A Microstructural Analysis of the Effects of News on Order Flows and on Price Discovery in Foreign Exchange Markets

by

Ryan Love

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

March 2005

ABSTRACT

This thesis brings together a number of studies using high frequency foreign exchange (FX) data. The first part examines the effects of scheduled, publicly released macroeconomic news, while the final chapter considers another, related, aspect of FX microstructure. Chapter 1 provides an introduction to the thesis and reviews the literature in high frequency empirical FX research. In Chapter 2, I use up to ten months of FX transactions and quote data to analyse foreign exchange activity around times of scheduled news releases. The effects of news on exchange rate levels are examined, as well as the effects on spreads, trading volume and volatility. Chapter 3 extends this analysis, asking *how* public information enters prices. Under rational expectations and efficient markets hypotheses, the news contained in public information announcements should be impounded directly, with there being no role for trades in this process of information assimilation. However, the results suggest that up to two thirds of the price relevant information enters via trading (order flow in particular). Chapter 4 provides an explanation why order flow is so important around public news releases and also examines the effects of news on market depths. In Chapter 5 I examine how much information is carried in trades by looking at the price impact of order flow when feedback trading is allowed. The model that is often used in the literature is proved to be misspecified when temporally aggregated data are employed and Chapter 5 introduces a method to estimate the otherwise unidentified model. Using impulse response functions, I show that trades actually carry more information than previous estimates suggest.

Contents

List of Tables

List of Figures

Acknowledgements

Many people have helped me during my PhD studies. First and foremost, I would like to thank my supervisor, Charles Goodhart, for his continued support, advice and encouragement. A student's research is often only as good as the supervision he/she receives and I can only hope that the work presented here is half as good as the supervision that Charles has given.

I am also hugely thankful to the members of the Financial Markets Group, where I have been based for the duration of my graduate work. In particular, I would like to thank the Director, David Webb, for giving me the chance to work in such an exciting and stimulating environment. Interaction with other researchers and academics is vital for any PhD student and the positive working atmosphere, from which I and numerous other students have benefitted, is largely due to David's work. Other members of the FMG, past and present, who deserve thanks are: Bob Nobay, Richard Payne, Jon Daníelsson, Paolo Vitale, Margaret Bray, Antoine Faure-Grimaud, Amil Dasgupta, Urs Haegler, Miguel Segoviano, Jinhui Luo, Dennis Kristensen, Robert Kosowski, Felix Muennich and Heski Bar-Isaac. I would also like to thank Giuseppe Bertola, Dagfinn Rime, Richard Lyons, Carol Osler, Mike Melvin, Philipp Hartmann, Michael Moore, William Killeen and Geir Bjønnes, all of whom have contributed to this thesis, either by giving comments on my research (and co-authoring a paper in the case of Dagfinn Rime) or simply taking the time to discuss interesting aspects of FX microstructure. I would also like to thank the Economic and Social Research Council (ESRC) for financial support, specifically their three year PhD scholarship.

There are a number of other people who have, sometimes unknowingly, helped me over the last few years. The highs and the (all too often) lows of my PhD studies have required much emotional support, and this has come unreservedly from my family and friends. The love and support from my Mom and Dad, and my brother, Jon, have helped me more than they could ever know, and I am forever grateful to them. Unfortunately, there is not enough room to list all my friends from whom I have taken so much, but a few who deserve mention are: (Vesey beers); Guy, Marc, Rob, Stuart, Wigs, Pete and Clive, (LSEish); Sami, Caroline, Tammy, John and Robbie, (JT); Nicole, Wilko, Av, Eddie and Big Gay, and of course not forgetting The Mullett and Kevin. You all know who you are, and I thank you for forgiving an all too often absent friend.

Finally, I would like to thank my partner, my beautiful Jo. More than anyone, she has seen and experienced the highs and lows of my PhD career. Without her love, support and patience, I would surely be a much lesser man.

Chapter 1 Introduction

This thesis brings together a number of research papers, all of which consider microstructural aspects of the foreign exchange market. Standard macroeconomic models of exchange rate determination have shown themselves to be remarkably unsuccessful at explaining movements in floating foreign exchange (FX) rates in the short term. See the seminal papers by Meese and Rogoff (1983a,b) for example, and also the surveys by Taylor (1995) and Frankel and Rose (1994). Essentially, the short term forecasts made from numerous macroeconomic models perform worse than a random walk, i.e. the prediction that tomorrow's exchange rate will be the same as that seen today will give a more accurate picture of exchange rate changes compared to predictions based on macroeconomic models. Consequently, greater emphasis has been placed on microstructure factors. O'Hara (1995) defines market microstructure as *"... the study of the process and outcomes of exchanging assets under explicit trading rules"* and as such gives an important role to traders and their actions, namely the quotes/prices they post and the trades they execute. By examining exchange rates at higher and higher frequencies and by rationalising the actions taken by those who actually determine the exchange rate (FX traders in a world of floating rates), it is hoped that microstructure analyses will help explain FX prices more accurately.

Despite the poor performance of macroeconomic models when explaining short run exchange rate movements, these theories are generally successful when explaining changes in rates at longer horizons.[1] On the other hand, microstructure models are usually used to explain very short term exchange rate changes, with the majority of studies using data sampled at daily or higher frequencies. This segregation of microstructure and macroeconomic approaches is not entirely appropriate (Lyons 2001, Sarno and Taylor 2001) and it is the purpose of the first part of this thesis to examine some of the interactions between them. In particular, I examine the process of price discovery following the release of scheduled macroeconomic news. I firstly analyse the effects of news releases on a number

[1]For example, the consensus view is that deviations of nominal exchange rates from those consistent with purchasing power parity have half lives of between three and five years. For empirical evidence on the validity of macro models when explaining exchange rate movements in the long run, see Mark (1995) and Flood and Taylor (1996), for example.

of microstructure variables, and then ask how information contained in these releases is incorporated into prices.

Chapter 2 focuses on the effects of scheduled US and UK news releases on three heavily traded floating rates: USD/EUR (US dollars per euro), GBP/EUR (pounds sterling per euro) and USD/GBP (US dollars per pound). By analysing the effects of news on both bid and ask prices and on the number of buyer and seller initiated trades separately, I am able to examine the effects on not only the level of the exchange rate, but also on spreads (the difference between the best prices at which traders can buy and sell currencies), volatilities, trading volume and order flow (the difference between the number of buyer and seller initiated trades). The main findings are that unanticipated announcements of macroeconomic data cause significant and permanent changes in the level of the exchange rate and also cause trading volume to increase. This is caused by an increase in both buyer *and* seller initiated trades, irrespective of whether the news was good or bad. Volume is also found to increase even when the news release has no effect on the level of the exchange rate, consistent with the hypothesis that traders differ in their interpretation of news. Exchange rate volatility also increases post release and this volatility surge persists for some time. Spreads are also found to increase after the data announcement, but the effects are small and statistically insignificant.

By examining the three markets simultaneously, I am also able to quantify any cross market effects of news and various other interrelationships between the three rates. These cross market effects are shown to be large, with US news affecting not only the dollar markets but also GBP/EUR activity. Asymmetric effects of trades are also demonstrated, with market buy orders having a greater impact on the ask price than on the bid, and vice versa for market sell orders. Impulse response functions and variance decompositions also show significant spill-overs from one market to another.

After asking *what* happens to various microstructure variables around periods of macroeconomic news releases, Chapter 3 asks *how* public information enters FX rates. Standard rational expectations and efficient markets hypotheses suggest that price relevant information contained in public news releases should be incorporated into prices without the need for trading. Indeed, French and Roll (1986) define public information to be that which is incorporated into prices before anyone can trade on it. Chapter 3 tests this assertion and finds that up to two thirds of the information contained in macroeconomic news is assimilated via the key microstructure variable — order flow. The models of Glosten and Milgrom (1985), Kyle (1985), Easley and O'Hara (1987) and Perraudin and Vitale (1996) all suggest that order flow is the mechanism through which *private* information is incorporated into price. If public information enters price via order flow then the segregation of public and private information, an assumption commonly made in the literature, may not be entirely valid. This is an avenue explored in Chapter 4, where, using the ideas of Harris and Raviv (1993) and Kandel and Pearson (1995), I present a simple explanation why order flow is so important in the process of public information assimilation. The basic idea is that traders receive noisy signals from a particular news announcement, but by also allowing traders to differ in their *abilities* to interpret news, this not only explains the increase in trading volume following the data release, but also explains the effects on order flow.

Chapter 4 also presents further empirical evidence of the effects of news announcements on the microstructure of FX markets, in particular, the effects of news on market depths (the quantities dealers are willing to trade at different prices). Using anecdotal evidence from a single release of US PPI data, the depth of the DEM/USD limit order book is found to fall, with dealers willing to trade smaller quantities for any deterioration of the bid and ask prices. This is argued to be a result of increased information asymmetry following the news release, and is not simply a result of market orders picking off limit orders which would naturally hollow out the book.

As such, the first part of this thesis focuses on the effects of scheduled, macroeconomic news announcements. Chapter 2 asks '*What* happens around these data releases?', Chapter 3 asks '*How* is the news incorporated into prices?' and Chapter 4 presents an explanation as to *why* macroeconomic news enters price this way. The final chapter of the thesis considers another, related, aspect of FX microstructure. Whereas Chapter 3 demonstrates how important trading is in the assimilation of *public* information, Chapter 5 shows how important trading is in general, when allowing for the possibility of feedback trading by FX dealers.

The effects of order flow on asset prices are commonly split into temporary and permanent components. The temporary effects of order flow are usually described as inventory or 'indigestion' effects and should have no bearing on the long run value of the asset being traded. On the other hand, order flow effects that result from traders having private information will have a permanent effect on price. Therefore, one way to ascertain how much information trades carry is to use the vector autoregressive (VAR) methodology introduced by Sims (1980), and implemented in a microstructure context by Hasbrouck (1991a,b). By using a two variable VAR in asset price returns and order flow, one can introduce a shock to the trade equation and examine the cumulative effect on price. This approach has been used in numerous studies[2] and assumes that the direction of causality runs explicitly from order flow to asset returns, i.e. returns depend on contemporaneous order flow but order flow can only depend on lagged price changes. This assumption is entirely reasonable when data are sampled at ultra-high frequencies, such as tick-by-tick, but as soon as data are aggregated at lower frequencies, the problem of contemporaneous feedback trading becomes apparent. If traders respond to price changes by executing trades themselves, then when using daily, hourly, or indeed minute by minute data if traders respond quickly enough to such price changes, it is possible that not only will returns depend on contemporaneous order flows, but order flow will also depend on contemporaneous returns. Chapter 5 proves that the standard model used in the literature is misspecified whenever aggregated data are used, and suggests a method to overcome these difficulties.

By using standard instrumental variables techniques, returns are not only found to be affected by contemporaneous order flows, but are also shown to be a determinant of these flows. By comparing impulse response functions with and without contemporaneous feedback, I show that trades carry more information (have a greater price impact) than previous estimates suggest.

The rest of this chapter gives a brief review of previous FX microstructure studies, describing some of the datasets that have been used and giving a more historical perspective of empirical FX research. For a comprehensive study of the issues involved in market microstructure, see O'Hara (1995) and for a more recent survey, see Madhavan (2000). Lyons (2001) gives more detail on the microstructure of the foreign exchange market and for an earlier description of the institutional workings of inter-dealer trading, see Flood (1991). Sarno and Taylor (2001) present a survey of the literature on FX microstructure and Vitale (2004) gives a more technical overview of some of the issues that have been studied in recent papers.

[2] See Section 5.1 for specific examples.

1.1 Literature Review of Empirical FX Market Microstructure Studies

The number of empirical studies of the inter-dealer foreign exchange market has grown considerably in the last few years, fuelled by the increasing availability of high quality datasets. Data availability has rarely been a problem for studies of equity markets, while the nature of the FX market has necessarily meant fewer datasets being on hand for academic research. Until recently the inter-bank foreign exchange market was split, roughly equally, between direct and brokered trading.[3] Direct trading is characterised by traders contacting each other directly (via computer terminals, such as the Reuters D2000-1 platform, or by telephone). The trader receiving the call/quote request gives bid and ask prices, which are subsequently accepted or declined by the initiating trader. As such, the direct market can be classified as a decentralised, continuous, multiple dealer structure. The decentralised nature of direct trading, where dealers interact in different physical locations with no requirement to notify an overseeing body of their quotes or trades, has resulted in relatively few of these datasets being used in academic work.

Brokered trading, on the other hand, is generally characterised as being quasi-centralised and continuous, with a limit order book structure. Brokers collect the limit orders of numerous dealers and place them within their limit order book.[4] Traders can then trade at these prices by executing market orders, or posting 'marketable' limit orders, where the posted bid price is higher than the current best ask for example. The broker automatically matches these orders and the trade occurs. Information from voice brokers has been difficult, if not impossible, to obtain due to the highly confidential nature of these trades. However, in recent years, brokered trading has tended to move to electronic limit order books such as the Reuters D2000-2 and EBS (Electronic Broking Services) platforms. These have largely replaced voice brokers and have in fact become the dominant form of inter-dealer trading. The BIS reported that in 2000, between 85 and 95% of all inter-bank trading took place using electronic brokers, increasing from about 50% in 1998 and 20-30% in 1995 (Bank for International Settlements 2001).

Reuters D2000-2 was launched in 1992 and by September 1993 it was used in 28 cities,

[3]See Flood (1991) for an early description of the makeup of the inter-dealer market.

[4]Limit orders are orders to buy (sell) X units of the asset at a specified bid (ask) price. Market orders, meanwhile, are orders to buy or sell X units of the asset at the best available price. As such, market orders execute immediately (assuming there is a willing counter-party) and transact against the limit orders in the book.

with 230 subscribers in 17 countries (Goodhart, Ito, and Payne 1996). September 1993 also saw the introduction of EBS as a direct competitor to the D2000-2 platform and, since these two trading structures have come to dominate inter-dealer trading, such spot market trading can be characterised as being quasi-centralised. Also, with all information on quotes and trades being stored electronically, they have naturally allowed the possibility for more thorough empirical investigation.

1.1.1 Indicative quotes in FX research

Until the direct or brokered inter-dealer data became available, high frequency empirical studies usually employed indicative quotes. These are essentially adverts posted by traders to show that they have an active interest in the market, and are shown on the screens of electronic data providers such as Reuters (known as Reuters FXFX pages) or Telerate for example. However, the bid/ask quotes are not firm prices at which dealers can trade. They are, by their very nature, 'indicative'.[5] One of the first high frequency datasets employed in the microstructure literature was that used in Goodhart and Giugale (1993). They used indicative quote data covering four exchange rates versus the US dollar (pound sterling, Deutsche mark, yen and Swiss franc) that were collected by MMS International from the Telerate electronic screens. The data were collected hourly, spanning 2nd January to 15th July 1986, and were used to examine a number of features of high frequency FX quotes: autocorrelation of the return series, cross-correlation between markets, the presence of a unit root, etc. This dataset was also used by Baillie and Bollerslev (1991), who examine the intra-day/hourly volatility patterns for these markets. Among other things, they found that when the diurnal patterns are taken into consideration, volatility no longer looks like an IGARCH process; the GARCH parameters sum to much less than unity.

One of the earliest studies that make use of data sampled at the one minute frequency was that of Goodhart and Figliuoli (1991). They use Reuters indicative quotes from three days in 1987: 14th and 15th September and 21st October. By using a number of exchange rates versus the US dollar, they analyse spreads and other time series characteristics of FX data. These data are also used in Goodhart and Figliuoli (1992), who test whether the different geographical locations of traders and country specific heterogeneities can explain the negative serial correlation of high frequency returns.

[5]The Reuters FXFX data have become common in high frequency FX research largely because of the work of Olsen and Associates (Zurich), who collect and store the quotes for a number of FX rates and sell them to academic researchers for a small fee.

Goodhart (1989) makes use of an even higher frequency dataset, where tick-by-tick indicative FXFX quotes are examined, covering the period from 9th April to 3rd July 1989.[6] The focus in Goodhart (1989) is to examine the effect of news releases on high frequency FX returns. This is extended in Goodhart, Hall, Henry, and Pesaran (1993) where, using the same dataset but concentrating on USD/GBP, a GARCH-M model is fitted, allowing the effects of news releases on the level and volatility of the exchange rate to be analysed simultaneously. These early studies of the effects of news are similar to the research presented in Chapters 2 to 4 of this thesis. However, due to the nature of the FXFX data, no data on transactions were available. Trading volume and order flow, meanwhile, are studied extensively in my research. Despite no transactions data being available in these three months of FXFX data, it has still been possible to address a number of other important issues. Goodhart and Demos (1991) present an extensive examination of weekly and intra-day patterns of quote frequency and Hsieh and Kleidon (1996) study the intra-day patterns of spreads and volatilities. This latter study compares the intra-day patterns between London and New York based traders, while the ideas of price leadership are tested explicitly using the same data in Wang (2001). Bollerslev and Domowitz (1993) look at similar issues by examining the intra-day patterns of quote frequency for large and small banks, also making reference to where the banks are located. Bollerslev and Melvin (1994), on the other hand, use these data to model and test the effect of FX volatility on spreads.

Another large FXFX dataset has been used extensively in academic research, this time covering up to three years of DEM/USD data from January 1992 to December 1994. Similar to a few of the research papers described above, de Jong, Mahieu, Schotman, and van Leeuwen (2001) test and confirm the ideas of price leadership between banks. Almeida, Goodhart, and Payne (1998) test the effects of US and German news releases on the level of the exchange rate, while Payne (1996), Andersen and Bollerslev (1997a,b) and Andersen and Bollerslev (1998) all study the effects of news on DEM/USD volatility when taking into account the intra-day/diurnal pattern of FX second moments. DeGennaro and Schrieves (1997) also study the effects of news on volatility in this period, but this time using JPY/USD FXFX quotes. The effects of news announcements on quote activity and volatility in the JPY/USD market are also analysed in Melvin and Yin (2000) when using a slightly later FXFX dataset from December 1993 to April 1995.

[6]This study also uses quotes from the Reuters FXFY pages, where quotes for a number of less heavily traded pairs are recorded.

Indicative quote data have also been used to address other issues. Dominguez (2003) uses FXFX data from August 1989 to August 1995 to examine the intra-day and daily volatility effects of interventions made by the Federal Reserve in the JPY/USD and DEM/USD markets. Fischer (2003) uses a similar data period to analyse the effects of intervention made by the SNB (Swiss National Bank), while Payne and Vitale (2003) exploit up to ten years of CHF/USD (Swiss francs per US dollar) indicative quotes from 1986 to 1995 to study the effects of SNB intervention. As well as these studies, FXFX data have been used to test market efficiency, as tested for by the presence of triangular arbitrage (Kollias and Metaxas 2001) and also to test models of price discreteness (Hasbrouck 1999).

An important set of papers that should be mentioned in this section are those that test for the presence of private information in the FX market. It is natural to assume that private information is only an issue for equity markets, where the fundamentals of the stock price include earnings data, on which some traders can have insider knowledge. Foreign exchange rates, meanwhile, should be determined by macroeconomic variables, which are in the public domain (Bessembinder 1994). To test for the presence of private information, one might expect to require transaction data.[7] However, until such data became available, researchers had to devise ways to test for this type of information using quote data only. When trying to explain the clustering of FX volatility, Engle, Ito, and Lin (1990) suggest that private information and the existence of traders with heterogenous priors will result in volatility continuing for some time after a particular shock. They use a multivariate GARCH model using opening and closing JPY/USD prices in different markets around the world (October 1985 to September 1986) to test whether volatility in one geographical location follows turbulent times in previously open geographical locations (meteor showers) or whether volatility increases following a turbulent domestic market yesterday (heat waves). They find evidence in favour of the meteor shower hypothesis while rejecting the heat wave explanation, suggesting that foreign news is more important than yesterday's domestic news. This, they argue, is consistent with the existence of private information in FX markets, or with stochastic coordination between the different countries.

Evidence of private information in FX markets is also presented in Ito, Lyons, and Melvin (1998) who use FXFX quotes from 29th July 1994 to 28th March 1995 in the JPY/USD and DEM/USD markets. This period saw the introduction of trading in the Tokyo lunch hour, after which, lunch-return variances were found to double. Using variance ratio

[7]This is discussed in Section 1.1.2.

methods, similar to French and Roll (1986), they argue that this increase in return variance is due to the presence of private information. Using the same dataset, Covrig and Melvin (2002) also find evidence for the existence of private information, suggesting that Japanese banks are better informed regarding the yen and tend to be price leaders.

This heterogeneity between banks is also found in Ben Omrane and Heinen (2003), who analyse the quotes of individual banks around macroeconomic news announcements. By using FXFX data on USD/EUR from August to October 2001, they find that some banks increase their quote activity following the data release while others reduce their participation in the market. This, they argue, is consistent with the existence of heterogenous traders who differ in their interpretation of news. These ideas are explored in more detail in Chapter 4.[8] Other recent studies that analyse the effects of macroeconomic news on rates using indicative data include Andersen, Bollerslev, Diebold, and Vega (2003), Faust, Rogers, Wang, and Wright (2003) and Melvin, Sager, and Taylor (2003). Andersen, Bollerslev, Diebold, and Vega (2003) study the effects of various US news releases on a number of currencies versus the US dollar (pound sterling, yen, Swiss franc, Deutsche mark and euro) using data sampled at the five minute frequency from January 1992 to December 1998. Faust, Rogers, Wang, and Wright (2003) use five minute FXFX data from 1987 to 2002. Using this long dataset they examine the effects of news on not only the USD/GBP and DEM/USD (USD/EUR) markets, but also on a number of bond markets, examining the interrelations between these different assets.[9] Melvin, Sager, and Taylor (2003) use indicative quotes from HSBC when examining the high frequency effects of interest rate announcements from the Bank of England's Monetary Policy Committee.

1.1.2 Transactions data in FX research

The first half of this thesis is similar in nature to a number of the studies described above, where the effects of scheduled news releases on prices and volatility are analysed. However, in order to study the process of price discovery and to analyse *how* new information is incorporated into exchange rates, an obvious variable to consider is trading activity. Unfortunately, data employing indicative quotes are entirely devoid of such information

[8]Trader heterogeneity is also shown up in the responses to a number of questionnaires. Traders appear to differ in their use of chartist/technical analysis, in their beliefs over whether exchange rate movements reflect changes in their fundamental value, and in their forecasts of future exchange rate changes for example. See Allen and Taylor (1990), Ito (1990), Taylor and Allen (1992) and Cheung and Chinn (1999).

[9]These interrelations between different markets are also examined in Andersen, Bollerslev, Diebold, and Vega (2004) and Hau and Rey (2002) to name but two.

and therefore these studies can only go so far when analysing the mechanics of information assimilation. Fortunately, a number of transaction based datasets have become available, such as that used in this thesis and described in detail in Chapter 2. Here I describe a number of the datasets that have been used, explaining where they come from and what they have been used to test.

One of the earliest examinations of trading activity is that of Goodhart and Giugale (1989). This uses the same indicative data as those used in Goodhart and Giugale (1993), covering 2nd January to 15th July 1986, but combines these data with the transactions records of two large banks operating in London. Using data on daily volumes, Goodhart and Giugale (1989) study the interaction between transactions and indicative prices, as well as the time series properties of these transactions data.

A more disaggregated transactions dataset was used by Richard Lyons, who was able to record all the quotes and transactions of a single DEM/USD trader for the week of 3rd to 7th August 1992. By observing all the trades of an individual dealer, he is able to compute and follow his inventory level, allowing a number of previously untestable microstructure hypotheses to be investigated. In particular, Lyons (1995) finds evidence for both inventory control and asymmetric information effects in the dealer's quotes. The inventory control effects, whereby the dealer would lower his quotes following an unwanted long position (in order to attract market buy orders from other traders), were found to be strong for this dealer, while previous studies had shown relatively weak inventory control effects in the quotes of equity traders, shown in Madhavan and Smidt (1991), for example. These strong inventory control effects suggest that the management of intra-day inventories is more important for FX traders than it is for NYSE specialists, a point illustrated by the fact that the trader Lyons tracked, traded approximately \$1 billion per day but still closed out all positions to leave overnight inventory near zero.

Lyons (1995) also found that the trader increased his spread to compensate for the risk of trading with more informed market participants. Therefore, for the first time using transaction data, Lyons (1995) found evidence for the existence of private information in FX markets. Private information was also illustrated by the significant effects of order flow on the level of the dealer's quotes, as explained by the models of Kyle (1985) and Glosten and Milgrom (1985). Order flow models have subsequently been used extensively in empirical FX research, including the studies presented in this thesis, and all have found significant effects of order flow on the level of the exchange rate. When examining the price impact of order flow, Lyons (1996) extends his original analysis to test between the 'event uncertainty' versus 'hot potato' hypotheses. Easley and O'Hara (1992) suggest that the price impact of order flow should be larger when trading volume is high since greater trading activity suggests the existence of more information 'events', i.e. more private information which needs to be incorporated into the price: the 'event uncertainty' hypothesis. Lyons (1996), on the other hand, suggests that the huge inter-dealer trading that is seen in FX markets largely results from traders passing on unwanted inventories: the 'hot potato' view. Since these trades are relatively uninformative, the price impact of any order flow should be small when volume is high. The data from Lyons' one dealer do in fact support the hot potato view, evidence for which is also presented in Luo (2002b).

Other studies that make use of data from individual dealers include Yao (1998) and Bjønnes and Rime (2003). Yao (1998) estimates a similar model to that of Lyons (1995) but uses all the transactions records for a large DEM/USD dealer from 1st November to 8th December 1995, i.e. the data sample is five times as long as that used in Lyons (1995). Whereas Lyons (1995) finds significant inventory effects in the dealer's quotes, such effects are not found by Yao (1998). Instead, Yao argues that the act of quote shading gives competing dealers information on your inventory position. This could prove to be very costly and therefore dealers refrain from trading in this way. This idea is confirmed in Bjønnes and Rime (2003), who argue that despite inventory control being of upmost importance, this is done by traders executing market orders, rather than by shading their limit order quotes. The data used by Bjønnes and Rime (2003) cover a range of currency pairs (Norwegian, Swedish and Danish krone, and Swiss franc all against the Deutsche mark, and both Deutsche mark and Norwegian krone against the US dollar) across four dealers from 2nd to 6th March 1998.

Whereas data from individual traders can allow analysis of dealers' inventories and therefore of a number of important microstructure hypotheses, such data have the disadvantage of lacking significant market coverage. Instead, one can use the aggregated trading information that can be taken from either the direct or brokered platforms. Fortunately, the majority of direct trading, which still accounted for a large share of inter-dealer trading in the mid 1990s, occurred via the Reuters D2000-1 system.[10] Data from this trading platform form the basis of a number of research papers by Martin Evans and Richard Lyons. Evans and Lyons (1999) use these data (4 months of DEM/USD and JPY/USD activity from 1st May to 31st August 1996) to test a model that gives an important role to order

[10]The Reuters D2000-1 system is an electronic form of direct trading and therefore allows relatively easy data capture.

flow.[11] Using daily data, a regression of FX returns on order flow and interest differentials (proxying macroeconomic fundamentals) produced an R^2 of 0.64 (DEM/USD) and only the coefficient on order flow was found to be significant. The importance of order flow in the determination of FX rates is attributable to the information it conveys about (non-dealer) customer trades. Each trader is argued to receive a signal from his/her customer base and these signals can only be aggregated via inter-dealer order flow. Even though each trader has a private signal of the currency's payoff, information is not concentrated, as implied by the informed trader models of equity markets, but rather it is dispersed among a large number of separate dealers. Order flow is therefore the proximate determinant of exchange rates as it is the mechanism through which all the dispersed pieces of information in the economy get aggregated into price.

These D2000-1 data have been used to analyse a number of issues in the foreign exchange market, all placing great importance on order flow. The high frequency dynamics of the data are discussed in Evans (2001) and the cross market effects of order flow are illustrated in Evans and Lyons (2002a). Evans and Lyons (2003) ask how macroeconomic news is incorporated into price, in particular, whether information enters directly or whether it enters via order flow. This is directly related to the work presented in Chapter 3, and similar conclusions are drawn; the majority of public information enters via order flow, inconsistent with standard rational expectations models. In a similar exercise, Evans and Lyons (2002c) examine the price impact of order flow and ask how this changes around public news announcements, and Evans and Lyons (2001) examine the effects of order flow around Central Bank intervention.

The earliest high frequency study of brokered inter-dealer trading is that of Goodhart, Ito, and Payne (1996). They were able to obtain the videotapes of the Reuters D2000-2 screen for 7 hours on 16th June 1993, when this trading platform was in its early stages of development. They analyse the characteristics of these brokered data, comparing them to the data from the FXFX screens, and also looking at the interaction between the different variables: quote frequency, spread, trading activity, etc. The analysis is extended in Goodhart and Payne (1996), who focus on the microstructural dynamics of bid/ask prices, executed orders, spreads and volatility.

A slightly longer D2000-2 dataset that has been used in academic studies covers the one week of DEM/USD activity from 6th to 10th October 1997. Despite only covering five days of trading, the dataset has the advantage that it includes every entry by every

[11] This is an extended version of Evans and Lyons (2002b).

trader, therefore allowing one to examine the entire limit order book, rather than purely the front end that has been available in other brokered data. These limit order data can therefore be used to analyse the depth of the DEM/USD market, as done in Chapter 4 of this thesis. An important paper that uses this dataset is that of Daníelsson and Payne (2002b). Similar to Goodhart, Ito, and Payne (1996), they examine the differences between the indicative FXFX data and the transactions based D2000-2 dataset, and find, among other things, that spreads and volatilities calculated from indicative data are not representative of those that are actually seen in the market. For example, the indicative spreads are much larger than those actually faced by traders and therefore any analysis on the cost of trading based on indicative quotes is likely to be misleading. Detailed investigations of liquidity supply on this trading platform are presented in Daníelsson and Payne (2002a) and Daníelsson and Payne (2001), while Payne (2003a) uses the data to examine the effects of order flow on FX returns. Using a VAR framework, Payne (2003a) finds a long run cumulative effect of order flow on DEM/USD price changes, consistent with the findings of Lyons (1995) and Evans (2001) where order flow was found to carry private information.[12] Whereas Payne (2003a) examines the effects of order flow (net buying pressure via executed market orders), the effects of net buying/selling pressure from unfilled limit orders left on the book are examined in Love and Goodhart (2004).[13] Chapter 5 fits into this branch of the literature perfectly. The conclusion that there exists private information in FX markets comes from the finding that order flow has a permanent effect on price. However, much of the current literature assumes that the direction of causality runs explicitly from order flow to asset return; returns depend on contemporaneous order flow but order flows do not depend on contemporaneous asset price changes. This assumption appears reasonable when data are sampled at very high frequencies, but the effects of contemporaneous feedback trading may be considerable when data are sampled more coarsely. This issue is addressed in Chapter 5.

Other studies that exploit D2000-2 data include Daníelsson, Luo, and Payne (2002), Luo (2002b), Goodhart, Love, Payne, and Rime (2002) and Payne (2003b), all of which use the data employed in this thesis, described in Chapter 2. Daníelsson, Luo, and Payne

[12] This is also supported by the results in Rime (2001), who uses weekly order flow data from January 1996 to May 1999, obtained from Norges Bank. When examining a number of currencies against the Norwegian krone, order flow was found to have a permanent effect on price, suggesting the presence of private information.

[13] Love and Goodhart (2004) represents work that was performed during my PhD studies and has been submitted to the Journal of International Money and Finance. It has not been included in this thesis since it does not fit neatly into the other research topics presented here; namely public information and the informativeness of order flow.

(2002) analyse the cross market order flow effects (the effects of USD/EUR order flow on USD/GBP and GBP/EUR returns for example) and investigate how the effects of order flow change when the data are sampled at lower and lower frequencies and Luo (2002b) examines how the effects of order flow change with different market conditions (low versus high trading volume, tight versus wide spreads, etc.). Goodhart, Love, Payne, and Rime (2002) presents a simple comparative statics exercise to see what happened to inter-dealer spreads in the FX market after the introduction of the euro, while Payne (2003b) explores a number of relationships between spreads, volatilities and trading volume.[14] These relationships are also investigated by Hartmann (1998) when using volume data from the BIS triennial surveys in 1989 and 1992, as well as futures volume from the Chicago International Monetary Markets and daily JPY/USD volumes from Tokyo brokers, which are published in the Japanese financial press. Hartmann (1999) extends this work by examining the relationships between volume and spreads in the JPY/USD market when decomposing trading volume into expected and unexpected components. Expected volume is found to lower spreads, consistent with the idea of economies of scale in the processing of FX transactions, while unexpected volume increases spreads, as it may suggest the arrival of new information. This exposes traders to the risk of asymmetric information and therefore traders post wider spreads to compensate.

Recently, however, a number of datasets from the EBS platform have become available. Analysis of these data are likely to give greater insights into how the FX market works, especially when one considers the anecdotal evidence that suggests much of the brokered trading has moved towards this system, especially for the EUR markets.

Ito and Hashimoto (2004) use EBS data, sampled at the one second frequency, from 1st January 1999 to 31st December 2001, which cover the JPY/USD and USD/EUR markets. They present a comprehensive description of the data, paying particular attention to the intra-day patterns. Killeen, Lyons, and Moore (2002) use the twelve months of FFR/DEM (French francs per Deutsche mark) data in 1998 to examine the effects of order flow on returns pre and post announcement of the EMU parities (made over the first weekend in May), while Lyons and Moore (2003) build and test an order flow model that explains why the US dollar acts as a vehicle currency between the yen and euro. Chaboud et al (2004) use EBS tick data on both JPY/USD and USD/EUR markets from January 1999 to January 2004 to examine the effects of scheduled macroeconomic news on FX rates and

[14] Goodhart, Love, Payne, and Rime (2002) has been published in Economic Policy and is not included in this thesis for the same reason as Love and Goodhart (2004); the research does not fit neatly into the ideas of public information or trade informativeness.

trading activity. Therefore, this study fits very comfortably with the research presented in Chapters 2 and 3.

The order flow models that have been estimated with direct or brokered transactions data have generally performed well when explaining short term exchange rate movements. However, the order flow is calculated from inter-dealer trades only, which leaves one asking where this order flow comes from. Following the discussion earlier in this section, inter-dealer flows are, it is believed, largely a result of the (non-dealer) customer trades that each dealer receives. If inter-dealer order flows are a proximate determinant of exchange rates, then a study of customer trades therefore appears warranted. Unfortunately, data on end customer trades are extremely confidential, but such data have been employed in a small number of research papers. Evans and Lyons (2004a) examine all the USD/EUR customer trades that Citibank received from 11th April 1993 to 30th June 1999, and break up these flows into different components: non-financial companies, unleveraged financial institutions (primarily mutual funds) and leveraged financial institutions (mostly hedge funds). They examine the effects of news on these end customer trades, while Fan and Lyons (2003) use the same data to see how the price impact of these trades differs from one group to another. Fan and Lyons (2003) also examine these different customer trades around the collapse of the US dollar in October 1998, when the JPY/USD rate fell from 130 to 118 in just one day. The impacts of institutional investor flows are also examined in Froot and Ramadorai (2002) who use data from State Street Corporation from 1st January 1994 to 9th February 2001. They find that these FX flows largely determine short run returns, while long run exchange rate movements are determined more so by macroeconomic fundamentals. Again, this study fits in well with the work presented here, where the interaction between macroeconomic and microstructure factors are analysed explicitly.

Chapter 2 First and Second Moment Effects of Macroeconomic News in High Frequency Foreign Exchange Data

2.1 Introduction

"(S)*cheduled* (macroeconomic news) *releases occasionally induce large price changes, but the associated volatility shocks appear short-lived... Market participants may have different information sets, and thus differ in their interpretation of the news, but the market typically settles on a new equilibrium price after a brief period of hectic trading.*" Andersen and Bollerslev (1998), page 234.

A number of papers have investigated the effects of scheduled macroeconomic news announcements on exchange rates and their volatility. Hakkio and Pearce (1985), Goodhart, Hall, Henry, and Pesaran (1993), Almeida, Goodhart, and Payne (1998) and Andersen, Bollerslev, Diebold, and Vega (2003) all find significant and long-lasting effects of macroeconomic news releases on the level of the exchange rate. For the effects of news on FX volatility, see Goodhart, Hall, Henry, and Pesaran (1993), Ederington and Lee (1993), Payne (1996), Andersen and Bollerslev (1997a) and Andersen and Bollerslev (1998). However, due to the limited availability of transaction data, little has been done to examine the effects of news on trading activity. Evans and Lyons (2003) and Chaboud et al (2004) are two exceptions. The former asks how the information contained in macro news releases is incorporated into the exchange rate, an issue also studied in Chapter 3. However, before asking *how* macroeconomic news is incorporated into asset prices, this chapter first asks *what* happens in the FX market around announcements. This chapter therefore examines the effects of news releases on foreign exchange prices, but also on spreads, trading activity and FX volatilities.

A large literature has emerged that explains, primarily at very high frequencies, the link between FX prices and order flow; signed trading volume. Due to the inability of standard macroeconomic models to explain exchange rates, at least at anything other than at very low frequencies, following Meese and Rogoff (1983a), emphasis has shifted to microstructure models. These studies suggest that trades carry information from the more informed participants to the market as a whole. In this way (private) information is aggregated via the trading process and so order flow should have permanent effects on prices. Such ideas are not new; see for example the models of Kyle (1985) and Glosten and Milgrom (1985).[1] A second strand of literature has also developed that tries to explain the determination of exchange rates. Rather than using macroeconomic models to explain low frequency FX rate determination, a number of researchers have examined the high frequency effects of macroeconomic news announcements on the level of exchange rates and also on their second moments. The artificial separation of these two research strands has been largely due to the lack of high frequency transactions data. If order flow, and the transactions process in general, is an important determinant of exchange rates, and if scheduled macroeconomic releases have significant effects on FX prices too, then it seems likely that the links between order flow and FX prices, and between macro news releases and FX prices, are intimately related. Any model which tries to explain foreign exchange rates should ideally incorporate both mechanisms.

The price effects of public information that come via order flow are analysed in Chapter 3. However, before outlining the importance of order flow in the assimilation of public information, this chapter examines exactly what happens to various market statistics around releases of scheduled macroeconomic announcements.[2] Using eight months of transactions and quote data on three foreign exchange markets, together with data on scheduled macroeconomic news, I investigate the links between news releases, FX prices, spreads, trading volume, etc. However, if there are linkages between the various market statistics and one were to examine the effects of news on FX prices, volume and spreads each in isolation, then each model would be incorrectly specified, being prone to omitted variable bias for example. Instead, I model the effects of news on prices, spreads and trading volume simultaneously by using the VAR model of Hasbrouck (1991a). The model

[1] See also the models of Easley and O'Hara (1987), Madhavan and Smidt (1991) and Evans and Lyons (2002b). Empirical results showing the effects of order flow on FX rates include Lyons (1995), Yao (1998), Evans (2002) and Payne (2003a).

[2] The quote by Andersen and Bollerslev at the head of this chapter has not yet been adequately researched. Andersen and Bollerslev (1998) do not test their 'hectic trading' hypothesis due to the unavailability of transactions data in their study. As mentioned in Chapter 1, their study employs indicative FXFX quotes

also allows me to examine the interrelationships between three major floating rates, which are found to be considerable.

There are a number of effects that announcements of macroeconomic news may have in the FX market. These are described briefly below.

Trading activity A number of papers have proposed that contrary to the pure rational expectations models, the mapping of information to price is not common knowledge. The models of Varian (1989), Kim and Verrecchia (1991), Harris and Raviv (1993) and Kandel and Pearson (1995) all assume that market participants interpret public information differently, views echoed in empirical work by Evans (2002), Evans and Lyons (2003) and Green (2004). The model of Varian (1989), for example, suggests that trading volume increases following an announcement of news; differences in interpretation induce an increase in both buyer *and* seller initiated trades following a news release. However, this is not the same as the results of Evans and Lyons (2003) and those reported in Chapter 3, which both find that 'good' news induces positive order flow. Positive order flow can simply result from an increase in buyer initiated trades with no change in the number of seller initiated transactions. Also, trading volume can increase without any change in order flow.[3] By splitting up transactions into *buyer* and *seller* initiated trades, I can examine the effects of news on both order flow and trading volume simultaneously.[4]

Prices and volatility The documented increase in exchange rate volatility following a news release suggests that it takes some time for the information contained in the release to be incorporated fully. On the other hand, the first moment (level) effects of news have been found to be almost immediate. Cheung and Chinn (1999), for example, report the results from a questionnaire sent out to FX traders and find that the majority of traders believe the response of exchange rates to news takes place within one minute. The surge in volatility, however, can last up to a couple of hours (Andersen and Bollerslev 1998, Payne 1996). Payne (1996) suggests that the post release volatility partly originates from dealers trading towards their new desired positions, i.e. inventory control channels are present in the propagation of volatility, as well as information channels. However, thus far, little has been done to explore

[3]This illustrates the fundamental difference between order flow and trading volume in the microstructure literature. Order flow is *signed* trading volume, i.e. the number (or value) of buyer initiated trades *minus* seller initiated trades. Trading volume is buyer initiated trades *plus* seller initiated transactions.

[4]The effects of news on order flow, and the mechanism by which news is impounded into prices, are dealt with in more detail in Chapter 3.

these different volatility effects. It is therefore of interest to see whether an increase in exchange rate volatility occurs with or without an associated increase in trading volume. For example, heightened trading activity following news releases suggests that inventory channels may also be at work, in addition to possible information channels.[5]

Events pre-release It is also interesting to note what happens prior to the data announcements. Payne (1996) finds that exchange rate volatility declines significantly in the minutes leading up to the release. This is explained by the fact that dealers, knowing that an information event is imminent, cease trading and wait to see what information is contained in the data release. If traders pull out of the market pre-release, one would expect market spreads to widen in this interval as well, which I am also able to test. It is, however, possible for volatility to increase prior to release, either due to some traders discovering the details of the announcement before the scheduled release time and exploiting their informational advantage[6] or by dealers trading out of exposed positions. If exchange rate volatility increases before announcements, it will be interesting to see what happens to trading volume during these periods. On the other hand, if exchange rate volatility falls immediately before data announcements, one may expect trading volume at that time to be similarly reduced.

These hypotheses on the effects of scheduled news releases on exchange rates, trading activity, spreads and volatilities, both pre and post release can all be tested using the data available in this study. The main results of the chapter are given below.

1. Announcements of news that are unexpected cause significant changes in the level of the exchange rate. Consistent with rational expectations models, prices adjust very quickly, usually within a couple of minutes.

2. News releases also cause significant increases in trading volume. An announcement of either good or bad news for a currency leads to significant increases in both buyer *and* seller initiated trades. Even when the news has no effect on the price

[5]This may occur if transactions that restore traders' desired inventory levels are only worked through slowly.

[6]This is unlikely to be the case. Careful measures are put in place to prevent such informational leakages. See Ederington and Lee (1993) for a brief summary of the procedures used in the US to disseminate data. In the UK, Reuters disciplined two of its staff after they published UK retail sales figures one hour early in 2003. (Daily Telegraph, 6th November 2003). Such premature announcements are accepted as being very rare.

of FX, trading activity still increases significantly. However, trading volume is only significantly affected for a few minutes following the release.

3. If anything, news releases cause market inside spreads to widen, but any changes are small and statistically insignificant.

4. News causes exchange rate volatility to increase and these second moment effects die out relatively slowly; the half-life of any news effects is approximately 15 minutes.

5. As well as the effects of news, this chapter also reports asymmetric effects of buys and sells in FX markets. A market buy order, for example, increases both bid and ask prices but increases the ask more than the bid. Using impulse response analysis, a buy or sell shock is found to significantly increase the spread but these effects are short lived.

6. The cross market effects of trades are also documented. For example, a buy shock in USD/EUR causes significant changes in not only USD/EUR prices, but also GBP/EUR and USD/GBP rates. Variance decompositions also show the importance of these cross market effects.

The outline for the rest of the chapter is as follows. Section 2.2 describes the data, Section 2.3 introduces the model used to analyse the effects of macroeconomic news releases and Section 2.4 examines the first moment effects of news, as well as describing how we convert the model of Section 2.3 into one that analyses spreads and trading volume. Section 2.5 discusses how the second moment effects of news releases are investigated, Section 2.6 discusses the results, placing them within the existing finance literature and Section 2.7 concludes. Since the model used in Section 2.3 is quite involved, a number of issues surrounding model building and testing are relegated to an appendix, Section 2.A.

2.2 Data

Throughout this thesis, extensive use is made of data from two sources. Firstly, data on foreign exchange activity are used, and the second source contains scheduled macroeconomic data releases. In order to avoid repetition, the data are described at length here, and in subsequent chapters, the reader will be referred to this section when data are being introduced.

2.2.1 Foreign exchange data

The FX data used in this study, and also used in Chapters 3 and 5, are taken from the Reuters D2000-2 electronic trading system, one of the two dominant brokered trading platforms used in the inter-dealer spot FX markets.[7] Thus the data contain no information on customer-dealer FX trades or on direct (i.e. non-intermediated) trades between dealers. Moreover, it should be noted that the trades occurring on D2000-2 should be regarded as public in the sense that they are visible to anyone looking at a D2000-2 screen as they occur. For a full description of the segments of the spot FX market and the data available from each see the excellent descriptions contained in Lyons (2001).

The raw data is composed of two datasets for each of the three major floating rates: USD/EUR (dollars per euro), GBP/EUR (pounds per euro) and USD/GBP (dollars per pound). The first dataset contains transaction level information, including a time stamp for every trade, a variable indicating whether the trade was a market buy or sell and the transaction price. Thus unlike many papers in equity market analysis, I do not need to make use of potentially inaccurate, ad hoc algorithms to assign trade direction. The samples for USD/EUR and USD/GBP cover a period of ten months from 28th September 1999 to 24th July 2000. The GBP/EUR sample is somewhat shorter, covering the eight month period from 1st December 1999 to 24th July 2000. Unfortunately, no information on traded quantities is available. Thus the order flow (and trading volume) measures are simply the difference between (sum of) the number of buyer-initiated trades and seller-initiated trades in an interval, rather than being calculated from traded quantities. However, to the extent that earlier work has shown little size variation in trades on this dealing system (Payne 2003a) and that in other applications it is the number rather than aggregate size of trades that has been shown to matter for prices and volatility (Jones, Kaul, and Lipson 1994, Fleming 2001, Green 2004), it is expected that this limitation will not distort the results. Furthermore, even when both the number and size of trades have been available, research has often focussed on the former measure of trading activity (Hasbrouck 1991a).

The second dataset contains quote level information, where a new entry is included in the data every time there is a change in the front end of the limit order book, i.e. if there is a change in the best bid or ask price, or if the quantities available at these prices are altered. However, the dataset includes only the time stamp and the best bid and ask

[7]The other electronic trading system is that of EBS and together they account for between 85% and 95% of all interdealer trading. See Bank for International Settlements (2001).

prices, so again, no information on depths is available.[8]

In this chapter, and also in Chapters 3 and 5, I choose a 1 minute sampling frequency, i.e. at the end of each minute of my sample I record the last price (transaction, bid and ask) in each exchange rate plus the trade variables (number of buys and number of sells), from which order flow and volume can be derived. Certain sparse trading periods are also removed from the sample. These include weekends, the overnight period, defined as 1800 to 0600 GMT (BST in the summer months) where trading activity was found to be very thin and some public holidays including Christmas, New Year, Easter (Good Friday and Easter Monday) and the May Day bank holiday. Periods where the D2000-2 data feed broke down were also excluded. These periods are defined as those where no transactions (and no bid or ask price changes) occurred for at least thirty minutes during the day in any of the three FX markets.[9] This filtering process reduced the total number of observations to 124,259 for the USD/EUR, 97,158 for the GBP/EUR and 124,997 for the USD/GBP FX markets.[10] For the model of Section 2.3, where the prices in the three markets are determined jointly, only the periods where no breakdowns in any of the three exchange rate data feeds were considered. This resulted in 90,270 data observations. Table 2.1 contains statistical information on exchange rate returns, defined as 100 times the logarithmic difference in prices (bid, ask and transactions), transaction frequencies and order flows for the filtered data sample.

Of interest at this stage is the fact that for each of the three exchange rates, the variances of the one-minute returns calculated from transactions prices are much smaller than the variances of bid or ask returns. One may have expected the return series from transactions prices to be greater due to the natural volatility induced by the bid-ask bounce. This finding is consistent with Bollerslev and Domowitz (1993), who find that the variance of bid or ask returns is larger than that calculated from a series that proxies transaction price returns. However, it may be the case that the bid and ask return volatilities are greater due to the Reuters D2000-2 trading platform going 'off-market', whereby the bid (ask) price falls (rises) to uncompetitive levels and trading moves to the more competitive EBS platform (especially in the two EUR markets, where EBS is dominant) until liquidity

[8]The bid and ask quotes are firm/tradable, showing prices at which dealers can trade, rather than indicative quotes used in a number of previous studies. In Chapter 4, a more detailed Reuters data-set is considered, which does include information on depths. However, this is a much shorter dataset than that used here.

[9]A number of different windows were used when filtering the data but the 30 minute window appeared to work well, removing enough of the missing data periods, while including the genuinely inactive times. However, the results were not susceptible to changes in this window length.

[10]The substantial reduction in the number of observations for the GBP/EUR market is due to the fact only eight, rather than ten, months of data were available.

Table 2.1
Summary Statistics of Exchange Rate Returns and Trades

Variable	Mean	Variance	Skewness	Kurtosis	Autocorrelation (lags)			
					1	5	10	20
USD/EUR market								
Prices								
returns (transactions price)	-0.000153	0.000829	0.0621	31.77	-0.0254*	-0.0139*	-0.0122*	-0.00210
ask returns	-0.000191	0.00185	-1.10	269.90	-0.198*	-0.00412	0.00375	-0.00325
bid returns	-0.000123	0.00187	-0.136	213.62	-0.207*	-0.00498	0.00729*	0.0108*
absolute returns (trans price)	0.0164	0.000559	4.66	53.73	0.217*	0.128*	0.106*	0.0862*
absolute ask returns	0.0199	0.00145	12.65	409.45	0.261*	0.0939*	0.0632*	0.0571*
absolute bid returns	0.0201	0.00146	11.59	321.88	0.260*	0.116*	0.0893*	0.0553*
Trades								
buys	2.03	7.64	2.74	15.48	0.462*	0.282*	0.235*	0.199*
sells	1.99	7.62	2.74	15.34	0.440*	0.261*	0.214*	0.171*
volume	4.02	19.01	2.41	16.17	0.573*	0.420*	0.366*	0.304*
order flow	0.0361	11.86	-0.386	20.66	0.232*	0.0296*	-0.00235	0.000158

Continued over

Table 2.1
Summary Statistics of Exchange Rate Returns and Trades, (cont.)

Variable	Mean	Variance	Skewness	Kurtosis	Autocorrelation (lags)			
					1	5	10	20
GBP/EUR market								
Prices								
returns (transactions price)	$-0.0^{4}472$	0.000688	-0.00592	41.33	-0.0284*	-0.00316	-0.00563	-0.00447
ask returns	$-0.0^{4}310$	0.00145	2.50	211.84	-0.188*	-0.0208*	-0.00387	-0.00497
bid returns	$-0.0^{4}236$	0.00181	0.962	180.25	-0.229*	0.00288	-0.0155*	-0.0182*
absolute returns (trans price)	0.0150	0.000462	5.13	74.26	0.213*	0.112*	0.0946*	0.0770*
absolute ask returns	0.0187	0.00110	11.08	339.95	0.225*	0.0930*	0.0782*	0.0517*
absolute bid returns	0.0198	0.00142	10.92	268.02	0.258*	0.0936*	0.0960*	0.0595*
Trades								
buys	1.91	5.76	2.27	11.21	0.488*	0.309*	0.256*	0.214*
sells	1.76	5.12	2.26	10.92	0.467*	0.286*	0.243*	0.209*
volume	3.67	14.78	2.11	10.65	0.612*	0.438*	0.371*	0.316*
order flow	0.151	7.07	0.165	7.05	0.186*	0.0174*	0.00493	0.00558

Continued over

Table 2.1

Summary Statistics of Exchange Rate Returns and Trades, (cont.)

Variable	Mean	Variance	Skewness	Kurtosis	Autocorrelation (lags)			
					1	5	10	20
USD/GBP market								
Prices								
returns (transactions price)	$-0.0^{4}817$	0.000351	-0.151	37.10	0.00386	-0.00754*	-0.0123*	-0.00830*
ask returns	$-0.0^{4}944$	0.000495	-0.123	102.39	-0.132*	-0.0142*	-0.00418	0.0126*
bid returns	-0.000107	0.000429	-0.961	93.86	-0.0879*	-0.00440	-0.00681*	-0.00135
absolute returns (trans price)	0.0113	0.000223	4.84	73.74	0.231*	0.128*	0.0990*	0.0855*
absolute ask returns	0.0131	0.000325	8.44	210.33	0.268*	0.0806*	0.0604*	0.0579*
absolute bid returns	0.0125	0.000272	7.61	206.66	0.226*	0.0825*	0.0587*	0.0477*
Trades								
buys	2.25	7.62	2.45	12.93	0.520*	0.316*	0.253*	0.197*
sells	2.18	7.56	2.49	13.09	0.494*	0.281*	0.230*	0.178*
volume	4.43	21.54	2.43	13.06	0.645*	0.432*	0.355*	0.282*
order flow	0.0633	8.38	-0.0304	7.86	0.163*	-0.00104	-0.00181	-0.00372

Notes: The data cover the eight month period from 1st December 1999 to 24th July 2000. All statistics are based on data sampled at the one minute frequency. The USD/EUR exchange rate is defined as the number of dollars (numerator currency) per euro (denominator currency) and similarly for the other rates. Returns are defined as 100 times the first difference of the logarithm of the exchange rate. Buys (sells) refer to trades where the initiator is purchasing (selling) the donominator/commodity currency; the euro in USD/EUR and GBP/EUR and sterling in USD/GBP. * denotes significance at the 5% level or less.

is restored.[11] This not only explains why returns from bid and ask prices have greater variances than those calculated from transactions prices, but also explains the greater kurtosis in the bid and ask return series.[12] This is also consistent with Goodhart and Payne (1996), who suggest that the volatility of the ask or bid series may be greater since the removal of the best ask quote (through an executed market order or the cancellation of the best limit order) will naturally result in the next best ask price lying further down the book if the book is thin. Only when liquidity is restored on the book will anyone wish to trade, hence explaining why transaction prices are less volatile than quote prices.

2.2.2 Macroeconomic data releases

The second component of the dataset consists of euro-area, UK and US macroeconomic information announcements along with expectations data for each of these releases. The expectations data are used to construct the "news" or "surprise" component of each individual announcement and from these numbers, aggregated news variables are constructed, also on a 1 minute sampling frequency, for each of the euro-area, UK and US. The announcement information is aggregated since with a maximum of 10 months of data (and thus ten releases for each individual announcement) it would be difficult to identify statistically significant effects from the individual news series.[13] The aggregated news variable is constructed by scaling each series of surprises by that series' surprise standard deviation. So the standardised news for series v is given by

$$S_{v,t} = \frac{A_{v,t} - E[A_{v,t}]}{\hat{\sigma}_v} \tag{2.1}$$

where $A_{v,t}$ is the actual announcement of the data, $E[A_{v,t}]$ is the market expectation of the announcement obtained from Standard and Poors and $\hat{\sigma}_v$ is the sample standard deviation of the forecast error, $A_{v,t} - E[A_{v,t}]$.[14] Each surprise series is then signed (i.e. multiplied by

[11]For this point I thank Richard Lyons.

[12]The fat tails of the bid and ask returns may be due to the 'outliers' where the Reuters platform goes 'off' and 'on-market'.

[13]In this chapter, only 8 months of data are employed.

[14]An identical standardisation is performed in Balduzzi, Elton, and Green (2001), Andersen, Bollerslev, Diebold, and Vega (2003) and Green (2004). The unanticipated component of news in our data was created from consensus forecasts obtained from Standard and Poors. The data/shocks we use are therefore the average of shocks across participants, implying that there may exist a positive shock for one trader and a negative shock for another. Heterogeneous forecasts will therefore generate a motive for trade even when the announcement is exactly equal to the median forecast. Little can be done about this shortcoming in the analysis. However, Chaboud et al (2004) find that when expectations are dispersed, as measured by the standard deviation of practitioners' forecasts, this has little impact on trading activity relative to when expectations are common. Therefore this problem should not distort the results.

Table 2.2
The Effect of a Country's Data Releases on its Exchange Rate

News emanating from	News in the form of an increase in: Prices/Money	Output	CA surplus
Euro-area (effect on EUR)	App.	App.	
UK (effect on GBP)	App.	App.	App.
US (effect on USD)	Dep.	App.	App.
Theoretical predictions:			
Monetary (portfolio balance)	Dep.	App.	App.
Keynesian model		Dep.	App.
Central Bank reaction function	App.	App.	

Notes: 'App.' refers to an appreciation of that region's currency. 'Dep' refers to a depreciation. 'CA' = current account.

+1 or -1) depending on its effect on the exchange rate.[15] The series is multiplied by +1 if greater than expected news, positive $S_{v,t}$, causes the domestic currency to appreciate, and by -1 if it causes a depreciation. This thesis is not aimed at evaluating competing models of exchange rate determination via examination of the effects of news on rates. Nor do I try to evaluate the importance of different news releases via their effect on exchange rates. The focus is simply on the effects of news on FX activity, and in Chapter 3, the focus is on the role played by order flow in the reaction of exchange rates to news. Therefore, I accept the way the markets appear to interpret the data releases without trying to explain the direction or strength of response to individual releases. However, to place the findings within the extant literature on exchange rate determination, I summarise the effects from each type of announcement (monetary/price, real/output and trade) in Table 2.2 and also give the predictions from a number of theoretical exchange rate models. For a more complete discussion of these issues, see Hoffman and Schlagenhauf (1985).

From the preliminary empirical analysis, there does not appear to be one theory which explains all the exchange rate responses to the different data releases. Indeed one might suggest that there is a contradiction in the way the markets appear to interpret price/money

[15]These signs were determined via standard event-study type analysis of exchange rate reactions to individual announcement surprises. In this way, the signing was done on a series by series basis, rather than observation by observation.

Table 2.3
US PPI Announcements and 'Good' and 'Bad' News

Date	Actual	Expected	'News'	St'ised	Signed	Good	Bad
15th Oct 1999	1.1	0.5	0.6	1.98	-1.98	-	-1.98
10th Nov 1999	-0.1	0.1	-0.2	-0.66	0.66	0.66	-
10th Dec 1999	0.2	0.2	0.0	0.00	0.00	-	-
13th Jan 2000	0.3	0.2	0.1	0.33	-0.33	-	-0.33
17th Feb 2000	0.0	0.2	-0.2	-0.66	0.66	0.66	-
16th Mar 2000	1.0	0.6	0.4	1.32	-1.32	-	-1.32
13th Apr 2000	1.0	0.6	0.4	1.32	-1.32	-	-1.32
12th May 2000	-0.3	-0.2	-0.1	-0.33	0.33	0.33	-
9th Jun 2000	0.0	0.3	-0.3	-0.99	0.99	0.99	-
14th Jul 2000	0.6	0.6	0.0	0.00	0.00	-	-

Notes: In all cases, US PPI announcements are made at 0830 ET (1330 GMT) and are defined as month on month percentage changes. 'News' is the difference between Actual and Expected figures. St'ised is standardised news ($S_{PPI,t}$) and is equal to 'News' divided by its standard deviation. Since greater than expected PPI figures tend to cause the US dollar to depreciate, the news is signed by multiplying by -1. Therefore, good news is that which causes the US dollar to appreciate, and bad news, that which leads to a depreciation.

data emanating from the US compared to that from the UK and euro-land. An announcement of higher than expected inflation in the US tends to cause the dollar to depreciate, whereas similar data coming from the UK or euro-land leads to an appreciation of that region's currency. The US case is consistent with a monetary model of the exchange rate while the other two cases are consistent with a Central Bank reaction function where greater than expected inflation or output increases the likelihood of a rise in domestic interest rates causing an increase in the demand for the domestic currency. Although surprising, it is not inconceivable that different rates respond in different ways to the same type of news. Different macroeconomies *may* behave differently and it is quite possible that markets expect policy makers in different regions to respond differently to the same "shock".[16]

In this chapter, each series is split up into 'good' or 'bad' news, denoted $S^G_{v,t}$ and $S^B_{v,t}$ respectively. Good (bad) news is therefore defined as any individual macro announcement which when released from one region, causes that region's domestic currency to appreciate (depreciate). Then, to obtain the aggregated good and bad news variables for each region, the sum is taken over the signed, good or bad standardised surprise numbers across

[16]Indeed the monetary authorities in the UK and euro-area have explicit inflation targets. This is not so for the US. Hence the differing effects of inflation/monetary data on the regions' currencies may not be too surprising. For this point I thank Paolo Vitale.

announcements, i.e. good UK news is the sum of the $S_{v,t}^G$ terms over all the v series relating to the UK economy. Of course, this technique will lead to the effects of more important announcements being diluted through the inclusion of less important releases, but this dilution is necessary to generate a meaningful sample size for news in each geographical area. To give an example of the signing procedure, consider the 10 US PPI announcements from October 1999 to July 2000, given in Table 2.3. This demonstrates how the good and bad US PPI data are calculated, which help form the aggregated good and bad US news variables.[17] In this exercise, even though the triangle of rates between the euro, dollar and UK pound are examined, only macroeconomic news emanating from the UK and US are considered.

The reason for excluding euro-area data releases is because, as documented in Chapter 3, only two euro-area series appeared to have significant first moment effects; industrial production and M3, giving a total of 15 data points. Splitting these into 'good' and 'bad' news would give too small a number of observations for valid inference. Hence, I concentrate on news from only the UK and US.[18]

In Section 2.5 the effects of macroeconomic news on FX volatility are considered. Since it may be possible for some news to have no first moment (level) effects but to have significant effects on volatility, two types of news are created for each of the UK and US. Primary news will be made up of those releases that have significant first moment effects, while secondary news will include the other data series that were found to have no effects on the level of the exchange rate. As volatility effects will be examined in Section 2.5, it is only the size of the news surprise that is of interest. Therefore, news is not split into good or bad, rather the absolute value of the relevant $S_{v,t}$ series is used. The sets of macroeconomic announcements that are included in the primary and secondary categories, for both the UK and US are given in Table 2.4.[19]

2.3 Empirical Methodology

The framework used to analyse the effects of scheduled macroeconomic news is a standard structural vector autoregression. The VAR methodology, introduced by Sims (1980), was

[17]Table 2.3 also shows two instances (in December and July) when the news announcement was equal to the market expectation. In these cases, for the purpose of the model presented in Section 2.3, there is no news. Surprisingly, this happened very rarely.

[18]Euro-area news is considered in Chapter 3, but it is not split into good or bad.

[19]US nonfarm payroll employment was also considered, which is released at the same time as the US unemployment rate, but was found to have no first moment effects.

Table 2.4
Description of 'Primary' and 'Secondary' Macroeconomic Data

		UK announcements		
Announcement[1]	Sign[2]	Reported as[3]	Dates	Local time GMT
Primary news				
RPIX	+1	Y/Y % change	12/99 - 7/00	08:30/09:30
Retail Sales	+1	M/M % change	12/99 - 7/00	08:30/09:30
Global Trade	+1	GBP (bn.)	12/99 - 6/00	08:30/09:30
Prov. M4	+1	M/M % change	12/99 - 7/00	08:30/09:30
Secondary news				
PPI		M/M % change NSA	12/99 - 7/00	08:30/09:30
Ind. Prod.		M/M % change	12/99 - 7/00	08:30/09:30
Unemployment		000's	12/99 - 6/00	08:30/09:30
Curr. Acc.		GBP (bn.)	12/99 - 6/00	08:30/09:30

		US announcements		
Announcement[1]	Sign[2]	Reported as[3]	Dates	Local time ET
Primary news				
CPI	-1	M/M % change	12/99 - 7/00	08:30
PPI	-1	M/M % change	12/99 - 7/00	08:30
Unemployment	-1	%	12/99 - 7/00	08:30
Trade bal.	+1	USD (bn.)	12/99 - 7/00	08:30
Secondary news				
Retail Sales		M/M % change	12/99 - 7/00	08:30
Ind. Prod.		M/M % change	12/99 - 7/00	09:15
M3		Change USD (bn.)	12/99 - 7/00	16:30

Notes:

1. All announcements are made monthly, except UK current account figures, which are released quarterly.

2. When forming the signed aggregate news variable, each series was multiplied by +1 (-1) if greater than expected news causes the domestic currency to appreciate (depreciate). When splitting the series into good and bad news, good news was defined as any macro announcement which caused that country's currency to appreciate.

3. M/M % change: month on month percentage change. 3M/3M % change: three month on three month percentage change. Y/Y % change: year on year percentage change. NSA: non-seasonally adjusted.

implemented in the microstructure literature by Hasbrouck (1991a) and has become a common tool used to analyse the effects of trades on asset prices. For recent examples, see Dufour and Engle (2000) and Engle and Patton (2004) for stocks, Evans (2002), Evans and Lyons (2002b) and Payne (2003a) for currencies and Cohen and Shin (2003) and Green (2004) for treasuries.

2.3.1 Cointegrating structural VAR

The standard framework is extended in this paper by examining the bid and ask returns separately, as in Engle and Patton (2004), and by considering the number of buys and sells separately. This therefore allows one to examine exchange rate returns and buying and selling pressure, as well as bid-ask spreads, order flows and trading volume simultaneously once an appropriate rotation is implemented, used in Engle and Patton (2004) and also described in Section 2.4.5. The effects of news are analysed by including a series of exogenous right hand side regressors, splitting the announcements into good and bad news as described above. The basic model to be estimated is given in (2.2).

$$
\begin{aligned}
\underbrace{\begin{bmatrix} \Delta P_t \\ T_t \end{bmatrix}}_{\Delta Y_t} &= c + \kappa t + \alpha\gamma' \underbrace{\begin{bmatrix} P_{t-1} \\ \sum_{r=1}^{t-1} T_r \end{bmatrix}}_{Y_{t-1}} + \begin{bmatrix} 0 & \beta \\ 0 & 0 \end{bmatrix} \underbrace{\begin{bmatrix} \Delta P_t \\ T_t \end{bmatrix}}_{\Delta Y_t} + \sum_{i=1}^{m} \theta_i \underbrace{\begin{bmatrix} \Delta P_{t-i} \\ T_{t-i} \end{bmatrix}}_{\Delta Y_{t-i}} \\
&+ \sum_{R}\sum_{j=1}^{n} \phi_j D_{t+j}^R + \sum_{R}\sum_{k=0}^{p} \delta_k N_{t-k}^R + \epsilon_t \qquad R = US,\ UK \qquad (2.2) \\
\underbrace{\begin{bmatrix} \Delta P_t \\ T_t \end{bmatrix}}_{12\times 1} &= \Delta Y_t = \begin{matrix} [\Delta ED_t^{ask} & \Delta ED_t^{bid} & \Delta ES_t^{ask} & \Delta ES_t^{bid} & \Delta SD_t^{ask} & \Delta SD_t^{bid} \\ B_t^{ED} & S_t^{ED} & B_t^{ES} & S_t^{ES} & B_t^{SD} & S_t^{SD}]' \end{matrix}
\end{aligned}
$$

Where ΔP_t is a 6×1 vector containing the return series and T_t is a 6×1 vector of trade variables. The model in (2.2) incorporates a number of elements which are discussed in turn below. ED_t^{ask} is the log of the USD/EUR ask price (dollars per euro) at the end of minute t and ED_t^{bid} is the log of the USD/EUR bid price in the same interval. ES_t^{ask} and ES_t^{bid} are the logs of the ask and bid prices for the GBP/EUR exchange rate (pound sterling per euro) and similarly for the USD/GBP market, denoted SD_t^{ask} and SD_t^{bid}. ΔED_t^{ask} is therefore the log first difference of the USD/EUR ask price; the USD/EUR ask return for minute t. B_t^{ED} and S_t^{ED} are the (deseasonalised) number of buys and sells in minute t respectively in the USD/EUR market and similarly for the GBP/EUR and USD/GBP markets, denoted B_t^{ES}, S_t^{ES} and B_t^{SD}, S_t^{SD} respectively. The reason for deseasonalising the number of buys and sells is to take into consideration the obvious 'M' shaped pattern over the trading day, which will be discussed in Section 2.3.3 below. ΔY_t is then a 12×1 vector of returns and trade variables (buys and sells) and Y_{t-1} is therefore a 12×1 vector of log exchange rate levels and (deseasonalised) cumulative buys and sells. Y_{t-1} is included in the regression to allow for any cointegration that may be present in the system. γ is then a $12 \times h$ matrix of the h cointegrating vectors and α is a $12 \times h$ matrix of speed of adjustment coefficients. See Section 2.3.4. D_t^R is a dummy variable that takes the value of unity if there is an announcement of news from region R in minute t. These dummy variables are only included as leads, which allows one to see what happens leading up to the announcement of news. N_t^R is a 2×1 vector of quantified 'good' and 'bad' news from region R in minute t, described in Section 2.2. Model (2.2) therefore allows me to examine what happens leading up to an announcement of news and also allows me to see the effects of a data release that was unanticipated by the market.

2.3.2 Impacts of buy and sell orders separately

Standard microstructure models suggest that it is one way buying or selling pressure, order flow, which should cause asset prices to change. In the model of Kyle (1985) for example, positive order flow (net buying pressure) causes the market maker to revalue the asset price since it suggests the presence of informed trading; another market participant has private information that the asset is undervalued and hence wishes to buy. These models, however, say nothing about the effects of buys and sells separately. One may expect buying pressure to cause both the ask and bid prices to increase.[20] But is it the case that buying pressure has a symmetric effect on buys and sells? Engle and Patton (2004) find that buys have a greater effect on the ask price than on the bid and vice versa for sells. This is also demonstrated in Biais, Hillion, and Spatt (1995), who examine the microstructure of the Paris Bourse, and can be explained by the dynamic limit order book models of Parlour (1998) and Foucault (1999). Simple 'barrier' arguments can also be used to explain the phenomenon; since a market buy order is executed at the ask price in the limit order book, the price at which the specialist dealer is willing to sell the asset,

[20] However, if the trade increased asymmetric information and if this effect was great enough, the specialist dealers in the FX markets may respond to an increase in buying pressure by increasing the ask price considerably but *lowering* the bid price, i.e. in the presence of greater asymmetric information, the dealer increases the mid-quote (his/her expected value of the asset) and widens the spread.

then buy orders drain liquidity at the ask side of the book and effectively should have no effect on liquidity at the bid side. The draining of ask side liquidity will naturally raise the ask price, explaining the asymmetric effect of buy (sell) orders on ask and bid prices. Such dynamic effects in the foreign exchange market have been documented by Daníelsson and Payne (2002a) but have not, as yet, been examined in a VAR framework.

The asymmetric effects of buy and sell orders will be shown in the 6×6 β submatrix that premultiplies the ΔY_t vector on the right hand side of (2.2). The ask and bid price returns in each of the three FX markets will therefore be allowed to depend on the number of buy and sell orders separately. Since the returns (bid and ask) in market k are allowed to depend, not only on contemporaneous buys and sells in market k, but also on the buys and sells in market l ($\neq k$), then model (2.2) also allows me to examine the cross market effects of buys and sells. The cross market effects of order flow have been documented in previous research. Evans and Lyons (2002a) document the role that order flow in one currency has in determining exchange rates in other markets. In particular DEM/USD (Deutsche marks per dollar) and CHF/USD (Swiss francs per dollar) order flows have significant effects on a number of other dollar exchange rates. These cross market effects are also documented in Daníelsson, Luo, and Payne (2002), which considers USD/EUR, GBP/EUR, USD/GBP and JPY/USD markets. Theoretical explanations as to why cross market effects of order flow exist are also presented in Lyons and Moore (2003), which examines the triangle of rates between the US dollar, euro and yen.

However, in (2.2) asset returns are allowed to depend on contemporaneous buy and sell orders but buy and sell orders are not allowed to depend on contemporaneous returns. This is intuitive considering the one minute sampling frequency considered in this exercise but it is possible for intra-minute feedback trading to exist, whereby traders buy or sell in minute t based on asset price changes experienced within that period. Evidence for such high frequency trading strategies is found in Chapter 5, where I model and test for contemporaneous feedback trading. The effect of this feedback trading, although not significant in a statistical sense, is economically large, even when sampling at the one minute frequency. However, little can be done about this problem in the recursively ordered structural VAR in (2.2).[21] For this reason, the recursive ordering of the VAR is used here but the results should be accepted with caution.

[21]In Chapter 5 I estimate the structural (contemporaneous feedback trading) parameter in the USD/EUR market using variables obtained from the GBP/EUR and USD/GBP markets as instruments. This cannot be done here since the cross market effects in the three currencies are modelled explicitly, i.e. GBP/EUR and USD/GBP statistics are used to instrument for themselves in the VAR.

