

An Empirical Analysis of Exchange Rates using High-Frequency Data

by

RICHARD PAYNE

London School of Economics and Political Science

Thesis submitted in partial fulfilment of the requirements of
the degree of Doctor of Philosophy at the University of London

September 1997

UMI Number: U111694

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U111694

Published by ProQuest LLC 2014. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

690822

7528

E

THESES

Abstract

This collection of studies examines several aspects of the behaviour of exchange rates using high-frequency data. Chapter 1 provides an introduction to the thesis, details the structure of what follows and presents a selective review of relevant literature. Chapters 2 and 3 focus on issues related to the microstructure of the foreign exchange market. Employing a short, transactions based data set derived from an electronic, inter-dealer broking system they test hypotheses related to the linkage between inside spreads and foreign exchange market volatility, the existence of private information on foreign exchange markets and the effects of 'thin markets' on the dynamics of quotations. Chapter 4 concentrates on the behaviour of intra-day foreign exchange volatility. The chapter presents and estimates a stochastic volatility model, extended to include intra-daily seasonality and the effects of U.S. macroeconomic announcements. Results demonstrate the importance of explicitly modelling the seasonal and show very large effects of public information releases. Chapter 5 builds on observations from Chapter 4 and estimates a long range dependent specification for volatility. Significant long memory is found in the volatilities of three major exchange rates and the long memory specification is shown to be superior to traditional short memory volatility specifications. Chapter 6 also extends the work in Chapter 4, examining the effect of U.S. and German macroeconomic news on the level of exchange rates. Estimations show that there are significant effects of U.S. and German 'news' at very high-frequencies but that these impacts are drowned very swiftly in subsequent exchange rate fluctuations. Finally, Chapter 7 uses daily data in order to examine the behaviour of the foreign exchange forward premium, proposing and estimating an explicitly non-linear model of forward premium behaviour. Results demonstrate that the forward premium is mean reverting, but in a non-linear fashion which cannot be captured in standard econometric specifications.

Acknowledgements

First and foremost I would like to express my gratitude to Charles Goodhart, my supervisor. His guidance has been invaluable. I would also like to extend my thanks to David Webb for providing me with a place in the Financial Markets Group. Access to the facilities and personnel therein have played a very large part in my completion of this project.

Throughout my (extended) tenure at the FMG I have drawn upon the wisdom and intellect of numerous researchers who have passed through or, in some cases, set down roots. A very incomplete list includes Alvaro Almeida, Sylvain Friederich, Dominique Guillaume, Phillip Hartmann, Andrew Harvey, Marc Henry, Bob Nobay, Gleb Sandmann and Ian Tonks. Also, I must acknowledge the tolerance and fortitude demonstrated by Ward Brown and Alex Stremme over the last few years through their tenure as my office-mates.

Finally I would like to thank my family and partner, Fiona, for their support (both financial and emotional) and encouragement. It is appreciated very deeply.

Contents

1	Introduction	23
1.1	Microstructure Research in Foreign Exchange	25
1.2	The FX Volatility Process	31
1.3	Announcement Effects and Exchange Rate Evolution	35
1.4	Spot and Forward Exchange Rates	38
2	Microstructural Dynamics in an Electronic Foreign Exchange Broking System	43
2.1	Introduction	43
2.2	Theoretical Motivations	47
2.3	The Data	53
2.4	Empirical Work	56
2.4.1	The Determinants of Quote Revisions	57
2.4.2	The Thin Market Hypothesis	61
2.4.3	Spreads and Volatility.	66

2.5 Conclusion	70
Tables and Figures	71
3 Information Transmission in Inter-Dealer Foreign Exchange Transactions	79
3.1 Introduction	79
3.2 Motivation and Methodology	81
3.2.1 Motivation	81
3.2.2 The Dealing System and the Data	85
3.2.3 Methodology	87
3.3 Results	93
3.3.1 Do trades cause Permanent Quote Revisions?	93
3.3.2 A Subsample Analysis	99
3.3.3 The Permanent/Transitory Decomposition	100
3.4 Conclusion	102
Tables and Figures	103
4 Announcement Effects and Seasonality in Intra-Day Foreign Exchange Rate Volatility	113
4.1 Introduction	113
4.2 Background: News and Seasonality	115
4.2.1 News and Announcements	116

4.2.2 Seasonality	121
4.3 The Empirical Methodology	126
4.3.1 The Basic SV Model	126
4.3.2 Modifications	127
4.3.3 Individual Announcements	130
4.3.4 Estimation	131
4.4 Results	133
4.4.1 Non-seasonal SV models	133
4.4.2 Seasonal SV Models	136
4.4.3 The Combined Model	138
4.4.4 Individual Announcements	140
4.5 Conclusions	143
Tables and Figures	145
5 An Investigation of Long Range Dependence in Intra-Day Foreign Exchange Rate Volatility	157
5.1 Introduction	157
5.2 Methodology	161
5.2.1 Long Memory Specifications	161
5.2.2 Semiparametric initial estimates	164
5.2.3 Fully parametric specification	167

5.2.4	Testing for persistence, long range dependence and stationarity	169
5.2.5	Deseasonalisation	171
5.3	The Data	173
5.4	Results	176
5.4.1	Testing for Long Range Dependence	176
5.4.2	Semiparametric Estimations	178
5.4.3	Specification Tests on the Fully Parametric Model	180
5.5	Conclusions	183
	Tables and Figures	185

6 The Effects of Macroeconomic ‘News’ on High Frequency Exchange Rate Behaviour 193

6.1	Introduction	193
6.2	Data and Methodology	197
6.2.1	The Data	197
6.2.2	Methodology	198
6.3	Analysis of American Macroeconomic Announcements	201
6.3.1	The Impact of ‘News’	202
6.3.2	Market efficiency	203
6.3.3	The Persistence of U.S. ‘news’ effects	204
6.3.4	Structural stability	206

6.4 Analysis of German Macroeconomic Announcements	208
6.4.1 The Impact of ‘News’	208
6.4.2 Market efficiency	213
6.4.3 The Persistence of German ‘news’ effects	214
6.4.4 Structural stability	216
6.5 Are the Impacts of German and US Announcements Similar?	217
6.6 Conclusions and Proposals for Further Work	222
Appendices	224
Tables and Figures	231
7 Nonlinear Mean Reversion in the Foreign Exchange Forward Premium	243
7.1 Introduction	243
7.2 Empirical Methodology	247
7.2.1 Kernel Regression	249
7.2.2 Threshold ADF Specifications	251
7.3 Data and Results	253
7.3.1 The Data	253
7.3.2 Results	254
7.4 Conclusions	263
Tables and Figures	265

List of Tables

2.1	Quote Revisions in Tick-time	72
2.2	Ask Revisions in Tick-time: Sub-sample study	72
2.3	Bid Quote Sub-sample Analysis	72
2.4	Ask Quote Sub-sample Analysis	73
2.5	First Period Bid Quote Sub-sample Analysis	73
2.6	First Period Ask Quote Sub-sample Analysis	73
2.7	Bid quote Return Autoregressions	74
2.8	Ask quote Return Autoregressions	74
2.9	Quote Revisions: MA(1)-GARCH(1,1)	75
2.10	Quote Revisions: MA(1)-GARCH(1,1)	75
3.1	Summary Statistics	103
3.2	Results from the Indicator/Quote Return VAR	104
3.3	Basic VAR: Group Significance tests	104
3.4	Size Effects in the Revision Equation	105

3.5	Size Effects: Group Significance tests	105
3.6	Revision Equation Parameters from the VAR using both Direction Indicators and Exhausts	106
3.7	VAR using Deal Indicator and Exhaust Indicator Group Significance tests of sets of Variables in all Equations	106
3.8	Revision Equation Parameters from the final VAR Specification in Subsample 1	107
3.9	Subsample 1: Group Significance tests of sets of Variables in all Equations	107
3.10	Revision Equation Parameters from the final VAR Specification in Subsample 2	108
3.11	Subsample 2: Group Significance tests of sets of Variables in all Equations	108
3.12	Results from the Permanent/Transitory Decompositions	108
4.1	Summary Statistics	145
4.2	Baseline SV Models	145
4.3	Estimates from SV/News Model	146
4.4	Estimates from Seasonal SV Model	147
4.5	Announcement Effects from Combined Model	148
4.6	Individual Announcement Impacts using Dummy Specification	149
4.7	Individual Announcement Impacts using Absolute Forecast Errors	149

5.1	Summary Statistics for Exchange Rate Returns	185
5.2	Summary Statistics for the Logarithm of Squared Returns	185
5.3	Tests of the significance of the Fractional Differencing Parameter for Raw Exchange Rate Returns	185
5.4	Tests of the significance of the Fractional Differencing Parameter for Exchange Rate Volatility ($\log(r^2)$)	186
5.5	Tests of the significance of the Fractional Differencing Parameter for Deseasonalised Exchange Rate Volatility	186
5.6	Estimation of the Fractional Differencing Parameter (\hat{d}) for Exchange Rate Returns	186
5.7	Estimation of the Fractional Differencing Parameter (\hat{d}) for Exchange Rate Volatility	187
5.8	Estimation of the Fractional Differencing Parameter (\hat{d}) for Deseasonalised Exchange Rate Volatility	187
5.9	Extensions of various LMSV specifications and Likelihood Ratio Testing using Deseasonalised Volatility	188
6.1	U.S. Announcements	224
6.2	German Announcements	224
6.3	Unit Root Tests - US announcements	227
6.4	Unit Root Tests - German announcements	228
6.5	Unbiasedness Regressions - US announcements	229
6.6	Unbiasedness Regressions - German announcements	230

6.7 The impact of US News on DEM/USD returns	231
6.8 Market Efficiency tests - US data	231
6.9 The persistence of the effect of US News on the DEM/USD	232
6.10 Predictive power of 15 minute return sign - US data	233
6.11 Structural stability tests - US News	234
6.12 The impact of German News on DEM/USD returns	234
6.13 The impact of German News on DEM/USD returns (weighted data) .	235
6.14 Market Efficiency tests - German data	235
6.15 The persistence of the effect of German News on the DEM/USD (original data)	236
6.16 The persistence of the effect of German News on the DEM/USD (weighted data)	237
6.17 The impact of German News on DEM/USD 3 hour returns	238
6.18 Predictive power of 15 minute return sign - German data	238
6.19 Structural stability tests - German data (original data)	239
6.20 Structural stability tests - German data (weighted data)	239
7.1 Summary Statistics for the Six Forward Premium Series	266
7.2 Summary Statistics for the First Difference of the Six Forward Premium Series	266
7.3 Results from the Linear ADF models for the change in the Forward Premium	267

7.4	Estimates of the Fractional Differencing Parameter for Forward Premia	267
7.5	<i>F</i> -Tests for the Delay Parameter in the TADF Specification	268
7.6	Results from the Modified TAR model for the change in the Forward Premium	268

List of Figures

2.1	Spreads on D2000-2	76
2.2	GARCH Return Volatility on D2000-2	77
3.1	D2000-2 Quote Midpoint Evolution	109
3.2	Quote Impulse Response Function: Basic VAR	109
3.3	Quote Impulse Response Functions to Exhausts and Non-Exhausts .	110
3.4	Quote Impulse Response Functions to Exhausts and Non-Exhaust, Sub-sample 1	110
3.5	Quote Impulse Response Functions to Exhausts and Non-Exhaust, Sub-sample 2	111
4.1	Average Daily Log Squared Return Pattern	150
4.2	Daily Autocorrelation Function	150
4.3	Average Weekly Log Squared Return Pattern	151
4.4	Weekly Autocorrelation Function	151
4.5	Average Daily Log Squared Return Pattern in Summer and Winter .	152

4.6	Comparison of Absolute Returns and Derived Volatility	152
4.7	Average State Value from Non-seasonal Model and Log Squared Returns	153
4.8	News Impacts from Non-Seasonal Model	153
4.9	Seasonal Periodogram	154
4.10	Estimated Seasonal in Volatility	154
4.11	Weekly ACF for Deseasonalised Log Squared Returns	155
4.12	News Effects from Seasonal and Non-Seasonal SV Models	155
5.1	Periodogram for JPY/USD Log Squared Returns	189
5.2	Log Periodogram for JPY/USD Log Squared Returns	189
5.3	JPY/USD Log Squared Returns: Sample Autocorrelations 1 to 1000 .	190
5.4	Periodogram for Deseasonalised JPY/USD Log Squared Returns . .	190
5.5	Comparison of Actual and Implied Correlograms for DEM/USD . . .	191
6.1	US Announcements: Average Scaled Impact over time	240
6.2	US Announcements: Average T-statistics over time	240
6.3	German Announcements: Average Scaled Impact over time	241
6.4	German Announcements: Average T-statistics over time	241
7.1	USD/GBP: Kernel and Linear ADF plots	269
7.2	DEM/GBP: Kernel and Linear ADF plots	270
7.3	CAD/GBP: Kernel and Linear ADF plots	271

7.4 JPY/GBP: Kernel and Linear ADF plots	272
7.5 FRF/GBP: Kernel and Linear ADF plots	273
7.6 SWK/GBP: Kernel and Linear ADF plots	274
7.7 USD/GBP TADF Model: Responses to 1 and 5 s.d. shocks	275
7.8 JPY/GBP TADF Model: Responses to 1 and 5 s.d. shocks	275
7.9 FRF/GBP TADF Model: Responses to 1 and 5 s.d. shocks	276
7.10 USD/GBP Long Memory Model: Responses to 1 and 5 s.d. shocks . .	276
7.11 JPY/GBP Long Memory Model: Responses to 1 and 5 s.d. shocks . .	277
7.12 FRF/GBP Long Memory Model: Responses to 1 and 5 s.d. shocks . .	277

Chapter 1

Introduction

This thesis brings together a number of studies written over the last three years, all of which focus on the behaviour and determination of exchange rates sampled at high-frequencies. Following the empirical ‘failure’ of the macroeconomic models of exchange rates presented and evaluated in the 1970’s and 1980’s,¹ researchers in the area have paid increasing attention to the manner in which exchange rates evolve in data sampled at higher and higher frequencies. Further, in recent years this increase in interest has been accompanied by the greater availability of high quality, high frequency data. This collection of studies hopes to provide a contribution to the burgeoning literature on intra-daily and daily exchange rate determination. It covers several different aspects of exchange rate behaviour, some of these being: how trading and exchange rate quotations are affected by the microstructure of the foreign exchange (FX) market, the assimilation of information, both private and public, into exchange rate levels, the behaviour of volatility and, in the final essay, the relationship between spot and forward exchange rates.

The thesis is structured as follows. Chapters 2 and 3 focus on issues of market

¹See, for example Meese and Rogoff (1983).

microstructure, examining the dynamics of foreign exchange quotes and trades emanating from an electronic, inter-dealer FX broking system run by Reuters and known as D2000-2. A specific issue treated in this analysis is the possibility of information asymmetries between FX market participants and the effect this private information has on prices. Following on from this, Chapters 4 and 6 examine the manner in which public information (specifically the announcement of key national macroeconomic data) impacts exchange rates. Chapter 4 analyses the response of volatility to public information announcements whilst Chapter 6 treats the effect on the levels of exchange rates. Chapter 5 is placed between the former pair as it is also a study of intra-day FX volatility. Specifically, some of the results from the analysis in Chapter 4 indicate that the autocorrelation in volatility extends to extremely long lags, behaviour which cannot be mimicked by standard ARMA-based models of conditional heteroskedasticity. Hence, Chapter 5 presents testing and estimation results relevant to the hypothesis of long memory in intra-day FX volatility. The final study, Chapter 7, is the only part of the thesis using daily data and is an analysis of the relationship between spot and forward foreign exchange rates. The work is motivated by the results of recent studies which suggest that FX forward premia follow non-stationary processes, a result which runs contrary to standard intuition. Chapter 7 provides evidence that forward premia are mean-reverting but that they follow non-linear stochastic processes. The chapter goes on to discuss several potential sources for non-linearities in forward premia.

The rest of this section provides a brief review of the development of the literature on high-frequency exchange rate behaviour.² The studies which comprise this thesis focus on four main issues, namely: empirical market microstructure research in foreign exchange, the behaviour of the FX volatility process, the impact of the announcements of public information on FX rates and, finally, the behaviour of the foreign exchange forward premium. In what follows we briefly summarise previous

²Excellent surveys of the properties and applications of high-frequency financial data, with particular reference to FX markets, can be found in Goodhart and O'Hara (1997) and Guillaume, Dacorogna, Davé, Müller, Olsen, and Pictet (1997).

research in each of these areas and motivate the research contained in subsequent chapters.

1.1 Microstructure Research in Foreign Exchange

The study of the effects of the microstructure of the FX market on trading and price formation is an area which is still very much in its infancy. Whilst in equities markets, the availability of high quality, intra-day data sets, coupled with extensive theoretical research on equity market structures, has led to a wealth of empirical microstructure research, for other asset markets (including FX) neither of these resources have been widely available.³

The stumbling block to both FX microstructure theory and data availability is essentially the decentralised nature of the market. The fact that, unlike the NYSE for example, the FX market consists of hundreds of dealers in a given exchange rate, interacting from physically different locations and with (possibly) differing information sets implies that the task of building a realistic market model based on the microeconomics of dealer behaviour is difficult. Deriving an estimable reduced form from such a model would be a further hurdle. This decentralisation also implies that well-formed transaction based data sets are very difficult to come by. Specifically, as FX dealers are under no obligation to report their trades immediately to an overseeing body, there is no central source for transactions data. Data on exchange rate quotations, taken from the Reuters FXFX service have become relatively widely available, but without disaggregated data on deals many interesting microstructural hypotheses are untestable.⁴

³The most heavily studied market has been the NYSE. Here, the specialist structure has permitted the construction of tractable theoretical models with estimable reduced forms and the exchange has collated and distributed transactions based data sets for individual stocks to academic researchers.

⁴The increased availability of the FXFX quotation data has essentially been due to the willingness of Olsen and Associates (Zurich) to collate and store these data and distribute them to

Some FX microstructure research has taken place however. It is mostly empirical in nature and employs the aforementioned FXFX data as a proxy for actual dealt exchange rates. This research includes analysis of the intra-day patterns in volatility and spreads (Baillie and Bollerslev (1991) and Hsieh and Kleidon (1996) for example) and the link between these patterns and FX volumes and dealer behaviour. Another large area of analysis has been the nature of the dependencies in FX volatility (see the review in Bollerslev, Chou, and Kroner (1992)) whilst further research examines the theoretical linkage between bid-ask spreads, volume and volatility.⁵ An interesting recent paper, Ito, Lyons, and Melvin (1996), explicitly tests the key microstructural proposition of informational asymmetries in intra-day FX markets using only FXFX quotation data. Employing a variance ratio framework (as in French and Roll (1986)) they provide evidence that informational asymmetries exist.

The above studies, although interesting, suffer from the problem that without disaggregated transaction data it is very difficult to evaluate many of the key propositions of microstructure research and dealer behaviour. Is inventory control an important factor in the determination of FX dealer behaviour, are there informational asymmetries between the agents on the intra-day FX market and, in relative terms, how important are these considerations? It is only really when one has access to trade by trade data that one can effectively test these propositions.

Thankfully then, very recently a couple of small transaction-based data sets on FX have appeared. A first, collected and collated by Richard Lyons, comprises the quotes and trades from a single broker-dealer pair over a single trading week and has been examined in Lyons (1995) and Lyons (1996). In the former of these papers Lyons tests for evidence of an inventory control motive in his dealer's price setting and also looks for signs that the dealer incorporates an element in his spread to

academics for a small fee.

⁵Bollerslev and Melvin (1994) model the linkage between volatility and spreads using a GARCH model for return volatility and the FXFX quotations data whilst Hartmann (1996) employs a new data set including twice daily aggregates of Japanese volumes in the JPY/USD spot market to explicitly include volumes in the model of spread determination.

protect himself against trading with an agent better informed than himself. To this end Lyons employs an empirical model of single dealer behaviour, that of Madhavan and Smidt (1991), originally used for equity analysis but which has a structure that fits Lyon's single dealer data set perfectly. The results contained in Lyons (1995) lend credence to both the inventory control and asymmetric information hypotheses. Indeed, the inventory control result which Lyons (1995) obtains is far stronger than that usually found in analysis of equities data (see the original paper by Madhavan and Smidt (1991).) This may be rationalised as due to the fact that the majority of intra-day FX dealers are forced to close out all positions overnight so that control of inventory intra-day is likely to be more important for FX dealers as compared to NYSE specialists. The more surprising result of the two is that Lyons (1995) finds evidence that his dealer protects himself against the possibility of informed trade by widening his spread.⁶ Whilst in equities markets the possibility of an agent gaining an informational advantage about the evolution of an individual stock is quite plausible, for an FX rate the mechanism by which an individual might gain superior information is unclear. The mechanism which Lyons (1995) promotes is that each dealer receives a signal over future exchange rate movements from the inflow of (non-dealer) customer orders to that bank. Heterogeneities in the size and composition of the customer base will then imply differing information sets across dealers which can be exploited by the better informed in inter-dealer trade.

Lyons (1996) builds upon Lyons (1995) by considering how the state of market activity might affect the informativeness of trade. This is done through the evaluation of two competing hypotheses. The first of these Lyons (1996) labels the 'event uncertainty' proposition. It states that, when trading activity is low, dealers place a smaller weight on the probability of an information event having occurred (i.e. they believe the probability of trading with an informed agent is low,) such that trades, as they are viewed as less informative, have a small affect on equilibrium quotations. The alternative view is the 'hot potato' hypothesis. It builds on intuitions from Ad-

⁶This result mirrors that from the aforementioned paper by Ito, Lyons, and Melvin (1996).

mati and Pleiderer (1988) which suggest that liquidity traders who have discretion over the timing of their trades will tend to clump together in their trading activity to exploit smaller transaction costs. If then the informed traders advantage is short lived, this mechanism implies that at times of high volume, the share of information based trades is low so that prices react less than in low volume times. Lyons (1996) builds a model which extends Madhavan and Smidt (1991) in order to incorporate the timing of trades and from this derives an estimable reduced form specification for quote revisions. The results derived are in favour of the 'Hot Potato' hypothesis, that information content is low when inter-dealer trading volumes are high.

The second intra-day transactions based data set which has recently emerged was provided by Reuters to Charles Goodhart. It is a seven hour transcript of the activity over their electronic, inter-dealer broking system D2000-2, providing data on all of the trades and the best quotes entered on the system over this period plus the quantities outstanding at the best quotes and quantities traded. The data then clearly encompass the activity of multiple dealers, and as such has greater coverage of the market than that of Lyons (1995), but obviously span a smaller amount of calendar time.

The primary examination of this data set was undertaken in Goodhart, Ito, and Payne (1996). An important first exercise undertaken in this study was to validate the use of the FXFX data employed in previous works as a proxy for the true evolution of dealt prices. It was found that the midpoint of the FXFX quotations tracked the firm quotations on D2000-2 very closely, but that in many other respects (such as spread behaviour and frequency of quote entry) the FXFX series were a poor proxy for actual market statistics.⁷ Goodhart, Ito, and Payne (1996) then progressed to give a preliminary examination of the inter-relationships between the variables available from the D2000-2 data set (i.e. deals, quote revisions quote

⁷This implies that the conclusions from previous work employing FXFX spreads, for example, are likely to be questionable whilst studies which have employed FXFX quotes or volatility will have results which are likely to very closely resemble those from actual dealt prices.

frequency and spreads.) Chapters 2 and 3 build upon the work of Goodhart, Ito, and Payne (1996), providing a more in depth examination of a number of market microstructure hypotheses.

Chapter 2 is a (slightly) revised version of Goodhart and Payne (1996). Using the D2000-2 data set it presents three distinct empirical exercises. The first of these is an analysis of the determinants of quote revisions for the bid and ask side of the market separately. Microstructure theory indicates a number of sources for quote revisions. The first, and most obvious, of these is transactions. Secondly, information revealed on the opposite side of the market will be impounded and, lastly, situations of uncompetitively large spreads will lead to spread improving revisions in general. Section 2.4.1 contains results from specifications for quote revisions reflecting these hypotheses.

The second exercise in Chapter 2 seeks to reconcile the negative autocorrelation found in the D2000-2 quote revision series (on a given side of the market) with the lack of dependence in transaction returns. A possible reason for this is that the market on D2000-2 is ‘thin’ in the sense that it contains few orders subsidiary to those at the touch.⁸ This will imply that after a cancellation of the best quote or transaction which fills the extant best order on a given side of the book, the best quote will jump a long way from its previous value. At the new level, it is likely that the spread is uncompetitively large, such that transactions are unlikely to occur. Hence one would expect the next action on the book to be a quote entry which takes the level of the quotation back towards the previous value. This will induce negative autocorrelation in quote returns, whilst transaction returns are uncorrelated due to deals only occurring at good prices. Analysis of this ‘thin market hypothesis’ is given in Section 2.4.2.

Lastly, in Section 2.4.3 we re-examine the link between spreads and volatility using

⁸The definition of ‘thin-ness’ given here is a statement about the depth of the limit order book and, as such, does not imply that trading volumes on D2000-2 are low.

the transactions based data of D2000-2. Both inventory control and asymmetric information hypotheses predict increases in spreads when markets are more volatile. For the former this is due to an increase in the risk of holding a given position whilst in the latter the link is between increased information arrival, implying higher volatility and greater spreads. The results given support these theoretical implications.

Chapter 3 takes a different tack, focussing purely on asymmetric information theories of dealer behaviour. As mentioned previously, both Lyons (1995) and Ito, Lyons, and Melvin (1996) find evidence of the existence of information asymmetries in FX markets. Chapter 3 presents a new set of results on this issue using the D2000-2 data. An implication of asymmetric information theories is that quotations are permanently updated in response to a trade, this effect being due to the possibility of the trade being information motivated and hence altering the equilibrium beliefs of dealers. The permanence of this impact distinguishes the effects of asymmetric information from other microstructure effects, such as inventory control and order processing cost impacts, which imply that quotations respond to trades, but only temporarily. This distinction can be employed fruitfully in empirical analysis. It predicts a correlation between unexpected trade activity and permanent changes in quotations which can be captured in a bivariate vector autoregression (VAR). This approach was originally suggested by Hasbrouck (1991a) and subsequently refined in Hasbrouck (1991b) and permits one to evaluate quantitatively the information content of trades, plus the proportion of all information which is trade related. Chapter 3 employs an extended version of this methodology which is applicable to the D2000-2 data. Results corroborate those given in Lyons (1995) and Ito, Lyons, and Melvin (1996), demonstrating the importance of asymmetric information effects in FX markets.

As detailed above, Chapters 2 and 3 provide a contribution to the growing literature on FX microstructure. They are limited, however, in that the data set employed is very small. Covering only a single European trading day on D2000-2, the data is really only a snapshot of electronic, brokered inter-dealer trade. Similarly, Lyon's

data covers only a single broker-dealer pair for one U.S. trading week. Despite this limitation, these studies have uncovered reliable evidence of both asymmetric information and inventory control effects in FX markets and, as such, provide a solid foundation for further work in this area as and when the required data become available.

1.2 The FX Volatility Process

An issue mentioned in the previous discussion of empirical FX microstructure research was the nature of the FX volatility process. This topic has recently become one of the most widely researched areas in exchange rate behaviour.

The motivating factors for the examination of volatility are twofold. First, there is the observation that, like all other financial markets but perhaps to an even greater degree, markets for FX have large temporal variations in volatility. Alongside this, recent years have seen the development of models of conditional heteroskedasticity which have proven to be highly satisfactory tools to describe such behaviour. The combination of these two factors has led to the plethora of work on FX volatility and it is to this strand of research that Chapters 4 and 5 belong.

A preliminary issue is the potential source of time variation in volatility. A popular class of microstructure models which delivers conditional heteroskedasticity and can be used to underpin our analysis are the Mixture of Distribution Models (MODM) presented in Tauchen and Pitts (1983), Epps and Epps (1976) and Clark (1973). The basis of this class of models is that the evolution of market activity is governed by a latent, random information arrival process. Within a given interval of calendar time, a stochastic number of information events occur and it is this number of events occurring in the interval which governs the volatility of the asset price and the volume of trade. Indeed, in the model of Tauchen and Pitts (1983), both volume and volatility are directly proportional to the information event count. Therefore,

the MODM yields the prediction that one will observe conditional heteroskedasticity in asset prices just as is observed in actual data. Several authors restrict the information arrival process to be a finite order autoregression such that the volatility process derived from the MODM has a similar structure to that derived from standard models of time-varying volatility.

Empirically, two families of models have been proposed to represent time-varying volatility, these being the Autoregressive Conditional Heteroskedasticity (ARCH) family (introduced by Engle (1982)) and Stochastic Volatility (SV) models (proposed in, for example, Harvey, Ruiz, and Shephard (1994) and Taylor (1994).) A survey of the former class can be found in Bollerslev, Chou, and Kroner (1992) and a review of the latter is contained in Ghysels, Harvey, and Renault (1996).

The original ARCH(p) structure contained in Engle (1982) models the conditional variance of an innovation as a linear function of p lags of the square of the innovation. The most popular member of the ARCH class, the GARCH(p,q) model introduced by Bollerslev (1986), also includes q lags of the conditional variance on the right hand side. None of the ARCH class of models permits a contemporaneous shock to the volatility process. Since the formulation of the ARCH model and its extension to GARCH, there have been numerous other variations on the theme, although, for the modelling of daily and intra-daily FX volatility, the GARCH(1,1) model has remained the most popular specification.⁹ Applications of such models to the FX market include Hsieh (1989) who finds strong GARCH effects in daily exchange rate data, Baillie and Bollerslev (1991) who examine hourly exchange rate returns and Andersen and Bollerslev (1997b) who present evidence of conditional heteroskedasticity in returns sampled every five minutes during the trading week.

The basic SV model has a quite different formulation to the ARCH class of models. The representation consists of two distinct equations. The first models asset returns

⁹Extensions include the GARCH-M model of Engle, Lilien, and Robins (1987), the EGARCH model of Nelson (1990) and the IGARCH formulation of Engle and Bollerslev (1986).

as the product of a white noise innovation and a function of an unobserved volatility process. The second equation is a law of motion for the latent volatility component, the most popular structure for which is an AR(1). This setup can be seen as directly analogous to that of the standard structural time-series models presented in, for example, Harvey (1989) and, as such, estimation can be performed in the same way. The representation for returns is linearised in the volatility process and, along with the law of motion for volatility, these two equations can be treated as a state space form, the parameters of which are retrieved by Quasi-Maximum Likelihood via the Kalman filter.¹⁰ An application of the standard SV model to daily foreign exchange data can be found in Harvey, Ruiz, and Shephard (1994), whilst a time-deformed intra-day SV model is presented in Ghysels, Gouriéroux, and Jasiak (1995).

Although the ARCH and SV classes of model have become very popular, they share the disadvantage that they are entirely univariate i.e. in the standard formulations one does not relate changes in volatility to external factors. The relationship between volatility and observable 'fundamentals' is, however, clearly of key importance as it provides a benchmark by which measures of 'excess' volatility can be judged. Thus, in recent times a small literature has arisen which attempts to explain the evolution of volatility, both in terms of its own history and the realisation of exogenous events. One strand of this literature relates intra-day volatility to flow measures of public information derived from, for example, the Reuters news headline pages (see Chang and Taylor (1996) and Melvin and Yin (1996).)¹¹ An alternative approach examines the impact of specific informational announcements on exchange rate evolution with the most common set of informational events employed in this literature being national macroeconomic announcements. A study which examines the effect of U.S. macroeconomic data announcements on minute-by-minute DeutscheMark futures volatility is Ederington and Lee (1993) and a similar study using hourly

¹⁰A number of alternative estimation methods have been proposed. For references see Section 4.3.4.

¹¹For studies of the impact of public information flows on equity indices see Berry and Howe (1994) and Mitchell and Mulherin (1994).

futures data is contained in Harvey and Huang (1991).

Chapter 4 examines the impact of U.S. macroeconomic announcements on the volatility of the spot DEM/USD, sampled at a five minute frequency, using an extended SV formulation. The SV structure is preferred to the more common GARCH model for two reasons. First the SV framework is closer in spirit to the MODM's introduced earlier. A key feature of the MODM is that the flow of information is stochastic in nature, implying that volatility and volume are also stochastic. The SV modelling strategy involves a stochastic representation for underlying volatility whilst in the GARCH model the conditional variance specification is entirely deterministic. The second, and more practical, reason is that pretesting and comparison of the GARCH and SV derived conditional variances with actual, instantaneous FX volatility demonstrated the SV model to be superior.

Our empirical analysis demonstrates strong announcement effects in intra-day exchange rate volatility, quantitatively similar to those in Ederington and Lee (1993). A complication in the analysis in Chapter 4 is the existence of a strong intra-day seasonal pattern in volatility. The model presented simultaneously captures this seasonal pattern and the announcement effects.¹²

Chapter 5 also focusses on intra-day FX volatility.¹³ As mentioned earlier, standard models for the analysis of volatility are the GARCH(1,1) and AR(1)-SV models. Estimations of both of these models generally demonstrate that intra-day volatility is well represented by a process with a near unit root (for the GARCH case this is known as an IGARCH representation).¹⁴ The implications of this are uncomfortable. Volatility will have a tendency to wander off to either infinity or zero. Chapter 5

¹²These seasonal patterns have themselves been the subject of much recent study. See Dacorogna, Müller, Nagler, Olsen, and Pictet (1993) and Andersen and Bollerslev (1997b). Seasonality can be embedded in the MODM framework by introducing a deterministic component, with a daily period, into the information flow process.

¹³This chapter is a revised version of Henry and Payne (1997).

¹⁴See the results in Chapter 4 and those in Andersen and Bollerslev (1997b) for a couple of examples.

investigates whether this result is an effect of mis-specification of traditional volatility models which essentially have ARMA structures. Building on observations from Chapter 4, Chapter 5 proposes and estimates models for intra-day volatility which explicitly permit long range dependence.¹⁵ Results demonstrate that volatility exhibits strong dependence and further that the IGARCH result seems to be due to mis-specification of traditional models in that they do not permit long range dependence.

1.3 Announcement Effects and Exchange Rate Evolution

The focus of Chapters 4 and 6 is the impact of macroeconomic announcements on exchange rate evolution. The essence of these studies and other works in the field is the event study methodology which attempts to relate changes in asset prices to specific informational events, precisely as the innovations to the public information set are made available to agents.¹⁶

The event study methodology is another tool which has a long history in studies of equity prices, with researchers in both Accounting and Finance having used the event study to investigate the effects of earnings and dividend announcements, for example, on equity price changes. The methodology has also been quite extensively applied to exchange rate evolution where the typical set of announcement data consists of releases of macroeconomic figures by relevant authorities.

One issue which must be addressed when using these tools is the actual quantity of ‘news’ contained in the announcement. Market participants will have formed expectations of the data contained in the announcement prior to its release and, hence, one needs a measure of these expectations in order to estimate the informa-

¹⁵ Examinations of long range dependence in daily exchange rate volatility can be found in Harvey (1993) and Baillie, Bollerslev, and Mikkelsen (1996).

¹⁶ As there are fairly extensive treatments of the literature relevant to this topic given in Chapters 4 and 6, the exposition here is kept relatively short.

tion content of the announcement. Fortunately, there are a number of commercial enterprises which collect and distribute such market expectation data for a number of macroeconomic announcements in a number of countries such that researchers can employ the median expectation in order to extract the information content of a given release.¹⁷

Once the information contained in the announcement has been extracted then, given the date and time of the release, one can proceed to characterise precisely how the asset price altered in response to the information. The advantage of high-frequency data in these analyses is that the increased sampling frequency permits a far sharper characterisation of the response to information.

Examples of the application of this methodology to exchange rates are as follows. Early studies concentrated on the response of daily USD exchange rates, primarily to U.S. money supply innovations. Engel and Frankel (1984), Cornell (1982) and Hardouvelis (1988) all uncover significant effects of these innovations on dollar exchange rates. Other U.S. aggregates have also gained attention. The Merchandise Trade figures have been shown to have significant effects in, for example, Harris (1995b). More recently, innovations in the U.S. non-farm employment statistics have been shown to have strong effects on daily exchange rates (see Harris (1995a), Harris and Zabka (1995) and Edison (1997).)

Later studies used exchange rate data sampled four times daily.¹⁸ Ito and Roley (1987) demonstrate that U.S. money supply innovations are again important but, in one of the few studies which examines non-U.S. announcements, that Japanese money supply surprises do not have a significant impact on JPY/USD evolution. Hogan, Melvin, and Roberts (1991) confirm the U.S. money supply result and provide evidence that trade balance news is important intra-day.

¹⁷An alternative method employed in the analysis of exchange rates has been to use innovations in a given series of macroeconomic data implied by a standard ARMA time-series model.

¹⁸The four intra-daily observations corresponding to the opening and closure of the four main regional FX markets (these being the Pacific, Japan, Europe and North America.)

In general, outside the three U.S. aggregates indicated above, there is little or no evidence that other U.S. announcements or any non-U.S. announcement affects dollar based exchange rates. Chapter 6 re-examines this issue in a high-frequency setting.¹⁹ Using data on the DEM/USD sampled at a five minute frequency and macroeconomic announcements from the U.S. and Germany we characterise the intra-day exchange rate reaction to these various 'news' series. Results show that in an intra-day setting, both U.S. and German releases have significant impacts on the DEM/USD, but that the effects are quickly drowned in subsequent exchange rate fluctuations for the majority of the announcements. Further results demonstrate that the response of the DEM/USD to surprises in a given variable is time-varying. Taking the above two results together helps explain the limited findings of previous studies at coarser exchange rate samplings.

As also detailed in Section 1.2, another strand of research has investigated the impact of the same type of macroeconomic releases on the volatility of FX rates. Harvey and Huang (1991) study hourly futures volatility and demonstrate that the hours containing the majority of U.S. macroeconomic releases are on average more volatile than the same hours on non-announcement days. Ederington and Lee (1993) and Ederington and Lee (1995) characterise the reactions in volatility and transaction prices in futures markets to the same macroeconomic data. They demonstrate that prices adjust within a very short space of calendar time (i.e. a few minutes) and that the majority of the volatility response is also very short term. Significantly increased volatility is felt for up to an hour post-announcement however. A further study which uses hourly data on the JPY/USD and examines the response of volatility to Japanese and U.S. 'news' is DeGennaro and Shrieves (1995). This study also finds significant post-release volatility effects. The contribution of Chapter 4 is the examination of impacts on spot volatility characterised at very high-frequencies as in Ederington and Lee (1993). Results demonstrate very strong intra-hourly increases in volatility in response to 'news', with the effects of the Employment

¹⁹This chapter is based on Almeida, Goodhart, and Payne (1997).

report statistics and trade figures being strongest. As most of the effect is intra-hourly, the results cast doubt on the findings of DeGennaro and Shrieves (1995).

1.4 Spot and Forward Exchange Rates

The relationship between spot and forward exchange rates is undoubtedly one of the most heavily researched areas in international finance. As such, the research has attracted two excellent survey articles in the last ten years, these being Hodrick (1987) and Engel (1996). In this section I provide a very brief review of the empirical research in the area, going on to motivate the study contained in Chapter 7.

A first area of study in this field is the empirical validity of the Covered Interest Parity (CIP) condition. CIP is an arbitrage condition which states that the foreign exchange forward premium should be identical to the differential between the nominal rates of return on foreign and domestic bonds.²⁰ Empirical evaluations of CIP using observed data on forward premia and interest differentials are, in general, supportive of its validity. Taylor (1987) and Taylor (1989), for example, demonstrates that there are no profits available from covered interest arbitrage, even in times of FX market turbulence.

A far larger body of work has sought to test the hypothesis of forward market efficiency. This hypothesis states that, under the conditions that agents are risk neutral and form rational expectations, the forward rate should be identical to the expected future spot exchange rate. Early tests of this hypothesis exploited the fact that under rational expectations the expected future spot could be written as the realised spot rate plus an error uncorrelated with information available at the time at which expectations were formed. This led to applied researchers regressing

²⁰The two investment strategies which underlie the condition are, first, investing funds in the domestic bond and, second, converting funds to the foreign currency at the spot rate, investing in foreign bonds and selling the proceeds forward.

realised spot rates on forward rates, with the null hypothesis that the coefficient on the forward rate be unity and the constant term zero. Studies based on this specification, e.g. Frenkel (1976), reported results which were very supportive of the 'unbiasedness' hypothesis. However, subsequently it was realised that, due to the $I(1)$ nature of exchange rates, standard inference was inapplicable to this regression framework. Hence, the results of tests based on the framework were brought into question.

This led to researchers positing alternative regression specifications in which to test the unbiasedness proposition. The most popular of these, examined in Bilson (1981) and Fama (1984) for example, estimates the relationship between realised spot rate changes and forward premia. It was assumed that the differencing implicit in the formulation would rid the specification of any problems associated with non-stationarity such that standard inference could be employed. The null corresponding to the 'unbiasedness' hypothesis is identical to that in the levels specification mentioned above. Results from this type of test were, however, uniformly poor.²¹ The estimates contained in Fama (1984) are typical of the literature. Estimated slope coefficients are found to be, on average, closer to negative unity than their hypothesised value of unity.

The negative results from the aforementioned tests have led researchers to explore several alternative possibilities for the rejection of the 'unbiasedness hypothesis.' One natural source of the rejection is the failure of the risk neutrality assumption. As Fama (1984) demonstrates, the existence of a time-varying risk premium introduces an omitted variable problem in the differenced regression framework which could lead to the negative results. Further research has attempted to specify the form of the risk premium and include it in estimations but with limited success. A second strand of research has focussed on learning and peso problems as the source of rejections. Finally, some studies have employed survey data on exchange rate

²¹See Engel (1996) for a number of references.

changes to examine the rational expectations component of the model, generally finding that failures of both rational expectations and risk neutrality account for the rejections (see Froot and Frankel (1989) and Peel and Pope (1991).)

The area of research into which Chapter 7 fits examines the more general time-series properties of spot and forward rates and, in other cases, of the forward premium. In this strand, researchers have concentrated on cointegration of spot and forward rates and the stationarity of the forward premium. The ‘unbiasedness hypothesis’ implies that spots and forwards should be cointegrated with vector (1,-1), so that the forward premium is a covariance stationary stochastic process. Applied research has found mixed results on both of these issues. Papers by Hakkio and Rush (1989) and Mark, Wu, and Hai (1993) present evidence of cointegration in spot and forward rates with cointegrating vector close to that implied by theory. Other studies, such as Evans and Lewis (1993), find the reverse result. Goodhart, McMahon, and Ngama (1997) also derive estimates of the cointegrating vector for spots and forwards which are very close to (1,-1), however they emphasise that their estimates are significantly different to the theoretical cointegrating vector. Similar mixed results are found when the covariance stationarity of the forward premium is explicitly examined.²² Whilst several authors report that premia are stationary, Crowder (1994) presents evidence that premia are I(1) and Baillie and Bollerslev (1994) conclude that forward premia are non-stationary, fractionally integrated processes. The possibility of non-stationarity in forward premia clearly has an impact on the Fama-type difference based regressions. Whilst the right hand side of the specification is stationary due to the I(1) nature of exchange rates, the left hand side is not. This implies misspecification, so that determining the order of integration of forward premia is vital for interpretation of the results of previous tests of forward market efficiency.

Chapter 7 contributes to this literature by conducting a closer examination of the time-series dynamics of the forward premium, focussing on its stationarity. The

²²This test amounts to imposing the theoretical cointegrating vector on spots and forwards and then testing the stationarity of the residual.

methodology departs from traditional linear econometrics, employing tools which permit non-linearities in the mean reversion of the premium. Results for a set of daily exchange rates suggest that forward premia are globally mean reverting, a result which fits more comfortably with theory.

Chapter 2

Microstructural Dynamics in an Electronic Foreign Exchange Broking System

2.1 Introduction

Financial markets exhibit high-frequency dynamic interaction between quotes and deals. Such interaction has recently been most intensively studied by Hasbrouck (1988), Hasbrouck (1991a), Hasbrouck (1991b) and Hasbrouck (1993) using equities data from the New York Stock Exchange. The NYSE differs from the majority of other markets by employing monopoly specialists to make markets in quoted stocks. Where there are competing, multiple market makers, as in the foreign exchange market, the dynamics are even richer and more complex. For example, when a single specialist market maker's quote has been hit, she will have to decide simultaneously on how to adjust both her bid and ask quote, and hence the spread between them. With multiple, competing market makers the decision on how to respond, say to

the effect of a deal at the ask, must be independently made - in principle by all the competing market makers - and (usually) by separate market makers on the bid and ask sides. Consequently the spread is not a choice variable, as it is in the case of a single market maker, but is endogenously determined as a consequence of the separate decisions of competing market makers. This induces a greater degree of interaction into the model, as those market makers setting ask (bid) prices have to react to changes in separately set bid (ask) prices, and to the resulting size of the spread, as well as to the effect of deals. We shall describe the effect of such, more complex, interactions, and show what effects it has on the estimated structure of the system.

The data used in this study emanate from Reuters D2000-2 electronic broking system, a system which caters only for dealers from the major banks, which are the main participants in the wholesale spot FX market. Limit orders input by some participants to D2000-2 are automatically matched with the market orders of others, yielding data on the best current bid and ask quotations, the inside spread and transaction prices and quantities.¹ In June 1993, when the record of our data was taken, the proportion of the FX market represented by such electronic broking systems was small, see Goodhart, Ito, and Payne (1996) for details. Inter-dealer trades can, of course, also be carried out via conventional brokers or, as is most common, via direct dealer-to-dealer contact. Finally, a significant proportion of trades occur between dealers and customers, a situation where the dealer again chooses the spread. We believe D2000-2 to be representative of the brokered segment of the inter-dealer market, but not, in all likelihood, of the direct dealer-dealer or customer-dealer trades.²

A particular incentive to pursue this line of research is that earlier work by Goodhart was innovative in finding strong negative autocorrelation, a first order negative mov-

¹The inside spread is defined as the difference between the highest bid and lowest ask quotations on the system at a point in time.

²For a more detailed description of the segmentation of the spot FX market see Flood (1991).

ing average, in extremely high frequency FXFX spot exchange rate quotation returns provided by Reuters (Goodhart (1989), Goodhart and Figliuoli (1991), Goodhart and Figliuoli (1992).) This finding, at frequencies higher than 10 minute intervals, has been confirmed by several other authors (Bollerslev and Domowitz (1993), Zhou (1992b), Zhou (1992a)), and seems also to appear in the NYSE data (Hasbrouck (1991a)). Most of these other research workers have tended to dismiss these findings as owing to noise (Zhou), discontinuous trading (Hasbrouck, p195, Lo and Mackinlay (1990), Bollerslev and Domowitz, p1432), being a consequence of the indicative nature of the FXFX series with consecutive quotes coming from banks with differing inventory positions, (Bollerslev and Melvin (1994), p364) or from banks in differing countries/continents with differing information sets (Goodhart and Figliuoli (1992)).

When we, with T. Ito, obtained data from Reuters for the D2000-2 system³, one of the features of the study of that data set was that the firm quotes still had (attenuated) negative first order autocorrelation in returns, whereas the transaction price returns showed no such autocorrelation, a result presented in Goodhart, Ito, and Payne (1996).⁴ One of the purposes of this exercise is to try to understand how this conjunction, of negatively autocorrelated quote returns and serially uncorrelated transaction returns, could occur. We believe that we now understand this phenomenon. It is largely due, we think, to forces which cause the limit order book, i.e. the bids (asks) within the system lying behind the best bid and ask, in a competing market maker system to be thin: so, if the best bid (ask) is exhausted by a deal, or removed for any other reason, the price will typically fall (rise) a comparatively long way until it hits, or provokes, a lower bid (higher ask) entry. That sharp fall in the bid (rise in the ask) is likely, however, to take the bid (ask) below (above) the level where a deal can reasonably be expected; in other words the spread will be-

³As indicated in the previous study the data remains the copyright of Reuters. Anyone wishing to use these data should refer to them, not to us, and hence we cannot ensure that this study can be replicated.

⁴Transaction price returns were defined as the first difference in prices of consecutive deals at the bid (or ask) respectively.

come excessively and uncompetitively large. This will then encourage a competing market maker wanting to increase (decrease) her inventory to reverse the prior fall (rise). This process can be seen graphically in the charts, showing the time path of the DM/Dollar data, which are contained in Goodhart, Ito, and Payne (1996). We set out the theoretical analysis of this, largely developed by Ho and Stoll (1983), with some additional flourishes by us, in Section 2.2.

This paper has the advantage that it is a tightly focused study of a part of the largest financial market in the world, the foreign exchange market, with a multiple, competing market maker structure, a structure which has not been so comprehensively studied as the single market maker set-up on the NYSE. The disadvantage is that the part of the forex market, for which we have data, is very small. Although the number of observations, on both quotes and deals, is quite large enough to obtain a statistically significant estimate of how this segment of the market operated at this particular juncture, it is a small segment of the market, i.e. the DEM/USD spot exchange rate as intermediated through Reuters D2000-2 electronic broking system, observed over a very short period of time, i.e. seven hours on one day in June 1993. Hence, the results obtained here may not be representative of the working of Reuters D2000-2 at other times, and Reuters D2000-2 may not be representative of the market as a whole. The data are exactly as described in Goodhart, Ito, and Payne (1996). Nevertheless the salient features are restated in Section 2.3.

Data similar to those which we employ in this work can be found in two recent studies. Biais, Hillion, and Spatt (1995) use a data set constructed from the electronically maintained limit order book on the Paris Bourse, whilst Hedvall and Niemeyer (1994) analyze the dynamics of the book and the order flow on the Helsinki Stock Exchange. In both of these works the data employed is far more complete than that used here. A first point is that both of these data sets represent more-or-less the entire market for trade in the stocks in question, whereas, as previously mentioned, our data covers only a small portion of the brokered inter-dealer FX market. Furthermore, Biais, Hillion, and Spatt (1995), as well as having the best quotes on each

side of the market, also have the five closest subsidiary orders on both sides. This obviously allows them to analyze explicitly the thickness of the book in the immediate vicinity of the 'touch'. The data employed in Hedvall and Niemeyer (1994) contains a further refinement, in that the identity of each dealer who performs any action in the book is available. Hence, again these data allow a far more detailed analysis of dealer behaviour and book dynamics than do ours. Further, in both cases the time span which the data encompasses is far greater than that covered in our work. Biais, Hillion, and Spatt (1995) work on a period of 19 trading days whilst Hedvall and Niemeyer (1994) analyze 147 trading days in total.

In Section 2.4, we report our empirical studies of the determinants of quote revisions and spreads. We start by exploring how bid and ask quotes respond to deals, to quote revisions on the opposite side and to the level of the spread. Next we examine the 'thin-market hypothesis' proposed, utilising a subsample analysis. Finally, we analyze the determination of spreads over D2000-2. Following the work of Bollerslev and Melvin (1994), we examine the relationship between volatility and spreads, before extending the specification explicitly to permit the impact of transaction activity on spreads. Section 2.5 concludes.

2.2 Theoretical Motivations

In this section we attempt to provide some theoretical motivations for our 'thin market' hypothesis. Unfortunately, however, we know of no prior theoretical work which treats a market structure as complex as that of D2000-2. Hence this section consists, in the main, of arguments and intuitions which we believe carry over from the two main strands of microstructure research (i.e. asymmetric information models and inventory control models,) rather than a cogent, rigorous model.

The first potential channel for the 'thin market' hypothesis is that of inventory control trading and to illustrate this mechanism we present a descriptive account

of the model contained in Ho and Stoll (1983).⁵ In this work, the authors analyze the conditions of dealer competition in a single asset market, in the absence of informational asymmetries. Dealers are sought out by potential traders in this asset, who, by interrogating all of the active dealers in the market, always elicit the lowest ask and highest bid prices possible. The objective of the dealers is to maximise their expected terminal wealth, but face a situation in which both transaction and return uncertainties exist. Casting the problem in a dynamic programming framework, the authors derive the optimal bid and ask fees for the dealer, fees which depend on the dealer's attitude to risk, his effective inventory, transaction quantities and the asset's return variance. From this starting point the market quotations can be calculated. In particular the dealer with the lowest fee has no incentive to base his quote, say at the ask, on that fee. His quote can embody a fee (very) marginally smaller than that of the dealer whose fee is next smallest (the second dealer) and still be certain that the next (sell) transaction will come to him, if a transaction opportunity actually arises. In this manner the authors demonstrate that the market spread is set by the fees of the second dealers on each side.

The next step in the analysis focuses on the incentives for inter-dealer trading. For this, the model is expanded to allow for an instant of inter-dealer trading prior to investor arrival: also each dealer is assumed to have identical risk preferences and to differ only in terms of inventories. Hence the incentive for inter-dealer trading is determined simply by the tradeoff between submitting a limit order, taking into account the transaction uncertainties this embodies, and hitting another dealer's quote, but then having to pay the fee this entails. By comparing the terminal wealth arising from these two alternatives, one can derive a condition on the tolerable inventory discrepancies which can persist if no inter-dealer trades are to occur. Ho and Stoll demonstrate that if transactions are of fixed magnitudes, then these tolerable discrepancies are relatively small. What does this say about the depth of the limit order book then? Clearly if there are large inventory discrepancies then the

⁵The formal model is set out on pages 1054-1069 of their paper.

incentive to inter-dealer trade exists. If we examine a second dealer in this market who desires, for example, to rid himself of some inventory he has two options. He could deal directly at another trader's prices, hence paying the fee which is set, but reducing his inventory with certainty. Alternatively he could post a limit order, but as a second dealer he knows that he has no chance of receiving the next incoming trade. In practice, depending on the size of the spread, the optimal course of action is usually to execute the market order. Given this process and the move towards equilibrium inventory positions it entails, uncompetitive limit orders are cleared out (actually never posted) as the uncompetitive traders are executing market orders rather than posting limit orders.

Hence, the inventory control channel gives us the possibility of 'thin market' effects. We now abstract away from the above mechanism and concentrate on the role played by asymmetric information in determining the depth of the market.

The classic asymmetric information models are based on the specialist structure of the NYSE e.g. Glosten and Milgrom (1985) and Copeland and Galai (1983), essentially addressing the adverse selection problem which a specialist faces. This arises as some agents who wish to trade are better informed than him, the rest sharing his information set. The implication is that the specialist will set quotations such that the spread incorporates an element over and above that due to order processing costs. Further, any trades which occur will cause a permanent alteration in subsequent quotations, arising from the possibility that the customer was informed and trading in order to profit from his superior information.

Our use of the idea of informational asymmetries is based on the properties of the price process and the trade-off between liquidity provision and the informational content of trades. Let us first assume that the price process is completely mean reverting and that trades carry no information. In this case dealers have no worries about informational trades and can submit orders from which they derive a profit from liquidity provision. This is true for all dealers so that one might expect a very

dense order book around the inside spread. A second reason for a thick order book in this case is that dealers can place limit orders away from the 'touch', hoping subsequently to profit from the mean reversion in quotations (see Harris (1995b)).

If, however, the price process is characterised by potentially permanent movements and market orders have a significant information content, then the order book is likely to be thin. This comes about because second dealers, in the sense introduced earlier, know that their quotations will only become transactable after a potentially information revealing market order has removed the best limit order. This will induce a desire to update their optimal quotations. Hence, there is an incentive not to submit subsidiary limit orders as this yields a protection against the price moving through the order permanently, as an informed market order trader walks up or down the book. In other words the option provided by writing the limit order is too expensive. This may be particularly so if the dealer submitting a secondary quote cannot monitor it continuously and/or adjust it instantaneously. Hence a price process with a large permanent component and the existence of informed market order traders are likely to bring about a thin limit order book.

There are obviously qualifications to the above argument. The first is the extent of the informational asymmetries which may be present in the forex market. In the market for single stocks, it is plausible to assume that some agents have insider information and base their trading activities on this. The application of this idea to the forex market is far less plausible. However, one manner in which players in this market may glean private information is through their direct customer order flow, a channel stressed by Lyons (1994), Lyons (1995), Perrauidin and Vitale (1996) and Chapter 3 of this thesis. If a certain bank and it's market maker observe a large imbalance in the orders received from customers, this may be taken as a sign of impending appreciation/depreciation. Given this signal, and the fact that it is unobservable to any other player the bank can formulate it's new expectation of the fundamental and optimal trading strategy. This phenomenon is modelled in Perrauidin and Vitale (1996) and empirical evidence for the existence of informational

asymmetries in the FX market can be found in Lyons (1995).

A further consideration is the manner in which informed traders participate in the market. The above discussion has assumed that they act via market orders, although they also have the option of posting limit orders. A simple argument would point out that the informed traders whole advantage rests on the privacy of the information carried, hence posting a limit order, which in the absence of any inventory control motives, can be simply seen as an advert of his beliefs, signals to others the information he carries and hence erodes his advantage. In the simple model of Perrauidin and Vitale (1996) the informed bank exploits it's advantage by trading at other banks prices, in line with the above argument.

However, recent research integrating market and limit orders points to situations where the optimal strategy may not be as simple as just executing a market order for example. Chakravarty and Holden (1994) show that a single, informed, risk-neutral trader can, in a world in which he can submit both limit and market orders, derive greater profit from the use of a combination of limit and market orders. This is driven by the fact that, in their model, a market order executes at an uncertain price and hence by simultaneously submitting a limit order on the other side of the market some bad states of nature can be cut out. In our example however, it can be argued that a market order executes with more-or-less certain price, as the best quotes are visible at all times.

A further study on optimal order placement is that of Harris (1995b). This author treats the manner in which three types of trader, who have differing amounts of private information, should submit orders under varying market structures. Some intuitive and general conclusions which can be drawn from Harris' work are the following. The order placement of informed traders should be more aggressive when their information is both pertinent and short-lived. Secondly, in volatile and low transaction cost markets there should be an incentive for traders to be more aggressive also. Lastly, a trader should be more aggressive as he becomes less patient.

In the terms of our segment of the FX market it seems that all the above reasons for aggressive order placement are present. In a market as liquid, and with such huge size and participation as the FX market, it seems certain that any privately informed trader can only hope to be so for a very short while, and as such, should trade very aggressively (i.e. through the use of market orders.) This reasoning is compounded in the inter-dealer segment of the market which we treat where, although being very volatile, spreads and transactions costs are, in general, very slight. Lastly, the bulk of activity in intra-day FX trading consists of individuals taking very short-term positions. This observation characterises most FX dealers as having very short time horizons and hence being far from patient.

