Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

Essays on Exchange Rates

A Dissertation Submitted in Fulfilment of the Requirements for the Degree of Doctor of Philosophy in Finance at Massey University

Vincent Kleinbrod

School of Economics and Finance
Massey University
September 2016

Abstract

This dissertation presents three essays on exchange rates. The reported work builds on the market microstructure approach to exchange rate determination and extends this approach to modelling and forecasting multivariate exchange rate movements, and to a multi-currency trading application.

The first study investigates the role of order flow in explaining joint movements of exchange rate returns, thereby building an original bridge between exchange rate co-movement and the market microstructure literature. We document that absolute order flow differentials have a significant negative effect on future joint currency movements at intraday frequencies. The analysis also shows that other intraday variables, such as the bid–ask spread, have no explanatory power for the co-movements after the absolute order flow differential is accounted for, thereby confirming the robustness of order flow as the driving force for exchange rate correlation. Further analysis demonstrates that absolute order flows also affect conditional variance dynamics.

The second study adds to the findings of the first study. It evaluates the information content of order flow for accurate predictions of exchange rate co-movement. In line with the first study, we find that order flow information substantially enhances the accuracy of covariance forecasts. Moreover, the interest rate differential has a limited role in explaining and predicting correlation dynamics once the order flow differential is accounted for. The study concludes by showing the economic value of the order-flow-based covariance predictions, namely the value of order flow information for covariance predictions beyond return predictions.

The third study focuses on the practical relevance of order flow information in foreign exchange trading. Given the dominance of technical trading among forex professionals, the study evaluates the value of order flow information for technical traders. Our initial investigation questions the accuracy of trading signals if these are derived directly from order flow. We conjecture that the reason for this is that order flow should first be used to generate exchange rate predictions, which can then be used to derive profitable trading signals. We examine this conjecture empirically, and the affirmative results highlight the value of order-flow-based return predictions for technical analysis. Further, we propose a multivariate trading strategy to boost the benefits of using order flow in technical analysis, which is shown to be a highly profitable.

Acknowledgements

I would like to express my sincere gratitude to my supervisors Professor Xiao-Ming Li and Professor David Ding, for their unreserved encouragement and support during my PhD study. I am especially indebted to my chief supervisor, Professor Xiao-Ming Li, for his constructive guidance and help. Thank you for having confidence in and patience with me. Your enthusiasm and dedication to research will continue to inspire me to become a better researcher.

Special thanks go out to the staff in the School of Economics and Finance who has been very helpful and supportive. The general staff have always made me and other PhD students feel at home. I would like to further acknowledge the support of my fellow PhD students, who have made this journey more colourful and memorable.

Thank you to Massey University for providing financial support for conferences. This thesis benefited from valuable comments from participants of the 2015 Asian FMA Doctoral Consortium, and the 2013–2016 New Zealand finance colloquia. In addition to the conferences, this thesis benefited from valuable comments given by faculty members from the School of Economics and Finance.

Last but not least, I would like to dedicate this thesis to my partner and my parents, in grateful thanks for their patience, support and continued enthusiasm during the project. Without them, this dissertation would never been feasible.

Table of Contents

List of Tables

List of Figures

Chapter One: Motivation and Overview

1.1 Introduction

The market microstructure approach to exchange rate determination has become a standard tool for analysing univariate exchange rate movements that overcomes the inability of traditional macroeconomic models to explain and predict reasonably accurate short-run exchange rate fluctuations (see, for example, Meese and Rogoff, 1983 or Rogoff, 2002). This approach, pioneered by Evans and Lyons (2002a), is based on different information sets across all market participants that result from different expectations about the future fundamentals of an exchange rate. The central element aggregating and transmitting all the dispersed information about future fundamentals into prices is order flow, defined as the number of buyer-initiated minus seller-initiated trades (Evans and Lyons, 2002a), arising from the asymmetries in information and expectations about future exchange rate fundamentals.

Although the literature links order flow to exchange rate movements, the relationships between order flow and joint appreciation or depreciation in exchange rates (co-movement) have not yet been evaluated. This relationship arises because co-movements in prices result due to co-movements in fundamentals (Barberis, Shleifer, and Wurgler, 2005). As order flow reflects the expected changes in fundamentals, we conjecture that order flow can proxy for the co-movements of the fundamentals that affect the co-movements of currency prices. Chapter 2 theoretically demonstrates this possibility in detail and undertakes an intraday examination of this issue.

Although the emphasis of the market microstructure approach to exchange rate determination has primarily been on explaining exchange rate movements via order flow, several authors highlight the value of order flow information for exchange rate predictions (e.g. Evans and Lyons, 2005b). The predictive power of current order flows for future

exchange rate movements results from the gradual incorporation of order flow information into prices (Berger, Chaboud, Chernenko, Howorka and Wright, 2008) thereby affecting exchange rate movements both instantly and with a substantial lag. Yet the predictive content of current order flow for future exchange rate co-movements has been neglected in the literature. Building upon the link between order flow and co-movement outlined in Chapter 2, Chapter 3 evaluates whether order flow information helps to enhance co-movement forecasting. Chapter 3 concludes by addressing the economic value of more accurate covariance predictions for portfolio optimisation.

An argument for the economic value of order-flow-based return predictions is that order flow is best suited to predicting short-term exchange rate fluctuations, as the impact of order flow on exchange rate fluctuations is stronger for more frequent (intraday and daily) movements and gradually decreases for lower frequencies (Berger, Chaboud, Chernenko, Howorka and Wright, 2008). However, the economic literature has ignored this argument and has evaluated the economic value of order flow information particularly in the context of portfolio optimisation (e.g. King, Sarno and Sojli, 2010) which is typically a long-term investment strategy. Upholding this argument, Chapter 4 turns away from order-flow-based portfolio optimisation and asks whether order flow information can be used to generate profitable technical trading signals. Further justification for doing so arises because the specific price patterns generated by order flow is what technical trading indicators try to exploit (Neely and Weller, 2011). Therefore, Chapter 4 explores this question by first demonstrating several problems associated with deriving trading signals directly from order flow. The chapter then goes further to show that using order flow to generate return predictions and then linking these predictions to changes in the underlying market volatility can generate a highly profitable trading strategy.

1.2 Main Findings and Contribution to the Literature

This section outlines the main findings and contributions of Chapters 2, 3 and 4 that comprise the core part of this thesis.

Chapter 2 links order flow to joint movements (co-movements) of exchange rates. The multivariate framework enables us to study how well the relative changes in the order flows of different national currencies explain the joint movements of exchange rates, which adds to both the market microstructure literature and the literature on modelling asset co-movements. With respect to the latter, we focus on the role of exogenous variables in driving conditional correlation dynamics, which links to the works of Benediktsdottir and Scotti (2009), Schoppen (2012) and Li (2011).

We document a significant negative relationship between the absolute order flow differential and exchange rate co-movement. This result relates to the market microstructure literature suggesting that order flow aggregates all dispersed information about the expected future fundamentals of an economy. Again in line with the market microstructure literature, our intraday analysis suggests that the relationship is strongest for the highest intraday frequencies evaluated and gradually decreases as frequencies increase (Berger et.al, 2008).

All conditional correlation studies suggest that correlation is time–dependent. This time dependency raises the question of whether the driving forces of correlation dynamics remain unchanged across financially turbulent and calm periods. A further, specific question is whether order flows and their impacts on price movements are the same across such different periods. To the best of our knowledge, no study has asked this research question with respect to asset price co-movements. Dijk, Munandar and Hafner (2011) provide some preliminary evidence of a structural break at the advent of the Euro and argue that neglecting these structural breaks would lead to over- or underestimation of conditional correlations for

the different time periods studied. Our findings not only confirm the presence of structural changes in conditional correlation dynamics caused by the 2007 global financial crisis (GFC) and the 2009 European sovereign debt crisis (EDC), but, more importantly, reveal that the impact of the absolute order-flow differential on exchange rate co-movements is not the same across financially turbulent and tranquil time periods.

By studying which type of asymmetry –positive or negative– dominates intraday correlation dynamics, Chapter 2 makes another contribution to the literature. The literature documents that either type of asymmetry can dominate. In stock or bond markets, negative-type asymmetry typically dominates, which is in line with leverage considerations (Cappiello, Engle, and Sheppard, 2006; De Goeij and Marquering, 2004). However, in foreign exchange markets positive-type asymmetry can dominate (Li, 2011). While all the aforementioned studies have focused on daily or lower frequencies, we provide an argument based on investor sentiment causing positive-type asymmetry to prevail in FX return correlations. This argument is confirmed by data, suggesting that joint positive shocks have a stronger impact on co-movements than joint negative shocks of the same magnitude.

The main contribution of Chapter 2 is to establish a link between order flow and exchange rate return co-movements. Our rigorous intraday examination of this link follows directly from theoretical considerations.

The study in Chapter 3 evaluates the accuracy of order-flow-based covariance predictions compared with several alternative benchmarks commonly used in the literature. It adds to the literature on correlation (covariance) forecasting, which mainly focuses on different estimation techniques (see Bauwens, Laurent and Rombouts, 2006, for a detailed review) , by examining how exogenous variables enhance co-variance predictions. This kind

of analysis has previously been limited to how the interest rate differential (IRD) affects future exchange rate co-movements (Benediktsdottir and Scotti, 2009; Li, 2011).

Our results strongly suggest that the IRD alone is not a suitable predictor of future correlation dynamics. Instead, we show that order-flow-based covariance predictions provide the highest statistical accuracy out of several competing benchmark models, including those in the aforementioned studies that use the interest rate differential to model exchange rate co-movements.

The performance gains from order-flow-based forecasting are particularly large for financially unstable time periods, confirming the incremental information content of order flow during a financial downturn (Danielsson and Saltoglu, 2003; Rime and Tranvåg, 2012). Furthermore, we find that order flow information enhances the accuracy of volatility predictions during economically instable time periods. The reason for this result is that volatility spikes tend to occur in tandem with order flow spikes.

This chapter further contributes to the literature by examining the accuracy of covariance forecasts using different estimation window sizes. This is warranted, as the literature on order-flow-based forecasting directly uses a medium-sized rolling window, without making any comparisons (e.g. Rime, Sarno, and Sojli (2010)). From a theoretical point of view, long and short estimation windows both have advantages and disadvantages, whereas a medium-sized window can minimise the disadvantages of both. We empirically verify this theoretical conjecture and demonstrate that a medium-sized window is indeed preferable to long and short windows.

While it assesses the statistical accuracy of our proposed methodology, Chapter 3 is also related to the studies of order-flow-based portfolio optimisation (Rime, Sarno, and Sojli, 2010; Della Corte, Sarno, and Tsiakas, 2011). However, these studies typically do not predict

future joint movements (co-variance): they assume the covariance between the assets to be constant. By contrast, we find that the order-flow-based covariance predictions lead to substantially higher Sharpe ratios than any non-order-flow-based covariance prediction or the constant covariance model used by the studies mentioned above.

The main aim of the study in Chapter 3 is the out-of-sample evaluation of order-flow-based co-variance forecasting. It offers a deeper examination of the link as discussed in Chapter 2, which confirms the importance of taking order flow into account in modelling and forecasting correlation dynamics.

Chapter 4 critically investigates the financial value of order flow information. We evaluate a trading strategy which utilises order flow information to generate profitable technical trading signals. So far, the use of order flow as a buy or sell indicator has been explored only for stock markets (Yamamoto, 2012). Chapter 4 starts by replicating Yamamoto's (2012) analysis but with a focus on the FX market, where order flow should have a stronger impact on price fluctuations because of the strong zero inventory preference of dealers (Bjonnes and Rime, 2005). We report similar results to those of the Yamamoto, suggesting that order flow cannot be used as a proxy for making profitable buy or sell decisions. This finding, however, does not imply that order flow has no value for technical trading; instead it raises an issue related to deriving trading signals from order flow. As shown in the literature, order flow is commonly used to generate exchange rate predictions, which in turn act as trading signals (Gradojevic, 2007; Chordia and Subrahmanyam, 2004). This indicates a potential non-linear relationship between order flow and exchange rate movements.

Addressing this issue, we use a non-linear setting to generate exchange rate return forecasts. The non-linear setting uses order flow to forecast future exchange rate movements-

and then generates a trading recommendation based on these forecasts. The initial setting is the same as that outlined in Gradojevic (2007). However, our analysis involves a set of exchange rates, not just the Canadian dollar (CAD) – US dollar (USD) rate only, thereby confirming the robustness of the results in Gradojevic (2007) and the profitability of this approach.

As a novel addition to the literature, we link the generated return predictions to the underlying market volatility. Our proposed trading strategy combines exchange rate return predictions and volatility information using a fuzzy logic inference setting to generate the trading recommendation. This setting relates to the work of Bekiros (2011), which studied the relationship between volatility and trading performance, and Christoffersen and Diebold (2006), who argued that volatility changes would alter the probability of a negative return sign. We show that our proposed trading approach is remarkably profitable and significantly outperforms the approach adopted by Gradojevic (2007), which only takes the order-flow-based exchange rate predictions into account, for several currency pairs.

Chapter 4's main objective is to critically evaluate whether order flow can indeed provide economic value for investors. Apart from the documented trading profits, our results also reveal that there is a link between the volatility and technical trading rule profitability. Thus- the results presented in Chapter 4 are of both academic and practical relevance. Academically, they suggest that technical trading rule profitability may depend on the underlying assets volatility. Practically, they provide currency traders with a useful guide for forming trading strategies.

1.3 Structure of the Dissertation

The core part of this dissertation embraces three essays, each building upon the market microstructure approach to exchange rate determination. In order to organise the dissertation in a methodical manner, the three essays will appear as three independent chapters. Specifically, the structure of this dissertation is briefly described as follows.

Chapter 2 theoretically and empirically examines the link between order flow and exchange rate co-movements. The analysis focuses on intraday dynamics, although the link applies to any data frequency. Chapter 3 builds upon this link and assesses its relevance for accurate forecasting of exchange rate co-movements. Critically reviewing the common practice of evaluating the economic value of order flow being in portfolio optimisation, Chapter 4 turns to exploring the economic value of order flow in deriving profitable trading signals. Chapter 5 concludes and discusses the intended future research direction.

2. Chapter Two: Order Flow and Exchange Rate Co-movement

2.1 Introduction

How assets co-move is crucially important for various financial and economic applications, such as portfolio diversification- and rebalancing decisions (Beine, 2004; Wu, Chung and Chang, 2012; Salmon and Schleicher, 2006) and monetary policy decisions (Benediktsdottir and Scotti, 2009). Thus, extensive research work has been devoted to explaining and modelling the dynamics of co-movements between different asset classes. Regarding foreign exchange markets, the literature has placed emphasis on examining changes in the macroeconomic fundamentals (Li, 2011) and central bank interventions (Beine, 2004; Beine, Grauwe and Grimaldi, 2009) as being the determinants of joint currency movements. By contrast, this chapter turns to order flow, a market -microstructure variable, as a possible driving force behind joint appreciation or depreciation in the exchange rates, using the intraday frequency for the first time in the literature.

The primary motivation for our study stems from the strong predictive power of order flow for single exchange rate movements (Evans and Lyons, 2002a). Order flow is defined as signed trading volume, or the net of buyer- and seller-initiated transactions, showing the direction of trade. What makes currency order flow so powerful in exchange rate forecasting is the fact that underneath order flow lie disperse public expectations about current and future economic fundamentals that are relevant to exchange rates (Rime et al., 2010). This notion suggests that changes in expectations will cause changes in order flow, which, in turn, will affect the price of the currency. Thus, order flow, though a market-microstructure variable, carries information about macroeconomic variables, but has data available at higher frequencies than the latter.

However, one question left unanswered is whether and how well the relative changes in the order flows of different national currencies predict the joint movements of exchange rates. If asset co-movements are important for financial and economic applications, then it is of interest to go from a univariate to a multivariate framework for studying the link between currency order flows and currency prices. In doing so, we consider the absolute differentials of order flows and evaluate their roles in determining the dynamic conditional correlations of exchange rates.

We also examine whether the relationship between order flow differentials and currency co-movements has undergone any structural change as a result of large events such as the 2007 GFC or the 2010 EDC. The motivation comes again from the univariate market-microstructure literature, namely the finding that order flows are more informative during an economic crisis (see, Rime and Tranvåg, 2012).

Yet another issue worth investigating is the asymmetric responses of exchange rate co-movements to information shocks. On the empirical level, Patton (2006) provides evidence in support of negative-type asymmetry, whereas Li (2011) shows evidence in favour of positive-type asymmetry (-i.e., stronger exchange rate correlations during joint appreciations than during joint depreciations). Whereas previous studies have explored the dominating type of asymmetry at daily or lower frequency, the present study is the first to examine the issue using intraday data. Note that the information shocks associated with the asymmetry issue may or may not overlap with those contained in order flows. In any event, allowing for asymmetry enables us to isolate the potential effects of investor sentiment from the effects of order flows on exchange rate co-movements.

In investigating the aforementioned issues, we focus on the co-movements of, respectively, eight exchange rate pairs: Euro (EUR)–British pound (GBP), EUR–Japanese

yen (JPY), EUR–CAD, GBP–JPY, GBP–CAD, JPY–Australian dollar (AUD), New Zealand dollar (NZD)–AUD and JPY–CAD. All the exchange rates measure the USD prices of the respective currencies (i.e. EUR denotes USD per EUR, GBP denotes USD per GBP, etc.). Our investigations yield several interesting results, outlined below:

First, there is overwhelming evidence that larger differences between order flows considerably reduce co-movements. That is, the higher the absolute order flow differential, the lower the probability of joint appreciation or depreciation. With intraday frequencies, we find that this negative impact is stronger at a higher intraday frequency and gradually decreases as the intraday frequency increases. The result is in line with the study of Berger et al., (2008) on single exchange rate movements. Furthermore, the impact is stronger (weaker) for exchange rate pairs with a higher (lower) level of unconditional correlation. However, for the AUD-NZD correlation, there is some evidence that a larger absolute order flow differential increases the co-movement.

Second, we find significant evidence that the 2007 GFC has caused structural changes in the relationship between order flow differentials and conditional correlations for all the eight exchange rate pairs. During the 2007 GFC, the relationship changed for the EUR–GBP and EUR–CAD correlation, as a larger absolute order flow differential suddenly started to increase the co-movement of these two pairs. A potential explanation for this finding is that the differences in trading intensities dramatically changed during the 2007 GFC, thereby allowing the order flow of one exchange rate to dominate the correlation dynamics.

Third, all exchange rate pairs, except AUD–NZD, exhibit strong positive-type asymmetry: joint positive news subsequently leads to a higher correlation than joint negative news. This result runs counter to stock market evidence that negative-type asymmetry prevails. These results basically confirm Li's (2011) study, which, nevertheless, evaluates

interest rate differentials, rather than order flow differentials, regarding their impacts on exchange rate co-movements. Since order flows contain information not only on interest rates but also on other economic fundamentals, the result should have farther-reaching implications, more credibility and greater practical relevance.

2.2 Related Literature

The market microstructure approach to exchange rate determination has become a standard tool for explaining and predicting exchange rate fluctuations, surpassing traditional macroeconomic models especially with respect to short-term frequencies (Evans and Lyons (2005b)). This approach has proven to be quite successful in explaining (Evans and Lyons,2002a, 2002b, 2002c; Payne, 2003; Killeen, Lyons, and Moore, 2006; Carlson and Lo, 2006) and predicting (Evans and Lyons, 2005b; Rime et al., 2010; Cerrato, Kim and McDonald, 2015)[1] single exchange rate movements.

The strong explanatory power of order flow (-defined as the net of buyer- and seller initiated trades -) for single exchange rate movements was first reported by Evans and Lyons (2002a). Their path-breaking study demonstrated that order flow can explain up to 60% of daily exchange rate variations in the Deutsche Mark (DEM) –USD spot rate, which opened a door for the development of modern microstructure theory. Central to the theory is order flow, which is believed to contain the dispersed information possessed by all agents or to represent information heterogeneity. Thus the emphasis is placed on the aggregate of traders willing to pay transaction costs in order to settle their trades. The difference between buying and selling orders arises from information heterogeneity, is captured by order flow, constitutes general buying or selling pressure on a currency's prices, and ultimately leads to its appreciation or depreciation.

[1] See King, Osler, and Rime (2013) for a review.

The information heterogeneity amongst traders is central in the market microstructure theory. It implies that order flow is the vehicle via which dispersed information about the expected future fundamentals is aggregated and impacts on current and future prices. This "proxy" view allows to explain why the relationship between order flow and exchange rates seems to be dependent upon market conditions (Luo, 2001) and why some order flow is more informative (Osler, 2005; Cerrato, Sarantis, and Saunders, 2011) than the other. Furthermore, this "proxy" view allows one to explain how macroeconomic news is incorporated into exchange rates (Evans and Lyons, 2005a; Love and Payne, 2008; Dunne, Hau, and Moore, 2010) via order flow and to link order flow to movements of joint return dynamics.[2] Note that, even when order flow is not related to changes in the underlying fundamentals, order flow can still have a persistent impact on exchange rate movements, as suggested by Bacchetta and Wincoops's (2006) "scapegoat" theory.

Given the strong explanatory power of order flow for univariate movements, it is surprising that no attempts so far have linked market microstructure theory with asset co-movements. Studies on how assets co-move are manifold in the literature, such as Cappiello et al. (2006) on international bond and equity returns, Albuquerque and Vega (2009) on interational stock markets, Silvennoinen and Thorp (2013) on commodities and Kuper and Lestano (2007) on interest rates and stocks, among others. Recent studies highlighting the economic value of modelling asset co-movement include the studies of Christoffersen, Errunza, Jacobs, and Jin (2014), Kalotychou, Staikouras, and Zhao (2014) and Billio and Caporin (2009) .

In addition to individual return shocks driving co-movement, the notion that joint positive or negative news shocks have an incremental effect on correlation dynamics has been investigated in detail in the literature (Patton, 2006; Hyde, Bredin, and Nguyen, 2007;

[2] This link will be discussed in more detail the next section.

Savva, 2009; Savva, Osborn, and Gill, 2009; Li, 2011), suggesting that both positive- and negative- type asymmetry can drive correlation dynamics, a finding which provides several implications for international diversification (Hyde et al. 2007). Although negative asymmetry (joint negative shocks having an additional impact on correlation dynamics) is widely found in stock and bond markets (e.g. Cappiello et al., 2006), the contrary is found in foreign exchange markets (Li, 2011).

Further to return shocks and return asymmetries, the literature has evaluated the role of exogenous variables as an additional driving force for correlation dynamics. For exchange rate co-movements, the literature has evaluated the role of macroeconomic variables (Li, 2011; Benediktsdottir and Scotti, 2009) or monetary policy interventions in affecting the volatility of exchange rates (Beine, 2004; Beine et al., 2009). Li (2011) evaluated the role of the IRD in driving correlation dynamics, finding that both widening and narrowing IRDs considerably lower the time-varying exchange rate correlation between five inflation-targeting countries. Schoppen (2012) evaluated several exogenous variables and concluded that gross domestic product growth and market turbulences drive conditional correlations for international bond co-movements.

As correlation is usually time-varying, several papers have attempted to test for structural breaks in correlation dynamics. This is to allow for different dynamics in times of economic distress or to allow for breaks in correlation dynamics caused by significant episodes such as the advent of the Euro in 1999 (Dijk et al., 2011; Kearney and Potì, 2006). Neglecting these structural breaks would lead to under- or overestimation of future co-movement, thereby reducing the benefits of covariance forecasts when they are most needed in times of distress. Again, there is a link with order flows, as Rime and Tranvåg (2012) suggested that order flows are particularly informative (i.e., they have the highest explanatory

content) during an economic downturn. This finding calls for an in-depth analysis of the relationship between order flow and co-movement across financially stable and turbulent periods, which we will investigate in this chapter.

In a previous study of the market microstructure effect on exchange rate correlations, Vargas (2008) looked at the link between foreign exchange and equity returns, using capital flows and the IRD as exogenous variables in a dynamic conditional correlation model. Our approach is in the same spirit but has several noteworthy differences. First, we consider the return correlations of exchange rate pairs, whereas Vargas (2008) examined how the returns of a stock index co-move with depreciation and appreciation of the country's exchange rate. Second, we employ high-frequency data and order flow only, whereas Vargas (2008) used microstructure variables as well as macroeconomic variables (the so-called hybrid model) with a monthly frequency. Finally, we use the order flow differential instead of the difference in equity purchases.

2.3 Theories and Hypotheses

Our primary research question is: "Does order flow drive the intraday co-movements between exchange rate returns?" Addressing this question should fill two voids in the literature: the nexus between order flow and currency correlations, and the intraday frequency for the nexus. To motivate the question, let us discuss the theoretical model of Rime et al. (2010):

$$\Delta s_{t+1} = \frac{1-b}{b}\left[E_t\left(f_t\right) - s_t\right] + \varepsilon_{t+1},\ \ \varepsilon_{t+1} \equiv (1-b)\sum_{q=0}^{\infty} b^q \left[E_{t+1}\left(f_{t+1+q}\right) - E_t\left(f_{t+1+q}\right)\right], \quad (2.1)$$

where $s_t \equiv$ log of nominal exchange rate of a currency against the USD (defined as the USD price of a currency) at time *t*, $\Delta s_{t+1} \equiv s_{t+1} - s_t$, b $(0 < b < 1) \equiv$ the discount factor, $f_t \equiv$ economic fundamentals at time *t*, $E_t(f_t) \equiv$ market expectations about the current fundamentals conditional on the information available at *t*, and $E_t\left(f_{t+q}\right)$ or $E_{t+1}\left(f_{t+1+q}\right) \equiv$ market expectations about future (q periods ahead or $q+1$ periods ahead) fundamentals conditional on the information available at t or $t+1$. This equation says that order flow[3] reflects both $[E_t(f_t) - s_t]$ and ε_{t+1} which involve changes in expectations about fundamentals, and thus can predict exchange rate movements, described by Δs_{t+1}. Note that Equation (2.1) applies to different countries by changing the discount factor *b* accordingly.

Rime et al. (2010) provided empirical evidence that macroeconomic news is an important determinant of order flow. This evidence is consistent with the theoretical prediction of Equation (2.1). To fix the idea and without loss of generality, assume that $E_t(f_t) - s_t = 0$. Now suppose that at $t+1$, news that indicates improvements in Country A's fundamentals is released. This positive news will make $E_{t+1}(f_{t+1+q}) > E_t(f_{t+1+q})$, which, given

[3] Order flow is defined as the net of buyer- and seller-initiated trades for each exchange rate pair.

everything else, will make $\Delta s_{t+1}^{A} > 0$ (an appreciation in Currency A in the USD price) via an increase in the order flow of (or the net buying pressure on) the currency. The opposite will occur if negative news about A's fundamentals is released.

In this case, how do changes in the difference between the order flows of different currencies affect their exchange rate correlations? Let $x^{A}(x^{B})$ be the order flow of Currency A (B), where $x^{A} > 0$ and $x^{B} > 0$ ($x^{A} < 0$ and $x^{B} < 0$) represent the net buying (selling) pressure on Currencies A and B respectively. Let ρ_t be the conditional correlation between Δs_t^{A} and Δs_t^{B} at *t*, resulting from $x_{t-1}^{A} - x_{t-1}^{B}$ (the difference between the two order flows at *t* – 1). In what follows below, we begin with the initial condition that $x_{t-1}^{A} - x_{t-1}^{B} = 0$.

Consider the case where $x_t^{A} > x_{t-1}^{A}$, $x_t^{B} > x_{t-1}^{B}$ and either (i) $x_t^{A} - x_t^{B} = x_{t-1}^{A} - x_{t-1}^{B} = 0$ (i.e., the order flows of Currencies A and B rise by the same amount) or (ii) $x_t^{A} - x_t^{B} > x_{t-1}^{A} - x_{t-1}^{B} = 0$ (i.e., the order flow of Currency A rises more than that of Currency B). Assume the normal circumstance that b^{A} and b^{B} in Equation (2.1) are close enough to each other in value. Under this circumstance, (i) suggests that s_t^{A} and s_t^{B} would be equally likely to appreciate (i.e., $\Delta s_t^{A} > 0$ and $\Delta s_t^{B} > 0$ are equally likely), while (ii) suggests that s_t^{A} is more likely to appreciate than s_t^{B} (i.e., $\Delta s_t^{A} > 0$ is more likely than $\Delta s_t^{B} > 0$). As a result, we expect $\rho_t^{(\mathrm{i})} > \rho_t^{(\mathrm{ii})}$.

Next consider the case where $x_t^{A} < x_{t-1}^{A}$, $x_t^{B} < x_{t-1}^{B}$ and (i) $x_t^{A} - x_t^{B} = x_{t-1}^{A} - x_{t-1}^{B} = 0$ (i.e., the order flows of currencies A and B fall by the same amount) or (ii) $x_t^{A} - x_t^{B} < x_{t-1}^{A} - x_{t-1}^{B} = 0$ (i.e., the order flow of currency A falls more than that of currency B). Under the normal circumstance defined above, (i) suggests that s_t^{A} and s_t^{B} would be equally likely

to depreciate (i.e., $\Delta s_t^A < 0$ and $\Delta s_t^B < 0$ are equally likely), while (ii) suggests that s_t^A is more likely to depreciate than s_t^B (i.e., $\Delta s_t^A < 0$ is more likely than $\Delta s_t^B < 0$). As a result, we expect $\rho_t^{(i)} > \rho_t^{(ii)}$.

Without exhausting all possible cases, we can make the following statement: under the normal circumstance that the effect of order flow on the exchange rate is approximately similar across two currencies, the absolute difference between order flows will have a negative effect on the exchange rate correlation of the two currencies.

However, the discount factor *b* can be significantly different between different countries, as it varies over time. Evans and Lyons (2008) proposed an empirical model for testing the theoretical model like Equation (2.1), as follows:

$$\Delta s_t = D(L, S_t)x_t + e_t \ ; \tag{2.2}$$

$$x_t = C(L)\xi_{t-m}, \tag{2.3}$$

where $D(L, S_t)$ is the state-dependent polynomial in the lag operator L; S_t is the market state, which is assumed to depend on trading intensity at t and $C(L)$ is the lag polynomial describing the dynamic responses of order flow to the dispersed information shocks ξ_{t-m} that arrived m periods prior to t. Equation (2.2) says that the effect of order flow on an exchange rate can change from time to time via changes in $D(L, S_t)$ and thus can differ from currency to currency. This is because the trading intensity of any currency changes over time. Comparing equation (2.2) with equation (2.1), it appears that the discount factor b in (2.1) is related to $D(L, S_t)$. As S_t hence $D(L, S_t)$.changes over time, so does b (hence b_t).

Based on the findings of Evans and Lyons (2008), let us re-examine the reasoning about the sign of the order flow's effect on exchange rate correlations. In the discussion

below, we revert to the theoretical model in Equation (2.1), with reference to the empirical model in Equation (2.2). This is to acknowledge the obvious fact that both x_t and b_t jointly determine the likelihood with which an exchange rate appreciates or depreciates.

Consider again the case where $x_t^A > x_{t-1}^A$, $x_t^B > x_{t-1}^B$ and (i) $x_t^A - x_t^B = x_{t-1}^A - x_{t-1}^B = 0$ or (ii) $x_t^A - x_t^B > x_{t-1}^A - x_{t-1}^B = 0$. Assume, however, an abnormal situation where the difference in trading intensity between A and B is so vast as to result in the difference between $D(L, S_t^A)$ and $D(L, S_t^B)$ that is so large to make $b_t^B << b_t^A$ happen. In (i), although x_t^B and x_t^A increase by the same amount, that b_t^B is much smaller than b_t^A implies that s_t^A would be much less likely to appreciate than s_t^B. As a result, $\rho_t^{(i)}$ should be low. Turning to (ii), although the increase in x_t^B is smaller than in x_t^A, this could be offset by the larger b_t^A than b_t^B. As a consequence, s_t^A and s_t^B may become close to being equally likely to appreciate, leading to a high $\rho_t^{(ii)}$.

Next, consider the case where $x_t^A < x_{t-1}^A$, $x_t^B < x_{t-1}^B$ and (i) $x_t^A - x_t^B = x_{t-1}^A - x_{t-1}^B = 0$ or (ii) $x_t^A - x_t^B < x_{t-1}^A - x_{t-1}^B = 0$. Conducting reasoning analogous to the above, we can state the following. In (i), s_t^A would be much less likely to depreciate than s_t^B, so $\rho_t^{(i)}$ should be low. In (ii), s_t^A and s_t^B may become close to being equally likely to depreciate, leading to a high $\rho_t^{(ii)}$.

Again, without exhausting all possible cases, the following statement holds: under the *abnormal* circumstance that the effect of order flow on an exchange rate is vastly different between two currencies, the absolute order flow differential will have a *positive* effect on the their correlations. It is worth reiterating that if the trading intensities of two currencies change significantly enough to make the effects of the order flows on their exchange rates vastly

different from each other, then the sign of the effects of order flows on their correlations will change. In conclusion from this discussion, we propose to test the following two hypotheses:

H1: Under normal circumstances, a greater absolute order flow differential between two currencies will subsequently reduce their exchange rate return co-movements.

H2: Under abnormal circumstances, a greater absolute order flow differential between two currencies will subsequently increase their exchange rate return co-movements.

Some additional remarks are in order here regarding empirical tests of the two hypotheses. Although "circumstance" is defined with respect to trading intensity, we do not attempt to explicitly investigate its role in influencing correlation dynamics for the following reasons. First, Evans and Lyons (2008) have produced convincing and relevant results, and our study will simply take these as a starting point. This will make the correlation models tractable and more focused. Second, we evaluate the effect of order flow on exchange rate correlations based on a constant model parameter. Our theoretical analyses imply that the sign of the parameter should not change as frequently as the trading intensity differential. If a structural break in the parameter is significant enough to alter its sign, we take this to mean that the trading intensity differential between two currencies exceeds a certain critical value.

We conjecture that a structural change occurs following a big event such as a global financial crisis. Accordingly, if we combine H1 and H2, this leads to the third testable hypothesis:

H3: The effect of the absolute order flow differential on exchange rate co-movements is not the same across financially tranquil and turbulent periods.

A financially tranquil (turbulent) period corresponds to normal (abnormal) circumstances as defined earlier. In testing H3, we choose two major economic events within

the sample period: the 2007 global financial crisis (GFC) and the 2010 European sovereign debt crisis (EDC). Note that a structural break, if detected, does not necessarily alter the sign of the order flow effect: the relevant parameter may change its value but not its sign, or it may turn from being negative to being positive or vice versa. We offer some data-based evidence to show the possible connections between trading intensity differentials and the strength of structural change for the exchange rate pairs under investigation.

Order flows are unlikely to fully capture macroeconomic news as described by the difference between the information set in $E_{t+1}(f_{t+1+q})$ and that in $E_t(f_{t+1+q})$ in Equation (2.1). However, the information shocks omitted by order flows will also move exchange rates, which is suggested by the error term e_t in Equation (2.2). Therefore, apart from order flows, we also examine how the omitted news shocks impact asymmetrically on exchange rate correlations or, more specifically, the issue of positive and the negative asymmetry. The former (latter) is defined as stronger exchange rate correlations during joint appreciation (depreciation) than during joint depreciation (appreciation). Li (2011) maintains that inflation-targeting currencies should demonstrate positive asymmetry in the correlations of their exchange rates against a world currency (e.g., USD) and provides empirical evidence at daily frequency to support his conjecture. However, the inflation-targeting argument may not be applicable to intraday data and thus some other factors need to be considered.

In this chapter, we use investor sentiment as a possible explanation for positive type asymmetry. Negative news shocks regarding the US economy may create fears of a further decline in the USD (Hibbert et al., 2008).[4] These fears would then cause FX traders to constantly sell USD for other major currencies such as EUR, GBP and JPY, for example, leading to a prolonged "upward spiral" in the demand for and hence the USD price of these

[4] Although Hibbert et al. (2008) focused on the S&P 500 and changes in the CBOE VIX, the theory is equally likely to hold in FX markets because of the large connection between both markets.

currencies. The positive exchange rate returns to EUR, GBP or JPY then would raise their co-movements during joint appreciations against USD. Conversely, when positive news shocks hit the US economy, no fears would be present. Therefore, the increased demand for USD would not be as persistent and strong as the increased supply of USD in the case of negative news shocks to the US economy. Accordingly, the negative exchange rate returns to EUR, GBP or JPY may also raise their co-movements during joint depreciation against USD, but to a lesser degree. We therefore conjecture the effects of positive asymmetry, described in the fourth hypothesis below:

H4: Joint positive shocks affect intraday correlation dynamics more strongly than joint negative shocks of a similar magnitude.

The fifth hypothesis arises from our intraday focus. The gradual information theory suggests that order flow information is gradually incorporated into prices (Marsh and Teng, 2012; Berger et al., 2008). This means that most of the information is incorporated directly after the orders are submitted, whereas the remaining small amount of information is incorporated in the long run, because dealers usually adjust their positions according to trades by other dealers and investors. Therefore, we hypothesise the following:

H5: The negative relationship between the absolute order flow differential and exchange rate correlations is strongest at the highest intraday frequency and weakens for lower intraday frequencies.

Some evidence exists in favour of the frequency-dependent impact of order flows on single exchange rate movements. For example, Berger et al. (2008) found that although lagged order flows affect future exchange rates, the impact gradually lessens as the frequency lowers. Our fifth hypothesis extends this work to exchange rate co-movements.

2.4 Data & Methodology

2.4.1 Data

Our empirical tests use the exchange rates of AUD, NZD, CAD, EUR, GBP and JPY against USD as a common denominator.[5] Under this definition, changes in (logarithmic) exchange rates can be interpreted as the USD return for holding one unit of foreign currency. The order flow of a currency measures the difference between buyer-initiated trades and seller-initiated trades of the currency. Absolute order flow differentials calculate the absolute values of the differences between two currencies' order follows.

We obtain data on the raw currency transactions and quotes from Thomson Reuters Tick History (TRTH) provided by the Securities Industry Research Centre of Asia-Pacific (SIRCA). TRTH covers all FX trades conducted in the D2000-2 electronic FX brooking system, a brokered inter-dealer trading platform run by Thomson Reuters. D2000-2 is one of the two main electronic brokers in this market, the other being Electronic Brooking Services (EBS). EBS is the dominant trading platform used for the USD–EUR and USD–JPY rates, but most of the trading under the USD–AUD, USD–NZD and USD–GBP rates is done via Thompson Reuters D2000-2 (Smyth, 2009). However, although we only have access to the D2000-2 data on the USD–EUR and USD–JPY rates as well as other rates, previous studies (e.g., Danielsson and Payne, 2012) show that both databases have the same patterns and prices as well as order flows moving in the same direction.[6]

[5] We exclude the three Nordic currencies (Danish krone, Swedish krona and Norwegian krone) and the Swiss franc because of a lack of data availability over the sample period. The definition of order flow is equal to the definition used in Evans and Lyons (2002a).

6 Due to the decentralized nature of the foreign exchange market, order flow data is available from end-user transactions (proprietary customer transactions from major banks), interdealer transactions (Thomson Reuters D2000-1), or brokered interdealer transactions (Thomson Reuters D2000-2 and EBS). In theory, the data from all sources should have similar characteristics and predictive content, as customer demand is the primary reason for trading. For a review of the different electronic trading systems we refer to Bjonnes and Rime (2005). Other studies using order flow data from Thomson Reuters D2000-2 include Payne (2003), Carlson and Lo (2006) and

Each transaction record contains a time stamp for the trade and a variable indicating the trade as a market buy or sell and the transaction price. This makes it unnecessary to use potentially inaccurate algorithms to assign the direction of the trade. A limitation of the data is the lack of information about the monetary value of each trade, as only the sign of each trade is given. However, studies on the information content of order flow data (e.g., Lyons and Moore, 2009) show that the signed order flow volume leads to the same conclusions as the absolute monetary volume of trades. This might be caused by the high degree of standardisation in the interdealer market, with a minimum trade volume of 1 million units of a base currency.

In addition to individual transactions, TRTH also provides intraday summary statistics. The database reports an opening (closing) bid and an opening (closing) ask price over certain time intervals ranging from 1 minute to 1 hour, which TRTH constructs using the bid and ask quote closest to the beginning (end) of the fixed time interval. After filtering out outliers, we calculate the exchange rate return as the difference between two log-mid quotes with the mid-quote calculated as the average of the closing bid and the closing ask quote. We add up tick-by-tick buying orders within a predetermined intraday time interval to obtain one aggregated value for that interval and do the same for selling orders. We then take the difference between the two values to yield the order flow for the time interval.

It is well documented that trading activity in FX markets slows down remarkably during weekends' and certain holidays' (Andersen, Bollerslev, Diebold, and Vega, 2003; Bauwens, Omrane, & Giot, 2005). Thus, we exclude a number of sparse trading periods. Following Danielsson, Luo and Payne (2012) and Frömmel, Kiss, and Pintér (2011), we

Danielsson and Payne (2012). The studies of Kileen, Lyons and Moore (2006) and Berger et al. (2008) use data from EBS. The initial studies of Evans and Lyons (2002a, 2002b) use data from Thomson Reuters D2000-1, a direct interdealer trading platform. Proprietary data from end-user trades are used in the studies of Evans and Lyons (2005b) and Cerrato et al. (2015).

remove the overnight period[7], weekends and some world–wide public holidays. Furthermore, following Selçuk and Gençay (2006) we eliminated any days that indicated no market activity. Specifically, we eliminated the days within which there were at least 12 consecutive 5-minute zero returns.

The total sample period spans from 3 January 2002 to 29 December 2013. It contains 361,977 observations for the 5-minute frequency, 124,000 observations for 15-minute frequency and over 30,000 observations for 60-minute frequency on the log returns as well as order flows for each exchange rate. Our sample period starts January 3, 2002 for two reasons: it was when the euro was physically introduced and it is the first date in Thomson Reuters D 2000-2 on which the required intraday data on the euro become available. Table 2.1 provides descriptive statistics for the return series.[8]

2.4.2 Methodology

Let $r_t = [r_{1t}, r_{2t}]$ be a 2×1 vector, containing two exchange rate return series ($r_{1t} \equiv \Delta s_{1t}$ and $r_{2t} \equiv \Delta s_{2t}$, where $s \equiv$ logarithm of the US dollar price of a currency). We filter each of the return series model to get its demeaned series z_{it} via an ARMA (m,n) model of the form:

$$r_{it} = c_i + \sum_{j=1}^{m} \varphi_{ij} r_{it-j} + z_{it} + \sum_{k=1}^{n} \kappa_{ik} z_{it-k}, \quad i = 1,\ 2\ . \tag{2.4}$$

[7] The overnight period is defined to be between 17:00 and 7:00 the following day. It should be noted that this definition is only appropriate for traders in London and New York, but not for traders in Asian markets. Love and Payne (2008) also exclude the overnight period in their study on the intraday relationship between order flow and exchange rate fluctuations. The main conclusions drawn from our analysis do not alter if we include the overnight period in the sample.

[8] In previous specifications, we also used 1-minute observations. However, at this frequency, there are many missing bid and ask quotes as well as lapses in the data feed, which biases the statistical analysis.

We assume that $z_t = [z_{1t}, z_{2t}]'$ is subject to a bivariate normal distribution conditional on the information set Ω_{t-1}:[9]

$$z_t \mid \Omega_{t-1} \sim N(0, H_t) \tag{2.5}$$

where H_t is the covariance matrix in which the diagonal elements h_{1t} and h_{2t} are the two conditional variances of z_{1t} and z_{2t} respectively. The covariance matrix H_t is modelled as:

$$H_t = D_t R_t D_t, \tag{2.6}$$

where $D_t = diag(h_{1t}^{1/2}, h_{2t}^{1/2})$, which is the 2×2 diagonal matrix of conditional standard deviations; $R_t = (diag(Q_t))^{-1} \cdot Q_t \cdot diag(Q_t))^{-1}$, which is the 2×2 conditional correlation matrix of $\varepsilon_t = [\varepsilon_{1t}, \varepsilon_{2t}]'$ which contains two standardized residuals defined as $\varepsilon_{1t} = z_{1t} / h_{1t}^{1/2}$ and $\varepsilon_{2t} = z_{2t} / h_{2t}^{1/2}$. $Q_t = \begin{pmatrix} q_{11t} & q_{12t} \\ q_{12t} & q_{22t} \end{pmatrix}$ is the 2×2 conditional covariance matrix of ε_t, whose elements, q_{11t}, q_{12t} and q_{22t}, are used to compute $\rho_{12t} = q_{12t} / (q_{11t} q_{22t})^{1/2}$.

We use a simple GARCH (1,1) model to estimate the conditional variances h_{it}:

$$h_{it} = \omega_i + \delta_i z_{it-1}^2 + \theta_i h_{it-1}, \qquad i = 1, 2. \tag{2.7}$$

For investigating how order flows affect exchange rate correlations, we use the asymmetric dynamic conditional correlation model of Engle (Engle, 2002) incorporating the absolute order flow differential as an exogenous variable affecting Q_t. We let Q_t have the following evolution:

[9] Even if conditional normality is absent, the results have a standard quasi-maximum likelihood (QMLE) interpretation.

$$Q_t = (\overline{Q} - A'\overline{Q}A - B'\overline{Q}B - G'\overline{N}G - \eta\overline{X}) + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + Gn_{t-1}n'_{t-1}G + B'Q_{t-1}B + \eta\,|x_{1t-1} - x_{2t-1}| \quad (2.8)$$

The parameter matrices are defined as:

$$A = \begin{pmatrix} \alpha_1 & 0 \\ 0 & \alpha_2 \end{pmatrix}, B = \begin{pmatrix} \beta_1 & 0 \\ 0 & \beta_2 \end{pmatrix}, G = \begin{pmatrix} g_1 & 0 \\ 0 & g_2 \end{pmatrix} \text{ and } K = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}.$$

The sample means are calculated as $\overline{Q} = T^{-1}\sum_{t=1}^{T}\varepsilon_t\varepsilon_t^{'}$, $\overline{N} = T^{-1}\sum_{t=1}^{T}n_t n_t^{'}$ ($n_t = \varepsilon_t^{-}$ or ε_t^{+}) and $\overline{X} = T^{-1}\sum_{t=1}^{T}|x_{1t-1} - x_{2t-1}|$.[10] . The absolute order flow differential $|x_{1t} - x_{2t}|$ as an exogenous variable enters q_{11t}, q_{12t} and q_{22t} in Q_t to affect ρ_{12t} indirectly.

We refer to the model characterised by Equation (2.8) as ADCCXS, which captures asymmetry via differing slopes across joint negative and joint positive values of past shocks. For the initial analysis following in Section 2.5.2, we set n_t equal to ε_t^{-}.

As Li (2011) suggests, Q_t can also be assumed to evolve according to what we call ADCCXE, to be differentiated from ADCCXS. The ADCCXE model captures asymmetry by appealing to "eccentricity", i.e., by re-centering the news impact surface away from zero. Specifically,

$$Q_t = (\overline{Q} - A'\overline{Q}A - B'\overline{Q}B - \eta\overline{X}) + A'e_{t-1}e'_{t-1}A + B'Q_{t-1}B + \eta\,|x_{1t-1} - x_{2t-1}| \quad (2.9)$$

where $\overline{Q} = \begin{pmatrix} 1 & \overline{\rho}_{12} \\ \overline{\rho}_{12} & 1 \end{pmatrix}$, $e_t = \begin{pmatrix} \varepsilon_{1t} + \gamma_1 \\ \varepsilon_{2t} + \gamma_2 \end{pmatrix}$, and other parameter matrices are defined as in the ADCCXS model (2.8) above.

[10] Note: Neither inclusion nor exclusion of the mean of the absolute order flow differential in the intercept alters the results and conclusions drawn from the analysis.

Some remarks regarding Equations 2.8 and 2.9. are in order. First, It is ready to see that ADCCXS and ADCCXE regress to respectively ADCCS and ADCCE if η = 0, and further to standard DCC if $g_1 = g_2 = 0$ and $\gamma_1 = \gamma_2 = 0$. We impose restrictions $\alpha_1^2 + \beta_1^2 + g_1^2 < 1$ and $\alpha_2^2 + \beta_2^2 + g_2^2 < 1$ when estimating ADCCXS, and impose $\alpha_1^2 + \beta_1^2 < 1$ and $\alpha_2^2 + \beta_2^2 < 1$ when estimating ADCCXE.

Second, the exogenous variable affects the conditional correlation coefficient p_{12t} in the following manner[11]. According to the definition of p_{12t} we have

$$p_{12t} = \frac{q_{12t}}{\sqrt{q_{11t}\ q_{22t}}} = \frac{C_{12t} + \eta y_{t-1}}{\sqrt{[C_{11t} + \eta y_{t-1}][C_{22t} + \eta y_{t-1}]}} \tag{2.10}$$

where, for the ADCCXE, [12]

$$y_{t-1} \equiv |x_{1t-1} - x_{2t-1}|$$

$$C_{11t} \equiv (1 - \alpha_1^2 - b_1^2) + \alpha_1^2(\varepsilon_{1t-1} + \gamma_1)^2 + \beta_1^2 q_{11t-1}$$

$$C_{12t} \equiv \bar{\rho}_{12}(1 - \alpha_1\alpha_2 - \beta_1\beta_2) + \alpha_1\alpha_2(\varepsilon_{t-1} + \gamma_1)(\varepsilon_{2t-1} + \gamma_2) + \beta_1\beta_2 q_{12t-1}$$

$$C_{22t} \equiv (1 - \alpha_2^2 - b_2^2) + \alpha_2^2(\varepsilon_{2t-1} + \gamma_2)^2 + \beta_2^2 q_{22t-1}$$

Note, C_{11t}, C_{12t} and C_{22t} do not depend on y_{t-1}. From these relations, we can derive $\frac{\delta \rho_{12t}}{\delta y_{t-1}} = \frac{\eta\left(1 - \frac{q_{12t}}{2}\right)}{\sqrt{[C_{11t} + \eta y_{t-1}][C_{22t} + \eta y_{t-1}]}}$. Clearly both the sign of η and the sign of $1 - \frac{q_{12t}}{2}$ determine the sign of $\frac{\delta \rho_{12t}}{\delta y_{t-1}}$, as $\sqrt{[C_{11t} + \eta y_{t-1}][C_{22t} + \eta y_{t-1}]}$ is >0. The estimation results show that $q_{12t} < 2$ for all *t* and all currency pairs. Thus, if $\eta < 0$, a rise (fall) in the absolute order flow differential in the previous period would reduce (increase) exchange rate correlation in the current period, given all other variables and parameters in (2.9); if $\eta > 0$, the

[11] In the following derivation we start with the initial condition that $\bar{X}$ is equal to zero.
[12] The same relationship between the absolute order flow differential and conditional correlation dynamics apply to the ADCCXS.

reverse is true; and if η = 0, the absolute order flow differential will have no effect on the correlation.

Third, the impact of joint positive (negative) news shocks - estimated via the parameters (γ_1, γ_2) in the ADCCXE specification - affect the conditional correlation as follows. The parameter estimates of (γ_1, γ_2) capture asymmetry by allowing the news impact surface to be centred away from zero. This implies a greater response (larger effect) to joint negative (positive) returns depending on where the surface is centred. When $\gamma_1, \gamma_2 > 0$ ($\gamma_1, \gamma_2 < 0$) positive-type (negative-type) asymmetry will result as the surface is centred at negative (positive) values of $\varepsilon_{1t-1}, \varepsilon_{2t-1}$ on ρ_{12t}. In other words, when $\gamma_1, \gamma_2 > 0$, joint positive shocks will increase correlation more than joint negative shocks of similar magnitude, and vice versa if $\gamma_1, \gamma_2 < 0$.

