competition, if market makers do not all perceive the same risk at the same time, or if the market is so liquid that unwanted positions are quick and easy to offload, it is conceivable that the bid-ask spread would not need to be regret-free. If the regret-free assumption is dropped then the modified trade indicator model is more appropriate than the original. I propose to relax the regret-free quotes assumption for the quote-driven STIR futures market and use the modified trade indicator model.

There is one other aspect about the Bloomfield et al(2003) paper that remains to be addressed. As I stated above, these researchers found that informed traders prefer to place limit orders most of the time. However, they also found that, infrequently, when price deviates greatly from fair value, informed traders favour market orders. This finding confirms earlier predictions of Angel(1994) and Harris(1998). As illustrated above, the modified trade indicator model accounts for informed trading in the same way, regardless of whether it is transmitted through limit orders or market orders. So, α may be generally interpreted as a measure both of predictability and of informed trader profitability. Similarly, β will still encapsulate the impact of inventory. Finally, the interpretation of the residual, $(1-\alpha-\beta)$, still fits better with the price clustering explanation than with order processing.

An uncomfortable aspect of my amendments to the trade indicator model is that they appear to depend on a transmogrification of the established cornerstones of market microstructure bid-ask spread theory. However, there is precedent in the microstructure literature for bid-ask spreads with the attributes that I describe. Black(1991) proposed a model for transaction costs in FX markets whereby the bid-ask spread is endogenous, positively related to volatility and negatively related to volume. Hartmann(1998) formalised and embellished this simple model. Taking a lead from Thomas and Wickins(1991), Hartmann(1998) brought together two branches of economic literature: vehicle currency theory from monetary microeconomics and currency substitution from empirical macroeconomics. He went on to produce an international (vehicle) currency based general theory of the foreign exchange market. Hartmann's definition of the bid-ask spread is:

$$s_{ji} = \frac{\sigma_{ji}^{0} + \gamma_{ji}}{\overline{x}_{ji}}$$

Equation 7.18: Hartmann's bid-ask spread equation.

where s is the bid-ask spread, σ is volatility, γ is a fixed order processing cost, x is volume, and j and i are two currencies. Hartman argues that while volume and volatility are positively related, i.e. that volatility rises as volume rises, the volume effect will outweigh the volatility effect. In other words, an exogenous increase volume should both lower transactions costs and increase volatility. Hartmann(1998) describes in detail how price adjusts to accommodate order flow. His bid-ask spread model fits the above observations about both the STIR futures market data, and the spot FX market data.

The preceding bid-ask spread model fits the quote-driven STIR data just as well as the order-driven STIR data. The reason that this model fits the quote-driven market may well be that quote-driven futures markets operate more like order-driven markets than they do like quote-driven equity markets. For a start, their bid-ask spreads are observed to be substantially lower. Also, they are more like order-driven markets in that liquidity is provided by scalpers, who do not offer two-way prices and, in effect, only broker deals between buyers and sellers, rather than sustain any long-term capital exposure.

7.5 Data

The data used to analyse the components of the bid-ask spread and the drivers of price consist of both months of EBS tick data and 4 years STIR tick data (front month only). Both quote and trade price data are required. As a result, only the STIR futures data from LIFFE can be used because the other futures exchanges do not supply quote data from which side may be derived. The 5 currency pairs used here are: EUR(DEM)/USD, USD/JPY, USD/CHF, EUR(DEM)/JPY and EUR(DEM)/CHF. The futures contracts used are: Euribor/Euromark, Euroswiss and Short Sterling

As in previous chapters, the data is collected into two main groups for the pre- and post-euro periods. In the STIR data, a 4 period structure emerges which reflects 1) the change in the Euromark minimum tick size from 1 to 0.5, 2) EMU and 3) the transition from floor to electronic trading. For the purpose of comparison, all STIR contracts have been grouped into by these dates, even though a particular instrument may not be directly affected by the change. The period from 23/08/99 to 20/09/99 is excluded from the analysis because the STIR instruments used here switched to electronic trading during this period.

Each STIR trade price is assigned a side by comparing with the nearest preceding bid and ask prices within the same day. The trade indicator variable, Q_t, equals +1 when it is between the ask and the mid-quote, -1 when it is between the mid-quote and the bid and 0 when the trade price exactly equals the mid-quote, as prescribed in the Huang and Stoll(1997) model. The EBS data comes with side already identified. So, here, ask-side or "paid" trades are assigned a value of +1 and bid-side or "given" trades are assigned a value of -1.

The change in trade price variable, ΔP_t, is defined as $P_t - P_{t-1}$, where P_t is the transaction price and where successive prices occur within contiguous trading periods. The latter is defined as the trading day for STIR futures, to avoid problems with overnight periods and roll-overs. For spot FX the trading week is used because this market operates 24-hours a day, globally.

7.6 Methodology

Both Huang and Stoll(1997) and Madhavan, Richardson and Roomans(1997) utilise the Generalised Method of Moments (GMM) of Hansen(1982). As mentioned previously, I follow closely the methodology of the former. Both studies choose the GMM because its very weak distributional assumptions make it good at capturing unspecified errors.

The stage one model is implemented in the GMM structure by the expression:

$$f(x_t,\omega) = \begin{bmatrix} e_t Q_t \\ e_t Q_{t-1} \end{bmatrix}$$

Equation 7.19: GMM orthogonality conditions for the basic trade indicator model.

where $\omega=[S\ \lambda]'$ is the vector of parameters. The orthogonality conditions are expressed as $E[f(x_t,\omega)]=0$.

The GMM procedure minimises the quadratic function

$$J_T(\omega) = g_T(\omega)' S_T g_T(\omega)$$

Equation 7.20: GMM quadratic function to be minimised.

where $g_T(\omega)$ is the sample mean of $f(x_t,\omega_t)$ and S_T is the sample symmetric weighting matrix. Hansen(1982) shows that, under weak regularity conditions, the GMM estimator $\hat{\omega}_T$ is consistent and

$$\sqrt{T}(\hat{\omega}_T - \omega_0) \to N(0, \Omega)$$

where

$$\Omega = (D'_0 S_0^{-1} D_o)^{-1}$$

$$D_0 = E\left[\frac{\partial f(x,\omega)}{\partial \omega}\right]$$

$$S_0 = E[f(x_t,\omega) f(x_t,\omega)']$$

Equation 7.21: The asymptotic distribution of $\hat{\omega}$

The stage one original model and the modified trade indicator model are both exactly identified using this method.

For the stage two model, the methodology is the same, except that the $f\ (x_t,\omega)$ vector is now expressed as

$$f(x_t,\omega) = \begin{bmatrix} e_t Q_t \\ e_t Q_{t-1} \\ e_t Q_{t-2} \\ u_t Q_{t-1} \end{bmatrix}$$

Equation 7.22: GMM orthogonality conditions for the extended trade indicator model.

where $\omega=[S\ a\ \beta\ \pi]'$ is the vector of parameters of interest. The second stage model is also exactly identified, since the number of orthogonality conditions equals the number of parameters to be estimated.