Combining these arguments with the observation noted above, that over D2000-2 there is little or no price uncertainty associated with a market order, leads us to believe that informed FX traders should trade via market order whilst the non-informed should trade with limit orders at or in the market. An uninformed trader, however, who is not prepared to place a limit order at (or in) the market, should not place a subsidiary order close to the best price. If he did so, he would risk being hit by an informed trader, without having a large enough probability of execution with an uninformed trader. In other words, the option value of limit orders in a volatile market with informed traders is so large when execution probabilities are low that traders refuse to submit them, e.g. they will only do so if at, or in, the market. The addition of absentee traders, who operate using longer time horizons and place their limit orders quite far from the market, provides a clear picture of what we believe represents our 'Thin Market Hypothesis.'

Clearly the above intuitions from asymmetric information and inventory control models are highly stylized. It is our belief, however, that they contain elements and mechanisms which are highly pertinent to intra-day FX traders and are hence very relevant to our analysis of D2000-2. We recognise that these intuitions are no substitute for a rigorous and empirically testable model of our segment of the FX market, but, to our knowledge, no such model exists.

2.3 The Data

The empirical exercises performed in the following sections of this work concentrate on the Deutsche Mark/ Dollar (DEM/USD) spot rate. The data employed are those observed over Reuters D2000-2 electronic broking system for a seven hour period on June 16th 1993.⁶ Reuters D2000-2 is essentially a closed, electronic limit order book. At any point, any participant can input, or withdraw a previously entered, limit order (a price quotation for one side of the market along with an associated quantity for which the quote is good.) The system collates the existing orders and displays the best bid and ask along with their associated quantities. From here, a participant can instantly trade on the market quote, specifying on which side he desires to trade and for what quantity (up to the maximum on offer,) at which point all affected quotes and quantities are immediately updated. Each participant on the system can be fully aware of the current best market quotes and quantities (though inferior orders are not visible,) plus any transaction, and can easily deduce the amount transacted on each occasion.

Hence, for both sides of the market, the records consist of every quotation entered and the associated quantity, plus each transaction (and its quantity), all of which are timed, relative to each other, to the second. From these raw data we construct our basic variables: these are bid revisions, ask revisions and the revisions in the quote midpoint, the spread, a transaction indicator, signed transaction volume and a variable indicating whether a transaction exhausted the quantity posted on the market.

This gives us a sample of 1643 quote revisions in total, coupled with 437 transactions. In purely numerical terms, therefore, our sample is large enough for meaningful econometric work. However we are conscious that the sample we have may well be unrepresentative of the market as a whole. Our seven hour snapshot may have

⁶For a complete description of both the workings of D2000-2 and the data collected, see Goodhart, Ito, and Payne (1996).

been an atypical period in the DEM/USD market, and further our data are only that business conducted over the D2000-2 system: whether the trade over D2000-2 is representative of the market as a whole is also questionable.

As previously mentioned, the data we employ in this study was previously analyzed in Goodhart, Ito, and Payne (1996). In the current work we attempt to build on the preliminary findings of that study. To help tie the two papers together, we provide a brief summary of the main findings of that work below.

A first set of exercises undertaken in Goodhart, Ito, and Payne (1996) was to examine the relationship between the D2000-2 data and the widely available FXFX data from the same period. Unsurprisingly we found the DEM/USD levels to be very highly correlated, with very similar moments, a finding which recurred in the examination of the two quote revision series. The negative autocorrelation widely apparent in the FXFX quote revision data was present in the D2000-2 sample, but appeared to be less strong. One marked difference was in the spread data of the two samples. FXFX spreads can be seen to cluster around certain ‘conventional’ values e.g. 5, 7, 10 and 15 basis points, whereas the D2000-2 spreads are distributed continuously over a range from zero to around 20 basis points. In most of the other comparative exercises, e.g. of GARCH characteristics, the FXFX and D2000-2 series were found to be very similar.

When examining the transaction data available over D2000-2, two facts became apparent. First, after controlling for the effect of a bid-ask bounce, the series of transaction prices apparent was a martingale, consistent with market efficiency. Second, transaction sign was highly autocorrelated, implying that the market at distinct points in time is hit by a succession of deals on one side.

As described in Section 2.1, one of the incentives of this present study is to seek to understand how a martingale transaction price series and negatively autocorrelated quote revisions coexist. The theoretical analysis in Section 2.2 leads us to believe that the negative autocorrelation comes from a thin limit order book, whereas the

martingale property of transactions is driven by participants only dealing at 'good' prices.

Turning to the determination of transactions, it was found in Goodhart, Ito, and Payne (1996) that the single most important set of variables for explaining transactions was quote entry frequency. Hence transactions are closely linked to the 'pace' of the market: intense market activity and the high transactions frequency, in general, coexist, perhaps because of a correlation between the release of news information and incentives to transact e.g. owing to differential interpretations of, and responses to, 'news'.

A study of the determination of quote revisions demonstrated that the major impact was from own lagged values, consistent with the findings of numerous other studies using the FXFX data. However, using dimensions of our data unavailable to previous studies, we also found strong impacts from (exhaustive) transactions and the market spread: signed transactions were shown to have a positive impact on revisions, whilst large spreads caused upward bid revisions and downward ask revisions. When we examined the determination of the spread itself, the dominant impact was again the lagged spread, but (foreshadowing the results of Section 2.4,) we also were able to detect a positive effect from a crude volatility measure.

A final exercise undertaken in Goodhart, Ito, and Payne (1996) was to examine whether transactions have a direct impact on volatility. For this we estimated an AR(1)-GARCH(1,1) specification for quote midpoint revisions, going on to examine whether using a transaction indicator or an exhaust indicator in the conditional variance equation yielded a significant improvement in volatility prediction. Our finding was that there was little discernable direct impact of transactions on volatility.

In the following empirical sections we build upon the work described from Goodhart, Ito, and Payne (1996), with slightly different emphasis. Goodhart, Ito, and Payne (1996) was essentially an attempt, first to investigate the interactions between the D2000-2 and FXFX data, and second between the individual variables yielded by

D2000-2. In this work we are concerned more with microstructural issues, the determination of quotations and spreads and the effects of the depth of the order book on the statistical properties of quotation and transaction returns.

2.4 Empirical Work

Our first set of empirical exercises concentrates upon the factors which cause the quotations placed in the system to be revised. There are many potential reasons why quote revisions should occur in this competing market maker setup: public information may become apparent which causes all participants to revise quotations. A transaction may (partially) reveal previously hidden information. The inventory positions of those submitting limit orders may alter, perhaps from direct trading with their customer base, etc.⁷ Note that in this market, as can be seen in the charts in Goodhart, Ito, and Payne (1996), over very short horizons the bid and ask sides of the market move in a less than perfectly correlated manner. The bid and ask quotations, over short periods of time, may diverge and only later re-converge, thereby affecting the spread. As noted in Goodhart, Ito, and Payne (1996) the bid and ask quotes behave like a pair of cointegrated variables, with the spread providing the error correction mechanism.

These observations give us a starting point for our empirical exercises. If we are to characterise the behaviour of quote revisions, for example on the bid side of the market, then transactions, spreads and revisions on the ask side should all be thought of as potential explanatory variables.

At this point, attention must be given to the manner in which we should scale

⁷Unfortunately, a direct test of the presence of an inventories effect would require knowledge of some participants' inventory positions and the identities of the agents involved in quote submissions and transactions. These data are not present in our snapshot of the system. For a forex study which does test for the presence of an inventory control motive (using a dealer's derived inventory position) see Lyons (1995).

the data for the empirical exercises. With such high-frequency observations on the forex market (and similarly for intra-day stock data,) there are two general approaches. The first is the more conventional: to scale the observations in calendar time i.e. to take an observation every k seconds, where k is a fixed positive integer. This gives a normal time-series with a fixed observation window. The drawback to this approach is that activity in these markets is generally irregular: a period when nothing happens for 10 minutes can be followed by a 30 second interval in which there are numerous revisions and transactions. This leads to a time-series which is empty for many of the observations, and only sporadically non-null. The alternative approach counters this problem. Instead of using time as a scaling rule, one can construct the data along an activity scale for which time is irrelevant. Using this approach, an observation is recorded at every point when there is an event in the market i.e. when a quote revision occurs or when a transaction takes place. Obviously this removes the problem of the null-data but it leaves a data set without any real reference point in time so that, for example, forecasting may be questionable.

This work, in the main, employs the latter approach. All tables shown at the end of the paper report activity time results, clock time results are mentioned only to the extent that they differ from those reported.

2.4.1 The Determinants of Quote Revisions

As stated above, lagged revisions, transactions, spreads and revisions on the other side of the market should impact upon, for example, bid revisions. In this sub-section we analyze this phenomenon with two data sets, one constructed using bid side activity as the observation criterion and a similarly constructed set for the ask. Starting with the construction of the ask data, an observation is recorded at every point where a revision or transaction occurs on the ask side of the market. To these two variables we add the spread resulting at each observation point and the

revision (if any) occurring on the bid side of the market in the time elapsed between the current observation and its antecedent. This process is repeated to generate a similar set of data for the bid side of the market. Our basic equations are shown below;

$$\Delta a_t = \alpha_0 + \alpha_1 \Delta a_{t-1} + \alpha_2 ad_t + \alpha_3 ad_{t-1} + \alpha_4 s_{t-1} + \alpha_5 \Delta b_t + v_{1t} \quad (2.1)$$

$$\Delta b_t = \beta_0 + \beta_1 \Delta b_{t-1} + \beta_2 bd_t + \beta_3 bd_{t-1} + \beta_4 s_{t-1} + \beta_5 \Delta a_t + v_{2t} \quad (2.2)$$

In the above representations Δa_t denotes ask revisions, Δb_t bid revisions, ad_t and bd_t are ask and bid deals respectively and s_t is the inside spread. We include lagged transactions to account for the possibility of delayed adjustment to the new information contained in such deals.⁸ Our hypotheses, motivated by the arguments above and in preceding sections, are outlined below.

- $\alpha_1, \beta_1 < 0$: negative autocorrelation is induced by the thin market.
- $\alpha_2, \alpha_3 > 0$: transactions at the ask cause upward market movements
- $\beta_2, \beta_3 < 0$: transactions at the bid cause downward market movements.
- $\alpha_4 < 0, \beta_4 > 0$: large spreads encourage more competitive quotes.
- $\alpha_5, \beta_5 > 0$: movements on the other side of the market are mimicked on this side.

The results for the two tick-by-tick data sets are shown in Table 2.1.⁹ The estimations support the validity of our hypotheses. All coefficients are of the expected sign

⁸We did examine whether deals on the bid side might directly affect ask quote revisions and vice versa, but the effects always appeared insignificant and these variables (bd_t in the ask equation and ad_t in the bid equation) are omitted.

⁹Note that all regressions are run using White heteroskedasticity-consistent standard errors in an attempt to mitigate the effects of conditional heteroskedasticity.

and most are highly significant. Revisions are negatively autocorrelated, in line with results from the widely used FXFX data sets. Transactions cause outward jumps in quotes on each side of the market (both instantaneously and in the subsequent period,) a result which might be seen as an indication of the existence of private information in FX markets.¹⁰ Finally, larger spreads engender more competitive quotes. However, a clear asymmetry occurs with spreads. Bid revisions were clearly affected by spreads, but this did not apply to the ask side of the market. One partial explanation, perhaps, is that much of the ask activity occurred in the latter part of the data set, at a time when the ask was rapidly rising towards a new 'equilibrium' but the bid did not respond instantaneously. This led to a period when rising spreads and upward revisions co-existed, hence potentially weakening the impact of the spread as a downward force on ask quotes. A further asymmetry occurred in the response of revisions to those on the opposite side of the market. Bid revisions are significantly positively affected by those on the ask side, but this is not as evident in reverse. Again one may put this down to the ask being the dominant (and leading) side of the market in the latter half of the data, rising quite sharply away from the bid for a fairly long period before the bid reacts and closes the gap after a point.

In order to test this hypothesis we split our time series for ask revisions into two sub-samples. The first of these represents the opening part of the data, a period of relative calm in the market when the level of the DEM/USD varied only marginally around its mean. The second sub-sample contains the period of sharp appreciation in the rate, the hypothesised source of the asymmetries pointed out above, when the ask broke sharply away from the bid quotations for a fairly long period of time.¹¹ Our sub-sample results for the ask side are given in Table 2.2. The results from the first sub-sample support our beliefs. The spread has a significant negative impact on ask revisions and there is a significant positive response of ask quotes to bid revisions. For the second sub-sample the results are similar to those for all observations. The

¹⁰See Chapter 3 for a detailed examination of this issue.

¹¹The first sub-sample contains observations 1-429 of the ask side data points and the second contains observations 430-790.

effects of both the spread and bid revisions are insignificant.

Hence, after controlling for the jump in the rate in the latter half of the data set, we obtain the hypothesised impact of spreads on quote revisions. This result is mirrored in both Biais, Hillion, and Spatt (1995) and Hedvall and Niemeyer (1994), both finding that 'excessive' spreads lead to spread-improving revisions in general, and also that transactions tend to occur when the spread is tight. This, as pointed out by Biais, Hillion, and Spatt (1995), corresponds to liquidity being supplied to the market when it is needed, (through the submission of fresh limit orders,) and drained from the market when it is abundant (as market orders hit the existing quotes.)

Examination of the same system in clock-time confirmed the tick-time results, in terms of the signs of the estimated parameters.¹² One discrepancy, however, was noted in the magnitude of both own autocorrelation impacts and effects from the opposite side of the market. This may be due simply to the dilution of the data i.e. the spacing of irregular observations over a far larger sample size than in the activity based data. Alternatively, the drop in these coefficients may be due to the fact that within a given interval there are instances of quote revisions on both sides of the market, implying the estimations may be biased due to a simultaneity problem. We investigated this possibility by estimating the same system using Two Stage Least Squares.¹³ Little innovation was gained from using 2SLS, none of the parameter values altered significantly, the only estimates which changed at all were the effects from the opposite side of the market which, as expected, rose. The final impacts remained far less than in the activity-based data, probably due to the dilution effect mentioned above.

¹²Our data set for these exercises is simply constructed by taking an observation from the market every five-seconds. With this small observation window we are clearly exacerbating the 'zero's' problem indicated earlier, but to balance this we are obscuring as little of the activity as possible. The variables constructed using this criterion are as in the previous sub-section: bid revisions, ask revisions, bid deals, ask deals and spreads.

¹³Contemporaneous revisions on the other side of the market were instrumented using their own first lags.

2.4.2 The Thin Market Hypothesis

The results of the previous sub-section, while not inconsistent with our hypothesis that the negative autocorrelation in returns is caused by the market being thin, do not, however, bear directly upon it. We sought, therefore, to define and test our ‘thin market’ hypothesis more directly. To this end we categorised the revisions on a single side of the market in terms of their determinants. A quote, say at the bid, may be changed either following a deal which has exhausted the quantity on offer, or after a change in the quote on the other side (in this case the ask), or when there has been neither a bid deal exhaust, nor an ask revision on which to condition. The ask revision, in turn, can be either positive or negative. There are, therefore, five possible alternatives, conditioned on events since the last change in the bid quote.

- B1. Quote revision following bid deal exhaust, which must be negative.
- B2. Quote revision to bid following positive ask quote revision.
- B3. Quote revision to bid following negative ask quote revision.
- B4. Quote revision to bid, positive, without deal exhaust or ask revision since last quote.
- B5. Quote revision to bid, negative, without deal exhaust or ask revision since last quote.

By construction B1 and B5 must be negative, and B4 positive. We would expect B2 to be positive and B3 negative, although these latter effects may be weakened by the non-instantaneous reaction of one side of the market to information being revealed on the opposite side.

As pointed out in Section 2.1, our data do not contain quotes subsidiary to the best bid and ask, and hence these subsidiary orders must be inferred. Using our classification, the only categories which will signal the next subsidiary order and are

free of the immediate impact of information revelation are B1 and B5 i.e. quote cancellations without a corresponding ask alteration and exhaustive transactions. Clearly if we were able to observe the subsidiary orders directly, as is possible in the studies of Biais, Hillion, and Spatt (1995) and Hedvall and Niemeyer (1994) then the picture we give could be greatly clarified. Nonetheless we continue by using the revisions associated with the two categories mentioned above as our indications of the 'thinness' of the market.

If there is a 'thin' market, then we would expect the size of the decline in B1 and B5 to be large, since the hypothesis is that there are few close subsidiary orders. Note that the size of movement immediately following a deal exhaust depends only on the position of the next lowest quote in the order book, not on information; hence there is no reason to expect the size of jump in B1 and B5 to differ. By contrast revisions to bid quotes following information from ask quote adjustments should be smaller, since we hypothesize that these will often be quote revisions from the bank already making the best bid quote. Note that revisions to bid quotes will incorporate the information in prior ask quote revisions. Since this information will include ask transactions that exhaust the quantity offered, and force ask prices upwards, we would expect bid revisions following upwards ask revisions to be stronger (than downwards bid revisions following downwards ask revisions).

If, following the removal of a bid quote via an exhaustive transaction or quote cancellation (cases B1 and B5), the next bid quote level is so low that it needs to elicit a strong rebound in order to reach a level where a transaction is feasible, then case B4 may show a larger (positive) jump than cases B2 or B3. Hence our hypothesis for the absolute size of mean bid revisions within each of the preceding categories, is as follows:

$$B1 = B5 > B4 > B2 > B3$$

Using an activity-time, revisions-based data set, we examine this hypothesis by constructing the mean basis-point revisions associated with each sub-sample.¹⁴ These are shown in the first two columns of Table 2.3.

In all but the sign of the mean revision associated with category B3 our hypothesis is supported. As a guide to the scale of these revisions, the modal level of the spread in our data is 3 basis points. This implies that, as the revisions associated with deal exhausts and quote cancellations are greater than 4 basis points, they more than double the usual spread.¹⁵

We can clearly construct a similar set of definitions for ask revisions, (corresponding exactly to those previously given for the bid and which we label A1-A5.) Then using similar arguments to those posited for the bid side of the market we obtain the following expectation for absolute mean changes:

$$A1 = A4 > A5 > A3 > A2$$

The results corresponding to this hypothesis are given in columns one and two of Table 2.4. The results for the ask broadly conform with our hypotheses, with again deal exhausts and quote cancellations engendering around a doubling of the spread.

Comparing the above results with those found in Biais, Hillion, and Spatt (1995) one can see a sharp distinction. Whereas our results indicate that the spreads between best bid and best ask, and between first and second orders (be they bid or ask,) are

¹⁴These sub-samples were constructed as follows. First, regardless of any other market activity, all revisions associated with a transaction were taken to form sub-sample B1. We also removed their immediate successors, classing them as reactions to the deal revision rather than revisions caused by any other source. We then removed all revisions which were associated with a change in the rate on the opposite side of the market, dividing them, depending on the direction of change on the other side, into classes B2 and B3. Again the immediate successors of these revisions were removed. Finally, the remaining observations were divided into positive and negative revisions in order to form classes B4 and B5. Precisely the same procedure was followed for the ask side of the market.

¹⁵This assumes there is no immediate movement in the quote of the other side of the market. Hence the original spread is 3 b.p.'s and a 4 b.p. downward bid revision then creates a 7 b.p. spread.

approximately equal, in their results the bid-ask spread is, on average, over double the size of the spread between consecutive quotes on one side of the market. This seems to indicate that the electronic order book of D2000-2 is indeed thinner than that associated with the Paris Bourse.

Next our ‘thin market’ hypothesis suggests that if the sharp jumps (down following cases B1 and B5 for bid revisions, and up following cases A1 and A4 for ask revisions) take the quotes too far off the ‘true’ market level, then there will normally be a rebound towards the ‘fundamental’ valuation. Note, however, that a jump following a deal exhaust does contain information that the ‘true’ market level may have shifted in that direction, whereas a quote removal without a deal does not do so to the same extent. Hence we would expect the reversal following case B5 for the bid and case A4 for the ask to be greater than for case 1. When a reversal has occurred back to the ‘true’ market level, cases B4 and A5, then we would not expect any subsequent regular change. Similarly, if information is processed efficiently, then following a quote revision, say to the bid, conditioned on a prior ask quote change we would not expect the subsequent bid quote revision to be of any particular sign or magnitude. Again though, processing inefficiencies may lead to the subsequent revision being of the same sign as its antecedent.

Put another way, our hypothesis is that the significant negative autocorrelations and reversals in the quote series are predominantly the result of large jumps (downwards for bids and upwards for asks) caused by a thin market being reversed. Hence our assumption is that the absolute mean levels for the subsequent change, the extent of negative autocorrelation between the change and the subsequent change, and the percentage reversal of the prior change in the subsequent change will have the ordering, for the bid:

$$B5 > B1 > B4 = B2 = B3$$

The estimates of the three statistics mentioned above are given in the final four columns of Table 2.3 and the results following the jumps were close to our expectations. However the extent of negative autocorrelation and the size of the subsequent percentage reversal following a positive bid change (no deal exhaust or ask revision intervening), case B4, and also following a bid change conditioned on a negative ask revision, case B3, were considerably larger than we had expected.

A similar exercise can be run for the ask side, where our expectations were that the ordering would be;

$$A4 > A1 > A5 = A2 = A3$$

Referring to Table 2.4, one can see that the results were less successful. The reversal of a positive revision to ask prices (case A4) was much less, and the reversals of negative revisions (case A5) and following adjustments conditional on prior negative bid change (case A3) were much greater than we had expected.¹⁶

Again, a possible explanation for these, somewhat mixed, results is the aforementioned DEM/USD jump towards the end of our short period. This may have led to a bias whereby negative changes to the ask, or bid, (cases 5 and 3), were shown as subject to exaggerated reversal, and vice versa for positive changes (cases 4 and 2). We proceeded to test this hypothesis by employing a similar data split, for both bid and ask sides, to that used for the ask data in Section 2.4.1. Tables 2.5 and 2.6 replicate the analysis of the full revision data sets for these pre-jump revisions. The results demonstrate that our intuition about the effects of the pronounced volatility in the latter half of the data are borne out. It is now clear, for both sides of the market, that the largest jumps and reversals are those associated with exhaustive transactions and quote removals.

¹⁶The results corresponding to category A3 are strongly influenced by a single observation, which entails a positive revision of 50 b.p.'s followed by a reversal of 45 b.p.'s. Removal of this one observation causes the correlation to drop to -0.32 and the % reversal to fall to -0.21.

Hence, we claim to have provided evidence that the negative autocorrelation inherent in the quote returns of these data is largely due to the effects of a ‘thin market’. Our final examination of this phenomenon explicitly focuses on the autocorrelation structure of returns. We compare the autoregressive structure of the raw revisions series with a pair of filtered series which omit all those revisions due to deal exhausts and quote cancellations, plus their immediate successors. From the previous results of this subsection, and according to our hypothesis, we expected a sharp decline in the significance of the autoregression parameters as this filtering amounts to removing the ‘thin market’ effects on the revision series. The estimates of the autoregressions for both sides are given in Tables 2.7 and 2.8.

Examining the results it seems clear that the removal of the thin market impact causes a sharp drop in the explanatory power of the autoregressions, but that this is more pronounced for the ask side of the market. On this side the R^2 drops from around 0.3 to less than 0.05, leaving no significant autoregressive terms. On the bid side the results are slightly less impressive, although still supportive of our hypothesis. The R^2 declines by approximately 66% to around 0.1 and as is visible from Table 2.7, the first and fourth lags retain significance at 5%. Hence, although there again seems to be an asymmetry between results for the two sides of the market, the removal of the ‘thin market’ effect can be seen to have an important impact on the characteristics of the revisions series, explaining a large proportion of these series’ negative autocorrelation.

Consequently while the thin market hypothesis appears to have validity and be responsible for much of the observed negative autocorrelation, it does not appear to represent a complete explanation of that.

2.4.3 Spreads and Volatility.

The last phenomenon we analyze in this section is the relationship between spreads and volatility for our data. A recent paper by Bollerslev and Melvin (1994) does

exactly this, but using the indicative quotes from the widely available FXFX data. As mentioned previously, one facet of the FXFX data is that the realised spreads tend to cluster around certain values, virtually the only recorded spreads being of 5, 7, 10 or 15 basis points. Further, these spreads are the input of bid and asks from a single bank (i.e. the prices at which one bank is willing to buy/sell currency,), and hence need not bear any relationship to the inside spread realised in the market. To overcome the first problem, Bollerslev and Melvin (1994) employ an Ordered Probit procedure, with spreads as the dependent and volatility as an independent variable.

Our data, on the other hand, exhibit neither of these problems. The distribution of spreads is fairly smooth over positive values up to around 20 b.p.'s (with a few very large spreads realised,) such that we could comfortably use OLS as an estimation technique. Also the spreads we have are the market's inside spreads i.e. the spread resulting from the combination of the lowest quoted ask and highest bid. These values should bear a closer relationship to the volatility in the market than the spreads quoted by individual dealers. A final point in comparing the data employed is that the quotes given over D2000-2 are firm quotes, that is prices at which dealers are committed to trade. The FXFX data provide only indicative quotes, prices which do not wholly represent a firm offer to trade.

To estimate this system we employ (yet) another set of data. Using an activity time scale, we record spreads, signed revisions and an unsigned, binary transaction indicator at every point where a revision or transaction occurs. Note that, unlike the previous data sets, in this analysis bid and ask revisions are combined at each observation point to give the net upward or downward revision in the quote midpoint.

We proceed by estimating the same system as Bollerslev and Melvin (1994), but here using standard Maximum Likelihood techniques, rather than the Ordered Probit they employ. The specification is demonstrated below;

$$s_t = \lambda_0 + \lambda_1 \sigma_{rt}^2 + \lambda_2 s_{t-1} + \epsilon_t, \quad \epsilon_t | I_{t-1} \sim N(0, \sigma_{kt}^2) \quad (2.3)$$

$$\sigma_{kt}^2 = \exp(\gamma_0 + \gamma_1 \sigma_{rt}^2 + \gamma_2 s_{t-1}) \quad (2.4)$$

The rationale for this specification is as follows: spreads in this market are strongly autocorrelated, hence the lagged spread is included as a regressor. The link between quote return volatility (σ_{rt}^2) and spreads comes, theoretically, from the greater return uncertainty and information arrival in the market. A trader faced with increased volatility and faster information arrival will widen his spread so as to protect himself against more extreme price revisions in future periods. This effect is apparent in a multitude of microstructure models e.g. the simple specification in Bollerslev and Melvin (1994), Ho and Stoll (1983) and Copeland and Galai (1983). Lastly, the volatility of the spread (σ_{kt}^2) is allowed to depend on the same set of conditioning variables. The quote return volatility we employ here is based on the estimation of an MA(1) process for the conditional mean of quote revisions, with a GARCH(1,1) conditional variance specification.

$$r_t = \gamma_0 + \nu_t + \theta \nu_{t-1}, \quad \nu_t | I_{t-1} \sim N(0, \sigma_{rt}^2) \quad (2.5)$$

$$\sigma_{rt}^2 = \delta_0 + \delta_1 \nu_{t-1}^2 + \delta_2 \sigma_{rt-1}^2 \quad (2.6)$$

To complement the parametric estimation of the above system, in Figures 2.1 and 2.2 we present the time-plots of the spread and revision volatility in our D2000-2 snapshot, the latter measure being the conditional variance series retrieved from our MA(1)-GARCH(1,1) specification for revisions. The graphs clearly demonstrate a high contemporaneous correlation between the inside spread and GARCH volatility in the market.

The results from ML estimation of the above system are displayed in Table 2.9. As the table indicates, the effects outlined above come through very clearly. The dominant impact on the conditional mean of the spread is its own first lag, although there is also a strong impact from volatility. Within the conditional variance esti-

mation there is a significant impact from return volatility (γ_1) and a definite effect from the lagged level of the spread to its volatility (γ_2).

Hence our analysis of the spreads between firm quotes bears out the results derived in Bollerslev and Melvin (1994). There is, however, in our case an alternative explanation of this empirical relationship, based on the substitution of dealer activity from brokered to direct trade. Essentially, one can view posting a limit order as writing a free option. In times of high volatility this becomes a less attractive prospect and hence dealers migrate from the brokered market towards direct trade. The effect of this is to thin the brokered segment of the market out endogenously as volatility rises, such that the spread increases.¹⁷ However, it is not possible to distinguish whether this effect or those outlined earlier underlie the spread/volatility relationship observed in our data.

We extend the analysis of Bollerslev and Melvin (1994) by examining the impact of transaction activity on spreads. Intuitively, when transaction volume is high and there is the possibility of informed trading, the spread in the market should widen and transaction activity should impact upon the spread through the conditional mean equation. Hence we add to equation (3) by including a transaction indicator, (taking the value unity when a deal occurs and zero at all other times,) current and lagged, to the equation.

The results in Table 2.10 demonstrate the impact of transactions on spreads. Both the current transaction and first lag are strongly significant (parameters λ_3 and λ_4 in Table 2.10), indicating that the spread widens when a deal occurs and that this widening persists through time, (through the impact of the lagged dependent variable and the lagged transaction variable.) These effects tie in both with our intuition and our a priori knowledge of the quote dynamics of the system. In many cases the impact of a transaction on the quote for one side of the market is to

¹⁷We would like to thank an anonymous referee for emphasising this alternative interpretation of our results to us.

exhaust the associated quantity. This entails an immediate quote revision, as the system posts the next most competitive quote on that side of the market, explaining the contemporaneous effect of transactions on spreads. The further increases in the spread can be justified with respect to the theoretical arguments given above i.e. that the transaction induces market makers to increase their subjective probabilities over the occurrence of an information event and hence they widen their spread.

2.5 Conclusion

In this study we have examined the dynamic interactions in the foreign exchange market between trades, quotes and spreads, using data from Reuters electronic D2000-2 system. A particular incentive for this study was to attempt to reconcile findings of negative, first-order autocorrelation in quote revisions, while trades follow a martingale. We believe that in this market the negative autocorrelation in quote returns is owing to it being 'thin', without subsidiary orders lying close behind the best bid or ask. That such a market would be thin had been analytically hypothesized by Ho and Stoll (1983) (and our results are consistent with that.) When the quantity offered at the best bid (or ask) is fully taken up (exhausted) by a transaction, or the quote is removed for some other reason, then in such a thin market the quote is likely to jump some distance from the effective going rate at which a transaction can be made. Hence the quote will have to be improved before a trade is likely to be made. [N.B. This rationale for negative autocorrelation cannot be carried directly over also as an explanation for the same phenomenon in the indicative FXFX return series, although it may have the same root i.e. that off-market quotations are recognised as such and quickly corrected.]

The foreign exchange market, and the players on D2000-2, consist primarily of competing market makers in the big banks. In this respect it both differs from, and has a more complex pattern of dynamic interactions than, the specialist market at the

NYSE, which has been most extensively studied heretofore. In particular, quotes will (usually) be entered by different agents at different times, so quote revisions will respond to (prior) quote changes on the other side of the market, and the spread becomes an endogenous, and not a choice, variable. We demonstrate the interrelationships between trades, quotes and spreads, repeating for example the exercise performed by Bollerslev and Melvin (1994).

The obvious limitation of this exercise is the relatively small size and particularity of the data set. While it is a large enough sample numerically to be reasonably confident of assessing the nature of these interactions within the data set itself, we cannot be at all sure whether these seven hours of activity over D2000-2 are representative of D2000-2 at other times, nor of course whether activity over D2000-2 is representative of activity in the wider market. To assess the microstructural influences which affect the intra-day forex market more fully, there is a clear need for more data of this type.

TABLES

Table 2.1: Quote Revisions in Tick-time

Ask	Coeff.	T-stat.	Bid	Coeff	T-stat.
α_0	0.055	0.195	β_0	-1.677	-3.088
α_1	-0.519	-5.63	β_1	-0.393	-5.319
α_2	2.238	5.316	β_2	-2.227	-4.074
α_3	1.001	2.23	β_3	-1.107	-3.648
α_4	-0.049	-1.296	β_4	0.384	4.556
α_5	0.09	1.378	β_5	0.293	2.769
R^2		0.34	R^2		0.46

Table 2.2: Ask Revisions in Tick-time: Sub-sample study

Ask(1)	Coeff.	T-stat.	Ask(2)	Coeff	T-stat.
α_0	1.179	2.356	α_0	0.283	0.586
α_1	-0.369	-3.084	α_1	-0.527	-4.747
α_2	1.486	3.88	α_2	3.082	3.924
α_3	0.452	1.086	α_3	1.485	1.841
α_4	-0.309	-3.01	α_4	-0.052	-1.233
α_5	0.299	3.283	α_5	0.06	0.724
R^2		0.376	R^2		0.357

Notes: In the above table Ask(1) represents the first sub-sample of ask side revisions (observations 1-429) and Ask(2) represents the second sub-sample (observations 430-790.)

Table 2.3: Bid Quote Sub-sample Analysis

Cat.	Mean Δb_t	s.e.	Mean Δb_{t+1}	s.e.	ρ	% rev.
B1	-4.378	0.446	2.605	0.451	-0.676	-0.684
B2	1.95	0.802	0.067	1.042	-0.421	-0.547
B3	0.175	0.714	-1.111	0.785	-0.601	-0.66
B4	3.54	0.536	-1.77	0.621	-0.777	-0.9
B5	-4.615	0.536	2.282	0.847	-0.748	-1.18

Notes: in the above table ρ is the correlation between each of the bid revision series (Δb_t) with the series of their immediate successors (Δb_{t+1}), %rev. is the slope coefficient from a linear regression of Δb_{t+1} on Δb_t , hence representing the proportion of a revision which is immediately reversed.

FIGURES

Figure 2.1: Spreads on D2000-2

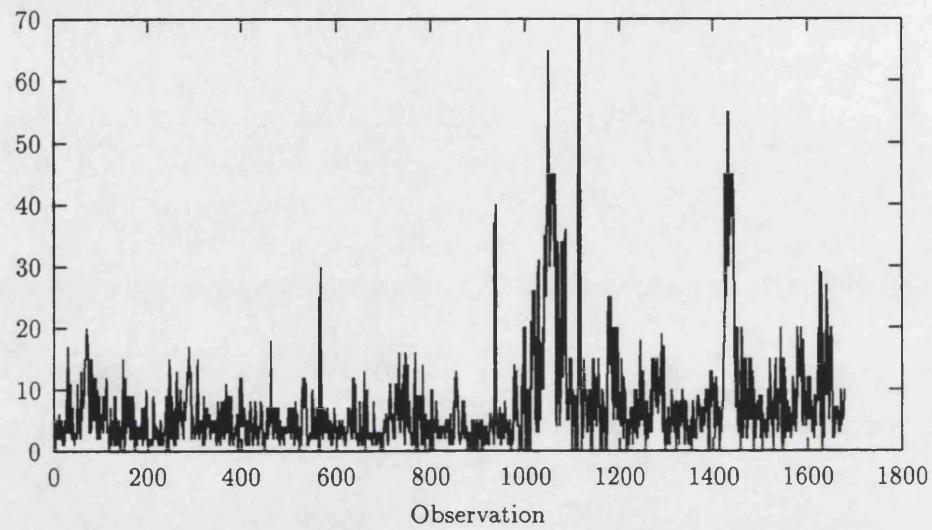


Table 2.4: Ask Quote Sub-sample Analysis

Cat.	Mean Δa_t	s.e.	Mean Δa_{t+1}	s.e.	ρ	% rev.
A1	4.81	0.494	-2.39	0.512	-0.797	-0.826
A2	0.62	0.58	-0.22	0.558	-0.209	-0.201
A3	-0.81	1.234	-1.00	0.955	-0.715	-0.553
A4	4.36	0.584	-1.81	0.836	-0.479	-0.687
A5	-3.27	0.243	2.05	0.556	-0.456	-1.04

Notes: in the above table ρ is the correlation between each of the ask revision series (Δa_t) with the series of their immediate successors (Δa_{t+1}), %rev. is the slope coefficient from a linear regression of Δa_{t+1} on Δa_t , hence representing the proportion of a revision which is immediately reversed.

Table 2.5: First Period Bid Quote Sub-sample Analysis

Cat.	Mean Δb_t	s.e.	Mean Δb_{t+1}	s.e.	ρ	% rev.
B1	-3.373	0.43	1.8	0.479	-0.826	-0.921
B2	0.947	0.695	0.184	0.852	-0.496	-0.608
B3	-0.976	0.611	-0.976	0.946	-0.435	-0.674
B4	2.217	0.321	-0.76	0.397	-0.316	-0.39
B5	-2.952	0.356	1.26	0.721	-0.586	-1.186

See notes to Table 2.3.

Table 2.6: First Period Ask Quote Sub-sample Analysis

Cat.	Mean Δa_t	s.e.	Mean Δa_{t+1}	s.e.	ρ	% rev.
A1	3.114	0.308	-1.743	0.351	-0.694	-0.791
A2	0.926	0.773	-0.63	0.529	-0.363	-0.248
A3	-2.294	1.1	0.882	0.503	-0.308	-0.141
A4	3.317	0.812	-2.39	1.018	-0.904	-1.134
A5	-2.344	0.281	0.766	0.393	-0.416	-0.581

See notes to Table 2.4

Figure 2.2: GARCH Return Volatility on D2000-2




Table 2.7: Bid quote Return Autoregressions

Variable	Raw Data		Filtered Data	
	Coefficient	t-stat	Coefficient	t-stat
Constant	0.167	1.464	1.188	6.594
Lag 1	-0.578	-16.192	-0.299	-5.859
Lag 2	-0.111	-2.700	-0.017	-0.327
Lag 3	-0.081	-1.964	-0.023	-0.443
Lag 4	0.02	0.496	-0.184	-3.507
Lag 5	0.12	3.362	-0.06	-1.177
R^2		0.304		0.112

Table 2.8: Ask quote Return Autoregressions

Variable	Raw Data		Filtered Data	
	Coefficient	t-stat	Coefficient	t-stat
Constant	0.178	1.765	-0.475	-1.215
Lag 1	-0.572	-15.997	-0.007	-0.143
Lag 2	-0.001	-0.014	0.124	2.437
Lag 3	-0.015	-0.372	0.037	0.717
Lag 4	-0.123	-2.992	-0.033	-0.647
Lag 5	-0.057	-1.587	0.061	1.199
R^2		0.322		0.022

Table 2.9: Quote Revisions: MA(1)-GARCH(1,1)

Spread	Coeff.	T-stat
λ_0	1.849	12.006
λ_1	0.135	6.878
λ_2	0.58	23.695
γ_0	2.067	41.392
γ_1	0.042	12.114
γ_2	0.067	15.329
LogLik	-3288.6	
$Q^2(10)$	7.01	

Note: *LogLik* is the maximised value of the likelihood function, $Q^2(10)$ is the Box-Ljung statistic for up to 10th order serial correlation in the standardised squared residuals.


Table 2.10: Quote Revisions: MA(1)-GARCH(1,1)

Spread	Coeff.	T-stat
λ_0	1.237	7.006
λ_1	0.134	6.804
λ_2	0.608	26.98
λ_3	1.228	5.766
λ_4	0.591	2.888
γ_0	2.061	42.47
γ_1	0.043	12.24
γ_2	0.063	14.01
LogLik	-3243.4	
$Q^2(10)$	4.34	

See notes to Table 2.9.

FIGURES

Figure 2.1: Spreads on D2000-2



Chapter 3


Information Transmission in Inter-Dealer Foreign Exchange Transactions

3.1 Introduction

In the last ten years there has been a significant shift in the structure of the inter-dealer market for foreign exchange (FX). Whereas traditionally, more-or-less all trading activity was carried out directly over the telephone or through conventional voice brokers, more recently electronic trading systems, such as those provided by Reuters and EBS/Minex, have captured significant market shares, especially from their human counterparts. As inter-dealer activity accounts for around 70% of spot volume it therefore seems worthwhile to examine these electronically brokered inter-dealer trades and the properties of the resultant quotations and transactions series.¹ This is precisely the focus of the current work.

¹See Bank for International Settlements (1993) for a detailed breakdown of FX volumes.

Figure 2.2: GARCH Return Volatility on D2000-2



Microstructure theory gives two possible explanations for the huge proportion of inter-dealer FX activity. The first is that the inter-dealer market has developed in order to allow dealers to control their inventories (see for example Ho and Stoll (1983).) Market makers experiencing shocks in their (non-dealer) customer order flow will desire to pass the unwanted inventory off to other dealers. This 'hot-potato' trading, as Lyons (1995) terms it, will then imply large volumes being passed around the inter-dealer market.

A second explanation abstracts from this inventory motivation, focussing on informational asymmetries as the driving force behind inter-dealer trades. The spot FX market is decentralised and almost completely opaque. In particular, the customer order flow entering a given bank is completely unobservable to any other participant. This implies that one rôle which the inter-dealer market might fulfil is as a mechanism for the transmission of information between participants, as modelled in Perraudin and Vitale (1996) and Lyons (1995). Dealers who experience a large and informative customer trade can exploit the information embodied and profit through trading on the inter-dealer market.

The latter explanation is the focus of the current work. Using the VAR approach, developed in Hasbrouck (1991a) and Hasbrouck (1991b), we attempt to quantify the information content of transactions by examining the long run impact of trades on quotations. Note that only information effects, be they from private or public sources, should alter equilibrium quotations, the effects of inventory rebalancing and other microstructure effects should disappear as they will not affect equilibrium beliefs. This focus does not imply that we believe that inventory effects are unimportant; on the contrary we believe that in a competitive, decentralised market such as that for FX, inventory effects are likely to be very significant.² Rather, we wish to test the commonly made assertion, contained in Bessembinder (1994) and Huang and Masulis (1995) for example, that informational asymmetries are unlikely to be

²Indeed, Lyons (1995) finds far stronger inventory effects in a sample of FX transactions data from a single dealer/broker pair than are typically found in stock market studies.

Chapter 3

Information Transmission in Inter-Dealer Foreign Exchange Transactions

3.1 Introduction

In the last ten years there has been a significant shift in the structure of the inter-dealer market for foreign exchange (FX). Whereas traditionally, more-or-less all trading activity was carried out directly over the telephone or through conventional voice brokers, more recently electronic trading systems, such as those provided by Reuters and EBS/Minex, have captured significant market shares, especially from their human counterparts. As inter-dealer activity accounts for around 70% of spot volume it therefore seems worthwhile to examine these electronically brokered inter-dealer trades and the properties of the resultant quotations and transactions series.¹ This is precisely the focus of the current work.

¹See Bank for International Settlements (1993) for a detailed breakdown of FX volumes.

Microstructure theory gives two possible explanations for the huge proportion of inter-dealer FX activity. The first is that the inter-dealer market has developed in order to allow dealers to control their inventories (see for example Ho and Stoll (1983).) Market makers experiencing shocks in their (non-dealer) customer order flow will desire to pass the unwanted inventory off to other dealers. This 'hot-potato' trading, as Lyons (1995) terms it, will then imply large volumes being passed around the inter-dealer market.

A second explanation abstracts from this inventory motivation, focussing on informational asymmetries as the driving force behind inter-dealer trades. The spot FX market is decentralised and almost completely opaque. In particular, the customer order flow entering a given bank is completely unobservable to any other participant. This implies that one rôle which the inter-dealer market might fulfil is as a mechanism for the transmission of information between participants, as modelled in Perraudin and Vitale (1996) and Lyons (1995). Dealers who experience a large and informative customer trade can exploit the information embodied and profit through trading on the inter-dealer market.

The latter explanation is the focus of the current work. Using the VAR approach, developed in Hasbrouck (1991a) and Hasbrouck (1991b), we attempt to quantify the information content of transactions by examining the long run impact of trades on quotations. Note that only information effects, be they from private or public sources, should alter equilibrium quotations, the effects of inventory rebalancing and other microstructure effects should disappear as they will not affect equilibrium beliefs. This focus does not imply that we believe that inventory effects are unimportant; on the contrary we believe that in a competitive, decentralised market such as that for FX, inventory effects are likely to be very significant.² Rather, we wish to test the commonly made assertion, contained in Bessembinder (1994) and Huang and Masulis (1995) for example, that informational asymmetries are unlikely to be

²Indeed, Lyons (1995) finds far stronger inventory effects in a sample of FX transactions data from a single dealer/broker pair than are typically found in stock market studies.

present on the spot FX market.

The analysis of the asymmetric information hypothesis is complicated by the nature of the data set. The system from which our data is derived is a multi-lateral broking system and, as such, has multiple participants quoting prices and consummating trades. This implies that many standard empirical microstructure models used in the analysis of other markets are inapplicable as they are based on single dealer structures. This motivates our use of the VAR framework as it is a general, multivariate time-series model within which the asymmetric information hypothesis can be tested.

The rest of the paper is organised as follows. In Section 3.2 we set out the motivation for this work and the methodology employed. The data which we utilise in this paper emanates from Reuters D2000-2 electronic trading system and has previously been examined in Goodhart, Ito, and Payne (1996) and Chapter 2 of this thesis. Nonetheless, a brief description of the workings of the system and a presentation of the basic features of the data are given in Section 3.2.2. In Section 3.3 we present the results from various specifications of the VAR system before the work is concluded in Section 3.4.

3.2 Motivation and Methodology

3.2.1 Motivation

As stated in Section 3.1, the objective of this study is to uncover evidence of information asymmetries in the inter-dealer market for foreign exchange. To this end we use tools developed in the examination of stock markets to quantify the impact this information has upon quotations and the proportion of all relevant information which such trading reveals. Before detailing the methodology employed we present a motivation for the asymmetric information hypothesis and a brief description of

the structure of the spot FX market.

Essentially the market for FX is a decentralised competing market maker system.³ Spot volume can be classified as emerging from two major segments. The first encompasses trade between dealers and (non-dealer) customers. Trading activity between these two types of agent is performed bilaterally and there is no formal trade disclosure regulation, such that this activity is completely unobservable to any market participants bar the counterparties themselves

The second, and by far the largest, portion of volume is inter-dealer. Here too, much of the activity occurs bilaterally between dealers (so called direct trades) but a significant proportion is brokered and hence at least partially observable to other agents. The data employed in this study is drawn from the brokered inter-dealer market. It emanates from an electronic inter-dealer broking system, Reuters D2000-2, which functions as a closed, electronic, limit order book.

Microstructure theory relevant to the asymmetric information hypothesis (see for example Glosten and Milgrom (1985) and Easley and O'Hara (1987)), however, has been modelled in general around the single dealer (specialist) market structure observed, for example, on the NYSE. The dominant paradigm in this area identifies three types of agent. First, there is the single market maker in a given asset, who quotes bid and ask prices for that asset and then stands ready to trade given amounts at these prices. Second there is a mass of 'noise' traders, agents whose exogenous needs create supplies of and demand for the asset in question and agents from whom the specialist can extract rents through his spread. Lastly, there is a group of informed agents, individuals who have received private signals concerning the future asset value and hence trade strategically in order to profit from their knowledge. Market makers will always lose out in their trades with these informed agents, balancing these losses with the profits gained from dealing with the noise traders.

³For a more complete description of the structure of the FX market see Flood (1991).

It is clear that the stylised structure outlined above and the setup of D2000-2 have important differences. D2000-2 operates as a continuous auction whereas traditional models assume a single dealer structure. Also all participants on D2000-2 are market-makers, blurring the distinctions between the three agent classification given above. Finally, given the fact that all participants are market makers, it becomes difficult to pinpoint the source of informational asymmetries.

A mechanism which has been posited to motivate the existence of informational asymmetries between dealers (adopted in Lyons (1994), Lyons (1995) and Perraudin and Vitale (1996),) relies on the lack of transparency of customer order flow. Each bank, in each trading interval, is posited to receive a noisy signal of the asset value through the net demand for currency from their customers. As these orders are entirely unobservable to any other bank, the signal represents private information. The banks then come together in the inter-dealer market where any informational advantages are exploited and, in equilibrium, information is disseminated and fully incorporated into quotations. Differentials in the amount and quality of information held across banks may simply be caused by differences in the size and type of their customer bases. A mechanism similar to the above is also utilised in Vogler (1993) who examines the incentives for inter-dealer trade in a competing market maker system and in the stylised model of the London Stock Exchange contained in Naik, Neuberger, and Viswanathan (1994).

Given the above mechanism, an inter-dealer trade will cause a permanent revision in equilibrium inter-dealer quotations (upwards for an inter-dealer market buy and downwards for a market sell,) just as the traditional single dealer models would predict. This effect is due to the possibility that the trade is information motivated and it is this phenomenon which is utilised in our empirical analysis of D2000-2. An assumption implicit in the above is that dealers submit market orders to D2000-2 in order to exploit information. This can be justified by noting that submitting a

limit order will noisily reveal ones information and hence erode ones advantage.⁴

The asymmetric information hypothesis motivates the analysis of the following questions. First, are quotations permanently (and systematically) revised following market orders? Using the VAR model described in subsection 3.2.3 we tackle this issue by calculating the quote impulse responses to transactions. From this we can examine which types of transactions have the greatest impact on revisions. For example, is the ultimate impact related, possibly non-linearly, to volume as the theoretical model of Easley and O'Hara (1987) would suggest? The second major point we cover is; how much do transactions contribute to the evolution of the efficient price path i.e. do trades reveal a significant proportion of total market information?

Due to the lack of available transaction based FX data other work which examines the existence of asymmetric information in FX markets is sparse. Lyons (1995) uncovers evidence of information based trading (as well as inventory control effects) using a data set derived from the activity of a single U.S. dealer and broker pair over a period of five trading days. The nature of these data allows Lyons to employ an extension of a well-known model from the extant microstructure literature, that of Madhavan and Smidt (1991). Unfortunately, this is not possible in the current context, since the data used here are derived from the activities of multiple, unidentifiable dealers. The only other paper of which we are aware which tackles this issue is Ito, Lyons, and Melvin (1996). Their work focuses on the removal of restrictions on inter-dealer trading in Tokyo during lunch hours. Utilising a variance ratio framework, à la French and Roll (1986), they find that the lunchtime Yen/Dollar return variance rose sharply after the abolition of the aforementioned restriction, evidence which is interpreted as supporting the hypothesis of the existence of informational asymmetries between dealers. Further results in their paper support this hypothesis.

Finally, we take another look at the ‘thin market hypothesis’ studied in Chapter

⁴For a comprehensive analysis of the incentives to submit limit versus market orders see Harris (1995b).

2. This phenomenon can be re-examined, in a dynamic context, using the current framework. Further, one potential reason for this lack of depth is that movements in the quotation series are largely permanent. If this is the case, then dealers, assuming that orders cannot be instantaneously adjusted, are loathe to enter orders subsidiary to the best quotes as prices may move through them and they may be executed at a loss. Hence by decomposing the quote time-series into permanent and transitory components we can assess whether this explanation for 'thinness' is potentially valid.

3.2.2 The Dealing System and the Data

The data we use is derived from Reuters D2000-2, a multi-lateral electronic FX broking system.⁵ As such, its essential function is to act as an electronic limit order book, compiling orders received from individual participants, removing those which have been cancelled and handling the effects of transactions on the book. At any moment a dealer linked to the system can observe, for up to 6 currencies, the most competitive bid and ask quotes plus the quantities available for trade on each side of the book. Further, there is an indicator of the last recorded transaction, and both the side of the market on which this transaction occurred and the associated quantity can be deduced easily from available information. With regard to limit order entry, strict rules of price and time priority apply. Transactions can be initiated virtually instantaneously by any participant, specifying the quantity involved and whether one is selling or buying, at which point all related information is updated on the screen.

The data only encompasses seven hours of trading in DM/Dollar, recorded on June 16th 1993. This is clearly a small sample, implying that any results must be interpreted with due caution. From the transcriptions of the screens the available data series are the best bid and ask quotes, hence the inside spread, quantities on

⁵For a more complete examination of the workings of D2000-2 and the characteristics of these data refer to Goodhart, Ito, and Payne (1996) and Chapter 2.

either side and signed transactions, including the associated volume. In this study we combine the bid and ask quotes into a single midpoint series, also aggregating transactions on each side of the market to give a signed deal time-series (i.e. a series in which a market buy is denoted +1, a market sell -1 and no transaction is assigned zero value,) plus a signed volume series. We combine the variables by constructing an activity-scale data set, an observation being recorded whenever a midpoint revision or transaction occurs. The basic statistical features of the revision, transaction and spread series are given in Table 3.1.

The mean return ⁶ for the DM/Dollar is positive, reflecting the fact that the exchange rate traded upward during this data period. Also, as in many other studies of financial market returns, these returns display excess kurtosis, a fact which can be explained, at least partially, by the existence of conditional heteroskedasticity in the return series. The mean spread over D2000-2 is approximately seven basis points, a very similar figure to the mean spread in the widely used FXFX data. However, the spread distribution is highly skewed on D2000-2, the modal spread being three basis points as compared to a figure of 5 b.p.'s for FXFX spreads. This demonstrates that the inside spreads on D2000-2 are generally bracketed by those on FXFX screens. The volatility of the return series is, as expected, time-varying, with estimation of a standard GARCH(1,1) process for returns demonstrating approximate IGARCH behaviour.

Lastly in Table 3.1 we report the summary statistics for three series which represent transactions. The first, x_0 , a simply a signed, binary indicator. There are 440 transactions in the seven hour period and slightly more market buys than sells, as evinced by the positive mean value for x_0 . The same observations hold for the second transaction series, x_s , which represents signed volume. Finally, x_e is a signed exhaust indicator. We define an 'exhaust' as a transaction which completely fills

⁶The return is defined as the first difference of the midquote. Results presented later in the paper are unaffected by the use of percentage returns instead of price changes.

the extant best order on one side of the book.⁷ There are 296 of these orders, with roughly equal numbers at the bid and ask. We separate these deals from the rest as, a priori, we know that the contemporaneous midquote response to exhausts must differ from the response to non-exhausts, as for the former group the electronic limit order book will immediately post the next most competitive quote, whilst for the latter group there is no contemporaneous midquote alteration.

A final point to note about these data is the following. As detailed in Chapter 2 the quotation series can be seen to be subject to two distinct regimes. In the opening half of the data set the market appeared to be subdued, with little volatility and small spreads. However, approximately two-thirds of the way through the period this picture changed radically with the exchange rate jumping sharply upwards, bringing on a period of high spreads and volatility which persisted to the end of the data set.⁸ The behaviour of the quotation process is shown in Figure 3.1. In later sections of the paper we estimate the VAR specifications for these two regimes separately in order to test the robustness of our findings.

3.2.3 Methodology

As discussed previously, the nature of the D2000-2 data implies that standard empirical microstructure models based on single dealer structures are not suitable in the current context. The solution to the above problem which is adopted in the current work is to utilise a general, dynamic model of quotations and trade characteristics. We employ the VAR formulation proposed by Hasbrouck (1991a) and refined in Hasbrouck (1991b) in order to quantify the information content of trades. In the equations which follow, r_t is the revision in the quote midpoint at each observation

⁷Note also that on D2000-2 a trader is not permitted to 'walk up one side of the book' by entering, for example, a market buy order with execution price above the most competitive limit sell and filling all limit orders with prices less than or equal to the market buy. A trader can only hit the extant best limit order.

⁸See Goodhart, Ito, and Payne (1996) for a linkage between this appreciation in the dollar and the macroeconomic information being released at the time.

and x_t is a vector of trade characteristics (e.g. direction and size of trades.) The structure of the model is as follows.

$$r_t = \sum_{i=1}^p \alpha_i r_{t-i} + \sum_{i=0}^p \beta_i x_{t-i} + \epsilon_{1t} \quad (3.1)$$

$$x_t = \sum_{i=1}^p \gamma_i r_{t-i} + \sum_{i=1}^p \delta_i x_{t-i} + \epsilon_{2t} \quad (3.2)$$

The following restrictions are placed on the innovations,

$$E(\epsilon_{1t}) = E(\epsilon_{2t}) = E(\epsilon_{1t}\epsilon_{2t}) = 0$$

$$E(\epsilon_{1t}^2) = \sigma^2, E(\epsilon_{2t}^2) = \Omega$$

$$E(\epsilon_{1t}\epsilon_{2s}) = E(\epsilon_{1t}\epsilon_{1s}) = E(\epsilon_{2t}\epsilon_{2s}) = 0, \quad \forall t \neq s$$

From equations (3.1) and (3.2) one can see that we allow quite general patterns of cross and serial correlation in both r_t and x_t . By analysing the Granger Causality patterns apparent in estimation one can give first-pass answers to questions regarding the impacts of trades on quotes and vice-versa. Note that the VAR is not entirely conventional as the *current* trade characteristic, x_t , is assumed to have a direct impact on quote revisions through the parameter β_0 . The model also allows non-instantaneous influence of trades on quotations which may reflect inefficiency in information processing. The trading pattern is also allowed to be endogenous, an outcome which may be due to microstructure imperfections. Identification of the VAR is achieved by setting the contemporaneous correlation between the innovations to zero such that all influence of x_t on r_t is directly via β_0 .⁹ Finally, the two innovations are assumed to be both mutually and serially uncorrelated at all leads and lags.

⁹This is equivalent to a Sims/Choleski orthogonalisation of an unidentified VAR system.

Information, in the context of the above model, can come from one of two sources. Innovations to the public information set (as well as transitory quote returns) enter the system through ϵ_{1t} whereas the trade innovation, ϵ_{2t} , reflects unpredictable transaction activity and hence the possibility of trades arising from private information. The information content of a trade is defined as the average, long-run response in the quote-midpoint which a transaction causes and this is calculated via the impulse response representation of the VAR. Rewriting equations (3.1) and (3.2) in matrix notation we get the following;

$$\begin{pmatrix} 1 & -\beta_0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} r_t \\ x_t \end{pmatrix} = \begin{pmatrix} \alpha(L) & \beta(L) \\ \gamma(L) & \delta(L) \end{pmatrix} \begin{pmatrix} r_t \\ x_t \end{pmatrix} + \begin{pmatrix} \epsilon_{1t} \\ \epsilon_{2t} \end{pmatrix} \quad (3.3)$$

where the elements of the right hand side matrix which multiplies the vector of r_t and x_t are polynomials in the lag operator formed from the coefficients in the original representation (equations (3.1) and (3.2).) and

$$VCV \begin{pmatrix} \epsilon_{1t} \\ \epsilon_{2t} \end{pmatrix} = \begin{pmatrix} \sigma^2 & 0 \\ 0 & \Omega \end{pmatrix}$$

This system can then be inverted to give,

$$\begin{pmatrix} r_t \\ x_t \end{pmatrix} = \begin{pmatrix} a(L) & b(L) \\ c(L) & d(L) \end{pmatrix} \begin{pmatrix} \epsilon_{1t} \\ \epsilon_{2t} \end{pmatrix} \quad (3.4)$$

where the new lag polynomials are formed by the inversion of the autoregressive system in equation (3.3). Equation (3.4) is the Vector Moving Average (VMA) representation of the original system. The coefficients on the disturbance terms give precisely the impulse responses desired i.e. the impact of a unit innovation in the quote revision equation ($\epsilon_{1t} = 1$) on the quote revision, k periods ahead, is given by

a_k whilst the effect of a unit trade innovation on the quote midpoint revision is b_k at a k -period horizon (see, for example, Hamilton (1994).) Effects on quotations are then the sum of these impacts, with the long-run response simply being the infinite sum. Hence, we can quantify the information content of a trade as the cumulative response of the revision to the trade innovation i.e. $\sum_{i=0}^{\infty} b_i$.¹⁰

This enables us to derive answers to the first set of questions. The adjustment path of quotations and the equilibrium impact of a trade can be derived by calculating the sums of the b coefficients over given horizons. Examination of the type of trades, as further explained in Section 3.3, which engender revisions can be performed by altering the make-up of the vector x_t . Further, dynamic evidence on the depth of the order book can be gauged from the shape of the impulse response functions.