The estimation of all the GARCH and DCC parameters and the computation of their standard errors follow the two-stage procedure proposed in Engle and Sheppard (2001).

2.5 Empirical Results

2.5.1 Descriptive statistics

Table 2.1 presents the summary statistics for six exchange rate return (Panel A) and order flow (Panel B) series calculated for the highest (5-minute) frequency available in our sample.[13]. Each exchange rate is quoted as the US dollar price of one unit of a currency. Thus, a positive return for an exchange rate implies an appreciation of the currency against the US dollar (or a depreciation of the US dollar relative to the currency). One can see that the mean returns are positive for all exchange rates, suggesting that all the six currencies (EUR, GBP, JPY, AUD, NZD and CAD) appreciated on average relative to USD during the

[13] Descriptive statistics for the other frequencies are outlined in Table A.1 in Appendix A.

whole sample period[14]. The maximum appreciation and depreciation of the six currencies against USD occurred during the US subprime crisis and the following GFC. Last, the summary statistics exhibit high kurtosis. Thus the Jarque–Bera test statistics (JB-Prob.) decisively reject the null of a Gaussian distribution for all return series at all frequencies.

Table 2.1 Descriptive statistics

	USD–EUR	USD–GBP	USD–JPY	USD–AUD	USD–NZD	USD–CAD
Panel A. Descriptive statistics for 5-minute returns (n=361,977)						
Maximum	0.024	0.035	0.029	0.039	0.037	0.027
Minimum	-0.014	-0.031	-0.041	-0.057	-0.037	-0.022
Mean	1.18e-04	3.25e-05	6.4e-05	1.54e-04	1.87e-04	1.12e-04
Std.Dev.	0.059	0.056	0.062	0.082	0.087	0.059
Skewness	0.91	0.13	0.89	-1.57	-0.95	0.67
Kurtosis	103.24	142.66	259.35	285.9	144.55	79.32
JB-Prob.	0.00	0.00	0.00	0.00	0.00	0.00
Panel B. Descriptive statistics for 5-minute order flows (n=361,977)						
Maximum	109	214	75	229	110	278
Minimum	-93	-212	-95	-238	-87	-177
Mean	0.13	0.31	0.03	0.12	0.08	0.39
Std.Dev.	6.98	13.20	1.70	14.34	5.88	12.34
Skewness	0.10	0.08	-0.14	-0.16	-0.08	0.26
Kurtosis	8.39	8.19	17.31	11.97	9.91	14.97

Note: This table shows the descriptive statistics for 5-minute- returns (Panel A) and order flows (Panel B) for six exchange rates. Returns are calculated as the difference between the log mid-quote. Order flow is defined as the net of buyer- and seller transactions within a certain time interval. We measure order flow by adding up tick-by-tick buying orders within an intraday time interval (e.g. five-minute intervals) to obtain one aggregated value for the interval and do the same for selling orders. We then take the difference between the two values to yield the order flow for the time interval. The means and standard deviations for exchange rate returns are expressed in percentage points.

[14] Take two most frequently traded currencies as an example: The USD–EUR and USD–GBP rates at the beginning of our sample were USD 0.893 per EUR and USD 1.452 per GBP respectively, and rose to US$1.369 per EUR and US$1.636 per GBP in December 2013. This is equivalent to depreciations in USD of 53% against the EUR and of 13% against the GBP.

Panel B in Table 2.1 sets out the descriptive statistics for the order flows associated with each spot rate[15]. For all six exchange rates, the mean order flows are positive, indicating buying pressures on average for the six currencies. This is consistent with the positive mean returns on the six exchange rates. The mean order flow is largest (0.39) for the USD–CAD rate, followed by the USD–GBP rate (0.31), and is smallest but still positive (0.03) for the USD–JPY spot rate. A similar finding relates to the volatility of the order flows, with the volatility being largest (14.34) for USD–AUD followed by USD–GBP (13.20) and USD–CAD (12.34) and smallest for order flow related to USD–JPY (1.70). There are two reasons for this finding. First, Thomson Reuters D2000-2 is the main database used for USD–CAD and USD–GBP inter-dealer trading, whereas for USD–JPY spot trades, dealers more commonly use EBS. As we rely on Thomson-Reuters data, the smaller trading intensity for USD–JPY is likely to result in lower order flows on average than the high trading intensity for USD–CAD and USD–GBP trading. Second, for the USD–JPY rate, its average lower buying pressure than other exchange rates could be caused by carry trade activities. In a carry trade, investors borrow low-interest currencies (such as JPY) and buy high-interest currencies (such as AUD and NZD). The carry trade will then generate frequent sell orders for the USD–JPY rate, thereby offsetting the buying pressure on JPY arising from good news shocks to JPY. For the analyses following in the reminder of the chapter, we measure order flow in units of a hundred.

2.5.2 Order flow and correlation dynamics

Testing the five hypotheses developed in Section 2.3, we look at eight exchange rate pairs, for the following reasons. The EUR–JPY pair is the primary hedging instrument for trade between Asia and Europe, so it is relevant for multi-national companies. The EUR–

[15] Summary statistics for the 15- and 60-minute return and order flow series used in the analysis can be found in Table A1 in Appendix A.

GBP and GBP–JPY pairs are both actively traded crosses and thus they are relevant to FX traders. The EUR–CAD, JPY–CAD and GBP–CAD pairs rely heavily on oil price movements and so they are highly related to the underlying fundamentals. The AUD–NZD pair has high levels of historical intraday correlation in the sample, as does the JPY–AUD pair. More importantly, they are among the most popular pairs for carry trade because of the large IRDs between New Zealand or Australia and Japan.

Table 2.2 summarises the estimation results of the GARCH-ADCCXS parameters and the unconditional correlations for the eight exchange rate pairs at the 5-minute frequency.[16] In this table, all of the GARCH parameters are statistically significant at the 1% level and their sums are close to unity, indicating a high degree of persistence in the volatility. Meanwhile, all the eight exchange rate pairs have modified Box–Pierce Q-statistics for ε_{1t}, ε_{2t}, ε_{1t}^2 and ε_{2t}^2 that are statistically insignificant at the 5% level, implying that the standardised residuals are i.i.d.

Turning to the ADCCXS parameters, we can see that the dynamic parameters (α_i *and* β_i) are economically and statistically significant, showing that the correlations are time-varying and highly persistent. The resulting χ^2 statistics allow us to reject the null of a constant correlation in favour of a dynamic correlation model (such as the ADCCXS model) for all the eight pairs. Correlation persistence is high in several pairs, as reflected by large estimates of β_1 and β_2: 0.994 and 0.9464 for EUR–GBP, 0.9965 and 0.9237 for EUR–JPY, 0.9862 and 0.9402 for EUR–CAD, 0.9961 and 0.9934 for NZD–AUD, and 0.9793 and 0.9773 for JPY–CAD. The large persistence in correlation is a common finding for exchange rate returns, both at the intraday frequency (as in our study here) or the daily frequency (e.g., Li, 2011).

[16] We also used other univariate volatility models and obtained qualitatively similar results.

Regarding asymmetric responses to past negative shocks, the ADCCXS model suggests that joint bad news has an economically insignificant effect on correlations, as the g_i estimate ranges between 0.002 and 0.0794 for all exchange rate combinations except for the GBP–CAD return correlation. This is expected, since unlike equity returns there are no theoretical reasons for the presence of the asymmetrical negative effect of joint past bad news on currency return correlations. Thus we will test for another type of asymmetry, namely joint positive asymmetry, for which we have theoretical justification (see the next section).

Table 2.2 also indicates that unconditional correlation varies. The EUR–GBP and NZD–AUD pairs have $\bar{\rho}_{12}$ equal to 0.6559 and 0.6595 respectively, which is quite high. The EUR–CAD, EUR–JPY, GBP–CAD and GBP–JPY pairs show medium levels of unconditional correlation (0.3680, 0.2788, 0.2939, 0.3018 and 0.2394, respectively). The lowest levels are found for the JPY–AUD and JPY–CAD pairs (0.1453 and 0.0770, respectively). Different unconditional correlations may be useful information for analysis the different impacts of order flow on correlation dynamics, to which we now turn.

As stated in the introduction, the parameter of main interest is η, which is associated with the absolute order flow differentials. For all the pairs except NZD–AUD, the negative sign of η is in line with the expectations of H1: Under normal circumstances where the effect of order flow on an exchange rate is close enough between two currencies, a greater absolute order flow differential lowers the correlation between two exchange rate returns[17]. This negative effect is strong for some currency pairs with high and medium levels of unconditional correlation, such as the EUR–GBP (η = -0.0649), EUR–JPY (η = -0.0821), EUR–CAD (η = -0.0880) and GBP–JPY pairs (η = -0.1575). On the other hand, the negative

[17] See Section 2.3 for the theory leading to H1, and section 2.4.2, especially equation (2.10), for further explanations of the sign of η in terms of the effect of the absolute order flow differential on correlation.

effect is weaker for pairs with low levels of unconditional correlation, such as the JPY–CAD pair (η = -0.0350) and the JPY-AUD pair (η = -0.0459).

For the NZD–AUD pair, the effect of order flow on the correlation is positive (η = 0.0011) and statistically significant, albeit not economically significant. We take this as evidence to support H2. The hypothesis concerns abnormal circumstances where trading intensities are vastly different between AUD and NZD, and so the effect of the AUD's order flow on AUD fluctuations (b^{AUD}) and the effect of NZD's order flow on NZD fluctuations (b^{NZD}) differ from each other significantly. Smyth (2009) finds evidence for this argument, as he finds that the order flows of AUD have a stronger price impact on NZD than the order flows of NZD *per se*. This evidence can be used to explain why η is positive for the NZD–AUD pair, according to our earlier theoretical analysis that led to H2.

Table 2.2 Parameter estimates of the GARCH-ADCCXS model

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
GARCH model parameters								
ω_1	0.0137 **	0.0137 **	0.0137 **	0.0000 **	0.0000 **	0.0002 **	0.0000 **	0.0002 **
δ_1	0.0502 **	0.0502 **	0.0502 **	0.0009 **	0.0009 **	0.0398 **	0.0020 **	0.0398 **
θ_1	0.9190 **	0.9190 **	0.9190 **	0.9990 **	0.9990 **	0.9255 **	0.9980 **	0.9255 **
ω_2	0.0000 **	0.0002 **	0.0000 **	0.0002 **	0.0000 **	0.0000 **	0.0000 **	0.0000 **
δ_2	0.0009 **	0.0398 **	0.0014 **	0.0400 **	0.0014 **	0.0010 **	0.0010 **	0.0014 **
θ_2	0.9990 **	0.9255 **	0.9985 **	0.9252 **	0.9985 **	0.9988 **	0.9988 **	0.9985 **
ADCCXS model parameters								
α_1	0.0587 **	0.0642 **	0.0652 **	0.1294 **	0.2250 **	0.1625 **	0.0885 **	0.1796 **
α_2	0.2836 **	0.3494 **	0.3030 **	0.1127 **	0.0309 **	0.1037 **	0.0945 **	0.0714 **
β_1	0.9940 **	0.9965 **	0.9862 **	0.9861 **	0.9241 **	0.9824 **	0.9961 **	0.9793 **
β_2	0.9464 **	0.9237 **	0.9402 **	0.9843 **	0.9152 **	0.9888 **	0.9934 **	0.9773 **
g_1	0.0479	0.0536 **	0.0020	0.0270 **	0.1512 **	0.0273 **	0.0601 **	0.0062
g_2	0.0582 **	0.0642 **	0.0794 **	0.0790 **	0.1875 **	0.0671	0.0461 **	0.0424 **
η	-0.0649 **	-0.0821 **	-0.0880 **	-0.1575 **	-0.0158 **	-0.0459 **	0.0011 **	-0.0350 **
$\bar{\rho}$	0.6559	0.2788	0.3680	0.2394	0.2939	0.1453	0.6595	0.0770
Statistics								
LRT	430.07 **	150.27 **	358.86 **	69.35 **	507.85 **	147.86 **	224.79 **	672.38 **
Q_1^*	40.18 **	40.18 **	40.18 **	54.16 **	54.16 **	22.03 **	75.60 **	22.31 **
Q_{11}^*	66.14 **	66.14 **	66.14 **	179.59 **	179.59 **	83.57 **	232.31 **	83.75 **
Q_2^*	54.16 **	22.31 **	32.82 **	22.31 **	32.82 **	95.74 **	95.74 **	32.82 **
Q_{22}^*	179.59 **	83.75 **	111.64 **	83.75 **	111.64 **	64.90 **	64.90 **	111.64 **
$\chi^2 - \text{test}$	896.59 **	1993.65 **	991.16 **	880.43 **	600.99 **	2179.17 **	1244.39 **	716.72 **
LLF	1199720.7	1133318.2	1099094.6	1089038.9	1137824.5	964495.4	998992.4	1061799.9

Note: This table presents the parameter estimates of Equations 2.7 and 2.8 for eight exchange rate pairs evaluated at the 5-minute frequency. The estimation (and computation of the modified standard errors) uses the two-stage procedure proposed in Engle and Sheppard (2001). LRT denotes the likelihood ratio test statistic, which tests the null hypothesis of g_1=g_2 =0. LLF is the quasi-log-likelihood function evaluated at the maximum. ** and * denote significance at the 1% level and the 5% level, respectively. Q^* represents the Box–Pierce Q-statistic adjusted for heteroskedasticity (see Lobato, Nankervis and Savin, 2001). Q_1^*, Q_{11}^*, Q_2^* and Q_{22}^* are associated with $\varepsilon_{1,t}$ ${\varepsilon^2}_{1,t}$, $\varepsilon_{2,t}$ and ${\varepsilon^2}_{2,t}$ respectively. The χ^2-test proposed by Engle and Sheppard tests the null hypothesis of constant conditional correlation during the sample period ($R_t = \bar{R}$) against the alternative of dynamic conditional correlation ($vech\ (R_t) = vech(\bar{R}) + \iota_1\ vech(R_{t-1}) + \iota_p\ (vech(R_{t-p})$).

2.5.3 Positive-type asymmetry

In this section, we further evaluate the impact of joint positive and joint negative news shocks in driving exchange rate co-movement. Specifically we test for positive-type asymmetry against symmetry or negative-type asymmetry. The ADCCXE model described in Equation (2.9) serves our purpose.

Panel A in Table 2.3 sets out the parameter estimates of γ_1 and γ_2 in the model. We can see that they are all positive and statistically significant at the 1% level for all exchange rate pairs excluding NZD–AUD. This indicates that joint positive shocks increase future co-movements more than joint negative shocks of the same magnitude reduce future co-movements. This result is confirmed by the corresponding likelihood ratio tests: the seven likelihood ratio test (LRT) statistics allow us to reject the null of $\gamma_1 = \gamma_2 = 0$ for the EUR–GBP, EUR–JPY, EUR–CAD, GBP–JPY, GBP–CAD, JPY–AUD and JPY–CAD correlations. It is important to bear in mind that a positive (negative) shock as represented by $\varepsilon_t > 0$ ($\varepsilon_t < 0$) corresponds to bad (good) news about the US, because our definition of the exchange rate is the USD price of a currency.

The impact of joint positive shocks on co-movement is strong for some major exchange rate pairs. For example, the EUR–GBP pair has parameter estimates of γ_1 and γ_2 of 0.258 and 0.1937 respectively, and the JPY–EUR pair has parameter estimates of γ_1 and γ_2 of 0.2329 and 0.4274 respectively. A similar observation applies to the JPY–GBP correlation, where γ_1 and γ_2 are estimated to be 0.2848 and 0.1841 respectively. The strong positive-type asymmetry for the three major pairs can be explained by strong links among the EUR, JPY, GBP and USD, and, more importantly, by the argument of investor sentiment, which leads to the fourth hypothesis (H4) (see Section 2.3 for a detailed discussion). The results provided in

the analysis thus offer overwhelming evidence in support of this hypothesis at the intraday frequency.

However, the co-movements of the intraday AUD-NZD rates seem to be characterised by negative-type asymmetry persisting (γ_1 = -0.0103 and γ_2 = -0.0073 with statistical significance at the 1% level). In other words, the correlation responds to joint negative shocks ($\varepsilon_{1t-1} < 0$ and $\varepsilon_{2t-1} < 0$; good news about the US) more strongly than joint positive shocks ($\varepsilon_{1t-1} > 0$ and $\varepsilon_{2t-1} > 0$; bad news about the US). A possible explanation is that bad news about the US mainly affects the major currencies such as EUR, JPY and GBP, and not the minor currencies like NZD and AUD. So, although the correlations between major currencies respond significantly to bad news (joint positive shocks), the correlation between NZD and AUD does so to a much lesser degree. On the other hand, when good news about the US arrives (joint negative shocks), investors may transfer massive funds from the two minor currencies to USD, causing significant joint depreciation of AUD and NZD against USD. Accordingly, joint negative shocks (good news about the US) increase NZD–AUD correlations more than joint positive shocks (bad news about the US) of the same magnitude do.

It is worth noting that that the two γ estimates are smaller than their daily counterparts reported in Li (2011). For instance, Li reported that γ_1 = 0.5865 and γ2=0.6237 for the EUR–GBP correlation. This suggests that, although positive asymmetry dominates at the intraday frequency, this dominance is even stronger at the daily frequency.

Note that ADCCXE is not the only model that is suitable for testing for positive asymmetry, as the ADCCXS can be tailored to estimate the impact of joint positive shocks (by setting $n_t = \varepsilon_t{}^+$). Panel B in presents the estimates of the ADCCXS model testing for

positive asymmetry. Briefly, the estimates confirm the economic significance of joint positive return shocks, as the asymmetry parameters are significantly higher than the ADCCXS specification testing for negative-type asymmetry via the change in slopes. Again, the positive –type asymmetry is confirmed to be strong for some major exchange rate pairs. For example, the EUR–GBP pair has g_1 and g_2 parameter estimates of 0.1641 and 0.1761 respectively, and the GBP–JPY pair has g_1 and g_2 parameter estimates of 0.2831 and 0.1095 respectively. However, the ADCCXS specification for positive asymmetry does not nest the specification for negative asymmetry, although both nest the symmetry specifications. Therefore, it is unclear whether positive or negative asymmetry is preferred, although both specifications are preferred over symmetry for all exchange rate combinations tested, as indicated by the LRT statistics. By contrast, the ADCCXE model, which focuses on eccentricity, confirms the dominance of positive asymmetry over negative asymmetry in intraday exchange rate co-movements. We will therefore use the ADCCXE specification in the reminder of the chapter.

Table 2.3 Parameter estimates of the GARCH-ADCCXS and GARCH-ADCCXE models

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
Panel A. ADCCXE model parameters ($e_t = [\varepsilon_{1t} + \gamma_1;\ \varepsilon_{2t} + \gamma_2]$)								
α_1	0.0562 **	0.0572 **	0.0562 **	0.0997 **	0.3217 **	0.0766	0.1121**	0.0853 **
α_2	0.2816 **	0.3692 **	0.2816 **	0.0810 **	0.0390 **	0.0460	0.2215 **	0.0584 **
β_1	0.9939 **	0.9953 **	0.9939 **	0.9928 **	0.9085 **	0.9971 **	0.9961 **	0.9963 **
β_2	0.9534 **	0.9188 **	0.9534 **	0.9966 **	0.9554 **	0.9980 **	0.9930 **	0.9962 **
γ_1	0.2580 **	0.2329 **	0.2576 **	0.2848 **	0.0482	0.1938 **	-0.0103 **	0.2232 **
γ_2	0.1937 **	0.4274 **	0.1942 **	0.1841 **	0.2365 **	0.1520 **	-0.0073 **	0.2327 **
η	-0.0498 **	-0.0652 **	-0.0525 **	-0.1432 **	-0.0141 **	-0.0077 **	0.0215 **	-0.0143 **
LRT	120.39 **	101.85 **	666.89 **	271.80 **	1438.56 **	294.92 **	478.10 **	1398.43 **
LLF	1199897.3	1133787.4	1099200.0	1089727.2	1137469.2	964760.1	999634.5	1061856.5
Panel B. ADCCXS model parameters ($n_t = {\varepsilon_t}^+$)								
α_1	0.0491 **	0.0684 **	0.0588 **	0.0253 **	0.2003 **	0.0725 **	0.0899 **	0.0766 **
α_2	0.2797 **	0.2620 **	0.2782 **	0.0374 **	0.0380 **	0.0592 **	0.2410 **	0.0460 **
β_1	0.9609 **	0.9962 **	0.9957 **	0.9997 **	0.9617 **	0.9964 **	0.9921 **	0.9971 **
β_2	0.9024 **	0.9193 **	0.9603 **	0.9993 **	0.9415 **	0.9972 **	0.9901 **	0.9980 **
g_1	0.1641 **	0.0001	0.0019	0.2832 **	0.0743 **	0.0045	0.0563 **	0.1938 **
g_2	0.1761 **	0.1435 **	0.1158 **	0.1095 **	0.2700 **	0.0099	0.1865 **	0.1520 **
η	-0.0466 **	-0.0794 **	-0.0846 **	-0.1230 **	-0.0171 **	-0.0853 **	-0.1032 **	-0.0077 **
LRT	359.59 **	55.60 **	85.03 **	232.92 **	644.67 **	27.48 **	253.04 **	918.54 **
LLF	1199220.7	1133758.2	1099720.7	1089472.5	1138143.9	964275.9	999392.4	1061764.2

Note: Panel A of this table presents the parameter estimates of Equations (2.7) and (2.9) for eight possible combinations of exchange rate pairs using 5-minute return and order flow intervals. Panel B reports the parameter estimates for the ADCCXS model given by Equations (2.7) and (2.8). The estimation (and computation of the modified standard errors) uses the two-stage procedure proposed in Engle and Sheppard (2001). LRT denotes the likelihood ratio test statistic testing the null hypothesis of γ_1=γ_2 =0. LLF is the quasi-log-likelihood function evaluated at the maximum. ** and * denote significance at the 1% and 5% level respectively.

For a graphical representation of our estimation results, highlighting the greater impact of joint positive shocks over joint negative shocks on future co-movement, we present the news impact surface of both the initial ADCCXS model which targets negative type asymmetry and the ADCCXE model. We treat current exchange rate shocks (i.e., ε_{it}) as ‘‘news” – with positive and negative shocks as good and bad news – and examine the impact of such news on exchange rate correlation using the notion of the ‘‘news impact surface” of Kroner and Ng (1998). The news impact surface graphically highlights how different combinations of shocks ($\varepsilon_{1t}\ and\ \varepsilon_{2t}$) impact p_{12t}[18]. News impact surfaces are evaluated in the domain ε_i=[-3,3] (i=1,2).

Note that for the graphical representation outlined in Figures 2.1 and 2.2, we focus on how news shocks ($\varepsilon_{1,t\text{-}1}$, $\varepsilon_{2,t\text{-}1}$) affect correlation dynamics, given the absolute order flow differential ($|x_{1,t\text{-}1} - x_{2,t\text{-}1}|$), following Li (2011). The news impact surface is therefore based on the actual estimates from ADCCXE and ADCCXS, where γ1 and γ2 are jointly estimated with η although the estimate of η becomes redundant in this exercise.

Figure 2.1 shows the impact of news on the EUR–GPB correlations, using 5-minute return observations and the parameter estimates from the ADCCXS model targeting negative type asymmetry via the changes in the slopes (see Table 2.2). As expected, given the small parameter estimates (g_1=0.0479 and $g_2 = 0.0582$) associated with negative asymmetry, the news impact surface is almost symmetrical (if we compare the +,+ and -,- quadrants in Panel A) and centred at zero (Panel B). In other words, the likelihood that EUR and GBP tend to depreciate together against the USD, given that they have already depreciated together, is almost similar to the likelihood that they tend to appreciate together against the USD given that they have already appreciated together.

[18] The news impact surface for the ADCCXS (ADCCXE) is given by Equation 2.15 (2.16) in Appendix A.2.

Figure 2.1 The ADCCXS news impact surface

Panel A. Panel B.

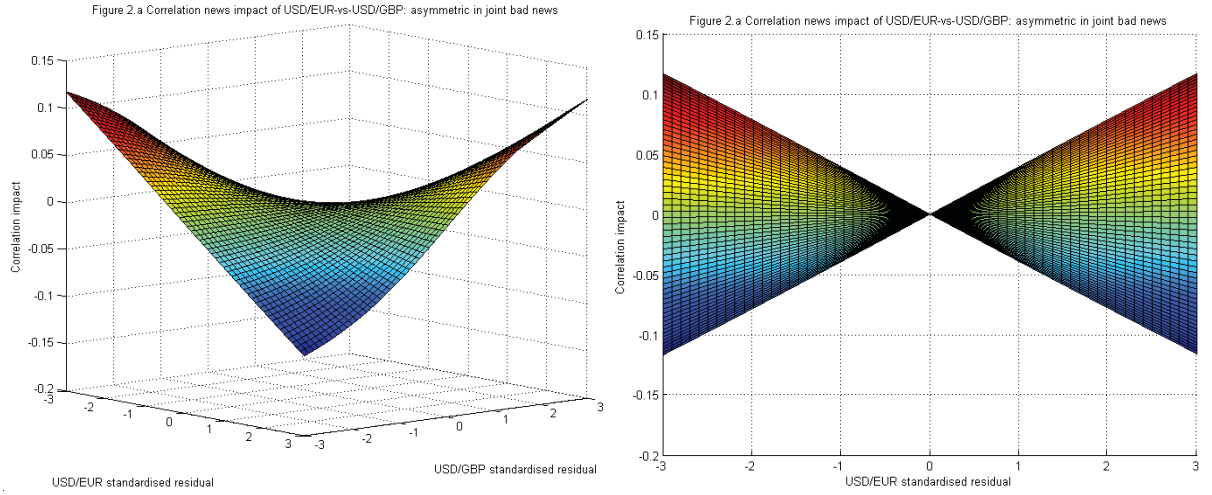

This figure presents the news impact surface for the EUR–GBP correlation. Note: this figure depicts the news impact surface by giving two different looks for the EUR–GBP correlation. We use the parameter estimates of the ADCCXS model taken from Table 2.2 for constructing the surface, with the absolute order flow differential given (by, e.g., assuming $|x_{1,\,t-1} - x_{2,\,t-1}| = 0$) to focus on the analysis of asymmetry.

By contrast, the news impact surface presented for the ADCCXE model (Figure 2.2) demonstrate that the surface is centred at a point (0.258, 0.1937)[19] in the ‘‘-,- ” standardised residual quadrant, away from the origin (0, 0), resulting in a greater surface value for joint positive than for joint negative standardised residuals of equal magnitudes. This implies a larger response to joint good news (in the +, +quadrant) than to joint bad news (in the -,- quadrant) of the correlation between the GBP and the EUR against the USD.

To be more concrete, let’s consider two scenarios, one where joint positive shocks ($\varepsilon_1 = \varepsilon_2 = 3$) hit the FX markets, and another scenario for joint negative shocks ($\varepsilon_1 = \varepsilon_2 = -3$). For the ADCCXS news impact surface function f(.,.), this would yield *f* (3,3) = 0.1152, whereas the latter yields *f*(-3,-3) =0.1162. The difference between joint positive shocks and joint negative shocks is hence a mere 0.001. By comparison, for the ADCCXE news impact

[19] The estimates for the asymmetry parameters are displayed in Table 2.2 for negative asymmetry and Table 2.3 for positive asymmetry.

surface, f (3, 3) = 0. 0.1292, whereas that for joint negative shocks is f (-3, -3) = 0.0964, about two-thirds of the former. This comparison further highlights the main differences between the ADCCXS and ADCCXE estimates, suggesting that joint positive returns have a much stronger effect on intraday correlation dynamics as opposed to joint negative shocks.

Figure 2.2 The ADCCXE news impact surface

Panel A. . Panel B.

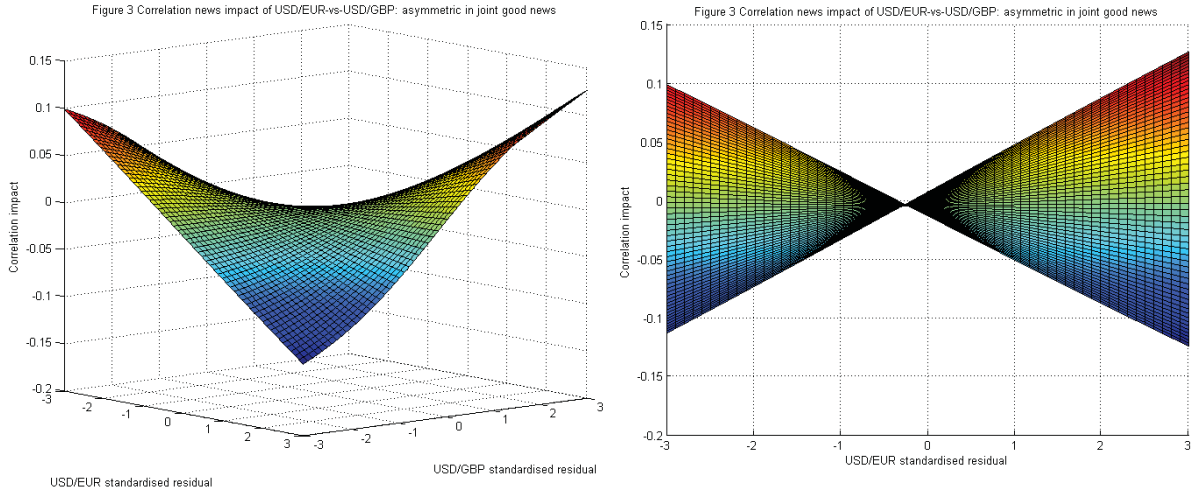

This figure presents the news impact surface for the EUR–GBP correlation using the parameter estimates of the ADCCXE model (except η) taken from Panel A in Table 2.3. Note: this figure depicts the news impact surface by giving two different looks for the EUR–GBP correlation. Note that the news impact surface is not centred at zero for the ADCCXE model.

2.5.4 Intraday Comparison

Table 2.4 compares three different intraday frequencies (5 minutes, 15 minutes and 60 minutes) in terms of the parameters for the ADCCXE model. This is to see if the impacts of order flow differentials on exchange rate correlations change as the frequency changes.

Compared to the 5-minute results in Panel A (reproduced from Panel A in Table 2.3), the estimates of η at the other two intraday frequencies (reported in Panels B and C) decline

in absolute terms. Specifically, the ADCCXE estimate of η for the EUR–GBP pair is -0.0498 for the 5-minute frequency, but shrinks to -0.0155 for the 15-minute frequency and to -0.0064 for the 60-minute frequency. For the EUR–JPY and GBP–JPY dynamic correlations, the same pattern also emerges, as the η estimated in the ADCCXE specification diminishes from -0.0652 and -0.1432 for the 5-minute frequency to -0.0048 and -0.0063 respectively for the 60-minute frequency. The same pattern applies to the other five exchange rate pairs, such as JPY–CAD, where η diminishes from -0.0143 to -0.0052 as the frequency lowers from 5-minute to 60-minute.

This pattern in exchange rate co-movements is consistent with the pattern in single exchange rate movements regarding the relationship between the impact of order flows and the data frequency (see Berger et al., 2008). Our bivariate evidence adds to top of Berger et al.'s univariate evidence confirming the gradual information hypothesis that the major part of information contained in order flows is incorporated into prices immediately, within minutes rather than hours.

The second observation to make from Table 2.4 is that persistence in correlation remains similarly strong across the different frequencies used, as demonstrated by the estimates of the nine β_1 values and the nine β_2 values all being above 0.91, with many above 0.99. This suggests that, even at the intraday frequency, shocks will have long-lasting effects on co-movements via persistent correlation dynamics.

With respect to the asymmetry parameter estimates, the documented positive-type asymmetry persists and hence dominates over negative-type asymmetry and symmetry across all the intraday frequencies evaluated. The magnitude of the positive asymmetry, however, does not follow a clear pattern across frequencies. For the EUR–GBP pair, γ_1 (γ_2) increases

from 0.2580 (0.1937) for the 5-minute frequency to 0.4246 (0.2034) for the 15-minute frequency and further increases to 0.5125 (0.330). By comparison for the EUR–JPY pair, γ_1 (γ_2) is initially 0.2329 (0.4274) for the 5-minute frequency and then decreases (increases) to 0.0940 (0.4776) for the 15-minute frequency, whereas the opposite is seen for the GBP–CAD pair.

To sum up, the negative relationship between the exchange rate correlation and the absolute order flow differential only slightly decreases as the intraday frequency increases, and thus the relationship is robust. This makes our main recommendations drawn from 5-minute-frequency data applicable to other intraday frequencies.

Table 2.4 Comparisons of intraday frequencies

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
Panel A. 5 minute frequency ADCCXE model estimates								
α_1	0.0562 **	0.0572 **	0.0562 **	0.0997 **	0.3217 **	0.0766	0.1121**	0.0853 **
α_2	0.2816 **	0.3692 **	0.2816 **	0.0810 **	0.0390 **	0.0460	0.2215 **	0.0584 **
β_1	0.9939 **	0.9953 **	0.9939 **	0.9928 **	0.9085 **	0.9971 **	0.9961 **	0.9963 **
β_2	0.9534 **	0.9188 **	0.9534 **	0.9966 **	0.9554 **	0.9980 **	0.9930 **	0.9962 **
γ_1	0.2580 **	0.2329 **	0.2576 **	0.2848 **	0.0482	0.1938 **	-0.0103 **	0.2232 **
γ_2	0.1937 **	0.4274 **	0.1942 **	0.1841 **	0.2365 **	0.1520 **	-0.0073 **	0.2327 **
η	-0.0498 **	-0.0652 **	-0.0525 **	-0.1432 **	-0.0141 **	-0.0077 **	0.0215 **	-0.0143 **
LLF	1199897.3	1133787.4	1099200.0	1089727.2	1137469.2	964760.1	999634.5	1061856.5
$\bar{\rho}$	0.6559	0.2788	0.3680	0.2394	0.2939	0.1453	0.6595	0.0770
Panel B. 15 minute frequency ADCCXE model estimates								
α_1	0.0599 **	0.0629 **	0.0492 **	0.0395 **	0.0548 **	0.2716 **	0.0627 **	0.0539 **
α_2	0.0497 **	0.0633 **	0.0738 **	0.0518 **	0.0531 **	0.0946 **	0.0614 **	0.0512 **
β_1	0.9982 **	0.9980 **	0.9988 **	0.9992 **	0.9985 **	0.9470 **	0.9977 **	0.9983 **
β_2	0.9984 **	0.9980 **	0.9973 **	0.9986 **	0.9986 **	0.9600 **	0.9981 **	0.9984 **
γ_1	0.4246 **	0.0940 **	0.0851 **	0.0025	0.6788 **	0.2152 **	0.7125	0.2409 **
γ_2	0.2034 **	0.4776 **	0.3286 **	0.1846 **	0.0609 **	0.4636 **	0.3661	0.5253 **
η	-0.0155 **	-0.0540 **	-0.0371 **	-0.0105	-0.0122 **	-0.0116 **	-0.0020	-0.0063
LLF	293016.9 **	245612.2**	246612.5 **	255833.5**	249095.5 **	210206.1**	220502.5 **	214571.5**
$\bar{\rho}$	0.6888	0.3298	0.4531	0.2424	0.3839	0.1429	0.7195	0.1141
Panel C. 60 minute frequency ADCCXE model estimates								
α_1	0.1849 **	0.1372 **	0.1274 **	0.0883 **	0.0886 **	0.1270 **	0.0639 **	0.1081 **
α_2	0.1073 **	0.1167 **	0.0753 **	0.1393 **	0.0616 **	0.1069 **	0.0757 **	0.0713 **
β_1	0.9822 **	0.9905 **	0.9909 **	0.9956 **	0.9961 **	0.9919 **	0.9979 **	0.9941 **
β_2	0.9834 **	0.9918 **	0.9939 **	0.9902 **	0.9973 **	0.9928 **	0.9971 **	0.9970 **
γ_1	0.5125 **	0.1353 **	0.3662 **	0.0840 **	0.0403	0.2146 **	0.2145 **	0.5731 **
γ_2	0.3340 **	0.4576 **	0.1372 **	0.2325 **	0.0000	0.0767 **	0.4433 **	0.0770 **
η	-0.0064 **	-0.0048	-0.0275 **	-0.0063 **	-0.0006	-0.0052 **	-0.0025	-0.0052
LLF	31054.3**	26040.7**	23087.3**	28055.8**	18874.5**	18597.5**	37367.2**	31899.5**
$\bar{\rho}$	0.6919	0.3317	0.4578	0.2340	0.4032	0.1493	0.7783	0.1229

Note: This table contrasts the parameter estimates of Equations (2.7) and (2.9) for eight exchange rate pairs evaluated at different intraday frequencies. The estimation (and computation of the modified standard errors) uses the two-stage procedure proposed in Engle and Sheppard (2001). LLF is the quasi-log-likelihood function evaluated at the maximum. ** and * denote significance at the 1% level and the 5% level, respectively.

2.5.5 Simulation

This section presents the simulation results based on the ADCCXE. These simulations aim to reveal the effects of order flow differentials on the evolution of time-varying correlation between exchange rate pairs more intuitively.

Figure 2.3 is based on the ADCCXE model[20], using the 15-minute-frequency EUR–GBP exchange rate correlation as an example. The figure shows the simulated "factual" and "counterfactual" correlation series and their differences. The uppermost solid line, labelled ''factual", uses the historical absolute order flow differentials, whereas the lower dotted line, labelled "counterfactual", uses the perturbed absolute order flow differentials. The perturbation is created by a permanent increase of 10 units[21] right from the beginning of the sample period. Though arbitrarily chosen, the size of the increase must be plausible enough to ensure the positive definiteness of Q_t or the non-negativity and nonzero of $q_{11,t}$ and $q_{22,t}$. Provided that this requirement is met, slightly changing the size of the increase does not qualitatively change the conclusion. This permanent increase in the absolute order flow differential, allows to see the degree to which increases/decreases in the historical absolute order flow differential would affect conditional correlation dynamics, as predicted by the ADCCXE model.

From Figure 2.3, one can see that the counterfactual line (green) always lies below the factual line (blue). In other words, had the absolute order flow differentials between the USD–GBP and USD–EUR been greater than their historical values, the USD–GBP-vs-USD–

[20] For the simulations, we ignore the mean of the absolute order flow differential for the ADCCXE estimation to make our graphical representation (albeit at the intraday frequency) comparable to that of Li (2011) (albeit at the daily frequency). As stated before, excluding the mean of the absolute order flow differential in the intercept does not change the sign or significance of the parameter estimate associated with the absolute order flow differential.

[21] This perturbation is equivalent to one standard deviation of the absolute order flow differentials for the EUR–GBP correlation at the 5-minute frequency.

EUR correlation would have been lower than its historical levels. Therefore, Figure 2.3 confirms visually that the absolute order flow differential has a negative effect on correlation for all the exchange rate pairs, again with the notable exception of the AUD–NZD correlation.

An economic implication follows directly from these findings: A trader witnessing a temporary increase in the absolute order flow differential, resulting in a temporarily weaker correlation between two assets, could enter a pair trade, betting on the correlation increasing again. A pair trade is defined as buying at an exchange rate that has previously depreciated and selling at an exchange rate that has previously appreciated, betting that the "spread" between the two will converge again.

Figure 2.3 Factual and counterfactual representation of the ADCCXE estimates

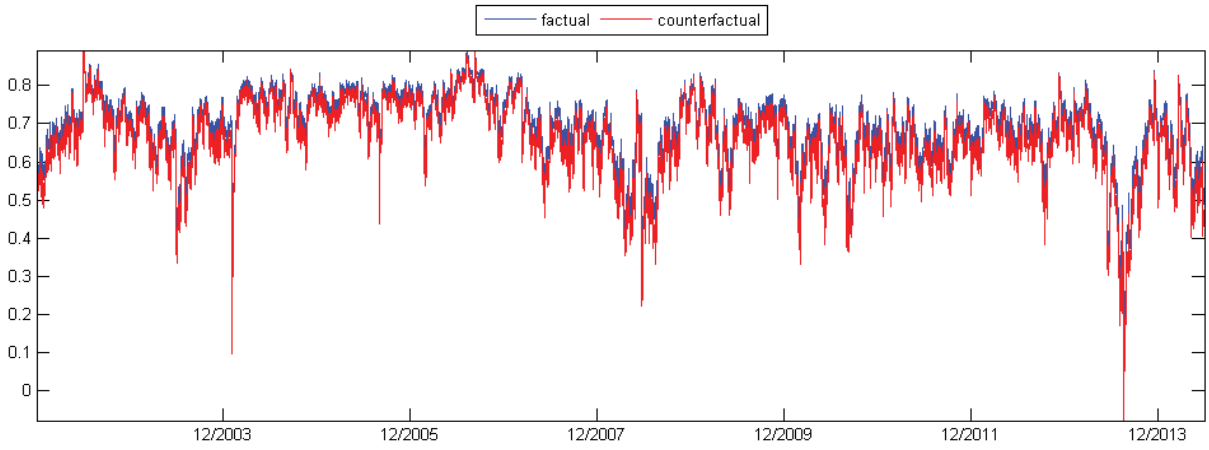

Note: This figure plots the simulated dynamic conditional correlations for the GBP–EUR pair using all the parameter estimates of the GARCH-ADCCXS model under two different scenarios: (1) we consider historical data on all the variables involved; and (2), we increase the historical absolute change in the absolute order flow differential by 10 units while retaining the historical data for all other variables involved. The blue ''factual" line represents the correlation series for (1) and the red ''counterfactual" line corresponds to the correlation series for (2).

Figure 2.4 is based on the ADCCXE model, using the 15-minute EUR–GBP exchange rate correlation as an example. We consider three simulations that were not seen in Figure 2.3. First, we simulate the model with all parameter estimates taken from Panel B in Table 2.4. This gives rise to the blue dotted line, labelled ADCCXE. Second, we impose a zero restriction on both $\gamma 1$ and $\gamma 2$ and leave the other parameters intact. This leads to the blue, solid line labelled DCCX. Third, we further impose a zero restriction on η, alongside the restrictions $\gamma 1 = \gamma 2 = 0$, and leave the other parameters unchanged. This yields the green dashed line labelled DCC. What is new about Figure 2.4 compared with previous studies such as that of Li (2011) is our focus on intraday covariance dynamics.

Figure 2.4 Unrestricted vs. restricted ADCCXE

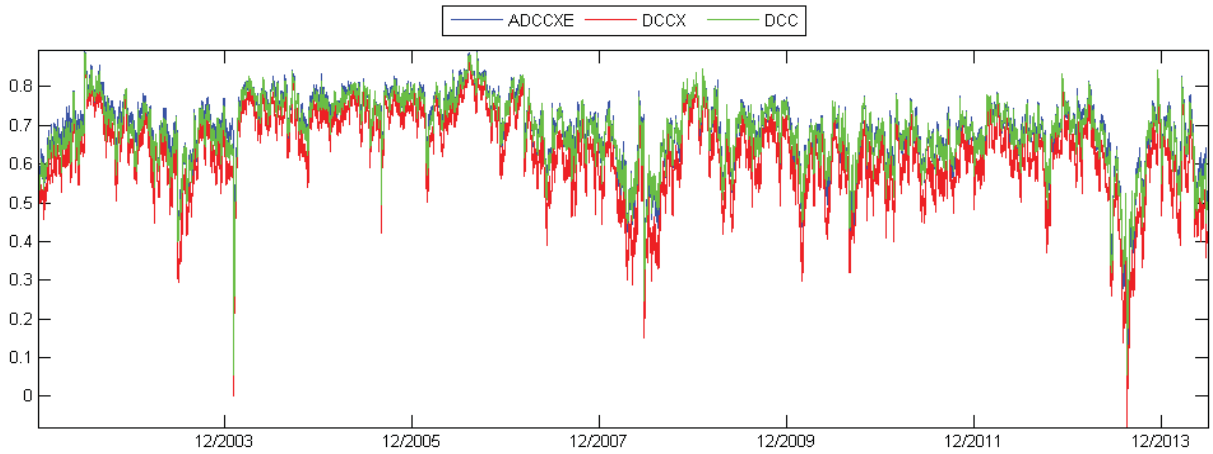

Note: This figure plots the simulated dynamic conditional correlations for the GBP–EUR pair based on the historical data on all the variables involved, under three different scenarios: (1) consideration of all the parameter estimates of the GARCH-ADCCXE models taken from Table 2.4 (Panel B), (2) a zero restriction on $\gamma 1$ and $\gamma 2$ while retaining all other parameter estimates and (3) another zero restriction on η.

Figure 2.4 illustrates that the ADCCXE (blue line) always lies above the DCCX (red line). The differences measure the contribution of positive asymmetry to the historical total correlation. In other words, ignoring positive asymmetry would lead us to underestimate the correlation. For example, the unrestricted ADCCXE model indicates a minimum correlation

between EUR-GBP of -0.0119, whereas the DCCX estimates a minimum correlation between EUR-GBP of -0.0819 towards the end of 2012.

With respect to the differences between ADCCXE (blue line) and DCC (red line) with the restrictions $\gamma_1 = \gamma_2 = 0$ and $\eta = 0$), ignoring both positive asymmetry and the effect of the absolute order flow differential will result into both over-and underestimation of the correlation. In case of a large order flow differential, the DCC will overestimate the conditional correlation dynamics, whereas in the case of a small order flow differential, the DCC will underestimate the conditional correlation between two currency pairs.

2.6 Structural Change

As the sample period contains major financial and economic events, testing for structural changes in correlation is indispensable. To ensure proper tractability with the large amount of intraday data used, we take an exogenous approach by pre-determining two break points. The first marks the start of the 2007 GFC on 9 August 2007. This is the date when there was an unusual jump in spreads between the overnight inter-bank lending rate and the London inter-bank offer rates. The event is often referred to as a "black swan" in the money market and has been used as a proxy for the beginning of the GFC (Taylor and Williams, 2008). The second breakpoint chosen is 9 March 2010. This is the date when the focus of global concern switched from the private sector to the public sector, as the International Monetary Fund and the European Union jointly announced that they would provide financial help to Greece. We take this date to be the start of the EDC. To perform the structural break tests, we modify the basic DCC model in a similar way to Dijk et al. (2011).[22]

[22] The ADCCXE specification taking two structural breaks into account is given in Appendix A (Section A2).

Figure 2.5 shows a graphical representation of the dynamic conditional correlation estimates over the sample period for the correlation dynamics of the major exchange rate combinations, such as the EUR–GBP pair. For the EUR–GBP pair, the unconditional correlation is significantly lower during the 2007 GFC (p=0.6850) than before this period (p= 0.7322). In the aftermath of the GFC, the correlation for the EUR–GBP pair decreases further (p= 0.6478)[23]. The same pattern applies to the EUR–JPY and GBP–JPY pairs, where the conditional correlation drops significantly during the latter stages of our dataset and swings more widely after the start of the 2007 GFC than before it. This graphical analysis therefore confirms the possibility of two structural changes arising in our dataset caused by the massive financial turmoil caused by the 2007 GFC and the 2010 EDC.

Figure 2.5 Conditional correlation dynamics for EUR–GBP, EUR–JPY and GBP–JPY

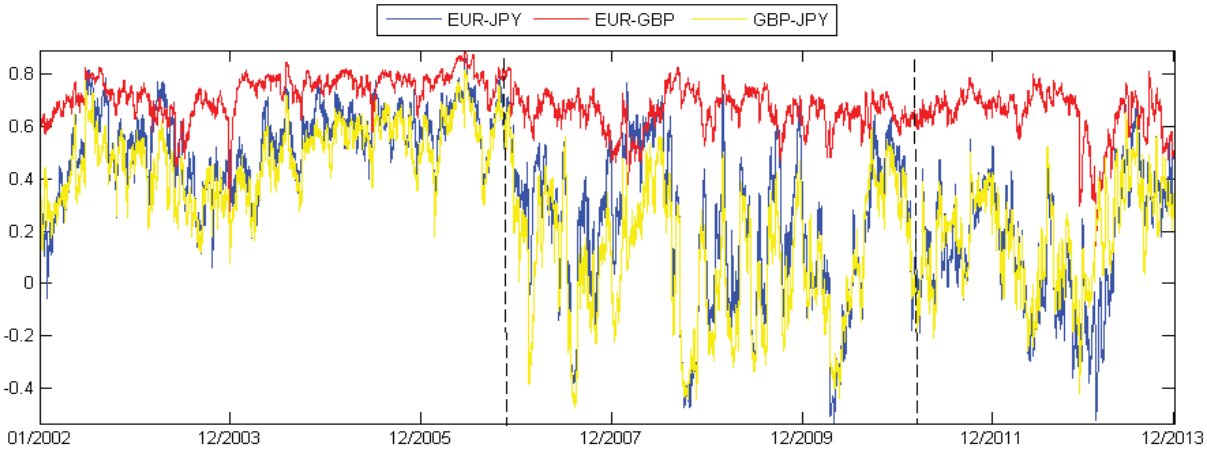

Note: This figure displays the conditional correlation dynamics over the total sample period for the EUR–GBP, EUR–JPY and GBP–JPY correlation dynamics. As before, EUR–GBP represents the correlation between USD–EUR and USD–GBP returns. The horizontal dashed lines indicate the two hypothesised structural breaks. For better visibility, we plot the 15-minute correlation dynamics; however, the same pattern arises for the other intraday frequencies.

[23] These correlations displayed here are the mean levels of the dynamic conditional correlation estimates for each subperiod.

Panel A in Table 2.5 sets out the specific parameter estimates for the ADCCXE model incorporating two structural breaks evaluated at the 5-minute frequency. Although we previously concluded that an abnormal case (a larger absolute order flow differential increasing exchange rate co-movement) is only seen for the AUD-NZD correlation dynamics, the structural break estimates suggest that the abnormal case was also prominent during the 2007 GFC for both the EUR–CAD and EUR–GBP correlation. For the EUR–CAD and EUR–GBP correlations, the order flow estimates were significantly positive during the 2007 GFC ($\eta_{\mathrm{dur-GFC}}$ = 0.0325 and 0.0288 respectively), whereas there is a strong negative relationship before ($\eta_{\mathrm{pre-GFC}}$ =-0.1070 and -0.0274) and after the 2007 GFC ($\eta_{\mathrm{post-EDC}}$ = -0.0405 and -0.0265). [24]

The key question is: what causes the relationship to be positive during the GFC for both EUR–GBP and EUR–CAD correlations? We again resort to the trading intensity-based explanation used to explain the positive relationship between absolute order flow differentials and the AUD–NZD correlation. Between 2007 and 2010, the trading intensity decreased drastically for the USD–EUR spot rate, compared to its pre-crisis level. However, average trading intensities increased for all other exchange rates, and in particular for the USD–GBP and USD–CAD rates each being the other party in the correlation pair with EUR.