2.3.3 Intra-day patterns and diurnality

As has been documented in a number of studies, there are clear seasonal/intra-day patterns present in high frequency financial data. See Bollerslev and Domowitz (1993), Andersen and Bollerslev (1998) and Hsieh and Kleidon (1996) for examples of seasonal patterns in foreign exchange data. Two notable features of Reuters D2000-2 data are, firstly, the 'M' shaped pattern of trading activity and the 'U' shaped pattern of bid-ask spreads from 0600 to 1800 GMT (BST in the summer months). The reason for only considering data from 0600 to 1800 is because outside of this interval, trading activity is close to zero, suggesting that the vast majority of Asian trading (during the European overnight period) occurs on the competing EBS platform.[22] Simply using the number of buys and sells in the VAR of (2.2), without allowing for the intra-day seasonality will induce a severe mis-specification and bias any estimates of the impact of news.[23]

A simple method used to account for the intra-day patterns is the Fourier Flexible Form (FFF) introduced by Gallant (1981) and used in applications to FX data by Payne (1996), Andersen and Bollerslev (1997a), Andersen and Bollerslev (1998), Andersen, Bollerslev, Diebold, and Vega (2003) and Dominguez (2003).[24] This simply fits the data to a set of deterministic trigonometric terms. The FFFs for series X, average number of buys and sells, and Z, spreads, are constructed as shown in (2.3) and (2.4) respectively.

$$\begin{aligned}
\psi_t^X &= e^+ + \sum_{q=1}^{Q} \left(\zeta_q \cos\lambda_q t + \zeta_q^+ \sin\lambda_q t\right) + \tau W_t \qquad \lambda_q = \frac{2\pi q}{S} \qquad (2.3)\\
X &= \frac{\tilde{B}_t^x + \tilde{S}_t^x}{2} \qquad x = ED,\ ES,\ SD\\
\psi_t^Z &= \tilde{e}^+ + \sum_{q=1}^{Q} \left(\tilde{\zeta}_q \cos\lambda_q t + \tilde{\zeta}_q^+ \sin\lambda_q t\right) \qquad (2.4)\\
Z &= z_t^{ask} - z_t^{bid} \qquad z = ED,\ ES,\ SD
\end{aligned}$$

where S is the number of time intervals in each day over which the FFF is constructed.

[22] See Ito and Hashimoto (2004) for example.

[23] Such biases are particularly pervasive when examining the effects of scheduled news releases, as demonstrated in Payne (1996). When examining the effects of macroeconomic news on FX volatility, Payne (1996) found that estimates of news effects were significantly biased if the diurnal patterns were ignored. Since this thesis is focussed on the effects of news releases, it is vital that these intra-day patterns be taken into consideration.

[24] Alternatively, one can simply deseasonalise by fitting piece-wise linear splines, used by Engle and Patton (2004) or divide each observation by the average value for that time period across all days, used in Payne (2003b). Due to the popularity of using the FFF, this method is used here.

(2.3) therefore creates the Fourier Flexible Form for the trade variables, defined as the average number of buys and sells in each market. For example $\psi_t^{ED\,trades}$ is the intra-day pattern of $\frac{\tilde{B}_t^{ED}+\tilde{S}_t^{ED}}{2}$, where $\tilde{B}_t^{ED}$ is the total number of buyer initiated trades in the USD/EUR market in minute t. Once the intra-day pattern has been created, the series is deseasonalised by subtracting ψ_t^X. For example B_t^{ED} and S_t^{ED} in (2.2) are the deseasonalised series created from (2.5) below.

$$
\begin{aligned}
B_t^{ED} &= \tilde{B}_t^{ED} - \psi_t^{ED\,trades} \\
S_t^{ED} &= \tilde{S}_t^{ED} - \psi_t^{ED\,trades}
\end{aligned} \tag{2.5}
$$

In (2.3), the FFF is created for the average number of buys and sells and means that we are imposing the same pattern for buys as we are for sells. This is done, firstly because their is no intra-day seasonality in order flow, as one would expect, and secondly, when defining the FFF to be identical for buys and sells, the difference between B_t^{ED} and S_t^{ED} will be equal to actual order flow and will not be contaminated by any differences in the buy and sell FFFs. In a similar fashion, (2.4) creates the intra-day pattern for each of the three spreads; the difference between the log ask and log bid prices, which will be used in the cointegration analysis of Section 2.3.4. However, in (2.3), a 4×1 vector of dummy variables, W_t, is included, each taking the value of unity for the different days of the week from Tuesday to Friday, i.e. the first dummy variable takes the value of unity if the observation occurred on Tuesday, the second takes the value of unity for Wednesday, etc. This therefore allows me to account for the day of the week effects in trading activity, discussed in Section 2.3.3.1.

Even though the VAR of (2.2) considers data from 0600 to 1800, implying 720 minutes per trading day, the FFF is constructed from 0500 to 1900 (840 minutes) in order to allow for any possible discontinuities over the overnight period. This is particularly noticeable when constructing the intra-day spread patterns. These spread patterns, together with the FFFs and intra-day patterns for the average of buys and sells in each of the USD/EUR, GBP/EUR and USD/GBP markets are shown in Figure 2.1. The figure also shows the 25th and 75th percentiles in order to get an idea of the dispersion around these averages. Due to the possible effects that news announcements may have on buys, sells and spreads, the FFF for these series were only calculated from non-announcement days, in order to

Figure 2.1
Intra-day Patterns of Spreads and the Average of Buys and Sells in the USD/EUR, GBP/EUR and USD/GBP Markets

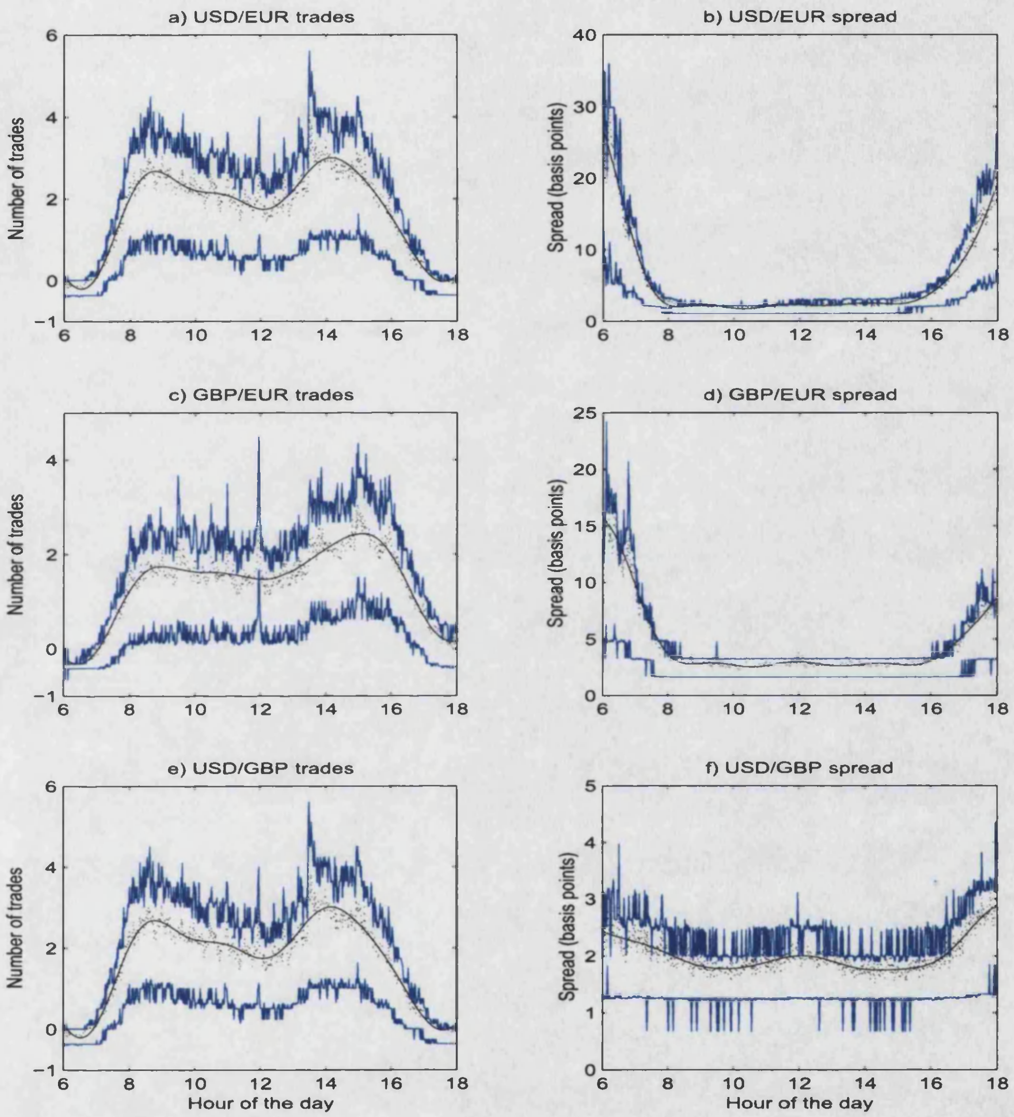

Notes: The figures plot the intra-day patterns of trades (defined as the average of buys and sells) and the spread from 0600 to 1800 GMT (BST in the summer months) for the three exchange rates. The black dotted line shows the intra-day average (mean), the solid black line shows the FFF (calculated using (2.3) and (2.4)) and the blue lines show the 75th and 25th percentiles. The day of the week effect is subtracted from each of the trade variables in order to give the intra-day pattern for a representative day. The subtraction of the day of the week effect results in the 25th percentile often becoming negative, especially in the early morning and late evening. All patterns were calculated using non-announcement days.

wipe out the distortionary effects that announcement times may have on these variables.[25] For spreads, the intra-day average (mean) occasionally lies above the 75th percentile and is purely due to the high skewness present in this variable.[26] The 'M' shaped pattern of trades is notable in each market. Trading picks up when London opens and there is a further increase in the early European afternoon, when New York opens. In both the euro markets, USD/EUR and GBP/EUR, the low levels of trades in the early morning and late afternoon are associated with large spreads. This is consistent with the fact that the Reuters D2000-2 system is not the dominant platform in Asian trading; during peak European trading, there are approximately 3 buys (sells) per minute in USD/EUR (Figure 2.2a) implying one trade (buy *or* sell) every 10 seconds, and spreads are consistently low at 2 to 2.5 basis points (Figure 2.1b). However, at 0600 there is, on average, one buy (sell) every 7 minutes and the spread is approximately 25 basis points. On the other hand, despite the similar 'M' shaped pattern of USD/GBP trades (2.2e) there is still one trade (buy or sell) every minute at 0600 and the spread is still less than 3 basis points. This suggests that Reuters still holds a commanding share of the spot USD/GBP market in Asia, even if the size of this market is relatively small in Asia/Pacific trading.

2.3.3.1 Day of the week effects

In order to show how trading activity changes over the course of the week, the results of the trade variable regression in (2.3) are reported in Table 2.5. For convenience, only the parameter values and t-stats are reported for the constant and the day of the week dummies, as little intuition can be gained from the parameters on the trigonometric terms. The results show that for each of the USD/EUR, GBP/EUR and USD/GBP markets, trading activity is significantly higher from Tuesday to Thursday. Trading is also significantly higher on Friday for the two GBP markets but in the USD/EUR market, trading is approximately the same as that seen on Monday. This is also shown in Figure

[25]We want the effects of news to be picked up in the ϕ and δ coefficients in (2.2) rather than to be picked up in the construction of the B_t^{ED} series for example. Since we only consider non-announcement days to construct the FFFs, the number of observations used in their construction fell to 27,370. In Figure 2.1 the day of the week effect is subtracted from each of the trade variables in order to give the intra-day pattern for a representative day. The subtraction of the day of the week effect results in the 25th percentile often becoming negative, especially in the early morning and late evening. The spike in GBP/EUR trading at 1200 GMT is likely to be due to scheduled news being released at this time, which is not included in the list of announcements in Table 2.4.

[26]Spreads have a lower bound at zero and so when the average spread is very low (such as for FX spreads), high skewness will cause the mean to lie well above the median. The step jumps seen in the 25th percentile, especially for USD/GBP, is indicative of the discrete nature of the pricing grid and the very small spreads in FX. As an alternative to the model presented in (2.2), one could model the discrete nature of FX prices using an ordered probit model, such as that used in Hausman, Lo, and Mackinlay (1992).

Figure 2.2
Day of the Week Patterns of the Average of Buys and Sells in the USD/EUR, GBP/EUR and USD/GBP Markets

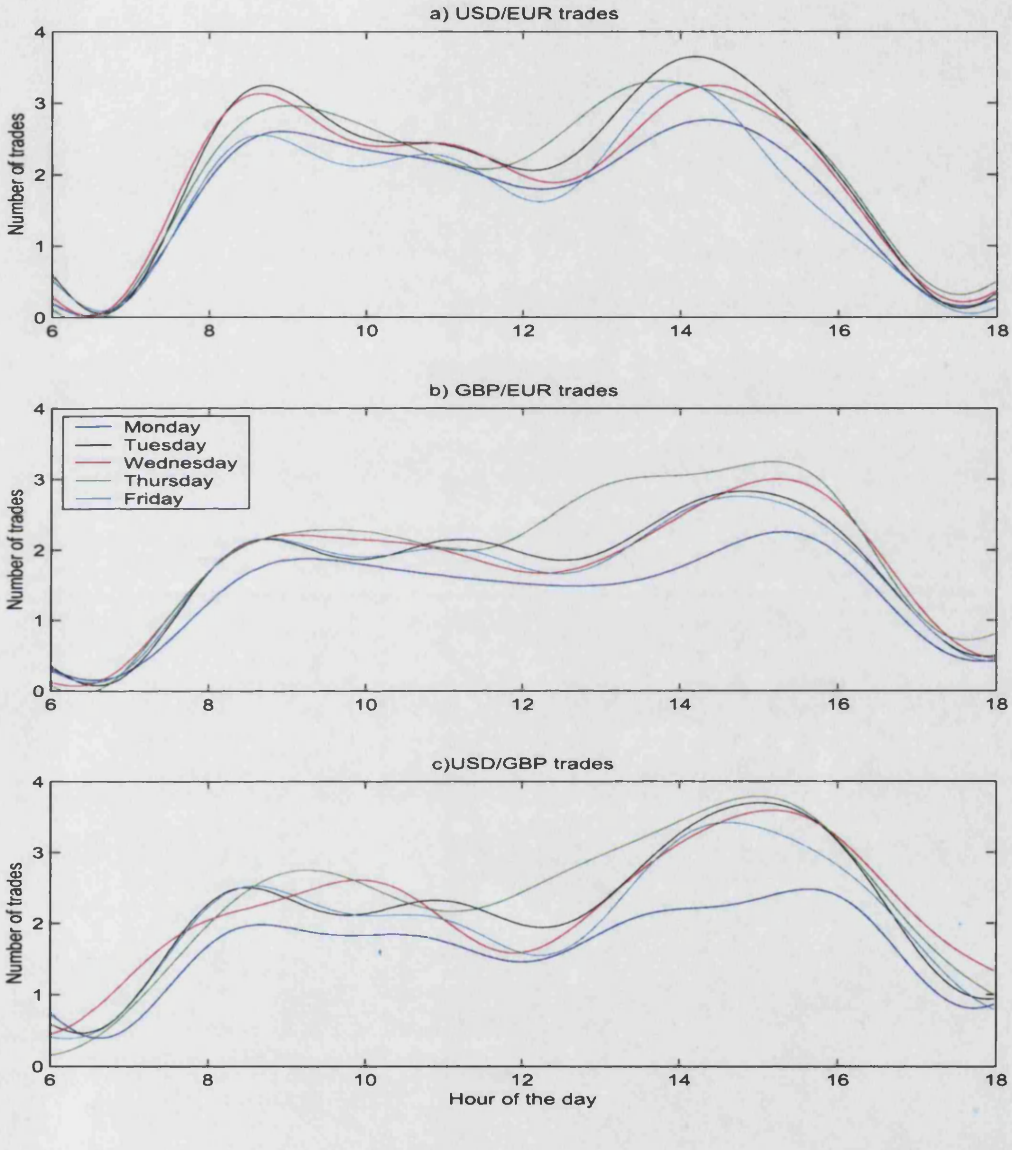

Notes: The figures plot the intra-day patterns of the average number of buys and sells for each day of the week (Monday to Friday) for the three exchange rates. All patterns were calculated using non-announcement days.

Table 2.5
Day of the Week Effects in FX Trading

	USD/EUR trades	GBP/EUR trades	USD/GBP trades
c	1.539	1.201	1.430
	(61.90)	(56.68)	(48.82)
Tuesday	0.395	0.385	0.601
	(10.88)	(12.38)	(14.92)
Wednesday	0.265	0.425	0.625
	(8.16)	(14.78)	(16.07)
Thursday	0.362	0.659	0.773
	(10.17)	(19.94)	(18.56)
Friday	-0.00770	0.328	0.452
	(-0.23)	(10.93)	(11.87)
R^2	0.184	0.140	0.101

Notes: The table shows the regression results from the Fourier Flexible Form model, model (2.3). In each regression, the dependent variable is the average number of buys and sells, and for convenience only the coefficient and t-stats (in parentheses) are given for the constant and the day of the week dummy variables. This is because no intuition can be gained from the coefficients on the trigonometric terms.

2.2, which shows the intra-day patterns of the average of buys and sells on each of the five week days.[27] It is clear that trading activity is lowest on Monday, and also Friday in the USD/EUR market, and this justifies the use of day of the week dummies in the construction of the intra-day patterns.[28]

2.3.4 Cointegrating relationships

In the VAR model of (2.2), each of the variables that make up Y_t were found to be I(1).[29] However, it is quite possible that some of the variables are cointegrated. Intuitively, one would expect the log bid and log ask prices of each exchange rate to be cointegrated with a cointegrating vector of $[1 \ \ -1]'$ implying a stationary log spread and leading to,

[27]Each is calculated from a FFF regression for that day only. For example, when calculating the intra-day pattern of $\frac{\bar{B}_t^{ED}+\bar{S}_t^{ED}}{2}$ for Monday, only Monday's data are used in the regression, hence explaining the slight variation in each of the daily patterns.

[28]No day of the week pattern was found for any of the spread series and therefore no day dummies are included in (2.4). Tests were also performed to see whether spreads were higher on Monday mornings and Friday afternoons, but no significant effects were found.

[29]Tests were performed using the standard Augmented Dickey-Fuller (ADF) procedure (Dickey and Fuller 1979) but are not reported here.

at least, three cointegrating vectors in the system.[30] However, there is also an absence of arbitrage condition that should hold, causing the three exchange rates to be cointegrated; in absence of arbitrage the USD/EUR rate should equal the USD/GBP rate multiplied by the GBP/EUR rate. Consider the case where I have \$1. With this \$1 I can buy euros from a dealer at the price at which the dealer is willing to sell euros, i.e. $USD/EUR : ask$, then buy sterling from a dealer at the price at which he is willing to sell sterling (or buy euros), $GBP/EUR : bid$. Then with this amount of sterling I can buy dollars at the price at which the dealer is prepared to buy sterling, $USD/GBP : bid$. In the absence of arbitrage and assuming a competitive FX market, this amount should be close to (but strictly less than) my initial \$1. Assume for the time being that this relationship holds with equality.

$$
\begin{aligned}
1 &= \frac{1}{USD/EUR : ask} \times GBP/EUR : bid \times USD/GBP : bid \\
\Rightarrow ED^{ask} &= ES^{bid} + SD^{bid}
\end{aligned}
\tag{2.6}
$$

where $ED^{ask} = \log(USD/EUR : ask)$, etc. This then generates a fourth theoretical cointegrating relationship.[31] However, since the vector system, (2.2), has twelve variables, each being I(1), there could be up to eleven cointegrating relationships. Killeen, Lyons, and Moore (2002) and Bjønnes and Rime (2003) suggest that the exchange rate and cumulative order flow should also be cointegrated and find evidence to support this hypothesis. Killeen, Lyons, and Moore (2002) use daily data on FFR/DEM (French francs per Deutsche mark) for 1998 while Bjønnes and Rime (2003) use tick data covering 5 days in March 1998 for the NOK/DEM (Norwegian krone per Deutsche mark) and DEM/USD (Deutsche marks per dollar) markets. The intuition behind these ideas is simple; using Bayesian updating, the quotes given by a dealer at time t will be a weighted average of his/her prior and the current signal (order flow). This quote then becomes the prior when forming the quotes at time $t+1$. In which case all previous order flows will be determinants of the current price and implying that some linear function of price and cumulative order flow will be stationary.

[30]This method is employed by Engle and Patton (2004) who use a VECM in bid and ask prices for US stocks.

[31]One could also demonstrate the absence of arbitrage relationship when making the round trip series of trades from dollars to sterling to euros and back to dollars. This would generate a cointegrating relationship between ED^{bid}, ES^{ask} and SD^{ask}. However, this cointegrating vector would not be independent of the other absence of arbitrage vector and the three cointegrating vectors between bid and ask prices.

Evidence presented in the appendix shows that using the maximum likelihood methods of Johansen (1995), the rank of $\alpha\gamma'$ in (2.2) is found to be equal to 7, i.e. there are 7 linearly independent cointegrating vectors. However, when testing whether these vectors are associated with the 7 hypothesised relationships (3 for the bid/ask spreads, 1 for the absence of arbitrage and the remaining 3 between the level of each exchange rate and cumulative order flow), results are not particularly supportive. ADF tests suggest that (log) bid and ask prices are cointegrated and that the three exchange rates are cointegrated, as one would expect. However, using system methods of cointegration leads one to reject these hypotheses, although the 1% critical values are only just breached. Tests also suggest that the level of the exchange rate (log ask price) and cumulative order flow are not cointegrated. However, despite this, I impose the theoretical cointegrating relationships discussed above. This was done for two reasons. Firstly, if the 7 cointegrating vectors were estimated freely, it becomes impossible to interpret the speed of adjustment coefficients. Instead of showing how each variable responds to the 7 true cointegrating vectors, the 12×7 α matrix will show how each variable responds to a linear combination of these 7 vectors. Therefore, any inference is without meaning. Pesaran and Smith (1998) indeed advocate the imposition of theoretical cointegrating vectors as this allows one to interpret the speed of adjustment coefficients. This is especially so when considering the relationships between bid and ask prices and between the three rates. Efficient markets hypotheses and the well-functioning/liquid nature of FX markets dictate that these variables be cointegrated. Secondly, in the appendix I show that whether the 7 theoretical cointegrating vectors are imposed, or whether the vectors are estimated freely, the dynamics of the system are not significantly affected. Therefore, imposing the 7 cointegrating vectors that were rejected using likelihood ratio tests, is not likely to distort the results.

2.3.4.1 Specification of the cointegrating relationships

The theoretical cointegrating relationships, between the bid and ask prices, between the three rates themselves and between the level of each exchange rate and cumulative order flows, are imposed on the model of (2.2). (2.2) also includes a time trend, κt, but this is constrained so as to only enter the cointegration space. Allowing a time trend to enter unrestrictedly in the VAR in first differences implies quadratic trends in the levels of the variables, which, on purely intuitive grounds, is infeasible. Since a time trend in the spread is also unintuitive, the coefficients on the time trend are only allowed to be non-zero in the three cointegrating relationships between the level of the exchange rate and cumulative order flow. Therefore κ is restricted to equal $\alpha\gamma_0$, where α is the 12×7 matrix of speed of adjustment

$$
\kappa t + \alpha\gamma' Y_{t-1} = \underbrace{\alpha}_{12\times 7} \begin{bmatrix} \underbrace{\gamma_0}_{7\times 1} & \underbrace{\gamma'}_{7\times 12} \end{bmatrix} \underbrace{\begin{bmatrix} t \\ Y_{t-1} \end{bmatrix}}_{13\times 1}
$$

$$
= \alpha \begin{bmatrix}
0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\
\gamma_{time}^{ED} & 1 & 0 & 0 & 0 & 0 & 0 & -\gamma_{OF}^{ED} & \gamma_{OF}^{ED} & 0 & 0 & 0 & 0 \\
\gamma_{time}^{ES} & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -\gamma_{OF}^{ES} & \gamma_{OF}^{ES} & 0 & 0 \\
\gamma_{time}^{SD} & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -\gamma_{OF}^{SD} & \gamma_{OF}^{SD}
\end{bmatrix} \begin{bmatrix} t \\ Y_{t-1} \end{bmatrix}
$$

$$
= \alpha \begin{bmatrix}
ED_{t-1}^{ask} - ED_{t-1}^{bid} \\
ES_{t-1}^{ask} - ES_{t-1}^{bid} \\
SD_{t-1}^{ask} - SD_{t-1}^{bid} \\
ED_{t-1}^{ask} - ES_{t-1}^{bid} - SD_{t-1}^{bid} \\
\gamma_{time}^{ED} t + ED_{t-1}^{ask} - \gamma_{OF}^{ED} \left(\sum_{r=0}^{t-1} B_r^{ED} - S_r^{ED} \right) \\
\gamma_{time}^{ES} t + ES_{t-1}^{ask} - \gamma_{OF}^{ES} \left(\sum_{r=0}^{t-1} B_r^{ES} - S_r^{ES} \right) \\
\gamma_{time}^{SD} t + SD_{t-1}^{ask} - \gamma_{OF}^{SD} \left(\sum_{r=0}^{t-1} B_r^{SD} - S_r^{SD} \right)
\end{bmatrix} \tag{2.7}
$$

coefficients in (2.2) and γ_0 is 7×1 but with the first four elements restricted to equal 0. The cointegrating relationships are given in (2.7).

Further structure was imposed on (2.2) by taking into consideration the intra-day patterns of spreads. If the bid and ask prices in the USD/EUR market are cointegrated then only if $ED_{t-1}^{ask} - ED_{t-1}^{bid} > \psi_{t-1}^{EDspread}$ will the ask price tend to fall and the bid price rise in period t, where $\psi_{t-1}^{EDspread}$ is the intra-day pattern of the spread at time $t-1$ calculated from the FFF. For example, around 0600 a reasonably large spread at time $t-1$ will induce a *widening* of the spread at time t so that it moves towards a value one would expect at that time of day. For this reason, the intra-day patterns of spreads are included in the cointegrating vectors. $\gamma_0 t + \gamma' Y_{t-1}$ then becomes:

$$\gamma_0 t + \gamma' Y_{t-1} = \begin{bmatrix} ED^{ask}_{t-1} - ED^{bid}_{t-1} - \psi^{EDspread}_{t-1} \\ ES^{ask}_{t-1} - ES^{bid}_{t-1} - \psi^{ESspread}_{t-1} \\ SD^{ask}_{t-1} - SD^{bid}_{t-1} - \psi^{SDspread}_{t-1} \\ ED^{ask}_{t-1} - ES^{bid}_{t-1} - SD^{bid}_{t-1} \\ \gamma^{ED}_{time} t + ED^{ask}_{t-1} - \gamma^{ED}_{OF} \left(\sum_{r=0}^{t-1} B^{ED}_r - S^{ED}_r \right) \\ \gamma^{ES}_{time} t + ES^{ask}_{t-1} - \gamma^{ES}_{OF} \left(\sum_{r=0}^{t-1} B^{ES}_r - S^{ES}_r \right) \\ \gamma^{SD}_{time} t + SD^{ask}_{t-1} - \gamma^{SD}_{OF} \left(\sum_{r=0}^{t-1} B^{SD}_r - S^{SD}_r \right) \end{bmatrix} \tag{2.8}$$

where $\psi^{EDspread}_{t-1}$, $\psi^{ESspread}_{t-1}$ and $\psi^{SDspread}_{t-1}$ are the intra-day patterns of spreads constructed from the Fourier Flexible Forms in Section 2.3.3. One would then expect the coefficient on $ED^{ask}_{t-1} - ED^{bid}_{t-1} - \psi^{EDspread}_{t-1}$ to be negative in the ΔED^{ask}_t equation and positive in the ΔED^{bid}_t equation; a larger than normal spread causes the ask price to fall and the bid to rise in the subsequent period, and similarly for the other two markets. A larger than normal spread in one market may also have effects on buys and sells in subsequent periods. Since the spread can be interpreted as the cost of trading, a large spread at time $t-1$ may cause the number of buys and sells in the market to reduce at time t. We may then expect the coefficient on $ED^{ask}_{t-1} - ED^{bid}_{t-1} - \psi^{EDspread}_{t-1}$ to be negative in the B^{ED}_t and S^{ED}_t equations. Further analysis of the dynamics in the triangular system can be performed quite easily and hence justifies the use of a rather large dynamic system with 12 endogenous variables.

2.3.5 Good and bad news

The model in (2.2) allows returns and trades (buys and sells) to be affected by news both pre and post release. The leads of the dummy variables, D^R_t, allow us to see what happens to bid and ask returns and buys and sells in the minutes leading up to announcements of scheduled news. If, as has been documented in Payne (1996), volatility declines leading up to the release of data, one may expect the coefficients on the leads of D^R_t to be negative in both buy and sell equations as traders pull out of the market. If one expects spreads to increase in the minutes leading up to a data release, as traders pull out of the market and wait to see what the informational content of the announcement is, then one would expect the coefficients on the leads of D^R_t to be positive in the ask return equation and negative in the bid return equation, since this will cause the spread to rise without any change in the mid-quote. However, to analyse the effects on spreads more accurately, one can transform the model of (2.2) by making the rotation used in Engle and Patton (2004). See Section 2.4.5.

Post release, we can examine the effects of good and bad news separately. One would expect good (bad) UK news to cause the returns of the ask and bid prices to increase (decrease) in USD/GBP for example. However, the effects on buys and sells are not as straightforward. Under the theoretical models of Varian (1989), Harris and Raviv (1993) and Kandel and Pearson (1995) for example, an announcement of either good or bad news will cause both buying *and* selling pressure to increase as market participants disagree over the mapping of information to price. The empirical results of Evans and Lyons (2003) suggest that good (bad) news induces positive (negative) order flow for the corresponding currency. See also the results of Chapter 3. Hence, an announcement of good UK news may cause an increase in both buys and sells of sterling in the USD/GBP market but more buys of sterling than sells. When examining the effects of news on order flow in Chapter 3, I find that order flow is only affected for, at most, a couple of minutes on average. However, the effects on buys and sells could possibly still persist for some time afterwards due to possible inventory rebalancing effects or further information effects associated with heightened FX volatility post release. Model (2.2) therefore allows us to examine such effects.

2.3.6 Estimation

The VAR model presented in (2.2) is estimated in two steps. The first step estimates the cointegrating vectors between the level of each exchange rate and cumulative order flow, and does so using the algorithm described in Johansen (1995) and discussed in more detail in the appendix, Section 2.A.4. Once the cointegrating vectors in (2.8) are determined, they are imposed on the model and estimation of the remaining parameters takes place. Since (2.2) is a recursively ordered structural VAR, it could be estimated equation by equation using OLS. However, it is highly likely that the error in the ΔED_t^{ask} equation is correlated with the error, not only in the ΔED_t^{bid} equation, but also in the other return equations due to the obvious cointegration between rates. If we restrict the errors in each return equation to be uncorrelated with the errors in each trade equation, then the variance/covariance matrix of the 12×1 error vector, ϵ_t, denoted by Ω, can be written as

$$Var(\epsilon_t) = \Omega = \begin{bmatrix} \Omega_{11} & 0 \\ 0 & \Omega_{22} \end{bmatrix} \tag{2.9}$$

where Ω_{11} and Ω_{22} are 6×6 matrices of variance/covariance terms, which can be freely estimated. The 36 restrictions used to estimate the 6×6 β matrix of structural parameters come from the null off-diagonal blocks in Ω. Due to the cross correlation of the errors, we have to estimate (2.2) as a system and because of the recursive ordering of the structural VAR (trades do not depend on contemporaneous returns) it can be estimated as a seemingly unrelated regression (SUR) system.

$$\underbrace{\begin{bmatrix} \Delta ED^{ask} \\ \vdots \\ S^{SD} \end{bmatrix}}_{\substack{\Delta Y \\ 12T \times 1}} = \underbrace{\begin{bmatrix} I_6 \otimes Z_1 & 0 \\ 0 & I_6 \otimes Z_2 \end{bmatrix}}_{\substack{Z \\ 12T \times (6g_1 + 6g_2)}} \underbrace{\begin{bmatrix} \pi_1 \\ \vdots \\ \pi_{12} \end{bmatrix}}_{\substack{\Pi \\ (6g_1+6g_2) \times 1}} + \underbrace{\begin{bmatrix} \epsilon^{ask}_{\Delta ED} \\ \vdots \\ \epsilon^{SD}_{S} \end{bmatrix}}_{\substack{\epsilon \\ 12T \times 1}} \tag{2.10}$$

$$\Rightarrow \Delta Y = Z\Pi + \epsilon$$

ΔED^{ask} is the $T \times 1$ vector of $\log(USD/EUR : ask)$ returns, etc. Z_2 is the $T \times g_2$ matrix of regressors for each trade equation (including constant, error correction terms, news and lagged returns and flows). Z_1 is the matrix of regressors for each return equation which is therefore the same as Z_2 but includes contemporaneous buys and sells in each market. This does not pose a problem since none of the trades (buys or sells) are correlated with the errors in the return equations due to the block diagonality of the variance/covariance matrix of the residuals. Creel and Farell (1996) allow the variance/covariance matrix of ϵ, denoted Σ_ϵ, to be decomposed into a part specified parametrically, P_ϵ, and a part which is not modelled, Ξ_ϵ, i.e. $\Sigma_\epsilon = P_\epsilon + \Xi_\epsilon$. If $\hat{P}_\epsilon$ is an estimator of P_ϵ based on OLS residuals obtained by estimating (2.2) equation by equation, then the proposed estimator of Π, denoted $\hat{\Pi}$ is given by

$$\hat{\Pi} = \left(Z'\hat{P}_\epsilon^{-1}Z\right)^{-1} Z'\hat{P}_\epsilon^{-1}\Delta Y \tag{2.11}$$

where, in this exercise, $\hat{P}_\epsilon = \left(\hat{\Omega} \otimes I_T\right)$, implying $\hat{P}_\epsilon^{-1} = \left(\hat{\Omega}^{-1} \otimes I_T\right)$. Creel and Farell (1996) show that the distribution of $\hat{\Pi}$ is given by

$$\sqrt{T}\left(\hat{\Pi} - \Pi\right) \xrightarrow{d} N(0, V) \tag{2.12}$$

where the estimated variance of $\hat{\Pi}$ is given by

$$V\hat{a}r\left(\hat{\Pi}\right) = \underbrace{\left(Z'\hat{P}_\epsilon^{-1}Z\right)^{-1}}_{A}\underbrace{\left(Z'\hat{P}_\epsilon^{-1}\left(\hat{P}_\epsilon + \hat{\Xi}_\epsilon\right)\hat{P}_\epsilon^{-1}Z\right)}_{B}\underbrace{\left(Z'\hat{P}_\epsilon^{-1}Z\right)^{-1}}_{A} \tag{2.13}$$

B is simply the estimated variance of the $(6g_1 + 6g_2) \times 1$ vector, $Z'\hat{P}_\epsilon\epsilon$, so by using the Newey-West correction on this vector, we can obtain HAC standard errors for the estimator $\hat{\Pi}$. Therefore I correct for heteroscedasticity both *across* and *within* equations. If there were no heteroscedasticity within each equation then Ξ_ϵ would simply be the 12T $\times$ 12T null matrix and the variance of $\hat{\Pi}$ would collapse to the GLS case associated with the standard SUR estimator.

2.4 Estimation Results

Table 2.6 gives a summary of the VAR estimations and for convenience, only the parameter estimates are given for the constant, the error correction terms, contemporaneous trades and news.[32] A number of results are noteworthy and are given below.

2.4.1 Error correction terms

In a large number of cases, the coefficients on the error correction terms have the expected signs. A large spread at time $t-1$, relative to what one would expect in the market at that time, causes the ask price to fall and the bid price to rise in the subsequent minute, as one would expect if the spread is mean reverting. This is the case for each of the three FX markets with the negative effect of large spreads on ask returns and positive effects on bid returns being significant at the 1% level at least. Not only this, a wide

[32]The lag length of returns and trades were chosen using the Schwartz information criterion and were found to be 8 and 5 respectively. Two leads of both UK and US news dummies were included despite the Schwartz criterion suggesting that none should be. They were included simply to allow us to examine the effects of news pre release. One lag of US good and bad news were included in all return equations and two lags of good UK and one lag of bad UK news were included in the return equations. Six lags of both UK and US good and bad news were included in all trade equations.

spread in market k at time $t-1$ tends to cause the spreads in the other two markets to *widen* at time t, via an increase in the ask and a decline in the bid.[33] If the large spread in one market is the result of asymmetric information effects, then information spill-overs may induce dealers in the other markets to widen their bid and ask prices. If there are a preponderance of informed traders in one market, resulting in a large spread, dealers of other currency pairs may expect a large number of informed traders in their markets, hence causing them to widen their spreads. However, despite the coefficients on these cross market error correction terms having the correct sign associated with the above information spill-over hypothesis, only 4 out of the 12 cases are significant at the 5% level. Large spreads at time $t-1$ also tend to reduce trading activity at time t in that market. As explained in Section 2.3.4.1, if the spread represents the cost of trading, then a large spread (relative to what one would expect at that time) will cause both buys and sells to fall. These effects on buys and sells are significant at the 1% level or more in each of the three markets.

However, it is also the case that a large spread in one market at date $t-1$ causes trading at date t to fall in other markets. This is especially true for the case of GBP/EUR spreads. Large spreads in this market cause trading activity (buyer *and* seller initiated trades) to fall in the USD/EUR and USD/GBP markets. Large USD/GBP spreads also cause USD/EUR trading activity to be reduced. Again, this is indicative of the information spill-over effects from one market to another; high spreads in one market at time $t-1$ cause spreads in other markets to widen at date t, which in turn discourage trades.

The cointegration terms associated with absence of arbitrage between the three rates also have significant effects on returns and trading activity. If $ED_{t-1}^{ask} - ES_{t-1}^{bid} - SD_{t-1}^{bid}$ is large and positive, then one would expect ΔED_t^{ask} to be negative and ΔES_t^{bid} and ΔSD_t^{bid} to be both positive, in order for the three rates to be brought back in line. Although this is true for ΔED_t^{ask} (-ve, 10% level) and ΔSD_t^{bid} (+ve, 1% level), the sign on ΔES_t^{bid} is negative and significant at the 1% level. The absence of arbitrage cointegration term also has significant effects on trading activity. A positive value of $ED_{t-1}^{ask} - ES_{t-1}^{bid} - SD_{t-1}^{bid}$ causes buyer and seller initiated trades in all three markets to increase and four out of these six coefficients are positive at the 1% level (five are positive at the 5% level). Theoretical reasons why such trading results from this absence of arbitrage term are discussed in more detail in Chapter 3 and are generally associated with traders trying to exploit arbitrage

[33]The exception to this is the effect of USD/GBP spreads on GBP/EUR spreads. However, the effects that USD/GBP spreads at time $t-1$ have on GBP/EUR ask and bid prices at time t are small and only statistically significant in the bid return equation (10% level).

Table 2.6
Estimation of Multivariate VAR Model in Returns and Trades

	USD/EUR		GBP/EUR		USD/GBP	
	ask	bid	ask	bid	ask	bid
c	1.187	0.907	1.929[b]	-1.022	1.122[b]	0.820[c]
$EDspr_{t-1}$	-0.109[a]	0.0873[a]	0.00451	-0.00337	0.00573[a]	-0.00195
$ESspr_{t-1}$	0.0309[c]	-0.0317[b]	-0.147[a]	0.233[a]	0.00654	-0.0108[a]
$SDspr_{t-1}$	0.0452[b]	-0.0243	-0.00780	0.0264[c]	-0.204[a]	0.152[a]
COA_{t-1}	-0.00564[c]	0.00110	0.00280	-0.0124[a]	-0.00141	0.00587[a]
$ED_{OF,t-1}$	0.00505	-0.00113	-0.00308	0.0124[a]	0.00132	-0.00558[a]
$ES_{OF,t-1}$	-0.00501	0.00125	0.00306	-0.0122[a]	-0.00116	0.00551[a]
$SD_{OF,t-1}$	-0.00559[c]	0.00100	0.00274	-0.0124[a]	-0.00150	0.00576[a]
B_t^{ED}	0.348[a]	0.289[a]	0.197[a]	0.168[a]	0.0965[a]	0.0821[a]
S_t^{ED}	-0.278[a]	-0.342[a]	-0.184[a]	-0.210[a]	-0.0717[a]	-0.0854[a]
B_t^{ES}	0.173[a]	0.151[a]	0.286[a]	0.250[a]	-0.0578[a]	-0.0715[a]
S_t^{ES}	-0.173[a]	-0.176[a]	-0.283[a]	-0.329[a]	0.0863[a]	0.0714[a]
B_t^{SD}	0.131[a]	0.119[a]	-0.0688[a]	-0.0965[a]	0.245[a]	0.230[a]
S_t^{SD}	-0.120[a]	-0.124[a]	0.121[a]	0.105[a]	-0.265[a]	-0.270[a]
US news						
$N_{US,t}^{Good}$	-11.864[a]	-11.437[a]	-5.494[b]	-9.359[b]	-3.326[a]	-4.769[a]
$N_{US,t-1}^{Good}$	-5.801[b]	-7.402[a]	-4.854[a]	-5.217[a]	-1.663	-1.833
$N_{US,t}^{Bad}$	2.053[c]	1.931[c]	2.548[a]	2.586[a]	-0.252	0.855
$N_{US,t-1}^{Bad}$	1.295	2.155	2.381[b]	1.970[c]	0.465	1.107
$D_{US,t+1}$	0.631	0.777	0.781	1.468[b]	-0.896[c]	-0.197
$D_{US,t+2}$	1.927	0.143	1.728[b]	-0.0655	1.367[b]	0.978
UK news						
$N_{UK,t}^{Good}$	0.155	0.652	0.185	0.159	1.945[b]	1.476[c]
$N_{UK,t-1}^{Good}$	0.174	-0.199	-1.372	-0.965	1.559[b]	1.889[b]
$N_{UK,t-2}^{Good}$	-0.555[c]	-0.541	-1.948[b]	-1.496	0.860	0.589
$N_{UK,t}^{Bad}$	0.0926	0.379	2.768[b]	1.531[c]	-0.839[c]	-0.684
$N_{UK,t-1}^{Bad}$	-0.0464	-0.0102	2.285	1.139	0.826	-0.572
$D_{UK,t+1}$	-0.271	-0.126	0.579	-0.0914	0.186	0.342
$D_{UK,t+2}$	-0.140	-0.0882	1.252	-0.405	1.092[c]	0.387
R^2	0.236	0.233	0.211	0.217	0.257	0.243

Continued over

Table 2.6
Estimation of Multivariate VAR Model in Returns and Trades (cont.)

	USD/EUR		GBP/EUR		USD/GBP	
	buys	sells	buys	sells	buys	sells
c	-0.693	-1.267[b]	3.228[a]	2.610[a]	3.095[a]	1.914[a]
$EDspr_{t-1}$	-0.00444[a]	-0.00355[a]	0.0^3904	-0.0^3575	0.0^3954	0.0^3610
$ESspr_{t-1}$	-0.0125[a]	-0.0124[a]	-0.0161[a]	-0.0204[a]	-0.00375	-0.00859[a]
$SDspr_{t-1}$	-0.0193[a]	-0.0207[a]	0.00729	0.00279	-0.0427[a]	-0.0541[a]
COA_{t-1}	0.0131[a]	0.0107[a]	0.00630[a]	0.00200	0.00408[b]	0.00664[a]
$ED_{OF,t-1}$	-0.0131[a]	-0.0107[a]	-0.00655[a]	-0.00233	-0.00432[b]	-0.00700[a]
$ES_{OF,t-1}$	0.0130[a]	0.0107[a]	0.00635[a]	0.00219	0.00418[b]	0.00676[a]
$SD_{OF,t-1}$	0.0132[a]	0.0108[a]	0.00624[a]	0.00196	0.00400[b]	0.00669[a]
US news						
$N_{US,t}^{Good}$	3.184[a]	12.289[a]	0.250	3.752[b]	0.905	1.356
$N_{US,t-1}^{Good}$	-0.593[c]	9.214[a]	-0.0581	2.719[a]	-0.863[b]	4.055[a]
$N_{US,t-2}^{Good}$	3.854[c]	-0.331	3.669[a]	-1.080	3.547[c]	1.249
$N_{US,t-3}^{Good}$	3.270	-0.666	2.458[a]	0.462	0.501	-0.141
$N_{US,t-4}^{Good}$	1.195	-0.0240	0.312	-1.515	-0.317	-1.305[c]
$N_{US,t-5}^{Good}$	0.731	1.753	-1.309[b]	-0.220	-0.103	0.710
$N_{US,t-6}^{Good}$	-2.354[a]	-0.524	-0.433	-0.318	-1.167	1.447
$N_{US,t}^{Bad}$	5.781[a]	3.607[a]	3.695[a]	0.283	2.110[a]	0.613
$N_{US,t-1}^{Bad}$	4.368[b]	1.842[c]	0.776	-0.324	1.090[b]	0.864
$N_{US,t-2}^{Bad}$	0.751	2.274[b]	1.128	0.921	1.062	1.469[b]
$N_{US,t-3}^{Bad}$	0.422	0.888	-0.114	0.253	-0.351	1.136
$N_{US,t-4}^{Bad}$	0.748	-0.100	0.222	0.126	-0.202	-0.235
$N_{US,t-5}^{Bad}$	1.614	0.629	1.110	-0.479	0.819	0.635
$N_{US,t-6}^{Bad}$	1.939[c]	-0.327	0.729	-0.501	1.539[a]	0.0575
$D_{US,t+1}$	-0.134	-0.739[a]	-0.241	-0.419	0.00253	-0.317
$D_{US,t+2}$	0.372	-0.327	0.603	-0.487	-0.0684	-0.314

Continued over

Table 2.6
Estimation of Multivariate VAR Model in Returns and Trades (cont.)

	USD/EUR		GBP/EUR		USD/GBP	
	buys	sells	buys	sells	buys	sells
UK news						
$N^{Good}_{UK,t}$	-0.369	-1.528^{a}	0.0180	0.701	4.956^{b}	1.551
$N^{Good}_{UK,t-1}$	-0.766^{c}	-0.262	1.514	2.112	2.462^{a}	1.264^{b}
$N^{Good}_{UK,t-2}$	-0.125	1.886	1.268	2.246^{c}	1.727	0.836
$N^{Good}_{UK,t-3}$	-1.335^{a}	-0.661	0.433	-1.041^{c}	-1.657^{a}	-0.403
$N^{Good}_{UK,t-4}$	0.0711	0.376	-0.162	-1.295^{a}	-1.140^{b}	0.170
$N^{Good}_{UK,t-5}$	-0.186	0.257	-0.482	-1.140^{b}	-0.993^{a}	0.649
$N^{Good}_{UK,t-6}$	0.0327	-0.871^{a}	0.541	-0.182	-0.416	-0.389
$N^{Bad}_{UK,t}$	0.762	-0.212	1.400	-0.334	1.780^{b}	4.311^{b}
$N^{Bad}_{UK,t-1}$	0.0772	2.135	2.229^{a}	0.783	1.075	5.244^{a}
$N^{Bad}_{UK,t-2}$	0.443	-1.560^{a}	2.872^{c}	2.542^{b}	1.099	2.480
$N^{Bad}_{UK,t-3}$	1.491	0.236	2.025	-0.0778	1.429^{c}	2.902
$N^{Bad}_{UK,t-4}$	0.970^{c}	-1.795^{b}	-0.159	1.120	0.866	-1.125
$N^{Bad}_{UK,t-5}$	-0.0430	-0.208	0.281	-0.00507	1.320	1.017
$N^{Bad}_{UK,t-6}$	-1.832^{a}	-0.880^{c}	0.479	-0.146	0.477	0.624
$D_{UK,t+1}$	-0.809^{a}	0.591	-0.118	0.431	-0.124	-0.00903
$D_{UK,t+2}$	-0.0515	1.018^{c}	0.0969	0.313	0.517	-0.0939
R^2	0.214	0.195	0.275	0.253	0.319	0.288

Notes: The data cover the eight month period from 1st December 1999 to 24th July 2000. The USD/EUR exchange rate is defined as the number of dollars (numerator currency) per euro (denominator currency). Buys (sells) refer to trades where the initiator is purchasing (selling) the donominator/commodity currency; the euro in USD/EUR and GBP/EUR and sterling in USD/GBP. All returns are defined as 10000× the log first difference of the rate and ^{a}, ^{b}, ^{c} denote significance at the 1, 5 and 10% levels respectively. $EDspr_{t-1}$ is the equilibrium error associated with the USD/EUR spread at date $t-1$, defined in (2.8), i.e. taking into account the intra-day pattern of the spread during the trading day, and similarly for $ESspr_{t-1}$ and $SDspr_{t-1}$. COA_{t-1} is the equilibrium error associated with the cost of arbitrage between the three rates. $ED_{OF,t-1}$ is the equilibrium error associated with the difference between the (log) level of the USD/EUR exchange rate and cumulative order flow, and similarly $ES_{OF,t-1}$ and $SD_{OF,t-1}$ for the GBP/EUR and USD/GBP rates. Coefficients are to be interpreted as follows; a one basis point increase in the USD/EUR spread (relative to normal) in minute $t-1$ causes the $USD/EUR : ask$ rate to fall by 0.109 basis points in minute t and $USD/EUR : buys$ to fall by 0.00444 (relative to normal). An extra buy in USD/EUR (relative to normal) causes the $USD/EUR : ask$ price to rise by 0.348 basis points in that minute. A one standard deviation announcement of good US news in minute t causes a fall in the $USD/EUR : ask$ price of 11.864 basis points and for $USD/EUR : buys$ to increase by 3.184 (relative to normal) in that minute.

opportunities.[34]

The results for the cointegration terms associated with the level of the exchange rate and cumulative order flow are not as easy to interpret, however. A positive value of $\gamma_{time}^{ED}t + ED_{t-1}^{ask} - \gamma_{OF}^{ED}\left(\sum_{r=0}^{t-1} B_r^{ED} - S_r^{ED}\right)$ should result in ΔED_t^{ask} being negative and should also lead to positive USD/EUR order flow at date t. As one can see from Table 2.6, this is not the case. The lack of cointegration between these variables reported in Section 2.A.4 and the poor results reported here are puzzling in light of the results from Killeen, Lyons, and Moore (2002) and Bjønnes and Rime (2003).[35]

2.4.2 Asymmetric effects of buys and sells

As documented in Engle and Patton (2004) for stocks, and predicted by the dynamic limit order book models of Parlour (1998) and Foucault (1999), there appear to be strong asymmetric effects of buyer and seller initiated market orders. Indeed, as all order flow models, such as Kyle (1985), Glosten and Milgrom (1985), etc., would predict, buyer initiated trades cause both ask and bid prices to increase and seller initiated trades cause both ask and bid prices to fall. However, in all cases, a buy order in one market has a greater effect on the ask price than on the bid. In the case of the USD/EUR market, for example, a buy order (of euros) causes the ask price to increase by 0.35 basis points and the bid price to increase by only 0.29 basis points. A market sell order (of euros) causes the ask price to fall by 0.28 basis points and the bid price to fall by 0.34 basis points. Since market buy orders are executed at the ask price, such trades drain liquidity on that side of the book and hence have a larger impact on the ask than on the bid. Wald tests show that these asymmetric effects of buys and sells are statistically significant. Buys (sells) have a greater impact on the ask (bid) in both USD/EUR and GBP/EUR markets (1% level), while only buyer initiated trades have asymmetric effects in USD/GBP (10% level).

The asymmetric effects of buys and sells can also be examined by impulse response analysis. If a market buy order increases the ask price more than the bid, then the spread will necessarily widen. This is shown in Figure 2.3, which gives the IRFs following or-

[34]In Chapter 3 I model these effects using a basic threshold cointegration model, which is more appropriate than the model presented above. Only if $ED_{t-1}^{ask} - ES_{t-1}^{bid} - SD_{t-1}^{bid}$ was greater than the cost of exploiting the arbitrage opportunity, will the rates be brought back in line.

[35]However, as reported in the appendix, these cointegrating relationships are not found to be important when analysing the high frequency dynamics that are of interest in this chapter.

Figure 2.3
Impulse Response Functions of Spreads Following 'Buy' Shocks in each Market

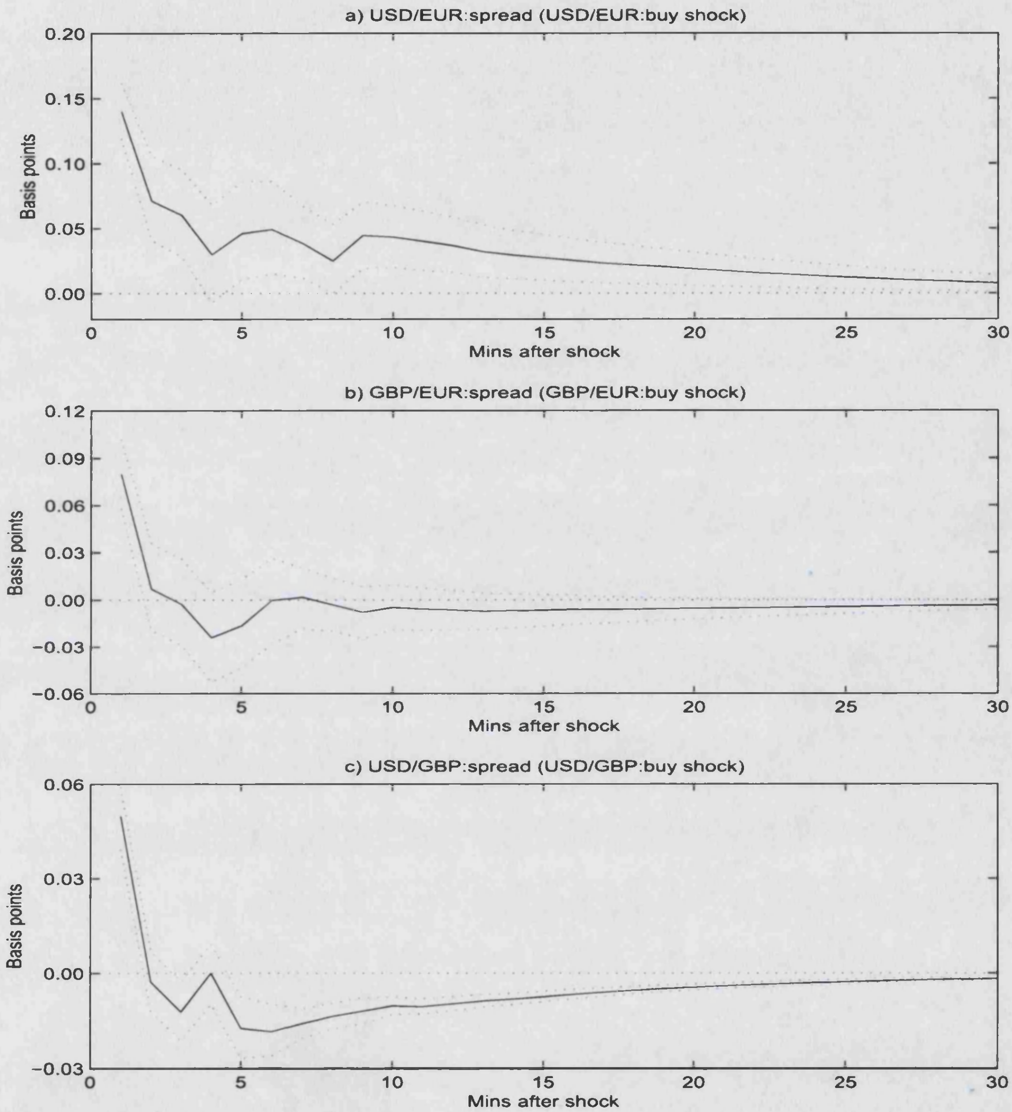

[redacted]

thogonalised, one standard deviation buy shocks in each market.[36] After a shock to $USD/EUR : buys$, for example, the USD/EUR spread increases, but despite being significant in a statistical sense, the widening of the spread is small in economic terms. Note that the figures show the effects on spreads over and above what one would expect. In all three cases, the spread reverts to zero, implying a buy shock has no long run effect on the spread, as one would expect after imposing the cointegrating vectors in (2.8).

The cross market effects of order flow, and buys and sells in particular, are also present. Evans and Lyons (2002a) and Daníelsson, Luo, and Payne (2002) document the effects that order flow in one market have on exchange rates in another.[37] A simple explanation for such effects runs as follows. A market buy order (of euros) in the USD/EUR market can be interpreted as the result of an informed trader receiving favourable news about the euro or having negative news about the dollar.[38] Good news for the euro should cause the dealer to revalue the euro against all currencies, including sterling. Bad news for the dollar, on the other hand, should cause the dealer to devalue the dollar against all currencies, including sterling. Since the dealer does not know whether the initial customer order of euro purchases is based on (private) good euro information or (private) bad dollar information, the dealer responds by revaluing the euro against sterling slightly and devaluing the dollar against sterling as well. Hence a buy order of euros in USD/EUR will cause a positive return in GBP/EUR at both the bid and ask, and a positive return in USD/GBP, again in both the bid and ask series. This logic can be applied to explain every cross market effect observed in Table 2.6.[39]

The spillover effects can be shown more clearly by way of impulse response functions and variance decompositions. Figure 2.4 shows the IRFs following a one standard deviation, orthogonalised buy shock in the USD/EUR market, i.e. purchase of euro. As expected, this causes a permanent increase in the dollar price of the euro, of approximately 1 basis point in both the bid and ask series. However, this euro buy shock causes the euro to appreciate by 0.7 basis points against sterling and also leads to an appreciation of sterling

[36]This orthogonalisation procedure and the impulse response analysis is explained in more detail in the appendix, Section 2.A.5.

[37]Although the Daníelsson, Luo, and Payne (2002) results were obtained using the same dataset as that used here.

[38]Such private information could come from a dealer's customer order as in Lyons (1995) or Evans and Lyons (2002b).

[39]How likely this story is is another matter. In these large markets, liquidity supplying dealers tend to trade only one currency. In which case, a market buy order in one currency pair is not likely to have much effect on prices in another market. However, on the Reuters D2000-2 trading platform, traders can view up to 6 currencies at one time and so even if they do not supply liquidity in more than one market, they can still observe the other market order flows and make inferences from this on the valuation of their own currencies.

Figure 2.4
Impulse Response Functions of Exchange Rates Following a USD/EUR 'Buy' Shock

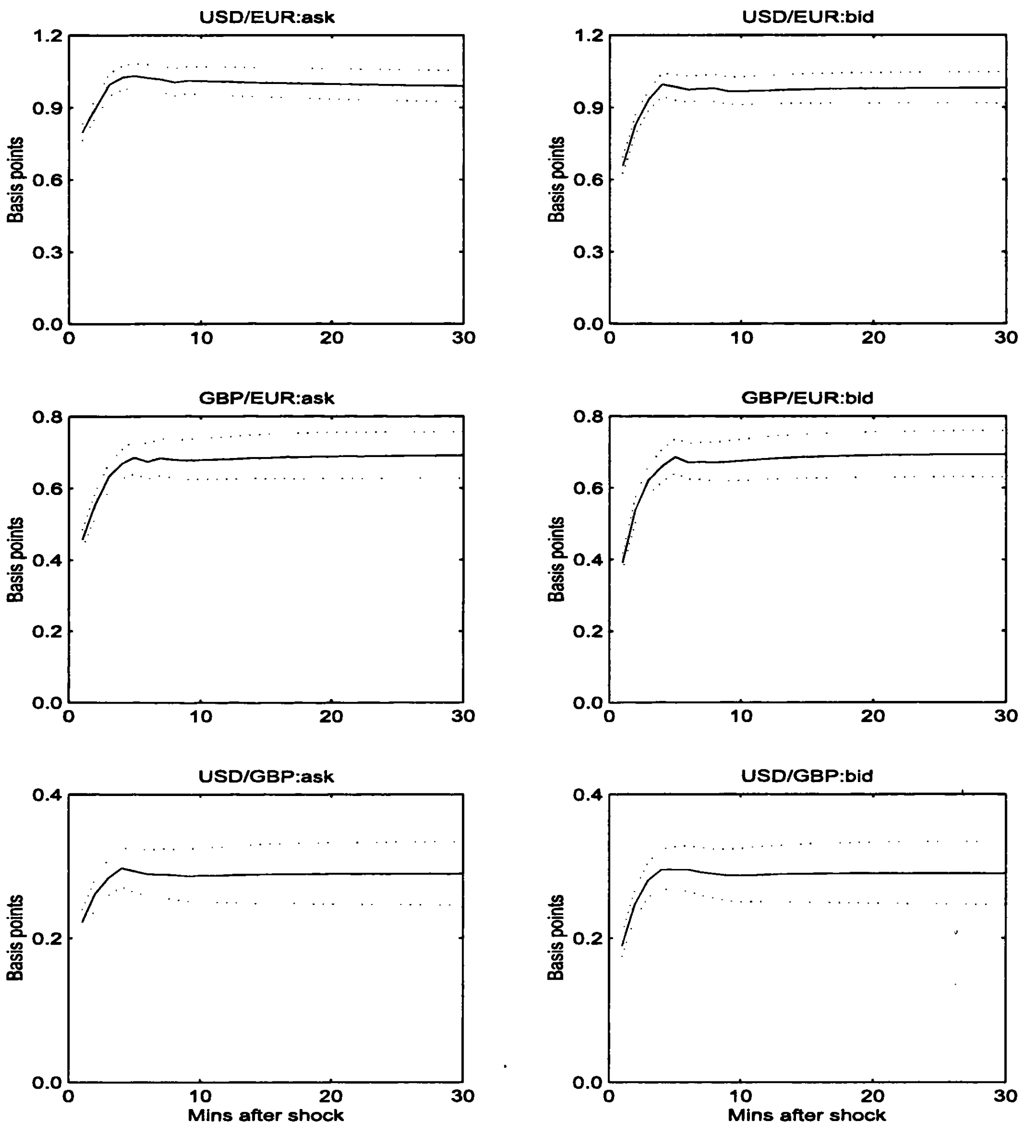

Notes: The figures plot the impulse response functions for exchange rates following an orthogonalised USD/EUR 'buy' shock. For each figure the solid black line shows the response of the exchange rate, while the dotted blue lines bound the 95% confidence interval found by bootstrapping over 1000 iterations.

against the dollar of approximately 0.3 basis points. Similar cross market effects are also present following trade shocks in the other two markets, and are shown in the appendix, Section 2.A.5.

Variance decompositions, showing the share of the mean square error of n step ahead forecasts attributable to trading in each of the three markets is shown in Figure 2.5. Figure 2.5a shows the variance decomposition for the $USD/EUR : ask$ return series. At the 60 step ahead forecast, just under a quarter of the MSE is attributed to USD/EUR trading (both buys and sells). However, approximately 11% and 7% of the variance of the USD/EUR ask rate is attributed to trading in GBP/EUR and USD/GBP respectively, implying significant spillovers from one market to another. It is always the case that the share of the variance of exchange rate k attributed to trading in market k is greater than that attributed to trading in market l ($\neq k$), but the cross market effects of trading, as defined by these variance decompositions, are significant (especially from USD/EUR trading). No matter what the causes of such cross market effects are, the information spillovers from one market to another appear considerable.

2.4.3 The effects of news on rates

The main focus of this chapter is to analyse the effects of scheduled macroeconomic news on FX prices and transaction activity. As can be seen from Table 2.6 both good and bad news from the US and the UK have significant effects on exchange rate returns (bid and ask) and on buying and selling pressure. For example, following a one standard deviation announcement of good US news, this causes the ask price in the USD/EUR market to fall by 11.86 basis points and the bid price to fall by 11.44 basis points after 1 minute.[40] However, such news also has a lagged effect; the ask and bid prices in the USD/EUR market fall by a further 5.80 and 7.40 basis points respectively in the following minute. The effects of good US news on USD/GBP are also significant in the minute post release. A one standard deviation announcement of good news causes the dollar to appreciate by 3.33 and 4.77 basis points in the ask and bid prices respectively. The lagged effects of this news are also present (a further 1.66 and 1.83 basis point appreciation) but they are

[40]The effects on exchange rate returns are to be expected since the news variables were signed so that good news caused an appreciation of that country's money. Since USD/EUR and USD/GBP are defined as the number of dollars per euro (or sterling) then good US news causes a negative return in both markets. Good UK news, on the other hand, causes a negative return in GBP/EUR (pounds per euro) and a positive return in USD/GBP (dollars per pound). Also, since news affects order flows in the minute of release and this order flow has an impact on prices, the total effect of news will be greater than these reported values. The total effects are shown in the impulse response functions documented in Section 2.4.6.

Figure 2.5
Variance Decompositions, Showing the Share of the Mean Square Error of n Step Ahead Forecasts Attributable to Trading in each of the Three Markets

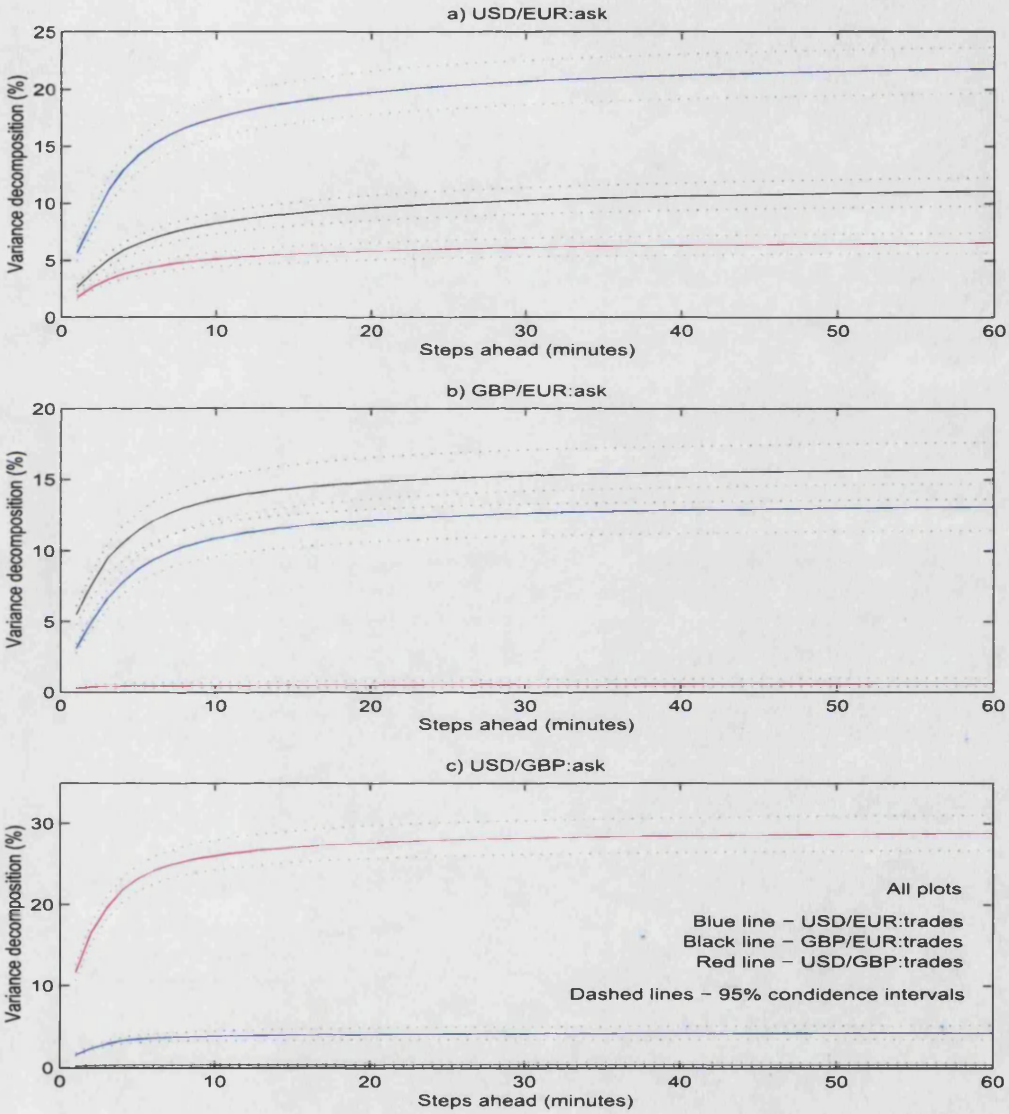

Notes: The figures plot the share of the mean square error of n step ahead forecasts attributable to trading in each of the three markets. The solid blue line shows the share attributed to USD/EUR trading (both buys and sells), the black line shows that attributed to GBP/EUR trading and the red line represents USD/GBP trading. Dotted lines bound the 95% confidence intervals, found by bootstrapping over 1000 iterations.

not statistically significant. The effects of bad US news are also present but they do not appear to be as pronounced as good news. A one standard deviation announcement of bad US news causes the dollar to depreciate against the euro (positive euro return) by 2.05 and 1.93 basis points in the ask and bid prices respectively. The effects of such news on USD/GBP are, however, quantitatively small and insignificant.

The effects of UK news are also significant. A one standard deviation announcement of good UK news, although not having a significant effect in the minute of release in the GBP/EUR market, causes an appreciation of sterling of 1.37 and 0.96 basis points one minute later in the ask and bid prices respectively and causes a further appreciation of 1.95 and 1.50 basis points in the following minute. The effects of good UK news in the USD/GBP market are also significant and persist for two further minutes post release. The cumulative direct effect of a one standard deviation announcement of good UK news in USD/GBP is a sterling appreciation of 4.36 and 3.95 basis points in the ask and bid prices. The effects of bad UK news only appear to be significant in the GBP/EUR market, where a one standard deviation announcement of such news causes a cumulative direct depreciation of sterling of 5.05 (ask) and 2.67 (bid) basis points. The effects of bad news emanating from the UK appear to be small and insignificant in USD/GBP.[41]

2.4.4 The effects of news on trades

Of equal importance are the effects of news on buys and sells. As mentioned earlier, if good US news causes positive dollar order flow then we would expect to see such news lead to increased dollar buying, i.e. an increase in the number of (euro) sells in USD/EUR. If good US news also results in an increase in dollar sales (euro purchases) then this would be consistent with the idea that market participants disagree on the interpretation of the news, i.e. the mapping from information to price is not common knowledge. As can be seen in Table 2.6, a one standard deviation announcement of good US news causes an increase of 3.18 euro purchases (B^{ED}) over and above what one would expect at that time of day.[42] The same news also causes an increase of 12.29 euro sales (dollar purchases) over and above what one would expect. These increases are significant at the 1% level and are consistent with the idea that traders disagree over the mapping of information

[41]The cross market effects are also present for US news. Good US news causes a cumulative direct appreciation of sterling in the GBP/EUR market of 10.35 (14.58) basis points in the ask (bid) series. These cross market effects are discussed in more detail in Chapter 3 and also show the importance of the US economy in the determination of 'non-US-related' asset prices.

[42]Remember that B^{ED} is deseasonalised euro buys, the number of buys minus the intra-day pattern calculated from the Fourier Flexible Form.

to price. Note also that good US news causes negative USD/EUR order flow, i.e. more euro sells (12.29) than euro buys (3.18). To put these numbers into perspective, on non-announcement days, the normal level of buys and sells at this time of day (0830 ET, 1330 GMT/BST) is approximately 3 per minute (see Figure 2.1a). Therefore a one standard deviation announcement of good US news causes euro buys to double and euro sells to increase by over 400% compared to non-announcement times.

A similar story holds when considering bad US news. A one standard deviation release of such news causes an increase of 5.78 euro buys (dollar sells) and also of 3.61 euro sells (dollar purchases). Bad US news therefore causes net euro buying (positive order flow) but an increase in both dollar buying *and* selling. The results are similar when examining the effects of good and bad US news on USD/GBP and also when considering the effects of good and bad UK news. However, in a number of cases, the increase in buying or selling pressure that is induced often does not occur immediately but happens in the subsequent minutes. This is consistent with the idea that traders wait to see how others interpret the news before deciding how to react. In which case, the initial order flow that results from the data release, which may be quantitatively small, becomes information itself. This, together with the initial price change, is used by other market participants and triggers more trades, buys *and* sells, depending on how *this* information is interpreted.