The second set of questions posed related to the characteristics of the permanent component of the quotation process. Specifically we want to examine to what extent the quotation process is permanent and, second, what proportion of the variation in this permanent component is explicable through trading. To this end we employ the methodology contained in Hasbrouck (1991b). The model begins with a representation for quote midpoints (q_t) as shown below.

$$q_t = m_t + s_t \quad (3.5)$$

This equation decomposes the quote series into a permanent component (m_t), which is assumed to follow a random walk; plus mean zero, transitory noise (s_t). These conditions are formalised as follows,

$$m_t = m_{t-1} + \omega_t,$$

¹⁰Obviously in practice this sum will not be infinite but truncated at a point at which the innovations have become negligible.

$$\omega_t \sim N(0, \sigma_w^2), \quad E(\omega_t \omega_s) = 0 \quad \text{for } t \neq s,$$

$$\lim_{k \rightarrow \infty} E_t s_{t+k} = \lim_{k \rightarrow \infty} E s_{t+k} = 0,$$

and also ω_t and s_t are assumed to be jointly covariance stationary. These conditions can be interpreted as follows. The I(1) process, m_t , follows a simple random walk. It can, hence, be thought of as the efficient price process, which by the efficient markets hypothesis is well represented by a random walk. The component s_t can have an arbitrary I(0) specification and, as such, can clearly have no impact on long-run quotations. It can hence represent any kind of transitory microstructure imperfection such as inventory control effects or price discreteness, although the latter phenomenon is likely to be far less important in FX markets than in stock markets.¹¹

Next returns are defined (simply as $r_t = q_t - q_{t-1}$) and the same VAR specified as above in equation (3.3) is proposed in quote revisions and transactions. After retrieving the VMA representation (as in equation (3.4)), one can construct the following two measures. The first is the variance of the permanent component, whilst the second is the contribution of trades to this random walk variance (see Hasbrouck (1991b), Appendix A.)

$$\sigma_w^2 = \left(\sum_{i=0}^{\infty} b_i \right) \Omega \left(\sum_{i=0}^{\infty} b'_i \right) + \left(1 + \sum_{i=1}^{\infty} a_i \right)^2 \sigma^2 \quad (3.6)$$

$$\sigma_{wx}^2 = \left(\sum_{i=0}^{\infty} b_i \right) \Omega \left(\sum_{i=0}^{\infty} b'_i \right) \quad (3.7)$$

¹¹This is due to the fact that the inter-dealer FX market has no pre-specified minimum tick size, unlike most stock markets, such that quotes can theoretically be entered on as fine a basis as one wishes.

An intuitive explanation of these equations is as follows. The quotation series incorporates two distinct sources of information, public information via the ϵ_{1t} innovation in equation (3.4) and the information revealed through trading activity. The cumulative effect of a unit innovation in ϵ_{1t} on quote revisions, and hence the long-run impact on quotes, is given by the sum of the a_i coefficients from equation (3.4) plus unity (the contemporaneous impact.) Hence the variance of the permanent impact of ϵ_{1t} is given by the square of this measure (i.e. the square of $(1 + \sum_{i=1}^{\infty} a_i)$), multiplied by the variance of ϵ_{1t} . This variance statistic represents the variation in returns caused by the revelation of public information. Similarly, the sum of the b_i coefficients from equation (3.4) represents the long-run impact of trades on quotations. The variance of the trade-correlated permanent component is therefore, in the case where x_t is scalar, equal to the variance of trade innovations multiplied by the square of $\sum_{i=0}^{\infty} b_i$. The representation in equation (3.7) is simply a matrix formulation of this, where x_t is a given vector of trade attributes.

The results presented in equations (3.6) and (3.7) rest on a number of assumptions. As noted by Quah (1992), however, the crucial assumption here is the restriction that m_t follows a random walk, rather than a more general I(1) process. In terms of the frequency domain analysis of quote returns, the effect of this decomposition is to equate the spectra of r_t and Δm_t at frequency zero and to require the spectral density of Δm_t to be constant at this level.¹² This allows one to tie down the variance of the permanent component of one's series precisely as it is simply the integral under the spectrum of Δm_t . Quah (1992) criticises the random walk assumption as both ad hoc and restrictive, but in the current context the efficient market hypothesis provides a good rationale for its use.

From this analysis we obtain three statistics which summarize the behaviour of the

¹²These conditions come from the following facts. First, s_t is a covariance stationary process and hence its first difference has spectral density zero at zero frequency. Second, any cross-spectral elements will have zero spectral density at the origin due to the first differencing operator. Finally, as m_t is a random walk, its first difference is white-noise and hence has a constant spectral density from zero to π .

permanent component of the quotation series. The first two are the variance of the random walk component (σ_w^2) and the trade-correlated component (σ_{wx}^2) defined in equations (3.6) and (3.7), plus the ratio of the standard deviation of the permanent component to the standard deviation of r_t . The second of these measures allows us to address our earlier question about the importance of trading in defining the efficient price path, whilst the third allows us to examine to what extent quote fluctuations have a significant permanent element.

3.3 Results

This section presents the estimations of the model presented in 3.2.3. The section divides into three parts; the first covers the basic VAR estimations, the second checks the robustness of findings via a sub-sample analysis and the third presents the variance decompositions.

3.3.1 Do trades cause Permanent Quote Revisions?

In this subsection we present the estimations associated with the VAR structures outlined in Section 3.2.3. First the parameter estimates are shown and interpreted, along with the causality patterns which the results imply and we then go on to derive the VMA representation and graphically show the Impulse Response functions.¹³

The first VAR specification we estimate corresponds identically to equations (3.1) and (3.2) but with x_t , the vector of trade characteristics containing the trade direction indicator only. The parameter estimates from this model are given in Table 3.2, with χ^2 -tests of the hypotheses that groups of right hand side variables are zero given in Table 3.3. The lag length of the VAR was set at five, a figure chosen by

¹³All results presented in this subsection use heteroskedasticity consistent standard errors in order to mitigate the problem of GARCH disturbances in the revision equation.

using conventional Akaike and Schwartz information criteria and also by noting that the inclusion of five lags is sufficient to remove any residual serial correlation.

In Table 3.3 the following causality patterns appear. In the return equation both the own lags and the trade direction impacts are highly significant. In fact the trade direction effects have greater significance than lagged returns, due to the inclusion of the contemporaneous transaction indicator. In the trade equation, however, the impact of the group of lagged returns is not significant at the 10% level, identifying an asymmetry in causality and contradicting the results presented in Hasbrouck (1991a). A further finding, discussed below, is that lagged trades impact significantly on current trade determination.

Turning next to an examination of the direction and significance of the individual variable impacts, Table 3.2 shows first that the returns series is negatively auto-correlated. This finding is mimicked in many other studies on high-frequency FX rates, (see for example Goodhart (1989), Bollerslev and Domowitz (1993) and Zhou (1992b),) and its appearance in this data set was studied and explained in Chapter 2 using the hypothesis that the order book is thin. Second, the transaction indicator has strong and positive impacts on quote revisions, both instantaneously and up to the fifth lag. These positive responses are in line with the asymmetric information hypothesis in that they predict that a market buy will lead to an increase in the DM/Dollar rate (and vice-versa), reflecting the trade revealing information. The strength of the contemporaneous impact of trades on revisions is also as expected, as many of the deals exhaust the quantity on offer, leading to the system posting the next most competitive quote.

Table 3.2 also indicates strong *positive* autocorrelation in trade sign, demonstrating that in this data set buys and sells cluster. The lack of causality from returns to trades is shown in the t-statistics for lagged returns in the trade equation. The only significant impact is a negative effect of the first lag of returns, implying that market buys, for example, are more likely after price reductions. This effect could be due to

inventory control price-setting by traders (who reduce quotes in order to elicit sales and vice versa.) As noted above, Hasbrouck (1991a) finds a significant, negative impact of the group of returns on trades in the NYSE data and it is the inventory control hypothesis to which he appeals in order to explain this. It is possible that in our data set, with quote entry from multiple market makers, this effect is largely muted.

Hence the causality patterns and directions of impact in the estimated model are largely explicable. We next examine the impulse response functions the system implies. Throughout, we concentrate on the impact of trades on quote revisions as the other impulse responses contain little interesting information. Figure 3.2 presents the quote impulse response to a market buy order for our basic VAR specification. Its shape is consistent with the earlier analysis of the individual parameters. The dominant revision is immediate and around 1.25 basis points (b.p.'s), a result which is clearly due to the effects of exhaustive trades. From then on the quote moves steadily upward until it reaches the neighbourhood of the final equilibrium after a five observation horizon. From then on only small, fine-tuning alterations in the quote level occur, implying an equilibrium impact of approximately 2 b.p.'s. Hence a preliminary answer to the question in the title of this subsection is that trades do cause a significant permanent quote revision and hence can be characterised as information revealing.¹⁴

A limitation of the above results, however, is the paucity of the description of trades. By utilising only a signed trade indicator we omit potentially important information on the characteristics of the transaction. Hence we proceed to test whether trades differ in their impacts on quote revisions across certain identifiable dimensions. The first of these trade characteristics is size. As is well known from Easley and O'Hara (1987) and Lyons (1995), for example, market makers may place a greater posterior

¹⁴It would be interesting to examine how the information content of trading alters over the length of the trading day. Unfortunately in the current study, due to the limitations of the data set, this is not possible.

probability on the revelation of information if they have witnessed a large transaction. The same intuition may be at work here. In order to examine the impact of size we, therefore, run a simple OLS regression, shown in equation (3.8), explaining returns in terms of their own lags, a set of transaction indicator variables (x_{0t}), a set of transaction size variables (x_{st} , calculated as signed volume) and signed squared volume (x_{st}^2). The motivation for the inclusion of the squared variable is the possibility of non-linearities in the size/quote impact relationship.

$$r_t = \sum_{i=1}^5 r_{t-i} + \sum_{i=0}^5 x_{0t-i} + \sum_{i=0}^5 x_{st-i} + \sum_{i=0}^5 x_{st-i}^2 + v_t, \quad v_t \sim N(0, \sigma_{vt}^2) \quad (3.8)$$

The results of the estimation of equation (3.8) do not support any relationship between the size of trades and their impact on returns, either linearly or non-linearly. As shown in Table 3.5, the χ^2 -tests for the influence of the different groups of trade variables demonstrate that only the indicators have significant explanatory power and the the regression R^2 rises by less than 0.01 when compared to that reported in Table 3.2.

Hence, unlike the results of Hasbrouck (1991a) for the NYSE and Lyons (1995) work on the FX market, our FX data do not support the hypothesis that the information content of trades is linked to transaction size. This, however, may be due to the brief period our data cover. Only a small proportion of our transactions differ greatly in size from the mean quantity transacted so that the data may not have the power to uncover an impact from size. Comparing the median transaction quantity on D2000-2 with that of the data used in Lyons (1995) confirms this observation. On D2000-2 median transaction quantity is \$1m. whereas in Lyons data it is \$3m. This feature may also be due to the D2000-2 system only attracting relatively small transactions in general. It is possible that market makers choose to execute large transactions bilaterally in order to take advantage of the lack of transparency this ensures. Only a longer span of data will allow examination of this hypothesis.

There is one dimension, however, along which we know, *a priori*, that the dynamic responses of returns to transactions must differ. As indicated earlier in this subsection, a transaction which exhausts the outstanding quantity immediately entails the system posting the next most competitive limit order. Hence for these exhaustive transactions there must be a significant contemporaneous quote revision, whereas for non-exhaustive deals the immediate impact should be zero. We embody this knowledge into our VAR specification by constructing the vector of trade characteristics (x_t) from the signed transaction indicator (x_{0t}) and a signed exhaust indicator (x_{et}).¹⁵ The results from the VAR estimated using this specification of the trade characteristic vector are given in Tables 3.6 and 3.7.¹⁶

The χ^2 -tests in Table 3.7 bear out our intuition about the differing dynamic responses of returns to the two sets of transactions. Both the indicator and exhaust groups of variables are significant at the 1% level, although it is clear that the far stronger causality runs from exhausts to returns. The estimates and significance of single parameters, shown in Table 3.6, shows that, as expected, the bulk of the explanatory power of the set of exhaust variables comes from the contemporaneous variable, whereas the contemporaneous trade indicator has an insignificant and very small impact on returns. Also note that, in general, the effects from lagged exhausts on revisions are negative whereas those from the lagged indicator variables are positive. This suggests that the contemporaneous exhaust response overshoots the long-run impact attributable to this type of trade, whilst the reverse is true for non-exhausts.

A clearer picture of the dynamic responses to exhaustive deals and simple indicated transactions is given in Figure 3.3, which graphs the Impulse Response functions of returns to both these sets of transactions.

¹⁵In order to ensure orthogonal innovations in the VAR specification the contemporaneous trade indicator variable is included in the equation for exhaust determination.

¹⁶From here on, only the parameters from the revision equation are reported in order to save space, χ^2 -tests are reported for the significance of groups of variables in all equations however.

This shows that exhausts cause quotations to overshoot their equilibrium response. The immediate impact is approximately 3.25 b.p.'s, but this is reduced by approximately 50% in the following period. From then, the midpoint settles at a level around 2 b.p.'s higher than originally. The equilibrium response to an indicated deal is virtually identical, although the manner in which it is reached is quite different. The immediate impact is 1.25 b.p.'s and from here there is a gradual upward adjustment to the long run quote impact of 2 b.p.'s.¹⁷ To scale the results, the mean half-spread in the data is 3.5 b.p.'s, implying that approximately 60% of the spread can be accounted for as compensation for asymmetric information.

These results also highlight the relatively thin nature of the limit order book on D2000-2. Clearly the large initial response of the midquote to an exhaustive transaction is not reflective of extra information content over indicated deals, but is due to a shortage of liquidity as the equilibrium response to the two types of transaction is identical. This phenomenon was the focus of Chapter 2 which attempted to reconcile the finding of negatively autocorrelated quote returns with serially uncorrelated transaction returns using the above hypothesis.

Figure 3.3 graphically corroborates the hypothesis of a thin limit order book on D2000-2. The fact that the original quote response to exhausts is around the same size as the half spread implies that, unlike results from other trading systems (e.g Biais, Hillion, and Spatt (1995) and de Jong, Nijman, and Röell (1995)), there is little clustering of quotations around the inside spread. Further, the large reversal of the immediate response in the subsequent period supports the thin market hypothesis of Chapter 2 as the source of much of the negative quote return autocorrelation.

¹⁷The large, immediate impact of indicated deals on quotations is due to the fact that these deals might exhaust the outstanding quote, this effect being propagated in the empirical specification by the inclusion of the contemporaneous deal indicator in the exhaust equation.

3.3.2 A Subsample Analysis

The final set of quote/transaction VARs which we estimate are essentially a check on our previous findings. As documented in Section 3.2.2 the quote data can be classified as coming from two distinct regimes, a rather flat opening period and a fairly volatile finish. In order to confirm the findings of the previous subsection we re-estimate the final VAR specification for each regime separately. Tables 3.8 through 3.11 give the estimated coefficients and group significance tests, whilst the estimated impulse responses are shown in Figures 3.4 and 3.5.

The results from the sub-sample VARs broadly confirm our previous finding. For the first, quieter sub-period the dynamic responses of quotes to the two types of deal are similar to those for the entire period. There is now a discrepancy, though, between the responses to exhaustive and indicated deals. The equilibrium response to an exhaust is now just under 2 b.p.'s whilst for indicated deals it is 1.5 b.p.'s. Moreover, the immediate response to an exhaust in this sub-sample is 2.5 b.p.'s, around 0.75 b.p.'s less than for the full sample. In the second sub-sample the picture is slightly closer to that of the full sample. We have the same patterns of response for the two deal classes, with both leading to an equilibrium response of around 2.5 b.p.'s. The immediate responses to both classes are higher than for the full sample, with exhausts leading to an immediate quote response of over 4 b.p.'s. As a proportion of the mean half-spread in each sub-period the quote response to transactions is between 60 and 80% in the first subsample and around 50% in the second subsample.

One potential reason for the difference in immediate responses to deals in the two sub-periods can also be formed in terms of the idea of a thin limit order book. One can essentially view the submission of a limit order as the writing of a free option. In times of high volatility the value of the option increases so that submission becomes less profitable. Hence, the increase in the immediate response to an exhaust in the more volatile, second sub-period may reflect an endogenous thinning of the market.

Our subsample VAR's then, in general, corroborate the earlier results. There are systematic quote changes in response to transactions which are, on average, around 2 b.p.'s. This is approximately 60% of the mean half-spread observed in the data, implying that the group of market makers using D2000-2 charge a significant proportion of their spread in order to compensate themselves for the possibility of informational asymmetries motivating trades. Hence our results agree with the findings in Hasbrouck (1991a) and de Jong, Nijman, and Röell (1995) both of which use a VAR structure to uncover evidence of informational asymmetries, the former for stocks traded on the NYSE and the latter for stocks traded on the electronic CAC system on the Paris Bourse.

3.3.3 The Permanent/Transitory Decomposition

In this subsection we present the results from the decomposition of quote returns via the basic VAR equations, equations (3.1) and (3.2), and the structure in equations (3.5) through (3.7). Three measures are presented: the variance of the permanent component (σ_w^2), the ratio of σ_w to the standard deviation of quote returns (σ_w/σ_r) and the the proportion of σ_w^2 which is trade-correlated (σ_{wx}^2/σ_w^2).¹⁸ We compute all three for VARs in which x_t contains only a trade direction indicator, then for x_t containing both an indicator and exhaust dummy and finally for the two sub-sample periods using the indicator and exhaust specification. Results are shown in Table 3.12.

The measure of the variance of the permanent component (column 1 in Table 3.12) varies quite significantly across the two sub-samples, more than doubling as we move from the quieter opening period to the more volatile final subsample. Hence, in absolute terms, the variance of the permanent component indicates the revelation of far more information (both private and public) in the latter period.

¹⁸ σ_r^2 is computed as the average of volatility retrieved from an estimated GARCH volatility model for the entire return series.

This rise in the variance of the permanent component in subsample 2 is balanced by an overall rise in return variance, so that the ratios of permanent to return standard deviations are similar for both sub-samples at around 0.6. Nevertheless, the permanent component accounts for a large proportion of overall return variation, such that this is one reason which could underlie the thin market on D2000-2. To repeat the reasoning of Section 3.2.3, this effect is propagated by a large permanent component implying that the option value of orders outside the 'touch' increases such that their submission is less profitable.

The basic indicator VAR, and the VAR augmented by the inclusion of exhausts in the trade description vector, both indicate that approximately 32% of the variance of the permanent component can be explained by trades. This suggests that around one third of all information entering D2000-2 during our sample comes via transactions. This is a quite surprisingly high number, given the reservations many authors express when discussing the possibility of information based trading on the FX market. Indeed, it is only slightly lower than the comparable results reported in both Hasbrouck (1991b) and de Jong, Nijman, and Röell (1995), both of which find values for σ_{wx}^2/σ_w^2 between 30 and 40%.

There is, however, quite a large discrepancy between the estimated size of the trade-correlated component for the two subsamples. For the first subsample, just over half of the information flow is revealed by trading whereas in the latter half of the data the comparable figure drops to 20%. This, however, is unsurprising. The first half of the data may be characterised as a period of little public information and low volatility, so the impact of trades may be expected to be far greater. In contrast, the second period is very volatile. Hence although these subperiods give very different indications of the extent to which trades reveal information, the results are explicable. What is clearly needed here is a longer span of data over which to assess the extent to which the electronic inter-dealer FX market is used to intermediate information.

In summary, the permanent/transitory decomposition gives the following indications. On average, around 30% of information incorporated into the DM/Dollar rate is transmitted via trading, although this proportion is sensitive to the state of the market at a point in time. A further point to note is that the permanent component of FX returns seems to account for a high proportion of overall quote return variation. This result is in line with the thin limit order book observed on D2000-2.

3.4 Conclusion

This work examines the commonly made assertion that there is little scope for asymmetric information dealing in FX markets (see for example Bessembinder (1994) and Huang and Masulis (1995).) We find that this assertion is not supported in our seven hour, intra-day data set, confirming the results of Lyons (1995). Transactions are shown to have information content via the VAR formulation of Hasbrouck (1991a). A transaction at the ask, for example has, on average, an upward equilibrium impact on quotes of approximately 2 basis points and vice-versa for a deal at the bid.

We also demonstrate the size of the permanent component of the quotation series. Results show that the ratio of the standard deviation of the permanent component of quote returns to the total quotation variation is approximately 0.6. This fairly large proportion gives a basis for the thin-market hypothesis propounded in Chapter 2. Finally we examine the total information content of the transaction series by calculating the trade-correlated component of the permanent quotation series. This decomposition demonstrates that approximately 30% of the information which enters the market in our seven hours can be attributed to trading activity.

TABLES

Table 3.1: Summary Statistics

Series	No. of Obs.	Mean	s.e.	Variance	Skew	Kurtosis
r	1683	8×10^{-6}	7×10^{-6}	8×10^{-8}	0.28	14.09
x	440	0.039	0.012	0.26	0.06	0.82
x_s	440	0.072	0.027	1.22	0.83	18.06
x_e	296	0.0095	0.01	0.18	0.06	2.69
s	1683	7×10^{-4}	2×10^{-5}	7×10^{-7}	3.19	13.77
v	1683	8×10^{-8}	3.5×10^{-9}	2×10^{-14}	6.60	61.48

Note: The above tables gives the first four sample moments for each data series along with the standard error of the mean (s.e.). r is the midquote revision series. x denotes the transaction indicator series, x_s denotes the signed transaction quantity series and x_e denotes the exhaust dummy. s is the inside spread measured in basis points and v represents the GARCH estimated volatility series.

Table 3.2: Results from the Indicator/Quote Return VAR

Parameter	Value	T-stat	Parameter	Value	T-stat
α_1	-0.4	-6.98	γ_1	-96.43	-2.21
α_2	-0.13	-2.47	γ_2	-20.11	-0.43
α_3	-0.08	-1.89	γ_3	1.63	0.03
α_4	-0.13	-2.47	γ_4	86.28	1.71
α_5	-0.12	-3.24	γ_5	62.73	1.36
β_0	1.25×10^{-4}	11.55	-	-	-
β_1	4.8×10^{-5}	3.36	δ_1	0.118	5.26
β_2	2.3×10^{-5}	1.65	δ_2	0.077	3.22
β_3	9.3×10^{-6}	0.7	δ_3	0.071	2.85
β_4	2.5×10^{-5}	1.76	δ_4	0.019	0.73
β_5	3.3×10^{-5}	2.46	δ_5	0.026	0.94
R^2	0.19		R^2	0.04	
$D.W.$	2.00		$D.W.$	2.00	

Note: Parameter estimates from basic VAR model in equations (3.1) and (3.2). All t -statistics are heteroskedasticity robust. $D.W.$ is the Durbin-Watson statistic

Table 3.3: Basic VAR: Group Significance tests

Equation	r_{t-i}	x_{0t-i}
Return	59.87	182.53
Deal Indicator	8.55	59.08

Note: Each cell is a χ^2 -squared test statistic of the hypothesis that the coefficients on the group of variables indicated in the column heading are zero in the equation indicated in the row heading.

Table 3.4: Size Effects in the Revision Equation

Var.	Value	T-stat	Var.	Value	T-stat	Var.	Value	T-stat
x_{0t}	1.7×10^{-4}	5.11	x_{st}	-3×10^{-5}	-1.18	x_{st}^2	2.4×10^{-6}	0.79
x_{0t-1}	2.4×10^{-5}	0.59	x_{st-1}	1.8×10^{-5}	0.54	x_{2t-1}^2	-1.3×10^{-6}	-0.26
x_{0t-2}	5.3×10^{-5}	1.42	x_{st-2}	-3.2×10^{-5}	-1.09	x_{2t-2}^2	5.1×10^{-6}	1.34
x_{0t-3}	2.4×10^{-5}	0.63	x_{st-3}	-9.7×10^{-6}	-0.31	x_{2t-3}^2	1.1×10^{-6}	0.03
x_{0t-4}	4.4×10^{-6}	0.11	x_{st-4}	1.2×10^{-5}	0.37	x_{2t-4}^2	3.4×10^{-6}	0.01
x_{0t-5}	1.1×10^{-5}	0.30	x_{st-5}	7.2×10^{-6}	0.25	x_{2t-5}^2	2.2×10^{-6}	0.49
R^2	0.20		RSS	1.2×10^{-4}		$D.W.$	2.01	

Note: Parameter estimates of the effects of trade size on quote returns (equation (3.8).) All t -statistics are heteroskedasticity robust. $D.W.$ is the Durbin-Watson statistic and RSS is the regression residual sum of squares.

Table 3.5: Size Effects: Group Significance tests

Equation	x_{0t-i}	x_{st-i}	x_{st-i}^2
Return	30.82	3.40	2.86

Note: Each cell is a χ^2 -squared test statistic of the hypothesis that the coefficients on the group of variables indicated in the column heading are zero in the equation indicated in the row heading.

FIGURES

Figure 3.1: D2000-2 Quote Midpoint Evolution

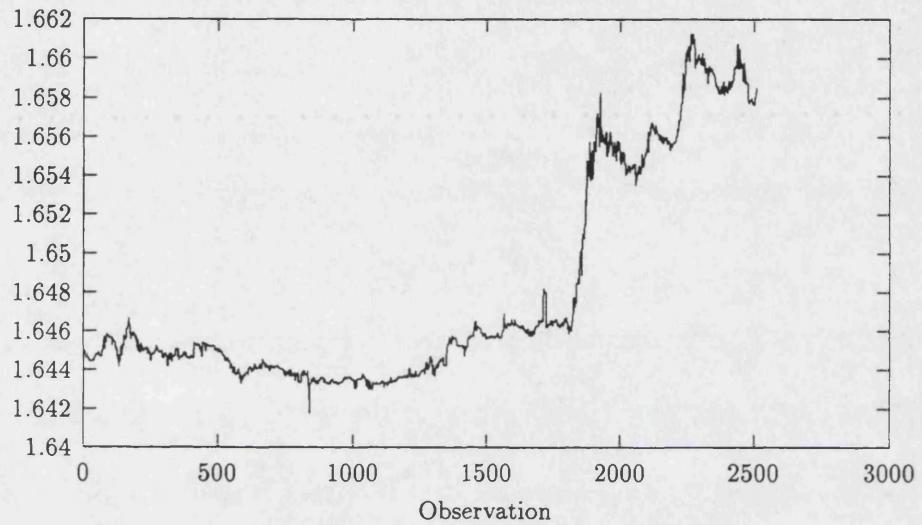


Figure 3.2: Quote Impulse Response Function: Basic VAR

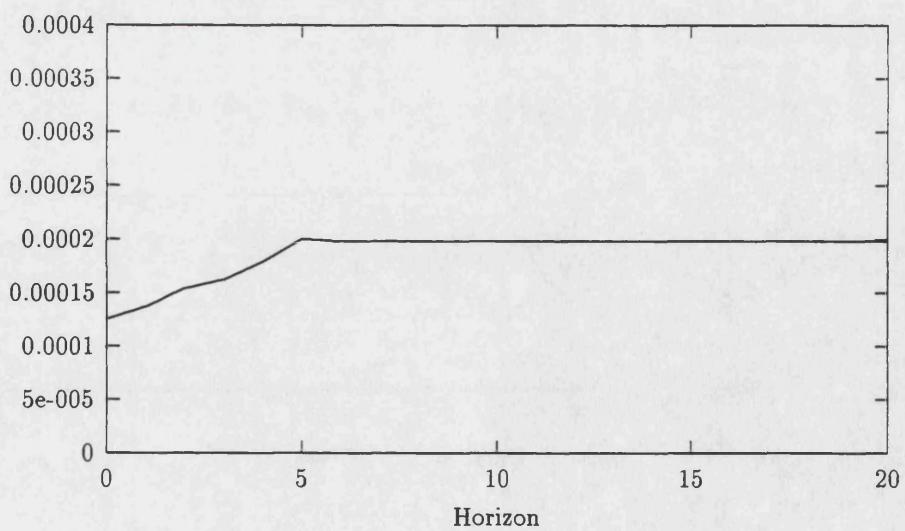


Table 3.6: Revision Equation Parameters from the VAR using both Direction Indicators and Exhausts

Var.	Value	T-stat	Var.	Value	T-stat	Var.	Value	T-stat
-	-	-	x_{0t}	-6×10^{-6}	-0.6	x_{et}	1.95×10^{-4}	10.89
r_{t-1}	-0.37	-6.48	x_{0t-1}	1×10^{-4}	4.04	x_{et-1}	-7.8×10^{-5}	-2.56
r_{t-2}	-0.13	-2.37	x_{0t-2}	4.1×10^{-5}	1.73	x_{et-2}	-2.7×10^{-5}	-0.96
r_{t-3}	-0.08	-1.83	x_{0t-3}	3.8×10^{-5}	1.54	x_{et-3}	-4.4×10^{-5}	-1.54
r_{t-4}	-0.10	-1.94	x_{0t-4}	5.2×10^{-5}	2.56	x_{et-4}	-4.2×10^{-5}	-1.53
r_{t-5}	-0.13	-3.40	x_{0t-5}	2.5×10^{-5}	1.06	x_{et-5}	1×10^{-5}	0.37
R^2	0.23		RSS	1.1×10^{-4}		$D.W.$	2.00	

Note: estimation of revision equation parameters in the VAR utilising trade direction and exhaust indicators in the vector of trade attributes. All t -statistics are heteroskedasticity robust. $D.W.$ is the Durbin-Watson statistic and RSS is the regression residual sum of squares.

Table 3.7: VAR using Deal Indicator and Exhaust Indicator Group Significance tests of sets of Variables in all Equations

Equation	r_{t-i}	x_{0t-i}	x_{et-i}
Return	55.2	27.5	129.2
Deal Indicator	7.1	19.3	1.1
Deal Exhaust	17.1	1.2	1.5

Note: Each cell is a χ^2 -squared test statistic of the hypothesis that the coefficients on the group of variables indicated in the column heading are zero in the equation indicated in the row heading.

Figure 3.3: Quote Impulse Response Functions to Exhausts and Non-Exhausts




Figure 3.4: Quote Impulse Response Functions to Exhausts and Non-Exhaust, Sub-sample 1




Table 3.8: Revision Equation Parameters from the final VAR Specification in Subsample 1

Var.	Value	T-stat	Var.	Value	T-stat	Var.	Value	T-stat
-	-	-	x_{0t}	-2×10^{-5}	-2.07	x_{et}	1.6×10^{-4}	9.7
r_{t-1}	-0.37	-5.21	x_{0t-1}	7.1×10^{-5}	3.91	x_{et-1}	-5.3×10^{-5}	-2.14
r_{t-2}	-0.11	-1.42	x_{0t-2}	1.3×10^{-5}	0.58	x_{et-2}	-1.2×10^{-5}	-0.40
r_{t-3}	-0.12	-1.74	x_{0t-3}	3.5×10^{-5}	1.75	x_{et-3}	-2.7×10^{-5}	-1.13
r_{t-4}	-0.02	-0.32	x_{0t-4}	4.7×10^{-5}	2.76	x_{et-4}	-6×10^{-5}	-2.56
r_{t-5}	-0.03	-0.59	x_{0t-5}	-2.8×10^{-5}	-1.61	x_{et-5}	6.5×10^{-5}	2.79
R^2	0.25		RSS	2.6×10^{-5}		$D.W.$		2.00

Note: estimation of revision equation parameters in the VAR utilising trade direction and exhaust indicators in the vector of trade attributes. All t -statistics are heteroskedasticity robust. $D.W.$ is the Durbin-Watson statistic and RSS is the regression residual sum of squares.

Table 3.9: Subsample 1: Group Significance tests of sets of Variables in all Equations

Equation	r_{t-i}	x_{0t-i}	x_{et-i}
Return	28.6	29.7	102.4
Deal Indicator	5.3	11.0	1.8
Deal Exhaust	10.6	1.7	4.3

Note: Each cell is a χ^2 -squared test statistic of the hypothesis that the coefficients on the group of variables indicated in the column heading are zero in the equation indicated in the row heading.

Figure 3.5: Quote Impulse Response Functions to Exhausts and Non-Exhaust, Sub-sample 2

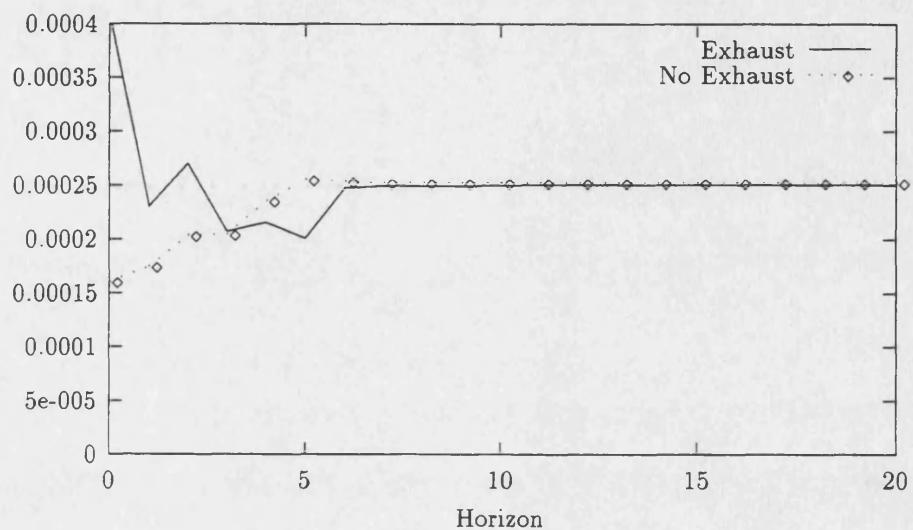


Table 3.10: Revision Equation Parameters from the final VAR Specification in Subsample 2

Var.	Value	T-stat	Var.	Value	T-stat	Var.	Value	T-stat
-	-	-	x_{0t}	2×10^{-6}	0.1	x_{et}	2.5×10^{-4}	7.82
r_{t-1}	-0.37	-5.27	x_{0t-1}	1.3×10^{-4}	3.08	x_{et-1}	-1×10^{-4}	-1.86
r_{t-2}	-0.14	-2.13	x_{0t-2}	5.3×10^{-5}	1.38	x_{et-2}	-2.7×10^{-5}	-0.56
r_{t-3}	-0.08	-1.33	x_{0t-3}	3.5×10^{-5}	0.84	x_{et-3}	-5.6×10^{-5}	-1.10
r_{t-4}	-0.14	-2.10	x_{0t-4}	5.7×10^{-5}	1.62	x_{et-4}	-2.2×10^{-5}	-0.44
r_{t-5}	-0.15	-3.33	x_{0t-5}	7.4×10^{-5}	1.81	x_{et-5}	-5×10^{-5}	-1.06
R^2	0.24		RSS	8.5×10^{-5}		$D.W.$	2.01	

Note: estimation of revision equation parameters in the VAR utilising trade direction and exhaust indicators in the vector of trade attributes. All t -statistics are heteroskedasticity robust. $D.W.$ is the Durbin-Watson statistic and RSS is the regression residual sum of squares.

Table 3.11: Subsample 2: Group Significance tests of sets of Variables in all Equations

Equation	r_{t-i}	x_{0t-i}	x_{et-i}
Return	39.4	17.5	66.3
Deal Indicator	7.8	12.4	1.25
Deal Exhaust	11.4	1.0	1.3

Note: Each cell is a χ^2 -squared test statistic of the hypothesis that the coefficients on the group of variables indicated in the column heading are zero in the equation indicated in the row heading.

Table 3.12: Results from the Permanent/Transitory Decompositions

	σ_w^2	σ_{wz}^2/σ_w^2	σ_w/σ_r
Basic VAR	3.15×10^{-8}	0.32	0.62
+ Exhausts	3.16×10^{-8}	0.32	0.62
Subsample 1	1.94×10^{-8}	0.51	0.66
Subsample 2	4.54×10^{-8}	0.22	0.60

Note: Column 1 contains the estimated variance of the permanent component of the revision process. Column two relates the proportion of the permanent component variance attributable to trades. The final column gives the ratio of the standard deviation of the permanent component to the standard deviation of returns.

FIGURES

Figure 3.1: D2000-2 Quote Midpoint Evolution




Figure 3.2: Quote Impulse Response Function: Basic VAR

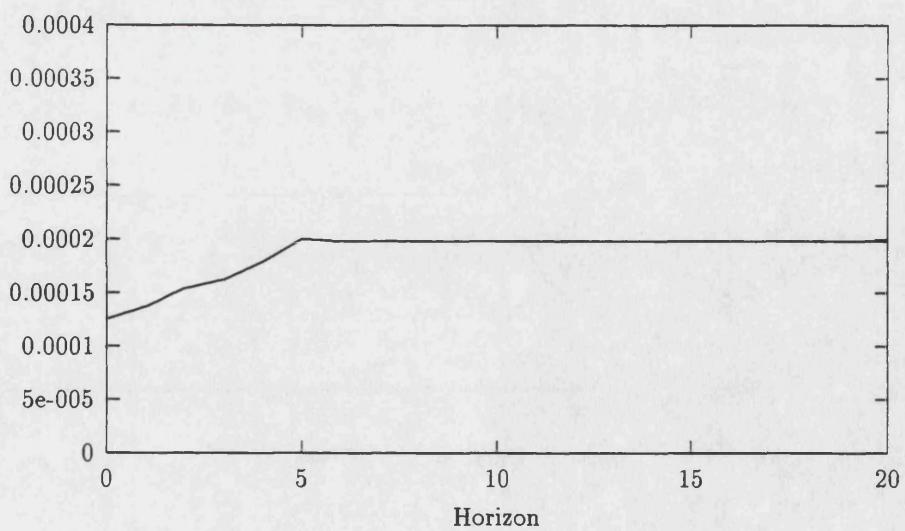


Figure 3.3: Quote Impulse Response Functions to Exhausts and Non-Exhausts





Figure 3.4: Quote Impulse Response Functions to Exhausts and Non-Exhaust, Sub-sample 1



Chapter 4


Announcement Effects and Seasonality in Intra-Day Foreign Exchange Rate Volatility

4.1 Introduction

In recent times it has become a well established fact that intra-day volatility in financial markets is subject to pronounced deterministic seasonality. Volatility effects around market closures, over weekends and within lunch hours have all been shown to be subject their own specific patterns. Works in this area include French and Roll (1986) for the NYSE, Barclay, Litzenberger, and Warner (1990) for the Tokyo Stock Exchange and the series of papers emanating from Olsen and Associates (Zurich) which concentrate on the FX market.

Perhaps the most striking examination of seasonality comes from Andersen and Bollerslev (1997b). They examine two series of intra-daily financial market returns, the first from the foreign exchange market and the second from the Standard and

Figure 3.5: Quote Impulse Response Functions to Exhausts and Non-Exhaust, Sub-sample 2



Poors 500. As the figures at the end of their paper indicate, both markets display pronounced seasonality in volatility; the familiar U-shaped pattern is apparent for the Standard and Poor data, whilst a much less regular seasonal is estimated for the FX data. The authors go on to pre-filter the data for deterministic seasonality using a Flexible Fourier Form, before estimating GARCH specifications for both series over varying observation frequencies. Results demonstrate that removing the seasonal component gives estimated GARCH parameters which conform much more closely with the predictions of theory as data are aggregated over time.

Over the same period of time, a second strand of literature has begun to examine the effects of scheduled news announcements on financial market volatility. The work of Goodhart, Hall, Henry, and Pesaran (1993) concentrates on the FX market, taking two specific announcements and investigating their effects on both the level and variance of exchange rates. Ederington and Lee (1993) and Ederington and Lee (1995) examine the impact of regular, scheduled US announcements on the volatility of T-bond, Eurodollar and Deutsche Mark futures prices. Both conclude that there is a significant impact from these announcements, both at the announcement instant and for an extended length of time afterwards.

The focus of this paper is to combine the above. We too seek to examine and estimate intra-day seasonality in volatility, but do so simultaneously with the estimation of announcement effects for certain US macroeconomic news. Our empirical model is an extension of the Stochastic Volatility models presented in, for example, Harvey, Ruiz, and Shephard (1994) and Taylor (1994). The model incorporates an unobserved AR(1) component, (included to pick up volatility clustering effects,) deterministic seasonal effects and deterministic news dummies. The data employed are quotations for the DEM/USD exchange rate, recorded continuously over the period from October 1992 to September 1993.¹

The main objective of the work is to provide an accurate assessment of the im-

¹The data were provided by Olsen and Associates (Zurich), to whom the author is most grateful.

Chapter 4

Announcement Effects and Seasonality in Intra-Day Foreign Exchange Rate Volatility

4.1 Introduction

In recent times it has become a well established fact that intra-day volatility in financial markets is subject to pronounced deterministic seasonality. Volatility effects around market closures, over weekends and within lunch hours have all been shown to be subject their own specific patterns. Works in this area include French and Roll (1986) for the NYSE, Barclay, Litzenberger, and Warner (1990) for the Tokyo Stock Exchange and the series of papers emanating from Olsen and Associates (Zurich) which concentrate on the FX market.

Perhaps the most striking examination of seasonality comes from Andersen and Bollerslev (1997b). They examine two series of intra-daily financial market returns, the first from the foreign exchange market and the second from the Standard and

Poors 500. As the figures at the end of their paper indicate, both markets display pronounced seasonality in volatility; the familiar U-shaped pattern is apparent for the Standard and Poor data, whilst a much less regular seasonal is estimated for the FX data. The authors go on to pre-filter the data for deterministic seasonality using a Flexible Fourier Form, before estimating GARCH specifications for both series over varying observation frequencies. Results demonstrate that removing the seasonal component gives estimated GARCH parameters which conform much more closely with the predictions of theory as data are aggregated over time.

Over the same period of time, a second strand of literature has begun to examine the effects of scheduled news announcements on financial market volatility. The work of Goodhart, Hall, Henry, and Pesaran (1993) concentrates on the FX market, taking two specific announcements and investigating their effects on both the level and variance of exchange rates. Ederington and Lee (1993) and Ederington and Lee (1995) examine the impact of regular, scheduled US announcements on the volatility of T-bond, Eurodollar and Deutsche Mark futures prices. Both conclude that there is a significant impact from these announcements, both at the announcement instant and for an extended length of time afterwards.

The focus of this paper is to combine the above. We too seek to examine and estimate intra-day seasonality in volatility, but do so simultaneously with the estimation of announcement effects for certain US macroeconomic news. Our empirical model is an extension of the Stochastic Volatility models presented in, for example, Harvey, Ruiz, and Shephard (1994) and Taylor (1994). The model incorporates an unobserved AR(1) component, (included to pick up volatility clustering effects,) deterministic seasonal effects and deterministic news dummies. The data employed are quotations for the DEM/USD exchange rate, recorded continuously over the period from October 1992 to September 1993.¹

The main objective of the work is to provide an accurate assessment of the im-

¹The data were provided by Olsen and Associates (Zurich), to whom the author is most grateful.

portance of each of these three volatility components. Firstly, we aim to confirm the conjecture contained in Ederington and Lee (1993), that the response of spot FX rates to news should be similar to the behaviour they demonstrate for the DM futures contract. However, as the work of Andersen and Bollerslev (1997b) shows, modelling the seasonal component is essential for an accurate analysis of intra-day volatility. The covariation of the seasonal with short-run volatility elements, such as news effects, will imply that ignoring the seasonal will bias estimated news effects. Therefore we estimate a model containing both seasonality and announcement effects. This further allows us to confirm the findings of Anderson and Bollerslev, also providing a simple alternative method of intra-daily seasonal estimation. We compare the estimated parameters of various models combining various different volatility components, demonstrating how the omission of one component or another may lead to mis-estimation of those components which remain in the specification.

In a final set of empirical estimations we examine the impact which individual announcements have on volatility. Using both announcement specific dummies and the median forecast errors associated with particular releases we construct a ranking of the importance of the U.S. macroeconomic data based on their estimated effect on quote variation.

The remainder of the paper is set out as follows. Section 4.2 provides background information and a review of existing studies on both seasonality and announcement effects. Our empirical model is set out in Section 4.3 and the results are presented and discussed in Section 4.4. Section 4.5 concludes.

4.2 Background: News and Seasonality

In this section we review the established facts regarding the components of, and changes in, intra-daily volatility, focussing on market activity and scheduled announcements. With regard to the impact of news, we seek to define clearly the

relationship of the announcement data to news and go on to present a discussion of the possible impacts before, at and after announcement

4.2.1 News and Announcements

In this sub-section we describe the composition of our announcement data and our hypotheses regarding the behaviour of volatility around announcement times. Before doing so, we present a summary of some of the previous work regarding news and volatility, and go on to clarify the relationship between announcements and news.

One of the pioneering studies on the impact of news on FX volatility was performed by Goodhart, Hall, Henry, and Pesaran (1993). In this work the authors employ a tick-by-tick data set of USD/GBP quotations, spanning a calendar time interval of eight weeks, and examine how two specific events affect volatility.² Employing a GARCH-M framework, they find that a system without news effects indicates that the level of the exchange rate has a unit root and also that the conditional variance is very nearly integrated. This latter conclusion alters dramatically when news effects are incorporated; the persistence of the conditional variance equation drops dramatically, and large and highly significant announcement effects are indicated. The authors model the announcement effects via dummy variables in the conditional variance equation, an approach we broadly follow in our empirical specifications. A further point to note regarding this study is that their usage of tick-by-tick data renders their results incomparable to those we present in the following sections.

Another work which examines this issue is Ederington and Lee (1993). They, however, do not examine the spot FX market, rather the reaction of the prices of three nearby futures contracts, the DEM/USD exchange rate, T-bond and Eurodollar. Using a news data set of scheduled US macroeconomic (and survey) releases, they investigate how both volatility and price-adjustment respond to ‘news’. As our fo-

²These events are the unexpectedly good US Trade figures announced on 17/5/89 and a 1% rise in UK base rates revealed on 24/5/89.

cus is on the spot FX market, we discuss only their results on the volatility of the nearby Deutsche Mark contract. The authors construct a set of five-minute transaction price returns from their tick-by-tick data and take the standard deviation of returns for each interval across all the days in their sample. This clearly demonstrates that the interval from 8:30 to 8:35 EST, the interval which is immediately after most of the announcements, is by far the most volatile of the day. The authors further show that this spike in volatility is only apparent for days on which announcements occur. As regards the persistence of this abnormal volatility, they show that the immediate impact of announcements is to increase volatility by five times, this drops to twice normal volatility over the next ten minutes and finally decays over the next few hours. Significantly higher volatility is felt up to 40 minutes after the initial impact. The authors also analyse which announcements are most influential for the DEM future. Results suggest that those which matter are the Employment report, Mercantile Trade, Retail Sales, GNP, PPI and Durable Goods orders, in declining order of significance.

Along the same lines, Harvey and Huang (1991) analyse the volatility patterns of currency futures traded on IMM and LIFFE. They show that the first hour of Thursday and Friday mornings are exceptionally volatile, a consequence which they assert is due to the release of US macroeconomic data of the type Ederington and Lee study, rather than any private information concentration at the market open.

A more recent study, and one which employs the same data period as this, is De-Gennaro and Shrieves (1995). This work utilises hourly quotation returns on the JPY/USD exchange rate, investigating how three types of news impact the market. These categories are scheduled macroeconomic announcements, unscheduled policy news and unscheduled interest rate reports. Results for the scheduled macroeconomic reports suggest that volatility is significantly higher in the hour prior to release, is insignificantly greater in the hour of release, before rising once more and then decaying gently over time.

A relevant issue, for all the above studies, is the relationship between announcements and news. Until now we have used these terms interchangeably, although their interpretations are quite different. All of the 'news' data used in the following empirical analysis, and mostly used in those papers above, are US macroeconomic and survey statistics whose release date and time are known in advance with certainty. This allows market participants to form expectations over their content. Clearly then, there is not an identity between announcements and news, announcements only being newsworthy to the extent that they are unexpected. Further, successive releases of the same data item are likely to differ in the amount of information carried, some having been largely predictable whilst others contain entirely unexpected results. Best practice in this situation is to create a news measure from the announcements by employing the consensus expectation of the market. This approach has been employed in the investigation of news effects on the level of exchange rates by, for example, Hakkio and Pearce (1985). A further consideration is touched on in Ederington and Lee (1993). It is most likely that certain announcements are more important for certain assets e.g. the Trade figures being vitally important for FX rates but, apparently, for little else. Ederington and Lee's demonstration of the influential items shows that if one were to construct the type of news metric mentioned above, not only would the size of the unexpected element be important, but also the type of announcement itself. We address both of these points in our final set of estimations. There we examine the effect of individual announcement series on volatility, utilising both a dummy variable specification and a more precise 'news' measure derived from the difference between actual announced figures and the consensus expectations of the market. Finally, changing market sentiment will be important here also. Over time, the market's belief as to which series are important may alter. At a certain point markets may believe unemployment figures are the key indicator of economic performance, although a year further down the road their focus may have shifted to the Producer Price Index. Hence it must be recognised that the impact of an individual announcement series is likely to be time-varying, as market sentiments shift.

Nevertheless, all the papers mentioned above show that a prominent role is played by announcement effects in short run volatility determination. Here we hope to demonstrate the same kind of impact which Ederington and Lee show, but for the spot FX rate. We depart from previous work on the spot market by employing a very fine, calendar time sampling frequency and explicitly modelling the seasonal component of volatility which is discussed in the sub-section below. A relative disadvantage of focussing on the spot market is that transaction prices, over a time period long enough to examine the issues addressed, are unavailable. In what follows we employ quotation returns as a proxy for transaction returns.³

Our hypotheses about the effects of 'news' on volatility in the pre- and post- announcement periods are as follows. In the periods immediately preceding announcements there are two possible effects. Firstly there is the possibility that volatility is abnormally high. This in turn could have at least two sources. The first, a channel stressed by DeGennaro and Shrieves, is that the details of the announcement are uncovered by some market participants. This creates an informational asymmetry between agents, the informed proceeding by trading on their information and gradually disseminating it to the rest of the market. This will engender high, pre-release volatility. The second possibility is based more on an inventory control idea. Dealers, knowing that an information event will occur at a precise point in the future may desire to trade out of exposed positions towards an equilibrium inventory. This generates a spate of inventory rebalancing trades just prior to the announcement, again causing higher than normal volatility.

It is also possible that in the periods immediately before announcements, volatility will be abnormally low. Again this is based on the fact that the dealers know a news release is occurring shortly and effectively cease all activities whilst they wait to see what the information content of the release is. Hence, pre-announcement volatility

³However, the comparison of transaction prices and FXFX quotations contained in Goodhart, Ito, and Payne (1996) demonstrates that the two series are virtually identical such that the use of quotations should not be a problem.

could be lower or higher than normal. Which effect predominates is clearly an empirical issue. Note that the interval containing the actual announcement should be treated as a pre-announcement period. This is due to the fact that all the news data examined are released on the hour, the quarter-hour or the half-hour, hence each announcement occurs precisely at the end of an interval.

In the post-announcement period we would expect to observe a surge in volatility as the information contained in the release is incorporated into quotations and dealers trade towards their new desired positions. What is interesting in the post-announcement period is the length of time over which this process takes place. How quickly do markets react to information? Again this is an empirical issue. Ederington and Lee (1993) stress that persistent increased volatility after announcement can come from two sources: firstly the price formation process can be inherently slow, quotations taking time to reach their equilibrium level, or secondly the information contained in announcements is only disseminated slowly, such that the market is reacting to a flow of information which is instantaneously incorporated into quotes. Which of these effects dominates will dictate whether 'news' opens the possibility of profitable trading opportunities, although from a volatility perspective they are indistinguishable.

High post-release volatility could also be propagated through an inventory control channel. After an announcement, traders will update their beliefs over the fundamental value of the asset, engendering a change in desired inventory holdings. If the transactions which restore each trader's inventory to equilibrium are not worked through immediately then one will again notice persistently high, post-release volatility. One way to distinguish between the effects of information revelation in the FX market and impacts from inventory trading to give persistently high volatility after a release is to examine simultaneously the impact of 'news' both on the level and volatility of exchange rates. In this chapter, this question is left unaddressed with the analysis of price changes in response to 'news' conducted in Chapter 6.

The above discussion defines our empirical hypotheses. In the pre-announcement period, the effect of news on volatility is, *a priori*, indeterminant, whereas in the post-release period one should expect an increase in volatility as long as there is an information content to the announcement. As indicated above, the focus in the post-release period is the persistence of the volatility surge.

Before moving on to an examination of seasonality in volatility we present the announcement data actually employed in estimations. All are monthly, US, macroeconomic or survey statistics, the list being: the Unemployment rate, Merchandise Trade Deficit, Producer Price Index, Consumer Price Index, Retail Sales, Consumer Confidence Index, Leading Indicators, Durable Goods Orders, Industrial Production and Capacity Utilisation⁴ and finally the NAPM survey. All but four of these announcements are made at 8.30 EST. The Industrial Production and Capacity Utilisation (IP/CU) figures are announced together at 9.15 EST whilst the NAPM survey and Consumer Confidence figures are released at 10.00 EST.⁵ Market participants know both the time of announcement and the date on which the release will occur in advance.

4.2.2 Seasonality

The major movements of intra-daily return volatility can be attributed to the passage of market activity around the globe and it is this which underlies the seasonal pattern we observe. One can regard the global FX market as being composed of three major regional centres, the Far East, Europe and North America, all of which have their own activity pattern. To begin to interpret the seasonal pattern one needs a feel for the openings and closures of the three components. Roughly one can say that the Far East is open from 21:00 GMT to 7:00 GMT, Europe trades between 6:00 GMT and 16:00 GMT, whilst trading occurs in North America from 14:00 GMT to 21:00

⁴These two announcements are paired here as they are always announced simultaneously.

⁵See Appendix 1 of Chapter 6 for more details on the timing of these releases.

GMT. The accumulation of activity in open markets gives the level of the seasonal at any point of the day, hence interpretation of the seasonal pattern is performed by appealing to the conditions in these open markets.

First we give a brief presentation of the exchange rate data. As indicated in Section 4.1 the data employed are observations on the DEM/USD exchange rate over the period 1/10/92 to 30/9/93. The data are essentially a filtered transcription of the activity on Reuters FXFX pages; the original form of the data is a set of tick-by-tick quotation inputs from the banks which participate on the system. We convert these data into a calendar time-series by imposing a five-minute observation grid, an observation being taken at the end of each interval.⁶ A modification to this structure is the omission of weekends (defined as 21:00 GMT on Friday to 21:00 GMT on Sunday,) as these are periods of little or no activity. The five minute observation window implies that a day spans 288 observations and the omission of weekends leaves us with 261 trading days. This yields a time-series of 75168 observations.⁷ The basic statistical features of the return series are given in Table 4.1 and analysed in Section 4.4

We show in Figure 4.1 the pattern of our volatility measure, the logarithm of squared returns, averaged over the 261 trading days in our sample, for each of the 288 five minute intervals of the day.⁸ The seasonal pattern which emerges for this measure is more-or-less identical to that demonstrated for average absolute returns in Andersen and Bollerslev (1997b) and that shown in Dacorogna, Müller, Nagler, Olsen, and Pictet (1993). Further evidence of the daily seasonal structure shows up in the autocorrelation function of our volatility proxy. In Figure 4.2 we show the auto-

⁶This sampling interval was chosen on the following basis. First, as earlier work shows, there are prominent intra-hourly effects from macroeconomic releases, necessitating a short observation window. Second, the computational tractability of the problem decreases very quickly with the sampling interval. Our five minute window was chosen to balance these two effects.

⁷At various points in the series, no quote is entered in a 5 minute interval. These data holes were filled by linear interpolation between the nearest preceding and succeeding quote.

⁸Our choice of $\log(r^2)$ as a volatility proxy is motivated by the empirical model in Section 4.3. The pattern demonstrated in Figure 4.1 is robust to the use of alternative volatility measures such as absolute returns and squared returns.

correlation function for $\log(r^2)$ over a span of 288 lags i.e. over one full trading day. What is notable is the peak in the representation at precisely the daily frequency, demonstrating that the memory of volatility is most closely attuned with what occurred precisely one day ago rather than at any point between now and then.

We also present the average pattern and autocorrelation of our volatility proxy over a week's span (Figures 4.3 and 4.4). The volatility averages over the trading week demonstrate the stability of the daily pattern, with little evolution or alteration of the pattern evident over different days. A point which the weekly autocorrelation function demonstrates is that there seems to be some seasonal structure at the weekly frequency. This can be seen through the local maximum at a lag of precisely 5 trading days (lag 1440) dominating those maxima attributable to lags of two, three and four trading days. However, in what follows we ignore this weekly structure, concentrating solely on the daily pattern.

Using the discussion of market openings and closures presented at the beginning of this subsection we can break down the volatility seasonal as follows.

The first interval of the day corresponds to the five minutes between 0:00 and 0:05 GMT, a time when the Far Eastern market has already been trading for around 3 hours and market activity is high. From this point until approximately 3:20 GMT (interval 40) volatility and activity levels remain high. At this point we come across the most prominent feature of the series, lunchtime in Tokyo. Volatility drops sharply, flooring at near zero levels and only regaining its former value at about 5:00 GMT. The next market to open is Europe, generally beginning to contribute to activity at around 6:30 GMT as activity in the Far Eastern market begins to wane. This gives a small peak in volatility between intervals 80 and 120, before activity and volatility both experience a slight lull during European lunch hours. The most active period of the day is clearly the interval when both the European and North American markets are open (intervals 160-180), after which volatility starts to decline as first the European, and then the US markets, wind down. Finally at around 21:00, the

Pacific market begins to trade again and the daily cycle is repeated after midnight.

Daylight Saving Time also has an effect on the seasonal pattern. In summer months, both North America and Europe lose one hour relative to GMT. This implies that, as Far Eastern local time is unchanged, the seasonal pattern alters in composition (rather than simply shifting laterally relative to the GMT hour of the day.) In estimation this phenomenon is handled by the introduction of two seasonal regimes, one relating to the winter months and the second for summer, parameterised using a simple dummy variable formulation.⁹ A comparison of the average daily $\log(r^2)$ pattern in summer and winter, showing precisely the effect which DST has on the seasonal, is presented graphically in Figure 4.5.