The trading intensity of USD–EUR changed in opposite direction to those of USD–GBP and USD–CAD. Such opposite movements must have drastically increased the *differences* between the former one and the latter three trading intensities. As discussed in section 3, if the difference between $D(L, S_t^A)$ and $D(L, S_t^B)$ (with S depending on trading intensity) is large enough, this would cause a large enough difference in the discount factor b

[24] The main conclusions derived from the 5-minute analysis apply to the other frequencies. Please see Table A3 in Appendix A for the analyses using 15- and 60-minute intervals for all exchange rate pairs.

between two currencies (e.g., $b_t^B << b_t^A$), which could then lead to a change of the sign of η capturing the effects of order flows on their correlations. Based on those theoretical discussions and on the above observations about trading intensities, we infer that the discount factors for the exchange rates USD–EUR, USD–GBP and USD–CAD became drastically apart, such that $b^{EUR} << b^{GBP}$ and $b^{EUR} << b^{CAD}$. Accordingly, the absolute order flow differentials started to have positive effects on the EUR–CAD and EUR–GBP co-movements, in line with the conjecture about abnormal trading circumstances in hypothesis H2.

Another question is whether alternative models, without break or with one break or with two breaks, confirm the existence of two breaks in the correlation dynamics. The LRT displayed in Panel A in Table 2.5 tests the null hypothesis of no structural break against the alternative of two structural breaks. They are all statistically significant at a higher than 1% level, thus rejecting the null. The LRT statistics in Panel B (C) test the null hypothesis of one structural break associated with the 2007 GFC (the 2010 EDC) against the alternative of two breaks. The statistics reject the null mostly at a higher than 1% level, and in two cases at a higher than 5% level, in favour of the alternative of two breaks. Thus, we conclude that the model with two structural breaks due to the 2007 GFC and the 2010 EDC significantly outperforms the model with no break, or a single break due to the 2007 GFC, or a single break due to the 2010 EDC. This amounts to lending support to hypothesis H3.

Table 2.5 Parameter estimates of the GARCH-ADCCXE model with structural change

	EUR-GBP	EUR-JPY	EUR-CAD	GBP-JPY	GBP-CAD	JPY-AUD	JPY-CAD
Panel A ADCCXE parameter estimates with 2 breaks (2007 GFC and 2010 EDC)							
α_1	0.0538 **	0.0666 **	0.0624 **	0.0283 **	0.2357 **	0.0606 **	0.0916 **
α_2	0.3462 **	0.3927 **	0.2915 **	0.0296 **	0.0360 **	0.0556 **	0.0552 **
β_1	0.9521 **	0.9968 **	0.9891 **	0.9996 **	0.9303 **	0.9209 **	0.9957 **
β_2	0.9015 **	0.9169 **	0.9472 **	0.9996 **	0.2678 **	0.9214 **	0.9963 **
γ_1	0.5112 **	0.0710 **	0.1661 **	0.4506 **	0.0137 **	0.3100 **	0.0974 **
γ_2	0.2020 **	0.2578 **	0.1506 **	0.0871 **	0.0179 **	-0.0991 **	0.4857 **
$\eta_{pre-GFC}$	-0.0274 **	-0.0879 **	-0.1070 **	-0.0128 **	-0.0669 **	-0.0108 **	-0.0235 **
$\eta_{dur-GFC}$	0.0288**	0.0113 **	0.0325 **	-0.0093 **	-0.0966 **	-0.0554 **	-0.0083 **
$\eta_{post-EDC}$	-0.0265 **	-0.0715 **	-0.0405**	-0.0015 **	-0.0511 **	-0.0114 **	-0.0022
$\bar{\rho}_{pre-GFC}$	0.690	0.302	0.364	0.434	0.312	0.327	0.229
$\bar{\rho}_{dur-GFC}$	0.618	0.337	0.391	0.006	0.337	-0.137	-0.085
$\bar{\rho}_{post-EDC}$	0.615	0.217	0.325	0.143	0.197	0.015	-0.003
LLF	1199963.2	1134267.9	1099413.2	1089927.2	1138064.5	965320.5	1062176.5
LRT	131.8 **	961.9 **	426.3 **	400.0 **	1190.8 **	1120.4 **	640.6 **
Panel B Log likelihood ratio test for H_0: 1 break (2007 GFC) against H_1: 2 breaks (2007 GFC and 2010 EDC)							
LLF	1199943.9	1134168.4	1099410.8	1089857.3	1138012.1	965290.9	1062167.2
LRT	38.0 **	199.0**	4.80 *	139.8 **	104.8 **	59.2 **	18.6 **
Panel C Log likelihood ratio test H_0: 1 break (2010 EDC) against H_1: 2 breaks (2007 GFC and 2010 EDC)							
LLF	1199926.2	1134163.4	1099389.8	1089887.2	1138042.1	965273.3	1062174.5
LRT	74.0 **	209.0 **	46.8 **	80.0 **	44.88 **	94.4 **	4.0 *

Note: This table presents the parameter estimates of the GARCH-ADCCXE model with two structural breaks and the log-likelihood test results for structural breaks. The break dates are, respectively, August 9, 2007 for the starting of the 2007 global financial crisis (GFC), and March 9, 2010 for the 2010 European sovereign debt crisis (EDC). The parameterization of the model is given in full detail in Appendix A.2. The estimation and computation of the modified standard errors use the two-stage procedure proposed in Engle and Sheppard (2001). LLF is the quasi-log-likelihood function evaluated at the maximum. LRT denotes the likelihood ratio test statistic. "pre-GFC" denotes the first subsample period prior to the 2007 GFC. "dur-GFC" denotes the second subsample period from the starting of the 2007 GFC to the starting of the 2010 EDC. "post-EDC" denotes the third subsample period between the starting of the 2010 EDC and the end of the whole sample period. ** and * denote significance at the 1% level and 5% level, respectively.

2.7 Robustness

2.7.1. Bid–ask-spread

The first robustness check is to see whether our results will change if we incorporate the bid–ask spread as an additional driving force for correlation dynamics. Traders are sensitive to actual bid–ask spreads, as the spreads determine their transaction costs and hence they may substitute trades in one currency for another that has a correlated exchange rate but smaller bid–ask spreads. This is particularly relevant for intraday (technical) traders, as they trade much more frequently than a long-term trader, thereby making transaction costs an essential consideration in their trading decisions. The sensitivity of traders to an increase in the bid–ask spread can also be used to infer the relationship between the bid–ask spread and exchange rate co-movements.

We conjecture that the relationship is negative: a rise (fall) in the difference between the bid-ask spreads across currencies leads to lower (higher) co-movements of their exchange rates. The intuition is this. Consider a currency trader who wishes to trade USD for either EUR or GBP as EUR and GBP are highly correlated in terms of their exchange rates against USD. Whether the trader should ultimately choose EUR or GBP for trading depends further on another important consideration – the bid-ask spread differential between the two currencies. If the spreads associated with EUR and GBP rise/decline by the same amount, so do their transaction costs, then the trader will be indifferent between the two trading decisions. However, if the spread increases for (say) EUR and/or declines for GBP, the trader should switch from trading USD for EUR to trading USD for GBP, to reduce transaction costs. As a result, GBP is likely to experience more frequent trading than EUR with USD, implying that GBP would experience greater appreciation/depreciation than EUR. In this

case, the magnitude of joint appreciation or joint depreciation of EUR and GBP would become smaller than when their bid-ask spreads increase (decreases) by the same amount.

Bid-ask spreads are related to market liquidity, which is a fundamental aspect of any market. Although there may be a strong link between bid-ask spreads and exchange rate co-movements, the market microstructure literature suggests that all the information about future fundamentals (including bid-ask spreads) will be incorporated in current order flows. Current order flows in turn transfer the information into price movements. To test this conjecture, we modify the ADCCXE model in Equation (2.9) as follows:

$$Q_t = (\overline{Q} - A'\overline{Q}A - B'\overline{Q}B - \eta_1\overline{X} - \eta_2\overline{Y}) + A'e_{t-1}e'_{t-1}A + B'Q_{t-1}B + \eta_1 \mid x_{1t-1} - x_{2t-1} \mid + \eta_2 \mid y_{1t-1} - y_{2t-1} \mid \quad (2.11)$$

where,

$\overline{X} = T^{-1}\sum_{t=1}^{T}\left|x_{1t-1} - x_{2t-1}\right|$, $\overline{Y} = T^{-1}\sum_{t=1}^{T}\mid y_{1t-1} - y_{2t-1} \mid$. $\mid x_{1t-1} - x_{2t-1} \mid$ and $\mid y_{1t-1} - y_{2t-1} \mid$ represent the absolute differences between the order flows and bid-ask spreads, respectively. All other variables are defined as in Equation 2.9.

Note that the bid-ask spread affects the conditional correlation dynamics as outlined in equation 2.10, i.e. if $\eta_2 < 0$, a rise (fall) in the absolute bid-ask spread differential in the previous period would reduce (increase) exchange rate correlation in the current period, given all other variables and parameters in (2.09); if $\eta_2 > 0$, the reverse is true; and if $\eta_2 = 0$, the absolute bid-ask spread differential will have no effect on the correlation.

Table 2.6 shows the results of our ADCCXE (Equation 2.11) with the bid–ask spread as exogenous variable in addition to the absolute order flow differential. We centre our analysis in depth on the 5-minute correlation dynamics, with the results for the other frequencies reported in Appendix A.

Although the bid–ask spread is statistically significant overall for seven out of the nine testing specifications, the parameter estimate of the absolute order flow differential is still significantly negative for all exchange rate pairs with the exception of the AUD–NZD pair in the sample. The parameter estimate increases for EUR–GBP, EUR–CAD, GBP–CAD, JPY–AUD and JPY–CAD, whereas the estimate decreases for EUR–JPY and GBP–JPY.

The increase in the parameter estimate is small for both the EUR–CAD and the JPY–CAD pair, where η_1 increases from -0.0525 to -0.0552 for EUR–CAD and from -0.0143 to -0.0145 for EUR–CAD. Large increases in the estimate related to the absolute order flow differential arise for EUR–GBP (-0.0498 to -0.0778) and GBP–CAD (-0.0141 to -0.0346). By contrast, the estimate related to the order flow differential drops significantly for GBP–JPY (-0.1432 to -0.0532) and EUR–JPY (-0.0652 to -0.0520). Note that the relationship between absolute order flow differential and co-movement was particularly strong for these two pairs before controlling for the underlying bid–ask spread.

For several spot correlations (such as EUR–CAD and GBP–JPY), the bid–ask spread parameter has the wrong sign, indicating that a larger bid–ask spread increases co-movement. This counterintuitive finding, coupled with the insignificance for one-third of our sample, makes us conclude (again in line with the market microstructure literature) that all information about changes in the underlying market fundamentals (including spreads and liquidity by definition) are already incorporated into order flows. [25]

An interesting field for future research would be to investigate whether order flows can be linked to spread dynamics itself (i.e., to investigate the effect of the absolute order flow differentials on the co-movement of bid–ask spreads from several currencies). We leave this aspect for future research.

[25] One may suggest to further incorporate IRDs as an additional exogenous variable in the correlation process. However, at an intraday frequency, the interest rate observations are constant throughout a day.

Table 2.6 Robustness of results: order flow and bid–ask spread

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
ADCCXE parameter estimates								
α_1	0.2496 **	0.0521 **	0.0605 **	0.0650 **	0.0971 **	0.1215 **	0.1241**	0.0372 **
α_2	0.2319 **	0.2643 **	0.2840 **	0.2179 **	0.0908 **	0.3738 **	0.3721 **	0.0198 **
β_1	0.9057 **	0.9584 **	0.9928 **	0.9853 **	0.9932 **	0.9000 **	0.9000 **	0.9993 **
β_2	0.9629 **	0.9514 **	0.9523 **	0.9723 **	0.9959 **	0.9000 **	0.9000 **	0.9998 **
γ_1	0.4917 **	0.0967 **	0.2064 **	0.0071 **	0.6387 **	0.1646 **	0.1339	0.2318 **
γ_2	0.4671 **	0.1055 **	0.1672 **	0.0229 **	0.1725 **	0.0855 **	0.0735	0.1325 **
η_1	-0.0778 **	-0.0520 **	-0.0552 **	-0.0532 **	-0.0346 **	-0.0198 **	0.0199 **	-0.0145 **
η_2	-0.0231 **	-0.0004 **	0.0023 **	0.0490 **	-0.0009	0.0266 **	-0.0091 **	-0.0000
LLF	1199970.7	1133789.4	1099219.6	1091313.4	1137477.0	964768.4	999682.4	1063864.0
$\bar{\rho}$	0.6522	0.2818	0.3549	0.2404	0.2645	0.1402	0.6591	0.1038

Note: This table presents the parameter estimates of the ADCCXE model of Equation (2.11), focusing on 5-minute frequencies. LLF is the quasi log-likelihood function evaluated at the maximum. The parameter η_1 is associated with the absolute order flow differential, whereas η_2 corresponds to the bid–ask spread differential. ** and * denote significance at the 1% and 5% levels respectively.

2.7.2 Standardised measures of order flow

Yet another question to address is whether our findings are robust to different standardisations of order flows. One potential concern is that individual order flows (referred to below as "un-standardised order flows") are of different magnitudes across different currency markets. For example, the mean imbalance between buyer- and seller-initiated order flows and the trading intensity are much greater in the trades of EUR with USD than in those of AUD with USD: the EUR market's order flows overwhelmingly dominate the AUD market's. So, our analysis based on un-standardized order flows might be biased when comparing the transactions of EUR and of AUD with USD. In fact, this concern is relevant to the use of daily net order flows for analysis. However, we still want to address it at the intraday frequency.

In doing so, we follow Menkhoff et al. (2012) and divide order flows by their standard deviations. This standardisation aims to make markets more comparable in terms of mean order flow levels and their fluctuations. Another approach is to relate the net order flows to the trading intensity in each market, thereby giving more weight to the net order flows arising from lower trading intensities. This setting will penalise larger net order flows due to larger trading volumes, thus making markets more comparable. The standardisations are done in the following manners:

$$\tilde{x}^{\sigma}_{i,t} = \frac{of_{i,t}}{\sigma(of_i, t:T)} \text{ for the volatility based standardisation;} \qquad (2.12)$$

$$\tilde{x}^{ti}_{i,t} = \frac{of_{i,t}}{ti_{i,t}} \text{ for the trading intensity based standardisation}^{26} \text{ ;} \qquad (2.13)$$

[26] In case the trading intensity (and hence the order flow) is zero during a particular interval we set $\tilde{x}^{ti}_{i,t}$ equal to 0.

where $\tilde{x}_{i,t}^{\sigma}$ denotes the standardised order flow of currency *i* adjusted by the standard deviation of its order flow at *t*; $\tilde{x}_{i,t}^{ti}$ the standardised order flow of currency *i* adjusted by its trading intensity at *t*; $of_{i,t}$ the order flow of currency *i* at *t*; $\sigma(of_i, t{:}T)$ the standard deviation of the order flow of currency *i*; and $ti_{i,t}$ the trading intensity of currency *i*. We then replace the un-standardised order flows with the standardised order flows in the ADCCXE model.

Table 2.7 reports the results of this robustness test. Panel A displays the parameter estimates for the ADCCXE model using the volatility-adjusted standardised order flows $\tilde{x}_{i,t}^{\sigma}$. We can see that the conclusions drawn from previous analysis using un-standardised order flows remain unchanged. Specifically, for all exchange rate pairs under investigation, the relationship between the absolute order flow differential and co-movement is still significantly negative. As with the analysis conducted in section 5.2, the negative relationship is strongest for the EUR-JPY and GBP-JPY co-movements (η = -0.0116 and -0.0215 respectively). The relationship is the weakest for the GBP-CAD and JPY-CAD correlation dynamics (η = -0.0032 and -0.0033 respectively).

Panel B reports the estimation results of the ADCCXE model using the trading-intensity-adjusted standardised order flows $\tilde{x}_{i,t}^{ti}$. We can see that the similar pattern to that in Panel A emerges: although one η parameter (with GBP-CAD) has a statistically and economically insignificant estimate (-0.0005), all others remain statistically and economically significant. More importantly, all the η estimates are still negative, confirming the robustness of the baseline estimation results presented in Table 2. These results of robustness check suggest that, even controlling for different trading intensities or order flow volatilities, the negative relationship stated in hypothesis H1 generally does not change. This conclusion also holds true for other intraday frequencies (See Table A2 in Appendix A.1).

Table 2.7 Robustness of results: Standardised measures of order flow

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
Panel A. ADCCXE model parameters using the volatility-adjusted standardised order flows								
α_1	0.2016 **	0.1138 **	0.0617 **	0.0325 **	0.2752 **	0.1300 **	0.3605 **	0.0853 **
α_2	0.2723 **	0.0930 **	0.2740 **	0.0148 **	0.0397 **	0.1503 **	0.1457 **	0.0584 **
β_1	0.9695 **	0.9522 **	0.9936 **	0.9995 **	0.9580 **	0.9741 **	0.9000 **	0.9963 **
β_2	0.9000 **	0.9528 **	0.9560 **	0.9999 **	0.9926 **	0.9885 **	0.9625 **	0.9962 **
γ_1	0.2176 **	0.1497 **	0.1580 **	0.1995 **	0.0698 **	0.3874 **	0.0747 **	0.5232 **
γ_2	0.2645 **	0.1582 **	0.1521 **	0.0152	0.4309 **	0.2739 **	0.0962 **	0.2327 **
η	-0.0054 **	-0.0116 **	-0.0034 **	-0.0215 *	-0.0032 **	-0.0039 **	0.0657 *	-0.0033 *
LLF	1196417.7	1133852.1	1099103.9	1089591.2	1137893.5	968444.7	999942.4	1061005.8
$\bar{\rho}$	0.6575	0.2801	0.3546	0.2414	0.3017	0.1483	0.6608	0.1128
Panel B. ADCCXE model parameters using the trading-intensity-adjusted standardised order flows								
α_1	0.1971 **	0.1082 **	0.0561 **	0.0833 **	0.0296 **	0.1383 **	0.2533 **	0.2081 **
α_2	0.2715 **	0.0762 **	0.2687 **	0.2193 **	0.0180 **	0.1175 **	0.1410 **	0.0710 **
β_1	0.9708 **	0.9939 **	0.9947 **	0.9724 **	0.9996 **	0.9816 **	0.9528 **	0.9780 **
β_2	0.9000 **	0.9971 **	0.9581 **	0.9747 **	0.9998 **	0.9897 **	0.9667 **	0.9906 **
γ_1	0.2068 **	0.0969 **	0.2312 **	0.2767 **	0.0853 **	0.4776 **	0.2087 **	0.5702 **
γ_2	0.2534 **	0.2552 **	0.1884 **	0.5642 **	0.4036 **	0.3932 **	0.3182 **	0.0959 **
η	-0.0124 **	-0.0025 *	-0.0016 **	-0.0018 **	-0.0005	-0.0167 **	0.0761 **	-0.0005 **
LLF	1196572.1	1133658.1	1099173.4	1089002.5	1137739.6	968222.9	999647.7	1061018.1
$\bar{\rho}$	0.6504	0.2779	0.3549	0.2424	0.2989	0.1428	0.6610	0.1088

Note: This table presents the parameter estimates of the ADCCXE model, using respectively the volatility-based standardisation of order flow (Panel A), and the trading-intensity-based standardisation of order flow (Panel B). The intraday frequency used in estimation is 5- minute. The estimation and computation of the modified standard errors use the two-stage procedure proposed in Engle and Sheppard (2001). LLF denotes the log-likelihood value evaluated at the maximum. ** and * denote significance at the 1% and 5% levels respectively.

2.8 Order Flow and Volatility Dynamics

A final question related to modelling co-movement is whether order flow also drives conditional variances, thereby affecting the results of testing hypothesis H1. This question is relevant because, based on the dynamic conditional correlation models, the covariance between two assets is a function of both conditional correlation and conditional variance. Since a larger order flow represents a greater buying or selling pressure for an asset, it should increase the conditional variance of the asset's returns. Thus, one needs to control for this volatility effect of order flow while investigating the relation between order flow and dynamic correlation.

Our approach to incorporating order flow as a driving force for variance dynamics is different from studies that relate the trading volume to volatility. Such studies apply theoretical market-microstructure models, involving strategic interactions among asymmetric informed agents (Payne, 2003). These models suggest that, due to information asymmetry, higher trading volumes can result in price jumps. However, we conjecture that it is the buying or selling pressures from order flows that are responsible for price jumps. This conjecture comprises a simpler explanation for two reasons. First, it does not require agents to be rational. Second, it does not rely on complex models to explain price changes, but simply and more intuitively maintains that a larger imbalance between demand and supply would result in greater price variations/volatility.

We assume that the exchange grate return follows a GARCH-X process (X = "exogenous"):

$$r_{it} = c_i + \sum_{j=1}^{m} \Phi_{ij} r_{it-j} + z_{it} + \sum_{k=1}^{n} \kappa_{ik} z_{it-k} \; ; \qquad (2.14)$$

$$z_t | \, \Omega_{t-1} \sim N(0, H_t); \qquad (2.15)$$

$$h_{it} = w_i + \delta z_{it-1}^2 + \theta h_{it-1} + \xi |X_{t-1}|, \quad (2.16)$$

where X_{t-1} is defined as the order flow for each exchange rate. Note that we use absolute order flows, as we want to relate the absolute buying or selling pressure to the variance. The second-stage correlation dynamics follows the ADCCXE representation outlined before. Note that the GARCH model in Equation (2.7) is a nested version of the GARCH-X in Equation (2.16), with ξ=0 being imposed.

Again, we base our analysis on the highest available frequency (i.e., 5 minutes)[27]. Table 2.8 shows the parameter estimates of the GARCH-X–ADCCXE model. We are particularly interested in the exogenous parameters of the volatility equation (ξ). At first sight, for all the exchange rates evaluated, the relationship between absolute order flow and volatility is positive and highly significant, which is in line with our hypothesis. In other words, a larger absolute order flow increases the volatility of an exchange rate, which is consistent with the idea that larger absolute order flows represent greater buying or selling pressure, resulting in larger price variations. This finding is statistically significant for all exchange rates and the relationship is particularly large for USD–GBP (0.0074) and USD–JPY (0.0032) volatility dynamics, whereas the effect is rather modest for the USD–AUD (0.0005), USD–NZD (0.0006) and USD–EUR (0.0007) volatility dynamics.

Apart from these GARCH-X estimates, we have to note that the first-stage GARCH-X estimation does not alter the conclusions drawn in the second-stage estimation. For all exchange rate pairs except the AUD–NZD correlation (i.e., the relationship between absolute order flow differentials and the correlation dynamics) is still significantly negative, with the strongest relationship estimated for the EUR–GBP (η = -0.0451), EUR–CAD (η = -0.0490) and GBP–JPY(η = -0.0516) correlation dynamics.

[27] Chapter 3 will evaluate the forecasting performance of this GARCH-X approach using daily frequencies; results for the other intraday frequencies are outlined in Table A5 in the Appendix A.

Table 2.8 Parameter estimates of the GARCH-X–ADCCXE model

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
GARCH model parameters								
ω_1	0.0006 **	0.0006 **	0.0006 **	0.0000 **	0.0000 **	0.0000 **	0.0000 **	0.0000 **
δ_1	0.0026 **	0.0026 **	0.0026 **	0.0053 **	0.0053 **	0.0025 **	0.0038 **	0.0025 **
θ_1	0.9954 **	0.9954 **	0.9954 **	0.9927 **	0.9927 **	0.9955 **	0.9942 **	0.9955 **
ξ_1	0.0007 **	0.0007 **	0.0007 **	0.0074 **	0.0074 **	0.0032 **	0.0006 **	0.0032 **
ω_2	0.0000 **	0.0000 **	0.0000 **	0.0000 **	0.0000 **	0.0000 **	0.0000 **	0.0000 **
δ_2	0.0053 **	0.0025 **	0.0519 **	0.0025 **	0.0519 **	0.0028 **	0.0027 **	0.0518 **
θ_2	0.9927 **	0.9955 **	0.9143 **	0.9955 **	0.9143 **	0.9952 **	0.9953 **	0.9143 **
ξ_2	0.0074 **	0.0032 **	0.0017 **	0.0032 **	0.0017 **	0.0005 **	0.0005 **	0.0017 **
ADCCXE model parameters								
α_1	0.0530 **	0.0307 **	0.0381 **	0.0623 **	0.2010 **	0.0457 **	0.2946 **	0.1717 **
α_2	0.0299 **	0.0459 **	0.0191 **	0.1537 **	0.0653 **	0.0264 **	0.0713	0.0757 **
β_1	0.9983 **	0.9995 **	0.9993 **	0.9972 **	0.9782 **	0.9990 **	0.9000 **	0.9850 **
β_2	0.9996 **	0.9989 **	0.9997 **	0.9873 **	0.9770 **	0.9997 **	0.9568 **	0.9899 **
γ_1	0.0001	0.0999 **	0.1637 **	0.2266 **	0.1362 **	0.1013 **	-0.0427	0.2642 **
γ_2	0.1155 **	0.2100 **	0.0123	0.1291 **	0.1374 **	0.4469 **	-0.1635 **	0.3200 **
η	-0.0053 **	-0.0101 **	-0.0490 **	-0.0516 **	-0.0195 **	-0.0181 **	0.0087 **	-0.0262 **
LLF	1199912.9	1133990.0	1099289.1	1098920.9	1138200.0	974600.9	993064.8	1061312.6
$\bar{\rho}$	0.6631	0.2774	0.3412	0.2398	0.3017	0.1194	0.6660	0.1224

This table presents the parameter estimates of Equations (2.9) and (2.16) for eight exchange rate pairs evaluated at the 5-minute frequency. The estimation (and computation of the modified standard errors) uses the two-stage procedure proposed in Engle and Sheppard (2001). LLF is the quasi-log-likelihood function evaluated at the maximum. ** and * denote significance at the 1% level and the 5% level, respectively.

Whereas this analysis documents the positive relationship between order flow and exchange rate volatility, Figure 2.6 graphically presents the differences in the evolution of the underlying volatility between the GARCH-X and GARCH parameterisations. This graphical analysis documents the nature of the conditional return series, as the volatility is governed by a high average persistence, as well as the implications of the 2007 GFC, which resulted in an enormous spike in conditional volatility. Based on this graphical representation, one might argue that testing for a structural break in the volatility dynamics akin to Dijk et al. (2011) might be in order. We leave this for future research.

It is clearly visible that the volatility spikes estimated by the GARCH-X procedure are higher volatility spikes, a pattern that is not as high in the in the normal GARCH process. The reason for those spikes is that order flows spike after a news shock, thereby generating price jumps and hence an increase conditional volatility. Given the fluent spikes in volatility predicted by the GARCH-X model, this approach could be beneficial for predicting conditional volatility during financial turmoil, a question we evaluate in the next chapter.

Figure 2.6 GARCH-X and GARCH variance dynamics

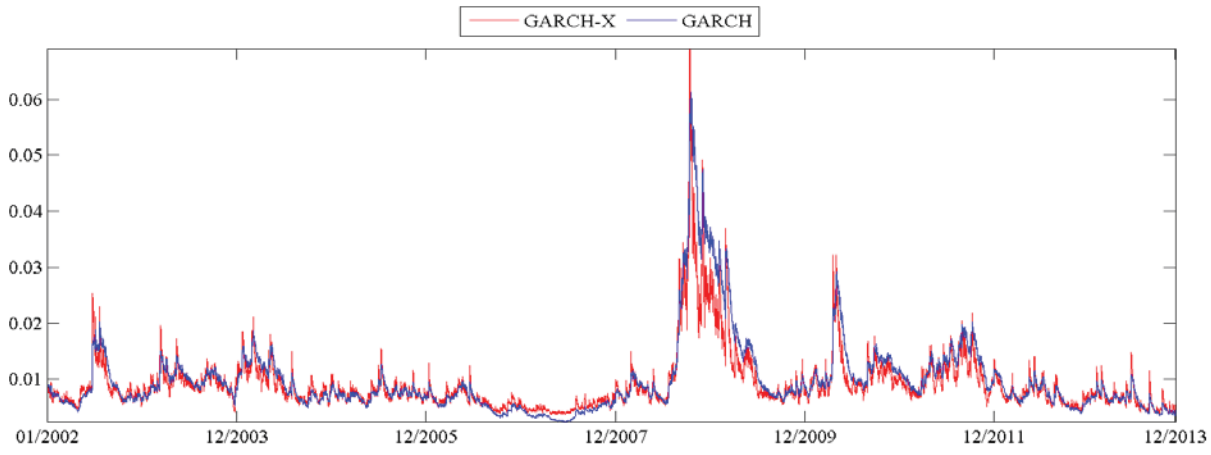

Note: This figure displays the estimated conditional variances for the USD–EUR exchange rate returns using two different GARCH specifications. The red line displays the evolution of volatility following from the GARCH-X estimates, given by Equation (2.16). The blue line displays the estimated variance dynamics from the GARCH (1,1) model, given by Equation (2.7).

2.9 Conclusion

This chapter extends traditional microstructure theory to analyse multivariate exchange rate movements in a high-frequency setting. We theoretically link and empirically test the link between the absolute order flow differential and the correlation dynamics of eight exchange rate pairs. In what follows, we summarise the main findings and hypotheses tested in this chapter.

First, this chapter confirms that correlations are dynamic and time-varying at intraday frequencies, which is in line with the existing literature. Though this finding is not necessarily unique to financial returns, this chapter evaluates how exchange rates co-move during intraday intervals for the first time in the literature. We find that intraday correlation dynamics are governed by an enormous level of persistence in correlation. We also find that intraday exchange rates co-move more closely after joint appreciation than after joint depreciation. This finding was previously reported for daily exchange rate correlation dynamics. We attribute this finding to investor sentiment and, in particular, to over-reactions to negative news for the US economy, which can explain why positive asymmetry dominates negative asymmetry for intraday correlation dynamics.

Second, as a unique result, the research provides strong evidence for the link between order flow and exchange rate correlation dynamics, finding that a higher order flow differential is negatively related to return co-movement. This effect is particularly strong for the highest intraday frequency evaluated and lowers as the frequency decreases, confirming the gradual incorporation of order flow information into prices. Apart from the gradual decrease in the relationship, we also find that the relationship is stronger for exchange rate pairs with higher levels of unconditional correlation. We argue that this is because exchange rates with low levels of unconditional correlation exhibit large absolute order flow differentials more frequently, thereby lowering the impact of an additional increase in the

absolute order flow differential. The negative relationship is highly significant for all exchange rate pairs evaluated, with the notable exception of the AUD–NZD correlation dynamics, which we argue arises from the vast differences in trading intensity between USD/AUD and USD/NZD.

Third, incorporating two structural breaks caused by the 2007 GFC and the dramatic worsening of the crisis in 2010 considerably improves the fit of our dynamic conditional correlation model. In addition to that, we document that the effect of the order flow differential changed for both the EUR–GBP and EUR–CAD pairs during the crisis, as a larger order flow differential suddenly increased the future co-movement for both pairs.

Our findings dictate the direction for the second essay, as testing the forecasting performance of our proposed models – thereby further examining the practical benefits of this research – is a natural extension.

Appendix A

A.1 Additional Tables

Table A.1 Further descriptive statistics

	USD–EUR	USD–GBP	USD–JPY	USD–AUD	USD–NZD	USD–CAD
Panel A. 15 minute returns						
Maximum	0.024	0.031	0.031	0.043	0.036	0.028
Minimum	-0.017	-0.027	-0.042	-0.059	-0.039	-0.023
Mean	3.50e-04	1.0e-04	1.87e-04	4.45e-04	5.35e-04	3.31e-04
Std.Dev.	0.09	0.09	0.11	0.14	0.14	0.09
Skewness	0.51	-0.14	0.42	-0.90	-0.67	0.47
Kurtosis	35.43	43.88	99.53	99.75	58.54	31.45
JB-Prob.	0.00	0.00	0.00	0.00	0.00	0.00
Panel B. 15 minute order flow						
Maximum	179	300	95	326	119	369
Minimum	-136	-252	-132	-368	-148	-330
Mean	0.38	0.87	-0.06	0.36	0.25	1.15
Std.Dev.	13.7	24.76	3.27	26.20	10.57	23.36
Skewness	0.20	0.14	0.61	-0.15	-0.11	0.297
Kurtosis	9.78	8.99	82.11	12.13	10.63	13.87
Panel C. 60 minute returns						
Maximum	2.77	3.59	3.19	4.13	4.03	3.28
Minimum	-1.88	-2.83	-4.25	-5.82	-4.02	-2.57
Mean	1.38e-03	3.31e-04	7.50e-04	1.80e-03	2.10e-03	1.31e-03
Std.Dev.	0.19	0.19	0.15	0.26	0.28	0.17
Skewness	0.20	-0.44	0.17	-0.67	-0.65	0.25
Kurtosis	11.57	14.80	28.20	35.62	19.83	12.79
JB-Prob.	0.00	0.00	0.00	0.00	0.00	0.00
Panel D. 60 minute order flow						
Maximum	276	489	99	637	156	536
Minimum	-241	-501	-117	-609	-264	-500
Mean	1.53	3.50	-0.02	1.44	0.98	4.60
Std.Dev.	29.3	52.1	7.2	57.6	22.70	50.9
Skewness	0.17	0.05	0.05	-0.21	-0.23	0.25
Kurtosis	4.73	6.89	25.30	8.80	8.52	9.40

Note: This table shows the descriptive statistics for 15 and 60-minute- returns- (Panels A and C) and order flows (Panels B and D) for six exchange rates. Returns are calculated as the difference between the log mid-quote. Order flows are aggregated tick-by-tick over the corresponding for 15 and 60-minute intervals. The means and standard deviations for exchange rate returns are expressed in percentage points.

Table A.2 Parameter estimates of the GARCH-ADCCXE model with structural change – further frequencies

	EUR-GBP	EUR-JPY	EUR-CAD	GBP-JPY	GBP-CAD	JPY-AUD	JPY-CAD
Panel A. 15-minute observations							
$\eta_{pre\text{-}GFC}$	-0.0065 **	-0.0325 **	-0.0089 **	-0.0109 **	-0.0015 **	-0.0142 **	-0.0018 **
$\eta_{dur\text{-}GFC}$	0.0064 **	0.0223 **	0.0048 **	-0.0019 **	-0.0004 *	-0.0556 **	-0.0035 **
$\eta_{post\text{-}EDC}$	-0.0086 **	-0.0184 **	-0.0150 **	-0.0014 **	-0.0007 **	-0.0589	-0.0025 **
$\bar{\rho}_{Pre-GFC}$	0.7359	0.2847	0.3145	0.4471	0.318	0.3719	0.1918
$\bar{\rho}_{Dur-GFC}$	0.6572	0.3032	0.3567	0.0082	0.338	-0.1339	0.0480
$\bar{\rho}_{Post-EDC}$	0.6326	0.2475	0.2824	0.1898	0.248	0.1130	0.0080
$\mathrm{LLF}_{GFC\ \&\ EDC}$	293154.2	245761.31	246810.8	256065.3	249484.4	211219.6	214950.2
LRT_1	266.1 **	298.2 **	396.6 **	463.9 **	777.8 **	2027.1 **	757.2 **
LLF_{GFC}	293119.2	245668.9	246651.0	255871.6	249452.3	211039.5	214834.2
LRT_2	70.1 **	184.7 **	319.6 **	387.4 **	64.2 **	320.8 **	232.7 **
LLF_{EDC}	293115.6	245615.64	246631.2	255838.5	249441.5	211059.55	214902.8
LRT_3	77.26 **	291.3 **	359.6 **	453.7 **	86.6 **	360.8 **	98.4 **
Panel B. 60-minute observations							
$\eta_{pre\text{-}GFC}$	-0.0069 **	-0.0047 **	-0.0012 **	-0.0038 **	-0.0038 **	-0.0068 **	-0.0013 **
$\eta_{dur\text{-}GFC}$	0.0057 **	0.0087 *	0.0018 **	-0.0012 **	-0.0042	-0.0042**	-0.0019 **
$\eta_{post\text{-}EDC}$	-0.0043 **	-0.0016 **	-0.0014 **	-0.0010 **	-0.0107	-0.0099 **	-0.0011 **
$\bar{\rho}_{Pre-GFC}$	0.7639	0.2876	0.3250	0.3833	0.3567	0.2813	0.1237
$\bar{\rho}_{Dur-GFC}$	0.6555	0.3037	0.3811	0.0358	0.3582	0.0358	0.0271
$\bar{\rho}_{Post-EDC}$	0.6423	0.2287	0.3380	0.1727	0.1987	0.0713	0.0419
$\mathrm{LLF}_{GFC\ \&\ EDC}$	31124.9	26280.5	23137.9	28089.5	18934.8	18610.7	31913.1
LRT_1	338.6 **	479.5 **	54.06 **	65.95 **	120.75 **	26.311 **	27.13 **
LLF_{GFC}	31117.8	26270.4	23130.4	28066.9	18916.8	18609.5	31912.8
LRT_2	14.13 **	20.27 **	36.18 **	45.01 **	36.18 **	2.30	0.6
LLF_{EDC}	31077.3	26212.9	23122.2	28065.3	18912.9	18599.6	31903.1
LRT_3	95.28 **	135.1 **	22.55 **	48.04 **	43.90 **	26.08 **	20.04 **

Note. This table presents the parameter estimates of the GARCH-ADCCXE model with two structural breaks and the log-likelihood test results for structural breaks. The data used for estimation are, respectively, 15-minute (Panel A) and 60-minute (Panel B). The break dates are, respectively, August 9, 2007 for the starting of the 2007 global financial crisis (GFC), and March 9, 2010 for starting of the 2010 European sovereign debt crisis (EDC). The parameterization of the model is given in full detail in equation (A2) in Appendix A.2. The estimation and computation of the modified standard errors use the two-stage procedure proposed in Engle and Sheppard (2001). LRT denotes the likelihood ratio test statistic. LRT_1 tests the null hypothesis of no break against the alternative of 2 breaks. LRT_2 tests the null of 1 break (the 2007 GFC break) against the alternative of 2 breaks. LRT_3 tests the null of 1 break (the 2010 EDC break) against the alternative of 2 breaks. $\mathrm{LLF}_{GFC\ \&\ EDC}$ denotes the log-likelihood function of the model with 2 structural breaks evaluated at the maximum. LLF_{GFC} (LLF_{EDC}) the log-likelihood function of the model with one structural break due to the 2007 GFC (2010 EDC) evaluated at the maximum. ** denotes significance at the 1% level.

Table A.3 Robustness of results: Standardised measures of order flow – further intraday frequencies

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
Panel A. ADCCXE model parameter estimate of η using 15-minute observations and volatility-adjusted standardised order flows								
η	-0.0050 **	-0.0002 **	-0.0076 **	-0.0006	-0.0018 **	-0.0035 **	0.0006**	-0.0015 **
LLF	292972.9	249565.2	262607.5	256132.7	249268.5	208214.0	220590.8	221002.2
Panel B. ADCCXE model parameter estimate of η using 60-minute observations and volatility-adjusted standardised order flows								
η	-0.0024 **	-0.00074**	-0.0025 **	-0.0009 **	-0.0022 *	-0.0008 **	-0.0140**	-0.0015 **
LLF	31362.2	26140.7	23190.3	28093.7	24793.3	18594.5	37164.2	31899.2
Panel C. ADCCXE model parameter estimate of η using 15-minute observations and trading-intensity-adjusted standardised order flows								
η	-0.0635 **	-0.040 *	-0.0110 **	-0.0004 *	-0.0011 **	-0.0007	-0.0092**	-0.0421**
LLF	293168.9	249648.2	262622.0	255959.2	249377.6	215137.5	220149.5	219593.0
Panel D. ADCCXE model parameter estimate of η using 60-minute observations and trading-intensity-adjusted standardised order flows								
η	-0.0223 *	-0.0060 **	-0.0102 **	-0.0003 **	-0.0001	-0.0001**	-0.0000	-0.0154**
LLF	31181.7	26134.6	23172.8	28111.9	27489.2	18513.9	37413.4	31887.3

This table presents the η parameter estimates of the ADCCXE model, using respectively the volatility-based standardised order flows for the 15- (Panel A) and 60-munite data (Panel B), and using the trading-intensity-based standardised order flows for the 15- (Panel C) and 60-minute data (Panel D). The estimation and computation of the modified standard errors use the two-stage procedure proposed in Engle and Sheppard (2001). LLF denotes the log likelihood function evaluated at the maximum. ** and * denote significance at the 1% and the 5% level respectively.

Table A.4 Robustness of results: order flow and bid-ask spread- further intraday frequencies

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
Panel A. 15 minute observations								
η_1	-0.0019 ***	-0.0209 **	-0.0030 ***	-0.0003	-0.0001	-0.0027 **	0.0318 **	-0.0026 **
η_2	-0.0030***	0.0448 **	-0.0020	-0.0020	-0.0045 **	0.0457	-0.0078 **	0.0040 **
LLF	293215.1	245522.9	246657.4	255912.7	249435.3	211001.1	221694.4	214419.9
Panel B. 60 minute observations								
η_1	-0.0025 **	-0.0037 **	-0.0092 **	-0.0051	-0.0008 **	-0.0026 **	-0.0014	-0.0016
η_2	0.0429 **	0.0790 **	-0.0138	-0.0015	0.0076 **	0.0136 **	0.0002	0.0004
LLF	31060.5	26135.7	23111.7	28072.3	24796.8	18597.8	37408.3	31906.1

This table presents the parameter estimates of the ADCCXE model in Equation (2.11), focusing on 15- and 60-minute frequencies. LLF is the quasi-log-likelihood function evaluated at the maximum. The parameter η_1 is associated with the absolute order flow differential, whereas η_2 corresponds to the bid–ask spread. ** and * denote significance at the 1% and 5% levels respectively.

Table A.5 Order flow and volatility dynamics – further intraday frequencies

	USD–EUR	USD–JPY	USD–CAD	USD–GBP	USD–AUD	USD–NZD
Panel A. 15 minute observations						
ω_1	1.31e-05 **	1.41e-05 **	1.05e-05 **	6.75e-06 **	1.21e-05	1.61e-05
δ_1	0.0041 **	0.0040 **	0.0038 **	0.0048 **	0.0044 **	0.0041 **
θ_1	0.9939 **	0.9940 **	0.9932 **	0.9932 **	0.9936 **	0.9939 **
ξ_1	5.28e-05 **	3.63e-04 **	1.23e-04 **	3.71e-05 **	7.68e-05 **	2.38e-04**
Panel B. 60 minute observations						
ω_1	3.93e-05 **	6.02e-05 **	4.98e-05 **	1.06e-05 **	8.50e-05 **	2.53e-05 **
δ_1	0.0074**	0.0141 **	0.0096 **	0.0068 **	0.0098 **	0.0105 **
θ_1	0.9905 **	0.9834 **	0.9883 **	0.9891 **	0.9882 **	0.9875 **
ξ_1	1.05e-04 **	2.98e-06	1.29e-05**	8.71e-05 **	4.97e-06 **	4.82e-05

This table presents the parameter estimates of the GARCH-X model in Equation (2.16), focusing on 15- and 60-minute frequencies. The parameter ξ_1 is associated with the absolute order flow ** and * denote significance at the 1% and 5% levels respectively.

A.2 News Impact Surface and Structural Break Specification

The news impact surface for the ADCCXS correlation estimates is given by the following formulae from Cappiello, Engle, and Sheppard (2006):

$$f(\varepsilon_1, \varepsilon_2) = \frac{\bar{c}_{12}+(\alpha_1\alpha_2+\gamma_1\gamma_2)\varepsilon_1\varepsilon_2+\beta_1\beta_2\bar{\rho}_{12}}{\sqrt{[\tilde{c}_{11}+\alpha_1^2+\gamma_1{}^2\,\varepsilon_1{}^2+\beta_1{}^2][\tilde{c}_{22}\alpha_2^2+\gamma_2{}^2\,\varepsilon_2{}^2+\beta_2{}^2]}} \text{ for } \varepsilon_1 \geq 0, \varepsilon_2 \geq 0$$

$$f(\varepsilon_1, \varepsilon_2) = \frac{\bar{c}_{12}+(\alpha_1\alpha_2)\varepsilon_1\varepsilon_2+\beta_1\beta_2\bar{\rho}_{12}}{\sqrt{[\tilde{c}_{11}+\alpha_1^2+\gamma_1{}^2\,\varepsilon_1{}^2+\beta_1{}^2][\tilde{c}_{22}\alpha_2^2\varepsilon_2{}^2+\beta_2{}^2]}}, \text{ for } \varepsilon_1 \leq 0, \varepsilon_2 \geq 0$$

$$f(\varepsilon_1, \varepsilon_2) = \frac{\bar{c}_{12}+(\alpha_1\alpha_2)\varepsilon_1\varepsilon_2+\beta_1\beta_2\bar{\rho}_{12}}{\sqrt{[\tilde{c}_{11}+\alpha_1^2\varepsilon_1{}^2+\beta_1{}^2][\tilde{c}_{22}\alpha_2^2+\varepsilon_2{}^2+\beta_2{}^2]}}, \text{ for } \varepsilon_1 \leq 0, \varepsilon_2 \leq 0 \qquad (2.15)$$

where ε_i and ε_j are standardised residuals for both markets; c_{ii}, c_{ij} and c_{jj}, are the corresponding elements of the intercept in the correlation equation; *a* and *b* are the diagonal elements of matrices A and B; and $p_{ij,t}$ is the corresponding element of the unconditional correlation matrix. These equations clearly highlight that asymmetry is captured by the differences in slopes, as embodied in $(\alpha_1\alpha_2 + \gamma_1\gamma_2)\varepsilon_1\varepsilon_2$. By comparison, the news impact surface of the ADCCXE model (Li, 2011) is given by:

$$f(\varepsilon_1, \varepsilon_2) = \frac{\bar{c}_{12}+\alpha_1\alpha_2(\varepsilon_1+\gamma_1)(\varepsilon_2+\gamma_2)+\beta_1\beta_2\bar{\rho}_{12}}{\sqrt{[\tilde{c}_{11}+\alpha_1^2(\varepsilon_1+\gamma_1)^2+\beta_1{}^2][\tilde{c}_{22}+\alpha_2^2(\varepsilon_2+\gamma_2)^2+\beta_2{}^2]}}, \qquad (2.16)$$

where the asymmetries are captured via "eccentricity" (i.e., $\alpha_1\alpha_2(\varepsilon_1 + \gamma_1)(\varepsilon_2 + \gamma_2)$).

The ADCCXE model that takes into account two structural breaks both in the long-run mean and the dynamic parameters related to the absolute order flow differential has following evolution for Q_t:

$$Q_t = I[t \leq \tau_1]\left(\overline{\mathrm{Q}}_1 - \mathrm{A}'\overline{\mathrm{Q}}_1\mathrm{A} - \mathrm{B}'\overline{\mathrm{Q}}_1\mathrm{B} - \eta_1\overline{\mathrm{X}}_1\right) + I[\tau_1 \leq t \leq \tau_2]\left(\overline{\mathrm{Q}}_2 - \mathrm{A}'\overline{\mathrm{Q}}_2\mathrm{A} - \mathrm{B}'\overline{\mathrm{Q}}_2\mathrm{B} - \eta_2\overline{\mathrm{X}}_2\right) + I[\tau_2 \leq t]\left(\overline{\mathrm{Q}}_3 - \mathrm{A}'\overline{\mathrm{Q}}_3\mathrm{A} - \mathrm{B}'\overline{\mathrm{Q}}_3\mathrm{B} - \eta_3\overline{\mathrm{X}}_3\right) + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + B'\mathrm{Q}_{t-1}B + I[t \leq \tau_1]\eta_1\left|\mathrm{x}_{1,t-1} - \mathrm{x}_{2,t-1}\right| + I[\tau_1 \leq t \leq \tau_2]\eta_2\left|\mathrm{x}_{1,t-1} - \mathrm{x}_{2,t-1}\right| + I[\tau_2 \leq t]\eta_3|(\mathrm{x}_{1,t-1} - \mathrm{x}_{2,t-1})| \quad , \quad (2.17)$$

where I[A] is the indicator function for Event A ,which can take on the values zero or one, and τ_1 and τ_2 denote the break points , with $\tau_1 < \tau_2$. $\overline{\mathrm{Q}}_1 = {\tau_1}^{-1}\sum_{t=1}^{\tau_1}(\varepsilon_t\varepsilon'_t)$, $\overline{\mathrm{Q}}_2 = (\tau_2 - \tau_1)^{-1}\sum_{t=\tau_1}^{\tau_2}(\varepsilon_t\varepsilon'_t)$ and $\overline{\mathrm{Q}}_3 = (T - \tau_2)^{-1}\sum_{t=\tau_2}^{T}(\varepsilon_t\varepsilon'_t)$. $\overline{\mathrm{X}}_1 = {\tau_1}^{-1}\sum_{t=1}^{\tau_1}(|(\mathrm{of}_{1,t-1} - \mathrm{of}_{2,t-1})|)$, $\overline{\mathrm{X}}_2 = (\tau_2 - \tau_1)^{-1}\sum_{t=\tau_1}^{\tau_2}(|(\mathrm{of}_{1,t-1} - \mathrm{of}_{2,t-1})|)$ and $\overline{\mathrm{Q}}_3 = (T - \tau_2)^{-1}\sum_{t=\tau_2}^{T}|(\mathrm{of}_{1,t-1} - \mathrm{of}_{2,t-1})|$.

The break points are fixed *a priori*, reflecting the start of the 2007 GFC (λ_1) as well as the dramatic worsening of the EDC crisis with the European Union's announcement of providing financial help to Greece in 2010. The models with one structural break either due to the 2007 GFC or 2010 EDC are nested versions of the model given in 2.17.

To evaluate the fit of the structural break model, we use standard likelihood ratio tests, testing the null hypothesis of no structural break (Equation 2.9) against the alternative of two structural breaks (Equation 2.17) and against the alternative of one structural break either due to the 2007 GFC or 2010 EDC.

3. Chapter Three: Forecasting FX Co-Movements via Order Flow

3.1 Introduction

Although the emphasis of the market microstructure approach to exchange rate determination has primarily been on explaining exchange rate movements via order flow (e.g. Evans and Lyons, 2002a) the predictive content of order flows for future exchange rate movements has been studied in much less detail, and neglected the predictive content of order flow for joint exchange rate movements.

This chapter adds to the studies on order-flow-based forecasting by examining how well order flow can forecast the joint variations (co-movements) of exchange rate returns. The intuition is that, although a single order flow explains and possibly predicts single exchange rate movements as a result of the net buying pressure on the currency concerned, the difference between the net buying pressures in two currency markets should determine how both markets move together[28]. The reason why order flow is likely to have a persistent effect on exchange rate fluctuations is the gradual incorporation of order flow information into prices, as suggested by Berger et al. (2008).

Using daily data throughout this chapter, we find that the absolute order flow differential between two currencies – defined as the absolute difference between two individual order flows – is negatively related to future co-variations in returns. These "daily" results and implications are therefore in line with those of the intraday analyses reported in Chapter 2. A new finding from this chapter is that, as far as daily correlation dynamics are concerned, the IRD loses its significance once the order flow differential is accounted for.[29] This finding contributes to the literature in two ways. First, it directly confirms the market

[28] As in Chapter 2, order flow is defined as the net of buyer- and seller-initiated currency transactions.
[29] Because intraday data on interest rates are not available or show no variations within a day, Chapter 2 did not include the IRD in intraday analyses.

microstructure literature which suggests that order flow contains information about economic fundamentals, including changes in interest rates (Evans and Lyons, 2009; Cerrato, Sarantis, and Sauders, 2011). Second, the finding challenges the previous studies which use the interest rate differential as an exogenous determinant of exchange rate correlations (Benediktsdottir and Scotti, 2009; Li, 2011).