The cumulative direct effect of news on buys and sells is presented in Figure 2.6. This sums the direct effect of good US news on USD/EUR buys and sells, for example, for minutes t to $t+2$ and also for minutes $t+3$ to $t+6$. Good US news causes an increase of 3.18 euro purchases (in the USD/EUR market) in the minute of release, a decline of 0.59 euro buys in the subsequent minute and an increase of 3.85 buys, compared to normal, two minutes after release. The cumulative direct effect is then an increase of 6.45 euro purchases, shown in the black bar of Figure 2.6a. In each of the four panels, there are four bars for each of the three exchange rates. The two darker bars (black and blue) show the effects of the news on buyer initiated trades, while the two lighter bars (white and yellow) show the effects on sells. The black and white bars show the cumulative effects on buys and sells respectively, from minutes t to $t+2$ and the blue and yellow bars show the sum of trades from minutes $t+3$ to $t+6$. This then allows us to see what happens to trading in the first 3 minutes post release and also how the news affects trading in the subsequent 4 minutes. Significance of the cumulative effects were determined using F-tests. In the case of good US news, the increase in buys and sells, in minutes t to $t+2$, in each of the three FX markets is significant at the 5% level at least. In five of the six cases, the increase

Figure 2.6
Effects of Good and Bad US and UK News Releases on Buyer and Seller Initiated Trades in the USD/EUR, GBP/EUR and USD/GBP Markets

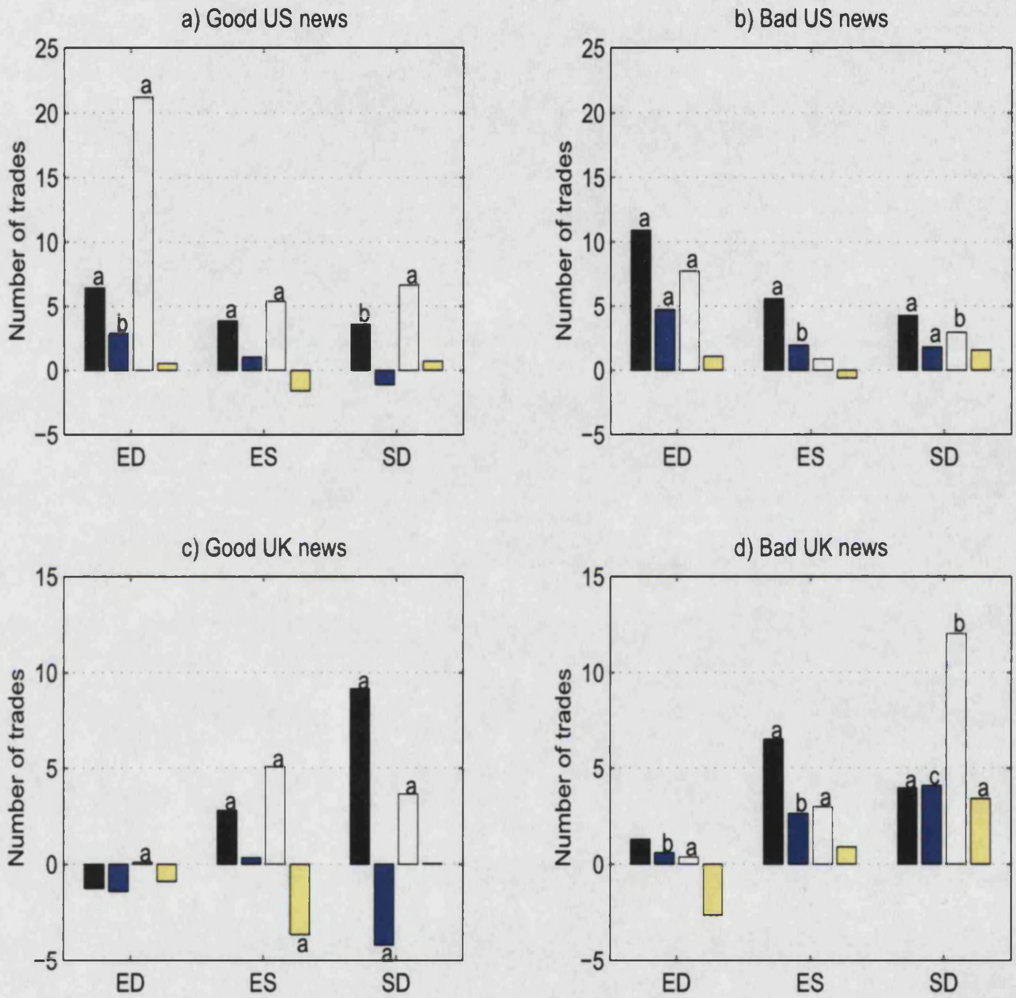

Notes: The figures show the effects of good and bad news releases from the US and UK on buyer (black and blue bars) and seller (white and yellow bars) initiated trades. The black (white) bar shows the cumulative effect on buys (sells) from minutes t to $t+2$ and the blue (yellow) bar shows the sum of buys (sells) from minutes $t+3$ to $t+6$. Buyer (seller) initiated trades are purchases (sales) of the commodity currency; the euro in USD/EUR and GBP/EUR and sterling in USD/GBP. ED is USD/EUR, ES is GBP/EUR and SD is USD/GBP. a, b and c denote significance at the 1, 5 and 10% levels respectively. The number of trades shown in the figures are the cumulated number of buys (sells) relative to what one would expect in the market at that time (from the FFF analysis). For example, in the three minutes following an announcement of good US news, there are 6.45 more euro buys (black bar) compared to normal/non-announcement times, and there are 2.84 more euro buys (blue bar) from minutes $t+3$ to $t+6$. The same news generates 21.17 more euro sells (white bar) in USD/EUR compared to normal and an increase of 0.54 euro sells (yellow bar) from minutes $t+3$ to $t+6$.

in buys or sells is significant at the 1% level. The fact that such news has significant effects on trading in GBP/EUR reflects the importance of US data figures in the world economy. However, the effect this news has on trading appears to be short lived. Trading in minutes $t+3$ to $t+6$ is only significantly increased for $USD/EUR: buys$ (5% level). Bad US news has similar effects on FX trading activity although the cross market effect on cumulative GBP/EUR sells is not significant. The effect of bad US news on trading in minutes $t+3$ to $t+6$ is significant, at least for buyer initiated orders, although in every case the effect is smaller than the immediate effects of news on trades. Both good and bad UK news cause significant increases in GBP/EUR and USD/GBP buying and selling pressure, but the effect on USD/EUR trading appears to be small. These smaller cross market effects of UK news on USD/EUR trading, compared to the effects of US news on GBP/EUR activity are to be expected considering the smaller, and less important, UK economy. Note also that good and bad UK news does not have significant effects on USD/EUR exchange rates, as documented in Table 2.6.

The effects of quantified news on buys, sells and exchange rate returns post release are quite clear. However, there only appears to be a limited systematic effect on trading activity leading up to the scheduled data releases. If traders pull out of the market and wait to see the extent of the information contained in the news release, consistent with the finding that FX volatility declines pre announcement, (Payne 1996), then one would expect the coefficients on D^R_{t+1} and D^R_{t+2} to be significantly negative in the trade equations, i.e. if there is an announcement due in the next minute or in two minutes time, the number of buyer or seller initiated transactions should be significantly reduced. In the case of the US, although five of the coefficients on D^{US}_{t+1} are negative, only that in the S^{ED} equation is significant. In the minute prior to a release of UK news, buys and sells in all markets are not significantly affected, with the exception of the B^{ED} equation, where trades appear to fall in the minute pre release. However, little should be inferred from this, considering the limited impact UK news has on USD/EUR rates and trades post release. Neither US nor UK news have significant effects on any of the rates pre release, as one may expect, consistent with there being no leakages of news. However, if traders pull out of the market pre release, one may expect the spread to widen, in which case one would expect to see a positive return in the ask price equation and a negative return in the bid price equation in the minutes leading up to the announcement. It is these effects on spreads, and also trading volume, to which I now turn.

2.4.5 The effects of news on spreads

Model (2.2) looks explicitly at the effects of news announcements on exchange rate returns and buys and sells separately. Following Engle and Patton (2004) we can transform the model of (2.2) into one that considers spreads, mid-quote returns, order flows and transaction volumes by making an appropriate rotation. Premultiplying (2.2) by the 12×12 matrix G in (2.14) will give a model in the first differences of spreads and mid-quotes, as well as order flows and volumes. This can be transformed easily into a model which examines the *level* of the spread.

$$\underbrace{\begin{bmatrix} I_3 \otimes \begin{bmatrix} 1 & -1 \\ 0.5 & 0.5 \end{bmatrix} & 0_{6\times 6} \\ 0_{6\times 6} & I_3 \otimes \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} \end{bmatrix}}_{G} \underbrace{\begin{bmatrix} \Delta ED_t^{ask} \\ \Delta ED_t^{bid} \\ \vdots \\ B_t^{SD} \\ S_t^{SD} \end{bmatrix}}_{\Delta Y_t} = \begin{bmatrix} \Delta SP_t^{ED} \\ \Delta MQ_t^{ED} \\ \vdots \\ OF_t^{SD} \\ Vol_t^{SD} \end{bmatrix} \tag{2.14}$$

Where SP_t^{ED} is the level of the spread in the USD/EUR market at date t, defined as the difference between log ask and bid prices, MQ_t^{ED} is the USD/EUR mid-quote at date t. OF_t^{ED} is the order flow, defined as the difference between the number of buyer and seller initiated trades and Vol_t^{ED} is trading volume in USD/EUR in minute t, defined as the number of buyer plus seller initiated trades in minute t. The effects of good and bad news, as well as the effects of an imminent release of data, from the US and UK on spreads and trading volume are given in Table 2.7. The effects that such news have on order flows and exchange rate returns (transactions prices) are described in detail in Chapter 3. Since we are only interested in the effects of the news, rather than the dynamics of spreads and transactions volumes, only the coefficients and t statistics are given for news effects.

The effects of news on trading volume are as expected. If news causes buys and sells to increase, then it obviously causes trading volume to rise.[43] As can be seen in the table, the increase in trading volume increases for approximately three minutes post release. There is also limited evidence that trading activity falls in the minutes leading up to the announcement. In the case of US news, volume does fall in the minute pre release in all markets but the effects are not statistically significant. In the USD/EUR market, for example, there are normally 2.96 buys (sells) at 1329 on non-announcement days,

[43]However, for correct inference, one needs to take into account the covariance terms from the matrix V in (2.12).

Table 2.7
Effects of News Releases on Spreads and Trading Volume

	USD/EUR		GBP/EUR		USD/GBP	
	spread	t-stat	spread	t-stat	spread	t-stat
US news						
$N_{US,t}^{Good}$	-0.426	(-0.13)	3.864	(0.81)	1.444	(1.27)
$N_{US,t-1}^{Good}$	1.601	(0.52)	0.363	(0.17)	0.170	(0.09)
$N_{US,t}^{Bad}$	0.122	(0.08)	-0.0378	(-0.03)	-1.106	(-0.80)
$N_{US,t-1}^{Bad}$	-0.859	(-0.29)	0.411	(0.27)	-0.642	(-0.48)
$D_{US,t+1}$	-0.145	(-0.10)	-0.687	(-0.66)	-0.699	(-1.00)
$D_{US,t+2}$	1.785	(0.86)	1.793	(1.25)	0.389	(0.46)
UK news						
$N_{UK,t}^{Good}$	-0.496	(-0.44)	0.0257	(0.03)	0.469	(0.40)
$N_{UK,t-1}^{Good}$	0.372	(0.55)	-0.407	(-0.27)	-0.331	(-0.30)
$N_{UK,t-2}^{Good}$	-0.0142	(-0.02)	-0.452	(-0.35)	0.270	(0.26)
$N_{UK,t}^{Bad}$	-0.286	(-0.31)	1.237	(0.85)	-0.155	(-0.21)
$N_{UK,t-1}^{Bad}$	-0.0362	(-0.02)	1.147	(0.59)	1.398	(1.44)
$D_{UK,t+1}$	-0.145	(-0.29)	0.670	(0.93)	-0.156	(-0.27)
$D_{UK,t+2}$	-0.0518	(-0.09)	1.657[c]	(1.66)	0.706	(0.88)

Notes: The spread is defined as the difference between the log ask and bid prices (less the intra-day pattern). Volume is defined as the number of buyer plus the number of seller initiated trades in a minute (less the intra-day pattern). Values and standard deviations were calculated using the appropriate rotation on (2.2); specifically, the G matrix in (2.14). t-stats in parentheses. [a], [b], [c] denote significance at the 1, 5 and 10% levels respectively. Coefficients are to be interpreted as follows; a one standard deviation announcement of good US news in minute t causes a fall in the USD/EUR spread of 0.426 basis points (relative to normal/non-announcement times) and causes an increase of 15.473 trades in the USD/EUR market in minute t.
Continued over

Table 2.7
Effects of News Releases on Spreads and Trading Volume (cont.)

	USD/EUR		GBP/EUR		USD/GBP	
	volume	t-stat	volume	t-stat	volume	t-stat
US news						
$N^{Good}_{US,t}$	15.473[a]	(6.63)	4.003[c]	(1.72)	2.261	(0.97)
$N^{Good}_{US,t-1}$	8.621[a]	(7.78)	2.661[b]	(2.40)	3.192[a]	(2.88)
$N^{Good}_{US,t-2}$	3.523	(1.55)	2.589	(1.14)	4.797[b]	(2.10)
$N^{Good}_{US,t-3}$	2.603	(0.96)	2.919	(1.08)	0.360	(0.13)
$N^{Good}_{US,t-4}$	1.171	(0.41)	-1.203	(-0.42)	-1.622	(-0.57)
$N^{Good}_{US,t-5}$	2.484	(1.26)	-1.529	(-0.78)	0.608	(0.31)
$N^{Good}_{US,t-6}$	-2.878	(-1.33)	-0.751	(-0.35)	0.280	(0.13)
$N^{Bad}_{US,t}$	9.388[a]	(6.21)	3.979[a]	(2.63)	2.722[c]	(1.80)
$N^{Bad}_{US,t-1}$	6.210[a]	(2.58)	0.451	(0.19)	1.954	(0.81)
$N^{Bad}_{US,t-2}$	3.025	(1.58)	2.049	(1.07)	2.531	(1.32)
$N^{Bad}_{US,t-3}$	1.309	(1.20)	0.139	(0.13)	0.785	(0.72)
$N^{Bad}_{US,t-4}$	0.648	(0.43)	0.348	(0.23)	-0.437	(-0.29)
$N^{Bad}_{US,t-5}$	2.242	(1.24)	0.631	(0.35)	1.454	(0.80)
$N^{Bad}_{US,t-6}$	1.612	(1.22)	0.228	(0.17)	1.597	(1.21)
$D_{US,t+1}$	-0.873	(-1.44)	-0.660	(-1.08)	-0.315	(-0.52)
$D_{US,t+2}$	0.0454	(0.08)	0.116	(0.20)	-0.382	(-0.67)
UK news						
$N^{Good}_{UK,t}$	-1.897[b]	(-2.30)	0.719	(0.87)	6.507[a]	(7.90)
$N^{Good}_{UK,t-1}$	-1.028	(-1.53)	3.627[a]	(5.41)	3.726[a]	(5.56)
$N^{Good}_{UK,t-2}$	1.761	(1.22)	3.514[b]	(2.43)	2.563[c]	(1.77)
$N^{Good}_{UK,t-3}$	-1.996[a]	(-3.10)	-0.608	(-0.95)	-2.060[a]	(-3.20)
$N^{Good}_{UK,t-4}$	0.449	(0.80)	-1.457[a]	(-2.60)	-0.970[c]	(-1.73)
$N^{Good}_{UK,t-5}$	0.0709	(0.08)	-1.622[c]	(-1.72)	-0.344	(-0.36)
$N^{Good}_{UK,t-6}$	-0.838	(-1.37)	0.359	(0.59)	-0.805	(-1.31)
$N^{Bad}_{UK,t}$	0.551	(0.40)	1.066	(0.77)	6.091[a]	(4.40)
$N^{Bad}_{UK,t-1}$	2.213	(1.29)	3.011[c]	(1.76)	6.319[a]	(3.69)
$N^{Bad}_{UK,t-2}$	-1.116	(-0.95)	5.413[a]	(4.59)	3.579[a]	(3.03)
$N^{Bad}_{UK,t-3}$	1.727	(1.15)	1.947	(1.29)	4.331[a]	(2.88)
$N^{Bad}_{UK,t-4}$	-0.825	(-0.85)	0.960	(0.99)	-0.259	(-0.27)
$N^{Bad}_{UK,t-5}$	-0.251	(-0.26)	0.276	(0.29)	2.337[b]	(2.47)
$N^{Bad}_{UK,t-6}$	-2.712[a]	(-4.40)	0.334	(0.54)	1.101[c]	(1.79)
$D_{UK,t+1}$	-0.218	(-0.32)	0.313	(0.46)	-0.133	(-0.19)
$D_{UK,t+2}$	0.967	(1.53)	0.410	(0.65)	0.423	(0.67)

implying 5.92 trades. An imminent announcement of US news causes trading volume in this market to fall by 0.87, i.e. a reduction in trading volume of 15% in the minute pre release. An imminent announcement of UK news also tends to cause trading volume to fall in the two dollar markets but causes GBP/EUR trading to rise. Again, none of these figures are statistically significant.

The effects of news announcements on spreads are not strong. Following an information event, one may expect the spread to widen. However, this does not appear to be the case; neither US nor UK news releases cause the spread to change significantly. This may be because the spread is measured as the difference between the best bid and ask in the market; the touch. If each individual dealer widened his/her spread but dealers disagreed on the mapping of information to price, as demonstrated by the fact that good and bad news cause an increase in both buying *and* selling pressure, then the wide spreads of individual dealers are likely to overlap only imperfectly, having an ambiguous effect on the touch. See Chapter 4. The effects of news on spreads pre release are also limited. Two minutes prior to a release of US news sees the USD/EUR spread widening by 1.78 basis points, implying an increase of nearly three quarters compared to non-announcement times. However this is not significant at any reasonable level. Two minutes prior to a release of UK data, spreads widen in the GBP/EUR and USD/GBP markets by 1.66 and 0.71 basis points respectively, the former being significant at the 10% level and represents an increase of nearly two thirds compared to non-announcement times. The 0.71 basis point increase in USD/GBP spreads prior to a UK news release represents an increase of approximately 40%. Market liquidity, if defined by spreads and trading volume, therefore appears to be reduced in the period leading up to scheduled news announcements, but not hugely so. However, in order to assess the impact of news releases on liquidity, one can also examine their effects on volatility. This is considered in Section 2.5.

2.4.6 Impulse response analysis

The effects of news releases on foreign exchange activity are shown in Table 2.6, for returns and buyer and seller initiated trades, and in Table 2.7 for spreads and trading volume. However, these only show the direct effects of news on each of the variables of interest. Due to the dynamic structure and interaction between variables, the total effects of news can be more involved and can be obtained using impulse response analysis. For example, a one standard deviation announcement of good US news causes a negative return (US

Figure 2.7
Exchange Rate and Trade Impulse Response Functions Following Good US News

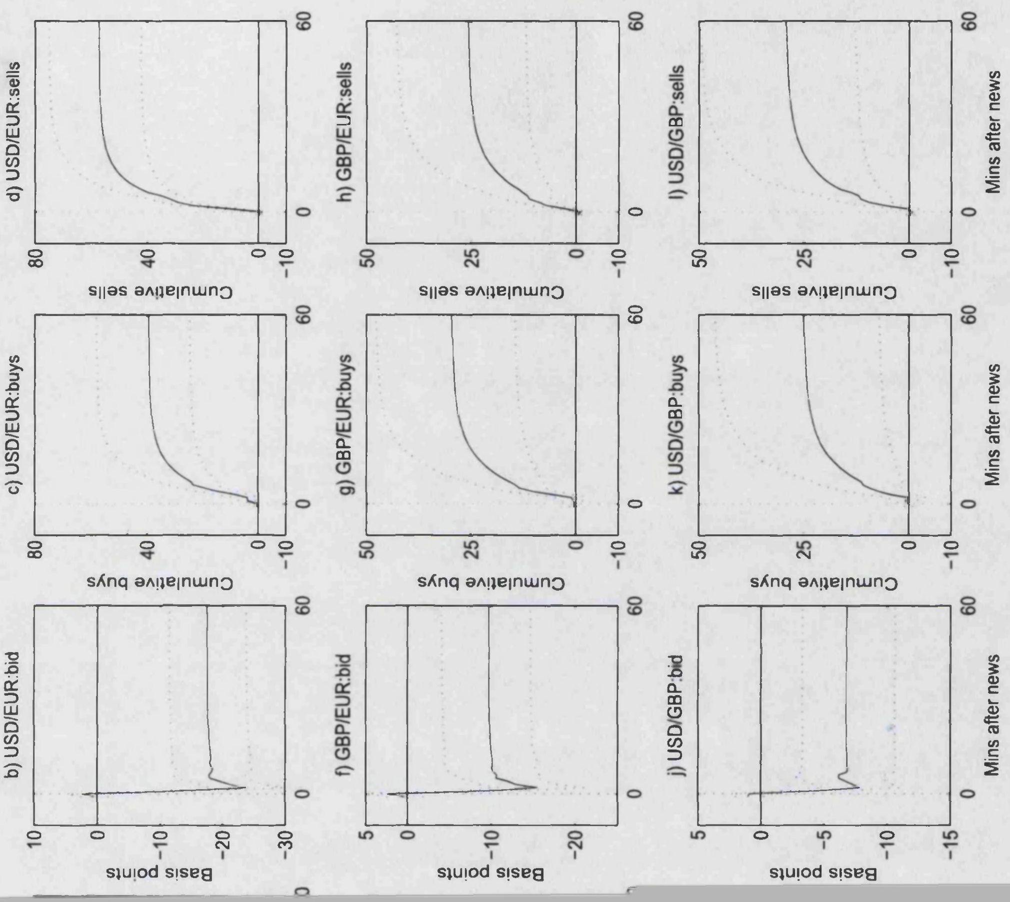

Figure 2.8
Exchange Rate and Trade Impulse Response Functions Following Bad US News

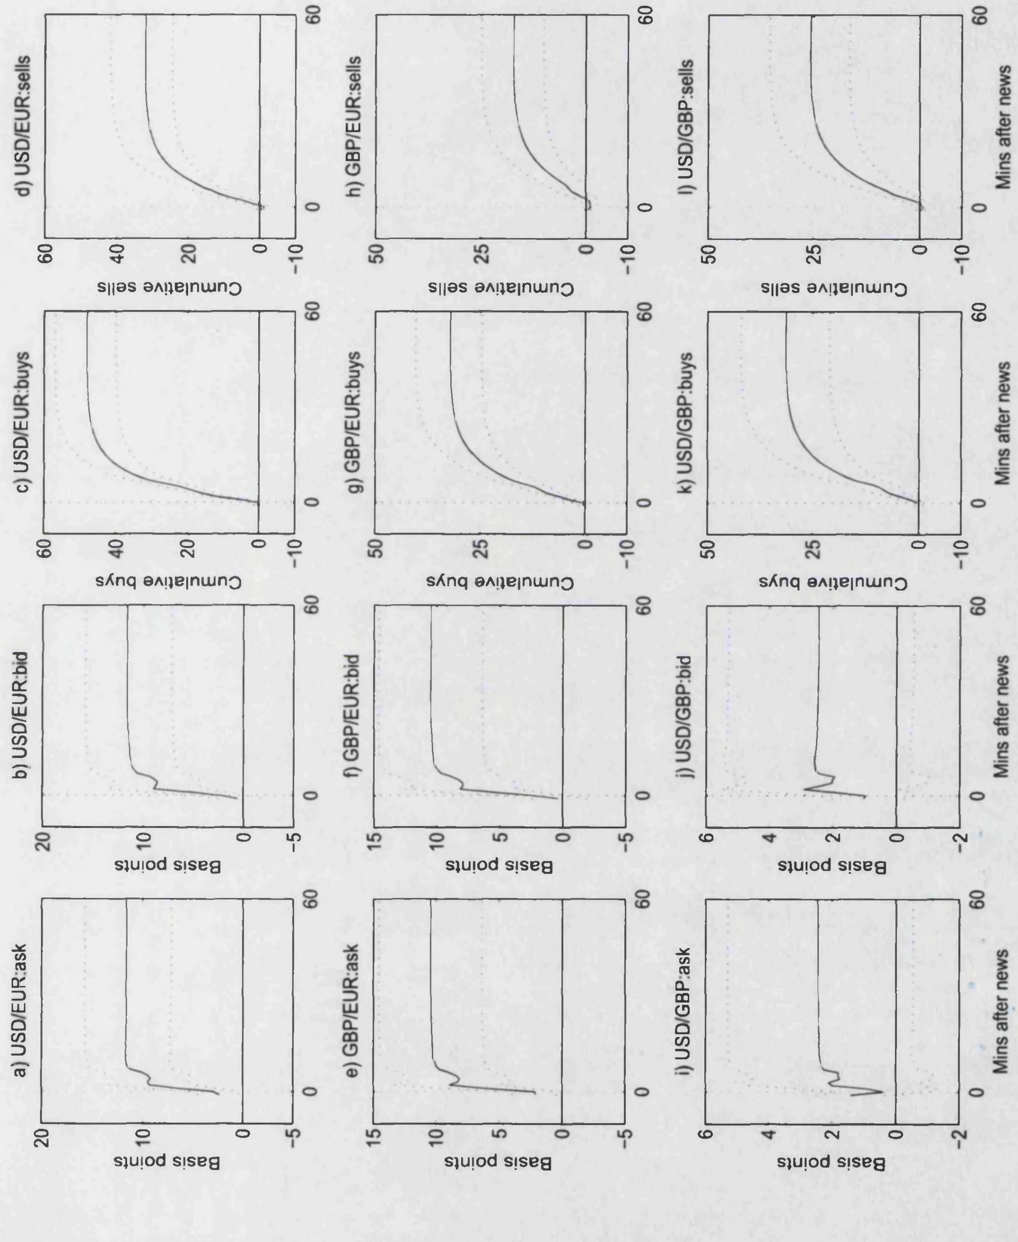

Notes: The figure shows the effects on the bid and ask exchange rates, and the number of buys and sells following a one standard deviation announcement of bad US news. The trade impulse response functions show the effects of the news on cumulative (deseasonalised) buys and sells, i.e. the effect on trades over and above what one would expect. The dotted lines bound the 95% confidence interval, found by bootstrapping model (2.2) over 1000 iterations.

Figure 2.9
Exchange Rate and Trade Impulse Response Functions Following Good UK News

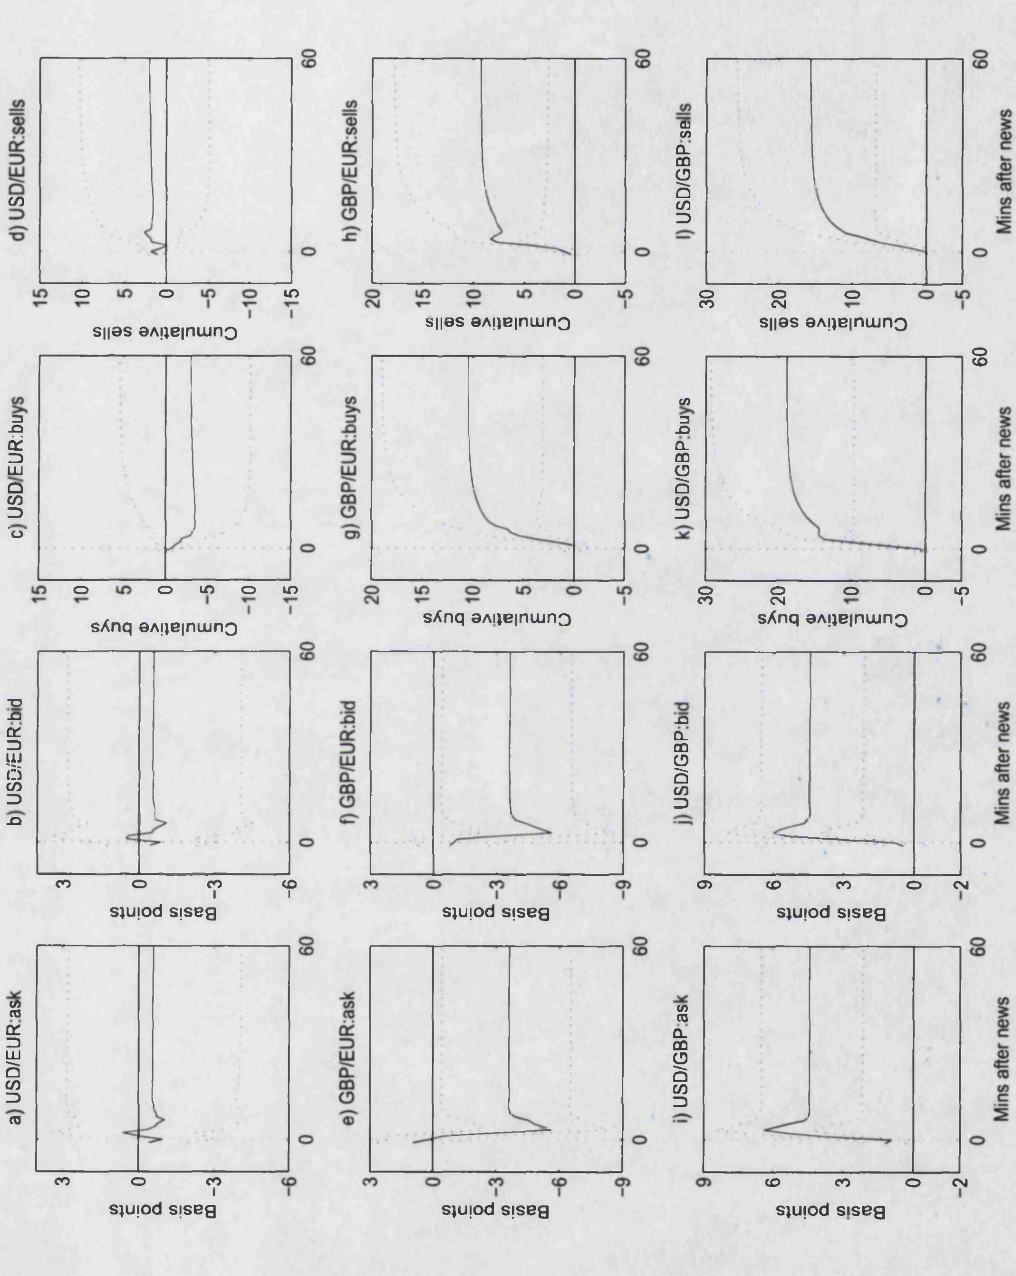

Notes: The figure shows the effects on the bid and ask exchange rates, and the number of buys and sells following a one standard deviation announcement of good UK news. The trade impulse response functions show the effects of the news on cumulative (deseasonalised) buys and sells, i.e. the effect on trades over and above what one would expect. The dotted lines bound the 95% confidence interval, found by bootstrapping model (2.2) over 1000 iterations.

Figure 2.10
Exchange Rate and Trade Impulse Response Functions Following Bad UK News

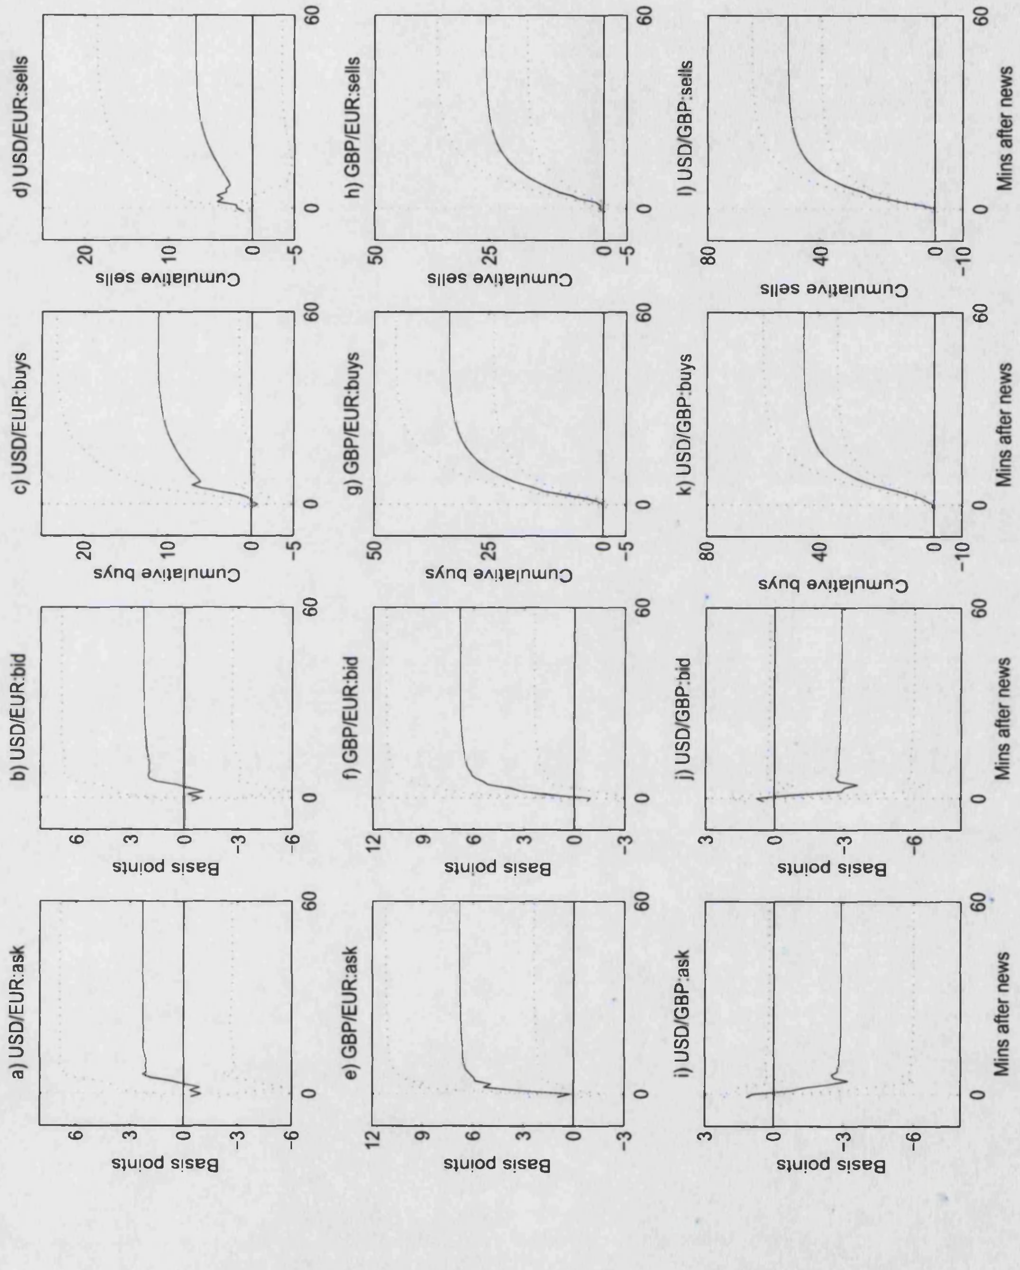

Notes: The figure shows the effects on the bid and ask exchange rates, and the number of buys and sells following a one standard deviation announcement of bad UK news. The trade impulse response functions show the effects of the news on cumulative (deseasonalised) buys and sells, i.e. the effect on trades over and above what one would expect. The dotted lines bound the 95% confidence interval, found by bootstrapping model (2.2) over 1000 iterations.

dollar appreciation) of 11.9 basis points in the USD/EUR : *ask* rate. (See Table 2.6). In the minute post release, such news causes buyer and seller initiated trades in this market to increase by 3.18 and 12.29 trades respectively. This negative order flow will therefore have a further negative effect on the return via the β matrix in (2.2), which shows the contemporaneous effects of trades on returns. The total effect of news on returns (and also on trades and spreads) is therefore likely to be more involved than that suggested by Tables 2.6 and 2.7.

In order to more accurately assess the effects of scheduled news releases on FX activity, I introduce a one standard deviation announcement of each type of news to the model in (2.2) and compute the effects on returns, buys and sells (Figures 2.7 to 2.10) and spreads and trading volume (Figures 2.11 to 2.14). For example, Figure 2.7 shows the effects of a one standard deviation announcement of good US news on each of the 12 variables in Y_t, i.e. it shows the effect on each exchange rate (cumulative return) and on cumulative trades (over and above what one would expect to observe). The panels show the effects from 2 minutes pre announcement to 60 minutes post release. As one would expect, and consistent with efficient markets hypotheses, the effects of news releases on the levels of the exchange rates are almost immediate. Good US news causes the US dollar to appreciate against the euro (negative euro return) by approximately 17 basis points[44] and also causes an appreciation against sterling of approximately 7 basis points. It also causes buying and selling activity in each market to increase significantly, over and above what one would expect, as defined by the intra-day diurnal patterns documented in Section 2.3.3.[45] The negative returns that good US news create in the USD/EUR and USD/GBP markets, are also associated with net selling pressure in these markets, i.e. good US news causes USD/EUR : *sells* to increase by more than USD/EUR : *buys*. These effects on order flow and the process by which news is incorporated into FX rates is discussed in more detail in Chapter 3. In the vast majority of cases, both good and bad news emanating from the UK and the US cause significant exchange rate returns and also causes buying *and* selling pressure to increase in each market. UK news has little effect on USD/EUR trading or FX returns, as one may expect considering the size of the UK economy. However, the cross market effects of US news are significant, as discussed above. US news causes significant buying and selling activity in GBP/EUR and also generates significant returns in this market.

[44]Note that this is substantially greater than the direct 11.9 basis point return reported in Table 2.6.

[45]The dashed lines in Figures 2.7 to 2.14 bound the 95% confidence limits, found by bootstrapping model (2.2) over 1000 iterations.

Figure 2.11
Spread and Volume Impulse Response Functions following Good US News

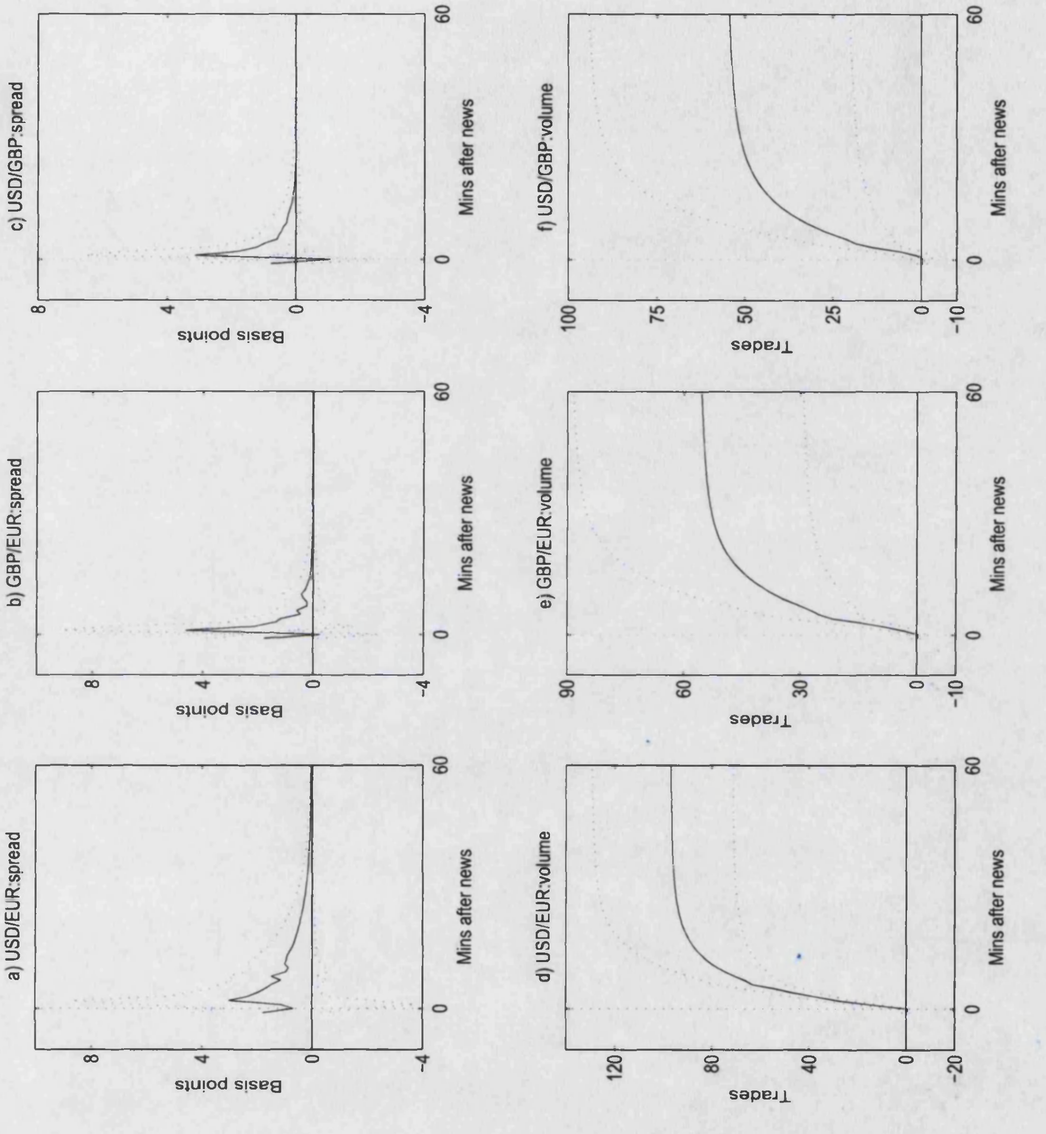

Notes: The figure shows the effects on the spreads, and the number of trades following a one standard deviation announcement of good US news. The spread IRFs show how the spreads change compared to normal; a negative value of the response does not imply a negative spread, rather the spread is smaller than one would expected to observe at that time on non-announcement days. The trade impulse response functions show the effects of the news on cumulative (deseasonalised) trades, i.e. the effect on trades over and above what one would expect. The dotted lines bound the 95% confidence interval, found by bootstrapping model (2.2) over 1000 iterations.

Figure 2.12
Spread and Volume Impulse Response Functions following Bad US News

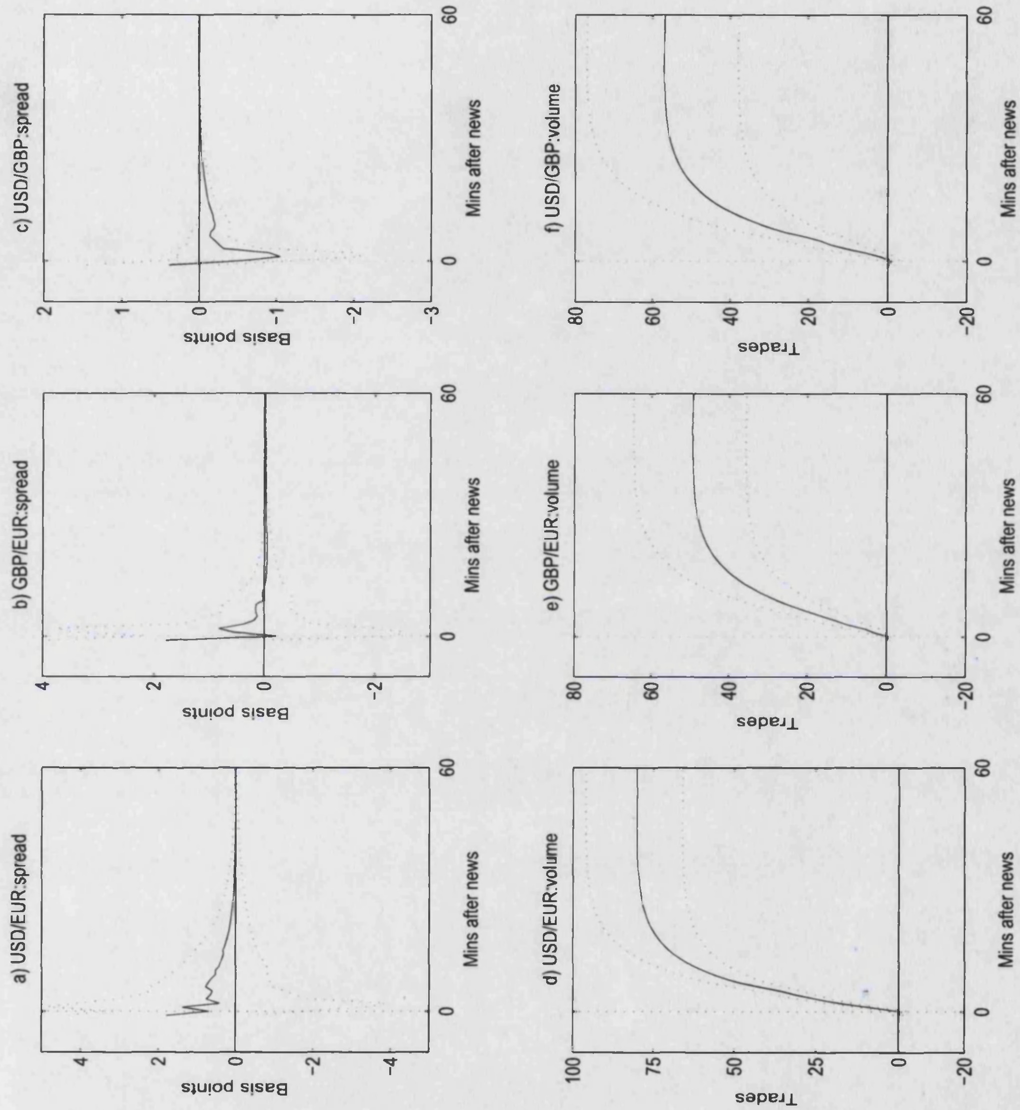

Notes: The figure shows the effects on the spreads, and the number of trades following a one standard deviation announcement of bad US news. The spread IRFs show how the spreads change compared to normal; a negative value of the response does not imply a negative spread, rather the spread is smaller than one would expected to observe at that time on non-announcement days. The trade impulse response functions show the effects of the news on cumulative (deseasonalised) trades, i.e. the effect on trades over and above what one would expect. The dotted lines bound the 95% confidence interval, found by bootstrapping model (2.2) over 1000 iterations.

Figure 2.13
Spread and Volume Impulse Response Functions Following Good UK News

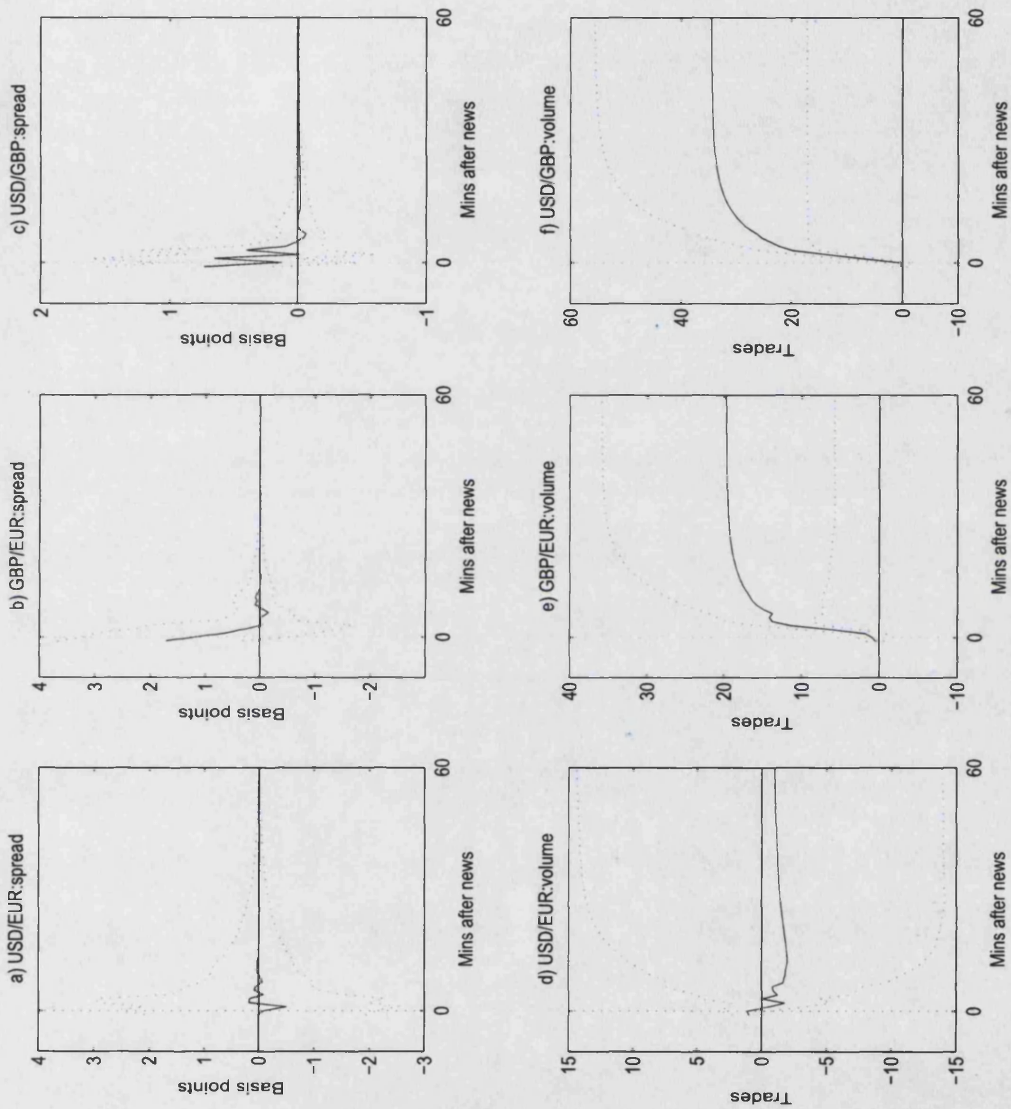

Notes: The figure shows the effects on the spreads, and the number of trades following a one standard deviation announcement of good UK news. The spread IRFs show how the spreads change compared to normal; a negative value of the response does not imply a negative spread, rather the spread is smaller than one would expected to observe at that time on non-announcement days. The trade impulse response functions show the effects of the news on cumulative (deseasonalised) trades, i.e. the effect on trades over and above what one would expect. The dotted lines bound the 95% confidence interval, found by bootstrapping model (2.2) over 1000 iterations.

Figure 2.14
Spread and Volume Impulse Response Functions Following Bad US News

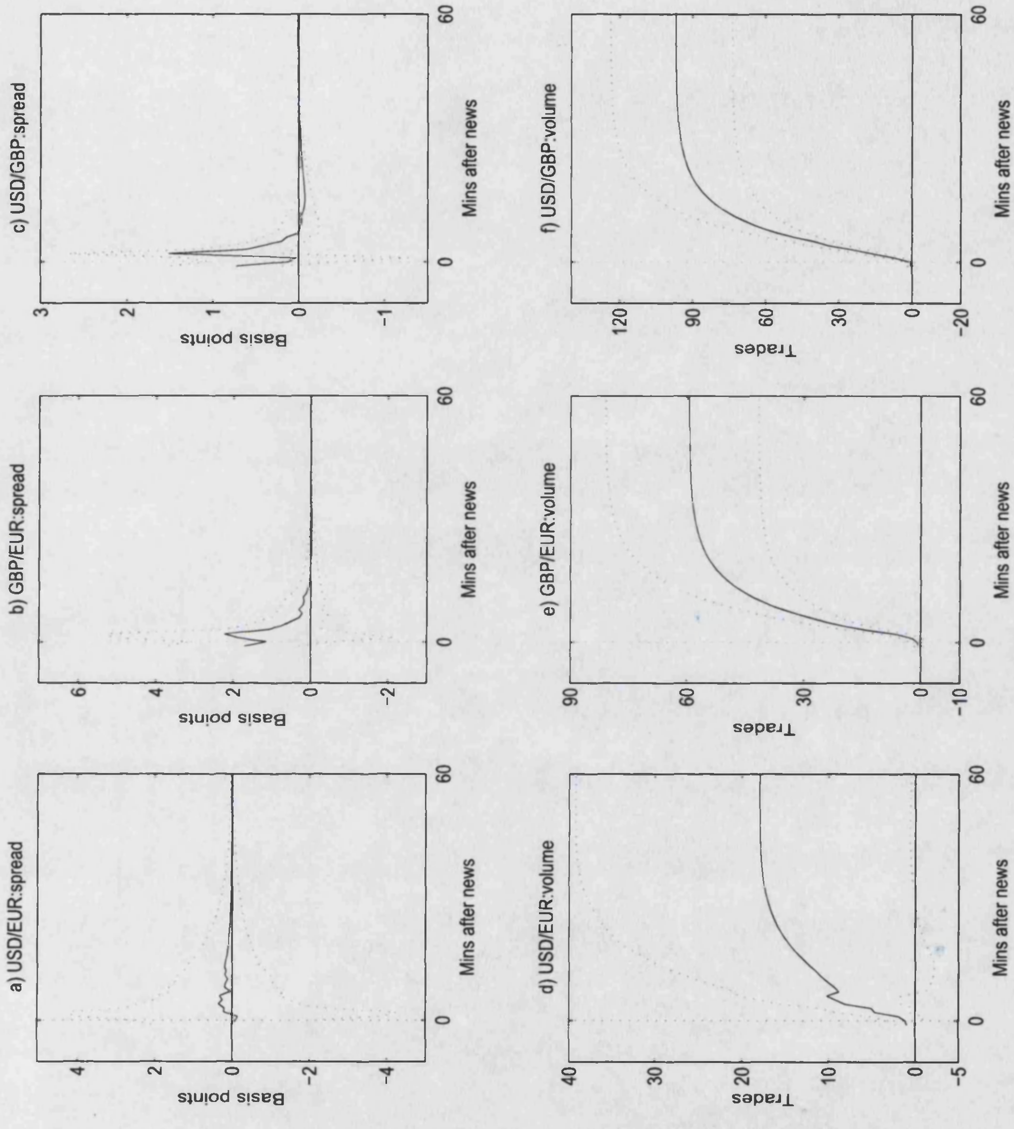

Notes: The figure shows the effects on the spreads, and the number of trades following a one standard deviation announcement of bad UK news. The spread IRFs show how the spreads change compared to normal; a negative value of the response does not imply a negative spread, rather the spread is smaller than one would expected to observe at that time on non-announcement days. The trade impulse response functions show the effects of the news on cumulative (deseasonalised) trades, i.e. the effect on trades over and above what one would expect. The dotted lines bound the 95% confidence interval, found by bootstrapping model (2.2) over 1000 iterations.

Figures 2.11 to 2.14 show how spreads and trading volume change from their expected values following a one standard deviation announcement of each type of news. As discussed in Section 2.4.5, the effects of news on spreads pre and post release tend to be small. If anything, spreads do rise following news announcements, especially following good US news (by approximately 3 basis points) but the effects are quantitatively small and are not statistically significant. Note that the IRFs show how the spreads change compared to normal. A negative value of the response does not imply a negative spread, rather the spread is smaller than one would expect to observe at that time on non-announcement days. In all cases, the spread 'difference' reverts to zero, implying there is no long term effect of news on the spread.

As expected, following the above discussion for the effects of news on buys and sells separately, UK and US news announcements have significant effects on trading volume. However, only bad UK news has an effect on USD/EUR trading activity but this is barely significant at the 5% level.

2.4.7 An illustrative example

To demonstrate the effects of news on FX activity, the effects of an individual announcement are shown in Figure 2.15. This shows the effects of an unanticipated announcement of low US CPI figures on 16th May 2000 in the USD/EUR market. At 1330 BST, an announcement was made that CPI inflation was zero for the month of April, whereas the median forecast was for a month on month increase of 0.1%.[46] This caused an appreciation of the dollar, consistent with our signing of news in Section 2.2 and this is seen in the figure. In the minute following the announcement, both bid and ask rates in USD/EUR fell by 15 and 16 pips (one pip is 0.0001 USD/EUR) respectively, representing a dollar appreciation of approximately 17 basis points. In the minute of the data release, both buying and selling activity increased dramatically. There were 7 purchases of euros (sales of dollars) and 17 sales of euros (purchases of dollars) in the minute of release, compared to 2.98 buys (sells) that were normally seen at 1330 on non-announcement days. See Figure 2.1a. Therefore the CPI announcement caused euro buying to increase by over 130% and dollar buying to increase by 470%. Trading activity remained high for another

[46]To put this difference into perspective, this is equal to a 0.8 standard deviation announcement of news. This is therefore not a large surprise but was deliberately chosen since it had a noticeable impact on FX activity. Also, CPI figures are released a week after PPI data. Analysts are more likely to correctly forecast the impending CPI figures from the PPI data just released, suggesting more 'news' from a 0.1% difference in CPI announcements than from a similar error in PPI forecasts.

Figure 2.15
Effects on USD/EUR Activity Following a Release of US CPI at 1330 BST on 16th May 2000

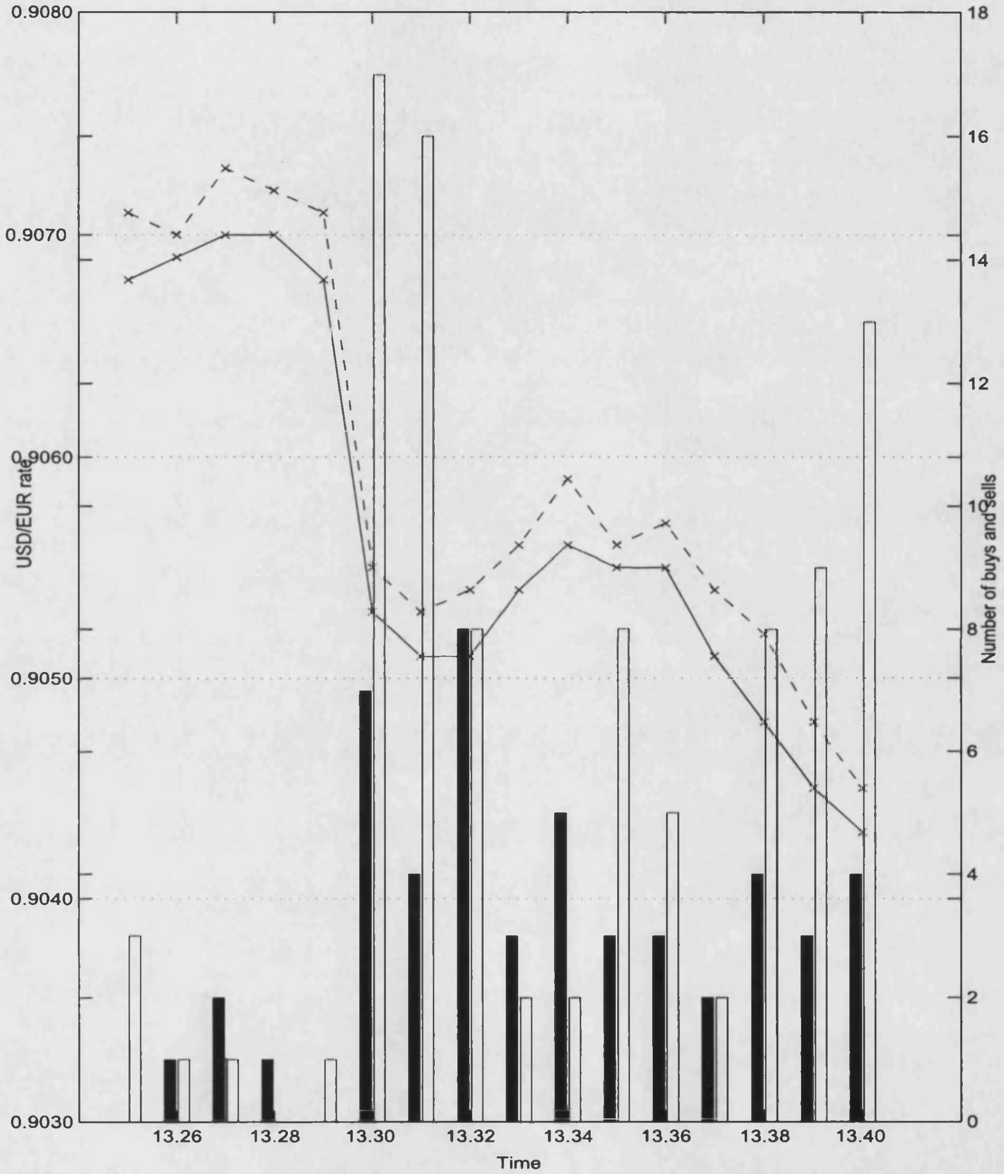

Notes: The figure shows the bid (solid line) and ask (dashed line) at the end of each minute from 1325 to 1340 (left axis). Buyer initiated trades (of euros) are shown as black bars and seller initiated trades (of euros) are shown as white bars (right axis). The announcement of CPI was of a 0% change, month on month, in April. The consensus forecast was of a 0.1% increase.

2 minutes before falling to more normal levels, 3 buys, 3 sells per minute. However, as seen in the figure, euro selling did pick up again 8 minutes after the release. Note that the announcement of good dollar news (that which causes the dollar to appreciate) leads to net dollar buying, evidenced by there being more euro sells than euro buys, following the data release. In the minute leading up to the announcement, the intensity of trading is seen to be very low. In the minute pre release, there was only one trade (a purchase of dollars). This drop off in trading is seen for each of the five minutes pre release; there are, at most, 3 trades per minute (buy or sell), compared to an average of nearly 6 per minute in non-announcement days. The effect of the CPI announcement on spreads, however, is negligible. Pre and post release, the spread never increases above 3 pips (approximately 3.3 basis points), compared to a normal/non-announcement time average of approximately 2.5 basis points. Quite often the spread is only 2 pips or even 1 pip (at 1326 and 1335).

2.5 Second Moment Effect of News

Section 2.4 demonstrates the first moment (level) effects of news on both exchange rates and trading activity. However, as found in a number of studies, including Ederington and Lee (1993), Payne (1996) and Andersen and Bollerslev (1998), such news can have significant and persistent effects on volatility. In order to measure the effects of news on the second moments, we first require a measure of FX volatility. Simply using the absolute value of the exchange rate return or the squared return will bias upwards the effects of news on volatility since doing so will cause the results to incorporate the jump (first moment) effects that such news generates. For this reason, a similar procedure to that used by Andersen, Bollerslev, Diebold, and Vega (2003) will be implemented. This essentially takes the residual from the regression model (2.2) and defines the absolute value of this residual as a proxy for the volatility process. In which case, all the first moment effects and expected changes in returns (and trades) will be purged from the volatility series. The effects of scheduled news releases on volatility can be examined using similar event-study regression techniques as used in (2.2).

Since one can define the absolute value of the residual in the ΔED^{ask} equation as $USD/EUR : ask$ volatility, one can similarly take the absolute value of the residual in the S^{ED} equation and use this as a proxy for the volatility of the trading process. However, the economic interpretation of the absolute value of the trade error is not clear.

When referring to the volatility of the trading process, it is more natural to think of the *level* of trading, since a large number of trades will suggest a volatile market. For this reason, I only consider the effects of news on the absolute value of the error in the return equations. The model to be estimated for each series, ΔED_t^{ask}, ΔED_t^{bid}, etc., is given in (2.15) below.

$$\begin{aligned}
\eta_{i,t} &= c + \rho\,\eta_{i,t-1} + \sum_{j=-3}^{h} \mu_j^{R,pri} D_{t-j}^{R,pri} + \nu_{i,t} \\
\eta_{i,t} &= c + \rho\,\eta_{i,t-1} + \sum_{k=-3}^{h} \mu_k^{R,sec} D_{t-k}^{R,sec} + \nu_{i,t} \\
& \qquad\qquad R = US,\ UK \\
& \qquad\qquad i = \Delta ED_t^{ask},\ \Delta ED_t^{bid}, \ldots,\ \Delta SD_t^{bid} \\
\eta_{i,t} &= \frac{|\hat{\epsilon}_{i,t}|}{\psi_t^{|\hat{\epsilon}_{i,t}|}}
\end{aligned} \tag{2.15}$$

The scalar, $\eta_{i,t}$, the dependent variable in the volatility regression, is the absolute value of the fitted residual of the ith equation in (2.2), but deseasonalised by its intra-day pattern. For example, $|\hat{\epsilon}_{\Delta ED^{ask},t}|$ is the absolute value of the fitted residual in the ΔED^{ask} equation in (2.2) at time t. $\psi_t^{|\hat{\epsilon}_{i,t}|}$ is the intra-day pattern of the $|\hat{\epsilon}_{i,t}|$ series, constructed using the FFF method of Section 2.3.3, and again calculated from the residuals in non-announcement days so as to negate the possibly distorting effects that announcement days may have on volatility. The intra-day volatility patterns for the ask return series and the daily autocorrelation functions are shown in Figure 2.16.[47] A notable difference to the findings in other studies is the presence of 'U' shaped return volatility patterns in two of the exchange rates.[48] Payne (1996) and Andersen and Bollerslev (1998) document 'M' shaped volatility patterns (from 0600 to 1800) using indicative quote data in the DEM/USD market. In contrast, volatility tends to be high at the start and end of the European trading day in this study for USD/EUR and GBP/EUR. This can be reconciled when one considers the fact that the indicative quote data used in the above studies are for a 24 hour market. On the other hand, the Reuters D2000-2 platform, although permitting trades from 1800 to 0600, is effectively 'closed' in the two euro markets, shown by the huge drop off in trading that occurs and very high spreads, see Figure 2.1. If

[47]The volatility patterns for the bid return series look very similar to those for the ask return series and for brevity are not shown.

[48]This is still the case when absolute returns or squared returns are considered instead of the absolute value of the residuals from (2.2).

Figure 2.16
Intra-day Volatility Patterns for Ask Returns and Daily Autocorrelation Functions for Absolute Ask Returns in the USD/EUR, GBP/EUR and USD/GBP Markets

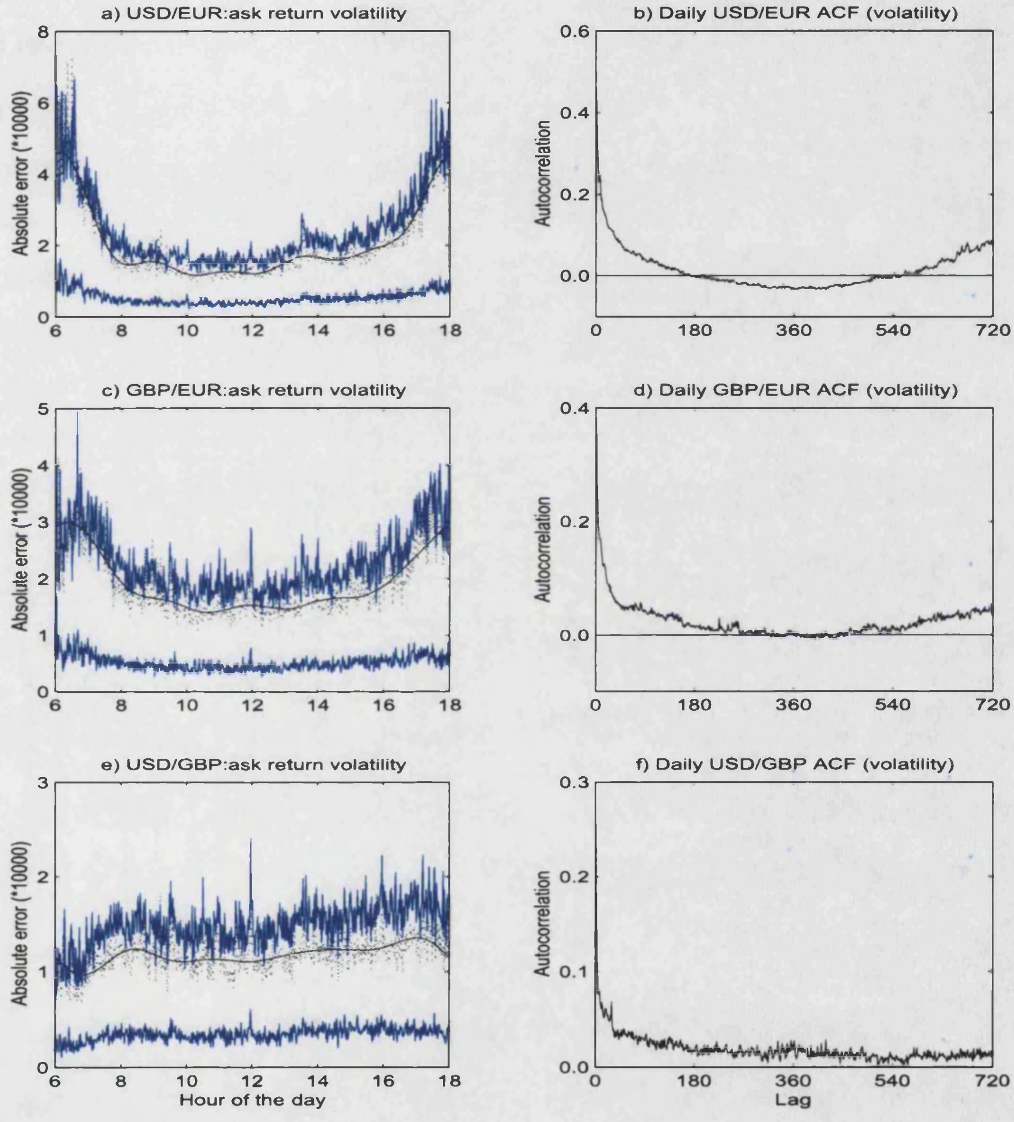

Notes: Panels a, c and e plot the intra-day patterns of the absolute value of the error in the ask return equations from 0600 to 1800 GMT (BST in the summer months) for the three exchange rates. The dotted line shows the intra-day average, the smooth solid line shows the FFF (calculated in a similar way to (2.4)) and the top and bottom blue lines show the 75th and 25th percentiles respectively. All patterns were calculated using non-announcement days. Panels b, d and f show the daily autocorrelation functions for the three rates' absolute errors.

only European and US traders use the Reuters platform for USD/EUR and GBP/EUR, 'U' shaped volatility patterns may not be too surprising. Hsieh and Kleidon (1996) find that return volatility (in terms of standard deviations) is found to be 'U' shaped in the FX market when quotes from dealers in the same geographical location are used. 'U' shaped intra-day volatility patterns are also present in S&P 500 returns where the market is open from 0930 to 1600 (Andersen and Bollerslev 1997b). The absence of a 'U' shaped volatility pattern for USD/GBP returns may be because the Reuters D2000-2 system is still used in Asia/Pacific trading, as discussed earlier.

In (2.15), $D_t^{R,pri}$ takes the value of unity if there is an announcement of 'primary' news from region R at time t. Primary news is that which was found to have first moment effects, described in Section 2.2. $D_t^{R,sec}$ again takes the value of unity but only if there was an announcement of 'secondary' news at time t. These news releases were not found to have any first moment (level) effects but may well cause changes in volatility.[49] Volatility persistence, usually modelled by ARCH type models, is picked up by the inclusion of one lag of $\eta_{i,t}$. Three leads of the news variable are included in (2.15) in order to examine the effects of news on volatility pre release. Due to the persistent effects that news has been found to have on volatility post release, a geometric decay of the volatility response is imposed. $\mu_0^{R,pri}$ ($\mu_0^{R,sec}$) and $\mu_1^{R,pri}$ ($\mu_1^{R,sec}$) are estimated freely but $\mu_j^{R,pri}$ is restricted so that $\mu_j^{R,pri} = \vartheta^{R,pri}\mu_{j-1}^{R,pri}$ for $j \geq 2$, and similarly for $\mu_k^{R,sec}$ ($\mu_k^{R,sec} = \vartheta^{R,sec}\mu_{k-1}^{R,sec}$ for $k \geq 2$), $R = US, UK$.[50] The persistence of the volatility response is therefore picked up by the $\vartheta^{R,pri}$, $\vartheta^{R,sec}$ coefficients.

2.5.1 Estimation results

The estimation results of the volatility equation (2.15) are given in Table 2.8. For ease of exposition, only the results for the ask return volatility models are presented due to the similarity between the ask and bid results. Consistent with a number of previous studies, announcements of macroeconomic news cause an increase in foreign exchange volatility. Consider the effects of news releases on $USD/EUR: ask$ volatility, $\eta_{\Delta ED^{ask},t}$.

[49] A second specification of (2.15) was also estimated, where $D_t^{R,pri}$ ($D_t^{R,sec}$) was replaced by a quantified news variable defined as the absolute value of the news surprise for primary (secondary) releases. This was done since one may expect a large news shock to have a greater effect on volatility than a data release that was largely anticipated.

[50] Due to the geometric decay of the volatility response, (2.15) was estimated by NLLS, with $\vartheta^{R,pri}$, $\vartheta^{R,sec}$ being chosen in the first stage by grid search. For this reason the effects of each type of news from each area were considered separately. If all the effects of news on volatility were modelled simultaneously, and if a grid size of ten points was used, this would imply 10^4 regressions in each step of the grid search, i.e. 10 points for each of $\vartheta^{US,pri}$, $\vartheta^{US,sec}$, $\vartheta^{UK,pri}$ and $\vartheta^{UK,sec}$ implies 10^4 possible combinations.

In the minute of a 'primary' US news release, i.e. US news which also has significant first moment effects, volatility increases by 170% compared to non-announcement days ($\mu_0^{US,pri} = 1.70$). The increase in volatility lasts for some time, shown by the fact that $\mu_1^{US,pri}$ (=1.09) is also significant and the persistence parameter, $\vartheta^{US,pri}$ is also close to unity. In order to give a more intuitive feel for the persistence of the volatility response, Table 2.8 also gives the number of minutes it takes for the maximum effect on volatility to be reached. The number of minutes it takes for half of this effect to dissipate is also given.[51] For example it takes 2 minutes for the maximum effect of primary US news on USD/GBP : ask volatility to be reached and half of this effect has dissipated after a further 17 minutes. The effects of primary news on volatility are perhaps not surprising when one considers the results of previous studies; large market movers, such as the US unemployment rate, have been found to cause the level of the exchange rate to change and also causes exchange rate volatility to increase. However, the effects of secondary news on volatility may not be expected. Secondary US news causes USD/EUR : ask volatility to increase by 76% in the minute of announcement and by 92% in the following minute, with a persistence parameter of 0.935; it takes 12 minutes for half the maximum effect on volatility to dissipate.

Results for GBP/EUR : ask and USD/GBP : ask volatilities are similar. Primary and secondary UK news have significant impacts on both of these volatilities, as do announcements of primary and secondary US news. There is also limited evidence that volatility *increases* in the minutes pre release. The vast majority of the coefficients on news dummies pre release are positive, although only a few of these are significant. This may, at first glance, suggest information leakage but this is unlikely since such news was found to have no first moment (level) effects pre release. More likely, it is evidence of heightened uncertainty around announcement times, as jittery traders pull out of the market, knowing an announcement is imminent. This is also suggested by the fall in trading volume leading up to the data release, and also by the increase in spreads pre release, although the statistical significance of these effects are limited. See Section 2.4.5.