Hence the seasonal pattern which emerges seems fully explicable. What is less obvious, however, is the way in which the omission of this component in estimation might impinge upon examination of other volatility components. As long as these components are not orthogonal it is likely that mis-specifying the model, by omitting the intra-daily seasonal will lead to biased estimation of the parameters in other components.

The works mentioned in the previous subsection dealt with the seasonal as follows. Ederington and Lee (1993) base their results on a comparison of announcement and non-announcement days, implicitly filtering the seasonal by examining the announcement to non-announcement volatility ratio. Further, volatility in the futures markets they treat may be less seasonal than that in the spot FX market. DeGennaro and Shrieves (1995) treat the seasonal explicitly. In certain of their specifications they add opening and closure dummies to account for weekend effects and employ a quotation frequency variable as a proxy for the seasonal pattern. This proxy is, in general, very good. However, one might expect quotation frequency to rise deterministically around ‘news’ announcements also. If this is the case, then their

⁹In the strictest sense there are actually four regimes here, as North America and Europe alter their times about one week apart, but we ignore these short periods, subsuming them into one of the major regimes.

seasonal proxy will absorb part of the news effects which they are attempting to estimate, biasing the 'news' coefficients downwards. Finally, Goodhart, Hall, Henry, and Pesaran (1993) include no seasonal effects whatsoever, although, as they use tick-by-tick data, the effects of this omission are likely to be less severe.

Some other studies which examine the seasonal patterns in the intra-day FX market are as follows. Baillie and Bollerslev (1991) presented one of the pioneering examinations of this phenomenon in their analysis of four spot FX rates. Employing hourly returns on the USD/GBP, DEM/USD, CHF/USD and JPY/USD they estimate seasonal GARCH models, the seasonality captured with a dummy variable specification. Results demonstrate very similar seasonal patterns across currencies and, clearly, the dummy specification for seasonality works well in this context. For the data set we employ, however, a dummy variable formulation would require the estimation of 287 parameters, an obvious drawback.

Andersen and Bollerslev (1997b), as already mentioned, use a flexible Fourier form to pre-filter the data for seasonality. The implication of this technique is that seasonality is treated as a nuisance component in the data which simply obscures the behaviour of underlying volatility. If, however, one wishes to examine announcement effects, which are also periodic, then pre-filtering the seasonal pattern will necessarily pre-filter a portion of the 'news' impact. Also neither this study nor Baillie and Bollerslev (1991) address the effect of DST in their estimations, implying that their estimated seasonal pattern is 'blurred' i.e. a combination of the patterns apparent in two distinct regimes.

Lastly, a fairly large literature has arisen which concentrates on the concept of time-deformation in order to explain and estimate the seasonality patterns apparent. The motivation behind these models can be found in Stock (1988) and it is, essentially, that markets work on a time scale which differs from simple clock time. Relevant variables evolve in market time and their behaviour in clock time is derived by applying a non-linear transformation between the two time scales. An application of

this approach can be found in Ghysels, Gouriéroux, and Jasiak (1995) who estimate a time-deformed SV model. In related work, Dacorogna, Müller, Nagler, Olsen, and Pictet (1993) model intra-daily seasonality by introducing a market activity variable, the integral of which defines a market time-scale. Examination of a regularly spaced price change series in this market time-scale demonstrates the removal of the seasonal in volatility.

4.3 The Empirical Methodology

4.3.1 The Basic SV Model

The starting point for our empirical model can be found in Harvey, Ruiz, and Shephard (1994). The basis of their model is a representation of asset returns as shown in equations (4.1) and (4.2) below,

$$r_t = \sigma \epsilon_t e^{h_t/2}, \quad \epsilon_t \sim N(0, 1) \quad (4.1)$$

$$h_t = \phi h_{t-1} + \nu_t, \quad \nu_t \sim N(0, \sigma_\nu^2) \quad (4.2)$$

Here, r_t is the return on the asset in question, σ is a volatility scale parameter, ϵ_t is a white noise term and h_t is a time varying volatility component. As equation (4.2) demonstrates, h_t is assumed to follow a first order autoregressive process, a specification chosen to parallel the volatility clustering motivation behind the GARCH literature. The disturbance term in this equation is also assumed white noise, with given variance, σ_ν^2 , and independent of ϵ_t .

The return process is clearly stationary if the process generating h_t is stationary, a situation which occurs when $|\phi| < 1$. Using this approach, h_t is treated as an

unobserved component which is retrieved via the Kalman Filter. In order to apply the filter to the above specification, the measurement equation, equation (4.1), must be linearized in the state (h_t). This is achieved by transforming the return equation into that shown below.

$$\log(r_t^2) = \log(\sigma^2) + h_t + \log(\epsilon_t^2) \quad (4.3)$$

Given the standard normal distribution posited for ϵ_t , the transformed error term is known to have a mean of -1.27 (and a variance of $\pi^2/2$), such that creating a term, $\xi_t = \log(\epsilon_t^2) + 1.27$, gives the following specification,

$$\log(r_t^2) = -1.27 + \log(\sigma^2) + h_t + \xi_t \quad (4.4)$$

in which the new disturbance term has zero mean. Combining equations (4.2) and (4.4) gives a time-invariant state-space form whose parameters can be estimated via the Kalman Filter. If estimation demonstrates that ϕ is approximately unity, a situation which parallels the IGARCH specification, then a unit root can be imposed upon the transition equation. In order to pre-test for the presence of a unit root in $\log(r_t^2)$ one can employ an ADF test, although, as mentioned in Harvey, Ruiz, and Shephard (1994), the power of this test in this situation may be questionable due to the possible near non-invertibility of the volatility representation.

4.3.2 Modifications

We generalise the above model as follows. The first modification is the addition of a set of deterministic trigonometric components used to model the seasonality inherent in volatility. These terms are constructed as shown below (see Harvey (1989)),

$$\psi_t = \sum_{j=1}^{s/2} (\gamma_j \cos \lambda_j t + \gamma_j^* \sin \lambda_j t), \quad \lambda_j = 2\pi j/s \quad (4.5)$$

with $s = 288$, as described in Section 4.2. If all of the seasonal frequencies were to be included this would lead to the estimation of 287 parameters for the seasonal, the same number of parameters that one would expend on a dummy variable seasonal specification. However, as the seasonal is likely to change fairly smoothly, subject to a qualification below, some of the high frequency components may be omitted without sacrificing a great deal of accuracy in estimation.

In order to gauge which of the components are most important, in Section 4.4 we examine the seasonal periodogram for our volatility proxy. The periodogram simply plots the amplitude of each Fourier component against its frequency, with each amplitude constructed as shown below,

$$P_j = \frac{2}{T} \left[\left(\sum_{t=1}^T y_t \cos \lambda_j t \right)^2 + \left(\sum_{t=1}^T y_t \sin \lambda_j t \right)^2 \right] \quad (4.6)$$

where y_t is the series of interest. Examination of the periodogram then indicates which of the frequencies are dominant, allowing one to trim the number of harmonics used in estimation.

An addition to the seasonal specification is made to cope with the sharp drop in volatility around lunch in the Far East. The smooth seasonal generated from the dominant Fourier terms is unlikely to cope well with this discontinuity, in all likelihood the drop will be largely underestimated. We therefore employ a set of dummy variables to cope with this phenomenon. Also, as mentioned earlier, Daylight Saving Time in North America and Europe will engender an alteration in the form of the seasonal. Hence we estimate two seasonal regimes (indexed by m) over which the parameters of the trigonometric terms are allowed to alter. Note that, as there is no Daylight Saving Time in the Far East, we can leave the lunch dummy unchanged

across the two regimes. This gives the final form for the seasonal which is presented below,

$$\psi_{mt} = \sum_{j=1}^{s/2} (\gamma_{mj} \cos \lambda_j t + \gamma_{mj}^* \sin \lambda_j t) + \sum_{i=0}^k \mu_i l_{t-i}, \quad \lambda_j = 2\pi j/s, \quad m = 1, 2 \quad (4.7)$$

where k is the number of intervals which lunch encompasses and l_t is a variable defined to be unity only in the first interval of the Japanese lunch hour. By adding lagged values of l_t to the specification we allow the dummy coefficients in each interval of lunch to differ. In the most general case the coefficients on the dummies are unconstrained, although a fixed lunchtime effect can be imposed by setting $\mu_i = \mu$ for all i .

If desired, the seasonal can also be made stochastic, allowing the parameters γ_{mj} and γ_{mj}^* to evolve over time (see Harvey (1989),) but this route is not followed in the current work as it seems that the pattern observed is quite stable. An alternative to the trigonometric formulation presented above would be to model the seasonal using a periodic spline formulation, à la Harvey and Koopman (1993) and Harvey, Koopman, and Riani (1995), a formulation which can also be deterministic or stochastic.

The second extension of the basic model is centered around the addition of a further dummy variable, d_t . This indicator is unity for only those five minute intervals which contain one of the news announcements in our data set. In practice, as all the included announcements occur either on the hour, the quarter-hour or the half-hour, this implies that the instant of release is precisely at the end of a period. We allow for the possibility of news effects prior to, at the time of, and post-announcement by including leads of, the contemporaneous value of and lags of the news dummy in our specification. Appropriate trimming of the lead and lag specifications gives us the approximate impact intervals before and after announcement. Hence the final

specification for volatility is as below,

$$\log(r_t^2) = -1.27 + \log(\sigma^2) + h_t + \psi_{mt} + \sum_{i=-p}^q a_i d_{t-i} + \xi_t, \quad p, q \geq 0, \quad m = 1, 2 \quad (4.8)$$

where ξ_t is as specified as above, ψ_{mt} is as in equation (4.7) and d_{t-i} is the indicator of an announcement at lag i .¹⁰

Combining equations (4.2) and (4.8) gives our final specification. As noted earlier, it involves elements which allow for volatility clustering, a smooth, deterministic seasonal pattern and an extended impact of announcements on volatility.

4.3.3 Individual Announcements

A last set of empirical exercises examines whether the volatility responses differ across announcement types. This is done by splitting our announcement data into 10 distinct variables, one for each announcement type. We then estimate a restricted version of the final model for each announcement type. The restrictions imposed in estimation are as follows. First, we impose a geometric decay on the post-release volatility response. In terms of the parameters of equation (4.8) we allow a_1 and a_2 to vary freely but restrict all subsequent impacts to decay at rate ρ , i.e. $a_i = a_{i-1}(1-\rho)$, for $i \geq 3$.

Secondly, these estimations are run with the seasonal pre-filtered. We subtract the seasonal estimated from the final specification (i.e. equations (4.8) and (4.2),) from our $\log r^2$ series in order to form the dependent variable in the measurement equation. This procedure should minimize any systematic bias in the estimation of the impacts for each announcement as the seasonal employed was originally estimated

¹⁰The effect of news on volatility before announcement is represented by i taking the requisite negative value.

in the presence of a full complement of 'news' dummies.

We run these estimations using both a simple dummy specification for announcements (i.e. a variable taking the value unity at the point when a specific announcement occurs only) and using the forecast errors associated with announcements. The forecast errors are created as the difference between the actual announced figure and a median survey expectation.¹¹ By using both of the above measures, we can get the general impact of a certain announcement from the first estimation and a more precise economic impact, e.g. the impact of an unexpected 1% rise in unemployment on FX volatility, from the latter.

Hence, from these results, we can examine which announcements really move the market. Are all pieces of macroeconomic data equally important for the determination of the DEM/USD rate, or do some announcements dominate? Note, however, that for each announcement we only have 12 observations within our one year span of data. Moreover, given the possibility of markets altering their opinion on the most important indicator of economic performance, these results may not generalise to other time periods.

4.3.4 Estimation

As previously mentioned, the basis for estimation of the class of models examined above is the Kalman Filter. The final specification for $\log(r_t^2)$ in equation (4.8) serves as the measurement equation, whilst equation (4.2) provides the transition equation. The unobserved component in our model is h_t , a measure which one might refer to as 'underlying' market volatility.

The prediction and updating equations of the Kalman Filter permit one to recursively build up the conditional expectations of the state (h_t). These can then be

¹¹The survey medians employed were obtained from MMS International, who survey around 40 major market watchers for their forecasts on the Friday previous to the release of each piece of information. See Appendix 2 of Chapter 6 for a treatment of the rationality of these expectations.

used to describe the conditional distribution of the volatility proxy at each time-series observation. Combining these conditional distributions yields the prediction error decomposition formulation of the log-likelihood function which can then be maximised using numerical methods. This procedure is by no means the only way in which to estimate SV models. Scott (1987), Chesney and Scott (1989) and Melino and Turnbull (1990) use GMM type estimations whilst others advocate Bayesian estimation methods, e.g. Jacquier, Polson, and Rossi (1994), pointing to the increased efficiency these methods bring. In our case however, with such a large data set, efficiency losses are likely to be minimised and, further, these alternative estimation techniques become computationally intractable.

The problem in using the Kalman Filter in this case is the non-normality of ξ_t . This implies that the filter delivers only MMSLE's, such that the approach must be treated as a QMLE estimation, i.e. working as if ξ_t was distributed normally, with mean zero and variance $\pi^2/2$. The filter is initialised using the unconditional distribution of h_t , except in the case where h_t is constrained to be a random walk, when the initial observation is used. From then on, the quasi-likelihood function is built observation by observation. The quasi-likelihood is maximised over the parameter space using the optimisation algorithm of Broyden, Fletcher, Goldfarb and Shanno. Standard errors for the hyperparameters, ϕ and σ_ν^2 , are calculated using the results presented in Ruiz (1994).

Once the parameters have been estimated, the smoothing algorithm of Koopman (1993) is employed in order to retrieve the within-sample values of the state, h_t . This allows one to examine, after accounting for the seasonal and announcement impacts, the time-series volatility of the FX market, from which a clarified picture of the volatility profile, and perhaps its determinants, can be drawn.

4.4 Results

Before presenting the empirical estimations of the models outlined in the previous section we briefly introduce the variables used in estimation, examining their basic statistical behaviour. Summary statistics for these variables are presented in Table 4.1.

Examining first the behaviour of the return series, two facts are immediately apparent. First of all the series displays pronounced excess kurtosis, confirming the findings of many previous studies which demonstrate that the distribution of financial market returns tend to have very thick tails. A second point is that the Box-Ljung statistic demonstrates that there is some serial dependence in the returns series. This motivates our examination of the second series, a set of residual returns constructed after the removal of a first order moving average from the original return series. One can note that the dependence in this series is substantially lower than that of the raw returns (although still statistically significant) and again excess kurtosis is apparent. The real series of interest, however, is our volatility proxy. This is constructed as the logarithm of squared residual returns. It is clearly apparent from a comparison of BL statistics that it is in volatility that the real temporal dependence lies. The BL statistic for $\log(r^2)$ is over 250 times that of raw returns and it is this phenomenon which we seek to address and explain in the estimations below.

4.4.1 Non-seasonal SV models

Our first empirical application consists of fitting the basic SV model, demonstrated in equations (4.2) and (4.4).¹² As described in Section 4.3.4 the model is estimated by Quasi-ML via the Kalman Filter, the procedure containing the restrictions that ϕ

¹²Before the estimation of all specifications, the scale factor in the measurement equation, $\log(\sigma^2)$, is removed.

be between zero and one and σ_ν^2 be positive. Results of this specification are shown in Table 4.2.

The estimate of ϕ conforms with the results of many other volatility studies (mainly of GARCH form,) which demonstrate that, at this fine, calendar time sampling frequency, the variance process is approximately integrated. In particular, our estimate of ϕ is 0.96, a value which is almost on the boundary of the parameter space. As previously indicated, to test for the presence of a unit root in volatility we employ an ADF test for $\log(r^2)$, (the ADF specification embodying 12 lags of the difference of $\log(r_t^2)$ as dictated by the Akaike and Bayes information criteria.) The outcome of the test demonstrates that one can strongly reject the hypothesised unit root.

Despite this test result and for comparison, we go on to impose the restriction that $\phi = 1$, the results being shown in Panel 2 of Table 4.2. The variance of h_t can be seen to drop quite sharply, as does the maximised Quasi-Log Likelihood, and due to this latter fact, in all further specifications, we drop the unit root imposition, allowing ϕ to take any value between zero and one.

The behaviour of the derived standard deviation in the unrestricted model is shown in Figure 4.6, alongside the behaviour of absolute residual returns.¹³ From the figure it is apparent that h_t tracks the underlying volatility of returns very well, but this observation masks a weakness of the specification. As outlined in previous sections of this chapter, probably the most distinctive feature of the intra-daily volatility process is its seasonality. In this basic specification there is no explicit treatment of that seasonal. Hence, when we examine the average behaviour of h_t over the intervals of one day, it becomes clear that all seasonal variation is picked up by the state variable. The comparison of the rescaled behaviour of h_t and that of $\log(r_t^2)$ is demonstrated in Figure 4.7. This has mixed effects: on the positive side it implies that the bias to parameters in the ‘news’ specification is likely to be reduced, but negatively, the state variable, h_t , now indicates little about the volatility clustering

¹³This is calculated as $\sigma e^{h_t/2}$ as in equation (4.1).

behaviour which the data may exhibit.

Despite this, and to provide a comparison later in the paper, we next estimate a non-seasonal SV specification which incorporates 'news' effects. The system consists of equation (4.2) and a version of equation (4.8), in which all seasonal parameters are set to zero.

The results of the SV/News specification are given in Table 4.3. First note that there is little change in the basic SV parameters, ϕ and σ^2 . Both are at almost precisely their level from the original model, ϕ is highly significant and still very close to unity. Of more interest are the 'news' coefficients. In the pre-announcement period the coefficients show that a significantly reduced measure of volatility in the period 15-30 minutes before announcement is followed by a positive significant impact in the 15 minutes immediately prior to release. The post-announcement period is characterised by an immediate large spike in volatility which dies away, non-monotonically, over at least the next hour. All of the estimated coefficients are highly significant. A graphical exposition of these announcement effects can be found in Figure 4.8.

Hence our preliminary results suggest that in the minutes immediately before announcement there may be either information leakage or closing-out trading, both of which could generate the higher than normal volatility. However, in the period from 15 to 30 minutes before announcement, markets are significantly less active, reflecting the possibility that traders slow their activity in anticipation of the upcoming news. On balance, the dominant pre-release impact seems to be greater volatility, both in terms of magnitudes and significance.

The response after release demonstrates the information contained in these data and the importance the FX market places upon it. The fact that volatility is persistently high for the hour post-announcement suggests that either the price formation process is slow, or that the information contained in the release data is only gradually extracted. Of course these results must be cautiously accepted, given the potential

bias in estimation which the lack of a seasonal component may bring about.

4.4.2 Seasonal SV Models

As indicated at the end of the previous sub-section there is a danger in taking the results of the non-seasonal specifications at face value. We go on now to incorporate explicit seasonal elements into the specification as described in Section 4.3.2. The primary problem is the choice of which of the set of 144 Fourier terms to incorporate, hence, as outlined previously, we examine the seasonal periodogram which is shown in Figure 4.9.

As expected, the fundamental frequency is dominant, it's amplitude dwarfing that of any of the harmonics. The very high frequency elements add very little variation to the seasonal, only the first fifteen or so components being visible on the chart. We employ the six Fourier terms with the greatest associated amplitudes. These are the first, third, fourth, fifth, ninth and tenth elements. As shown later on in this section the inclusion of only these elements, along with the dummies for lunchtime in Japan, gives a perfectly satisfactory estimated seasonal representation.

We can now progress to the estimation of the seasonal SV models. The first results, which are presented in Table 4.4, are for a model which incorporates the basic SV structure plus the seasonal only. Examining first the estimated parameters from the transition equation, there is little change in the autoregression coefficient; it rises only very slightly and remains highly significant. There is, however, a significant drop in the transition equation error variance. This is due to the removal of seasonal effects from the time-varying volatility component, h_t . All the coefficients of the seasonal representation are significant, (a LR test of their joint significance indicates that the hypothesis that all seasonal coefficients are identically zero can be strongly rejected,) although their numerical interpretation is not straightforward. Hence we reconstruct the daily seasonal they represent, in both regimes, and both are shown in Figure 4.10. In comparison to Figure 4.5, which shows the average value of $\log(r^2)$

over the 288 daily intervals for summer and winter, there seems little difference. Of course the estimated seasonal patterns are smooth relative to the average intra-daily patterns but they pick up the general seasonal shape quite well. Figure 4.5 also shows some interesting features. First, as expected, the summer and winter seasonal coincide for the Far Eastern portion of activity, this being due to the lack of DST in Asia. Next our representations pick up a small spike in activity just after Japanese lunch. This may represent volatility engendered by trading on information which has been revealed over the lunch hour closure. Finally, European and North American trading activity can be seen to shift approximately 12 intervals to the right during winter months, reflecting their time changes relative to GMT.

The adequacy of our seasonal representation is emphasised in the autocorrelation function of the deseasonalised $\log(r^2)$ series.¹⁴ Figure 4.11 plots this function and it is immediately apparent that the decaying, repeated U-shaped structure which can be seen in Figure 4.2 is almost completely removed. This compares favourably with the plot of deseasonalised returns from Andersen and Bollerslev (1997b), which shows that their procedure does not completely remove the seasonal. One feature which does become more apparent though, is the seasonal structure at the weekly period. Also, there seems to be a very slow decline in the correlogram, although masked by the weekly seasonal, an indication of long memory in volatility. An examination of this feature of the data is beyond the scope of the current study but is the focus of Chapter 5.

It seems, therefore, that our trigonometric treatment of the seasonal is quite successful. This conclusion is re-affirmed by the average intra-daily behaviour of the state variable, h_t , derived from this model. A graphical examination of the average intra-daily behaviour of the state shows that the only distinctive feature of average 'underlying' volatility is a small, residual effect from Japanese lunch.

¹⁴Deseasonalised volatility is simply calculated as the values of $\log(r^2)$ after filtering the deterministic trigonometric and dummy variables.

4.4.3 The Combined Model

We can now move on to the final, combined model of seasonality and news. This consists of estimation of equations (4.2) and (4.8), allowing a maximum window for ‘news’ effects of 30 minutes pre- and 1 hour 15 minutes post-announcement. The results from this model are presented in Table 4.5.

As the seasonal coefficients directly tell us little about the seasonal pattern, and their magnitudes and significances are little changed from those in the previous subsection, we treat only the transition equation and announcement dummy parameters explicitly. A cursory examination of the former shows that little has changed here also. Both the autoregressive parameter and the transition equation error variance are very marginally reduced in magnitude from their levels in the previous estimation, although the significance of the autoregressive parameter has risen.

A far more noticeable alteration is apparent in the estimated ‘news’ coefficients. Earlier in the paper we referred to the potential bias to the announcement effects which may occur due to the omission of seasonal effects from the specification. Here, we see that this is precisely true. All news coefficients are now around 0.25 lower than in the previous estimation.¹⁵ The downward change is exactly what we would expect, although the magnitude of the bias is quite small. Again, as mentioned in Section 4.4.1 this is likely to be due to the absorption of the seasonal by h_t in the models estimated without an explicit seasonal component.

Now, in the pre-announcement period our results of a distinctive quietening in the market are re-inforced. In the period 10 to 30 minutes before announcement the coefficients become more negative and significant, whereas there is a drop in both magnitude and significance for the positive impacts in the 10 minute period before release. Hence, the dominant impact is now one of reduced volatility pre-release.

¹⁵This 0.25 decline in the parameters gives an approximate 15% reduction in the impact on derived standard deviation across the board.

In the post-announcement interval, there is again a general reduction in the magnitude of the coefficients of the order of about 0.25. What is still true, however, is that the positive volatility impact of announcements lasts for at least one hour, in line with the results derived in Ederington and Lee (1993). In fact the volatility effect is consistently positive until 1 hour 15 minutes after announcement when the first negative coefficient is encountered. A plot of the new announcement coefficients, alongside those derived from the non-seasonal news specification is shown in Figure 4.12.

Hence our results conform qualitatively with those of Ederington and Lee (1993). The quantitative comparison is, however, not as close. As mentioned in Section 4.2, their estimations demonstrate that the standard deviation of returns rises fivefold in the interval immediately after an announcement, dropping to double the normal standard deviation in the following five minutes. Constructing the derived standard deviation from the corresponding parameters in our estimations gives a different picture. We predict an immediate post-announcement response of less than 3 times normal standard deviation which drops to just under twice the baseline level after ten minutes have elapsed. Hence, the impact of these announcements on the spot FX market is less pronounced than for the currency futures market.

So, the conclusions from the combined specification are as follows. The pre-release period seems to be characterised by overall volatility reduction, an effect which corresponds to dealers scaling down activity in the face of the impending announcement. In the post-announcement period, the conclusion is that there is a prominent, immediate volatility impact which is eroded quite quickly over the following 10 minutes and from then on, dies out quite slowly. Again there is evidence of persistent and important informational effects from these macroeconomic and survey releases. The most prominent effects of announcements are intra-hourly. The study of DeGennaro and Shrieves (1995), employing an hourly observation window, necessarily misses all of this structure. Also, as most announcements are on the half-hour, their use of hourly observations implies that the observation containing a news item will contain

both pre- and post-announcement periods. Disentangling the effects in these two periods would seem to be impossible. In order to examine announcement effects properly, an analysis at a very fine sampling frequency is essential.

4.4.4 Individual Announcements

Finally we present the volatility estimations in which the different announcements are treated individually. Which of our macroeconomic releases has most effect on the DEM/USD spot rate? As previously noted, Ederington and Lee (1993) find that, of the present set of announcements, the Employment report, Trade figures, Retail Sales, Durable Goods and the PPI have the most prominent impact on DEM future's volatility in descending order. This also ties in with Harris (1995a) who comments that the Employment report has become increasingly viewed as the key U.S. indicator by the markets.

Table 4.6 gives our results for the estimations where the news variable employed in each case is a simple announcement-specific dummy. It presents, for each piece of data, the volatility impacts 5 and 10 minutes post-release, the estimated persistence parameter for this increased volatility over the following hour and finally the implied percentage increase in return standard deviation for the five minutes immediately post-release (i.e. $e^{a_1/2}$.)

The Employment Report clearly has the greatest effect on volatility, the coefficient a_1 being almost one half as great again as that for the next most important and implying an instantaneous volatility jump of over 1000%. The report contains two key figures: payroll employment and the unemployment rate and hence this one large response incorporates two smaller impacts. Next come the Trade figures which engender, on average, a rise in return standard deviation of over 500%. Again this is unsurprising given the intuitive impact of exports and imports on the demand for currency. The persistence of volatility is greater than average for both of the above announcements, possibly because both of these monthly documents contain

multiple statistics and are lengthy and time-consuming to digest.

The next most important pair of announcements are the Retail Sales figures and the PPI report. Hence, the four most influential releases we derive from the spot FX data are also represented in Ederington and Lee (1993) in their top five for the DM future, demonstrating the similarity in importance across spot and derivative markets.

The Consumer Confidence figures, Durable Goods Orders and CPI figures then form a medium impact sub-group. These data show very similar immediate impacts, a standard deviation rise of around 300%, but the 10 minute response to the Durables figures is much greater than for the other pair and the persistence stronger. This announcement is the fifth of the group Ederington and Lee (1993) cite.

Finally there is a group of low impact announcements which comprises the Leading Indicators figures, the NAPM survey and the Industrial Production/Capacity Utilisation (IP/CU) results. It is clear that the IP/CU announcement has the smallest volatility impact across all dimensions, with the Leading Indicator figures being fairly important. The most surprising result of estimation is a ten minute impact for the NAPM which is small and negative, coupled with a negative persistence parameter, a result which implies damped oscillations in the NAPM volatility response!

Lastly, in Table 4.7, the results from the same individual specifications but using the absolute announcement forecast errors as our news data are presented. Rather than reporting the percentage standard deviation impact here, we include a different set of figures, containing the estimated immediate impacts, a_1 , multiplied by the mean absolute forecast error for each particular release over our 12 month period.¹⁶ Note that we also have two extra pieces of data which can be analysed here. This is due to the fact that we can split the Employment report into the unemployment rate and the payroll employment figures via their forecast errors, and can similarly separate

¹⁶This rescaling gives us a basis for the direct comparison of the 'news' impacts across announcements measured in differing units.

the Industrial Production and Capacity Utilisation figures.

The impact ranking which emerges broadly corroborates our results from the previously reported dummy variable specifications. The two main components of the Employment report are ranked first and second, although the difference between either of these immediate impacts and the next highest ranked is considerably diminished from the dummy results. This is to be expected as it is the combination of the unemployment rate and payroll figure volatility impacts which gives the total Employment report effect. Again, the next most important effect on volatility comes from the Mercantile Trade figures.

In comparison with the results of Ederington and Lee (1993) we again find that the four highest ranked from this study are represented in their top five, these being the unemployment rate, trade figures, durables and retail sales. The member of their five highest ranked announcements which does not conform with the results of this part of our study is the PPI release. In comparison to the dummy results, the scaled PPI impact has plummeted in rank. This may be due to the dummy results being dominated by a couple of very large price movements which are associated with large forecast errors. In this case, the linear specification in forecast errors deals well with this variation in the data.

The NAPM, Leading Indicator and IP/CU figures are again towards the bottom end of the spectrum. There are, however, some anomalous estimation results associated with these releases. The leading indicator announcement now shows a negative persistence figure, the NAPM has almost zero persistence whilst the IP figures have a persistence parameter which is insignificantly different from unity.

So we can draw the following broad conclusions. The announcements which cause greatest post-release volatility are those associated with the Employment report and the Mercantile Trade report. Next in line come a group of releases including Retail Sales, Durable Goods orders and Consumer Confidence, all of which have large impacts on volatility also. Finally, the NAPM, Leading Indicators and IP/CU

figures have consistently the smallest post-release impact.

4.5 Conclusions

In this work we have examined the importance of certain components of intra-daily FX volatility. Using a SV framework, based on that contained in Harvey, Ruiz, and Shephard (1994), we estimated seasonal patterns, announcement effects and an unobserved autoregressive component. Our results corroborate those of previous studies on seasonality, e.g. Andersen and Bollerslev (1997b), which point to the prominence of this phenomenon in this market and the necessity of its inclusion in any intra-day volatility examination. The results from the announcement data show that these are also an important element in volatility determination, confirming that the results of Ederington and Lee (1993) carry over to the spot FX market in a qualitative sense, although the impact here is quantitatively smaller. Our final specification shows that markets seem to quieten down in anticipation of news releases, but that post-release there is a pronounced and persistent volatility impact. If the seasonal is omitted from the specification, then it is shown that the estimated 'news' parameters are overstated in magnitude, as one would expect.

Examination of individual announcements points to the Employment report and Trade figures being associated with extremely large volatility impacts. There are also consistent, large responses to Consumer Confidence, Retail Sales and Durable Goods order figures, whilst the NAPM, Leading Indicator and IP/CU releases have the smallest effect.

There are several obvious extensions to the current study which could be undertaken. A first would be to examine how these results are affected when one looks at other exchange rates and the macroeconomic announcements of other countries. Also, in the current work we have only examined the link between 'news' and volatility. An interesting further study would be to examine how these announcements affect the

level of the exchange rate on an intra-daily data sampling. We go on to study the link between ‘news’ and exchange rate levels in Chapter 6 of this thesis. In Section 4.4, it was indicated that the correlogram of deseasonalised volatility demonstrated long memory tendencies of the type introduced by Granger and Joyeux (1980) and Hosking (1981). We explicitly address the possibility of long memory in volatility in Chapter 5. Finally, although this study has demonstrated that at least some of the movements in exchange rate volatility are explicable, there is still a large amount of volatility which is unexplained by the current model. It is possible that an examination of the effect of broader information flow variables on the exchange rate process will yield a more complete picture of the way in which FX volatilities are determined.

TABLES

Table 4.1: Summary Statistics

	Return	Residual	$\log(r^2)$
Mean	1.9×10^{-6}	0	-17.49
Variance	2.4×10^{-7}	2.4×10^{-7}	8.887
Skew	0.344	0.36	-1.27
Kurtosis	16.09	16.32	2.15
Q(20)	544.18	207.02	153715.8

Table 4.2: Baseline SV Models

Panel 1
Unrestricted SV Model

	Coeff	s.e.	T-stat
ϕ	0.9546	0.0132	72.32
σ_v^2	0.457	0.1051	4.35
LogL			-179817

Panel 2
Restricted SV Model: $\phi=1$

	Coeff	s.e.	T-stat
ϕ	1	-	-
σ_v^2	0.278	0.0066	42.22
LogL			-180486

Table 4.3: Estimates from SV/News Model

	Coeff	s.e.	T-stat
ϕ	0.954	0.0133	71.73
σ_v^2	0.456	0.1053	4.33
a_{-5}	-0.186	0.045	-4.16
a_{-4}	-0.273	0.047	-5.79
a_{-3}	-0.262	0.073	-3.55
a_{-2}	0.019	0.004	5.04
a_{-1}	0.415	0.078	5.29
a_0	0.493	0.078	6.32
a_1	2.376	0.142	16.72
a_2	1.533	0.173	8.86
a_3	1.390	0.219	6.36
a_4	1.184	0.052	22.59
a_5	1.004	0.157	6.38
a_6	0.869	0.225	3.86
a_7	0.730	0.133	5.49
a_8	0.979	0.146	6.72
a_9	0.317	0.025	12.49
a_{10}	0.664	0.177	3.76
a_{11}	0.345	0.012	28.05
a_{12}	0.355	0.068	5.22
a_{13}	0.804	0.099	8.13
a_{14}	0.371	0.046	8.01
a_{15}	-0.084	0.013	-6.45
Loglik			-179704

FIGURES

Figure 4.1: Average Daily Log Squared Return Pattern




Figure 4.2: Daily Autocorrelation Function




Table 4.4: Estimates from Seasonal SV Model

Panel 1: Trigonometric Parameters

	Coeff	s.e.	T-stat		Coeff	s.e.	T-stat
ϕ	0.962	0.0122	78.86	σ_v^2	0.309	0.0785	3.94
$\gamma_{1,1}$	-0.846	0.118	-7.18	$\gamma_{2,1}$	0.1	0.025	4.05
$\gamma_{1,3}$	0.520	0.026	20.33	$\gamma_{2,3}$	0.039	0.009	4.39
$\gamma_{1,4}$	0.133	0.007	19.34	$\gamma_{2,4}$	-0.082	0.023	-3.61
$\gamma_{1,5}$	-0.012	0.001	-9.18	$\gamma_{2,5}$	-0.003	0.001	-2.79
$\gamma_{1,9}$	0.137	0.009	15.19	$\gamma_{2,9}$	0.005	0.001	5.32
$\gamma_{1,10}$	0.139	0.003	49.66	$\gamma_{2,10}$	-0.006	0.001	-7.33
$\gamma_{1,1}^*$	-0.213	0.011	-19.05	$\gamma_{2,1}^*$	-0.156	0.015	-10.58
$\gamma_{1,3}^*$	-0.205	0.016	-13.06	$\gamma_{2,3}^*$	0.216	0.003	62.71
$\gamma_{1,4}^*$	0.360	0.013	26.7	$\gamma_{2,4}^*$	-0.179	0.011	-15.51
$\gamma_{1,5}^*$	0.196	0.032	6.1	$\gamma_{2,5}^*$	-0.05	0.007	-6.92
$\gamma_{1,9}^*$	-0.09	0.006	-14.62	$\gamma_{2,9}^*$	0.04	0.006	7.18
$\gamma_{1,10}^*$	0.032	0.003	10.89	$\gamma_{2,10}^*$	0.013	0.003	3.6

Panel 2: Japanese Lunch Dummies

	Coeff	s.e.	T-stat
μ_1	-0.905	0.121	-7.49
μ_2	-1.012	0.327	-3.12
μ_3	-1.243	0.148	-8.41
μ_4	-1.430	0.145	-9.89
μ_5	-1.555	0.130	-11.97
μ_6	-1.580	0.090	-17.55
μ_7	-1.429	0.109	-13.13
μ_8	-1.206	0.139	-8.65
μ_9	-1.103	0.224	-4.92
μ_{10}	-1.270	0.385	-3.30
μ_{11}	-0.834	0.052	-16.05
μ_{12}	-0.319	0.022	-14.37
μ_{13}	0.949	0.110	8.61
μ_{14}	1.146	0.076	15.09
μ_{15}	0.975	0.095	10.30
LogLik			-179227.3

Figure 4.3: Average Weekly Log Squared Return Pattern




Figure 4.4: Weekly Autocorrelation Function




Table 4.5: Announcement Effects from Combined Model

	Coeff	s.e.	T-stat
ϕ	0.962	0.0122	78.85
σ_v^2	0.31	0.0786	3.94
a_{-5}	-0.29	0.050	-5.78
a_{-4}	-0.406	0.027	-14.96
a_{-3}	-0.426	0.035	-12.15
a_{-2}	-0.213	0.011	-19.59
a_{-1}	0.192	0.038	5.12
a_0	0.269	0.054	4.95
a_1	2.127	0.117	18.24
a_2	1.270	0.10	12.22
a_3	1.127	0.112	10.05
a_4	0.916	0.085	10.78
a_5	0.744	0.224	3.31
a_6	0.598	0.046	12.90
a_7	0.459	0.201	2.28
a_8	0.711	0.021	34.10
a_9	0.056	0.006	9.59
a_{10}	0.416	0.053	7.77
a_{11}	0.121	0.027	4.43
a_{12}	0.155	0.010	14.78
a_{13}	0.622	0.101	6.09
a_{14}	0.210	0.008	26.51
a_{15}	-0.205	0.015	-14.06
LogLik			-179174.6

Figure 4.5: Average Daily Log Squared Return Pattern in Summer and Winter

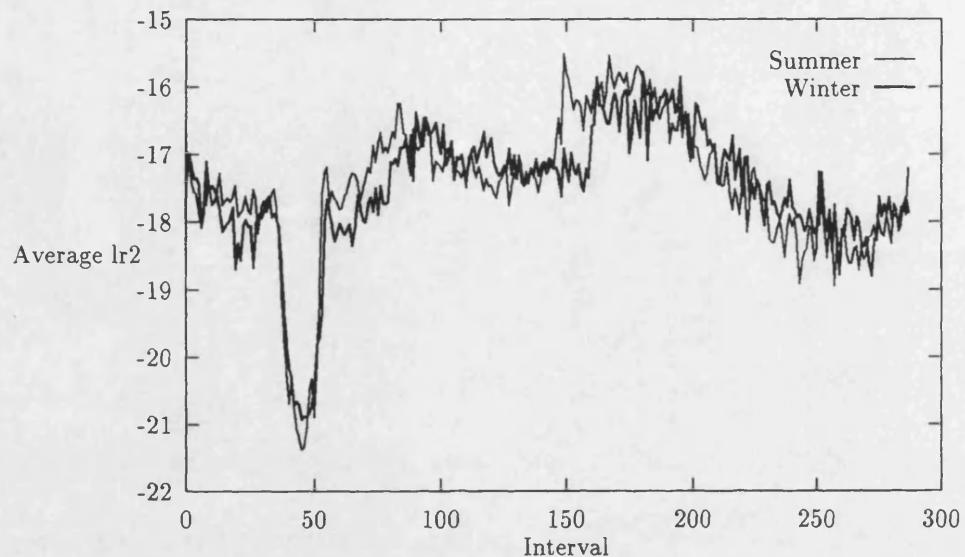


Figure 4.6: Comparison of Absolute Returns and Derived Volatility

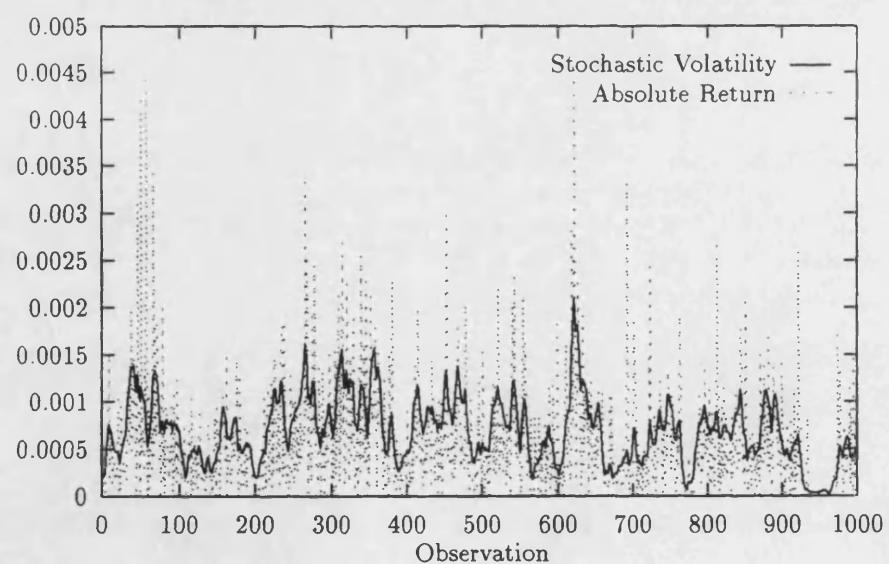


Table 4.6: Individual Announcement Impacts using Dummy Specification

Announcement	a_1	a_2	Persistence	% s.d. Response
Employment	4.66	3.02	0.85	1000+
Trade	3.59	1.68	0.86	600+
Retail Sales	3.20	2.51	0.80	450+
PPI	3.04	1.95	0.76	450
Cons. Conf.	2.38	1.10	0.80	330
Durables	2.37	1.65	0.88	330
CPI	2.35	2.01*	0.74	325
Lead. Ind.	1.89	1.49	0.67	260
NAPM	1.86	-0.49*	-0.48	250
IP/CU	0.72	0.71	0.9	140

Note : as more-or-less all coefficients are significant, only those insignificant at 5% are indicated, with an asterisk. 'Persistence' is calculated as $1 - \rho$.

Table 4.7: Individual Announcement Impacts using Absolute Forecast Errors

Announcement	a_1	a_2	Persistence	Scaled Impact
Unemp. Rate	19.72	13.20	0.87	2.96
Payroll Emp.	0.038	0.026	0.51	2.8
Trade	2.33	1.30	0.83	2.74
Retail Sales	4.47	2.93	0.88	1.75
PPI	5.83	4.42	0.79*	1.262
Cons. Conf.	0.45	0.17	0.74	1.84
Durables	0.88	0.37	0.90	2.09
CPI	12.34	8.41	0.70	1.65
Lead. Ind.	7.26	4.13	-0.92	1.06
NAPM	0.794	0.08	0.02*	1.56
Ind. Prod.	8.37	4.48	1.05	0.70
Cap. Util.	1.26	1.65	0.93	0.35

Note : as more-or-less all coefficients are significant, only those insignificant at 5% are indicated, with an asterisk. 'Persistence' is calculated as $1 - \rho$.

FIGURES

Figure 4.1: Average Daily Log Squared Return Pattern

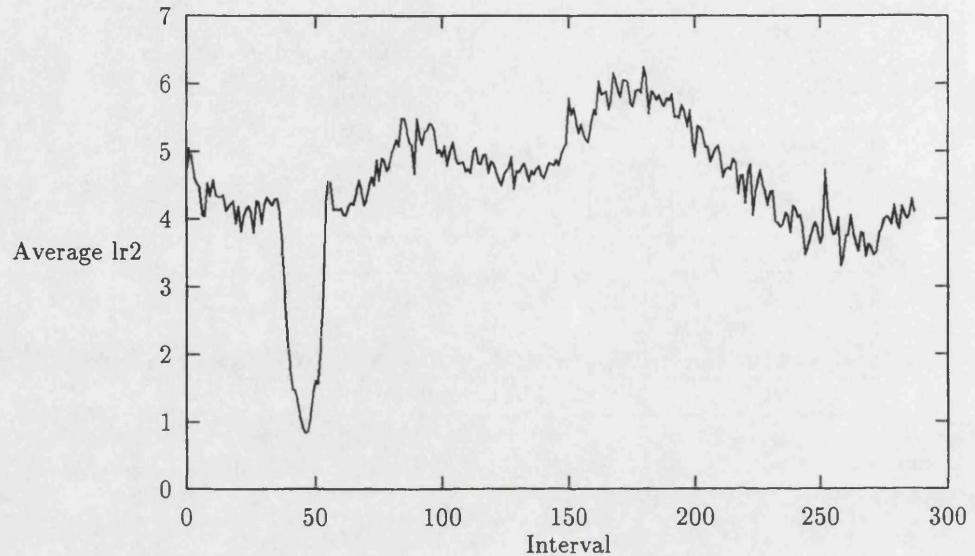


Figure 4.2: Daily Autocorrelation Function




Figure 4.3: Average Weekly Log Squared Return Pattern




Figure 4.4: Weekly Autocorrelation Function




Figure 4.5: Average Daily Log Squared Return Pattern in Summer and Winter




Figure 4.6: Comparison of Absolute Returns and Derived Volatility

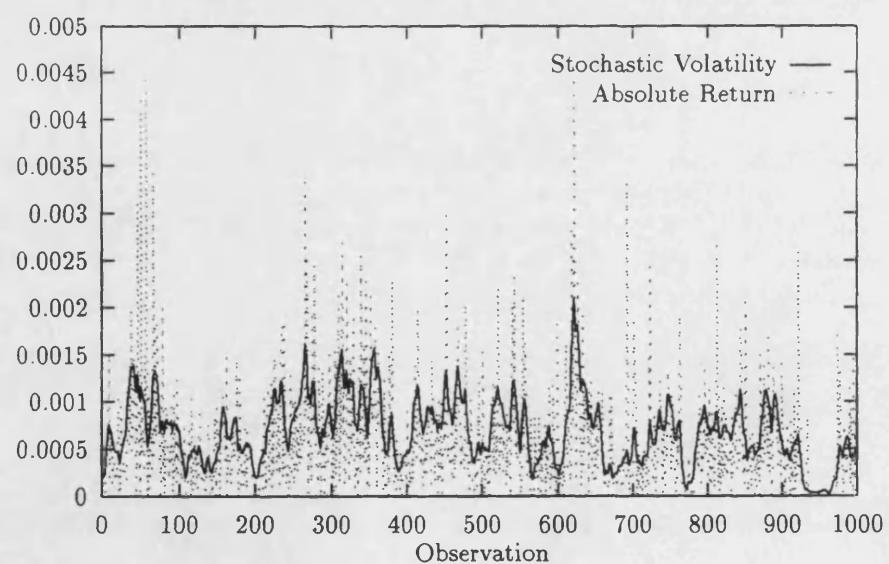


Figure 4.7: Average State Value from Non-seasonal Model and Log Squared Returns

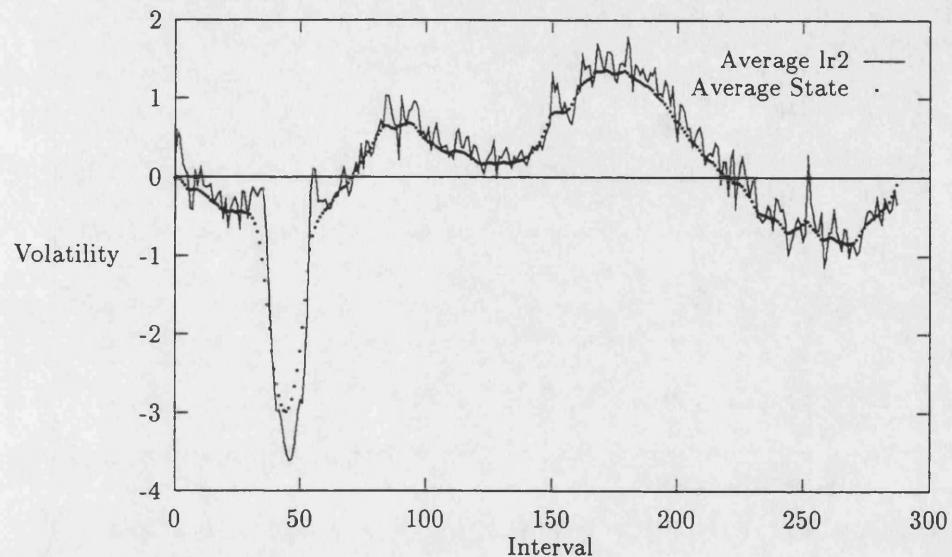


Figure 4.8: News Impacts from Non-Seasonal Model




Figure 4.9: Seasonal Periodogram

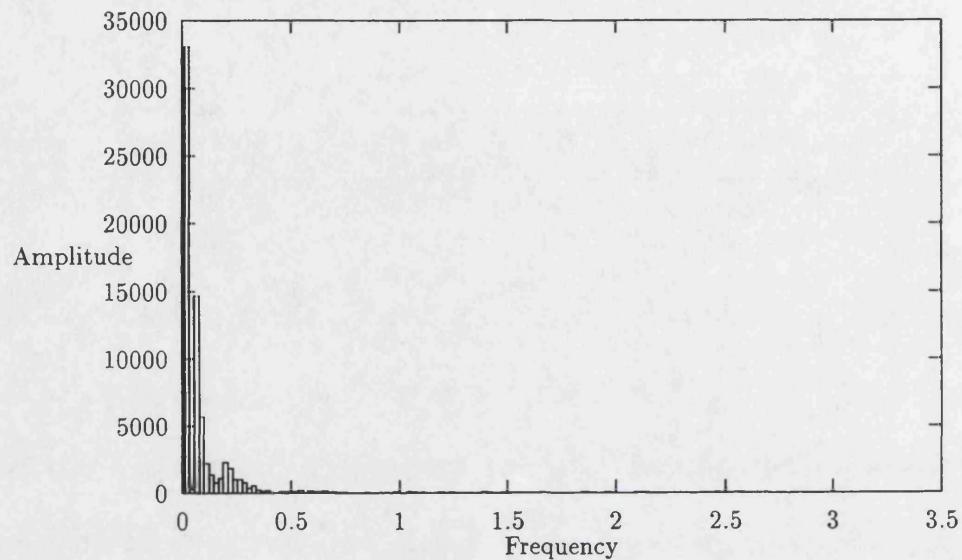



Figure 4.10: Estimated Seasonal in Volatility



Chapter 5

An Investigation of Long Range Dependence in Intra-Day Foreign Exchange Rate Volatility

5.1 Introduction

In recent years a vast amount of empirical work has been devoted to the characterisation of the temporal dependence in financial time series. Many authors have examined the time-series structure in asset returns, trading volumes and, perhaps most extensively, return volatility. Such studies are valuable in that they yield insights into issues such as the discrimination between regular and irregular market activity, the nature of information flows into financial markets, the way in which this information is assimilated into asset prices and the manner in which information is transmitted between markets.

The current study extends the research in this area. We examine the relative performance of two classes of econometric models, ARMA-type representations and

Figure 4.11: Weekly ACF for Deseasonalised Log Squared Returns

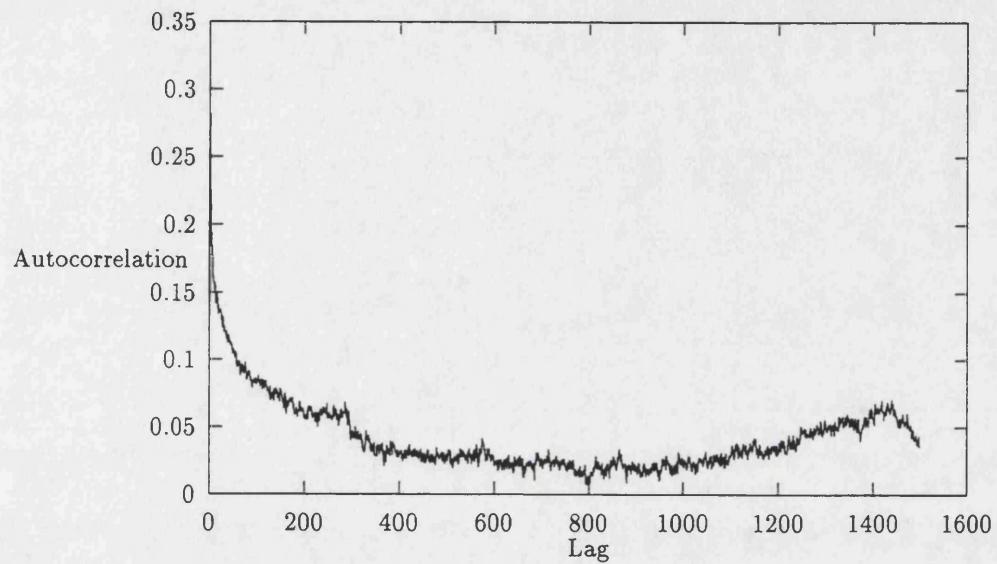



Figure 4.12: News Effects from Seasonal and Non-Seasonal SV Models



long range dependent specifications, in explaining the time-variation in financial time-series. The methodology employed yields first a robust characterisation of the degree of long memory in a given time-series, via a set of efficient tests for long range dependence and two semi-parametric estimates of the degree of fractional integration. We go on to estimate a fully parametric model which includes long and short memory structure and hence allows one to separate and assess the contributions of these two types of behaviour to overall dependence. At this second stage, the robust semi-parametric estimates can be used as a benchmark against which the adequacy of the fully parametric specification is judged. Our empirical analysis concentrates on modelling the volatility process associated with a year long intra-daily sample of three major exchange rates.

The standard methodologies used to assess the degree of temporal dependence in a given time-series are the ARIMA class of models for series such as returns and volumes and GARCH or Stochastic Volatility specifications for modelling conditional heteroskedasticity.¹ Both of the above classes of model imply similar restrictions on the nature of the temporal dependence in the series in question. Their covariance stationary representations imply that the process is completely mean reverting and that this mean reversion is exponential in nature. In terms of the properties of the infinite MA representation, the series of Wold coefficients is both absolutely and square summable. The usual alternative to covariance stationarity is to assume the series is first order integrated i.e. an ARIMA($p,1,q$) in the conditional mean or an IGARCH conditional variance process. This representation implies vastly different behaviour for the time-series in question. Shocks are completely persistent and hence the process diverges.

Clearly, then, there is a very large distinction between the behaviours implied by $I(0)$ and $I(1)$ processes. On the one hand, one has a process which mean reverts at an exponential rate whilst on the other hand, the process is completely persistent. A far

¹Bollerslev, Chou, and Kroner (1992) and Ghysels, Harvey, and Renault (1996) give excellent reviews of the GARCH and Stochastic Volatility modelling approaches respectively.

Chapter 5

An Investigation of Long Range Dependence in Intra-Day Foreign Exchange Rate Volatility

5.1 Introduction

In recent years a vast amount of empirical work has been devoted to the characterisation of the temporal dependence in financial time series. Many authors have examined the time-series structure in asset returns, trading volumes and, perhaps most extensively, return volatility. Such studies are valuable in that they yield insights into issues such as the discrimination between regular and irregular market activity, the nature of information flows into financial markets, the way in which this information is assimilated into asset prices and the manner in which information is transmitted between markets.

The current study extends the research in this area. We examine the relative performance of two classes of econometric models, ARMA-type representations and

long range dependent specifications, in explaining the time-variation in financial time-series. The methodology employed yields first a robust characterisation of the degree of long memory in a given time-series, via a set of efficient tests for long range dependence and two semi-parametric estimates of the degree of fractional integration. We go on to estimate a fully parametric model which includes long and short memory structure and hence allows one to separate and assess the contributions of these two types of behaviour to overall dependence. At this second stage, the robust semi-parametric estimates can be used as a benchmark against which the adequacy of the fully parametric specification is judged. Our empirical analysis concentrates on modelling the volatility process associated with a year long intra-daily sample of three major exchange rates.

The standard methodologies used to assess the degree of temporal dependence in a given time-series are the ARIMA class of models for series such as returns and volumes and GARCH or Stochastic Volatility specifications for modelling conditional heteroskedasticity.¹ Both of the above classes of model imply similar restrictions on the nature of the temporal dependence in the series in question. Their covariance stationary representations imply that the process is completely mean reverting and that this mean reversion is exponential in nature. In terms of the properties of the infinite MA representation, the series of Wold coefficients is both absolutely and square summable. The usual alternative to covariance stationarity is to assume the series is first order integrated i.e. an ARIMA($p,1,q$) in the conditional mean or an IGARCH conditional variance process. This representation implies vastly different behaviour for the time-series in question. Shocks are completely persistent and hence the process diverges.

Clearly, then, there is a very large distinction between the behaviours implied by $I(0)$ and $I(1)$ processes. On the one hand, one has a process which mean reverts at an exponential rate whilst on the other hand, the process is completely persistent. A far

¹Bollerslev, Chou, and Kroner (1992) and Ghysels, Harvey, and Renault (1996) give excellent reviews of the GARCH and Stochastic Volatility modelling approaches respectively.

more general description of temporal dependence can be achieved by using more recent models which allow for long range dependence or long memory. One such model is the ARFIMA representation proposed by Granger and Joyeux (1980). Similar to the ARIMA class they permit non-zero orders of integration but, importantly, the order of integration is allowed to be non-integer.² This generalisation allows one to model covariance stationary processes which have non-summable autocovariances. Hence the effect of a shock on the process is completely transitory although the process of mean reversion is very prolonged (the rate of decay in the Wold coefficients is hyperbolic rather than exponential as in the $I(0)$ case.) The model also permits non-stationary processes which mean revert. Hence, the ARFIMA model can be viewed as a more flexible alternative to the standard ARIMA, permitting a far a more general characterisation of the temporal dependencies in a given time-series.

We propose a comprehensive methodology for assessing the nature of temporal dependence. The methodology comprises the following steps. First we test the degree of integration in the process using the methodology presented in Robinson (1996). The test is based on an underlying ARFIMA structure for the series in question and permits any degree of integration (integer or fraction) as a null hypothesis. Next, we gain a precise estimate of the degree of integration using two robust semiparametric long memory estimators. The first is the Log Periodogram regression (due to Geweke and Porter-Hudak (1983)) and the second is the Gaussian estimate proposed by Künsch (1987). Both of these estimates have been proven, in Robinson (1995b) and Robinson (1995a), to have standard asymptotics and are robust to any short range dependent features in the data. At this point, and analogous to the $I(1)$ case, one could filter the long range dependence from the series and fit a covariance stationary ARMA to the residuals using traditional model selection procedures. We, however, go on to fit a fully parametric model to the series. The model is a generalisation of the Long Memory in Stochastic Volatility (LMSV) model presented in

²This is achieved by raising the first differencing operator to a given power which is permitted to take any real value. Hence, these models are also termed 'fractionally integrated.'