7.7 *Huang and Stoll Model Results*

	Short Sterling				Euroswiss				Euromark		Euribor	
Time Period	01/01/97 - 19/01/98	20/01/98 - 31/12/98	01/01/99 - 22/08/99	20/09/99 - 31/12/00	01/01/97 - 19/01/98	20/01/98 - 31/12/98	01/01/99 - 22/08/99	20/09/99 - 31/12/00	01/01/97 - 19/01/98	20/01/98 - 31/12/98	01/01/99 - 22/08/99	20/09/99 - 31/12/00
No. of Obs.	23,160	22,086	14,035	67,683	30,320	38,258	18,090	35,875	23,605	23,334	16,452	155,308
(A) Stage 1 model												
λ	24%	30%	32%	8%	44%	40%	35%	24%	10%	24%	26%	8%
S.E.(λ)	*(0.0124)*	*(0.0108)*	*(0.0138)*	*(0.0039)*	*(0.0092)*	*(0.007)*	*(0.0189)*	*(0.0134)*	*(0.0069)*	*(0.0106)*	*(0.0118)*	*(0.0023)*
(1-λ)	76%	70%	68%	92%	56%	60%	65%	76%	90%	76%	74%	92%
S	0.6024	0.6647	0.7333	0.6161	0.6091	0.7666	0.7258	0.6452	0.7545	0.3666	0.3961	0.3494
S.E.(S)	*(0.0165)*	*(0.013)*	*(0.0114)*	*(0.0087)*	*(0.0102)*	*(0.0073)*	*(0.032)*	*(0.01)*	*(0.0065)*	*(0.0044)*	*(0.0052)*	*(0.0017)*
(B) Stage 2 model												
α	-39%	-52%	-62%	-26%	-23%	-22%	-31%	-35%	-37%	-57%	-64%	-35%
S.E.(α)	*(0.0458)*	*(0.0799)*	*(0.0817)*	*(0.0138)*	*(0.0327)*	*(0.0391)*	*(0.0525)*	*(0.0445)*	*(0.0237)*	*(0.0485)*	*(0.0886)*	*(0.0087)*
β	61%	30%	91%	25%	64%	61%	64%	57%	42%	77%	89%	33%
S.E.(β)	*(0.0465)*	*(0.0797)*	*(0.074)*	*(0.0097)*	*(0.0360)*	*(0.0274)*	*(0.0491)*	*(0.0408)*	*(0.0237)*	*(0.0436)*	*(0.0812)*	*(0.0067)*
(1-α-β)	78%	72%	70%	101%	58%	62%	67%	80%	94%	79%	75%	102%
S	0.5765	0.6590	0.7272	0.6007	0.5876	0.7833	0.7192	0.6378	0.7429	0.3525	0.3930	0.3385
S.E.(S)	*(0.0103)*	*(0.0133)*	*(0.0114)*	*(0.0066)*	*(0.0103)*	*(0.0073)*	*(0.0316)*	*(0.0101)*	*(0.0062)*	*(0.0043)*	*(0.0051)*	*(0.0017)*
π	0.3859	0.4095	0.4184	0.2287	0.3514	0.3553	0.3662	0.3432	0.3567	0.3926	0.4782	0.2604
S.E.(π)	*(0.0015)*	*(0.0015)*	*(0.0018)*	*(0.0013)*	*(0.0014)*	*(0.0012)*	*(0.0018)*	*(0.0013)*	*(0.0017)*	*(0.0014)*	*(0.0017)*	*(0.0007)*
(C) Stage 3 model												
α(ll)	-40%	-79%	-89%	-6%	-24%	-50%	-42%	15%	-19%	-46%	-37%	-51%
S.E.(α(ll))	*(0.1061)*	*(0.2152)*	*(0.2675)*	*(4.5712)*	*(0.0741)*	*(0.0952)*	*(0.1474)*	*(0.3624)*	*(0.0855)*	*(0.0898)*	*(0.1463)*	*(0.0566)*
α(ls)	-15%	-50%	-82%	-55%	-26%	-42%	-96%	-7%	-35%	-36%	-115%	-91%
S.E.(α(ls))	*(0.2886)*	*(0.3366)*	*(0.4843)*	*(1.1993)*	*(0.2107)*	*(0.138)*	*(0.4473)*	*(0.3051)*	*(0.1956)*	*(0.5219)*	*(0.7823)*	*(0.059)*
α(sl)	-234%	-210%	-144%	-35%	-53%	-30%	-21%	-53%	-141%	-319%	-274%	-51%
S.E.(α(sl))	*(0.4575)*	*(0.3357)*	*(0.4958)*	*(0.3703)*	*(0.4944)*	*(0.16)*	*(1.021)*	*(0.1971)*	*(0.1229)*	*(0.3713)*	*(0.5447)*	*(0.0272)*
α(ss)	-82%	-111%	-146%	-31%	-35%	8%	-65%	-34%	-93%	-144%	-315%	-35%
S.E.(α(ss))	*(0.1767)*	*(0.3547)*	*(0.2687)*	*(0.0675)*	*(0.1842)*	*(0.1036)*	*(0.6296)*	*(0.0824)*	*(0.0839)*	*(0.1755)*	*(0.8223)*	*(0.0106)*
β(ll)	65%	102%	115%	26%	72%	95%	85%	27%	22%	88%	67%	65%
S.E.(β(ll))	*(0.1149)*	*(0.2281)*	*(0.2512)*	*(3.8738)*	*(0.0752)*	*(0.0935)*	*(0.1464)*	*(0.3697)*	*(0.0748)*	*(0.083)*	*(0.1347)*	*(0.049)*
β(ls)	30%	102%	110%	69%	77%	89%	123%	53%	32%	49%	135%	103%
S.E.(β(ls))	*(0.2837)*	*(0.328)*	*(0.4778)*	*(1.0294)*	*(0.2045)*	*(0.192)*	*(0.4114)*	*(0.3043)*	*(0.1997)*	*(0.4906)*	*(0.7442)*	*(0.0597)*
β(sl)	263%	251%	183%	40%	65%	73%	56%	60%	155%	347%	304%	57%
S.E.(β(sl))	*(0.4555)*	*(0.3377)*	*(0.477)*	*(0.2875)*	*(0.4846)*	*(0.1576)*	*(1.042)*	*(0.1748)*	*(0.1227)*	*(0.3547)*	*(0.543)*	*(0.0284)*
β(ss)	114%	146%	184%	31%	80%	35%	104%	53%	106%	173%	353%	38%
S.E.(β(ss))	*(0.1709)*	*(0.3706)*	*(0.2595)*	*(0.0134)*	*(0.1787)*	*(0.101)*	*(0.6074)*	*(0.0764)*	*(0.0822)*	*(0.1795)*	*(0.8118)*	*(0.0126)*
S(l)	0.5807	0.4997	0.5622	0.5102	0.5056	0.6039	0.4833	0.6079	0.5334	0.2837	0.3214	0.3035
S.E.(S(l))	*(0.0169)*	*(0.0125)*	*(0.0161)*	*(0.1158)*	*(0.0101)*	*(0.0101)*	*(0.0144)*	*(0.0143)*	*(0.0109)*	*(0.0056)*	*(0.0082)*	*(0.0018)*
S(s)	0.7535	0.7342	0.7876	0.6084	0.8112	0.8693	0.8840	0.6348	0.7901	0.4031	0.4349	0.3439
S.E.(S(s))	*(0.0116)*	*(0.0188)*	*(0.0133)*	*(0.005)*	*(0.0182)*	*(0.0089)*	*(0.0478)*	*(0.0119)*	*(0.0091)*	*(0.0055)*	*(0.0055)*	*(0.0014)*
π(ll)	0.3873	0.4036	0.4203	0.4071	0.3563	0.3382	0.3720	0.4344	0.3767	0.3915	0.4053	0.4080
S.E.(π(ll))	*(0.0002)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0.0001)*	*(0.0001)*	*(0)*
π(ls)	0.4567	0.4579	0.4069	0.4340	0.4369	0.4463	0.4423	0.4602	0.4423	0.4520	0.4086	0.4400
S.E.(π(ls))	*(0.0001)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*	*(0)*
π(sl)	0.4707	0.4555	0.4545	0.3070	0.4650	0.4392	0.4487	0.3951	0.4316	0.4657	0.4764	0.3395
S.E.(π(sl))	*(0.0001)*	*(0)*	*(0.0002)*	*(0.0021)*	*(0)*	*(0)*	*(0.0001)*	*(0.0095)*	*(0)*	*(0)*	*(0)*	*(0.0014)*
π(ss)	0.4559	0.4552	0.4689	0.2997	0.4377	0.4512	0.4642	0.3727	0.4450	0.4629	0.4894	0.3247
S.E.(π(ss))	*(0.0001)*	*(0.0001)*	*(0.0002)*	*(0.0042)*	*(0)*	*(0)*	*(0.0001)*	*(0.0001)*	*(0)*	*(0)*	*(0)*	*(0)*

Table 7.1: The results from the Huang and Stoll(1997) model applied to STIR futures data.
a=adverse selection component
β=inventory component
(1-a-β)=price clustering component
λ=a+β
S=bid-ask spread
π=probability of reversal

	USDJPY		USDCHF		EURUSD		EURJPY		EURCHF	
Time Period	01.01.98 - 04.09.98	01.01.99 - 03.09.99	[illegible] - 04.09.98	01.01.99 - 03.09.99	01.01.98 - [illegible]	01.01.99 - [illegible]	[illegible] - 04.09.98	[illegible] - 03.09.99	01.01.98 - [illegible]	[illegible] - 03.09.99
No. of Obs.	399,124	225,825	42,952	72,939	484,005	310,300	139,054	42,743	73,698	29,654
(A) Stage 1 model										
λ	72%	64%	86%	81%	56%	45%	91%	79%	74%	67%
S.E.(λ)	(0.0309)	(0.0048)	(0.0199)	(0.0080)	(0.0053)	(0.0036)	(0.0039)	(0.0113)	(0.0135)	(0.0123)
(1-λ)	28%	36%	14%	19%	44%	55%	9%	21%	26%	33%
S	0.006751	0.009698	0.000160	0.000148	0.000060	0.000067	0.006563	0.020216	0.000036	0.000055
S.E.(S)	(0.000005)	(0.000054)	(0.000003)	(0.000002)	(0)	(0)	(0.000073)	(0.000319)	(0.000001)	(0.000001)
(B) Stage 2 model										
α	68%	-142%	-2%	-41%	100%	-215%	5%	8%	-33%	-6%
S.E.(α)	(0.0537)	(0.0716)	(0.1097)	(0.0547)	(0.064)	(0.0744)	(0.1498)	(0.0456)	(0.1211)	(0.1316)
β	5%	206%	87%	121%	-44%	260%	86%	71%	107%	73%
S.E.(β)	(0.060607)	(0.059797)	(0.100570)	(0.062519)	(0.050773)	(0.073154)	(0.147768)	(0.043594)	(0.110317)	(0.128033)
(1-α-β)	27%	36%	14%	20%	44%	55%	9%	21%	26%	33%
S	0.006714	0.009490	0.000160	0.000147	0.000060	0.000065	0.006686	0.020267	0.000036	0.000055
S.E.(S)	(0.000005)	(0.000053)	(0.000003)	(0.000002)	(0)	(0)	(0.000074)	(0.000323)	(0.000001)	(0.000001)
π	0.5573	0.4655	0.4038	0.4259	0.5404	0.4736	0.4659	0.3745	0.4416	0.4524
S.E.(π)	(0.0303)	(0.0005)	(0.0011)	(0.0008)	(0.0003)	(0.0004)	(0.0005)	(0.0012)	(0.0008)	(0.0012)

Table 7.2: The results from the Huang and Stoll(1997) model applied to spot FX data.
- Note the EUR in 1999 refers to the euro, but in 1998 it refers to the deutschmark. In order to compare the 1998 and 1999 EUR values, a conversion rate must be introduced – the appropriate rate is the official fixed EUR/DEM conversion exchange rate of 1.95583.
α=adverse selection component
β=inventory component
(1-α-β)=price clustering component
λ=α+β
S=bid-ask spread
π=probability of reversal

In the stage one model, λ represents the combined adverse selection and inventory management bid-ask spread components. My analysis reveals very high λ values for the inter-dealer spot FX market and very low λ values for STIR futures. In other words, under the Huang and Stoll(1997) / Glosten and Harris(1988) model, price clustering makes up the bulk of the bid-ask spread for STIR futures but plays only a small part in the spot FX market. However, the size of the price clustering component in FX bid-ask spreads appears to have risen considerably since monetary convergence. Spot FX bid-ask spreads also appear to have appreciated notably over the same period. At first, these may seem inconsistent, until you factor in that spot FX volume has fallen drastically in this period also. The post-euro fall in λ pervades all 5 FX rates studied, while the increased bid-ask spread is evident in 4 of the 5. USD/CHF is the only exchange rate bid-ask spread which does not rise. Instead, it shows a fall of over 7%. Then again, this is the only one of the 5 currency pairs to experience any increase in volume in 1999.

The biggest single contributor to the change in STIR bid-ask spreads was the change in minimum tick size of the Euromark future. By the measure used in this model, this event caused Euromark bid-ask spreads to fall by 50%. At the same time, the order-processing portion of this bid-ask spread fell from 90% to 76%. Lower bid-ask spreads may have permitted greater levels of inventory management and/or greater exploitation of informed trading. However, it is impossible to be conclusive on this because expectations about the euro would also have had a big influence on investment strategies around this time. Alongside that, Europe's futures exchanges would have been competing to dominate the budding pan-European business.

At stages two and three, the model blows apart. In most cases, the stage 2 model produces strongly negative values for α, while the values for β are frequently well in excess 100%. Without aggregating trades together, π values are persistently below 0.5. In the expanded model, many of the t-statistics are below the 95% critical value. Still following closely in the footsteps of Huang and Stoll(1997), for quote-driven markets, I aggregate sequential STIR futures trades of the same side and price, which occurred within a contiguous two minute time segment. This should re-combine any broken-up large deals that were pre-negotiated in an upstairs market. For order-driven STIR trades, I aggregate trades which occurred at the same second, side and price. The rationale behind is that, in an order-driven market, a market order can take out a number of existing limit orders. On the other hand, if the trader feared that his trade was too large to go through without moving the price, he would be more likely to place a number of limit orders. The spot FX data could not be aggregated in this way because only one observation per timestamp is reported.