More importantly, this chapter shows that order flow can enhance covariance forecasts, as shown by their higher statistical accuracy compared with a variety of competing benchmark models commonly used in the literature. This result applies to all combinations of the three major exchange rates with different out-of-sample periods being evaluated. We have to note that the order-flow-based covariance predictions also generate higher statistical accuracy than any other interest-rate-based covariance prediction, again questioning the value of the IRD for determining co-movement. Again in line with the market microstructure literature, we find that order-flow-based covariance forecasts perform particularly well for the out-of-sample period of the second subsample embracing the 2007 subprime mortgage crisis (the GFC).

As in Chapter 2, we document that all exchange rate pairs exhibit strong positive asymmetry on average across the sample period. In the context of predicting future co-movements, however, positive asymmetry specifications outperform negative asymmetry specifications only during financially tranquil periods in terms of accuracy. By contrast, for financially turbulent time periods, such as the 2007 GFC, negative asymmetry models dominate their counterparts in terms of predictive accuracy. This is in line with the literature, suggesting that during economic downturns, assets start to co-move more closely (e.g. Kasch and Caporin, 2013; Antonakakis, 2012).

The methodology evaluated and the results presented in this chapter are also relevant to the discussions on optimal window size selection (see, for example, Clark and McCracken

(2009)) for return and covariance predictions. In line with the market microstructure literature, we find that a medium-sized in-sample estimation window provides the most accurate co-variance predictions. However, our conclusion is based on a comparison of the performances of different window sizes, unlike most of the literature, which simply assumes that a medium-sized window provides the highest accuracy.[30] We argue that the reason for this is that a medium-sized window combines the benefits (or the reduced disadvantages) of both a short and a long window being used for parameter estimation. Specifically, a short estimation window has a comparative advantage over a longer estimation window by allowing the parameters to be more flexible to recent observations, since parameters are less stable in long window estimations. As such, by choosing a medium-sized window, we get the best balance in the trade-off between having too much and having too few data.

To further highlight the differences between the competing forecasting approaches, we depict the distribution of the forecasting errors incurred by each competing forecasting approach. This graphical representation confirms that the order-flow-based forecast approach has smaller forecasting errors on average than other benchmark models. Furthermore, it also reveals that the order-flow-based covariance predictions result in lower mean square errors during extreme spikes in realised covariance. This result is particularly relevant to two types of investors: investors needing accurate day-by-day covariance predictions, such as mean-variance investors facing frequent rebalance decisions, and investors who are concerned with extreme movements or tail risk. [31]

Alongside from statistical accuracy, we also assess the benefits of the more accurate covariance forecast in a portfolio management application. Specifically, we evaluate the economic value in a classic mean-variance setting in the spirit of Rime et al. (2010) and Della

[30] See, for example, King et al. (2010) or Rime et al. (2010)
[31] The copula approaches, as opposed to the multivariate GARCH approach adopted in this study, however, are more commonly used to evaluate tail risk. An overview of copula approaches can be found in Patton (2009).

Corte et al. (2011). Whereas the two other studies investigated the benefits of order-flow-based return forecasts, we turn to the incremental gain of using order flows to forecast not only returns but also future joint variations in returns. We find that more accurate covariance forecasts lead to substantial trading profits for the mean-variance investor. Thus, risk-averse investors will be willing to pay a substantial performance fee of up to 477 basis points to move from a constant covariance forecasting model to our proposed order-flow-based covariance forecasting model. Note that this setting does not allow for transaction costs, an aspect commonly ignored in the literature.

To address this aspect, we consider a more dynamic trading application that incorporates transaction costs. This setting confirms the benefits of the order-flow-based covariance forecasts. As in the classic mean-variance trading strategy, the order-flow-based covariance forecasts generate substantially higher Sharpe ratios than several other benchmark models commonly used in the literature. However, in contrast to the classic mean-variance setting, Sharpe ratios remain positive even after taking into account transaction costs.

3.2 Related Literature

Apart from the explanatory power of order flow for exchange rate movements[32], previous studies have also highlighted the role of order flow for predicting future exchange rate movements. Arguably, Evans and Lyons (2002a) provided the first "forecast", allowing for a lagged impact of order flow on exchange rate returns. This inter-temporal relationship follows from the argument that the information conveyed by order flows is not instantly, but gradually, incorporated into prices. This gradual incorporation also conforms to the structure of FX markets, where customer order flows (primary demand for a currency) triggers inter-dealer trading (inter-dealer order flows), which then affects spot currency prices. At a

[32] Chapter 2 offers a detailed literature review regarding the explanatory content of order flow, the transmission of news into prices via order flow, the sources of investor heterogeneity and the foreign exchange environment.

theoretical level, Evans and Lyons (2008) formalize the notion that order flow conveys fundamental information about exchange rates, which are not yet fully available to all agents and therefore have an impact on future price movements.

Berger et al. (2008) documented one key mechanism of this gradual incorporation, finding that order flow itself has a high degree of persistence and autocorrelation; that is, a positive (negative) order flow is likely to be followed by a positive (negative) order flow. In other words, buying (selling) pressure is likely to be followed by buying (selling) pressure, thereby leading to a persistent effect on prices. Using brokered inter-dealer data, Berger et al. (2008) confirmed this pattern for intraday frequencies as well as daily and monthly frequencies. Danielsson et al. (2012) quantified the lag in the gradual incorporation from order flow into prices, finding that although most of the content of order flow information is incorporated at intraday frequencies, the "leftover" still affects daily or even weekly spot rate movements. Both articles used order flow data from brokered inter-dealer transactions, the same data source as used in this chapter.

The initial forecasting results from Evans and Lyons (2002a) have been readily confirmed by the literature (Evans and Lyons, 2005b; Gradojevic and Yang, 2006; Marsh and Teng, 2012). Contrary evidence has been provided only by Sager and Taylor (2008), which we discuss in a moment. Whereas all of the aforementioned literature focused on one-step-ahead exchange rate predictions, the study of Evans and Lyons (2006) found that forecasts based on an order flow model outperformed those based on the random walk model at various forecast horizons ranging from 1 to 20 trading days.

A concern related to the study of Evans and Lyons (2005b) and Sager and Taylor (2008) points to the short sample size and the simple forecasting methods employed in those articles. In a short sample, good forecasting performance could be due to chance or the

specific time period used. Cerrato et al. (2015) and Della Corte (2011) joined this discussion with enhanced models and longer datasets. Cerrato et al.(2015) adopted a non-linear approach to predicting future exchange rate movements with current order flows, and found that this approach produces smaller forecast errors than a random walk forecast. Indirectly addressing the concerns in Sager and Taylor's (2008) study, he concludes that the forecasting performance depends on the modelling approach rather than the sample size. By the same token, Della Corte et al. (2011) followed a predictive regression approach using a substantially longer dataset than the datasets used in Evans and Lyons (2005b) and Sager and Taylor (2008), and confirmed the advantages of applying order-flow-based models in predicting future exchange rate movements. In a nutshell, all of the abovementioned literature indicates that order flow indeed transmits expectations about future fundamentals with a lag. This enables one to relate it to co-movement, as noted in Chapter 2.

Note that many previous studies rely on statistical criteria, focusing on the forecasting accuracy while ignoring the economic benefits of order-flow-based return forecasts. Targeting at economic benefits, however, Rime et al. (2010) evaluated the gains for mean-variance investors, using order-flow-based return forecasts for the rebalancing decisions. He finds portfolios formed by using order-flow-based return forecasts yielded substantially higher Sharpe ratios than portfolios formed using macro-economic variables (such as the interest rate) to predict future returns.[33] In line with Rime et al.'s (2010) findings, King et al. (2010) also explored the benefits of employing order-flow-based return forecasts in portfolio optimisation and reported substantial improvements in the Sharpe ratios compared with macroeconomic forecasting models. A key shortcoming of the aforementioned studies, however, is that they do not predict the future joint variations in returns, which is a key aspect

[33] The findings state that a risk-averse investor will be willing to pay a significantly higher performance fee to switch from a random walk forecast to an order-flow-based forecast compared with switching from a random walk forecast to a macroeconomic forecast (labelled 'forward bias' in Rime et al., 2010). The performance fee calculated in this study follows the methodology outlined in Della Corte et al. (2011).

affecting portfolio rebalancing, portfolio stability and profitability (Engle and Colacito, 2006; Della Corte et al. 2012).

Due to the importance of accurate multivariate volatility predictions for portfolio diversification, several different attempts have been made to predict joint variations in returns (see Bauwens, Laurent and Rombouts, 2006, for a detailed review), resulting into a variety of different modelling approaches. The most commonly used models are multivariate GARCH approaches which include the well-known dynamic conditional correlation model (Engle, 2002) and its extensions[34], the constant conditional correlation model (Bollerslev, 1990), and models focusing on covariance dynamics directly, such as the BEKK model (Engle and Kroner, 1995). Other popular modelling approaches include moving average based predictors, such as the exponentially weighted moving average model (Riskmetrics, 1996) and copula approaches (see Patton, 2009, and the references therein).

Yet, it is unclear which of the competing modelling approaches generates the most accurate predictions and largest economic value. The studies of Laurent et al. (2012, 2013) evaluated the statistical accuracy of multivariate GARCH predictions. Laurent et al.'s (2012) study on volatility dynamics between NASDAQ and NYSE stocks compared 125 multivariate GARCH configurations based on 1-, 5-, and 20- day-ahead predictions over a period of 10 years, concluding that covariance predictions generated by dynamic correlation models (DCC) models provide the most accurate predictions during unstable economic conditions. In a subsequent study, Laurent et al. (2013) demonstrated that CCC models provide the most accurate forecasts for intraday exchange rate co-movement.

Targeting economic applications, Wang and Wu (2012) concluded that the co-variance predictions from a BEKK model are the best choice for hedging crude oil. By contrast,

[34] See Engle (2009) for an overview.

Chong's (2005) study on using covariance predictions for option trading, portfolio hedging decisions, and value-at-risk calculations highlighted that an exponentially weighted moving average model (EWMA) provides the largest gains overall. For portfolio optimisation, however, Della Corte et al. (2012) concluded that DCC based covariance predictions generate the largest Sharpe ratios. The contrasting findings suggest that apart from the modelling approach, the intended applications determine the preferred forecasting model.

In contrast to the manifold modelling approaches, the role of exogenous variables in generating accurate covariance predictions has not been evaluated in the literature, a gap this chapter aims to fill.

3.3 Motivation and Testable Hypothesis

The existing studies on order flow and exchange rate determination have focused on single exchange rate movements instead of exchange rate co-movements. Chapter 2 filled this void by modelling and explaining intraday correlation dynamics via order flow. Going on from there, this chapter investigates the predictive content of order flow for future exchange rate covariance dynamics. In accomplishing this task, we use daily data. If order flow can improve the accuracy of covariance forecasts, this will have significant practical relevance to financial applications, such as currency option pricing and correlation timing for investment decisions or portfolio optimisation.

Our main research question is: Can one predict future correlation and covariance based on currently known order flow? As Chapter 2 already pointed out, the motivation for linking order flows to co-movements arises from the market microstructure studies on univariate movements driven by order flows. The studies suggest that order flow represents buying or selling pressure on an asset leading to its price appreciation or depreciation. However, due to the gradual incorporation of order flow information into prices, not all order flow information

is incorporated into spot rates instantly. In other words, part of the information in current order flows will be transmitted into prices with a substantial lag, so current order flows ought to predict future exchange rate movements.

As for linking order flows to future correlation dynamics, the motivation is similar to that presented in Chapter 2. Here, we restate it, albeit in a more intuitive way. Consider a positive order flow (an excess of buyer-initiated trades over seller-initiated trades) for a currency at its prevailing exchange rate. This will cause an increase in demand and an appreciation of the exchange rate. However, not all of the information contained in the positive order flow will be incorporated in the price of a currency immediately, as the markets learn the details of and reasons behind the net demand gradually. Therefore, part of the information is incorporated into prices with a significant lag, implying that current order flows affect future movements. The hypothesis on how current order flow affects future co-movements follows the same argument as mentioned in Chapter 2. If two currencies experience a similar buying or selling pressure, their exchange rates against a world currency (say, USD) are likely to move in the same direction. However, in the case where the order flows are of different signs, this will reduce the probability of both currencies rolling in the same direction, thereby reducing co-movements. Going further from here, if the buying or selling pressure is incorporated over time into the spot rates, then the current differences in buying or selling pressures will influence the future correlation of the two exchange rates. Based on this analysis, our first two testable hypotheses are therefore:

H1: Absolute differences in current order flows will negatively affect future co-movements of exchange rates.

H2: Incorporating individual order flows into multivariate forecasting models will improve their forecasting accuracy relative to the standard multivariate GARCH approaches.

Note that although the first hypothesis follows Hypothesis 1 in Chapter 2, the underlying research question is different. In this chapter, we investigate the out-of-sample forecasting ability of current order flows for future covariance dynamics at the daily frequency, whereas in Chapter 2, we studied how differences in order flows affect in-sample correlation dynamics at the intraday frequency. The second hypothesis is based on the intuition given above (i.e., differences in order flows between currencies predict differences in their respective price movements), an argument following from univariate market microstructure, which suggests that order flow enhances return predictions.

Another finding in the market microstructure literature is that order-flow-based return forecasts dominate interest-rate-based return forecasts in terms of accuracy (Evans and Lyons, 2005b). The theoretical justification for this is that order flows contain all the dispersed information available in the market, including the information content of interest rates. This is an interesting aspect, as several previous studies have investigated the effect of the IRD on correlation dynamics. However, we maintain that order flow is more powerful for predicting future exchange rate co-movements, as order flow contains all the information about macroeconomic fundamentals, including changes in the interest rate of a currency.[35] In other words, multivariate GARCH models incorporating IRDs will yield less accurate predictions than their counterparts using order flow. Therefore, the third testable hypothesis is:

H3: Any interest-rate-based covariance forecasting model will underperform compared with the order-flow-based multivariate forecasting models in terms of forecasting accuracy.

This hypothesis, if supported by the data, implies that interest rates, though widely used and deemed to be a very important macroeconomic variable, contain less information about economic fundamentals than order flows.

[35] The reason for this is that any changes in the interest rates (and hence in the profitability of carry trades) would result in investors adjusting their portfolios, hence changing order flow, which then transmits information about changes in IRDs into prices via order flows.

The last relevant question arising from the market microstructure literature is whether order flows can be used to predict volatilities directly, as volatilities, in turn, affect covariance dynamics (See Chapter 2 for more detail). Opschoor, Taylor, Wel, and Dijk (2014) argued that larger order flows predict increases in future volatility, because volatility does not mean-revert immediately after news shocks. The news shocks are transmitted via order flows thereby affecting the volatility of exchange rates. Thus the link between order flows and exchange rate volatilities actually reflects the indirect link between news shocks and exchange rate volatilities. This chapter evaluates this finding for out-of-sample forecasting purposes. Based on the frequent spikes in conditional volatility attributable to order flow documented in Chapter 2, we conjecture that order-flow-based volatility predictions will enhance covariance predictions during financially turbulent periods. The resulting hypothesis is:

H4: Using order flows as an exogenous variable in a GARCH-X model will provide more accurate volatility forecasts than a standard GARCH model, particularly when the underlying financial time period is characterised by frequent spikes in realised volatility.

The GARCH-X model is the standard GARCH model augmented with an exogenous variable, which is absolute order flow in our case. If order flows generate more accurate forecasts of the entire covariance matrix of returns, we will deem H2 and H4 to be valid. In this case, we are able to claim that order flows not only help better forecast returns, as already seen in previous studies, but also help forecast variance, correlation and covariance more accurately, which is a novel contribution to the existing literature.

In addition to showing the higher statistical accuracy of our order-flow-based approach, we also investigate the economic gains of this approach for predicting future covariance series. Although statistical accuracy allows us to compare the competing covariance forecasts in terms of MSE, forecasts with smaller MSEs do not necessarily result

in higher economic values for an investor. For this reason, we attempt to verify that our approach is not only statistically more accurate but also economically more valuable. We consider an indirect evaluation of models, focusing on a mean-variance investor who uses order flows to predict future returns and joint variations of returns for determining the optimal investment weights. We conjecture that a more accurate covariance forecast should provide more economic benefits to a mean-variance investor, since an accurate covariance forecast is the key input in the classic mean-variance portfolio optimisation problem. We thus test the following hypothesis:

H5: More accurate order-flow-based, covariance predictions enhance the performance of a portfolio management strategy that relies on covariance predictions.

We employ the mean-variance setting for the following two reasons. First, several studies (e.g., Della Corte et al., 2012; Zakamulin, 2015) have evaluated covariance forecasts via mean-variance optimisation, so our results are comparable to theirs. Second, recent studies on using order flow for mean-variance portfolio optimisation (e.g., King et al., 2010) concern themselves only with forecasting univariate return movements. Thus using the same setting makes our multivariate results comparable with, and, more importantly, an addition to their univariate results in terms of the economic gains for mean-variance investors.

3.4 Data and Methodology

3.4.1 Data

The empirical work in this chapter continues to use data on the exchange rates of AUD, NZD, CAD, EUR, GBP and JPY, all quoted in USD (i.e., the USD price of one unit of a currency). Thus the change rate in a spot rate measures the rate of returns to an investor purchasing one unit of a currency with USD and then selling the currency for USD. We model the bivariate return correlations (i.e., the correlations between the returns on two exchange rates), as we did in Chapter 2. However, the data frequency used in this chapter is daily, not intraday, as in Chapter 2.

Daily order flow is calculated as the difference between buyer- and seller-initiated transactions[36] aggregated over the most active part of a trading day between 7:00 and 17:00 GMT. Hence, a positive (negative) order flow represents a net buying (selling) pressure in a trading day. The daily rate of return on an exchange rate is measured as the difference between the log mid-quote at 17:00 GMT of day *t* and 17:00 GMT of day *t* -1. As in Chapter 2, we excluded weekends and some world–wide public holidays.

Data on interest rates (3-month interbank interest rates) are obtained from the Thomson Reuters DataStream. Our realised volatility measures are constructed by using intraday observations from TRTH: we use bid and ask quotes to calculate the intraday returns, and intraday variance and covariance. The methodology for the construction of the realised volatility measures is outlined in the next section. The sample period for the empirical investigation spans from 3 January 2002 to 29 December 2013, a total of 3112 observations for each exchange rate.[37]

[36] Note: A seller-initiated trade origins from exchanging (selling) the non-USD currency for (buying) USD. A buyer-initiated trade is defined the other way around.
[37] Table 3.1 provides descriptive statistics for the return and order flow series.

3.4.2 Methodology

We employ the same ADCCXS model using the absolute order flow differential as an additional input affecting correlation dynamics as in Chapter 2. However, the aim here is to evaluate the predictive content of our proposed model, which uses the order flow differential as an exogenous input. Our initial analysis starts with a comparison of the ADCCXS estimates for daily return co-movement and the intraday estimates reported in Chapter 2. Furthermore, we evaluate whether interest rates affect co-movement additionally to order flow.

To test the predictive power of current order flow for future joint return movements, we forecast one-day-ahead covariance series for the three major cross-rates (EUR– GBP, EUR–JPY and JPY–GBP), using an ensemble of competing covariance forecasting approaches. The competing models are based on the ADCCXS approaches captured in Equations 2.7 and 2.8, as well as its nested versions (ADCC, DCC and CCC). Furthermore, because of its popularity, we evaluate the performance of the risk-metrics moving average (EWMA) as well as a BEKK type covariance prediction. Last, we evaluate the predictive performance of a variant of the ADCCXS model that uses the IRD as a driving force for correlation dynamics, which again follows Equations (2.7) and (2.8). To avoid possible confusion with notation, we denote the ADCCXS taking into account the absolute order flow differential as ADCCXS–OF and the version taking into account the interest rate differential as ADCCXS–IR. The competing models are presented in detail in Section B.2 of Appendix B.

To measure out-of-sample forecasting performance, we compare the model based-forecasts to ex post realisations as they become available. To do so, we need to select a loss function and a proxy for the true conditional covariance. Following Laurent et al. (2012), we

use an estimate for the realized covariance[38] as a proxy for the "true" covariance. The construction of the realised covariance follows Andersen, Bollerslev, Diebold, and Labys (2003), using ad-hoc calendar-time sampling, described as below:

Let $s_t = \left[s_{1,t}, s_{2,t}\right]$ be a 2×1 vector of the log prices of 2 currencies at time *t*, where each currency's log price is aligned to a 5-minute interval. As trading takes place 24 hours a day in the FX market, there are 288 5-minute intervals per trading day. For notational simplicity, let Δ denote the 5-minute interval, and $m = 1/\Delta$ the number of samples taken per day. The 2×2 realized covariance matrix (referred to as RCM) of daily returns[39] on day *t* is given as:

$$RCM_t = \sum_{j=1}^{m} r_{t-1+j\Delta}\, r_{t-1+j\Delta}{}', \tag{3.1}$$

where $r_{t+\Delta} = s_{t+\Delta} - s_t$

The loss function used evaluates the mean squared error between the realized covariance and the model based prediction. This allows ranking the competing models directly, based on the criterion of the lowest loss. In this paper we employ a robust matrix-wise multivariate MSE (denoted MMSE) loss function[40], given by:

$$\mathrm{Lf}_{\mathrm{M}_{u,}\mathrm{t}}^{\mathrm{MMSE}} = \mathrm{Tr}[\left(RCM_t - \mathrm{H}_t^{M_u}\right)'\left(RCM_t - \mathrm{H}_t^{M_u}\right)] ,$$

$$\bar{L}f_{\mathrm{M}_{u,}\mathrm{t}}^{MMSE} = \frac{1}{\mathrm{T}}\sum_{\mathrm{t}=1}^{\mathrm{T}} \mathrm{Lf}_{\mathrm{M}_u,\mathrm{t}}^{\mathrm{MMSE}} , \tag{3.3}$$

where $\mathrm{H}_t^{M_u}$ represents the predicted covariance matrix produced by model M_u with M_u =[ADCCXS, ADCC, DCC, CCC, BEKK and EWMA] and $\hat{\Sigma}_t$ represents the realised covariance matrix.

[38] Apart from realized covariance, the literature outlines several other proxies for the "true" conditional co-variance, such as bi-power co-variation (Barndorff-Nielsen and Sheppard, 2004) or moving- average-based estimators for realized covariance (Hansen, Large, and Lunde, 2008).
[39] This is the end-of-day return from the end of 17:00 GMT at day t − 1 until 17:00 GMT day t.
[40] We choose this particular loss function due to its robustness to the proxy of the true covariance matrix (Patton, 2011). We are aware that this particular version double counts the loss associated to the conditional covariance.

The UMSE related to the covariance prediction only, is given by:

$$Lf_{\mathrm{M}_{u,}\mathrm{t}}^{UMSE} = (RCM_{i,j,t} - h_{ijt}^{Mu})^2$$

$$\bar{L}f_{\mathrm{M}_{u,}\mathrm{t}}^{UMSE} = \frac{1}{T}\sum_{t=1}^{T} Lf_{\mathrm{M}_{u,}\mathrm{t}}^{UMSE} \quad (3.4)$$

where $RC_{i,j,t}$ denotes the realized covariance taken from the off-diagonal elements in RC_t and h_{ijt}^{Mu} denotes the predicted covariance from one of the competing models.

In order to assess the relative accuracy of each of the competing forecasts, we implement Diebold–Mariano (DM) statistics (Diebold and Mariano, 1995). The DM test compares the predictive powers of competing covariance forecasting models. The test is based on the loss differential between two competing forecasts and is given by: $d_{\mathrm{M}_{u,},\mathrm{M}_{v,}t} = LF_{\mathrm{M}_{u,t}} - LF_{\mathrm{M}_{v,t}}, for\ v \neq u$. The null hypothesis of equal predictive accuracy tested is $H_0: E\left[d_{\mathrm{M}_{u,},\mathrm{M}_{v,}t}\right] = 0$. The test statistic is:[41]

$$S = \frac{\bar{d}}{a\ var(\bar{d})^{1/2}}. \quad (3.5)$$

For the initial analysis presented in Sections 3.6.1 to 3.6.3 we will rank the competing models based on the univariate mean squared error (UMSE), i.e. on their losses resulting from covariance predictions. In Section 3.6.4, where we compare the predictive accuracy of the GARCH-X and GARCH-based class of dynamic conditional correlation models, however, we focus on the matrix-wise multivariate MSE.

[41] *S* is asymptotically normally distributed and hence can be interpreted in the same way as a *t*-statistic.

3.5 Empirical Results

3.5.1 Descriptive statistics

Table 3.1 shows the descriptive statistics for the daily exchange rates investigated. As in Chapter 2, the exchange rate is defined as the USD price of one unit of a currency. This means that a deprecation in USD relative to a currency results in a positive return on holding one unit of the currency such as EUR, GBP or JPY. As in the case of intraday returns, the daily mean returns are small but positive for all six exchange rates. A comparison of the daily descriptive statistics with the intraday descriptive statistics presented Chapter 2 indicates that temporary intraday exchange rate shocks are smoothed out during a trading day, as shown by the lower skewness and kurtosis statistics for daily rather than intraday exchange rate returns.[42]

The positive mean returns to the six currencies can be linked to the positive means of the order flows (see Panel B) for each currency. A positive mean order flow implies an average net demand for a non-USD currency over the entire sample, causing USD to deprecate against the currency overall. Panel B of Table 3.1 shows that the mean order flow differentials are much larger than those with intraday order flows (referring back to Chapter 2), which is caused by the aggregation of small but positive intraday order flows. For the analyses following in the reminder of the chapter, we measure order flow in units of a thousand.

[42] However, although the skewness and kurtosis statistics are smaller than those for intraday returns, the Jarque–Bera statistic still rejects the null hypothesis of a normal distribution.

Table 3.1 Descriptive statistics

	USD/EUR	USD/GBP	USD/JPY	USD/AUD	USD/NZD	USD/CAD
Panel A. Descriptive statistics for daily returns						
Maximum	0.046	0.045	0.037	0.067	0.058	0.042
Minimum	-0.038	-0.039	-0.046	-0.088	-0.088	-0.038
Mean	0.013	0.004	0.007	0.017	0.021	0.013
Std.Dev.	0.46	0.53	0.57	0.87	0.89	0.53
Skewness	0.08	-0.07	0.26	-0.96	-0.47	0.12
Kurtosis	5.81	7.33	7.05	15.99	8.40	7.01
JB-Prob.	0.00	0.00	0.00	0.00	0.00	0.00
Panel B. Descriptive statistics for daily order flows						
Maximum	478	990	152	986	447	986
Minimum	-453	-765	-176	-1196	-496	-829
Mean	15.39	35.19	0.26	14.35	9.88	46.10
Std.Dev.	110.6	213.3	25.9	233.9	85.08	186.58
Skewness	0.09	0.05	0.04	-0.43	-0.18	0.24
Kurtosis	1.14	1.41	5.79	1.87	2.42	1.55

Note: This table shows the descriptive statistics for 5-minute- returns (Panel A) and order flows (Panel B) for six exchange rates. Daily returns are calculated as the difference between the log mid-quote at 17:00 GMT of day *t* and 17:00 GMT of day *t* -1. Daily order flows are aggregated tick-by-tick between 7:00 and 17:00 GMT. The means and standard deviations for exchange rate returns are expressed in percentage points.

3.5.2 Comparison of daily and intraday correlation dynamics

The estimation results for the ADCCXS model targeting negative type asymmetry on daily data are presented in Table 3.2.[43] Akin to the intraday analysis reported in Chapter 2, all of the GARCH parameters are statistically significant at the 1% level and their sums are close to unity, indicating a high degree of persistence in volatility. Furthermore, the modified Box–Pierce Q-statistics for ε_{1t}, ε_{2t}, ε_{1t}^2 and ε_{2t}^2 for all exchange rates are statistically insignificant at the 5% level, implying that the standardised residuals are i.i.d.

Turning to the second-stage estimation – the estimation of the DCC parameters – we can see that for all six exchange rate pairs, the α and β parameters for, respectively, shocks and persistence are statistically significant. As for the intraday analysis presented in Chapter

[43] As in Chapter 2 we follow Cappiello et al. 2006 in the initial parameterization of the ADCCXS, setting $n_t = \varepsilon_t^-$

2, Correlation persistence is high in several pairs as reflected by the large estimates of β_1 and β_2, such as 0.9773 and 0.9735 for EUR–GBP, 0.9868 and 0.9630 for EUR–JPY and 0.9795 and 0.9810 for GBP-JPY.

Another similarity to the intraday analysis is that the estimates related to negative asymmetry are very small and sometimes insignificant, consistent with the results of the intraday analyses in Chapter 2. For the effect of order flow differentials on correlation, the sign of the parameter estimates (η) is consistently and significantly negative across seven out of eight pairs, but not for AUD–NZD. Thus, for the seven pairs, a rise in the absolute order flow differential will reduce their future daily co-movements, supporting our hypothesis H1. For the AUD–NZD pair, a higher order flow differential increases its future co-movements, contrary to H1. The possible reason for this was already mentioned in Chapter 2, citing Smyth (2009) . Note, however, that our research objective here is to investigate the predictive content of order flow for future co-movements, not the sign of the eta parameter.

As in Chapter 2, we test for positive type asymmetry, which can dominate in FX markets, by using both the ADCCXS (using joint positive returns to estimate the asymmetry parameters) and ADCCXE models described in Chapter 2. The results are reported in Table B.2 in Appendix B. In a nutshell, our findings indicate the same conclusions as drawn those for intraday co-movements and are therefore in line with the results reported in Li (2011). We find that both negative- and positive type asymmetry are preferred over the null hypothesis of symmetry when employing the ADCCXS specification. The ADCCXE model focusing on eccentricity as a measure of asymmetry confirms, however, that positive asymmetry dominates negative asymmetry for all exchange rate combinations. As the focus of this study is on predicting correlation dynamics, we do not discuss the differences in the parameter estimates in detail, but rather evaluate which type of asymmetry results in more accurate out-of-sample predictions.

Table 3.2 Parameter estimates of the GARCH-ADCCXS model

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
GARCH model parameters								
ω_1	0.0000 **	0.0000 **	0.0013 **	0.0024 **	0.0024 **	0.0059 **	0.0058 **	0.0059 **
δ_1	0.0309 **	0.0309 **	0.0311 **	0.0448 **	0.0448 **	0.0367 **	0.0594 **	0.0367 **
θ_1	0.9663 **	0.9663 **	0.9656 **	0.9483 **	0.9483 **	0.9499 **	0.9325 **	0.9499 **
ω_2	0.0000 **	0.0000 **	0.0019 **	0.0059 **	0.0019 **	0.0058 **	0.0074 **	0.0019 **
δ_2	0.0439 **	0.0347 **	0.0486 **	0.0366 **	0.0486 **	0.0594 **	0.0442 **	0.0485 **
θ_2	0.9507 **	0.9544 **	0.9466 **	0.9500 **	0.9466 **	0.9325 **	0.9459 **	0.9467 **
ADCCXS model parameters								
α_1	0.1904 **	0.1518 **	0.1492 **	0.1989 **	0.1477 **	0.2512 **	0.1733 **	0.1897 **
α_2	0.1173 **	0.2496 **	0.0836 **	0.1517 **	0.0931 **	0.2282 **	0.1421 **	0.1611 **
β_1	0.9733 **	0.9868 **	0.9885 **	0.9795 **	0.9880 **	0.9648 **	0.9795 **	0.9713 **
β_2	0.9735 **	0.9630 **	0.9936 **	0.9810 **	0.9916 **	0.9636 **	0.9861 **	0.9858 **
γ_1	0.0342 **	0.0568 **	0.0000	0.0316 **	0.0000	0.0821 **	0.0030	0.0766 **
γ_2	0.0460 **	0.0464 **	0.0001	0.0330 **	0.0000	0.0394 **	0.0009	0.0441 **
η	-0.0124 **	-0.0180 **	-0.0164 *	-0.0372 **	-0.0044 *	-0.0279 **	0.0078 **	-0.0054 **
$\bar{\rho}$	0.699	0.286	0.489	0.239	0.355	0.154	0.853	0.098
Statistics								
LRT	15.27 **	10.271 ***	1.062	7.642 *	8.092 *	3.504	14.939 ***	8.236 *
Q_1^*	11.78	11.781	22.461	17.89	17.889	18.655	8.236	35.037
Q_{11}^*	26.68	26.688	18.359	13.257	13.257	40.217	35.037	11.781
Q_2^*	17.88	11.07	11.781	11.07	18.654	11.07	11.07	26.688
Q_{22}^*	3.25	10.03	26.688	10.35	40.217	10.3	10.3	39.013
$\chi^2 - \text{test}$	77.03	197.055 ***	46.547 ***	270.752 ***	2.783	279.070 ***	248.902 ***	57.837 ***
LLF	-4174.85	-5463.45	-4895.42	-5308.65	-4764.6	-6190.85	-5354.63	-5468.23

Note: This table presents the parameter estimates of Equations (2.7) and (2.8) for eight possible combinations of exchange rate pairs using daily return and order flow intervals. The estimation (and computation of the modified standard errors) uses the two-stage procedure proposed in Engle and Sheppard (2001). LRT denotes the likelihood ratio test statistic testing the null hypothesis of γ_1=γ_2 =0. LLF is the quasi-log-likelihood function evaluated at the maximum. ** and * denote significance at the 1% level and 5% level, respectively. Q^* represents the Box–Pierce Q-statistic adjusted for heteroskedasticity (see Lobato et al., 2001) . Q_1^*, Q_{11}^*, Q_2^* , Q_{22}^* are associated with $\varepsilon_{1,t}$ ${\varepsilon^2}_{1,t}$, $\varepsilon_{2,t}$ and ${\varepsilon^2}_{2,t}$ respectively. The χ^2-test proposed by Engle and Sheppard tests the null hypothesis of constant conditional correlation during the sample period ($R_t = \bar{R}$) against the alternative of a dynamic conditional correlation ($vech\ (R_t) = vech(\bar{R}) + \iota_1\ vech(R_{t-1}) + \iota_p\ (vech(R_{t-p}))$.

3.5.3 The role of the interest rate differential (IRD)

Previous studies suggest that, at the daily frequency, higher IRDs lower future exchange rate co-movements (see, e.g., Li, 2011). However, the market microstructure literature argues that order flow already contains all the information about macroeconomic fundamentals, including changes in the interest rate of a currency.[44] Therefore, the interest rate variable should not have a significant effect on dynamic conditional correlations in addition to the effect of order flow.

To test this hypothesis, we use the ADCCXE model given by Equation (2.10), replacing the bid–ask spread with the absolute IRDs. The absolute IRDs are constructed as the absolute difference between the 3-month interbank interest rates of two economies other than the US.

The test results are provided in Table 3.3. One can see that controlling for the absolute IRD as an additional variable does not qualitatively alter the η_1 parameter estimates associated with the order flow differential: the η_1 estimates remain significantly negative, ranging between -0.046 (GBP–JPY) and -0.0045 (GBP–CAD), again with the exception of the AUD–NZD pair, where η_1 is positive (0.0297) but insignificant. On the other hand, the η_2 parameter estimates on the IRD are not statistically different from zero for several cross-rates (EUR–GBP, GBP–JPY, GBP–CAD and JPY–CAD). This means that the IRD has no effect on their conditional correlations once the order flow differential is allowed for. In addition, although the η_2 parameter estimate for the EUR–JPY pair is significant, its sign is wrong (0.010), inconsistent with the theoretical prediction. In other words, statistical significance does not necessarily guarantee economic plausibility.

[44] The reason for this is that any changes in the interest rates (and hence in the profitability of carry trades) would result in investors adjusting their portfolios, hence changing order flow, which then transmits information about changes in IRDs into prices via order flows.

In general, these results are in line with the predictions of the market microstructure literature[45]. They also raise doubt about the usefulness of a macroeconomic variable such as the interest rate for predicting future exchange rate correlations. We evaluate this doubt in the next section in detail.

[45] As in Chapter 2, we evaluate whether changes in the absolute bid–ask spread affect the parameter estimates relating to the order flow differential. We report the corresponding estimates in Appendix B. In a nutshell, this analysis further confirms the robustness of the documented negative relationship between the absolute order flow differential and exchange rate co-movement, even after changes in the bid–ask spread are accounted for, with similar patterns found to those in the discussion presented in section 3.5.2.

Table 3.3 Parameter estimates of the GARCH-ADCCXE model (with the IRDs variable added)

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
Panel A. ADCCXE (2.11) estimates – absolute order flow differential (η_1) and IRD (η_2)								
α_1	0.2025 **	0.1686 **	0.1554 **	0.1965 **	0.1477 **	0.3011 **	0.1786 **	0.1996 **
α_2	0.1165 **	0.2154 **	0.0823 **	0.1518 **	0.0959 **	0.1911 **	0.1283 **	0.1504 **
β_1	0.9787 **	0.9842 **	0.9878 **	0.9779 **	0.9878 **	0.9536 **	0.9760 **	0.9683 **
β_2	0.9871 **	0.9645 **	0.9939 **	0.9810 **	0.9909 **	0.9761 **	0.9893 **	0.9886 **
γ_1	0.5646 **	0.0535 **	0.1306 **	0.0711 **	0.0013 **	0.2576 **	-0.1909 **	0.1869 **
γ_2	0.3132 **	0.1528 **	0.0006	0.0726 **	0.0124 **	0.0865 **	-0.1227 **	0.1362 **
η_1	-0.0111 **	-0.0344 **	-0.0073 **	-0.0465 **	-0.0045 **	-0.0288 **	0.0297	-0.0081 **
η_2	0.0036	0.0103 **	-0.0002 **	-0.0001	0.0001	-0.0016 **	-0.0095 *	0.0005
LLF	-4172.5	-5418.6	-4886.7	-5308.1	-4757.5	-6151.9	-5352.1	-5459.6

This table presents the parameter estimates of Equation 2.11 for eight possible combinations of exchange rate pairs using daily return and order flow intervals. The estimation (and computation of the modified standard errors) uses the two-stage procedure proposed in Engle and Sheppard (2001). The parameter η_1 is associated with the absolute order flow differential, whereas η_2 corresponds to the interest rate differential. LLF is the quasi-log-likelihood function evaluated at the maximum. ** and * denote significance at the 1% level and 5% levels respectively.

3.6 Forecasting

The forecast methodology is as follows. We split the whole sample period into three sub periods, each containing a total of 1000 observations. Out of each subsample, the first 800 observations are used as in-sample observations, whereas the remaining 200 observations represent the out-of-sample observations used to evaluate the forecasting performances of each competing model. Dividing the total sample into three subsamples of equal length is justified by the following consideration: the first subperiod corresponds to the relatively calm period before the 2007 GFC, the second covers the GFC period prior to the EDC, and the third corresponds to the recovery period of the US but the deteriorating period of Europe.[46] [47]

The initial estimation window consists of the 800 observations closest to the out-of-sample period. We then predict one-step-ahead[48] covariance of two exchange rate returns for the remaining 200 observations. Throughout this exercise, we re-estimate the model parameters for all competing models after every 25 one-day-ahead covariance forecasts using a rolling window of 500 observations. During the 25-day window between the re-estimations, the parameters are kept fixed and only the data are updated.

3.6.1 Statistical accuracy

Table 3.4 outlines the univariate mean squared errors (UMSEs) from each competing forecasting model for three distinctively different out of sample periods (Panel A, Panel B and Panel C). The UMSE represent the squared difference between the forecasted and the realised covariance. Next to the UMSEs are the corresponding DM statistics, testing whether

[46] The choice of 200 out-of-sample observations to evaluate predictive accuracy is based on the portfolio optimisation setting used later on. As Bailey and Lopez de Prado (2014) indicate, short track records tend to inflate the Sharpe Ratios. To avoid this problem, we evaluate a track record of 200 out-of-sample observations and use this out-of-sample size throughout this chapter to make the analysis coherent and consistent.
[47] The three distinctive periods evaluated are 01/2002–12/2005, 01/2006–12/2009 and 01/2010–12/2013.
[48] The reason for forecasting one-day-ahead rather than several-days- ahead is that a longer forecast horizon would bias the models towards converging to unconditional correlation.

the differences in the UMSE between each competing forecast model and the benchmark ADCCXS–OF model are statistically different from zero. A negative value for the DM statistic will indicate that the ADCCXS–OF is more accurate in terms of "loss" (i.e., in terms of the UMSEs between the predicted and the realised covariance). Initially, we use the parameterisation for negative- type asymmetry via the change in slopes for the ADCC class of models.

For the out-of-sample period of the first subsample (Panel A), the ADCCXS–OF provides the lowest UMSE and hence its predictions are closest to realised covariance for the EUR–GBP, EUR–JPY and GBP–JPY exchange rates. The differences among the forecasting approaches are substantial: For the EUR–GBP joint return predictions, the order-flow-based covariance prediction results in a UMSE of 1.24e-10, which is significantly lower than its next competitor, the BEKK model, with an UMSE of 1.31e-10. The same holds for the EUR–JPY and GBP–JPY pairs where the UMSEs of the ADCCXS–OF are equal to 5.82e-11 and 1.51e-10, and the closest competing models have UMSEs of 6.45e-11 (ADCCS) and 1.56e-10 (BEKK). Looking at the ADCCXS and its nested variants, the ADCCXS–OF significantly outperforms its nested variants, including the ADCCXS–IR for all three exchange rate pairs in terms of forecasting accuracy during this subsample (1.24e-10 vs. 1.33e-10 for EUR–GBP, 5.82e-11 vs. 8.14e-11 for EUR–JPY and 1.51e-10 vs. 1.92e-10 for GBP–JPY).

An interesting question is whether the first subsample's results will hold for a more turbulent financial period, as the first subsample is characterised by small fluctuations in the daily realised volatilities. We conjecture that the order-flow-based forecasts will perform better than the other competing models for the more turbulent period. This conjecture is based on the argument of Rime and Tranvåg (2012) that during financial meltdowns, order flows are more informative than other macroeconomic and price variables, as the latter will

lose part of their information content. However, there is another argument: during an economic downturn, frequent spikes characterise the realised volatility, so it is likely that multivariate models in general will suffer from large imprecision. This seems to suggest that multivariate models would not work well in forecasting for turbulent periods.

Panel B in Table 3.4 sets out the forecasting results for the second sub-period, which includes the GFC. As conjectured above, the differences between the order-flow-based dynamic conditional correlation predictions and its competitors become larger during this period, highlighting the incremental value of order flow information in predicting future joint movements during an economic downturn.

Looking at the individual results, the ADCCXS–OF forecasts are advantageous over the forecasts of its nested dynamic versions for the EUR–GBP, EUR–JPY and GBP–JPY joint return variations. For all of these pairs, the UMSEs from the ADCCXS–OF (7.47e-10, 9.13e-10 and 1.00e-09) are considerably lower than the losses from its nested variants, such as the ADCCS (7.61e-10, 1.09e-09 and 1.20e-09) and the DCC (7.90e-10, 1.57e-09, 1.30e-09). Furthermore, the ADCCXS–OF predictions are much more accurate than the competing BEKK (9.72e-10 and 1.69e-09) and EWMA (2.10e-09 and 9.32e-10) forecasts for the EUR–GBP and EUR–JPY covariance forecasts. Note, however, that all multivariate models encounter a larger loss than the relatively tranquil first out-of-sample period, indicating that the volatility predictions failed to predict the frequent spikes in volatility present during this sub-period [49].

Comparing the interest-rate-based covariance forecasting model (i.e., ADCCXS–IR) against the order-flow-based forecasting model (i.e., ADCCXS–OF) again highlights the improved information content of order flows over the IRD in predicting future correlations. For all

[49] This finding provides the first intuition for the GARCH-X approach evaluated later on.

three exchange rate pairs, the UMSEs of the ADCCXS–OF model are significantly smaller than the predictive errors of the ADCCXS–IR model, with the difference being the largest for the EUR–JPY covariance forecasts (9.13e-10 vs. 3.92e-09) during the second out-of-sample period. Note that the loss generated by the ADCCXS–OF is higher than the losses incurred by the CCC and EWMA for GBP–JPY covariance predictions for this subsamples out-of-sample period.

Panel C shows that in the third period, the differences among the competing forecasting models become much smaller and that the ADCCXS–OF only significantly outperforms its nested dynamic variants for the EUR–JPY covariance forecasts. For this exchange rate pair, the ADCCXS–OF model has an error of 1.11e-10, which is substantially lower than the errors of the ADCCS (1.26e-10), DCC (1.26e-10) and ADCCXS–IR (1.36e-10) models. However, the ADCCXS–OF model generates a larger UMSE than the CCC (9.64e-11) and BEKK (9.49e-11) model based predictions.

For the EUR–GBP covariance predictions, the accuracy of the ADCCXS–OF approach (1.59e-10) is quickly followed by that of the ADCCXS–IR (1.61e-10), ADCCS (1.64e-10) and DCC (1.62e-10) models during this period. The same pattern applies to the GBP–JPY covariance predictions, where the accuracy of the ADCCXS–OF model is not statistically different from that of the ADCCXS–IR (4.02e-10) or CCC (4.00-10) models, suggesting that the increase in benefits of using order flows is lower in the third subsample, although the ADCCXS–OF approach still has a lower UMSEs than the nested dynamic models for all three exchange rate pairs, consistent with the findings from the first and the second subsamples. Note, however, that during the third period, volatility decreased substantially. This decrease in volatility leads to overestimations of the conditional covariance in each forecasting model. We will show later on (Section 3.6.4) that a GARCH–X model can

capture the underlying volatilities better for the second and third subsample period, giving rise to a more accurate prediction of the entire covariance matrix.

Table 3.4 Out-of-sample UMSEs for competing covariance forecasts

	EUR–GBP		EUR–JPY		GBP–JPY	
Panel A. January 2002 – December 2005						
	UMSE	DM	UMSE	DM	UMSE	DM
ADCCXS–OF	1.24E-10		5.82E-11		1.51E-10	
ADCCXS–IR	1.33E-10***	-2.52	8.14E-11***	-3.74	1.92E-10 ***	-5.54
ADCCS	3.49E-10***	-8.58	6.45E-11*	-1.58	1.76E-10 ***	-3.38
DCC	1.32E-10**	-1.83	6.54E-11**	-1.78	1.75E-10 ***	-3.25
CCC	1.33E-10**	-1.95	8.11E-11***	-4.41	1.67E-10**	-2.16
EWMA	1.47E-10***	-3.24	1.73E-10***	-8.99	2.32E-10***	-4.52
BEKK	1.31E-10**	-2.14	7.63E-11***	-3.65	1.56E-10	-0.68
Panel B. January 2006 – December 2009						
	UMSE	DM	UMSE	DM	UMSE	DM
ADCCXS–OF	7.47E-10		9.13E-10		1.00E-09	
ADCCXS–IR	7.89e-10 **	-2.25	3.92E-09***	-4.59	1.33E-09 ***	-3.44
ADCCS	7.61E-10 ***	-5.09	1.09E-09 ***	-2.46	1.20E-09 ***	-2.96
DCC	7.90E-10 ***	-3.37	1.57E-09***	-2.83	1.30E-09 ***	-3.89
CCC	9.38E-10 ***	-4.89	9.21E-10	-1.24	9.79E-10	1.15
EWMA	2.10E-09 ***	-10.43	9.32E-10*	-1.70	9.85E-10	1.27
BEKK	9.72E-10 ***	-4.43	1.69E-09***	-3.41	1.59E-09 ***	-3.30
Panel C. January 2010 – December 2013						
	UMSE	DM	UMSE	DM	UMSE	DM
ADCCXS–OF	1.59E-10		1.11e-10		3.85E-10	
ADCCXS–IR	1.61E-10*	-1.39	1.36E-10**	-2.24	4.02E-10	-0.67
ADCCS	1.64E-10*	-1.51	1.26E-10 ***	-3.75	5.44E-10 ***	-2.37
DCC	1.62E-10	-1.23	1.26E-10 ***	-3.75	6.81E-10 ***	-2.53
CCC	2.35E-10***	-6.72	9.64E-11	0.98	4.00E-10	-0.80
EWMA	3.39E-10***	-6.14	1.15E-10	-0.22	4.40E-10 ***	-2.49
BEKK	1.89E-10***	-4.37	9.49E-11*	1.88	4.95E-10 ***	-2.56

Note: This table shows the univariate mean square errors (UMSEs) from the various competing forecasting models evaluated at three different out-of-sample periods. The UMSEs are calculated based on the squared difference between the forecasted covariance and the realised covariance. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. DM refers to the Diebold–Mariano test statistic, evaluating whether the difference between two competing models' forecasts is statistically different from zero. The asterisks indicate that the ADCCXS–OF covariance forecasts are statistically more accurate than the forecasts made by the competing forecasting approach. *** denote significance at the 1%, ** at the 5% and * at the 10% levels.

Summarising the results presented above, we find a lot of evidence confirming both H2 and H3: order flow information enhances covariance predictions (H2) whereas, by comparison, interest-rate-based methods generate significantly less accurate covariance predictions (H3).

Figure 3.1 shows a graphical representation of the different performances. The performance differences refer to the differences in the UMSEs between the ADCCXS–OF model and its nested variants for the EUR–GBP covariance forecasts. This analysis will shed further light on how the ADCCXS–OF model performs relative to its main competitors. If the ADCCXS–OF model indeed performs better than its competitors on average, this should be indicated by smaller UMSEs for most of the out-of-sample period; any superior predictive ability for covariance spikes should also be shown in the time-series plots.

Two observations can be made from Figure 3.1. First of all, the ADCCXS–OF-based forecasts provide lower UMSEs (represented by the blue line) on average than any of its nested models during the whole out-of-sample period (represented by the red and yellow lines). This indicates that the ADCCXS–OF model provides more accurate covariance estimates on average than its nested variants. Second, when looking at the extreme spikes of the diagram, the UMSEs of the ADCCXS–OF model usually fall below the nested forecasting models' UMSEs, suggesting that when the predicted forecast was imprecise, the imprecision was the smallest for the ADCCXS–OF forecast. This pattern is particularly notable during the out-of-sample period of the first subsample period. However, a similar conclusion also applies to the out-of-sample period of the second subperiod (Panel B) [50,51]

[50] As the differences in accuracy among the competing covariance prediction models are very small for the out-of-sample period of the third subsample, we do not provide a graphical representation for this subsample.
[51] Please note that the pattern is more visible in the first graph, as the maximum of the *y*-axis is equal to 4.5e-10, whereas because of the larger MSEs in the second subperiod, the maximum of the *y*-axis is equal to 4.5e-9. If we look at the differences in the spikes, the main conclusions are the same for all out-of-sample periods.