Harvey (1993), allowing for both short and long range dependence, and is sensitive to any mis-specification due to its fully parametric nature. The reason for fitting the fully parametric model is to permit one to assess the contributions of the short and long memory components to the overall dependence in the series. This is accomplished via a set of quasi-Likelihood Ratio statistics for the fully parametric model. Hence, ultimately, we can discriminate between the long and short range dependent features of the process.

The data we examine in this work are the volatilities associated with three intra-day foreign exchange (FX) returns series (the exchange rates being the DEM/USD, JPY/USD and JPY/DEM.) A pervasive result from previous work on this type of data is that the volatility process can be characterised as very nearly non-stationary (see, *inter alia*, Andersen and Bollerslev (1997b), DeGennaro and Shrieves (1995), Guillaume (1995) and the analysis in Chapter 4.) There are however, some indications that this result may be due to mis-specification of the volatility models employed. First, the temporal aggregation results for GARCH processes do not hold when applied to FX data. The degree of persistence one identifies in daily data, for example, far exceeds that which would be implied by the results of estimations from data sampled at 1 hour intervals.³ Second, as noted in Chapter 4 and Dacorogna, Müller, Nagler, Olsen, and Pictet (1993), the correlograms of these intra-day volatility series decay far more slowly than the exponential decay which is associated with conventional GARCH or SV models. The combination of these two points serves as the motivation for our investigation of long memory in volatility.

A theoretical motivation for the presence of long memory in asset price volatility can be generated by combining the mixture of distributions model in Tauchen and Pitts (1983) and the results on aggregation in Granger (1980). The former demonstrate, in a highly stylized framework, that both the volume and volatility in asset markets inherit the temporal dependencies associated with the latent flow of information into

³See Andersen and Bollerslev (1997b), for example.

the market. Now assume that information flows are heterogeneous. Specifically, assume that there are an infinity of information arrival processes, each of which follows a stationary autoregression. The heterogeneity is modelled by variation in the AR parameters, which we assume follow a beta distribution. As Granger (1980) demonstrates, the aggregate information flow process will then exhibit long range dependence and, hence, so will volatility.

The paper is set out as follows. In Section 5.2 we present a more detailed account of our empirical methodology. Section 5.3 introduces the data employed in our study. As previously mentioned we focus on intra-daily FX volatilities, which are sampled at a ten minute calendar interval. Section 5.4 presents our estimation and testing results. We find that FX returns are well characterised by an $I(0)$ process, in line with the efficient markets hypothesis. Results for the three volatility series demonstrate that all are covariance stationary and exhibit significant long memory. Further, estimation and testing of the fully parametric model demonstrates that the finding of non-stationarity in intra-day volatility is due to mis-specification. When one permits the possibility of long memory in volatility all specifications strongly indicate covariance stationarity. Section 5.5 concludes the paper and points to some directions for further work.

5.2 Methodology

5.2.1 Long Memory Specifications

Let $\{r_t\}^T$ be our time series of raw returns and define $y_t \equiv \log r_t^2$ as an estimate of the volatility of the process.⁴ Assume y_t to be covariance stationary with autocovariance γ_j at the j^{th} lag. We require the autocovariances to satisfy the basic long range

⁴The use of this proxy for volatility is motivated by the assumption that returns follow the following process, $r_t = \sigma e^{y_t/2} \epsilon_t$ where $\epsilon_t \sim N(0, 1)$. This follows Harvey (1993), Harvey, Ruiz, and Shephard (1994) and is presented and discussed more fully in Section 4.3.

dependent specification,

$$\gamma_j \sim cj^{2d-1} \text{ as } j \rightarrow \infty \quad (5.1)$$

where c is a finite positive constant and d is the slope parameter. Processes following 5.1 are generally labelled $I(d)$ processes. An $I(d)$ process is invertible when $d > -1/2$, in which case it can be reparameterized in an infinite moving average or Wold representation. The long lag coefficients in the Wold representation of a fractionally integrated process can be shown to be of the form given below;

$$\psi_j \sim kj^{d-1} \text{ as } j \rightarrow \infty \quad (5.2)$$

where k is another finite, positive constant. The effects which the value of d has on the properties of an $I(d)$ process are directly obtainable from equation 5.2. When $d = 0$ we are in the standard $I(0)$ paradigm so that the process is short memory, covariance stationary and mean reverting. For $0 < d < 1/2$, the coefficients in the linear representation become non-summable although they remain square-summable. The latter fact implies that the process is still covariance stationary, but, as equation 5.2 demonstrates, the coefficients decay hyperbolically and it is this property which defines the long memory characteristic. When $1/2 < d < 1$ the coefficients in the Wold representation are neither absolutely nor square summable and hence the process is non-stationary, long memory. Finally, note that for any d less than unity the process mean reverts.

Hence, the main feature of specifications 5.1 and 5.2 is the smooth description of temporal dependence from the short range dependent $I(0)$ case through the boundary of stationarity $I(1/2)$ to the unit root $I(1)$ case. The parameter d provides a measure of the strength of temporal dependence in the process. The focus of our attention is the range $d \in (0, 1/2)$ within which the process is covariance stationary

and long range dependent.

Specification 5.1 translates to a related specification in the frequency domain.⁵ We assume that y_t has spectral density $f(\lambda)$ satisfying

$$\gamma_j = \int_{-\pi}^{\pi} f(\lambda) \cos(j\lambda) d\lambda.$$

The basic long range dependent specification in the frequency domain is

$$f(\lambda) \sim g\lambda^{-2d} \text{ as } \lambda \rightarrow 0^+ \quad (5.3)$$

where g is a finite positive constant and d is the slope parameter.

Equations 5.1 and 5.3 deliver the essential features of long range dependent processes in the time and frequency domains respectively. Specification 5.1 demonstrates that the long memory characteristic is described by long-lag autocovariances which decay hyperbolically whilst equation 5.3 shows that, in the frequency domain, long range dependence implies the spectral density of the process contains a hyperbolic pole around zero frequency. The most commonly used class of fully-parametric time-series models which display long memory is the aforementioned ARFIMA family, introduced by Granger and Joyeux (1980), which are constructed as below,

$$(1 - L)^d y_t = \epsilon_t, \quad (5.4)$$

where ϵ_t has a stationary ARMA(p, q) representation and, in general, d is taken to be less than unity in absolute value. The fractional noise model is the simplest long memory model in the class and corresponds to ϵ_t being a serially independent

⁵See Robinson (1994) for details on the correspondence between the frequency domain specification 5.3 and the time domain specification 5.1.

normal variate. Models of the ARFIMA type conform with equations 5.1, 5.2 and 5.3.

5.2.2 Semiparametric initial estimates

In Section 5.4 we report the results from two semiparametric, frequency domain estimates of the long memory parameter d in equation 5.3. These are the Gaussian semiparametric estimate proposed by Künsch (1987) and the log-periodogram estimate proposed by Geweke and Porter-Hudak (1983). The asymptotic properties of these estimates were derived in Robinson (1995a) and Robinson (1995b) respectively. Both estimation procedures are based on the periodogram,

$$I(\lambda) = \frac{1}{2\pi T} \left| \sum_{t=1}^T y_t \exp(it\lambda) \right|^2$$

computed at the harmonic frequencies $\lambda_j = 2\pi j/T$ for $j \in [1, T/2]$.

The log-periodogram estimate is often preferred for its intuitive appeal. The long memory property is characterized by the existence of a hyperbolic pole in the spectral density around zero frequency as in 5.3. The log-periodogram estimate works by linearizing the expression for the spectral density in the parameter d such that d can be retrieved via a linear regression.

Specifically, assume that our time series, y_t , is generated by an ARFIMA model (equation 5.4.) The spectral density of y_t (denoted $f_y(\lambda)$) can be shown to be,

$$f_y(\lambda) = \frac{\sigma^2}{2\pi} (1 - e^{i\lambda})^{-d} (1 - e^{-i\lambda})^{-d} f_\epsilon(\lambda), \quad (5.5)$$

where $f_\epsilon(\lambda)$ is the spectral generating function of ϵ_t and σ^2 is the variance of the innovations to the ϵ_t process. This expression can be shown to be equivalent to,

$$f_y(\lambda) = \frac{\sigma^2}{2\pi} (4 \sin^2(\frac{\lambda}{2}))^{-d} f_\epsilon(\lambda). \quad (5.6)$$

From 5.6 it is clear that, after taking logarithms of both sides of the expression, the log spectral density is linear in d . Hence, as equation 5.6 would suggest, Geweke and Porter-Hudak (1983) proposed employing the log periodogram in place of the log spectral density and retrieving d from the following OLS regression,

$$\log(I(\lambda_j)) = c - d \log(4 \sin^2(\frac{\lambda_j}{2})) + u_j, \quad j = l, \dots, m \quad (5.7)$$

where m is the bandwidth and l is a trimming parameter such that $\frac{1}{l} + \frac{l}{m} + \frac{m}{T} \rightarrow 0$ as $T \rightarrow \infty$ and u_j is a Gaussian error term.⁶ The bandwidth parameter (m) is introduced such that only the low frequency components are included in the estimation and therefore the estimator is robust to any short memory (high frequency) characteristics of the process. Comte and Hardouin (1995) further demonstrated the pathological behaviour of the extremely low frequency periodogram ordinates, necessitating that they be removed from the regression analysis and hence introducing a further parameter, l . Finally, the assumption of unconditional Gaussianity of the regression errors (u_j) is a drawback in many applications but it is extremely hard to relax due to the high degree of non-linearity involved.

Note that in what follows in Section 5.4 the log-periodogram estimator is employed using $m = \sqrt{n}$ and $l = m^{0.25}$. The results presented later are robust to alternative choices of the bandwidth and trimming parameters.

The other semiparametric estimate which we employ is the Gaussian semiparametric estimate, proposed by Künsch (1987), and is preferred generally for its efficiency and robustness properties.⁷ Like the log-periodogram estimate, it is semiparametric in

⁶The above condition on m and l guarantees the consistency and asymptotic normality of the log-periodogram estimator.

⁷Log-periodogram estimation results will be reported for comparison.

the sense that it only relies on the low frequency harmonics of the periodogram, but in this case, the lowest harmonics need not be trimmed out, such that it is consistent under the minimal bandwidth requirement

$$\frac{1}{m} + \frac{m}{T} \rightarrow 0 \quad \text{as } T \rightarrow \infty, \quad (5.8)$$

and an automatic bandwidth selection procedure is available in Henry and Robinson (1996).⁸ Estimation is based on the maximisation of a local form of frequency domain log-likelihood,

$$-\frac{1}{m} \sum_{j=1}^m \left\{ \log g \lambda_j^{-2d} + \frac{I(\lambda_j)}{g \lambda_j^{-2d}} \right\}$$

for $0 < m < [T/2]$ which, after concentrating out the constant g , is equivalent to minimizing,

$$R(d) = \log \left\{ \frac{1}{m} \sum_{j=1}^m \lambda_j^{2d} I(\lambda_j) \right\} - \frac{2d}{m} \sum_{j=1}^m \log \lambda_j.$$

The most appealing features of this estimate are its asymptotic normality (under a slightly stronger requirement than 5.8) with asymptotic variance $1/4m$ independent of d and g (see Robinson (1995a)) and the robustness of this asymptotic normality result to non-Gaussianity and conditional heteroskedasticity of general form in the Wold innovations of the process under investigation (including long memory GARCH and other specifications introduced in Robinson (1991b)).

⁸This bandwidth selection procedure is based on the minimisation of mean squared estimation errors.

5.2.3 Fully parametric specification

The semi parametric estimates of the fractional differencing parameter presented in the previous subsection are robust to the short memory characteristics of any process due to the fact that they consider the behaviour of the spectral density only in the neighbourhood of zero frequency. A fully parametric estimate, however, must specify the short memory components correctly in order for the estimate of d to be consistent. Our rationale for presenting the semi-parametric estimates is, therefore, twofold. First they yield a robust characterisation of the degree of long memory in the process. Second, the semi-parametric estimates can be used to validate the results from a fully parametric specification i.e. the semi-parametric estimates provide a benchmark against which the results of the fully parametric model can be inspected for obvious signs of bias.

This second point is important in the current context as our motivation for estimating a fully parametric model is to enable one to discriminate between the long memory and short memory features. Comparison of a general, fully parametric model with certain restricted alternatives allows one to assess the likelihood contributions and significance of each of the posited components of the volatility model.

The parametric model we adopt is an extension of the LMSV model of Harvey (1993):

$$\left\{ \begin{array}{l} y_t = c + h_t + \xi_t \\ (1 - L)^d (1 - \phi L) h_t = \eta_t \end{array} \right. \quad (5.9)$$

where c is a constant, ξ has mean zero and variance $\pi^2/2$, $\eta \sim N(0, \sigma_\eta^2)$ and d lies within the stationarity and invertibility range $(-1/2, 1/2)$.⁹ This framework

⁹This specification for the log squared returns is derived from the formulation in footnote 4,

is consistent with the semiparametric specification described above insofar as the spectral density of a process thus specified follows 5.3.¹⁰

The estimation procedure is a frequency domain log likelihood maximisation.¹¹ Exact maximum likelihood has been proven (Fox and Taqqu (1986)) to be asymptotically most efficient in the case of a correctly specified normal error distribution η_t . The estimates of the fractional differencing parameter \hat{d} and of the autoregressive parameter $\hat{\phi}$ maximize

$$\mathcal{L}(d, \phi) = - \sum_{j=1}^{T-1} \log(2\pi g_j) - \sum_{j=1}^{T-1} \frac{I(\lambda_j)}{g_j} \quad (5.10)$$

where

$$g_j = \frac{\sigma_\eta^2 (4 \sin^2(\frac{\lambda_j}{2}))^{-d}}{2\pi(1 - 2\phi \cos \lambda_j + \phi^2)} + \frac{\sigma_\xi^2}{2\pi}.$$

is the spectral density of y_t . Note that the maximisation of specification 5.10 must be treated as a Quasi-Maximum Likelihood estimation due to the non-Gaussianity of ξ_t .

The advantage which the fully parametric specification yields is that it nests the standard autoregressive SV representation and the LMSV model of Harvey (1993). Hence, the importance of the AR(1) and long memory components can be evaluated with a Quasi-Likelihood Ratio test. Specifically, the unconstrained model (specification 5.9) is compared to two nested alternatives,

hence $\xi_t = \log(\epsilon_t^2) + 1.27$ is distributed as a $\log \chi^2$ variate. See also equation 4.4.

¹⁰This follows from the facts that h_t is an ARFIMA(1, d , 0) process and is independent of ξ_t .

¹¹Note that in this case, the maximisation is performed over the whole range of harmonic frequencies. In the semiparametric case, only a degenerate band of harmonic frequencies was used. As mentioned earlier, the present estimate is therefore sensitive to any short range dependent misspecification. It is likely to be sensitive to the seasonal component in the series discussed in Section 5.2.5, but it is nonetheless reported before as well as after deseasonalisation for completeness.

$$(LM) \left\{ \begin{array}{l} y_t = c + h_t + \xi_t \\ (1 - L)^d h_t = \eta_t \end{array} \right. \quad \text{and} \quad (AR) \left\{ \begin{array}{l} y_t = c + h_t + \xi_t \\ (1 - \phi L) h_t = \eta_t \end{array} \right.$$

obtained for $\phi = 0$ and $d = 0$ respectively. The frequency domain likelihood is computed as in 5.10 for each of these nested models and the following tests are performed using the likelihood ratio principle:

$$H_0 : d = 0 \quad \text{against} \quad H_a : d > 0.$$

$$H_0 : \phi = 0 \quad \text{against} \quad H_a : 0 < |\phi| < 1.$$

If the subscript u denotes the unconstrained estimates, the likelihood ratio statistic is $2 [\mathcal{L}(\hat{d}_u, \hat{\phi}_u) - \mathcal{L}(\hat{d}, \hat{\phi})]$ where \mathcal{L} is the concentrated form of the quasi-likelihood in 5.10.¹²

5.2.4 Testing for persistence, long range dependence and stationarity

Given that the Gaussian semiparametric estimate is justified for all values of d in a compact subset of $(-\frac{1}{2}, \frac{1}{2})$, a preliminary test of stationary and invertibility of the log squared returns (y_t) is required. The testing procedure presented here is fully parametric, and hence sensitive, as indicated above, to misspecified short range dynamics. Therefore the conclusions of the test need to be confirmed after the model selection stage. The procedure relies on efficient tests of long range dependence (Robinson (1996) and Gil-Alaña and Robinson (1995)) which permit a wide class of null hypotheses. The object of the test is the hypothesis of persistence in foreign ex-

¹²In Section 5.4 the significance of the quasi-LR statistics is evaluated against a χ^2 -distribution. The assumption that the limiting distribution of these statistics is χ^2 is, at this point, still a conjecture.

change volatility. As was mentioned above, most of the available methods for testing for unit roots (see Diebold and Nerlove (1989) for a review on the subject) have non standard limiting distributions and lack Pitman efficiency.¹³ Unit root tests against autoregressive alternatives, in particular, are based on the Wald, Likelihood ratio and Lagrange multiplier principles, but they lack the sufficient degree of smoothness across the parameter of interest that would yield null χ^2 limiting distributions and Pitman efficiency.

Moreover, these tests give only one possible persistence null hypothesis. The testing procedure used here, on the other hand, allows one to postulate any value of d (integer or fraction) as a null hypothesis and possesses efficiency and a null χ^2 limiting distribution. Consider an ARFIMA(1, d , 0) model for y_t , i.e. $(1 - L)^d(1 - \phi L)y_t = \eta_t$ which can be rewritten as $(1 - L)^d y_t = u_t$ where u_t is a stationary AR(1), therefore I(0), process. u_t has spectral density,

$$f_u(\lambda; \phi) = \frac{\sigma_\eta^2}{2\pi} \left[\frac{1}{1 - 2\phi \cos \lambda + \phi^2} \right].$$

Suppose we want to test the hypothesis $H_0 : d = d_0$. Let $I_u(\lambda)$ be the periodogram of the residuals $\tilde{u}_t = (1 - L)^{d_0} y_t$. The frequency domain quasi-likelihood is

$$\mathcal{L}(\sigma_\eta^2, \phi) = - \sum_{j=1}^{T-1} \log(2\pi f_{u_j}) - \sum_{j=1}^{T-1} \frac{I_u(\lambda_j)}{f_{u_j}} \quad (5.11)$$

where $f_{u_j} = f_u(\lambda_j; \phi)$. Concentration of this likelihood yields \sqrt{n} -consistent estimates $\hat{\phi} = \arg \min_\phi \sigma_\eta^2(\phi)$ and $\hat{\sigma}_\eta^2 = \sigma_\eta^2(\hat{\phi})$ where $\sigma_\eta^2(\phi) = \frac{2\pi}{T} \sum_{j=1}^{T-1} \frac{I_u(\lambda_j)}{f_{u_j}}$.

The test statistic is constructed on the score principle. Let ω be the $(T-1) \times 1$ vector with j -th element $\log(4 \sin^2(\frac{\lambda_j}{2}))$, let \hat{f}_u be the $(T-1) \times 1$ vector with j -th element $f_u(\lambda_j; \hat{\phi})$, and let M be the projector on the space orthogonal to the $(T-1) \times 1$

¹³These tests are improved in Elliot, Stock, and Rothenberg (1994)

vector with j -th element $\frac{\partial}{\partial \phi} \log f_u(\lambda_j; \hat{\phi})$. The test statistic is

$$\hat{S} = -\frac{\pi}{\hat{\sigma}_\eta^2} \frac{\omega' \hat{f}_u}{\|M\omega\|}.$$

Under suitable regularity conditions (Robinson (1996)), $\hat{S} \xrightarrow{D} N(0, 1)$ as $T \rightarrow \infty$.

The resulting testing rules for H_0 are summarized in the table below:

Alternative Hypothesis	Reject H_0 when
$H_1 : d > d_0$	$\hat{S} > z_\alpha$
$H_1 : d < d_0$	$\hat{S} < -z_\alpha$
$H_1 : d \neq d_0$	$\hat{S} > z_{\alpha/2}$

Note: Rules for α -level tests of $H_0 : d = d_0$ against various alternatives. z_α is the quantile of a standard normal variate.

This testing procedure provides us with two efficient tests of persistence: the null $H_0 : d = 1$ against the alternative $H_1 : d < 1$, which is the unit root test, and the null $H_0 : d = 1/2$ against the alternative $H_1 : d < 1/2$, which is a test of non-stationarity.

5.2.5 Deseasonalisation

A salient feature of the intra-day volatility data considered in this work is the seasonality in volatility described in more detail in Andersen and Bollerslev (1997b) and Chapter 4. Seasonal components appear in the periodogram as peaks at certain harmonic frequencies. These peaks affect all periodogram based estimation and the fully parametric estimation becomes invalid.

If one thinks of the spectrum of the process with strong seasonal components as a

mixed spectrum, there is a need for spectral estimation methods which remove the Dirac mass points at the seasonal frequencies and smooth out the leakage from these peaks into the neighbouring frequencies. The spectral estimate used here is a Double-Window smoother proposed by Priestley (1981) which is designed to remove seasonal components and the leakage around the seasonal frequency. Suppose the volatility series is decomposed into two uncorrelated components $y_t = z_t + \varsigma_t$ where z_t has a continuous spectral density and $\varsigma_t = \sum_{r=1}^K A_r \cos(\omega_r t + \phi_r)$. The examination of this mixed spectrum is greatly simplified by the knowledge of the seasonal harmonics ω_r , which correspond to the weekly frequency and multiples of the daily frequency.¹⁴ The amplitudes of the seasonal harmonics can be estimated through a regression of y_t against $(\cos(\omega_r t + \phi_r))_{r=1}^K$, and the spectrum of z_t is consistently estimated with a Double-Window smoother. The spectral window adopted is the Bartlett-Priestley window

$$W(\theta; M) = \begin{cases} \frac{3M}{4\pi} \left\{ 1 - \left(\frac{M\theta}{\pi} \right)^2 \right\}, & |\theta| \leq \frac{\pi}{M}, \\ 0, & |\theta| \geq \frac{\pi}{M}, \end{cases}$$

where M is the bandwidth.¹⁵ Define $\hat{f}_M(\omega) = \int_{-\pi}^{\pi} I(\theta)W(\omega - \theta; M)d\theta$ the spectral estimate using $W(\theta; M)$. The Double Window spectral estimate is constructed as follows:

¹⁴Hence we can avoid employing tests to detect harmonic components (Whittle, Bartlett, Hannan or Priestley, in Priestley (1981)).

¹⁵This spectral window is a smoothed version of the Daniell (or rectangular) window and it is chosen for its compact support.

$$\hat{f}_{DW}(\omega) = \begin{cases} \hat{f}_n(\omega), & |\omega - \omega_r| > \frac{\pi}{n}, \\ (\hat{f}_m(\omega) - c\hat{f}_n(\omega))/(1 - c), & |\omega - \omega_r| \leq \frac{\pi}{n}, \end{cases}$$

where $n > m$, $c = W(0; m)/W(0; n)$ and the ω_r 's are the harmonics of the seasonal components defined above. A cross-validated likelihood maximizing procedure for the determination of both bandwidths (see Hurvich (1985), Beltrão and Bloomfield (1987) and Robinson (1991a) for the asymptotics) proved computationally too expensive and gave poor results. An ad hoc choice of bandwidths $n = \sqrt{T}$ and $m = n/10$ was preferred.

5.3 The Data

As indicated in Section 5.1, the focus of this work is the behaviour of volatility in the intra-day Foreign Exchange (FX) market. We study three sets of FX returns, on the DEM/USD, JPY/USD and JPY/DEM, covering the period from the beginning of October 1992 to the end of September 1993.¹⁶ These return series are filtered transcriptions of the tick-by-tick quotation series which appear on the Reuters FXFX page. Each quote encompasses a timestamp, bid and ask quotation pair, plus identifiers which allow one to determine the inputting bank and its location. In this study we ignore the identification of the inputting institution, using the tick-by-tick data solely to construct a regularly spaced calendar time-series.

The basic horizon over which we calculate returns is 10 minutes.¹⁷ This yields, for

¹⁶These data were supplied by Olsen and Associates (Zurich), to whom we are most grateful.

¹⁷Returns are determined exactly as outlined in Chapter 4 except using a 10 minute sampling.

each currency, a time-series with 37583 observations. The basic summary statistics of the returns are shown in Table 5.1.

Table 5.1 illustrates the following facts. First, all three return series have a mean which is insignificantly different from zero. A point which conforms with many earlier studies is that there is pronounced excess kurtosis in the returns distribution. This, as pointed out by Bollerslev and Domowitz (1993), is a natural feature of time-series which display conditional heteroskedasticity, although their analysis shows that after correcting for the conditional heteroskedasticity much of the kurtosis remains. Finally, the autocorrelation coefficients show that there is some temporal dependence in the return series, the DEM/USD and JPY/USD demonstrating negative autocorrelation whilst the JPY/DEM displays positive first-order autocorrelation. The significance of these autocorrelation coefficients is confirmed in the Box-Ljung statistics, which demonstrate that one cannot reject the hypothesis of up to tenth order serial correlation.

In Table 5.2 we present identical sets of statistics for our volatility proxy. We employ the logarithm of squared returns as our volatility measure, a choice which is motivated by the Long Memory in Stochastic Volatility model which was presented in Section 5.2.3. The main feature of these results lies in the correlation structure of volatility. As is visible from comparing Tables 5.1 and 5.2, there is far larger dependence in volatility than in returns. The first-order autocorrelations are between 3 and 5 times greater for volatility than for returns, whilst the Box-Ljung statistics are at least an order of magnitude greater. The characterisation of this temporal dependence is the focus of this work.

In order to clarify the nature of the dependence in volatility, in Figures 5.1 to 5.3 we present the first 1000 periodogram and logged periodogram ordinates for the

Note also that the results presented in this paper carry over to the analysis of percentage returns.

JPY/USD volatility plus the first 1000 sample autocorrelations.¹⁸

Examining first the correlogram, one feature which is immediately apparent is the existence of a pronounced daily seasonal in volatility. This seasonal has recently been the subject of many papers, including Dacorogna, Müller, Nagler, Olsen, and Pictet (1993), Andersen and Bollerslev (1997b) and Chapter 4. It is generated by the 24 hour activity in the foreign exchange market and the alterations in market activity which occur as trading shifts from the Far East to Europe to North America and so on. There is also evidence of seasonality at the weekly frequency. In the current context, however, this component is of no intrinsic interest and simply masks the underlying temporal structure of volatility. Hence, when estimating our long memory specifications we filter this component using the double window smoother presented in Section 5.2.5.

In the periodogram of the data this seasonal component is represented by peaks at integer multiples of the fundamental seasonal frequency.¹⁹ A feature of the periodogram which is more relevant to the current study is the behaviour of the periodogram in a neighbourhood of zero frequency, where the peak (visible on Figures 5.1 and 5.2) can be viewed as tentative evidence for the presence of long memory in our volatility series.

¹⁸Throughout the work we present graphical examples for this exchange rate only as those for the other two are qualitatively similar.

¹⁹As there are 144 ten minute intervals in one day, the seasonal frequency is $\frac{2\pi}{144}$, corresponding, approximately, to harmonic 228.

5.4 Results

Our results are presented in two stages. In the first sub-section we document the results of pre-testing for long range dependence, using the procedure developed in Robinson (1996). We then go on to present the estimations of the fractional differencing parameter, for both returns and volatility, using the two semiparametric procedures described in Section 5.2 and the fully parametric AR(1)-LMSV model.²⁰ Finally, a set of specification tests of the AR(1)-LMSV model is presented.

5.4.1 Testing for Long Range Dependence

In Table 5.3 we present the test results for the raw exchange rate returns of the DEM/USD, JPY/USD and JPY/DEM. As indicated in Section 5.2.4 the test statistic (\hat{S}) proposed by Robinson (1996) has a limiting standard normal distribution under the specified null hypothesis, implying a two-sided rejection region of 2.575 at 1%.²¹ Our null hypotheses are formulated as follows. First, standard efficient markets theory indicates that asset prices should follow a random walk, implying that returns should be $I(0)$. This defines one hypothesis as $H_0 : d = 0$. Second, we employ the theoretical bounds for stationarity and invertibility of the fractionally integrated representation for returns as hypotheses, yielding $H_0 : d = -\frac{1}{2}$ and $H_0 : d = \frac{1}{2}$.

Results yield the following observations. For all three currencies one can strongly

²⁰Whilst conducting this study we were made aware of similar research, on the long memory features of intra-day DEM/USD volatility, being carried out by Torben Andersen and Tim Bollerslev. Their results, now published as Andersen and Bollerslev (1997a) support those presented here.

²¹As also indicated in Section 5.2.4, the test is fully parametric and based on an underlying ARFIMA representation for y_t . We conducted these tests using various different ARFIMA specifications, finding that the results were robust. Those presented in this section are based on an ARFIMA(1,d,0) model.

reject the hypothesis that $d = 0.5$ in favour of $d < 0.5$ implying that returns are covariance stationary. Similarly one can conclude that the return processes are invertible given the sign and magnitude of the test statistics corresponding to the hypothesis that $d = -0.5$. A more interesting result appears when examining the second column of Table 5.3. Whereas for the JPY/DEM the test statistic indicates that one cannot reject the hypothesis of FX *quotations* following an I(1) process, for the DEM/USD and JPY/USD there is evidence that the degree of fractional integration in returns is negative. This then implies that the degree of integration for the *quotation* series of these two currencies is between one half and unity, such that these rates are non-stationary but not I(1). The implication of this result is that exchange rates may be mean reverting over extremely long horizons, rather than completely persistent as the usual random walk representation would suggest.

Tables 5.4 and 5.5 present the results from the same testing framework on raw and deseasonalised FX volatility (computed, as indicated in Sections 5.2 and 5.3 as $\log(r_t^2)$.) As an indication of the efficacy of our deseasonalisation procedure, in Figure 5.4 we present the periodogram of deseasonalised JPY/USD volatility. Comparison with Figure 5.1 demonstrates that the amplitudes at the seasonal frequencies are greatly reduced, although not completely eliminated.²²

The hypotheses of interest in our examination of volatility are as follows. First, given the many previous studies which have demonstrated that intra-day FX volatility has an IGARCH or almost-integrated SV representation, is there a random walk in volatility? Second, can one characterise volatility as being long range dependent? The former dictates examination of $H_0 : d = 1$ whilst the hypotheses pertinent to the latter are $H_0 : d = 0$ and $H_0 : d = \frac{1}{2}$.

²²As an alternative to our deseasonalisation procedure, we also computed all estimations for the volatility of a time-scale transformed series of midquotes. We used the theta-time scale proposed by Dacorogna, Müller, Nagler, Olsen, and Pictet (1993). Results from these estimations were very similar to those for deseasonalised volatility and are available upon request from the authors.

The final columns of Tables 5.4 and 5.5 present the evidence relevant to the hypothesis of a random walk in volatility. There is strong evidence that, for all currencies, this hypothesis can be strongly refuted in favour of a degree of integration in volatility of less than unity. Further, this conclusion is stable across both raw and deseasonalised volatility. There is still, however, the possibility of non-stationarity in volatility if $d \in [0.5, 1]$. Column 2 of the tables indicates that the non-stationarity hypothesis can be rejected also, with the test statistics indicating that $d < 0.5$ for all three currencies. Finally, evidence on the long range dependence of volatility is shown in the first column. From both tables one can draw the conclusion that the true value of d lies between zero and one half, evidence of long memory in the volatility of all currencies.

Hence, the testing procedure indicates the following. Returns can be characterised as short range dependent, covariance stationary processes, with some indication of negative degrees of fractional integration for the DEM/USD and JPY/USD. The volatility results indicate covariance stationarity also, with the non-stationarity and $I(1)$ hypotheses convincingly refuted, although there is consistent evidence of long range dependence. The volatility processes are therefore completely mean reverting, a result which is more comfortable than that of $I(1)$ volatility from a theoretical point of view.

5.4.2 Semiparametric Estimations

The testing procedure contained in Section 5.4.1 indicates that the fractional differencing parameter is between positive and negative one half for both exchange rate returns and the instantaneous volatility measure. This implies that the semiparametric estimates can be applied to the series directly to give robust point estimates of the fractional differencing parameter. We first present estimates for the three

returns series. The semiparametric procedures of Robinson (1995a) and Geweke and Porter-Hudak (1983) are employed in estimation with Table 5.6 giving the results.

Examining first the GPH estimates it is quite clear that the negativity of d indicated in the previous subsection is a very minor economic phenomenon, indicating that returns display very small anti-persistent tendencies. For no currency does \hat{d} exceed 0.05 in absolute value. This implies that the *quotation* series may be regarded as following I(1) processes to more-or-less any degree of precision, in line with the efficient markets hypothesis.

This conclusion becomes less clear when one examines the Gaussian estimates. Whilst the results for the JPY/USD and JPY/DEM are very similar to their GPH counterparts, the value of \hat{d} derived for the DEM/USD is now greatly negative. This implies a covariance-stationary and invertible representation for DEM/USD returns which displays non-negligible anti-persistence. Given the confluence between the testing and GPH results, however, we are inclined to treat this feature as an anomaly and describe the return generating process as approximately I(0)

The estimation results for the volatility series are presented (for raw and deseasonalised volatility) in Tables 5.7 and 5.8. Here we complement the semiparametric procedures used in the analysis of returns with the fully parametric AR(1)-LMSV model.

The results of both tables demonstrate that the testing procedures contained in the previous subsection deliver the correct inference regarding the value of d . Across currencies and estimators there is consistent evidence that the value of \hat{d} for the volatility series is between 0.2 and 0.3. This indicates that volatility can be char-

acterised as covariance-stationary, invertible and long range dependent.²³ The only real difference in estimation results for raw and deseasonalised volatility is that the AR(1)-LMSV estimates tend to be slightly greater for the former and greater than the results delivered by the semiparametric estimators. This is due to misspecification of the short-range dependence in the series i.e. omission of an explicit seasonal in the fully parametric model. As one might expect, the difference in estimated \hat{d} between the semi and fully parametric procedures is far smaller for deseasonalised volatility.

5.4.3 Specification Tests on the Fully Parametric Model

The final step in our empirical methodology involves a series of estimations and specification tests on the fully parametric model outlined in Section 5.2.3. These tests allow us to examine the relative contributions of short memory and long memory components to the temporal dependence in the volatility process. As indicated in Section 5.2.3 our fully parametric model nests a pure LMSV model (obtained by setting $\phi = 0$ in equation 5.9) and a standard AR(1)-SV model (obtained by restricting $d = 0$ in 5.9.) By estimating the unrestricted model and these two restricted alternatives we can employ a frequency-domain Quasi-Likelihood Ratio test to gauge the significance of the long memory and AR(1) components.

It is important, at this point, to note the confluence between the estimates of d from the fully and semi parametric estimates presented in the previous subsection. This demonstrates a lack of systematic bias in the fully parametric estimations so that the extended Harvey (1993) model can be regarded as well specified.²⁴ Hence we can

²³These conclusions also hold when examining instantaneous volatility derived from exchange rates sampled at intra-day frequencies other than the 10 minute sampling employed here.

²⁴Note also that the performance of the fully parametric model vindicates our use of an underlying ARFIMA(1, d , 0) structure in the test procedure contained in Section 5.4.1.

proceed to compare the unrestricted fully parametric model with the two restricted versions which permit, in the first case, long memory structure only and, secondly, a simple AR(1) underlying volatility structure. Results of these estimations and the associated tests, for deseasonalised volatility only, are given in Table 5.9.

Examining first the results which correspond to the simple AR(1)-SV models one can note the appearance of the common result that the 'underlying' volatility process has an autoregressive parameter very close to unity in all cases. This conforms with the results of many earlier studies. Moving on to the pure LMSV models it is quite clear that the conclusions of the previous estimations still hold and further that the long memory specification gives a far better fit than does the autoregression (as evinced by the lower minimised log likelihood.)

A comparison with the unrestricted model yields the following observations. First, the improvement in fit of the combined model over the pure LMSV model is marginal, as the comparison of Log Likelihoods demonstrates. Second, and more importantly, the value of the autoregressive parameter is far lower in the unrestricted model, in all cases between 0.1 and 0.2. Given the pure AR(1)-SV results, this may be taken as evidence that the previous findings of near unit-root behaviour in intra-day FX volatility are caused by model mis-specification through the omission of long-range dependent components. This result confirms the evidence from daily data found in Baillie, Bollerslev, and Mikkelsen (1996). These conclusions are reinforced by the quasi-LR test statistics. In the unrestricted model one can convincingly reject the hypothesis that the degree of fractional integration in volatility is zero for all currencies, whereas the hypothesis that the autoregressive parameter is zero cannot be rejected in two of the three cases and is only marginally rejected in the third.

A clear comparison of the AR-SV and LM-SV models, in terms of how well they fit the data, is shown in Figure 5.5. The figure graphs the actual correlogram of

the deseasonalised volatility process for the DEM/USD alongside those implied by the estimated AR-SV and LM-SV models. It is immediately apparent that the long memory specification gives a far better approximation of the true volatility dynamics than does the autoregressive model. The autoregressive model greatly overstates the low order autocorrelations but dies out too quickly to mimic the persistently positive high-order autocorrelations in the DEM/USD data. Intuitively, the estimated autoregressive parameter is driven very close to unity in order to try to approximate the long memory in volatility, but this only results in the low order autocorrelations being far too high whilst the exponential decay in the correlogram ensures that, even if the estimated autoregressive parameter is arbitrarily close to unity, the autoregressive specification cannot match the persistence in volatility exhibited by the data.

In summary, the empirical work delivers the following findings. First there is evidence that exchange rate returns are $I(0)$ processes such that, as efficient markets theory would suggest, exchange rate levels are $I(1)$. An examination of instantaneous exchange rate volatility, however, provides strong evidence of long memory and covariance stationarity. The finding of stationarity is at odds with the results of Baillie, Bollerslev, and Mikkelsen (1996) and Harvey (1993), both of which present estimates of d which imply non-stationarity in daily volatility, and also with the outcomes of numerous GARCH and AR(1)-SV estimations for intra-day exchange rate volatility. Further, the now standard IGARCH or integrated SV result can be seen to be due to mis-specification of the conditional variance model through the omission of long memory structure.

5.5 Conclusions

We have presented in this work a study of the long range dependent properties of three major exchange rates sampled at an intra-day frequency. Evidence on the returns series tends to imply that they are $I(0)$ processes, a result which supports the $I(1)$ hypothesis for exchange rates. We however do not reject strongly the hypothesis of antipersistence ($I(d)$ characteristics with $d < 0$) in the returns, which would imply mean reversion (however slow it may be) in the exchange rate series. In contrast, testing and estimation results for volatility are consistent across currencies and give a clear indication that the volatility of exchange rates series follows a covariance stationary long range dependent process. Further results suggest that the commonly found near unit-root behaviour of intra-day FX volatility is driven by the mis-specification of traditional models which do not allow for long range dependence. The existence of long range dependence in intra-day volatility is in line with similar findings for daily exchange rate volatility found in Baillie, Bollerslev, and Mikkelsen (1996) and Harvey (1993). A point to note, however, is that whereas our results indicate a covariance stationary volatility process ($0 < \hat{d} < 0.5$), those of both Baillie, Bollerslev, and Mikkelsen (1996) and Harvey (1993) indicate that \hat{d} exceeds 0.5 and, therefore, volatility is non-stationary.

As far as methodology is concerned, the present paper has emphasized the performance of several individual econometric methods in the treatment of a stationary long memory component in a time series. It needs to be stressed that a practitioner should also rely on the strength of the sequential methodology. The pretesting procedure enabled us to identify the range in which the dependence parameter lies, in this case $0 < d < 1/2$, and therefore to identify the basic characteristics of the process within the chosen framework, namely its long range dependence, covariance stationarity and mean reversion. Robust estimation of d is then carried out to identify the main feature of interest in the data, namely the degree of fractional

integration. The robust estimates can then be used as a benchmark against which one can assess the adequacy of a fully parametric model involving both long and short memory structure. If the fully parametric specification is shown to be satisfactory, one can proceed to assess the relative importance of the long memory and short memory components for overall dependence in the process.

In future research we intend to address the following issues. First, we have shown that the volatilities of all three exchange rates have very similar long memory properties. An interesting examination would therefore be to examine whether the long memory structure has a single, common source. This would entail assessing the adequacy of a single long memory factor structure for the vector of volatilities. Also, in Section 5.1, we motivated the possibility of long memory in volatility with a variant of the mixture of distributions model of Tauchen and Pitts (1983). The long memory structure in volatility might then imply similar features for data on volumes. This is an issue we also intend to address.

TABLES

Table 5.1: Summary Statistics for Exchange Rate Returns

Rate	Mean	s.d.	Skew	Kurtosis	ρ_1	ρ_2	ρ_3	$Q(10)$
DEM/USD	6×10^{-6}	0.001	0.16	9.61	-0.076	-0.040	-0.005	306.8
JPY/USD	-4×10^{-4}	0.074	-0.06	13.85	-0.09	-0.015	0.0035	334.8
JPY/DEM	-5×10^{-4}	0.045	-0.25	7.93	0.0066	-0.0004	0.0009	13.5

Notes: the coefficients ρ_1 , ρ_2 and ρ_3 represent the first through third sample autocorrelations respectively. The $Q(10)$ statistic is the Box-Ljung test statistic for up to tenth order serial correlation. The Box-Ljung statistic is distributed χ^2_{10} and has critical value 23.2 at 1%.

Table 5.2: Summary Statistics for the Logarithm of Squared Returns

Rate	Mean	s.d.	Skew	Kurtosis	ρ_1	ρ_2	ρ_3	$Q(10)$
DEM/USD	-14.67	1.28	0.828	0.072	0.281	0.244	0.220	15948.03
JPY/USD	-7.68	3.36	-1.313	1.24	0.24	0.183	0.166	8433.88
JPY/DEM	-8.66	3.16	-0.99	0.45	0.353	0.29	0.264	20463.21

Notes: the coefficients ρ_1 , ρ_2 and ρ_3 represent the first through third sample autocorrelations respectively. The $Q(10)$ statistic is the Box-Ljung test statistic for up to tenth order serial correlation. The Box-Ljung statistic is distributed χ^2_{10} and has critical value 23.2 at 1%.

Table 5.3: Tests of the significance of the Fractional Differencing Parameter for Raw Exchange Rate Returns

Rate	d=-0.5	d=0	d=0.5
DEM/USD	113.48	-5.02	-18.51
JPY/USD	112.55	-4.68	-19.71
JPY/DEM	122.09	0.46	-18.02

Note: Testing of the value of the fractional differencing parameter is carried out via the test procedure developed in Robinson (1996).

FIGURES

Figure 5.1: Periodogram for JPY/USD Log Squared Returns
Ordinates 1 to 1000

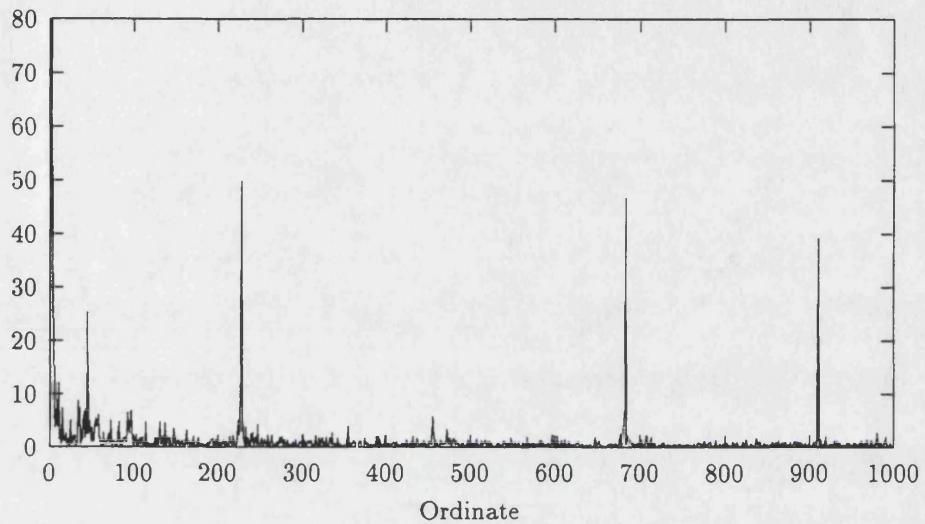


Figure 5.2: Log Periodogram for JPY/USD Log Squared Returns
Ordinates 1 to 1000




Table 5.4: Tests of the significance of the Fractional Differencing Parameter for Exchange Rate Volatility ($\log(r^2)$)

Rate	d=0	d=0.5	d=1
DEM/USD	33.23	-13.15	-20.53
JPY/USD	21.12	-14.05	-20.70
JPY/DEM	31.06	-12.40	-20.38

Note: Testing of the value of the fractional differencing parameter is carried out via the test procedure developed in Robinson (1996).

Table 5.5: Tests of the significance of the Fractional Differencing Parameter for Deseasonalised Exchange Rate Volatility

Rate	d=0	d=0.5	d=1
DEM/USD	22.38	-14.26	-20.62
JPY/USD	18.54	-14.63	-20.28
JPY/DEM	24.26	-13.65	-20.55

Note: Testing of the value of the fractional differencing parameter is carried out via the test procedure developed in Robinson (1996).

Table 5.6: Estimation of the Fractional Differencing Parameter (\hat{d}) for Exchange Rate Returns

Rate	\hat{d}	
	Rob95	GPH92
DEM/USD	-0.38 (0.02)	-0.01 (0.05)
JPY/USD	-0.02 (0.02)	-0.03 (0.05)
JPY/DEM	-0.01 (0.02)	-0.01 (0.05)

Note: Estimation of the long memory models is carried out via the semiparametric procedures of Robinson (1995) (Rob95) and Geweke and Porter-Hudak (1992) (GPH92). Standard errors in parentheses.

Figure 5.3: JPY/USD Log Squared Returns: Sample Autocorrelations 1 to 1000
Horizontal line is 95% confidence interval for White Noise

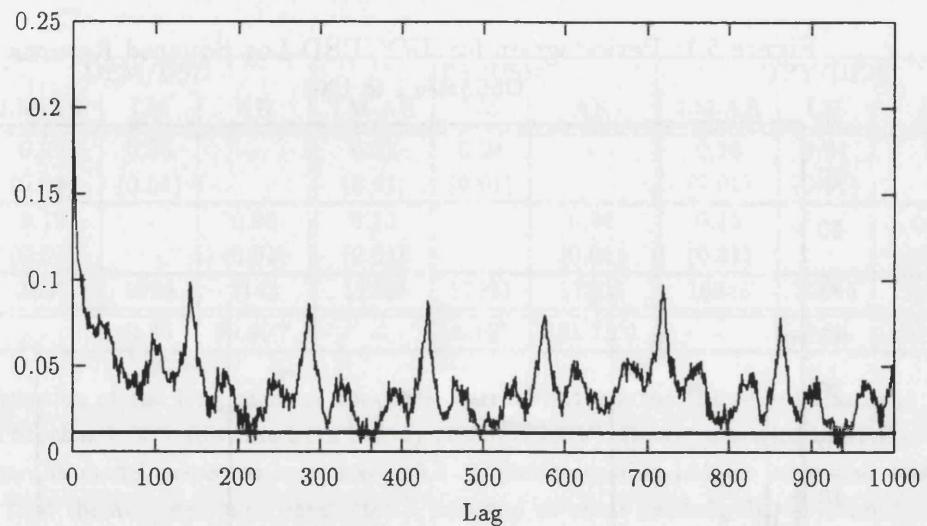


Figure 5.4: Periodogram for Deseasonalised JPY/USD Log Squared Returns
Ordinates 1 to 1000




Table 5.7: Estimation of the Fractional Differencing Parameter (\hat{d}) for Exchange Rate Volatility

Rate	\hat{d}		
	Rob95	GPH92	LMSV
DEM/USD	0.29 (0.02)	0.21 (0.06)	0.37 (0.01)
JPY/USD	0.27 (0.02)	0.19 (0.06)	0.24 (0.01)
JPY/DEM	0.27 (0.02)	0.28 (0.06)	0.32 (0.01)

Note: Estimation of the long memory models is carried out via the semiparametric procedures of Robinson (1995) (Rob95) and Geweke and Porter-Hudak (1992) (GPH92), plus the fully parametric Long Memory in Stochastic Volatility model of Harvey (1993) (LMSV). Standard errors in parentheses.

Table 5.8: Estimation of the Fractional Differencing Parameter (\hat{d}) for Deseasonalised Exchange Rate Volatility

Rate	\hat{d}		
	Rob95	GPH92	LMSV
DEM/USD	0.29 (0.02)	0.19 (0.04)	0.26 (0.01)
JPY/USD	0.30 (0.02)	0.18 (0.04)	0.22 (0.01)
JPY/DEM	0.30 (0.02)	0.25 (0.04)	0.26 (0.01)

Note: Estimation of the long memory models is carried out via the semiparametric procedures of Robinson (1995) (Rob95) and Geweke and Porter-Hudak (1992) (GPH92), plus the fully parametric Long Memory in Stochastic Volatility model of Harvey (1993) (LMSV). Deseasonalisation is carried out via a frequency domain Double-Window Smoother. Standard errors in parentheses.

Figure 5.5: Comparison of Actual and Implied Correlograms for DEM/USD

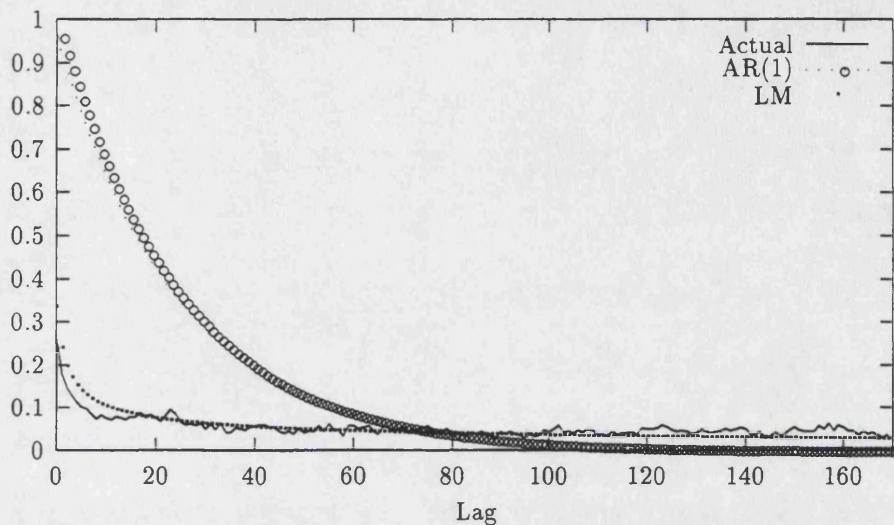


Table 5.9: Extensions of various LMSV specifications and Likelihood Ratio Testing using Deseasonalised Volatility

Rate	DEM/USD			JPY/USD			JPY/DEM		
	LM-AR	LM	AR	LM-AR	LM	AR	LM-AR	LM	AR
d	0.26 (0.01)	0.25 (0.01)	-	0.22 (0.01)	0.24 (0.01)	-	0.26 (0.01)	0.24 (0.01)	-
ϕ	0.12 (0.01)	-	0.96 (0.01)	0.15 (0.01)	-	0.96 (0.01)	0.15 (0.01)	-	0.93 (0.01)
LogL	3098	3098	3143	17229	17231	17305	15646	15646	15726
LR	-	0.85	90.90**	-	4.10*	151.75**	-	0.88	161.50**

Note: Estimation of the long memory models is carried out via the the fully parametric Long Memory in Stochastic Volatility model of Harvey (1993) (LMSV). Columns headed LM-AR present results from the model presented in equation 5.9. Columns headed LM are estimated with the restriction that the autoregressive parameter is zero and columns headed AR are estimates from the model where the fractional differencing parameter is set to zero. The final row of the table gives Likelihood Ratio Statistics relevant to the omission of the given parameter. * denotes the test is significant at 5%, ** denotes significance at 1%. Deseasonalisation is carried out via a frequency domain Double-Window Smoother . Standard errors in parentheses.

FIGURES

Figure 5.1: Periodogram for JPY/USD Log Squared Returns
Ordinates 1 to 1000

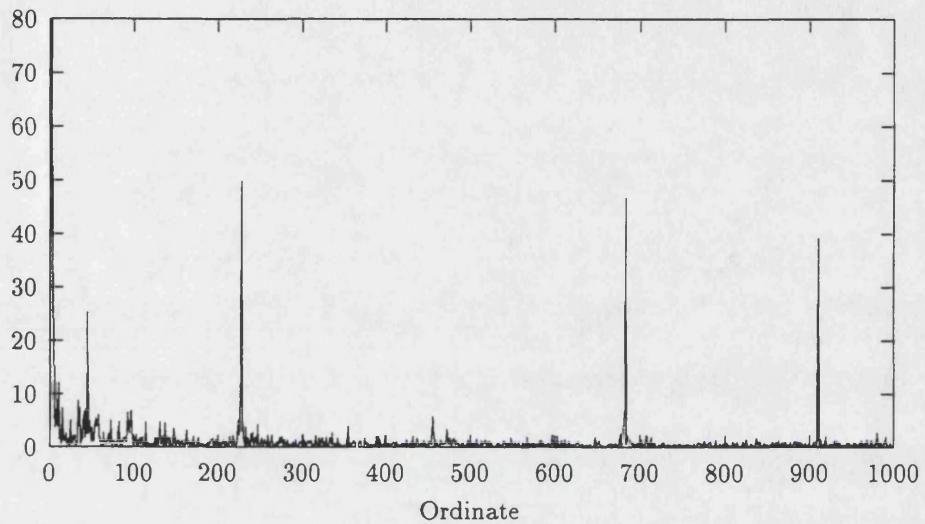


Figure 5.2: Log Periodogram for JPY/USD Log Squared Returns
Ordinates 1 to 1000

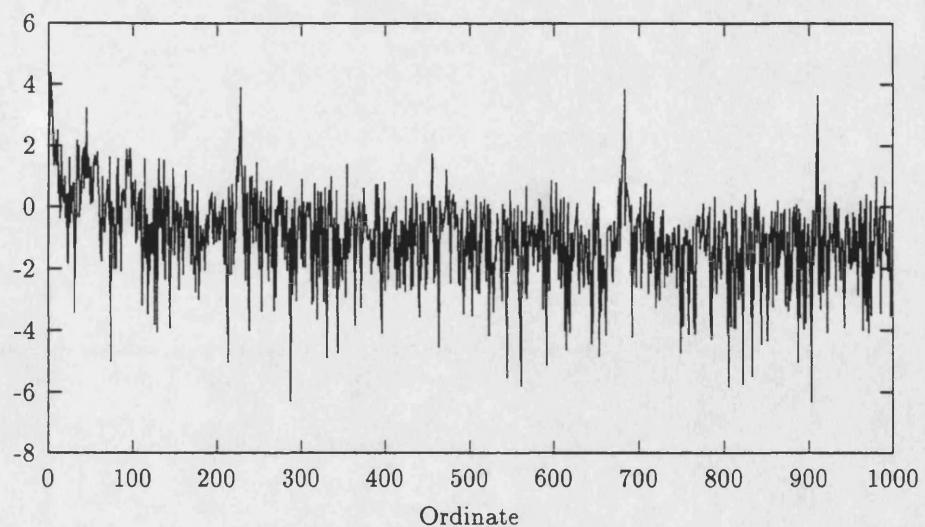


Figure 5.3: JPY/USD Log Squared Returns: Sample Autocorrelations 1 to 1000
Horizontal line is 95% confidence interval for White Noise





Figure 5.4: Periodogram for Deseasonalised JPY/USD Log Squared Returns
Ordinates 1 to 1000




Chapter 6

The Effects of Macroeconomic 'News' on High Frequency Exchange Rate Behaviour

6.1 Introduction

This paper studies the high-frequency reaction of the DEM/USD exchange rate to macroeconomic information emanating from Germany and the U.S. Specifically, we utilise exchange rate data covering the period 1/1/92 to 31/12/94, sampled at a five minute frequency, to investigate how the major monthly macroeconomic releases from these two countries affect the level of the exchange rate. The information contained in announcements over this three year period is extracted via a set of

Figure 5.5: Comparison of Actual and Implied Correlograms for DEM/USD



market expectation series supplied by Money Market Services International (MMS).¹ Our analysis improves on previous work in this area in two main respects. First, our study is conducted using very high-frequency data, whereas most earlier work has used exchange rate data sampled at a frequency of a number of hours or more. This allows us to construct a very precise characterisation of the reaction of the exchange rate to macroeconomic information. Second, this is, to our knowledge, the first study which includes German data releases. Most work in this area has focused on U.S. (and to a lesser extent Japanese) macroeconomic announcements.

The major issue under examination is whether one can trace systematic effects of economic 'news' on the evolution of the exchange rate. The direction in which 'news' will push the exchange rate is, however, *a priori* unknown. It will depend on the market's belief about both the model of exchange rate determination and the manner in which monetary authorities respond to new information. Take an unexpected increase in U.S. real activity for example. A Monetarist model of exchange rate determination would imply that the dollar should appreciate as domestic money demand rises, whilst a Keynesian model would predict the opposite due to increased import demand by U.S. citizens. These two mechanisms do not, however, factor in the possibility of a reaction to the 'news' by the U.S. monetary authority. Assuming the Fed to have a preference for low inflation, they should raise short term interest rates in order to cool the inflationary pressures in the economy, implying U.S. dollar appreciation. We label this latter mechanism the 'reaction function' effect. Which of the above mechanisms pertain to the determination of the DEM/USD and which dominates will determine the sign of the exchange rate response to 'news'.²

The impact of macroeconomic 'news' on exchange rates has previously been the sub-

¹We would like to thank MMS International in London and Belmont, CA, for the provision of the expectations data and Olsen and Associates in Zurich for providing the exchange rate data.

²See Hoffman and Schlagenhauf (1985) for a more complete treatment of the alternative theories of exchange rate determination and an empirical test of these theories.

Chapter 6

The Effects of Macroeconomic 'News' on High Frequency Exchange Rate Behaviour

6.1 Introduction

This paper studies the high-frequency reaction of the DEM/USD exchange rate to macroeconomic information emanating from Germany and the U.S. Specifically, we utilise exchange rate data covering the period 1/1/92 to 31/12/94, sampled at a five minute frequency, to investigate how the major monthly macroeconomic releases from these two countries affect the level of the exchange rate. The information contained in announcements over this three year period is extracted via a set of

market expectation series supplied by Money Market Services International (MMS).¹ Our analysis improves on previous work in this area in two main respects. First, our study is conducted using very high-frequency data, whereas most earlier work has used exchange rate data sampled at a frequency of a number of hours or more. This allows us to construct a very precise characterisation of the reaction of the exchange rate to macroeconomic information. Second, this is, to our knowledge, the first study which includes German data releases. Most work in this area has focused on U.S. (and to a lesser extent Japanese) macroeconomic announcements.