	Short Sterling				Euroswiss				Euromark		Euribor	
Time Period	01.01.97 - 19.01.98	20.01.98 - 31.12.98	01.01.99 - 22.08.99	20.09.99 - 31.12.00	01.01.97 - 19.01.98	20.01.98 - 31.12.98	01.01.99 - 22.08.99	20.09.99 - 31.12.00	01.01.97 - 19.01.98	20.01.98 - 31.12.98	01.01.99 - 22.08.99	20.09.99 - 31.12.00
No. of Obs.	15,623	19,749	8,903	63,146	11,146	11,451	7,696	32,717	9,727	14,444	8,740	180,389
(A) Stage 1 model												
λ	103%	110%	111%	7%	112%	109%	109%	20%	86%	105%	99%	7%
S.E.(λ)	*(0.0264)*	*(0.0187)*	*(0.0242)*	*(0.0036)*	*(0.0116)*	*(0.0096)*	*(0.0303)*	*(0.0132)*	*(0.0352)*	*(0.0229)*	*(0.0221)*	*(0.0022)*
(1-λ)	-3%	-10%	-11%	93%	-12%	-9%	-9%	80%	14%	-5%	1%	93%
S	1.2364	1.2618	1.3933	0.6249	1.2238	1.3420	1.3201	0.6524	1.2642	0.6790	0.6787	0.3528
S.E.(S)	*(0.0346)*	*(0.0244)*	*(0.0368)*	*(0.0057)*	*(0.0161)*	*(0.0127)*	*(0.0396)*	*(0.0102)*	*(0.0391)*	*(0.0166)*	*(0.0142)*	*(0.0017)*
(B) Stage 2 model												
α	-21%	-27%	-24%	-29%	-33%	-36%	-27%	-45%	-8%	-14%	-13%	-37%
S.E.(α)	*(0.0189)*	*(0.0515)*	*(0.0199)*	*(0.0134)*	*(0.0296)*	*(0.0165)*	*(0.0494)*	*(0.0413)*	*(0.0263)*	*(0.0236)*	*(0.0198)*	*(0.0087)*
β	113%	126%	123%	26%	139%	138%	129%	61%	84%	111%	105%	34%
S.E.(β)	*(0.0307)*	*(0.0277)*	*(0.0249)*	*(0.0095)*	*(0.0268)*	*(0.0159)*	*(0.0605)*	*(0.0391)*	*(0.0384)*	*(0.0234)*	*(0.0277)*	*(0.0066)*
(1-α-β)	7%	1%	0%	103%	-6%	0%	-1%	65%	24%	3%	8%	103%
S	1.2530	1.2918	1.4022	0.6069	1.2342	1.3488	1.3195	0.6430	1.2464	0.6738	0.6869	0.3408
S.E.(S)	*(0.035)*	*(0.0274)*	*(0.0345)*	*(0.0066)*	*(0.0161)*	*(0.0124)*	*(0.0393)*	*(0.0102)*	*(0.0383)*	*(0.0163)*	*(0.0143)*	*(0.0017)*
π	0.8599	0.8557	0.8581	0.2253	0.7274	0.7491	0.7061	0.3381	0.9025	0.8748	0.8538	0.2589
S.E.(π)	*(0.0027)*	*(0.0027)*	*(0.0031)*	*(0.0013)*	*(0.0022)*	*(0.0019)*	*(0.003)*	*(0.0013)*	*(0.0021)*	*(0.0023)*	*(0.003)*	*(0.0007)*
(C) Stage 3 model												
α(ll)	-36%	-30%	-32%	-12%	-34%	-45%	-23%	27%	-32%	-16%	-22%	-52%
S.E.(α(ll))	*(0.0448)*	*(0.1328)*	*(0.0717)*	*(0.2444)*	*(0.0544)*	*(0.0389)*	*(0.0959)*	*(0.3272)*	*(0.0606)*	*(0.0415)*	*(0.0582)*	*(0.0707)*
α(ls)	-30%	-46%	-53%	-51%	-39%	-22%	-11%	-5%	-14%	-33%	-15%	-92%
S.E.(α(ls))	*(0.0964)*	*(0.0997)*	*(0.0997)*	*(0.133)*	*(0.0792)*	*(0.0492)*	*(0.0752)*	*(0.2935)*	*(0.1027)*	*(0.0761)*	*(0.0953)*	*(0.1239)*
α(sl)	-215%	-132%	-121%	-40%	-532%	-447%	-284%	-71%	-70%	-136%	-141%	-52%
S.E.(α(sl))	*(0.1909)*	*(0.1094)*	*(0.096)*	*(0.0424)*	*(0.6197)*	*(0.2307)*	*(0.6742)*	*(0.1821)*	*(0.1502)*	*(0.1438)*	*(0.1346)*	*(0.034)*
α(ss)	-74%	-131%	-51%	-32%	-217%	-105%	-118%	-41%	-31%	-25%	-55%	-35%
S.E.(α(ss))	*(0.1219)*	*(0.7524)*	*(0.1164)*	*(0.0231)*	*(0.3524)*	*(0.1018)*	*(0.1509)*	*(0.0776)*	*(0.1171)*	*(0.3879)*	*(0.1151)*	*(0.0133)*
β(ll)	127%	144%	147%	28%	146%	157%	140%	12%	113%	127%	121%	64%
S.E.(β(ll))	*(0.0497)*	*(0.0831)*	*(0.0536)*	*(0.2062)*	*(0.0455)*	*(0.0339)*	*(0.1136)*	*(0.3473)*	*(0.0629)*	*(0.0328)*	*(0.0458)*	*(0.0613)*
β(ls)	137%	154%	154%	73%	154%	143%	135%	48%	101%	133%	120%	102%
S.E.(β(ls))	*(0.0877)*	*(0.0808)*	*(0.0828)*	*(0.1292)*	*(0.0611)*	*(0.0372)*	*(0.0605)*	*(0.2926)*	*(0.0832)*	*(0.0648)*	*(0.0768)*	*(0.1121)*
β(sl)	261%	207%	203%	45%	601%	517%	345%	95%	123%	197%	200%	59%
S.E.(β(sl))	*(0.1718)*	*(0.0985)*	*(0.0947)*	*(0.0375)*	*(0.6161)*	*(0.2267)*	*(0.6919)*	*(0.1603)*	*(0.1229)*	*(0.127)*	*(0.1162)*	*(0.0366)*
β(ss)	145%	189%	129%	33%	279%	183%	189%	58%	50%	100%	125%	38%
S.E.(β(ss))	*(0.1196)*	*(0.5413)*	*(0.1174)*	*(0.0168)*	*(0.3226)*	*(0.0917)*	*(0.131)*	*(0.0715)*	*(0.0933)*	*(0.2572)*	*(0.1245)*	*(0.0108)*
S(l)	1.1757	1.2055	1.3222	0.5380	1.1756	1.2519	1.1093	0.5994	1.1223	0.6220	0.6189	0.3021
S.E.(S(l))	*(0.0388)*	*(0.0238)*	*(0.0312)*	*(0.0097)*	*(0.0174)*	*(0.0164)*	*(0.0411)*	*(0.0142)*	*(0.0343)*	*(0.0177)*	*(0.0171)*	*(0.0022)*
S(s)	1.3158	1.3393	1.4349	0.6132	1.3177	1.3974	1.4065	0.6425	1.3709	0.7095	0.7209	0.3453
S.E.(S(s))	*(0.0366)*	*(0.0489)*	*(0.0453)*	*(0.0063)*	*(0.0329)*	*(0.016)*	*(0.1098)*	*(0.0119)*	*(0.0435)*	*(0.0143)*	*(0.0134)*	*(0.0018)*
π(ll)	0.7592	0.7192	0.6948	0.4049	0.5962	0.6771	0.6894	0.4316	0.7506	0.7398	0.7319	0.4062
S.E.(π(ll))	*(0.0003)*	*(0)*	*(0.0002)*	*(0)*	*(0)*	*(0)*	*(0.0001)*	*(0)*	*(0.0002)*	*(0.0001)*	*(0.0001)*	*(0)*
π(ls)	0.7223	0.7192	0.7259	0.4330	0.7025	0.7126	0.7415	0.4581	0.7041	0.7503	0.7365	0.4400
S.E.(π(ls))	*(0.0002)*	*(0.0001)*	*(0.0002)*	*(0)*	*(0)*	*(0.0001)*	*(0.0001)*	*(0)*	*(0.0002)*	*(0.0001)*	*(0.0001)*	*(0)*
π(sl)	0.5732	0.6055	0.6227	0.3356	0.5254	0.5341	0.5528	0.3908	0.6423	0.6035	0.5980	0.3547
S.E.(π(sl))	*(0.0001)*	*(0)*	*(0.0001)*	*(0)*	*(0)*	*(0)*	*(0.0001)*	*(0)*	*(0.0001)*	*(0.0001)*	*(0.0001)*	*(0)*
π(ss)	0.6458	0.6501	0.6529	0.2959	0.5752	0.6074	0.6131	0.3666	0.6711	0.6443	0.6559	0.3219
S.E.(π(ss))	*(0.0001)*	*(0.0001)*	*(0.0002)*	*(0)*	*(0)*	*(0.0002)*	*(0.0001)*	*(0.0001)*	*(0.0002)*	*(0.0001)*	*(0.0001)*	*(0)*

Table 7.3: The results from the Huang and Stoll(1997) model applied to STIR futures data, with trades aggregated where they may be components of a single large trade.
a=adverse selection component
β=inventory component
(1-a-β)=price clustering component
λ=a+β
S=bid-ask spread
π=probability of reversal

The order-driven aggregation makes little difference to the results. Aggregation in the quote-driven market drives π over 0.5, while the values of a remain stubbornly negative and the magnitudes of both a and β become even more extreme, particularly in the expanded models. However, it is impossible to determine whether this π>0.5 is meaningful or not, because the bunching together of same-signed sub-sequences and labelling of these as single occurrences, induces negative serial correlation in any time series regardless of the underlying economics. Therefore, it is impossible to discern whether π>0.5 is real and is confirmed with evidence of recovered big trades, or whether it is spurious, the outcome of an unwarranted bunching exercise.

7.8 Modified Model Results

The first thing to note about the results in table 7.4 is that a and β now behave as the theory predicts. In contrast to the results for the original model, a is positive in every case, and the sum of a and β is always less than 1. Also, all t-statistics are well above the 95% critical value.