Figure 3.1 Graphical UMSEs of the ADCCXS, ADCC and DCC covariance predictions

Panel A. **Panel B.**

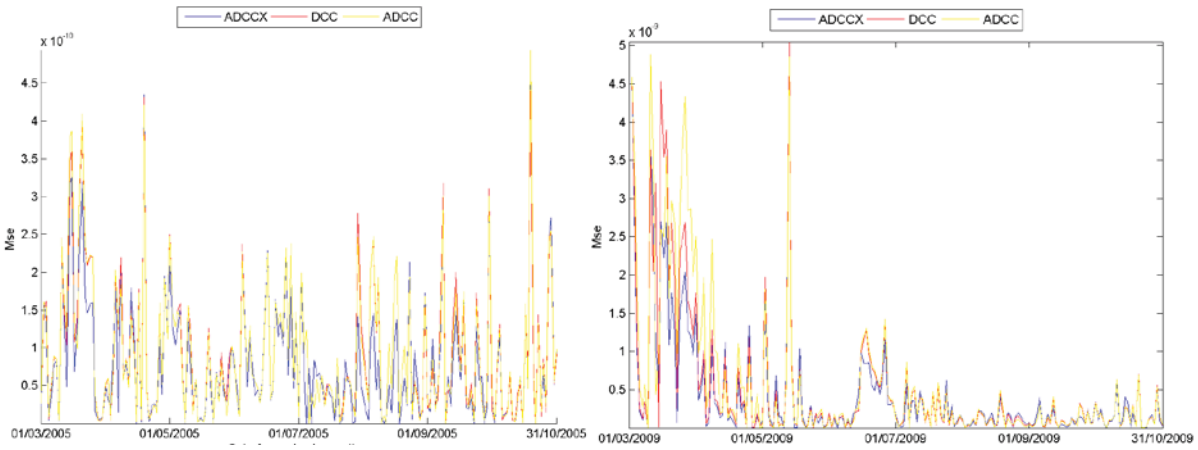

Note: This graph compares the univariate mean square errors (UMSEs) of the ADCCXS-OF model and its nested versions. The UMSEs are associated with the EUR–GBP covariance forecasts. The blue line represents the UMSEs incurred by the ADCCXS–OF and the red (yellow) line represents UMSEs incurred by the DCC (ADCC). Panel A represents the UMSEs for the out-of-sample period of the first subsample (15/03/2005–31/12/2005) and Panel B for the second subsample (15/03/2009–31/12/2009) respectively.

3.6.2 Positive-type asymmetry

A question that arises is whether the dominance of positive-type asymmetry over negative type asymmetry and symmetry documented in daily exchange rates (see Panel B in Table B.1 in Appendix B) can enhance the predictions of our proposed ADCCXS–OF model even further. To test this possibility, we use the ADCCXS–OF specification given in Equation (2.8); however, we allow for positive asymmetry via the change in the slope ($n_t = I(\varepsilon_t > 0) \circ \varepsilon_t.$) and also use the ADCCXE–OF specification given in Equation (2.8). Note that the ADCCXE–OF specification allows for either positive or negative asymmetry to arise. However, given that positive asymmetry dominates on average across the sample period (see Panel A in Table B.1 in Appendix B), we consider it as an additional specification

test for positive asymmetry. For robustness purposes, we also evaluate the ADCCXS–IR model, taking positive asymmetry into account. As before, we choose the ADCCXS–OF for modelling negative type asymmetry via the change in the slope as the benchmark for our comparison. To avoid confusion, we specifically label the ADCCXS model incorporating joint positive returns as ADCCXP and as ADCCXN for the version incorporating joint negative returns in the parameterization.

Table 3.5 sets out the empirical results for this exercise. Based on the results, the answer to the question on whether positive asymmetry or negative asymmetry enhances covariance predictions seems to depend on the time period studied. For the financially stable initial out-of-sample period (Panel A), positive asymmetry models such as the ADCCXP–OF dominate their negative asymmetry counterparts such as the ADCCXN-OF for all three exchange rate combinations (EUR–GBP, EUR–JPY and GBP–JPY) in terms of accuracy. The differences in the UMSEs between ADCCXP–OF and ADCCXN–OF, however, are not statistically significant for any exchange rate pair (1.22 e-10 vs. 1.24e-10 for EUR–GBP, 5.38e-11 vs. 5.82e-11 for EUR–JPY and 1.49e-10 vs. 1.51e-10 for GBP–JPY). Further evidence of the higher accuracy of models that rely on positive asymmetry during this sample period can be found by comparing the ADCCXP–IR and ADCCP models in Table 3.5 with their counterparts in Table 3.4. Note that the ADCCXE–OF model provides significantly lower UMSEs for the GBP-JPY covariance predictions (1.41e-10) than the ADCCXN–OF benchmark (1.51e-10).

Panel B documents a different pattern. For the EUR–GBP covariance dynamics, the ADCCXN–OF model provides a significantly lower UMSE (7.47e-10) than the ADCCXP–OF (7.77e-10) and outperforms any other specification with positive asymmetry. A similar pattern is documented for the EUR–JPY and GBP–JPY pair, where the ADCCXN–OF model

has a lower UMSE (9.13e-10 and 1.009e-10) than its closest competitor with positive asymmetry, namely the ADCCXP–OF model (1.04e-09 and 1.17e-09). The accuracy of the ADCCXE–OF-based forecasts (7.8-10, 1.14e-09 and 1.02e-09) closely follows the accuracy of the ADCCXP–OF predictions (7.77e-10, 1.04e-09 and 1.17e-09).

Panel C further confirms the main findings in Panel B for the out-of-sample period of the third subsample. Specifically the ADCCXN–OF model significantly outperforms all other competing approaches for all exchange rate pairs, except the ADCCXE–OF specification evaluated for GBP–JPY covariance predictions. The differences are not as large as they are during the second subperiod for EUR-GBP and EUR-JPY covariance predictions, but are still significant, thereby again showing the importance of taking negative asymmetry into account when predicting joint variations in exchange rate returns during a financially turbulent period. As our dataset contains two financially turbulent periods, we use the ADCC class of models taking negative asymmetry into account via the change in a slope for the remainder of this chapter.

Table 3.5 Out-of-sample UMSEs for ADCCXN, ADCCXP, and ADCCXE

	EUR–GBP		EUR–JPY		GBP–JPY	
Panel A. January 2002 – December 2005						
	UMSE	DM	UMSE	DM	UMSE	DM
ADCCXN–OF	1.24E-10		5.82E-11		1.51E-10	
ADCCXP–OF	1.22E-10	1.15	5.38E-11	1.43	1.49E-10	0.79
ADCCXP–IR	1.28E-10 ***	-2.29	7.14E-11 ***	-4.29	1.72E-10 ***	-3.25
ADCCP	1.29E-10*	-1.53	5.71E-11	0.54	1.72E-10 ***	-3.42
ADCCXE–OF	1.22E-10	0.82	5.74E-11	0.34	1.41E-10	3.19
Panel B. January 2006 – December 2009						
	UMSE	DM	UMSE	DM	UMSE	DM
ADCCXN–OF	7.47E-10		9.13E-10		1.00E-09	
ADCCXP–OF	7.77E-10 **	-1.92	1.04E-09	-1.31	1.17E-09*	-1.57
ADCCXP–IR	8.67E-10 ***	-2.37	3.81E-09 ***	-9.62	1.33E-09 ***	-3.44
ADCCP	7.85E-10 **	-2.17	1.19E-09 ***	-2.46	1.20E-09 ***	-2.96
ADCCXE–OF	7.80E-10 **	-2.22	1.14E-09 **	-2.18	1.02E-09	-0.60
Panel C. January 2010 – December 2013						
	UMSE	DM	UMSE	DM	UMSE	DM
ADCCXN–OF	1.59E-10		1.11e-10		3.85E-10	
ADCCXP–OF	1.62E-10 **	-1.69	1.15E-10 **	-2.04	6.05E-10 ***	-2.50
ADCCXP–IR	1.65E-10 ***	-2.73	1.17E-10 **	-2.24	5.97E-10 **	-2.15
ADCCP	1.64E-10 ***	-2.51	1.20E-10 **	-1.84	5.44E-10 **	-2.30
ADCCXE-OF	1.64E-10 ***	-2.83	1.29E-10 ***	-4.15	3.94E-10	-0.68

Note: This table shows the univariate mean square errors (UMSEs) from the ADCC class of competing forecasting models evaluated at three different out-of-sample periods. The UMSEs are calculated based on the squared difference between the forecasted covariance and the realised covariance. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. DM refers to the Diebold–Mariano test statistic, evaluating whether the difference between two competing models' forecasts is statistically different from zero. The asterisks indicate that the ADCCXSN–OF covariance forecasts are statistically more accurate than the forecasts made by the competing forecasting approach. *** denote significance at the 1%, ** at the 5% and * at the 10% levels.

3.6.3 Choice of rolling estimation window

In the baseline forecasting exercises, we have used a rolling window size of 500 observations. No previous studies on predicting co-movement have ever considered what would happen if different window sizes were used. This question is practically relevant, so

we empirically examine it by using three different rolling window sizes: a small one of 250 observations, a long one of the total number of observations in the sample period and a medium one of 500 observations, as is commonly used. Again, the parameters are re-estimated every 25 observations during evaluation of the performances associated with all three window sizes.

The choice of 250 observations as the smallest window size is suggested by our methodology presented above, where the first-stage estimation needs at least 250 observations to generate reasonably unbiased GARCH parameter estimates. We then select 500 observations for a medium-sized rolling window, which should provide a good trade-off between having too far and having too much data. The largest possible estimation window size contains 800 observations, the total in-sample period.[52] We argue that a medium-sized window should provide us with the most accurate forecasting results. The reason for using this medium-sized window is that it should reduce the disadvantages of either a long rolling window or a short rolling window. For long rolling windows, the problem that arises is that a lot of observations used for parameter estimation are potentially unrelated to the present (or, even worse, not related to the data generating process itself), possibly resulting in biased parameter estimates and forecasts (Clark and McCracken, 2009). Unfortunately, a very small rolling window increases the variance of the parameter estimates. This increase in variance leads to forecasting errors and causes the mean square forecasting error to increase. Hence, when making a forecast, there needs to be a balance between using too much or too few data for estimating the model parameters, which justifies a medium-sized rolling window.

[52] We are aware that, apart from the discussion in the literature on the optimal window length, there is a growing discussion as to whether expanding or rolling windows are preferred for forecasting applications (Giacomini and White, 2006). We follow the finance literature in this chapter, where rolling windows are frequently used, and compare the effects of different rolling windows.

Table 3.6 presents the results of our experiments. To make the interpretation easier, we report the differences between the short and medium, and between the long and medium estimation windows, in terms of the MSE differences. Thus any positive values will indicate that the medium-sized window is more accurate in terms of "loss" (i.e., in terms of the UMSEs between the predicted and the realised covariance).

Looking at the first period in Table 3.6 (Panel A), a picture emerges. For all covariance forecasts, the medium-sized window produces smaller UMSEs than the small estimation window; in other words, the predicted covariance from the medium-sized window better matches the realised covariance. For the EUR–JPY pair, the differences are statistically significant for all models. For the ADCCXS–OF model, the additional loss caused by a small or long moving window compared with a medium-sized window is 2.43e-11 and 1.66e-11 respectively. This is a substantial difference, as the MSE of the ADCCXS-OF model with a medium-sized window is 1.24e-10, which indicates that the small estimation window generates a ~10% higher UMSE, with the long estimation window generating a ~19% greater UMSE. However, for other exchange rate pairs, no window size significantly outperformed another, although the medium-sized window provided the lowest UMSEs on average.

Turning to the second subsample period, Panel B indicates that the differences between the medium window approach and both the short and long window approaches begin to widen drastically. For this subperiod, the medium-sized estimation window results in significantly more accurate forecasts, regardless of which forecasting model is run for the EUR–GBP covariance forecasts. For the second subsample, the differences between different window sizes are substantial. For the ADCCXS–OF model, the differences between the medium and short windows reach 5.11e-10, which is equal to about two thirds of UMSE of

the medium-sized window. In other words, the UMSE of the short window is more than 66% larger than the UMSE of the medium window.

Table 3.6 Out-of-sample UMSE differences for competing estimation windows

	EUR–GBP		EUR–JPY		GBP– JPY	
Panel A. January 2002 – December 2005						
	SW- MW	LW- MW	SW- MW	LW- MW	SW- MW	LW- MW
ADCCXS–OF	9.19E-12	1.58E-11*	2.43E-11 ***	1.66E-11 ***	3.79E-12	3.47E-12
ADCCXS– IR	7.89E-12	1.55E-11	8.09E-11 ***	1.30E-11 ***	4.03E-13	3.06E-12
ADCCS	1.27E-12	1.22E-11	3.63E-11 ***	1.47E-11	8.64E-12	2.26E-11
DCC	2.06E-12	1.21E-11	3.32E-11 ***	8.24E-12 *	1.22E-12	2.25E-11
Panel B. January 2006 – December 2009						
	SW- MW	LW- MW	SW- MW	LW- MW	SW- MW	LW- MW
ADCCXS–OF	5.11E-10 ***	2.48E-10 ***	1.69E-11	2.57E-12	1.16E-11	1.71E-10
ADCCXS–IR	5.60E-11 *	1.37E-10 ***	8.49E-11 ***	4.00E-12 **	7.88E-11 ***	8.25E-11 ***
ADCCS	3.48E-10 ***	1.11E-10 ***	3.93E-11 ***	2.95E-12	3.81E-10 ***	5.07E-10 ***
DCC	4.14E-10 ***	1.87E-10 ***	3.85E-11 ***	5.27E-12	2.60E-10 ***	4.70E-10 ***
Panel C. January 2010 – December 2013						
	SW- MW	LW- MW	SW- MW	LW- MW	SW- MW	LW- MW
ADCCXS–OF	1.69E-11	2.17E-12	4.27E-11 *	4.48E-11	9.95E-11	1.58E-11
ADCCXS–IR	1.45E-11	3.00E-13	8.53E-11 *	4.91E-11 **	3.53E-10	6.40E-11
ADCCS	1.59E-11	7.35E-13	5.32E-12	1.25E-11 *	4.93E-11	2.20E-11
DCC	1.77E-11	1.17E-13	3.35E-12	1.25E-11	4.71E-11	1.97E-11

Note: This table shows the differences in univariate mean squared error (UMSE) between the different rolling estimation windows investigated. We evaluate the differences for the ADCCXS and its nested variants. LW denotes the long window estimation using 800 in-sample observations for parameter estimation. Similarly, MW denotes a medium-sized window using 500 in-sample observations for parameter estimation and SW denotes the short estimation window which uses 250 in-sample observations for parameter estimation. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. DM stands for the Diebold–Mariano test statistic, evaluating whether the difference between two competing estimation windows models is statistically different from zero. The asterisks indicate that the covariance forecasts based on a medium sized window are statistically more accurate than the forecasts made by a short or long estimation window. *** indicates significance at the 1%, ** indicates significance at the 5% level and * indicates significance at the 10% level.

The same pattern applies to the long window, with a MSE difference of 2.48e-10, which is about 150% larger than the medium window's. The medium window's forecasts are also statistically more accurate than those of the small or the long windows, for four out of five competing models of forecasting GBP–JPY covariance. Though the differences are not statistically significant for the EUR–JPY covariance predictions, the medium-sized window still provides the highest accuracy for this pair. The final period's results are shown in Panel C. For each multivariate GARCH model and currency pair studied, the medium-sized window yields the lowest UMSEs. However, the differences are not statistically significant on average.

Figure 3.2 provides a graphical representation of the performance differences arising from different rolling window choices, including a short, a medium-sized and a long window. We evaluate the performance of each window on the accuracy of the ADCCXS-OF model's predictions for EUR–GBP covariance. The graphical representation helps to highlight how the small and the large rolling windows' forecasts result in larger inaccuracy than the medium window's forecasts.[53]

The smallest rolling window forecast produces extreme spikes, indicating that the parameter flexibility leads to extreme covariance forecasts, resulting in large over- or underestimations of the forecasted covariance evaluated against the realised covariance. By comparison, the largest rolling window's forecast generates higher MSEs over the total sample period on average, as it never outperforms the medium rolling window's forecast because of the parameter inflexibility and hence the somewhat constant covariance predictions arising from a large dataset. A second aspect is worth mentioning. During the financially turbulent period, the short-term-window-based predictions suffer from re-

[53] We do not show the results for the third subperiod in the graph, as these are represents a mix between the findings from the out-of-sample periods of the first and second subsamples.

estimation issues, as shown by the spikes in the UMSEs during the 90th and 160th out-of-sample evaluations. This again suggests the problem of having a short estimation window, where the estimated parameters only depend on the most recent observations and thus do not generalise well to new observations. Although the choice of the re-estimation window would be another aspect worth investigating, we leave this for further research and follow the study of Laurent et al. (2013), re-estimating the parameters after 25 out-of-sample predictions.

Figure 3.2 Graphical UMSEs of different window sizes

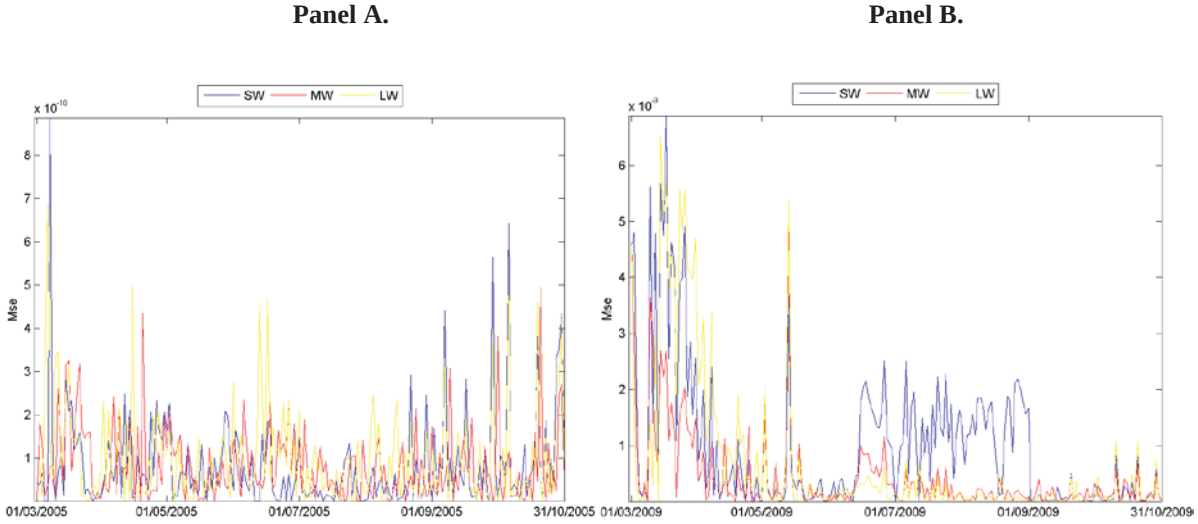

Note: This graph compares the mean square errors (UMSEs) of different window sizes for, respectively, the first subsample's and the second subsample's out-of-sample period. The UMSEs are associated with the EUR–GBP covariance forecasts resulting from the ADCCXS-OF model using different estimation window sizes. The blue line represents the MSEs from a rolling window size of 250 observations (denoted short window or SW), the red line is for 500 observations (MW) and the yellow line is for 800 observations (LW). Panel A represents the MSEs for the out–of-sample period of the first subsample (01/03/2005–31/10/2005), Panel B is for the second subsample (01/03/2009–31/10/2009).

3.6.4 Volatility Predictions

A natural extension of the work reported in the previous sections is to explore whether order flows can predict future volatilities and so further improve covariance forecasting. To this end, we replace the standard GARCH with the GARCH-X model in the first-stage

estimation for all conditional correlation models. The GARCH-X–ADCCXS follows Equations 2.15 and 2.8 outlined in the second chapter, using order flow to predict both variance and correlation dynamics.

We conjecture that, under the GARCH-X model, order flows will be shown to be beneficial for predicting the entire out-of-sample covariance matrix of a currency pair. The underlying story, as told in Chapter 2, is that order flows create buying or selling pressure that results in larger exchange rate fluctuations or volatility. Some initial evidence for the link between order flows and the volatility dynamics has been provided in Chapter 2, as the GARCH-X estimation results showed that order flows have positive effects on variance dynamics. However, the estimates presented previously were in-sample in nature. In this chapter, we turn to out-of-sample forecasting problems.

The approaches evaluated include the GARCH-X–ADCCXS, GARCH-X–ADCC ($\eta = 0$) and GARCH-X–DCC (restricting $\gamma_1 = \gamma_2 = 0$) models. We evaluate their performances against the performance of the GARCH-based dynamic correlation models, based on the matrix-wise MSE (Equation 3.3) criterion. However, the matrix-wise and univariate errors are ultimately linked, as the covariance forecasts depend on the correlation and variance forecasts. Table 3.7 sets out the results of the various GARCH-X–and GARCH-based forecasting approaches, using a medium-sized estimation window of 500 observations.

In Table 3.7, the results of the GARCH-X-based forecasting models support hypothesis H4: the MMSEs are on average lower for this class of models than for the GARCH-based dynamic conditional correlation models. This finding holds for all the three exchange rate pairs investigated and is robust to all subsamples' out-of-sample periods, except the covariance matrix prediction for the GBP–JPY pair over the third subsample's out-of-sample periods. In particular, for the second and third subsample period, the GARCH-X

specification results in significantly lower MMSEs than the GARCH specification. This confirms our two conjectures made before. First, the rather modest forecasting accuracy in the third subsample may be caused by inaccurate variance forecasts, which affects the covariance predictions. Second, the larger volatility spikes predicted by the GARCH-X model (captured by spikes in the order flows) are conducive to volatility forecasting, especially for periods of high fluctuation in the realised volatility (such as the second subsample period, which covers the GFC).

Table 3.7 also demonstrates that, for the first subsample period, the GARCH-X–ADCCXS model significantly outperforms four out of the five competing conditional correlation models by showing the lowest MMSEs overall. It is followed in terms of accuracy by the GARCH-X–ADCC model, thus revealing again the higher forecasting accuracy of the GARCH-X-based models. Although the differences in MMSEs between the GARCH-X–ADCCXS and GARCH-ADCCXS models are modest for this subsample period (ranging between 4e-12 and 1.4e-11), the differences become substantially larger for the out-of-sample period of the second subsample period.[54]

Looking at the covariance matrix predictions of the EUR–GBP exchange rate pair during the second subsample period, the MMSE of the GARCH-X–ADCCXS model is 4.75E-09, whereas the GARCH-ADCCXS MMSE equals 5.05E-09. The difference between the two forecasting models reaches 3.0e-10, which is 6.3% of the average difference. For the EUR–JPY and GBP–JPY pairs, again the GARCH-X–ADCCXS model results in slightly lower MMSEs (4.78E-09 and 6.26E-09 respectively) than the GARCH-ADCCXS specification (4.79E-09 and 6.30E-09).

[54] Even though the differences appear to be small, they are statistically significant.

Table 3.7 Out-of-sample MMSEs for GARCHX-ADCCXS and GARCH-ADCCXS forecasts

	EUR–GBP		EUR–JPY		GBP–JPY	
Panel A. January 2002 – December 2005						
	MMSE	DM	MMSE	DM	MMSE	DM
GARCH-X–ADCCXS	6.41e-10		8.00e-10		8.29e-10	
GARCH-X–ADCCS	6.49e-10 **	-2.23	8.03e-10	-0.56	8.32e-10	-0.61
GARCH-X–DCC	6.45e-10 **	-1.92	8.17e-10 ***	-2.38	8.43e-10 **	-2.09
GARCH-ADCCXS	7.68e-10 **	-2.01	8.08e-10 **	-1.74	9.35e-10 ***	-5.70
GARCH-ADCCS	7.79e-10 **	-1.70	8.14e-10 **	-1.77	9.71e-10 ***	-5.82
GARCH-DCC	7.81e-10 **	-1.75	8.16e-10 **	-2.27	9.67e-10 ***	-5.48
Panel B. January 2006 – December 2009						
	MMSE	DM	MMSE	DM	MMSE	DM
GARCH-X–ADCCXS	4.75e-09		4.78e-09		6.26e-09	
GARCH-X–ADCCS	4.83e-09 ***	-3.35	5.09e-09 ***	-2.58	6.30e-09***	-3.75
GARCH-X–DCC	4.81e-09*	-1.39	5.41e-09 ***	-3.81	6.52e-09 ***	-3.88
GARCH-ADCCXS	5.05e-09***	-4.15	4.79e-09	-0.23	6.30e-09	-1.07
GARCH-ADCCS	5.04e-09 ***	-3.56	5.25e-09 ***	-2.65	6.48e-09***	-2.93
GARCH-DCC	5.06e-09 ***	-3.68	5.11e-09 **	-2.25	6.33e-09	-0.09
Panel C. January 2010 – December 2013						
	MMSE	DM	MMSE	DM	MMSE	DM
GARCH-X–ADCCXS	9.28e-10		1.93e-09		2.26e-09	
GARCH-X–ADCCS	9.46e-10 **	-1.82	1.99e-09 ***	-8.20	2.30e-09 **	-2.15
GARCH-X–DCC	9.43e-10 **	-2.18	1.99e-09 ***	-8.20	2.31e-09 ***	-2.46
GARCH-ADCCXS	9.86e-10 **	-2.02	2.02e-09 ***	-4.12	2.19e-09	9.65
GARCH-ADCCS	9.90e-10 **	-1.87	2.04e-09 ***	-4.24	2.29e-09*	-1.31
GARCH-DCC	9.90e-10 **	-2.17	2.04e-09 ***	-4.24	2.28e-09	-1.08

Note: This table shows the multivariate mean square errors (MMSEs) from GARCH and GARCH-X based dynamic conditional correlation models. Multivariate mean square errors (MMSEs) are calculated as outlined in Equation (3.3). The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. DM stands for the Diebold–Mariano test statistic, evaluating whether the difference between two competing forecasting models is statistically different from zero. The asterisks indicate that the GARCH-X–ADCCXS covariance forecasts are statistically more accurate than the forecasts made by the competing forecasting approach. *** indicates significance at the 1%, ** indicates significance at the 5% level and * indicates significance at the 10% level.

As far as the third subsample period is concerned, Table 3.7 shows that the GARCH-X-based correlation models significantly outperform the GARCH-based correlation models for two out of the three major exchange rates. The GARCH-X–ADCCXS model provides the

highest accuracy for the exchange rate pairs that involve EUR (9.28e-10 for EUR–GBP and 1.93e-09 for EUR–JPY), whereas the GARCH-ADCCXS model shows the highest accuracy for predicting the GBP–JPY covariance matrix.

Figure 3.3 shows the distribution of the MMSEs for the entire out-of-sample period. The GARCH-X-based dynamic conditional correlation models have similar MMSEs to the GARCH-based correlation models on average. However, in the presence of spikes, the MMSEs are substantially lower in the former than in the latter. This indicates that incorporating order flows into a GARCH model is important, especially if volatility is high, again confirming H4 visually in addition to the results presented in Table 3.7.

Figure 3.3 Graphical MMSEs of GARCH-X and GARCH-based predictions

Panel A. **Panel B.**

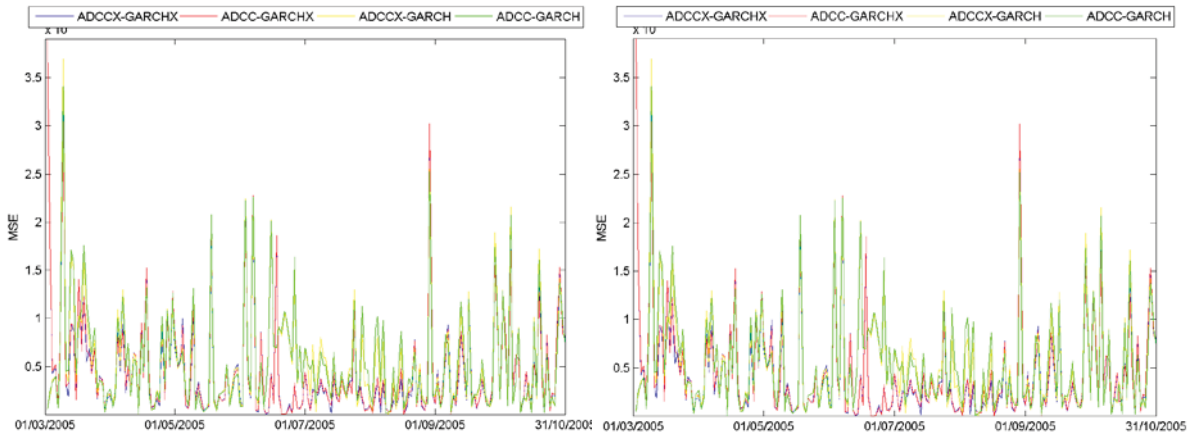

Note: This figure highlights the distribution of the multivariate mean squared errors (MMSEs) from both the GARCH- and GARCH-X-based dynamic conditional correlation models during the first (Panel A) and the second (Panel B) subsample periods' out-of-sample period. Both periods are distinctively different in terms of their mean returns and volatilities, as the former is a financially stable period with low volatility, whereas the latter contains the GFC and frequent volatility spikes. ADCCXS refers to the model incorporating order flow in the correlation driving process. Both graphs depict the MMSEs between the forecasted and realised covariance matrix for the EUR–GBP pair.

3.7 Portfolio Optimisation

As noted in introduction, this section aims to evaluate the economic gain for investors arising from the work reported in the previous sections. The previous section showed that the ADCCXS–OF model outperforms any other models in terms of predictive accuracy for the EUR–GBP covariance forecasts. However, a further question is whether or not a more accurate covariance forecast would bring investors practical gains in trading. To address this question, we consider a representative US investor who allocates his or her wealth every period between two risky currency assets (EUR and GBP deposits) [55] and one riskless currency asset (USD deposits). The optimal investment weights are based on predictions of the future returns and the future joint variations of the risky assets returns.

3.7.1 Notation and setup

Taking the perspective of a US investor, we forecast the daily exchange rate returns for USD–EUR (USD–GBP) using predictive regressions including order flow as an explanatory variable. The regressions are not new and therefore the explanation will be short (see Rime et al., 2010). The predictive regressions have the following linear structure:

$$P_t = C + \Lambda X_{t-1} + \Gamma P_{t-1} + U_t, \qquad (3.4)$$

where $P_t = [\Delta s_t]$ is a vector of the exchange rate returns, X_{t-1} is a vector of order flow, Λ and Γ are the coefficients to be estimated, C is the constant and U_t is the residual[56].

[55] Our choice of EUR and GBP is because these are most actively traded currencies worldwide. However, we will show that the main conclusions also apply when EUR and JPY are considered as the risky assets.

[56] We decided to use this particular forecasting model for two reasons. First, the results in Rime et al. (2010) seem to be in favour of this model (i.e., it has the highest predictive accuracy). Second, we experimented with other models; however, only the order-flow-based return predictions generated positive Sharpe ratios for the portfolio optimisation problem in what follows. For further information about the forecasting approach, see Rime et al. (2010). We show the parameter estimates, and potential shortcomings of this regression in the next chapter.

Suppose the return from holding EUR and GBP is given by the following 2x1 vector of risky returns, $r_{t+1|t}$:

$$r_{t+1|t} = [E_t(\Delta s_{t+1}^{EUR} + i_t^{EUR}), E_t(\Delta s_{t+1}^{GBP} + i_t^{GBP})], \quad (3.5)$$

and the return from holding the risk-free asset is the domestic (US) interest rate. The investor dynamically rebalances his/her portfolio to maximise the expected return conditional on a given target volatility. Thus, at each period *t*, the investor solves the following problem:

$$\frac{\max}{wt}\left\{r_{p,t+1} = w_t^{'\prime} r_{t+1|t} + (1 - w_t^{'\prime}\iota) i_t^{USD}\right\} \quad (3.6)$$

$$s.t.\left(\sigma_p^*\right)^2 = w_t^{'\prime}\sum_{t+1|t} w_t,$$

where $r_{p,t+1}$ is the portfolio return at time $t+1$; $\left(\sigma_p^*\right)$ is the target level of risk (volatility) for the portfolio, ι is a 2x1 matrix of ones, $\Sigma_{t+1|t}$ is the predicted variance-covariance matrix of the risky assets at time t+1 and $w_t^{'}$ represents the 2x1 matrix of portfolio weights. The solution to the optimization problem in (3.6) gives the following risky-asset weights:

$$w_t = \frac{\sigma_p^*}{\sqrt{C_t}}\Sigma_{t+1|t}^{-1}(r_{t+1|t} - i_t^{US}) \quad (3.7)$$

where,

$$C_t = (r_{t+1|t} - i_t^{US})\Sigma_{t+1|t}^{-1}(r_{t+1|t} - i_t^{US}). \quad (3.8)$$

The weight assigned to the risk-free asset is $(1 - w_t)$. Clearly, the solution and hence the weights (profitability) of portfolio depends on the future appreciations or depreciations of the EUR and GBP relative to the USD ($r_{t+1|t}$) as well as the future joint variations of the exchange rate returns $\Sigma_{t+1|t}^{-1}$, all of which have to be forecasted.

We construct the portfolio weights based on Equation 3.7 for each covariance forecast, so different forecasts lead to different portfolios. We then evaluate the relative performances of the portfolios by using their Sharpe ratios. The portfolio that yields the highest Sharpe ratio is the one whose underlying covariance forecast will be deemed to be the "best", as all other parameters are the same for each portfolio. Following Rime et al. (2010), we set the target volatility to 10% p.a. and compute the Sharpe Ratio (SR) , Sortino ratio (*SO*) and performance fee a risk-averse investor will be willing to pay for switching from a constant covariance model (denoted STATIC[57]) to a dynamic (ADCCXS, ADCC, DCC, CCC, BEKK, EWMA) model. Further to the Sharpe Ratio (SR) we evaluate the Sortino ratio (*SO*) and the performance fee noted above.

The Sharpe and Sortino ratios[58] are given as:

$$SR = \frac{\overline{r_p - \iota_f^{US}}}{\sigma_p} \quad ,\text{SO} = \frac{\overline{r_p - \iota_f^{US}}}{\sigma_d}, \tag{3.9}$$

where $\overline{\mathrm{r_p} - \iota_{\mathrm{f}}^{\mathrm{US}}}$ represents the annualised average excess return from investment and $\sigma_{\mathrm{p}}$is the annualised standard deviation of portfolio returns. σ_{d} represents the standard deviation of the negative returns.

Note that setting the target volatility at 10% p.a. in Equation 3.7 does not imply that the volatility of the constructed portfolio is in fact equal to 10% p.a., nor does it imply that a covariance prediction provides more economic value if the actual portfolio volatility is closer to the target volatility. The target volatility is only used as a benchmark for the portfolio construction. The best covariance prediction will lead to the highest risk adjusted returns for

[57] The constant covariance approach "predicts" that the future covariance matrix is equal to unconditional covariance matrix of returns, i.e. the sample the sample average. This approach is commonly referred to as "static" covariance prediction or constant covariance model.
[58] The Sortino ratio measures the return to downside risk, whereas the Sharpe ratio measures the return to overall risk.

the portfolio, regardless of whether the target volatility is reached or not. A common method to evaluate risk-adjusted returns is the Sharpe Ratio, a measure also employed in Rime et al. (2010) and King et al. (2010) .

Different co-variance predictions might result into slightly lower or higher portfolio volatility than the targeted volatility. In order to take this into account in our analysis, we follow Della Corte et al. (2013) and present evaluate the trading strategies in terms of Sharpe

With regard to the performance fee criterion, we follow Goetzmann, Ingersoll, Spiegel and Welch (2007). These authors define the so-called manipulation-proof performance measure as:

$$M(r_p) = \frac{1}{(1-\delta)} \ln \left\{ \frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{1+r_{p,t+1}}{1+r_f} \right)^{(1-\delta)} \right\}, \quad (3.10)$$

where $M(r_p)$ is an estimate of the portfolio's premium return after adjusting for risk. It can be interpreted as the certainty equivalence of excess portfolio returns. This is an attractive criterion, since it is robust to the distribution of portfolio returns and does not require the assumption of a particular utility function when ranking portfolios. The δ parameter denotes the investor's degree of relative risk aversion. We select a risk aversion coefficient of δ = 6. We then evaluate the differences between the risk premium of the portfolio constructed using a particular covariance forecast and the risk premium of a portfolio constructed using a constant variance covariance, as follows:

$$\Phi = M(r_p^*) - M(r_p), \quad (3.11)$$

where Φ is the maximum performance fee an investor will pay to switch from the constant covariance forecast to the time-varying alternative. In other words, this utility-based criterion

measures how much a mean-variance investor is willing to pay for conditioning his investment decisions on better covariance forecasts. We report Φ in annualised basis points.

3.7.2 Results

Table 3.8 presents the results from the portfolio optimisation setting, with annualised Sharpe and Sortino ratios displayed in Columns 3 and 4. For the out-of-sample period of the first subsample (Panel A), the following conclusions can be drawn. First, simple covariance forecasting models such as the EWMA, CCC and STATIC approach provide the lowest Sharpe (Sortino) ratios of 0.11 (0.31), 0.18 (0.39) and 0.14 (0.33). By contrast, for our portfolios that consist of EUR, GBP and USD, the covariance forecasts using order flows provide the investor with the highest Sharpe (Sortino) ratio of 0.55 (0.93), followed by the BEKK forecasts with a Sharpe (Sortino) ratio equal to 0.41 (0.72), and the interest-rate-based version of our forecast (ADCCXS–IR), with Sharpe and Sortino ratios of 0.37 and 0.70. Although the difference in the Sharpe ratios is small, we can still confirm that covariance forecasts based on order flow can improve the Sharpe ratio and therefore bring economic value to a mean-variance investor. The Jobson and Korkie (1981)[59,60] test for differences in the Sharpe ratio suggests that the ADCCXS–OF-based covariance forecast results in a significantly higher Sharpe ratio (0.55) than the covariance "prediction" generated by a constant covariance model (0.14). By contrast, the Sharpe ratios for all the other forecasting approaches fall into the 90% confidence interval of the constant covariance prediction, which means that they are not statistically different from the Sharpe ratios generated by a constant covariance model.

[59] The methodology tests whether the Sharpe Ratio from any dynamic covariance forecast is significantly greater than that from the historical covariance "prediction". Since King et al. (2010) and Rime et al. (2010) use this particular benchmark, we do so too in order to link our findings to theirs.

[60] Note that the Jobson and Korkie (1981) test for differences in Sharpe ratios is not the only possible test for differences in Sharpe ratios. In fact, Lo (2002) proposed a methodology that takes differences in skew and kurtosis between the portfolios' returns into account. Although we argue that the Lo (2002) is more accurate and will use it in the next chapter, in this chapter, we follow the literature on portfolio management strategies and apply the more popular Jobson and Korkie (1981) test.

The performance fees for each competing covariance forecast, except for the EWMA based covariance prediction, are positive, indicating that a risk-averse trader will be willing to give up part of his/her return to switch to a dynamic covariance forecast. The positive performance fees indicate the economic value of all the dynamic covariance models (i.e., all models except the constant covariance model). For the ADCCXS-OF-based forecasts, a risk-averse investor would be willing to give up 222 basis points to switch from a static/ constant covariance prediction to an order-flow-based covariance prediction. By comparison, the performance fee given up to switch to the BEKK based forecast, the closest competitor in terms of performance fee, is 149 basis points, a difference of 73 basis points[61].

For the second subsample's out-of-sample period, shown in Panel B, the profitability of the portfolio management strategy reduces significantly. The performance fees related to all dynamic covariance predictions except for the CCC model are negative, confirming the low Sharpe ratios associated with the dynamic models (ranging between -0.02 and 0.27). The low Sharpe and Sortino ratios on average, however, cast doubt on the accuracy of the order-flow-based spot rate predictions. Furthermore, we argue that a simple GARCH estimation does not capture the underlying volatility dynamics particularly well during this financially turbulent subperiod, an argument in line with the dominance of the CCC and STATIC model during this subsample's out-of-sample period.

For the third subsample period, Panel C shows that the ADCCXS-OF model incorporating order flows dominates all other models in terms of the Sharpe ratio arising from the portfolio optimisation strategy. What is astonishing is that during this subsample period, the ADCCXS-IR covariance forecasts result in a substantially lower Sharpe ratio than the

[61] We exclude the GARCH-X-based models outlined in Section 3.6.4, so that any performance difference between the competing dynamic conditional correlation models is caused by different covariance predictions only. For GARCH-X as opposed to GARCH-based correlation models, performance differences could arise as a result of both variance and covariance predictions.

ADCCXS-OF model (1.43 vs. 1.9). This finding is further confirmed by the Sortino ratios (3.01 vs. 3.5) and the estimated performance fees (174.01 vs. 476.94 bps). The latter indicates that a risk-averse trader would be willing to pay up to 476.94 percentage points annually to switch from a static covariance prediction (constant covariance model) to order-flow-based dynamic covariance forecasting. On the other hand, the performance fee the investor is willing to give up to switch to interest-rate-based forecasting is substantially smaller, a finding in line with the other out-of-sample periods. This confirms H5.

Table 3.8 Out-of-sample economic evaluation of covariance forecasts (EUR–GBP–USD portfolio)

	μ_p	σ_p	SR	SO	Φ
Panel A. January 2002 – December 2005					
ADCCXS-OF	7.25%	8.96%	0.55 *	0.93	222.46
ADCCXS-IR	5.57%	8.67%	0.37	0.70	127.29
CCC	4.15%	10.00%	0.18	0.39	23.20
ADCCS	4.69%	9.29%	0.25	0.51	62.31
DCC	5.20%	9.35%	0.30	0.58	91.03
EWMA	3.37%	9.46%	0.11	0.31	-15.88
BEKK	6.21%	9.36%	0.41	0.72	149.40
STATIC	3.71%	9.91%	0.14	0.33	
Panel B. January 2006 – December 2009					
ADCCXS-OF	3.19%	8.99%	0.22	0.40	-107.05
ADCCXS-IR	2.33%	10.14%	0.11	0.18	-448.44
CCC	3.98%	10.34%	0.27	0.44	41.33
ADCCS	2.63%	9.15%	0.15	0.28	-309.34
DCC	2.68%	9.34%	0.16	0.28	-302.49
EWMA	1.06%	8.54%	-0.02	-0.03	-846.80
BEKK	2.98%	9.98%	0.17	0.29	-240.34
STATIC	3.98%	10.88%	0.25	0.40	
Panel C. January 2010 – December 2013					
ADCCXS-OF	17.01%	8.33%	1.90 ***	3.50	476.94
ADCCXS-IR	11.68%	7.29%	1.43	3.01	174.01
CCC	16.72%	8.33%	1.86***	3.44	451.16
ADCCS	16.19%	8.29%	1.81 **	3.35	418.94
DCC	15.58%	8.54%	1.68 **	3.04	335.15
EWMA	13.42%	8.41%	1.45*	2.65	186.90
BEKK	15.96%	8.85%	1.66**	2.93	322.25
STATIC	11.49%	8.88%	1.16	2.03	

This table presents the mean portfolio return (μ_p), volatility (σ_p) and Sharpe ratios (SRs) for the portfolio optimisation strategy embodied in Equation (3.6). Return forecasting is based on Equation (3.4) and the competing covariance forecasting models are described as previously. All returns and SRs are before the transaction costs and have been annualised to make our results comparable to King et al.'s work (2010). The Sortino ratio (SO) measures returns to downside risk and Φ refers to the risk premium a risk-averse investor is willing to give up to switch from a static covariance prediction (denoted STATIC) to a dynamic covariance forecast. The risk premium displayed is in annualised basis points. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. Panel A is concerned with the out-of-sample period of the first subsample, and Panels B. and C. with the second and third subsamples' out-of-sample period accordingly. We use the Jobson–Korkie (1981) test to evaluate whether the differences in SRs generated by the dynamic and static predictions are statistically different from zero. *** represents significance at the 1% level, ** at the 5% level and * at the 10% level.

3.7.3 Robustness

A natural question is whether the benefits of our proposed model still hold for other risky assets. Table 3.10 concerns this question for the US trader who chooses to invest in EUR and the JPY. The exchange rates of the two currencies against the USD normally move in opposite directions, which is in contrast to the portfolio consisting of the EUR and GBP (EUR and GBP tend to move in the same direction relative to the USD).

The answer to the question is affirmative. The results here follow the results presented in Table 3.9. For a trader choosing to trade the EUR and the JPY as two risky assets, the order flow- based covariance forecasts (ADCCXS-OF) yield higher Sharpe ratios than its nested dynamic variants (ADCC, DCC), as well as higher Sharpe ratios than the ADCCXS-IR model for all the three subsamples' out-of-sample periods.

Looking at the differences in detail, the ADCCXS-OF generates the highest Sharpe ratio of 0.78, closely followed by the EWMA specification (0.75) and the ADCCXS-IR specification (0.72) for the first subsample's out-of-sample period. The smallest Sharpe ratio is generated by the constant covariance model (0.51). In contrast to the out-of-sample period of the first subsample, for the second subsample, the risk-metrics EWMA predictions produce the largest trading gains (0.27), followed by the CCC (0.13) and ADCCXS-OF covariance predictions (0.12). Note that again the Sharpe ratios are very small for this time period, similar to the findings presented earlier for a portfolio containing EUR, GBP and USD. The interest-rate-based covariance prediction generates a very low Sharpe ratio for this out-of-sample period, which is in fact negative (-0.01) even before considering the transaction costs. For the out-of-sample period of the third subsample, the largest gains go to the ADCCXS-OF (0.64), followed by the CCC and ADCCS models (0.62). The smallest Sharpe ratio is generated by the constant covariance model (0.32). As for the first and second subsample, the

Sharpe ratio of the ADCCXS-OF prediction is higher than the Sharpe ratio generated by the ADCCXS-IR model (0.58).

Table 3.9 Out-of-sample economic evaluation of covariance forecasts (EUR–JPY–USD portfolio)

	Panel A. January 2002 – December 2005			Panel B. January 2006 – December 2009			Panel C. January 2010 – December 2013		
	μ_p	σ_p	SR	μ_p	σ_p	SR	μ_p	σ_p	SR
ADCCXS-OF	9.10%	10.19%	0.78	2.21%	11.34%	0.12	7.34%	9.62%	0.64
ADCCXS –IR	8.30%	9.92%	0.72	0.68%	11.00%	-0.01	7.0 9%	9.92%	0.58
CCC	8.92%	10.64%	0.73	2.36%	11.24%	0.13	7.01%	9.40%	0.62
ADCCS	8.69%	10.70%	0.70	1.83%	11.44%	0.08	7.02%	9.42%	0.62
DCC	8.69%	10.70%	0.70	2.21%	11.56%	0.11	7.09%	9.79%	0.60
EWMA	8.88%	10.35%	0.75	4.01%	11.53%	0.27	6.72%	9.75%	0.55
BEKK	7.72%	9.94%	0.66	2.12%	11.42%	0.10	6.80%	9.93%	0.52
STATIC	5.90%	9.30%	0.51	2.11%	11.34%	0.11	5.79%	10.10%	0.32

Note: This table presents the mean portfolio return (μ_p), volatility (σ_p) and Sharpe ratios (SRs) for the portfolio optimisation strategy embodied in Equation (3.6), for a EUR-JPY-USD portfolio. All returns and SRs are before the transaction costs and have been annualised to make our results comparable to King et al.'s work (2010). The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. Panel A is concerned with the out-of-sample period of the first subsample, and Panels B and C with the second and third subsamples' out-of-sample period accordingly. We use the Jobson–Korkie (1981) test to evaluate whether the differences in SRs generated by the dynamic and static predictions are statistically different from zero. *** represents significance at the 1% level, ** at the 5% level and * at the 10% level.

Our analysis so far has not taken transaction costs into account.[62] However, with daily portfolio rebalancing, the transaction costs are fairly large and could range between a 5% and 7% decrease in profit. This may lead to a 0.3–0.5 decrease in the Sharpe ratio on average. Comparing this with the actual returns and the Sharpe ratios generated by the strategy in Equation (3.6), the economic value of our analysis is questionable for the first subsample's and the second subsample's out-of-sample period. This is because the returns may be equal to

[62] Excluding the transaction costs from analysis is common practice in the literature on order-flow-based return forecasting. Although excluding transaction costs allows us to highlight the differences between the competing (co-) variance predictions, as it excludes differences arising from changes in the investment weights, any analysis neglecting these costs will overestimate the Sharpe ratio.

the transaction costs. In contrast, the returns generated by the strategy during the third subsample's out-of-sample period are more than sufficient to generate positive Sharpe ratios after taking transaction costs into account.

We therefore follow an approach derived from Garleanu and Pedersen (2013), where we trade partially towards the portfolio weights given by Equation 3.6, again considering an US trader who chooses to invest in EUR and the GBP. The story behind this strategy is that, because of transaction costs, it is not optimal to rebalance portfolio holdings fully. In other words, transaction costs make it optimal to slow down trading and to trade only partially to the optimal weights. Garleanu and Pedersen (2013) suggest that the optimal trading rate is based on the risk aversion of the investor, the transaction costs, and the discount factor of the future. As outlined by the authors, for a static solution assuming a risk aversion of $\gamma = 6$ and that the transaction costs are equal to 2 basis points, the solution is to rebalance ~10% of the current portfolio towards the optimal portfolio weights every period and to maintain the previous 90% weights. This setting is denoted as "aim-for-the target" approach in Garleanu and Pedersen (2013).

Note that in this setting, the return and covariance predictions affect the portfolio weights via two different channels. First, the return and covariance predictions are used to generate the optimal portfolio weights the portfolio trades partially towards. Second, as a result of partially trading towards the optimal portfolio every period, the portfolio weights also depend on the previous weights of the optimal portfolio.

Table 3.9 shows the transaction cost-adjusted returns and Sharpe ratios of this strategy. The order-flow-based covariance forecasting approach results in the highest Sharpe ratios for the trader for the first and third subsamples' out of sample periods (0.21 and 1.01 respectively). However, the Sharpe ratios are positive even after deducting the transaction

costs over the three out-of-sample periods. This is different from the traditional portfolio optimisation strategy mentioned before.[63] To be more precise, for the first two subsamples' out-of-sample periods, the Sharpe ratios are positive, but in the previous portfolio optimisation strategy, they turn negative after taking transaction costs into account.

Specifically investigating the performance of the ADCCXS-IR model, we can see from Table 3.9 that, in the first and the third subsample, the portfolio constructed using the IRD results in significantly lower returns than the order-flow-based approach (2.62% vs. 2.85% and 8.70% vs. 8.05%), higher volatility (14.18% vs. 13.6% and 9.50% vs. 8.63%) and hence a lower Sharpe ratio (0.21 vs. 0.18 and 1.11 vs. 0.85). Looking at the performance between the ADCCXS-OF model and its nested variants, the same pattern remains.

Consistent with the portfolio optimisation setting outlined before, this aim-for-the-target approach provides substantially higher Sharpe ratios for the first and third subsamples' out-of-sample periods, with the Sharpe ratios for the second subsample's out-of-sample period being modest. Note that Sharpe ratios can be negative even when the mean return is positive, as the risk free rate is deducted from the mean returns in calculating the Sharpe ratios.

[63] This aim-for-the-target strategy must not be confused with no trading in case the return forecast is smaller than the transaction costs. We implemented this strategy and all Sharpe ratios were negative after including the transaction costs.

T**able 3.10** Out-of-sample economic evaluation of competing covariance forecasts (Aim portfolio)

	Panel A. January 2002 – December 2005			Panel B. January 2006 – December 2009			Panel C. January 2010 – December 2013		
	μ_p	σ_p	SR	μ_p	σ_p	SR	μ_p	σ_p	SR
ADCCXS-OF	2.85%	13.60%	0.21	0.74%	9.86%	0.07	8.70%	8.63%	1.01
ADCCXS –IR	2.62%	14.18%	0.18	0.57%	11.09%	0.05	8.05%	9.50%	0.85
CCC	0.61%	13.64%	0.04	1.39%	12.12%	0.11	4.83%	7.71%	0.63
ADCCS	1.38%	14.22%	0.09	0.42%	10.04%	0.04	8.14%	8.63%	0.94
DCC	1.75%	14.18%	0.12	0.48%	10.03%	0.04	7.49%	8.56%	0.88
EWMA	1.98%	14.38%	0.13	1.09%	12.08%	0.09	7.52%	9.86%	0.76
BEKK	2.85%	13.60%	0.20	0.64%	9.86%	0.06	8.17%	8.83%	0.93
STATIC	2.62%	14.18%	0.18	0.57%	11.09%	0.05	8.05%	9.50%	0.85

Note: This table evaluates the aim-for-the-target approach where the investor trades towards the optimal portfolio (at a rate of ~ 10%) derived from equation 3.6. Mean portfolio return (μ_p), volatility (σ_p) and Sharpe ratios (SRs) are annualized and do take into account transaction costs. We employ the Jobson-Korkie (1981) test to evaluate whether the differences in Sharpe Ratios generated from the dynamic and static predictions are statistically different from zero. *** represents significance at 1% level, ** at the 5% and * at the 10% level.