The major issue under examination is whether one can trace systematic effects of economic 'news' on the evolution of the exchange rate. The direction in which 'news' will push the exchange rate is, however, *a priori* unknown. It will depend on the market's belief about both the model of exchange rate determination and the manner in which monetary authorities respond to new information. Take an unexpected increase in U.S. real activity for example. A Monetarist model of exchange rate determination would imply that the dollar should appreciate as domestic money demand rises, whilst a Keynesian model would predict the opposite due to increased import demand by U.S. citizens. These two mechanisms do not, however, factor in the possibility of a reaction to the 'news' by the U.S. monetary authority. Assuming the Fed to have a preference for low inflation, they should raise short term interest rates in order to cool the inflationary pressures in the economy, implying U.S. dollar appreciation. We label this latter mechanism the 'reaction function' effect. Which of the above mechanisms pertain to the determination of the DEM/USD and which dominates will determine the sign of the exchange rate response to 'news'.²

The impact of macroeconomic 'news' on exchange rates has previously been the sub-

¹We would like to thank MMS International in London and Belmont, CA, for the provision of the expectations data and Olsen and Associates in Zurich for providing the exchange rate data.

²See Hoffman and Schlagenhauf (1985) for a more complete treatment of the alternative theories of exchange rate determination and an empirical test of these theories.

ject of extensive research, but most of the work has concentrated on announcements of U.S. economic data. Research based on several USD exchange rates sampled at a daily frequency (Hardouvelis (1988), Aggarwal and Schirm (1992), Harris (1995a), Edison (1997)) finds significant positive relationships between dollar appreciation and US M1 and non-farm employment (and in some cases also the merchandise trade balance) news, but no significant impact of any other macroeconomic news. Other papers utilizing spot quotations from the opening and closing of the main regional FX markets (these being North America, the Pacific, Tokyo and Europe), reach similar conclusions. Hakkio and Pearce (1985), Ito and Roley (1987), Hogan, Melvin, and Roberts (1991) and Hogan and Melvin (1994) demonstrate that USD exchange rates respond quite rapidly to US money supply and trade balance surprises, but not to other types of US news. Ito and Roley (1987) also find that the JPY/USD does not respond to macroeconomic news from Japan. Hence, a general result from these studies is that few economic announcements have systematic impacts on exchange rates when rates are sampled at a frequency of a number of hours or more. We conjecture that other announcements may have discernible impacts on exchange rates when examined in a higher frequency setting, with the disappearance of the effects at lower frequencies due to their being drowned in subsequent exchange rate fluctuations. By examining the impact of U.S. and German news on exchange rate returns measured over different time horizons (from 5 minutes to 12 hours), we are able to test this conjecture.³

We extend our analysis by also considering a number of other issues. Following Hakkio and Pearce (1985) we examine the efficiency of the intra-day FX market via

³Previous research on the reaction of exchange rate volatility to economic announcements (Ederington and Lee (1993), Chapter 4 of this thesis, Almeida (1997)) has shown that large increases in volatility are apparent at the times of these announcements, but most of the increased volatility is dispersed in around 10 minutes. Using high frequency data from the foreign currency futures market, Ederington and Lee (1995) also demonstrate that transaction prices react very swiftly (i.e. within a couple of minutes) to the announcement of U.S. economic data. This evidence suggests that the impact of macroeconomic announcements may only be identified using exchange rate returns derived from very high-frequency data.

the traditional equation linking exchange rate changes to anticipated and unanticipated macroeconomic data. We also examine the temporal stability of our results, since most of the previous literature finds that there are significant structural breaks in the response of exchange rates to news (usually associated with shifts in the Fed's policy).

Lastly, the use of German announcement data allows us to address two further interesting questions. This is due to the fact that, unlike U.S. announcements, German releases do not have pre-set and pre-advertised release dates and times. Hence first, we can examine how pre-scheduling of announcements affects the impact of 'news' on the DEM/USD by comparing the dynamic responses to U.S. and German data. A hypothesis relevant here is that the response to scheduled announcements is completed more quickly than that associated with a non-scheduled release. Using the German data we can also examine how the proximity of the announcement to other events, specifically the biweekly Bundesbank council meeting, affects the reaction of the exchange rate. We hypothesise that announcements which occur closer to Bundesbank council meetings have a larger exchange rate impact due to the belief that they will carry more weight in Bundesbank policy deliberations.

The structure of the paper is as follows. Section 6.2 describes our data and methodology. We then turn to a discussion of our empirical findings, first for US news, Section 6.3, and then German news, Section 6.4. We have a Section reporting some comparative statistics, Section 6.5, and the paper closes with Conclusions and ideas for further study.

6.2 Data and Methodology

6.2.1 The Data

Our exchange rate data covers the period 1/1/92-31/12/94. The data were originally received as an irregularly spaced, continuous-time set of DEM/USD quotations, published on the screens of Reuters' information system. We then converted the data to an equally spaced, calendar time-series by imposing a five minute observation grid, taking the last quotation in a five minute period as effective.⁴ Finally, the average of bid and ask quotations was taken as our basic quotation variable.

The other portion of our data set consists of US and German macroeconomic announcements covering the same period, plus a market expectation series for each type of announcement obtained from Money Market Services International.⁵ The list of series employed in our 'news' analysis are as follows: for the US we use the Employment Report, Trade figures, PPI and CPI announcements, Retail Sales, Durable Goods Orders, Consumer Confidence figures, Leading Indicators, the NAPM survey and the Industrial Production and Capacity Utilisation announcements; for Germany the series used are the CPI, Industrial Output, M3, Industrial Orders, PPI, Retail Sales, the Trade Balance, Unemployment and the Wholesale Price Index. Appendix 1 gives a list of the definitions of these announcements along with their usual release timing and the identifiers assigned to each series. Since all of these announcements are monthly, for each series we have a maximum of 36 observations, although irregularity in some of the releases and missing expectations data reduce the number of available observations in some cases.

⁴At points where there was no observation in a five minute interval, we linearly interpolated between the nearest preceding and succeeding quotation.

⁵To determine the time of the German announcements, a data set consisting of news headlines published on the screens of Reuters' information systems, with the corresponding date and time stamp, was also used.

The MMS expectations we employ are calculated as the median from a survey of forecasts made by leading practitioners and academics. These data then allow us to separate the series relevant to each type of announcement into an unexpected and expected portion. Efficient markets theory would imply that only the unexpected part of an announcement should have any impact on the DEM/USD, the expected portion having already been impounded into quotations. The adequacy of the MMS expectations series is an issue which has been examined in previous work. We re-examine this issue in Appendix 2. Utilising a simple test of rational expectations our results suggest that, overall, the MMS expectations series are unbiased. Some indications of systematic biases in expectations are present but these are relatively scarce. Hence, whilst other authors have cast doubt on the rationality of the MMS expectations series (notably Aggarwal, Mohanty, and Song (1995)) our analysis supports the use of these data, in line with the results of Pearce and Roley (1985).

6.2.2 Methodology

Define q_t to be our 5 minute quotation series for the DEM/USD. Define $x_{i,t}$ to be the actual announced value for series i at moment t , $x_{i,t}^e$ to be the corresponding expected value from the MMS data, and $x_{i,t}^{ne} = x_{i,t} - x_{i,t}^e$ the unexpected part of the announcement. The basic equation which underlies most of our empirical analysis, equation 6.1, is derived from the efficient markets hypothesis,

$$r_{i,t+k} = \alpha + \beta_1 x_{i,t}^{ne} + \beta_2 x_{i,t}^e + u_{i,t} \quad (6.1)$$

where $r_{i,t+k}$ is defined as $q_{t+k} - q_t$ and $u_{i,t}$ is an error term. Note that estimations of the above equation are not time-series regressions as the observations are not tem-

porally consecutive. An observation for a given series is added at every point when a new release of data occurs and the associated return is then constructed. Market efficiency would dictate that the expected portion of an announcement should have no impact on the return and, further, that the constant term be zero. Impacts from the 'news' contained in announcements on the evolution of the DEM/USD will show up as significant coefficients on the unexpected portion of the release i.e. β_1 being significantly different from zero.

Equation 6.1 is utilized in the empirical work in the following ways. First we test for systematic effects of each type of 'news' on the DEM/USD. As mentioned above, this entails testing the difference of β_1 from zero. However, as argued in Section 6.1, the signs of the coefficients are a priori uncertain, depending on which of the fundamental or reaction function responses dominates for a given release. To conduct these tests we initially assume market efficiency and restrict the parameters α and β_2 to zero i.e., we estimate equation 6.2 for each series i ,⁶

$$r_{i,t+k} = \beta_1 x_{i,t}^{ne} + u_{i,t} \quad (6.2)$$

We employ the return in the fifteen minutes immediately post-announcement as the dependent variable (i.e. we took $k=15m$), the fifteen minute window chosen with reference to earlier work (Ederington and Lee (1993), Ederington and Lee (1995) and Chapter 4) which suggests that the major impacts on both prices and volatility occur within a fifteen minute span post-announcement. The results of these tests are described in Sections 6.3.1 and 6.4.1. Next we explicitly test the efficiency of the FX market. As argued above this entails estimating equations of the type of 6.1, and demonstrating that the coefficients α and β_2 are not significantly different

⁶All the tests in this paper were computed using White (1980) heteroskedasticity-consistent standard errors for the OLS estimates.

from zero.⁷ ⁸

Our next estimations concern the persistence of the impact from announcements on the DEM/USD. If the ‘news’ contained in these macroeconomic data is considered ‘fundamental’ for the determination of the DEM/USD exchange rate then there should not only be a significant impact in the fifteen minutes post-release, but also, for example, in the twelve hours immediately after announcements. This is tested in Sections 6.3.3 and 6.4.3 by varying the window over which the return variable is calculated, i.e., estimating equations similar to 6.2, but with k assuming different values. The minimum return considered is five minutes and the maximum is twelve hours. Another test of persistence is to ask whether one can predict the sign of, for example, the twelve hour return from the sign of the fifteen minute return i.e. does the almost immediate reaction of the exchange rate to news dominate the direction in which the exchange rate moves in the twelve hours immediately post announcement. In order to examine this we employ a sign test based on the binomial distribution.⁹

In the exercises above, we have assumed that the response to macroeconomic ‘news’ is invariant over time. In Sections 6.3.4 and 6.4.4 we examine whether the market reaction to unexpected information varies across our sample. To this end we employ three dummy variables: the first taking the value unity only in the first year of

⁷Note that if the forecast series for any announcement proved to be $I(1)$, the forecasts were first differenced before inclusion in this regression.

⁸Another exercise we performed was to examine the effects of the announced data on returns prior to the official release time in order to test for the possibility of information leakage. For both the U.S. and German data these results indicated that the effect of ‘news’ on pre-announcement returns was largely insignificant. Further, these results for Germany vindicate our use of the Reuters reporting time as the time at which the ‘news’ hits the market.

⁹This test is based on the comparison of the sign of the returns in the 15 minutes and in a longer period (we used 6, 12 or 24 hours) after the announcement. The test statistic is the percentage of returns of the same sign in both periods. Under the null hypothesis of the initial return having no influence on the long term, this percentage is 50%. Using the binomial distribution, we can compute the probability of the number of predictions being equal or greater than the observed value under the null hypothesis the true distribution is binomial with $p=0.5$. If this probability is lower than the usual significance levels, one may reject the null hypothesis of the immediate return having no influence over the longer term returns.

our sample and zero otherwise, the second being unity in the second year of the sample and zero everywhere else and the final dummy taking the value unity in the last sample year only. These dummies were then interacted with the forecast error series, giving three regressors in the basic 'news' regressions, $x92_{i,t}^{ne}$, $x93_{i,t}^{ne}$ and $x94_{i,t}^{ne}$. The test consisted of the estimation of equation 6.3 for each series i ,

$$r_{i,t+k} = \beta_2 x92_{i,t}^{ne} + \beta_3 x93_{i,t}^{ne} + \beta_4 x94_{i,t}^{ne} + u_{i,t} \quad (6.3)$$

Significant differences in the coefficients in this regression will indicate that the response to 'news' is not invariant over the entire three year announcement sample.

6.3 Analysis of American Macroeconomic Announcements

In this section we present the results from testing the impact of U.S. data releases on the DEM/USD exchange rate. As detailed in Section 6.2.1, our analysis concentrates on the following monthly U.S. announcements: the statistics contained in the Employment Report, the Mercantile Trade Report (Trade), PPI and CPI figures, Retail Sales (RS), Durable Goods Orders (DG), Consumer Confidence (CC), Leading Indicators (LI), the NAPM survey and finally the Industrial Production/Capacity Utilisation figures (IP/CU).¹⁰

¹⁰Note that in our analysis of individual announcements we can decompose the Employment Report responses into the effects due to its two principal components, the unemployment rate (Unemp) and payroll employment figures (PAY), via their forecast errors. For the analysis of this pair of announcements, and for the simultaneously released IP and CU figures, we use a multiple regression version of equation 6.2.

6.3.1 The Impact of ‘News’

The first question we address in our empirical analysis of the US data is do markets react systematically to good (and bad) news about the state of the US economy? Given the list of announcements above, one would expect the coefficients on the forecast error series for all real activity indicators to have the same sign aside from that on the unemployment series. The actual sign which the coefficients take will depend on which of the fundamental or reaction effects outlined in Section 6.1 dominate. *A priori*, the signs of the price indicator series and their relationship with those of the real indicators is unknown. The results from estimations of equation 6.2 are given in Table 6.7.

A first point to note is that the coefficients on all indicators have the sign predicted by the reaction function response to news. Unexpected Retail Sales growth, for example, entails an appreciation in the dollar whereas unexpected Unemployment shocks have the opposite effect. The coefficients on announcements of real activity also conform with a Monetarist model of exchange rates. However, there is marginal evidence that price shocks tend to cause dollar appreciations also, an effect which is inconsistent with the Monetarist model. Hence our preferred interpretation of these results is that the response of exchange rates to macroeconomic ‘news’ is indicative of expected Federal Reserve reactions in domestic money markets.

Moving to the inference on these coefficients the *t*-statistics demonstrate that most are significant at, at least, 10%. Exceptions to this are the coefficients associated with PPI, CPI, Leading Indicators, Industrial Production and Capacity Utilisation figures. Whilst one might expect these results for the final three of this group, due to their being fairly unimportant indicators, the lack of impact from the price series is a notable result. The poor results for the price series may, however, be due to the relative inadequacy of the forecast series for these announcements (see Appendix 2).

In the final column of Table 6.7 we give a common scale to the results by forming the product of the estimated coefficient and the average absolute forecast error. These figures yield a ranking of announcements, in terms of the mean impact on the DEM/USD, as follows. The Payroll Employment data is clearly most influential, entailing in excess of a 30 b.p. revision on average. There is then a group of indicators, comprising the Unemployment Rate, Trade figures, Retail Sales, Durable Goods orders, Consumer Confidence and the NAPM survey, which give, on average, at least a 10 b.p. impulse to the DEM/USD. A low impact group of indicators includes the PPI, CPI, LI, IP and CU.¹¹

The dominance of the Employment Report statistics links with the comments in Harris (1995b) who expounds the view that they are regarded as the key indicator of U.S. performance by the markets. Also, the influence of the Trade figures on exchange rates is unsurprising. Again, the presence of the price indicators in the low impact group is surprising.

6.3.2 Market efficiency

We now move on to analysing the efficiency of the DEM/USD market around these announcements. The test of efficiency we employ is a simple regression, of the type of equation 6.1. In terms of the significance of the coefficients on forecast errors, similar results and a similar ranking of announcements to that in the previous sub-section obtain. Efficient markets would imply that the coefficients on the constant terms and forecast values should be zero. Table 6.8 is generally supportive of efficiency. Only one of the constants, that in the Consumer Confidence equation, is significantly different from zero. Further, for 10 of the 12 releases the forecast data

¹¹These results are also broadly consistent with the scaled impact of absolute forecast errors on the volatility of the DEM/USD presented in Table 4.7 of Chapter 4.

has no significant impact on the return. For the 2 releases for which the forecasts are significant (PPI and DG), the coefficients are negative. These significant negative coefficients may be linked to the results on the adequacy of the expectations series contained in Appendix 2. As Table 6.5 shows, the two announcements mentioned above are those for which the slope coefficients in the unbiasedness regressions are greatest in absolute value and both coefficients are significantly above unity. This implies that the MMS expectations we employ tend to consistently overpredict the actual data releases. Given this fact, the negative coefficients on the forecast series derived from our market efficiency testing can be rationalised as the market discounting the overprediction of the MMS expectations.

6.3.3 The Persistence of U.S. ‘news’ effects

Thus far we have examined the effect of the information revealed by announcements on the DEM/USD in the fifteen minutes immediately post-release, finding that the major U.S. macroeconomic announcements have significant impacts. In this subsection we test the persistence of the effect of ‘news’ on the exchange rate by examining the impact of unanticipated information on exchange rate returns measured over various intervals post-announcement.

As Table 6.9 demonstrates, the general picture that emerges is that the impact of these macroeconomic releases on the DEM/USD is a very short run phenomenon. Examining first the group of seven announcements which were found to be fairly influential in the previous subsection, the following results emerge.¹² An encouraging result is that of the 77 regressions over differing horizons, more or less all yield a correctly signed coefficient (only 4 of the 77 coefficients are ‘wrongly’ signed).¹³

¹²These seven being the Employment statistics, Trade figures, RS, CC, DG and NAPM figures.

¹³By ‘correctly’ signed we mean that the signs of the coefficients agree with those from Table 6.7 and our inference that the Fed follows ‘reaction function’ policy in response to ‘news’.

The pattern of significance, however, is not strong. In general the impacts are significant only up to around 2 hours after release, after which, it seems, the effect of unanticipated macroeconomic information is drowned in the subsequent random fluctuations of the exchange rate. The only exceptions to this picture are the patterns for the Payroll employment and Consumer Confidence figures. These retain significance until the 12 hour horizon, confirming the earlier results and, with regard to the former, in line with the earlier remarks that these figures are by far the most influential in the market.

The persistence patterns associated with the less important announcements are, unsurprisingly, even less impressive. Only 29 of the 55 estimated coefficients are correctly signed and there is no consistent impact of 'news' on returns from any of the individual announcements.

Hence our results suggest that, aside from the Payroll and CC figures, one can trace very little long term effect of 'news' on quotations. This seems to indicate that, although one can find very short term impacts from forecast errors on the DEM/USD, the longer run impact is minimal, with the reaction to 'news' being drowned in subsequent 'noise'.

An alternative way to examine this issue, however, is the following. Can one predict the direction of the exchange rate change over the day following an announcement, for example, from the direction of the 15 minute return? In order to examine this question we employ the sign test outlined in Section 6.2.2. The results of this testing procedure are given in Table 6.10. Note that we cannot now separate the effects of the Payroll and Unemployment figures, and the IP and CU figures due to their simultaneous release.

Whereas Table 6.9 demonstrates that, over longer horizons, one cannot predict the

magnitude of the revision using the forecast error, Table 6.10 indicates that it is possible to predict the sign of six, twelve and twenty four hour exchange rate changes based on the sign of the immediate response. Examining each release separately, one can see that for the majority of announcements at least two of the three p -values are less than 10%, indicating predictability of the longer return direction given the direction of the immediate return. The strongest predictabilities in sign are associated with the Employment statistics, RS, DG, CC, and NAPM figures. Surprisingly poor results, indicating a lack of sign predictability at all horizons, were obtained for the Trade figures. For the 6 hour horizon, p -values of greater than 20% were obtained for PPI, CPI, LI and IP/CU, although results improved over twelve and twenty four hour horizons for the PPI and IP/CU figures. In order to increase the power of the test, we also aggregated across announcements and re-computed the test statistics, the results reinforcing the general predictability of the direction of returns.

Hence to conclude this subsection it seems that the unexpected proportion of most macroeconomic announcements has a very short-run impact on exchange rates. Forecast errors can be seen to help predict exchange rate returns over horizons of around one hour rather than days or weeks. There is, however, a statistically significant relationship between the direction of the exchange rate change in the 12, or even 24, hours following announcement and the change in the fifteen minutes immediately after release, as evinced in the sign tests above.

6.3.4 Structural stability

Finally we perform the structural stability tests described in Section 6.2.2. In examining these results we split the announcements into two groups, the first consisting of those series with significant results from the basic ‘news’ regression and the second

containing all other series. With regard to the first group, all coefficients are of the sign implied by the reaction function response to news, in line with the results of Table 6.7. A regularity in the results, excepting the Unemployment and CC coefficients, is that the absolute value and significance of the coefficients declines over the three years, suggesting a greater impact of these news variables in 1992 relative to 1994. This diminution of the size of the coefficients is significant for the Payroll and Trade data and especially marked for the durable goods news. A possible explanation for this is the fact that global FX markets were far more turbulent in 1992, perhaps leading to greater emphasis being placed on macroeconomic announcements than in the relatively quiet times of 1994. Overall, for these announcements there are strong patterns in significance and sign which confirm their, at least short-term, importance in FX markets.

The picture for the latter group is quite different. Only the coefficients associated with the CPI figures are all of the sign predicted by the reaction function response. Results for the LI, IP and CU news, in particular, are very erratic. One point to note is the correct sign and significance of the PPI news in 1992 and 1993, coefficients which are significantly greater than their 1994 counterpart. This shows that the previous poor showing of PPI forecast errors as a source of news may be attributed to their behaviour in 1994, where an insignificant negative coefficient obtains.¹⁴

Hence, in line with earlier results, there is robust evidence that news associated with Unemployment, Payroll employment, Trade, Retail Sales, Durable Goods orders, Consumer Confidence and the NAPM survey have a strong short-term impact on the DEM/USD spot rate. The direction of these impacts corresponds to a reaction function response to 'news'.

¹⁴Further examination reveals that the results for the PPI data are greatly affected by a single observation in which a positive forecast error coincided with a very large negative DEM/USD return. Omission of this one observation results in the PPI forecast errors having a significantly positive effect on the subsequent fifteen minute return for the 1992 to 1994 period.

6.4 Analysis of German Macroeconomic Announcements

In this section we present the results obtained for the German macroeconomic announcements. Our analysis concentrates on the following monthly German macroeconomic series: Consumer Prices (CPI), Industrial Production (IP), Money supply M3, Manufacturing ORDERS, Producer Prices (PPI), Retail Sales (RS), TRADE Balance, UNEMployment and Wholesale Prices (WPI).

6.4.1 The Impact of ‘News’

The Explanatory power of ‘News’

The results of estimations of equation 6.2 for the German announcement series are given in Table 6.12. Note that a positive coefficient in these regressions implies that a positive surprise in the announcement is associated with a depreciation of the DEM (i.e. the USD appreciates). Of the 9 series analysed, only 3 (UNEM, IP and CPI) have a significant impact at the 10% significance level. The R^2 is small in all the regressions, including those with significant coefficients. These results suggest that German macroeconomic announcements have low explanatory power for exchange rate changes in the 15 minutes after their release. Note that in some cases (for instance, TRADE) the low explanatory power may be a consequence of the bias in the expectations series.

In the final column of Table 6.12 we give a common scale to the results by forming the product of the estimated coefficient and the average absolute forecast error. These figures show that IP and UNEM are the series causing the largest response (each announcement entails a 3 b.p. revision, on average), followed by M3, CPI

and RS (which entail a 2 b.p. revision, on average). The average response yields a ranking of the impact of the announcements which is similar, but not identical, to the ranking provided by the significance levels. For instance, the significance level of CPI is higher than for M3, but its average response is smaller, because the CPI is forecast more accurately than the other series, and thus the 'news' content of a CPI announcement is smaller on average.

The impact of 'News' and Bundesbank council meetings

Previous work on the reaction of DEM/USD exchange rate volatility to central bank actions (Almeida (1997)), has shown that the DEM/USD exchange rate is very sensitive to statements issued after Bundesbank council meetings. These results suggest the possibility that Bundesbank decisions, and expectations about those decisions, are the key German factor driving the DEM/USD exchange rate. If this is true, then the importance of German macroeconomic announcements should be a function of the extent to which they influence the outcome of the Bundesbank council meetings. Our hypothesis is that this is a function of the period of time between the release and the meeting: if an announcement is made just after a Bundesbank council meeting (long before the next) the markets will not pay much attention to it; the Bundesbank will not act on that information for two weeks,¹⁵ and during that period new information may arrive that could be more relevant for their decision.

¹⁵The Bundesbank meets every other Thursday, with some exceptions: it did not meet over the Christmas or Easter holidays, or in the beginning of August, so there was a 3 or 4 weeks interval between meetings at those times; also, on the 4 times when the Thursday was a holiday, the meeting was held on the Wednesday (3 times), or the Friday (once). As such, there is a weekly regularity to Bundesbank meetings. The macroeconomic data are announced monthly, but there is no precise monthly regularity to these announcements. The combination of these two facts implies that for a particular series, the proximity to the next Bundesbank meeting varies considerably across announcements. Overall the average number of days to the next Bundesbank meeting is 7.9, with a standard deviation of 5.7, and these results hold roughly for each individual series. Note that in order not to have a number of days to the next announcement equal to zero, we included both the day of the announcement and the day of the meeting in the difference. Thus a figure of 1 means the announcement occurred on the day of a Bundesbank meeting, and each meeting is 15 days before the next.

This hypothesis implies that an announcement will be more important the closer it is to the next Bundesbank council meeting, so we weighted our observations according to this. The test consisted of running regressions similar to the ones used in subsection 6.4.1 but with each observation weighted by the inverse of the number of days to the next meeting, i.e. we ran cross-section regressions of the form

$$R_{t+k} = \delta X_{i,t}^{ne} + u_{i,t} \quad (6.4)$$

where $R_t = r_t/d$, $X_{i,t}^{ne} = x_{i,t}^{ne}/d$, and d is the number of days to the next Bundesbank council meeting (including the day of the announcement).

The results for these regressions, presented in Table 6.13, support our hypothesis, since there is an increase in the significance of the coefficient on ‘news’ (and in the R^2) for most of the announcements (the only exceptions being WPI and ORDERS, but neither of these are significant). When the proximity to the Bundesbank council meetings is considered, 6 of our 9 series have coefficients significant at the 10% level. The scaled impact is, on average, doubled and at least for M3, UNEM, IP and PPI, the R^2 is relatively high, suggesting the German macroeconomic announcements may indeed have an important impact on the behaviour of the exchange rate, at least over the short term.

The interpretation of the sign of the coefficients

As described in Section 6.1, the sign of the coefficient on ‘news’ depends on the market’s belief about both the appropriate model of exchange rate determination and the likely reaction of the monetary authorities. In this subsection, we try to identify those beliefs based on the results presented in Tables 6.12 and 6.13. Note that the

signs of the coefficients are the same in both Tables (although the significance of the coefficients changes when we weight by the proximity to Bundesbank council meetings): negative for IP, M3, PPI, RS, TRADE and UNEM (i.e. a positive surprise in these announcements causes a DEM appreciation), positive for CPI, ORDERS and WPI.¹⁶ As explained below, these results suggest that the markets believe the DEM behaves according to a model where international capital flows dominate trade flows, i.e. the key variable for exchange rate determination is the interest rate differential, and where the monetary authorities set interest rates according to their expectations of future inflation (and eventually, growth).

The negative sign on M3 is not consistent with the direct effects of any of the 'fundamentals' models, but is consistent with a 'reaction function' interpretation: faced with higher than expected monetary growth, the Bundesbank will raise interest rates to reduce M3 growth; higher interest rates will cause an appreciation of the DEM. The signs on IP, PPI and RS are also consistent with a 'reaction function' interpretation:¹⁷ higher than expected producer prices or real activity may be seen as indicators of future inflation; to curb these inflationary pressures, the Bundesbank will raise interest rates, causing an appreciation of the DEM. However, the signs on CPI and UNEM are not consistent with the 'reaction function' interpretation, but they are consistent with the Keynesian model: lower unemployment will increase consumer spending and increase imports; higher domestic prices should also increase imports and reduce exports; in both cases, the balance of payments will deteriorate, causing a depreciation of the DEM.

The Bundesbank is usually described as a 'money targeting' central bank. Each year the Bundesbank council sets a (band) target for M3 growth, and monetary

¹⁶In the following analysis, we will ignore the variables that do not have significant coefficients in any of the Tables 6.12 and 6.13, i.e., ORDERS, TRADE and WPI.

¹⁷Note that the negative sign on these three variables is not consistent with the direct effects in the Keynesian model, but the signs on IP and RS (although not on PPI) are consistent with the direct effects in the Monetarist model.

policy actions are justified with reference to the attainment of this target. A strict interpretation of this framework would imply that the Bundesbank would only react to unexpected changes in M3, and not to any other information. However, some authors claim that in practice the Bundesbank does not follow money targets strictly. Von Hagen (1995) claims that the Bundesbank follows an inflation target framework, with expected future inflation being the main factor driving policy actions, and M3 growth just an (important) indicator of future inflation. Clarida and Gertler (1996) estimate a policy reaction function for the Bundesbank, and conclude that German monetary policy is conditioned on inflationary pressures and the state of real economy, but since the Bundesbank behaves in a forward-looking manner, it only reacts to future (consumer price) inflation and growth. In particular, they show that the Bundesbank reacts first of all to M3 shocks, also to past growth as measured by the IP, or to past commodity price shocks, but not to past consumer price inflation, i.e., the Bundesbank reacts to announcements of variables that are leading indicators of future inflation or of current or future growth, but not to indicators of current consumer price inflation. Although the variables used by Clarida and Gertler (1996) do not match exactly the variables used in this paper, their results suggest that the Bundesbank should react to unexpected changes in M3, IP, RS (as indicators of current and future growth), or PPI (as proxy for commodity prices), but not to unexpected changes in CPI or UNEM.¹⁸ Our results are consistent with this.¹⁹

¹⁸Note that Clarida and Gertler (1996) do not include an unemployment variable in their study. However, current unemployment may be seen more as a consequence of low growth in the past, than as an indicator of future growth, i.e., it is more a backward-looking variable that should not enter the reaction function of a forward-looking central bank.

¹⁹This interpretation implies that some of the variables used in this paper should be good predictors of future inflation (M3, IP, PPI and RS), and others not (CPI and UNEM). As a rough test of this hypothesis, we ran a regression of the change in the CPI from t to $t+12m$ on the changes from $t-12m$ to t in the Lombard rate and the 9 macroeconomic series used in this paper, using monthly data from 1993:3 to 1994:12. Results demonstrated that the signs of coefficients were consistent with a reaction function interpretation of the results of Table 6.12, but the significance of the coefficients on most of the variables was poor. Due to the lack of degrees of freedom in this regression and the simplistic structure imposed, we do not report the results, leaving a more careful and detailed analysis for further work.

This interpretation is also consistent with the results of the weighting procedure. Scaling by the proximity to the next Bundesbank council meeting will only make a significant difference if the market believes that the Bundesbank will react to the information released, i.e., if the reaction effect is dominant. The significance and magnitude of the coefficients (and also the R^2) for M3, IP, PPI, and RS are all increased, the biggest increase being for M3, the main variable in the Bundesbank's reaction function (according to Clarida and Gertler (1996)). However, although there is some increase in the significance of both coefficients, the effects seem to be quantitatively smaller for the CPI and UNEM and some measures of the importance of the effects actually deteriorate (the size of the coefficient for UNEM and the R^2 for CPI). These results suggest that M3, IP, PPI, and RS affect the DEM/USD exchange rate through their effects on Bundesbank policy decisions, whereas the effect of CPI and UNEM may be felt through other channels.

Since the 'reaction function' effect dominates for the majority of the variables, it seems that expectations about future interest rates are the main factor driving the DEM/USD reaction to German macroeconomic news. The effect of backward-looking variables (like the CPI and UNEM) suggest that the markets see the DEM behaving according to a model in which the elasticity of the exchange rate to the current account is small relative to the elasticity of the exchange rate to interest rate differentials.

6.4.2 Market efficiency

Table 6.14 presents the results from running the market efficiency regressions (equation 6.1). The coefficients on the forecast errors (and their significance) are not significantly affected by the use of equation 6.1 instead of equation 6.2, as may be seen by comparing Table 6.14 with Table 6.12. With regard to the coefficients on

constant terms and forecast values (expectations), the results are generally supportive of the efficiency hypothesis, i.e., in general these coefficients are not significantly different from zero. Only one of the coefficients on the expectation (CPI), and one of the constants (M3) are significantly different from zero at the 5% level. This does not seem to constitute enough evidence to reject the market efficiency hypothesis, and thus, supports our choice of equation 6.2 as the basis for the analysis in the previous section.

6.4.3 The Persistence of German ‘news’ effects

In this section, we extend our analysis to assess if and how the effect of German macroeconomic announcements is felt over longer periods of time using the regression framework and non-parametric sign tests as in Section 6.3.3.

Table 6.15 presents the results from estimations of equation 6.2 (the ‘original’ data) using returns over different periods of time, from 5 minutes to 12 hours, while Table 6.16 presents similar results for regressions of the type of equation 6.4 (the ‘weighted’ data). Figure 6.3 plots the group average.

For periods longer than the 15 minute interval already considered (see Tables 6.12 and 6.13,) the significance levels tend to fall, although for the ‘weighted’ data the significance levels are relatively high even after 3 hours. For the ‘original’ data only PPI has consistently significant coefficients up to 3 hours, and of the 3 series identified as significant in section 6.4.1, only UNEM extends its significance to the 30 minute interval. For the ‘weighted’ data the significant effects extend up to 6 hours for some series. However, the full impact of the announcement (measured by the highest value of the coefficient) is, on average, only felt after 3 hours, for the ‘original’ and ‘weighted’ data, although the pattern is more clearly defined for

the latter. The relatively long period that the market takes to adjust is probably associated with the fact that the German announcements are not scheduled. This also explains the relatively small scaled effect found in Section 6.4.1. If one takes the 3 hour instead of the 15 minute return, then the impact of German announcements on the DEM/USD exchange rate becomes higher, as shown in Table 6.17. For the 'original' data, the scaled response reaches an exceptional 19 b.p. for PPI, but is lower than 7 b.p. for the other series. For most of the series, the scaled response is even higher for the 'weighted' data, reaching 16 b.p. for TRADE and CPI, 12 b.p. for M3 and 11 b.p. for IP.

Table 6.18 reports the results of the sign test described in Section 6.2.2. For most of the series, the direction of the 15 minute returns seems to have some predictive power, although weak, over the direction of the longer horizon exchange rate movements. The number of predictions is generally above 50%, but only in 5 cases is it high enough for the null to be rejected at the 5% significance level. We should expect the predictive power to be declining over time, i.e. to be greater for the 6 hour period than for the 12 or 24 hour period. Given this assumption, the only consistent and significant results are for ORDERS, since we reject the null hypothesis at 10% significance level for the 6 h and 12 h periods, and for M3 over 6 h, where we reject the null at the 5% significance level. However, the tests for individual announcements are not very powerful, since they are based only on 36 observations.

One of the advantages of the sign test is that it allows for aggregation across announcements. Thus we construct tests based on a larger number of observations and on a broader class of series than in the individual series tests. We aggregate over 4 (overlapping) groups of series (see Table 6.18), defined using the results from the news test or on the basis of the nature of the macroeconomic variables involved: the first with all 9 series, the second ('Top 6') comprises the 6 series significant at the 10% level in the 'weighted' regressions (see Table 6.13), the third ('Prices') with the 3 price indices and finally, the 4 series describing the real sector ('Real'). Given the

large number of observations, and the diversity of the announcements, the test for the first group is the one with most power. For this group the results suggest the sign of the 15 minute return has a significant influence on the sign of the 6 h and 12 h returns, and a smaller but non negligible influence on the sign of the 24 hour return. The significance levels are even higher for the 'Top 6' and 'Real' groups, with the German macroeconomic announcements having a significant influence even over 24 h. It seems that even though the impact of the German macroeconomic announcements on the DEM/USD rate is quantitatively small, it is sufficiently important to be one of the major factors influencing the direction of exchange rate movements on the day of the announcements.

As would be expected, the relative importance of the German announcements is higher in the 6 hours following the announcement, than over 12 or 24 hours. Since most of the announcements occur between 7:00 and 13:00 GMT, we can interpret the 6 hour period as corresponding (roughly) to European trading, and the 12 hour period to European and American trading. Then we can interpret the results from this subsection as suggesting that German macroeconomic announcements have a significant effect over European trading on the day of their release, but that this effect is dwarfed by the impact of the arrival of information from the US when the American dealers enter the market, and thus gradually ignored by market participants in the American and Asian markets.

6.4.4 Structural stability

The results of the structural stability tests for German data are presented in Table 6.19, for the regressions using the original values, and Table 6.20, for the regressions using the values weighted by the proximity to the following Bundesbank council meeting. The main feature of the results in Tables 6.19 and 6.20 is the consistency

in the signs of the coefficients across time. For all of the series with significant coefficients in Tables 6.12 and 6.13 (CPI, IP and UNEM in both tables, M3, PPI and RS only in Table 6.13), the sign of the coefficient is the same in the 3 years. However, the size of the coefficient varies across the sample, even though the differences across periods are seldom statistically significant. Given the small number of observations in each subsample, it is not surprising that the significance levels vary considerably across years. The consistency of the signs of the coefficients strongly suggests that the results are not spurious, and that there is strong evidence that some German macroeconomic announcements have a significant (although small) impact on the DEM/USD exchange rate.

6.5 Are the Impacts of German and US Announcements Similar?

In this section we compare the results obtained for the German and US announcements and try to extract some general conclusions. The main feature of the results is that the FX markets' primary concern is with the future likely reaction of the monetary authorities, both in Germany and the USA. The results for Germany are somewhat mixed, since they suggest that the exchange rate reacts to news on CPI and UNEM according to the Keynesian model, but the effects of other variables follow the 'reaction function' hypothesis. However, the quantitative effects of the latter variables dominate the effects of the former, which lead us to prefer the reaction function hypothesis as the main force driving the DEM/USD reaction to German macroeconomic news. The reaction of the DEM/USD exchange rate to US macroeconomic announcements is less ambiguous as all series have the sign predicted by the reaction function hypothesis, with the announcements with the largest and most significant impact in the USD being related to the real economy, in particular

with payroll employment.

In both countries the implied reaction function has, however, some curious and interesting features. In Germany, for example, the Bundesbank has been usually described as basing its policy on monetary targets, whereas Von Hagen (1995) and Clarida and Gertler (1996) have recently argued that the Bundesbank seldom complies with its M3 targets, and actually reacts to divergences of inflation and output from their desired values, using a modified Taylor (1993) rule. Our results suggest that the FX market does not believe that the Bundesbank will only react to monetary shocks, but that it will also react to other macroeconomic variables. However, the market still sees the Bundesbank reacting primarily to monetary surprises. Does the former result mean that the market places more belief in the Bundesbank's rhetoric than is actually justified?

By contrast, in the US the FX market primarily reacts to unexpected shocks emanating from the real economy. How does this square with the greater weight which Central Banks, including the Fed, are supposedly now giving to the primacy of price stability as an objective? One, perhaps slightly cynical, answer is that the switch to awarding price stability much greater weight as an objective has been much more pronounced in Central Bank rhetoric than in their actions; there is an emerging academic literature, (Taylor (1993), Goodhart (1996), Almeida and Goodhart (1997), Muscatelli and Tirelli (1996)), documenting how much continuity there has been in Central Bank actions in recent years, and how little these have altered in response to supposed regime changes, e.g. Independence, Inflation Targets, etc. A kindlier interpretation is that the FX market's stronger reaction to (US) data on real shocks (than on price shocks) is that real shocks provide a better forward guide to future inflationary pressures than do price shocks (which may be more backwards looking).²⁰

²⁰One should also note that our sample period (1992-94) is mostly a period of low inflation and

If the main determinant of the foreign exchange market's response to news is their view of how the authorities will react to such news, as we suggest here, these such effects are likely to shift over time. Specifically, changes in the policy priorities of monetary authorities and the market's perceptions of those priorities will imply that the coefficients derived from our 'news' regressions will alter over time. The results from our structural stability regressions corroborate the above intuition. These results also suggest one reason why stable relationships between FX rates and economic 'news' have been hard to uncover. Moreover, as detailed earlier, most economic news items can be interpreted in different ways. A stronger real economy can, under the Keynesian model, be regarded as bad news for the exchange rate, since imports and future inflation will rise, but good news if the authorities react by raising domestic interest rates. We find that the reaction function effect tends to dominate but this does not imply that all market participants favor a reaction function interpretation. If some follow a Keynesian model, this will also weaken the discernible impact of 'news' on exchange rates.

Although most of the macroeconomic announcements considered in this paper have a significant (short term) impact on the DEM/USD exchange rate, this impact might be seen as quantitatively small. For the series with the largest impact, the US payroll figures, the exchange rate change caused by the average announcement is 31 b.p., a mere 0.2% change, and the impact of the other announcements is even smaller. Apparently, the DEM/USD is driven more by US than German announcements, since the impact of the former is much larger, even taking the results of the weighted regressions for Germany (which reflect a bigger impact for the German announcements). For the series with significant coefficients, the average revisions after 15 minutes caused by the US announcements are between 10 and 17 b.p, whereas for the German announcements the average revisions are between 2 and 6 b.p. How-

low growth in the USA. In such circumstances, primacy of price stability may be consistent with the Fed focusing on growth indicators: if inflation is thought to be under control, then the Fed will react mostly to the state of the real economy.

ever, this difference could be caused by the different time pattern of response for the US and German announcements:²¹ if instead of taking the 15 minute returns, one compares the periods with the maximum average revisions, then the impact for most of the German significant series is of similar magnitude (between 10 to 15 b.p.) as for the US series.

One of the major differences between US and German announcements of economic data is that the US data are announced at regular pre-arranged times, fixed to the minute, whereas German releases are unscheduled and irregular. Pre-arranged announcements are likely to encourage contingency planning beforehand, in terms of consultation with support staff, economists and analysts. In contrast, the variation of German announcement times will potentially discourage concentrated pre-planning of reactions. Hence it is plausible that there will be longer lags in the exchange rate assimilating German information, relative to that from the U.S.

Our results support this view. As shown in Figures 6.1 and 6.3, the maximum impact (as measured by the average scaled response across announcements) is reached after 15 minutes for the US announcements, but only after 3 hours for the German announcements. Note, however, that the evidence in Figures 6.2 and 6.4 suggests that the significance levels are maximum at the 15 minute period, for both the German and US announcements.²² Hence, market reaction can be seen to be concentrated in a short post-announcement period for the U.S., with the response to German 'news' being far more protracted. Like the German data, the majority of public news, e.g. on political events, market developments, arrives at unexpected times. The timing differences we found between responses to the US and German announcements suggest that the response of exchange rates to these events is likely

²¹As is described below, the maximum impact (measured by the size of the coefficient) of the German announcements is only reflected in the DEM/USD after 3 hours, whereas the maximum impact of the US announcements is felt after 15 minutes.

²²This vindicates our use of the 15 minute return as the basis for our analysis.

to be somewhat slower than earlier studies which concentrate on scheduled U.S. announcements have suggested, and hence there may be some profit opportunities available to those agents that are able to respond more quickly to such unpredictable events.

A priori, one might have thought that the data we examine incorporate a significant number of releases which are 'fundamental' to the DEM/USD. However, although there is a clear effect on very high frequency DEM/USD returns from most of our 'news' series, their influence on lower frequency returns is surprisingly weak. On average, U.S. 'news' is only significant for at most a couple of hours. German 'news' releases retain significance for a marginally longer period after which subsequent exchange rate fluctuations drown their effects. The sign tests performed indicate that, for both U.S. and German releases, the sign of the immediate response to news is correlated with the signs of longer horizon returns. However, although statistically significant, in purely numerical terms these results are not striking. Hence, our results suggest that, whilst announcements have a significant impact on short-run DEM/USD determination, they cannot be considered the key factors driving exchange rates.

Finally, in agreement with earlier research, our testing shows the FX market to be broadly efficient.²³ The expected portion of an announcement does not affect post-release exchange rate returns and, pre-release, there are no signs of 'news' affecting exchange rates. These results are in line with our expectations. The notion that macroeconomic information is leaked to or discovered by some traders before the official release time is, we believe, implausible. Although there may still be some sources of private information in the FX market, e.g. the order book of a bank with a large position in the market is private knowledge to that bank's traders,

²³There is now, however, an intriguing academic literature which suggests that following some forms of technical trading rule can generate excess profits (see, for example, Brock, Lakonishok, and LeBaron (1992) and Curcio, Goodhart, Guillaume, and Payne (1997).)

(see, for example, Lyons (1995), Perraудин and Vitale (1996), and Chapter 3 of this thesis), informational asymmetries based on these macroeconomic announcements are unlikely to arise.

6.6 Conclusions and Proposals for Further Work

In this paper we have studied the impact of macroeconomic news on exchange rates, using high frequency data for the DEM/USD. As described in Section 6.5, the overall picture is one of a strong, quick effect of macroeconomic ‘news’, with the exchange rate reflecting the anticipated policy reaction by the monetary authorities to the piece of news just released. However, this impact may be seen as quantitatively small, and the overall effect of the macroeconomic news on lower frequency exchange rate changes decays quite rapidly towards insignificance.

Although the main features of the ‘news’ effects are common to German and US announcements, there are some peculiarities and interesting features in the former group. First, the ‘news’ from German announcements tend to be incorporated in the exchange rate more slowly than the ‘news’ emanating from the US, due to differences in the timing arrangements. Second, the impact on the exchange rate is, on average, quantitatively smaller for the German announcements. Finally, the effect of the German announcements depends on the proximity to the next Bundesbank council meeting: when the observations were weighted by this proximity, the impact of the variables seen to be entering the Bundesbank’s reaction function (especially M3) was significantly increased.

We have argued above that the most sensible explanation of the set of coefficients showing the markets’ response to unexpected news is that these are, primarily, driven

by their interpretation of the monetary authorities' reactions. Those reactions, when triggered, affect short-term money market interest rates. Consequently the finding of a dominant response in a reaction function model would seem to suggest that exchange rates would in turn predominantly respond to, unexpected, changes in such interest rates. This latter relationship has not yet been intensively studied using high-frequency data. It is to this that we shall turn in future research.

Appendix 1

U.S. and German Macroeconomic Data

Table 6.1: U.S. Announcements

SYMBOL	ANNOUNCEMENT	REPORTED AS	TIME	WEEK	OBS.
CC	Consumer Confidence	level	10.00EST	4	36
CPI	Consumer Price Index	m/m % change	8.30EST	2/3	36
CU	Capacity Utilisation	%	9.15EST	3	36
DG	Durable Goods Orders	m/m % change	8.30EST	4	36
IP	Industrial Production	m/m % change	9.15EST	3	36
LI	Index of Leading Indicators	m/m % change	8.30EST	1	35
NAPM	U.S. N.A.P.M. survey	% level	10.00EST	1	36
PAY	Nonfarm Payrolls	thousands	8.30EST	1	36
PPI	Producer Price Index	m/m % change	8.30EST	2	36
RS	Advance Retail Sales	m/m % change	8.30EST	2	36
TRADE	Merchandise Trade Deficit	\$ billions	8.30EST	3	33
UNEM	Civilian Unemployment rate	m/m % change	8.30EST	1	36

Table 6.2: German Announcements

SYMBOL	ANNOUNCEMENT	REPORTED AS	WEEK	OBS.
CPI	Consumer Prices	m/m % change	4	36
IP	Industrial Production	m/m % change	1	35
M3	Money Stock M3	% change from Q4	4	33
ORDERS	Manufacturing Orders	m/m % change	2	35
PPI	Producer Prices	m/m % change	4	33
RS	Retail Sales, volume	y/y % change	3	34
TRADE	Trade Balance	DEM billions	-	29
UNEM	Number of unemployed, s/a	m/m change, thousands	1	35
WPI	Wholesale Price index	m/m % change	3	35

NOTES:

-week refers to the trading week of each month within which the given announcement is released usually e.g. in general, the US PPI figures are published in the second trading week of the month;

-the German data refer only to Western Germany (Federal Republic of Germany before the unification), except for M3 and TRADE which refer to unified Germany;

-the timing of the German announcements is not regular, but they usually occur between 7:00 and 13:00 GMT; in the absence of other sources of information, we take the time of a German announcement to be the time of its report by Reuters' news service; since the majority of the market participants receive their information through Reuters' or similar news services, it is reasonable to assume that the time of the Reuters' report is the time the information reaches the market; this would not be true if other news services reported the announcement before Reuters, but given the periodicity of our data (5 minutes) it is not likely that this is a serious problem;

-the regularity in the week of the month of the German releases is not precise; there is no regularity for the Trade Balance announcements, due to the fact that during our sample period the German Federal Statistics Office was implementing the transition from West German to unified German data; until the end of 1992 the emphasis was on data for Western Germany only, although data for unified Germany was also published, but with some delay; from January 1993, the emphasis shifted to unified German data, which was the first to be announced; this shift originated the irregular data releases;

-we have a maximum of 36 observations; however, as the publication of the US Mercantile Trade figure was discontinued in late 1994 there are only 33 data points for this announcement and within this 3 year span there are only 35 US LI announcements; also, the irregularity in the German TRADE releases (there are only 35 TRADE releases in our data period) and some missing expectations data reduce the number of usable observations for most of the German series.

Appendix 2

Testing the Rationality of Expectations

Testing is conducted in two phases. First, all series are examined for the presence of a unit root using the Augmented Dickey-Fuller methodology. If they appear to be covariance stationary then a simple regression of the time-series of announced figures on expected figures is run. Under the null of rationality of expectations the intercept in these regressions should be zero and the slope coefficient unity. If a unit root is found in the actuals and expectations then the cointegrating relationship is estimated via the Engle-Granger two-step procedure.²⁴ As above, the cointegrating vector on actuals and expectations should be (1,-1) with no constant present. To verify the existence of a cointegrating relationship the residuals from these regressions were checked for non-stationarity using the ADF test.

A2.1 Unit root tests

US announcements

Table 6.3 presents the results from ADF testing for the US announcements. The main conclusions from these tests are:

1. The PPI, CPI, RS, DG, LI and IP announcements and expectation series all appear to be covariance stationary; this is unsurprising as all are announced in terms of (percentage) changes of underlying series.
2. Unemployment, Payroll Employment, the Trade figures, the NAPM, Consumer Confidence and Capacity Utilisation all appear to have unit roots in both the actual and expected series;

²⁴The results from the two-step procedure were checked against those from the Engle-Yoo three-step procedure, demonstrating that the two-step estimators were sufficient.

this motivates the estimation of cointegrating relationships in order to test the rationality of expectations.

Table 6.3: Unit Root Tests - US announcements

Series	Forecast		Δ Forecast		Actual		Δ Actual	
	DF	Lags	DF	Lags	DF	Lags	DF	Lags
Pay	-2.47	0	-6.13 *	1	-4.38 *	0	-6.94 *	3
U	0.84	0	-5.81 *	1	0.91	0	-8.19 *	0
Trade	0.43	12	-3.26 *	8	0.11	8	-6.44 *	1
PPI	-5.35 *	0	-	-	-4.64 *	0	-	-
CPI	-3.97 *	0	-	-	-7.28 *	0	-	-
RS	-3.95 *	2	-	-	-7.62 *	0	-	-
DG	-6.96 *	0	-	-	-10.4 *	0	-	-
CC	-0.67	0	-3.99 *	1	-0.86	0	-7.56 *	0
LI	-3.09 *	1	-	-	-5.41 *	0	-	-
NAPM	-2.49	2	-3.60 *	7	-1.94	2	-8.22 *	0
IP	-5.10 *	0	-	-	-4.74 *	0	-	-
CU	0.38	0	-5.97 *	1	0.36	0	-8.47 *	0

Notes: *DF* is the value of the Dickey-Fuller statistic and *Lags* is the number of lags used in the ADF test. The null hypothesis for the Dickey-Fuller test is the existence of a unit root in the series under examination. The critical value for the *DF* tests is -2.62 at 1%. A * indicates that the null can be rejected at 1%.

German announcements

Table 6.4 displays the results of the unit root tests for the German announcements. The main features of Table 6.4 are:

1. The existence of a unit root, either in the actuals or the expectations, is rejected for the CPI, IP, ORDERS and PPI series.
2. The M3 and UNEM series appear to have unit roots in both the actual and the expectations; for these series, the second stage consists of the estimation of cointegrating relationships.
3. For the RS, TRADE and WPI series, the existence of a unit root is rejected for the actuals, but not for the expectations, at the 5% level (although it is rejected at the 10% level for RS and WPI); if the actuals and expectations are integrated of different orders, expectations are

not rational; however, the Dickey-Fuller methodology has been criticised for its lack of power, and thus it is possible that there is no unit root in the expectations of these series, even though we could not reject its existence; given that if the actuals and the expectations were integrated of different orders, the residuals of a regression of the actuals on the expectations would not be stationary, one can proceed assuming that both are stationary, regressing the actuals on the expectations, and test the stationarity of the residuals of such a regression; stationarity of the residuals would ensure that both series are $I(0)$.

Table 6.4: Unit Root Tests - German announcements

Series	Forecast		Δ Forecast		Actual		Δ Actual	
	DF	Lags	DF	Lags	DF	Lags	DF	Lags
CPI	-4.95 *	0	-	-	-4.74 *	0	-	-
IP	-4.07 *	0	-	-	-7.71 *	0	-	-
M3	-1.38	0	-3.60 *	0	-2.40	0	-5.95 *	0
ORDERS	-3.93 *	0	-	-	-6.73 *	0	-	-
PPI	-4.84 *	0	-	-	-4.40 *	0	-	-
RS	-2.88	1	-4.23 *	3	-5.76 *	1	-	-
TRADE	-1.37	0	-4.30 *	0	-3.26 *	0	-	-
UNEM	-0.58	0	-3.81 *	1	-2.49	0	-6.71 *	1
WPI	-2.80	0	-6.71 *	0	-4.43 *	0	-	-

Notes: DF is the value of the Dickey-Fuller statistic and $Lags$ is the number of lags used in the ADF test. The null hypothesis for the Dickey-Fuller test is the existence of a unit root in the series under examination. The critical value for the DF tests is -2.95 at 5%. A * indicates that the null can be rejected at 5%.

A2.2 Unbiasedness regressions

US announcements

The second stage of the test of the rationality of expectations consists of regressing the actuals on the expectations. For the US announcements, the following results emerge from Table 6.5:

1. Most of the regression results support the hypothesis of rationally formed expectations; the series which fall into this category are Unemployment, Payroll Employment, Trade, RS, CC,

NAPM, IP and CU; for all the constant is insignificantly different from zero and the slope not different from unity.

2. The coefficients on PPI, LI and DG expectations are significantly greater than unity whilst the constants are negative; this implies they tend to over predict the actual realisation.
3. The converse is true for the CPI expectations, with a slope coefficient much lower than unity and a positive intercept, implying under-prediction of the actual figures.
4. The DW statistics in all regressions indicate that residuals are serially uncorrelated, again as market efficiency would dictate.

Table 6.5: Unbiasedness Regressions - US announcements

Series	Constant		Expected		Residuals		R^2
	Coefficient	$t : a = 0$	Coefficient	$t : b = 1$	DW	DF test	
U	-0.54	-1.56	1.07	1.37	1.62	-5.15 *	0.93
Pay	-5.83	-0.16	1.03	0.14	2.46	-7.52 *	0.38
Trade	-1.03	-1.09	0.93	-0.60	2.18	-6.01 *	0.72
PPI	-0.19	-3.16 *	1.39	1.77	1.94	-	0.39
CPI	0.14	1.54	0.35	-2.03	2.57	-	0.03
RS	-0.02	-0.14	1.03	0.02	2.65	-	0.28
DG	-0.81	-1.85	1.84	2.47 *	2.65	-	0.40
CC	0.01	0.00	1.01	0.19	1.71	-5.02 *	0.88
LI	-0.02	-0.75	1.10	2.49 *	2.06	-	0.92
NAPM	6.09	1.19	0.88	1.33	2.23	-6.69 *	0.72
IP	0.01	0.12	1.13	1.12	1.69	-	0.81
CU	0.26	0.14	1.00	0.00	1.99	-5.80 *	0.98

Notes: The t -statistics for the constant terms are for the hypothesis that the intercept is zero, whilst those for the slope are relevant to the hypothesis that the slope is unity. The final three columns give the Durbin-Watson tests and Dickey-Fuller tests on residuals and the regression R^2 . A * indicates that one can reject the null hypothesis at the 5% level.

German announcements

The results of the second stage for the German announcements, presented in Table 6.6, show that:

1. Most of the series satisfy the rational expectations hypothesis restrictions of a zero intercept and a unit slope coefficient.

2. The only evidence against rationality are the intercepts in the RS and TRADE regressions; the negative and significant intercept for RS shows that the expectations tend to overestimate the actuals, while the opposite (underestimation of the actuals) occurs for TRADE.
3. The DW statistics and the DF tests indicate that the residuals are serially uncorrelated and stationary, as market efficiency would dictate.

Table 6.6: Unbiasedness Regressions - German announcements

Series	Constant		Expected		Residuals		R^2
	Coefficient	$t : a = 0$	Coefficient	$t : b = 1$	DW	DF test	
CPI	0.02	0.72	0.94	-0.68	1.91	-6.66 *	0.75
IP	-0.15	0.29	0.45	-1.16	2.53	-7.58 *	0.03
M3	-0.38	1.20	1.11	0.82	1.64	-4.51 *	0.68
ORDERS	0.18	0.40	1.32	0.59	2.14	-6.42 *	0.15
PPI	-0.01	-0.31	0.65	-1.37	2.12	-6.41 *	0.17
RS	-1.78	-2.19 *	0.49	-1.45	2.69	-8.75 *	0.06
TRADE	1.79	1.83 *	0.71	-1.36	1.99	-3.89 *	0.29
UNEM	1.32	0.24	0.92	-0.43	1.77	-4.92 *	0.44
WPI	-0.03	-0.53	0.96	-0.15	2.12	-6.14 *	0.27

Notes: See Table 6.5.

TABLES

Table 6.7: The impact of US News on DEM/USD returns

Series	Coeff.	T-stat	R^2	Scaled
PAY	0.00004	5.77	0.44	0.00310
DG	0.00090	5.39	0.44	0.00170
NAPM	0.00087	4.94	0.35	0.00140
RS	0.00385	3.48	0.23	0.00160
TRADE	0.00110	3.41	0.24	0.00130
CC	0.00030	2.74	0.20	0.00130
UNEM	-0.00730	-1.77	0.44	-0.00100
CPI	0.00450	1.48	0.06	0.00046
PPI	0.00290	1.29	0.04	0.00057
IP	0.00280	0.76	0.03	0.00037
CU	0.00046	0.28	0.03	0.00010
LI	0.00060	0.25	0.01	0.00007

Notes: The second column displays the estimated coefficients on the forecast error series. The next column gives the t-statistics relating to the hypothesis that the coefficients are zero. The critical values for the t-statistics are 2.04 at 5% and 1.70 at 10%. The final column displays the product of the coefficient with the average absolute forecast error.