The fact that primary news has a greater impact on volatility than secondary news is to be expected and this is the case for the vast majority of volatility responses; in virtually all the exchange rate equations, $\mu_0^{R,pri}$ ($\mu_1^{R,pri}$) is greater than $\mu_0^{R,sec}$ ($\mu_1^{R,sec}$). Primary UK news also appears to have an effect on $USD/EUR : ask$ volatility, which is surprising

[51]These numbers were calculated by simulating model (2.15). The half life of primary US news is not simply $\log(0.5)/\log(\vartheta^{US,pri})$. This is due to the first order autocorrelation of volatility, shown by the value of ρ, together with the other μ coefficients, which complicate the calculation.

Table 2.8
Volatility Effects of News Releases

	c	ρ	μ_{-3}	μ_{-2}	μ_{-1}	μ_0	μ_1	ϑ	minutes after release max vol	half life
USD/EUR : *ask* volatility										
Primary US news	0.702^a	0.293^a	-0.0644	0.317	0.449	1.696^a	1.095^a	0.959^a	0	14
Secondary US news	0.704^a	0.295^a	0.084	0.602^c	0.809^b	0.755^b	0.915^a	0.935^a	1	12
Primary UK news	0.703^a	0.295^a	-0.0335	0.0498	-0.0774	-0.00696	0.186^a	1.027		
Secondary UK news	0.706^a	0.296^a	-0.107	0.0171	-0.198	-0.0401	-0.122	0.795^a		
GBP/EUR : *ask* volatility										
Primary US news	0.711^a	0.286^a	0.182	0.130	0.448	0.220	0.876^a	0.939^a	2	13
Secondary US news	0.712^a	0.287^a	0.0505	0.429	0.640^b	-0.206	0.868^a	0.905^a	2	8
Primary UK news	0.711^a	0.286^a	0.0433	0.790^a	0.0712	0.483^c	0.481^a	0.962^a	2	19
Secondary UK news	0.713^a	0.287^a	-0.156	0.463^c	0.493^b	-0.0670	0.0933^c	0.985^a	2	48
USD/GBP : *ask* volatility										
Primary US news	0.764^a	0.233^a	0.407	0.253	0.196	0.150	0.696^a	0.956^a	2	17
Secondary US news	0.765^a	0.234^a	0.288	0.612^b	0.194	0.407	0.352^a	0.950^a	0	14
Primary UK news	0.764^a	0.233^a	-0.193	0.0390	0.716^a	-0.0224	0.884^a	0.925^a	2	10
Secondary UK news	0.765^a	0.234^a	-0.181	-0.0540	1.023^a	0.941^a	0.259^a	0.943^a	0	1

Notes: The dependant variable used in the volatility regression is the absolute value of the fitted residual in (2.2) but deseasonalised by its intra-day pattern. Estimation was performed by NLLS, estimating ϑ in the first stage by grid search. ^{a}, ^{b}, ^{c} denote significance at the 1, 5 and 10% levels respectively. The μ coefficients are to be interpreted as follows; on announcement of primary US news, USD/EUR : *ask* volatility increases by a factor of 1.696, shown by μ_0, i.e. volatility increases by 169.6% compared to non-announcement times. One minute after release, volatility is 109.5% greater than normal. t minutes after release, volatility is $109.5 \times \vartheta^{t-1}$ % greater than normal. However, due to the dynamics of the volatility process shown by the ρ parameter, the total effect on volatility is more involved. To give a more intuitive feel for the volatility response, 'max vol' shows the number of minutes it takes for the maximum effect on volatility to be reached (a value of zero indicates that the effect on volatility of the news release is immediate). 'half life' shows the number of minutes it takes (after 'max vol') for half of this maximum volatility response to dissipate. These times were obtained by simulating (2.15). No 'max vol' or 'half life' values are given for UK news effects on USD/EUR : *ask* volatility because in the case of primary news, $\vartheta > 1$, implying an unstable volatility response, and for secondary news, volatility was found to decrease, implying no 'max vol'.

considering that no UK news is found to have significant first moment effects on rates (see Section 2.4.3), although the coefficients on the UK news dummies are quantitatively small relative to those on US news.

Therefore, scheduled macroeconomic news announcements have significant effects on FX volatility. The volatility of the exchange rate increases significantly following announcements of news and this volatility surge persists for some time. When model (2.15) was estimated with quantified news shocks on the right hand side, instead of dummy variables, very few coefficients were found to be significant, and indeed some took the wrong sign. This suggests that the occurrence of news is what induces volatility, rather than the size of the unanticipated news release. One may have thought that a large news shock should cause greater volatility than an announcement that was largely anticipated. This does not appear to be the case.

2.6 Discussion

Under rational expectations and efficient markets hypotheses, the information contained in publicly announced macroeconomic data should be incorporated into the price both immediately and without the need for the trading process. As documented in Evans and Lyons (2003), this does not seem to be true. However, before going on to examine *how* macroeconomic news is incorporated into prices in Chapter 3, this chapter asks *what* happens in FX markets around times of scheduled macroeconomic data releases. By considering bid and ask prices separately, together with buys and sells, I examine the effects of good and bad news from the UK and US on not only exchange rates but also spreads and trading volume. I also examine the effects of news on exchange rate volatility. As one may expect, and consistent with the efficient markets hypotheses, the first moment (level) effects of news releases appear to be complete within one or two minutes. However, there is most certainly an increase in trading activity immediately following a data release, as traders 'hunt' towards the new equilibrium price (Goodhart and Figliuoli 1992). The quote by Andersen and Bollerslev at the head of this chapter, and in particular, the 'hectic trading' hypothesis, appears to be confirmed. Traders appear to interpret a release of macroeconomic data differently, consistent with the theoretical models of Varian (1989), Harris and Raviv (1993) and Kandel and Pearson (1995). After a release of good US news, for example, i.e. news which should cause the dollar to appreciate, there is a significant increase in both buying *and* selling pressure, but the news also leads to net dollar purchases demonstrated in Figures 2.6 and 2.15.[52]

Leading up to the announcement of macroeconomic news, there is limited evidence that traders pull out of the market as they wait until the informational content of the data release is disclosed. Trading volume does appear to fall leading up to the scheduled announcement time and spreads also tend to widen. However, these phenomena are not found to be significant in a statistical sense. The decline in exchange rate volatility leading up to the announcement, as found in Payne (1996), is not found here and there is limited evidence to suggest that volatility is higher pre release, compared to non-announcement days. However, this is likely to be due to jittery traders pulling out of the market, resulting in the volatility of quotes rising.

As has been found in a number of previous studies (Goodhart, Hall, Henry, and Pesaran 1993, Ederington and Lee 1993, Payne 1996, Andersen and Bollerslev 1997a, Andersen and Bollerslev 1998) macroeconomic data releases cause exchange rate volatility to increase and these volatility surges last for some time post release; the half life of the volatility responses to primary UK and US news varies from 10 to 19 minutes in this study. Despite the surge in exchange rate volatility following scheduled macroeconomic releases, trading appears to settle down quite quickly. The level of FX trading increases substantially post release but the effects appear to die out within ten minutes. This suggests that there is indeed a period of hectic trading following data announcements as traders hunt for the new equilibrium price. The differing interpretations of the same public news is a factor that explains this but trading does tend to calm down fairly quickly. The persistence in exchange rate volatility could then be due to initial differences in interpretation of news, which explains the huge increase in trading following the shock, and also by dealers trading towards their new desired portfolio positions. However, once trading calms down, the volatility of the exchange rate continues for some time, perhaps due to the fact that some of the documents that contain the data are quite lengthy and it takes some time for the information to be digested.

Not only does this chapter examine the effects of scheduled macroeconomic news on FX activity, it also examines the dynamics of high frequency spot exchange rates. To my knowledge, this is the first time that the asymmetric impact of buys and sells in FX have been examined in a VAR framework. The results, consistent with those of Engle and Patton (2004) and Biais, Hillion, and Spatt (1995), both of which examine equity markets, can be explained by simple 'barrier' theories of the limit order book. Market

[52]These effects on order flow are left to Chapter 3.

buys (sells) appear to have a greater impact on the ask (bid) price in the book since buys (sells) drain liquidity at the ask (bid) side only. The cross market effects of trades are also documented, consistent with the results of Evans and Lyons (2002a) and Daníelsson, Luo, and Payne (2002) and the effects of spreads/costs of trading on the number of buys and sells are also seen. As expected, a large spread in one market causes the number of buyer and seller initiated trades in that market to fall in the subsequent minute. Not only this, a large spread in one market at date t causes spreads in other markets to widen in the following minute, possibly due to information spill-overs from one market to another.

2.6.1 On the mapping of information to price

As mentioned in the introduction, a common explanation for the rise in trading activity following releases of public information is that data announcements are interpreted differently by different traders/market practitioners. The models of Varian (1989), Kim and Verrecchia (1991), Harris and Raviv (1993) and Kandel and Pearson (1995) all explain rises in trading intensity following announcements of public information by appealing to the idea that traders differ in the way they interpret the news. In this chapter I have shown that scheduled, publicly released information causes trading activity to rise. Following releases of good and bad news, this leads to an increase in both buying *and* selling pressure, perfectly consistent with the differences of opinion hypothesis. Kandel and Pearson (1995) also test this differences of opinion hypothesis by comparing trading volume in announcement and non-announcement times while conditioning on the actual price impact of the news. In particular, if trading activity increases following the news release but the news causes no change in the asset price, then this, they argue, is evidence in favour of the hypothesis that traders differ in their interpretation of publicly released data.

If the data release implied a new equilibrium price for the asset, then an increase in trading activity following the news announcement could simply be a result of portfolio rebalancing trades. On the other hand, if the news announcement results in an unchanged asset price, then heightened trading activity is perfectly consistent with the idea that traders differ in their interpretation of the news release.

To complement the results in Section 2.4.4, where good and bad news causes buying *and* selling pressure to increase, I also present the results of a test similar to that performed by Kandel and Pearson (1995). I essentially test to see whether trading activity is higher during announcement times than during non-announcement times, when the returns observed in those periods are small. Since it was found in Section 2.4.4 that trading activity increases in the first three minutes following a data release, I sample the data at the 3 minute frequency, computing the return in each of the three markets, defined as the log first difference in transaction price, and also noting the number of transactions in each interval. After removing weekends, bank holidays, the overnight period and times when the Reuters data-feed failed, I group all the data for each exchange rate into three categories. The first contains all observations where the return was less than -2 basis points, i.e. where a large negative return is seen. The second contains data where the return is greater or equal to -2 basis points and less than +2 basis points, and the third contains all the observations where the return was greater than or equal to +2 basis points, i.e. a large positive return. For each exchange rate and for each return category I compare the trading volume in periods immediately following a data release with the volume seen on days when no scheduled announcements were made. However, for each FX market I only consider data announcements from the two relevant geographical areas. For example, in the USD/EUR market, I take the observations when US and euro-area news is released, 1330 and 1100 GMT respectively, and compare the trading volume in the three minute intervals starting at these times, to the trading volume at 1330 and 1100 GMT on non-announcement days.[53] In order to account for the intra-day pattern in trading activity, documented in Section 2.3.3, all trading volumes have the intra-day pattern subtracted from them, where the intra-day pattern is calculated from the FFF methods in Section 2.3.3 and using only non-announcement days.[54] A Mann-Whitney rank test is then performed to see whether trading volume on announcement days is significantly greater than that seen on non-announcement days, for each of the return categories and for each exchange rate. The results are presented in Table 2.9.

The results suggest that for each exchange rate, trading volume is greater during the three minute periods following data releases than when no announcement is made, even when the news does not result in a change in price. This is therefore consistent with the hypothesis that traders differ in their interpretation of macroeconomic news, a hypothesis which is further supported by the finding that both good and bad news causes buying *and* selling pressure to increase.

[53] The euro-area news releases considered are those used and described in Chapter 3.

[54] In order to compare like with like, we have to consider deseasonalised trading volume, since it may be the case that trading at 1330 GMT is significantly different from trading at 1100 GMT.

Table 2.9
Mann-Whitney Rank Test of Differences in Trading Volume (Announcement Versus Non-Announcement Times)

		Return, R, (basis points)		
		$R < -2$	$-2 \leq R < +2$	$+2 \leq R$
USD/EUR	U	814.5	957	870
	μ_U	595	650	540
	σ_U^2	6941.67	8450	6030
	t-stat	**2.63**[a]	**3.34**[a]	**4.25**[a]
GBP/EUR	U	317.5	1151	352.5
	μ_U	272	799	256.5
	σ_U^2	2312	10919.67	2009.25
	t-stat	**0.95**	**3.37**[a]	**2.14**[b]
USD/GBP	U	849	1554	578.5
	μ_U	567	1000.5	378
	σ_U^2	6615	16508.25	3654
	t-stat	**3.47**[a]	**4.31**[a]	**3.32**[a]

Notes: The table shows the statistics for the Mann-Whitney rank test that trading volume is the same during announcement and non-announcement times. For each exchange rate, the 3 minute return is calculated following news announcements and trading volume in those periods is compared to the volume seen at the same time on non-announcement days. For each exchange rate, only the news emanating from the two associated geographical areas are considered. For example, only US and euro-area news are considered when testing the differences in USD/EUR trading volume. Trading volume is deseasonalised by subtracting the intra-day average for that time interval, calculated by fitting a Fourier Flexible Form. The test statistic is calculated as $\frac{(U-\mu_U)}{\sigma_U}$ and compared to the normal distribution.

$$U = n_1 n_2 + \frac{n_1 (n_1 + 1)}{2} - R_1$$
$$\mu_U = \frac{n_1 n_2}{2}$$
$$\sigma_U^2 = \frac{n_1 n_2 (n_1 + n_2 + 1)}{12}$$

n_2 (n_1) is the number of observations in (non) announcement times and R_1 is the sum of the ranks for announcement period trading volume. Large positive values of the test statistic imply trading volume is larger following news announcements. [a], [b], [c] denote significance at the 1, 5 and 10% levels respectively.

2.6.2 Good and bad news

In this study, I also find asymmetric responses to good and bad news emanating from the UK and the US. Following a one standard deviation announcement of good US news, this causes the USD/EUR and USD/GBP prices to adjust by approximately 11 and 4 basis points in the minute of release respectively. Following a similar sized announcement of bad US news, these rates adjust by approximately 2 and 1 basis points in the minute of release. See Table 2.6. The results for UK releases appear to be more symmetric, although USD/GBP rates still react more to good UK news than to bad news. Andersen, Bollerslev, Diebold, and Vega (2003) suggest that bad news in good times (economic expansions) will have greater impacts than good news in good times, and their evidence, using high frequency/5 minute data from January 1992 to December 1998, a period of continued US expansion, supports this hypothesis. Such findings can be explained, they argue, using behavioural theories, whereby traders are fooled into thinking that the good state, when prevailing for long enough, is the norm. Hence good news in good times simply confirms their beliefs while bad news in good times will have a large impact since it comes more as a surprise. This reasoning can be used (tentatively) to explain the findings here. Five months of the data in this study occur after the bursting of the NASDAQ stock market bubble and hence could be considered 'bad times', although the NBER only define the US recession to start in March 2001. However, this explanation as to why good US news has larger impacts than similar sized bad news is purely a conjecture and cannot be tested using this short (eight months) dataset available here. Galati and Ho (2003) also find that the USD/EUR rate responds more to good US data than to bad, when considering data sampled at the daily frequency from January 1999 to December 2000 and find that the response to news is greater when data change from good to bad or vice versa.

2.7 Conclusions

According to pure rational expectations and efficient markets hypotheses, the information contained in publicly released macroeconomic information should be incorporated into prices with no need for the trading process. Using eight months of data on three major floating rates, USD/EUR, GBP/EUR and USD/GBP, this chapter shows the effects of UK and US data releases on FX activity. An announcement of good US news, for example, causes both buying *and* selling activity to increase, consistent with the ideas that traders disagree over the mapping of information to price (Varian 1989, Harris and Raviv 1993, Kandel and Pearson 1995). Macroeconomic news releases cause trading activity to increase dramatically post release but trading tends to settle down quite quickly. This suggests that there is indeed a period of hectic trading following news releases as traders hunt for the new equilibrium price. In contrast, the volatility surge seen in the exchange rate series lasts for some time, with half lives up to 20 minutes. The initial volatility increase is likely to be due to the differences in opinion, which generates the high trading volume, and also because of the trades initiated by dealers as they move towards their new desired portfolio positions. The continued exchange rate volatility that persists for the next hour or so, perhaps due to the fact that some of the documents containing the data are quite lengthy, does not appear to coincide with heightened trading activity.

Pre release, there is only limited evidence that traders pull out of the market. Trading volume does fall and spreads do widen in the minutes leading up to the announcement, as one would expect if traders waited until the informational content of the release was disclosed, but these effects are quantitatively small and statistically insignificant. By using the VAR model of (2.2) I have been able to show *what* happens in the FX markets around periods of scheduled news releases. The question of *how* this information enters prices, whether it is direct or intermediated by order flow, is addressed in the next chapter.

2.A Appendix

2.A.1 Cointegration analysis of high frequency foreign exchange data

In Chapters 2, 3 and 4 I examine the effects of scheduled macroeconomic news on various foreign exchange market statistics, including FX returns, order flows, trading volume, spreads and depths. Chapters 2 and 3 make explicit use of the vector autoregression (VAR) framework and it is the purpose of this appendix to give more detail on this model, explaining how it is estimated and used to analyse the high frequency foreign exchange data in this study. The VAR provides a very flexible statistical model which can be employed to analyse the interrelationships between a number of variables. By using impulse response functions and variance decompositions, one can obtain a very clear picture as to how the variables interact and how the dynamics, of what can become a complex and cumbersome model, evolve. In Section 2.A.2 I introduce the basic setup and show how this can be extended to analyse quite complex asset market dynamics between bid and ask prices, buyer and seller initiated trades and between different markets. In Section 2.A.3 I test for the cointegrating rank of the model presented in (2.2) and in Section 2.A.4 I test the theoretical cointegrating relationships proposed in Section 2.3.4. Impulse response functions for the cointegrating VAR are discussed in Section 2.A.5 and Section 2.A.6 analyses variance decompositions.

2.A.2 Empirical methodology

The basic framework used to analyse these FX data is a cointegrating structural VAR. This is given in (2.A.1) below. In its simplest form, ΔY_t is a 2×1 vector containing the asset return, ΔP_t, and order flows, F_t, i.e. $\Delta Y_t = [\Delta P_t \; F_t]'$. The contemporaneous effect of flows on returns is captured by the scalar β parameter that appears in the matrix that pre-multiplies ΔY_t on the right hand side of (2.A.1). It is common in the literature to model asset returns and order flows using a recursively ordered structural VAR, whereby asset returns depend on contemporaneous order flows but where the converse is not true. This framework is used in this chapter and also in Chapter 3, but a model which allows order flow to depend on contemporaneous returns is analysed in detail in Chapter 5.[55]

$$
\begin{aligned}
\Delta Y_t &= c + \kappa t + \alpha\gamma' Y_{t-1} + \begin{bmatrix} 0 & \beta \\ 0 & 0 \end{bmatrix} \Delta Y_t + \sum_{i=1}^{m} \theta_i \Delta Y_{t-i} + \epsilon_t \\
Var\left[\epsilon_t\right] &= \Omega = \begin{bmatrix} \Omega_{11} & 0 \\ 0 & \Omega_{22} \end{bmatrix}
\end{aligned}
\tag{2.A.1}
$$

In order to estimate the structural parameter, β, the contemporaneous effect of flows on returns, a restriction has to be made in the model and it is common in the literature to impose this restriction in the variance/covariance matrix of the residuals, by making Ω diagonal. The parameters in (2.A.1) can therefore be estimated by firstly estimating the reduced form (including only the lags of ΔY_t in the regression) and backing out the structural parameter from the variance/covariance matrix of the residuals, see Enders (1995) and Hamilton (1994) for example. Alternatively, one can estimate the structural form equation by equation since the triangular representation used in (2.A.1), together with a diagonal variance/covariance matrix, Ω, implies F_t, the second element of ΔY_t, is not correlated with the error in the ΔP_t equation, even though F_t is endogenous. These two approaches are equivalent and so produce the same results. See Greene (2000) for example.

The two equation VAR in (2.A.1) can easily be extended to examine the interaction between a number of returns and trades. The model presented below in (2.A.2) is used in this chapter and a simplified version is used in Chapter 3. It essentially allows me to examine the dynamics between six returns, from bid and ask prices in the three FX markets under consideration, and the six trade variables, the buys and sells in each market.

[55]It is also common to see models similar to (2.A.1) being estimated without any cointegrating terms, especially if the VAR only examines returns and order flows. See Hasbrouck (1991a) and Payne (2003a) for example. However, in Section 2.A.3 I test for, and indeed find evidence of, cointegration in a larger system, and therefore allow for the possibility of cointegration in (2.A.1).

Hence the model becomes a 12 variable structural VAR.[56]

$$\underbrace{\begin{bmatrix} \Delta P_t \\ T_t \end{bmatrix}}_{\Delta Y_t} = c + \kappa t + \alpha\gamma' \underbrace{\begin{bmatrix} P_{t-1} \\ \sum_{r=1}^{t-1} T_r \end{bmatrix}}_{Y_{t-1}} + \begin{bmatrix} 0 & \beta \\ 0 & 0 \end{bmatrix} \underbrace{\begin{bmatrix} \Delta P_t \\ T_t \end{bmatrix}}_{\Delta Y_t} + \sum_{i=1}^{m} \theta_i \underbrace{\begin{bmatrix} \Delta P_{t-i} \\ T_{t-i} \end{bmatrix}}_{\Delta Y_{t-i}} + \epsilon_t$$

$$Var\left[\epsilon_t\right] = \Omega = \begin{bmatrix} \Omega_{11} & 0 \\ 0 & \Omega_{22} \end{bmatrix}$$

$$\underbrace{\begin{bmatrix} \Delta P_t \\ T_t \end{bmatrix}}_{12\times 1} = \Delta Y_t = \begin{matrix} [\Delta ED_t^{ask} & \Delta ED_t^{bid} & \Delta ES_t^{ask} & \Delta ES_t^{bid} & \Delta SD_t^{ask} & \Delta SD_t^{bid} \\ B_t^{ED} & S_t^{ED} & B_t^{ES} & S_t^{ES} & B_t^{SD} & S_t^{SD}]' \end{matrix} \tag{2.A.2}$$

The standard framework is extended in this model by examining the bid and ask returns separately, as in Engle and Patton (2004), and by considering the number of buys and sells separately. This therefore allows us to examine exchange rate returns and buying and selling pressure, as well as bid-ask spreads, order flows and trading volume simultaneously.

2.A.3 Testing for cointegrating rank

As explained earlier, even though each of the variables in Y_t, the 6 exchange rates and the 6 cumulative trade variables were found to be $I(1)$, it is quite possible that some linear combination(s) of these variables is (are) stationary, i.e. that they are cointegrated. For example, one would expect the log bid and ask prices of each exchange rate to be cointegrated, with a cointegrating vector of $[1 \ -1]'$, implying a stationary log spread. Tests of this and other theoretical cointegrating relationships will be left until Section 2.A.4. However, firstly we need to determine the dimension of the cointegration space, i.e. find the cointegrating rank of $\alpha\gamma'$ and hence determine how many linearly independent vectors there are that cointegrate the system. To do this I use the likelihood methods of Johansen (1995), which are also used to test the hypothesised cointegrating vectors that one would expect to find in the system.[57] As discussed above, the time trend is restricted so as to only enter the cointegration space, ruling out quadratic time trends. However,

[56]The model presented in (2.A.2) is the same as that given in (2.2) but without the exogenous macroeconomic news variables as right hand side regressors.

[57]Hargreaves (1994) suggests that Johansen's maximum likelihood approach works well when the sample size is large (> 100) and so is used here.

the constant in (2.A.2) is left unrestricted. The cointegrating terms are therefore given in (2.A.3) below.

$$\kappa t + \alpha\gamma' Y_{t-1} = \underbrace{\alpha}_{12\times h} \begin{bmatrix} \underbrace{\gamma_0}_{h\times 1} & \underbrace{\gamma'}_{h\times 12} \end{bmatrix} \underbrace{\begin{bmatrix} t \\ Y_{t-1} \end{bmatrix}}_{13\times 1} \tag{2.A.3}$$

α is a $12\times h$ matrix of speed of adjustment coefficients, γ is a $12\times h$ matrix of cointegrating vectors and κ is 12×1 but is restricted to have rank equal to h. The likelihood methods used here are described at length in Johansen (1995) but the notation below is largely borrowed from Hamilton (1994).[58] To test for cointegrating rank, the structural VAR presented in (2.A.2) is first converted to a reduced form, shown in (2.A.4) below.

$$\begin{aligned} \underbrace{\begin{bmatrix} I_6 & -\beta \\ 0 & I_6 \end{bmatrix}}_{B} \Delta Y_t &= c + \alpha\gamma_0 t + \alpha\gamma' Y_{t-1} + \sum_{i=1}^{m} \theta_i \Delta Y_{t-i} + \epsilon_t \\ \Rightarrow \Delta Y_t &= \underbrace{B^{-1}c}_{c^+} + \underbrace{B^{-1}\alpha}_{\alpha^+}\gamma_0 t + \underbrace{B^{-1}\alpha}_{\alpha^+}\gamma' Y_{t-1} + \sum_{i=1}^{m} \underbrace{B^{-1}\theta_i}_{\theta_i^+} \Delta Y_{t-i} + \underbrace{B^{-1}\epsilon_t}_{\epsilon_t^+} \\ Var\left[\epsilon_t^+\right] &= \Sigma = B^{-1}\Omega\left(B^{-1}\right)' \end{aligned} \tag{2.A.4}$$

This is done to simplify the mechanics of the test since a reduced form VAR estimated as a SUR system, can be estimated equation by equation using OLS.[59] Converting (2.A.2) into a reduced form VAR will have no effect on either the cointegrating vectors or the test for cointegrating rank. Since $\alpha\gamma'$ has rank h, pre-multiplying this by the inverse of a full rank 12×12 matrix of structural parameters will have no effect on the rank of $\alpha^+\gamma'$, where α^+ is the $12 \times h$ matrix of reduced form speed of adjustment coefficients, $B^{-1}\alpha$.

Since the time trend is restricted so as to only enter the cointegrating vectors, the first stage of the test for cointegrating rank includes a battery of three auxiliary regressions,

[58]The procedure discussed below and the 'λ max' and 'trace' tests are based on the VAR specification where the constant, c, in (2.A.2) is unrestricted and the time trend is restricted so as to only enter via the cointegrating relationships. λ max and trace tests were also performed for a number of other specifications, in particular 'c restricted, no time time trend', 'c unrestricted, no time trend' and 'c unrestricted, t unrestricted'.

[59]This greatly simplifies the auxilliary regressions presented below.

which partial out the constant and lagged dependent variables.

$$
\begin{aligned}
\Delta Y_t &= \hat{\pi}_0 + \sum_{i=1}^{m} \hat{\pi}_i \Delta Y_{t-i} + \hat{u}_t \\
t &= \hat{\chi}_0 + \sum_{i=1}^{m} \hat{\chi}_i \Delta Y_{t-i} + \hat{w}_t \\
Y_{t-1} &= \hat{\omega}_0 + \sum_{i=1}^{m} \hat{\omega}_i \Delta Y_{t-i} + \hat{v}_t
\end{aligned}
\tag{2.A.5}
$$

Letting $\hat{x}_t = [\hat{w}_t \ \hat{v}_t']'$ then from the fitted residuals, $\hat{u}_t$, $\hat{w}_t$ and $\hat{v}_t$, I construct the sample variance/covariance matrices $\hat{\Sigma}_{uu}$, $\hat{\Sigma}_{xx}$, $\hat{\Sigma}_{ux}$ and $\hat{\Sigma}_{xu}$.

$$
\begin{aligned}
\underbrace{\hat{\Sigma}_{uu}}_{12\times 12} &= \frac{1}{T}\sum_{t=1}^{T} \hat{u}_t \hat{u}_t' \\
\underbrace{\hat{\Sigma}_{xx}}_{13\times 13} &= \frac{1}{T}\sum_{t=1}^{T} \hat{x}_t \hat{x}_t' \\
\underbrace{\hat{\Sigma}_{ux}}_{12\times 13} &= \frac{1}{T}\sum_{t=1}^{T} \hat{u}_t \hat{x}_t' \qquad\qquad \hat{\Sigma}_{xu} = \hat{\Sigma}_{ux}'
\end{aligned}
\tag{2.A.6}
$$

The motivation for performing these auxiliary regressions is most articulately described in Hamilton (1994), chapter 20 and the reader is referred to this chapter for a more detailed discussion.[60] Hamilton (1994) shows that, using canonical correlations, the maximum value of the log likelihood function, subject to there being h cointegrating relations and no quadratic trends, is ℓ_h.

$$
\ell_h = -\left(\frac{Tn}{2}\right)\log(2\pi) - \frac{Tn}{2} - \left(\frac{T}{2}\right)\log\left|\hat{\Sigma}_{uu}\right| - \left(\frac{T}{2}\right)\sum_{i=1}^{h}\log\left(1-\hat{\lambda}_i\right) \tag{2.A.7}
$$

where $\hat{\lambda}_i$ $(i = 1, \ldots, h)$ are the h largest eigenvalues of the $(n+1) \times (n+1)$ matrix $\hat{\Sigma}_{xx}^{-1}\hat{\Sigma}_{xu}\hat{\Sigma}_{uu}^{-1}\hat{\Sigma}_{ux}$.[61] Likelihood ratio tests of the cointegrating rank can then be performed quite easily using (2.A.7). To test the null hypothesis that the rank of $\alpha^{+}\gamma' = h$, i.e. that there are h cointegrating vectors, against the alternative that there are 12 cointegrating

[60]See also Johansen (1995) and Banerjee, Dolado, Galbraith, and Hendry (1993).

[61]Maximum likelihood estimates of $[\gamma_0 \ \gamma']'$ are given by $[z_1 \ldots z_h]$ where $z_1 \ldots z_h$ are the normalised eigenvectors associated with the h largest eigenvalues, $\hat{\lambda}_i$.

relationships, i.e. that $\alpha^{+}\gamma'$ has full rank and that every element of Y_t is stationary, the 'trace' test statistic can be used, shown in (2.A.8). Alternatively, one can test the null hypothesis that there are h cointegrating relationships against the alternative that the rank equals $h+1$. This is the 'λ max' test statistic.

$$
\begin{array}{lll}
'trace' & H_0 \; rank = h \; vs \; H_A \; rank = n & \\
 & likelihood \; ratio \; statistic \; = -T \displaystyle\sum_{i=h+1}^{n} \log\left(1-\hat{\lambda}_i\right) & \\
'\lambda \; max' & H_0 \; rank = h \; vs \; H_A \; rank = h+1 & \\
 & likelihood \; ratio \; statistic \; = -T \log\left(1-\hat{\lambda}_{h+1}\right) &
\end{array}
\tag{2.A.8}
$$

Under the null hypothesis, all the variables entering (2.A.2) are stationary, but since the alternative hypothesis includes non-stationary combinations of Y_t, then the critical values of these tests are non-standard and therefore the values tabulated in Osterwald-Lenum (1992) are used. Results for the trace test are reported in Panel A of Table 2.A.1.[62]

The null hypothesis of there being 6 cointegrating relationships against the alternative of there being 12 is rejected but the null hypothesis of there being 7 is not. Panel B gives the results of the 'λ max' test. Again, the null that there are 6 cointegrating vectors is rejected against the alternative that $h = 7$, but the null of 7 cointegrating vectors against the alternative of there being 8 such relationships is rejected. Both tests therefore suggest that there are 7 cointegrating relationships between the 12 elements of Y_t.[63]

2.A.4 Tests of the cointegrating vectors

The previous section suggested that there are 7 cointegrating relationships between the 12 variables in Y_t and as explained in Section 2.3.4, economic theory can help identify

[62]Since the constant in (2.A.2) is unconstrained but the time trend is restricted to only enter the cointegration space, hence ruling out quadratic trends in the levels of the series, then the critical values from Osterwald-Lenum (1992), Table 2* are used. However, critical values are only given for up to 11 random walks under the null, and therefore the null hypothesis that $h = 0$ against the alternative that $h = 12$, 'trace', or $h = 1$, 'λ max', cannot be tested.

[63]In the 'c unrestricted, no time trend' (Osterwald-Lenum (1992), Table 1) and 'c unrestricted, t unrestricted' (Osterwald-Lenum (1992), Table 2) specifications, both the trace and λ max tests suggested a cointegrating rank equal to 7. On the other hand, the 'c restricted, no time trend' (Osterwald-Lenum (1992), Table 1*) specification suggested 8 cointegrating vectors. However, in this case the 1% critical values were only just breached and since this specification implies no trend in the levels of the series, i.e. no drift in any of the exchange rates, this specification appears overly restrictive. For this reason, the following empirical analysis assumes the existence of 7 cointegrating vectors.

Table 2.A.1

Johansen Test for Cointegrating Rank: Bid and Ask Prices

Panel A 'trace' test				Panel B 'λ-max' test			
H_0 rank ($\alpha\gamma'$)$= h$	H_A	test statistic	1% critical value	H_0 rank ($\alpha\gamma'$)$= h$	H_A	test statistic	1% critical value
$h = 0$	$h = 12$	19818.52		$h = 0$	$h = 1$	6984.96	
$h = 1$	$h = 12$	12833.56	327.45	$h = 1$	$h = 2$	5431.91	79.23
$h = 2$	$h = 12$	7401.65	279.07	$h = 2$	$h = 3$	4026.30	73.73
$h = 3$	$h = 12$	3375.35	234.41	$h = 3$	$h = 4$	2590.75	67.88
$h = 4$	$h = 12$	784.60	196.08	$h = 4$	$h = 5$	456.59	62.46
$h = 5$	$h = 12$	328.02	158.49	$h = 5$	$h = 6$	176.88	54.71
$h = 6$	$h = 12$	151.14	124.75	$h = 6$	$h = 7$	61.38	49.51
$h = 7$	$h = 12$	89.76	96.58	$h = 7$	$h = 8$	41.06	42.36
$h = 8$	$h = 12$	48.70	70.05	$h = 8$	$h = 9$	26.03	36.65
$h = 9$	$h = 12$	22.66	48.45	$h = 9$	$h = 10$	13.24	30.34
$h = 10$	$h = 12$	9.43	30.45	$h = 10$	$h = 11$	7.67	23.65
$h = 11$	$h = 12$	1.76	16.26	$h = 11$	$h = 12$	1.76	16.26

Notes: Critical values are taken from Osterwald-Lenum (1992), Table 2* (c unrestricted, t restricted) but statistics are only tabulated for up to 11 random walks. Hence we cannot test the null that $h = 0$ against either $h = 12$, 'trace' test, or $h = 1$, 'λ-max' test.

which vectors cointegrate the system. Intuitively, one would expect the log bid and log ask prices of each exchange rate to be cointegrated with a cointegrating vector of $[1 \;\; -1]'$ implying a stationary log spread and leading to, at least, three cointegrating vectors in the system. However, there is also an absence of arbitrage condition that should hold, causing the three exchange rates to be cointegrated and as explained in Killeen, Lyons, and Moore (2002) and Bjønnes and Rime (2003), the level of each exchange rate and cumulative order flow should also be cointegrated.

This conveniently gives us 7 theoretical cointegrating vectors, matching the number of relationships found in Section 2.A.3. However, it still remains to be shown that these 7 relationships are indeed the same as those which cointegrate the system. In this section I test whether the 7 cointegrating vectors are of the form suggested by economic theory; 3 cointegrating relationships between bid and ask prices, 1 relationship between the three exchange rates and a further 3 relationships between the level of the exchange rate and cumulative order flow.

2.A.4.1 ADF tests

The results of the ADF tests are reported in Table 2.A.2. Here, the tests for the theoretical cointegrating relationships between bid and ask prices are given, along with the tests for cointegration between prices and cumulative order flows. When testing for cointegration between log USD/EUR bid and ask prices, ED_t^{bid} is regressed on ED_t^{ask} and ADF tests are performed on the fitted residuals. However, in the first stage regression, due to the presence of the spread between ask and bid prices, a number of different specifications are used. In the first specification, ED_t^{bid} is regressed only on ED_t^{ask}. In the second specification, ED_t^{bid} is regressed on a constant and ED_t^{ask}. Next, ED_t^{bid} is regressed on ED_t^{ask} and the intra-day pattern of the spread (the FFF constructed in Section 2.3.3) in order to allow for the 'U' shaped pattern of the difference in log bid and ask prices that naturally occurs during the trading day, and finally, ED_t^{bid} is regressed on a constant, the intra-day pattern and ED_t^{ask}. A similar procedure was used to test for cointegration between bid and ask prices in the GBP/EUR and USD/GBP markets and for convenience only the results of specification 4 are given in the table, i.e. a constant and the intra-day pattern of spreads are included in the first stage regression. When testing for cointegration between ED_t^{ask}, ES_t^{bid} and SD_t^{bid}, ED_t^{ask} is regressed on ES_t^{bid} and SD_t^{bid} (spec 1), and a constant is included in specification 2. Evidence suggests that the bid and ask prices are cointegrated and that the three exchange rates are also cointegrated.[64] For example, testing for cointegration between ED_t^{ask} and ED_t^{bid} produces a t statistic of -5.17, which is significant at the 5% level using Phillips and Ouliaris tables. Hence non-stationarity of the residuals is rejected, suggesting cointegration between log ask and bid prices. Also note that the β coefficients in the regression of the bid price on the ask price is very close to unity, as one would expect. The coefficients β_1 and β_2 in the regression of ED_t^{ask} on ES_t^{bid} and SD_t^{bid} are also close to unity, as expected.

The ADF tests for cointegration between log ask prices and cumulative order flow, on the other hand, are not supportive of the cointegration hypothesis. For each FX market the log ask price is regressed on cumulative order flow in specification 1, a constant and cumulative order flow in specification 2 and in the final version a time trend is included. However, in each case the fitted residuals appeared to be non-stationary, implying prices and cumulative order flows are not cointegrated. In the case of USD/EUR, the ADF test statistic was -1.91, insignificant at the 5% level, again using Phillips and Ouliaris tables. However, when testing for cointegration between the level of the exchange rate and cumulative order flow, there is a significant problem associated with the missing data periods, described in Section 2.2.1. Due to the nature of the data, only the exchange rate and the number of buys and sells in each period are given. Therefore if the Reuters data-feed collapsed from periods $t+1$ to $t+10$, the exchange rate movement is still picked up when the data-feed continues at time $t+11$, but the change in cumulative order flow from $t+1$ to $t+10$ is entirely missed. In order to compensate for this shortcoming, a dummy variable was included for each of the minutes where the data-feed was brought back on line after a collapse of more than 3 hours. In this way, the large jumps in the exchange rate that may occur in these intervals, but where the change in cumulative order flow is zero, will not affect the test for cointegration. However, for each of the three exchange rates, these results are still not supportive of the cointegration hypothesis. Therefore, even when allowing for breaks in the data-feed, the hypothesis that the level of the exchange rate and cumulative order flow are cointegrated is still rejected.

This may, however, be a result of the fact that single equation methods are used in these cointegration tests. In order to take advantage of the information available in the system, the full information maximum likelihood methods of Johansen (1995) can be used. These are presented next.

[64]The conclusions are not affected by the choice of specification in the first stage regression.

Table 2.A.2

ADF Tests for Cointegration

Cointegration between bid and ask prices

First stage $X_t^{bid} = c + \alpha\psi_t^{Xspread} + \beta X_t^{ask} + \epsilon_t \qquad X = ED,\ ES,\ SD$

ADF test $\Delta\hat{\epsilon}_t = \gamma_0\hat{\epsilon}_{t-1} + \gamma_1\Delta\hat{\epsilon}_{t-1} + \ldots + \gamma_p\Delta\hat{\epsilon}_{t-p} + \eta_t$

	c	α	β	ADF statistic	5% critical value
USD/EUR	0.0^4339	-1.001	1.0009	-5.17	-3.37
GBP/EUR	0.000642	-1.002	1.0013	-11.55	-3.37
USD/GBP	0.0^4786	-0.997	1.0002	-27.67	-3.37

Cointegration between the three rates

First stage $ED_t^{ask} = c + \beta_1 ES_t^{bid} + \beta_2 SD_t^{bid} + \epsilon_t$

ADF test $\Delta\hat{\epsilon}_t = \gamma_0\hat{\epsilon}_{t-1} + \gamma_1\Delta\hat{\epsilon}_{t-1} + \ldots + \gamma_p\Delta\hat{\epsilon}_{t-p} + \eta_t$

c	β_1	β_2	ADF statistic	5% critical value
0.000174	0.9984	0.9990	-5.18	-3.77

Cointegration between prices and cumulative order flow

First stage $X_t^{ask} = c + \beta_1 t + \beta_2 \sum_{k=1}^{t}\left(B_k^X - S_k^X\right) + \epsilon_t \qquad X = ED,\ ES,\ SD$

ADF test $\Delta\hat{\epsilon}_t = \gamma_0\hat{\epsilon}_{t-1} + \gamma_1\Delta\hat{\epsilon}_{t-1} + \ldots + \gamma_p\Delta\hat{\epsilon}_{t-p} + \eta_t$

	c	β_1	β_2	ADF statistic	5% critical value
USD/EUR	0.0343	-0.0^5218	0.0^4222	-1.96	-3.42
GBP/EUR	-0.4514	-0.0^5646	0.0^4390	-1.53	-3.42
USD/GBP	0.5005	-0.0^5132	0.0^5320	-2.25	-3.42

Notes: ED_t^{bid} (ED_t^{ask}) is the log of the $USD/EUR : bid$ (ask) price in minute t and similarly for the other two rates. $\psi_t^{Xspread}$ is the intra-day pattern of the spread in market X in minute t, $X = ED,\ ES,\ SD$, calculated from the Fourier Flexible Form in Section 2.3.3. B_k^X (S_k^X) is the number of buyer (seller) initiated trades in market X in minute k. In the ADF tests, the lag length was increased until the η_t error resembled white noise (using an LM test). Only rarely did this happen and so a maximum lag length of 15 was used.

2.A.4.2 System methods of cointegration tests

In this section a number of tests are performed to test the structure of the 7 cointegrating vectors using the full information maximum likelihood methods of Johansen (1995). Three particular tests are performed, each imposing a different set of restrictions. The first imposes the restrictions that 4 of the cointegrating vectors are those associated with bid and ask prices and between the levels of each of the three rates, while leaving the 3 remaining cointegrating vectors unrestricted. The second test imposes the restrictions that 3 of the cointegrating vectors are those between the level of each exchange rate and cumulative order flow, while leaving the remaining 4 vectors unrestricted. Finally, the 7 theoretical cointegrating vectors are tested jointly.

2.A.4.3 Test 1: cointegration between bid and ask prices and between the three rates

These restrictions imply that the matrix of cointegrating vectors in (2.A.4) is of the form given in (2.A.9) below.

$$\Delta Y_t = c^+ + \alpha^+\gamma_0 t + \alpha^+\gamma' Y_{t-1} + \sum_{i=1}^{m} \theta_i^+ \Delta Y_{y-i} + \epsilon_t^+$$

$$\begin{aligned} \alpha^+\gamma_0 t + \alpha^+\gamma' Y_{t-1} &= \alpha^+ \left[\gamma_0 \ \gamma'\right] \begin{bmatrix} t \\ Y_{t-1} \end{bmatrix} \\ &= \alpha^+ \left[\begin{array}{cc} \multicolumn{2}{c}{\underbrace{\underbrace{\begin{array}{ccccccccc} 0 & 1 & -1 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 & \dots & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 & -1 & \dots & 0 \end{array}}_{\Upsilon'}}_{4\times 13}} \\ \multicolumn{2}{c}{\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots} \\ \underbrace{\psi_0'}_{3\times 1} & \underbrace{\psi_1'}_{3\times 12} \end{array} \right] \begin{bmatrix} t \\ Y_{t-1} \end{bmatrix} \end{aligned} \tag{2.A.9}$$

However, even though the time trend is restricted to only enter the cointegration space, in order to rule out quadratic trends in the levels of the series, it is infeasible that a time trend enters the spread in each market, and also between the levels of the three rates.[65]

[65]The spread and the difference between the three rates, associated with absence of arbitrage, do not

The time trend is therefore restricted to only enter the three unrestricted cointegrating vectors.

Υ is therefore the matrix of known cointegrating vectors associated with the bid-ask spread and the absence of arbitrage between the three rates. As in Section 2.A.3, the test continues by partialling out the constant and the lagged dependent variables, as in (2.A.5). By construction, the errors $\hat{u}_t$, $\hat{w}_t$ and $\hat{v}_t$ are related by the cointegration terms. This is shown in (2.A.10), where the α^+ matrix of speed of adjustment coefficients has been partitioned into α_1^+ (12×4), relating the 4 known cointegrating vectors, and α_2^+ (12×3), relating to those that are unrestricted.

$$\begin{aligned} \hat{u}_t &= \left[\alpha_1^+ \; \alpha_2^+\right] \begin{bmatrix} \Upsilon' \\ \psi_0' \; \psi_1' \end{bmatrix} \underbrace{\begin{bmatrix} \hat{w}_t \\ \hat{v}_t \end{bmatrix}}_{\hat{x}_t} + error_t \\ \hat{u}_t &= \alpha_1^+ \Upsilon' \hat{x}_t + \alpha_2^+ \left[\psi_0' \; \psi_1'\right] \hat{x}_t + error_t \end{aligned} \tag{2.A.10}$$

The model is then concentrated with respect to α_1^+ by regressing $\hat{u}_t$ and $\hat{x}_t$ on $\Upsilon'\hat{x}_t$. The errors from these regressions are related through the unrestricted cointegration terms.

$$\begin{aligned} \hat{u}_t &= \tilde{\alpha}_1^+ \Upsilon' \hat{x}_t + \hat{U}_{\Upsilon,t} \\ \hat{x}_t &= \hat{\alpha}_1^+ \Upsilon' \hat{x}_t + \hat{X}_{\Upsilon,t} \\ \Rightarrow \hat{U}_{\Upsilon,t} &= \alpha_2^+ \left[\psi_0' \; \psi_1'\right] \hat{X}_{\Upsilon,t} + \tilde{error}_t \end{aligned} \tag{2.A.11}$$

where $\hat{U}_{\Upsilon,t}$ is the 12×1 vector of residuals from the regression of $\hat{u}_t$ on $\Upsilon'\hat{x}_t$ and $\hat{X}_{\Upsilon,t}$ is the 13×1 vector of residuals from the regression of $\hat{x}_t$ on $\Upsilon'\hat{x}_t$. In a similar way to (2.A.6), the variance/covariance matrices of these errors can be obtained and are given

show any time trend over the data sample used in this study.

in (2.A.12).

$$
\begin{aligned}
\underbrace{\hat{\Sigma}_{UU_\Upsilon}}_{12\times 12} &= \frac{1}{T}\sum_{t=1}^{T} \hat{U}_{\Upsilon,t}\hat{U}'_{\Upsilon,t} \\
\underbrace{\hat{\Sigma}_{XX_\Upsilon}}_{13\times 13} &= \frac{1}{T}\sum_{t=1}^{T} \hat{X}_{\Upsilon,t}\hat{X}'_{\Upsilon,t} \qquad\qquad (2.A.12)\\
\underbrace{\hat{\Sigma}_{UX_\Upsilon}}_{12\times 13} &= \frac{1}{T}\sum_{t=1}^{T} \hat{U}_{\Upsilon,t}\hat{X}'_{\Upsilon,t} \qquad\qquad \hat{\Sigma}_{XU_\Upsilon} = \hat{\Sigma}'_{UX_\Upsilon}
\end{aligned}
$$

Johansen (1995) shows that the likelihood ratio test that the 4 cointegrating vectors are of the form in Υ comes from solving two eigenvalue problems. The first involves finding the 9 $(= n + 1 - r_1)$ eigenvalues of[66]

$$\left(G'\hat{\Sigma}_{XX_\Upsilon}G\right)^{-1} G'\hat{\Sigma}_{XU_\Upsilon}\left(\hat{\Sigma}_{UU_\Upsilon}\right)^{-1}\hat{\Sigma}_{UX_\Upsilon}G \qquad (2.A.13)$$

where the eigenvalues are of the form $1 > \tilde{\lambda}_1 > \tilde{\lambda}_2 > \ldots > \tilde{\lambda}_9 > 0$, and where G is the matrix whose columns span the null space of Υ, i.e. the columns of G are orthogonal to Υ.[67] The second eigenvalue problem involves finding the n $(= 12)$ eigenvalues of

$$\hat{\Sigma}^{-1}_{UU_\Upsilon}\hat{\Sigma}_{UX_\Upsilon}\Upsilon\left(\Upsilon'\hat{\Sigma}_{XX_\Upsilon}\Upsilon\right)^{-1}\Upsilon'\hat{\Sigma}_{XU_\Upsilon} \qquad (2.A.14)$$

where the eigenvalues are of the form $\tilde{\rho}_1 > \ldots > \tilde{\rho}_4 > \tilde{\rho}_5 = \ldots = \tilde{\rho}_{12} = 0$. The likelihood ratio test that the 4 cointegrating vectors are of the form given in Υ is shown in Johansen (1995) to be

$$\ell = T\left\{\sum_{i=1}^{4}\log\left(1-\tilde{\rho}_i\right) + \sum_{i=1}^{3}\log\left(1-\tilde{\lambda}_i\right) - \sum_{i=1}^{7}\log\left(1-\hat{\lambda}_i\right)\right\} \qquad (2.A.15)$$

[66] r_1 is the number of known cointegrating vectors, $= 4$.

[67] From the definition of Υ in (2.A.9), it can be shown that the null space is given by

$$null\left(\Upsilon\right) = \begin{bmatrix}
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}'$$

and this is orthogonalised using the Gram Schmidt procedure to obtain G (13×9).

Table 2.A.3
Test 1: Cointegration between Bid and Ask Prices and between the Three Rates

Test statistic ℓ	degrees of freedom	1% critical value	5% critical value
42.99	24	42.98	36.42

Notes: The table shows the likelihood ratio test calculated from (2.A.15). ℓ is compared to $\chi^2_{(24)}$ critical values, where the degrees of freedom are obtained from Johansen (1995), lemma 7.1.

where $\hat{\lambda}_i$ $(i = 1, \ldots, 7)$ are the 7 largest eigenvalues from the unrestricted model, used to determine the cointegrating rank in Section 2.A.3.[68] ℓ is then compared to $\chi^2_{(24)}$ critical values and the results are reported in Table 2.A.3.[69] As can be seen in the table, the null hypothesis that 4 of the cointegrating vectors are those associated with the bid-ask spreads and with the absence of arbitrage between rates is only just rejected at the 1% level; the test statistic of 42.99 only just breaches the $\chi^2_{(24)}$ 1% critical value of 42.98.[70] This brings into doubt the null hypothesis of cointegration between rates and is at odds with the ADF tests reported in Section 2.A.4.1.[71]

2.A.4.4 Test 2: cointegration between the exchange rate and cumulative order flow

Here the likelihood methods are used to test the hypothesis that the exchange rate (log ask price) and cumulative order flow are cointegrated, while leaving the remaining 4 cointegrating vectors unrestricted. The restrictions under the null imply that the cointegrating

[68]Maximum likelihood estimates of $[\psi_0' \; \psi_1']'$ are given by $G\,[E_1 \ldots E_3]$, where $E_1 \ldots E_3$ are the three eigenvectors associated with the three largest eigenvalues $\tilde{\lambda}_1 \ldots \tilde{\lambda}_3$.

[69]The degrees of freedom are calculated using Johansen (1995), lemma 7.1 and are equal to $r_1(n+1-h)$ where r_1 is equal to the number of known cointegrating vectors under the null (= 4). The +1 term in the bracket comes from the time trend which is restricted to enter the cointegration space.

[70]Two other specifications were also tested; the first allowed the time trend to enter the model of (2.A.4) unrestrictedly, and the second did not include a time trend at all. Both, however, rejected the null of cointegration between bid and ask prices and between the three rates.

[71]Even when the diurnal patterns of the bid-ask spreads were included in the model of (2.A.4) and restricted to lie within the cointegration space, the null hypothesis was still rejected. This test was carried out in a similar way to that described above, whereby a 3×1 vector of diurnal spread patterns, P_{t-1}, was included in (2.A.4). The cointegration terms are therefore of the form

$$\alpha \left[\gamma_0 \; \gamma_P \; \gamma'\right] \begin{bmatrix} t \\ P_{t-1} \\ Y_{t-1} \end{bmatrix}$$

and a fourth set of auxilliary regressions is performed, regressing P_{t-1} on a constant and the lagged dependent variables. The residuals are stacked in $\hat{x}_t$, along with $\hat{w}_t$ and $\hat{v}_t$, and the testing precedure continues as above.

vectors are of the form given in (2.A.16).

$$\underbrace{\begin{bmatrix} \gamma_0' \\ \gamma \end{bmatrix}}_{13\times 7} = \begin{bmatrix} \underbrace{\gamma^U}_{13\times 4} & \underbrace{H_1\gamma^{ED}}_{13\times 1} & \underbrace{H_2\gamma^{ES}}_{13\times 1} & \underbrace{H_3\gamma^{SD}}_{13\times 1} \end{bmatrix}$$

$$\underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{H_1} \underbrace{\begin{bmatrix} \gamma^{ED}_{time} \\ \gamma^{ED}_{ask} \\ \gamma^{ED}_{OF} \end{bmatrix}}_{\gamma^{ED}} \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{H_2} \underbrace{\begin{bmatrix} \gamma^{ES}_{time} \\ \gamma^{ES}_{ask} \\ \gamma^{ES}_{OF} \end{bmatrix}}_{\gamma^{ES}} \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 0 & 1 \end{bmatrix}}_{H_3} \underbrace{\begin{bmatrix} \gamma^{SD}_{time} \\ \gamma^{SD}_{ask} \\ \gamma^{SD}_{OF} \end{bmatrix}}_{\gamma^{SD}} \quad (2.A.16)$$

γ^U is the 13×4 matrix containing the 4 unrestricted cointegrating vectors, while the three known cointegrating vectors are contained in $H_1\gamma^{ED}$, $H_2\gamma^{ES}$ and $H_3\gamma^{SD}$. γ^{ED} contains the 3 unknown elements in the vector which cointegrates the $USD/EUR : ask$ rate and cumulative order flow,[72] and similarly for γ^{ES} and γ^{SD}. The algorithm used to estimate γ^{ED}, γ^{ES} and γ^{SD} is that used in Johansen (1995) and is described in more detail below. As before, the constant and lagged dependent variables are concentrated out of the model, in a similar fashion to (2.A.5). The errors, $\hat{u}_t$, $\hat{w}_t$ and $\hat{v}_t$ are related by the cointegration terms, as shown in (2.A.17) and again the α^+ matrix of speed of adjustment coefficients

[72] γ^{ED}_{ask} is later normalised to unity.

is partitioned into 4 matrices corresponding to the 4 cointegrating terms in (2.A.16).

$$\begin{aligned}
\hat{u}_t &= \left[\underbrace{\alpha_0^+}_{12\times 4} \ \underbrace{\alpha_1^+}_{12\times 1} \ \underbrace{\alpha_2^+}_{12\times 1} \ \underbrace{\alpha_3^+}_{12\times 1} \right] \left[\gamma^U \ H_1\gamma^{ED} \ H_2\gamma^{ES} \ H_3\gamma^{SD}\right]' \underbrace{\left[\begin{array}{c} \hat{w}_t \\ \hat{v}_t \end{array} \right]}_{\hat{x}_t} + error_t \\
&= \alpha_0^+\gamma^{U'}\hat{x}_t + \alpha_1^+\left(H_1\gamma^{ED}\right)'\hat{x}_t + \alpha_2^+\left(H_2\gamma^{ES}\right)'\hat{x}_t + \alpha_3^+\left(H_3\gamma^{SD}\right)'\hat{x}_t + error_t
\end{aligned} \tag{2.A.17}$$

The estimation and testing procedure requires an algorithm to recursively estimate each of the cointegrating terms, γ^U, $H_1\gamma^{ED}$, $H_2\gamma^{ES}$ and $H_3\gamma^{SD}$ and is discussed in detail in Johansen (1995). Initial estimates of $H_1\gamma^{ED}$, $H_2\gamma^{ES}$ and $H_3\gamma^{SD}$ are obtained from the unrestricted estimates of γ ($\hat{\gamma}$),[73] and then these are concentrated out of the model, allowing γ^U to be estimated by reduced rank regression. This, together with the initial estimates of $H_2\gamma^{ES}$ and $H_3\gamma^{SD}$, is used to estimate γ^{ED} in a similar fashion. By repeating this algorithm a sufficient number of times, maximum likelihood estimates of γ^U, $H_1\gamma^{ED}$, $H_2\gamma^{ES}$ and $H_3\gamma^{SD}$ can be obtained, together with the eigenvalues to perform the likelihood ratio test.[74]

The results of the test that 3 of the 7 cointegrating vectors are associated with the level of each exchange rate and cumulative order flow are presented in Table 2.A.4. The null hypothesis is clearly rejected, with a test statistic of 555.17 being much greater than the $\chi^2_{(12)}$ 1% critical value of 26.22.[75]

2.A.4.5 Test 3: 7 theoretical cointegrating vectors tested jointly

The test that the 7 cointegrating vectors found in Section 2.A.3 are associated with the 3 bid-ask spreads, absence of triangular arbitrage and the 3 relationships between the level of the exchange rate and cumulative order flow, is carried out in a similar way to Test 2. The restrictions under the null hypothesis imply that the cointegrating vectors look as

[73] Since the ordering of the vectors in the unrestricted estimates of γ need not correspond to the ordering presented in (2.A.17), the initial estimates are obtained from finding combinations of the vectors in $\hat{\gamma}$ that best span H_1, and similarly for H_2 and H_3.

[74] 20 iterations were sufficient to achieve convergence.

[75] Again, the degrees of freedom are calculated from Johansen (1995), lemma 7.1.

Table 2.A.4
Test 2: Cointegration between the Level of the Exchange Rate and Cumulative Order Flow

Test statistic ℓ	degrees of freedom	1% critical value	5% critical value
555.17	12	26.22	21.03

γ^{ED}_{time}	$0.0^{5}166$	γ^{ES}_{time}	$-0.0^{5}201$	γ^{SD}_{time}	$0.0^{5}192$
γ^{ED}_{ask}	1	γ^{ES}_{ask}	1	γ^{SD}_{ask}	1
γ^{ED}_{OF}	$-0.0^{4}268$	γ^{ES}_{OF}	$0.0^{5}901$	γ^{SD}_{OF}	$-0.0^{4}133$

Notes: The table shows the likelihood ratio test for the null hypothesis that the level of each exchange rate is cointegrated with cumulative order flow. ℓ is compared to $\chi^2_{(12)}$ critical values, where the degrees of freedom are obtained from Johansen (1995), lemma 7.1. The estimates of the γ^{ED}, γ^{ES} and γ^{SD} parameter vectors are also presented, where the γ_{ask} coefficients have been normalised to unity.

follows.

$$\underbrace{\begin{bmatrix} \gamma_0' \\ \gamma \end{bmatrix}}_{13\times7} = \begin{bmatrix} \underbrace{\Upsilon}_{13\times4} & \underbrace{H_1\gamma^{ED}}_{13\times1} & \underbrace{H_2\gamma^{ES}}_{13\times1} & \underbrace{H_3\gamma^{SD}}_{13\times1} \end{bmatrix} \tag{2.A.18}$$

These restrictions are the same as those presented in (2.A.16, with the exception that Υ replaces γ^U. Υ is the 13×4 matrix given in (2.A.9), which imposes the restrictions that the bid and ask prices are cointegrated and that the three rates are cointegrated. After concentrating out the constant and lagged dependent variables, a similar expression to (2.A.17) is obtained. This is given in (2.A.19) below.

$$\hat{u}_t = \alpha_0^+ \Upsilon' \hat{x}_t + \alpha_1^+ \left(H_1\gamma^{ED}\right)' \hat{x}_t + \alpha_2^+ \left(H_2\gamma^{ES}\right)' \hat{x}_t + \alpha_3^+ \left(H_3\gamma^{SD}\right)' \hat{x}_t + error_t \tag{2.A.19}$$

α_0^+ is then concentrated out of the model by regressing $\hat{u}_t$ and $\hat{x}_t$ on $\Upsilon'\hat{x}_t$. The residuals from these regressions are related by the 3 cointegrating vectors between the level of the (log ask) exchange rate and cumulative order flow.

$$\begin{aligned}\hat{u}_t &= \tilde{\alpha}_0^+ \Upsilon' \hat{x}_t + \hat{U}_{\Upsilon,t} \\ \hat{x}_t &= \hat{\alpha}_0^+ \Upsilon' \hat{x}_t + \hat{X}_{\Upsilon,t} \\ \Rightarrow \hat{U}_{\Upsilon,t} &= \alpha_1^+ \left(H_1 \gamma^{ED}\right)' \hat{X}_{\Upsilon,t} + \alpha_2^+ \left(H_2 \gamma^{ES}\right)' \hat{X}_{\Upsilon,t} + \alpha_3^+ \left(H_3 \gamma^{SD}\right)' \hat{X}_{\Upsilon,t} + \tilde{error}_t\end{aligned} \tag{2.A.20}$$

A similar switching algorithm to that used in Test 2 is then used, whereby $H_1\gamma^{ED}$ is estimated given the initial estimates of $H_2\gamma^{ES}$ and $H_3\gamma^{SD}$. Each of the cointegrating vectors are estimated in turn in a recursive manner until convergence is achieved. The likelihood ratio test that the 7 cointegrating vectors are those associated with the bid-ask spreads, absence of triangular arbitrage and prices and cumulative order flows, is presented in Table 2.A.5. Unsurprisingly, given the rejection of the hypothesis that prices and cumulative order flows are cointegrated in Test 2, the hypothesis that the system is cointegrated by the 7 theoretical relationships, is also rejected. The likelihood ratio test statistic is 503.48, clearly breaching the 1% critical value of 58.62 from a $\chi^2_{(36)}$ distribution.[76]

The cointegration tests therefore give mixed evidence on the form of the 7 cointegrating vectors. ADF tests suggest that (log) bid and ask prices are cointegrated and that the three exchange rates are cointegrated, as one would expect. However, using system methods of cointegration leads one to reject these hypotheses, although the 1% critical values are only just breached. Tests also suggest that the level of the exchange rate (log ask price) and cumulative order flow are not cointegrated. Due to these mixed results, I estimate a number of different specifications of the VAR in (2.A.2), each placing a different set of restrictions on the cointegration space.

- The first model imposes no restrictions on the cointegration space at all, i.e. all 7 cointegrating vectors are estimated freely. However, the time trend is only allowed

[76] Again, the degrees of freedom are found using Johansen (1995), lemma 7.1. To explain where these degrees of freedom come from, consider (2.A.19), rewritten below.

$$\hat{u}_t = \underbrace{\alpha_0^+}_{12\times4} \Upsilon' \hat{x}_t + \underbrace{\alpha_1^+}_{12\times1} \underbrace{\gamma^{ED'}}_{1\times3} H_1' \hat{x}_t + \underbrace{\alpha_2^+}_{12\times1} \underbrace{\gamma^{ES'}}_{1\times3} H_2' \hat{x}_t + \underbrace{\alpha_3^+}_{12\times1} \underbrace{\gamma^{SD'}}_{1\times3} H_3' \hat{x}_t + error_t$$

Using the result that the dimension of the product of two matrices, M and N (where M is $a \times b$ and N is $b \times c$, $a > b$, $c > b$) is equal to $ac - (a-b)(c-b)$, then the dimension of $\alpha_1^+\gamma^{ED'}$, for example, is equal to $12 \times 3 - (12-1)(3-1) = 14$. The dimension of the above model is therefore $12 \times 4 + 3(14) = 90$. The unrestricted model, with cointegrating terms $\alpha^+\gamma'$, has dimension equal to $12\times13-(12-7)(13-7) = 126$. The degrees of freedom is the difference between the two, giving 36. Notice that the 36 degrees of freedom used here is equal to the sum of the degrees of freedom used in Tests 1 and 2 (24 and 12). This follows since Test 3 is simply the joint test of Tests 1 and 2.

Table 2.A.5
Test 3: Joint Test of the 7 Theoretical Cointegrating Vectors

Test statistic ℓ	degrees of freedom	1% critical value	5% critical value
503.48	36	58.62	51.00

γ^{ED}_{time}	0.0^6265	γ^{ES}_{time}	-0.0^6710	γ^{SD}_{time}	0.0^5105
γ^{ED}_{ask}	1	γ^{ES}_{ask}	1	γ^{SD}_{ask}	1
γ^{ED}_{OF}	0.0^5102	γ^{ES}_{OF}	-0.0^6373	γ^{SD}_{OF}	0.0^6456

Notes: The table shows the likelihood ratio test for the null hypothesis that the 7 cointegrating vectors are those associated with the bid-ask spreads, the absence of triangular arbitrage and the level of each exchange rate and cumulative order flow. ℓ is compared to $\chi^2_{(36)}$ critical values, where the degrees of freedom are obtained from Johansen (1995), lemma 7.1. The estimates of the γ^{ED}, γ^{ES} and γ^{SD} parameter vectors are also presented, where the γ_{ask} coefficients have been normalised to unity.

to enter the cointegration space, in order to rule out quadratic trends.

- The second model imposes the restriction that bid and ask prices are cointegrated and that the three rates are themselves cointegrated. The three remaining cointegrating vectors are estimated freely. The cointegration terms are therefore of the form given in (2.A.9).

- The third model imposes all 7 theoretical cointegrating restrictions, i.e. cointegration is imposed via the bid-ask spreads, absence of triangular arbitrage and also via the relationships between the levels of the exchange rates and cumulative order flow.

- The final model assumes only 4 cointegrating vectors and that these relationships are between bid and ask prices and between the three rates.

By comparing different model specifications, we can see how important the cointegration terms are. The fourth specification, which only allows 4 cointegrating relationships, is included simply to see how important the 5th, 6th and 7th cointegrating vectors are. The 4 relationships between bid and ask prices and between the three rates are intuitive and should cointegrate the system in a well-functioning and liquid market. Any differences (in the impulse response functions, etc.) that occur when a further 3 cointegrating vectors are allowed, will show the relative importance of these additional relationships.

However, in specifications 2 to 4, further structure was imposed on the VAR in (2.A.2) by taking into consideration the intra-day patterns of spreads. If the bid and ask prices in the USD/EUR market are cointegrated then only if $ED^{ask}_{t-1} - ED^{bid}_{t-1} > \psi^{EDspread}_{t-1}$ will the ask price tend to fall and the bid price rise in period t, where $\psi^{EDspread}_{t-1}$ is the intra-day pattern of the spread at time $t-1$ calculated from the FFF in Section 2.3.3.

2.A.5 Impulse response analysis

In order to examine the dynamics of vector autoregressive models, such as (2.A.2), it is common in the literature to use impulse response functions (IRFs) and variance decompositions. By introducing a shock to one of the elements of ϵ_t in (2.A.2), one can observe how Y_t responds, and through the lags in the model, how Y_{t+1}, $Y_{t+2}, \ldots,$ are affected. As explained in Lütkepohl and Reimers (1992), when a cointegrating VAR is being considered, the impulse responses are most easily computed from the VAR in levels. From the reduced form VAR of (2.A.4), it is clear that the VAR representation in levels is as given in (2.A.21).

$$
\begin{aligned}
Y_t = c^+ + \alpha^+\gamma_0 t + \phi_1 Y_{t-1} + \phi_2 Y_{t-2} + \ldots + \phi_{m+1} Y_{t-m-1} + \underbrace{B^{-1}\epsilon_t}_{\epsilon_t^+} \\
\textit{where} \quad \phi_{m+1} &= -\theta^+_m \qquad\qquad (2.A.21) \\
\phi_i &= -\theta^+_{i-1} + \theta^+_i \qquad i = 2, \ldots, m \\
\phi_1 &= I_{12} + \alpha^+\gamma' + \theta^+_1
\end{aligned}
$$

The levels VAR is still consistent with any of the cointegrating restrictions in Section 2.A.2 that one may wish to impose, since the coefficient on the time trend, and on Y_{t-1} depend on the restricted cointegrating vectors contained in $[\gamma_0 \ \gamma']'$. Since the variance/covariance matrix of ϵ, Ω, is only block diagonal (see (2.A.2)), introducing a shock to one element of ϵ_t will not represent a 'pure' shock to the corresponding element in Y_t. Therefore, orthogonalised impulse responses are preferred since this is the easiest way to analyse how 'new information' entering via USD/EUR : *buys*, for example, feeds through the system, while not being correlated with any new information from any of the other variables. Since Ω is block diagonal, one need only perform the Choleski decomposition on Ω_{11} and Ω_{22}. Therefore, let P_1 and P_2 be lower triangular 6×6 matrices, where $P_1 P_1' = \Omega_{11}$ and

$P_2P_2' = \Omega_{22}$. Then $Var\left[\epsilon_t^+\right]$ is given by

$$Var\left[\epsilon_t^+\right] = \Sigma = B^{-1}\Omega\left(B^{-1}\right)' = \begin{bmatrix} I_6 & \beta \\ 0 & I_6 \end{bmatrix} \begin{bmatrix} P_1P_1' & 0 \\ 0 & P_2P_2' \end{bmatrix} \begin{bmatrix} I_6 & 0 \\ \beta' & I_6 \end{bmatrix}$$
$$= \underbrace{\begin{bmatrix} I_6 & \beta \\ 0 & I_6 \end{bmatrix}}_{B^{-1}} \underbrace{\begin{bmatrix} P_1 & 0 \\ 0 & P_2 \end{bmatrix}}_{P} \underbrace{\begin{bmatrix} P_1' & 0 \\ 0 & P_2' \end{bmatrix}}_{P'} \underbrace{\begin{bmatrix} I_6 & 0 \\ \beta' & I_6 \end{bmatrix}}_{(B^{-1})'} \tag{2.A.22}$$

Define the vector v_t by $v_t = P^{-1}B\epsilon_t^+$, then the variance of v_t is given by

$$\begin{aligned} Var\left[v_t^+\right] &= P^{-1}BVar\left[\epsilon_t^+\right]B'\left(P^{-1}\right)' \\ &= P^{-1}BB^{-1}PP'\left(B^{-1}\right)'B'\left(P^{-1}\right)' \\ &= I_{12} \end{aligned} \tag{2.A.23}$$

The errors contained in v_t are therefore orthogonal and since ϵ_t^+ in (2.A.21) can be written as $B^{-1}Pv_t$, then orthogonalised impulse responses can be obtained by shocking v_t.[77] The impulse response functions following a shock to the 'buy' variables for each of the three markets are presented in Figures 2.A.1 to 2.A.6. Figures 2.A.1 to 2.A.3 show the effects on the level of each exchange rate and on the cumulative number of buys and sells, over and above what one would expect, while Figures 2.A.4 to 2.A.6 show the effects on the level of the spread and on trading volume, again over and above what one would expect. In every panel of each figure, a number of IRFs are plotted, associated with the different cointegrating specifications discussed in Section 2.A.4. The solid blue lines show the responses in the model where the 7 cointegrating vectors are those associated with the theoretical restrictions in Section 2.A.2. The dotted blue lines trace out the 95% confidence interval for this model, found by bootstrapping over 1000 iterations, see Section 2.A.7. The solid black lines show the responses when all of the cointegrating vectors are estimated freely. The red lines correspond to the model where 4 of the cointegrating vectors are associated with the bid-ask spreads and absence of arbitrage, and the green

[77]This is slightly different to the standard case, since ϵ_t^+ is already 'half way' to being orthogonalised before the Choleski decomposition is performed; $\epsilon_t^+ = B^{-1}\epsilon_t$, where B^{-1}, the matrix of structural parameters, is already a (block) triangular matrix. The variance matrix of ϵ_t is restricted to be block diagonal in order for (2.A.2) to be identified and this means only the sub-matrices of Ω need to be orthogonalised by the Choleski method.