Based on the analyses in Section 3.7, we conclude that the results have important implications for traders. A successful investor should use order flows to forecast (1) exchange rate movements,[64] (2) exchange rate co-movements and (3) volatility dynamics, especially when facing financial turbulence. All the three steps combined would ensure that the risk-adjusted returns from our proposed portfolio optimisation strategies will be the highest, thanks to the benefits of order-flow-based return and covariance forecasts.

[64] Note that order-flow-based return predictions are the key input in the portfolio optimisation strategy displayed in Sections 3.7.2 and 3.7.3. We experimented with interest-rate-based return predictions; however, they did not result in positive Sharpe ratios for any strategy.

3.8 Conclusion

Previous studies have applied the traditional market microstructure theory in the univariate framework to forecast single exchange rate movements. The objective of this chapter is to extend those studies to forecasting co-movements of exchange rates.

The chapter begins by empirically examining the link between the absolute order flow differential and the correlation dynamics of spot exchange rates. We find that the latter depend significantly negatively on the former. This result is robust to controlling for interest rates, a macroeconomic variable commonly used to predict to exchange rate co-movements. Moreover, our research also unambiguously suggests that the IRD has no significant impact on the co-movements after allowing for the order flow differential. Therefore, we conclude that all the information content about changes in the interest rate is already incorporated in the order flows.

We then evaluate the predictive accuracy of order flows for future covariance dynamics. We find that order flows can significantly improve the forecasting accuracy of a dynamic conditional correlation model. The order-flow-based conditional correlation model outperforms all the other models commonly used in the literature in terms of statistical forecasting accuracy, regardless of the time periods and/or exchange rates considered. We also show that our order-flow-based covariance forecasting approach has a particular advantage for periods afflicted with economic or financial instability. This is in line with Rime and Tranvåg's (2012) claim that order flows are most informative during an economic downturn.

Having compared alternative covariance forecasting approaches, we also examine the effects of window size on forecasting performance. We find that a medium-sized window provides the highest accuracy for the covariance predictions of all the models used. Although

the differences between the medium and either a short or a long rolling window are mostly insignificant for economically tranquil times, the differences are large and statistically significant for economically turbulent periods. This indicates that a medium-sized estimation window is particularly useful for covariance forecasting during troubled time periods.

Our work also highlights the fact that order flows can enhance univariate volatility forecasts, as order flows capture jumps in the underlying volatility. The model used to yield this result is labelled the GARCH-X model. Comparing GARCH-X- and GARCH-based conditional correlation models, the former considerably outperform the latter, as judged by the MSEs, for the three subsample periods and for the world's three major exchange rate pairs. Put differently, the information content incorporated in order flows can improve second-moment forecasts alongside covariance forecasts. This result is new, as it adds to previous results on return forecasts (first-moment forecasts) reported in the existing literature.

Finally, this chapter investigates the economic value of accurate covariance forecasts for portfolio optimisation, which is of practical relevance to investors. Our proposed approach leads to higher Sharpe ratios than the constant covariance models commonly used in the literature. Tackling the transaction costs associated with frequent rebalancing, we propose a variant of the mean-variance optimisation setting. This setting involves trading towards the optimal portfolio at a rate determined by the transaction costs. Again, the order-flow-based covariance forecasts outperform all other competing models in terms of the Sharpe ratios. This result indicates the robustness of our approach and highlights the economic value for investors, as the Sharpe ratio gains remain positive even after accounting for transaction costs.

In a nutshell, our results confirm that order flows have significant effects on future correlation dynamics and thus can be used to enhance the forecasting accuracy of

multivariate conditional correlation models. The high predictive accuracy brings nontrivial trading gains to a mean-variance investor. Although this chapter gives a first glance at the economic benefits of using order flows, the next chapter will further asses the practical value of order flows for investors.

Appendix B

B.1 Additional tables

Table B.1 Parameter estimates of the GARCH-ADCCXS and GARCH-ADCCXE models

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
Panel A. ADCCXE model parameters $\varepsilon_t = [\varepsilon_{1,t} + \gamma_1 \varepsilon_{2,t} + \gamma_2])$								
α_1	0.2223 **	0.1510 **	0.1625 **	0.1989 **	0.1645 **	0.2228 **	0.1735 **	0.1957 **
α_2	0.1201 **	0.3147 **	0.0947 **	0.1517 **	0.0936 **	0.1706 **	0.1366 **	0.1623 **
β_1	0.9727 **	0.9879 **	0.9864 **	0.9795 **	0.9864 **	0.9749 **	0.9791 **	0.9703 **
β_2	0.9817 **	0.9492 **	0.9913 **	0.9810 **	0.9906 **	0.9730 **	0.9890 **	0.9867 **
γ_1	0.5779 **	0.1738 **	0.1232 **	0.0316 **	0.0244	0.3611 **	-0.1256 **	0.0902 **
γ_2	0.3556 **	0.5952 **	0.0294	0.0330 **	0.0398	0.1118 *	-0.0735 **	0.0915 **
η	-0.0128 **	-0.0346 **	-0.0070 *	-0.0372 **	-0.0050	-0.0598 **	0.0206 **	-0.0152 **
LRT	16.77 **	1.30	15.51 **	36.58 **	21.81 **	271.60 **	5.52 **	12.48 **
LLF	-4173.24	-5418.6	-4887.7	-5302.9	-4764.6	-6159.4	-5351.9	-5462.2
Panel B. ADCCXS model parameters $(n_t = I(\varepsilon_t > 0) \circ \varepsilon_t)$								
α_1	0.1148 **	0.1513 **	0.1475 **	0.2447 **	0.1467 **	0.2511 **	0.1671 **	0.1881 **
α_2	0.1475 **	0.3164 **	0.0874 **	0.1704 **	0.0942 **	0.2282 **	0.1378 **	0.1618 **
β_1	0.9732 **	0.9878 **	0.9881 **	0.9695 **	0.9882 **	0.9675 **	0.9771 **	0.9734 **
β_2	0.9826 **	0.9486 **	0.9934 **	0.9708 **	0.9915 **	0.9602 **	0.9834 **	0.9845 **
γ_1	0.1995 **	0.0876 **	0.0431 **	0.4176 **	0.0164	0.0372 **	0.1071 **	0.0140
γ_2	0.0875 **	0.0636 **	0.0000	0.1334 **	0.0000	0.1611 **	0.1179 **	0.0399 **
η	-0.0125 **	-0.0394 **	-0.0057 *	-0.0704 **	-0.0044	-0.0313 **	0.0064 **	-0.0142 **
LRT	10.28 **	0.13	0.03	32.12 **	21.81 **	282.71 **	5.89 **	10.59 **
LLF	-4177.35	-5436.5	-4895.4	-5308.7	-4764.6	-6153.9	-5350.6	-5468.5

Note: Panel A of this table presents the parameter estimates of Equation 2.9 for eight possible combinations of exchange rate pairs using daily return and order flow data. Panel B reports the parameter estimates for the ADCCXS model given by Equation 2.8. The estimation (and computation of the modified standard errors) uses the two-stage procedure proposed in Engle and Sheppard (2001). LRT denotes the likelihood ratio test statistic testing the null hypothesis of γ_1=γ_2 =0. LLF is the quasi-log-likelihood function evaluated at the maximum. ** and * denote significance at the 1% level and 5% level respectively.

Table B.2 Robustness of results: order flow and bid–ask spread

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
ADCCXE (2.11) model parameters								
α_1	0.186 **	0.152 **	0.1558 **	0.2253 **	0.1485**	0.2970**	0.1708**	0.1931 **
α_2	0.122 **	0.315 **	0.0825 **	0.1726 **	0.0942 **	0.2049 **	0.1394 **	0.1654 **
β_1	0.982 **	0.987 **	0.9878 **	0.9743 **	0.9879 **	0.9549 **	0.9797 **	0.9705 **
β_2	0.989 **	0.949 **	0.9938 **	0.9716 **	0.9913 **	0.9706 **	0.9877 **	0.9857 **
γ_1	0.537 **	0.175 **	0.1208 **	0.3592 **	0.0013	0.2480 **	-0.0825**	0.0436 **
γ_2	0.248 **	0.593 **	0.0010	0.1127 **	0.0483 **	0.0681 **	-0.0359**	0.0677 **
η_1	-0.021 **	-0.045 **	-0.0071 **	-0.0725 **	-0.0052	-0.0226 **	0.0050	-0.0069 **
η_2	-0.012	-0.099 **	0.0003	-0.0100 **	0.0173	-0.0393 **	0.0091 **	-0.0099
LLF	-4171.5	-5418.4	-4886.5	-5303.7	-4757.6	-6152.8	-5353.1	-5459.5

This table presents the parameter estimates of the ADCCXE model of by Equation (2.11), focusing on daily return co-movements. LLF is the quasi-log-likelihood function evaluated at the maximum. ** and * denote significance at the 1% and 5% levels respectively. The parameter η_1 is associated with the absolute order flow differential, whereas η_2 corresponds to the bid–ask spread.

Table B.3 Robustness of findings: standardised measures of order flow

	USD–EUR USD–GBP	USD–EUR USD–JPY	USD–EUR USD–CAD	USD–GBP USD–JPY	USD–GBP USD–CAD	USD–JPY USD–AUD	USD–NZD USD–AUD	USD–JPY USD–CAD
Panel A. ADCCXE model parameters using the volatility-adjusted standardised order flows								
α_1	0.2180 **	0.1511 **	0.1783 **	0.2423 **	0.1474 **	0.3109 **	0.1732 **	0.1766 **
α_2	0.1211 **	0.3145 **	0.3185	0.2112 **	0.0936 **	0.2312 **	0.1364 **	0.1901 **
β_1	0.9739 **	0.9879 **	0.9663 **	0.9702 **	0.9882 **	0.8948 **	0.9786 **	0.9801 **
β_2	0.9890 **	0.9493 **	0.9472 **	0.9485 **	0.9916 **	0.9478 **	0.9885 **	0.9763 **
γ_1	0.5550 **	0.1739 **	0.1408 **	0.7462 **	0.0234 **	0.5924 **	-0.1506	0.1101
γ_2	0.2810 **	0.5974 **	0.3182	0.2379 **	0.0478	0.4387 **	-0.0991	0.0973
η	-0.0142 **	-0.0039 **	-0.0037 **	-0.0330 **	-0.0010 **	-0.0331 **	-0.0045	-0.0014 **
LRT	15.27 **	33.68 **	13.68 **	24.71 **	35.16 **	218.80 **	2.69	8.55 **
LLF	-4174.8	-5419.4	-5428.3	-5312.4	-4757.9	-6186.0	-5356.4	-5464.2
Panel B. ADCCXE model parameters using the trading-intensity-adjusted standardised order flows								
α_1	0.2969 **	0.1499 **	0.0924 **	0.2162 **	0.1440 **	0.3020 **	0.1757 **	0.1509 **
α_2	0.1597 **	0.3075 **	0.0687 **	0.1686 **	0.0998 **	0.2063 **	0.1299 **	0.2128 **
β_1	0.9477 **	0.9887 **	0.9565 **	0.9764 **	0.9886 **	0.9533 **	0.9787 **	0.9850 **
β_2	0.9592 **	0.9515 **	0.9538 **	0.9780 **	0.9912 **	0.9716 **	0.9901 **	0.9769 **
γ_1	0.4206 **	0.1666 **	0.4247 **	0.2812 **	0.0264 **	0.2418 **	-0.2316	0.1056 **
γ_2	0.2250 **	0.5981 **	0.4802 **	0.0526 **	0.0831	0.0529 **	-0.1464	0.0200 **
H	-0.2391 **	-0.0376 **	-0.0236 **	-0.0054 **	-0.0037 **	-0.0176 **	0.0330	-0.0016
LRT	12.85 **	33.18 **	15.76 **	34.12 **	32.58 **	281.43 **	5.98 *	11.44 **
LLF	-4187.4	-5419.6	-5426.3	-5307.6	-4758.4	-6154.5	-5351.8	-5462.9

Note: This table presents the parameter estimates of the ADCCXE model, using respectively the volatility-based standardisation of order flow (Panel A), and the trading-intensity-based standardisation of order flow (Panel B). Daily data is used in the estimation. The estimation and computation of the modified standard errors use the two-stage procedure proposed in Engle and Sheppard (2001). LLF denotes the log-likelihood value evaluated at the maximum. ** and * denote significance at the 1% and 5% levels respectively.

B.2 Competing forecasting approaches

This section outlines the competing models in detail. The covariance forecasts for all DCC-type models come from the basic DCC elements given in Engle (2002) , where the covariance is a function of the underlying volatilities and conditional correlations between assets:

$$H_t = D_t R_t D_t , \tag{3.13}$$

Both the conditional correlations R_t and the conditional volatilities D_t at time *t* depend on lagged variables, implying that R_t and D_t are one-step-ahead forecasts made at time *t*-1 for time *t*. For the diagonal elements (h_{it}) of the conditional variance matrix D_t, the evolution follows a GARCH process:

$$h_{it} = \omega_i + \delta_i z_{it-1}^2 + \theta_i h_{it-1}, \qquad i = 1, 2 \tag{3.14}$$

We assume that:

$$z_t \mid \Omega_{t-1} \sim N(0, H_t)$$ [65]

The time-varying correlation matrix R_t evolves according to:

$$R_t = (diag(Q_t))^{-1} \cdot Q_t \cdot diag(Q_t))^{-1} \tag{3.15}$$

where $Q_t = \begin{pmatrix} q_{11t} & q_{12t} \\ q_{12t} & q_{22t} \end{pmatrix}$ represents a 2×2 conditional covariance matrix of standardised residuals driving the time-varying correlation. Performance differences between the DCC and its extensions arise from different evolutions of Q_t which prompts us to evaluate these

[65] This assumption is not crucial, as the results have a standard quasi-maximum likelihood estimate interpretation in the absence of a normal distribution.

models to see how differently they perform in forecasting. Therefore, in what follows, we describe different ways of modelling Q_t The family of dynamic conditional correlation models evaluated consists of DCC, ADCCS and the ADCCXS models for which we show the dynamics below.

The DCC model assumes that Q_t evolves according to:

$$Q_t = (\overline{Q} - A'\overline{Q}A - B'\overline{Q}B) + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + B'Q_{t-1}B \quad (3.16)$$

where $A = \begin{pmatrix} \alpha_1 & 0 \\ 0 & \alpha_2 \end{pmatrix}$ and $B = \begin{pmatrix} \beta_1 & 0 \\ 0 & \beta_2 \end{pmatrix}$ and $\overline{Q} = T^{-1}\sum_{t=1}^{T}\varepsilon_t\varepsilon_t^{'}$, $\alpha_1^2 + \beta_1^2 < 1$ and $\alpha_2^2 + \beta_2^2 < 1$ are the restrictions used to guarantee positive definiteness, and ε_t represents the standardised (de-volatised) return residuals.

The ADCCS takes asymmetry into account in the response of correlation to joint shocks. Q_t is assumed to evolve according to:

$$Q_t = (\overline{Q} - A'\overline{Q}A - B'\overline{Q}B - G'\overline{N}G) + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + Gn_{t-1}n'_{t-1}G + B'Q_{t-1}B \quad (3.17)$$

where $G = \begin{pmatrix} g_1 & 0 \\ 0 & g_2 \end{pmatrix}$, $\overline{N} = T^{-1}\sum_{t=1}^{T} n_t n_t^{'}$ [66] and all other matrices are defined as before, and $\alpha_1^2 + \beta_1^2 + g_1^2 < 1$ and $\alpha_2^2 + \beta_2^2 + g_2^2 < 1$ are parameter restrictions.

For the constant conditional correlation model, the conditional correlation driving process is replaced by the sample average: $Q_t = \overline{Q}$ (3.18)

[66] n_t represents the standardised (de-volatised) negative or positive return residuals.

In the ADCCXS-OF model, conditional correlations are driven by order flow as an exogenous variable and the corresponding Q_t evolves according to:

$$Q_t = (\overline{Q} - A'\overline{Q}A - B'\overline{Q}B - G'\overline{N}G - \eta\overline{X}) + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + Gn_{t-1}n'_{t-1}G + B'Q_{t-1}B + \eta \mid x_{1t-1} - x_{2t-1} \mid \quad (3.19)$$

where $\mid x_{1t-1} - x_{2t-1} \mid$ denotes the absolute order flow differential, $\overline{X} = T^{-1}\sum_{t=1}^{T}\left|x_{1t-1} - x_{2t-1}\right|$, and all the other matrices and restrictions are the same as given above.

The ADCCXS-IR model has the same correlation dynamics as the ADCCXS-OF model, with the only difference being that order flow is replaced by the interest rate. Therefore, Q_t evolves according to:

$$Q_t = (\overline{Q} - A'\overline{Q}A - B'\overline{Q}B - G'\overline{N}G - \eta\overline{X}) + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + Gn_{t-1}n'_{t-1}G + B'Q_{t-1}B + \eta \mid x_{1t-1} - x_{2t-1} \mid \quad (3.20)$$

where $\mid x_{1t-1} - x_{2t-1} \mid$ denotes the absolute IRD, $\overline{X} = T^{-1}\sum_{t=1}^{T}\left|x_{1t-1} - x_{2t-1}\right|$,and all other matrices and restrictions are the same as given above.

The ADCCXE-OF model proposed by Li (2011) has the same correlation driving process as the ADCCXS model, with the only difference occurring in the way asymmetry is modelled:

$$Q_t = (\overline{Q} - A'\overline{Q}A - B'\overline{Q}B - \eta\overline{X}) + A'e_{t-1}e'_{t-1}A + B'Q_{t-1}B + \eta \mid x_{1t-1} - x_{2t-1} \mid \quad (3.21)$$

where, $e_t = \begin{pmatrix} \varepsilon_{1t} + \gamma_1 \\ \varepsilon_{2t} + \gamma_2 \end{pmatrix}$, and all the other matrices are the same as given above.

Estimation of all the GARCH and DCC parameters in the conditional correlation models is done using the two-stage maximum likelihood procedure proposed in Engle (2002).

Similar to the dynamic conditional correlation models, covariance dynamics in the BEKK-type models depend on lagged variables, which naturally generate one-step-ahead covariance forecasts. We consider a diagonal BEKK model. The evolution of the covariance matrix H_t is given by:

$$\mathrm{H_t} = \mathrm{C_0 C_0'} + \mathrm{A'}\,\varepsilon_{t-1}\varepsilon_{t-1}'\mathrm{A} + \mathrm{B'H_{t-1}B}\,, \qquad (3.22)$$

where C_0 is a lower triangular matrix, and *A* and *B* are set as the diagonal matrices:

$A = \begin{pmatrix} \alpha_1 & 0 \\ 0 & \alpha_2 \end{pmatrix}$ and $B = \begin{pmatrix} \beta_1 & 0 \\ 0 & \beta_2 \end{pmatrix}$.

The main difference between the BEKK and DCC covariance models is that the BEKK-type models model covariance dynamics directly, whereas in the DCC-type models, the covariance dynamics follow from the correlation and variance dynamics. Estimation of all the BEKK parameters is done via maximum likelihood.

In the risk metrics smoother or EWMA, covariance evolves according to:

$$\mathrm{H_t} = (1-\lambda)\,\varepsilon_{t-1}\varepsilon_{t-1}' + \lambda\,\mathrm{H_{t-1}}\,, \qquad (3.23)$$

where λ is the smoothing parameter, typically set at 0.94 and 0.96. Similar to the BEKK models, the risk metrics approach models covariance dynamics directly and focuses on return shocks and persistence. Unlike the BEKK, however, the parameters in the risk metrics approach are not estimated using maximum likelihood: λ is set by the researcher (as 0.94 in this paper).

4. Chapter Four: Order Flow as Technical Trading Signal

4.1 Introduction

Technical analysis involves the use of historical market data such as price, volume and other variables to predict future asset movements. This technique is widely used by practitioners, and almost all FX professionals use technical analysis as a tool in decision making to some degree (Menkhoff and Taylor, 2007). Given the large practical relevance of technical trading, it is not surprising that the performance of indicators that rely on past prices or volumes to predict future exchange rate movements have been studied in much detail in the literature (e.g., Qi and Wu, 2006; Hsu, Taylor and Wang, 2016).

In contrast to price- and volume-based trading indicators, this chapter evaluates the performance of order-flow-based technical trading indicators in FX trading. Furthermore, this chapter quantifies the economic value of monitoring the underlying market volatility to justify the trading decisions arising from order-flow-based trading indicators.

The primary motivation for this study arises from the common practice of evaluating the economic benefits of order-flow-based exchange rate predictions in the context of optimizing currency portfolios (e.g. King et al., 2010; Rime et al., 2010). This is somewhat puzzling, as order-flow information is incorporated into prices into prices rapidly – within minutes and days rather than weeks (Berger et al., 2008), making order flow a powerful variable for predicting short run exchange rate fluctuations. Since the benefits of using order-flow-based return forecasts are related to predicting short-term exchange rate fluctuations, one may cast doubts on their relevance to portfolio optimization which is typically a long-term investment strategy. By contrast, technical traders usually change position sizes at daily or even intraday intervals (Neely and Weller, 2003; Marshall, Cahan and Cahan, 2008).

Therefore, the economic value of order flow information should be particularly relevant for technical traders.

However, the literature on order-flow based trading indicators is scarce, reporting mixed results regarding the profitability of trading signals derived from order-flow for stock markets (Yamamoto, 2012; Chordia and Subrahmanyam, 2004). The same applies to forex applications, where, to the best of our knowledge, the only profitable order-flow-based technical trading indicators use order flow information to either derive a return- (Gradojevic, 2007) or momentum (Gradojevic and Lento, 2015) prediction rather than deriving a buy or sell indicator directly. The lack of applying other modelling approaches is puzzling. Neely and Weller (2012) note that "[A]n empirical challenge for technical analysis – one that has not been successfully met so far – would be to link order flow to technical trading signals and returns."

In view of this call, therefore, the first objective of this chapter is to analyse the case where one links order flow directly to technical trading signals; that is, to use order flow as a buy or sell indicator instead of using order flow to generate a return prediction as a trading indicator. This setting is akin to Yamamoto's (2012) analysis of selected Nikkei stocks and evaluates whether the signal conveyed in order flows is robust enough to generate profitable trading recommendations. We deem this analysis worthwhile, as the relationship between order flow and price fluctuations is stronger in FX markets than in stock markets.[67] Note, however, that there are several problems with using order flows to generate trading signals directly rather than using a return prediction based on order flow information as a trading signal. One reason for this is that the relationship between order flow and exchange rate movements is probably nonlinear (Cerrato et al., 2015), which means that a larger absolute order flow does not necessary result in a larger fluctuation in an exchange rate. Several other

[67] A detailed discussion on this follows in Section 4.3.

aspects, to be discussed later on, allow us to conjecture that deriving trading signals directly from order flow will not provide substantial trading profits. The analysis confirms the need for a more sophisticated setting, as no technical trading indicator, regardless of whether the indicator uses order-flow- or price-based information, can significantly outperform a buy-and-hold strategy for all exchange rates in the sample.

After demonstrating that using order flows to generate trading signals directly does not work out, we consider a multivariate trading application that incorporates an order-flow-based exchange rate change predication alongside a proxy for the underlying market volatility. The methodology used builds upon the neuro-fuzzy framework proposed in Gradojevic (2007), using order flows to forecast future exchange rate returns. This strategy is shown to be highly profitable for the USD–CAD spot rate. While confirming that the results of Gradojevic (2007) are valid for other exchange rate combinations as well, the novelty of our approach is that we use the change in the underlying volatility to support or contradict the order-flow-based return prediction. A fuzzy logic inference setting is used to link both variables to a trading recommendation.

Our proposed approach significantly outperforms the neuro-fuzzy approach that relies on order-flow-based exchange rate predictions only. Notably, the profitability of our approach is stunning, leading to annualised Sharpe ratios up to 5.43 for an USD–EUR and 5.52 for an USD–NZD trader. The incremental performance is solely caused by the adjustment in trading weights if the volatility change contradicts the magnitude or sign of the order-flow-based return prediction.

Moreover, this finding is highly robust. Neither the inclusion of transaction costs, nor evaluation of different time periods, affects the superiority of our approach for five out of the six currency pairs under investigation. The performance of the approach is further robust to

different shapes of the membership functions used in the fuzzy logic inference setting or differences in accuracy of the order-flow-based exchange rate change predictions.

As an extension of our approach, we explore the question of whether the performance of a classic price-based indicator can be boosted by linking its trading recommendation to the underlying volatility. This was in fact already implied by Gradojevic and Gençay (2013). The two authors note that a moving average (MA) indicator performs better during more volatile days. Exploring along this line, we show that there are significant trading gains arising from linking the trading signal generated by the moving average indicator to the change in the underlying market volatility. This indicates that the performance of both price-based and order-flow-based forex trading strategies can be boosted significantly by combining their trading signals with a proxy for the underlying volatility.

4.2 Related Literature

Technical analysis has been widely studied and evaluated in the literature because of its high practical relevance (Gehrig and Menkhoff, 2006; Menkhoff, 2010), yielding a variety of different conclusions regarding the actual value of technical analysis. While early studies typically document limited profitability of technical trading indicators (e.g. Fama and Blume, 1966), the findings from more recent studies are mixed, with several studies documenting that technical trading indicators can indeed generate substantial trading profits (Brock, Lakonishok and LeBaron, 1992; LeBaron, 1999;Qi and Wu, 2006).

Summarising all the contradictory findings, Park and Irwin (2007) , found that 56 out of 95 "modern" studies on technical analysis produced supportive evidence of its profitability. It must be noted, however, that technical trading indicators seem to be particularly profitable for foreign exchange markets (Neely and Weller, 2013). Although this documented profitability contradicts the efficient market hypothesis (Fama, 1970), the adaptive market hypothesis outlined by Lo and MacKinlay (2002) is suitable for explaining the profitability of using technical trading indicators in FX trading.[68]

Survey evidence suggests, however, that in FX exchange trading, technical trading indicators are seldom used in isolation and are often accompanied by order flow analysis (Gehrig and Menkhoff, 2004). The motivation for traders to monitor order flows arises from the strong predictive power of order flows for future exchange rate movements and, to some extent, from feedback trading affecting future price movements (Gradojevic and Lento, 2015). Further motivation can be drawn from the hypothesis that order flow drives price trends by gradually incorporating private information within prices (Neely, Weller and

[68] See Neely, Weller and Ulreich (2009) for a review of studies confirming the adaptive market hypothesis specifically related to technical trading strategies in FX.

Ulreich, 2009). This means that the underlying order flow results in the price patterns technical traders try to exploit.

In contrast to all the aforementioned literature implying the value of order-flow-based trading indicators, actual empirical applications of order-flow-based technical indicators have been somewhat neglected in the literature, except for the studies of Gradojevic (2007) and Yamamoto (2012), two studies our approach builds upon. Yamamoto (2012) reported that neither price-based nor order-flow-based trading indicators can outperform a buy-and-hold strategy after considering data-snooping biases. The findings might imply that while order flow is useful for predicting future exchange rate fluctuations, it is a too noisy proxy for deriving buy and sell recommendations directly. In line with this argument, Gradojevic (2007) used order flow to predict future exchange rate fluctuations and used a fuzzy logic control setting that assists in generating the buy and sell recommendations, an approach that has been shown to be highly profitable.

Nonlinear trading models – such as the neuro-fuzzy approach mentioned above – have gained popularity in academic research, as they are frequently used by professionals. Apart from evaluating the statistical accuracy of nonlinear asset return predictions (Gençay, 1998; Gencay and Stengos (1998); Cao, Leggio, Schniederjans, 2005) the literature has highlighted the profitability of trading models relying on nonlinear return predictions (Dunnis and Williams, 2002; Fernández-Rodríguez, González-Martel, and Sosvilla-Rivero 2000; Jasic and Wood, 2004).

The economic value of the nonlinear return predictions is commonly evaluated using a simple switching strategy: positive predicted returns are executed as long positions and negative returns as short positions. Fernández-Rodríguez et al. (2000) and Jasic and Wood (2004) proved the profitability of this trading rule for stock market index trading, where the

return signs are estimated by a feed forward neural network, a class of artificial neural networks (ANN). The inputs in the neural network correspond to the daily returns over the previous days, which are then used by the network to learn[69] the underlying relationship between lagged and current returns, and to generate return predictions. Using the same switching strategy, Dunnis and Williams (2002) confirm the profitability of neural network based return predictions for an USD–EUR trader.

A shortcoming of this simple regime switching approach, however, is that the trading indicator can only take three values: 1 which represents a long position, 0 which represents a neutral position, and -1 which represents a short position. This incurs frequent rebalancing and hence transaction costs, and further substantial losses in case the return prediction is slightly imprecise. As a remedy, Gradojevic (2007) proposed a fuzzy logic inference technique which we follow in this chapter. The fuzzy logic inference technique links the magnitude and sign of the network based return predictions to the trading recommendations, thereby reducing transaction costs and trading uncertainty. Gradojevic (2007) found that this strategy is highly profitable for a USD–CAD trader, a finding confirmed by Lee and Tsung (2007) for a trader specialising in USD–AUD.

In a closely related paper, Bekiros (2011) proposed a volatility-based-neuro-fuzzy model (VNF), which generates trading recommendations based on return and volatility information. This study builds up on earlier findings in Bekiros and Georgoutsos (2008) highlighting that S&P 500 volatility information enhances the directional accuracy of neural network based S&P 500 index return predictions. The author showed that a trading strategy based on the VNF model outperforms several other nonlinear trading models, with the performance gains arising solely due to incorporating a proxy for the underlying volatility.

[69] We will elaborate on the learning process in the methodology section.

The motivation for incorporating estimates of the underlying volatility in the trading strategy arises from the documented negative relationship between returns and conditional (implied) volatility (Bekaert and Wu, 2000; Christoffersen and Diebold (2006); Giot (2005); Hibbert et al. (2008)).

Our methodology is in the same spirit but has several noteworthy differences. The approach outlined in Bekiros (2011) uses return and volatility information as inputs in a fuzzy logic inference setting. The fuzzy logic setting then configures the membership functions for both returns and volatility, where the parameters of the membership functions are optimized via neural network training. By contrast, our approach uses a neural network based return prediction and the underlying market volatility as inputs in the fuzzy logic inference setting. The fuzzy logic rule base and membership functions are constructed in advance and do not adjust to the data. This setting allows to link order flows to exchange rate return predictions and trading recommendations, while controlling for changes in the underlying market volatility.

4.3 Research Question and Hypotheses

Based on the ability of order flow to both explain and forecast exchange rate movements, one would be tempted to think that a simple technical trading strategy deriving buy and sell signals from order flow may outperform other traditional price-based strategies. The intuition is that order flow creates the patterns that price-based trading strategies try to exploit. In its most simplest from, order flows could be used to generate a buy or sell recommendation based on whether the daily order flow is positive or negative.

However, there may be several potential problems associated with using order flows directly as a trading indicator. First and foremost, it is enormously difficult to find the

appropriate trigger value. The trigger value essentially specifies a critical value that leads to a buy, a hold or a sell recommendation. Clearly, to set a trigger value, one needs some sort of benchmark, such as the maximum or minimum price jump, for the entire sample period. If the trigger value is set too high, the indicator might miss profitable investment opportunities; if it is set too low, the signal might be noisy, thereby resulting in losses. Finding an appropriate trigger value is a challenge in choosing an order-flow-based indicator, as there are no clear benchmarks. The reason for this is that the distribution of order flows shows massive peaks (because of its non-normal distribution). Second, achieving significant trading gains requires information on both the directional accuracy and difference in magnitude, which poses a problem for any non-price-based indicator.[70]

The above considerations lead us to hypothesise that any technical trading strategy using the sign of order flow (i.e., the sign of the net of net buyer- and seller-initiated trades), as a trading indicator will not outperform traditional price-based strategies:

H1: Technical trading indicators deriving a buy or sell recommendation from order flows will not outperform price-based technical trading indicators.

Though we do not expect that order flows can be used directly as an indicator, recent literature on nonlinear modelling has shown that order flows can be used to obtain an accurate prediction of spot rate movements. An example of this approach was provided by Gradojevic (2007), who used a neuro-fuzzy trading strategy, as mentioned in the literature review. We also use this approach; however, we hypothesise that the performance of this approach can be substantially boosted from incorporating the underlying volatility in the

[70] By comparing a simple filter rule with an order-flow-based indicator, we can further illuminate this argument. In a filter rule, a trading signal is derived if the current price is above or below a certain percentage of the previous high or low. Therefore, this indicator will give a direction (positive or negative) and an indication of how big the future movement is going to be (measured by the percentage difference). On the other hand, an order-flow-based indicator will show the buying or selling pressure, but it is ambiguous as to whether a larger positive or negative order flow will indeed trigger a larger appreciation or depreciation.

fuzzy logic setting, a motivation arising from Bekiros (2011) and Christoffersen and Diebold (2006).

The "sign-dependence" theory of Christoffersen & Diebold (2006) suggests that volatility dependence produces sign dependence, and therefore directional predictability, as long as expected returns are non-zero. The intuition behind this theory is that volatility changes will alter the probability of observing negative or positive returns. Specifically, the ''higher'' the volatility, the ''higher'' is the probability of a negative return arising, given that the expected returns are positive[71].

Following this theory, we hypothesise about the probability of the exchange rate change[72] prediction having the correct sign based on current increases or decreases in the underlying volatility. As all our exchange rates are defined as the USD price of one unit of foreign currency, we hypothesise that an increase in the underlying market volatility[73] results in a lower future return from holding the foreign currency[74].

In other words, a positive (negative) predicted return, i.e. a depreciation (appreciation) of the USD relative to the foreign currency, is more likely to be correct (or to have the correct sign) if the underlying market volatility decreases (increases). The probability of the return prediction having the correct sign will be largest for very high positive (negative) predicted returns accompanied by large decreases (increases) in the underlying volatility and will be smallest for small values of positive (negative) predicted returns accompanied by large

Note that this "sign-dependence" theory applies to all financial markets, whereas other theories on the asymmetric return-volatility relationship typically apply to equity markets only.

[72] The exchange rate change prediction is equivalent to the predicted return (in USD) for holding one unit of foreign (non-USD) currency.

[73] The underlying market refers to the underlying market for a particular exchange rate.

[74] This hypothesis further links to the flight to quality or flight to liquidity phenomena associated with increases in risk or volatility.

increases (decreases) in the underlying volatility. Accordingly, the trading recommendations should be adjusted for very low probabilities of the return forecast having the "correct" sign.

Note that we do not aim to derive trading signals directly from volatility changes[75], but instead take both predicted returns and the underlying volatility into account. We therefore hypothesise:

H2: We expect the highest returns from order flow information in a neuro-fuzzy logic setting controlling for the underlying market volatility

The reason why we conjecture that this setting will provide the largest returns is as follows: The exact relationship and causality between volatility changes and future returns is not generally agreed upon. The fuzzy interaction allows us to empirically test different relationships by setting the rule base in advance and allowing us to specify the relationship in linguistic terms. Furthermore, the return prediction is possibly imprecise, incomplete or unreliable. The same applies to market volatility, which is hard to quantify in real terms. Fuzzy logic, by its very nature, tolerates uncertainty by defining variables as imprecise terms. .Other benefits include a smoother decision output, which is simply smoother trading recommendations, thereby reducing transaction costs.

A final key question is whether our proposed methodology can be used to boost the performance of simple price based trading strategies. Though we firmly believe that order-flow-based return forecasts are preferable over any other technical trading indicator, we are interested in whether even simple trading indicators could be improved by linking them to changes in the underlying market volatility. An interesting finding from Gradojevic and Gençay (2013) gives an initial indication of this claim, showing that a MA fuzzy logic control setting provides the higher intraday returns when the conditional volatility is higher. This

[75] As outlined in Gradojevic, Lento, and Wright (2007), the economic value of volatility based trading indicators, such as Bollinger bands, is questionable.

finding is in line with any risk premium considerations (Kho, 1996) and contradicts the findings in Reitz (2006), suggesting that increases in volatility deteriorate the performance of technical trading indicators. Based on the findings in Gradojevic and Gençay (2013) we hypothesise:

H3: Even the performance of a simple MA indictor can be boosted significantly if we link the signal of the former to changes in the underlying market volatility.

To test this hypothesis, we will use the same fuzzy logic methodology that was used to evaluate the second hypothesis. This hypothesis is an add-on, and if proven correctly will show the increased benefits of linking changes in a currency pair's volatility to simple price-based technical trading signals.

In the next section, we describe the trading rules. The dataset used is the same as that in Chapter 4, with the addition of the oil price, CBOE VIX and EPU data which we obtain from Thomson Reuters DataStream. The results are reported and discussed in Sections 4.5, 4.6 and 4.7. Section 4.8 concludes the chapter.

4.4 Methodology

In the first step, we evaluate the performance of classic technical trading strategies and simple order flow strategies, compared with a simple buy-and-hold strategy, which serves as a benchmark. The main question is whether the simple order-flow-based indicator will outperform the buy-and-hold strategy, a common benchmark used in the finance literature.[76] The reason why the buy-and-hold strategy is commonly used as benchmark is because it supports the efficient market hypothesis: as the security is valued fairly all the time, it is best to buy and hold the security for long-run gains.

This initial setting lays the foundation for our fuzzy logic modelling, as it confirms that after accounting for transaction costs, no simple technical trading strategy can outperform a simple buy-and-hold strategy, thus showing the need of a more sophisticated soft computing approach. The focus on simple and classic technical trading strategies in the beginning also permits us to tie our work to that of Qi and Wu (2006). These two authors considered the benchmark to be a zero-return strategy and concluded that price-based trading indicators can significantly outperform this benchmark. However, in our opinion, a benchmark of zero is not practically relevant. We start with a short description of the technical trading rules to be evaluated; for a complete description of the price-based rules we refer the reader to Sullivan, Timmermann, and White (1999) or Qi and Wu (2006). After that, we introduce the multi-fuzzy strategy to be evaluated in this chapter.

[76] See, for example, Brock et al. (1992) and Gradojevic (2007).

4.4.1 Price- and order–flow-based technical trading rules

In the standard filter rule, a buy (sell) signal is extracted if the daily closing price (in US dollars) of a foreign currency (EUR, GBP, JPY,CAD, NZD, AUD) increases (decreases) by at least x percent from the previous low. For all rules considered, a buy signal means that the trader borrows in US dollars and buys the foreign currency. Once traders have observed a signal, they execute their trade and hold the position until the opposite signal appears. We define the previous low (high) as the lowest (highest) price achieved while holding a particular short (long) position. In addition, we consider an alternative definition of high (low), by calculating the highest (lowest) price over the previous periods *e*. Another version defines the rule as one that that generates a buy signal when the price increases by at least *x* percent from the previous low, but the sell signal is observed when the price decreases by at least *y* percent from the previous high. We also consider the case where a given long or short position is held for *c* days, during which period, all other signals are ignored.[77.]

The standard MA rule generates a buy signal when the foreign currency price (in US dollars) is higher than the moving average of prices over the previous *n* periods. This is similar to the filter rule: once a buy signal is received, the position is kept until the next sell signal has been extracted and vice versa. We also use a fixed percentage band filter that requires the buy or sell signal to exceed the MA by a fixed multiplicative amount, *b*. A second version of the MA rule focuses on the differences between slow and fast MAs with different window sizes. A buy (sell) signal is observed when a fast MA is larger than a slow MA, where a slow MA refers to the average over longer periods compared with a fast MA[78].

[77] The parameter values are given as follows: *x* = 0.005,0.01,0.015,0.02,0.025,0.03,0.035,0.04, 0.045, 0.05, 0.06, 0.07, 0.08,0.09,0.1,0.12,0.14,0.16,0.18,0.2,0.25,0.3, 0.4 and 0.5 [24 values]; 0.005, 0.01, 0.015, 0.02, 0.025, 0.03, 0.04, 0.05, 0.075, 0.1, 0.15 and 0.2 [12 values]; e=1,2,5,20 [five values] and *c* = 5, 10, 15, 25, and 50 [five values]. Given that *y* is less than *x*, this yields a total of 449 filter rules.

[78] The parameter values for moving average rules are given as follows: *n* = 2, 5, 10, 15, 20, 25, 50, 75, 100, 150, 200 and 250 [12 values]; *m* = 66, which refers to the number of fast–slow combinations of moving averages

As for the filter rules, we consider the case where traders take a position for *c* periods, ignoring any signal during these periods.

In a trading break rule, a buy signal is produced when the price of a foreign currency (in US dollars) exceeds the maximum price over the previous *tb* periods; a sell signal is produced when the price falls below the minimum price over the previous *tb* periods[79]. We considered a fixed multiplicative band filter *b* and a fixed holding period *c*, during which traders ignore any other signals. Channel breakout rules allow traders to buy (sell) and hold a long (short) position when the price exceeds (falls below) the channel, where a channel is produced when the high over the previous *cb* periods is within *xb* percent of the low over the previous *cb* periods[80].

The order flow based rules use the order flow in a currency pair to derive a buy or sell recommendation for that currency pair. For this version, a buy (sell) signal for a foreign currency is observed when its most recent order flow is positive (negative) by a fixed amount *o*. Once a trading signal has been observed, investors are assumed to trade immediately and to keep the position until the opposite signal appears. As before, we consider two variations of the trading rule. First, we allow for a fixed holding period *c* during which traders ignore any other signals. Second, in the spirit of the filter rules, we define the rule that a buy signal arises when the most recent order flow is positive by a fixed amount *o*, whereas a sell signal arises when the most recent order flow is negative by a fixed amount k. [81]

The total amount of technical trading rules amounts to 2691. To test whether the best trading rule can outperform a buy- and hold strategy after accounting for possible data-

based on the window sizes *n*; *b* = 0.001, 0.005, 0.01, 0.015, 0.02, 0.03, 0.04 and 0.05 [eight values], yielding a total of 1222 moving average rules.

[79] *tb* = 5, 10, 15, 20, 25, 50, 100, 150, 200 and 250 [10 values]. The total number of trading break rules is 300.

[80] *cb* = 5, 10, 15, 20, 25, 50, 100, 150, 200 and 250 [10 values]; *xb* = 0.005, 0.01, 0.02, 0.03, 0.05, 0.06, 0.07 and 0.08 [eight values]. Assuming *b* to be less than *xb*, the total number of channel breakout rules is 480.

[81] *o*=1, 2, 5, 7, 10, 15, 20, 25, 30, 40, 50, 60, 75, 100, 125, 150, 175 and 200 [18 values]. *k*=1, 2, 5, 10, 15, 20, 25, 30, 40, 50, 60, 75 and 100 [13 values]. Assuming that *k*<*o*, the total number of order-flow based rules is 240.

snooping biases, we use White's reality check (White, 2000) and Hansen's test for superior predictive accuracy (Hansen, 2005), both of which are outlined in detail in Section C2 in Appendix C.

4.4.2 Multi-fuzzy trading strategy

After investigating the simple trading strategies, we propose and evaluate a non-linear fuzzy logic based trading strategy. This strategy builds up on the methodology outlined in Gradojevic (2007), which combines an artificial neural network (ANN) based exchange rate change prediction with a fuzzy logic inference technique (denoted neuro-fuzzy approach) to derive trading recommendations. Novel to our setting is that we take into account the change in the underlying volatility additionally. We denote our methodology as multivariate fuzzy (multi-fuzzy) trading strategy and multivariate fuzzy logic framework. Figure 4.1 depicts a graphical representation of the proposed methodology.

Figure 4.1 Graphical representation of the multivariate fuzzy trading strategy

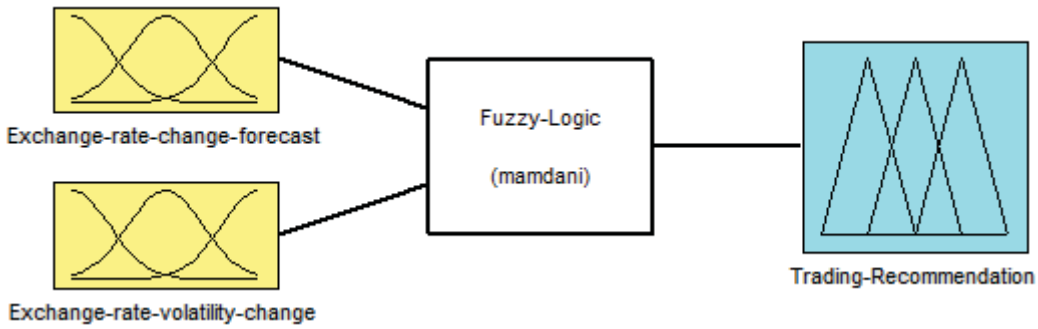

Note: This figure displays the architecture of our proposed multivariate fuzzy (multi-fuzzy) trading strategy. The fuzzy logic inference setting links two inputs, namely the predicted return and a proxy for the change in the underlying volatility (denoted "exchange-rate-volatility-change"), to a trading recommendation. The neuro-fuzzy approach outlined in Gradojevic (2007) generates the trading recommendation based on the exchange rate change prediction only and is hence a nested version of our proposed approach. The fuzzy logic inference setting will be outlined in detail later on.

The exchange rate change (return) prediction is based on an artificial neural network (ANN). ANNs have the advantage of detecting noisy, unstable and time-varying relationships between variables. This is relevant to providing accurate exchange rate forecasts and thus is part of the reason why ANNs have been used in virtually every research area. However, an ANN is, in some sense, a black box, as we can only set the input parameters and the design of the network (e.g., the neurons) but cannot explain how the system comes up with the optimal solution. By contrast, fuzzy logic operators have predefined membership functions, although it is possible to adapt the membership functions to the data (see the adaptive neuro-fuzzy logic models of Atsalakis and Valavanis (2009)). As such, it is clear how a decision is reached. However, their learning capability is limited compared with ANNs, as fuzzy logic operators can only learn about input–output patterns, whereas networks typically have a memory.

The ANN model produces exchange rate forecasts based on the following input/ output relationship:

$$\Delta s_{t=} f\left(\Delta of_{t-1,} \Delta oil_{t-1,} \Delta ir_{t-1,}\right), \qquad (4.1)$$

where Δs_t represents the change in the future exchange rate, Δof_{t-1} represents the order flow arising in a particular currency pair and Δir_{t-1} represents changes in the IRD between two economies. The order flows create buying or selling pressure for an exchange rate pair, thereby affecting changes in the spot price. The IRD is commonly used in predicting exchange rate movements, as differences in interest rates have an effect on the future exchange rate if the interest rate parity (IRP) holds. Akin to Gradojevic (2007), we incorporate the change in the oil price, $\Delta \mathrm{oil}_{t-1}$, in the ANN. The underlying rationale is that movements in commodity prices affect capital inflows and outflows and hence the exchange rate.

The networks trained and tested were the three-layer and four-layer feed-forward back-propagation networks. The number of input neurons was three (one neuron per input), while the number of hidden neurons varied between three and five. The last layer had one linear output neuron. To avoid overtraining, the ANN used an early stopping technique, where all the in-sample data was divided into three subsets: a training set (used for "learning", i.e. gradient calculation, weights and bias updating); a validation set (used to stop training when the error in this sets starts increasing) and a testing set (used to compare real and model output and to compare different network architectures). This division reduces the problem of overtraining and thus results in a network that can be generalised better to new unseen data. The training set consists of 60 percent of the total in-sample observations, with the other forty percent split equally between validation and testing set. The transfer functions (activation functions) in the hidden layers are set to be sigmoid and tan-sigmoid. Figure 4.2 graphically outlines the network architecture for the three layer feed-forward network tested.

Figure 4.2 A three-layer feed-forward neural network

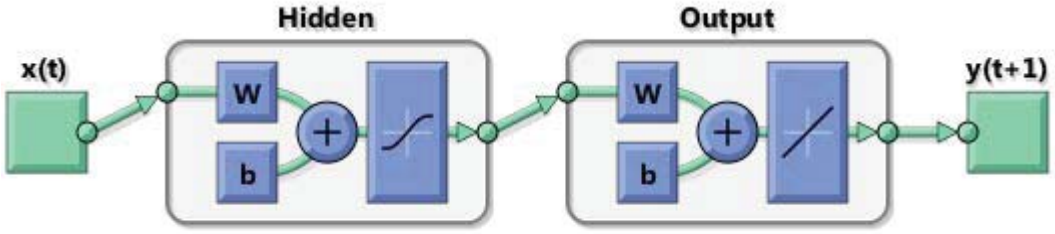

Note: This figure displays a common three layer feed-forward network employed. "x" represents the inputs in the input layer (order flow, changes in the oil price and IRDs). "Hidden" stands for the hidden layer, used by the neural network to memorize and "Output" represents the output layer. "w" and "b" represent the weights and biases assigned to the inputs in the hidden and output layers. The transfer (activation) functions displayed are sigmoid (in the hidden layer) and linear (in the output layer). "y" represents the output (the exchange-rate-change prediction) generated by the neural network.

The network development followed the following steps:

1) Setting the number of hidden layers, neurons, training algorithm (resilient backpropagation in our case), transfer functions, data division, initial connection weights, neuron biases and activation function for each neurons

2) Network training and validation. The first sixty percent of the data given to the network is used for training, i.e. to determine the connection weights and biases values within the network. The next 20 percent was used for validation. The network learns the relationship using the training set, and the validation set is used to monitor the network errors when the network is applied to new unseen data (i.e., data that are not used during training) and to halt training when the validation set's test error increases during training.

3) Estimation of the predicted output. The last 20 percent of the observations were then used by the trained ANN to generate exchange rate predictions. The accuracy of the exchange rate predictions is then compared to detect the optimal ANN architecture. Steps. 1-3 were repeated until the error goal (the minimum of the mean squared error in the testing set) is reached. The testing set is used to compare different network architectures, and the network architecture that provides the smallest MSE during the testing set is used to generate out-of-sample predictions. In order to evaluate a large number of initial connection weights, we evaluate a total of 30 trials with three to five hidden neurons. We then pick the one that yields the lowest MSE during the testing set.

As a proxy for the underlying volatility, we follow Bekiros (2011) and use a simple GARCH (1,1) model to estimate the conditional exchange rate volatility. The change in the

underlying volatility is then constructed as the difference between the conditional volatility estimates at time *t* and *t-1*.[82]

Fuzzy logic inference (denoted "Fuzzy-Logic" in Figure 4.1) works as follows: Fuzzy logic is built upon the notion of fuzzy sets (see Zadeh, 1965), which allow for partial memberships between variables. The fuzzy logic inference setting takes the ANN-based exchange rate prediction and the change in the underlying volatility as inputs (both scaled to fall into a [-1,1] interval) and evaluates assigned degrees of membership from each fuzzy set to the inputs. These fuzzy sets are governed by a membership function, which determines the degrees to which an input belongs to a certain set (labelled the "degree of truth"). We initially choose Gaussian membership functions for the order-flow-based return forecasts. For the change in the volatility, we choose triangular membership functions which have clearer boundaries between the different states.