	Short Sterling				Euroswiss				Euromark		Euribor	
Time Period	01/01/97 - 19/01/98	20/01/98 - 31/12/98	01/01/99 - 22/08/99	20/09/99 - 31/12/00	01/01/97 - 19/01/98	20/01/98 - 31/12/98	01/01/99 - 22/08/99	20/09/99 - 31/12/00	01/01/97 - 19/01/98	20/01/98 - 31/12/98	01/01/99 - 22/08/99	20/09/99 - 31/12/00
No. of Obs.	23,160	22,986	14,935	67,689	30,320	38,268	18,080	35,875	23,605	29,934	16,462	195,309
(A) Summary Statistics												
Av. Spread	0.9693	0.9680	0.9815	1.0228	1.0262	1.0172	0.9986	1.2083	0.9841	0.4907	0.4859	0.5195
Av. Daily Volume	15,316	22,218	23,008	17,434	7,507	12,322	11,284	7,808	27,458	45,087	28,186	50,373
Av. Volatility (Rtn)												
(B) % Breakdown of Bid-Ask Spread (with Standard Errors)												
α	31%	33%	26%	39%	33%	25%	28%	45%	24%	27%	21%	32%
S.E.(α)	(0.0106)	(0.0132)	(0.0117)	(0.0067)	(0.0099)	(0.0071)	(0.0319)	(0.0084)	(0.0086)	(0.0085)	(0.0103)	(0.0033)
β	16%	20%	23%	5%	29%	30%	25%	12%	7%	16%	20%	5%
S.E.(β)	(0.0097)	(0.008)	(0.0121)	(0.0022)	(0.0069)	(0.006)	(0.0098)	(0.0079)	(0.0054)	(0.0088)	(0.0105)	(0.0015)
(1-α-β)	52%	47%	51%	56%	38%	45%	47%	43%	69%	57%	59%	62%
(C) The Components of the Average Quoted Spread												
α	0.3036	0.3191	0.2593	0.3998	0.3410	0.2501	0.2799	0.5464	0.2392	0.1311	0.1009	0.1682
β	0.1569	0.1937	0.2229	0.0464	0.2987	0.3064	0.2488	0.1413	0.0672	0.0803	0.0961	0.0268
(1-α-β)	0.5088	0.4552	0.4993	0.5766	0.3865	0.4606	0.4699	0.5205	0.6777	0.2794	0.2889	0.3245
(D) % Change in the Components of the Average Quoted Spread												
α		5%	-19%	54%		-27%	12%	95%		-45%	51%	67%
β		23%	15%	-79%		3%	-19%	-43%		19%	134%	-72%
(1-α-β)		-11%	10%	15%		19%	2%	11%		-59%	102%	12%

Table 7.4: The components of the bid-ask spread for STIR futures. All inter-temporal comparisons of the Euribor with the Euromark futures contract have been adjusted by the fixed EUR/DEM conversion rate of 1.95583.
a=adverse selection component
β=inventory component
(1-a-β)=price clustering component

Part B of table 7.4 reveals the percentage components of the bid-ask spread / price innovation for STIR futures. Part C of table 7.4 shows the component percentages from part B multiplied by the average bid-ask spreads from Part A, which reveal the actual bid-ask spread components in amounts. In part D, the change in the bid-ask spreads from part C is shown. (Note that the EUR/DEM conversion exchange rate is applied to all cases where EUR denominated amounts are compared to DEM denominated amounts.)

The advent of electronic trading causes a big change in the respective sizes of the information and inventory bid-ask spread components. The sharp fall in the inventory component, from 21% to 7% on average, may reflect that scalpers did engage in some inventory management, while limit order traders, who actually want the position resulting from a trade, have no such activities at all. If the data is grouped into pre-EMU V post-EMU instead of floor V electronic, the average inventory percentages are 20% and 15% respectively for pre-EMU and post-EMU. The average size of the information component grows from of 28% to 39% as trading gravitates from floor to electronic. This contrasts with corresponding α values of 29% and 32% when the data are grouped as pre and post-EMU.

When comparing the roles of inventory and news as drivers of price innovation, the latter accounts for 85% in the post-electronic period. News had only accounted for only 57% in the floor-trading world. This may signify that informed traders can make better use of their information via limit orders than they were able to when they had to pay away bid-ask spread to scalpers. However, if information about inventory has become a feature of the market, like it is in the spot FX market, this interpretation of the numbers would not hold.

Price clustering is by far the largest component of STIR futures bid-ask spreads, accounting for more than half in most cases. This supports the findings of chapter 5 of the thesis. The fall from 69% to 57% when the minimum tick size was halved is particularly telling. The breakdown in part D reveals that this equates to a 59% fall in the value of the price clustering component itself.

	USDJPY		USDCHF		EURUSD		EURJPY		EURCHF	
Time Period	01/08/98 - 04/09/98	01/08/99 - 03/09/99	01/08/98 - 04/09/98	01/08/99 - 03/09/99	01/08/98 - 04/09/98	01/08/99 - 03/09/99	01/08/98 - 04/09/98	01/08/99 - 03/09/99	01/08/98 - 04/09/98	01/08/99 - 03/09/99
No. of Obs.	399,124	225,825	42,952	72,939	484,005	310,300	128,064	42,743	73,898	29,654
(A) Summary Statistics										
Av. Spread	0.009915	0.010257	0.000231	0.000177	0.000082	0.000073	0.010843	0.024639	0.000071	0.0001
Av. Daily Volume	399,124	225,825	42,952	72,939	484,005	310,300	128,064	42,743	73,898	29,654
Av. Volatility (Rtn)	0.0405%	0.0323%	0.0362%	0.0301%	0.0267%	0.0244%	0.0393%	0.0399%	0.0173%	0.0066%
(B) % Breakdown of Bid-Ask Spread (with Standard Errors)										
α	29%	7%	29%	17%	25%	9%	36%	18%	45%	33%
S.E.(α)	*(0.0046)*	*(0.0045)*	*(0.0133)*	*(0.0084)*	*(0.0039)*	*(0.0036)*	*(0.0069)*	*(0.012)*	*(0.0088)*	*(0.0098)*
β	43%	61%	56%	70%	38%	42%	50%	65%	37%	45%
S.E.(β)	*(0.0052)*	*(0.0056)*	*(0.0149)*	*(0.01)*	*(0.0044)*	*(0.0039)*	*(0.007)*	*(0.0117)*	*(0.008)*	*(0.0113)*
$(1-\alpha-\beta)$	28%	32%	15%	13%	37%	49%	13%	16%	18%	21%
(C) The Components of the Average Quoted Spread										
α	0.002900	0.000732	0.000068	0.000031	0.000020	0.000007	0.003950	0.004549	0.000032	0.0000
β	0.004233	0.006277	0.000130	0.000123	0.000031	0.000031	0.005460	0.016035	0.000027	0.0000
$(1-\alpha-\beta)$	0.002782	0.003247	0.000034	0.000023	0.000030	0.000036	0.001433	0.004055	0.000013	0.0000
(D) % Change in the Components of the Average Quoted Spread										
α		-75%		-54%		-40%		52%		2%
β		48%		-5%		81%		287%		68%
$(1-\alpha-\beta)$		17%		-32%		118%		273%		65%

Table 7.5: The components of the bid-ask spread for spot FX. In the column headings, the currency code EUR refers to the euro and to its predecessor, the deutschemark. The EUR rates comparison in part D has been adjusted by the fixed EUR/DEM conversion rate of 1.95583.
a=adverse selection component
β=inventory component
(1-a-β)=price clustering component

Compared with the original model results for spot FX, where β often accounted for more than 100% of the bid-ask spread, the modified model's β component accounts for an average of around 45% of the bid-ask spread in 1998 and 57% in 1999. In every spot FX case, the 1999 β value is larger than that in 1998. Furthermore, in all but one case, the inventory component is bigger than either the information component or the price clustering component. This corroborates the finding of the previous chapter that asymmetric information alone is not the main driver of price innovations or of the bid-ask spread.

The 1999 results show a sharp decline in *a* from 33% in 1999 to 17% in 1998. This suggests that "informed" FX limit-order traders are less well able to predict and profit from futures price moves since currency convergence than they were before. It seems unlikely that there would be a marked difference either in the macroeconomic fundamentals or in the ability to interpret those fundamentals, between the two sample periods. This suggest that the ability of informed traders to aggregate and interpret information linked to future accumulated inventory is the most likely source of the stark decline in the information component of the bid-ask spread, in the lower-volume post-EMU period. Equivalently, it also shows why the importance of inventory has risen after EMU.

Echoing a central conclusion from chapter 5, part D of table 7.5 shows that by far the largest change in the components of the EUR(DEM)/USD since EMU was a 118% increase in the price clustering component.

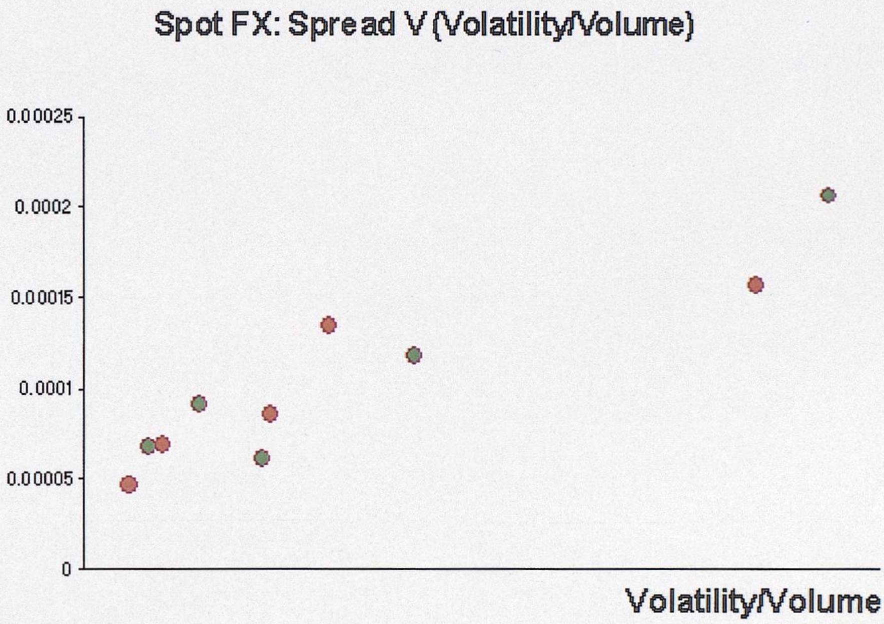

Figure 7.1: The relationship between FX bid-ask spreads (as a % of the FX rate) and FX Volatility/Volume is upward sloping.

Figure 7.1 provides independent evidence from the empirical data which supports my treatment of the bid-ask spread in order-driven markets as different from that of quote driven. This graph shows that the average spot FX bid-ask spreads are positively related to volatility/volume, as predicted by the Black-Hartmann bid-ask spread model for order-driven markets. The data comes from both 1998 and 1999. The 1998 observations are shown in orange and the 1999 observations are in green.

7.8.1 How to Interpret "Information" and "Inventory" for Spot FX

In a conventional trade indicator model, knowing the component percentages of price change or, equivalently of the bid-ask spread, which is attributable to inventory and to asymmetric information conveys two distinct pieces of information. First, it shows how much of price change is driven by news and by inventory respectively. Second, it shows how much of the profit generated from bid-ask spreads goes to market makers and to informed traders respectively. This is because informed traders can only act on private information about future news and because market makers react to inventory and adverse selection risk. The third component of the conventional trade indicator model, the price clustering cost, constitutes additional revenue for the market maker. In short, where informed trading is driven only by news, the 'who profits from trading?' question is perfectly aligned with the 'what drives price?' question.