Table 6.8: Market Efficiency tests - US data

Series	Constant	T-Const.	News	T-news	Expected	T-Exp.	R^2
CC	-0.00116	-2.34	0.00031	3.22	0.00001	0.15	0.26
CPI	0.00006	0.03	0.00523	1.82	0.00136	0.22	0.08
CU	-0.00042	-0.57	0.00172	1.31	0.00070	0.94	0.11
DG	0.00029	0.51	0.00103	6.05	-0.00092	-2.07	0.52
IP	-0.00042	-0.57	0.00275	0.82	-0.00227	-1.59	0.11
LI	-0.00021	-0.58	0.00077	0.30	-0.00015	-0.18	0.01
NAPM	-0.00018	-0.47	0.00090	5.58	0.00028	1.61	0.41
PAY	-0.00004	-0.05	0.00003	5.34	-0.00001	-0.75	0.47
PPI	0.00023	0.34	0.00096	0.50	-0.01125	-3.22	0.22
RS	0.00029	0.30	0.00392	3.61	-0.00151	-0.95	0.24
TRADE	0.00011	0.21	0.00110	2.79	0.00017	0.22	0.24
UNEM	-0.00004	-0.05	-0.00963	-2.89	-0.00144	-0.28	0.47

Notes: The column denoted 'News' displays the estimated coefficients on the forecast error series. The column denoted 'Expected' displays the coefficients on the forecast series. Note that if the forecast data for a given series contained a unit root the forecasts were differenced before inclusion in this regression. The 'T' columns give the t-statistics relating to the hypothesis that the coefficients are zero. The critical values for the t-statistics are 2.04 at 5% and 1.70 at 10%.

Table 6.9: The persistence of the effect of US News on the DEM/USD

Release	5min	15min	30min	45min	1h	1.5h
CC	0.00008*	0.00029*	0.00027*	0.00029*	0.00036*	0.00034*
CPI	0.00290*	0.0045	-0.0016	-0.0024	-0.00088	-0.002
CU	0.00016	0.00046	0.00028	0.0016	0.0011	-0.00021
DG	0.00065*	0.00091*	0.00083*	0.00075*	0.00077*	0.00089*
IP	0.00058	0.0029	-0.00021	-0.0018	-0.00147	-0.00015
LI	0.00240*	0.00058*	0.00049	-0.00048	0.0051	0.0015
NAPM	0.00044*	0.00087*	0.00087*	0.00056*	0.00080*	0.00048
PAY	0.00001	0.00004*	0.00005*	0.00005*	0.00005*	0.00005*
PPI	0.00042	0.0029	0.0021	0.0012	0.0021	-0.00043
RS	0.0011	0.00390*	0.00280*	0.0022	0.0022	0.0024
TRADE	0.00078*	0.00110*	0.00097*	0.00120*	0.00140*	0.00120*
UNEM	-0.0026	-0.00730*	-0.00850*	-0.0057	-0.0058	-0.0027

Table 6.9 cont.

Release	2h	2.5h	3h	6h	12h
CC	0.00042*	0.00041*	0.00041*	0.00031*	0.00044*
CPI	-0.0051	-0.0066	-0.0058	0.0032	0.0028
CU	-0.0019	-0.0024	-0.0021	-0.003	-0.0012
DG	0.00053	0.00043	0.00068	0.00043	0.0006
IP	-0.00023	0.0034	0.0043	0.01235	0.01783*
LI	0.0002	0.01000*	0.01270*	0.00597	0.01350*
NAPM	0.00025	-0.00006	0.00029	0.00012	0.00008
PAY	0.00005*	0.00006*	0.00006*	0.00006*	0.00005*
PPI	-0.0022	-0.0032	-0.0036	-0.0029	-0.00088
RS	0.0024	0.0033	0.0021	0.001	0.0034
TRADE	0.00130*	0.00048	0.00055	0.00018	-0.00022
UNEM	-0.001	0.0029	0.0012	-0.0001	-0.003

Notes: Each cell gives the slope coefficient from a linear regression of the return over the period displayed in the first row of the given column on the series of forecast errors from the series in the first cell of the given row. An asterisk denotes that the coefficient is significantly different from zero at a 5% level.

Table 6.10: Predictive power of 15 minute return sign - US data

Series	6 H		12 H		24 H		No. of obs.
	Pred.	P-value	Pred.	P-value	Pred.	P-value	
CC	22	12%	23	7%	24	3%	36
CPI	20	31%	21	20%	19	43%	36
DG	22	12%	24	3%	23	7%	36
IP/CU	19	43%	23	7%	26	1%	36
LI	20	25%	19	37%	15	84%	35
NAPM	23	7%	22	12%	24	3%	36
PPI	9	100%	23	7%	25	1%	36
RS	23	7%	25	1%	22	12%	36
TRADE	19	24%	18	36%	18	36%	33
UNEM/PAY	26	1%	26	1%	25	1%	36
All	192	0%	202	0%	202	0%	325
All 8.30	128	0%	134	0%	128	0%	217
All 10	45	2%	45	2%	48	0%	72

Notes: The column headed 'Pred.' displays the number of observations with the sign of the 15 m return equal to the sign of the return over the period in the first row. The next column gives the probability of observing a number of predictions equal or larger than the actual, under the null of the number of predictions being random. The row labelled 'All' reports the sign tests for the aggregate of all announcements. Similarly, the rows labelled 'All 8.30' and 'All 10' report the results for the aggregate of releases at 8.30 EST and 10.00 EST.

Table 6.11: Structural stability tests - US News

Series	1992	1993	1994	t:92	t:93	t:94	t:92=93	t:93=94	t:92=94
CC	0.00025	0.00052	0.00005	1.01	4.71	0.91	-1.03	3.73	0.77
DG	0.00144	0.00095	0.00019	5.32	4.64	1.47	1.44	3.10	4.13
NAPM	0.00103	0.00071	0.00034	5.49	2.08	0.73	0.84	0.64	1.38
PAY	0.00005	0.00003	0.00002	5.49	3.80	2.12	1.47	0.60	1.85
RS	0.00737	0.00327	0.00303	1.68	1.72	4.19	0.86	0.12	0.98
TRADE	0.00221	0.00129	0.00020	3.86	3.75	0.43	1.39	1.85	2.71
UNEM	-0.00939	-0.00371	-0.0072	-1.40	-0.67	-1.18	-0.65	0.42	-0.24
CPI	0.008548	0.00145	0.004938	1.01	0.58	2.02	0.80	-0.99	0.41
CU	-0.00674	0.0015	-0.011	-1.07	1.81	-1.96	-1.29	2.21	0.50
IP	0.020748	-0.00423	0.008972	2.34	-3.05	1.56	2.79	-2.23	1.11
LI	-0.00314	0.006629	-0.00469	-1.07	3.61	-1.79	-2.83	3.53	0.39
PPI	0.009596	0.005931	-0.00182	1.97	2.67	-0.55	0.68	1.94	1.93

Notes: The columns with number headings give the coefficients on the series created from interacting the forecast error series with that year's dummy. The following three columns give the t-statistics relevant to the test that the coefficient on the forecast error is zero in a given year. The final three columns give the t-statistics relating to the hypothesis that the coefficients in two given years are equal. The critical values for the t-statistics are 2.04 at 5% and 1.70 at 10%.

Table 6.12: The impact of German News on DEM/USD returns

Series	Coeff.	T-stat	\bar{R}^2	Scaled
UNEM	-0.00003	-3.06	0.16	-0.00029
IP	-0.00023	-2.07	0.05	-0.00034
CPI	0.00261	2.00	0.08	0.00020
RS	-0.00006	-1.35	0.01	-0.00018
ORDERS	0.00006	1.15	0.01	0.00010
PPI	-0.00113	-0.81	-0.08	-0.00014
M3	-0.00017	-0.75	-0.19	-0.00022
TRADE	-0.00004	-0.41	0.01	-0.00008
WPI	0.00017	0.35	-0.01	0.00005

Notes: See Table 6.7.

Table 6.13: The impact of German News on DEM/USD returns (weighted data)

Series	Coeff.	T-stat	\bar{R}^2	Scaled
M3	-0.00046	-5.02	0.35	-0.00060
UNEM	-0.00002	-3.43	0.29	-0.00027
IP	-0.00029	-2.80	0.24	-0.00043
CPI	0.00540	2.72	0.08	0.00042
PPI	-0.00358	-2.00	0.33	-0.00044
RS	-0.00007	-1.84	0.03	-0.00020
TRADE	-0.00014	-1.30	0.05	-0.00025
WPI	0.00012	0.23	-0.03	0.00003
ORDERS	0.00000	0.08	0.00	0.00001

Notes: See Table 6.7.

Table 6.14: Market Efficiency tests - German data

Series	Constant	T-Const.	News	T-news	Expected	T-Exp.	\bar{R}^2
CPI	-0.00022	-1.04	0.00281	2.36	0.00114	2.64	0.14
IP	0.00030	1.76	-0.00023	-2.33	-0.00030	-1.27	0.16
M3	-0.00098	-2.43	-0.00007	-0.35	-0.00005	-0.27	0.02
ORDERS	-0.00012	-0.67	0.00006	1.11	-0.00011	-0.74	0.03
PPI	0.00039	1.42	-0.00059	-0.53	-0.00023	-0.13	0.00
RS	0.00038	1.09	-0.00003	-0.68	0.00015	1.16	0.05
TRADE	0.00013	0.76	0.00006	0.49	0.00010	0.64	0.04
UNEM	-0.00001	-0.10	-0.00002	-2.93	0.00001	0.41	0.15
WPI	-0.00007	-0.40	0.00014	0.29	-0.00037	-0.36	0.01

Notes: See Table 6.8.

Table 6.15: The persistence of the effect of German News on the DEM/USD (original data)

PERIOD	5 m	15 m	30 m	45 m	1 h	1.5 h
CPI	0.0001	0.00261*	0.00023	0.00149	0.00073	0.00192
IP	-0.00008	-0.00023**	-0.00005	-0.00004	-0.00012	-0.00035
M3	0.00008	-0.00017	-0.00014	-0.00008	-0.0002	-0.00021
ORDERS	-0.00007**	0.00006	0.0001	0.00012	0.00021*	0.00005
PPI	-0.00005	-0.00113	-0.00343*	-0.00343	-0.00672**	-0.00786**
RS	-0.00002	-0.00006	-0.00008	-0.00004	-0.00003	-0.00002
TRADE	0.00003	-0.00004	0.0001	-0.00018	-0.00008	-0.00034
UNEM	-0.00001	-0.00003**	-0.00003**	-0.00002	-0.00002	-0.00001
WPI	-0.00006	0.00017	-0.00062	-0.00028	0.00004	-0.0006

Table 6.15 cont.

PERIOD	2 h	2.5 h	3 h	6 h	12 h
CPI	0.00101	0.00159	0.0023	0.00751	-0.00123
IP	-0.00022	-0.00029	-0.00046	-0.00041	-0.0008
M3	-0.00029	-0.00021	-0.00042	-0.00059**	-0.00052
ORDERS	0.00004	-0.00012	0.00014	0.0004	0.00027
PPI	-0.01175**	-0.01459**	-0.01500**	-0.0084	-0.01027
RS	-0.00004	-0.00004	-0.00017*	0.00021	0.00028
TRADE	-0.00041*	-0.00016	-0.00031	-0.00017	-0.00013
UNEM	-0.00002	-0.00002	0	-0.00004	-0.00009
WPI	0.00047	-0.00097	-0.00003	-0.00007	-0.00101

Notes: Each cell gives the slope coefficient from a linear regression of the return over the period displayed in the first row of the given column on the series of forecast errors from the series in the first cell of the given row. An * (**) denotes that the coefficient is significantly different from zero at a 10% (5%) level.

FIGURES

Figure 6.1: US Announcements: Average Scaled Impact over time

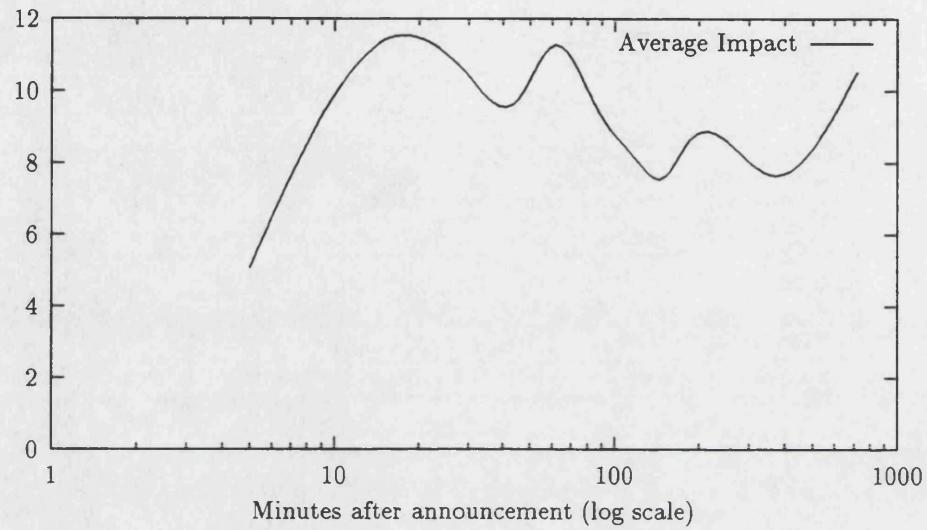


Figure 6.2: US Announcements: Average T-statistics over time

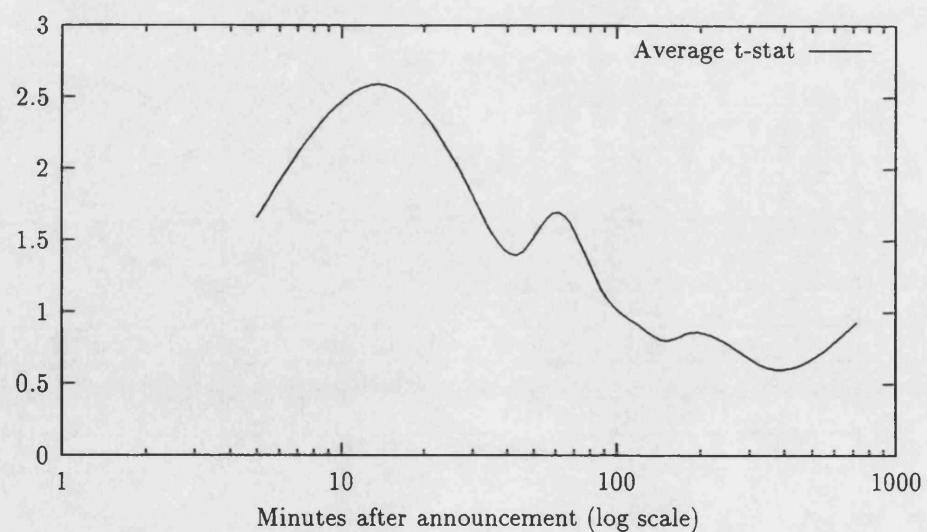


Table 6.16: The persistence of the effect of German News on the DEM/USD (weighted data)

PERIOD	5 m	15 m	30 m	45 m	1 h	1.5 h
CPI	0.00437*	0.00540**	0.00574	0.00812*	0.00978	0.00903*
IP	-0.00011	-0.00029**	-0.00025**	-0.00022**	-0.00037**	-0.00066
M3	-0.0001	-0.00046**	-0.00049	-0.00051**	-0.00085*	-0.00066
ORDERS	0	0	-0.00006	-0.0001	0.00009	-0.00012
PPI	-0.00027	-0.00358*	-0.00207	-0.00244*	-0.00380**	-0.00467**
RS	-0.00007*	-0.00007*	0.00001	0.00008	0.00003	0
TRADE	0.00001	-0.00014	0.00022	-0.00026**	-0.0001	-0.00033
UNEM	0	-0.00002**	-0.00002	-0.00002	-0.00001	0
WPI	0.00037	0.00012	-0.00016	-0.00043	0.00012	-0.00025

Table 6.16 cont.

PERIOD	2 h	2.5 h	3 h	6 h	12 h
CPI	0.01132	0.01613**	0.02003**	0.01423	0.00824
IP	-0.00032	-0.00062	-0.00077	0.00015	-0.0015
M3	-0.00095**	-0.00099**	-0.00095**	-0.00099**	-0.00023
ORDERS	-0.00004	-0.00027**	-0.00039**	-0.00124**	-0.00075
PPI	-0.00886**	-0.00719	-0.00697	0.01362	0.02167
RS	-0.00011	-0.00008	-0.00022	0.00005	0.00016
TRADE	-0.00039	-0.00035**	-0.00089**	-0.0001	-0.00071
UNEM	-0.00001	0	0.00001	0	0.00003
WPI	-0.00001	-0.00071	-0.00058	-0.00344	-0.00001

Notes: See Table 6.15.

Table 6.17: The impact of German News on DEM/USD 3 hour returns

Series	Original			Weighted		
	Coeff.	T-stat	Scaled	Coeff.	T-stat	Scaled
TRADE	-0.00031	-0.88	0.00055	-0.00089	-2.85	0.00159
CPI	0.00230	0.45	0.00018	0.02003	2.25	0.00156
M3	-0.00042	-1.11	0.00054	-0.00095	-3.13	0.00123
IP	-0.00046	-1.45	0.00067	-0.00077	-0.82	0.00112
PPI	-0.01500	-2.46	0.00186	-0.00697	-1.27	0.00087
ORDERS	0.00014	0.43	0.00025	-0.00039	-2.33	0.00067
RS	-0.00017	-1.84	0.00050	-0.00022	-1.28	0.00066
WPI	-0.00003	-0.02	0.00001	-0.00058	-0.33	0.00016
UNEM	0.00000	-0.15	0.00004	0.00001	0.29	0.00015

Notes: The column headed 'coefficient' displays the estimated coefficients on the forecast error series. The next column gives the t-statistics relating to the hypothesis that the coefficients are zero. The critical values for the t-statistics are 2.04 at 5% and 1.70 at 10%. The column headed 'Scaled' displays the product of the coefficient with the average absolute forecast error.

Table 6.18: Predictive power of 15 minute return sign - German data

Series	6 H		12 H		24 H		No. of obs.
	Pred.	P-value	Pred.	P-value	Pred.	P-value	
CPI	22	12%	24	3%	22	12%	36
IP	23	7%	25	1%	25	1%	36
M3	25	1%	21	20%	23	7%	36
ORDERS	24	3%	23	7%	21	20%	36
PPI	17	69%	16	80%	14	93%	36
RS	22	12%	20	31%	18	57%	36
TRADE	22	9%	19	37%	15	84%	35
UNEM	17	69%	18	57%	20	31%	36
WPI	19	43%	19	43%	17	69%	36
All series	191	0%	185	1%	175	7%	323
Top 6	126	1%	124	2%	122	3%	216
Prices	69	25%	64	19%	60	61%	107
Real	81	1%	82	1%	80	3%	144

Notes: See Table 6.10.

Table 6.19: Structural stability tests - German data (original data)

Series	1992	1993	1994	t:92=93	t:93=94	t:92=94
CPI	0.00136	0.00736	0.00150	-1.92	1.84	-0.07
IP	-0.00013	-0.00027	-0.00035	0.35	0.18	1.66
UNEM	-0.00003	-0.00004	-0.00000	0.20	-2.32	-1.85
M3	-0.00131	0.00036	-0.00034	-2.68	2.09	-1.82
ORDERS	0.00015	0.00001	0.00006	0.96	-0.35	0.48
PPI	-0.00206	-0.00252	0.00027	0.10	-0.76	-0.76
RETAIL	-0.00013	-0.00001	-0.00003	-1.15	0.22	-1.21
TRADE	0.00015	-0.00020	-0.00003	1.63	-1.29	0.81
WPI	-0.00065	0.00171	0.00003	-2.52	1.70	-0.86

Notes: see Table 6.11.

Table 6.20: Structural stability tests - German data (weighted data)

Series	1992	1993	1994	t:92=93	t:93=94	t:92=94
CPI	0.00197	0.00855	0.00604	-1.24	0.49	-2.28
IP	-0.00023	-0.00047	-0.00050	0.51	0.06	2.93
M3	-0.00131	-0.00071	-0.00036	-0.31	-4.57	-0.50
PPI	-0.00545	-0.00472	-0.00106	-0.15	-0.74	-4.12
RS	-0.00013	-0.00004	-0.00010	-1.02	0.70	-0.30
UNEM	-0.00002	-0.00003	-0.00001	0.52	-0.98	-1.02
ORDERS	0.00006	0.00000	-0.00028	0.98	1.06	1.29
TRADE	0.00014	-0.00033	-0.00022	2.74	-0.96	1.94
WPI	-0.00097	0.00168	-0.00014	-3.14	1.73	-0.96

Notes: See Table 6.11.

FIGURES

Figure 6.1: US Announcements: Average Scaled Impact over time

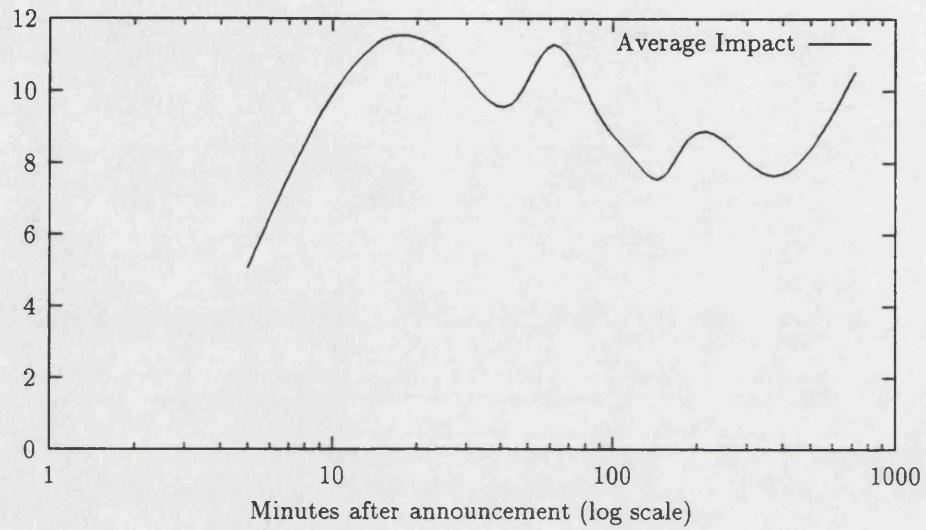
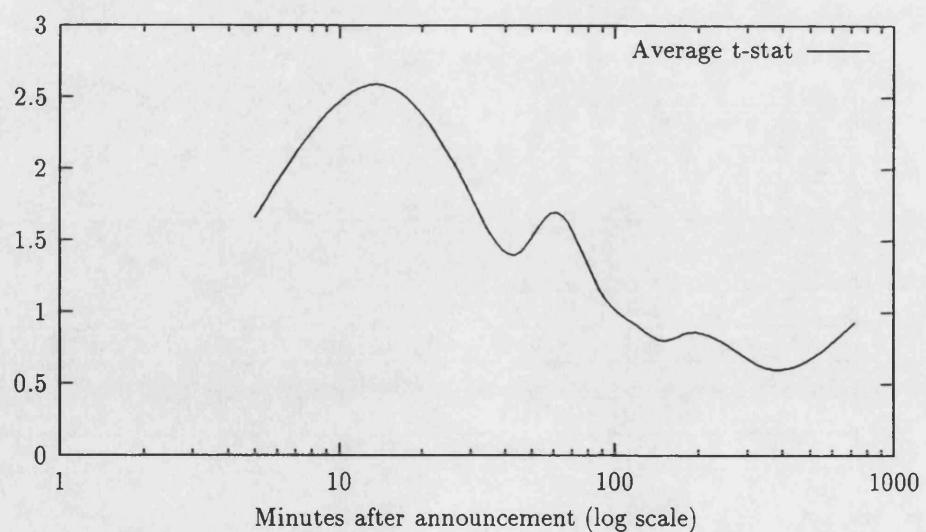


Figure 6.2: US Announcements: Average T-statistics over time



Chapter 7

Nonlinear Mean Reversion in the Foreign Exchange Forward Premium

7.1 Introduction

Conventional wisdom asserts that the foreign exchange forward premium should follow a stationary stochastic process. Indeed, covered interest rate parity implies that the stochastic process for the premium should be identical to that of the interest rate differential which, as Brenner and Kroner (1995) suggest, is very likely to be covariance stationary. However, recent research has provided evidence that the forward premium may be a non-stationary process. Results contained in Crowder (1994), for example, suggest that forward premia are $I(1)$ processes, whilst Baillie

Figure 6.3: German Announcements: Average Scaled Impact over time





Figure 6.4: German Announcements: Average T-statistics over time



Chapter 7

Nonlinear Mean Reversion in the Foreign Exchange Forward Premium

7.1 Introduction

Conventional wisdom asserts that the foreign exchange forward premium should follow a stationary stochastic process. Indeed, covered interest rate parity implies that the stochastic process for the premium should be identical to that of the interest rate differential which, as Brenner and Kroner (1995) suggest, is very likely to be covariance stationary. However, recent research has provided evidence that the forward premium may be a non-stationary process. Results contained in Crowder (1994), for example, suggest that forward premia are I(1) processes, whilst Baillie

and Bollerslev (1994) present estimations which provide evidence that forward premia are fractionally integrated processes. Importantly, for two of the three currency pairs studied by Baillie and Bollerslev (1994) the estimate of the fractional differencing parameter exceeds one half, implying non-stationary (although mean-reverting) forward premia.¹

The impact of the above results can be seen quite clearly if one examines the standard regression used in order to test the unbiasedness of forward rates. Popularised by Fama (1984), its structure is as follows,

$$\Delta s_{t+k} = \alpha_0 + \alpha_1(f_t - s_t) + \epsilon_t \quad (7.1)$$

where s_t is the logarithm of the spot exchange rate, f_t is the log of the k -period forward rate and ϵ_t is a covariance stationary, regression error term. Equation (7.1) is derived from the fact that in a world characterised by risk neutrality and rational expectations, the forward rate should only differ from the future spot rate by a rational expectations error. Hence, under risk neutrality and rational expectations, the value of α_0 should be zero and the value of α_1 should be unity. However, numerous applications of the above regression, see Engel (1996) for references, result in estimates of the slope coefficient which are significantly below unity and, in many cases, negative.² A possible reason for this finding, presented by Fama (1984), is that the forward rate contains a risk premium which is an omitted regressor in equation (7.1) and hence biases the estimation.³

¹The fractional differencing parameter for the third exchange rate studied by Baillie and Bollerslev (1994) is also quite close to the boundary of stationarity.

²Such studies include the aforementioned Fama (1984) paper, Bilson (1981), Backus, Gregory, and Telmer (1993) and Mark, Wu, and Hai (1993) to name but a few. Also, estimates of equation (7.1) for the data set employed in this study give exactly the results mentioned above.

³An alternative explanation for the result is the failure of rational expectations. Froot and Frankel (1989) and Peel and Pope (1991) both present results using survey expectations of exchange

The 'Fama' regression can be seen as a restricted version of a more general, cointegrating VAR representation for spot and forward rates. Evidence suggests that both spots and forwards are unambiguously $I(1)$ which, assuming risk neutrality or at least an $I(0)$ risk premium, implies cointegration between the pair with cointegrating vector $(1, -1)$. Hence, equation (7.1) can be interpreted as a restricted version of an error correction representation for the spot rate, with the error correction term given by the forward premium.

It is clear that for equation (7.1) to be well specified and for standard asymptotics to apply, both the right and left hand side variables should be covariance stationary. There is extensive evidence on the unit root behaviour of exchange rates, implying that the dependent variable in (7.1) is $I(0)$. However, the results of Crowder (1994) and Baillie and Bollerslev (1994) both indicate non-stationarity in the right hand side variable. Taken together, the above facts on the time-series dynamics of exchange rate returns and forward premia imply that the 'Fama' regression is mis-specified or, alternatively, that spot and forward exchange rates cannot be fully cointegrated.⁴

The possible lack of cointegration between spot and forward exchange rates is reflected in the empirical literature on this subject. Whereas several authors present evidence of cointegration with a near $(1, -1)$ cointegrating vector, others dispute the existence of cointegration. Examples of the former group are Hakkio and Rush (1989) and Mark, Wu, and Hai (1993). A paper which indicates a lack of cointegration between forward and spot rates is Evans and Lewis (1993). Goodhart, McMahon, and Ngama (1997) also present evidence that the cointegrating vector for a group of spot and forward exchange rates is very close to $(1, -1)$ but they go on to interpret this evidence very differently from the authors cited above. Using the Phillips-Hansen methodology, they show that the cointegrating vector is signif-

rate changes which suggest expectational inefficiency.

⁴By this we mean that spot and forward rates cannot be $CI(1,1)$ in the terminology of Engle and Granger (1987) although they may be fractionally cointegrated.

icantly different from that implied by theory and, hence, reject the unbiasedness hypothesis and question the validity of the ‘Fama’ regression.

The approach taken in this paper is to return to the univariate time-series properties of the forward premium and conduct a closer examination of its dynamic behaviour. Specifically we present two sets of results which bear on this issue. First, we employ a kernel regression technique to yield a nonparametric estimate of the relationship between changes in the forward premium and the lagged level of the forward premium. This corresponds to the conventional augmented Dickey-Fuller methodology for testing stationarity and permits non-linearity in the relationship between the change in the premium and the lagged level. These nonparametric characterisations of the dynamics of the forward premia lead on to our second set of estimations which are fully parametric models of non-linear mean reversion. The estimations are modified threshold autoregressions (TAR’s) for the change in the forward premium. The modification is the addition of the lagged level of the forward premium to the right hand side variable set such that the specification permits a fully parametric description of the non-linearities in the ADF regression, formalising the evidence from the kernel regressions.

Non-linearities in the dynamics of the forward premium could result from several sources. A first potential reason is that the nominal interest differential is a non-linear process. Assuming no hindrance to arbitrage between spot and forward markets implies that covered interest parity holds. Hence the forward premium should at all times be identical to the nominal interest differential and clearly then the dynamics of the two processes will be identical. Alternatively, if one relaxes the risk neutrality assumption, a risk premium is introduced into the spot-forward relationship and non-linear temporal dependence in this process may yield non-linear dynamics in the forward premium.

Frictions in covered interest arbitrage might also introduce non-linearities into the process for the forward premium. Assuming small transaction costs to spot and forward trades will imply a range of deviations between the forward and expected spot in which the deviation process contains a unit root.⁵ For large deviations from covered interest parity running the arbitrage becomes profitable, however, such that the difference between forward and expected spot mean reverts. If it is the case that the nominal interest differential follows a linear $I(0)$ process, for example, then this mechanism will produce non-linearities in the forward premium process.

Lastly, it may be the case that over the sample periods chosen in previous studies, there are structural breaks in the forward premium process. These shifts in the process will imply different dynamics for different time periods such that the entire process is better represented by a non-linear rather than linear process.

The rest of the paper is set out as follows. In Section 7.2 we discuss the empirical methodologies employed in the study, covering the kernel regression technique first and then examining the specification and estimation issues associated with TAR models. The data and its essential features are presented in Section 7.3.1, followed by the results of our empirical work in Section 7.3.2. The paper concludes in Section 7.4.

7.2 Empirical Methodology

As detailed in the previous section, the focus of this research is the stationarity and mean reversion of the forward premium. The starting point of our analysis is

⁵See Michael, Nobay, and Peel (1997) for a transactions costs based explanation of non-linear adjustment of real exchange rates towards purchasing power parity.

the standard Augmented Dickey-Fuller (ADF) regression generally employed to test for unit roots. Given a time-series y_t with an AR(1) representation a test for the presence of a unit-root in the series is available from the following regression;

$$\Delta y_t = (\phi - 1)y_{t-1} + \epsilon_t, \quad \epsilon_t \sim NID(0, \sigma^2) \quad (7.2)$$

Note that the error term in the above structure is restricted to be IID. If the true data generating process for y_t is an AR(p) process, the omission of higher order dependence in the above equation will induce autocorrelation in the innovations and invalidate the test statistics. A generalisation of equation (7.2) which takes account of the possibility of higher order dependence is given below;

$$\Delta y_t = \rho y_{t-1} + \sum_{j=1}^{p-1} \theta_j \Delta y_{t-j} + \epsilon_t, \quad \epsilon_t \sim NID(0, \sigma^2) \quad (7.3)$$

As in equation (7.2) the unit root test from equation (7.3) consists of testing the null that the coefficient on y_{t-1} is zero using a conventional t -statistic. Critical values for this test were first calculated by Dickey and Fuller (1979). For values of ρ significantly below zero, y_t is a short memory, covariance stationary, mean reverting process. If ρ is indistinguishable from zero the process contains a unit root and demonstrates complete persistence, whilst y_t is explosive for values of ρ significantly above zero. Hence ρ can be seen to summarise the mean reverting properties of the process. It is the possibility of variation in mean reversion for differing levels of the process that our alternative specifications are designed to address. Note also that the above characterisation of the process y_t does not permit long range dependence. The possibility of this type of behaviour in forward premia is the focus of Baillie and Bollerslev (1994).

7.2.1 Kernel Regression

The kernel regression methodology we employ provides a straightforward generalisation of the linear ADF approach outlined above.⁶ Consider the following, possibly non-linear, relationship between y and x ;

$$y_t = m(x_t) + u_t, \quad u_t \sim NID(0, \sigma_u^2) \quad (7.4)$$

where $m(x)$ is an unknown function which is fixed and continuous. A natural way in which to get an estimate of the function $m(x)$ is to use a localised weighted average of y_t i.e. given certain smoothness requirements on $m(x)$, an estimate of the value of the function at a given point (x_0) can be constructed by taking a weighted average of the y -values associated with observations in a neighbourhood of x_0 . The general form of this locally weighted estimator can be expressed as below;

$$\hat{m}(x) = \frac{1}{T} \sum_{t=1}^T \omega_t(x) y_t \quad (7.5)$$

where $\omega_t(x)$ is the weighting function and T the total sample size. Clearly, given the previous discussion, we would wish the weighting function to be large for realisations of the process close to x_0 and to decrease towards zero as we move further from x_0 .

The specific choice of weighting function is as follows. Consider a probability density function $f(u)$ (i.e. a function which is weakly greater than zero for all u and integrates to unity.) This density is modified by scaling with respect to a given

⁶For a more comprehensive explanation of non-parametric regression techniques see Härdle (1990) and Campbell, Lo, and Mackinlay (1997).

parameter h i.e.

$$K_h(u) = \frac{1}{h} f\left(\frac{u}{h}\right), \quad \int K_h(u) du = 1 \quad (7.6)$$

The weighting function employed is then constructed as follows;

$$\omega_t(x) = \frac{K_h(x - x_t)}{\frac{1}{T} \sum_{t=1}^T K_h(x - x_t)} \quad (7.7)$$

Substituting equation (7.7) into the original non-parametric regression formulation, equation (7.5), yields the Nadaraya-Watson kernel estimator;

$$\hat{m}_h(x) = \frac{1}{T} \sum_{t=1}^T \omega_t(x) y_t = \frac{\sum_{t=1}^T K_h(x - x_t) y_t}{\sum_{t=1}^T K_h(x - x_t)} \quad (7.8)$$

In the above representation of the kernel estimator there are two factors which are, as yet, unspecified. The first is the choice of the density function to employ as the basis for the weighting function. In what follows we use the standard normal density. Second we need a choice of the bandwidth parameter (h).⁷ As the bandwidth controls the effective spread of the weighting function, too large a bandwidth will imply a very smoothed function whilst too small a bandwidth will lead to an excessively variable estimate.⁸ In the empirical work which follows, the bandwidth is

⁷It is generally accepted that the choice of kernel affects estimation to a far smaller extent than the choice of bandwidth. Hence we use an optimal bandwidth selection procedure for the given standard normal kernel.

⁸This is the standard tradeoff between bias and variance: small bandwidths imply large variance whilst large bandwidths yield greater bias.

chosen optimally via the cross-validation procedure.⁹

The use of the above tool for analysis of the dynamics of the forward premium is analogous to the standard ADF regression presented in equation (7.2). Specifically, the change in the forward premium is used as the dependent variable in equation (7.8) whilst the lagged level of the premium is the right hand side variable. This yields an ADF function which is not constrained to be linear and can be compared to that derived from the conventional regression specification.

7.2.2 Threshold ADF Specifications

The kernel regressions yield non-parametric estimates of the relationship between the change in and lagged value of the premium and hence give a graphical summary of its mean reverting properties. We complement these non-parametric estimates with a fully parameterised model of non-linear mean reversion based on the Threshold Autoregressive (TAR) approach presented by, for example, Tong (1990). A standard two regime TAR model can be represented as follows;

$$y_t = \begin{cases} \kappa_1 + \sum_{i=1}^p \phi_{1i} y_{t-i} + \epsilon_{1t}, & |y_{t-d}| < \delta \\ \kappa_2 + \sum_{i=1}^p \phi_{2i} y_{t-i} + \epsilon_{2t}, & |y_{t-d}| > \delta \end{cases} \quad (7.9)$$

where y_t is our forward premium series, d is the delay parameter (i.e. the lag of the premium which determines switches in regimes) and δ is the inner regime cutoff parameter.¹⁰ Note that in the above specification the dynamics for large positive and

⁹See Härdle (1990) for a complete presentation of the cross-validation procedure.

¹⁰We refer to the above model as a two regime TAR as, although there are actually three regimes in terms of the level of the process, the regimes for large positive and negative realisations of the process are constrained to have identical dynamics.

negative realisations of the premium are constrained to be identical. This restriction is imposed both for reasons of symmetry and due to the relatively small number of extreme observations.

Analogous to the manipulations used to yield the ADF specification from the standard autoregressive model, the TAR model can be reparameterised to give the following Threshold ADF representation;

$$\Delta y_t = \begin{cases} \kappa_1 + \rho_1 y_{t-1} + \sum_{i=1}^{p-1} \phi'_{1i} \Delta y_{t-i} + \epsilon_{1t}, & |y_{t-d}| < \delta \\ \kappa_2 + \rho_2 y_{t-1} + \sum_{i=1}^{p-1} \phi'_{2i} \Delta y_{t-i} + \epsilon_{2t}, & |y_{t-d}| > \delta \end{cases} \quad (7.10)$$

Given values for both the delay parameter (d) and the threshold parameter (δ) the above specification can be estimated efficiently via ordinary least squares. The length of the autoregressive polynomials can be chosen via either the Akaike or Bayes Information Criterion.¹¹

Estimating the delay parameter was attempted via two methodologies. The first and most straightforward is to estimate the model over a range of different delays and choose the specification which minimises the residual sum of squares. An alternative, proposed by Tsay (1989), rests on the fact that if one estimates a linear specification one is mixing data from different regimes. Ordering the data by the threshold variable (y_{t-d}) and running a rolling linear regression will then yield correlation between the independent variables and the residuals at points when a change in regimes occurs. A regression of the residuals from the linear specification on the independent variables can then be used to form a standard F -test of the null of linearity versus a TAR alternative hypothesis. This procedure can be repeated over

¹¹Throughout the empirical work in this study we employ the Akaike criterion due to the fact that the Bayes criterion tends to yield more parsimonious models and we wish to ensure that we have eliminated any residual serial correlation.

a range of different delays and choice of the optimal delay corresponds to choosing the value of d for which the p -value of the test is minimised. Finally, the remaining free parameter, the threshold cutoff value (δ), is estimated via a grid search over values deemed to be reasonable after inspection of the kernel regressions.

7.3 Data and Results

7.3.1 The Data

In this study we employ a sample of spot and forward exchange rates covering the period 3/1/76 to 31/12/85.¹² Both rates are sampled at a daily frequency and the data set is constructed such that on each observation date we have the prevailing spot rate plus the forward rate for a contract maturing in 30 days.¹³ The full sample consists of the following currencies: USD/GBP, DEM/GBP, CAD/GBP, JPY/GBP, FRF/GBP and SWK/GBP. The basic series of interest is the forward premium for each currency pair. Summary statistics for the forward premium data are given in Table 7.1.

Perhaps the most striking statistics in Table 7.1 are the Box-Ljung Q -statistics. Across exchange rates there is evidence of huge temporal dependence with the Q -statistics being orders of magnitude greater than the 1% and 5% critical values for the null of no dependence. A further result is that most of the premia display platykurtosis, aside from that for the Swedish Krona which demonstrates strong

¹²We chose not to use data from the mid-1980's onward due to the possibility of structural shifts in the forward premia processes due to Sterling's entry into and exit from the ERM.

¹³Both spot and forward data were retrieved from the Datastream exchange rate database. The Datastream code mnemonics are available from the author.

leptokurtosis. The Swedish Krona is also anomalous in that it is strongly positively skewed.

As the focus of this work is the stationarity of the forward premia we also present summary statistics for the first difference of the premia. These are given in Table 7.2.

Examining the differenced data one first sees a large drop in the Q -statistics for the forward premia. However, despite a reduction of a couple of orders of magnitude, they still indicate strong serial dependence in the premia at horizons of up to 10 and 20 lags. Also, there is now more consistent evidence of leptokurtosis in the differenced premia, with all 6 kurtosis coefficients exceeding 3, two of them by a very large amount.¹⁴

7.3.2 Results

In this subsection we present the results from our empirical specifications. We start with the results from a standard ADF formulation for the set of forward premia. Then, for comparability with Baillie and Bollerslev (1994), we present estimates of the fractional differencing parameter in each forward premium series. Kernel regression results are presented next and, finally, the modified TAR estimates are displayed and discussed.

¹⁴This feature of the data is likely to be due to conditional heteroskedasticity in the differenced forward premium processes.

Linear ADF Specifications

Results from the standard linear ADF specification are given in Table 7.3. As indicated in the previous section we employ the Akaike criterion in order to determine the appropriate lag length in the ADF regression. The residual diagnostics indicate that the AIC chosen specifications are, in general, satisfactory. All of the Durbin-Watson statistics indicate a lack of first order serial correlation in the residuals from the ADF regression. Also, the majority of the Box-Ljung Q -statistics are very small and insignificantly different from zero, implying a lack of higher order dependence. The only exceptions are the $Q(20)$ statistics for the USD/GBP and SWK/GBP, which indicate that the hypothesis of up to twentieth order serial correlation in the residuals cannot be rejected.

Examining the Dickey-Fuller coefficients (and their respective heteroskedasticity-robust standard errors) demonstrates that for no exchange rate is the hypothesis of a unit root in the premium rejected. The robust t -statistics are generally greater than two but in all cases fall short of the 5% Dickey-Fuller critical level.¹⁵

Hence the results from testing for unit-roots in our daily forward premia entirely corroborate those obtained from monthly data by Crowder (1994). Forward premia seem to demonstrate $I(1)$ behaviour implying that spot and forward exchange rates cannot be $CI(1,1)$ with cointegrating vector $(1,-1)$. This in turn implies that the deviation between spot and forward rates can become arbitrarily large, an uncomfortable result. A potential explanation for this finding is that the result is driven by the linearity of the ADF formulation. If, for example, the risk premium follows a mean reverting but non-linear process, the forward premium will adopt this behaviour, implying that the linear ADF representation is mis-specified.

¹⁵The critical values for the ADF test are -2.86 at 5% and -3.12 at 1%. Note also that five of the six premia fall short of the 10% critical value of -2.57, the exception being the JPY/USD premium.

Long Memory Estimates

Baillie and Bollerslev (1994) present evidence that the forward premium is a fractionally integrated process. Using monthly data on the Canadian Dollar, German Mark and British Pound, all quoted against the US Dollar, they derive the degree of fractional integration in forward premia from a fully parametric ARFIMA model. For all three currencies the degree of integration (\hat{d}) is between zero and unity, implying long range dependent although mean-reverting behaviour. Importantly, for the Deutschemark and Sterling data the degree of integration exceeds one half, implying that the premia are non-stationary.

We attempt to corroborate their results using our daily sample of six exchange rates. Our estimate of the degree of integration is derived from the Gaussian semiparametric estimator proposed by Künsch (1987) with an automatic bandwidth selection procedure derived in Henry and Robinson (1996).¹⁶ The results are given in Table 7.4.

As is clear from Table 7.4, using daily forward premia does not alter the conclusion of Baillie and Bollerslev (1994) that the series are long range dependent. For all of the estimates the degree of integration is seen to be significantly above zero, implying that premia possess long memory and are mean reverting, albeit at a very slow rate. However, one feature of the results of Baillie and Bollerslev (1994) which does not hold up in our data set is the non-stationarity of the premia. All of the six long memory estimations yield values for \hat{d} which are significantly below one half implying that these daily forward premia can be characterised as covariance stationary processes. This is potentially good news for studies testing the unbiasedness of the

¹⁶For a more detailed presentation of the Gaussian estimator and an application to intra-day foreign exchange rate volatility see Chapter 5. Also note that the Gaussian estimate is robust to conditional heteroskedasticity in the process under study, a feature which is likely to be prominent for forward premium data.

forward foreign exchange market using the 'Fama' regression. Results from the daily data employed in the current study indicate that both the right and left hand sides of equation (7.1) are stationary such that the regression equation is well specified and parameter estimates consistent. However, as Engel (1996) points out, the fractional integration of the premium is inconsistent with the unbiasedness hypothesis. Rational expectations imply that the error term in the 'Fama' regression must be $I(0)$, such that exchange rate returns and forward premia must be integrated to the same degree. Given the empirical evidence that returns are $I(0)$, the fractional integration in premia violates this condition.

Kernel Regression

In the previous two subsections we have presented results similar to those found using monthly data by Crowder (1994) and Baillie and Bollerslev (1994). Standard ADF tests indicate a unit root in all six forward premium series, whilst estimates of the degree of fractional integration in premia demonstrate that they are covariance stationary, long memory processes. The rest of this section focuses on the mean-reversion of the premia, first using the kernel regression methodology presented in Section 7.2.1 and then presenting results from the TADF specifications introduced in Section 7.2.2.

As outlined in Section 7.2.1 the kernel regressions yield a non-parametric estimate of the relationship between the change in the forward premium and its first lag. If the forward premium were to follow a stationary $AR(1)$ process, this function should be linear with a significantly negative slope. Alternatively, an $I(1)$ process would imply no relationship between the change and first lag, implying that the function should be flat.

The estimated kernel regression lines are plotted in Figures 7.1 to 7.6. Alongside each kernel regression estimate we plot the linear relationship between changes in premia and lagged premia derived from the basic ADF results in Section 7.3.1.

It is immediately apparent that the relationship between changes in and lagged premia is not linear. In line with the transactions costs argument presented in Section 7.1, we find evidence that in a neighbourhood of the unconditional mean of the premium the function is almost flat. This implies approximate unit root behaviour in the premium at these levels. On the other hand, for extreme values of the premium, the slope of the regression line is negative. Hence in these regions the premium mean reverts and, therefore, the process is globally stationary.

A good example of this behaviour is the kernel regression estimate for the USD/GBP premium. Although it is not centred on zero, the function displays a fairly strong negative slope for extreme values of the premium, whilst the function is virtually flat in the centre of its domain. A comparison with the linear ADF function shows that the conventional methodology understates the degree of mean reversion for extreme values of the premium whilst the converse holds for values of the process close to the unconditional mean.

The above conclusions are quite robust across exchange rates. In many cases, however, the kernel regression estimates are ‘one-sided’ in that the forward premium is predominantly positive or negative. Also, the non-parametric estimate for the DEM/GBP contains a region over which the function is actually positively sloped. This aberration is due to a single, very large, instance of a large negative premium which subsequently becomes even more negative.

These results yield insights into the conclusions drawn by both Crowder (1994) and

Baillie and Bollerslev (1994).¹⁷ The former asserts that the forward premium is first order integrated. In the scenario described by our kernel regression estimates this conclusion may be erroneously drawn as over a fairly large range of premium values, a range in which most of the observed data falls, the premium does actually appear to be an approximately $I(1)$ process. Hence, as the testing methodology employed by Crowder (1994) has little power against the non-linear mean reversion apparent in the data he concludes that the premia are $I(1)$ when in fact they are globally mean reverting. Baillie and Bollerslev (1994) describe the forward premium as a fractionally integrated process with fractional differencing parameter greater than one half. This implies non-stationarity in the premium although the process mean reverts. The defining characteristic of this type of process is that reversion to the mean is an extremely lengthy process. Again this can be rationalised in terms of our description of the data by noting that, following a fairly large shock to the premium, mean reversion will at first be relatively fast but will then slow to almost zero as we enter the region of the premium around the unconditional mean. Hence one might conclude that the premium follows a long memory process when, for example, the transactions costs argument expounded earlier yields a far more economically straightforward description of the non-linearities in the premium process.

TAR Estimations

The TAR models build on the results of the previous subsection. Specifically, through the kernel regressions, we have demonstrated that there is non-linearity in the mean reversion of the forward premia of our six currency pairs. The modified TAR models, explained in Section 7.2.2, are essentially threshold ADF formulations such that the time series dynamics and mean reversion of the premia are dependent

¹⁷When applied to the data used by Baillie and Bollerslev (1994) and Crowder (1994) the kernel regression estimator gave results similar to those in Figures 7.1 to 7.6, although less precise due to the smaller number of observations.

on the level of lagged premia. Hence, as our kernel regression results suggest, we permit the possibility that mean reversion is stronger for more extreme premia and is largely attenuated when the premium is close to its unconditional mean.

One drawback of this approach is that, for a fully parametric description of the non-linearities in mean reversion, one needs a sufficient amount of extreme observations in order to correctly specify the outer regimes. Unfortunately, for the DEM/GBP, CAD/GBP and SWK/GBP, this condition is not satisfied. Estimations of the modified TAR models for these currencies identified less than twenty observations (in each case) in the outer regime. As the number of autoregressive terms contained in the TAR model was around 10, this yielded very imprecise estimates of the parameters in the regime containing the extreme observations. Hence, at this point, we omit all TAR results for these currencies and focus on the USD/GBP, FRF/GBP and JPY/GBP.¹⁸

Identification and estimation of these models is carried out as outlined in Section 7.2.2. The first step in identification is to calculate the appropriate threshold variable for the TAR. As described in Section 7.2.2, this is done by carrying out a series of F -tests for non-linear behaviour using different lagged premia as the threshold variables. The model corresponding to the lag which minimises the p -value of the F -statistic is then chosen. For our forward premium data we considered the first to the third lag of the respective premium as possible threshold variables. The outcomes of the testing procedure are contained in Table 7.5.

A first point to note from Table 7.5 is that eight of the nine tests performed indicate non-linearity in forward premia. The only case in which we fail to reject the null of linearity is $d = 3$ for the JPY/GBP. Table 7.5 also indicates that the correct

¹⁸In order for their dynamics to conform with those of the TADF model given in equation (7.10), the premium data for the USD/GBP and JPY/GBP were demeaned prior to estimation.

threshold variable is the the first lag of the forward premium for both the FRF/GBP and JPY/GBP whilst the tests indicate that the third lag of the forward premium is the appropriate threshold variable for the USD/GBP data.¹⁹ One should note, however, that the use of the first lag of the premium as the threshold variable for the USD/GBP does not qualitatively alter any of the results.

The results from the three TADF specifications are given in Table 7.6.²⁰ For reasons of space we do not report the autoregressive parameters from the models, concentrating on the parameter on the lagged premium in each regime.²¹

A first point is that in all three cases the Durbin-Watson and Box-Ljung Q -statistics generally indicate a lack of residual serial correlation, the only exception being the $Q(20)$ test for the USD/GBP. Further, comparison of the regression R^2 with those for the linear ADF models presented in Table 7.3 show a significant improvement in fit, although overall explanatory power is still low. This is especially marked for the French Franc, where the R^2 can be seen to more than treble.

Analysis of the Dickey-Fuller coefficients from the inner regimes (ρ_1) demonstrates that for all three exchange rates one cannot reject the hypothesis of a unit-root in the premium. The coefficients are very small in general and for the Dollar and Franc cases are exceeded in magnitude by their respective standard errors. The coefficient for the Yen is more marginal: the robust t -statistic is insignificant at the 5% level

¹⁹For the USD/GBP data we went on to consider values for d of 4 and 5 in order to check that a delay of 3 was optimal. The result was confirmed. Also, note that if one uses the alternative procedure for identifying d by estimating the TADF model across a number of different delays and choosing the specification with the greatest likelihood, the results are identical.

²⁰For comparability with the linear specifications we retained exactly the lag length chosen by the AIC for the linear models and which are reported in Table 7.3.

²¹Note that the inference carried out on the above regressions employs heteroskedasticity robust standard errors obtained via the White (1980) procedure. As, a priori, we recognised that conditional heteroskedasticity would be a problem with these data, we re-ran all the final TADF models with a GARCH(1,1) error specification and all the qualitative results obtained and discussed above were robust to this change in specification.

but is greater in magnitude than the absolute value of the 10% critical value. Also, for the Yen only, the inner regime coefficient is more negative than that derived from the linear specification, the reverse holding for the other two currency pairs.

The Dickey-Fuller coefficients for the outer regimes show a completely different picture. The coefficients are at least an order of magnitude greater than their inner regime counterparts and in all cases are significantly smaller than zero at the 5% level. The Dollar and Yen coefficients are very close to the 1% level for the ADF test. Hence, in all cases, there is strong evidence that the premium is mean reverting in the outer regime.

Putting together the results presented above entirely corroborates the outcomes from the kernel regressions. When using a straightforward linear ADF model (with robust standard errors), one cannot reject the hypothesis of a unit root in any of the six exchange rates in our sample. This is likely to be due to the vast majority of all observations on the premium falling into the inner regime. When, however, one conducts a more careful analysis of the data one finds that, whilst in the inner regime the process does contain a unit root, the premia are globally mean reverting as one can reject the presence of a unit root in the outer regime for all estimations.²²

A graphical exposition of the impact of these results can be seen in Figures 7.7 to 7.9, which present the impulse responses of the premia to large and small shocks. All three figures demonstrate that the effect of a small shock is persistent. For large shocks, however, the premium is originally driven into the outer regime such that it

²²The superiority of the TADF models over the conventional ADF specification is also demonstrated in an examination of their respective forecasting abilities. Utilising an extra two years of spot and forward exchange rates (i.e. covering the period 1/1/86 to 31/12/87) we find that the one step ahead mean squared prediction errors from the TADF models are lower than those from the ADF specification for the USD/GBP and JPY/GBP. Further, employing the test of Diebold and Mariano (1995) for equality of forecast accuracy, demonstrates that these improvements in accuracy are significant at 1% for the JPY/GBP and at 10% for the USD/GBP whilst there is no significant difference in accuracy for the FRF/GBP.

begins to mean revert quite swiftly before dropping into the inner regime at which point the effect on the premium persists. The features of these impulse response functions are qualitatively very different from those implied by the estimated long memory models (presented in Figures 7.10 to 7.12) which display complete but slow mean reversion.

7.4 Conclusions

Recent research has cast doubt on the conventionally accepted wisdom that the foreign exchange forward premium follows a stationary stochastic process. Crowder (1994) and Baillie and Bollerslev (1994) both provide evidence that the premium follows a non-stationary process, the former concluding in favour of an $I(1)$ representation whilst the latter authors argue that premia are best represented by non-stationary, fractionally integrated processes.

Using a sample of daily observations on spot and forward rates for six currencies we conduct a finer examination of the time-series properties of the forward premium. We employ non-parametric kernel regression plus a modified fully parametric TAR model in order to examine the mean reverting properties of our premia. Our results show that premia demonstrate approximate unit root behaviour for a fairly wide range around their unconditional mean but can be seen to mean revert when the absolute value of the premium is large. Hence the premium process is globally stationary. These results imply that the standard cointegrating relationship assumed to hold between spot and forward exchange rates is mis-specified. Only when the forward premium is large in absolute value does mean reversion occur and hence spots and forwards converge. Most of the data, however, lies in a regime in which the forward premium contains a unit root. Given the $I(1)$ nature of spot and forward

rates, this result directly implies that the pair cannot be globally cointegrated with the vector (1,-1).

Our results can be reconciled with those of Crowder (1994) and Baillie and Bollerslev (1994) in the following ways. First, the vast majority of observations on the premium fall within the unit root regime for all currencies. Indeed, running a standard ADF regression for our sample indicates that all six premia have unit roots although as discussed above all are globally mean reverting. Hence, as the methodology of Crowder (1994) has little power to pick up the non-linear mean-reversion apparent in our data set he erroneously concludes that the premium process is globally non-stationary. The defining characteristic of the long memory models which Baillie and Bollerslev (1994) estimate is that mean reversion does occur although it is a very lengthy process. This can be explained through our modelling by noting that the response to a large shock to the premium will be an initial phase of mean reversion (in the outer regime) followed by complete persistence (once the premium re-enters the inner regime.) The mixture of these two effects may yield dynamics which are observationally equivalent to those from a long memory specification.

The current study has not addressed the source of the non-linearities apparent in the forward premium which, as suggested in Section 7.1, could arise in a number of ways. An obvious possibility is that the dynamics of the forward premium reflect a non-linearly dependent, time-varying risk premium. In a risk-averse setting, the forward premium may be decomposed into observed spot rate changes, a risk premium plus an expectational error. Given the strong evidence on the lack of dependence in daily spot returns, a logical source of non-linear behaviour is then the risk premium term.

An alternative perspective may be given through examination of Covered Interest Parity. CIP equates the forward premium with the differential in nominal rates of return on domestic and foreign deposits. It is possible therefore that the forward

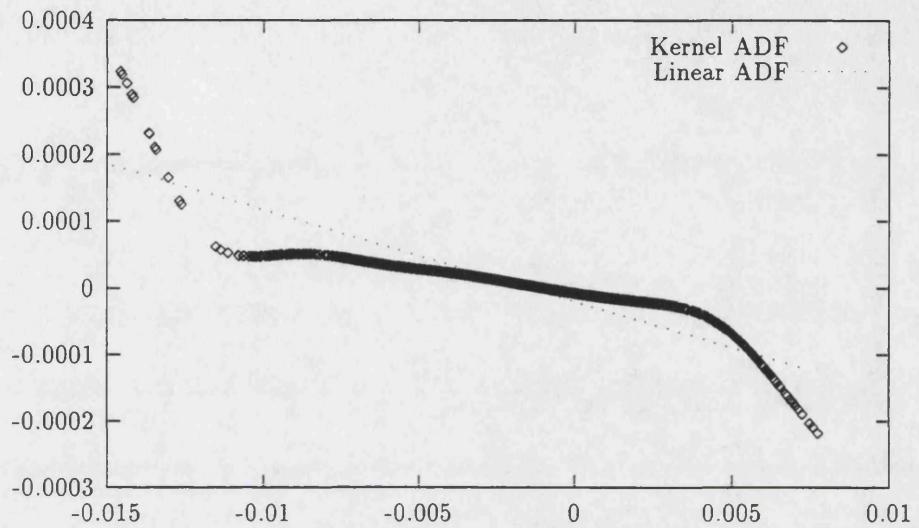
premium simply mimics non-linear behaviour in differentials. A further possibility is that frictions associated with covered interest arbitrage transmit non-linear dependence to the forward premium. Fixed, nominal transactions costs, for example, will introduce a range of deviations from CIP in which arbitrage is not profitable such that there is no tendency for premia and differentials to converge to CIP. Outside the transactions costs band, however, arbitrage strategies are run such that deviations from CIP are eroded. Assuming the differential to be a linear, covariance stationary process will then imply that non-linear behaviour in forward premia may be the result of non-linear dependence in deviations from CIP.