The following functions are assigned to each exchange-rate-change state (fuzzy set): "VERY NEGATIVE (V-NEG)," "NEGATIVE (NEG)," "WEAKLY NEGATIVE (W-NEG)," "STABLE," "WEAKLY POSITIVE (W-POS)," "POSITIVE (POS)," and "VERY POSITIVE (V-POS)". The functions assigned to the each volatility-change state are: "STRONG DECREASE (S-DEC)," "DECREASE (DEC)," "STABLE," "INCREASE (INC)" and "STRONG INCREASE (S-INC)". The FX trader's action or trading recommendation is also a fuzzy variable with the triangular fuzzy membership functions that represent five fuzzy sets between -1 and 1: "VERY STRONG SELL (VS-SELL)," "SELL (S-SELL)," "WEAK SELL (W-SELL)," "HOLD," "WEAK BUY (W-BUY)," "BUY (BUY),"

[82] Note: The main conclusions drawn in the analysis are robust to different GARCH type models. The GARCH (1,1) model used in the analysis is initially estimated using all of the in-sample observations, and afterwards re-estimated daily during the out-of-sample period. The GARCH (1,1) is outlined in Equation 2.7.

and "STRONG BUY (S-BUY)". Figure 4.3 provides a graphical representation of the membership functions (fuzzy sets)[83].

Figure 4.3 Graphical representation of the membership functions and fuzzy sets

Panel A.

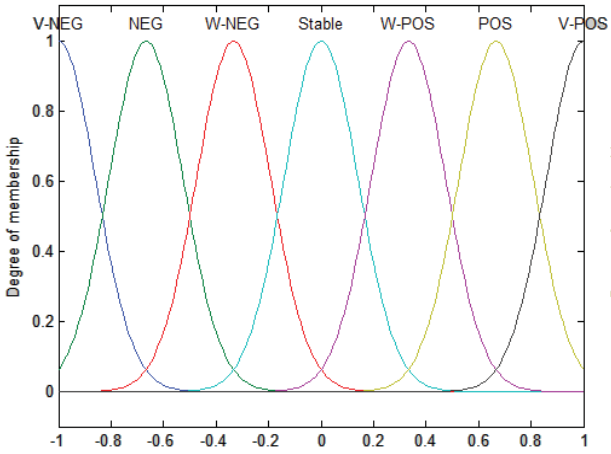

Exchange-rate-change-forecast

Panel B.

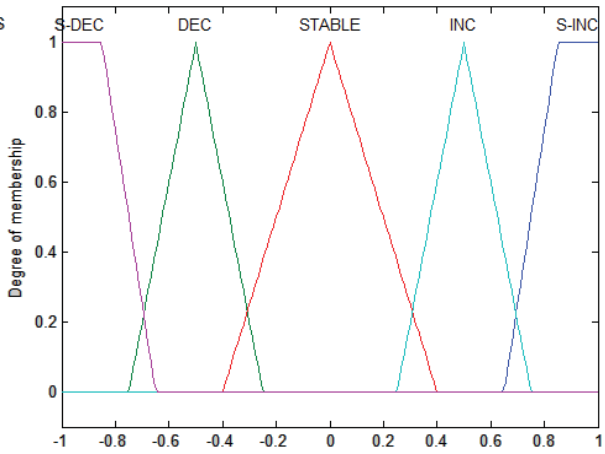

Exchange-rate-volatility-change

Panel C.

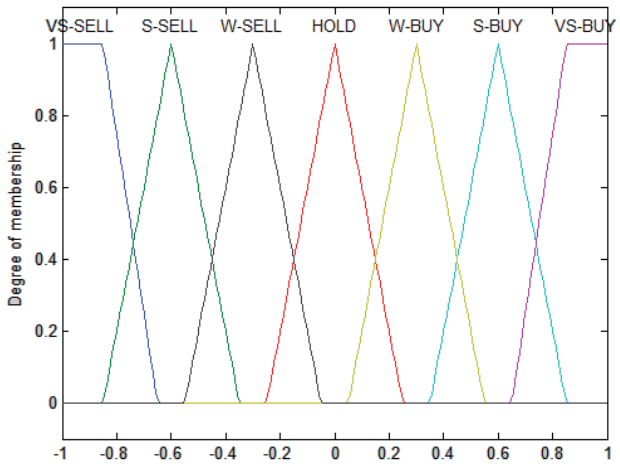

Trading-Recommendation

Note: This figure displays a graphical representation of the membership functions (fuzzy sets) used by the neuro-fuzzy and multi-fuzzy trading strategy. The multi-fuzzy approach uses both the proxy for the change in the underlying volatility (denoted "exchange-rate-volatility-change") and the exchange rate return predictions as inputs. By contrast, the neuro- fuzzy approach, which acts as a benchmark in the analysis presented later on, directly links the exchange rate change forecast to the trading recommendation.

[83] Note that increasing the number of membership functions for the input (output) states did not produce any significant improvements in the profitability of the multi-fuzzy and neuro-fuzzy models evaluated.

The point of fuzzy logic framework is to map an input space to an output space. The primary mechanism for doing this is a list of if–then statements called rules. All rules are evaluated in parallel and refer to variables and the adjectives that describe those variables. In our approach, the if–then rules are specified with a focus on the hypothesised negative relationship between increases in the underlying market volatility and future appreciations of the foreign currency relative to the USD. Given that the change in the volatility matches the sign of the predicted returns (i.e., when the underlying volatility decreases (increases) and the return forecast is positive (negative)), the trading recommendation will be based entirely on the magnitude of the return prediction. On the other hand, if the volatility changes contradict the sign of the predicted return (i.e., when the underlying volatility decreases (increases) and the predicted return is negative (positive)), the trading recommendation will be adjusted. Furthermore, if the volatility change contradicts the magnitude of the return prediction, the trading weights are adjusted. To underline this, consider the rules given below:

1. If the return forecast is “V-NEG” and the change in the volatility is a “S-INC”, “INC”, “STABLE”, or “DEC”, then “VS-SELL.”
2. If the return forecast is “V-NEG” and the change in the volatility is a “S-DEC”, then “S-SELL.”
3. If the return forecast is “V-POS” and the change in the volatility is a “S-DEC”, “DEC”, “STABLE”, or ”INC“, then “VS-BUY.”
4. If the return forecast is “V-POS” and the change in the volatility is a “S-INC” then “W-BUY”
5. If the return forecast is “NEG” and the change in the volatility is an “INC”, “STABLE”, or “DEC”, then “S-SELL.”

6. If the return forecast is “NEG” and the change in the volatility is a “S-DEC” (“S-INC”) then “W-SELL” (“VS-SELL”).
7. If the return forecast is “POS” and the change in the volatility is a “DEC”, “STABLE”, or ”INC“, then “S-BUY”.
8. If the return forecast is “POS” and the change in the volatility is a “S-INC” (“S-DEC”), then “W-BUY” (“VS-BUY”).
9. If the return forecast is “W-NEG” and the change in the volatility is an “INC”, “STABLE”, or “DEC”, then “W-SELL.”
10. If the return forecast is “W-NEG” and the change in the volatility is an “S-DEC” (“S-INC”), then “HOLD” (“S-SELL”).
11. If the return forecast is “W-POS” and the change in the volatility is a “DEC”, “STABLE”, or ”INC“, then “W-BUY”.
12. If the return forecast is “W-POS” and the change in the volatility is an “S-DEC” (“S-INC”), then “S-SELL” (“HOLD”).
13. If (the return forecast is stable) then (the trading recommendation is to hold)[84]

Rules 1, 3, 5, 7, 9 and 11 are the examples for when the volatility change is in line with the return predictions and give full weight to the return predictions. However, if the sign or magnitude of the return forecast contradicts the volatility change, as indicated in Rules 2, 4, 6, 8, 10 and 12, then the trading recommendation is adjusted. This is a novelty of our approach: we adjust the weights in the case when the fuzzy logic inputs contradict each other. By contrast, the neuro-fuzzy (univariate) framework (Gradojevic, 2007), which our approach builds up on, would use the following rules: If the return forecast is “V-NEG” then “VS-SELL”; if the return forecast is “NEG” then “S-SELL”; if the return forecast is “W-NEG”

[84] For any combination of a stable return forecast and changes in the underlying volatility, all rules result into a hold recommendation. Note that the “if-then” rules only allow for two “if” conditions at once per rule, i.e. “if the return forecast is “V-NEG” and the change in the volatility is “S-INC”, then “VS-SELL”; if the return forecast is “V-NEG” and the change in the volatility is “INC”, then “VS-SELL”.

then “W-SELL”; if the return forecast is “STABLE” then “HOLD”; if the return forecast is “W-POS” then “W-BUY”; if the return forecast is “POS” then “S-BUY”; and if the return forecast is “S-POS” then “VS-BUY”. Note that the rules of the neuro-fuzzy approach correspond to Rules 1,3,5,7,9,11 and 13 used in the multi-fuzzy approach. The output of the trading recommendation falls within the discrete intervals, shown in Table 4.1.

Table 4.1 Intervals for discrete trading recommendations (based on defuzzified output)

Trading strategy	Interval for z
SELL STRONG	$-1 < zt < -0.8$
SELL	$-0.8 < zt < -0.55$
SELL WEAK	$-0.55 < zt < -0.25$
HOLD	$-0.25 < zt < 0.25$
BUY WEAK	$0.25 < zt < 0.55$
BUY	$0.55 < zt < 0.8$
BUY STRONG	$0.8 < zt < 1$

Note: This table shows the boundaries between the overlapping membership functions for the variable’s trading recommendation, which falls between -1 (sell one unit of foreign currency) and +1 (buy one unit of foreign currency). Given that all exchange rates are denominated against the USD, a positive value suggests buying the foreign currency using the fraction of US dollars, whereas a negative value indicates a fraction of foreign currency to be sold.

The fuzzy logic is, therefore, essentially a look-up function that looks up the corresponding trading recommendation based on the inputs, membership functions and rules. The main difference is how the trading recommendation is looked up, as more than one rule can be active in the fuzzy logic setting (due to the partial memberships of the inputs). A “centroid method” is used to generate the trading recommendation based on the combined membership functions of the inputs.

To give further insight into the differences between the neuro-fuzzy and the multi-fuzzy framework, consider the following example. Suppose the ANN predicts an exchange rate change with an assigned value of -0.54 for the next period and the assigned value for the change in the underlying volatility is -0.7. This is graphically depicted in Figure 4.4. Panel A. displays the trading recommendation generated by our multivariate fuzzy logic framework, whereas Panel B displays the trading recommendation generated by the univariate fuzzy logic framework.

Figure 4.4 Comparison of multivariate and univariate fuzzy logic framework

Panel A. Panel B.

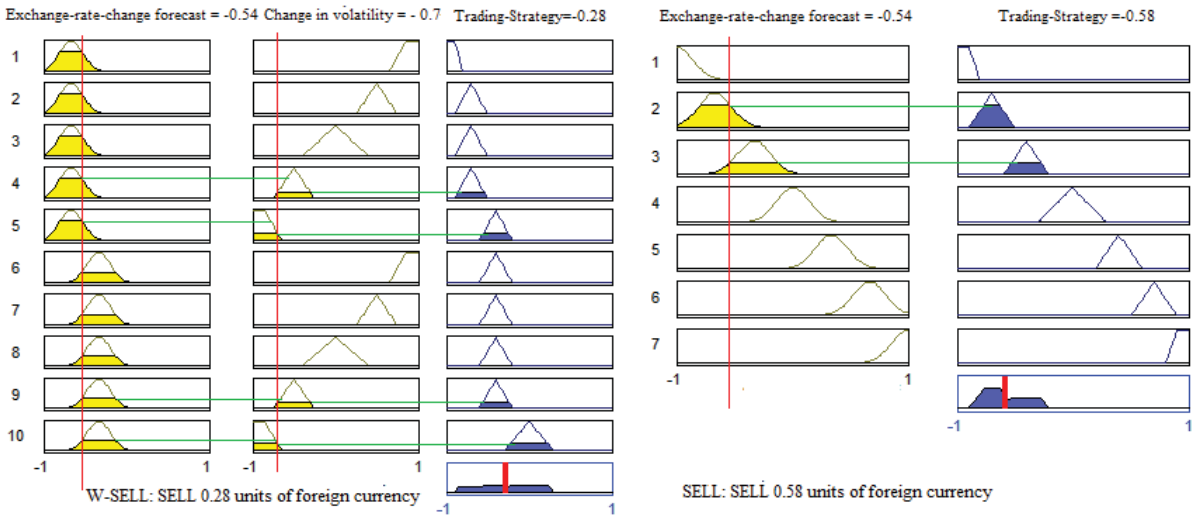

Note: This figure is for illustrative purposes only. Panel A displays parts of the fuzzy logic inference setting for the multi-fuzzy trading strategy. Panel B displays the fuzzy logic inference setting for the neuro-fuzzy strategy replicated from Gradojevic (2007). The yellow shaded areas represent the degree of membership from the inputs to each fuzzy set. The blue triangles represent the fuzzy sets for the output states (trading recommendation). The blue shaded areas represent the outputs generated by the fuzzy rules and their degree of membership to a particular output set. The red bar shows the centre of mass, which gives the trading recommendation.

The exchange rate change prediction of -0.54 corresponds to two fuzzy logic sets: the "negative" and "weakly negative" set. This is the same for both the neuro- and multi-fuzzy strategies. However, the neuro-fuzzy strategy would directly link the return prediction to the

trading recommendation via two "if-then" rules (the two active rules are "if the return forecast is "NEG" then "S-SELL" and if the return forecast is "W-NEG" then "W-SELL") resulting in a recommendation to sell a fraction of 0.58 units of the foreign currency. By contrast, in the multivariate fuzzy logic inference setting, four "if then" rules are active resulting in a recommendation to 0.28 units of the foreign currency. This difference in the trading weights between both approaches is due to the decrease in the underlying volatility, which the univariate fuzzy approach (Panel B) does not take into account.

4.5 Empirical Results

4.5.1 Initial assessment

Prior to investigating the profitability of the technical trading strategies, we examine the statistical properties of the returns and determine the statistical tests that are most appropriate for examining the performance of technical trading rules[85]. As our first step, we investigate the daily return predictability from order flow, using the same regression as in Section 3.7.1 in Chapter 3:

$$P_t = C + \Lambda X_{t-1} + \Gamma P_{t-1} + U_t, \qquad (4.2)$$

where $P_t = [\Delta s_t]$ is a vector of the exchange rate returns, X_{t-1} is a vector of order flow (measured in thousands), Λ and Γ are the coefficients to be estimated, C is the constant and U_t is the residual.

Table 4.2 shows the results of our regression. As we can see, the order flow parameter estimate Λ is only statistically significant for four (USD–GBP, USD–NZD, USD–AUD and USD–CAD) out of the six exchange rate pairs investigated. Although the parameter estimates

[85] We use the same dataset as in the previous chapter. Data on the CBOE VIX and EPU Index used in the robustness section is obtained from Thomsen Reuters DataStream.

are not significant for the USD–EUR and USD–JPY rates, the other exchange rates show that the current order flow has significantly positive effect on their future movements.

There are a variety of reasons for this result. The most obvious is the high autocorrelation of order flow itself (a positive order flow today is likely to be followed by a positive order flow tomorrow and similarly for a negative order flow today), leading to a persisting impact of the current order flow on future spot rate fluctuations[86]. However, the parameter estimates on the lagged returns are all insignificant (except for USD–CAD). This might indicate that either the lagged returns have no effect on future spot rate movements, which would contradict the popularity of price-based technical trading indicators, or that one should consider different lags in the relationship. The last observation is that the R^2 values of all regressions are very small: literally zero, casting doubts on the presence of a linear relationship among order flows, lagged returns and current spot rate movements.[87]

Table 4.2 Regression of return predictability from order-flow and lagged returns

	USD–EUR	USD–GBP	USD–JPY	USD–AUD	USD–NZD	USD–CAD
C	0.012	0.032 ***	0.007	0.013	0.014	0.082
Λ	0.086	0.174***	0.516	0.374***	0.735***	0.119 ***
Γ	-0.001	0.003	0.039	-0.062	-0.002	0.126 ***
R^2	0.002	0.048	0.015	0.088	0.048	0.062

Note: The regression presented uses the exchange rate return as the dependent variable and order flow and lagged exchange rate returns as the explanatory variable. Returns are multiplied by 100. *** denotes significance at the 1% level, ** at the 5% level and * at the 10% level.

4.5.2 Price- and order-flow-based technical trading indicators

Table 4.3 shows the performance of the price- and order flow indicators. The total return for each strategy is calculated as:

[86] Another explanation is the gradual incorporation of current information into future prices via order flow, as outlined in Chapter 3.
[87] As the dataset is mainly the same as used in Chapter 2, we do not report the descriptive statistics here.

$$R_k = \sum_{t=1}^{N-1} w_{t,k} r_{t+1} \quad (4.3)$$

where w_t represents the trading recommendations (-1,0 or +1 in this case). k=1,…,M refers to the number of technical trading rules (M=2691) and *N* is the number of trading periods (N= 2856)[88]. We report the performance of only the best price- and order-flow-based trading indicators for parsimony[89]. The White Reality Check and Hansen's test for superior predictive accuracy, however take into account all the 2691 trading rules in the universe to calculate the p-values.

Before accounting for transaction costs, we can see a fairly clear pattern. The best order flow based trading indicator generates the highest returns for the USD–GBP, USD–AUD and USD–NZD currency pairs, with a maximum of roughly 17% p.a. seen for the USD–NZD exchange rate. Conventional trading strategies seem to work fairly well, with the best MA and filter rules resulting in greater returns on average than channel breakouts and trading range breaks. This finding is in line with Qi and Wu (2006). Although the best order flow based technical indicator leads to a higher return than conventional strategies for USD–GBP, USD–AUD and USD–NZD (the difference is roughly 8.35% p.a. for the USD–NZD pair), our results might be biased by data snooping, which is suggested by the insignificant *p*-value of White's reality check and Hansen's test for superior predictive accuracy. The evidence in favour of achieving superior trading profits from simple order-flow (price-based) trading strategies is therefore very weak. Furthermore, the differences in profitability between order-flow and price-based trading strategies are small for four out of the six currency pairs evaluated a finding which supports H1.

[88] Note: In this analysis we follow the literature and use the entire dataset to evaluate the profitability of the competing technical trading indicators. The main conclusions drawn from this analysis, however, apply to the three specific subsamples used in the next section. The first trading signal is generated for the 256th observation for all specifications, as some rules require 255 days of previous data in order to generate a trading signal.
[89] We report the performance of the best strategy in terms of returns generated. The main conclusions drawn, however, do not alter when comparing the best strategies in terms of Sharpe ratios.

Table 4.3 Performance of the best trading rule among price- and order-flow-based strategies

	USD–EUR	USD–GBP	USD–JPY	USD–AUD	USD–NZD	USD–CAD
Panel A. Performance before transaction costs. January 2002 – December 2013						
FILT	7.24%	3.90%	4.40%	8.24%	8.67%	7.37%
MA	7.61%	6.40%	7.00%	8.36%	8.59%	7.27%
TB	5.47%	3.31%	5.75%	9.95%	8.30%	7.36%
CB	5.13%	5.00%	5.32%	6.65%	7.51%	6.60%
OF	7.24%	7.47%	3.38%	16.06%	17.02%	5.50%
Buy-and-hold	3.44%	1.04%	2.53%	5.04%	5.69	3.30%
White RC	0.79	0.46	0.42	0.30	0.38	0.91
Hansen SPA	0.79	0.46	0.42	0.30	0.38	0.90
Panel B. Performance after transaction costs. January 2002 – December 2013						
FILT	7.17%	2.86%	3.80%	8.07%	8.60%	7.34%
MA	7.54%	6.08%	5.77%	8.25%	8.58%	7.26%
TB	5.34%	2.40%	5.15%	9.33%	8.28%	7.32%
CB	4.97%	4.71%	5.29%	4.65%	4.51%	6.26%
OF	7.03%	6.74%	3.24%	10.20%	13.52%	5.45%
Buy-and-hold	3.44%	1.04	2.53%	5.04%	5.69%	3.30%
White RC	0.83	0.47	0.45	0.56	0.53	0.91
Hansen SPA	0.83	0.44	0.45	0.55	0.52	0.86

Note: This table presents the annualised percentage returns of the best price-based and order-flow-based technical indicators from the universe of 2,691 trading rules for each currency. Panel A. shows the returns in the absence of transaction cost, whereas Panel B. shows the transaction cost-adjusted returns. The buy-and-hold strategy (BH) acts as a benchmark. The annualised returns for each strategy correspond to the optimal critical values. Transaction costs are equal to 2 basis points. FILT represents the filter rule, SMA is the simple moving average indicator, TB and CB represent the trading range break and the channel break technical indicators respectively, OF represents the performance of the order-flow-based indicator. White's reality check (RC) and Hansen's test for superior predictive accuracy (SPA) are *p*-values obtained from the bootstrap approach outlined in Appendix C, and test whether any given trading strategy can outperform the buy-and-hold benchmark.

Turning to the results that allowing for transaction costs reported in Panel B, the pattern still persists: the best order flow based trading indicator still provides higher returns than the best price-based indicators for the USD–GBP (6.74%), the USD–AUD (10.20%) and the USD–NZD rate (13.52%). The best MA indicator provides the highest return for the USD-EUR spot rate of 7.54% and USD–JPY rate of 5.77% p.a. For the USD-CAD, the

largest return is generated by the best filter rule. Again, although the order flow based trading rules generate a substantially higher return for USD–GBP, USD–AUD and USD–NZD than price-based indicators, the performance differences are not significant after accounting for data snooping[90.]

All of the findings in section 4.5.2 confirm the hypothesis that implementing order flows directly as a trading indicator does not provide significant benefits after accounting for data snooping. We believe that this finding warrants a more sophisticated nonlinear approach, a possible one being the neuro-fuzzy setting to which we now turn.

4.5.3 Performance of the neuro-fuzzy and the multi-fuzzy strategy

The primary input in both the neuro–fuzzy and the multi– fuzzy trading strategy is the order-flow-based ANN forecast. The performance of both approaches hence relies critically on the accuracy of the ANN-based return prediction. Given that the return predictions would be highly inaccurate the profitability of the trading strategy would be limited. Therefore, we first asses the accuracy of the neural network based return predictions.

Figure 4.2 plots the relationship between the inputs (order flow, interest rate differential and oil price) and the targets (changes in the logarithmic exchange rate) of the ANN. This figure gives the first indication of the relationship between the order flows, changes in the oil price, changes in the IRD and the standardised change in the EUR–USD rate. There seems to be a relationship between the standardised inputs and targets, but the

[90] We also experimented with combinations of price- and order-flow-based technical indicators, following the documented profitability of combined trading indicators outlined by Lento (2009). Contrary to Lento (2009), however, we found that no combined trading indicator led to a significantly larger rate of return than a buy-and-hold strategy, regardless of the exchange rate pair evaluated. We do not display the results because of parsimony.

relationship is not linear. The aim of the ANN is to estimate this nonlinear relationship between exchange rate movements and the input variables.

Figure 4.5 Relationship between artificial neural network (ANN) inputs and targets

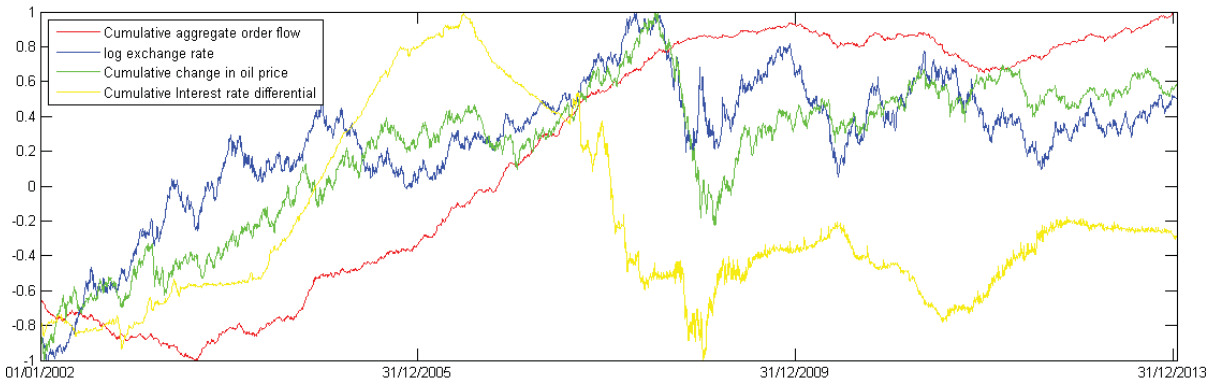

Note: This figure plots the relationship between normalised inputs (change in the oil price, order flows and the interest rate differential) and the logarithm of the USD-GBP exchange rate. All variables are normalised to the [-1,1] interval. The same data are given to the ANN to establish a nonlinear relationship between changes in the exchange rate and the input variables.

Although Figure 4.2 gives a visual justification for using the ANN forecast, a more formal justification for this relies on whether the ANN model can outperform a random walk prediction (in terms of mean squared errors arising from both predictions) and the directional accuracy of the ANN predictions, measured as the percentage of correctly predicted signs (PERC).

Mean squared errors (MSEs) are constructed as the squared differences between the return prediction and the actual returns and are a standard measure to evaluate the accuracy of competing return predictions. Yet, the primary indicator for forecasting power (trading rule profitability) is not necessarily the mean squared error, but the percentage of correctly forecasted directions of real exchange rate fluctuations. PERC measures the total number of correctly forecasted positive and negative movements over the out-of-sample period and are calculated as:

$$PERC = \frac{1}{N}\sum_{t=1}^{N} p_t \qquad (4.4)$$

where, $p_t = 1$ if $r_t \hat{r}_t > 0$ and 0 otherwise.

To see whether the PERCs are significantly different from the sample average and if the MSEs are significantly different from a random walk forecast, we apply the Pesaran and Timmermann (PT) test for directional- and the Diebold-Mariano (DM) test for predictive accuracy. A positive DM statistic indicates that the ANN based prediction resulted in a lower MSE than the random walk based return prediction. A higher PT statistic indicates a better directional accuracy. The null hypothesis is that there is no difference in the PERCs of the random walk and the ANN model. Both the PT and DM statistics are asymptotically distributed as a N(0,1).

Table 4.4 shows the summary statistics of the ANN-based return forecasts for three distinctive out-of-sample periods and the six exchange rates studied. The three out-of-sample periods are constructed following the same methodology as Chapter 3: each subsample contains 1000 observations, with the out-of-sample part of each subsample used to evaluate the ANN accuracy chosen to be the last 200 observations. The 800 in sample observation represent the training, validation and testing set for the neural network. Afterwards the network takes the data from the out-of-sample period to generate one-step-ahead exchange rate predictions.

The main finding from Table 4.4 is that, regardless of the exchange rate or time period studied, the ANN forecast always outperforms a random walk forecast, in terms of mean squared errors (MSE), as indicated by the corresponding DM test statistics. The PERC varies widely, depending on both the exchange rate under investigation as well as on the subsample period studied. To put things in perspective, for USD–CAD return sign

predictions, the PERC statistic ranges between 56% and 59.5%, showing the strongest sign detection ability during the out-of-sample period of the second subsample. The same observation can be made for USD–AUD and USD–NZD where the PERC statistics vary between 55% and 59.5%, showing the strongest sign detection ability during the out-of-sample period of the first subsample. For both exchange rates, the sign predictions are statistically different from the sample averages for all periods, as indicated by the PT statistics. The largest directional accuracy of the network based return predictions is found for the USD–EUR and USD–GBP exchange rates during the second subsample's out-of-sample period (61% and 62% correctly predicted signs). By contrast, the directional accuracy of the network predictions is questionable for USD–JPY predictions during this period.

Table 4.4 Out-of-sample accuracy of the Artificial Neural Network (ANN) predictions

	USD–EUR	USD–GBP	USD–JPY	USD–CAD	USD–AUD	USD–NZD
			Panel A. January 2002 – December 2005			
PERC	0.550	0.605	0.560	0.560	0.595	0.590
DM	5.75***	4.69 ***	2.43***	4.72 ***	4.445 ***	3.39 ***
PT	3.28***	2.97***	2.32**	1.41*	3.377***	2.98***
			Panel B. January 2006 – December 2009			
PERC	0.610	0.620	0.530	0.595	0.590	0.580
DM	3.46***	2.09 ***	1.43*	5.199 ***	2.77***	2.993 ***
PT	3.29***	3.74**	1.05	2.472 ***	2.56***	2.05**
			Panel C. January 2010 – December 2013			
PERC	0.570	0.590	0.59	0.560	0.550	0.550
DM	4.73***	4.50***	1.99**	3.546 ***	2.471 ***	3.26 ***
PT	2.24**	2.57***	2.09**	1.762 **	1.692 **	2.34***

Note: This table displays the key statistics of the accuracy of the ANN. PERC denotes the percentage of correctly predicted hits over the out-of-sample period. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. DM is the Diebold–Mariano (1995) test statistic for equal predictive accuracy. The benchmark prediction is the order-flow-based ANN forecast. The competing prediction is a random walk forecast. PT refers to the Pesaran–Timmermann (1992) test for sign prediction. The one sided critical values for both tests are and 2.33, 1.645 and 1.282 for confidence levels of 99%, 95% and 90%, respectively. *** denotes significance at the 1% level, ** at the 5% level and * at the 10% level.

Apart from screening the ANN inputs and evaluating the statistical accuracy of the network prediction, another starting point is to ask whether a simple trading indicator derived from an order-flow-based ANN forecast can indeed outperform a simple buy-and-hold strategy, as this would indicate that order flows can only be linked indirectly to technical trading rule profits, which further justifies its use in the fuzzy logic application below. In fact, such a strategy was outlined by Fernández-Rodríguez et al. (2000) and Jasic and Wood (2004), where the return sign from the neural network prediction is used as a buy or sell [1,-1] indicator. We follow this "regime-switching" strategy, i.e. a positive predicted return generates a buy- and a negative predicted return generates a sell recommendation.

In this analysis, we evaluate the profitability in terms of Sharpe ratios, which makes it easier for comparison against the fuzzy logic strategies (to be described later). The returns for each strategy are constructed as in Equation 4.4., and the Sharpe ratios are constructed as:

$$SR = \frac{\overline{\mu R_k}}{\sigma R_k}, \qquad (4.5)$$

where the numerator is the average return in each trading strategy (the buy-and-hold strategy, the regime-switching strategy which uses the ANN prediction directly as a trading indicator, or the neuro fuzzy-logic-based trading strategies outlined later on)[91], 17 and the denominator is the standard deviation of the returns from each trading strategy. We then test whether the Sharpe ratios generated by the three strategies are statistically different from each other, following the testing procedure outlined in Jobson and Korkie (1981) and Memmel

[91] The numerator of the Sharpe ratio is usually defined as the excess return from a particular trading strategy over some risk-free rate. In this paper, however, we follow Qi and Wu (2006) by setting risk-free rate to zero, because they suggest that the excess return (zero net investment) for an investor is the speculative position in the FX market (the benchmark is for the investor not to take a position in the FX market and hence a zero return). This practice of setting risk-free rate to zero can further allow us to link our findings directly to Gradojevic (2007) who also specifies the Sharpe ratio according to equation (4.5). More importantly, this practice also allows us to make the setting intuitive and practically relevant to small traders who do not necessarily get interest rates from their brokers.

(2003), denoted as JKM test. This test has been most frequently used for the equality of two trading strategies [92].

The test for differences between two competing Sharpe ratios follows from the asymptotic distribution of the differences in risk-adjusted returns between two trading strategies given by:

$$\sqrt{T}(\widehat{SR}_a - \widehat{SR}_b) \xrightarrow{a} N(0, \sqrt{V_{diff}}), \text{with } V_{diff} = 2 - 2\rho_{a,b} + \frac{1}{2}[SR_a^2 + SR_b^2 - 2SR_a SR_b \rho_{a,b}^2] \text{ , } [93]$$

(4.6)

Where, $\widehat{SR}_a$ and $\widehat{SR}_b$ represent the estimated Sharpe Ratios from the two competing trading strategies and $\rho_{a,b}$ represents the correlation in returns between both strategies. The standard error of the estimated Sharpe ratio difference is equal to the square root of V_{diff}. The null hypothesis tested is that the difference in Sharpe Ratios is equal to zero.

We compare different trading strategies in terms of Sharpe ratios only, for the following reasons. First, comparing raw returns to evaluate the fuzzy logic trading strategies, though more intuitive, is inappropriate, as it would neglect the differences in the trading weights between the competing strategies. Such neglecting would overlook the differences in volatility and hence riskiness between the strategies. For instance, smaller trading weights will typically result in smaller returns, but at the same time will also reduce the risk involved in the trading strategy. If using raw returns only, the strategies with smaller trading weights would be considered inferior strategies, whereas according to risk-adjusted returns (Sharpe ratio) they may not be. Second, all other financial ratios, with the exception of the Traynor ratio albeit more applicable for stock markets, can only gauge the profitability of a trading

[92] The test was initially outlined in Jobson and Korkie (1981) with Memmel (2003) correcting some shortcomings in the derivation of the important formulas.
[93] See Opdyke (2007) for further explanation on the testing procedure.

strategy, but do not enable one to judge whether one strategy is significantly better than another. This shortcoming of other financial ratios renders it necessary to apply the Sharpe ratio which can sever the purposes of comparisons. Finally, a common practice in the existing studies that evaluate the profitability of the trading strategies is to employ the Sharpe ratio, due to its practical relevance for market participants. We follow this argument and practice.

As outlined in Table 4.5, using an order-flow-based ANN forecast as a trading indicator typically generates substantially larger Sharpe ratios than a buy-and-hold strategy. This regime switching strategy provides significantly higher Sharpe ratios than the buy-and-hold strategy for trading USD–EUR, USD–GBP, USD–CAD and USD–NZD, regardless of time period studied. For a trader specialising in USD–JPY and USD–AUD, the regime switching strategy based on an ANN prediction provides significantly higher Sharpe ratios than the buy-and-hold strategy for the first and third sub-sample's out-of-sample periods.

For the financially tranquil first out-of-sample period, the Sharpe ratios are particularly large for trading the USD–EUR and USD–GBP currency pairs (4.45 and 4.64 before, and 3.89 and 4.17 after transaction costs). For the financially turbulent second out-of-sample period, the Sharpe ratios generated by the regime switching strategy are particularly high for trading USD–NZD (3.44 and 3.24) and again USD–EUR (3.25 and 3.00); however, the strategy does not produce significantly different Sharpe ratios compared to a buy-and hold strategy for a trader investing in USD–JPY and USD–AUD. For the third sub-sample, the simple regime switching strategy provides the highest Sharpe Ratios for an USD–EUR (3.78 and 3.01) and USD–JPY (3.07 and 2.72) trader; however, the strategy results in a negative risk-adjusted return for trading USD–AUD after transaction costs (-0.05) .

The findings from this analysis are, however, not unexpected, or else ANNs would not be so prominently used by traders, and a buy-and-hold strategy is not really a benchmark

in FX trading, as exchange rates do not pay a dividend. Furthermore, although a simple trading indicator based on an ANN return prediction indeed provides larger Sharpe ratios than a buy-and-hold strategy, the performance of this trading strategy suffers from frequent rebalancing, as transaction cost-adjusted returns remove about one-third of the Shape ratio on average.

Table 4.5. Out-of-sample performance of the regime switching strategy

	USD–EUR	USD–GBP	USD–JPY	USD–CAD	USD–AUD	USD–NZD
Panel A. January 2002 – December 2005						
BH	-1.42	-1.18	-1.79	0.70	-0.71	-0.47
BH*	-1.42	-1.18	-1.79	0.70	-0.72	-0.47
ANN	4.45**	4.64**	2.99**	2.04**	1.99***	2.89***
ANN*	3.89**	4.17**	2.59**	1.80**	1.63***	2.20 ***
Panel B. January 2006 – December 2009						
BH	1.06	0.83	0.14	1.21	1.75	1.58
BH*	1.06	0.83	0.14	1.21	1.75	1.58
ANN	3.25***	2.04**	3.02	2.80**	3.39	3.44 ***
ANN*	3.00***	1.80***	2.71	2.52**	3.16	3.25 ***
Panel C. January 2010 – December 2013						
BH	0.24	-0.92	-2.85	-1.29	-1.46	-0.34
BH*	0.24	-0.92	-2.85	-1.29	-1.46	-0.34
ANN	2.01***	3.78**	3.07***	1.39***	0.50***	3.34***
ANN*	1.41***	3.01***	2.72***	0.48***	-0.05***	3.12***

Note: This table presents the annualised Sharpe ratios generated by a buy-and-hold strategy (denoted BH) and a simple regime switching strategy using the ANN-based return forecast as a technical trading indicator (denoted ANN). Panel A reports the results for the out-of-sample period of the first subsample, and Panels B and C do the same for the out-of-sample periods of the other subsamples. An asterisk next to a strategy indicates that the Sharpe ratios for the trading strategy are calculated after the transaction costs, which are assumed to be equal to 2 basis points. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. We use the test outlined in Jobson and Korkie (1981) and Memmel (2003) to evaluate whether the differences in Sharpe ratios are statistically different from zero. The asterisks next to the Sharpe ratios indicate that the Sharpe ratios generated by the ANN-based trading strategy are statistically different from the buy-and-hold strategy. *** denotes significance at the 1% level, ** at the 5% level and * at the 10% level.

The natural solution, which allows us to reduce trading uncertainty and also generates a smoother decision surface, is to link the ANN forecast to the trading recommendations in a

fuzzy logic setting. This has been discussed above for the CAD–USD exchange rate, and our approach builds upon this. Table 4.7 compares the neuro-fuzzy approach (replicated from Gradojevic (2007)), which uses only the ANN based return predictions as an input, versus our proposed approach outlined in section 4.4.2. To avoid confusion we label the univariate fuzzy approach as NF to highlight the fact that this approach only uses the ANN prediction. Opposed to this, our proposed multi-fuzzy strategy takes the order-flow-based ANN return predictions into account as well as changes in the underlying market volatility, which we proxy with changes in the conditional exchange rate volatility. We therefore denote this approach MF.

As demonstrated in Table 4.6, both the neuro-fuzzy and multi-fuzzy strategy, on average, provide substantially higher Sharpe ratios for each exchange rate and over all the subsample periods compared to the strategy evaluated in Table 4.5, which uses the return prediction directly as a trading signal. This result becomes even stronger when we compare the Sharpe ratios after the transaction costs because of the smoother trading recommendations of the fuzzy logic approach.

What is truly striking is that the multi- fuzzy approach dominates the neuro-fuzzy approach, for 5 out of the six exchange rates regardless of the time period evaluated. The differences in the Sharpe ratios are statistically significant for USD–NZD across all out-of-sample periods evaluated, and for USD–EUR, USD-GBP, USD–AUD and USD–CAD for two out of the three periods evaluated.

Overall, the differences between both approaches are the largest for the first subsample's out-of-sample period, where the differences in Sharpe Ratios range up to 0.88 (0.80 after transaction costs) for a trader trading AUD with USD, and 0.7 (0.67 after

transaction costs) for a trader trading NZD with USD. To put the aspects into perspective, these differences in Sharpe Ratios correspond to a gain of around 25% and 32%.

This immense difference reveals the economic value of our proposed approach. Though the differences are particularly large for the first subsample period, the benefits are still economically significant for the second and third subsamples. For the out-of-sample period of the second subsample, the Sharpe ratios from the neuro-fuzzy approach range between 2.54 and 4.43, which are lower than those of the multi-fuzzy approach (ranging between 2.85 and 4.83); a similar pattern is also discernible for the third subsample's out-of-sample period, where the multi-fuzzy strategy generates an astonishing Sharpe ratio of 5.52 and 5.20 after transaction costs for trading the NZD-USD currency pair. The same conclusions apply to the transaction cost-adjusted Sharpe ratios, as both strategies incur similar shares of transaction costs, thereby confirming H2.

Table 4.6 Out-of-sample performance of the neuro- and the multi-fuzzy trading strategy

	USD–EUR	USD–GBP	USD–JPY	USD–CAD	USD–AUD	USD–NZD
			Panel A. January 2002 – December 2005			
NF	4.86	4.75	3.79	1.93	3.56	2.20
NF*	4.23	4.27	3.38	1.55	3.12	1.65
MF	5.43 **	4.95	3.79	2.22*	4.44***	2.90***
MF*	4.75 **	4.45	3.31	1.81*	3.92***	2.32***
			Panel B. January 2006 – December 2009			
NF	3.52	2.54	3.72	2.80	3.21	4.43
NF*	3.22	2.27	3.40	2.47	2.97	4.25
MF	3.93**	2.95***	3.49	2.85	3.48	4.83**
MF*	3.60**	2.66***	3.14	2.49	3.25	4.64**
			Panel C. January 2010 – December 2013			
NF	2.58	4.28	2.76	2.02	0.72	4.76
NF*	1.88	3.40	2.36	0.81	0.24	4.45
MF	2.71	5.08***	2.38	2.39*	0.98*	5.52***
MF*	2.01	4.25***	2.17	1.33**	0.50*	5.20***

Note: This table presents annualised Sharpe ratios generated by the neuro- and the multi-fuzzy trading strategy. NF refers to a trading strategy using the ANN-based exchange rate return prediction as a fuzzy logic input directly; The MF strategy uses the change in the underlying (conditional) volatility as an additional input. The MF setting is displayed in Figure 1. Panel A. reports the results for the out-of-sample period of the first subsample, and Panels B and C do the same for the out-of-sample periods of the other subsamples. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. An asterisk next to a strategy indicates that the Sharpe ratios for the trading strategy are calculated after the transaction costs, which are assumed to be equal to 2 basis points. We use the test outlined in Jobson and Korkie (1981) and Memmel (2003) to evaluate whether the differences in Sharpe ratios are statistically different from zero. The asterisks next to the Sharpe ratios indicate that the Sharpe ratios generated by the MF strategy are statistically different from the NF strategy *** denotes significance at the 1% level, ** at the 5% level and * at the 10% level.

4.6 Robustness

4.6.1 Choice of membership functions

The first robustness test evaluates two alternative shapes for the input space: triangular and trapezoidal membership functions. A graphical representation of the fuzzy logic setting using triangular and trapezoidal membership functions is given in Figure C1 in the Appendix.

In a nutshell, using triangular membership functions does not change the conclusions drawn before as outlined in Table 4.6. The multi-fuzzy strategy outperforms its univariate counterpart for 5 out of the six exchange rates and over all the subsample periods. As in the baseline analysis, the differences are statistically significant for USD–NZD across all out-of-sample periods evaluated and for USD–EUR, USD–AUD, USD–GBP and USD–CAD for two out of the three periods evaluated.

Scrutinising the differences among the periods, the first subsample demonstrates that the biggest difference is 0.21 (= 4.00 – 3.79) for the USD-JPY rate, followed by -0.13 (= 2.77 - 2.90) for the USD-NZD rate. The second subsample indicates that the biggest difference arises for the USD-EUR rate (-0.10 = 4.03- 3.93), followed by the USD-CAD rate (0.08 = 2.93 – 2.85). The third subsample shows that the biggest difference goes to the USD-GBP rate (-0.12 = 4.96 – 5.08) followed by the USD-CAD (0.04 = 2.39 – 2.35). As the Sharpe ratios are very similar to the original parameterisation, this analysis confirms the robustness of our setting to triangular membership functions[94].

[94] Note that the results do not alter when we use trapezoidal membership functions (see Table C.1 in Appendix C), which further confirms the robustness of the trading strategy.

Table 4.7 Out-of-sample performance of the neuro- and the multi-fuzzy trading strategy (triangular membership functions)

	USD–EUR	USD–GBP	USD–JPY	USD–CAD	USD–AUD	USD–NZD
Panel A. January 2002 – December 2005						
NF	4.87	4.76	3.62	1.93	3.56	2.20
NF*	4.25	4.27	3.14	1.55	3.12	1.66
MF	5.44**	4.82	4.00*	2.25*	4.44***	2.77***
MF*	4.77**	4.33	3.42	1.83*	3.93***	2.23***
Panel B. January 2006 – December 2009						
NF	3.52	2.54	3.72	2.80	3.21	4.43
NF*	3.22	2.27	3.40	2.47	2.97	4.25
MF	4.03***	2.97***	3.44	2.93	3.44	4.82**
MF*	3.70***	2.67***	3.10	2.56	3.21	4.63**
Panel C. January 2010 – December 2013						
NF	2.58	4.28	2.76	2.02	0.72	4.75
NF*	1.88	3.40	2.36	0.81	0.23	4.44
MF	2.74	4.96**	2.37	2.35*	0.99*	5.54***
MF*	2.04	4.22***	2.16	1.29**	0.46*	5.22***

Note: This table presents annualised Sharpe ratios generated by the neuro- and the multi-fuzzy trading strategy. NF refers to a trading strategy using the ANN-based exchange rate return prediction as a fuzzy logic input directly; The MF strategy uses the change in the underlying (conditional) volatility as an additional input. The MF setting is displayed in Figure 1. Panel A. reports the results for the out-of-sample period of the first subsample, and Panels B and C do the same for the out-of-sample periods of the other subsamples. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. An asterisk next to a strategy indicates that the Sharpe ratios for the trading strategy are calculated after the transaction costs, which are assumed to be equal to 2 basis points. We use the test outlined in Jobson and Korkie (1981) and Memmel (2003) to evaluate whether the differences in Sharpe ratios are statistically different from zero. The asterisks next to the Sharpe ratios indicate that the Sharpe ratios generated by the MF strategy are statistically different from the NF strategy *** denotes significance at the 1% level, ** at the 5% level and * at the 10% level.

4.6.2. Alternative volatility proxies

An interesting question is whether using changes in a proxy for global (exchange rate) volatility would generate a similar boost in trading rule performance. To test this possibility, we test the performance of our proposed multi-fuzzy strategy using proxies for the change in global risk aversion and uncertainty as the second input. As outlined below, increases in

global risk aversion and uncertainty should increase the volatility of the individual currency pairs; we therefore use this analysis as a further test for the robustness of our proposed approach.

We focus on changes in the Chicago Board Options Exchange volatility index (VIX) and economic policy uncertainty (EPU). The VIX measures the implied volatility of index options of the S&P 500 index, which is shown to be a close proxy for the volatility of the constituents of the S&P 500 (Fleming, Ostdiek, and Whaley, 1995), their price jumps (Becker, Clements, and McClelland, 2009) and their expected returns (Whaley, 2000). The VIX is, in itself, a forward-looking measure, as it is based on the options' implied volatility, which is the expected market volatility over the next 30 days and is commonly termed "fear index" (Whaley, 2000) as it reflects option implied volatility which rises given there is a higher demand for put options.

The EPU index consists of three underlying components: media coverage of policy-related economic uncertainty, reports by the Congressional Budget Office on tax code provisions and uncertainty about policy-related macroeconomic variables. Its applications in the literature have not only evaluated EPU-based volatility predictions, as outlined in Liu and Zhang (2015), but have also indicated the existence of a link between policy uncertainty and macroeconomic fundamentals (Baker, Bloom, and Davis, 2015).

The intuition is that all exchange rates display a larger volatility when global risk aversion or uncertainty increases, which is reflected by strong increases in the VIX and EPU. Based on the reasoning outlined before, increases (decreases) in the VIX (EPU) should result in an appreciation (depreciation) of the USD relative to all other foreign currencies. A behavioural explanation for this argument is that increases in uncertainty (EPU) or risk aversion (VIX) could cause a flight to quality or flight to liquidity, thereby causing an

increased demand for USD holdings, as the USD is the major reserve currency. Note that although the EPU and the VIX have a strong correlation to each other, conceptually, they measure something different. The VIX measures changes in risk aversion, whereas the EPU is a measure for (policy related) economic uncertainty.

The testing procedure is same as before: evaluating the risk-adjusted returns from each trading strategy, we test for differences in the Sharpe ratios generated by both proxies used as inputs in the multi-fuzzy strategy against the benchmark of the neuro-fuzzy approach, replicated from Gradojevic (2007). To avoid confusion, we will label the baseline multi-fuzzy strategy as MF–GARCH and the multi-fuzzy strategies using changes in the VIX and EPU as MF–VIX and MF–EPU.

Table 4.8 shows the results of this analysis. The MF–GARCH results are replicated from Table 4.6. At first glance, regardless of the volatility proxy used, the multi-fuzzy approaches (MF–GARCH, MF–VIX and MF–EPU) typically yield higher risk-adjusted returns than the neuro-fuzzy approach (NF). A second general observation is in order. The MF–VIX typically yields higher risk-adjusted returns than the MF–EPU and closely matches the MF–GARCH.

Looking at the differences in detail, we find that the MF–GARCH (MF–VIX) is always preferred over the MF–EPU strategy during the first out-of-sample period, regardless of the exchange rate pair evaluated. The differences between the MF–GARCH and MF–VIX are small by comparison, and no clear pattern emerges. For the USD–EUR and USD–NZD the MF–GARCH approach results in the highest Sharpe ratios generated (4.75and 2.32), whereas the MF–VIX approach results in the highest Sharpe Ratios for the USD–GBP, USD–AUD and USD–CAD pairs. As before, no multi-fuzzy strategy outperforms the neuro-fuzzy strategy for the USD–JPY pair.

For the second out-of-sample period, the same pattern remains, as the MF-GARCH and MF-VIX typically dominate the MF-EPU strategy. Taking the USD–NZD as an example, the Sharpe ratio generated by the MF–VIX (MF–GARCH) is equal to 4.69 (4.64) compared with 4.00 for the MF-EPU, a difference of 17.2 (16) %. A similar conclusion applies to USD–EUR and USD–AUD. Note that during this period, the MF–VIX provides a significantly larger Sharpe Ratio than the neuro-fuzzy strategy for an USD–JPY trader, whereas both MF-GARCH and MF-EPU result in a smaller Sharpe ratio than the NF trading strategy. Summarising the findings for this period, the MF–VIX is the preferred strategy for five out of the six exchange rate pairs, and the MF–GARCH for the other currency pair.

For the out-of-sample period of the third subsample, the MF-EPU strategy generates the highest Sharpe Ratio for trading USD–AUD (0.55), the MF-VIX generates the largest Sharpe ratio for trading USD–CAD (2.07) and the MF-GARCH is preferred for the other four exchange rates.

Although the fuzzy logic setting incorporating the EPU generates lower Sharpe ratios than using other volatility proxies on average, this does not mean that the EPU is less informative proxy for the change in the underlying volatility. The lower Sharpe ratios generated are a result of the distribution of the EPU index: the index shows massive spikes followed by calm periods, which makes the normalisation of EPU changes to any interval problematic[95]. By contrast, changes in the VIX and GARCH based volatility estimates are easier to model for an interval, as they are less volatile. Overall, it seems that the MF–GARCH is the most consistent approach; however the MF-VIX also provides substantially higher risk adjusted returns than the NF strategy. We take this finding as a confirmation for the robustness of our approach.

[95] See Section 4.4.2. The proxy for the global volatility change and the network based return prediction both fall within the range [-1,1].