Figure 2.A.1
Impulse Response Functions Following a USD/EUR 'buy' Shock

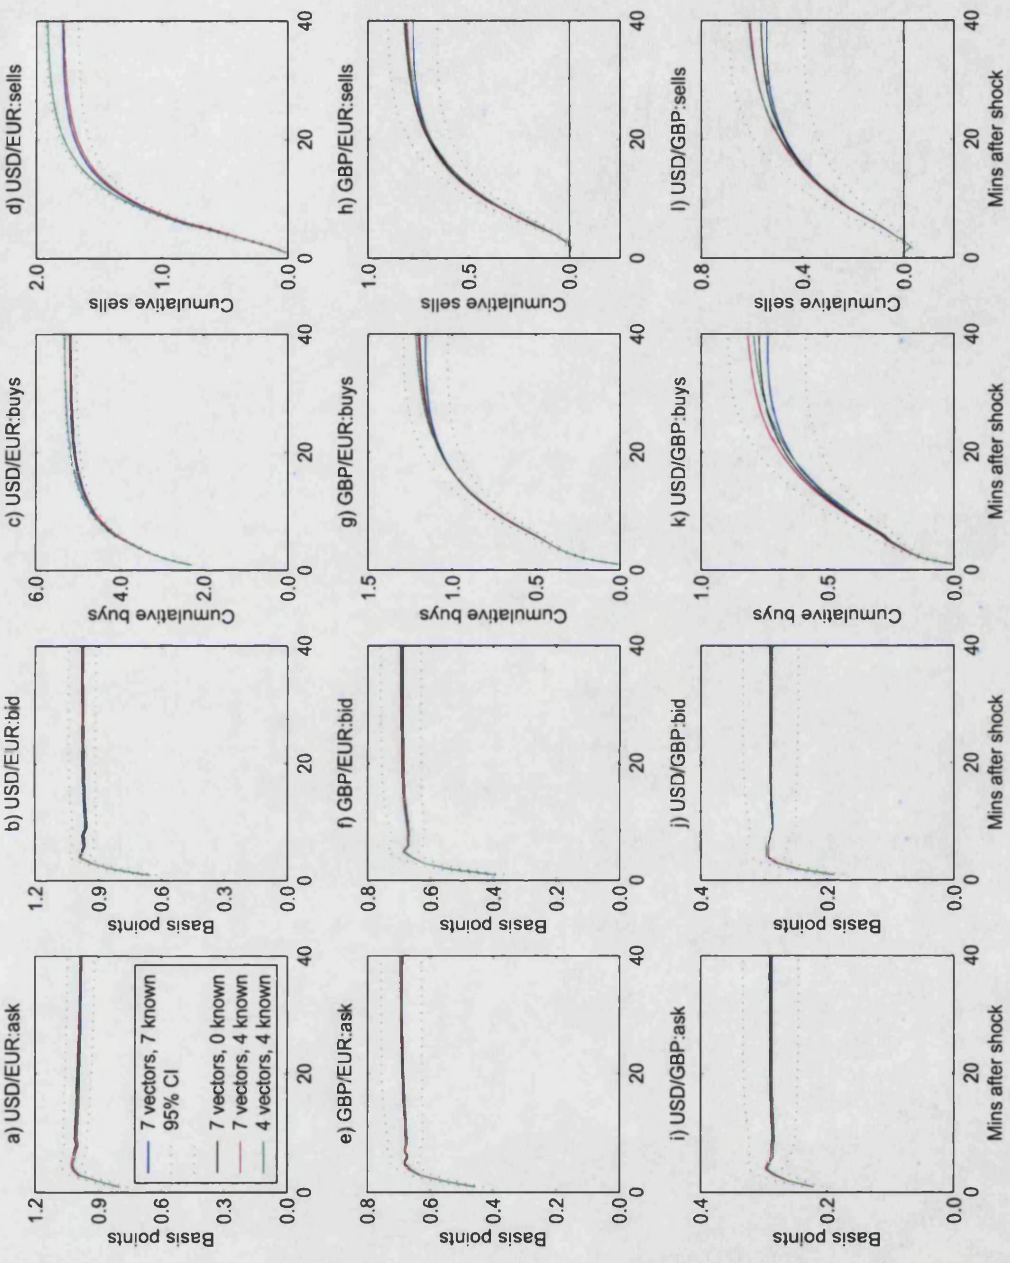

Notes: The figures plot the impulse response functions following a one standard deviation USD/EUR : *buy* shock. The IRFs for 4 model specifications are given. The solid blue lines show the responses in the model where the 7 cointegrating vectors are those associated with the theoretical restrictions in Section 2.A.2. The dotted blue lines trace out the 95% confidence interval for this model, found by bootstrapping over 1000 iterations. The solid black lines show the responses when all of the cointegrating vectors are estimated freely. The red lines correspond to the model where 4 of the cointegrating vectors are associated with the bid-ask spreads and absence of arbitrage, and the green lines correspond to the model where only 4 cointegrating vectors are allowed, associated with bid-ask spreads and absence of arbitrage.

Figure 2.A.2
Impulse Response Functions Following a GBP/EUR 'Buy' Shock

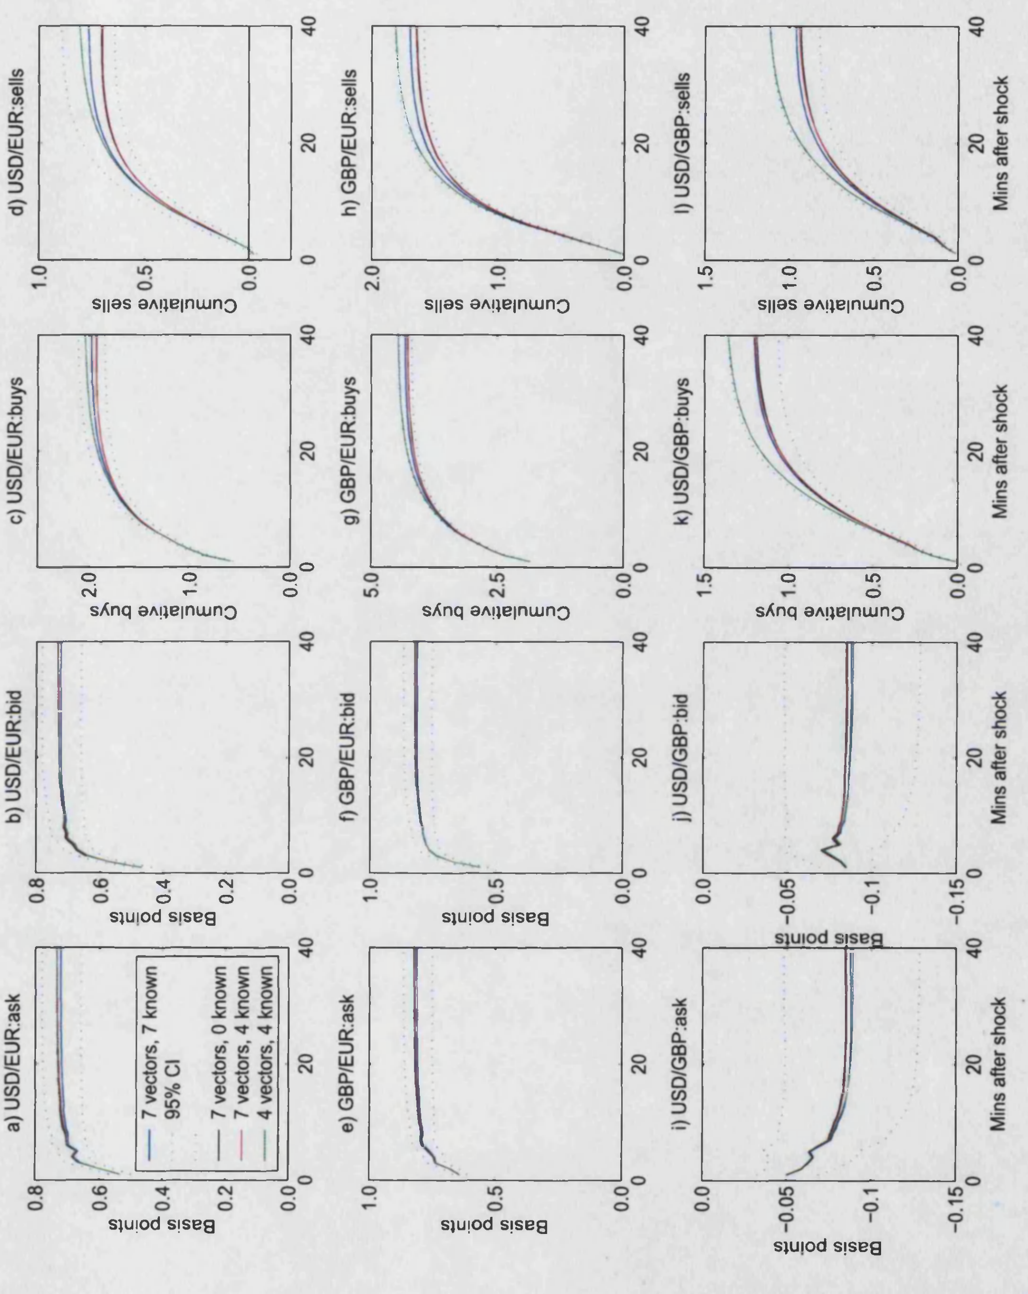

Notes: The figures plot the impulse response functions following a one standard deviation $GBP/EUR:$ *buy* shock. The IRFs for 4 model specifications are given. The solid blue lines show the responses in the model where the 7 cointegrating vectors are those associated with the theoretical restrictions in Section 2.A.2. The dotted blue lines trace out the 95% confidence interval for this model, found by bootstrapping over 1000 iterations. The solid black lines show the responses when all of the cointegrating vectors are estimated freely. The red lines correspond to the model where 4 of the cointegrating vectors are associated with the bid-ask spreads and absence of arbitrage, and the green lines correspond to the model where only 4 cointegrating vectors are allowed, associated with bid-ask spreads and absence of arbitrage.

Figure 2.A.3
Impulse Response Functions Following a USD/GBP 'Buy' Shock

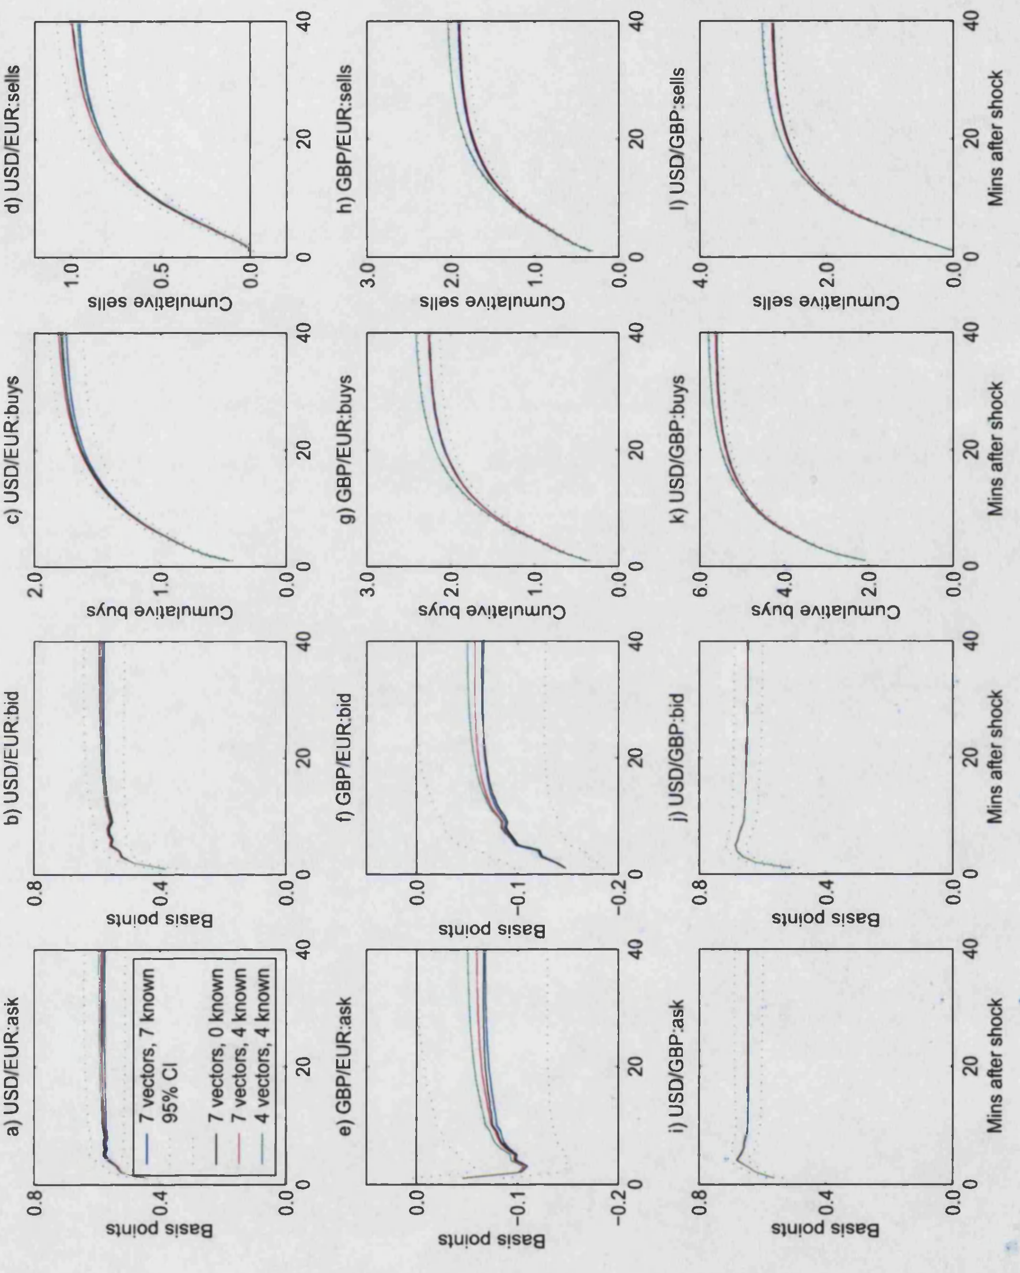

Notes: The figures plot the impulse response functions following a one standard deviation $USD/GBP:$ *buy* shock. The IRFs for 4 model specifications are given. The solid blue lines show the responses in the model where the 7 cointegrating vectors are those associated with the theoretical restrictions in Section 2.A.2. The dotted blue lines trace out the 95% confidence interval for this model, found by bootstrapping over 1000 iterations. The solid black lines show the responses when all of the cointegrating vectors are

Figure 2.A.4
Spread and Volume Impulse Response Functions Following a USD/EUR 'Buy' Shock

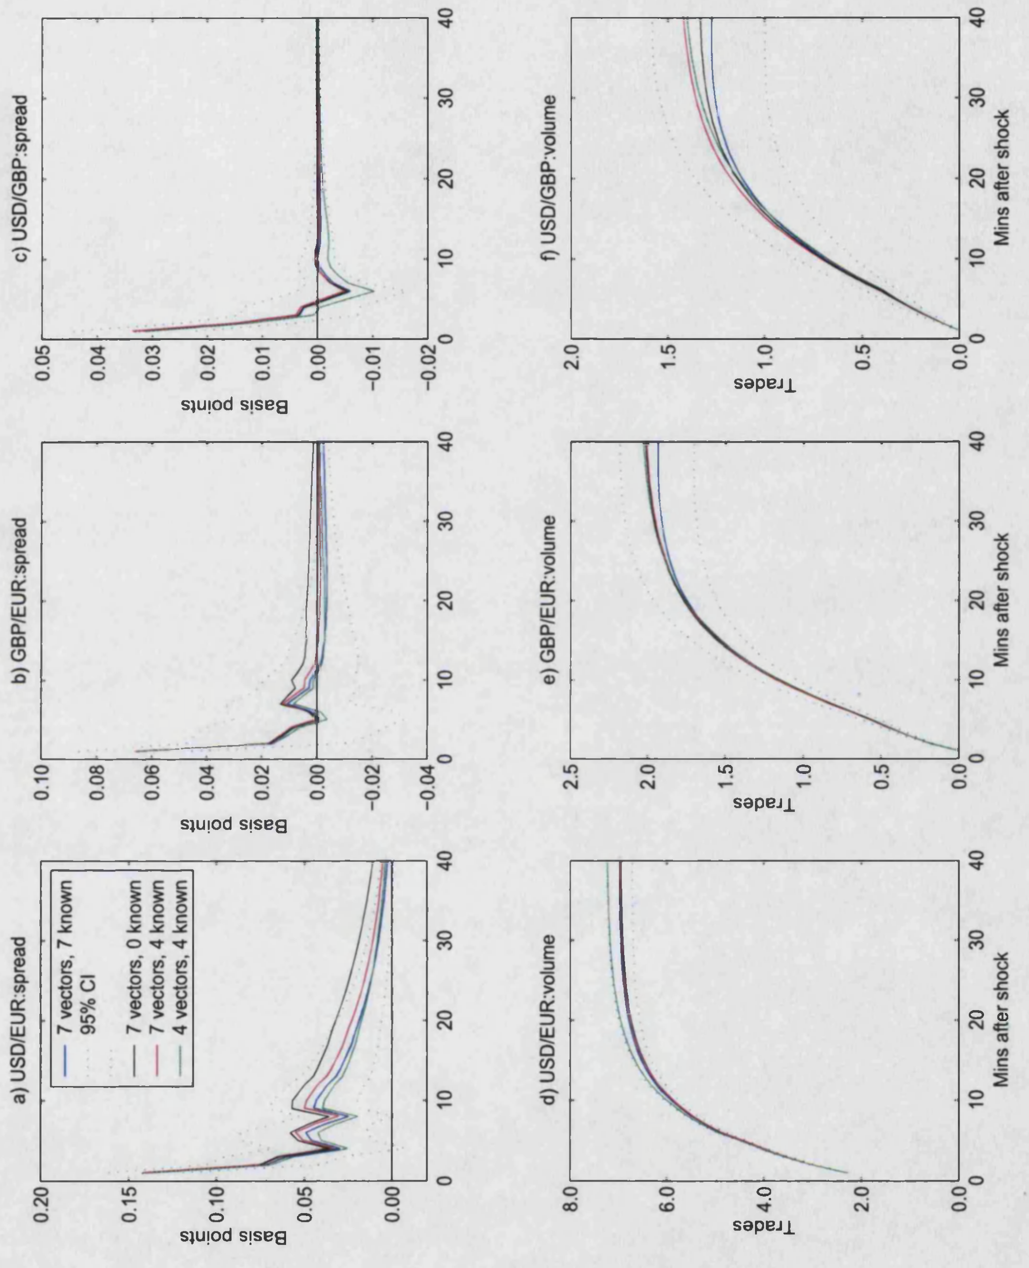

Notes: The figures plot the impulse response functions following a one standard deviation USD/EUR : *buy* shock. The IRFs for 4 model specifications are given. The solid blue lines show the responses in the model where the 7 cointegrating vectors are those associated with the theoretical restrictions in Section 2.A.2. The dotted blue lines trace out the 95% confidence interval for this model, found by bootstrapping over 1000 iterations. The solid black lines show the responses when all of the cointegrating vectors are estimated freely. The red lines correspond to the model where 4 of the cointegrating vectors are associated with the bid-ask spreads and absence of arbitrage, and the green lines correspond to the model where only 4 cointegrating vectors are allowed, associated with bid-ask spreads and absence of arbitrage.

Figure 2.A.5
Spread and Volume Impulse Response Functions Following a GBP/EUR 'Buy' Shock

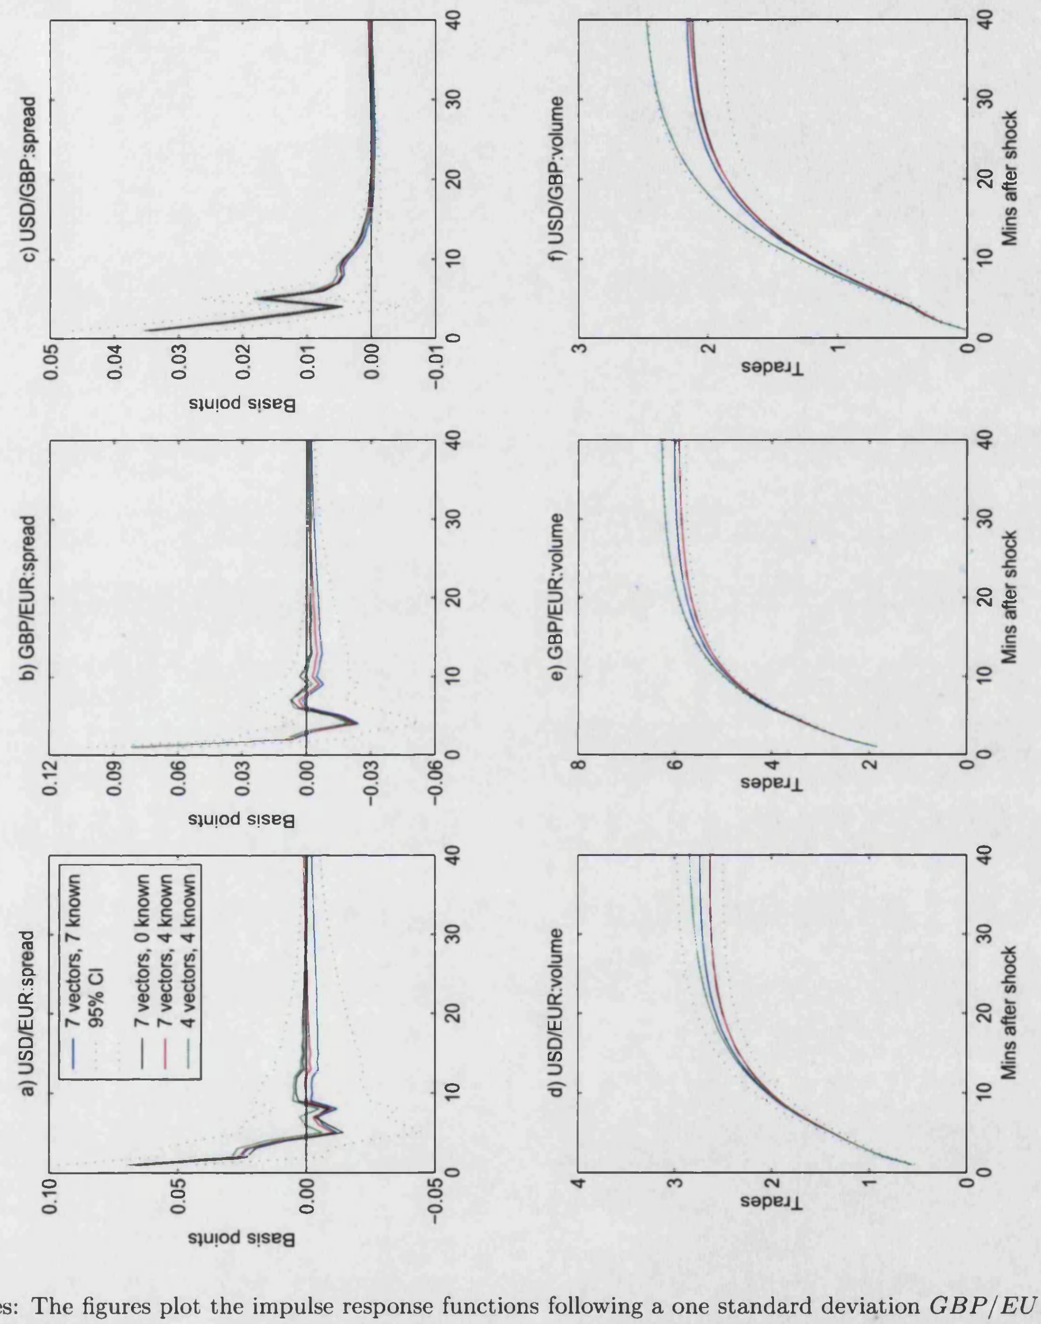

Notes: The figures plot the impulse response functions following a one standard deviation GBP/EUR : *buy* shock. The IRFs for 4 model specifications are given. The solid blue lines show the responses in the model where the 7 cointegrating vectors are those associated with the theoretical restrictions in Section 2.A.2. The dotted blue lines trace out the 95% confidence interval for this model, found by bootstrapping over 1000 iterations. The solid black lines show the responses when all of the cointegrating vectors are estimated freely. The red lines correspond to the model where 4 of the cointegrating vectors are associated with the bid-ask spreads and absence of arbitrage, and the green lines correspond to the model where only 4 cointegrating vectors are allowed, associated with bid-ask spreads and absence of arbitrage.

Figure 2.A.6
Spread and Volume Impulse Response Functions Following a USD/GBP 'Buy' Shock

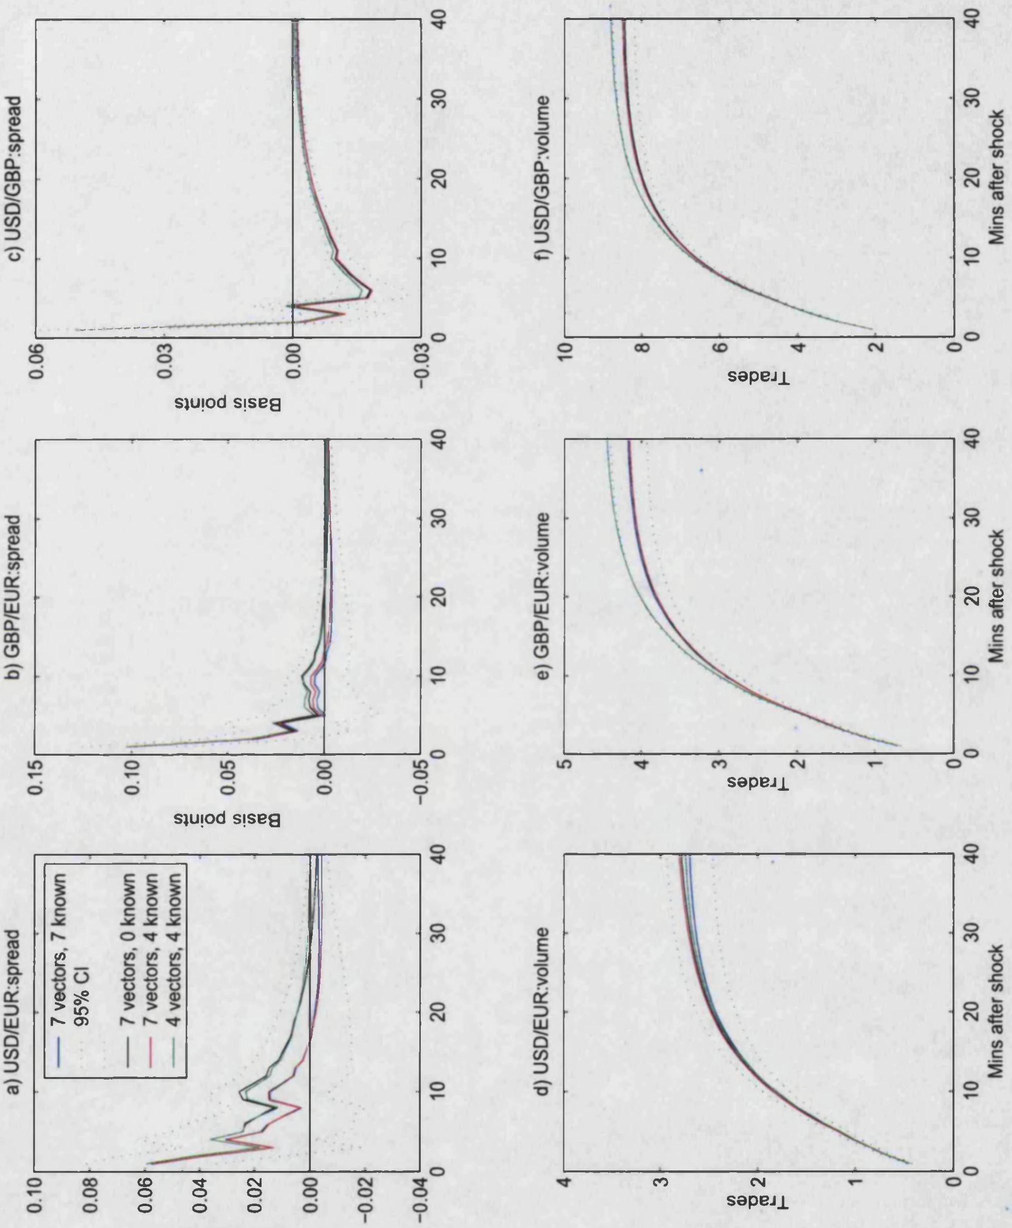

Notes: The figures plot the impulse response functions following a one standard deviation USD/GBP : *buy* shock. The IRFs for 4 model specifications are given. The solid blue lines show the responses in the model where the 7 cointegrating vectors are those associated with the theoretical restrictions in Section 2.A.2. The dotted blue lines trace out the 95% confidence interval for this model, found by bootstrapping over 1000 iterations. The solid black lines show the responses when all of the cointegrating vectors are estimated freely. The red lines correspond to the model where 4 of the cointegrating vectors are associated with the bid-ask spreads and absence of arbitrage, and the green lines correspond to the model where only 4 cointegrating vectors are allowed, associated with bid-ask spreads and absence of arbitrage.

lines correspond to the model where only 4 cointegrating vectors are allowed, associated with bid-ask spreads and absence of arbitrage.[78]

As can be seen in all of the figures, whether one imposes the cointegrating vectors associated with bid/ask prices and absence of arbitrage, or the vectors associated with the level of each exchange rate and cumulative order flow, the impulse response functions are very similar. All IRFs lie within the 95% confidence interval of that associated with all 7 theoretical cointegrating vectors. Therefore, even though the data reject the hypothesis that the level of each exchange rate and cumulative order flow are cointegrated, making this assumption has no significant effect on the IRFs. The IRF analysis used in Section 2.4 is therefore unlikely to be affected by the imposition of the 7 theoretical cointegrating vectors.[79]

2.A.6 Variance decompositions

Variance decompositions, showing the share of the mean square error of n step ahead exchange rate forecasts attributable to trading in each of the three markets can also be computed. Following Hasbrouck (1991b) this can be used to give an idea as to the information content of the trading process; the greater the variance of the exchange rate attributed to the trading process, the more information trades are argued to carry. Since the model in (2.2) allows cross-market effects, then it is possible to attribute the MSE of n step ahead exchange rate forecasts into trading in the different markets. This is shown in Figure 2.A.7. The top three panels show how the share of the forecast MSEs are split between both buys and sells in each market. The middle three panels show the share of the forecast MSEs attributed to trading (i.e. buys *and* sells) and are the same as those presented in Figure 2.5, while the bottom three panels show shares attributed to all trades.[80]

The information contained in trading that is relevant for the ask prices of each exchange rate, as defined by these variance decompositions, is large, between 30 and 40% (bottom three panels). However, the information contained in trading from other currency

[78]For each specification, the VAR was estimated by first estimating the (restricted) cointegrating vectors and then estimating the remaining parameters as described in Section 2.3.6.

[79]Impulse rsponse functions were also performed when the news variables were included as right hand side regressors. However, the IRFs following any particular news shock in all four sets of cointegrating assumptions were very similar to one another. Therefore, whether one imposes the theoretical cointegrating vectors or not in the system of (2.2), the results of this chapter are not affected.

[80]95% confidence bounds have not been plotted in the top three panels as this complicates the already 'busy' figures.

Figure 2.A.7
Variance Decompositions, Showing the Share of the Mean Square Error of n Step Ahead Forecasts Attributable to Trading in each of the Three Markets

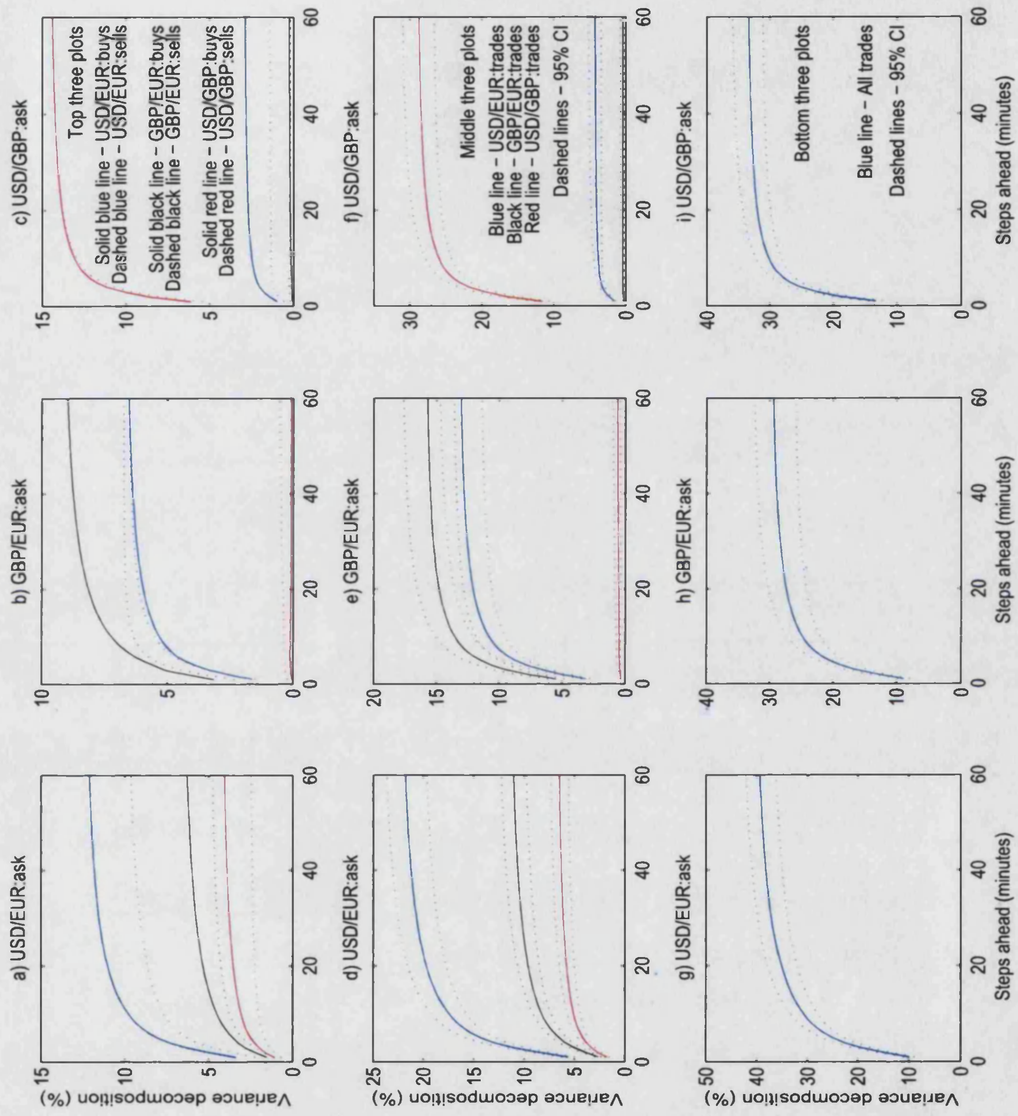

Notes: The figures plot the share of the mean square error of n step ahead forecasts attributable to trading in each of the three markets. The top three panels show how the share of the forecast MSEs are split between both buys and sells in each market. The middle three panels show the share of the forecast MSEs attributed to trading (i.e. buys *and* sells) in each market, while the bottom three panels show shares attributed to all trades. 95% confidence intervals, where shown, were found by bootstrapping model (2.2) over 1000 iterations.

pairs is significant, as described in Section 2.4.2. At the 60 step ahead forecast for the $USD/EUR : ask$ rate, just under a quarter of the MSE is attributed to USD/EUR trading (both buys *and* sells), see Figure 2.A.7d. However, approximately 11% and 7% of the variance of the USD/EUR ask rate is attributed to trading in GBP/EUR and USD/GBP respectively, implying significant spillovers from one market to another.

2.A.7 Bootstrapping

In Chapters 2 and 3, confidence bounds are given for a number of impulse response functions, such as those reported in Figures 2.7 to 2.10, and for the variance decompositions reported in Figure 2.5. These are calculated using the bootstrap, a numerical method that has the advantage of not requiring an assumption on the distributional form of the residuals. In all cases, the bootstrap was performed as follows:

1. The model, such as (2.2), was estimated and the model parameters and fitted residuals were saved.

2. Taking the first m observations as given, where m is the number of lags in the VAR, a single $n \times 1$ vector of residuals from the remaining $T - m$ vectors of residuals, all of which are $n \times 1$, is used to construct the first observation of an artificial sample of ΔY_t. n is the number of equations in the model, 12 in the case of (2.2), and each of these elements is taken from the same time observation.

3. This is repeated until a full artificial sample of ΔY_ts is created. Note that the $n \times 1$ vector of fitted residuals is drawn with replacement, so there is a $\frac{1}{T-m}$ probability that the kth $n \times 1$ vector of residuals is the same as the $k - 1$th $n \times 1$ vector of residuals.

4. With this new artificial sample, the model, such as (2.2), is estimated again and the relevant impulse response functions and variance decompositions are calculated. These are set to one side.

5. Steps 2. to 4. are repeated until a large number of impulse responses and variance decompositions have been obtained; 1000 times in this thesis. 95% confidence bounds for these functions can be computed from the range that contains 95% of the IRFs or variance decompositions.

The bootstrap was introduced by Efron (1979) and for a more detailed discussion, see Efron and Tibshirani (1993). However, in this thesis, the iid bootstrap is used, in that the $n \times 1$ vector of residuals is drawn with replacement in steps 2. and 3. This assumes that the vectors of errors are iid through time, which may be incorrect if there is serial correlation in the residuals or if the residuals are heteroscedastic. Although this is a deficiency, I do not believe it will have a significant impact on the results or conclusions in these chapters.

Chapter 3 Macroeconomic News, Order Flows and Exchange Rates

3.1 Introduction

Traditional asset market models of exchange rate determination, based on rational expectations and efficient markets, imply that announcements of public information are directly impounded in prices with there being no role for trades in this process of information assimilation.[1] More recent exchange rate analysis, based on microstructure considerations, stresses the role that trading plays in price formation via a concept called *order flow*. Order flow is defined to be the difference between buyer-initiated and seller-initiated trading interest in a given market and thus corresponds broadly to what practitioners might describe as aggressive buying or selling pressure. In the models of Lyons (1995), Perraudin and Vitale (1996) and Evans and Lyons (2002b) order flow explains contemporaneous exchange rate movements because it contains information, either about fundamentals or long-run risk premia, that was previously dispersed among market participants. Thus, one of the key differences between the microstructure level analysis and traditional exchange rate frameworks is that the same information is not shared by all market participants and/or is interpreted differently by participants.

This chapter seeks to test the hypothesis that public information announcements alter exchange rates with no role for order flow. The test of this is direct; using 10 months of transaction-level exchange rate information on USD/EUR (dollars per euro), GBP/EUR (pounds per euro) and USD/GBP (dollars per pound) and data on euro-area, UK and US macroeconomic announcements, I examine whether announcement surprises have a systematic and significant effect on both order flow and prices. I also decompose the price reactions to announcements into a part that is direct and a part intermediated by order

[1] This chapter is a revised version of Love and Payne (2003), which has been submitted to the Journal of Financial and Quantitative Analysis.

flow. Therefore, as opposed to Chapter 2, where I examine *what* happens around periods of scheduled macroeconomic news, this chapter asks *how* this information is incorporated into asset prices.

The results are unambiguous. At a 1 minute sampling frequency, macroeconomic information releases do have systematic effects on order flow and, as established in previous studies, on exchange rate transaction prices. After releases of "good news", not only does the exchange rate tend to appreciate but order flow tends to be positive, reflecting an excess of agents aggressively buying over agents aggressively selling. Moreover, I show that in periods just after macroeconomic announcements, the significance of order flow in exchange rate determination is much greater than in normal times. Finally, I estimate a multivariate VAR model in transaction price changes and order flow with a signed macroeconomic information surprise variable included as an exogenous right-hand side regressor. Via this model one is able to characterise the extent to which the final effect of macroeconomic information on prices is intermediated by order flow. The results suggest that nearly two thirds of the final price reaction to news comes via this order flow mechanism.

Thus, the analysis corroborates not only earlier results on the relevance of both macroeconomic news (Almeida, Goodhart, and Payne 1998, Andersen, Bollerslev, Diebold, and Vega 2003) and order flows (Lyons 1995, Yao 1998, Payne 2003a) for high frequency exchange rate determination but it also ties in with other recent work, Evans and Lyons (2003), which indicates that at least part of the response of exchange rates to news comes via order flow. These results are based on a four month sample of exchange rates and order flows and a news variable that is a count of the number of headline news items released that day. Methodologically, due to the inability to use their flow of news headlines to accurately construct a time-series of signed or sized news releases, Evans and Lyons (2003) use a variance decomposition to identify the role that order flow has to play in the assimilation of news into prices. Thus, their paper provides no *direct* evidence on the effects of news releases on order flows or exchange rates.

The key result of this chapter is that even macroeconomic information that is publicly and simultaneously released to all market participants is largely impounded into prices via the key micro-level price determinant — order flow. This is clearly at odds with rational expectations, efficient market models of asset price determination. However, despite the role for order flow in the assimilation of public information into prices, I do not suggest that FX markets are not efficient. Indeed I find that the average effects of news are impounded quickly; virtually all of the price changes associated with public/macro news announcements occur within the first two minutes of release.[2] Within the context of exchange rate determination the results suggest that the recent distinctions drawn between macroeconomic and microstructure models are not clear cut; the modelling of exchange rates should incorporate both elements of macro and microstructure. Further effort needs to be expended on theoretical and empirical work to merge the two sides of exchange rate determination in an attempt to explain more accurately how exchange rates are determined.

The rest of the chapter is set up as follows. Section 3.2 briefly describes the data, Section 3.3 introduces the univariate and multivariate models, Section 3.4 discusses the results and Section 3.5 concludes.

3.2 Data

The exchange rate dataset used in the paper comes from the brokered segment of the inter-dealer FX market, specifically from the Reuters D2000-2 system and is described in detail in Section 2.2.1. Rather than concentrating on bid and ask prices, as was done in Chapter 2, this chapter focuses on the transaction price series of the three rates (USD/EUR, GBP/EUR and USD/GBP). As explained in Chapter 2, these three exchange rates form a triangular set, linked by an obvious absence of arbitrage condition, which is exploited in the empirical analysis of Section 3.3.3. The samples for USD/EUR and USD/GBP cover a period of ten months from 28th September 1999 to 24th July 2000, while the GBP/EUR sample is somewhat shorter, covering the eight month period from 1st December 1999 to 24th July 2000.

In this analysis a 1 minute sampling frequency is chosen, i.e. at the end of each minute of the sample I record the last transaction price in each exchange rate and the order flow, defined as the number of buyer minus seller initiated trades in that minute. After removing weekends, public holidays, the overnight period and collapses in the Reuters data-feed, as described in Chapter 2, the total number of observations was reduced to 124,259 for the USD/EUR, 97,158 for the GBP/EUR and 124,997 for the USD/GBP FX markets.[3] For the cointegrating VAR model of Section 3.3.3 only the periods where no

[2]This is in contrast to the effects of news on price volatility. Ederington and Lee (1993), Andersen and Bollerslev (1998) and Payne (1996) show that return volatility can persist for some time following releases of macroeconomic data. See also the results of Chapter 2.

[3]The substantial reduction in the number of observations for the GBP/EUR market is due to the fact

breakdowns in any of the three exchange rate data feeds were considered. This resulted in 90,270 data observations. Statistical information on exchange rate returns, defined as 100 times the logarithmic difference in prices, transaction frequencies and order flows for the filtered data sample is given in Chapter 2, Table 2.1.

The second component of the dataset consists of euro-area, UK and US macroeconomic information announcements along with expectations data for each of these releases. Again these data are described at length in Section 2.2.2. However, since the news is not broken up into 'good' and 'bad', then there are sufficient news releases emanating from the euro-area to generate a meaningful sample size and allow hypotheses on the effects of such data on FX activity to be tested.[4] The expectations data, obtained from market practitioners, are used to construct the 'news' or 'surprise' component of each individual announcement and from these numbers I construct, also on a 1 minute sampling frequency, an aggregated news variable for each of the euro-area, UK and US respectively. As in Section 2.2.2, each surprise series, $S_{v,t}$, is signed (i.e. multiplied by +1 or -1) depending on its effect on the exchange rate, whereby the series is multiplied by +1 if greater than expected news causes the domestic currency to appreciate, and by -1 if it causes a depreciation. Then, to obtain the aggregated variable, I simply sum the signed, standardised surprise numbers across announcements. The sets of macroeconomic announcements that are included in this study for each of the areas are given in Table 3.1.

3.3 Empirical Analysis

3.3.1 The Effects of macroeconomic news on returns and flows Separately

I begin the empirical analysis by characterizing the effects that news surprises have on exchange rate returns and on order flows, each in isolation. As outlined in the introduction, whilst, given previous results, we would expect macroeconomic news to immediately and significantly move exchange rates, standard models of exchange rate determination suggest that this adjustment should occur without the occurrence of one-sided aggressive trading. I also test a number of other hypotheses. For example, when examining

only eight, rather than ten, months of data were available.

[4]Instead, only one signed news variable is created for each type of release. In the example of US PPI data in Table 2.3, the news variable is simply that in the column denoted 'Signed' rather than split into 'Good' and 'Bad'.

Table 3.1
Description of Macroeconomic Data Releases

Euro-area announcements[1]

Announcement	Sign[2]	Reported as[3]	Obs.[4]	Dates	Local time GMT
Ind. Prod.	+1	3M/3M % change[5]	10 (8)	10/99 - 7/00	10:00/11:00
M3	+1	Y/Y % change	9 (7)	10/99 - 6/00	08:00/09:00

UK announcements[6]

Announcement	Sign[2]	Reported as[3]	Obs.[4]	Dates	Local time GMT
RPIX	+1	Y/Y % change	10 (8)	10/99 - 7/00	08:30/09:30
Retail Sales	+1	M/M % change	10 (8)	10/99 - 7/00	08:30/09:30
Global Trade	+1	GBP (bn.)	9 (7)	10/99 - 6/00	08:30/09:30
Prov. M4	+1	M/M % change	10 (8)	10/99 - 7/00	08:30/09:30

US announcements[7]

Announcement	Sign[2]	Reported as[3]	Obs.[4]	Dates	Local time ET
CPI	-1	M/M % change	10 (8)	10/99 - 7/00	08:30
PPI	-1	M/M % change	10 (8)	10/99 - 7/00	08:30
Unemployment	-1	%	10 (8)	10/99 - 7/00	08:30
Trade bal.	+1	USD (bn.)	10 (8)	10/99 - 7/00	08:30

Notes:

1. Euro-area Harmonised CPI, PPI, Retail Sales, Unemployment Rate and Ex EMU Balance of Trade, both preliminary and final were also considered in pre-testing but not included when forming standardised news due to insignificant or inconsistent exchange rates effects.
2. When forming the aggregate news variable, each series was multiplied by +1 (-1) if greater than expected news causes the domestic currency to appreciate (depreciate).
3. M/M % change: month on month percentage change. 3M/3M % change: three month on three month percentage change. Y/Y % change: year on year percentage change.
4. The USD/EUR and USD/GBP data spans 10 months, implying 10 observations for each release. The bracketed figure gives the number of observations for the 8 months of GBP/EUR data.
5. Industrial Production was reported as 3M/3M % change for October 1999 - March 2000 and M/M seasonally adjusted % change for April 2000 - July 2000.
6. The following announcements were also considered: PPI, Industrial Production, Unemployment, Current Account, EX EU Trade. They were not included when forming the standardised UK news because they had either insignificant or inconsistent effects on sterling exchange rates. However, they were included when examining the effects on volatility in Section 2.5.
7. US Retail Sales, Industrial Production, Nonfarm Payroll Employment and Monthly M3 were also considered in the pre-testing, but found to have insignificant effects on the dollar rates. Again, with the exception of Nonfarm Payroll Employment, these variables were included in the volatility analysis of Section 2.5. Nonfarm Payroll Employment was excluded from Section 2.5 since the data are released at the same time as the US unemployment rate.

Euro-Sterling behaviour, one would naturally expect UK and euro-area macroeconomic information to have strong effects while news emanating from the US should have a much smaller or zero effect on the rate.[5] Finally, standard efficient markets reasoning would imply that the reaction to news items should be very swift, i.e. completed within seconds or minutes of the announcement itself.

To test the preceding hypotheses I make use of standard time-series methods. I estimate the following models for flows and returns for exchange rate k;

$$\Delta P_t^k = \alpha + \sum_{i=-m}^{m} \beta_i^k N_{t-i} + \epsilon_t^k \quad (3.1)$$

$$F_t^k = \lambda + \sum_{i=-m}^{m} \gamma_i^k N_{t-i} + \eta_t^k \quad (3.2)$$

where P_t^k is 100 times the logarithm of exchange rate k at time t,[6] F_t^k is order flow in the market for exchange rate k in the observation period ending at t, positive order flow in the USD/EUR market for example indicates net euro purchases/dollar sales, and N_t is a 3×1 vector consisting of standardized euro-area, UK and US news measures, respectively, for the interval ending at t. These specifications then simply explain exchange rate movements or flows in terms of news from all three regions. Estimation is performed using OLS but correcting the coefficient variance/covariance matrix for autocorrelation and heteroscedasticity using the Newey-West method.

Table 3.2 and Figure 3.1 present estimation results for all three currency pairs for the preceding equations.[7] Looking first at the return equations for each exchange rate, a number of results emerge. As expected, in the majority of cases returns are significant in the minute immediately following a news surprise.[8] Eight of the nine coefficients on

[5]Of course, in practice, US macroeconomic information might well move rates if it has information content for the state of other economies, perhaps because it indicates the well-being of the "global economy". These cross market effects of news are also reported in Chapter 2.

[6]Unsurprisingly each of the exchange rates were found to be I(1) so (100×) log first differences were used. ADF tests are not reported here.

[7]I estimate the specifications with m as ten minutes but for convenience Table 3.2 only gives the results for the minute pre and post announcement. The very low R^2s reported in the table are to be expected due to the small number of news announcements relative to the tens of thousands of return observations for each exchange rate. This is consistent with the whole event study literature on exchange rates.

[8]Recall that the news variables were signed so that greater than expected news causes an appreciation of the domestic exchange rate. Hence 'good' UK news causes a negative return in the GBP/EUR market but a positive return for USD/GBP. Since only announcements that had significant impacts in the pre-testing were chosen and appropriately signed, it is not surprising that significant coefficients were found at this stage. However, this will not affect the subsequent analysis in Sections 3.3.2 and 3.3.3.

current news are significant at 10% and seven at 5%. Leads of euro-area news are also significant for GBP/EUR and surprisingly also for USD/GBP but this is not believed to represent compelling evidence of any information leakage.[9] It is also the case that one of the first lags of news is significant, UK news on USD/GBP – possibly indicating a small amount of delayed reaction to information. It is also interesting to note that UK and US announcements are significant in all three return estimations such that UK (US) data releases have a systematic effect on the USD/EUR (GBP/EUR) rate.[10] Similarly, euro-area announcements significantly change the USD/GBP rate, although only at the 10% level. In purely quantitative terms, for two of the three exchange rates the largest coefficient is on UK news whilst US news has the largest impact in the USD/EUR market. This final observation can also be seen in the component graphs of Figure 3.1 which plots the cumulative returns and flows from ten minutes pre announcement to ten minutes post announcement. This is done for each of the three foreign exchange markets and examines the effects around announcements from each region.

Table 3.2 also presents results from the order flow estimations. Clearly, the statistical significance in these equations is much stronger than that in the return equations. All nine coefficients on current news are significant at the 5% level in the flow equations, five of which are significant at 1%. US news has the largest effect in the EUR currency markets and UK news has the largest impact on USD/GBP flows. See Figure 3.1. Three of the nine coefficients on the first lead of news are significant, although only one is at the 1% level. However, five of the nine coefficients on the first lag of flow are also significant at 10%, three of which are at 1%, providing fairly consistent evidence that news has both an instant and a slightly delayed effect on order flows.[11] This evidence is entirely novel and is at odds with the predictions of standard asset pricing or exchange rate determination theories.

[9] The significant coefficients on leads of euro-area news may be due to the smaller number of euro-area news releases available. As seen in Table 3.1, only 15 news releases are available for the GBP/EUR sample and 19 for the USD/EUR and USD/GBP samples. This is less than half the number of announcements from the UK and US.

[10] The significant impact of UK news on USD/EUR is surprising considering the insignificant effects found in Chapter 2. This is likely to be due to the fact that transaction prices are used here, while Chapter 2 focuses on bid and ask prices, the high frequency dynamics of which may differ from those associated with executed trades.

[11] It is perhaps not surprising that order flow reacts for more than a minute. The initial order flow itself becomes information, which triggers yet more order flow. News releases will then have both an immediate and a delayed effect on order flows. For this point I thank Carol Osler.

Figure 3.1
Effects of News Releases on Exchange Rate Returns and Flows

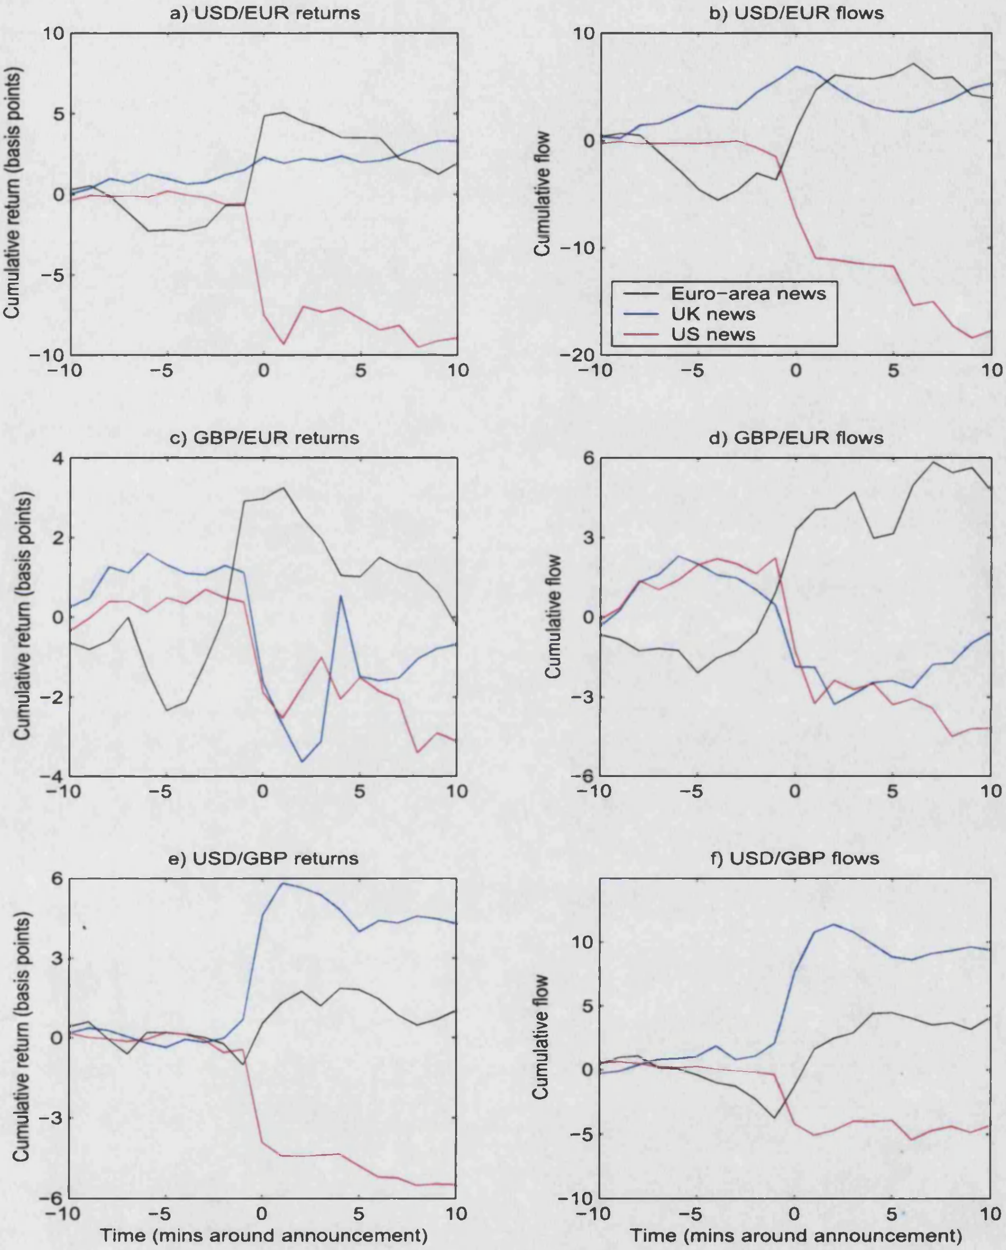

Notes: The figures plot the cumulative returns and flows from ten minutes pre announcement to ten minutes post announcement for each of the foreign exchange markets that are considered. The cumulative returns are plotted by firstly regressing returns on leads and lags of news from each region and then summing the coefficients on news over the twenty-one minute time period. In all plots, the black lines give the responses to euro-area data, the blue lines give responses to UK data and the red lines show the effects of US data releases.

Table 3.2
The Effects of Macroeconomic News on Returns and Flows Separately

	USD/EUR returns		GBP/EUR returns		USD/GBP returns		USD/EUR flows		GBP/EUR flows		USD/GBP flows	
constant	-0.000161^{b}	(-2.16)	$-0.0^{4}457$	(-0.58)	-0.000104^{b}	(-2.08)	0.0385^{a}	(3.17)	0.142^{a}	(14.10)	0.0767^{a}	(8.25)
Euro (1 lead)	0.0000146	(0.01)	0.0287^{a}	(3.45)	-0.00783^{a}	(-2.75)	-0.658	(-1.07)	1.51^{a}	(4.23)	-1.51^{c}	(-1.67)
Euro news	0.0554^{a}	(2.76)	0.000599	(0.07)	0.0155^{c}	(1.89)	4.75^{a}	(2.74)	2.36^{b}	(2.19)	2.51^{b}	(2.51)
Euro (1 lag)	0.00207	(0.34)	0.00280	(0.31)	0.00799	(1.12)	3.62^{a}	(2.84)	0.763	(1.58)	2.87^{a}	(2.87)
UK (1 lead)	0.00310	(0.97)	-0.00187	(-0.47)	0.00707	(1.15)	1.07	(1.37)	-0.567	(-1.09)	0.98	(1.44)
UK news	0.00785^{b}	(2.37)	-0.0275^{a}	(-3.15)	0.0393^{a}	(5.51)	1.30^{b}	(1.99)	-2.31^{b}	(-2.23)	5.59^{a}	(6.40)
UK (1 lag)	-0.00378	(-0.64)	-0.0107	(-1.17)	0.0120^{b}	(2.18)	-0.600	(-0.56)	-0.0246	(-0.01)	3.13^{a}	(3.00)
US (1 lead)	0.000657	(0.20)	-0.00117	(-0.41)	0.00126	(0.68)	-0.870^{c}	(-1.86)	0.588	(1.42)	-0.200	(-0.38)
US news	-0.0696^{a}	(-3.29)	-0.0226^{b}	(-2.33)	-0.0348^{a}	(-2.65)	-5.42^{a}	(-3.66)	-3.50^{a}	(-3.51)	-3.85^{a}	(-3.34)
US (1 lag)	-0.0181	(-1.54)	-0.00643	(-1.12)	-0.00489	(-0.89)	-4.03^{b}	(-2.22)	-1.99^{c}	(-1.81)	-0.874	(-1.24)
R^2	0.00237		0.00171		0.00236		0.00202		0.00136		0.00211	

Notes: The USD/EUR exchange rate is defined as the number of dollars (numerator currency) per euro (denominator currency) and similarly for the other rates. Returns are defined as 100 times the first difference of the logarithm of the exchange rate. Positive order flow in the USD/EUR market implies net purchases of euro, the denominator currency. a, b, c denote significance at the 1%, 5% and 10% levels respectively. T-stats in parentheses.

3.3.2 The role of order flow in exchange rate determination around announcements

The preceding analysis has demonstrated that announcements of macroeconomic information not only cause exchange rates to move, but also generate one-sided order flows. I now begin to focus on the manner in which flows and rates are jointly determined around announcement times. In this section a simple question is posed. Does order flow have a greater or smaller role to play in exchange rate determination when macroeconomic news is publicly released? Ex ante, one might have thought that the answer to this question was almost certainly negative; public information releases would be expected to move rates in the absence of flows leading to high-frequency disconnection between these two variables. However, the prior analysis has shown a strong reaction of flow to news and thus perhaps this simple intuition is not valid. To answer this question I estimate the following specification for the three rates;

$$\Delta P_t^k = \alpha + \beta F_t^k + \sum_R \sum_{i=-m}^{m} \gamma_{i,R} \, F_t^k \cdot I^R(i)_t + \epsilon_t^k \tag{3.3}$$

where $I^R(i)_t$ is an indicator variable taking the value unity if and only if there was an announcement surprise from region R in period $t - i$. Thus, the terms forming the summation in the equation above simply pick out intervals around news releases and test whether the coefficient on flow changes relative to its normal level. For example, the coefficient on the product of flow and $I^{UK}(0)_t$ tells us whether, in a minute that has begun with a UK news announcement, flow matters more or less than usual.

In Table 3.3 I present the results for the nine exchange-rate/news combinations as well as the benchmark order flow coefficient ($\hat{\beta}$). Again, while (3.3) is estimated for i = -10 to +10, for convenience only the results for i = -1 to +1 are presented. The results are clear cut, especially for US and UK news. Around the release of US (UK) information, order flow has a significantly larger effect on the determination of Dollar (Sterling) related exchange rates. Coefficients on the contemporaneous interaction terms are significant at the 5% level (at least) and positive. Their magnitudes are such that in the case of US news, the effect of order flow more than doubles at the time of release while for UK releases the order flow impact is almost doubled. Results are less impressive for euro-area news, however.[12]

[12]The poor results for euro-area news may come from the smaller number of news releases available.

Table 3.3
The Role of Order Flow in Exchange Rate Determination Around Announcements

	USD/EUR return	GBP/EUR return	USD/GBP return
constant	-0.000322[a]	-0.000613[a]	-0.000325[a]
	(-5.79)	(-9.19)	(-7.92)
Flow_t	0.00413[a]	0.00397[a]	0.00291[a]
	(72.80)	(91.49)	(119.53)
$F_t \cdot I^{euro}$ (1 *lead*)	-0.00100	0.00809[b]	0.000127
	(-0.90)	(2.45)	(0.10)
$F_t \cdot I^{euro}$	0.00626[c]	0.000544	0.00206
	(1.77)	(0.34)	(0.75)
$F_t \cdot I^{euro}$ (1 *lag*)	-0.00205[a]	0.00234	0.000533
	(-2.57)	(0.65)	(0.29)
$F_t \cdot I^{UK}$ (1 *lead*)	-0.00207[a]	-0.00170[b]	0.00181
	(-2.81)	(-2.27)	(1.36)
$F_t \cdot I^{UK}$	0.000531	0.00339[a]	0.00322[a]
	(0.69)	(3.27)	(3.89)
$F_t \cdot I^{UK}$ (1 *lag*)	-0.000583	-0.00330	-0.000447
	(-0.65)	(-1.59)	(-0.58)
$F_t \cdot I^{US}$ (1 *lead*)	-0.000952	0.000338	-0.000921
	(-0.98)	(0.20)	(-0.78)
$F_t \cdot I^{US}$	0.00701[a]	0.00204	0.00342[b]
	(4.09)	(0.89)	(2.17)
$F_t \cdot I^{US}$ (1 *lag*)	0.00127	-0.00142	0.00182
	(1.22)	(-1.20)	(1.41)
R^2	0.272	0.173	0.219

Notes: The USD/EUR exchange rate is defined as the number of dollars (numerator currency) per euro (denominator currency) and similarly for the other rates. Returns are defined as 100 times the first difference of the logarithm of the exchange rate. a, b, c denote significance at the 1%, 5% and 10% levels respectively. T-stats in parentheses.

Thus, contrary to what one might expect, strong evidence is derived showing that around US and UK macroeconomic announcements, exchange rates are more sensitive to order flow than at other times. Given this result and that of Section 3.3.1, that flows react strongly to announcements, it would seem that the effect of publicly released information on exchange rates is transmitted, at least partially, through order flow. The empirical analysis in the next section tests this assertion.

3.3.3 Multivariate VAR analysis of returns and flows with exogenous news variables

Finally, I move on to test whether any part of the exchange rate response to news can be characterised as intermediated by order flow. Thus far, the analysis has established that both rates and flows respond to news and also that flows are more important in exchange rate determination around news events. Now I seek to measure the contribution of order flows to the overall exchange rate response to news. For this task, since the data cover the triangle of exchange rates, USD/EUR, GBP/EUR and USD/GBP, it seems natural to estimate a VAR for rates and flows, imposing the obvious cointegrating vector for the exchange rates that is implied by absence of triangular arbitrage.

A similar model to that employed in (2.2) is used here, but only transaction price returns and order flows are included, rather than considering buys and sells separately. Each of the exchange rates and cumulative order flows were found to be $I(1)$ but as explained in Chapter 2, linear combinations of these 6 variables may be stationary, implying cointegration in the system. The absence of triangular arbitrage is one obvious cointegrating vector, but Killeen, Lyons, and Moore (2002) and Bjønnes and Rime (2003) both find evidence that the level of the exchange rate and cumulative order flow are cointegrated. This implies a further three cointegrating vectors. In order to determine the cointegrating rank of the 6 variable system in exchange rate returns and order flow, 'trace' and 'λ max' tests were performed based on Johansen (1995). The procedure is outlined in the appendix to Chapter 2, Section 2.A.3 and the results are reported in Table 3.4.

Despite 7 cointegrating vectors being found in the 12 variable system of (2.2), the trace test suggests there is only one cointegrating relationship when transaction price returns and order flows are considered. However, the λ max test rejects the null hypothesis that the cointegrating rank equals 1 against the alternative that there are two cointegrating vectors,

See footnote 9.

Table 3.4
Johansen Test for Cointegrating Rank: Transaction Prices

Panel A 'trace' test				Panel B 'λ-max' test			
H_0	H_A	t stat	1% critical	H_0	H_A	t stat	1% critical
rank $= h$			value	rank $= h$			value
$h=0$	$h=6$	204.44	103.18	$h=0$	$h=1$	135.18	45.10
$h=1$	$h=6$	69.26	76.07	$h=1$	$h=2$	40.11	38.77
$h=2$	$h=6$	29.15	54.46	$h=2$	$h=3$	13.79	32.24
$h=3$	$h=6$	15.36	35.65	$h=3$	$h=4$	8.42	25.52
$h=4$	$h=6$	6.94	20.04	$h=4$	$h=5$	6.41	18.63
$h=5$	$h=6$	0.53	6.65	$h=5$	$h=6$	0.53	6.65

Notes: Critical values are taken from Osterwald-Lenum (1992), Table 1 (constant unrestricted, no time trend).

although the 1% critical value is only just breached. I therefore choose to include only 1 cointegrating relationship, based on the absence of triangular arbitrage between rates.[13] As in Section 2.3, this arbitrage condition implies a cointegrating vector of $[1 \ \ -1 \ \ -1]'$ for the logarithm of the three rates. Therefore, similar to the model presented in Chapter 2, the model in (3.4) allows one to characterise the effects of news announcements on all three exchange rates and respective order flows simultaneously.

$$\begin{aligned} \begin{bmatrix} \Delta P_t \\ F_t \end{bmatrix} &= \alpha + \delta z_{t-1} + \begin{bmatrix} \beta \\ 0 \end{bmatrix} F_t + \sum_{i=1}^{m} \Gamma_{(i)} \begin{bmatrix} \Delta P_{t-i} \\ F_{t-i} \end{bmatrix} + \sum_{j=0}^{n} \Theta_{(j)} N_{t-j} + \varepsilon_t \\ Var\left[\epsilon_t\right] = \Omega &= \begin{bmatrix} \Omega_{11} & 0 \\ 0 & \Omega_{22} \end{bmatrix} \end{aligned} \tag{3.4}$$

where ΔP_t is the 3×1 vector of USD/EUR, GBP/EUR and USD/GBP exchange rate returns, F_t is the corresponding 3×1 vector of order flows and N_t is again the 3×1 vector of standardized euro-area, UK and US news. Hence I allow news from all three regions to affect a given exchange rate's returns. Finally z_{t-1} is an error correction term derived

[13]Intuition dictates that if there is one cointegrating vector, this has to be associated with absence of triangular arbitrage. ADF tests on the residuals (in the regression of one rate on the other two) suggested that these three rates were indeed cointegrated. However, system methods based on Johansen (1995) and similar to those presented in Section 2.A.4 suggested that the cointegrating vector was *not* associated with absence of arbitrage. Since any relationship other than the absence of arbitrage appears implausible, it is imposed on the system.

from the above theoretical cointegrating vector. Define d_{t-1} to be the discrepancy:

$$d_{t-1} = \ln(USD/EUR_{t-1}) - \ln(GBP/EUR_{t-1}) - \ln(USD/GBP_{t-1}) \qquad (3.5)$$

Then z_{t-1} is defined in the following way.

$$z_{t-1} = D_{t-1}\left(d_{t-1} - k_{t-1} \cdot sign\left(d_{t-1}\right)\right) \qquad (3.6)$$

where k_{t-1} is the intra-day average cost of arbitrage for minute $t-1$. D_{t-1} is a dummy variable that takes the value of unity if d_{t-1} is greater than k_{t-1} in absolute value and zero at all other times.[14] Essentially z_{t-1} takes into account the cost associated with the absence of triangular arbitrage condition implied by the bid-ask spreads in each market. The cost of making the round trip series of trades from dollars to euros to sterling and back to dollars at date $t-1$ is, in log terms, $\ln\left(USD/EUR^{ask}_{t-1}\right) - \ln\left(GBP/EUR^{bid}_{t-1}\right) - \ln\left(USD/GBP^{bid}_{t-1}\right)$. The cost of trading from dollars to sterling to euros and back to dollars, is similarly $\ln\left(USD/GBP^{ask}_{t-1}\right) + \ln\left(GBP/EUR^{ask}_{t-1}\right) - \ln\left(USD/EUR^{bid}_{t-1}\right)$. The average cost of arbitrage at date $t-1$ is therefore half the sum of the log bid-ask spreads. k_{t-1} is defined as the average of these arbitrage costs at time $t-1$ calculated over all days except weekends and public holidays. The intuition is that in equilibrium, d_{t-1} may not equal zero because of the existence of the bid-ask spreads. Only if $\ln(USD/EUR)$ wanders too far from the other two rates will the market drive the prices back in line. δ is then a 6×1 vector of speed of adjustment coefficients. Note that order flow at date t can be affected by the equilibrium error at date $t-1$ due to the absence of triangular arbitrage argument. To see why, assume for simplicity that all three exchange rates initially equal unity. Also assume the cost of arbitrage is zero, i.e. zero bid-ask spreads in each market. Now allow the euro to appreciate only against the dollar at date $t-1$ to 1.02 dollars per euro. To exploit the arbitrage opportunity a simple strategy would be to sell 1 euro, buying \$1.02. With this you can buy £1.02 and then buy 1.02 euros making a pure profit of 0.02 euros for every 1 euro traded. Note that exploitation of the existence of a positive equilibrium error at $t-1$ induces a net sale of euros for dollars (negative USD/EUR order flow), a net sale of dollars for sterling (positive USD/GBP order flow) and a net sale of sterling for euros (positive GBP/EUR order flow). Therefore I hypothesize that

[14]This is a simple threshold cointegration model (see Balke and Fomby (1997) and also Tsay (2002)) where the coefficients in the model are restricted to be the same in each regime but where $\delta = 0_{6\times1}$ in the intermediate region, i.e. where there are small deviations from equilibrium.

the six coefficients in the speed of adjustment vector, δ, will have the following signs: $[-, +, +, -, +, +]$.

The key assumptions in the VAR formulation are that order flows do not depend on contemporaneous exchange rate returns and also that the returns for exchange rate k do not depend on contemporaneous returns for exchange rate l ($\neq k$). Finally I also assume that order flows do not depend on the contemporaneous realisation of other order flows. With 9 structural parameters to estimate, the 3×3 β matrix of price impact of order flow coefficients, only 9 restrictions in the variance/covariance matrix of the 6×1 error vector, ϵ_t, are needed. These 9 restrictions come from the assumption that none of the errors in the return equations are correlated with any of the errors in the flow equations. This implies the variance/covariance matrix of the residuals is block diagonal, with the variance matrix of the 3×1 vector of return (flow) errors being freely estimated. So despite the restriction that the return for exchange rate k does not depend on the contemporaneous return for exchange rate l ($\neq k$), and similarly for order flows, the *error* in the return (flow) equation for exchange rate k *is* allowed to be correlated with the error in the return (flow) equation for exchange rate l ($\neq k$), as one may expect if the three rates are kept together very closely by the absence of triangular arbitrage.

The recursive ordering that allows flows to contemporaneously affect returns but which rules out the converse seems to be a reasonable economic restriction, especially at the very high frequency considered here. This recursive ordering of the VAR is also common in the literature. See Hasbrouck (1991a) and Payne (2003a) to name but two. However, in Chapter 5 I present a VAR model which does allow contemporaneous feedback trading, but find that at the one minute frequency, which is used in this chapter, the effects of contemporaneous feedback trading are not statistically significant, although in economic terms, the effects may be considerable.[15] Due to the cross correlation of the errors, one has to estimate (3.4) as a system and because of the recursive ordering of the structural VAR (flows do not depend on contemporaneous returns) we can estimate (3.4) as a SUR system. This allows heteroscedasticity across equations and I also correct for heteroscedasticity within each equation using the Newey-West technique. This procedure, of correcting for heteroscedasticity both *across* and *within* equations, is outlined in Chapter 2 and was introduced by Creel and Farell (1996). See also Fiebig (2001).

[15]As explained in Chapter 5, positive feedback trading suggests a greater role played by transactions in the assimilation of information into price. Since this chapter tries to estimate the share of public information entering via the trading process, the existence of positive feedback trading actually reinforces the results presented here, whereby a large proportion of public information enters via trading.