In the spot FX market, much, if not most, of asymmetric information is believed to refer to future inventory. This means that the numbers in table 7.5 illustrate how the profit from the bid-ask spread divides between limit-order trader and market order trader. The residual component, due to price clustering, represents a windfall gain for limit order liquidity providers. However, in this case, β can be interpreted as linking to unanticipated inventory. Therefore, the lower bound for the influence of inventory on price change may be computed as $\beta/(\beta + \alpha)$. The corresponding upper bound is 100% ($=(\beta + \alpha)/(\beta + \alpha)$), since α represents a combination of asymmetric information about inventory and asymmetric information about news and the presence of the latter is unconfirmed. By this measure, in the pre-EMU period, on average, inventory accounts for at least 58% of order-driven price innovation, but could actually be a lot higher. After EMU, the average minimum estimate rises to 77%. In both cases, inventory clearly dominates news as the principal driver of trading induced price change. Furthermore, if at least half of informed trading could be attributed to information about inventory, the lower bound for the influence of inventory on price would rise to $(\beta + 0.5\alpha)/(\beta + \alpha)$, which equates to 79% in the pre-EMU era and 88% post-EMU.

Changes in price consist of changes in the transaction price as a result of trading and revisions to the mid-quote. The original trade indicator model only explains the components of former, leaving changes in the mid-quote which are not linked to transactions to be captured by the error term. By contrast, the modified trade indicator model captures all of the price movement. Under the Bloomfield et al.(2003) model, informed traders release information by revising the mid-quote. The separation between mid-quote revision and transaction price revision no longer exists. This means that the numbers given in the preceding paragraph attributing the proportion of spot FX rate change to news and inventory refer to the total rate change, not just the transaction part.

One final issue that needs to be considered is the influence or meaning of price clustering for price innovations. Price clustering will clearly introduce error into the price innovation process. However, the most interesting issue is whether that error will dissipate or linger. As mentioned in chapter 5, some research links price discreteness to path dependence. This may be evidence of a lingering affect. However, there is nothing in the present analysis that can clarify this matter further.

7.8.2 Some Comments on How Inventory Drives Order Flow

The previous section shows that inventory is the primary driver of order flow and that news plays a small part, at most. The Bloomfield et al(2003) model provides an insight into how this occurs. They find that the informed trader's decision of whether to place a market order or a limit order is determined by the size of the expected price innovation, i.e. the difference between the current price and what the informed trader believes the true price to be. This feature means that we can think of informed traders as a single group which has two ways to implement its price revision information. For small future price innovations, the informed traders will revise their prices via limit orders because they do not wish to give away the spread. For large innovations, they will use market orders to engage limit orders from the uninformed and from those not quick enough to revise or withdraw them.

My own direct experience can corroborate the above. I observed firsthand that fundamental news is rare in the spot FX market and that traders are cautious about interpreting it. They will take a guess but they will also closely scrutinise the early order flow and price reaction very carefully to evaluate the significance of a macro event to the market. On the other hand, one source of incoming inventory is both frequent and easy to read. This is because of a common practice whereby customers who wish to do a trade to ring around the market to see where they can get the best price. This alerts several dealers to the fact that there is a customer order coming in. It is customary for the customer to disclose the size of the transaction they wish to effect, but not the side, i.e. the customer asks for a two-way price. However, the job of the FX salesperson who deals with that customer is to know the customer well and to know his/her investment objectives. This enables the salesperson to make an educated guess as to the side of the incoming trade. This prediction of the customer's trade is a source of private information about incoming inventory. If the prediction says that this inventory will drive the price up, then the informed trader can take out existing limit orders and stock up on a currency with the expectation of unloading it at a higher price as soon as the inventory arrives. This stocking up process drives the price up, closing the gap between the pre-information level and the level the inventory induces.

7.9 Conclusions

Results from the modified trade indicator corroborate the key finding of the previous chapter, that asymmetric information can not be the dominant determinant of prices or of bid-ask spreads. In every case, asymmetric information comprises less than 50% of price innovations and of the bid-ask spreads, while inventory and price clustering together always make up more than 50%.

The modified trade indicator model proved more appropriate for an order-driven environment that the original model. The former, which makes extensive use of new theory from Bloomfield et al(2003), produces reasonable results for all instruments and time periods. This contrasts with wildly implausible results produced by the original model. The principal reason for these extreme results proved to be a key assumption in quote-driven market microstructure models: mean-reverting price behaviour induced by inventory. However, this feature does not fit with the microstructure theory specific to order-driven markets.

In all cases, the inventory component is shown to have a large and lasting impact on price, flatly contradicting the notion that lasting price perturbations can only arise from the information component of the bid-ask spread. The contribution of the information component to both FX rates and bid-ask spreads is considerably lower in 1999, and the inventory component is correspondingly higher. This is probably because the disappearance of the intra-European crosses and reduced volume in the remaining currency pairs together make it harder to any single dealer to form a better overall picture of currency demand from aggregate order flow than any other dealer. Without that information advantage, it may be more difficult to price limit orders accurately and so to profit from them.

Aside from the Euribor, STIR futures bid-ask spreads and prices do not seem particularly perturbed by EMU. That said, the role of currency convergence in nudging the futures markets towards electronic trading can not be dismissed.

The price clustering bid-ask spread component showed little variation in its percentage contribution over the whole period studied. This component alone constitutes the major part of the STIR futures bid-ask spread but is noticeably less important for spot FX.

The Black-Hartmann bid-ask spread model fits the attributes of a bid-ask spread model that my changes to the trade indicator model require. A positive relationship between bid-ask spread and (volatility/volume) was shown for spot FX, confirming the appropriateness of this model for the spot FX and STIR futures markets and, in particular, for the datasets studied here.

The respective importance of inventory and news as drivers of spot FX rates is revealed by refining the definitions of the terms "inventory" and "information" in the context of a known feature of the market. This feature refers to the fact that, in this market, asymmetric information can often relate to future inventory. This reinterpretation shows that inventory dominates news far more completely than first appears, probably accounting for between 88% and 100% of total price innovation in the post-EMU period.

Chapter 8.

Conclusion

8.1 Contributions

The primary goal of this dissertation has been to answer two big questions in financial economics: 1) what drives price?, and 2) what determines the bid-ask spread? Chapters 5 and 6 analysed competing candidate explanations. Then chapter 7 revealed the contribution of each potential explanatory factor to each question. In addition, opportunities arose along the way to make some specific theoretical and methodological contributions which will hopefully have value beyond the immediate subject and markets studied here. My contributions in each area are summarised below.

8.1.1 Main Empirical Findings

Chapter 5 found that price discreteness exerts a strong influence over the bid-ask spread in both the spot FX and the STIR futures markets. In fact, the bid-ask spread's most frequent value was the minimum tick size, for most instruments, most of the time. For STIR futures, the extreme dominance of bid-ask spread values at the minimum tick size led me to conclude that the minimum tick size is probably too high and should be lowered. Confirming previous empirical research, I found that the spot FX price clustering evidence was consistent with the Ball et al(1985) price resolution hypothesis. I reveal the price clustering pattern of STIR futures for the first time and show that it is linked to the price concentration hypothesis. In addition, in contrast to the view proposed by Hau et al(2002) that higher inter-dealer bid-ask spreads had caused spot FX volume to fall, I conclude the opposite causality. Insofar as the inter-dealer bid-ask spread has risen on average, I find that non-synchronous pricing, which is a direct consequence of lower volumes, inflates the apparent bid-ask spread in an order-driven market. In contrast to the piecemeal view that had been presented in previous research, my extensive datasets enabled me to present a broad overview of what has happened to the bid-ask spread, volume and volatility since EMU. Finally, I demonstrate that re-denomination alone increases bid-ask spread costs by 74% because this is the amount by which the EUR/USD one-tick bid-ask spread exceeds that of USD/DEM.

In chapter 6, I showed that asymmetric information alone could not explain the much observed peculiar patterns in the intra-day data for several variables. Furthermore, this analysis showed that asymmetric information could not be the primary driver of spreads and that, if it is a factor at all, it is drowned out by competing forces at every instance. I concluded that the Brock and Kleidon(1992) model best explained observed U-shaped patterns in the bid-ask spread, but that something else is required to explain the other persistent intra-day features.
I showed that order flow and volume are very highly positively correlated, at least in these datasets, and that the observed order flow is probably the result of random imbalances in supply and demand levels. From this, I deduced that inventory was a strong potential candidate to explain the persistent intra-day irregularities.

Chapter 7 reveals the anatomy of prices and bid-ask spreads in terms of the inventory, asymmetric information and price clustering. This analysis confirmed a key conclusion of chapter 6 – that asymmetric information alone could not be the dominant driver of price. In every market, instrument and sample period, asymmetric information accounted for less than half of each price innovation and bid-ask spread. Inventory and price clustering together proved to be the main contributors to price innovations and bid-ask spreads. Contrary to conventional expectations, inventory is shown to have a large and lasting influence on price. Inventory is the biggest factor contributing to spot FX prices and bid-ask spreads, while price clustering is the biggest factor for STIR futures. The latter affirms one of the central conclusions from chapter 5. For spot FX, inventory is shown to be a much larger influence than news in motivating price change. Finally, the level of asymmetric information present in the spot FX market is much lower after EMU than it was before.

8.1.2 Theoretical Contributions

The price concentration hypothesis presented in chapter 5 is a totally new way of viewing price clustering. This hypothesis reflects the notion that the minimum tick size is too high, or, equivalently that the price resolution is too low to allow the natural patterns of final-digit price clustering to emerge. Evidence supporting the price concentration hypothesis was found in the STIR futures data.

In chapter 6, the synthesis of asymmetric information theories which explain intra-day patterns into a set of consistent hypotheses is new. Also in this chapter, the specification of how order flow erodes return is original. The latter arises because price innovations should be independent but the early release of some information has the effect of dispersing and mixing these uncorrelated increments, lowering the average price change.

I introduce a new explanatory factor, price clustering, into an established model, the trade indicator model, in chapter 7. This factor was mistakenly omitted from early versions of this model.

8.1.3 Methodological Contributions

In chapter 5, I present two new test statistics which measure the goodness of fit of observed patterns in final price digits with the patterns predicted respectively by two of the key competing explanations for the phenomenon, the attraction and resolution hypotheses.

Chapter 6 presents a new use for the correlation matrix as a means to testing multiple coincident hypotheses.

I develop a new "modified" trade indicator model in chapter 7, which accommodates acknowledged features of order-driven markets, producing a much better fit in the data that I analyse than the original model was able to achieve. Furthermore, this model appears appropriate for various other market types where "regret-free" bid and ask priced are not a realistic assumption.

8.2 Implications

Some implications of my finding are discussed below.

8.2.1 What it Means for the Big FX Puzzles

The work presented in this thesis contributes to the debate surrounding two of the three big puzzles in foreign exchange economics: the excess volatility puzzle and the determination puzzle.