Table 4.8 Out-of-sample performance comparisons of different volatility proxies

	USD–EUR	USD–GBP	USD–JPY	USD–CAD	USD–AUD	USD–NZD
Panel A. January 2002 – December 2005						
NF*	4.23	4.27	3.38	1.55	3.12	1.65
MF–GARCH*	4.75**	4.45	3.04	1.81*	3.93***	2.32***
MF–VIX*	4.54	4.32	3.08	2.02**	3.98***	1.87**
MF–EPU*	4.25	3.84	3.11	1.30	2.54	1.82
Panel B. January 2006 – December 2009						
NF*	3.22	2.27	3.40	2.47	2.97	4.25
MF–GARCH*	3.60**	2.66***	3.14	2.49	3.25	4.64**
MF–VIX*	3.77***	2.32	3.82**	3.25***	3.57**	4.69**
MF–EPU*	3.54**	2.53	3.24	2.72	2.73	4.00
Panel C. January 2010 – December 2013						
NF*	1.88	3.40	2.36	0.81	0.24	4.45
MF–GARCH*	2.01	4.25***	2.17	1.34**	0.50*	5.20***
MF–VIX*	2.00	3.48	2.50	1.02	0.48*	4.70
MF–EPU*	2.09	3.53	2.53	1.00	0.55*	4.49

This table presents annualised Sharpe ratios generated by the neuro- and the multi-fuzzy trading strategy. NF refers to a trading strategy using the ANN-based exchange rate prediction as a fuzzy logic input directly; The MF–GARCH uses the change in the underlying (conditional) volatility as an additional input. The MF–VIX and MF–EPU use changes in a proxy for global exchange rate volatility, measured by changes in the VIX and EPU, as an additional input. The performance of each strategy is net of transaction costs. The MF setting is displayed in Figure 1. Panel A. reports the results for the out-of-sample period of the first subsample, and Panels B and C do the same for the out-of-sample periods of the other subsamples. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. We use the test outlined in Jobson and Korkie (1981) and Memmel (2003) to evaluate whether the differences in Sharpe ratios are statistically different from zero. The asterisks next to the Sharpe ratios indicate that the Sharpe ratios generated by the MF strategy are statistically different from the NF strategy *** denotes significance at the 1% level, ** at the 5% level and * at the 10% level.

4.7. A MA Fuzzy Logic Approach

Based on the striking boost in trading rule performance achieved by linking the underlying volatility to order-flow-based exchange rate predictions, we now investigate whether the performance of moving average rules can be boosted by linking changes in the underlying volatility to the trading signals generated by the moving average rules.

Moving average rules commony compare the differences between a short (l_1) and long (l_2) moving average to derive the trading signal, which results in either a buy [1] or sell [-1] recommendation. While many combinations of short and moving averages are possible, we focus on the four combinations ([1,50], [1,200], [5,200] and [2,200]) used in Gradojevic and Gençay (2013) in this section. The buy and sell signals are calculated as $s_t^{l1,l2} = m_t^{l1} - m_t^{l2}$, with a buy signal generated when $s_t^{l1,l2} \geq 0$ and a sell signal when $s_t^{l1,l2} \leq 0$. As before, a buy signal indicates to buy the foreign currency using US dollars.

We compare the performance of a simple moving average trading strategy against two moving average based fuzzy logic trading strategies. The simple moving average trading strategy generates a trading signal based on the sign of the moving average difference directly. The two fuzzy logic trading strategies work as follows. The first strategy uses the differences between the long and short moving average as the only input, and links the magnitude of the moving average differences directly to the trading recommendation. This trading strategy is replicated from Gradojevic and Gençay (2013), and is in the spirit of the neuro-fuzzy approach, using the moving average difference in lieu of the neural network based return prediction.

The second strategy is based on the multi-fuzzy strategy outlined in section 4.4.2. This strategy uses two inputs, namely the difference between the long and short moving

average for a currency pair and the currrent change in underlying volatility to derive the trading recommendation. The if-then rules used by this version of the strategy are akin to the multi-fuzzy strategy using the neural network based exchange rate change prediction as primary input (e.g. “If the difference between the long and short MA is “POS” and the change in the volatility is “S–DEC” (“S–INC”) then the trading recommendation is a “VS–BUY” (“W–BUY”)).

To avoid confusion, we denote the simple moving average strategy as MA, the fuzzy logic strategy relying on the moving average differences only as MA–FUZZY and the multivariate counterpart as MULTI–FUZZY. As we use four different moving average combinations, this results in four different trading recommendations (and hence risk adjusted returns) generated by each strategy. We then average the risk adjusted returns generated by all four moving average combinations for each strategy.

As reported in Table 4.9, we find overwhelming evidence that the MULTI–FUZZY approach yields significantly higher Sharpe ratios than both the MA and MA–FUZZY trading strategies. The superior performance of the MULTI–FUZZY approach is solely due to adjusting the trading weights in case the volatility change does not match the moving average trading signal. This confirms H3.

Looking at the individual subsamples in detail, we find that for out-of-sample period of the first subsample, the MULTI–FUZZY generates a significantly higher Sharpe ratio than the MA–FUZZY and the simple MA indicator, regardless of the exchange rate studied, both before and after taking transaction costs into account. The trading gains are especially high for the USD–CAD rate: the Sharpe ratio of the MULTI–FUZZY equals 2.45 (2.24 after transaction costs) compared to 1.90 (1.82). The largest relative increase arises for a USD–

GBP trader, where the Sharpe ratio generated by the MULTI–FUZZY strategy is about 60% larger than the Sharpe ratio generated by the MA–FUZZY and MA strategy.

The superior performance of our strategy does not substantially alter for the second subsample. For this period and the USD–EUR exchange rate, the MULTI–FUZZY generates a gross risk adjusted return of 0.73, which is much larger than the Sharpe ratio of the MA-FUZZY (0.37) and MA trading strategy (0.02). The same applies for a trader investing in USD–GBP, USD–JPY, USD–CAD or USD–AUD. Only the USD–NZD rate does not follow this pattern, where the MA–FUZZY approach dominates the other approaches; the differences in performances between the competing approaches are, however, not statistically significant.

The out-of-sample period of the third subsample further confirms the main conclusions drawn from the analysis presented for the first and second subsamples' out-of-sample periods. The MULTI–FUZZY approach dominates the other approaches in terms of risk adjusted returns and provides significantly higher returns for a USD–GBP or USD–AUD trader than the other two approaches. Note that the fuzzy logic based trading strategies even result in positive risk adjusted returns for the USD–GBP (0.88 and 1.36) when the moving average trading strategy generates a loss (-0.22).

Comparing the analysis presented in Table 4.9 with the regime switching strategy or the neural network based fuzzy logic trading strategies in Table 4.5 and Table 4.6, however, highlights the larger economic value of order flow based return predictions over price-based trading indicators. In other words, the ANN-based (-fuzzy) trading strategies dominate the MA-based (-fuzzy) trading strategies on average. This explains why practitioners apply ANNs very widely.

T**able 4.9** Out-of-sample performance for simple and fuzzy-logic-based MA rules

	USD–EUR	USD–GBP	USD–JPY	USD–CAD	USD–AUD	USD–NZD
Panel A. January 2002 – December 2005						
MA	1.00	0.53	0.64	1.39	0.49	-0.66
MA *	0.95	0.50	0.60	1.32	0.39	-0.72
MA–FUZZY	1.20	0.55	0.88	1.90	0.75	-0.81
MA–FUZZY *	1.18	0.51	0.84	1.82	0.69	-0.85
MULTI–FUZZY	1.58***	0.88***	1.15***	2.45***	1.12***	-0.37***
MULTI–FUZZY *	1.35***	082***	1.08***	2.24*	0.95***	-0.54***
Panel B. January 2006 – December 2009						
MA	0.02	0.06	2.63	0.11	0.78	0.64
MA *	-0.03	0.03	2.56	0.09	0.70	0.60
MA–FUZZY	0.37	1.20	1.86	0.38	0.92	1.04
MA–FUZZY *	0.34	1.15	1.82	0.32	0.88	1.02
MULTI–FUZZY	0.73***	1.51***	2.04***	0.72***	1.14***	0.92
MULTI–FUZZY *	0.65***	1.30*	1.95	0.51**	1.06*	0.87
Panel C. January 2010 – December 2013						
MA	-1.12	-0.22	0.58	0.06	0.77	0.46
MA *	-1.21	-0.28	0.54	0.03	0.70	0.40
MA–FUZZY	-0.30	0.88	0.39	0.61	0.92	0.11
MA–FUZZY *	-0.34	0.81	0.34	0.54	0.88	0.04
MULTI–FUZZY	-0.21	1.36***	0.46	0.73	1.27**	0.15
MULTI–FUZZY *	-0.27	1.10***	0.32	0.33	1.10**	0.10

Note: This table presents annualised Sharpe ratios of the competing moving average based trading strategies. MA refers to the average performance of the four simple moving average trading rules. MA–FUZZY refers to a fuzzy logic based trading strategy only considering the differences between a slow and fast MA. MULTI–FUZZY refers to a trading strategy using the differences between the long and the short MA and the change in the underlying (conditional) volatility of the currency pair. As for the MA rule, the performance measure for the MA-FUZZY and MULTI–FUZZY refers to the average performance of the trading strategies for all four moving average combinations evaluated. An asterisk next to a strategy indicates that the Sharpe ratios for the trading strategy are net of transaction costs, which are assumed to be equal to 2 basis points. The out-of-sample period used for the analysis is equal to the last 200 observations of each subsample We use the test outlined in Jobson and Korkie (1981) and Memmel (2003) to evaluate whether the differences in Sharpe ratios are statistically different from zero. The asterisks next to the Sharpe ratios indicate that the Sharpe ratios generated by the MULTI–FUZZY strategy are statistically different from the MA-FUZZY strategy. *** denotes significance at the 1% level, ** at the 5% level and * at the 10% level.

4.8 Conclusion

This chapter proposes a multivariate fuzzy logic framework to boost the benefits of using order flows in technical analyses for FX trading. The approach incorporates and combines order-flow-based return predictions as well as the underlying market volatility. We show the superiority of our approach by comparing it against the performance of various strategies such as simple trading rules and neuro-fuzzy trading strategy. Our study yields conclusions related to the former and conclusions pertaining to the latter including our novel approach.

We begin with a simple but interesting research question regarding the practical value of order flow information. Given that order flows have been found to be useful for predicting future movements and that the forecasting accuracy has been found to be the highest for intraday movements, this chapter goes further to evaluate the practical benefits that may appeal to investors. We find that a simple trading indicator based on order flow yields a higher return than conventional strategies for several currency pairs. However, this result could be caused by data snooping, thus casting doubt on the profitability of simple modelling. The same argument also applies to combinations of price-based and order-flow-based technical trading rules.

By contrast, we show that order-flow-based return forecasts can produce statistically significant excess returns. This result provides the foundation for evaluating a neuro-fuzzy logic approach. Although order-flow-based return forecasts generate an excellent trading indicator, the trading recommendations of any simple strategy will be a binary integer, which does not and cannot take the magnitude of the return forecasts into account. We circumvent this problem by using the return forecasts as inputs in a fuzzy logic setting, which generates smoother trading recommendations by considering the magnitude of the return forecast in

question. This approach outperforms the direct strategy for all the exchange rates under investigation, a finding previously reported by Gradojevic (2007) for the USD–CAD rate only.

Although this fuzzy logic trading strategy can provide even higher excess returns than any simple trading strategy using order-flow-based return predictions, we also uncover that the performance of this strategy can be improved substantially by linking the exchange rate forecast to changes in the underlying volatility. The theory behind the strategy is based on a hypothesized positive relationship between appreciations of the USD relative to a foreign currency and increases in the underlying market volatility. The multi–fuzzy strategy makes use of this theory by adjusting the trading weights in the case where the exchange rate return prediction contradicts the change in the underlying market volatility.

Appendix C

C.1. Additional Tables and Graphs

Table C.1 Out-of-sample performance of the neuro- and the multi-fuzzy trading strategy (trapezoidal membership functions)

	USD–EUR	USD–GBP	USD–JPY	USD–CAD	USD–AUD	USD–NZD
Panel A. January 2002 – December 2005						
NF	4.92	4.76	3.62	1.93	3.56	2.20
NF*	4.33	4.27	3.14	1.55	3.12	1.82
MF	5.57 **	4.96	4.00*	2.25*	4.42***	2.85***
MF*	4.85 **	4.45	3.40	1.84*	3.90***	2.30***
Panel B. January 2006 – December 2009						
NF	3.52	2.54	3.72	2.80	3.21	4.43
NF*	3.22	2.27	3.40	2.47	2.97	4.25
MF	4.02***	2.99***	3.44	2.83	3.43	4.82**
MF*	3.68***	2.69***	3.10	2.49	3.20	4.62**
Panel C. January 2010 – December 2013						
NF	2.59	4.28	2.76	2.02	0.72	4.73
NF*	1.89	3.40	2.36	0.81	0.22	4.43
MF	2.75	4.99**	2.37	2.34*	0.89**	5.50***
MF*	2.05	4.22***	2.16	1.30**	0.45**	5.18***

Notes: This table presents annualised Sharpe ratios of the neuro-fuzzy and multi-fuzzy trading strategies. NF refers to a trading strategy using the artificial neural network (ANN) forecast as a fuzzy logic input directly; The MF strategy uses the change in the underlying (conditional) volatility as an additional input. The MF setting is displayed in Figure 4.1. Panel A. reports the results for the out-of-sample period of the first subsample, and Panels B and C do the same for the out-of-sample periods of the other subsamples. The out-of-sample period used in the analysis is equal to the last 200 observations of each subsample. An asterisk next to a strategy indicates that the Sharpe ratios for the trading strategy are calculated after the transaction costs, which are assumed to be equal to 2 basis points. We use the test outlined in Lo (2002) to evaluate whether the differences in Sharpe ratios are statistically different from zero. The asterisks next to the Sharpe ratios indicate that the Sharpe ratios generated by the MF strategy are statistically different from the NF strategy. *** denotes significance at the 1% level, ** at the 5% level and * at the 10% level.

Figure C.1 Triangular and trapezoidal membership functions

Panel A.

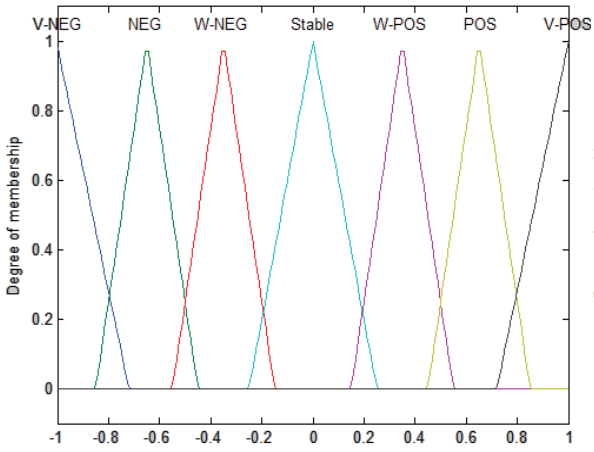

Exchange-rate-change-forecast

Panel B.

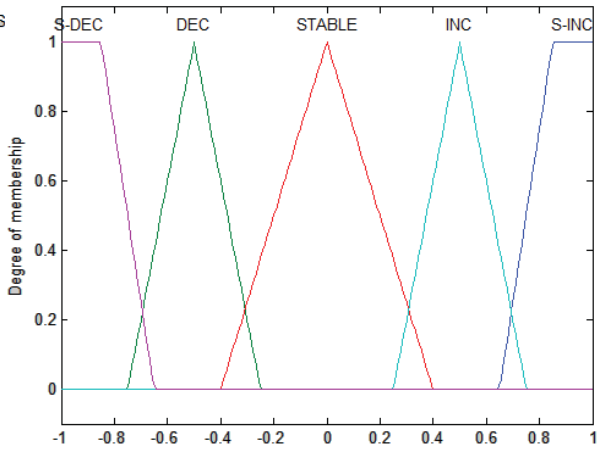

Exchange-rate-volatility-change

Panel C.

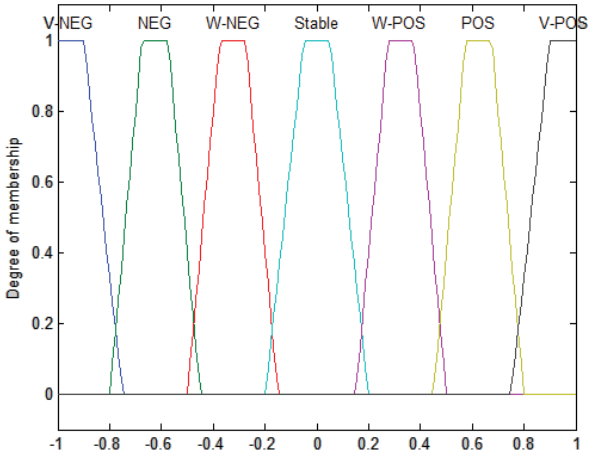

Exchange-rate-change-forecast

Panel D.

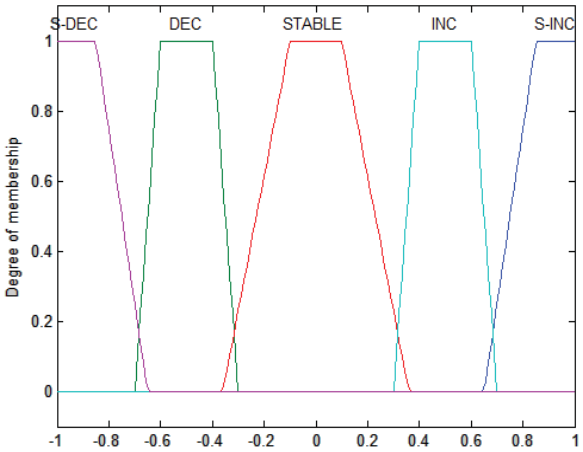

Exchange-rate-volatility-change

Note: Panels A. and B. display the triangular membership functions used in analysis in Section 4.6.1. Panels C and D display the trapezoidal membership functions used in the analysis presented in Table C.1.

C.2 White's Reality Check and Hansen's test for superior predictive ability

White's (2000) reality check allows us to test whether any given model has superior predictive accuracy over a benchmark model after accounting for data snooping effects. The null hypothesis is that no rule can outperform a predefined benchmark (in our case a buy-and-hold strategy) in terms of returns. The White reality check gives an estimate of the true *p*-value for the null hypothesis, adjusted for data snooping by taking the entire universe of rules from which the best rule is selected into account. Mathematically, this is described below.

The performance statistic for each trading rule is given by $\bar{f}_k = T^{-1} \sum_{t-1}^{T} f_{k,t+1}$, where $k = 1, \ldots, M$. M refers to the number of technical trading rules, *T* is the number of trading periods, and $f_{k,t+1}$ represents the performance measure at $t + 1$. $f_{k,t+1}$ is defined as the excess returns of a trading rule *k* over a benchmark return (a return from the buy-and-hold strategy). The null hypothesis is that the returns from the best technical trading rule will not be higher than those from the benchmark: $H_0: \max_{k=1,\ldots,M}\{\bar{f}_k\} \leq 0$. The rejection of the null hypothesis implies that the best trading rule results in a higher performance than the benchmark.

The null hypothesis is evaluated by applying the stationary bootstrapping of Politis and Romano (1994) to the observed value of $f_{k,t+1}$. The steps are as follows. First, for each trading rule, we resample $f_{k,t+1}$ *B* times. Following Yamamoto (2012) we set B equal to 500. Let us denote the re-sampled series as $f^*_{k,t+1,b}$ (*b*=1, …, *B*). Then we calculate the average of the bootstrap returns as the mean over the sample period, as follows: $\bar{f}^*_{k,b} = T^{-1} \sum_{t=1}^{T} f^*_{k,t,b}$ (b=1,…,*B*). We then compute the statistic $\bar{V} = max_{k=1,\ldots,M}[\sqrt{T}\,\bar{f}_k]$ and $\overline{V_b^*} = \max_{k=1,\ldots,M}[\sqrt{T}(\bar{f}^*_{k,b} - \bar{f}_k)]$. White's reality check *p*-value is obtained by comparing $\bar{V}$ with V_b^*. In particular, we sort $\overline{V_b^*}$ (*b*=1,…*B*) and find the number of times $\overline{V_b^*}$ is larger than $\bar{V}$.

Denoting N as the number of times $\overline{V_b^*}$ is larger than $\bar{V}$, the reality check *p*-value is given as: 1 - *N/M*.

The superior predictive ability test proposed by Hansen (2005) employs a studentised test statistic that reduces the influence of very inaccurate forecasts and evaluates the risk-adjusted performances of the competing trading rules. Hansen's superior predictive ability test evaluates the risk-adjusted performance by re-centring the null distribution. The average of the centred returns is calculated as follows: $\bar{R}_{k,b}^* = \bar{f}_{k,b}^* - \bar{f}_k \, I$, where *I* is an indicator function that takes 1 if the average is reasonably small. In this approach, the critical statistics are given as:

$$\bar{V}_b^* = \max\left(\max_{k=1,\ldots,M} \frac{\sqrt{T}\,\bar{R}_{k,b}^*}{\hat{\sigma}_k}, 0\right) \quad \bar{V} = \max\left(\max_{k=1,\ldots,M} \frac{\sqrt{T}\,\bar{f}_k^*}{\hat{\sigma}_k}, 0\right). \qquad (4.7)$$

Where, as before, the *p*-value is obtained by comparing $\bar{V}$ with V_b^*.

Chapter 5 Conclusion

5.1 Summary of Contributions

The market microstructure approach to determining exchange rates has proven to be highly successful for explaining and predicting single exchange rate movements. This thesis builds upon the foundations of order flow theory and links it to the explanation and prediction of multivariate exchange rate movements. A further multivariate application studied in this thesis links order flows to profitable trading signals by coupling them with the underlying volatility.

As anticipating correlations is at the heart of risk management, the driving forces of conditional correlation warrant a further investigation. We do so by employing a proxy that contains information about all other potential macroeconomic correlation drivers. The proxy is order flow, a market microstructure variable. It has been shown in the existing literature to be able to explain and predict univariate exchange rate movements. Order flow is simply signed trading volume, the net of buyer- over seller-initiated trades, and thus indicates whether a currency experiences buying or selling pressure. Although this buying or selling pressure drives single exchange rate movements, its relationship with joint movements of currencies has not yet been evaluated in the literature. This dissertation attempts to fill the void.

We find that absolute order flow differentials between spot rates have a significant negative effect on joint (co-)movements. This finding applies to all exchange rate pairs in our sample and remains valid when we use standardised proxies for the order flow differential. More specifically, in line with the gradual information hypothesis, the relationship is the strongest for the highest intraday frequency evaluated and gradually decreases as the frequency increases. This finding is robust when we control for the bid–ask spread, a key

variable in intraday price dynamics. In fact, the finding suggests that the bid–ask spread has only a very limited role to play in explaining correlation dynamics once the order flow differential is accounted for, in line with the market microstructure literature. Finally, we document that the relationship between absolute order flow differentials and exchange rate co-movements is not the same across financially stable and turbulent periods, with a special case arising from the different trading intensities between financially calm and financially turbulent periods.

As well as modelling correlation dynamics, anticipating correlations involves prediction of future fluctuations in the joint return movements of assets. Given that exogenous variables, such as the IRD, drive conditional correlation dynamics, a question arises as to whether this relationship would result in benefits from predicting future correlation dynamics. Surprisingly, the previous studies linking exogenous variables to co-movement have overlooked this question (e.g., Li, 2011). Chapter 3 of this thesis thus turns to addressing this omission in the literature. We find that covariance predictions taking order flow information into account trace realised covariance more closely than several competing benchmark models that do not take the fact that exogenous variables may affect covariance dynamics into account. In addition, our findings suggest that the IRD only has a limited role to play in predicting future joint return fluctuations.

Although order-flow-based predictions provide the highest predictive accuracy regardless of the time period evaluated, we document that the incremental gains are the largest for economically turbulent time periods, such as the 2007 GFC. In other words, the superior performance of our proposed order-flow-based models over other competing benchmark models becomes more apparent during an economic downturn. This finding again confirms the market microstructure literature, suggesting that the information content of order

flows is strongest during economic downturns, though from the perspective of co-movement predictions.

Alongside the modelling approach, we evaluate different estimation windows in terms of generating covariance predictions. We show that a medium-sized estimation window typically provides the highest accuracy, which is in line with the literature. The theoretical reason is that a medium-sized window mitigates the parameter inflexibility caused by a long estimation window while maintaining the parameter robustness that short estimation windows lack.

We also assess the economic value of more accurate covariance predictions by using the same portfolio optimisation setting as King et al. (2010) and Rime et al. (2010). The setting allows us to investigate the incremental value of order-flow-based covariance predictions, an aspect neglected in the order-flow-based portfolio optimisation studies. Our findings indicate that a risk-averse investor will be willing to give up a substantial risk premium to switch from a static covariance prediction to our proposed model. We deem this to be the economic value of order-flow-based covariance predictions. A more dynamic application that takes transaction costs into account confirms that our results are robust.

However, the portfolio optimisation setting commonly used to evaluate the benefits of order-flow-based return and covariance predictions entails frequent rebalancing because order flow information is strongest for short-term return predictions. The fourth chapter of this thesis therefore examines the benefits of order flow information, again in a multivariate setting, to generate trading profits before and after the transaction costs involved in portfolio rebalancing. By linking the probability of order-flow-based forecasts to have the correct sign to changes in the underlying volatility, we create a highly profitable trading strategy. This

trading application is novel and adds to the practical value of utilising order flow information, an aspect that has largely been neglected in the market microstructure literature.

Starting from a very simple approach, we derive technical trading signals directly from the observed order flow in a market, an analysis previously conducted for stock markets (Yamamoto, 2012). This approach, although intuitive and in the spirit of classic price-based trading indicators, does not outperform a buy-and-hold strategy after taking data snooping biases into account. Our explanation for this finding is that order flow is a problematic variable for any technical trading indicator because of its non-normal distribution and therefore is better suited for exchange rate predictions; the exchange rate predictions can then be used as a profitable proxy for trading decisions.

Using order flow to generate return predictions, we confirm the value of order flow information for a FX trader. We find that this approach indeed provides substantial trading profits for all exchange rate pairs evaluated, confirming the robustness of the findings of Gradojevic (2007) for all exchange rates under our investigation, whereas Gradojevic (2007) focuses on the CAD–USD exchange rate only. Furthermore, we propose a multivariate fuzzy-logic-based trading strategy which uses order-flow-based exchange rate change predictions and changes in the underlying market volatility to derive the trading recommendations. This novel methodology turns out to be highly profitable.

To sum up, this thesis contributes to the literature by going beyond the existing univariate framework to the multivariate framework in terms of the order flow–exchange rate nexus. Our results suggest that order flow differentials are a significant driver of the high-frequency co-movements of exchange rates. The bridge we build between market microstructure and dynamic currency co-movements transcends the traditional research on how such co-movements react to macroeconomic factors, including interest rates or central

banks' FX interventions. This bridge facilitates exchange rate predictions in terms of both their statistical accuracy and economic gains. The study conducted in Chapter 4 critically evaluates the value of order flow information in technical analysis as an investment strategy commonly used by professionals.

5.2 Future Research Agenda

This section discuses an intended agenda for future research. The present dissertation endeavours to investigate the important role of order flow in explaining and forecasting exchange rate co-movements and the economic value of information contained in order flow and volatility.

With respect to co-movement, an interesting question for future research is whether and how the co-movements of interest rates are affected by order flow. The intuition is that order flow should fully represent the aggregated expected changes in interest rates (Rime et al., 2010). Therefore, the relationship between order flow and interest rate co-movements should be even stronger than the relationship between order flow and exchange rate co-movement. Ferreira (2005) studied the forecasting of the co-movements of spot interest rates, but without incorporating exogenous variables into his analysis. Our proposed research work would help fill this void by allowing order flow to be a potential exogenous variable and by, presumably, improving forecasts of spot interest rate co-movements.

Yet another possible avenue of research would be to enhance the fuzzy logic setting outlined in Chapter 4. In the spirit of Atsalakis and Valavanis (2009), we plan to allow the fuzzy logic setting to adapt to data, rather than specifying its core elements in advance. In the fuzzy logic setting presented in this thesis, the membership functions were chosen arbitrarily and the rule structure was essentially predetermined by our interpretation of the

characteristics of the variables in the model. By contrast, the adaptive version of the fuzzy logic setting tailors the membership functions to the input/output data and the rule structure depends on the patterns learned by the fuzzy logic setting based on actual input/output data. The adaptive setting further tunes its rules towards new data, as opposed to our setting where the if–then rules are the same for the entire time period.

Although this implies that the fuzzy logic setting would be more of a black box, it assures us that the underlying membership functions and rule base fully reflect the actual data at all times, which should further enhance the profitability of our proposed multivariate fuzzy logic setting outlined in Chapter 4.

Bibliography

Aielli, G. P. (2013). Dynamic conditional correlation: on properties and estimation. *Journal of Business & Economic Statistics, 31*(3), 282–299.

Albuquerque, R., & Vega, C. (2009). Economic news and international stock market co-movement. *Review of Finance, 13*(3), 401–465.

Andersen, T., Bollerslev, T., Diebold, F., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica, 71*(2), 579-625.

Andersen, T., Bollerslev, T., Diebold, F., & Vega, C. (2003). Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange. *American Economic Review, 93*(1), 38–62.

Antonakakis, N. (2012). Exchange return co-movements and volatility spillovers before and after the introduction of euro. *Journal of International Financial Markets, Institutions and Money, 22*(5), 1091-1109.

Atsalakis, G., & Valavanis, K. (2009). Forecasting stock market short-term trends using a neuro-fuzzy based methodology. *Expert Systems with Applications, 36*(7), 10696-10707.

Bacchetta, P., & Van Wincoop, E. (2006). Can information heterogeneity explain the exchange rate determination puzzle? *The American Economic Review, 96*(3), 552–576.

Bailey, D., & Lopez de Prado, M. (2014). The deflated sharpe ratio: correcting for selection bias, backtest overfitting and non-normality. *Journal of Portfolio Management, 40*(5), 94–107.

Baker, S., Bloom, N., & Davis, S. (2015). *Measuring Economic Policy Uncertainty*. National Bureau of Economic Research: Cambridge MA.

Barberis, N., Shleifer, A., & Wurgler, J. (2005). Comovement. *Journal of Financial Economics, 75*(2), 283–317.

Barndorff-Nielsen, O., & Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of financial econometrics, 2*(1), 1-37.

Bauwens, L., Laurent, S., & Rombouts, J. (2006). Multivariate GARCH models: a survey. *Journal of applied econometrics, 21*(1), 79–109.

Bauwens, L., Omrane, W., & Giot, P. (2005). News announcements, market activity and volatility in the euro/dollar foreign exchange market. *Journal of International Money and Finance, 24*(7), 1108–1125.

Becker, R., Clements, A., & McClelland, A. (2009). The jump component of S&P 500 volatility and the VIX index. *Journal of Banking & Finance, 33*(6), 1033–1038.

Beine, M. (2004). Conditional covariances and direct central bank interventions in the foreign exchange markets. *Journal of Banking & Finance, 28*(6), 1385–1411.

Beine, M., Grauwe, P., & Grimaldi, M. (2009). The impact of FX central bank intervention in a noise trading framework. *Journal of Banking & Finance, 33*(7), 1187–1195.

Bekiros, S. (2011). Sign prediction and volatility dynamics with hybrid neurofuzzy approaches. *IEEE Transactions on Neural Networks, 22*(12), 2353–2362.

Bekiros, S., & Georgoutsos, D. (2008). Direction-of-change forecasting using a volatility-based recurrent neural network. *Journal of Forecasting, 27*(5), 407-417.

Benediktsdottir, S., & Scotti, C. (2009). *Exchange rates dependence: what drives it?* : Board of Governors of the Federal Reserve System (US), International Finance Discussion papers, 969.

Berger, D., Chaboud, A., Chernenko, S., Howorka, E., & Wright, J. (2008). Order flow and exchange rate dynamics in electronic brokerage system data. *Journal of International Economics, 75*(1), 93–109.

Billio, M., & Caporin, M. (2009). A generalized dynamic conditional correlation model for portfolio risk evaluation. *Mathematics and Computers in Simulation, 79*(8), 2566–2578.

Bjonnes, G. H., & Rime, D. (2005). Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics, 75*(3), 571–605.

Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *The Review of Economics and Statistics, 72*, 498–505.

Brock, W., Lakonishok, J., & LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *The Journal of Finance, 47*(5), 1731–1764.

Cao, Q., Leggio, K., & Schniederjans, M. (2005). A comparison between Fama and French's model and artificial neural networks in predicting the Chinese stock market. *Computers & Operations Research, 32*(10), 2499-2512.

Cappiello, L., Engle, R., & Sheppard, K. (2006). Asymmetric dynamics in the correlations of global equity and bond returns. *Journal of Financial Econometrics, 4*(4), 537–572.

Carlson, J., & Lo, M. (2006). One minute in the life of the DM/US $: public news in an electronic market. *Journal of International Money and Finance, 25*(7), 1090–1102.

Cerrato, M., Kim, H., & MacDonald, R. (2015). Microstructure order flow: statistical and economic evaluation of nonlinear forecasts. *Journal of International Financial Markets, Institutions and Money, 39*, 40–52.

Cerrato, M., Sarantis, N., & Saunders, A. (2011). An investigation of customer order flow in the foreign exchange market. *Journal of Banking & Finance, 35*(8), 1892–1906.

Chordia, T., & Subrahmanyam, A. (2004). Order imbalance and individual stock returns: theory and evidence. *Journal of Financial Economics, 72*(3), 485–518.

Christoffersen, P., & Diebold, F. (2006). Financial asset returns, direction-of-change forecasting, and volatility dynamics. *Management Science, 52*(8), 1273–1287.

Christoffersen, P., Errunza, V., Jacobs, K., & Jin, X. (2014). Correlation dynamics and international diversification benefits. *International Journal of Forecasting, 30*(3), 807–824.

Clark, T., & McCracken, M. (2009). Improving forecast accuracy by combining recursive and rolling forecasts. *International Economic Review, 50*(2), 363–395.

Danielsson, J., & Love, R. (2006). Feedback trading. *International Journal of Finance & Economics, 11*(1), 35–53.

Danielsson, J., Luo, J., & Payne, R. (2012). Exchange rate determination and inter-market order flow effects. *The European Journal of Finance, 18*(9), 823–840.

Danielsson, J., & Payne, R. (2012). Liquidity determination in an order-driven market. *The European Journal of Finance, 18*(9), 799–821.

Danielsson, J., & Saltoglu, B. (2003). Anatomy of a market crash: A market microstructure analysis of the Turkish overnight liquidity crisis. *EFA 2003 Annual Conference Paper*(965).

De Goeij, P., & Marquering, W. (2004). Modeling the conditional covariance between stock and bond returns: a multivariate GARCH approach. *Journal of Financial Econometrics, 2*(4), 531–564.

Della Corte, P., Sarno, L., & Tsiakas, I. (2011). Why does order flow forecast exchange Rates? *CFA Society Masterclass Series: London*.

Della Corte, P., Sarno, L., & Tsiakas, I. (2012). Volatility and correlation timing in active currency management. *Handbook of Exchange Rates. Hoboken, NJ: Wiley*.

Diebold, F., & Mariano, R. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics, 13*(3), 253–263.

Dijk, D. V., Munandar, H., & Hafner, C. (2011). The Euro introduction and noneuro currencies. *Applied Financial Economics, 21*(1–2), 95–116.

Dunis, C., & Williams, M. (2002). Modelling and trading the EUR/USD exchange rate: Do neural network models perform better? *Derivatives use, trading and regulation, 8*(3), 211-239.

Dunne, P., Hau, H., & Moore, M. (2010). International order flows: explaining equity and exchange rate returns. *Journal of International Money and Finance, 29*(2), 358–386.

Engle, R. (2002). Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics, 20*(3), 339–350.

Engle, R. (2009). *Anticipating correlations. A new paradigm for risk management*: Princeton University Press: Princeton NJ.

Engle, R., & Colacito, R. (2012). Testing and valuing dynamic correlations for asset allocation. *Journal of Business & Economic Statistics*.

Engle, R., & Kroner, K. (1995). Multivariate simultaneous generalized ARCH. *Econometric Theory, 11*(1), 122–150.

Engle, R., & Sheppard, K. (2001). *Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH*. National Bureau of Economic Research: Cambridge, MA.

Evans, M., & Lyons, R. (2002a). Order Flow and Exchange Rate Dynamics. *Journal of Political Economy, 110*(1), 170–180.

Evans, M., & Lyons, R. (2002b). Informational integration and FX trading. *Journal of International Money and Finance, 21*(6), 807–831.

Evans, M., & Lyons, R. (2002c). Time-varying liquidity in foreign exchange. *Journal of Monetary Economics, 49*(5), 1025–1051.

Evans, M., & Lyons, R. (2005a). Do currency markets absorb news quickly? *Journal of International Money and Finance, 24*(2), 197–217.

Evans, M., & Lyons, R. (2005b). Meese-Rogoff redux: Micro-based exchange-rate forecasting. *American Economic Review, 95*(2).

Evans, M., & Lyons, R. (2006). Understanding order flow. *International Journal of Finance & Economics, 11*(1), 3–23.

Evans, M., & Lyons, R. (2008). How is macro news transmitted to exchange rates? *Journal of Financial Economics, 88*(1), 26–50.

Evans, M., & Lyons, R. (2009). Forecasting exchange rates and fundamentals with order flow. *Georgetown University and UC Berkeley Working Paper*.

Fama, E. (1970). Efficient capital markets: a rReview of theory and empirical work. *The Journal of Finance, 25*(2), 383–417.

Fama, E., & Blume, M. (1966). Filter rules and stock-market trading. *The Journal of Business, 39*(1), 226–241.

Fernández-Rodríguez, F., González-Martel, C., & Sosvilla-Rivero, S. (2000). On the profitability of technical trading rules based on artificial neural networks: Evidence from the Madrid stock market. *Economics Letters, 69*(1), 89-94.

Fleming, J., Ostdiek, B., & Whaley, R. (1995). Predicting stock market volatility: a new measure. *Journal of Futures Markets, 15*(3), 265–302.

Frömmel, M., Kiss, M., & Pintér, K. (2011). Macroeconomic announcements, communication and order flow on the Hungarian foreign exchange market. *International Journal of Finance & Economics, 16*(2), 172–188.

Gehrig, T., & Menkhoff, L. (2004). The use of flow analysis in foreign exchange: exploratory evidence. *Journal of International Money and Finance, 23*(4), 573–594.

Gehrig, T., & Menkhoff, L. (2006). Extended evidence on the use of technical analysis in foreign exchange. *International Journal of Finance & Economics, 11*(4), 327–338.
Gençay, R. (1998). The predictability of security returns with simple technical trading rules. *Journal of Empirical Finance, 5*(4), 347–359.
Gencay, R., & Stengos, T. (1998). Moving average rules, volume and the predictability of security returns with feedforward networks. *Journal of Forecasting, 17*, 401-414.
Giacomini, R., & White, H. (2006). Tests of conditional predictive ability. *Econometrica, 74*(6), 1545–1578.
Giot, P. (2005). Relationships between implied volatility indexes and stock index returns. *Journal of Portfolio Management, 2005, 31*(3), 92–100.
Goetzmann, W., Ingersoll, J., Spiegel, M., & Welch, I. (2007). Portfolio performance manipulation and manipulation-proof performance measures. *Review of Financial Studies, 20*(5), 1503–1546.
Gradojevic, N. (2007). Non-linear, hybrid exchange rate modeling and trading profitability in the foreign exchange market. *Journal of Economic Dynamics and Control, 31*(2), 557–574.
Gradojevic, N., & Gençay, R. (2013). Fuzzy logic, trading uncertainty and technical trading. *Journal of Banking & Finance, 37*(2), 578–586.
Gradojevic, N., & Lento, C. (2015). Multiscale analysis of foreign exchange order flows and technical trading profitability. *Economic Modelling, 47*, 156–165.
Gradojevic, N., Lento, C., & Wright, C. (2007). Investment information content in Bollinger Bands? *Applied Financial Economics Letters, 3*(4), 263-267.
Gradojevic, N., & Yang, J. (2006). Non-linear, non-parametric, non-fundamental exchange rate forecasting. *Journal of Forecasting, 25*(4), 227–245.
Hansen, P. (2005). A test for superior predictive ability. *Journal of Business & Economic Statistics, 23*(4), 365–380.
Hansen, P. R., Large, J., & Lunde, A. (2008). Moving average-based estimators of integrated variance. *Econometric Reviews, 27*(1-3), 79-111.
Hibbert, A. M., Daigler, R., & Dupoyet, B. (2008). A behavioral explanation for the negative asymmetric return–volatility relation. *Journal of Banking & Finance, 32*(10), 2254–2266.
Hsu, P.-H., Taylor, M., & Wang, Z. (2016). Technical trading: Is it still beating the foreign exchange market? *Journal of International Economics, 102*, 188-208.
Hyde, S., Bredin, D., & Nguyen, N. (2007). Correlation dynamics between Asia-Pacific, EU and US stock returns. *International Finance Review, 8*, 39-61.
Jasic, T., & Wood, D. (2004). The profitability of daily stock market indices trades based on neural network predictions: Case study for the S&P 500, the DAX, the TOPIX and the FTSE in the period 1965–1999. *Applied Financial Economics, 14*(4), 285-297.
Jobson, J., & Korkie, B. (1981). Performance hypothesis testing with the Sharpe and Treynor measures. *The Journal of Finance, 36*(4), 889–908.
Kalotychou, E., Staikouras, S., & Zhao, G. (2014). The role of correlation dynamics in sector allocation. *Journal of Banking & Finance, 48*, 1–12.
Kasch, M., & Caporin, M. (2013). Volatility threshold dynamic conditional correlations: an international analysis. *Journal of Financial Econometrics*, 1–28.
Kearney, C., & Potì, V. (2006). Correlation dynamics in European equity markets. *Research in International Business and Finance, 20*(3), 305–321.
Kho, B.-C. (1996). Time-varying risk premia, volatility, and technical trading rule profits: Evidence from foreign currency futures markets. *Journal of Financial Economics, 41*(2), 249–290.

Killeen, W., Lyons, R., & Moore, M. (2006). Fixed versus flexible: Lessons from EMS order flow. *Journal of International Money and Finance, 25*(4), 551–579.
King, M., Osler, C., & Rime, D. (2013). The market microstructure approach to foreign exchange: looking back and looking forward. *Journal of International Money and Finance, 38*, 95–119.
King, M., Sarno, L., & Sojli, E. (2010). Timing exchange rates using order flow: The case of the Loonie. *Journal of Banking & Finance, 34*(12), 2917–2928.
Kroner, K., & Ng, V. (1998). Modeling asymmetric comovements of asset returns. *Review of Financial Studies, 11*(4), 817–844.
Kuper, G. (2007). Dynamic conditional correlation analysis of financial market interdependence: An application to Thailand and Indonesia. *Journal of Asian Economics, 18*(4), 670-684.
Laurent, S., Rombouts, J., & Violante, F. (2012). On the forecasting accuracy of multivariate GARCH models. *Journal of Applied Econometrics, 27*(6), 934–955.
Laurent, S., Rombouts, J., & Violante, F. (2013). On loss functions and ranking forecasting performances of multivariate volatility models. *Journal of Econometrics, 173*(1), 1-10.
LeBaron, B. (1999). Technical trading rule profitability and foreign exchange intervention. *Journal of International Economics, 49*(1), 125–143.
Lee, V., & Tshung-Wong, H. (2007). A multivariate neuro-fuzzy system for foreign currency risk management decision making. *Neurocomputing, 70*(4–6), 942-951.
Lento, C. (2009). Combined signal approach: evidence from the Asian–Pacific equity markets. *Applied Economics Letters, 16*(7), 749–753.
Li, X.-M. (2011). How do exchange rates co-move? A study on the currencies of five inflation-targeting countries. *Journal of Banking & Finance, 35*(2), 418–429.
Liu, L., & Zhang, T. (2015). Economic policy uncertainty and stock market volatility. *Finance Research Letters, 15*, 99–105.
Lo, A. (2002). The statistics of Sharpe ratios. *Financial Analysts Journal, 58*(4), 36–52.
Lo, A., & MacKinlay, A. (2002). *A non-random walk down Wall Street*: Princeton University Press: Princeton, NJ.
Lobato, I., Nankervis, J., & Savin, N. (2001). Testing for autocorrelation using a modified Box-Pierce Q- test. *International Economic Review, 42*(1), 187–205.
Love, R., & Payne, R. (2008). Macroeconomic news, order flows, and exchange Rates. *Journal of Financial and Quantitative Analysis, 43*(02), 467–488.
Luo, J. (2001). Market conditions, order flow and exchange rates determination.Typescript. London School of Economics.
Lyons, R., & Moore, M. (2009). An information approach to international currencies. *Journal of International Economics, 79*(2), 211–221.
Marsh, I., & Miao, T. (2012). High-frequency information content in end-user foreign exchange order flows. *The European Journal of Finance, 18*(9), 865– 884.
Marshall, B., Cahan, R., & Cahan, J. (2008). Does intraday technical analysis in the U.S. equity market have value? *Journal of Empirical Finance, 15*(2), 199–210.
Meese, R. A., & Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of international economics, 14*(1), 3-24.
Memmel, C. (2003). Performance hypothesis testing with the Sharpe ratio. *Finance Letters, 1*(1), 21-23.
Menkhoff, L. (2010). The use of technical analysis by fund managers: International evidence. *Journal of Banking & Finance, 34*(11), 2573–2586.
Menkhoff, L., Sarno, L., Schmeling, M., & Schrimpf, A. (2012). Currency momentum strategies. *Journal of Financial Economics, 106*(3), 660–684.

Menkhoff, L., & Taylor, M. (2007). The obstinate passion of foreign exchange professionals: technical analysis. *Journal of Economic Literature*, 936–972.
Neely, C., & Weller, P. (1999). Technical trading rules in the European monetary system. *Journal of International Money and Finance, 18*(3), 429–458.
Neely, C., & Weller, P. (2003). Intraday technical trading in the foreign exchange market. *Journal of International Money and Finance, 22*(2), 223–237.
Neely, C., & Weller, P. (2013). Lessons from the evolution of foreign exchange trading strategies. *Journal of Banking & Finance, 37*(10), 3783–3798.
Neely, C., Weller, P., & Ulrich, J. (2009). The adaptive markets hypothesis: evidence from the foreign exchange market. *Journal of Financial and Quantitative Analysis, 44*(02), 467–488.
Neely, C. J., & Weller, P. A. (2012). Technical analysis in the foreign exchange market. *In: James, J., Marsh, I., Sarno, L. (Eds.), Handbook of Exchange Rates, Wiley Publishing Inc., New Jersey, pp. 369–402.*
Opschoor, A., Taylor, N., van der Wel, M., & van Dijk, D. (2014). Order flow and volatility: an empirical investigation. *Journal of Empirical Finance, 28*, 185–201.
Osler, C. (2000). Support for resistance: technical analysis and intraday exchange rates. *Economic Policy Review, 6*(2), 53–68.
Osler, C. (2003). Currency orders and exchange rate dynamics: An explanation for the predictive success of technical analysis. *The Journal of Finance, 58*(5), 1791–1820.
Osler, C. (2005). Stop-loss orders and price cascades in currency markets. *Journal of international Money and Finance, 24*(2), 219-241.
Park, C. H., & Irwin, S. (2007). What do we know about the profitability of technical analysis? *Journal of Economic Surveys, 21*(4), 786–826.
Patton, A. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review, 47*(2), 527–556.
Patton, A. (2009). Copula–Based Models for Financial Time Series. In T. Mikosch, J.-P. Kreiß, R. A. Davis & T. G. Andersen (Eds.), *Handbook of Financial Time Series* (pp. 767-785): Springer: Berlin Heidelberg.
Patton, A. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics, 160*(1), 246-256.
Payne, R. (2003). Informed trade in spot foreign exchange markets: an empirical investigation. *Journal of International Economics, 61*(2), 307–329.
Pesaran, M. H., & Timmermann, A. (1992). A simple nonparametric test of predictive performance. *Journal of Business & Economic Statistics, 10*(4), 461–465.
Pesaran, M. H., & Timmermann, A. (2007). Selection of estimation window in the presence of breaks. *Journal of Econometrics, 137*(1), 134-161.
Politis, D., & Romano, J. (1994). The stationary bootstrap. *Journal of the American Statistical Association, 89*(428), 1303–1313.
Qi, M., & Wu, Y. (2006). Technical trading-rule profitability, data snooping, and reality check: evidence from the foreign exchange market. *Journal of Money, Credit and Banking*, 2135–2158.
Reitz, S. (2006). On the predictive content of technical analysis. *The North American Journal of Economics and Finance, 17*(2), 121-137.
Rime, D., Sarno, L., & Sojli, E. (2010). Exchange rate forecasting, order flow and macroeconomic information. *Journal of International Economics, 80*(1), 72–88.
Rime, D., & Tranvåg, H. (2012). Flows of the Pacific: Asian foreign exchange markets through tranquility and turbulence. *Pacific Economic Review, 17*(3), 434–466.
Riskmetrics. (1996). Riskmetrics technical document. *J.P. Morgan: New York.*

Rogoff, K. (2002). The failure of empirical exchange rate models: No longer new but still true. *Economic Policy Web Essay*.

Sager, M., & Taylor, M. (2006). Under the microscope: the structure of the foreign exchange market. *International Journal of Finance & Economics, 11*(1), 81–95.

Sager, M., & Taylor, M. (2008). Commercially available order flow data and exchange rate movements: Caveat emptor. *Journal of Money, Credit and Banking, 40*(4), 583–625.

Salmon, M., & Schleicher, C. (2006). Pricing multivariate currency options with copulas. *in: J. Rank (Ed.), Copulas: From Theory to Application in Finance.Risk Books: London*.

Schoppen, K. (2012). Exogenous Variables in Dynamic Conditional Correlation Models for Financial Markets. *Doctoral Thesis, University of Bremen*.

Selçuk, F., & Gençay, R. (2006). Intraday dynamics of stock market returns and volatility. *Physica A: Statistical Mechanics and its Applications, 367*, 375–387.

Silvennoinen, A., & Thorp, S. (2013). Financialization, crisis and commodity correlation dynamics. *Journal of International Financial Markets, Institutions and Money, 24*, 42–65.

Smyth, N. (2009). *Order flow and exchange rate changes : a look at the NZD/USD and AUD/USD*: Reserve Bank of New Zealand: Wellington.

Sullivan, R., Timmermann, A., & White, H. (1999). Data-snooping, technical trading rule performance, and the bootstrap. *The Journal of Finance, 54*(5), 1647–1691.

Taylor, J., & Williams, J. (2008). A Black Swan in the Money Market. *National Bureau of Economic Research: Cambridge, MA, No. 13943*.

Vargas, G. (2008). What drives the dynamic conditional correlation of foreign exchange and equity Returns? *Munich press archive research Paper: Munich*

Wang, Y., & Wu, C. (2012). Forecasting energy market volatility using GARCH models: Can multivariate models beat univariate models? *Energy Economics, 34*(6), 2167–2181.

Whaley, R. (2000). The investor fear gauge. *The Journal of Portfolio Management, 26*(3), 12–17.

White, H. (2000). A reality check for data snooping. *Econometrica, 68*(5), 1097–1126.

Wu, C.-C., Chung, H., & Chang, Y.-H. (2012). The economic value of co-movement between oil price and exchange rate using copula-based GARCH models. *Energy Economics, 34*(1), 270–282.

Yamamoto, R. (2012). Intraday technical analysis of individual stocks on the Tokyo Stock Exchange. *Journal of Banking & Finance, 36*(11), 3033–3047.

Zadeh, L. (1965). Fuzzy sets. *Information and control, 8*(3), 338–353.

Zakamulin, V. (2014). A Test of covariance matrix forecasting methods. *Journal of Portfolio Management, 215*.