Table 3.5 gives a summary of the VAR estimations and for convenience, only the parameter estimates and t-stats are given for the constant, the equilibrium error, contemporaneous flows and news.[16] The coefficients on the news variables are to be interpreted as follows: after a one standard deviation announcement of good US news for example, this causes a direct dollar appreciation (decrease in the USD/EUR rate) of 1.63 basis points and causes a decrease in net USD/EUR order flow (i.e. leads to a net purchase of dollars) of 2.91, i.e. 2.91 more purchases of dollars than sales of dollars.[17]

Some general comments are as follows. First there are indications of high-frequency negative autocorrelation in returns for all three rates and also a positive effect of contemporaneous order flow in the return equations, as expected. Interestingly, order flow exhibits high-frequency positive autocorrelation and also high frequency positive dependence on recent returns. These parameters might be interpreted as indicating high-frequency momentum trading by market participants. More importantly, news effects are strong in the return and flow equations. These are summarised below.

- USD/EUR exchange rate returns are only affected by euro-area data. Flows are weakly positively affected by EU news and weakly negatively affected by US news.
- UK news strongly affects GBP/EUR exchange rate returns (1% level) and flows (5% level) in the predicted directions but euro-area news causes GBP/EUR returns to move in the 'wrong' direction. However, this coefficient is only significant at the 10% level and may result from the small number of euro-area data announcements in the dataset.
- UK news causes significant changes in USD/GBP returns and flows (both at the 1% level) while US news only causes a significant change in USD/GBP order flows (5% level). The effect on returns is in the correct direction but not significant.

In addition to these effects, 'good' US data announcements tend to induce a net purchase of sterling in the GBP/EUR market and good euro-area data likewise in the USD/GBP market.[18] The fact that these effects are significant in the flow equations suggests exploitation of arbitrage opportunities in the seconds (minutes) following the data releases.

[16]The lag length on returns, flows and news were chosen using the Schwartz information criterion and were found to be 7, 3 and 0 respectively.

[17]The total effect on the USD/EUR rate of this news will be larger than 1.63 basis points due to the indirect effect of the news via the induced order flow. The total effect, as shown in Figure 3.2g, is a dollar appreciation of 3.71 basis points.

[18]Bivariate VARs were also estimated for each currency pair in a similar fashion to (3.4). However, the results, including the cross effects of news on 'non-relevant' returns, were, for all intents and purposes, the same.

Table 3.5
Multivariate VAR Analysis of Returns and Flows with Exogenous News Variables

	USD/EUR returns		GBP/EUR returns		USD/GBP returns		USD/EUR flows		GBP/EUR flows		USD/GBP flows	
constant	-0.0598[a]	(-8.49)	-0.0485[a]	(-7.09)	-0.0198[a]	(-3.94)	0.0368[a]	(3.54)	0.120[a]	(14.92)	0.0492[a]	(5.47)
δ	-0.0577[a]	(-3.21)	0.0520[a]	(3.71)	0.0132[a]	(3.00)	-0.441	(-1.00)	0.140	(0.40)	0.182	(0.51)
$F_{USD/EUR,t}$	0.383[a]	(50.44)	0.208[a]	(39.31)	0.101[a]	(35.92)						
$F_{GBP/EUR,t}$	0.205[a]	(40.19)	0.328[a]	(67.77)	-0.0783[a]	(-23.32)						
$F_{USD/GBP,t}$	0.146[a]	(36.60)	-0.101[a]	(-24.59)	0.287[a]	(100.32)						
$N_{euro,t}$	1.273[a]	(2.76)	-0.620[c]	(-1.77)	-0.170	(-0.62)	3.47[c]	(1.87)	2.42[b]	(2.16)	3.09[a]	(3.15)
$N_{UK,t}$	0.125	(0.57)	-1.582[a]	(-3.16)	2.119[a]	(3.42)	0.710	(1.14)	-2.15[b]	(-2.45)	4.98[a]	(5.07)
$N_{US,t}$	-1.628	(-1.27)	-0.614	(-0.74)	-0.922	(-1.55)	-2.91[c]	(-1.76)	-3.50[a]	(-3.64)	-1.73[b]	(-2.16)
R^2	0.367		0.277		0.271		0.084		0.067		0.054	

Notes: The data cover the eight month period from 1st December 1999 to 24th July 2000. The USD/EUR exchange rate is defined as the number of dollars (numerator currency) per euro (denominator currency) and similarly for the other rates. Positive order flow implies net purchases of the donominator/commodity currency; the euro in USD/EUR and GBP/EUR and sterling in USD/GBP. All returns are defined as 10000× the log first difference of the rate [a], [b], [c] denote significance at the 1, 5 and 10% levels respectively. T-stats in parentheses. Coefficients are to be interpreted as follows; one unit of USD/EUR order flow causes the USD/EUR rate to rise by 0.383 basis points in that minute. A one standard deviation announcement of good euro-area news in minute t causes a direct rise in the USD/EUR rate of 1.273 basis points and for USD/EUR order flow to increase by 3.47 units, i.e. 3.47 more purchases of euro than sales of euro.

The estimates of the speed of adjustment coefficients in the cointegrating VAR are also consistent with these arbitrage strategies. The sign of the δ coefficients are as expected (and significant) in the return equations and are as expected in the three flow equations, although none are statistically significant. This suggests that in the period immediately following an announcement of news, traders try to exploit the profitable opportunities that exist when the equilibrium error is non zero. For example, if following an announcement of good US data, the dollar appreciates more against the euro than it does against sterling,[19] this will cause the equilibrium error to become negative, inducing a net sale of euros for sterling as market participants try to exploit the arbitrage opportunities explained above, assuming the arbitrage strategy covers the cost of making the round trip series of trades, denoted by k in (3.6). This would explain the statistically significant negative coefficient of US news on GBP/EUR order flow. Similar arguments can be used to explain the significant coefficients of euro-area news in the USD/GBP flow equation.[20] However, it is interesting to note that the order flow generated by this 'non-relevant' news does not translate into the exchange rate changes that one would expect. For example, good US news causes negative GBP/EUR order flow but does not have any significant effect on GBP/EUR exchange rate returns.

In order to evaluate the role that order flow plays in the reaction of exchange rates to news I compare the actual impulse response of the exchange rate following news releases from each region to that which would occur if (counterfactually) all coefficients on news in the structural order flow equation were restricted to be zero, i.e. I replace the bottom 3 by 3 sub-matrix of each $\Theta_{(j)}$ in the estimated version of (3.4) by the null matrix. By doing this, flows are not given any role in the assimilation of public information into prices and hence it gives a measure of the extent to which the equilibrium exchange rate response is driven by the impact of news on flows. To calculate the IRF I introduce a one standard deviation surprise announcement of news in the N_t vector from one region only and examine the effects on each of the returns through time. The impulse response functions are plotted in Figure 3.2 for cumulated returns and Figure 3.3 for flows. The solid black line shows the IRF following the hypothetical announcement and the dashed blue lines bound the 95% confidence interval, found by bootstrapping over a thousand iterations. For example, following a hypothetical announcement of positive, one standard deviation,

[19]This is indeed the case. From the impulse response analysis, the dollar appreciates by 3.71 basis points against the euro (Figure 3.2g) and 1.44 basis points against sterling (Figure 3.2i).

[20]Good euro-area news causes an appreciation of 3.55 basis points against the dollar (Figure 3.2a) and an appreciation of 0.58 basis points against sterling (Figure 3.2b). Hence significant positive USD/GBP order flow is not surprising.

Figure 3.2
Multivariate VAR Impulse Response Analysis of News on Cumulative Returns

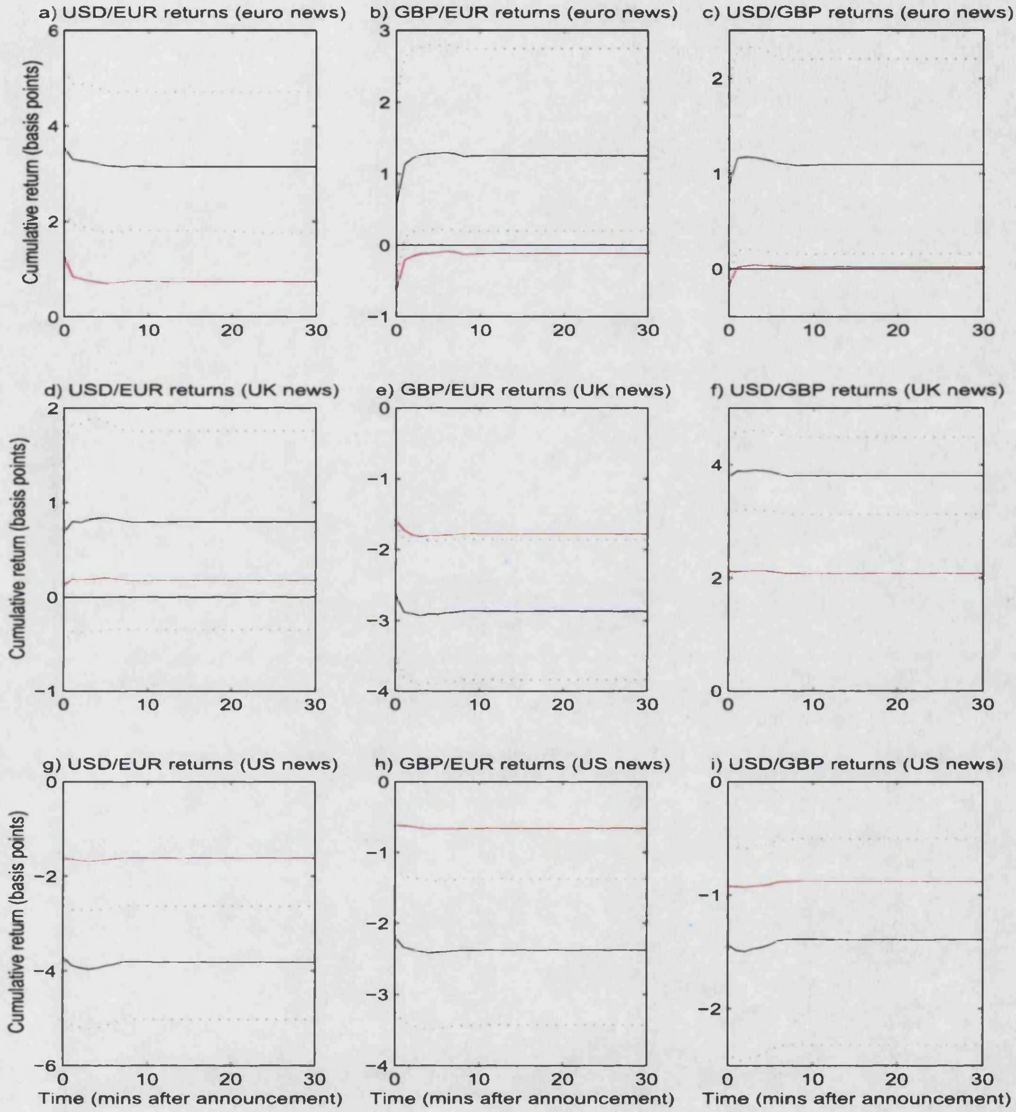

Notes: The figures plot the impulse response functions following a one standard deviation announcement of news from each region. The news shock was introduced into the estimated VAR of (3.4) and the cumulative return calculated. In all plots, the solid black lines give the actual impulse response function and the dashed blue lines trace out a 95% confidence interval for the IRF found by bootstrapping over a thousand iterations. The red lines give the implied impulse responses when all news coefficients in the order flow equations are set to zero, and hence show the impact of news with 'no flow' effects.

Figure 3.3
Multivariate VAR Impulse Response Analysis of News on Exchange Rate Flows

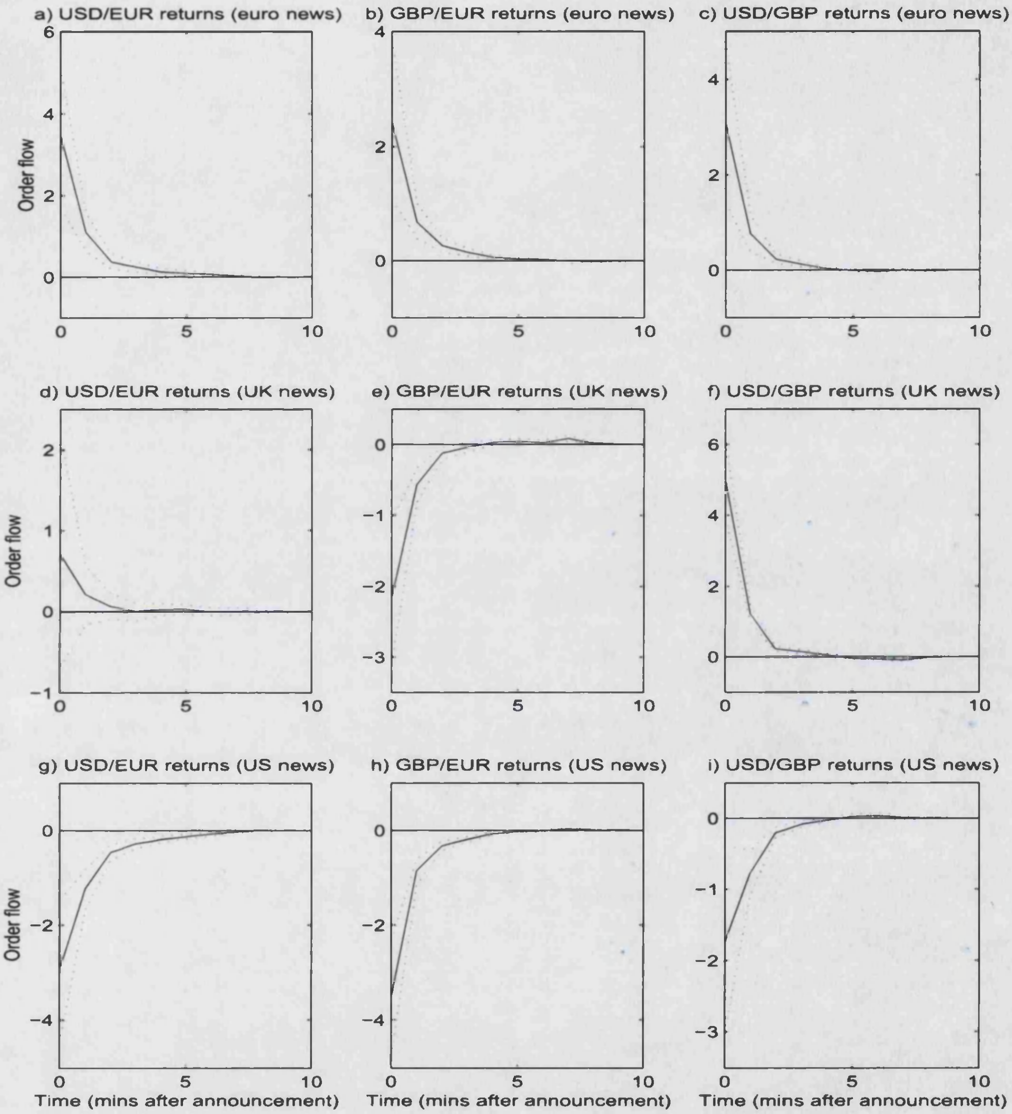

Notes: The figures plot the impulse response functions following a one standard deviation announcement of news from each region. The news shock was introduced into the estimated VAR of (3.4) and the order flows calculated. In all plots, the solid black lines give the actual impulse responses and the dashed blue lines trace out a 95% confidence interval found by bootstrapping over a thousand iterations.

euro-area news, the USD/EUR exchange rate increases (euro appreciates) by 3.55 basis points on announcement and is still above 3 basis points higher after 30 minutes. (See Figure 3.2a). In all cases where the news release comes from one of the regions associated with the currency pair, the cumulated exchange rate return is significant and in the direction expected. However, in two of the three cases where the news emanates from areas which are not associated with the currency pair, the hypothetical announcement causes a significant exchange rate change: positive US data causes sterling to significantly appreciate against the euro and positive euro-area data causes a significant appreciation of sterling against the dollar. The IRF figures also show the hypothetical exchange rate responses that would occur if all news coefficients in the structural order flow equations were set to zero, shown as red lines and denoted 'no flow'. Thus comparison of black and red lines give some indication of the importance of flow in transmitting news to rates. In all cases, the hypothetical response to news with 'no flow' is less than that where flows are allowed to be influenced by news, indicating that the mechanism through which information affects price directly does not explain the full story of (public) information assimilation. Order flows and the trading process therefore account for at least some of the price movements following releases of public information.

In order to give a quantitative assessment of the role order flow plays in the assimilation of news into foreign exchange prices, I break down the exchange rate change into the component coming directly from news into price (shown by the IRF when order flow is constrained not to be affected by news) and that coming via order flow, the difference between the original IRF (black) and the 'no flow' IRF (red). This is shown in Table 3.6, which also decomposes the effect for each of the three currency pairs, for each news announcement and also examines the breakdown at the time of the announcement, 5 minutes after announcement and finally 30 minutes after. Note, if the 'no flow' IRF is negative, and the original IRF is positive, the portion of the news announcement attributed to flow will be greater than 100% and that attributed to direct impounding into price will be negative since without order flow, the effect is to move the exchange rate in the 'wrong' direction. As can be seen in Figure 3.2b, this is only the case for euro-area news on GBP/EUR returns. Good news causes a euro appreciation but when we rule out the possibility of information entering via flows, a simulated depreciation is seen. This is because as seen in Table 3.5, euro-area news causes a positive flow in the GBP/EUR market but a negative return. The net effect, however, is a positive return since the indirect effect of news on returns coming via the contemporaneous flow variable outweighs the negative

Table 3.6

Breakdown of Information Assimilation into 'Flow' and 'Direct' Effects for the Multivariate VAR

	Time after announcement (mins)	USD/EUR 'Flow' (% share)	USD/EUR 'Direct' (% share)	GBP/EUR 'Flow' (% share)	GBP/EUR 'Direct' (% share)	USD/GBP 'Flow' (% share)	USD/GBP 'Direct' (% share)	Average over all pairs 'Flow' (% share)	Average over all pairs 'Direct' (% share)
Euro	0	64.13	35.87			119.35	-19.35	**91.74**	**8.26**
data	5	77.74	22.26			96.73	3.27	**87.23**	**12.77**
	30	76.43	23.57			97.60	2.40	**87.01**	**12.99**
UK	0	81.60	18.40	40.10	59.90	44.07	55.93	**55.26**	**44.74**
data	5	75.87	24.13	38.09	61.91	45.29	54.71	**53.08**	**46.92**
	30	77.37	22.63	38.11	61.89	45.24	54.76	**53.58**	**46.42**
US	0	56.13	43.87	72.01	27.99	35.97	64.03	**54.70**	**45.30**
data	5	57.80	42.20	72.56	27.44	37.16	62.84	**55.84**	**44.16**
	30	57.77	42.23	72.52	27.48	36.92	63.08	**55.74**	**44.26**
Average	0	**67.29**	**32.71**	**56.05**	**43.95**	**66.46**	**33.54**	**64.17**	**35.83**
over all	5	**70.47**	**29.53**	**55.33**	**44.67**	**59.73**	**40.27**	**62.66**	**37.34**
news	30	**70.52**	**29.48**	**55.32**	**44.68**	**59.92**	**40.08**	**62.75**	**37.25**

effect of news on returns directly. However, since the negative direct effect on the GBP/EUR rate is only significant at the 10% level (Table 3.5) and we only have 15 euro-area data announcements available, little should be taken or inferred from the fact that the share of price changes coming via flows is over 100% in the GBP/EUR market. For this reason I ignore the effect of euro-area news on GBP/EUR returns when decomposing the price movements. When averaging over all currency pairs and over all announcements but excluding the spurious euro-area news effect, the assimilation of information coming via the trading process is substantial. The share that flows have in the impounding of news into price is 64.17% on announcement, 62.66% after five minutes and 62.75% thirty minutes after announcement. This decomposition appears stable through time following the announcement but when we look at the breakdown for each country's news release and the effect on each currency pair, the results are more erratic. However, it certainly appears to be the case that the role played by order flow is substantial; up to twice as important, or at least just as important, as the direct effect of news being impounded into price.

3.4 Discussion and Interpretation

Theory suggests that under the assumptions of rational expectations and efficient markets, public information should be incorporated into asset prices immediately and without a need for trading activity. Indeed French and Roll (1986) define public information as that which is incorporated into prices before any market participant can trade on it. Whereas equity values are determined by both public (macro) and firm specific information, exchange rate determination, it has been argued, is primarily concerned with macroeconomic information (Bessembinder 1994). The release of unexpected, publicly announced, macroeconomic information should then be the major cause of exchange rate changes.[21] The results presented here are not inconsistent with this hypothesis. Unexpected macroeconomic data releases have significant effects on exchange rate levels. However, I also find that the same news has significant impacts on order flow and what is more, these impacts are in the same direction as the associated exchange rate changes. Further analysis demonstrates that the price impact of order flow is significantly increased around announcement times, especially for UK and US data releases. This suggests that at least part of the process whereby public information is incorporated into prices, comes

[21] However Lyons (1995) and Payne (2003a), amongst others, find evidence suggesting a large role played by asymmetric (private) information in foreign exchange markets.

via trading. This is explicitly tested in Section 3.3.3 and the hypothesis that public information is impounded into price without the need for trading is decisively rejected. On average over half of the price adjustment comes via order flow, a result that is entirely inconsistent with standard efficient markets, rational expectations hypotheses of asset price determination.[22] Although I demonstrate a role for the trading process in the impounding of public information into prices, no evidence is found to suggest that foreign exchange markets are not efficient. Virtually all of the associated price changes occur within two minutes of announcement.

Why order flow is so important in the formation of prices around releases of public information is a separate, important question and the results can be interpreted in a number of different ways. First, it could be argued that the agents trading on the system from which the data are drawn are slow to update the prices at which they are willing to trade and thus individuals who learn the implications of macroeconomic news more quickly can trade profitably on this information. In such a case, some public information will appear to get into prices via order flow as market orders are executed against 'stale' limit orders. This scenario seems unlikely given the frequency of trading and order updating that is observed on the system. The news announcements I consider are all scheduled releases and are generally accepted by market participants as being the main movers of rates. FX traders will then continuously monitor their limit order quotes around announcements of such data releases, otherwise they are certain to face losses as their stale quotes are hit by those who are monitoring and are active in the market.[23] Alternatively, perhaps it is the case that groups of agents, because they have differing views of exchange rate determination, disagree on the implications of a release for rates and this disagreement generates a motive for trade, as in the models of Harris and Raviv (1993), Kandel and Pearson (1995) and Varian (1989).[24] When the mapping of fundamentals/macroeconomic announcements to price is known and common across agents, the associated price movements should be independent of order flow. By relaxing the assumption that this mapping

[22]Evans and Lyons (2003), looking at four months of foreign exchange data, suggest that the channel through which public information is incorporated into prices via order flow is around twice as important as the direct channel whereby information is impounded into prices with no need for trading. Whereas this and Evans and Lyons' work in the foreign exchange market suggest a very important role for order flow in the assimilation of information into price, Fleming and Remolona (1999) find that prices in the US Treasury market incorporate macroeconomic information without the need for trading. This is also the case for stock returns. See for example Jain (1988), who finds that stock prices adjust to macro news announcements without significant increases in trading.

[23]Moreover, many FX dealers in these markets are quoting in only a single market, enabling them to monitor their quotes around announcement times more easily.

[24]As explained in Chapter 4, this disagreement can easily explain the increase in trading volume, reported in Chapter 2, but why this generates signed order flow is a little more puzzling.

is perfectly understood, perhaps because it is costly to discover the true mapping process, then after a news release the differing beliefs will induce order flow that will move prices to the new equilibrium. Green (2004) finds that the informational role (and price impact) of order flow increases in the government bond market around announcements of scheduled macroeconomic data releases, suggesting that the increase in information asymmetry is consistent with the idea that market participants have different abilities in interpreting the price implications of data releases. This view, that there may be no consensus on the price implications of macroeconomic data releases is also taken by Evans (2002) when examining the FX market.

Regardless of the reasons for these findings, the results are very strong. Information that is publicly and simultaneously released to all market participants is only fully assimilated into prices via the trading process. As such, the results suggest that the recent separation of macroeconomic and microstructure models for exchange rates is somewhat artificial. As shown in a number of papers since Meese and Rogoff (1983a), macroeconomic models have only been able to explain exchange rates in the long run, whereas the more recent microstructure models have only been considered successful at explaining high frequency exchange rate fluctuations, at the daily level or higher (Evans and Lyons 2002b). More realistic models of exchange rates that merge both macro and microstructure elements should be developed to more accurately explain how exchange rates are determined. The interaction between macroeconomic and microstructure variables is clearly visible in the results presented here. Furthermore, the results imply that the distinction made in the microstructure literature between public and private information is not clear cut. Since the (public) information studied here enters price primarily through the trading process, rather than entering directly, then it might be argued that public information announcements create informational asymmetries across the population of traders, as suggested in Evans (2002) and Green (2004). French and Roll's definition of 'public' information, that which is incorporated into prices before anyone can trade on it, may still be valid, but what constitutes 'public' information may differ from that which is commonly believed.

3.5 Conclusions

Under rational expectations, efficient markets hypotheses, there should be no role for order flow in the assimilation of public information into prices. This chapter shows that these ideas are not correct. I show that in the three floating foreign exchange markets, USD/EUR, GBP/EUR and USD/GBP, using data sampled at the one minute frequency, publicly announced macroeconomic information not only causes exchange rates to move but also causes order flow to change significantly in directions consistent with the exchange rate movements. Indeed, the main driver of exchange rate movements in the microstructure literature, namely order flow, is found to be more informative around macroeconomic data releases. The assertion that public information is impounded into prices without the need for order flow is tested and strongly rejected. Using impulse response analysis I find that up to two thirds of the price relevant information is impounded into prices via order flow. This is at odds with standard theory.

Chapter 4 Foreign Exchange Traders and their Reactions to Public Information

4.1 Introduction

Evidence that public information announcements cause increased trading activity is widespread. Berry and Howe (1994) find that public information releases cause trading activity in the New York Stock Exchange to increase significantly, as do Mitchell and Mulherin (1994) and Kandel and Pearson (1995). In the foreign exchange market, Melvin and Yin (2000) document an increase in quote frequency following releases of public information, which suggests an increase in trading activity.[1] The results of Chaboud et al (2004) and also those presented in Chapter 2 show that macroeconomic data releases lead to a significant increase in FX trading.

Such findings are not easily consistent with efficient markets hypotheses, which suggest that news that is publicly and simultaneously released should cause asset prices to change without the need for trading. Indeed, French and Roll (1986) define public information to be that which is incorporated into asset prices before anyone can trade on it. Any finding that public information causes trading activity to increase therefore suggests a violation of the basic theoretical setup.

However, a common explanation for this increase in trading is trader heterogeneity. Such heterogeneity can originate from differences in preferences, endowments or information and it is differences in information sets on which I concentrate here. These differences in information sets around publicly released data announcements are argued to be due

[1] The positive relationship between quoting activity and trade intensity is documented in Daníelsson and Payne (2002b).

to differences of assessment and analysis among market practitioners. An announcement that a firm's quarterly earnings is greater than expected is likely to be interpreted differently by traders, i.e. the mapping of information to price is not common knowledge. Theoretical models that predict an increase in trading activity following publicly released news include Varian (1989), Kim and Verrecchia (1991), Harris and Raviv (1993) and Kandel and Pearson (1995), all of which take the differences of opinion/differences in interpretation approach. Bessembinder, Chan, and Seguin (1996) test the differences of opinion hypothesis by examining a proxy for the cross-sectional divergence of opinions, defined as the 'open interest' in the S&P 500 index futures contract. They conclude that the differences in traders' interpretation of public information is a major determinant of trading. This is also found in the empirical tests of Kandel and Pearson (1995), where trading volume is found to increase following releases of firms' earnings announcements, even when the news has no effect on the stock's price. Similarly, in the foreign exchange market, announcements of CPI, unemployment or trade figures which are unexpected, and hence constitute 'news', were found in Chapter 2 to increase trading even when these releases had no effect on the level of the exchange rate. This suggests that traders differ in their interpretation of the same news release.

Recent evidence from studies of the foreign exchange market suggest that not only does public information generate trading volume, but trading activity is in fact the mechanism through which public information is assimilated into price. More precisely, public information mainly enters price via order flow, the difference between the number of buyer and seller initiated trades. The models of Kyle (1985), Glosten and Milgrom (1985), Easley and O'Hara (1987), Perraudin and Vitale (1996), among others, show that order flow carries information from the more (privately) informed to the market as a whole. However, Evans and Lyons (2003) and the results of Chapter 3 both show that *public* information leads to an increase in order flow, and what is more, this order flow is more informative during periods of public information releases, in that any given level of order flow results in a larger price change compared to non-announcement times. From this, both studies conclude that up to two thirds of the price relevant information contained in publicly released news is incorporated via the trading process.[2]

Since public information enters prices primarily via order flow (the mechanism by which *private* information is incorporated into asset prices) then segregation of public and pri-

[2]See also the note 'Is order flow correlated with public information?' by Richard Lyons, available at http://faculty.haas.berkeley.edu/lyons/public%20info.pdf

vate information, common in the literature, may not be entirely valid. Indeed, Evans (2002) allows macroeconomic news releases to be split into common knowledge (CK) and non common knowledge (NCK) components, if the data release is not interpreted in a homogenous fashion across market participants, a view echoed in other empirical work by Chari (2002), Green (2004) and Ben Omrane and Heinen (2003).

However, what is unclear is why public information releases should have any effects on *order flow* at all. Differences in interpretation can easily explain increases in *trading volume*; if, following the data release, a trader believes the value of the asset is high, v^H say, and another believes the value of the asset is low, v^L, then the trader who values the asset more highly will tend to want to buy, and the other trader will tend to want to sell, hence generating trading volume, see Varian (1989) for example. But why will such news generate signed order flow? In the above example, if the news was universally considered to be good so that the value of the asset was believed to be higher, but still $v^H > v^L$, if the trader who values it at v^L initiates the trade and sells to the trader who values the asset more highly, then the good news will be associated with negative order flow![3] Any news release (good or bad) should not systematically induce net buyer or seller initiated trades. The aim of the first half of this chapter is therefore to present an explanation as to why announcements of public information generate signed order flow. This is presented in Section 4.2 and is based on the ideas of Harris and Raviv (1993) and Kandel and Pearson (1995), where traders differ in their interpretation of news.

However, if traders disagree on the mapping of information to price, not only would one expect to see an increase in trading volume (and also, as will be demonstrated in Section 4.2, signed order flow) following an announcement of public information, but one would also expect the asymmetric information component of the spread to increase.[4] Suppose traders receive a noisy signal of the value of a currency from a data release, allowing dealers to disagree on the asset's true value. Each dealer then faces the risk of being on the wrong side of a more informed trade. For example, if a trader wishes to execute a market order at another dealer's prices then it may be the case that the aggressive trader has interpreted the announcement more accurately, in which case the recipient of the market order should post a wider spread to compensate. It is also likely that this latter

[3]This does, of course, assume that the trader with expected value, v^H, supplies liquidity by posting limit orders against which the lower value dealer can trade.

[4]The models of Kim and Verrecchia (1994,1997) suggest that spreads/information asymmetry increase following a release of public information. Their models are again based on traders receiving noisy signals of the asset's value from the public announcement and concentrate on the effects of news on spreads and trading volume. They do not address the issue of order flow.

trader will reduce his/her depths, i.e. be willing to trade less for any reduction (increase) in the bid (ask) price. This is indeed seen in the results of Lee, Mucklow, and Ready (1993), who find that depths deteriorate following publicly released announcements of firms' earnings.[5] The hypothesis that announcements of public information cause depths to decrease in the foreign exchange market is tested in the second half of this chapter by examining the effects of a scheduled announcement of US PPI data in the DEM/USD market. The announcement is found to decrease the depth of the market, especially in the first thirty seconds following the data release. What is more, depths are found to decrease even after taking into consideration the huge increase in market orders that naturally drain liquidity and hollow out the limit order book.

Such disagreement between traders over the mapping of information to price can therefore explain the increase in trading volume, documented in Chapter 2, the increase in signed order flow, seen in Chapter 3 and explained in the next section, and the reduction in depths documented later in this chapter. The rest of the chapter is therefore organised as follows. Section 4.2 provides an explanation why public news announcements generate signed order flow and Section 4.3 tests the hypothesis that public news releases lead to a reduction in market depth. Section 4.4 discusses the results and Section 4.5 concludes.

4.2 Differences of Opinion and Order Flow Following Public News Announcements

In this section I present a simple explanation why public news releases generate signed order flow. Unfortunately, there are a number of technical challenges that must be overcome before these ideas can be articulated in an explicit model. Such challenges are beyond the scope of this chapter, but instead I simply describe the intuition. The difficulties that prevent these ideas being laid down in an economic model are described briefly in Section 4.4.1.

As in Harris and Raviv (1993) and Kandel and Pearson (1995) I assume traders differ in their interpretation of news, but I also assume one type of trader is better/more confident when interpreting particular news items than another. The trader who considers himself better at interpreting a particular information release, trader i say, is more aggressive when updating his/her price (limit order) schedules, in that on announcement of 'good'

[5] See also Kavajecz (1999).

news, he/she increases the price of the asset by more than the other more uncertain trader, trader j. In the trading round, after observing each other's price schedules, the trader who is more confident when interpreting news, trader i, believes that trader j is underpricing the asset and so executes market buy orders. Trader j, on the other hand, believes trader i is overvaluing the asset and therefore executes market sell orders. Since trader i is more confident when interpreting public information than trader j, his/her conditional variance of the value of the asset is likely to be lower than that of the less certain trader. This causes trader i to want to trade more than trader j, i.e. trader i initiates more market buy orders than j wishes to initiate market sell orders. Therefore, trader i is likely to be more aggressive, not only when updating his/her price (limit order) schedules, but also when initiating market orders in the trading round. This generates the trading volume associated with public information releases, caused by differences in interpretation, but also causes positive (negative) order flow on announcement of good (bad) news, as found in Evans and Lyons (2003) and Chapter 3.

4.2.1 A simple model

Assume there are two assets, a risk-free asset, which is assumed to have a zero rate of return, and a risky asset, representing foreign exchange, FX. The end of period value of FX is denoted V and at the time of trading, this is a random variable. At the beginning of the period, the distribution of V is given by

$$V \sim N\left(\tilde{V}, \sigma_V^2\right) \tag{4.1}$$

and this is common knowledge across all traders. At date 0, all traders simultaneously observe an announcement of public information, S. This could be in the form of payroll employment figures, CPI or statements on the balance of trade. As in Harris and Raviv (1993) and Kandel and Pearson (1995) I assume the mapping of information to price is not common knowledge. In particular, I assume that a trader obtains a signal of the end of period value, V, from the announcement, but this signal is not perfect, i.e. the signal is observed with noise. Imagine there are two traders; i and j. i observes a signal $R^i = V + \epsilon^i$ where $\epsilon^i \sim N\left(0, \sigma_{\epsilon^i}^2\right)$ and j observes a signal $R^j = V + \epsilon^j$ where $\epsilon^j \sim N\left(0, \sigma_{\epsilon^j}^2\right)$. Without loss of generality, trader i is assumed to be more confident in interpreting the news announcement than trader j and this is shown by the variance of the subjective

Figure 4.1
Time Line of the Model and Sequencing of Events

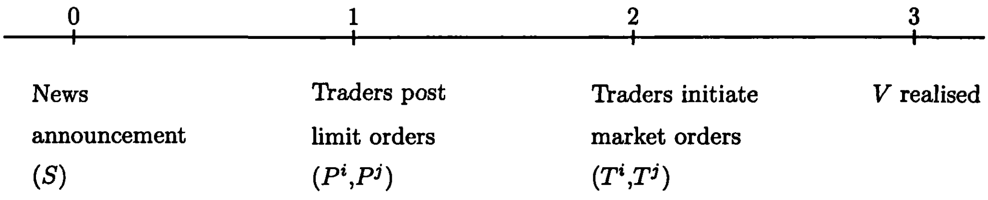

error in i's signal being less than that in j's signal, i.e. $\sigma^2_{\epsilon^i} < \sigma^2_{\epsilon^j}$.[6] Each trader is assumed to know their own type and also knows both parameter values, $\sigma^2_{\epsilon^i}$ and $\sigma^2_{\epsilon^j}$.

At date 1, after observing the announcement of public news and updating their beliefs of the value of the asset, both traders post schedules of limit order prices, denoted P^i for trader i and P^j for trader j. After posting their limit order price schedules, each trader, at date 2, initiates market orders, denoted T^i (T^j) for trader i (j). The market orders of trader i, T^i, are assumed to execute against the limit orders of the other trader, trader j, and vice versa. The sequencing of events is shown more clearly in Figure 4.1 below.

4.2.2 Limit order schedules

If traders centre their limit order quotes on their expectation of the value of FX then we can get an idea of what happens to traders' prices following the news release by examining how their expectations change.[7] Each trader has a common prior on the value of the asset, $\tilde{V}$, before the announcement of news.[8] When observing the public news announcement, each trader is assumed to update his/her beliefs using standard Bayesian methods. With

[6] I describe this situation as that where trader i is more confident than trader j. However, if $\sigma^2_{\epsilon^i} < \sigma^2_{\epsilon^j}$, this is the same as saying i is 'better' at interpreting news since he/she receives a signal with a smaller variance.

[7] I assume traders centre their quotes on their posterior expectation for simplicity. However, simple inventory models allow for a discrepancy between the posterior expectation and the mid-quote; an unexpectedly large inventory will cause traders to lower their quotes (relative to the expected value) in order to encourage market orders that will restore desired inventory levels. See, for example, the models used in Madhavan and Smidt (1991) and Lyons (1995). Assuming inventory disturbances before the news announcement are random, this will not affect the results in this simple setup. Also, with only two traders, i and j, it will almost certainly be the case that they act strategically and may not post limit order schedules exactly on their expected value. For simplicity, I assume this is not the case.

[8] The differences of opinion story told here, similar to Harris and Raviv (1993) and Kandel and Pearson (1995) is different to the Kim and Verrecchia (1991) explanation for why public information leads to greater trading activity. Kim and Verrecchia (1991) use a heterogenous priors story, or more specifically, traders have priors with different precisions. When the public information is released, traders who had very precise priors do not change their expectation of the value of the asset as much as those with very imprecise priors. Therefore the public information generates differences in opinions of the value of the asset and hence explains the increase in trading activity.

a signal R^i, trader i updates his/her prior, $\tilde{V}$, to obtain a posterior belief, denoted η^i. Trader j updates his/her beliefs similarly.

$$\begin{aligned} E\left[V|\Omega_1^i\right] &= E\left[V|\tilde{V}, R^i\right] = \eta^i = R^i\left(\frac{\sigma_{\epsilon^i}^{-2}}{\sigma_{\epsilon^i}^{-2}+\sigma_V^{-2}}\right) + \tilde{V}\left(\frac{\sigma_V^{-2}}{\sigma_{\epsilon^i}^{-2}+\sigma_V^{-2}}\right) \\ E\left[V|\Omega_1^j\right] &= E\left[V|\tilde{V}, R^j\right] = \eta^j = R^j\left(\frac{\sigma_{\epsilon^j}^{-2}}{\sigma_{\epsilon^j}^{-2}+\sigma_V^{-2}}\right) + \tilde{V}\left(\frac{\sigma_V^{-2}}{\sigma_{\epsilon^j}^{-2}+\sigma_V^{-2}}\right) \end{aligned} \tag{4.2}$$

η^i is therefore the expected value of FX for trader i, conditional on the information available at the time the quotes are posted, denoted Ω_1^i. η^i and η^j are both normally distributed with mean V and variances $\sigma_{\eta^i}^2$ and $\sigma_{\eta^j}^2$ respectively. Since trader i is more confident in interpreting the news than trader j, shown by $\sigma_{\epsilon^i}^2 < \sigma_{\epsilon^j}^2$, then i will be more aggressive when updating his beliefs, i.e. for any large positive signal, η^i will be greater than η^j in expectation. If trader j was very uncertain when interpreting the news, then the variance of the signal R^j will be very large, implying $\sigma_{\epsilon^j}^{-2}$ and the coefficient on R^j being very small. η^j will then be very close to the prior belief, $\tilde{V}$. On the other hand, when trader i is very confident when interpreting news, then $\sigma_{\epsilon^i}^{-2}$ will be large and hence η^i will tend to be quite different from $\tilde{V}$. If traders centre their quotes on their expected value of the asset then i will naturally be more aggressive when updating his prices than trader j.

4.2.3 Market orders and order flow

A standard functional form for risky asset demand is given in (4.3) below.

$$\begin{aligned} T^i &= \theta^i\left(\eta^i - P^j\right) + X^i \\ T^j &= \theta^j\left(\eta^j - P^i\right) + X^j \end{aligned} \tag{4.3}$$

T^i and T^j are the market orders (demands) that traders i and j wish to initiate. Positive (negative) T^i means market order purchases (sales) of FX which execute at the ask (bid) price of dealer j, P^j.[9] T^i is then a linear function of the difference between the expected

[9]The functional form shown in (4.3) is proved to be optimal if all random variables are normally distributed and if traders have CARA preferences. See the appendix to O'Hara (1995), Chapter 3, or the appendix to Lyons (2001), Chapter 4, for example.

value of the asset, η^i, and the price at which he/she can trade, P^j.[10] X^i and X^j are liquidity demand components, independent of V, and θ^i (θ^j) is a parameter which shows how aggressive i (j) trades for any discrepancy between i's (j's) expectation of the value of the asset and the price at which he/she can trade. θ^i is shown to be inversely related to the conditional variance of V. Therefore, if i is better at interpreting news than j, then at the time the market orders are executed, i will have a more precise expectation/lower conditional variance than trader j, implying $\theta^i > \theta^j$. Assuming i's price, P^i, is centred on his posterior expectation, η^i, and similarly for trader j, then it follows that i will wish to execute more market orders than j.

A release of good news will therefore generate positive order flow as well as an increase in trading volume. If traders not only differ in their interpretation of news but also differ in their *abilities* to interpret news announcements, then those who are more confident are more aggressive when updating their expectations and are also more aggressive when executing market orders in the trading round. Those who are unsure of the correct price implication following good news will tend to post a price close to their initial prior. Those who are more confident then think that they are underpricing the asset and therefore initiate market buy orders that transact against the relatively low limit prices. In a similar argument, those who are unsure of the correct price believe the other traders are overpricing the asset and therefore initiate sells. However, since those who are unsure of the mapping from information to price have a higher conditional variance of the value of the asset, they initiate fewer market orders (sells) than those who are more confident (buys). Therefore, on announcement of good news, positive order flow is generated.

Consider what happens if traders do not differ in their abilities to interpret news, i.e. the precision of each trader's signal is the same ($\sigma^2_{\epsilon^i} = \sigma^2_{\epsilon^j}$). In this case, both traders will be just as aggressive when updating their expectations since the coefficient on each trader's signal in (4.2) will be the same. In expectation, each trader will post the same mid-quote and there will be no systematic order flow following the news announcement. Trading volume will still increase, caused by non-zero realisations of ϵ^i and ϵ^j, which result in one trader believing the value of FX is higher than the other. However, there is no reason for

[10] The information sets available when traders post limit orders and when they execute market orders will be different, implying the conditional expectations, η^i in (4.2) and η^i in (4.3) will not be the same. When trader i executes market orders (date 2), he has seen j's limit order prices (posted at date 1) and can therefore use the information in j's price to update his expectation of V. The quotes of other dealers are an important source of information and this learning behaviour has been found to be common in FX markets. See Wang (2001) and de Jong, Mahieu, Schotman, and van Leeuwen (2001). For simplicity, and to help see the intuition, assume traders do not update their expectations between posting limit orders and executing market orders.

order flow to result since both traders will wish to execute equal (and offsetting) numbers of market orders (with equal conditional variances, neither trader has better/more precise information). In order to generate systematic/non-zero order flow, traders must differ in their *abilities* to interpret news.

It seems intuitive that some traders feel that they have greater abilities to interpret news than others. Different trading desks may have more accurate models of the economy or have larger research teams that can more accurately assess the implication of any news announcement. However, it should be noted that the intuition presented above is not one of trader *over*confidence. Traders do not overestimate their abilities to interpret news in this model, which would manifest itself in traders acting as though $\sigma^2_{\epsilon^i}$ (or $\sigma^2_{\epsilon^j}$) was lower than it actually was. In the above explanation, some traders simply receive more accurate signals from the same public announcement. Considering the myriad of factors that go into the 'black box' of exchange rate determination, it seems reasonable to suggest that some traders' boxes are not as black as others.

If traders differ in their interpretation of public news releases, this will generate trading volume, shown theoretically by Varian (1989) among others, and empirically by, inter alia, Berry and Howe (1994), Chaboud et al (2004) and in the results of Chapter 2. If traders differ in their *abilities* to interpret news, modelled very simplistically above by traders receiving signals with different precisions, then this should generate positive (negative) order flow following good (bad) news. This relationship is documented in Evans and Lyons (2003) and in Chapter 3. However, if traders receive noisy signals from public news announcements, then the problem of asymmetric information is also likely to increase, as shown in the models of Kim and Verrecchia (1994,1997). If it is possible for another trader to interpret the news announcement more accurately than you, then it will be optimal for you to post a wider spread and reduce your depth. This is tested in the next section.

4.3 Liquidity in the DEM/USD Market around a Scheduled Announcement of PPI Data

If traders receive noisy signals from the public news announcement, due to the interpretation errors, ϵ^i and ϵ^j, then dealers face the risk of being on the wrong side of a more informed trade (if the counter-party has interpreted the news more accurately for example). The increase in asymmetric information following the news release will naturally cause the depths of traders to fall; dealers are willing to trade smaller amounts for any deterioration in their bid and ask prices. As argued by Lee, Mucklow, and Ready (1993), one needs to know how both spreads *and* depths change in order to determine whether liquidity has fallen around periods of public news announcements; one cannot look at spreads in isolation. If a news release is associated with a higher spread, this could simply be due to a trader moving along their existing schedule of limit order prices, where they are willing to trade more of the asset. In which case one cannot claim that higher spreads are a result of reduced liquidity.

The effect of publicly released news on market depth has been difficult to ascertain, primarily because of a lack of data. Lee, Mucklow, and Ready (1993) find that the depths of specialist dealers in equity markets do fall around public news announcements, but thus far, little has been done in FX markets.[11] The one exception to this is the study by Carlson and Lo (2004). They study the effect of an increase in German interest rates on spot DEM/USD (Deutsche-marks per dollar) depth. At 11:30 GMT on 9th October 1997, the German Bundesbank decided to increase the repo rate from 3% to 3.3%. However, this announcement was unscheduled and came as a complete surprise to market participants. Carlson and Lo (2004) show that the announcement caused the exchange rate (mid-quote) to change and also caused depth to fall dramatically in the minute following the release. At 11:30 GMT the inside spread was 5 pips (the best bid in the market was 1.7530 DM per dollar and the best ask was 1.7535) and the spread for $10 million was 21 pips (0.0021 DM per dollar). In the following 35 seconds there was considerable trading with the transaction price increasing from 1.7530 to 1.7560 DM. At 11:30:35 the inside spread had increased to 39 pips and the spread for $10 million had increased to 52 pips. Even though an unexpected German interest rate rise should precipitate a fall in the price of the dollar, the US currency appreciated significantly in the initial 35 seconds post release. Carlson and Lo put this down to the existence of a large number of open short positions in the dollar, evidenced by a much deeper limit order schedule at the bid compared to the ask, i.e. there was greater buying support for the dollar at the time of the announcement. After the shock announcement of news, traders immediately tried to cover these short dollar positions by buying the currency (through market orders) which then led to its appreciation. However, in the following 25 seconds to 11:31:00, $86 million of liquidity at the bid were removed, $81 million of which were caused by seller initiated trades. Carlson and Lo suggest that traders did not remove their bid prices but simply let them be hit

[11]The data used so far in this thesis include information on the best quotes in the market. They do not contain any depth measures, even at the front end of the book.

because traders who still had short dollar positions tried to book their profits rather than posting much lower bid prices that might face the risk of not being executed. By 11:31:00 the depth of the market, especially at the bid, had collapsed, primarily because of seller initiated trades that had wiped out a large section of bid side liquidity. Therefore, depth (and also liquidity) on the Reuters D2000-2 limit order book had collapsed as a result of the public news.

However, Carlson and Lo's story relies heavily on the fact that the German interest rate announcement was unscheduled; the subsequent dynamics largely result from the fact that a number of traders held short positions and then tried to book profits by allowing existing limit orders to be hit. But what if the macroeconomic news announcement was scheduled? It is highly likely that leading up to a scheduled announcement of news, traders will close any positions they have.[12] If depth on the limit order book declines in the Carlson and Lo case because traders allow their existing limit orders to be picked off in an attempt to book profits from their short positions, then what happens if traders do not hold short positions at the time the data are released?

If traders interpret the same data release differently, modelled above by the inclusion of a stochastic 'interpretation' error (ϵ^i, ϵ^j) in the signal obtained from news, then the release of macroeconomic news will enhance the problems of asymmetric information. If you interpreted the news incorrectly, then you may find yourself on the wrong side of the resulting 'more informed' trades. To compensate for this, each trader will reduce the depth of their quotes, i.e. offer to buy (sell) less for any given deterioration of the bid (ask) price. One may then expect to find a reduction in the depth of the limit order book around an announcement of public news. It could well be the case that some individual traders pull out of the market altogether and wait until trading settles down following the news release. This too will cause the depth of the limit order book in the market to fall. However, as noted by Lyons (1997),

"... dealers who choose not to quote during regular hours are viewed as breaching the implicit contract of quote reciprocity, and are punished by other dealers (e.g., breaches are met with subsequent refusals to provide quotes, or by quoting large spreads)." Lyons (1997), page 282.[13]

[12]After conversations with FX traders, this does appear to be the case, at least for the 'large movers' such as the US employment report.

[13]Such 'punishment' is more likely in the direct inter-dealer market, such as the Reuters D2000-1 platform. However, in the brokered inter-dealer market considered in this thesis, where the quotes are anonymous, any punishment of those who do not follow the rules of the game may not be possible.

Whether 'regular hours' include periods of (often large) unanticipated macroeconomic news releases, is, however, an open question. It is also possible for depths to deteriorate in the minutes/seconds leading up to a pending news announcement. If traders know that a news item is to be released at a certain time, they may revise the depth of their quotes or pull out of the market before the news is released. For example, if it is possible for some traders to react more quickly than others, then those traders who effectively move second will be 'leaving money on the table' when they have their 'stale' quotes hit. In order to protect themselves from this risk, they should revise their quotes pre-release.[14] It is therefore the purpose of this section to examine what happens to the depth of the limit order book around periods of scheduled macroeconomic news announcements. Unfortunately, due to limited data on order book depth, what follows can only be considered as anecdotal evidence. Similar to Carlson and Lo (2004), I examine the effects of a single macroeconomic data release. However, the fundamental difference between this study and that of Carlson and Lo, is that I examine the effects of a scheduled macroeconomic release, whereas the timing of the German interest rate rise, studied by Carlson and Lo, came as a complete surprise to all market participants.

At this point it should be noted that the data used in this section, and also in Carlson and Lo (2004), are those from the limit order book and not of individual traders. Lee, Mucklow, and Ready (1993) and Kavajecz (1999) find that spreads of individual traders increase following news announcements, but if traders disagree on the mapping of information to price, then different traders may centre their quotes on different prices. This imperfect overlap of quotes may cause the 'touch', the difference between the best ask and best bid in the market, to fall, even though the spreads of individual traders may have risen. The net effect on the touch is therefore ambiguous; the increase in individual dealer spreads, combined with the negative effect on the touch caused by traders posting different mid-quotes, may result in a rise in the touch, as found in Chapter 2,[15] an unchanged market spread, as found in Morse and Ushman (1983), who consider earnings announcements on NYSE stocks, or a fall in the touch, as found in Acker, Stalker, and Tonks (2002) (London Stock Exchange). Even if the effects of differences in opinion on the touch are ambiguous, the effects on market depth are not. If individual traders reduce the depth of their quotes, then market depth will unambiguously fall. As such, even though we do not have data on the quotes of individual dealers, we can still test the hypothesis

[14]However, these quotes will still be centred on the same mid-quote since they do not know if the forthcoming news is good or bad.

[15]However, this increase in the touch was not found to be significant in a statistical sense.

that public data releases increase information asymmetry (resulting from differences of opinion/interpretation errors) by examining their effects on market depth.

4.3.1 Data

The dataset considered is exactly the same as that used in Carlson and Lo (2004), although the macroeconomic news considered is different. Carlson and Lo analyse the effect of a German interest rate increase at 11:30 GMT on Thursday 9th October 1997. However, the following day saw a release of scheduled US PPI data at 08:30 EST (12:30 GMT). A US Bureau of Labour Statistics report stated that the producer price index had increased by 0.5% in September 1997, compared to a consensus forecast of a 0.2% rise. The core index (which excludes food and energy prices) increased by 0.4% compared to an expected 0.2% increase.[16] Therefore US PPI increased unexpectedly in September 1997 and since US PPI news has been found to affect foreign exchange market activity (Ederington and Lee 1993, Almeida, Goodhart and Payne 1998 and the results of Chapters 2 and 3) then this announcement seems a prime candidate for examining the effects of scheduled news on order book depth. Although not large, this unexpected news component should be considered representative.[17]

The dataset used in this exercise is not the same as that described in Chapter 2. Elsewhere in this thesis, FX data from September 1999 to July 2000 are employed. However, these data only contain entries listing transactions and the best bid or ask prices. Therefore one is not able to determine what the depth of the market is like, i.e. the quantities of unfilled limit orders at prices away from the touch. Fortunately, an analysis of market depth can be made using a second, more comprehensive dataset from 1997. This gives every entry on the Reuters D2000-2 platform, and covers one week of DEM/USD trading and quoting activity from 6th to 10th October 1997. Features of the dataset are discussed at length in Daníelsson and Payne (2002a), Daníelsson and Payne (2002b) and Payne (2003a) and the reader is referred to these papers for a more detailed description. The dataset consists of every limit order entered onto the system and includes the time entered, the time it is removed,[18] the quoted price (at the bid or ask), the quantity associated with the quote ($

[16]The Bureau of Labour Statistics PPI reports are available at `http://stats.bls.gov/ppihome.htm`. The consensus forecasts were obtained from 'The Week Ahead' section of 'Business Week' (3rd October 1997). These (median) consensus forecasts are obtained from interviews with approximately 40 money managers and market practitioners held in the week before the announcement.

[17]From the ten months of US data releases analysed in Chapter 2 and given in detail in Table 2.3, this difference between actual and expected PPI is approximately equal to a one standard deviation announcement of news.

[18]As such, I can follow every individual order from when it is entered onto the system to when it is

Figure 4.2
Depths in the Spot DEM/USD Market Around the Announcement of US PPI at 12:30 GMT on Friday 10th October 1997

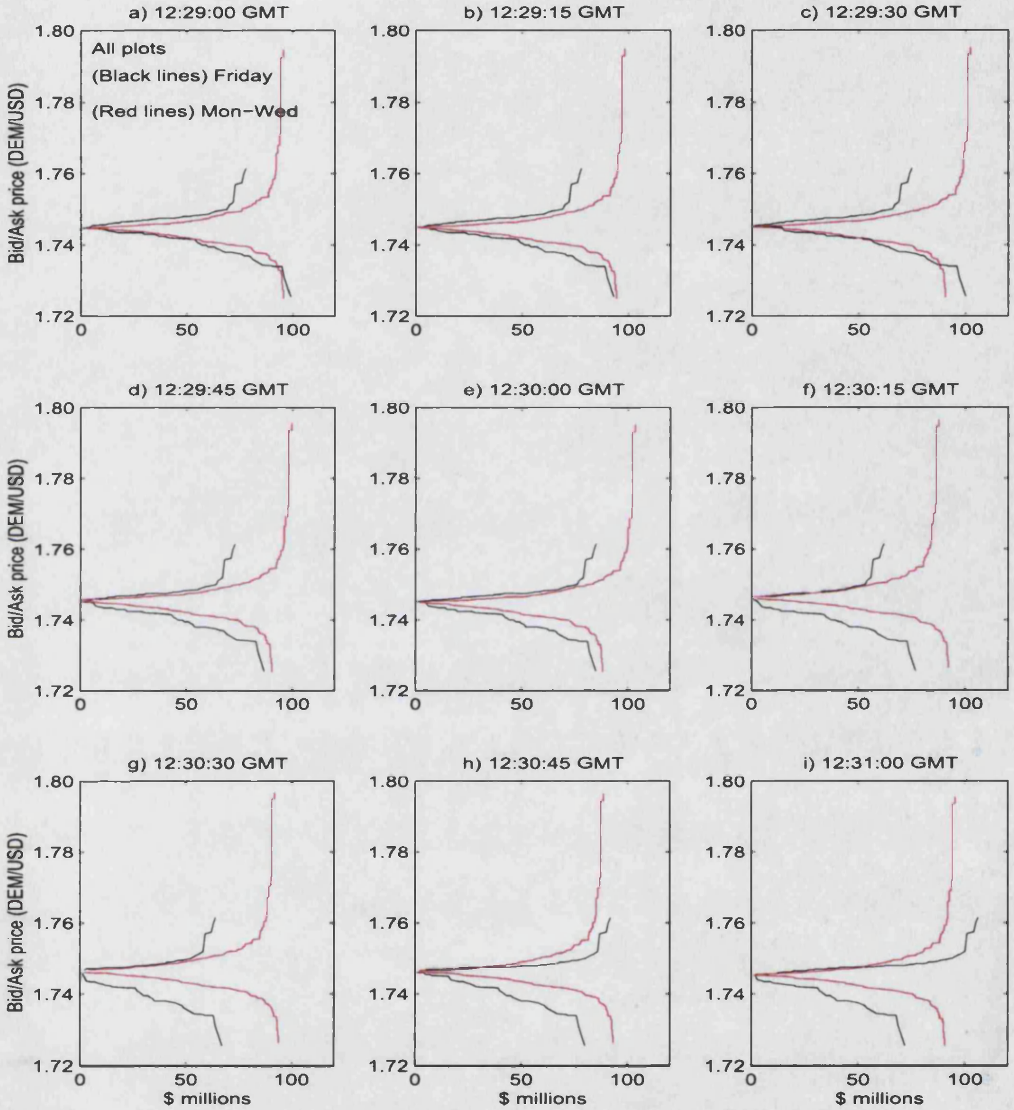

Notes: The figures plot the depths of the Reuters D2000-2 limit order book in the DEM/USD market, defined as the value of unfilled limit orders at the ask (upward sloping schedules) and bid (downward sloping schedules) sides of the market. Black lines show the depths on Friday (the day of the announcement), while the red lines show the average depth on Monday to Wednesday. These red lines show how depth in non-announcement times deteriorates *relative to the best prices on Friday*, hence showing how one would expect the depth to change as we move through the limit order book.

million) and the quantity traded, due to executed market orders, which are also included in the dataset.[19] Therefore the dataset allows me to construct the entire limit order book, defined by the quantity ($ millions) of unfilled limit orders at each price along the bid and ask schedules, at any point in time. Since the reaction to scheduled news has been found to be quite rapid (no more than two minutes in Chapters 2 and 3[20] then I examine the depths of the limit order book in 15 second intervals from 12:29:00 GMT (one minute before the scheduled announcement) to 12:31:00 GMT (one minute after the data release). Figure 4.2 shows the depth of the limit order book at these intervals for both the bid and ask sides of the market and in order to give a comparison to normal/non-announcement times, the average order book depths are also displayed, calculated from the order book at these times on Monday, Tuesday and Wednesday.[21]

The average order book depths at the bid and ask are calculated relative to the best prices seen on Friday 10th October and are computed as the average quantity ($ millions) of unfilled, tradeable limit orders at every price. For example, Figure 4.2a shows the depths of the spot DEM/USD market at 12:29:00 GMT. The best/lowest ask price (valid for $1 million) on Friday 10th October was 1.7447 DEM/USD and at a price of 1.7448 DEM/USD there were $3 million of unfilled limit orders at the ask. The top, black line therefore traces out how the value of unfilled limit sell orders increases as the price deteriorates. The average quantity available at the best ask price on Monday to Wednesday was $1.67 million. The average quantity available on these days at a 1 pip (0.0001 DEM/USD) worse price was $6.33 million and at a 2 pip worse price, the value of unfilled limit orders is $12.33 million. The top, red line therefore traces out how the value of unfilled limit sell orders changes as the price deteriorates in normal/non-announcement times.[22]

---

removed.

[19]Transactions can also occur due to crossed limit orders, whereby a limit order to buy is posted at a price that is higher than the current best ask, for example. The Reuters system automatically matches these quotes and the transaction occurs. For this reason, such limit orders are interpreted in exactly the same way as market orders.

[20]These are of course first moment effects. The second moment effects of macroeconomic news, i.e. volatility responses, last approximately an hour, see Ederington and Lee (1993), Andersen and Bollerslev (1998), Payne (1996) and also the results reported in Chapter 2.

[21]Thursday is excluded from the calculation because i) the effects of the German interest rate rise at 11:30 GMT lasted for more than an hour (Carlson and Lo 2004) and ii) US jobless claims figures were also released at 12:30 GMT. There were no *major* macroeconomic news announcements scheduled on Monday - Wednesday, as reported by Business Week. Consumer credit figures were released at 19:00 GMT on Tuesday 7th October and wholesale trade figures were announced at 14:00 GMT on Wednesday. These releases have not been found to affect FX markets significantly and the release times are also substantially later than 12:30 GMT.

[22]One criticism of this method of comparing Friday's depths to normal times is that it does not take into consideration how the inside spread changes on Friday. If the depths on Monday to Wednesday are all calculated *relative* to the best prices on Friday, then a large inside spread on this day will also imply a large inside spread on Monday to Wednesday, which may not be the case. However, as shown in Figure 4.2 the inside spread around the announcement on Friday 10th October was small and no greater

A more detailed description of how liquidity (defined by book depth) changes from 12:29:00 to 12:31:00 is also presented in Tables 4.1 and 4.2. This shows how the depth at both the bid and ask sides of the market are affected and also gives the reasons for the change in liquidity. Liquidity can be removed in any fifteen second interval for two reasons. Firstly, market orders can be initiated which cancel existing limit orders.[23] Secondly, traders can remove their existing unfilled limit orders. These liquidity draining commands on the Reuters D2000-2 trading platform are shown in Table 4.1, which splits the activity into bid and ask sides of the market and also shows the number of orders (market orders or removed limit orders) together with the value of removed liquidity. Table 4.2 shows how liquidity is added on both sides of the book in each fifteen second interval from 12:29:00 to 12:31:00. Liquidity is added by traders posting limit orders onto the book. However, due to the fast-paced activity of the book, many limit orders are only posted for a short period of time and if they are not hit by market orders within a couple of seconds, they are removed. Therefore, Table 4.2 shows the total number and value of limit orders posted onto the system and also shows how many of these are removed (voluntarily) before the end of the fifteen second interval. The last rows of Table 4.2 therefore show the net change in liquidity on both the bid and ask sides of the limit order book, calculated as the difference between 'Total liquidity added' and 'Total liquidity removed'.

4.3.2 One minute pre announcement

As shown in the first column of Tables 4.1 and 4.2, the first fifteen seconds of 12:29 GMT saw little activity. There was one trade (a sale of \$1 million) and another \$10 million of liquidity were removed through traders cancelling existing limit orders. However, \$8 million of liquidity were posted onto the book in this period (Table 4.2), explaining why the limit order schedules do not change significantly from 12:29:00 GMT (Figure 4.2a) to 12:29:15 GMT (Figure 4.2b). As these panels show, the depth of the book at 12:29:00 on Friday is not substantially different from the Monday to Wednesday average. The value of outstanding limit buy orders is \$99 million on Friday, compared with the Monday to Wednesday average of \$95.67 million. The value of outstanding limit sell orders is less than the Monday to Wednesday average (\$78 million compared to \$95.67 million) but

than usual. As shown in Chapter 2, Section 2.3.3.1, no day of the week effects are present in the spread. Therefore, comparing the depth on Friday to that seen earlier in the week should not distort the results.

[23]Alternatively, traders can cross limit orders, i.e. post a limit order to buy at a price higher than the current best limit sell price. The D2000-2 system matches these orders in exactly the same way as if the limit buy order was a market order. For this reason crossed limit orders are treated in the same way as traded market orders.

Table 4.1

Orders *Reducing* the Depth of the DEM/USD Limit Order Book Around a Release of US PPI Data

	Minute pre release				Minute post release			
	12:29:00 -12:29:14.99	12:29:15 -12:29:29.99	12:29:30 -12:29:44.99	12:29:45 -12:29:59.99	12:30:00 -12:30:14.99	12:30:15 -12:30:29.99	12:30:30 -12:30:44.99	12:30:45 -12:30:59.99
Liquidity removed								
Market orders								
Buys	- ($0m)	4 ($5m)	5 ($6m)	- ($0m)	7 ($16m)	12 ($17m)	4 ($13m)	2 ($2m)
Sells	1 ($1m)	- ($0m)	- ($0m)	- ($0m)	4 ($6m)	7 ($17m)	1 ($1m)	9 ($15m)
Total	1 ($1m)	4 ($5m)	5 ($6m)	- ($0m)	11 ($22m)	19 ($34m)	5 ($14m)	11 ($17m)
Removal of limit orders								
Ask	3 ($4m)	1 ($2m)	1 ($1m)	1 ($1m)	6 ($7m)	3 ($3m)	1 ($1m)	7 ($9m)
Bid	3 ($6m)	5 ($10m)	5 ($15m)	2 ($2m)	6 ($6m)	6 ($9m)	2 ($3m)	8 ($13m)
Total	6 ($10m)	6 ($12m)	6 ($16m)	3 ($3m)	12 ($13m)	9 ($12m)	3 ($4m)	15 ($22m)
Total liquidity removed								
Ask side	**($4m)**	**($7m)**	**($7m)**	**($1m)**	**($23m)**	**($20m)**	**($14m)**	**($11m)**
Bid side	**($7m)**	**($10m)**	**($15m)**	**($2m)**	**($12m)**	**($26m)**	**($4m)**	**($28m)**
Total	**($11m)**	**($17m)**	**($22m)**	**($3m)**	**($35m)**	**($46m)**	**($18m)**	**($39m)**

Notes: At 12:30:00 GMT on Friday 10th October 1997, a Bureau of Labour Statistics report announced that US PPI had increased by 0.5% in September. (The consensus forecast was for a 0.2% rise). Table 4.1 shows the number and value of orders entered onto the Reuters D2000-2 system that drained liquidity in each of the fifteen second intervals from 12:29:00 to 12:31:00. Market buy orders execute at limit ask prices, i.e. a trader can only buy at the price another trader is willing to sell. Therefore total liquidity removed at the ask side is equal to market buy orders plus removed limit ask orders. Table 4.2 shows the number and value of quotes posted onto the system. Due to the fast paced nature of quote activity at this time, a large number of quotes posted to the system were subsequently removed within seconds. Therefore, 'Total liquidity added' is calculated as the value of quotes added, less those that were removed before the end of the same fifteen second period.

Table 4.2

Orders *Increasing* the Depth of the DEM/USD Limit Order Book Around a Release of US PPI Data

	Minute pre release				Minute post release			
	12:29:00 -12:29:14.99	12:29:15 -12:29:29.99	12:29:30 -12:29:44.99	12:29:45 -12:29:59.99	12:30:00 -12:30:14.99	12:30:15 -12:30:29.99	12:30:30 -12:30:44.99	12:30:45 -12:30:59.99
Liquidity added								
Limit orders								
Ask	4 ($5m)	3 ($4m)	5 ($6m)	8 ($16m)	2 ($3m)	17 ($26m)	29 ($45m)	21 ($42m)
Bid	2 ($3m)	9 ($23m)	5 ($12m)	- ($0m)	8 ($13m)	21 ($62m)	11 ($17m)	9 ($20m)
Total	6 ($8m)	12 ($27m)	10 ($18m)	8 ($16m)	10 ($16m)	38 ($88m)	40 ($62m)	30 ($62m)
Less orders removed								
Ask	- ($0m)	2 ($3m)	2 ($3m)	1 ($3m)	2 ($3m)	5 ($6m)	9 ($12m)	5 ($13m)
Bid	- ($0m)	2 ($7m)	4 ($11m)	- ($0m)	6 ($10m)	20 ($61m)	- ($0m)	7 ($14m)
Total	- ($0m)	4 ($10m)	6 ($14m)	1 ($3m)	8 ($13m)	25 ($67m)	9 ($12m)	12 ($27m)
Total liquidity added								
Ask side	**($5m)**	**($1m)**	**($3m)**	**($13m)**	**($0m)**	**($20m)**	**($33m)**	**($29m)**
Bid side	**($3m)**	**($16m)**	**($1m)**	**($0m)**	**($3m)**	**($1m)**	**($17m)**	**($6m)**
Total	**($8m)**	**($17m)**	**($4m)**	**($13m)**	**($3m)**	**($21m)**	**($50m)**	**($35m)**
Net change								
Ask side	**(+$1m)**	**(-$6m)**	**(-$4m)**	**($12m)**	**(-$23m)**	**($0m)**	**($19m)**	**($18m)**
Bid side	**(-$4m)**	**($6m)**	**(-$14m)**	**(-$2m)**	**(-$9m)**	**(-$25m)**	**($13m)**	**(-$22m)**
Total	**(-$3m)**	**($0m)**	**(-$18m)**	**($10m)**	**(-$32m)**	**(-$25m)**	**($32m)**	**(-$4m)**

Notes: See Table 4.1.

this is unlikely to represent a significant reduction; the value of outstanding limit orders on Wednesday was only $77 million.[24] The story is similar between 12:29:15 and 12:29:30 in that equal values of liquidity were withdrawn and added ($17 million). However, there was a slight shift of liquidity to the bid side of the book. Between 12:29:30 GMT and 12:29:45 GMT there is a definite net fall in liquidity on the order book. $22 million of liquidity were removed ($6 million through market orders and $16 million through cancelled limit orders) while only $4 million of liquidity were added. However, in the fifteen seconds up to 12:30:00 there is a net increase in liquidity of $10 million and as shown in Figure 4.2e, the value of outstanding limit orders on the bid and ask sides of the book at 12:30:00 do not look substantially different from the Monday to Wednesday average. Therefore, liquidity of the book (defined by depth) at the time the US PPI news was released does not appear to have collapsed compared with normal/non-announcement times. However, it is still true that spreads were wider along the entire order book, shown by the bid/ask schedules for the Monday to Wednesday average lying within the schedules of Friday. As shown in Table 4.3, the spread for $10 million was 22 pips (0.0022 DEM/USD) on Friday compared to the Monday to Wednesday average of just 8 pips. The spreads for $40 million were 58 and 36 pips for Friday and Monday to Wednesday respectively.

In contrast to the findings of Carlson and Lo (2004), the order book appears reasonably symmetric at the time the news is about to be released, suggesting that traders cover their long/short positions prior to the data announcement. At the time of the German interest rate rise, Carlson and Lo find that depth on the bid side of the book is greater than that on the ask side, suggesting that there were a number of traders with open short dollar positions. However, at the time of the release of US PPI, which was known in advance by traders, the depth on the bid side was equal to that on the ask ($85 million). Any short dollar positions before 12:30:00, indicated by a greater value of outstanding bids than asks in the first four panels of Figure 4.2, were covered by traders initiating market buy orders in the thirty seconds between 12:29:15 GMT and 12:29:45 GMT.[25] The second and third columns of Table 4.1 show 9 buyer initiated trades amounting to $11 million, while no seller initiated trades occurred. It therefore appears as though traders covered their short positions prior to the PPI release, as one would expect if unexpected PPI data moves the market, increasing the risk of holding short or long positions.

[24]Due to the small dataset, it is impossible to make statistical inference on the difference between one observation (Friday) and the average of three observations (Monday to Wednesday).

[25]Unfortunately, data on inventory positions for individual traders are not available. The link between more unfilled bids than asks and net short dollar positions is, at this point, an untested conjecture but seems sensible in the light of the discussion in Carlson and Lo (2004).

Table 4.3
Spreads in the Spot DEM/USD Market Around a Release of US PPI

Book Depth		Time (GMT)				
		12:29:00	12:29:30	12:30:00	12:30:30	12:31:00
$5m	Friday	7	6	6	37	18
	(Mon-Wed)	(3)	(4)	(5)	(10)	(4)
$10m	Friday	13	10	22	46	26
	(Mon-Wed)	(4)	(5)	(8)	(11)	(7)
$20m	Friday	30	33	37	58	41
	(Mon-Wed)	(12)	(15)	(15)	(20)	(14)
$40m	Friday	51	50	58	106	93
	(Mon-Wed)	(31)	(36)	(36)	(43)	(36)
$60m	Friday	91	93	107	235	136
	(Mon-Wed)	(66)	(66)	(65)	(80)	(72)

Notes: Spreads are defined as the number of pips (units of the fourth decimal) between the best ask and bid prices in the market. 1 pip = 0.0001 DEM/USD. Spreads are given for a representative sample of book depths on the Reuters D2000-2 trading platform, given in $ millions.

However, the trading that occurred in the minute pre release to cover short positions was not large compared to normal times. Figure 4.3 shows the intra-day pattern of the number (top panel) and value (bottom panel) of trades occurring on the Reuters D2000-2 system. The average value of trades on Monday to Wednesday in afternoon trading was approximately $30 million per minute and equalled $20.67 million for the minute starting at 12:29:00. On Friday 10th October there were only ten trades amounting to $12 million in this minute.[26] It is also interesting to note that there were no market orders initiated in the fifteen seconds prior to 12:30:00, suggesting that traders, once covering their short positions, wait to see what the news announcement is before deciding how to react.