My research suggests that temporary imbalances between supply and demand perturb price and that this makes a significant contribution to observed price change volatility. The magnitude of these imbalances is observed to be highly correlated with volume. A plausible explanation for this observed fact is the idea that buy volume and sell volume might be independently but identically distributed as originally suggested by Garman(1976). Over time, these would average to zero but, at any individual instance, a large imbalance, and consequent price deviation, could arise. The empirical evidence that order flow volatility is positively related to return volatility rules out asymmetric information about future price innovations as a key driver of order flow, implying that informed trading can not be responsible for excess spot FX volatility. Furthermore, when the much quoted anecdote, that in the spot FX market information relates to future inventory, is taken into consideration, it appears that inventory is the principal driver of this volatility, possibly leaving no role at all for fundamental based news. However, the role of price discreteness and clustering can not be overlooked. These contribute to volatility by forcing prices to more extreme values than might be chosen if a finer price resolution was available. The component decomposition in chapter 7 reveals that the price clustering factor does indeed play a significant role.

What these combined forces predict for post-EMU spot FX volatility is unclear. On the one hand, lower volume means the variance of random order flow will be lower. This, in conjunction with lower price clustering, should mean lower volatility. On the other hand, the general decrease in the magnitude of information component should lead to higher return volatility. In contrast to every other commentator on this subject that I am aware of, my data shows spot FX volatility to be lower after EMU than it was before.

The exchange rate determination puzzle is obviously closely linked to the excess volatility puzzle in that prices contain noise. That said, the important issue is how the net-of-noise price evolves. We know from the research of Mark(1995), Flood and Taylor(1996) and Froot and Ramadorai(2002) that there is a link between long-run exchange rate changes and fundamentals. If the residuals in these models can be shown to have errors which are sustained for no more than short periods, then a combined fundamentals and inventory story, in conjunction with certain features of the market structure like stop-loss/take-profit orders and price discreteness induced path dependence, should prove sufficient to fully explain exchange rate determination. On the other hand, if these errors prove to be sustained for long periods, then another explanation may be needed. The fact that substantial parts of the pricing process that can be explained by inventory and fundamentals undermines the credibility of the random walk as an explanation of the remainder. Furthermore, the Froot and Ramadorai(2001) finding that offshore-quoted closed-ended country funds exhibit similar innovations in demand to those exhibited by the onshore equity market is offered by those authors as evidence that information, not inventory, must be at the root of changes in portfolio flows / order flow / exchange rates. However, as I pointed out in chapter 6, for private information to have meaning, it must be associated with a subsequent public information release. This implies a negative relationship between order flow and price change. The opposite is consistently observed empirically. There is one obvious alternative explanation which could accommodate the commonality of behaviour that Froot and Ramadorai(2001) detect but that also does not dissipate future returns through leakage as I require. This alternative is false news or widespread mistaken beliefs about future price innovations.

8.2.2 Discoveries about Order Flow

My work reveals two things about order flow that previously appear to have gone unrecognised. These are: 1) asymmetric-information-driven order flow erodes return because it takes price impact from one place and mixes it with other, uncorrelated, price impact, and 2) observed order flow is noisy and the variance of that noise is a positive function of traded volume. The last point is more of a re-discovery than a new discovery, since Garman's(1976) model did assume this kind of order flow behaviour. However, the point seems to have been overlooked in more recent literature.

8.2.3 What Difference Did EMU Make?

The obvious changes were that volume fell in the spot FX market and by more than just the legacy currency internal cross-rate volumes. Also, huge volume flooded into the Euribor. The bid-ask spread widened for most spot FX transactions but showed little movement for STIR futures contracts, except when the minimum tick size was adjusted. By most statistical measures, the level of price clustering has fallen since EMU. There are no remarkable changes in how variables relate to each other on an intra-day basis. However, the component study did indicate that asymmetric information is a smaller feature of price innovation after EMU than it was before. Furthermore, the evidence that inventory, not news, drives almost all exchange rate movement is greater after EMU.

8.2.4 How Did Quote-Driven Models Fit in an Order-Driven World?

The trade indicator required substantial modification in order to fit the order-driven world. On the other hand, the theory underlying the study of intra-day patterns did not clash with the fact that the markets were order-driven per se. It may first appear that the positive order flow volatility and positive volume-volatility show that the model does not fit the environment. However, a positive relationship between volume and volatility is the norm, even in equity markets, and the fact that the theory predicts a negative relationship between order flow and return was not previously recognised.

8.3 Avenues for Future Research

One of the most pressing things to establish in the near future is how the residual errors from the models of Mark(1995), Flood and Taylor(1996) and Froot and Ramadorai(2002) dissipate. If it can be shown that they do not dissipate quickly, then the question of what drives sustained error in the spot FX market takes centre stage. If it is necessary to go down this route, I expect a combination of false news and mistaken beliefs to provide an explanation for persistent errors. I can not see why inventory imbalances should accumulate in any very prolonged way. On the other hand, if the residual error in these models does dissipate quickly, this would support an inventory, price clustering plus fundamentals explanation.

The link between price clustering / discreteness and path dependence warrants further investigation. To the best of my knowledge, the Hausman et al(1992)ordered probit model has never been applied to either the spot FX market or the STIR futures markets.

Another interesting question relates to the comparative effect of market structure on price formation. In the US equity market, on which most market microstructure theory is based, market makers are not only present but have prolonged capital exposure because they carry inventory, often for sustained periods. This gives them a vested interest in price stability that is missing in markets like the spot FX market and the STIR futures. Could that vested interest lead market makers to dampen price volatility and to curb deviations from fundamental fair value? It is obvious that market maker capital will automatically absorb some price volatility because buy inventory and sell inventory will net off without coming to market. On the other hand, if inventory holding market makers are absent, each inventory imbalance would require its own price concession to be absorbed by the general market.

Also, an analysis of friction in spot FX and STIR futures markets, along the lines of Stoll's(2000) study of friction in US equity markets, may introduce a different perspective on the microstructural features of these markets.

Finally, my research indicates that much theoretical work remains to be done to fully attune market microstructure to markets other than the US equity markets.

When it comes, this work will necessitate new or adapted empirical methodologies which, in turn, will facilitate and elicit new empirical analyses.

References

Abhyankar, A.H., L.S. Copeland and W. Wong, 1995, "Liffe cycles: intraday evidence from the FT-SE 100 stock index futures market", European Journal of Finance, 5: 123-139

Abhyankar, A.H., D. Ghosh, E. Levin and R.J. Limmack, 1997, "Bid-ask spreads, trading volume and volatility: intra-day evidence from the London Stock Exchange", Journal of Business Finance and Accounting, 24(3/4), April: 343-362

Admati, A.R., and P. Pfleiderer, 1988, "A theory of intraday patterns: volume and price variability", Review of Financial Studies, 1(1), Spring: 3-40

Ahn, H-J., K-H. Bae and K. Chan, 1999, "Limit orders, depth and volatility: Evidence from the Stock Exchange of Hong Kong", Journal of Finance(forthcoming)

Aitken M., P. Brown, C. Buckland, H.Y. Izan and T. Walter, 1996, "Price clustering on the Australian Stock Exchange", Journal of Pacific Basin Finance, 4: 297-314

Allen, H., and M. Taylor, 1992, "The use of technical analysis in the foreign exchange market", Journal of International Money and Finance, 11: 304-314

Amihud, Y.and H. Mendelson, 1980, "Dealership market: market making with inventory", Journal of Financial Economics, 8, 31-53

Andersen, T., T. Bollerslev and A. Das, 2001, "Variance-ratio statistics and high-frequency data: Testing for changes in intraday volatility patterns", Journal of Finance, 56: 305-328

Andersen, T.G. and T. Bollerslev, 1997, "Intraday periodicity and volatility persistence in financial markets", Journal of Empirical Finance, 4(2-3), June: 115-158

Andersen, T.G. and T. Bollerslev, 1998, "Deutsche mark dollar volatility: intraday activity patterns, macroeconomic announcements and longer run dependencies", Journal of Finance, 53: 219-265

Andersen, T., T. Bollerslev, , F. Diebold, and P. Labys, 2001, "The distribution of realized exchange rate volatility", Journal of the American Statistical Association, 96, 42-55

Angel, J. J., 1994, "Limit versus market orders", Mimeo, School of Business Administration, Georgetown University

ap Gwilym, O. and E. Alibo, 2003, "Decreased price clustering in FTSE100 futures contracts following a transfer from floor to electronic trading", Journal of Futures Markets, 23: 647-659

ap Gwilym, O., M. Buckle, T. Foord, and S.H. Thomas, 1996, "The intra-day behaviour of European bond futures", Journal of Fixed Income, 6(2), Winter: 7-17

ap Gwilym, O., M. Buckle, and S.H. Thomas, 1999, "The intraday behaviour of key market variables for LIFFE derivatives". In Financial Markets Tick by Tick, edited by P. Lequeux, John Wiley and Sons, Chichester: 151-189

ap Gwilym, O., A. Clare and S. Thomas, 1998a, "Extreme price clustering in the London equity index futures and options markets", Journal of Banking and Finance, 22, 1193-1206

ap Gwilym, O., A. Clare and S. Thomas, 1998b, "Price clustering and bid-ask spreads in international bond futures", Journal of Banking and Finance, 8: 377-391

ap Gwilym, O., D. McMillan and A. Speight, 1999, "The intraday relationship between volume and volatility in LIFFE futures markets", Applied Financial Economics, 9: 593-604

ap Gwilym, O., and C. Sutcliffe, 1999, "High-frequency financial market data: sources, applications and market microstructure", Risk Books

Atkins, A.B., and S. Basu, 1995, "The effect of after-hours announcements on the intraday U-shaped volume pattern", Journal of Business Finance and Accounting, 22(6), Summer: 789-809

Bagehot, W. (a.k.a. Jack Treynor), 1971, "The only game in town", Financial Analysts Journal, 27, 12-14, 22.

Bagwell, L. S., 1992, "Dutch Auction Repurchases: An Analysis of Shareholder Heterogeneity." Journal of Finance, 47, 71–105.

Baillie, R.T. and T. Bollerslev, 1990, "Intraday and inter-market volatility in foreign exchange rates", Review of Economic Studies, 58: 565-585

Ball, C. A., and T. Chordia, 2001, "True spreads and equilibrium prices". Journal of Finance, 56, 1801–1835.

Ball C.A., W.A. Torous and A.E. Tschoegl, 1985, "The degree of price resolution: the case of the gold market", Journal of Futures Markets, 5: 29-43

Bank of England, 2001, "The foreign exchange and over-the-counter derivatives markets in the United Kingdom", (http://www.bankofengland.co.uk/statistics/bis-survey/fxotcsum01.pdf)

Becker, K.G., J.E. Finnerty and K.J. Kopecky, 1993, "Economic news and intraday volatility in international bond markets", Financial Analysts Journal, 49(5-6), May-June: 81-6, 65.