The preceding paragraphs point to several obvious extensions to the current study. A first extension would be to examine nominal interest rate differentials to see whether they display non-linear behaviour similar to that found in the forward premium. If non-linearities cannot be found in the differentials a next logical step would be to combine the premium and differential data to yield the deviations from covered interest parity. As stated above, trading frictions might impinge upon the profitability of arbitrage such that non-linearities are induced by the existence of, for example, small transaction costs.

On an econometric front, a further extension would be to explicitly propose and estimate a threshold cointegrating representation for spot and forward rates as the dynamics of the forward premium suggest. Specifically, we would expect a threshold error correction representation for spot and forward rates to indicate that the error correction term is significant for large absolute values of the premium (such that spot and forward rates tend to converge) but when premia are slight, no such convergence occurs and the parameter of the error correction variable is statistically indistinguishable from zero. Finally it is also possible that a more complex 'regime-switching' model for the premium is appropriate. A candidate model in this case would a reparameterisation of the Smooth Transition Autoregression (STAR) model contained in Granger and Teräsvirta (1993).

FIGURES

Figure 7.1: USD/GBP: Kernel and Linear ADF plots



TABLES

Table 7.1: Summary Statistics for the Six Forward Premium Series

Exchange Rate	Mean	Variance	Skew	Kurtosis	Q(10)	Q(20)
USD/GBP	-0.0015	0.000011	-0.46	1.04	23287	44805
DEM/GBP	-0.0049	0.000006	-0.19	0.73	22851	42756
CAD/GBP	-0.0006	0.000008	0.13	0.52	23425	44402
JPY/GBP	-0.0050	0.000005	-0.63	-0.22	17586	32633
FRF/GBP	0.0008	0.000021	1.41	3.05	20323	34916
SWK/GBP	0.0010	0.000017	2.82	14.84	19896	34677

Notes: in addition to the first four moments of the distribution the table also presents the Box-Ljung Statistics for up to 10th and 20th order serial correlation respectively.

Table 7.2: Summary Statistics for the First Difference of the Six Forward Premium Series

Exchange Rate	Mean	Variance	Skew	Kurtosis	Q(10)	Q(20)
USD/GBP	4.9×10^{-7}	2.9×10^{-7}	-0.33	6.31	510	795
DEM/GBP	-5.1×10^{-8}	2.9×10^{-7}	-0.16	6.86	216	235
CAD/GBP	2.9×10^{-7}	3.8×10^{-7}	-0.12	3.81	267	278
JPY/GBP	2.1×10^{-6}	1.8×10^{-7}	-0.22	3.82	103	113
FRF/GBP	1.2×10^{-6}	1.2×10^{-6}	-0.42	36.29	68	85
SWK/GBP	1.3×10^{-6}	1.1×10^{-6}	-3.38	85.87	105	196

Notes: in addition to the first four moments of the distribution the table also presents the Box-Ljung Statistics for up to 10th and 20th order serial correlation respectively.

Figure 7.2: DEM/GBP: Kernel and Linear ADF plots




Table 7.3: Results from the Linear ADF models for the change in the Forward Premium

Exchange Rate	ρ	$t : \rho = 0$	R^2	Lags	D.W.	Q(10)	Q(20)
USD/GBP	-0.0093	-2.19	0.154	10	1.99	10.13	42.82
DEM/GBP	-0.0131	-2.28	0.116	12	2.00	0.32	12.85
CAD/GBP	-0.0109	-2.27	0.137	4	1.99	6.36	17.67
JPY/GBP	-0.0122	-2.69	0.066	4	1.99	8.34	13.83
FRF/GBP	-0.0218	-2.21	0.038	10	1.99	0.15	18.92
SWK/GBP	-0.0231	-1.42	0.050	9	1.99	1.80	66.04

Notes: ρ represents the coefficient on the lagged premium. The reported standard errors are heteroskedasticity and serial correlation robust. The column headed *Lags* gives the number of lags in the ADF specification, chosen via the Akaike information criterion. *D.W.* is the Durbin-Watson statistic. *Q(10)* and *Q(20)* are the Box-Ljung statistics for up to 10th and 20th order serial correlation respectively.

Table 7.4: Estimates of the Fractional Differencing Parameter for Forward Premia

Exchange Rate	\hat{d}	s.e.
USD/GBP	0.215	0.046
DEM/GBP	0.147	0.045
CAD/GBP	0.190	0.046
JPY/GBP	0.213	0.051
FRF/GBP	0.258	0.046
SWK/GBP	0.236	0.045

Notes: The long memory estimates are derived from the semi parametric estimator proposed by Künsch (1987), utilising an optimal bandwidth selection procedure derived by Henry and Robinson (1996). Asymptotic standard errors are presented in the final column of the table.

Figure 7.3: CAD/GBP: Kernel and Linear ADF plots

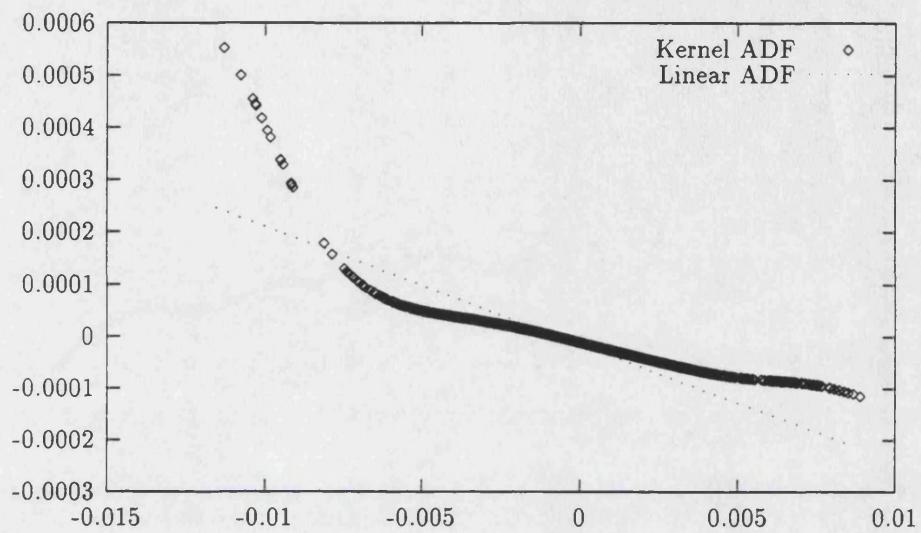


Table 7.5: *F*-Tests for the Delay Parameter in the TADF Specification

Delay Parameter	Exchange Rate		
	USD/GBP	JPY/GBP	FRF/GBP
1	2.62**	3.63**	4.75**
2	2.48**	2.35*	4.39**
3	3.06**	2.07	4.11**

Notes: The statistics in the table are *F*-statistics for threshold non-linearity derived from arranged autoregressions. The approach follows Tsay (1989). The degrees of freedom for the *F*-statistics are $p + 1$ where p is the order of the autoregression and $n - n_1 - 2p - 1$ where n_1 is the number of observations used to initialise the arranged autoregression. A * denotes significance at the 5% level and ** denotes significance at 1%.

Table 7.6: Results from the Modified TAR model for the change in the Forward Premium

Exchange Rate	ρ_1	$t : \rho_1 = 0$	ρ_2	$t : \rho_2 = 0$	R^2	D.W.	Q(10)	Q(20)
USD/GBP	-0.0009	-0.24	-0.0564	-3.02	0.176	2.01	12.10	44.94
JPY/GBP	-0.0137	-2.70	-0.4769	-3.07	0.080	2.00	8.98	14.75
FRF/GBP	-0.0050	-0.65	-0.4261	-2.88	0.126	1.98	4.31	19.68

Notes: ρ_1 and ρ_2 represent the coefficients on the lagged premium in the central and outer regimes respectively. The reported standard errors are heteroskedasticity and serial correlation robust. D.W. is the Durbin-Watson statistic. Q(10) and Q(20) are the Box-Ljung statistics for up to 10th and 20th order serial correlation respectively.

FIGURES

Figure 7.1: USD/GBP: Kernel and Linear ADF plots

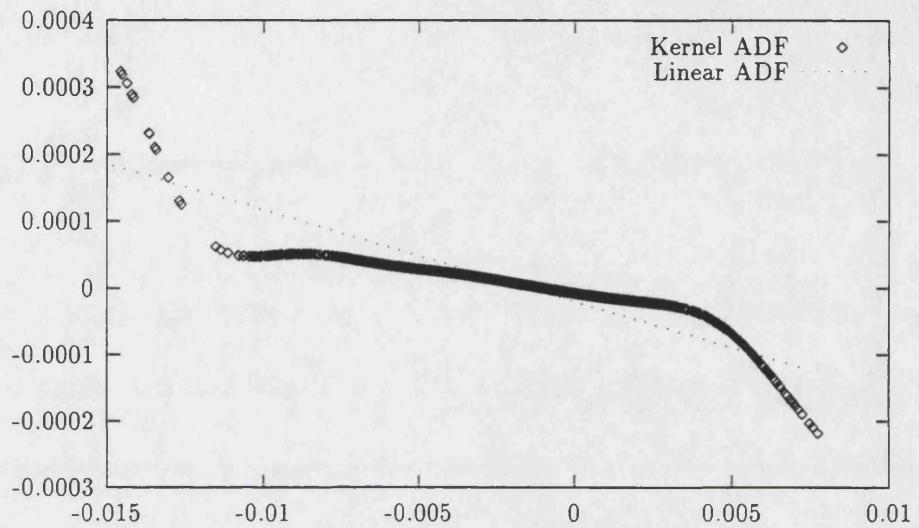


Figure 7.2: DEM/GBP: Kernel and Linear ADF plots

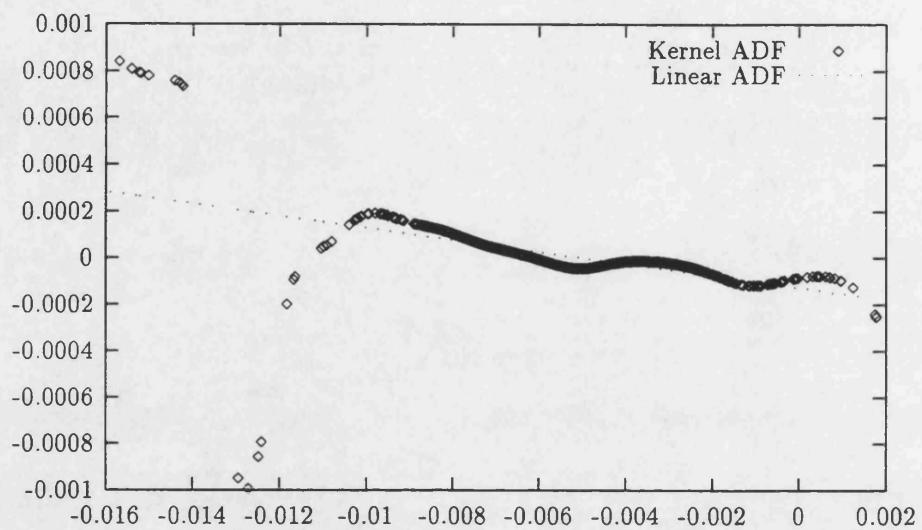


Figure 7.3: CAD/GBP: Kernel and Linear ADF plots

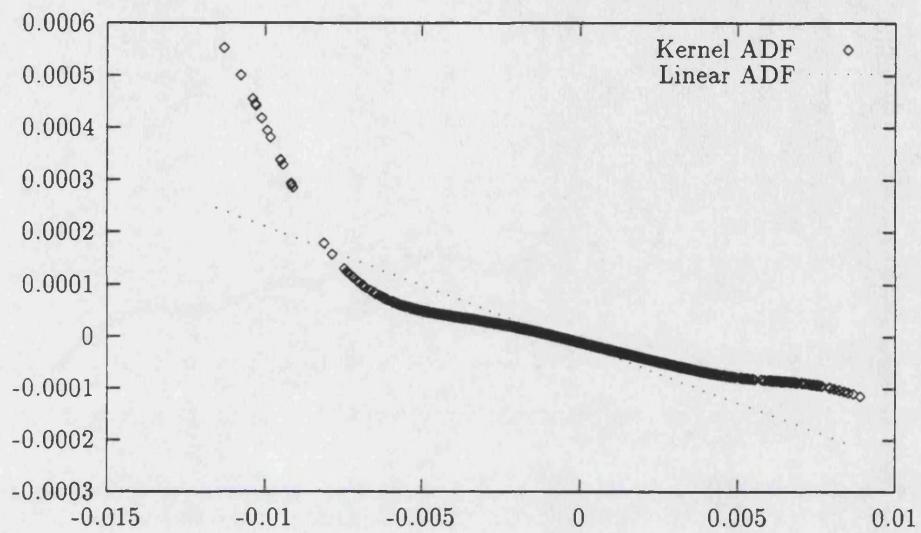


Figure 7.4: JPY/GBP: Kernel and Linear ADF plots

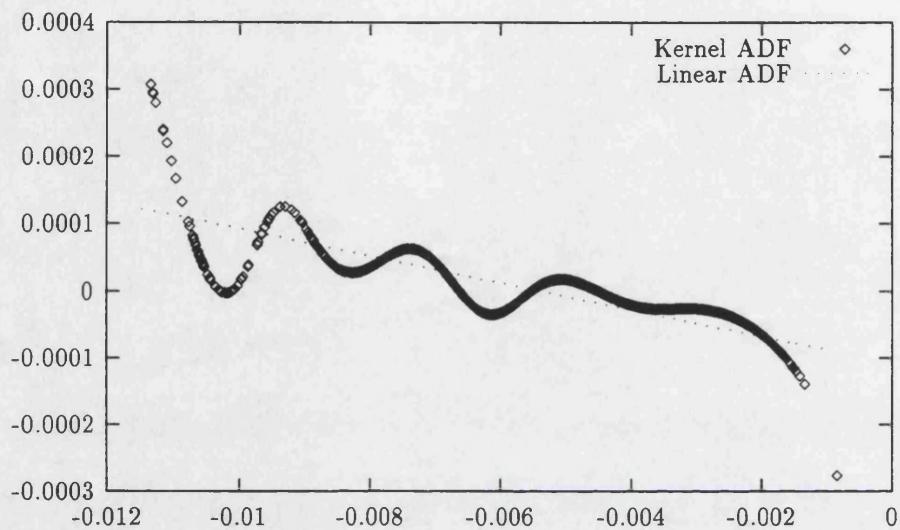


Figure 7.5: FRF/GBP: Kernel and Linear ADF plots

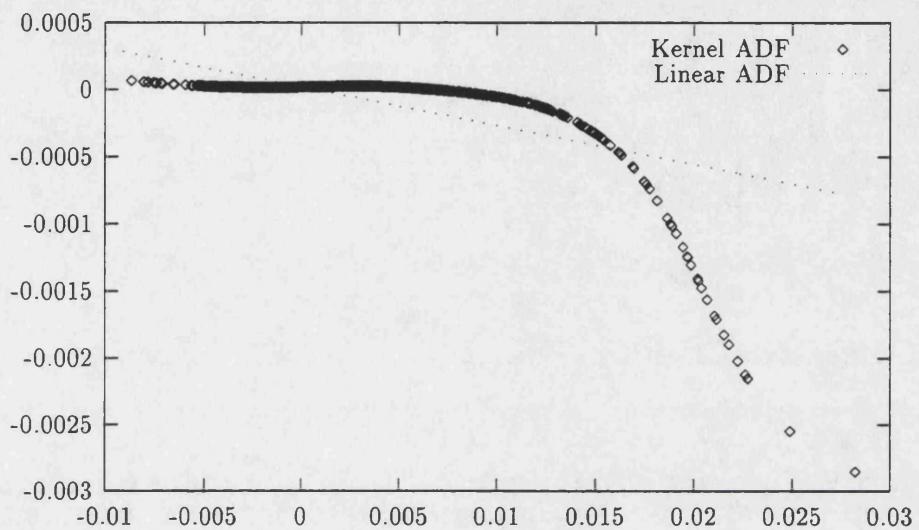


Figure 7.6: SWK/GBP: Kernel and Linear ADF plots

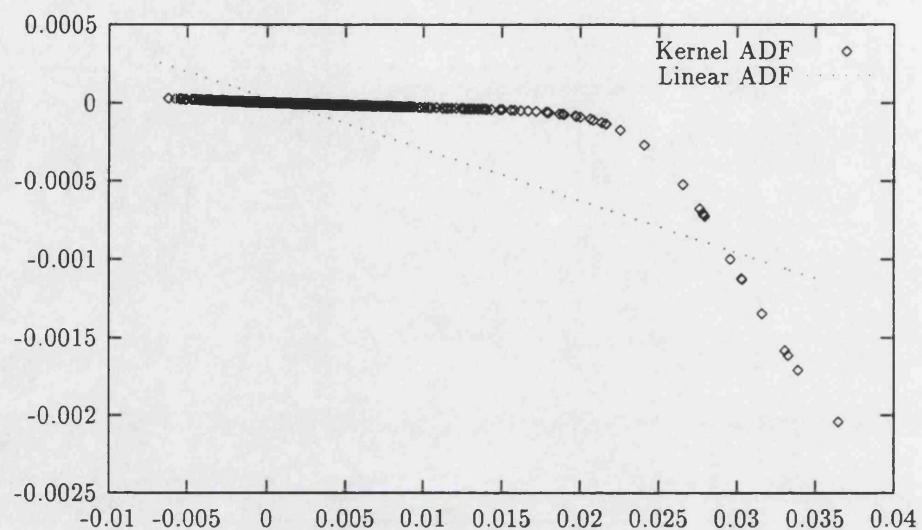


Figure 7.7: USD/GBP TADF Model: Responses to 1 and 5 s.d. shocks




Figure 7.8: JPY/GBP TADF Model: Responses to 1 and 5 s.d. shocks




Figure 7.9: FRF/GBP TADF Model: Responses to 1 and 5 s.d. shocks

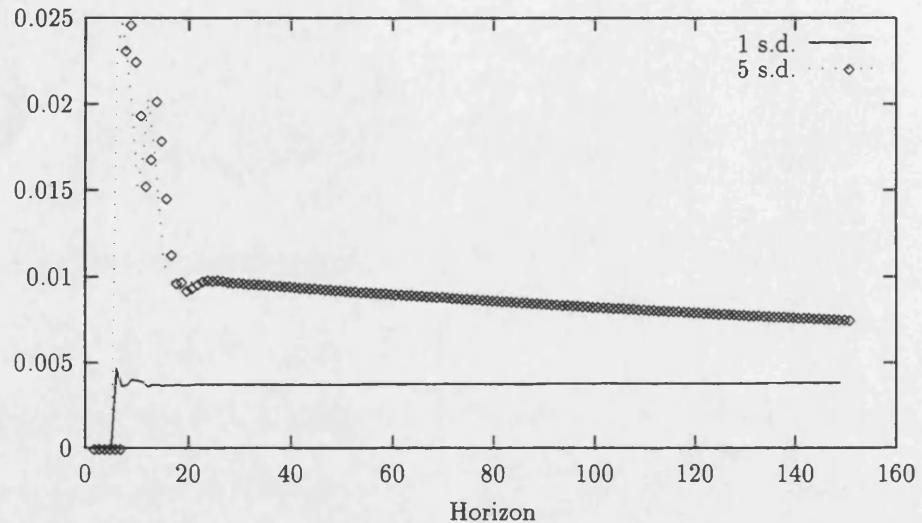
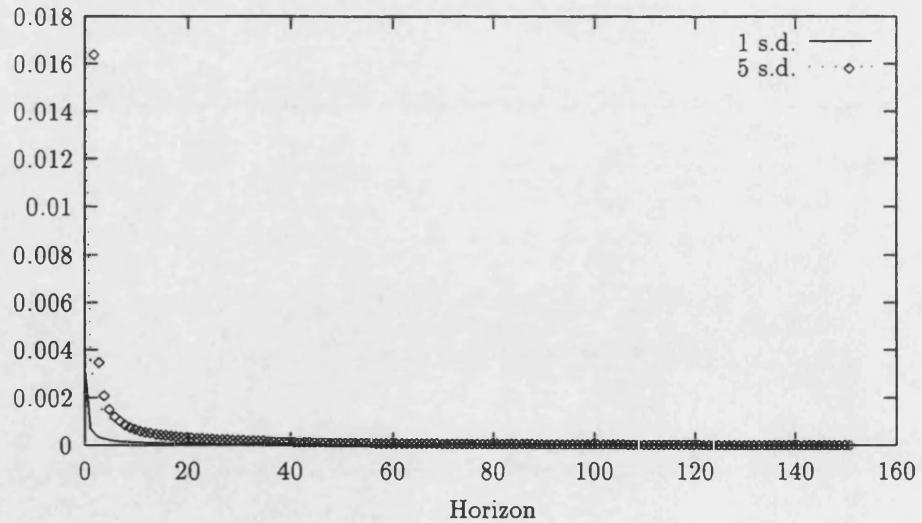


Figure 7.10: USD/GBP Long Memory Model: Responses to 1 and 5 s.d. shocks



Chapter 8

Conclusions

The empirical analysis of exchange rate behaviour and determination has a long history in applied economics. The collection of studies contained in the previous chapters contributes to this area by providing a number of examinations of the behaviour of exchange rates sampled at high-frequencies. Four main topics are covered, these being: the effect of the microstructure of the FX market on the properties of transactions and prices, the behaviour of intra-day exchange rate volatility, the effect of public information announcements on exchange rates and, finally, the relationship between spot and forward exchange rates. In this concluding section the results of the preceding studies are summarised and directions for further work are indicated.

The study of the impact of FX market microstructure on trading and liquidity is an area which is still very much in its youth. The lack of consolidated transactions databases and realistic theoretical models of the intra-day FX market have, thus far, inhibited work in this area. Nonetheless, it is to this strand of research that

Figure 7.11: JPY/GBP Long Memory Model: Responses to 1 and 5 s.d. shocks

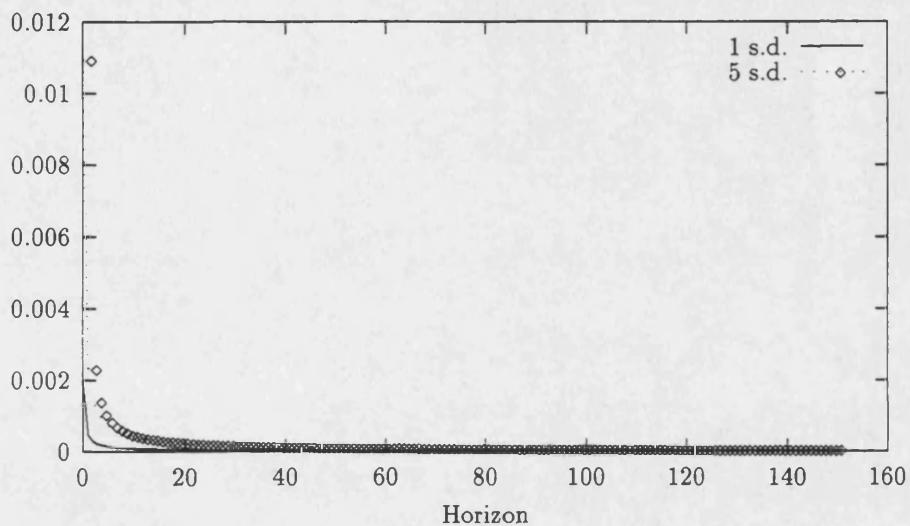



Figure 7.12: FRF/GBP Long Memory Model: Responses to 1 and 5 s.d. shocks



Chapters 2 and 3 belong. Using a short, transactions based data set from the Reuters inter-dealer broking system known as D2000-2, they present four separate empirical exercises.

The first of these is an analysis of the determinants of D2000-2 quote revisions for the bid and ask sides of the market separately. Results indicate that three main factors cause quotes to be revised: transactions, quote revisions on the opposite side of the market and spreads. The effects of all three of these impacts are in line with the hypotheses in Chapter 2. The second exercise undertaken in Chapter 2 investigated the source of negative autocorrelation in quote revisions. Results suggest that this autocorrelation is due to the existence of a 'thin' limit order book on D2000-2 i.e. there are few orders lying just behind the best quotation on a given side of the market, such that after an exhaustive transaction, for example, quotes jump a long way from their original value and this jump is subsequently reversed. The final exercise undertaken in Chapter 2 examined the linkage between spreads and volatility. In line with the results from previous work and intuitions from standard inventory control and asymmetric information models of dealer behaviour, spreads are found to be strongly positively affected by volatility. The final empirical microstructure exercise undertaken is presented in Chapter 3. This study examines the existence of informational asymmetries between intra-day FX dealers. Using the VAR approach developed in Hasbrouck (1991a) and Hasbrouck (1991b) we find consistent evidence that asymmetric information problems are present on the inter-dealer market. Quantitatively, the results suggest that around one half of the inside spread is charged to guard against the possibility of informed trade and that trading activity reveals, on average, 30% of all information relevant to the DEM/USD.

The shortcoming of the above studies, along with most other work in empirical FX microstructure, is the small coverage of the data set. The D2000-2 data covers a multi-dealer system for only seven hours whilst the other transactions based FX data set which has been studied, that of Lyons (1995), covers a single dealer for one

Chapter 8

Conclusions

The empirical analysis of exchange rate behaviour and determination has a long history in applied economics. The collection of studies contained in the previous chapters contributes to this area by providing a number of examinations of the behaviour of exchange rates sampled at high-frequencies. Four main topics are covered, these being: the effect of the microstructure of the FX market on the properties of transactions and prices, the behaviour of intra-day exchange rate volatility, the effect of public information announcements on exchange rates and, finally, the relationship between spot and forward exchange rates. In this concluding section the results of the preceding studies are summarised and directions for further work are indicated.

The study of the impact of FX market microstructure on trading and liquidity is an area which is still very much in its youth. The lack of consolidated transactions databases and realistic theoretical models of the intra-day FX market have, thus far, inhibited work in this area. Nonetheless, it is to this strand of research that

Chapters 2 and 3 belong. Using a short, transactions based data set from the Reuters inter-dealer broking system known as D2000-2, they present four separate empirical exercises.

The first of these is an analysis of the determinants of D2000-2 quote revisions for the bid and ask sides of the market separately. Results indicate that three main factors cause quotes to be revised: transactions, quote revisions on the opposite side of the market and spreads. The effects of all three of these impacts are in line with the hypotheses in Chapter 2. The second exercise undertaken in Chapter 2 investigated the source of negative autocorrelation in quote revisions. Results suggest that this autocorrelation is due to the existence of a 'thin' limit order book on D2000-2 i.e. there are few orders lying just behind the best quotation on a given side of the market, such that after an exhaustive transaction, for example, quotes jump a long way from their original value and this jump is subsequently reversed. The final exercise undertaken in Chapter 2 examined the linkage between spreads and volatility. In line with the results from previous work and intuitions from standard inventory control and asymmetric information models of dealer behaviour, spreads are found to be strongly positively affected by volatility. The final empirical microstructure exercise undertaken is presented in Chapter 3. This study examines the existence of informational asymmetries between intra-day FX dealers. Using the VAR approach developed in Hasbrouck (1991a) and Hasbrouck (1991b) we find consistent evidence that asymmetric information problems are present on the inter-dealer market. Quantitatively, the results suggest that around one half of the inside spread is charged to guard against the possibility of informed trade and that trading activity reveals, on average, 30% of all information relevant to the DEM/USD.

The shortcoming of the above studies, along with most other work in empirical FX microstructure, is the small coverage of the data set. The D2000-2 data covers a multi-dealer system for only seven hours whilst the other transactions based FX data set which has been studied, that of Lyons (1995), covers a single dealer for one

U.S. trading week. It is clear, therefore, that more comprehensive FX transactions data is required in order for this area of research to progress. Thankfully then, we have obtained very recently, access to a far more extensive data set. Reuters have made available to us the entire feed from their D2000-2 trading system for a number of weeks. This feed yields data on all limit orders entered into the system (not only the best limit order on each side of the market,) plus all transactions and each of the data items is timed to the hundredth of a second. As yet, we have not started the analysis of these data, but it is obvious that such an improvement in data quality will allow us to re-evaluate the results of previous studies and also to explore as yet untested microstructure hypotheses for the foreign exchange market.

The second main research area to which this thesis has contributed is the behaviour of intra-day exchange rate volatility, with our analysis of volatility contained in Chapters 4 and 5. At the time when these studies were being written, the analysis of intra-day FX volatility was centred, more or less exclusively, on the standard GARCH and AR(1)-SV models. There were very few treatments of the impact of external factors on volatility and divergences from the ARMA-type dynamics implicit in the above formulations were fewer still. Chapter 4 addresses the former of these points. It proposes an extended Stochastic Volatility specification which contains two extra components. The first of these is a component allowing deterministic seasonality, which has been shown recently, in for example Dacorogna, Müller, Nagler, Olsen, and Pictet (1993) and Andersen and Bollerslev (1997b), to be a pronounced feature of intra-day financial market volatility. Secondly, the baseline model is extended to assess the impact of U.S. macroeconomic announcements on volatility. Results demonstrate that both of these additional factors are extremely important in short-run volatility determination. The U.S. Payroll Employment news, for example, causes volatility to increase to over 1000% of its normal level. Chapter 5 moves outside the standard GARCH/SV paradigm to conduct an investigation of the long memory structure of volatility. Volatility is characterised as covariance stationary and long range dependent, in line with certain observations from Chap-

ter 4. Further, the common result from GARCH and SV representations, that the volatility process contains a unit root, is strongly rejected and shown to be due to mis-specification of these models in that they do not permit long range dependence.

There are several obvious extensions to the analyses of Chapters 4 and 5. First one would like to evaluate the effects of the macroeconomic announcements of countries other than the U.S. on exchange rates other than the DEM/USD. Also, the Mixture of Distribution Models (MODM's) presented in the Introduction provide another set of possible studies. These models imply that changes in volatility are caused by alterations in the flow of public information into the market. Hence once should be able to detect covariation between volatility and broader measures of information flow. Further, the MODM of Tauchen and Pitts (1983) results in both volume and volatility inheriting the dependencies in information flow. This implies that time-series of volumes might also exhibit long memory. Finally, a multivariate examination of a vector of volatilities could be performed through a fractional cointegration analysis. These latter two points are the focus on ongoing research.

Chapter 6 also links with Chapter 4 in that it provides an analysis of the impact of macroeconomic 'news' on exchange rate levels. It examines the effect of U.S. and German announcements on a three year span of DEM/USD quotations, sampled at a five minute frequency. The study provides evidence that the 'news' contained in announcements has a significant impact on exchange rates but that this impact is drowned very quickly in subsequent fluctuations. This result ties in with those from studies using exchange rate data sampled less finely, which show that no discernible effects of 'news' are apparent. Further results in Chapter 6 demonstrate the efficiency of the FX market with respect to these announcements and also show that the response to German 'news' is slower than that to U.S. information. This result is likely to be due to the fact that, unlike U.S. releases, German announcements are unscheduled.

The analysis in Chapter 6 explains the direction in which ‘news’ pushes the exchange rate with reference to a ‘reaction function hypothesis.’ This hypothesis states that the market’s interpretation of ‘news’ is driven by the likely interest rate response of monetary authorities e.g. higher than expected U.S. real activity induces the Fed to increase interest rates to head off future inflation and the market anticipates this, causing a Dollar appreciation. If this hypothesis holds, therefore, one should also be able to detect an impact from these ‘news’ data on interest rates. This is a topic which we intend to examine in the future.

The final study contained in the thesis, Chapter 7, looks at the relationship between spot and forward exchange rates. Using a sample of daily spots and forwards, for six currency pairs, it presents a close examination of the dynamics of the foreign exchange forward premium. Recent work on forward premia has indicated that these processes are non-stationary (Crowder (1994) and Baillie and Bollerslev (1994),) a result which runs contrary to standard intuition. Chapter 7 presents two sets of empirical estimations which focus on the stationarity of premia. The first is a set of kernel regressions which correspond to non-parametric analogues of the conventional ADF test for stationarity. These regressions indicate that forward premia are globally mean reverting, although this mean reversion is non-linear in nature. Specifically, premia display approximate unit root behaviour when close to their unconditional mean, however when the absolute value of the premia is large they mean revert quite strongly. These results are formalised in the second set of estimations which are fully parametric threshold ADF specifications.

The non-linear behaviour discovered in forward premia opens a number of avenues for future research. A first study which could be undertaken is to examine the time-series of interest rate differentials which, by covered interest parity, should have identical dynamics to forward premia. If differentials are found, however, to follow linear $I(0)$ processes then it may be the case that the non-linearities in premia are the result of trading frictions in covered interest arbitrage. Constructing

the deviations from covered interest parity and examining them for non-linearities would then provide a direct test of this hypothesis. Finally, the non-linearities in the mean reversion of the forward premium would suggest that spot and forward exchange rates could be well represented by a threshold cointegrating framework.

Bibliography

ADMATI, A., and P. PLEIDERER (1988): "A Theory of Intra-Day Patterns: Volume and Price Variability," *Review of Financial Studies*, 1, 3–40.

AGGARWAL, R., S. MOHANTY, and F. SONG (1995): "Are Survey Forecasts of Macroeconomic Variables Rational?," *Journal of Business*, 68, 99–119.

AGGARWAL, R., and D. SCHIRM (1992): "Balance of Trade Announcements and Asset Prices: Influence on Equity Prices, Exchange Rates and Interest Rates," *Journal of International Money and Finance*, 11, 80–95.

ALMEIDA, A. (1997): "The Reaction of Exchange Rates to Monetary Policy Signals," Working Paper, Financial Markets Group, London School of Economics.

ALMEIDA, A., and C. GOODHART (1997): "Afecta la Adopción de Objectivos Directos de Inflación al Comportamiento de los Bancos Centrales," in *La Política Monetaria y la Inflación en España*. Alianza Editorial, Madrid.

ALMEIDA, A., C. GOODHART, and R. PAYNE (1997): "The Effects of Macroeconomic 'News' on High-Frequency Exchange Rate Behaviour," Discussion Paper 258, Financial Markets Group, London School of Economics.

ANDERSEN, T., and T. BOLLERSLEV (1997a): "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High-Frequency Returns," *Journal Of Finance*, 52, 975–1005.

ANDERSEN, T., and T. BOLLERSLEV (1997b): "Intraday Seasonality and Volatility Persistence in Financial Markets," *Journal of Empirical Finance*, 4, 115–158.

BACKUS, D., A. GREGORY, and C. TELMER (1993): "Accounting for Forward Rates in Markets for Foreign Currency," *Journal of Finance*, 48, 1887–1908.

BAILLIE, R., and T. BOLLERSLEV (1991): "Intra-Day and Inter-Market Volatility in Foreign Exchange Markets," *Review of Economic Studies*, 58, 565–586.

BAILLIE, R., and T. BOLLERSLEV (1994): "The Long Memory of the Forward Premium," *Journal of International Money and Finance*, 13, 565–571.

BAILLIE, R., T. BOLLERSLEV, and H. MIKKELSEN (1996): "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity," *Journal Of Econometrics*, 74, 3–30.

BANK FOR INTERNATIONAL SETTLEMENTS (1993): *Central Bank Survey of Foreign Exchange Market Activity in April 1992*. B.I.S., Basle.

BARCLAY, M., R. LITZENBERGER, and J. WARNER (1990): "Private Information, Trading Volume and Stock-Return Variances," *Review of Financial Studies*, 3, 233–253.

BELTRÃO, K., and P. BLOOMFIELD (1987): "Determining the bandwidth of a kernel spectral estimate," *Journal of Time Series Analysis*, 8, 21–38.

BERRY, T., and K. HOWE (1994): "Public Information Arrival," *Journal of Finance*, 49, 1331–1346.

BESSEMBINDER, H. (1994): "Bid-Ask Spreads in the Interbank Foreign Exchange Market," *Journal of Financial Economics*, 35, 317–348.

BIAIS, B., P. HILLION, and C. SPATT (1995): "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse," *Journal of Finance*, 50, 1655–1689.

BILSON, J. (1981): "The 'Speculative Efficiency' Hypothesis," *Journal of Business*, 54, 435–452.

BOLLERSLEV, T. (1986): "Generalised Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307–327.

BOLLERSLEV, T., R. CHOU, and K. KRONER (1992): "ARCH modeling in finance," *Journal of Econometrics*, 52, 5–59.

BOLLERSLEV, T., and I. DOMOWITZ (1993): "Trading Patterns and Prices in the Interbank Foreign Exchange Market," *Journal of Finance*, 48, 1421–1444.

BOLLERSLEV, T., and M. MELVIN (1994): "Bid-ask Spreads and Volatility in the Foreign Exchange Market: An Empirical Analysis," *Journal of International Economics*, 36, 355–372.

BRENNER, R., and K. KRONER (1995): "Arbitrage, Cointegration and Testing the Unbiasedness Hypothesis in Financial Markets," *Journal of Financial and Quantitative Analysis*, 30, 23–42.

BROCK, W., J. LAKONISHOK, and B. LEBARON (1992): "Simple Technical Trading rules and the Stochastic Properties of Stock Returns," *Journal of Finance*, 47, 1731–1764.

CAMPBELL, J., A. LO, and C. MACKINLAY (1997): *The Econometrics of Financial Markets*. Princeton University Press, Princeton, New Jersey.

CHAKRAVARTY, S., and C. HOLDEN (1994): "An Integrated Model of Market and Limit Orders," Unpublished Paper.

CHANG, Y., and S. TAYLOR (1996): "Information Arrivals and Intra-day Exchange Rate Volatility," Working Paper, Lancaster University.

CHESNEY, M., and L. SCOTT (1989): "Pricing European Currency Options: A Comparison of the Modified Black-Scholes Model and a Random Variance Model," *Journal of Financial and Quantitative Analysis*, 24, 267–284.

CLARIDA, R., and M. GERTLER (1996): "How the Bundesbank Conducts Monetary Policy," NBER Working Paper 5581.

CLARK, P. (1973): "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices," *Econometrica*, 41, 135–155.

COMTE, F., and C. HARDOUIN (1995): "Regression on log-regularized periodogram for fractional models at low frequencies," Discussion Paper, CREST.

COPELAND, T., and D. GALAI (1983): "Information Effects on the Bid-Ask Spread," *Journal of Finance*, 38, 1457–1469.

CORNELL, B. (1982): "Money Supply Announcements, Interest Rates and Foreign Exchange," *Journal of International Money and Finance*, 1(2), 201–208.

CROWDER, W. (1994): "Foreign Exchange Market Efficiency and Common Stochastic Trends," *Journal of International Money and Finance*, 13, 551–564.

CURCIO, R., C. GOODHART, D. GUILLAUME, and R. PAYNE (1997): "Do Technical Trading Rules Generate Profits? Conclusions from the Intra-Day Foreign Exchange Market," *International Journal of Finance and Economics*, forthcoming.

DACOROGNA, M., U. MÜLLER, R. NAGLER, R. OLSEN, and O. PICTET (1993): "A Geographical Model for the Daily and Weekly Seasonal Volatility in the Foreign Exchange Market," *Journal of International Money and Finance*, 12, 413–438.

DE JONG, F., T. NIJMAN, and A. RÖELL (1995): "Price Effects of Trading and Components of the Bid-Ask Spread on the Paris Bourse," Discussion Paper 207, Financial Markets Group, London School of Economics.

DEGENNARO, R., and R. SHRIEVES (1995): "Public Information Releases, Private Information Arrival, and Volatility in the Foreign Exchange Market," Unpublished Paper.

DICKEY, D., and W. FULLER (1979): "Distributions of the estimators for autoregressive time series with a unit root," *Journal of the American Statistical Association*, 74, 427-431.

DIEBOLD, F., and R. MARIANO (1995): "Comparing Predictive Accuracy," *Journal of the American Statistical Association*, 13, 253-263.

DIEBOLD, F., and M. NERLOVE (1989): "Unit roots in economic time series: a selective survey," *Advances in Econometrics*, 8, 3-69.

EASLEY, D., and M. O'HARA (1987): "Price, Trade Size and Information in Securities Markets," *Journal of Financial Economics*, 19, 69-90.

EDERINGTON, L., and J. LEE (1993): "How Markets Process Information: News Releases and Volatility," *Journal of Finance*, 48(4), 1161-1191.

EDERINGTON, L., and J. LEE (1995): "The Short-Run Dynamics of the Price Adjustment to New Information," *Journal of Financial and Quantitative Analysis*, 30(1), 117-134.

EDISON, H. (1997): "The Reaction of Exchange Rates and Interest Rates to News Releases," *International Journal of Finance and Economics*, 2, 87-100.

ELLIOT, G., J. STOCK, and T. ROTHENBERG (1994): "Efficient tests of an autoregressive unit root," preprint.

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., and J. FRANKEL (1984): "Why do Interest Rates React to Money Announcements: An Answer from the Foreign Exchange Market," *Journal of Monetary Economics*, 13, 31-39.

ENGLE, R. (1982): "Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation," *Econometrica*, 50, 987-1007.

ENGLE, R., and T. BOLLERSLEV (1986): "Modelling the Persistence of Conditional Variances," *Econometric Reviews*, 5, 1–50.

ENGLE, R., and C. GRANGER (1987): "Co-integration and Error Correction: Representation, Estimation and Testing," *Econometrica*, 55, 251–276.

ENGLE, R., D. LILIEN, and R. ROBINS (1987): "Estimating Time Varying Risk Premia in the Term Structure: the ARCH-M model," *Econometrica*, 55, 391–407.

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., and K. LEWIS (1993): "Trends in Excess Returns in Currency and Bond Markets," *European Economic Review*, 37, 1005–1019.

FAMA, E. (1984): "Forward and Spot Exchange Rates," *Journal of Monetary Economics*, 14, 319–338.

FLOOD, M. (1991): "Microstructure Theory and the Foreign Exchange Market," *Federal Reserve Bank of St. Louis Review*, 73, 52–70.

FOX, R., and M. TAQQU (1986): "Large sample properties of parameter estimates for strongly dependent stationary Gaussian time series," *Annals of Statistics*, 14, 517–132.

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.

FRENKEL, J. (1976): "A Monetary Approach to the Exchange Rate: Doctrinal Aspects and Empirical Evidence," *Scandinavian Journal of Economics*, 78, 200–224.

FROOT, K., and J. FRANKEL (1989): "Forward Discount Bias: Is it an Exchange Risk Premium?," *Quarterly Journal of Economics*, 104, 139–161.

GEWEKE, J., and S. PORTER-HUDAK (1983): "The estimation and application of long memory time series models," *Journal of Time Series Analysis*, 4, 221-238.

GHYSELS, E., C. GOURIÉROUX, and J. JASIAK (1995): "Trading Patterns, Time Deformation and Stochastic Volatility in Foreign Exchange Markets," Unpublished Paper.

GHYSELS, E., A. HARVEY, and E. RENAULT (1996): "Stochastic Volatility," in *Handbook of Statistics*, vol. 14.

GIL-ALAÑA, L., and P. ROBINSON (1995): "Testing of unit root and other non-stationary hypotheses in macroeconomic time series," mimeo, LSE.

GLOSTEN, L., and P. MILGROM (1985): "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71-100.

GOODHART, C. (1989): "'News' and the Foreign Exchange Market," Pamphlet, Manchester Statistical Society.

GOODHART, C. (1996): "Why do the Monetary Authorities Smooth Interest Rates," Special Paper 81, Financial Markets Group, London School of Economics.

GOODHART, C., and L. FIGLIUOLI (1991): "Every Minute Counts in Financial Markets," *Journal of International Money and Finance*, 10, 23-52.

GOODHART, C., and L. FIGLIUOLI (1992): "The Geographical Location of the Foreign Exchange Market: A Test of an 'Islands' Hypothesis," *Journal of International and Comparative Economics*, 1, 13-27.

GOODHART, C., S. HALL, S. HENRY, and B. PESARAN (1993): "News Effects in a High-Frequency Model of the Sterling-Dollar Exchange Rate," *Journal of Applied Econometrics*, 8, 1-13.

GOODHART, C., T. ITO, and R. PAYNE (1996): "One Day in June, 1993: A Study of the Workings of Reuters D2000-2 Electronic Foreign Exchange Trading

System," in *The Microstructure of Foreign Exchange Markets*, ed. by J. Frankel, G. Galli, AND A. Giovannini. University of Chicago Press.

GOODHART, C., P. McMAHON, and Y. NGAMA (1997): "Why Does the Spot-Forward Discount Fail to Predict Changes in Future Spot Rates," *International Journal of Finance and Economics*, 2(2), 121–130.

GOODHART, C., and M. O'HARA (1997): "High-Frequency Data in Financial Markets: Issue and Applications," *Journal of Empirical Finance*, 4, 73–114.

GOODHART, C., and R. PAYNE (1996): "Microstructural Dynamics in a Foreign Exchange Electronic Broking System," *Journal of International Money and Finance*, 15(6), 829–852.

GRANGER, C. (1980): "Long memory relationships and the aggregation of dynamic models," *Journal of Econometrics*, 14, 227–238.

GRANGER, C., and R. JOYEUX (1980): "An introduction to long memory time series models," *Journal of Time Series Analysis*, 1, 15–29.

GRANGER, C., and T. TERÄSVIRTA (1993): *Modelling Nonlinear Economic Relationships*. Oxford University Press, Oxford.

GUILLAUME, D. (1995): "On the intra-day performance of GARCH processes," mimeo., Financial Markets Group, London School of Economics.

GUILLAUME, D., M. DACOROGNA, R. DAVÉ, U. MÜLLER, R. OLSEN, and O. PICTET (1997): "From the Bird's Eye to the Microscope: A Survey of New Stylized Facts of the Intra-Daily Foreign Exchange Market," *Finance and Stochastics*, 1, 31–39.

HAKKIO, C., and D. PEARCE (1985): "The Reaction of Exchange Rates to Economic News," *Economic Inquiry*, 23, 621–635.

HAKKIO, C., and M. RUSH (1989): "Market Efficiency and Co-integration: An Application to the Sterling and Deutschemark Exchange Rates," *Journal of International Money and Finance*, 8, 75-88.

HAMILTON, J. (1994): *Time Series Analysis*. Princeton University Press, Princeton, New Jersey.

HÄRDLE, W. (1990): *Applied Nonparametric Regression*. Cambridge University Press, Cambridge.

HARDOUVELIS, G. (1988): "Economic News, Exchange Rates and Interest Rates," *Journal of International Money and Finance*, 7, 23-35.

HARRIS, E. (1995a): "The Employment Report and the Dollar," in *Current Issues in Economics and Finance*, vol. 1. Federal Reserve Bank of New York.

HARRIS, E., and N. ZABKA (1995): "The Employment Report and the Dollar," *Current Issues in Economics and Finance*, 1(8), 1-6.

HARRIS, L. (1995b): "Optimal Dynamic Order Submission Strategies in some Stylized Trading Problems," Working Paper, University of Southern California.

HARTMANN, P. (1996): "Trading Volumes and Transaction Costs in the Foreign Exchange Market: Evidence from Daily Dollar-Yen Spot Data," Discussion Paper 232, Financial Markets Group, London School of Economics.

HARVEY, A. (1989): *Forecasting, Structural Time-Series Models and the Kalman Filter*. Cambridge University Press, Cambridge.

HARVEY, A. (1993): "Long Memory in Stochastic volatility," mimeo, London School of Economics.

HARVEY, A., and S. KOOPMAN (1993): "Forecasting Hourly Electricity Demand Using Time-Varying Splines," *Journal of the American Statistical Association*, 88, 1228-1236.

HARVEY, A., S. KOOPMAN, and M. RIANI (1995): "The Modelling and Seasonal Adjustment of Weekly Observations," Discussion Paper EM/95/284, STICERD, London School of Economics.

HARVEY, A., E. RUIZ, and N. SHEPHARD (1994): "Multivariate Stochastic Variance Models," *Review of Economic Studies*, 61, 247–264.

HARVEY, C., and R. HUANG (1991): "Volatility in the Currency Futures Market," *Review of Financial Studies*, 4(3), 543–569.

HASBROUCK, J. (1988): "Trades, Quotes, Inventories and Information," *Journal of Financial Economics*, 22, 229–252.

HASBROUCK, J. (1991a): "Measuring the Information Content of Stock Trades," *Journal of Finance*, 46(1), 179–206.

HASBROUCK, J. (1991b): "The Summary Informativeness of Stock Trades: An Econometric Analysis," *Review of Financial Studies*, 4(3), 571–595.

HASBROUCK, J. (1993): "Assessing the Quality of a Security Market: A New Approach to Transaction Cost Measurement," *Review of Financial Studies*, 6, 191–212.

HEDVALL, K., and J. NIEMEYER (1994): "Order Flow Dynamics: Evidence from the Helsinki Stock Exchange," Working Paper, Swedish School of Economics and Business Administration.

HENRY, M., and R. PAYNE (1997): "An Investigation of Long Range Dependence in Intra-day Foreign Exchange Rate Volatility," Discussion Paper 264, Financial Markets Group, London School of Economics.

HENRY, M., and P. ROBINSON (1996): "Bandwidth choice in Gaussian semiparametric estimation of long range dependence," in *Athens Conference on Applied Probability and Time Series Analysis*, ed. by P.M.Robinson, AND M.Rosenblatt, vol. 2. Springer Verlag, In memory of E.J. Hannan.

HO, T., and H. STOLL (1983): "The Dynamics of Dealer Markets under Competition," *Journal of Finance*, 38(4), 1053–1074.

HODRICK, R. (1987): *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*. Harwood, Chur.

HOFFMAN, D., and D. SCHLAGENHAUF (1985): "The Impact of News and Alternative Theories of Exchange Rate Determination," *Journal of Money, Credit and Banking*, 17, 328–346.

HOGAN, K., and M. MELVIN (1994): "Sources of Meteor Showers and Heat waves in the Foreign Exchange Market," *Journal of International Economics*, 37, 239–247.

HOGAN, K., M. MELVIN, and D. ROBERTS (1991): "Trade Balance News and Exchange Rates: Is There a Policy Signal?," *Journal of International Money and Finance*, 10, S90–S99.

HOSKING, J. (1981): "Fractional differencing," *Biometrika*, 68, 165–176.

HSIEH, D. (1989): "Modelling Heteroskedasticity in Daily Foreign Exchange Rates," *Journal of Business and Economic Statistics*, 7, 307–317.

HSIEH, D., and A. KLEIDON (1996): "Bid-Ask Spreads in Foreign Exchange Markets: Implications for Models of Asymmetric Information," in *The Microstructure of Foreign Exchange Markets*, ed. by J. Frankel, G. Galli, AND A. Giovannini. University of Chicago Press.

HUANG, R., and R. MASULIS (1995): "Spread, Dealer Competition and Market Regimes: A Market Microstructure Analysis of FX trading," Working Paper, Owen Graduate School of Management, Vanderbilt University.

HURVICH, C. (1985): "Data-driven choice of a spectrum estimate: extending the applicability of cross-validation methods," *Journal of the American Statistical Association*, 80, 933–940.

ITO, T., R. LYONS, and M. MELVIN (1996): "Is There Private Information in the FX market: The Tokyo Experiment," Working Paper, Haas School of Business, UC Berkeley.

ITO, T., and V. ROLEY (1987): "News from the US and Japan: Which Moves the Yen/Dollar Exchange Rate," *Journal of Monetary Economics*, 19, 255-277.

JACQUIER, E., G. POLSON, and P. ROSSI (1994): "Bayesian Analysis of Stochastic Volatility Models," *Journal of Economics and Business Statistics*, forthcoming.

KOOPMAN, S. (1993): "Disturbance Smoother for State Space Models," *Biometrika*, 80(1), 117-126.

KÜNSCH, H. (1987): "Statistical aspects of self-similar processes," in *Proceedings of the First World Congress of the Bernouilli Society*, pp. 67-74. VNU Science Press.

LO, A., and C. MACKINLAY (1990): "An Econometric Analysis of Non-Synchronous Trading," *Journal of Econometrics*, 45, 181-211.

LYONS, R. (1994): "Information Externalities in the Microstructure of the Foreign Exchange Market," Working Paper, Berkeley Business School, Berkeley, CA.

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

LYONS, R. (1996): "Foreign Exchange Volume: Sound and Fury Signifying Nothing?," in *The Microstructure of Foreign Exchange Markets*, ed. by J. Frankel, G. Galli, AND A. Giovannini. University of Chicago Press.

MADHAVAN, A., and S. SMIDT (1991): "A Bayesian Model of Intraday Specialist Pricing," *Journal of Financial Economics*, 30, 99-134.

MARK, N., Y. WU, and W. HAI (1993): "Understanding Spot and Forward Exchange Rate Regressions," Working Paper, Ohio State University, Columbus, OH.

MEESE, R., and K. ROGOFF (1983): "Empirical Exchange Rate Models of the Seventies: Do they Fit out of Sample," *Journal of International Economics*, 14, 3-24.

MELINO, A., and S. TURNBULL (1990): "Pricing Foreign Currency Options with Stochastic Volatility," *Journal of Econometrics*, 45, 239-265.

MELVIN, M., and X. YIN (1996): "Public Information Arrival, Exchange Rate Volatility and Quote Frequency," Working Paper, Arizona State University.

MICHAEL, P., A. NOBAY, and D. PEEL (1997): "Transactions Costs and Nonlinear Adjustment in Real Exchange Rates: An Empirical Investigation," *Journal of Political Economy*, forthcoming.

MITCHELL, M., and J. MULHERIN (1994): "The Impact of Public Information on the Stock Market," *Journal of Finance*, 49, 923-950.

MUSCATELLI, A., and P. TIRELLI (1996): "Institutional Change, Inflation Targets and the Stability of Interest Rate Reaction Functions in OECD Economies," Working Paper, University of Glasgow.

NAIK, N., A. NEUBERGER, and S. VISWANATHAN (1994): "Disclosure Regulation in Competitive Dealer Markets: Analysis of the London Stock Exchange," Working Paper.

NELSON, D. (1990): "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59, 349-370.

PEARCE, D., and V. ROLEY (1985): "Stock Prices and Economic News," *Journal of Business*, 58, 49-67.

PEEL, D., and P. POPE (1991): "Forward Foreign Exchange Rates and Risk Premia - A Reappraisal," *Journal of International Money and Finance*, 10, 443-456.

PERRAUDIN, W., and P. VITALE (1996): "Interdealer Trade and Information Flows in the Foreign Exchange Market," in *The Microstructure of Foreign Exchange*

Markets, ed. by J. Frankel, G. Galli, AND A. Giovannini. University of Chicago Press.

PRIESTLEY, M. (1981): *Spectral analysis and time series*, Probability and Mathematical Statistics. Academic Press.

QUAH, D. (1992): "The Relative Importance of Permanent and Transitory Components: Identification and Some Theoretical Bounds," *Econometrica*, 60(1), 107–118.

ROBINSON, P. (1991a): "Automatic frequency domain inference on semiparametric and nonparametric models," *Econometrica*, 89.

ROBINSON, P. (1991b): "Testing for strong serial correlation and dynamic conditional heteroskedasticity in multiple regression," *Journal of Econometrics*, 47.

ROBINSON, P. (1994): "Time series with strong dependence," in *Advances in Econometrics*, ed. by C. Sims, vol. 1. Cambridge University Press.

ROBINSON, P. (1995a): "Gaussian semiparametric estimation of long range dependence," *Annals of Statistics*, forthcoming.

ROBINSON, P. (1995b): "Log periodogram regression of time series with long range dependence," *Annals of Statistics*, 23, 1048–1072.

ROBINSON, P. (1996): "Efficient tests of nonstationary hypotheses," *Journal of the American Statistical Association*, forthcoming.

RUIZ, E. (1994): "Quasi-Maximum Likelihood Estimation of Stochastic Volatility Models," *Journal of Econometrics*, 63, 289–306.

SCOTT, L. (1987): "Option Pricing when the Variance Changes Randomly: Theory, Estimation and an Application," *Journal of Financial and Quantitative Analysis*, 22, 419–438.

STOCK, J. (1988): "Estimating Continuous-Time Processes subject to Time Defor-mation," *Journal of the American Statistical Association*, 83, 77–85.

TAUCHEN, G., and M. PITTS (1983): "The price variability-volume relationship on speculative markets," *Econometrica*, 51, 485-505.

TAYLOR, J. (1993): "Discretion Versus Policy Rules in Practice," *Carnegie Rochester Conference on Public Policy*, 39, 195-214.

TAYLOR, M. (1987): "Covered Interest Parity: A High-Frequency, High-Quality Data Study," *Economica*, 54, 429-438.

TAYLOR, M. (1989): "Covered Interest Arbitrage and Market Turbulence," *Economic Journal*, 99, 376-391.

TAYLOR, S. (1994): "Modelling Stochastic Volatility: A Review and Comparative Study," *Mathematical Finance*, 4(2), 183-204.

TONG, H. (1990): *Non Linear Time Series: A Dynamic System Approach*. Clarendon Press, Oxford.

TSAY, R. (1989): "Testing and Modeling Threshold Autoregressive Processes," *Journal of the American Statistical Association*, 84, 231-240.

VOGLER, K.-H. (1993): "Inter-Dealer Trading," Discussion Paper 174, Financial Markets Group, London School of Economics.

VON HAGEN, J. (1995): "Inflation and Monetary Targeting in Germany," in *Inflation Targets*, ed. by L. Leiderman, AND L. Svensson. Centre for Economic Policy Research, London.

WHITE, H. (1980): "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica*, 48, 817-838.

ZHOU, B. (1992a): "Forecasting Financial Time-Series subject to De-volatization," Working Paper, Sloan School of Management, M.I.T., MA.

ZHOU, B. (1992b): "High Frequency Data and Volatility in Foreign Exchange Rates," Working Paper, Sloan School of Management, M.I.T., MA.