4.3.3 One minute post announcement

The announcement that PPI had increased by 0.5% in September 1997, compared to consensus forecasts of a 0.2% rise, came at 08:30 EST (12:30 GMT). As shown in the fifth column of Table 4.1, the news was immediately followed by a wave of trading activity and

[26]However, as noted before, statistical testing of this difference is impossible due to the small sample size. This evidence should then be considered anecdotal. There may also be 'day of the week' effects, which may account for the low trading volume at 12:29 (see Chapter 2, Section 2.3.3.1, although such effects were found to be small). However, there is nothing we can do about this problem in this exercise.

Figure 4.3
Intra-day Patterns of the Number and Value of Trades in the Spot DEM/USD Market (per minute)

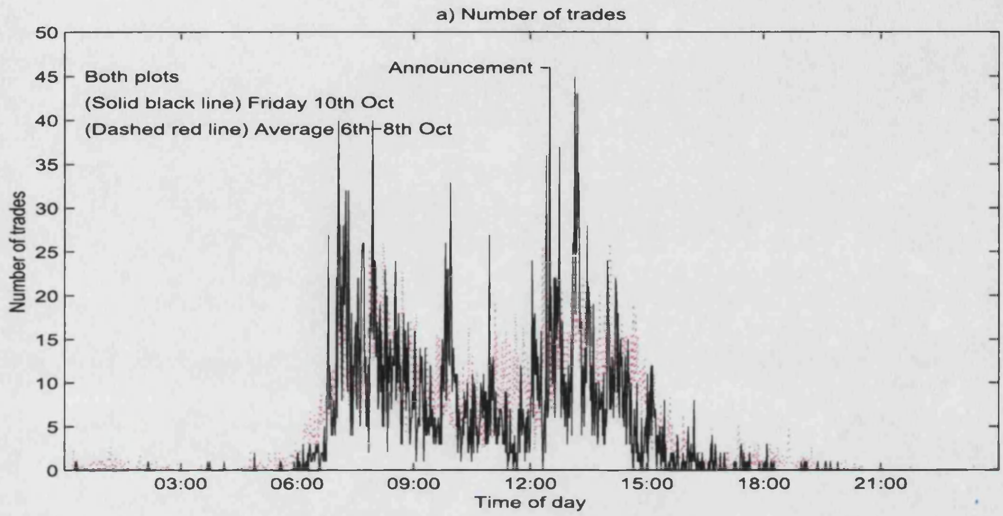

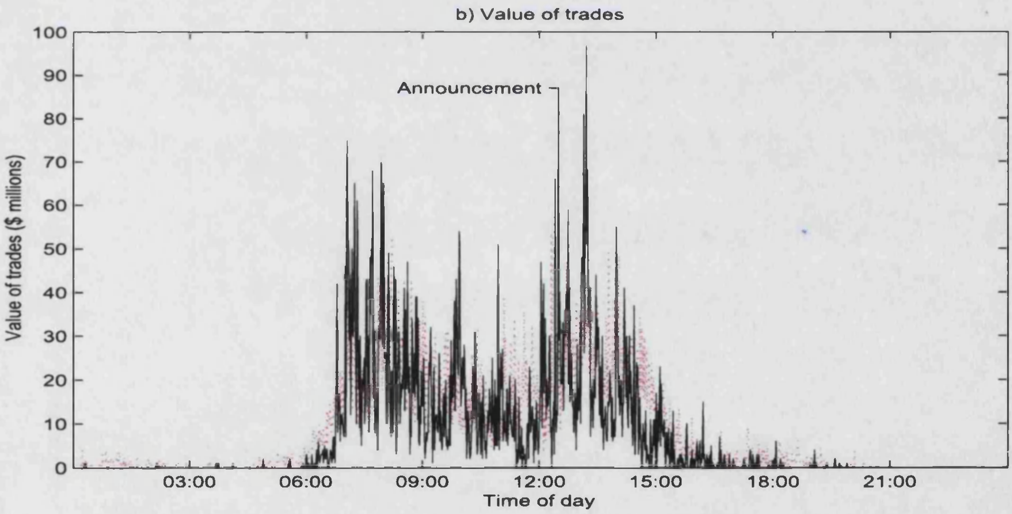

Notes: The figures show the intra-day patterns of the number (top panel) and value (bottom panel) of trades in the spot DEM/USD market, sampled at the 1 minute frequency. The solid black line shows the intra-day pattern for Friday (the day of the PPI announcement), while the dashed red line shows the pattern for Monday to Wednesday, hence allowing us to see what the trading activity is like in normal/non-announcement times.

repositioning of the limit order book. Even though the news release was of an unexpected increase in US PPI, the effect on the level of the exchange rate was minimal; Figure 4.4 shows that the prices on the Reuters D2000-2 trading platform were not significantly affected by the news. There is a spike at 1.7468 DEM/USD at 12:30:10 but the rate quickly falls back to levels seen pre release. In the fifteen seconds between 12:30:00 GMT and 12:30:15 GMT there were $22 million of market orders ($16 million buyer initiated and $6 million seller initiated). An additional $13 million of liquidity were also removed from the book as a result of traders cancelling existing limit orders. Despite $16 million of liquidity being posted to the book in these fifteen seconds, $13 million of these were voluntarily removed by 12:30:15. This implies a net fall in liquidity of $32 million in the fifteen seconds post release, ($22 million + $13 million - ($16 million - $13 million)).

This draining of liquidity continued in the following fifteen seconds. Between 12:30:15 and 12:30:30 there were $34 million of market orders split evenly between buyer and seller initiated trades. $12 million of existing limit orders were also cancelled, implying $46 million of liquidity being drained in this fifteen second interval. This was, however, accompanied by traders posting huge quantities of limit orders onto the book. $88 million of new quotes were posted ($62 million of which were at the bid side), but $67 million worth of these were removed before 12:30:30, implying only $21 million of liquidity being added. Therefore, in the second fifteen second period post release, there was a net fall in liquidity of $25 million, ($46 million - $21 million = $25 million). Combined with the $32 million fall in the initial fifteen seconds, this caused the bid and ask schedules to widen by 12:30:30, as shown in Figure 4.2g. The total value of unfilled limit orders at this time had fallen to $64 million at the ask and $67 million at the bid. Not only this, market spreads had also deteriorated. The spread at $10 million was 46 pips, four times greater than the normal/non-announcement value of 11 pips, while spreads for $40 million was 106 pips, compared to a normal value of 43 pips, see Table 4.3.

It was not simply the case that depths in the market had fallen because market orders were transacting against existing limit orders. Traders did have the chance to replace their executed limit orders, shown by the fact that $16 million of quotes were entered in the first fifteen seconds and $88 million of quotes were posted in the subsequent fifteen seconds. Of these quotes, however, $13 million and $67 million respectively were removed voluntarily within a few seconds. Therefore I suggest that traders reacted to the news announcement by deliberately reducing their depths in an attempt to protect themselves from traders who may have interpreted the news more accurately. The reduction in market depth was

Figure 4.4
Best Bid and Ask Prices in the Spot DEM/USD Market in the Hour Interval Containing the US PPI Announcement

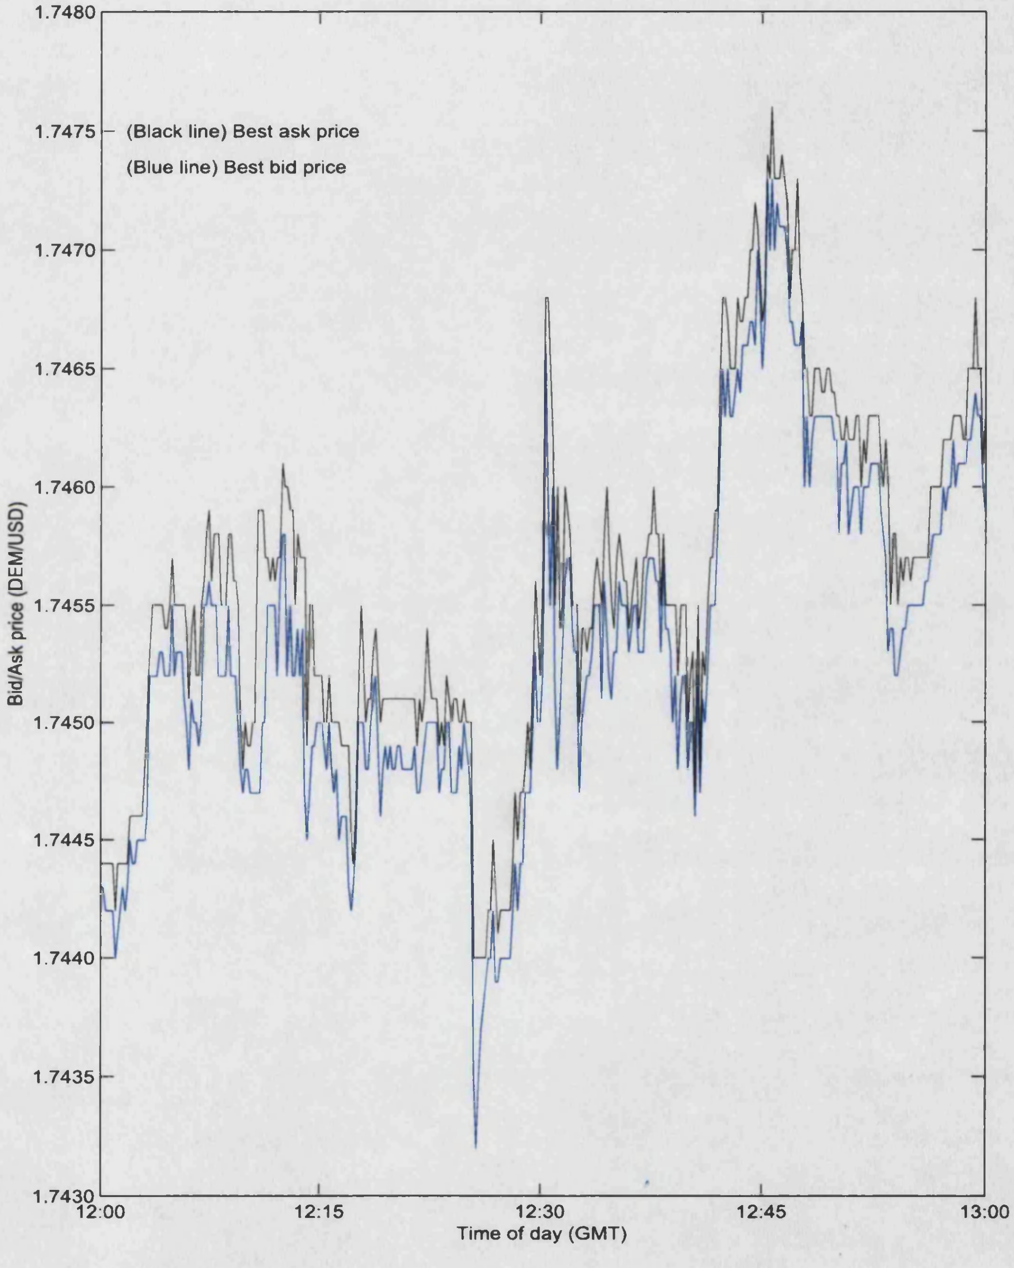

Notes: The figure shows the best bid and ask prices on the Reuters D2000-2 platform for the DEM/USD. The announcement of US PPI came at 12:30 GMT.

not simply a result of market orders picking off limit orders, which would naturally lead to a hollowing of book depth, but instead, the reduction in depth was an equilibrium market response.

In the following thirty seconds, between 12:30:30 GMT and 12:31:00 GMT, trading reduced slightly but remained high compared to non-announcement times. In the fifteen seconds between 12:30:30 and 12:30:45, \$18 million of liquidity were removed (\$14 million of which were through market orders) but \$50 million of liquidity were added, mostly at the ask side of the market. Market depth therefore increased in this fifteen second interval but spreads for large trade sizes still remained quite high, see Figure 4.2h.

Ultimately, the PPI announcement had little effect on the DEM/USD price. Despite huge trading activity in the minute following the news, \$87 million of trades compared to normal quantities of approximately \$30 million (see Figure 4.3), the order flow was small; net dollar buying amounting to only \$9 million. The best ask price at 12:30:00 was 1.7452 DEM/USD and this had only increased to 1.7455 DEM/USD by 12:31:00, implying a return of just 1.7 basis points. Even though the order flow, return and inside spread were not affected by the news release, trading volume and market depth did appear to change, and in ways consistent with the idea that traders differ in their interpretation of news. From Figure 4.4 it is clear that the unanticipated increase in the PPI figures had little effect on the value of the DEM/USD exchange rate. After an initial period of hectic trading, dealers appeared to accept that the news was not going to move the market.

4.4 Discussion

Following releases of public information, trading activity in financial markets has been found to increase relative to non-announcement times. This has often been put down to trader heterogeneities, such as differences in preferences (risk aversion), endowments or information, specifically the interpretation of information. If the mapping of information to price is not common knowledge, then an announcement that US PPI has increased unexpectedly by 0.3%, for example, is likely to lead to traders having different expectations of the 'fundamental' value of dollar denominated assets, including the currency itself. As in the models of Varian (1989), Kim and Verrecchia (1991), Harris and Raviv (1993) and Kandel and Pearson (1995), differences in the interpretation of news can easily lead to increased trading volume. However, what is not so clear is why such news generates signed order flow. Evans and Lyons (2003) and the results of Chapter 3 show that following a release of good (bad) news, order flow becomes positive (negative). Differences in interpretation can generate trading volume but systematic order flow is more difficult to explain.

In this chapter I present a simple explanation why public news releases generate signed order flow. If traders not only differ in their interpretation of the news (modelled by idiosyncratic interpretation errors) but also differ in their abilities to interpret news (modelled by traders receiving signals with different precisions), then those who are better/more confident in interpreting news will be more aggressive when updating their expectations and also when executing market orders. This, I suggest, will generate positive (negative) order flow following releases of good (bad) news.

If traders receive noisy signals from a news announcement then there is the possibility that a trader has interpreted the news incorrectly and may therefore find themselves on the wrong side of a more informed trade. Traders may naturally increase their spreads and reduce their depths in response to the data release. This is tested in Section 4.3, which analyses the depth of the DEM/USD market around a scheduled release of US PPI data.

The results, although anecdotal in nature, are consistent with the idea that traders differ when interpreting public news. Following a scheduled release of macroeconomic data, the depth of the limit order book decreases and trading volume (executed market orders) increases. However, the depth of the book decreases as an equilibrium market response, rather than simply due to market orders picking off 'stale' quotes, which would naturally drain liquidity. Traders can (and do) update their quotes and post limit orders onto the book following the announcement but due to the possibility of a trader interpreting the news incorrectly, the asymmetric information component of the spread increases and the optimal decisions of traders is to reduce their depths.

4.4.1 Technical challenges

In this chapter I have presented a simple explanation why public news releases generate, not only trading volume, but also signed order flow, consistent with the results of Evans and Lyons (2003) and with those presented in Chapter 3. However, there are a number of features, peculiar to the foreign exchange market, that prevent analytical solutions for order flow from being obtained. In particular, there is a fundamental difference between the microstructure setup that is common in the literature and the workings of the FX market. In the foreign exchange market, when forming the trader's wealth constraint, from which his/her optimal demands are derived, the end of period wealth will not only depend on the trades (market orders) that he/she initiates (at the limit order prices of other traders) but will also depend on the market orders that other traders initiate at his/her limit order prices. In previous models of quote driven markets such as Glosten and Milgrom (1985) and Madhavan (1992), the actions of the liquidity suppliers (in the form of one or more market makers) and those of the liquidity drainers (in the form of informed or liquidity traders) are completely separated.[27] Even in models where traders have a choice between limit orders and market orders, such as Cohen, Maier, Schwartz, and Whitcomb (1981), Handa and Schwartz (1996), Parlour (1998) and Foucault (1999), the decision is whether to add *or* remove liquidity.[28]

This, however, is not an accurate description of inter-dealer foreign exchange trading. As discussed in Lyons (1995) and Evans (2002), not only can a foreign exchange dealer trade by posting limit orders, he/she can also trade by executing market orders and lay off inventory at other dealers' prices. A trader in the inter-bank foreign exchange market is then both a potential liquidity supplier and a potential liquidity drainer/taker. The optimal demands (market orders) and limit order schedules when traders both supply *and* drain liquidity may therefore differ from the optimal schedules when traders post *either* limit *or* market orders *only.*[29]

To overcome this problem, Lyons (1997) and Evans and Lyons (2002b) present models that allow dealers to trade simultaneously. However, I choose not to use the more recent order flow model of Evans and Lyons (2002b), primarily because one needs to model the reactions of a number of traders individually if one is to explain both buyer *and* seller initiated trades following a release of public information. The simple explanation offered above considers the prices of different traders rather than looking at the market clearing price, which may be significantly different from the prices of individual dealers following a large shock to traders' information sets, especially if the public news is interpreted differently. In the model of Evans and Lyons (2002b) each dealer quotes a single price at which

[27]This is also the case in the various extensions that model limit order book markets with a number of market makers/liquidity supplying traders, such as Ho and Stoll (1983), Dennert (1993), Glosten (1994), Biais, Martimort, and Rochet (2000) and Viswanathan and Wang (2002).

[28]The choice between market or limit orders is also considered in empirical work by Harris and Hasbrouck (1996) and Luo (2002a) to name but two.

[29]The market order schedules in (4.3) are shown to be optimal only if the trader executes market orders and therefore may be invalid for traders who both supply and drain liquidity. One model that does allow traders to post limit *and* market orders is that of Chakravarty and Holden (1995). In this model, the informed trader posts limit orders as a form of insurance, to soak up any of his/her market orders that would otherwise transact against others' uncompetitive limit orders further down the book.

he/she is willing to trade any amount. In equilibrium all traders have to quote identical prices in order to prevent infinite arbitrage. Evans (2001) introduces a model where each trader can post different prices, arguing that in equilibrium there is a *distribution* of prices rather than a singleton price. No trader can exploit the arbitrage opportunities that result from the different prices of dealers due to the lack of transparency that characterises direct inter-dealer trading, the type of trading which is explicitly modelled in Evans' paper. However, Evans (2001) uses an overlapping generations type model to capture the FX dynamics, whereby in any trading round one half of the traders quote prices and the other half initiate trades. These roles are reversed in the subsequent trading round, and so the process continues. This approach is not chosen here since I wish to allow all traders to respond to the news at the same time and in the same order, as this is more likely to correspond to reality, i.e. after observing the public information all traders first change their prices and then trade by posting market orders. It may well be true that some traders respond to news events quicker than others, in which case some traders execute market orders that transact against others' 'stale' limit prices. However, in the simple model presented above, there is no 'game of snap' element to trading, i.e. the results do not rely on one trader being able to react faster than another.[30]

Even if one were to assume that the liquidity supplying and draining functions of traders were completely segregated, an approach commonly made in the literature - even in the FX market where it is least applicable,[31] then technical challenges still remain. As found in Wang (2001) and de Jong, Mahieu, Schotman, and van Leeuwen (2001), dealers learn from the quotes posted by other traders. Knowing that their quotes are giving information to competing traders for free, it appears reasonable to suggest that traders act strategically rather than centering their quotes on their posterior expectations. For these reasons, I do not build an economic model with thorough micro-foundations but simply describe the intuition why news generates systematic order flow. The ideas are based on traders not only differing in their interpretation of news but also differing in their *abilities* to interpret news.

One model that does allow heterogeneity in traders' information sets and which is based

[30]Following the release of public information, the optimal strategy for FX traders is to update their quote (limit order) schedule and trade at available prices of the other traders by executing market orders (if desired) as quickly as possible. This can be done very easily on the electronic trading systems available to dealers, such as EBS and Reuters D2000-2, simply by pressing a (number of) button(s) on a control panel. However, it makes sense that one dealer will press the button(s) quicker than another, even if by a second or so. In this way, the dealer who presses the button first is the dealer who wins (the game of snap) and the dealer who presses the button second, loses.

[31]See Derviz (2002) for example.

on thorough micro-foundations is that of Brennan and Cao (1997). In the context of the home-bias puzzle, they explain why foreign investors purchase domestic assets when the return on these assets has been relatively high. They suggest that domestic residents have an informational advantage over foreign investors and therefore have a more precise prior of the value of domestic assets, similar to the model of Kim and Verrecchia (1991).[32] If investors update their expectations using Bayesian methods, then,

"(Following good news) *the less well informed (i.e. foreign) investors revise the means of their distribution by more than do better-informed (i.e. local) investors. This implies... the less well informed foreign investors purchase more of the domestic market portfolio from the better informed domestic investor."* Brennan and Cao (1997), pages 1854-5.

However, despite the model of Brennan and Cao (1997) explaining *international portfolio flows*, it is unlikely to explain the systematic effects of news on *order flow*. Even though Brennan and Cao find that good news for domestic assets results in a net outflow of those assets to foreigners, this simply reflects a shift in relative demands between domestic and foreign investors. This is not order flow! For example, if the domestic residents initiated the trades and sold to the foreign investor, the news will be associated with negative flow. On the other hand, if the foreign investors initiated the trades, then order flow will be positive. The net outflow of assets to foreign investors can then be accompanied by either positive or negative order flow. This is a common criticism of the rational expectations models, such as Grossman (1976) and Grossman and Stiglitz (1980), around which the model of Brennan and Cao (1997) is based. Such models do not allow trades to be signed and hence can never explain order flow.[33]

4.5 Conclusions

The significant increase in trading volume around periods of publicly announced information has been explained in the literature by traders differing in their interpretation of news, i.e. the mapping of public information to price is not common knowledge. However, despite being able to explain the rise in trading volume, such stories cannot explain the systematic effects of news on order flow that have been found in Evans and Lyons (2003) and in the results reported in Chapter 3. As in the models of Glosten and Milgrom (1985), Easley and O'Hara (1987) and Perraudin and Vitale (1996), order flow is the mechanism

[32]As opposed to the above explanation where the more confident traders receive more precise *signals*.
[33]See the excellent discussion in Lyons (2001), especially Chapter 4.

by which *private* information is incorporated into price. Why *public* information is assimilated via this channel has been more difficult to explain.

In this chapter I present a simple explanation why good (bad) news for a currency results in positive (negative) order flow. As in Harris and Raviv (1993) and Kandel and Pearson (1995) I assume traders differ in their interpretation of news so that they receive a noisy signal from the commonly observed announcement. However, I also assume traders differ in their abilities to interpret news so that traders receive signals with different precisions. Those who are better at interpreting the news will be more aggressive when updating their quotes and will also be more aggressive when executing market orders in the trading round. Following good news, confident traders initiate more market buys than the less well informed initiate market sells, therefore resulting in positive order flow.

If traders differ in their interpretation of public news releases, then they face the risk of being on the wrong side of a more informed trade if another trader has interpreted the news more accurately. Traders should naturally reduce their depths following a news release and this is indeed found when examining the effects of a scheduled release of US PPI news. Market depth is shown to deteriorate, especially in the first 30 seconds following the release. This is an equilibrium market response and is not simply due to market orders picking off stale quotes.

Chapter 5 Feedback Trading

5.1 Introduction

Order flow, one way buying or selling pressure, has a contemporaneous impact on prices, while at the very highest frequencies, the converse is not true.[1] However, when aggregated over time, order flow and prices can be expected to impact on each other simultaneously, a phenomenon I call *contemporaneous feedback trading.* The existence of contemporaneous feedback trading implies that such models cannot be estimated with traditional techniques. This is unfortunate, since empirical models with feedback trading can be expected to give a more accurate assessment of theoretical models than models that ignore such trading strategies. In this chapter I argue that feedback trading is an inevitable consequence of time aggregation of order flow models, and propose an estimator of such models by using instrumental variable techniques. I find that the price impact of trades is much stronger when feedback trading is incorporated, further supporting market microstructure theories generally, and the validity of the order model specifically.

It is well known in the theoretical microstructure literature that order flow conveys private information to the market as a whole. In this way, information is aggregated via the trading process implying that order flow has permanent effects on prices. See Kyle (1985), Glosten and Milgrom (1985), Easley and O'Hara (1987), and Evans and Lyons (2002b) for examples of such models. These models imply that if trades carry private information then the informativeness of trades can be accessed by their price impact. Furthermore, in empirical models when data are employed at the very highest frequencies, order flow, by definition, can only be affected by the lags of price changes. However, when data are aggregated, transactions and order entry are simultaneous, frustrating empirical investigations.

The existence and profitability of feedback trading strategies has been considered in a number of papers. De Long, Shleifer, Summers, and Waldmann (1990) build a model of feedback trading with rational speculators who will buy (sell) when the price rises (falls).

[1]This chapter is a revised version of Daníelsson and Love (2004), which has been submitted to the Journal of International Money and Finance.

The profitability of a number of feedback trading strategies in stock markets is considered in Jegadeesh and Titman (1993) and the existence of high frequency positive feedback trading in the US treasury market is documented in Cohen and Shin (2003).

The most common methodology for empirically assessing the informativeness of order flow, is the vector autoregressive (VAR) model of Hasbrouck (1991a). His model was originally applied to data at the tick–by–tick frequency, where the direction of causality runs explicitly from order flow to asset price returns. Hasbrouck introduces a shock to the trading process, representing private information, and computes the cumulated effect on the asset return. The greater the cumulated effect, or impulse response, the more information trades are argued to carry. These VAR models have become standard in the microstructure literature; recent examples include Dufour and Engle (2000) and Engle and Patton (2004) for stocks, Evans (2002) and Payne (2003a) for currencies and Cohen and Shin (2003) and Green (2004) for treasuries. In the VAR framework, the asset returns in period t are regressed on contemporaneous order flow (date t) as well as lagged returns and order flows (dated $t-1$ or earlier), whereas order flows are only regressed on lagged returns and flows; order flows at date t do not depend on contemporaneous asset returns.

These models therefore rule out contemporaneous feedback trading, an assumption which is overly restrictive when the data are sampled at anything other than at the highest frequencies. If traders have the ability to respond to price changes and trade before the end of the time interval used in the empirical investigation, then order flows can indeed cause asset price changes within that period, but these price changes may then feed back into order flows in the same interval. A shock to order flows within period t may indeed cause a change in the asset price within that interval. However, if other traders react to this price change by buying or selling the asset themselves in that period, perhaps because they expect a wave of trading activity that pushes the price in one direction or another, this significantly affects market dynamics and estimates of trade informativeness. Therefore, a model which does not allow contemporaneous feedback trading will bias any estimates of how much information trades actually carry. If a positive order flow shock causes an increase in the asset price, which in turn causes an *increase* in order flows via feedback trading within that period, the total effect/price impact of the order flow shock will be higher than when feedback trading is ruled out. Alternatively, if there exists negative feedback trading, perhaps because of expected return reversals of the initial asset price change, the ultimate price impact of the trade will be smaller than the non-feedback case.

The reason for ruling out contemporaneous feedback trading in empirical models with aggregated data is that without such a restriction, the model cannot be estimated since the VAR becomes unidentified. By allowing returns to depend on order flow but ruling out the converse, the two equation VAR in returns and order flows becomes a recursively ordered structural VAR, which is just identified when the variance/covariance matrix of the residuals is restricted to be diagonal.[2] However, as is shown in Section 5.2, when data are aggregated from tick–by–tick to any lower frequency, contemporaneous feedback trading is inevitable. Imposing the restriction that order flows do not depend on contemporaneous asset returns then represents a mis-specification of the empirical model and is therefore likely to result in biased parameter estimates and incorrect inference.

However, it is not possible to estimate the simultaneous impact of order flows and prices on each other in the VAR model since, in the standard setup, not enough information is available for the VAR's identification. In order to identify both structural parameters, additional information is required and I suggest that sufficient information can be obtained from statistics (typically order flows or returns) from related assets or markets. This additional information can be used to identify and estimate the structural VAR in a very efficient manner through standard instrumental variables methods.[3] In so far as many assets, whether they are currency, stocks or bonds, have a number of assets related to them, it seems likely that information should be readily available. This could then be used by the econometrician in order to estimate an otherwise unidentified model and hence allow for contemporaneous feedback trading. In the case of stocks, when trying to estimate a feedback trading VAR for IBM, one could try using statistics based on Hewlett Packard flows and returns for example, or any stock in the same or related industry. The question of *which* instruments to use is simply an empirical one. One may simply choose those instruments found to be strongly correlated with the endogenous right hand side regressors.

I apply this instrumental variables VAR methodology to the spot USD/EUR (US dollars per euro) foreign exchange market since this is a very active market and provides a natural testing ground for the hypotheses of contemporaneous feedback trading. The data are

[2]Blanchard and Quah (1989) suggest another restriction for identification; that one type of shock has no long run effect on one of the other variables. This assumption cannot be used in this setting. Order flow shocks are likely to have permanent effects on the asset price, as in the models of Kyle (1985), etc., and asset return shocks are likely to have non-zero long run effects on cumulative order flow due to reasons of portfolio rebalancing for example.

[3]Rigobon (2003) has suggested identifying simultaneous equation models based on heteroscedasticity in the data. If the heteroscedasticity of the structural errors is not constant, then this changing variance can be used to identify the structural equations in a similar way to the IV methods used here. See also Rigobon and Sack (2003).

taken from the Reuters D2000-2 electronic trading system, one of the two dominant brokered trading platforms used in the inter-dealer spot FX market and cover the eight month period from 1st December 1999 to 24th July 2000. In order to estimate the structural, feedback VAR for the USD/EUR market I use instrumental variables, where the instruments are statistics obtained from the closely linked markets of USD/GBP (US dollars per pound sterling) and GBP/EUR (pounds sterling per euro). It is clear that the USD/EUR rate should, in the absence of arbitrage, equal the USD/GBP rate multiplied by the GBP/EUR rate, a relationship which is exploited in Chapters 2 and 3. Statistics obtained from these markets, in particular returns and order flows, may then be correlated with the endogenous variables for which we are trying to instrument, i.e. USD/EUR returns and flows.

In order to evaluate the importance of the feedback trading parameter, I consider two sampling frequencies; one minute and five minutes.[4] From both frequencies I estimated the VARs with and without feedback trading and calculated impulse response functions following an order flow shock, representing private information, in order to assess the informativeness of trades.

For the one minute frequency, the feedback trading parameter in the structural VAR is found to be positive and significant at the 1% level and the impulse response function following an order flow shock is larger when feedback trading is permitted. However the difference between the restricted and unrestricted IRFs is not significantly different, at the 5% level. At the five minute frequency, the feedback trading parameter in the VAR is quantitatively large and significant at the 5% level, and the impulse response without feedback trading is significantly below that when feedback trading is permitted. This suggests that in the case of spot FX markets, feedback trading is prevalent and has significant implications when examining the price impact/informativeness of order flows, especially when the data are sampled at the lower/five minute frequency. For the spot FX market considered in this exercise, feedback trading is positive, i.e. order flow in one period depends positively on the asset return in that period.

I demonstrate theoretically in Section 5.2 that omitting feedback trading in aggregated data will result in a mis-specified model and will bias the estimated impact of order flow on prices. This is confirmed in the empirical exercise where I show that feedback trading is a significant concern for empirical order flow models. By employing appropriate

[4]When the data were sampled at lower frequencies the abilities of other market statistics to instrument for USD/EUR endogenous variables deteriorated to such an extent that such analysis became pointless.

instruments, I show that in high frequency foreigh exchange data, there exists positive contemporaneous feedback trading: order flow in one period depends positively on the return experienced within that period. This results in the price impact of an order flow shock being significantly greater than when one imposes a recursive ordering of the VAR. Private information, in the form of unanticipated order flow shocks, then has a larger impact on returns than previously believed, i.e. trades carry more information than previous estimates suggest.

The rest of the paper is organised as follows. Section 5.2 motivates the need to model contemporaneous feedback trading when data are aggregated at any level. Section 5.3 introduces the model to be estimated and describes the standard techniques to be employed as well as explanations as to how to obtain analytical confidence bounds for the impulse response functions. Section 5.3 also discusses the choice of instruments used. Section 5.4 presents the regression results and reports the impulse response functions. Section 5.5 discusses the findings, placing them within the existing finance literature and Section 5.6 concludes.

5.2 The Inevitability of Contemporaneous Feedback Trading in Aggregated Data

In this section I show how contemporaneous feedback trading can result simply from considering aggregated data. By definition, contemporaneous feedback trading, whereby order flows at date t depend on date t asset returns, cannot occur in tick-by-tick data; a trader can only respond to a price change once the price has indeed changed. Assume, without loss of generality, that returns and flows can be characterised by a VAR with only 1 lag. The VAR model, originally introduced by Sims (1980) and implemented in the microstructure literature by Hasbrouck (1991a), is a simple statistical framework that allows one to examine the relationships between asset returns and trading activity; more specifically, order flows. Hasbrouck originally applied the model to US equity data where the data were sampled in 'transaction time', i.e. tick-by-tick data were used. In what follows, assume the data are sampled at a sufficiently high frequency so that contemporaneous feedback trading is ruled out. This could be tick-by-tick, ten second or even one second data for very active markets.[5] Since the possibility of contemporaneous

[5]For the purpose of this study I simply wish to examine data at a frequency where no contemporaneous feedback trading can occur and then see what happens when the data are aggregated at a lower frequency.

feedback trading is ruled out, the system of equations can be written as:

$$\begin{aligned} R_t &= \alpha_1 + \beta F_t + \phi_{11} R_{t-1} + \phi_{12} F_{t-1} + \epsilon_t^R \\ F_t &= \alpha_2 + \phi_{21} R_{t-1} + \phi_{22} F_{t-1} + \epsilon_t^F \end{aligned} \tag{5.1}$$

where R_t is the return on the asset in period t, defined as the first log difference of the price, F_t is the order flow in period t; the number of buyer less seller initiated trades in that interval and ϵ_t^R and ϵ_t^F are serially uncorrelated, independent errors with variances $\sigma^2_{\epsilon^R}$ and $\sigma^2_{\epsilon^F}$ respectively. The system can be written as a structural VAR:

$$\begin{gathered} \underbrace{\begin{bmatrix} 1 & -\beta \\ 0 & 1 \end{bmatrix}}_{Q} \underbrace{\begin{bmatrix} R_t \\ F_t \end{bmatrix}}_{y_t} = \underbrace{\begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix}}_{\alpha} + \underbrace{\begin{bmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{bmatrix}}_{\phi} \underbrace{\begin{bmatrix} R_{t-1} \\ F_{t-1} \end{bmatrix}}_{y_{t-1}} + \underbrace{\begin{bmatrix} \epsilon_t^R \\ \epsilon_t^F \end{bmatrix}}_{\epsilon_t} \\ \Rightarrow Qy_t = \alpha + \phi y_{t-1} + \epsilon_t \qquad Var(\epsilon_t) = \Omega = \begin{bmatrix} \sigma^2_{\epsilon^R} & 0 \\ 0 & \sigma^2_{\epsilon^F} \end{bmatrix} \end{gathered} \tag{5.2}$$

where $y_t = [R_t \ F_t]'$ and $\epsilon_t = \left[\epsilon_t^R \ \epsilon_t^F\right]'$, $t = 1, \ldots, 2T$. What I wish to do is to move from a structural VAR at a high frequency, t, to a structural VAR where the frequency is halved (frequency τ).[6] However, in order to do this, (5.2) must be converted into a reduced form and put into state space representation. Aggregation can then be performed using the methods in Harvey (1989). The reduced form of (5.2), at the frequency, t, is clearly

$$y_t = \mu + Q^{-1}\phi y_{t-1} + Q^{-1}\epsilon_t \qquad t = 1, \ldots, 2T \tag{5.3}$$

where $\mu = Q^{-1}\alpha$ and the corresponding reduced form at the frequency, τ, where $\tau = 2t$, is shown in the appendix, Section 5.A.1, to be

$$y_\tau = \mu^+ + Ay_{\tau-1} + e_\tau \qquad \tau = 1, \ldots, T \tag{5.4}$$

where $A = Q^{-1}\phi\,(I_2 + Q^{-1}\phi)$. Since returns and order flows are both flow variables, as opposed to stock variables, then the period τ (low frequency) return is simply the sum of

[6]For example, t could represent data sampled at the ten second frequency and τ would represent data sampled at the twenty second frequency.

the two t period returns in that interval, and similarly for order flows. The variance of e_τ is given by

$$Var\left(e_\tau\right) = G\Omega G' + Q^{-1}\Omega\left(Q^{-1}\right)' \tag{5.5}$$

where $G = \left(I_2 + Q^{-1}\phi\right)Q^{-1}$. To convert (5.4) into a structural VAR, one first needs to factorise the variance of e_τ. This is done in (5.6).

$$Var\left(e_\tau\right) = \underbrace{\left[\begin{array}{cc} G & Q^{-1}\end{array}\right]}_{P}\underbrace{\left[\begin{array}{cc}\Omega & 0\\ 0 & \Omega\end{array}\right]}_{\tilde{\Omega}}\underbrace{\left[\begin{array}{c} G' \\ (Q^{-1})'\end{array}\right]}_{P'} = P\tilde{\Omega}P' \tag{5.6}$$

If we premultiply (5.4) by $[I_2\ I_2]\,P^+$, where P^+ is the Moore-Penrose inverse of the 2×4 matrix P, $P^+ = \left(P'P\right)^{-1}P'$, then the structural form can be written as:

$$[I_2\ I_2]\,P^+y_\tau = [I_2\ I_2]\,P^+\mu^+ + [I_2\ I_2]\,P^+Ay_{\tau-1} + [I_2\ I_2]\,P^+e_\tau \tag{5.7}$$

The variance of the error vector in the structural form is now

$$Var\left([I_2\ I_2]\,P^+e_\tau\right) = 2\Omega \tag{5.8}$$

(5.7) therefore has the appealing property that the variance of the error (return or flow) at the τ period frequency is twice that of the error at the t period frequency.[7] In order to solve for the structural parameters in (5.7), note that if we let $x_\tau = P^+y_\tau$, then x_τ is the solution to

$$\underbrace{\left[\begin{array}{cccc} g_{11} & g_{12} & 1 & \beta\\ g_{21} & g_{22} & 0 & 1\end{array}\right]}_{[G\ Q^{-1}]=P}\underbrace{\left[\begin{array}{c} x_{1\tau}\\ x_{2\tau}\\ x_{3\tau}\\ x_{4\tau}\end{array}\right]}_{x_\tau} = \underbrace{\left[\begin{array}{c} R_\tau\\ F_\tau\end{array}\right]}_{y_\tau} \tag{5.9}$$

There are clearly an infinite number of solutions for x_τ and this is to be expected when considering the literature on simultaneous equation models; premultiply any structural

[7]Note that the errors in the return and flow equations at the t period frequency are serially uncorrelated and independent, as too are those at the τ period frequency.

form by a non-singular matrix and the reduced form, (5.4), will be unaffected. However, $[I_2 \;\; I_2]\, P^{+} y_\tau$, the right hand side of the structural equation, (5.7), will have the general form

$$\frac{1}{|G|}\begin{bmatrix} g_{22} & -g_{12} \\ -g_{21} & g_{11} \end{bmatrix}\begin{bmatrix} R_\tau \\ F_\tau \end{bmatrix} \tag{5.10}$$

$$+\begin{bmatrix} \left(1-\frac{1}{g_{11}}+\frac{g_{21}}{|G|}\right) & -\left(\frac{\beta}{g_{11}}+\frac{(g_{11}-g_{21}\beta)}{|G|}\right) \\ \left(\frac{g_{21}}{|G|}\right) & \left(1-\frac{(g_{11}-g_{21}\beta)}{|G|}\right) \end{bmatrix}\begin{bmatrix} m_\tau \\ n_\tau \end{bmatrix} \tag{5.11}$$

for any real values of m_τ and n_τ, $\tau = 1, \dots, T$. For a proof, see the appendix, Section 5.A.2. Therefore there are an infinite number of structural VARs at the τ period frequency that are consistent with the recursively ordered structural t period VAR in (5.2). The top rows of (5.10) and (5.11) can be interpreted as the left hand side of the structural return equation and the bottom rows can be interpreted as the left hand side of the structural flow equation. However, a common choice of structural form in (5.7) is the Choleski solution. This essentially chooses the arbitrary values of m_τ and n_τ so that the R_τ term in the structural flow equation drops out. The recursively ordered structural VAR which results is just identified if the variance/covariance matrix of the residuals is assumed to be diagonal. Therefore, by choosing the Choleski solution, standard econometric methods can be used to estimate the model (OLS for example). For all other solutions though, the coefficient on contemporaneous returns in the flow equation will depend on g_{21}. Only if g_{21} equals zero will there be no contemporaneous feedback trading. Assume t represents data sampled at the ten second frequency and τ represents data at the lower, twenty second frequency. In Section 5.A.1, g_{21} is shown to be equal to ϕ_{21}, the coefficient on lagged returns in the flow equation in the ten second VAR. This is perfectly intuitive. If the order flow from seconds 11 to 20 depends on the return from seconds 1 to 10, then part of the order flow from seconds 1 to 20 will depend on part of the return from seconds 1 to 20, i.e. contemporaneous feedback trading exists! The question then becomes, on what grounds should the Choleski solution be chosen? The justification for choosing the Choleski solution is always on the basis that it makes life easier; the model becomes just identified and estimation can take place simply using OLS. However, by imposing a recursive structure on the VAR, the Choleski solution rules out contemporaneous feedback trading. In this chapter, I suggest that one should use the data to tell us what the coefficient on returns in the flow equation should be, rather than assuming it to be zero. In the empirical application below I show that using data to calculate the effect of contemporaneous returns on flows, rather than assuming the coefficient to be zero, can have serious implications when calculating the price impact of trades and hence on estimates of their information content.

5.3 The VAR Model with Feedback Trading

The VAR model is a convenient statistical framework that allows one to analyse the relationships between, among other things, asset returns and order flow. By interpreting order flow shocks as private information, one can examine the price impact of such shocks, via impulse response functions, and therefore give a quantitative estimate of the information content of trades. The greater the price impact, the more information trades are argued to carry.

5.3.1 Model design

The standard VAR model allows asset returns to depend on contemporaneous order flows but not the converse. From Section 5.2 we saw that when data are aggregated, even at still very high frequencies, this recursive ordering may not be valid. Here I allow both returns and order flows to depend on each other contemporaneously. The model to be estimated can be written as:

$$y_t + By_t = c + \sum_{j=1}^{p} \phi_j y_{t-j} + \epsilon_t \qquad t = 1, \ldots, T \tag{5.12}$$

In this example, y_t is simply the 2×1 vector of endogenous variables at date t, i.e. returns and flows; $y_t = [R_t \; F_t]'$. In the appendix I generalise the analysis from the 2 to the n variable case.[8] c is a 2×1 vector of constants, the summation term contains the lags of the VAR and ϵ_t is a 2×1 vector of residuals with zero mean and variance matrix Ω, assumed to be diagonal. B is the 2×2 matrix of structural parameters with zeros along

[8] In the model of Engle and Patton (2004), returns of ask prices and returns of bid prices are considered separately, so y_t need not be restricted to be a 2×1 vector. In the original Hasbrouck (1991a) setting, y_t also contained a number of trade related variables including trade sign and the interactions between trade sign and volume and spread.

the main diagonal.[9]

$$B = \begin{bmatrix} 0 & -b_{12} \\ -b_{21} & 0 \end{bmatrix} \tag{5.13}$$

b_{12} represents the contemporaneous effect of flows on returns and b_{21} represents the contemporaneous feedback trading parameter. For each equation, $i = 1,\ 2$, the k exogenous and $n-1\,(=1)$ endogenous regressors can be stacked into a $(1+k)\times 1$ vector, z_{it}. Stacking these vectors across the T observations allows me to write

$$y_i = z_i\Pi_i + \epsilon_i \qquad i = 1,\ 2 \tag{5.14}$$

Π_i is a $(1 + k) \times 1$ vector of parameters. y_i is the $T \times 1$ vector of the scalar y_{it}s ($y_{it} \in \{R_t,\ F_t\}$) and similarly for ϵ_i. z_i is the $T \times (1 + k)$ matrix formed by stacking the T, $(1 + k) \times 1$, z_{it} vectors. In matrix form, (5.14) can be written:

$$\underbrace{\begin{bmatrix} R \\ F \end{bmatrix}}_{Y} = \underbrace{\begin{bmatrix} z_1 & 0_{T\times(1+k)} \\ 0_{T\times(1+k)} & z_2 \end{bmatrix}}_{Z} \underbrace{\begin{bmatrix} \Pi_1 \\ \Pi_2 \end{bmatrix}}_{\pi} + \underbrace{\begin{bmatrix} \epsilon^R \\ \epsilon^F \end{bmatrix}}_{\epsilon} \tag{5.15}$$

Y is $2T \times 1$, Z is $2T \times 2(1 + k)$, π is $2(1 + k) \times 1$ and ϵ is $2T \times 1$. Writing the system in this form will help to calculate the distribution of the impulse response functions in Section 5.3.3, since they will be functions of the distribution of the $2(1 + k) \times 1$ vector, π.

5.3.2 Instrumental variables

Since each of the equations in (5.14) contain endogenous variables on the right hand side, I estimate using instrumental variables. For the $1 + k$ variables in z_{it} I use the $g + k$ instruments w_{it}, $i = 1,\ 2$ and $g \geq 1$. For a greater explanation of the use of these variables as instruments, see Section 5.3.5. Using two stage least squares, the IV estimator for Π_i, $i = 1,\ 2$, is denoted $\hat{\Pi}_i$ and calculated as:

$$\hat{\Pi}_i = \left[z_i'w_i\left(w_i'w_i\right)^{-1}w_i'z_i\right]^{-1} z_i'w_i\left(w_i'w_i\right)^{-1}w_i'y_i \tag{5.16}$$

[9] I separate y_t and By_t in (5.12) as this simplifies the notation in the appendix when the distribution of the impulse response functions is calculated.

where w_i is simply the $T \times (g+k)$ matrix formed by stacking the T w_{it} vectors of instruments. Using standard instrumental variables methods, the distribution of $\hat{\pi}$, the $2(1+k) \times 1$ vector of parameters in (5.15) is given by (see Section 5.A.3):

$$\sqrt{T}\,(\hat{\pi} - \pi) \xrightarrow{d} N\,(0\ ,\ \Sigma_{\pi}) \tag{5.17}$$

5.3.3 Impulse response functions

In order to evaluate the informativeness of trades, a common approach is to use impulse response functions (IRFs). To calculate these impulse response functions, I convert the VAR of (5.12) into its MA(∞) representation. It is simple to show that (5.12) can be written as[10]

$$y_t = \mu + \Psi_0\,(I_2 + B)^{-1}\,\epsilon_t + \Psi_1\,(I_2 + B)^{-1}\,\epsilon_{t-1} + \Psi_2\,(I_2 + B)^{-1}\,\epsilon_{t-2} + \ldots \tag{5.18}$$

where μ is the unconditional mean of the vector y_t. B is given in (5.13) and Ψ_S, $S = 0, 1, 2, \ldots$, is given by

$$\Psi_S = (I_2 + B)^{-1}\,\phi_1 \Psi_{S-1} + (I_2 + B)^{-1}\,\phi_2 \Psi_{S-2} + \ldots + (I_2 + B)^{-1}\,\phi_P \Psi_{S-P} \tag{5.19}$$

where ϕ_j, $j = 1, \ldots, P$, are the coefficients on the lags in the VAR in (5.12) and $\Psi_0 = I_2$, the 2×2 identity matrix. $\Psi_k = 0_{2\times 2}\ \forall\ k < 0$. The impacts of ϵ_t on y_{t+S} are shown by the impulse response functions, H_S, where H_S is the 2×2 coefficient matrix on ϵ_{t-S} in (5.18).

$$H_S = \Psi_S\,(I_2 + B)^{-1} \tag{5.20}$$

I therefore introduce a one unit order flow shock to each of the VARs and examine the effect of the feedback trading parameter by comparing the IRFs with and without this feedback trading. In order to determine whether the non-feedback IRF is significantly different from the IRF of the unrestricted VAR, one has to calculate the distribution of this feedback impulse response and I do so analytically using the delta method.[11] From

[10] See Hamilton (1994) for example.

[11] In Chapter 2 and 3, confidence intervals for IRFs were found by using numerical methods (bootstrapping). However, this proved to be problematic in this exercise due to the difficulties of bootstrapping

(5.20), H_S depends only on the parameters in the structural VAR, the distribution of which is shown in (5.17). If $\hat{H}_S = \hat{\Psi}_S \left(I_2 + \hat{B}\right)^{-1}$, where the caret denotes parameter estimate, and if $\hat{h}_S = vec\left(\hat{H}'_S\right)$ then the distribution of the IRF parameters is given, using (5.17), as

$$\sqrt{T}\left(\hat{h}_S - h_S\right) \xrightarrow{d} N\left(0\ ,\ G_S \Sigma_\pi G'_S\right) \tag{5.21}$$

where G_S is $4 \times 2(1+k)$ and equals $\frac{\partial h_S}{\partial \pi'}$. However, to calculate G_S, one cannot use the results of Hamilton (1994) or Lütkepohl (1990). The structural VARs considered in Hamilton (1994), chapter 11, are estimated from the reduced form and the structural parameters are backed out from the variance/covariance matrix of the residuals. In that way, the distribution of the IRFs depends, not only on the distribution of the reduced form parameters, but also on the distribution of the variance/covariance matrix of the errors. Since I use instrumental variables to estimate the structural parameters directly, the distribution of the IRFs in (5.21) will depend only on the distribution of the π parameters and *not* on the distribution of the variance/covariance matrix of residuals. In Section 5.A.4, using methods similar to those of Hamilton (1994) and Lütkepohl (1990), I show that G_S can be written as

$$G_S = \frac{\partial h_S}{\partial \pi'} = \left[I_4 + (I_2 \otimes B')\right]^{-1}\left[\frac{\partial \psi_S}{\partial \pi'} - (H_S \otimes I_2)\, S_{B'} \frac{\partial \Theta_{B'}}{\partial \pi'}\right] \tag{5.22}$$

where $S_{B'}$ and $\frac{\partial \Theta_{B'}}{\partial \pi'}$ are shown to be matrices of zeros and ones, and $\frac{\partial \psi_S}{\partial \pi'}$ is given by

$$\begin{aligned}\frac{\partial \psi_S}{\partial \pi'} = &- \left(I_2 \otimes [\phi_1 \Psi_{S-1} + \ldots + \phi_P \Psi_{S-P}]'\right)\left((I_2+B)^{-1} \otimes \left((I_2+B)^{-1}\right)'\right) S_{B'} \frac{\partial \Theta_{B'}}{\partial \pi'} \\ &+ \sum_{j=1}^{P}\left\{\left((I_2+B)^{-1} \otimes \Psi'_{S-j}\right)\frac{\partial \Theta_{\phi'_j}}{\partial \pi'} + \left((I_2+B)^{-1}\phi_j \otimes I_2\right)\frac{\partial \psi_{S-j}}{\partial \pi'}\right\}\end{aligned} \tag{5.23}$$

where, again, $\frac{\partial \Theta_{\phi'_j}}{\partial \pi'}$, $j = 1, \ldots, P$, are matrices of zeros and ones. Using (5.21), (5.22) and (5.23) one can then calculate the distribution of the impulse response functions and therefore see whether the restricted/non-feedback IRF is significantly different from the

instruments - the procedure suggested by Freedman (1984) proved unsuccessful. Instead, analytical expressions for confidence intervals were found. Runkle (1987) shows that this technique works reasonably well. See also Watson (1994).

unrestricted impulse response.[12]

5.3.4 Data

The market considered in this exercise is that of the spot USD/EUR (US dollars per euro) inter-dealer foreign exchange market, taken from the Reuters D2000-2 electronic trading system. The data I consider as instruments are those from the USD/GBP (US dollars per pound sterling) and GBP/EUR (pounds sterling per euro) markets. As in Chapter 3, when sampling the data I record the last transaction price in each period (one minute or five minutes) and the order flow, defined as the number of buyer initiated trades minus the number of seller initiated trades. As explained in Section 2.2.1, no information on traded quantities is available. However, to the extent that earlier work has shown little size variation in trades on this dealing system (Payne 2003a) and that in other applications it is the number rather than aggregate size of trades that has been shown to matter for prices and volatility (Jones, Kaul, and Lipson 1994), I expect that this limitation will not distort the results. Furthermore, even when both the number and size of trades have been available, research has often focussed on the former measure of trading activity (Hasbrouck 1991a). Certain sparse trading periods are also removed from the sample. These include weekends, the overnight period, defined as 1800 to 0600 GMT (BST in the summer months) where trading activity was found to be very thin and some public holidays. See Chapter 2, Section 2.2, for more detail on this filtering process. However, if the Reuters data feed broke down on GBP/EUR or USD/GBP but not USD/EUR then those data are still excluded, purely because the GBP/EUR and USD/GBP data are needed in the construction of the instruments. As in Chapter 2, this filtering process reduced the total number of observations to 90,270 at the one minute frequency and 18,401 at the five minute frequency. Table 5.1 contains statistical information on exchange rate returns, defined as 100 times the logarithmic difference in prices, transaction frequencies and order flows for the filtered data sample.

5.3.5 Instrumenting the endogenous variables

The IV estimator and its distribution, reported in Section 5.3.2 are standard results. The main question at this point concerns what instruments one can use and how good they

[12]The distribution of the cumulative IRFs can be calculated quite easily from (5.21). See Lütkepohl (1990).

Table 5.1
Summary Statistics on USD/EUR Exchange Rate Returns and Flows

Variable	Mean	Variance	Skewness	Kurtosis	Autocorrelation (lags)			
					1	5	10	20
1 minute frequency								
returns	-0.000153	0.000829	0.0621	31.77	-0.0254*	-0.0139*	-0.0122*	-0.00210
absolute returns	0.0164	0.000559	4.66	53.73	0.217*	0.128*	0.106*	0.0862*
flows	0.0361	11.86	-0.386	20.66	0.232*	0.0296*	-0.00235	0.000158
trades	4.02	19.01	2.41	16.17	0.573*	0.420*	0.366*	0.304*
5 minute frequency								
returns	-0.000820	0.00374	-0.00431	7.84	-0.0448*	0.00631	-0.000700	0.00559
absolute returns	0.0422	0.00196	2.42	13.0	0.195*	0.101*	0.0738*	0.0553*
flows	0.190	94.2	-0.106	9.91	0.124*	0.0171*	-0.00752	0.00198
trades	19.9	295.1	1.99	11.1	0.685*	0.451*	0.312*	0.108*

Notes: The USD/EUR exchange rate is defined as the number of dollars (numerator currency) per euro (denominator currency). Returns are defined as 100 times the first difference of the logarithm of the exchange rate. Positive order flow in the USD/EUR market implies net purchases of euro. * denotes significance at the 5% level or less.

are at instrumenting the endogenous regressors. Since the data available in this exercise include not only USD/EUR returns and transactions but also those from the USD/GBP and GBP/EUR markets, the statistics from these other two markets seem prime candidates for use as instruments.[13] Previous research has documented the cross effects of order flow on exchange rates. Evans and Lyons (2002a) document the role that order flow in one currency has in determining exchange rates in other markets. In particular DEM/USD (Deutsche marks per dollar) and CHF/USD (Swiss francs per dollar) order flows have significant effects on a number of other dollar exchange rates. These cross market effects are also documented in Daníelsson, Luo, and Payne (2002) which considers the USD/EUR, GBP/EUR, USD/GBP and JPY/USD markets. Theoretical explanations as to why cross effects of order flow exist are also presented in Lyons and Moore (2003), which examines the triangle of rates between the US dollar, euro and yen.

Since the triangle of rates between the dollar, sterling and euro form a strict cointegrating system, using contemporaneous USD/GBP and GBP/EUR returns as instruments for USD/EUR returns is likely to be problematic. Following a shock to USD/EUR returns at date t for example, this will affect not only the USD/EUR rate but also one or both of the sterling rates, otherwise clear arbitrage opportunities would result. In which case date t USD/GBP and GBP/EUR returns will be correlated with the date t error in the USD/EUR return equation. The use of contemporaneous USD/GBP and GBP/EUR returns as instruments for the endogenous USD/EUR variables will then result in biased parameter estimates just as OLS estimates would. Instead, only lags of sterling returns, which should not be correlated with the errors, are considered and hence stand a good chance of being valid instruments.

Also, under conditions of no arbitrage, it is clear that the USD/EUR return at time t will equal the sum of the returns in the USD/GBP and GBP/EUR markets. Therefore, using lags of *both* sterling returns as instruments will be problematic since, unless the coefficients on the sterling returns are different in the first stage regression in the 2SLS procedure, one would essentially be using a 'synthetic' lagged USD/EUR return to instrument for contemporaneous returns. However, the lagged USD/EUR return is effectively being used as an instrument for itself, since it too is included in the VAR. For this reason I only use one of either USD/GBP or GBP/EUR returns as instruments. Which return series is used depends on how good they are at instrumenting for the endogenous regressors. Since the

[13] Also available in the dataset are the numbers of buyer and seller initiated trades separately. One may have thought that trading volume could be used as an instrument. Unfortunately, this is not the case and is proved in the appendix, Section 5.A.5.

no arbitrage problem does not place any restrictions on the order flow series; USD/EUR order flow in period t does not have to equal the sum of the USD/GBP and GBP/EUR flows, I consider both USD/GBP and GBP/EUR order flows as candidate instruments for the contemporaneous USD/EUR flow regressor.

Researchers have often pointed to the pitfalls of using weak instruments and the bias that such instruments introduce. See for example Buse (1992), Bound, Jaeger, and Baker (1995), Wang and Zivot (1998) and Staiger and Stock (1997). It is therefore vital that the quality of the instruments is examined. In the two equation case, in returns and flows, only one endogenous regressor is present and so the quality of the instruments is tested using the procedure discussed in Pagan and Robertson (1998). They suggest an easily implementable test for the quality of potential instruments by means of a Wald test. For cases with multiple endogenous regressors, see Shea (1996) and Hall, Rudebusch, and Wilcox (1996).[14] The Wald test is essentially a test of the overidentifying restrictions in the model. One could write model (5.12) with other right hand side explanatory variables, including GBP/EUR and USD/GBP returns and flows, the cross market effects of which were documented in Chapters 2 and 3. By imposing 'zero' restrictions on the variables, they can be used to instrument the endogenous regressors, and the Wald test is essentially testing these overidentifying restrictions in the reduced form VAR. The results of these Wald tests are reported in Table 5.2 for the one and five minute frequency models.

For the one minute frequency VAR, three lags of USD/GBP and GBP/EUR flows were chosen to instrument for USD/EUR flows in the return equation, while two lags of USD/GBP returns were chosen as instruments for USD/EUR returns in the flow equation. Testing the quality of the instruments for USD/EUR flows in the return equation produced a Wald test of 28.09. For returns in the flow equation, the instrument Wald test was 236.13. The 1% critical values for the corresponding χ^2 distributions are 16.81 and 9.21 respectively, suggesting that the chosen variables are good instruments for the endogenous regressors. When considering the five minute frequency VAR, contemporaneous and one lag of both USD/GBP and GBP/EUR flows were chosen to act as instruments for USD/EUR flows in the return equation, while in the flow equation, two lags of GBP/EUR

[14] To test the quality of the instruments I run a regression of the endogenous regressor in each of the structural equations in (5.12) on all the exogenous variables (lagged USD/EUR flows and returns) as well as the candidate instruments. A Wald test is then performed on the coefficients on those instruments. For the endogenous USD/EUR return regressor I begin by using the first lag of either USD/GBP or GBP/EUR returns and continue increasing the lag length of the instruments until no more explanatory power is added by their inclusion. The choice as to which sterling series to use, is made based on the overall fit of this first stage regression. For the USD/EUR flow regressor, I start by considering contemporaneous USD/GBP and GBP/EUR flows and increase the lag length in a similar fashion.

Table 5.2
Instrumenting USD/EUR Returns and Flows

Return equation

$$R_t^{USD/EUR} = c + b\,F_t^{USD/EUR} + \sum_{i=1}^{m} u_i R_{t-i}^{USD/EUR} + \sum_{j=1}^{n} v_j F_{t-j}^{USD/EUR} + \epsilon_t^R$$

	1 min freq.	5 min freq.
Instruments for $F_t^{USD/EUR}$ (flows)	$\underbrace{F_{t-1}^{USD/GBP}\ F_{t-1}^{GBP/EUR}\ F_{t-2}^{USD/GBP}\ F_{t-2}^{GBP/EUR}\ F_{t-3}^{USD/GBP}\ F_{t-3}^{GBP/EUR}}_{w_t}$	$\underbrace{F_t^{USD/GBP}\ F_{t-1}^{USD/GBP}\ F_t^{GBP/EUR}\ F_{t-1}^{GBP/EUR}}_{w_t}$

Frequency	Wald test	1% critical value	Degrees of freedom
1 min	28.09	16.81	6
5 min	7109.01	13.28	4

Flow equation

$$F_t^{USD/EUR} = c + b\,R_t^{USD/EUR} + \sum_{i=1}^{m} u_i R_{t-i}^{USD/EUR} + \sum_{j=1}^{n} v_j F_{t-j}^{USD/EUR} + \epsilon_t^F$$

	1 min freq.	5 min freq.
Instruments for $R_t^{USD/EUR}$ (returns)	$\underbrace{R_{t-1}^{USD/GBP}\ R_{t-2}^{USD/GBP}}_{w_t}$	$\underbrace{R_{t-1}^{GBP/EUR}\ R_{t-2}^{GBP/EUR}}_{w_t}$

Frequency	Wald test	1% critical value	Degrees of freedom
1 min	236.13	9.21	2
5 min	62.26	9.21	2

Notes: R_t^x is 100 × the log first difference of exchange rate x at date t. F_t^x is the order flow for exchange rate x, defined as the number of buyer less the number of seller initiated transactions in period t.

returns were chosen to instrument for USD/EUR returns. The regression of USD/EUR flows on all exogenous variables and instruments produced a Wald test of 7109.01, while testing how good the GBP/EUR returns are as instruments for returns in the flow equation produced a Wald test of 62.26. Again both of these are significant at the 1% level, suggesting that these variables make good instruments.[15] I also considered sampling the data at lower frequencies. However, at anything lower than the five minute frequency, the instrument Wald test became insignificant, even at the 5% level, suggesting that neither USD/GBP or GBP/EUR variables would be good at instrumenting for the endogenous USD/EUR returns and flows. Since only lagged returns are suggested as instruments, due to the problems of using contemporaneous USD/GBP and GBP/EUR returns explained above, as soon as one considers lower frequency data, the ability of these lagged variables to instrument for USD/EUR returns is likely to fall. At the hourly or daily frequency for example, returns and flows will no longer be serially correlated. If USD/EUR returns are not correlated with its own lag, it is highly unlikely that they will be correlated with the lags of USD/GBP or GBP/EUR returns. At the lower frequencies, even the fifteen minute level, the candidate instruments became very weak. Therefore the estimations were only performed for the one and five minute VARs.

To my knowledge, only one other paper has tried to examine contemporaneous feedback trading in the foreign exchange market. Evans and Lyons (2003) use a VAR model, as is done here, in returns and order flows but is not as general as the procedure outlined above. Evans and Lyons split order flows into two types, both having different roles. They assume returns depend on contemporaneous 'informational' trades while 'feedback' trades depend on contemporaneous returns. The order flow measure available from the data is simply the sum of these two components. This is more restrictive than the approach set out above since I do not, in any way, split trades into different motives. By using variables obtained from other FX markets as instruments one is able to estimate an otherwise unidentified model.[16]

[15]It may appear strange that different variables were chosen to act as instruments at the one and five minute frequencies. This may be reconciled when one considers the different FX market dynamics at the different frequencies. With different dynamics and cross correlations at one and five minute frequencies, it may be unsurprising to find different choices of instruments. The question of which variables to include as instruments is, after all, an empirical one. The huge Wald statistic on the instruments for USD/EUR flows (7109.01) comes primarily from the use of contemporaneous USD/GBP and GBP/EUR flows as instruments. Surprisingly, contemporaneous flows were not of any use for the one-minute VAR, suggesting possibly delayed information spill-overs from one market to another.

[16]At the daily frequency Evans and Lyons (2003) find evidence of negative feedback trading!

Table 5.3

USD/EUR VAR Results (1 Minute Frequency)

	Feedback VAR				Non-feedback VAR			
	R_t equation		F_t equation		R_t equation		F_t equation	
constant	-0.000519[a]	(-4.21)	0.0435[a]	(3.88)	-0.000376[a]	(-4.75)	0.0381[a]	(3.44)
$flow_t$	0.00841[a]	(4.17)			0.00468[a]	(57.61)		
$return_t$			27.02[a]	(3.13)				
$return_{t-1}$	-0.249[a]	(-5.28)	24.12[a]	(26.55)	-0.165[a]	(-15.89)	22.52[a]	(28.86)
$return_{t-2}$	-0.0684[a]	(-4.37)	7.31[a]	(12.83)	-0.0420[a]	(-6.01)	7.07[a]	(12.47)
$return_{t-3}$	-0.0258[a]	(-2.99)	2.91[a]	(4.09)	-0.0152[b]	(-2.19)	2.86[a]	(4.02)
$return_{t-4}$	-0.0346[a]	(-4.58)	2.83[a]	(5.05)	-0.0255[a]	(-4.85)	2.45[a]	(4.44)
$return_{t-5}$	-0.0366[a]	(-5.20)	2.65[a]	(5.47)	-0.0286[a]	(-5.41)	2.16[a]	(4.67)
$return_{t-6}$	-0.0267[a]	(-4.47)	2.16[a]	(4.64)	-0.0197[a]	(-4.82)	1.87[a]	(4.09)
$return_{t-7}$	-0.0171[a]	(-3.64)	0.994[b]	(2.04)	-0.0146[a]	(-3.74)	0.688	(1.44)
$flow_{t-1}$	-0.000399[c]	(-1.76)	0.0954[a]	(12.64)	0.00000878	(0.16)	0.109[a]	(17.94)
$flow_{t-2}$	-0.000191[a]	(-2.88)	0.0213[a]	(3.60)	-0.000113[a]	(-2.74)	0.0209[a]	(3.53)
$flow_{t-3}$	-0.000268[a]	(-3.37)	0.0313[a]	(3.33)	-0.000152[a]	(-2.88)	0.0311[a]	(3.31)
R^2	0.112		0.265		0.302		0.082	
$\hat{\sigma}^2$	0.0285		3.33		0.0239		3.33	

Notes: The data cover the eight month period from 1st December 1999 to 24th July 2000. The USD/EUR exchange rate is defined as the number of dollars (numerator currency) per euro (denominator currency). Returns are defined as 100 times the log first difference of the exchange rate. Positive order flow in the USD/EUR market implies net purchases of euro. a, b, c denote significance at the 1%, 5% and 10% levels respectively. Newey-West corrected T-stats in parentheses.

5.4 Estimation Results

The estimation results of the one minute VAR with and without feedback trading are shown in Table 5.3. The lag lengths of the VAR were chosen using the Schwartz Information Criterion and this resulted in seven lags of returns and three for order flows. Each VAR was estimated equation by equation and heteroscedasticity consistent standard errors were calculated using the Newey-West method.[17] The left panel shows the estimation results of the feedback VAR. There are a number of important findings. Firstly, as one would expect, returns depend positively on contemporaneous order flow. One way buying (selling) pressure causes positive (negative) returns intra minute. Returns also display negative serial correlation, as have been found in Payne (2003a) and Evans (2002), both of whom consider the Deutsche mark-dollar market. Returns also depend negatively on lagged own order flow, although the explanatory power of these variables in the determination of returns is surprisingly low compared to previous studies; the R^2 is only 11.2%.[18] Of more interest in this chapter are the results for the order flow equation. Order flow appears to depend positively on contemporaneous returns, with a coefficient that is significant at the 1% level. This suggests that following a positive return in one minute, traders 'buy into' the currency in that same period, possibly because they expect further price changes in the same direction. However, this is not consistent with the negative serial correlation observed in one minute returns. On the other hand, positive intra minute feedback trading is consistent with the positive effect of lagged returns on order flows, seen in both versions of the VAR. This lagged feedback trading phenomenon is considered in more detail by Cohen and Shin (2003) in the US treasury market. Without feedback trading, the R^2 in the flow equation is only 8.2%. However, when contemporaneous feedback trading is allowed, the R^2 increases to 26.5%, suggesting that contemporaneous price changes are an important determinant of order flows.

Table 5.4 gives the VAR results when the data are sampled at the five minute frequency. Again the VARs were estimated equation by equation and the Newey-West method was used to correct for heteroscedasticity. The Schwartz Information Criterion suggested using

[17]The feedback VAR was estimated equation by equation using instrumental variables, described above, while the recursively ordered/non-feedback VAR was estimated equation by equation using OLS. In this way, we compare the 'true', unrestricted VAR to that which would have been estimated using current best practice.

[18]This raises the question of how good lagged sterling flows are at instrumenting for contemporaneous USD/EUR flows at the one minute frequency. The dramatic reduction in R^2 from non-feedback to feedback VAR, along with the dramatic reduction in the t-stat on contemporaneous flows, suggests either a huge mis-specification in the non-feedback VAR, or the use of instruments which are not as strong as the Wald test suggests. However, this problem is not apparent in the 5 minute frequency VAR, in which the feedback trading parameter is larger and makes more of a difference.

Table 5.4

USD/EUR VAR Results (5 Minute Frequency)

	Feedback VAR				Non-feedback VAR			
	R_t equation		F_t equation		R_t equation		F_t equation	
constant	-0.00232[a]	(-6.30)	0.254[a]	(3.48)	-0.00181[a]	(-5.48)	0.201[a]	(2.85)
$flow_t$	0.00688[a]	(63.47)			0.00435[a]	(45.95)		
$return_t$			56.27[b]	(2.44)				
$return_{t-1}$	-0.183[a]	(-12.28)	21.23[a]	(9.41)	-0.138[a]	(-9.82)	17.83[a]	(10.33)
$return_{t-2}$	-0.0660[a]	(-7.72)	6.15[a]	(3.72)	-0.0559[a]	(-6.79)	3.98[a]	(2.90)
$return_{t-3}$	-0.0337[a]	(-4.53)	3.08[b]	(2.29)	-0.0288[a]	(-4.41)	1.94	(1.54)
$flow_{t-1}$	-0.000213[a]	(-2.78)	0.0398[b]	(2.57)	-0.0000979	(-1.20)	0.045[a]	(2.98)
R^2	0.311		0.375		0.469		0.022	
$\hat{\sigma}^2$	0.0501		9.60		0.0446		9.60	

Notes: The data cover the eight month period from 1st December 1999 to 24th July 2000. The USD/EUR exchange rate is defined as the number of dollars (numerator currency) per euro (denominator currency). Returns are defined as 100 times the log first difference of the exchange rate. Positive order flow in the USD/EUR market implies net purchases of euro. a, b, c denote significance at the 1%, 5% and 10% levels respectively. Newey-West corrected T-stats in parentheses.

three lags of returns and one of flows. The results are similar, but more pronounced, than those from the higher/one minute frequency VAR. Again, order flows have a positive and significant effect on contemporaneous returns, as one would expect, and returns display negative serial correlation and depend negatively on lagged flows. The explanatory power of these variables for five minute returns is quite high, with an R^2 of 31.1%. When examining the flow equation, I again find evidence of feedback trading. Flows depend positively, not only on lagged returns, but also on contemporaneous five minute returns. The coefficient on contemporaneous returns in the flow equation is significant at the 5% level and quantitatively very large; the size of the contemporaneous feedback trading parameter is more than two and a half times larger than that on the first lag of returns (56.27 versus 21.23).[19] Indeed, if there is positive feedback trading (lagged and contemporaneous) at the one minute frequency, this, by definition, will be shown as contemporaneous feedback trading at the five minute frequency, as demonstrated in Section 5.2. The results therefore suggest that positive feedback trading is present in the spot USD/EUR market and significant at high frequencies. Intra minute feedback trading is significant but not large, possibly because of the time it takes for traders to react to the price movements. At the five minute frequency, however, intra period feedback trading becomes much larger as any lagged feedback trading at higher frequencies gets incorporated into the contemporaneous feedback effect. One can interpret the effects of this feedback trading using standard VAR analysis, namely impulse response functions. This is the focus of the next section.

5.4.1 Implications of contemporaneous feedback trading: Impulse response functions

Following standard practice, I use IRFs to estimate the information content of trades. Shocking the system with an order flow shock, ϵ_{it}, where i corresponds to the order flow equation in (5.12), can be interpreted as examining the effect of private information. The larger the impact such a shock has on returns, the more informative order flows are argued to be. By comparing the impulse response functions following an order flow shock in the two VARs (feedback trading versus non-feedback trading) one can examine how important contemporaneous feedback trading is. Intuitively, by ignoring the positive feedback trading (in the recursively ordered structural VAR) it is likely that any order flow shock will have a smaller impact on returns. The existence of positive feedback

[19]For the one minute VAR, contemporaneous feedback trading had a coefficient of 27.02, compared with 24.12 for the first lag of returns.

trading will cause the price impact of order flow shocks/private information to be larger than when feedback trading is ignored, i.e. trades carry *more* information than previous estimates suggest.

The impulse response functions following a one unit order flow shock are shown in Figure 5.1a for the one minute frequency VAR. A number of features can be noted.