Benos, A., and M. Rockinger, 2000, "Market response to earnings announcements and interim reports: an analysis of SBF 120 companies", Annales d'Economie et de Statistique, 60, October/December,

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

BIS (Bank of International Settlements), 1998, "Central bank survey of foreign exchange and derivative market activity in April 1998"

BIS (Bank of International Settlements), 2001, "Central bank survey of foreign exchange and derivative market activity in April 2001"

Black, S., 1991, "Transaction costs and vehicle currencies", Journal of International Money and Finance, 56, 512-527

Bloomfield, R., M. O'Hara and G. Saar, 2003, "The 'make or take' decision in an electronic market: evidence on the evolution of liquidity", Mimeo, NYU

Brock, W.A., and A.W. Kleidon, 1992, "Periodic market closure and trading volume: a model of intraday bids and asks", Journal of Economic Dynamics and Control, 16(3&4): 451-489

Brockman, P., and D. Y. Chung, 1998, "Inter and intra-day liquidity patterns on the Stock Exchange of Hong Kong", Journal of International Financial Markets, Institutions and Money, 8(3-4), November: 277-298

Brown P., A. Chua and J. Mitchell, 2002, "The influence of cultural factors on price clustering: Evidence from Asia-Pacific stock markets", Pacific Basin Finance Journal, 10: 307-332

Brown, S., P. Laux and B. Schachter, 1991, "On the existence of an optimal tick size", Review of Futures Markets, 10: 50-72

Buckle, M., O. ap Gwilym, S.H. Thomas, and M.S. Woodhams, 1998, "Intraday empirical regularities in interest rate and equity index futures markets and the effect of macroeconomic announcements", Journal of Business, Finance and Accounting, 25(7/8), September-October: 921-944

Campbell, J., A. Lo, and C. MacKinley, 1992, "The econometrics of financial markets", Princeton University Press

Cassel, G., 1916, "The present situation in the foreign exchanges", Economic Journal, 62-65

Cassel, G., 1918, "Abnormal deviations in international exchanges", Economic Journal, 413-415

Chan, K., K.C. Chan and G.A. Karolyi, 1991, "Intraday volatility in the stock index and stock index futures markets", Review of Financial Studies, 4(4): 657-684

Chan, K., Y.P. Chung and H. Johnson, 1995, "The intraday behaviour of bid-ask spreads for NYSE stocks and CBOE options", Journal of Financial and Quantitative Analysis, 30(3), September: 329-346

Chan, K.C., W.G. Christie and P.H. Schultz, 1995, "Market structure and the intraday pattern of bid-ask spreads for NASDAQ securities", Journal of Business, 68(1), January: 35-60.

Chang, E.C., P.C. Jain and P.R. Locke, 1995, "Standard and Poor's 500 index futures volatility and price changes around the New York Stock Exchange close", Journal of Business, 68, 1: 61-84

Cheung, Y.W., and L.K. Ng, 1990, "The dynamics of S&P500 index and S&P500 futures intraday price volatilities", Review of Futures Markets, 9(2): 458-486

Choi, J. Y., D. Salandro, and K. Shastri, 1988, "On the Estimation of Bid-Ask Spreads: Theory and Evidence," Journal of Financial and Quantitative Analysis, 23, 219–230.

Christie, W.G. and P.H. Schultz, 1994, "Why do Nasdaq market makers avoid odd-eighth quotes?", Journal of Finance, 49: 1813-1840

Christie, W.G., J.H. Harris and P.H. Schultz, 1994, "Why did Nasdaq market makers stop avoiding odd-eighth quotes?", Journal of Finance, 49: 1841-1860

Clark, P.K., 1973, "A subordinated stochastic process model with finite variance for speculative prices", Econometrica, 41:. 1149-1168

Cohen, K. J., S. F. Maier, R. A. Schwartz and D. K. Whitcomb, 1981, "Transaction costs, order placement strategy, and existence of the bid-ask spread", Journal of Political Economy 89, 287-305.

Cook, T., and R. Laroche, ed.,1993, "Instruments of the Money Market", Federal Reserve Bank of Richmond.

Copeland, T., 1976, "A model asset trading under the assumption of sequential information arrival", Journal of Finance, 31: 1149-1168

Copeland, T.C. and D. Galai, 1983, "Information effects on the bid-ask spread", Journal of Finance, 38, 1457-1469

Cornell, B., 1978, "Determinants of the bid-ask spread on forward foreign exchange contracts under floating exchange rates", Journal of International Business Studies, 9: 33-41

Dacarogna, 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

Daigler, R., 1993, "Financial Futures Markets", Harper Collins

Daigler, R.T., 1997, "Intraday futures volatility and theories of market behaviour", Journal of Futures Markets, 17(1), February: 45-74

Danielsson, J., and R. Payne, 2001, "Measuring and explaining liquidity on an electronic limit order book: evidence from Reuters D2000-2", presented at European Finance Association, January, Barcelona Meetings

DeJong F. and M.W.M. Donders, 1998, "Intraday lead-lag relationships between the futures, options and stock market", European Finance Review, 1: 337-359

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

Detken, C. and P. Hartmann, 2002, "Features of the euro's role in international financial markets", Economic Policy

Diamond, D.W. and R.E. Verrecchia, 1987, "Constraints on short-selling and asset price adjustment to private information", Journal of Financial Economics, 18: 277-311

Ding, D.K., 1999, "The determinants of bid-ask spreads in the foreign exchange futures market: a microstructure analysis", Journal of Futures Markets, 19(3), May: 307-324

Docking, D.S., I.G. Kawaller and P.D. Koch, 1999, "Mid-day volatility spikes in US futures markets", Journal of Futures Markets, 19(2), April: 195-216

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

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(2): 577-605

Ekman, P.D., 1992, "Intraday patterns in the S&P500 index futures market", Journal of Futures Markets, 12(4): 365-381

Engel, C., 1996, "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence", Journal of Empirical Finance, 3:. 123-192

Engel, C., 2000, "Long Run PPP May Not Hold After All", Journal of International Economics

Epps, T., and M. Epps, 1976, "The stochastic dependence of security price changes and transaction volumes: Implications for the mixture of distributions hypothesis", Econometrica, 44: 305-321

Evans, M., 2001, "Foreign exchange trading and exchange rate dynamics", National Bureau of Economic Research, working paper.

Evans, M., and R. Lyons, 2002, "Order flow and exchange rate dynamics", Journal of Political Economy, 110, 170-180

Federal Reserve Bank of New York, 2001, "The foreign exchange and interest rate derivatives market survey: turnover in the United States", (http:// www.ny.frb.org/pihome/triennial/fx_survey.pdf)

Fisher, I., 1930, "The theory of interest", Macmillan, NewYork

Fleming, J. M., 1962, "Domestic financial policies under fixed and under floating exchange rates", IMF Staff Papers, 369-379

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

Flood, M., and A. Rose, "Fixed exchange rates: A virtual quest for fundamentals", Journal of Monetary Economics, 36: 3-37

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

Foster, F.D., and S. Viswanathan, 1993, "Variations in trading volume, return volatility and trading costs: evidence on recent price formation models", Journal of Finance, 48(1), March: 187-211

Foucault, T., 1999, "Order flow composition and trading costs in a dynamic limit order market", Journal of Financial Markets, 2:. 99-134

Foucault, T., O. Kadan, and E. Kandel, 2001, "Limit order book as a market for liquidity", Mimeo, HEC School of Management

Franses, P.H., R.van Ieperen, M. Martens and B. Menkveld, 1997, "Volatility transmission and patterns in bund futures", Journal of Financial Research, 20: 459-482

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

French, K., and R. Roll, 1986, "Stock return variances; the arrival of information and the reaction of traders", Journal of Financial Economics, 17: 5-26

Froot, K.A., J.F. Gammill and A.F. Perold, 1990, "New trading practices and the short-run predictability of the S&P500", in Market Volatility and Investor Confidence, New York Stock Exchange

Froot, K., and T. Ramadorai, 2001, "The information content of international portfolio flows", NBER working paper no. 8472

Froot, K., and T. Ramadorai, 2002, "Currency returns, institutional investor flows and exchange rate fundamentals", NBER working paper no. 9080

Galati, G., and K. Tsatsaronis, 2001, "The impact of the euro on Europe's financial markets", Bank of International Settlements, working paper.

Gannon, G.L., 1994, "Simultaneous volatility effects in index futures", Review of Futures Markets, 13(4): 1027-1068

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

George, T. J., G. Kaul, and M. Nimalendran, 1991, "Estimation of the Bid-Ask Spreads and its Components: A New Approach," Review of Financial Studies, 4, 623–656.

Gerety, M.S. and J.H. Mulherin, 1992, "Trading halts and market activity: an analysis of volume at the open and the close", Journal of Finance, 47(5), December: 1765-1784

Glosten, L., 1994, "Is the electronic open limit order book inevitable?", Journal of Finance, 49: 1127-1161

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

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

Goldstein, M., and K. Kavajecz, 2000, "Eights sixteenths and market depth: Changes in tick size and liquidity provision on the NYSE", Journal of Financial Economics, 56: 125-149

Goodhart, C., and R. Curcio, 1991, "The clustering of bid-ask prices and the spread in the foreign exchange market", London School of Economics, Financial Markets Group Discussion Paper 110.

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

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

Goodhart, C., R. Love, R. Payne and D. Rime, 2002, "Analysis of spreads in the dollar/euro and deutschemark/dollar foreign exchange markets", Economic Policy

Gottlieb, G. and A. Kalay, 1985, "Implications of the discreteness of observed stock prices", Journal of Finance, 40: 135-153

Grossman, S.J., M.H. Miller, D.R. Fischel, K.R. Cone and D.J. Ross, 1997, "Clustering and competition in asset markets", Journal of Law and Economics, 40: 23-60

Harris, L.E., 1990, "Estimation of stock variance and serial covariancess from discrete observations", Journal of Financial and Quantitative Analysis, 25: 291-306

Harris, L.E., 1991, "Stock price clustering and discreteness", Review of Financial Studies, 4: 389-415

Harris, L., 1994, "Minimum price variations, discrete bid-ask spreads and quotation sizes", Review of Financial Studies, 7, 149-178

Harris, L., 1998, "Optimal dynamic order submission strategies in some stylised trading problems", Financial Markets, Institutions and Instruments, 7, 1-75

Harris, L., and E. Gurel, 1986, "Price and volume effects associated with changes in the S&P500 list: New evidence for the existence of price pressures", Journal of Finance, 41: 815-829

Hartmann, P., 1998, "Currency Competition and Foreign Exchange Markets", Cambridge University Press.

Hasbrouck, J., and G. Sofianos, 1993, "The trades of market-makers: an analysis of NYSE specialists", Journal of Finance, 48, 1565-1593

Hau, H., W. Killeen and M. Moore, 2000, "The euro as an international currency: explaining puzzling first evidence", Centre for Economic Policy Research, working paper.