- The impact of the order flow shock is almost immediate. Following the one unit shock in the feedback VAR, this causes a 1.09 basis point return and after ten minutes the cumulative return is 1.06 basis points.
- When contemporaneous feedback trading is allowed, the effect of the order flow shock is larger than when contemporaneous feedback trading is prohibited; the feedback impulse response is more than double that in the non-feedback VAR. However, the non-feedback VAR impulse response function is not significantly different from the feedback IRF, i.e. the non-feedback IRF lies within the 95% confidence bound of that from the unrestricted VAR.[20]

This suggests that at the one minute frequency, the difference between the two impulse response functions is *economically* significant, if not statistically so (at the 5% level). Evaluating the informativeness of order flows by considering their price impact will result in a bias if a recursively ordered VAR is considered. However, the statistical significance of this bias is questionable.[21]

The results from the five minute frequency VAR are more pronounced and suggest a much more important role for feedback trading in the interpretation of IRFs. The impulse response functions following a one unit order flow shock are shown in Figure 5.1b. The notable features are given below.

- On impact of the order flow shock in the feedback VAR, this causes a return of 1.12 basis points. The cumulative return is 1.02 basis points after thirty minutes.
- Again, as in the one minute VAR results, the effect of a one unit order flow shock is larger when feedback trading is allowed. On impact of the order flow shock in the

[20]Impulse responses were also done using a one standard deviation order flow shock but this had no effect on the results. This is because the standard deviations of order flow shocks were very similar in both VAR specifications; see Table 5.3.

[21]As reported in footnote 18, the results for the one minute VAR may be open to question if the instruments are not as good as the Wald test suggests.

Figure 5.1
Impulse Response Functions for Feedback and Non-feedback VARs

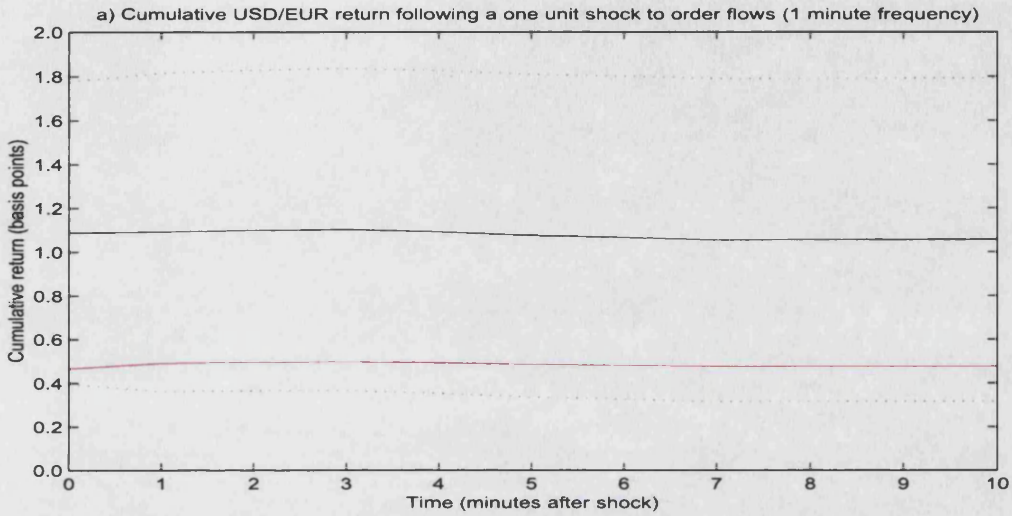

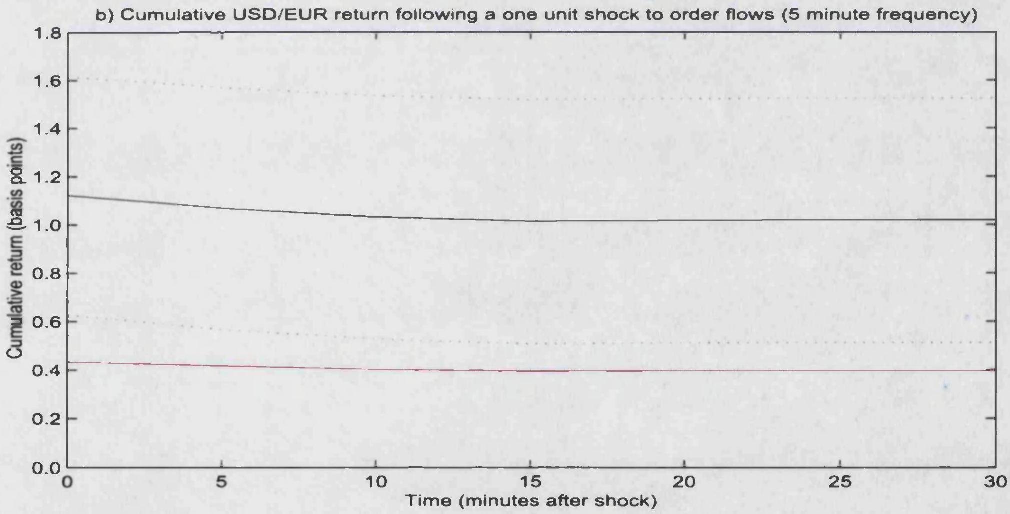

Notes: The figures plot the impulse response functions following a one unit order flow shock. The shock was introduced into the estimated VAR of (5.12) and the cumulative return calculated. In both plots, the solid black line gives the impulse response function from the feedback VAR and the dashed blue lines trace out a 95% confidence interval for the IRF derived from (5.21). The red lines give the impulse response from the non-feedback VAR.

non-feedback VAR, the return is only 0.43 basis points, i.e. the feedback IRF is over two and a half times that of the non-feedback VAR. However, for the five minute frequency case, the non-feedback IRF is significantly different from the unrestricted impulse response, i.e. it lies outside the 95% confidence bound.

Therefore, at the five minute frequency, feedback trading appears to have important consequences when trying to calculate the price impact/informativeness of order flow. The IRF that is commonly computed (that does not allow contemporaneous feedback trading) is significantly below the 'true' IRF which does allow such trading strategies. The price impact of order flow, and hence proxies for the informativeness of such trades, is therefore larger than is commonly believed, implying that trades carry more information than previous studies suggest. The feedback trading that occurs both contemporaneously and also with lags at the one minute frequency, has significant repercussions when modelling five minute data without feedback trading, as is commonly done.

5.5 Discussion and Interpretation

This chapter has shown that feedback trading in the USD/EUR spot FX market does exist even at high frequencies, specifically the one and five minute sampling frequencies. At the one minute frequency, even though the non-feedback impulse response is not different from the unrestricted IRF in a statistical sense, it is different in an economic sense; the feedback impulse response is more than twice that of the non-feedback IRF, implying trades carry over twice as much information than previous estimates suggest. At the five minute frequency, the contemporaneous effect of returns on order flows is significant and causes the IRFs with and without feedback trading to differ significantly, economically and statistically. Indeed, if feedback trading occurs at the one minute frequency in the lags and also contemporaneously, then by definition, such trading strategies will appear contemporaneous at the five minute frequency. This positive feedback trading causes the price impact of unanticipated order flow shocks (representing private information) to be larger compared to when contemporaneous feedback trading is ruled out. The price impact of private information/order flow shocks that is commonly calculated will then be biased downwards compared with the true impact.

In this chapter I have labelled the effects of contemporaneous asset returns on order flows as feedback trading effects. However, there may be other reasons why date t order flows depend on contemporaneous asset returns. Firstly, traders wishing to trade large quantities may break up their trade into a number of smaller sized transactions. By walking up and down the limit order book and splitting a large buy order into a number of smaller trades, to be executed within a short time interval, this will be shown up in the VAR as order flows depending on contemporaneous asset returns. Even though traders are not ultimately wishing to trade based on previous price changes (the decision to trade was made some time earlier) this will still manifest itself as date t order flow depending (statistically) on date t asset returns. Imposing a recursively ordered structural VAR will still be a mis-specification. Another reason why order flow may appear dependent on contemporaneous returns in aggregated data is because of the existence of stop-loss orders (Osler 2002). If the price of an asset falls to a certain level, traders may initiate sell orders in order to *stop losses* from getting any larger. In which case, negative (positive) returns induce negative (positive) order flow immediately. Since traders have the option of posting market or limit orders on the Reuters D2000-2 trading system considered here, positive effects of asset returns on flows will only come out in the data if such stop-loss orders are executed via market orders.[22] It is likely that stop-loss orders will be executed via market, as opposed to limit, orders since if the price reaches the stop-loss level, only market orders are guaranteed to be executed, bringing with it the cost of transacting; the spread. If a limit buy order was placed once the price rose to a certain stop-loss level, the transaction would only occur if another trader wanted to sell to you, i.e. initiate a market sell order. If there was a wave of market buy orders from other traders with stop-loss orders for example,[23] and these buy orders pushed the price up further, then your limit buy order will become increasingly unattractive to anyone wishing to sell. Posting limit stop-loss orders, although avoiding the cost of the spread, are very risky and hence tend to be initiated via market orders. In any case, this chapter has demonstrated that order flow at date t depends positively on date t asset returns in the spot FX market. Whether this is due to the splitting of dealers' trades, stop loss orders or due to 'pure' feedback trading based on extrapolative expectations of future price changes, is irrelevant. In all cases, the assumption of a recursively ordered structural VAR will result in a misspecified model and in a bias in any estimate of the price impact/informativeness of trades.

This chapter does not try to explain *why* feedback trading may occur in the foreign

[22] Buy market orders generate $+1$ of order flow while buy limit orders, once executed, will result in -1 of order flow since a market sell had to be matched against this limit buy for the trade to occur. Order flow is defined as buyer initiated trades (market buy orders) minus seller initiated trades (market sell orders). Stop-loss orders therefore only generate positive feedback if executed via market orders.

[23] Stop-loss orders tend to be clustered at 'round' prices. See Osler (2002).

exchange market, or indeed how profitable such strategies may or may not be. The purpose of the chapter is to analyse the effects of contemporaneous feedback trading on estimates of the price impact of order flows. Such measures are commonly used as proxies for the informativeness of trades (how much information trades carry). As shown in this study, these measures are underestimated if feedback trading is ignored and the data are sampled at anything other than at the highest frequencies. Following a positive order flow shock, representing private information, this causes a positive return due to asymmetric information channels (Kyle 1985, Glosten and Milgrom 1985).[24] If such price changes induce further trades, which in turn cause price changes, etc., then the total effect of the order flow shock will be greater than when contemporaneous feedback trading is prohibited.

At the five minute frequency I find that after a one unit order flow shock, this causes a return of approximately 1.12 basis points whereas when feedback trading is ignored, the return is only 0.43 basis point. The non-feedback IRF is significantly different from that of the unrestricted VAR, implying that feedback trading makes a difference when calculating the price informativeness of trades.

5.6 Conclusions

Microstructure theory suggests that trades carry information and hence have permanent effects on prices. The information content of these trades is normally quantified by examining their price impact (Hasbrouck 1991a) after fitting the data to a vector autoregression. However, common practise is to allow returns to depend on contemporaneous order flows but not the converse. The recursively ordered structural VAR that results can then be estimated quite easily. Although intuitive at ultra-high frequencies, such as tick-by-tick, as soon as one starts aggregating the data, any feedback trading (that by definition can only occur in the lags of tick-by-tick data) will appear contemporaneous. The recursively ordered VAR then becomes misspecified and can have important repercussions when examining the price impact of order flow shocks. In this chapter I use standard instrumental variables techniques in order to estimate a VAR model that allows contemporaneous feedback trading. Feedback trading is found to be significant and positive at the one and five minute frequencies, with the implication that the price impact

[24]Such trades can also generate price changes due to inventory effects (Ho and Stoll 1983, Lyons 1995) but the long run impact of these trades is argued to be zero as traders' inventory positions are restored.

of order flows is underestimated when such trading strategies are not allowed. Trades, in the form of order flow shocks, therefore carry more information/have a larger impact on asset prices than previously believed.

5.A Appendix

5.A.1 Aggregation of the ultra-high frequency VAR in model (5.1)

The methods used here are taken from Harvey (1989). Let τ denote timing at the twenty second frequency, $\tau = 1, \ldots, T$, and t denote timing at the ten second frequency, $t = 1, \ldots, 2T$. The reduced form of the ten second frequency VAR is then given in (5.3). For convenience, this is given in (5.A.1) below.

$$y_t = Q^{-1}\alpha + Q^{-1}\phi y_{t-1} + Q^{-1}\epsilon_t \qquad t = 1, \ldots, 2T \tag{5.A.1}$$

Let μ denote the unconditional mean of the stationary vector process, y_t. In which case one can write

$$y_t - \mu = Q^{-1}\phi\left(y_{t-1} - \mu\right) + \underbrace{Q^{-1}\epsilon_t}_{u_t} \tag{5.A.2}$$

Putting (5.A.2) into state space form, gives the state and observation equations as (5.A.3) and (5.A.4) respectively.

$$\begin{gathered}
\underbrace{\begin{bmatrix} y_{t+1} - \mu \\ y_t - \mu \end{bmatrix}}_{\xi_{t+1}} = \underbrace{\begin{bmatrix} Q^{-1}\phi & 0 \\ I_2 & 0 \end{bmatrix}}_{F} \underbrace{\begin{bmatrix} y_t - \mu \\ y_{t-1} - \mu \end{bmatrix}}_{\xi_t} + \underbrace{\begin{bmatrix} u_{t+1} \\ 0 \end{bmatrix}}_{v_{t+1}} \\
Var\left(v_{t+1}\right) = \Sigma = \begin{bmatrix} Q^{-1}\Omega\left(Q^{-1}\right)' & 0 \\ 0 & 0 \end{bmatrix}
\end{gathered} \tag{5.A.3}$$

$$y_t = \mu + \underbrace{[I_2 \; 0]}_{H'} \begin{bmatrix} y_t - \mu \\ y_{t-1} - \mu \end{bmatrix} \tag{5.A.4}$$

If y_t^f denotes the cumulator variable, i.e. $y_t^f = y_t$ for the first ten seconds of a twenty second period and $y_t^f = y_t + y_{t-1}$ for the second ten second period, then:

$$
\begin{aligned}
y^f_{2(\tau-1)+1} &= y_{2(\tau-1)+1} \\
y^f_{2(\tau-1)+2} &= y_{2(\tau-1)+1} + y_{2(\tau-1)+2}
\end{aligned}
\tag{5.A.5}
$$

The cumulator variable at the ten second frequency, but at times $t = 2, 4, 6, \ldots,$ can therefore be given, using the observation equation, (5.A.4), as

$$
y_t^f = y_{t-1}^f + \mu + H'F\xi_{t-1} + H'v_t \tag{5.A.6}
$$

At times $t = 1, 3, 5, \ldots,$ the cumulator variable is given by

$$
y_{t-1}^f = \mu + H'\xi_{t-1} \tag{5.A.7}
$$

Substituting y_{t-1}^f into (5.A.6) and using $\xi_{t-1} = F\xi_{t-2} + v_{t-1}$ gives

$$
y_t^f = 2\mu + H'\left(F + F^2\right)\xi_{t-2} + \underbrace{H'\left(I_2 + F\right)v_{t-1} + H'v_t}_{H'\bar{v}_t^f} \tag{5.A.8}
$$

Since the cumulator function, y_t^f, at times $t = 2, 4, 6, \ldots,$ is the same as y_τ, the data sampled at the twenty second frequency, then one can write

$$
y_\tau = 2\mu + H'\left(F + F^2\right)\begin{bmatrix} y_{\tau-1} - 2\mu \\ y_{\tau-2} - 2\mu \end{bmatrix} + H'\bar{v}_t^f \tag{5.A.9}
$$

Noting that $H'\left(F + F^2\right) = \left[Q^{-1}\phi\left(I_2 + Q^{-1}\phi\right) \quad 0\right]$ then the reduced form model at the twenty second frequency can be written as

$$
y_\tau = \underbrace{\left(I_2 - Q^{-1}\phi\left(I_2 + Q^{-1}\phi\right)\right)2\mu}_{\mu^+} + \underbrace{Q^{-1}\phi\left(I_2 + Q^{-1}\phi\right)}_{A}y_{\tau-1} + e_\tau \tag{5.A.10}
$$

where $e_\tau = [I_2 \ 0]\, \bar{v}_t^f$. $Var\,(e_\tau)$ is therefore given by

$$\begin{aligned} Var\,(e_\tau) &= [I_2 \ 0]\, Var\left(\bar{v}_t^f\right) \begin{bmatrix} I_2 \\ 0 \end{bmatrix} = [I_2 \ 0] \left[(I_2+F)\,\Sigma\,(I_2+F)' + \Sigma\right] \begin{bmatrix} I_2 \\ 0 \end{bmatrix} \\ &= \left(I_2 + Q^{-1}\phi\right) Q^{-1}\Omega \left(Q^{-1}\right)' \left(I_2 + Q^{-1}\phi\right)' + Q^{-1}\Omega \left(Q^{-1}\right)' \\ &= G\Omega G' + Q^{-1}\Omega \left(Q^{-1}\right)' \end{aligned} \tag{5.A.11}$$

where $G = \left(I_2 + Q^{-1}\phi\right) Q^{-1}$. Expanding this gives us the elements of G.

$$G = \begin{bmatrix} g_{11} & g_{12} \\ g_{21} & g_{22} \end{bmatrix} = \begin{bmatrix} 1+\phi_{11}+\beta\phi_{21} & (1+\phi_{11}+\beta\phi_{21})\,\beta + \phi_{12} + \beta\phi_{22} \\ \phi_{21} & \phi_{12}+\beta\phi_{22} \end{bmatrix} \tag{5.A.12}$$

5.A.2 Derivation of the structural form of the twenty second VAR

The structural VAR at the twenty second frequency is given in (5.7) and $x_\tau, = P^+ y_\tau$, is the solution to the equation in (5.9). The general solution for $x_{i\tau}$, $i = 1, \ldots, 4$, can be found using Gaussian elimination, resulting in

$$\begin{aligned} x_{1\tau} &= \left(\frac{g_{22}}{|G|}\right) R_\tau - \left(\frac{g_{12}}{|G|}\right) F_\tau - \left(\frac{1}{g_{11}} - \frac{g_{21}}{|G|}\right) m_\tau - \left(\frac{\beta}{g_{11}} + \frac{(g_{11}-g_{21}\beta)}{|G|}\right) n_\tau \\ x_{2\tau} &= -\left(\frac{g_{21}}{|G|}\right) R_\tau + \left(\frac{g_{11}}{|G|}\right) F_\tau + \left(\frac{g_{21}}{|G|}\right) m_\tau - \left(\frac{g_{11}-g_{21}\beta}{|G|}\right) n_\tau \\ x_{3\tau} &= m_\tau \\ x_{4\tau} &= n_\tau \end{aligned} \tag{5.A.13}$$

for any real values of m_τ and n_τ, $\tau = 1, \ldots, T$, implying an infinite number of solutions. The right hand side of (5.7) can be written as $[I_2 \ I_2]\, x_\tau$, and the coefficients on R_τ and F_τ are the structural parameters of interest. Noting that $[I_2 \ I_2]\, x_\tau = [(x_{1\tau}+x_{3\tau}) \quad (x_{2\tau}+x_{4\tau})]'$ gives us (5.10).

5.A.3 Distribution of $\hat{\pi}$

The notation and methods used here are similar to those of Hamilton (1994) but I allow IV estimation rather than simple OLS. In this appendix I also generalise the 2 variable case, in returns, R_t, and flows, F_t, to the n variable case, since one may wish to model a number of returns and trade characteristics. From (5.16)

$$\sqrt{T}\left(\hat{\Pi}_i - \Pi_i\right) = \left[\left(\frac{1}{T}\sum_{t=1}^{T} z_{it}w_{it}'\right)\left(\frac{1}{T}\sum_{t=1}^{T} w_{it}w_{it}'\right)^{-1}\left(\frac{1}{T}\sum_{t=1}^{T} w_{it}z_{it}'\right)\right]^{-1} \left(\frac{1}{T}\sum_{t=1}^{T} z_{it}w_{it}'\right)\left(\frac{1}{T}\sum_{t=1}^{T} w_{it}w_{it}'\right)^{-1}\frac{1}{\sqrt{T}}\sum_{t=1}^{T} w_{it}\epsilon_{it} \tag{5.A.14}$$

$$\text{Let} \quad \underset{(g+k)\times(g+k)}{Q^{ii}_{ww}} = \frac{1}{T}\sum_{t=1}^{T} w_{it}w_{it}' \quad \text{and also} \quad Q^{ii}_{ww} \longrightarrow E\left[w_{it}w_{it}'\right] \tag{5.A.15}$$

$$\underset{(n-1+k)\times(g+k)}{Q^{ii}_{zw}} = \frac{1}{T}\sum_{t=1}^{T} z_{it}w_{it}' \qquad Q^{ii}_{zw} \longrightarrow E\left[z_{it}w_{it}'\right] \tag{5.A.16}$$

$$\underset{(g+k)\times(n-1+k)}{Q^{ii}_{wz}} = \frac{1}{T}\sum_{t=1}^{T} w_{it}z_{it}' \qquad Q^{ii}_{wz} \longrightarrow E\left[w_{it}z_{it}'\right] \tag{5.A.17}$$

Also assume that these expectations exist and are finite. This basically states that the instruments are correlated with the right hand side regressors in (5.14). The instruments, w_{it}, have the property that

$$plim\frac{1}{T}\sum_{t=1}^{T} w_{it}\epsilon_{it} = 0 \qquad i = 1,\ldots,n \tag{5.A.18}$$

so that the instruments are uncorrelated with the errors in the original VAR. Let $Q^{ii} = \left[Q^{ii}_{zw}Q^{ii-1}_{ww}Q^{ii}_{wz}\right]^{-1} Q^{ii}_{zw}Q^{ii-1}_{ww}$, then from (5.A.14),

$$\sqrt{T}\left(\hat{\Pi}_i - \Pi_i\right) = Q^{ii}\frac{1}{\sqrt{T}}\sum_{t=1}^{T} w_{it}\epsilon_{it} \qquad i = 1,\ldots,n \tag{5.A.19}$$

Stacking these up for $i = 1,\ldots,n$, noting that $\pi = [\Pi_1' \; \Pi_2' \; \ldots \; \Pi_n']'$, i.e. $\pi = vec(\Pi')$

where $\Pi' = [\Pi_1 \; \Pi_2 \; \ldots \; \Pi_n]$, then one obtains

$$\sqrt{T}\left(\hat{\pi} - \pi\right) = \begin{bmatrix} Q^{11}\frac{1}{\sqrt{T}}\sum_{t=1}^{T} w_{1t}\epsilon_{1t} \\ Q^{22}\frac{1}{\sqrt{T}}\sum_{t=1}^{T} w_{2t}\epsilon_{2t} \\ \vdots \\ Q^{nn}\frac{1}{\sqrt{T}}\sum_{t=1}^{T} w_{nt}\epsilon_{nt} \end{bmatrix} \tag{5.A.20}$$

which can be written

$$\sqrt{T}\left(\hat{\pi} - \pi\right) = \underbrace{\begin{bmatrix} Q^{11} & 0 & & 0 \\ 0 & Q^{22} & & \\ & & \ddots & \\ 0 & & & Q^{nn} \end{bmatrix}}_{Q} \begin{bmatrix} \frac{1}{\sqrt{T}}\sum_{t=1}^{T} w_{1t}\epsilon_{1t} \\ \frac{1}{\sqrt{T}}\sum_{t=1}^{T} w_{2t}\epsilon_{2t} \\ \vdots \\ \frac{1}{\sqrt{T}}\sum_{t=1}^{T} w_{nt}\epsilon_{nt} \end{bmatrix} \tag{5.A.21}$$

$$= Q\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\zeta_t \tag{5.A.22}$$

where

$$\zeta_t = \begin{bmatrix} w_{1t}\epsilon_{1t} \\ w_{2t}\epsilon_{2t} \\ \vdots \\ w_{nt}\epsilon_{nt} \end{bmatrix} \tag{5.A.23}$$

Let $S = \frac{1}{T}\sum_{t=1}^{T}\zeta_t\zeta_t'$ and let $S \longrightarrow E\left[\zeta_t\zeta_t'\right] = S^P$ which is assumed to exist and to be finite. If the population analogues of Q and S are denoted with superscript Ps then

$$\sqrt{T}\left(\hat{\pi} - \pi\right) \xrightarrow{d} N\left(0 \; , \; \underbrace{Q^P S^P Q^{P'}}_{\Sigma_\pi}\right) \tag{5.A.24}$$

In the empirical application of the paper, S^P is estimated using the Newey-West method, i.e.

$$\hat{\pi} \approx N\left(\pi \; , \; \frac{Q\hat{S}Q'}{T}\right) \tag{5.A.25}$$

where Q is as in (5.A.21) and $\hat{S}$ is given by

$$\hat{S} = \hat{S}_0 + \sum_{v=1}^{q} \left[1 - \frac{v}{q+1}\right] \left(\hat{S}_v + \hat{S}_v'\right) \tag{5.A.26}$$

where

$$\hat{S}_v = \begin{bmatrix} \frac{1}{T}\sum_{t=v+1}^{T} w_{1t}\epsilon_{1t}\epsilon_{1t-v}w_{1t-v}' & 0 & \cdots & 0 \\ 0 & \frac{1}{T}\sum_{t=v+1}^{T} w_{2t}\epsilon_{2t}\epsilon_{2t-v}w_{2t-v}' & & \\ \vdots & & \ddots & \\ 0 & & & \frac{1}{T}\sum_{t=v+1}^{T} w_{nt}\epsilon_{nt}\epsilon_{nt-v}w_{nt-v}' \end{bmatrix} \tag{5.A.27}$$

This is the same as estimating the system using instrumental variables equation by equation and making the Newey-West correction on each equation in turn.

5.A.4 Distribution of the impulse response functions

The results outlined here are similar to those given in Lütkepohl (1990) and Hamilton (1994). However, the structural forms of the VAR that they consider are estimated from reduced forms, with the structural parameters backed out of the variance/covariance matrix of the residuals. In which case the distribution of the IRFs will be functions of the distribution of the parameters of the VAR but also of the distribution of the variance/covariance matrix of the errors. Since the structural VAR in Section 5.3 is estimated directly, using instrumental variables, the distribution of the IRFs will depend only on the distribution of the VAR parameters. The differences between these results and those in Lütkepohl (1990) and Hamilton (1994) are non-trivial.

The impulse response functions are given by H_S in (5.20). Using $\hat{h}_S = vec\left(\hat{H}_S'\right)$, the distribution of the IRFs can be calculated from (5.21). For convenience, this is given below in (5.A.28).

$$\sqrt{T}\left(\hat{h}_S - h_S\right) \xrightarrow{d} N\left(0\ ,\ G_S \Sigma_\pi G_S'\right) \tag{5.A.28}$$

G_S is the matrix formed from the derivatives of each of the elements in the vector h_S with respect to each of the elements in the vector π, i.e. $G_S = \frac{\partial h_S}{\partial \pi'}$. Below, I show how these are calculated using $H_S = \Psi_S \left(I_n + B\right)^{-1}$ and noting that Ψ_S and B depend on the

π parameters, the distribution of which is given in (5.17).

$$\begin{aligned} H_S\left(I_n+B\right) &= \Psi_S \\ \Rightarrow \left(I_n+B'\right) H_S' &= \Psi_S' \end{aligned} \tag{5.A.29}$$

Letting η denote an element of π, then differentiating (5.A.29) with respect to η gives

$$\left(I_n+B'\right) \frac{\partial H_s'}{\partial \eta} + \frac{\partial B'}{\partial \eta} H_S' = \frac{\partial \Psi_S'}{\partial \eta} \tag{5.A.30}$$

Using the result that $vec\,(ABC) = (C' \otimes A)\, vec(B)$ and letting $\psi_S = vec\,(\Psi_S')$, then (5.A.30) can be written as

$$\frac{\partial h_S}{\partial \eta} = \left[I_{n^2} + \left(I_n \otimes B'\right)\right]^{-1} \left[\frac{\partial \psi_S}{\partial \eta} - \left(H_S \otimes I_n\right) \frac{\partial vec\,(B')}{\partial \eta}\right] \tag{5.A.31}$$

One then needs expressions for $\frac{\partial \psi_S}{\partial \eta}$ and $\frac{\partial vec(B')}{\partial \eta}$. Start by considering $\frac{\partial vec(B')}{\partial \eta}$. From (5.13), and using the n variable case,

$$B = \begin{bmatrix} 0 & -b_{12} & \cdots & -b_{1n} \\ -b_{21} & 0 & \cdots & -b_{2n} \\ \vdots & & & \vdots \\ -b_{n1} & -b_{n2} & \cdots & 0 \end{bmatrix} \Rightarrow B' = \begin{bmatrix} 0 & -b_{21} & \cdots & -b_{n1} \\ -b_{12} & 0 & \cdots & -b_{n2} \\ \vdots & & & \vdots \\ -b_{1n} & -b_{2n} & \cdots & 0 \end{bmatrix} \tag{5.A.32}$$

Write

$$vec\,(B') = S_{B'} \Theta_{B'} \tag{5.A.33}$$

where $vec\,(B')$ is the $n^2 \times 1$ vector formed by stacking the n^2 elements of B'. Since B has zeros along the diagonal, there are only $n^2 - n$ structural parameters that are estimated. Hence $\Theta_{B'}$ is the $(n^2-n) \times 1$ vector of the $-b_{ij}$ parameters, $i, j = 1, \ldots, n$, $i \neq j$. $S_{B'}$ is the $n^2 \times (n^2-n)$ matrix of zeros and ones that maps the elements of $\Theta_{B'}$ onto $vec\,(B')$. Therefore

$$\frac{\partial vec\,(B')}{\partial \eta} = S_{B'} \frac{\partial \Theta_{B'}}{\partial \eta} \tag{5.A.34}$$

Since $\Pi' = [\Pi_1 \; \Pi_2 \; \ldots \; \Pi_n]$, then Π' can be written as

$$\Pi' = \underbrace{\begin{bmatrix} b_{12} & b_{21} & \cdots & b_{n1} \\ b_{13} & b_{23} & & b_{n2} \\ \vdots & \vdots & & \vdots \\ b_{1n} & b_{2n} & \cdots & b_{nn-1} \\ & c' & & \\ & \phi_1' & & \\ & \vdots & & \\ & \phi_P' & & \end{bmatrix}}_{n\ columns} \begin{matrix} \left.\vphantom{\begin{matrix} a\\a\\a\\a \end{matrix}}\right\} n-1\ rows \\ \left.\vphantom{\begin{matrix} a\\a\\a\\a \end{matrix}}\right\} k\ rows \end{matrix} \tag{5.A.35}$$

Noting that $\pi = vec(\Pi')$, it then becomes clear that $\frac{\partial \Theta_{B'}}{\partial \pi'}$ is an $(n^2 - n) \times n(n-1+k)$ matrix of zeros and ones, i.e.

$$\frac{\partial \Theta_{B'}}{\partial \pi'} = \begin{bmatrix} -I_{n-1} & & & & 0_{(n-1)\times(n-1)(n-1+k)+k} \\ 0_{(n-1)\times(n-1+k)} & -I_{n-1} & & & 0_{(n-1)\times(n-2)(n-1+k)+k} \\ 0_{(n-1)\times 2(n-1+k)} & & -I_{n-1} & & 0_{(n-1)\times(n-3)(n-1+k)+k} \\ \vdots & & & & \vdots \\ 0_{(n-1)\times(n-1)(n-1+k)} & & & -I_{n-1} & 0_{(n-1)\times k} \end{bmatrix} \tag{5.A.36}$$

Since $\frac{\partial \Theta_{B'}}{\partial \pi'}$ is formed by stacking the $n(n-1+k)$ $\frac{\partial \Theta_{B'}}{\partial \eta}$ vectors, each of which are $(n^2 - n) \times 1$ and $\frac{\partial h_S}{\partial \pi'}$ is formed by stacking the $n(n-1+k)$ $\frac{\partial h_S}{\partial \eta}$ vectors, then we have our first expression, the matrix analogue of $\frac{\partial vec(B')}{\partial \eta}$. We now need an expression for $\frac{\partial \psi_S}{\partial \eta}$. From (5.19) one can write

$$\Psi_S' = \Psi_{S-1}'\phi_1' \left((I_n + B)^{-1}\right)' + \Psi_{S-2}'\phi_2' \left((I_n + B)^{-1}\right)' + \ldots + \Psi_{S-P}'\phi_P' \left((I_n + B)^{-1}\right)' \tag{5.A.37}$$

Differentiating this with respect to η, an element of π, and rearranging results in

$$\begin{aligned}\frac{\partial \Psi_S'}{\partial \eta} &= [\phi_1 \Psi_{S-1} + \phi_2 \Psi_{S-2} + \ldots + \phi_P \Psi_{S-P}]' \frac{\partial \left((I_n + B)^{-1} \right)'}{\partial \eta} \\ &+ \Psi_{S-1}' \frac{\partial \phi_1'}{\partial \eta} \left((I_n + B)^{-1} \right)' + \ldots + \Psi_{S-P}' \frac{\partial \phi_P'}{\partial \eta} \left((I_n + B)^{-1} \right)' \\ &+ \frac{\partial \Psi_{S-1}'}{\partial \eta} \left((I_n + B)^{-1} \phi_1 \right)' + \ldots + \frac{\partial \Psi_{S-P}'}{\partial \eta} \left((I_n + B)^{-1} \phi_P \right)'\end{aligned} \tag{5.A.38}$$

and implementing the vec operator, where $\psi_S = vec\,(\Psi_S')$, gives

$$\begin{aligned}\frac{\partial \psi_S}{\partial \eta} &= \left(I_n \otimes [\phi_1 \Psi_{S-1} + \phi_2 \Psi_{S-2} + \ldots + \phi_P \Psi_{S-P}]' \right) \frac{\partial vec \left((I_n + B)^{-1} \right)'}{\partial \eta} \\ &+ \left((I_n + B)^{-1} \otimes \Psi_{S-1}' \right) \frac{\partial vec\,(\phi_1')}{\partial \eta} + \ldots + \left((I_n + B)^{-1} \otimes \Psi_{S-P}' \right) \frac{\partial vec\,(\phi_P')}{\partial \eta} \\ &+ \left((I_n + B)^{-1} \phi_1 \otimes I_n \right) \frac{\partial \psi_{S-1}}{\partial \eta} + \ldots + \left((I_n + B)^{-1} \phi_P \otimes I_n \right) \frac{\partial \psi_{S-P}}{\partial \eta}\end{aligned} \tag{5.A.39}$$

Stack the elements of the $n \times n$ matrix ϕ_j' in an $n^2 \times 1$ vector, $\Theta_{\phi_j'}$, i.e. $\Theta_{\phi_j'} = vec\,(\phi_j')$ then from (5.A.35), it can be seen that $\frac{\partial \Theta_{\phi_j'}}{\partial \pi'}$ can be written as[25]

$$\frac{\partial \Theta_{\phi_j'}}{\partial \pi'} = \begin{bmatrix} 0_{n \times nj} & I_n & 0_{n \times (n-1)(n-1+k)+(P-j)n} \\ 0_{n \times (n-1+k)+nj} & I_n & 0_{n \times (n-2)(n-1+k)+(P-j)n} \\ \vdots & & \vdots \\ 0_{n \times (n-1)(n-1+k)+nj} & I_n & 0_{n \times (P-j)n} \end{bmatrix} \tag{5.A.40}$$

Again, since $\frac{\partial \psi_S}{\partial \pi'}$ is formed by stacking the $n(n-1+k)$ $\frac{\partial \psi_S}{\partial \eta}$ vectors, then (5.A.40) gives us the matrix analogue terms of $\frac{\partial vec(\phi_j')}{\partial \eta}$. The last term we need is $\frac{\partial vec\left((I_n+B)^{-1}\right)'}{\partial \eta}$. Using the results of Magnus and Neudecker (1988), pages 96, 148 and 151, one can see that

$$\frac{\partial vec \left((I_n + B)^{-1} \right)'}{\partial \eta} = - \left[(I_n + B)^{-1} \otimes \left((I_n + B)^{-1} \right)' \right] S_{B'} \frac{\partial \Theta_{B'}}{\partial \pi'} \tag{5.A.41}$$

where $S_{B'}$ and $\frac{\partial \Theta_{B'}}{\partial \pi'}$ are defined as before. Substituting (5.A.41) into (5.A.39) and sum-

[25]If we have different lag lengths of returns and flows, as in the empirical model presented in Section 5.3, these $\frac{\partial \Theta_{\phi_j'}}{\partial \pi'}$ matrices need to be altered slightly. Also $vec\,(\phi_j') = S_{\phi_j'} \Theta_{\phi_j'}$ and $S_{\phi_j'}$ will not in general equal the identity matrix.

ming the elements in the second and third rows will give us our definition of $\frac{\partial \psi_S}{\partial \pi'}$ when we stack the $n(n-1+k)$ $\frac{\partial \psi_S}{\partial \eta}$ vectors. This is (5.23) in the text. Stacking the $n(n-1+k)$ $\frac{\partial h_S}{\partial \eta}$ vectors in (5.A.31) will give the result in (5.22), which we can then use to calculate the distribution of the IRFs.

5.A.5 Extra information in buys and sells separately

The structural VAR in (5.12) which allows contemporaneous feedback trading is clearly unidentified. Compared to the just identified reduced form model, there are two extra parameters to estimate; contemporaneous effects of order flow on FX returns and returns on FX order flows, while only one restriction is available, that the errors in the two equations are uncorrelated.[26] Therefore the model can only be estimated via the use of instrumental variables. However, since we have information on the number of buyer initiated and seller initiated trades separately, could we not use this extra information to instrument for order flow? Order flow is defined as the number of buyer *less* seller initiated trades, but could we use trading volume, the number of buyer *plus* seller initiated trades as an instrument? The answer to this question is 'no'! This is proved below. Suppose our two equation contemporaneous feedback VAR is given by (5.A.42).

$$\begin{aligned} R_t &= \alpha_1 + \beta_1 F_t + \phi_1 z_t + \epsilon_t^R \\ F_t &= \alpha_2 + \beta_2 R_t + \phi_2 z_t + \epsilon_t^F \end{aligned} \tag{5.A.42}$$

As before, R_t is the log first difference in prices in time t, F_t is order flow in that period and z_t contains the lags in the VAR. ϵ_t^R and ϵ_t^F are the errors in the return and flow equations respectively, with variances $\sigma^2_{\epsilon^R}$ and $\sigma^2_{\epsilon^F}$ and are assumed to be uncorrelated. In matrix form (5.A.42) can be written

$$\begin{bmatrix} 1 & -\beta_1 \\ -\beta_2 & 1 \end{bmatrix} \begin{bmatrix} R_t \\ F_t \end{bmatrix} = \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} + \begin{bmatrix} \phi_1 \\ \phi_2 \end{bmatrix} z_t + \begin{bmatrix} \epsilon_t^R \\ \epsilon_t^F \end{bmatrix} \tag{5.A.43}$$

Since we have information on buys and sells separately, we can write order flow, F_t, and

[26]One could make exclusion restrictions in the lags of the VAR but such restrictions were described as being 'incredible' by Sims (1980), and such practice is not performed in empirical work.

volume, V_t, as

$$F_t = B_t - S_t$$
$$V_t = B_t + S_t \tag{5.A.44}$$

B_t and S_t are the number of buyer and seller initiated trades in period t respectively.[27] We can augment (5.A.43) with the definition of volume to obtain

$$\underbrace{\begin{bmatrix} 1 & -\beta_1 & \beta_1 \\ -\beta_2 & 1 & -1 \\ 0 & 1 & 1 \end{bmatrix}}_{B} \begin{bmatrix} R_t \\ B_t \\ S_t \end{bmatrix} = \underbrace{\begin{bmatrix} \alpha_1 \\ \alpha_2 \\ 0 \end{bmatrix}}_{\alpha} + \underbrace{\begin{bmatrix} \phi_1 \\ \phi_2 \\ 0 \end{bmatrix}}_{\phi} z_t + \underbrace{\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}}_{\lambda} V_t + \begin{bmatrix} \epsilon_t^R \\ \epsilon_t^F \\ 0 \end{bmatrix} \tag{5.A.45}$$

It then becomes clear that (5.A.45) is not identified and hence one cannot use volume as an instrument to estimate the structural VAR. Consider the order and rank conditions for the identification of the first equation of (5.A.45), the return equation. Imposing the restriction that the coefficient on buys is equal to minus the coefficient on sells, implied by (5.A.42), together with the restriction that volume does not enter the return equation, gives 2 restrictions $\geq N - 1$ where $N = 3$ equations. Therefore the order condition is passed. However, for the rank condition to be passed, the matrix in (5.A.46) must be non-singular, i.e. have a non-zero determinant. The matrix is formed from the elements of the columns in (5.A.45) that are used as restrictions, i.e. for the first equation, the second and third elements of B sum to zero. Therefore the first column of the matrix in (5.A.46) is made up of the sum of the second and third elements of the second and third rows of B. The second column contains the second and third elements of the vector premultiplying V_t in (5.A.45).

$$\text{Rank condition} \Rightarrow \begin{vmatrix} 0 & 0 \\ 2 & 1 \end{vmatrix} \neq 0 \qquad \text{Contradiction!} \tag{5.A.46}$$

Therefore the rank condition is not satisfied and hence the return equation is not identi-

[27]The intra-day pattern in trading volume can also be purged from these series by subtracting the diurnal pattern from both buy and sell series in a similar fashion to Chapter 2. B_t and S_t will then be deseasonalised buys and sells.

fied. Similar arguments can be used to show that the flow equation, the second equation in (5.A.45), is not identified, again due to the failure of the rank condition. Therefore, even though there is extra information available in the number of buyer and seller initiated trades separately, this information cannot be used to estimate the contemporaneous feedback VAR of (5.A.42). Instruments must be found elsewhere and in Section 5.3 they are obtained from statistics in other, related markets.

Chapter 6 Conclusions

This thesis has explored a number of different issues in foreign exchange market microstructure. The first part considered the effects of scheduled macroeconomic news releases on a number of statistics, including the level of the exchange rate, spreads, depths, trading activity and volatility, and also asked *how* public information is incorporated into price.

Following the disappointing results from macroeconomic models when explaining short run exchange rate movements, emphasis has shifted towards microstructure and higher frequency investigations. However, as emphasised in Lyons (2001), this does not suggest that macroeconomic and microstructure explanations are totally separated. The link between them is demonstrated in Chapters 2 to 4, when the high frequency effects of scheduled macroeconomic news are studied. Contrary to the standard rational expectations and efficient markets hypotheses, whereby public information is incorporated into prices without the need for trading, I find that the transactions process is vital when impounding new macroeconomic data. Not only does a release of news cause trading (buyer *and* seller initiated) to increase, even when the announcement has no discernable effects on the level of the exchange rate, but trading (more precisely — order flow) is in fact the mechanism through which this information gets incorporated into price.

Order flow is therefore extremely important in the assimilation of information. The microstructure theories of Kyle (1985), Glosten and Milgrom (1985), etc. all suggest that order flow is the route through which private information enters price, and this is confirmed in the empirical results of Lyons (1995), Rime (2001), Evans (2002) and Payne (2003a) among others. If order flow is also the route through which *public* information enters price, then this suggests a much more important role played by order flow models, and microstructure theories in general.

However, the importance of order flow is often measured by the price impact of trades; the greater the price impact, the more information trades are argued to carry. These information measures are commonly obtained by fitting the return and order flow data to a Vector Auto-Regression model and calculating the cumulative price change following an (orthogonalised) order flow shock. The models that are estimated assume that the direction of causality runs explicitly from order flow to returns, i.e. returns depend on contemporaneous order flow but the converse is ruled out. This assumption is reasonable if data are sampled at ultra high frequencies, such as tick-by-tick, but as soon as data are aggregated, the problem of contemporaneous feedback trading may become pervasive. Chapter 5 proves that the commonly estimated model is misspecified whenever aggregated data are considered and shows that when allowing for contemporaneous feedback trading, the price impact of order flow shocks is greater than when such trading strategies are prohibited. Therefore, not only does order flow help in the process of *public* information assimilation, it also carries more *private* information than previous estimates suggest. The importance of order flow is therefore paramount.

In Chapter 3, up to two thirds of the price relevant information contained in public news announcements was found to enter via order flow, argued to be the mechanism through which *private* information enters price. This suggests that the distinction between public and private information may not be entirely accurate. It may well be the case that traders disagree on the price implications of a given news announcement, and theoretical models have shown that such differences of opinion can generate the observed increases in trading volume. However, despite being able to explain the increase in volume, the 'differences of opinion' story cannot explain the systematic effects on order flow. In Chapter 4, I present a simple argument, *based* on differences of opinion, that can explain why good (bad) news leads to positive (negative) order flow. If traders not only differ in their interpretation of news releases, but also differ in their abilities to interpret news (by traders receiving signals with different precisions) then this is likely to generate signed order flow.

Also, if traders differ in their interpretation of news, this will enhance the problem of asymmetric information as traders face the risk of being on the wrong side of a more informed trade. To compensate, traders will reduce their depths (the quantities they are willing to trade for any deterioration of the bid and ask prices) and this is indeed found in Chapter 4, where the effects of an unexpected increase in US PPI causes the depth of the DEM/USD limit order book to fall.

This thesis not only looks at the role played by order flow around times of public news announcements, it also examines the effects that these data releases have on other microstructure variables. In Chapter 2, news is found to increase volatility and trading volume significantly. However, whereas the effects on volume appear to be relatively short lived, the effects on volatility persist for some time, with half lives of approximately

15 or 20 minutes. Spreads are also found to increase immediately before and after a news release, but these effects are not statistically significant. By analysing the triangle of rates between the US dollar, the euro and sterling simultaneously, I am also able to examine the cross market effects of news and the complex dynamics and interrelationships between these three markets. For example US news is found to have significant effects, not only on the dollar currency markets, but also on the GBP/EUR market. UK news, on the other hand, is found to have no significant effects on the USD/EUR market, which is not surprising considering the size of the UK economy.

When considering the possibilities for future research, the interactions between microstructure and macroeconomics, demonstrated in Chapters 2 to 4, is an obvious candidate. Theoretical advances are required that combine order flow (commonly believed to be a determinant of exchange rates at high frequencies) and macroeconomic fundamentals, and work in this area is already beginning. Evans and Lyons (2004b) propose a new model of exchange rate determination, adding dispersed information in a dynamic general equilibrium setting. Therefore Evans and Lyons (2004b) attempt to combine the microfoundations associated with market microstructure (information and institutions) with the microfoundations associated with the new open economy macro models (tastes and technology).[1] Bacchetta and van Wincoop (2003) also combine the microstructure and macro approaches by introducing investor heterogeneity (dispersed information about fundamentals) to the standard monetary model of exchange rate determination. A similar model is presented in Breedon and Vitale (2004). Not only do Breedon and Vitale (2004) combine microstructure and macroeconomic elements in a model that allows a role for order flow, they also test the model using all USD/EUR transactions from both EBS and Reuters D2000-2 platforms. Their data then offer the broadest market coverage to date in any study of order flow effects and should therefore give a clearer picture of how exchange rates evolve.

However, the availability of data is just as important for future research as new, more complete, models of exchange rate determination. The ten months of high frequency transactions data used in this thesis represents a very small sample when looked at in a macroeconomic perspective. Much larger datasets that span longer time periods are therefore needed if we are to successfully combine microstructure and macroeconomic models. The fifteen years of high frequency data used in Faust, Rogers, Wang, and

[1]This terminology is borrowed from Lyons (2001), where the differences between the 'two Is' and the 'two Ts' are described in more detail.

Wright (2003) allow a number of macroeconomic hypotheses to be tested, but since these data are indicative, no information on transactions (and therefore on order flow) are available. As such, this lengthy dataset can only go so far when merging microstructure and macroeconomic hypotheses.

Longer datasets could be used to test/confirm the time varying effects of macroeconomic news. It has been suggested that traders focus on different sets of macroeconomic announcements, depending on the current phase of the business cycle, and perhaps on current fashionable theories, and although preliminary work in this area has been done in Andersen, Bollerslev, Diebold, and Vega (2003), Galati and Ho (2003) and Andersen, Bollerslev, Diebold, and Vega (2004), none of these studies employ transactions data, which are shown to be of upmost importance in this thesis. Unfortunately it is unlikely that comprehensive datasets that cover long time intervals will become available for academic research. The majority of inter-dealer trades occur on the Reuters D2000-2 and EBS platforms and these systems have only been around for little over 10 years. As such, studies using end customer data may be a useful avenue for future research, especially if these trades are the main driver of inter-dealer flows. Work in this area is also starting, see for example, Fan and Lyons (2003) and Evans and Lyons (2004a).

The results in Chapters 2 and 3 show important high frequency interactions between the US dollar, euro and pound sterling, and also the cross market effects of macroeconomic news. It would therefore be interesting to see whether such interrelationships exist between different asset classes, i.e. between foreign exchange, bond and equity markets in different countries. A number of research papers have examined the cross market effects between assets, such as Rigobon and Sack (2003), Faust, Rogers, Wang, and Wright (2003) and Andersen, Bollerslev, Diebold, and Vega (2004). However, again, none of these studies employ transactions data. The dynamics of larger systems that allow cross market effects of FX, bond and equity order flows may therefore be a promising line of research.[2]

[2]Hau and Rey (2002) build a model in which exchange rates, stock prices and *capital flows* are determined endogenously. However, as mentioned in Chapter 4 when discussing the portfolio model of Brennan and Cao (1997), there is a fundamental difference between capital flows and order flow. Capital flows show changes in relative demands between two (sets of) agents and this could be associated with either positive, negative, or indeed zero order flow, depending on who initiated the trades.

References

Acker, D., M. Stalker, and I. Tonks, 2002, Daily inside closing spreads and trading volumes around earning announcements, Discussion Paper 404, Financial Markets Group, London School of Economics.

Allen, H., and M. Taylor, 1990, Charts, noise and fundamentals in the London foreign exchange market, Economic Journal, 100, 49–59.

Almeida, A., C. Goodhart, and R. Payne, 1998, The Effects of Macroeconomic 'News' on High-Frequency Exchange Rate Behaviour, Journal of Financial and Quantitative Analysis, 33(3), 383–440.

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.

———, 1997b, Intraday periodicity and volatility persistence in financial markets, Journal of Empirical Finance, 4, 115–158.

———, 1998, Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies, Journal of Finance, 53(1), 219–265.

Andersen, T., T. Bollerslev, F. Diebold, and C. Vega, 2003, Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange, American Economic Review, 93, 38–62.

———, 2004, Real-Time Price Discovery in Stock, Bond and Foreign Exchange Markets, Working Paper.

Bacchetta, P., and E. van Wincoop, 2003, Can information heterogeneity explain the exchange rate determination puzzle?, Working Paper.

Baillie, R., and T. Bollerslev, 1991, Intra-Day and Inter-Market Volatility in Foreign Exchange Rates, Review of Economic Studies, 58, 565–586.

Balduzzi, P., E. Elton, and T. Green, 2001, Economic news and the Yield Curve : Evidence from the US Treasury Market, Journal of Financial and Quantitative Analysis, 36, 523–545.

Balke, N., and T. Fomby, 1997, Threshold cointegration, International Economic Review, 38, 627–645.

Banerjee, A., J. Dolado, J. Galbraith, and D. Hendry, 1993, Co-integration, Error Correction, and the Econometric Analysis of Non-stationary Data. Oxford University Press.

Bank for International Settlements, 2001, 71st annual report, section 5, 'Foreign exchange markets'. B.I.S., Basle.

Ben Omrane, W., and A. Heinen, 2003, The response to individual FX dealers' quoting activity to macroeconomic news announcements, Working Paper.

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.

Bessembinder, H., K. Chan, and P. Seguin, 1996, An empirical examination of information, differences of opinion, and trading activity, Journal of Financial Economics, 40, 105–134.

Biais, B., P. Hillion, and C. Spatt, 1995, An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse, Journal of Finance, 50, 1655–1689.

Biais, B., D. Martimort, and J. Rochet, 2000, Competing mechanisms in a common value environment, Econometrica, 68, 799–837.

Bjønnes, G., and D. Rime, 2003, Dealer behaviour and trading systems in foreign exchange markets, Working Paper.

Blanchard, O., and D. Quah, 1989, The dynamic effects of aggregate demand and supply disturbances, American Economic Review, 79, 655–673.

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.

Bound, J., D. Jaeger, and R. Baker, 1995, Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak, Journal of the American Statistical Association, 90, 443–450.

Breedon, F., and P. Vitale, 2004, An empirical study of liquidity and information effects of order flow on exchange rates, Working Paper.

Brennan, M., and H. Cao, 1997, International portfolio investment flows, Journal of Finance, 52, 1851–1880.

Brunnermeier, M., 2001, Asset Pricing Under Asymmetric Information: Bubbles, crashes, technical analysis and herding. Oxford University Press.

Buse, A., 1992, The bias of instrumental variable estimators, Econometrica, 60, 173–180.

Carlson, J., and M. Lo, 2004, One Minute in the Life of the DM/$: Public Information in an Electronic Market, .

Carlson, W., P. Newbold, and B. Thorne, 2003, Statistics for Business and Economics. Prentice Hall, Upper Saddle River, N.J., U.S.A.

Chaboud, A., S. Chernenko, E. Howorka, R. Iyer, D. Liu, and J. Wright, 2004, The high frequency effects of U.S. macroeconomic data releases on prices and trading activity in the global interdealer foreign exchange market, Working Paper, Federal Reserve.

Chakravarty, S., and C. Holden, 1995, An integrated model of market and limit orders, Journal of Financial Intermediation, 4, 213–241.

Chari, A., 2002, An analysis of macroeconomic announcements across locations in the DM/dollar market, Working Paper.

Cheung, Y., and M. Chinn, 1999, Macroeconomic implications of the beliefs and behaviour of foreign exchange traders, NBER Working Paper 7417.

Cohen, B., and H. Shin, 2003, Positive feedback trading under stress: Evidence from the US Treasury securities market, Working Paper, London School of Economics.

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

Covrig, V., and M. Melvin, 2002, Asymmetric information and price discovery in the FX market: does Tokyo know more about the yen?, Journal of Empirical Finance, 9, 271–285.

Creel, M., and M. Farell, 1996, SUR estimation of multiple time series models with heteroscedasticity and serial correlation of unknown form, Economics Letters, 53, 239–245.

Daníelsson, J., and R. Love, 2004, Feedback trading, Discussion Paper 510, Financial Markets Group, London School of Economics.

Daníelsson, J., J. Luo, and R. Payne, 2002, Exchange rate determination and inter-market order flow effects, mimeo, LSE.

Daníelsson, J., and R. Payne, 2001, Measuring and explaining liquidity on an electronic limit order book: evidence from Reuters D2000-2, Working Paper, Financial Markets Group, London School of Economics.

——— , 2002a, Liquidity Determination in an Order Driven Market, Working Paper, Financial Markets Group, London School of Economics.

——— , 2002b, Real Trading Patterns and Prices in Spot Foreign Exchange Markets, Journal of International Money and Finance, 21, 203–222.

de Jong, F., R. Mahieu, P. Schotman, and I. van Leeuwen, 2001, Price discovery on foreign exchange markets with differentially informed traders, Working Paper.

De Long, J., A. Shleifer, L. Summers, and R. Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation, Journal of Finance, 45, 379–395.

DeGennaro, R., and R. Schrieves, 1997, Public information releases, private information arrival and volatility in the foreign exchange market, Journal of Empirical Finance, 4, 295–315.

Dennert, J., 1993, Price Competition between Market Makers, Review of Economic Studies, 60, 735–751.

Derviz, A., 2002, Dealer quotes, order flow and indirect foreign currency utility in a multiple dealership market, Working 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.

Dominguez, K., 2003, When do Central Bank interventions influence intra-daily and longer-term exchange rate movements?, NBER Working Paper 9875.

Dufour, A., and R. Engle, 2000, Time and the Price Impact of a Trade, Journal of Finance, 55, 2467–2498.

Easley, D., and M. O'Hara, 1987, Price, Trade Size and Information in Securities Markets, Journal of Financial Economics, 19, 69–90.

———, 1992, Time and the process of security price adjustment, Journal of Finance, 47, 577–605.

Ederington, L., and J. Lee, 1993, How Markets Process Information: News Releases and Volatility, Journal of Finance, 48(4), 1161–1191.

———, 1995, The Short-Run Dynamics of the Price Adjustment to New Information, Journal of Financial and Quantitative Analysis, 30(1), 117–134.

Efron, B., 1979, Bootstrapping methods: Another look at the Jackknife, Annals of Statistics, 7, 1–26.

Efron, B., and R. Tibshirani, 1993, An Introduction to the Bootstrap. Chapman and Hall, New York.

Enders, W., 1995, Applied Econometric Time Series. New York: Wiley.

Engle, R., T. Ito, and W. Lin, 1990, Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market, Econometrica, 58, 525–542.

Engle, R., and A. Patton, 2004, Impacts of trades in an error-correction model of quote prices, Journal of Financial Markets, 7, 1–25.

Evans, M., 2001, FX Trading and Exchange Rate Dynamics, NBER Working Paper 8116.

———, 2002, FX trading and exchange rate dynamics, Journal of Finance, 57, 2405–2447.

Evans, M., and R. Lyons, 1999, Order flow and Exchange Rate Dynamics, .

Evans, M., and R. Lyons, 2001, Portfolio balance, price impact and secret intervention, NBER Working Paper 8356.

———, 2002a, Informational integration and FX trading, Journal of International Money and Finance, 21, 807–831.

Evans, M., and R. Lyons, 2002b, Order flow and Exchange Rate Dynamics, Journal of Political Economy, 110, 170–180.

Evans, M., and R. Lyons, 2002c, Time-varying liquidity in foreign exchange, Journal of Monetary Economics, 49, 1025–1051.

———, 2003, How is Macro News Transmitted to Exchange Rates?, Working Paper.

———, 2004a, Do currency markets absorb news quickly?, Journal of International Money and Finance, forthcoming.

———, 2004b, A new micro model of exchange rate dynamics, NBER Working Paper 10379.

Fan, M., and R. Lyons, 2003, Customer trades and extreme events in foreign exchange, in P. Mizen , eds., Central banking, monetary theory and practice : essays in honour of Charles Goodhart. Edward Elgar.

Faust, J., J. Rogers, S. Wang, and J. Wright, 2003, The high-frequency response of exchange rates and interest rates to macroeconomic announcements, Working Paper.

Fiebig, D. G., 2001, Seemingly unrelated regressions, in B. H. Baltagi , eds., A companion to theoretical econometrics. Blackwell.

Fischer, A., 2003, Reuters news reports versus official interventions: A cautionary warning, Working Paper.

Fleming, M., 2001, Measuring treasury market liquidity, staff report, Federal Reserve Bank of New York.

Fleming, M., and E. Remolona, 1999, Price Formation and Liquidity in the U.S.Treasury Market: The Response to Public Information, Journal of Finance, 54, 1901–1915.

Flood, M., 1991, Microstructure Theory and the Foreign Exchange Market, Federal Reserve Bank of St. Louis Review, 73, 52–70.

Flood, R., and M. Taylor, 1996, Exchange rate economics: What's wrong with the conventional macro approach, in J. Frankel, G. Galli, and A. Giovannini , eds., The Microstructure of Foreign Exchange Markets. University of Chicago Press, Chicago, IL., pp. 261–294.

Foucault, T., 1999, Order flow Composition and Trading Costs in a Dynamic Limit Order Market, Journal of Financial Markets, 2, 99–134.

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

Frankel, J., and A. Rose, 1994, A survey of empirical research on nominal exchange rates, NBER Working Paper 4865.

Freedman, D., 1984, On bootstrapping two-stage least-squares estimates in stationary linear models, The Annals of Statistics, 12, 827–842.

French, K., and R. Roll, 1986, Stock Return Variances: The Arrival of Information and the Reaction of Traders, Journal of Financial Economics, 17, 5–26.

Froot, K., and T. Ramadorai, 2002, Currency returns, institutional investor flows and exchange rate fundamentals, NBER Working Paper 9101.

Galati, G., and C. Ho, 2003, Macroeconomic news and the euro/dollar exchange rate, Economic Notes, 32, 371–398.

Gallant, R., 1981, On the bias in flexible functional forms and an essentially unbiased form: The Fourier flexible form, Journal of Econometrics, 15, 211–245.

Glosten, L., 1994, Is the Electronic Open Limit Order Book Inevitable?, Journal of Finance, 49, 1127–1161.

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., and A. Demos, 1991, Reuters' screen images of the foreign exchange markets: The Yen/Dollar and Sterling/Dollar spot market, Journal of the International Securities Markets, 5, 35–64.

Goodhart, C., and L. Figliuoli, 1991, Every Minute Counts in Financial Markets, Journal of International Money and Finance, 10, 23–52.

———, 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., and M. Giugale, 1989, Some evidence on daily trading activity in the London foreign exchange market, Journal of the International Securities Markets, pp. 137–146.

———, 1993, From hour to hour in the foreign exchange market, The Manchester School Journal, 61, 1–34.

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 J. Frankel, G. Galli, and A. Giovannini , eds., The Microstructure of Foreign Exchange Markets. University of Chicago Press, Chicago, IL., pp. 107–179.

Goodhart, C., R. Love, R. Payne, and D. Rime, 2002, Analysis of spreads in the dollar/euro and deutschemark/dollar foreign exchange markets, Economic Policy, 35, 537–552.

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.

Green, C., 2004, Economic news and the impact of trading on bond prices, Journal of Finance, 59, 1201–1233.

Greene, W., 2000, Econometric analysis. Prentice Hall, Upper Saddle River, N.J., U.S.A.

Grossman, S., 1976, On the efficiency of competitive stock markets where trades have diverse information, Journal of Finance, 31, 573–585.

Grossman, S., and J. Stiglitz, 1980, On the impossibility of informationally efficiency markets, American Economic Review, 70, 393–408.

Hakkio, C., and D. Pearce, 1985, The Reaction of Exchange Rates to Economic News, Economic Inquiry, 23, 621–635.

Hall, A., G. Rudebusch, and D. Wilcox, 1996, Judging instrument relevance in instrumental variables estimation, International Economic Review, 37, 283–298.

Hamilton, J., 1994, Time Series Analysis. Princeton University Press, Princeton, New Jersey.

Handa, P., and R. Schwartz, 1996, Limit Order Trading, Journal of Finance, 51, 1835–1861.

Hargreaves, C., 1994, Á review of methods of estimating cointegrating relationships, in C. Hargreaves , eds., Nonstationary Time Series Analysis and Cointegration. Oxford University Press, pp. 87–131.

Harris, L., and J. Hasbrouck, 1996, Market vs. Limit Orders: The SuperDOT Evidence on Order Submission Strategy, Journal of Financial and Quantitative Analysis, 31, 213–231.

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

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

———, 1999, Trading Volumes and Transaction Costs in the Foreign Exchange Market: Evidence from Daily Dollar-Yen Spot Data, Journal of Banking and Finance, 23, 801–824.

Harvey, A., 1989, Forecasting, Structural Time-Series Models and the Kalman Filter. Cambridge University Press, Cambridge.

Harvey, C., and R. Huang, 1991, Volatility in the Currency Futures Market, Review of Financial Studies, 4(3), 543–569.

Hasbrouck, J., 1991a, Measuring the Information Content of Stock Trades, Journal of Finance, 46(1), 179–206.

———, 1991b, The Summary Informativeness of Stock Trades: An Econometric Analysis, Review of Financial Studies, 4(3), 571–595.

———, 1999, Security bid/ask dynamics with discreteness and clustering: Simple strategies for modelling and estimation, Journal of Financial Markets, 2, 1–28.

Hau, H., and H. Rey, 2002, Exchange rate, equity prices and capital flows, NBER Working Paper 9398.

Hausman, J., A. Lo, and A. Mackinlay, 1992, An ordered probit analysis of transaction stock prices, Journal of Financial Economics, 13, 317–379.

Hendry, D., 1995, Dynamic Econometrics. Oxford University Press, Oxford.

Ho, T., and H. Stoll, 1983, The Dynamics of Dealer Markets under Competition, Journal of Finance, 38(4), 1053–1074.

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.

Hsieh, D., and A. Kleidon, 1996, Bid-Ask Spreads in Foreign Exchange Markets: Implications for Models of Asymmetric Information, in J. Frankel, G. Galli, and A. Giovannini , eds., The Microstructure of Foreign Exchange Markets. University of Chicago Press.

Huang, R., and H. Stoll, 1997, The Components of the Bid-Ask Spread: a General Approach, Review of Financial Studies, 10, 995–1034.

Ito, T., 1990, Foreign exchange rate expectations: Micro survey data, American Economic Review, 80, 434–449.

Ito, T., and Y. Hashimoto, 2004, Microstructure of the yen/dollar foreign exchange market: Patterns of intra-day activity revealed in the electronic broking system, NBER Working Paper 10856.

Ito, T., R. Lyons, and M. Melvin, 1998, Is There Private Information in the FX market: The Tokyo Experiment, Journal of Finance, 53(3), 1111–1130.

Jain, p., 1988, Response of Hourly Stock Prices and Trading Volume to Economic News, Journal of Business, 61, 219–231.

Jegadeesh, N., and S. Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, Journal of Finance, 48, 65–91.

Johansen, S., 1988, Statistical analysis of cointegration vectors, Journal of Economic Dynamics and Control, 12, 231–254.

———, 1991, Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models, Econometrica, 59, 1551–1580.

———, 1995, Likelihood-based inference in cointegrated vector auto-regressive models. Oxford University Press.

Jones, C., G. Kaul, and M. Lipson, 1994, Transactions, Volume and Volatility, Review of Financial Studies, 7, 631–651.

Kandel, E., and N. Pearson, 1995, Differential interpretation of public signals and trade in speculative markets, Journal of Political Economy, 103, 831–872.

Kavajecz, K., 1999, A Specialist's Quoted Depth and the Limit Order Book, Journal of Finance, 54, 747–771.

Killeen, W., R. Lyons, and M. Moore, 2002, Fixed versus flexible: Lessons from EMS order flow, Working Paper.

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

———, 1994, Market liquidity and volume around earnings announcements, Journal of Accounting and Economics, 17, 41–67.

———, 1997, Pre-announcement and event-period private information, Journal of Accounting and Economics, 24, 395–419.

Kollias, C., and K. Metaxas, 2001, How efficient are FX markets? Empirical evidence of arbitrage opportunities using high-frequency data, Applied Financial Economics, 11, 435–444.

Kyle, A., 1985, Continuous Auctions and Insider Trading, Econometrica, 53, 1315–1335.

Laakkonen, H., 2004, The impact of macroeconomic news on exchange rate volatility, Bank of Finland Discussion Paper.

Lee, C., B. Mucklow, and M. Ready, 1993, Spreads, Depths and the Impact of Earnings Information: an Intraday Analysis, Review of Financial Studies, 6(2), 345–374.

Love, R., and C. Goodhart, 2004, Order book depth and foreign exchange returns, Working Paper, Financial Markets Group, London School of Economics.

Love, R., and R. Payne, 2003, Macroeconomic news, order flows and exchange rates, Discussion Paper 475, Financial Markets Group, London School of Economics.

Luo, J., 2002a, Limit order execution probability: Evidence from the London Stock Exchange, mimeo, LSE.

——— , 2002b, Market conditions, order flow and exchange rate determination, mimeo, LSE.

Lütkepohl, H., 1990, Asymptotic distributions of impulse response functions and forecast error variance decompositions of vector autoregressive models, The Review of Economics and Statistics, 72, 116–125.

Lütkepohl, H., and H. Reimers, 1992, Impulse response analysis of cointegrated systems, Journal of Economic Dynamics and Control, 16, 53–78.

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

——— , 1996, Foreign Exchange Volume: Sound and Fury Signifying Nothing?, in J. Frankel, G. Galli, and A. Giovannini , eds., The Microstructure of Foreign Exchange Markets. University of Chicago Press.

——— , 1997, A simultaneous trade model of the foreign exchange hot potato, Journal of International Economics, 42, 275–298.

——— , 2001, The Microstructure Approach to Exchange Rates. MIT Press, Boston, U.S.A.

Lyons, R., and M. Moore, 2003, An information approach to international currencies, Working Paper.

Madhavan, A., 1992, Trading mechanisms in securities markets, Journal of Finance, 47, 607–641.

———, 2000, Market microstructure: asurvey, Journal of Financial Markets, 3, 205–258.

Madhavan, A., and S. Smidt, 1991, A Bayesian Model of Intraday Specialist Pricing, Journal of Financial Economics, 30, 99–134.

Magnus, J., and H. Neudecker, 1988, Matrix Differential Calculus with Applications in Statistics and Econometrics. New York: Wiley.

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

Meese, R., and K. Rogoff, 1983a, Empirical Exchange Rate Models of the Seventies: Do they Fit out of Sample?, Journal of International Economics, 14, 3–24.

——— , 1983b, The Out-of-Sample Failure of Empirical Exchange Rate Models, in J. Frankel , eds., Exchange Rates and International Macroeconomics., vol. 14. University of Chicago Press, Chicago, IL.

Melvin, M., M. Sager, and M. Taylor, 2003, Can an old lady keep a secret? A microstructural study of monetary policy committee announcements and the foreign exchange market, .

Melvin, M., and X. Yin, 2000, Public Information Arrival, Exchange Rate Volatility and Quote Frequency, Economic Journal, 110, 644–661.

Mitchell, M., and J. Mulherin, 1994, The Impact of Public Information on the Stock Market, Journal of Finance, 49, 923–950.

Morse, D., and N. Ushman, 1983, The effect of information announcements on the market microstructure, The Accounting Review, 58, 247–258.

O'Hara, M., 1995, Market Microstructure Theory. Blackwell Publishers, Oxford, U.K.

——— , 2003, Presidential address: Liquidity and price discovery, Journal of Finance, 58, 1335–1354.

Osler, C., 2002, Stop-loss orders and price cascades in currency markets, working paper, Federal Reserve Bank of New York.

Osterwald-Lenum, M., 1992, A note with quantiles of the asymptotic distribution of the maximum likelihood cointegration rank test statisics, Oxford Bulletin of Economics and Statistics, 54, 461–471.

Pagan, A., and J. Robertson, 1998, Structural models of the liquidity effect, The Review of Economics and Statistics, 80, 202–217.

Parlour, C., 1998, Price Dynamics in Limit Order Markets, Review of Financial Studies, 11(4), 789–816.

Payne, R., 1996, Announcement Effects and Seasonality in the Intra-day Foreign Exchange Market, Discussion Paper 238, Financial Markets Group, London School of Economics.

———, 2003a, Informed trade in spot foreign exchange markets: an empirical investigation, Journal of International Economics, 61, 307–329.

———, 2003b, Trading activity, volatility and transactions costs in spot FX markets, in P. Mizen , eds., Central banking, monetary theory and practice : essays in honour of Charles Goodhart. Edward Elgar.

Payne, R., and P. Vitale, 2003, A transaction level study of the effects of Central Bank intervention on exchange rates, Journal of International Economics, 61, 331–352.

Perraudin, W., and P. Vitale, 1996, Interdealer Trade and Information Flows in the Foreign Exchange Market, in J. Frankel, G. Galli, and A. Giovannini , eds., The Microstructure of Foreign Exchange Markets. University of Chicago Press.

Pesaran, M., and R. Smith, 1998, Structural analysis of cointegrating VARs, Journal of Economic Surveys, 12, 471–505.

Rigobon, R., 2003, Identification through heteroskedasticity, Review of Economic and Statistics, 85, 777–792.

Rigobon, R., and B. Sack, 2003, Spillovers across U.S. financial markets, NBER Working Paper 9640.

Rime, D., 2001, Private or Public Information in Currency Markets: an Empirical Investigation, Working Paper.

Runkle, D., 1987, Vector autoregressions and reality, Journal of Business and Economic Statistics, 5, 437–442.

Sarno, L., and M. Taylor, 2001, The microstructure of the foreign exchange market: a selective survey of the literature, Princeton Studies in International Economics number 89.

Shea, J., 1996, Instrument relevance in multivariate linear models: a simple measure, unpublished manuscript, University of Wisconsin.

Sims, C., 1980, Macroeconomics and reality, Econometrica, 48, 1–48.

Staiger, D., and J. Stock, 1997, Instrumental variables regression with weak instruments, Econometrica, 65, 557–586.

Taylor, M., 1995, The economics of exchange rates, Journal of Economic Literature, 33, 13–47.

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

Tsay, R., 2002, Analysis of Financial Time Series. Wiley-Interscience.

Varian, H., 1989, Differences of opinion in financial markets, in C. Stone , eds., Financial risk: theory, evidence and implications, proceedings of the eleventh annual economic policy conference of the Federal Reserve Bank of St. Louis. Kluwer, Boston, pp. 3–37.

Viswanathan, S., and J. Wang, 2002, Market architecture: limit-order books versus dealership markets, Journal of Financial Markets, 5, 127–167.

Vitale, P., 2004, A guided tour of the market microstructure approach to exchange rate determination, Working Paper.

Wang, J., 2001, Quote revision and information flow among foreign exchange dealers, Journal of International Financial Markets and Money, 11, 115–136.

Wang, J., and E. Zivot, 1998, Inference on structural parameters in instrumental variables regression with weak instruments, Econometrica, 66, 1389–1404.

Watson, M., 1994, Vector autoregressions and cointegration, in Handbook of Econometrics.

Yao, J., 1998, Market Making in the Interbank Foreign Exchange Market, New York University Salomon Center Working Paper #S-98-3.

Yohn, T., 1998, Information asymmetry around earnings announcements, Review of Quantitative Finance and Accounting, 11, 165–182.