Hau, H., W. Killeen and M. Moore, 2002, "How has the euro changed the foreign exchange market?", Economic Policy

Hausman, J., A. Lo, and C. MacKinley, 1992, "An ordered probit analysis of transaction stock prices", Journal of Financial Economics, 31: 319-379

Hiraki, T., E.D. Maberly and N. Takezawa, 1995, "The information content of end of the day index futures returns: international evidence from the Osaka Nikkei 225 futures contract", Journal of Banking and Finance, 19(5): 921-936

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

Hseih, D.A., and A.W. Kleidon, 1996, "Bid-ask spreads in foreign exchange markets: implications for models of asymmetric information", in J.A. Frankel, G. Galli and A. Giovannini (eds), The Microstructure of Foreign Exchange Markets, University of Chicago Press: 183-206

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

Hume, D., 1741, "Of the balance of trade, in Essays: Moral, Political and Literary", Oxford University Press, (1963 reprint)

Ito, T. and W.L. Lin, 1992, "Lunch break and intraday volatility of stock returns", Economic Letters, 39(10, May, 1992: 85-90

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

Jain, P.C. and G.H. Joh, 1988, "The dependence between hourly prices and trading volume", Journal of Financial and Quantitative Analysis, 23(3), September: 269-83

Jennings, R., L. Starks, and J. Fellingham, 1981, "An equilibrium model of asset trading with sequential information arrival", Journal of Finance, 36: 143-161

Jones, C., and M. Lipson, 2001, "Sixteenths: Direct evidence on institutional trading costs", Journal of Financial Economics, 59: 253-278

Jorion, P., 1996, "Risk and turnover in the foreign exchange market", in J.A. Frankel, G. Galli and A. Giovannini (eds), The Microstructure of Foreign Exchange Markets, University of Chicago Press: 19-37

Kandel, S., O. Sarig, and A. Wohl, 2001, "Do investors prefer round stock prices? Evidence from Israeli IPO auctions", Journal of Banking and Finance, 25: 1543-1551

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.

Kawaller, I.G., P.D. Koch and T.W. Koch, 1990, "Intraday relationships between volatility in S&P500 futures prices and volatility in the S&P index", Journal of Banking and Finance, 14(2/3): 373-397

Kawaller, I.G., P.D. Koch and J.E. Peterson, 1994, "Assessing the intraday relationship between implied and historical volatility", Journal of Futures Markets, 14(3): 323-346

Keynes, J.M., 1930, "A treatise on money", Macmillan, London

Killeen, W., R. Lyons and M. Moore, 2003, "Fixed versus floating exchange rates: Lessons from order flow", Journal of International Money and Finance, forthcoming.

Killeen, W., F. McGroarty and M. Moore, 2000, "Early evidence on the success of the euro", mimeo., U.C. Berkeley, January.

Kleidon, A. and R. Willig, 1995, "Why do Christie and Schultz infer collusion from their data?" Working paper, Cornerstone Research, New York.

Kofman, P. and M. Martens, 1997, "Interaction between stock markets: an analysis of the common trading hours at the London and New York Stock Exchange", Journal of International Money and Finance, 16(3): 387-414

Kyle, A.S., 1984, "Market structure, information, futures markets and price formation", in International Agricultural Trade: Advanced Readings in Price Formation, Market Structure, and Price Instability, ed. by Gary G. Storey, Andrew Schmitz and Alexander H. Sarris. Boulder and London: Westview Press, 45-64.

Kyle, A.S., 1985, "Continuous auctions and insider trading", Econometrica, 53: 1315 - 1335

Lahtinen, M., 2003, "The purchasing power parity puzzle. Nonlinear and asymmetric aspects of goods-market arbitrage", University of Tampere working paper.

Lee J.H. and S.C. Linn, 1994, "Intraday and overnight volatility of stock index and stock index futures returns", Review of Futures Markets, 13(1): 1-38

Lee, C.M.C., B. Mucklow and M.J. Ready, 1993, "Spreads, depth and the impact of earnings information: an intraday analysis", Review of Financial Studies, 6(2): 345-374

Lehmann, B.N. and D.M. Modest, 1994, "Trading and liquidity on the Tokyo Stock Exchange: a bird's eye view", Journal of Finance, 49(3), July: 951-984

Lequeux, P. and E. Acar, 1996, "Tick by tick: an empirical study of BTP/bund spread trading", Smart spreads (LIFFE), no. 5, June: 2-4

Lequeux, P., 1999, "A practical approach to information spillover at high frequency: empirical study of the gilt and FTSE LIFFE contracts", in P. Lequeux (ed.), Financial Markets Tick by Tick, John Wiley & Sons, Chichester: 207-26

Levin, E.J. and R.E. Wright, 1999, "Explaining the intraday variation in the bid-ask spread in competitive dealership markets: a research note", Journal of Financial Markets, 2(2), May: 179:191

Liu, Y., and D. Mole, 1998, "The use of fundamental and technical analyses by foreign exchange dealers", Journal of Finance, 54: 1705-1765

Lo, A., and C. MacKinley, 1999, "A non-random walk down Wall Street", Princeton University Press

Low, A.H.W. and J. Muthuswamy, 1996, "Information flows in high frequency exchange rates", in C. Dunis (ed), Forecasting Financial Markets: Exchange Rates, Interest Rates and Asset Management, John Wiley: 3-32

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

Lyons, R., 2001, "The microstructure approach to exchange rates", MIT Press

Ma, C.K., R.L. Peterson, and R.S. Sears, 1992, "Trading noise, adverse selection and intraday bid-ask spreads in futures markets", Journal of Futures Markets, 12(5), October: 519-538

Madhavan, A., M. Richardson and M. Roomans, 1997, "Why do security prices change? A transaction-level analysis of NYSE stocks", Review of Financial Studies, 10(4), Winter: 1035-1064

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

Madhavan A. and S. Smidt, 1993, "An analysis of changes in specialist inventories and quotations", Journal of Finance, 48(5), December: 1595-1628

Manaster, S. and S.C. Mann, 1996, "Life in the pits: competitive market making and inventory control", Review of Financial Studies, 9(3), Fall: 953-975

Mark, N., 1995, "Exchange rates and fundamentals: evidence on long-horizon predictability", American Economic Review, 85: 210-218

Markham, J., 1987, "The history of commodity futures trading and its regulation", Praeger, New York

McInish, T.H. and R.A. Wood, 1990, "An analysis of transactions data for the Toronto Stock Exchange", Journal of Banking and Finance, 14(2/3), August: 441-458

McInish, T.H. and R.A. Wood, 1992, "An analysis of intraday patterns in bid/ask spreads for NYSE stocks", Journal of Finance, 47(2), June: 753-764

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

Mundell, R., 1962, "The appropriate use of monetary and fiscal policy for internal and external stability", IMF Staff Papers, 70-77

Neiderhoffer, V., 1965, "Clustering of stock prices", Operations Research, 13: 258-265

Neiderhoffer, V., 1966, "A new look at clustering of stock prices", Journal of Business, 39: 309-313

Neiderhoffer, V. and M. Osborne, 1966, "Market making and reversals on the stock exchange", Journal of the American Statistical Association, 61: 897-916

Niemeyer, J. and P. Sandas, 1993, "An empirical analysis of the trading structure at the Stockholm Stock Exchange", Journal of Multinational Financial Management, 3: 63-101.

Osborne, M., 1962, "Peridoic structure in the Brownian motion of stock prices", Operations Research, 10: 345-379

Osler, C.L., 2003a, "Currency orders and exchange rate dynamics: an explanation for the predictive success of technical analysis", forthcoming Journal of Finance

Osler, C.L., 2003b, "Stop-loss orders and price cascades in currency markets", Mimeo, Brandeis University

Parlour, C., 1998. "Price dynamics in limit order markets", Review of Financial Studies, 11, 789- 816.

Payne, R., 1997, "Announcement effects and seasonality in the intraday foreign exchange market", Financial Markets Group, Discussion Paper No.238

Piccinato, B., G Ballocchi and M. Dacorogna, 1998, "A closer look at the Eurofutures market: intraday statistical analysis", Working paper, Olsen and Associates

Rime, D., 2001, "Private or public information in the foreign exchange markets? An empirical analysis", Stockholm Institute for Financial Research, working paper.

Rock, K., 1990. "The specialist's order book and price anomalies", Mimeo, Graduate School of Business, Harvard University.

Roll, R., 1984, "A simple implicit measure of the effective bid-ask spread in an efficient market", Journal of Finance, 39, 1127-1139

Schwager, J., 1984, "A complete guide to the futures markets", Wiley, London

Scholes, M. S., 1972, "The Market for Securities: Substitution versus Price Pressure and the Effects of Information on Share Prices." Journal of Business, 45, 179–211.

Seppi, D. J., 1997. "Liquidity provision with limit orders and a strategic specialist", Review of Financial Studies, 10, 103-150.

Shleifer, A., 1986, "Do Demand Curves for Stocks Slope Down?" Journal of Finance, 41, 579–90.

Sopranzetti, B.J. and Datar, V., "Price clustering in foreign exchange spot markets", Journal of Financial Markets, 5: 411-417

Stephan J.A. and R.E. Whaley, 1990, "Intraday price change and trading volume relationships in the stock and stock options markets", Journal of Finance, 45(1), March: 191-220

Stoll, H.R., 1978, "The supply of dealer services in securities markets", Journal of Finance, 33, 1133-1151

Stoll, H.R., 2000, "Friction", Journal of Finance, 55: 1479-1514

Subrahmanyam, A., 1991, 1989, "Risk aversion, market liquidity and price efficiency", Review of Financial Studies, 4(3): 416-441

Tauchen, G., and M. Pitts, 1983, "The price-variability volume relationship on speculative markets", Econometrica, 51: 485-505

Thomas, S.H., and M. Wickins, 1991, "Currency substitution and vehicle currencies: tests of alternative hypotheses for the dollar, DM and yen", CEPR Discussion Paper, 507

Tse, Y., 1999, "Market microstructure of FTSE 100 index futures: an intraday empirical analysis", Journal of Futures Markets, 19(1), February: 31-58

US Securities and Exchange Commission, 1996, "Report pursuant to section 21(a) of the Securities Exchange Act of 1934 regarding the NASD and the Nasdaq stock market", US Government Printing Office, Washington, D.C.

Vila, A.F. and G. Sandmann, 1995, "Floor trading versus electronic screen trading: an empirical analysis of market liquidity and information transmission in the Nikkei stock index futures market", Financial Markets Group, London School of Economics, Discussion paper no. 218, 49 pages

Wang, G.H.K., R.J. Michalski, J.V. Jordan and E.J. Moriarty, 1994, "An intraday analysis of bid-ask spreads and price volatility in the S&P500 index futures market", Journal of Futures Markets, 14(7): 837-859

Werner I.M. and A.W. Kleidon, 1996, "UK and US trading of British cross-listed stocks: an intraday analysis of market integration", Review of Financial Studies, 9(2), Summer: 619-664

Weston, J., 2000, "Competition on the Nasdaq and the impact of recent market reforms", Journal of Finance, 55: 2565-2596

Woodward, S.A., 1998, "Price Fixing at Nasdaq? A Reconsideration of the Evidence", prepared for the Congressional Budget Office, presented at AEA meetings, January

Yule, U., 1927, "On reading a scale", Journal of the Royal Statistical Society, 90: 570-579

Zhou, B., 1996, "High frequency data and volatility in foreign exchange rates", Journal of Business and Economic Statistics 14:45-52